

MRI

IN PRACTICE

Catherine Westbrook

John Talbot



WILEY Blackwell

MRI in Practice

MRI in Practice

Fifth Edition

Catherine Westbrook

Edd, MSc, FHEA, PgC (Learning & Teaching), DCRR, CTC

Senior Lecturer

Anglia Ruskin University

Cambridge

UK

John Talbot

Edd, MSc, FHEA, PgC (Learning & Teaching), DCRR

Senior Lecturer

Anglia Ruskin University

Cambridge

UK

WILEY Blackwell

This edition first published 2019
© 2019 John Wiley & Sons Ltd

Edition History

Blackwell Science (1e, 1993, 2e, 1998); Blackwell Publishing Ltd (3e, 2005, 4e 2011).

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, except as permitted by law. Advice on how to obtain permission to reuse material from this title is available at <http://www.wiley.com/go/permissions>.

The right of Catherine Westbrook and John Talbot to be identified as the authors of this work has been asserted in accordance with law.

Registered Offices

John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, USA
John Wiley & Sons Ltd, The Atrium, Southern Gate, Chichester, West Sussex, PO19 8SQ, UK

Editorial Office

9600 Garsington Road, Oxford, OX4 2DQ, UK

For details of our global editorial offices, customer services, and more information about Wiley products visit us at www.wiley.com.

Wiley also publishes its books in a variety of electronic formats and by print-on-demand. Some content that appears in standard print versions of this book may not be available in other formats.

Limit of Liability/Disclaimer of Warranty

The contents of this work are intended to further general scientific research, understanding, and discussion only and are not intended and should not be relied upon as recommending or promoting scientific method, diagnosis, or treatment by physicians for any particular patient. In view of ongoing research, equipment modifications, changes in governmental regulations, and the constant flow of information relating to the use of medicines, equipment, and devices, the reader is urged to review and evaluate the information provided in the package insert or instructions for each medicine, equipment, or device for, among other things, any changes in the instructions or indication of usage and for added warnings and precautions. While the publisher and authors have used their best efforts in preparing this work, they make no representations or warranties with respect to the accuracy or completeness of the contents of this work and specifically disclaim all warranties, including without limitation any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives, written sales materials or promotional statements for this work. The fact that an organization, website, or product is referred to in this work as a citation and/or potential source of further information does not mean that the publisher and authors endorse the information or services the organization, website, or product may provide or recommendations it may make. This work is sold with the understanding that the publisher is not engaged in rendering professional services. The advice and strategies contained herein may not be suitable for your situation. You should consult with a specialist where appropriate. Further, readers should be aware that websites listed in this work may have changed or disappeared between when this work was written and when it is read. Neither the publisher nor authors shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

Library of Congress Cataloging-in-Publication Data

Names: Westbrook, Catherine, author. | Talbot, John (Writer on magnetic resonance imaging), author.

Title: MRI in practice / by Catherine Westbrook, John Talbot.

Description: Fifth edition. | Hoboken, NJ : Wiley, 2018. | Includes bibliographical references and index. |

Identifiers: LCCN 2018009382 (print) | LCCN 2018010644 (ebook) | ISBN 9781119391999 (pdf) | ISBN 9781119392002 (epub) | ISBN 9781119391968 (pbk.)

Subjects: | MESH: Magnetic Resonance Imaging--methods | Magnetic Resonance Imaging--instrumentation

Classification: LCC RC78.7.N83 (ebook) | LCC RC78.7.N83 (print) | NLM WN 185

| DDC 616.07/548--dc23
LC record available at <https://lccn.loc.gov/2018009382>

Cover image: Courtesy of John Talbot
Cover design by Wiley

Set in 10/12pt TradeGothicLTStd by Aptara Inc., New Delhi, India

10 9 8 7 6 5 4 3 2 1

Contents

<i>Preface to the Fifth Edition</i>	<i>ix</i>
<i>Acknowledgments</i>	<i>xi</i>
<i>Acronyms</i>	<i>xiii</i>
<i>Nomenclature</i>	<i>xvii</i>
<i>About the Companion Website</i>	<i>xix</i>

Chapter 1	Basic principles	1
Introduction	1	Precession and precessional (Larmor) frequency
Atomic structure	2	10
Motion in the atom	2	Precessional phase
MR-active nuclei	4	Resonance
The hydrogen nucleus	5	MR signal
Alignment	6	The free induction decay (FID) signal
Net magnetic vector (NMV)	8	Pulse timing parameters
		22

Chapter 2	Image weighting and contrast	24
Introduction	24	Relaxation in different tissues
Image contrast	25	T1 contrast
Relaxation	25	T2 contrast
T1 recovery	26	Proton density contrast
T2 decay	27	Weighting
Contrast mechanisms	31	Other contrast mechanisms
		51

Chapter 3 Spin-echo pulse sequences		58
Introduction	58	Inversion recovery (IR)
RF rephasing	59	Short tau inversion recovery (STIR)
Conventional spin-echo	65	Fluid attenuated inversion recovery (FLAIR)
Fast or turbo spin-echo (FSE/TSE)	68	

Chapter 4 Gradient-echo pulse sequences		89
Introduction	89	Incoherent or spoiled gradient-echo
Variable flip angle	90	Reverse-echo gradient-echo
Gradient rephasing	91	Balanced gradient-echo
Weighting in gradient-echo pulse sequences	94	Fast gradient-echo
Coherent or rewind gradient-echo	106	Echo planar imaging

Chapter 5 Spatial encoding		128
Introduction	128	Frequency encoding
Mechanism of gradients	129	Phase encoding
Gradient axes	134	Bringing it all together – pulse sequence timing
Slice-selection	135	

Chapter 6 <i>k</i>-Space		158
Introduction	158	Part 3: Some important facts about <i>k</i> -space!
Part 1: What is <i>k</i> -space?	159	
Part 2: How are data acquired and how are images created from these data?	165	Part 4: How do pulse sequences fill <i>k</i> -space?
		Part 5: Options that fill <i>k</i> -space

Chapter 7 Protocol optimization		209
Introduction	209	Scan time
Signal-to-noise ratio (SNR)	210	Trade-offs
Contrast-to-noise ratio (CNR)	226	Protocol development and modification
Spatial resolution	232	

Chapter 8 Artifacts		242
Introduction	242	Shading artifact 276
Phase mismapping	243	Moiré artifact 277
Aliasing	253	Magic angle 279
Chemical shift artifact	261	Equipment faults 280
Out-of-phase signal cancellation	265	Flow artifacts 280
Magnetic susceptibility artifact	269	Flow-dependent (non-contrast-enhanced) angiography 298
Truncation artifact	272	Black-blood imaging 303
Cross-excitation/cross-talk	273	Phase-contrast MRA 304
Zipper artifact	275	

Chapter 9 Instrumentation		311
Introduction	311	Shim system 328
Magnetism	313	Gradient system 330
Scanner configurations	315	RF system 337
Magnet system	318	Patient transport system 343
Magnet shielding	326	Computer system and graphical user interface 344

Chapter 10 MRI Safety		346
Introduction (and disclaimer)	346	Time-varying gradient magnetic fields 363
Definitions used in MRI safety	347	Cryogenics 365
Psychological effects	350	Safety tips 367
The spatially varying static field	351	Additional resources 368
Electromagnetic (radiofrequency) fields	357	

Glossary
Index

370
387

Preface to the fifth edition

The *MRI in Practice* brand continues to grow from strength to strength. The fourth edition of *MRI in Practice* is an international best-seller and is translated into several languages. At the time of writing, the accompanying *MRI in Practice* course is 26 years old. We have delivered the course to more than 10 000 people in over 20 countries and have a large and growing *MRI in Practice* online community. Our readers and course delegates include a variety of professionals such as radiographers, technologists, radiologists, radiotherapists, veterinary practitioners, nuclear medicine technologists, radiography students, postgraduate students, medical students, physicists, and engineers.

The unique selling point of *MRI in Practice* has always been its user-friendly approach to physics. Difficult concepts are explained as simply as possible and supported by clear diagrams, images, and animations. Clinical practitioners are not usually interested in pages of math and just want to know how it essentially “all works.” We believe that *MRI in Practice* is so popular because it speaks your language without being oversimplistic.

This fifth edition has had a significant overhaul and specifically plays to the strengths of the *MRI in Practice* brand. We have created a synergy between the book and the course so that they are best able to support your learning. We purposefully focus on physics in this edition and on essential concepts. It is important to get the fundamentals right, as they underpin more specialist areas of practice. There are completely new chapters on MRI equipment and safety, and substantially revised and expanded chapters on gradient-echo pulse sequences, *k*-space, artifacts, and angiography. The very popular learning tips and analogies from previous editions are expanded and revised. There is also a new glossary, lots of new diagrams and images, and suggestions for further reading for those who wish to delve deeper into physics. The accompanying website includes new questions and additional animations. We also include some equations in this edition, but don't worry: they are there only for those who like equations, and we explain what they mean in a user-friendly style.

However, probably the most significant change in this edition is the inclusion of scan tips. Throughout the book, your attention is drawn to how theory applies to practice. Scan tips are specifically used to alert you to what is going on “behind the scenes” when you select a parameter in the scan protocol. We hope this helps you make the connection between theory and practice. Physics in isolation is of little value to the clinical practitioner. What is important is how this knowledge is applied. We stand by the *MRI in Practice* philosophy that physics does not have to be difficult, and we hope that our readers, old and new, find these changes helpful. Richard Feynman, who is considered one of the finest physics teachers of all time, was renowned for his ability

to transfer his deep understanding of physics to the page with clarity and a minimum of fuss. He believed that it is unnecessary to make physics more complicated than it need be. Our aspiration is that the fifth edition of *MRI in Practice* emulates his way of thinking.

We hope that the many fans of *MRI in Practice* around the world continue to enjoy and learn from it. A big thank you for your continued support and happy reading!

Catherine Westbrook
John Talbot
November 2017
United Kingdom

Acknowledgments

Many thanks to all my loved ones for their continued support, especially Maggie Barbieri (my mother, whose brain scans feature many times in all the editions of this book and in the *MRI in Practice* course for the last 26 years. She must have the most viewed brain in the world!), Francesca Bellavista, Amabel Grant, Adam, Ben and Maddie Westbrook.

Catherine Westbrook

I'd like to thank my family Dannie, Joey, and Harry for bringing coffee, biscuits, and occasionally gin and tonic. I would also like to take the opportunity to acknowledge the work of a great MRI pioneer, Prof. Sir Peter Mansfield, who died this year. Prof. Mansfield's team created the first human NMR image in 1976, and he kindly shared all of his most important research papers with me when I first started writing about this amazing field.

John Talbot

Acronyms

Generic	Siemens	GE	Philips	Hitachi	Toshiba
<i>Pulse sequences</i>					
Conventional spin-echo (SE)	SE	SE	SE	SE	SE
Turbo spin-echo (TSE)	TSE	FSE	TSE	FSE	FSE
Single-shot TSE (SS-TSE)	HASTE	SS-FSE	SS-TSE	SS-FSE	FASE
TSE (with restoration pulse)	RESTORE	FRFSE	DRIVE	driven equilibrium FSE	T2 Puls FSE
Inversion recovery (IR)	IR	IR/MPIR	IR	IR	IR
Fast inversion recovery	TIR	Fast IR	IR-TSE	IR	IR
Short tau IR (STIR)	STIR	STIR	STIR	STIR	fast STIR
Fluid-attenuated IR (FLAIR)	turbo dark fluid	FLAIR	FLAIR	FLAIR	fast FLAIR
Gradient-echo (GRE)	GRE	GRE	FFE	GE	field echo
Coherent gradient-echo	FISP	GRASS	FFE	rephased SARGE	SSFP
Incoherent gradient-echo	FLASH	SPGR	T1 FFE	spoiled SARGE	fast FE
Reverse-echo gradient-echo	PSIF	SSFP	T2 FFE	time-reversed SARGE	—
Balanced gradient-echo	true FISP	FIESTA	BFFE	balanced SARGE	true SSFP
Echo-planar imaging (EPI)	EPI	EPI	EPI	EPI	EPI
Double-echo steady state	DESS	—	—	—	—
Balanced dual excitation	CISS	FIESTA-C	—	phase balanced SARGE	—
Multi-echo-data-image-combination	MEDIC	MERGE	MFFE	—	—
Fast gradient-echo	turbo FLASH	fast GRE, fast SPGR	TFE	RGE	Fast FE
Hybrid sequence	TGSE	—	GRASE	—	Hybrid EPI

Generic	Siemens	GE	Philips	Hitachi	Toshiba
<i>Contrast parameters</i>					
Repetition time (TR)	TR	TR	TR	TR	TR
Time to echo (TE)	TE	TE	TE	TE	TE
Time from inversion (TI)	TI	TI	TI	TI	TI
Flip angle	flip angle	flip angle	flip angle	flip angle	Flip angle
Number of echoes (in TSE)	turbo factor	ETL	turbo factor	shot factor	ETL
<i>b</i> factor/value	<i>b</i> factor	<i>b</i> factor	<i>b</i> factor	<i>b</i> factor	<i>b</i> factor
<i>Geometry parameters</i>					
Field of view (FOV)	FOV (mm)	FOV (cm)	FOV (mm)	FOV (mm)	FOV (mm)
Rectangular FOV	FOV phase	PFOV	rectangular FOV	rectangular FOV	rectangular FOV
Slice gap	distance factor	gap	gap	slice interval	gap
<i>Data acquisition parameters</i>					
Averages	average	NEX	NSA	NSA	NSA
Bandwidth	bandwidth (Hz/pixel)	receive bandwidth (KHz)	fat water shift (pixel)	bandwidth (KHz)	bandwidth (KHz)
Variable bandwidth	optimized bandwidth	variable bandwidth	optimized bandwidth	variable bandwidth	matched bandwidth
Partial averaging	half Fourier	fractional NEX	half scan	half scan	AFI
Partial echo	asymmetric echo	partial echo	partial echo	half echo	matched bandwidth
Parallel imaging (image based)	mSENSE	ASSET	SENSE	RAPID	SPEEDER
Parallel imaging (<i>k</i> -space based)	GRAPPA	ARC	—	—	—
<i>Artifact reduction techniques</i>					
Radial <i>k</i> -space filling	BLADE	PROPELLOR	multiVane	RADAR	JET
Gradient moment rephasing	GMR/flow comp	flow comp	flow comp/FLAG	GR	FC
Presaturation	pre SAT	Sat	REST	Pre SAT	Pre SAT
Moving sat pulse	travel SAT	walking SAT	travel REST	Sequential pre SAT	BFAST
Fat saturation	fat SAT	chem SAT	SPIR	Fat Sat	MFOFT
Out-of-phase imaging	DIXON	IDEAL	ProSET	Water excitation	PASTA
Respiratory compensation	respiratory gated	respiratory compensation	PEAR	MAR	respiratory gated
Antialiasing (frequency)	oversampling	antialiasing	frequency oversampling	frequency oversampling	frequency wrap suppression
Antialiasing (phase)	phase oversampling	no phase wrap	fold-over suppression	antiwrap	phase wrap suppression

Generic	Siemens	GE	Philips	Hitachi	Toshiba
<i>Special techniques</i>					
Volume TSE variable flip angle	SPACE	CUBE	VISTA	—	—
Volume gradient-echo	VIBE	LAVA-XV	THRIVE	TIGRE	—
Dynamic MRA	TWIST	TRICKS-SV	keyhole (4d Trak)	—	—
Noncontrast MRA gradient-echo	NATIVE – true FISP	inhance inflow IR	B-TRANCE	VASC ASL	TIME-SLIP
Noncontrast MRA spin- echo	NATIVE-SPACE	—	TRANCE	VASC FSE	FBI
Susceptibility weighting	SWI	SWAN	Venous BOLD	—	—
High-resolution breast imaging	IEWS	VIBRANT-XV	BLISS	—	RADIANCE
Diffusion-weighted imaging	DWI	DWI	DWI	DWI	DWI
Diffusion tensor imaging	DTI	DTI	diffusion tensor imaging	—	DTI
Body diffusion imaging	REVEAL	—	DWIBS	—	body vision

Nomenclature

S	spin quantum number	
N^+	number of spins in the high-energy population (Boltzmann)	
N^-	number of spins in the low-energy population (Boltzmann)	
ΔE	energy difference between high- and low-energy populations (Boltzmann)	J
k	Boltzmann's constant	J/K
T	temperature of the tissue	K
ω_0	precessional or Larmor frequency	MHz
γ	gyromagnetic ratio	MHz/T
B_0	external magnetic field strength	T
E	energy of a photon	J
h	Planck's constant	J/s
θ	flip angle	°
ω_1	precessional frequency of B_1	μ T
B_1	magnetic field associated with the RF excitation pulse	mT
τ	duration of the RF excitation pulse	ms
ϵ	emf	V
N	number of turns in a coil	
$d\Phi$	changing magnetic flux in a single loop	V/s
dt	changing time	s
Mz_t	amount of longitudinal magnetization at time t	
Mz	full longitudinal magnetization	
Mxy_t	amount of transverse magnetization at time t	
Mxy	full transverse magnetization	
SI	signal intensity in a tissue	
ΔB_0	variation in magnetic field	ppm
G	gradient amplitude	mT/m
δ	gradient duration	ms
Δ	time between two gradient pulses	ms

b	b value or b factor	s/mm ²
ST	scan time	s
E_s	echo spacing in turbo spin-echo (TSE)	ms
t	time from inversion (TI)	ms
Ernst	Ernst angle	°
TE_{eff}	effective TE	ms
TE_{act}	TE set at the console	ms
B_p	magnetic field strength at a point along the gradient	T
SI_t	slice thickness	mm
TBW	transmit bandwidth	KHz
ω_{sampling}	digital sampling frequency	KHz
ΔT_s	sampling interval	ms
ω_{Nyquist}	Nyquist frequency	KHz
RBW	receive bandwidth	KHz
W_s	sampling window	ms
$M(f)$	frequency matrix	
$M(p)$	phase matrix	
N_s	number of slice locations	
$G(p)$	max amplitude of the phase encoding gradient	mT/m
ϕ	incremental step between each k -space line	
$G(f)$	amplitude of the frequency encoding gradient	mT/m
FOV(f)	frequency FOV	cm
σ	standard deviation of background signal or noise	
S_p	separation between ghosts due to motion p	pixels
T_m	period of motion of something moving in the patient	ms
Re	Reynolds number	
d	density of blood	g/cm ³
v	velocity of flow	cm/s
m	diameter of a vessel	cm
vis	viscosity of blood	g/cm s
f_p	perceived frequency	KHz
f_t	actual frequency	KHz
ω_{csf}	chemical shift frequency difference between fat and water	Hz
C_s	chemical shift (3.5 ppm or 3.5×10^{-6})	ppm
CS_p	pixel shift	mm
H_0	magnetic intensity	A/m
q	charge of a particle	C
F	Lorentz force (total emf on a charged particle)	V
E	electric field vector	
B	magnetic field vector	

About the companion website

This book is accompanied by a companion website:

www.wiley.com/go/westbrook/mriinpractice



The website includes:

- Brand new 3D animations of more complex concepts from the book
- 100 short-answer questions to aid learning and understanding.

1

Basic principles

Introduction	1	Precession and precessional (Larmor) frequency	10
Atomic structure	2		
Motion in the atom	2	Precessional phase	13
MR-active nuclei	4	Resonance	13
The hydrogen nucleus	5	MR signal	18
Alignment	6	The free induction decay (FID) signal	20
Net magnetic vector (NMV)	8	Pulse timing parameters	22

After reading this chapter, you will be able to:

- Describe the structure of the atom.
- Explain the mechanisms of alignment and precession.
- Understand the concept of resonance and signal generation.

INTRODUCTION

The basic principles of magnetic resonance imaging (MRI) form the foundation for further understanding of this complex subject. It is important to grasp these ideas before moving on to more complicated topics in this book.

There are essentially two ways of explaining the fundamentals of MRI: classically and via quantum mechanics. **Classical theory** (accredited to Sir Isaac Newton and often called Newtonian theory) provides a mechanical view of how the universe (and therefore how MRI) works. Using classical theory, MRI is explained using the concepts of mass, spin, and angular momentum on a large or bulk scale. **Quantum theory** (accredited to several individuals including Max Planck, Albert Einstein, and Paul Dirac) operates at a much smaller, subatomic scale and refers to the energy levels of protons, neutrons, and electrons. Although classical theory is often used to describe physical principles on a large scale and quantum theory on a subatomic level, there is evidence that all physical principles are explained using quantum concepts [1]. However, for our purposes, this chapter mainly relies on classical perspectives because they are generally easier to understand. Quantum theory is only used to provide more detail when required.

In this chapter, we explore the properties of atoms and their interactions with magnetic fields as well as the mechanisms of excitation and relaxation.

ATOMIC STRUCTURE

All things are made of **atoms**. Atoms are organized into **molecules**, which are two or more atoms arranged together. The most abundant atom in the human body is **hydrogen**, but there are other elements such as oxygen, carbon, and nitrogen. Hydrogen is most commonly found in molecules of water (where two hydrogen atoms are arranged with one oxygen atom; H_2O) and fat (where hydrogen atoms are arranged with carbon and oxygen atoms; the number of each depends on the type of fat).

The atom consists of a central nucleus and orbiting **electrons** (Figure 1.1). The nucleus is very small, one millionth of a billionth of the total volume of an atom, but it contains all the atom's mass. This mass comes mainly from particles called **nucleons**, which are subdivided into **protons** and **neutrons**. Atoms are characterized in two ways.

- The **atomic number** is the sum of the protons in the nucleus. This number gives an atom its chemical identity.
- The **mass number** or **atomic weight** is the sum of the protons and neutrons in the nucleus.

The number of neutrons and protons in a nucleus is usually balanced so that the mass number is an even number. In some atoms, however, there are slightly more or fewer neutrons than protons. Atoms of elements with the same number of protons but a different number of neutrons are called **isotopes**.

Electrons are particles that spin around the nucleus. Traditionally, this is thought of as analogous to planets orbiting around the sun with electrons moving in distinct shells. However, according to quantum theory, the position of an electron is not predictable as it depends on the energy of an individual electron at any moment in time (this is called Heisenberg's Uncertainty Principle).

Some of the particles in the atom possess an electrical charge. Protons have a positive electrical charge, neutrons have no net charge, and electrons are negatively charged. Atoms are electrically stable if the number of negatively charged electrons equals the number of positively charged protons. This balance is sometimes altered by applying energy to knock out electrons from the atom. This produces a deficit in the number of electrons compared with protons and causes electrical instability. Atoms in which this occurs are called **ions** and the process of knocking out electrons is called **ionization**.

MOTION IN THE ATOM

Three types of motion are present within the atom (Figure 1.1):

- Electrons spinning on their own axis
- Electrons orbiting the nucleus
- The nucleus itself spinning about its own axis.

The principles of MRI rely on the spinning motion of specific nuclei present in biological tissues. There are a limited number of spin values depending on the atomic and mass numbers. A nucleus has no spin if it has an even atomic and mass number, e.g. six protons and six neutrons, mass number 12. In nuclei that have an even mass number caused by an even number of protons and neutrons, half of the nucleons spin in one direction and half in the other. The forces of rotation cancel out, and the nucleus itself has no net spin.





Figure 1.1 The atom.

However, in nuclei with an odd number of protons, an odd number of neutrons, or an odd number of both protons and neutrons, the spin directions are not equal and opposite, so the nucleus itself has a net spin or **angular momentum**. Typically, these are nuclei that have an odd number of protons (or odd atomic number) and therefore an odd mass number. This means that their spin has a half-integral value, e.g. $\frac{1}{2}$, $\frac{5}{2}$. However, this phenomenon also occurs in nuclei with an odd number of both protons and neutrons resulting in an even mass number. This means that it has a whole integral spin value, e.g. 1, 2, 3. Examples are ^6Li (which is made up of three protons and three neutrons) and ^{14}N (seven protons and seven neutrons). However, these elements are largely unobservable in MRI so, in general, only nuclei with an odd mass number or atomic weight are used. These are known as **MR-active nuclei**.

Learning tip: What makes a proton spin and why is it charged?

On a subnuclear level, individual protons are made up of quarks, each of which possesses the characteristics of alignment and spin. The net charge and spin of a proton are a consequence of its quark composition. The proton consists of three spinning quarks. Two quarks spin up and the other spins down. The net spin of the proton ($1/2$) is caused by the different alignment of the quarks. The net charge of the proton is caused by each spin-up quark having a charge of $+2/3$, while the spin-down quark has a charge of $-1/3$ (total charge $+1$) [2].

MR-ACTIVE NUCLEI

MR-active nuclei are characterized by their tendency to align their axis of rotation to an applied magnetic field. This occurs because they have angular momentum or spin and, as they contain positively charged protons, they possess an electrical charge. The law of electromagnetic induction (determined by Michael Faraday in 1833) refers to the connection between electric and magnetic fields and motion (explained later in this chapter). Faraday's law determines that a moving electric field produces a magnetic field and vice versa.

MR-active nuclei have a net electrical charge (electric field) and are spinning (motion), and, therefore, automatically acquire a magnetic field. In classical theory, this magnetic field is denoted by a **magnetic moment**. The magnetic moment of each nucleus has vector properties, i.e. it has size (or magnitude) and direction. The total magnetic moment of the nucleus is the vector sum of all the magnetic moments of protons in the nucleus.

Important examples of MR-active nuclei, together with their mass numbers are listed below:

- ^1H (hydrogen)
- ^{13}C (carbon)
- ^{15}N (nitrogen)
- ^{17}O (oxygen)
- ^{19}F (fluorine)
- ^{23}Na (sodium).

Table 1.1 Characteristics of common elements in the human body.


Element	Protons	Neutrons	Nuclear spin	% Natural abundance
^1H (protium)	1	0	1/2	99.985
^{13}C (carbon)	6	7	1/2	1.10
^{15}N (nitrogen)	7	8	1/2	0.366
^{17}O (oxygen)	8	9	5/2	0.038

THE HYDROGEN NUCLEUS

The isotope of hydrogen called **protium** is the most commonly used MR-active nucleus in MRI. It has a mass and atomic number of 1, so the nucleus consists of a single proton and has no neutrons. It is used because hydrogen is very abundant in the human body and because the solitary proton gives it a relatively large magnetic moment. These characteristics mean that the maximum amount of available magnetization in the body is utilized.

Faraday's law of electromagnetic induction states that a magnetic field is created by a charged moving particle (that creates an electric field). The protium nucleus contains one positively charged proton that spins, i.e. it moves. Therefore, the nucleus has a magnetic field induced around it and acts as a small magnet. The magnet of each hydrogen nucleus has a north and a south pole of equal strength. The north/south axis of each nucleus is represented by a magnetic moment and is used in classical theory.

In diagrams in this book, the magnetic moment is shown by an arrow. The length of the arrow represents the magnitude of the magnetic moment or the strength of the magnetic field that surrounds the nucleus. The direction of the arrow denotes the direction of alignment of the magnetic moment as in Figure 1.2.

**Figure 1.2** The magnetic moment of the hydrogen nucleus.

Learning tip: Use of terms – MRI active nuclei

From now on in this book, the terms spin, nucleus, or proton are all used when we refer to the ^1H nucleus, protium. However, it is important to remember that the other types of MR-active nuclei behave in a similar way when exposed to an external magnetic field. Some of these, phosphorous, sodium, and carbon, are used in certain MRI applications, but the majority use protium.

Table 1.2 Things to remember – basics of the atom.

Hydrogen is the most abundant element in the human body
Nuclei that are available for MRI are those that exhibit a net spin
As all nuclei contain at least one positively charged proton, those that also spin have a magnetic field induced around them
An arrow called a magnetic moment denotes the magnetic field of a nucleus in classical theory

ALIGNMENT

In the absence of an applied magnetic field, the magnetic moments of hydrogen nuclei are randomly orientated and produce no overall magnetic effect. However, when placed in a strong static external magnetic field (shown as a white arrow on Figure 1.3 and termed B_0), the magnetic moments of hydrogen nuclei orientate with this magnetic field. This is called **alignment**. Alignment is best described using classical and quantum theories as follows.




Figure 1.3 Alignment – classical theory.

Classical theory uses the *direction of the magnetic moments* of spins (hydrogen nuclei) to illustrate alignment.

- **Parallel alignment:** Alignment of magnetic moments in the same direction as the main B_0 field (also referred to as spin-up).
- **Antiparallel alignment:** Alignment of magnetic moments in the opposite direction to the main B_0 field (also referred to as spin-down) (Figure 1.3).

After alignment, there are always more spins with their magnetic moments aligned parallel than antiparallel. The net magnetism of the patient (termed the **net magnetic vector, NMV**) is therefore aligned parallel to the main B_0 field in the **longitudinal plane** or **z-axis**.

Learning tip: Magnetic moments vs hydrogen nucleus

A very common misunderstanding is that when a patient is exposed to B_0 , the hydrogen nucleus itself aligns with the external magnetic field. This is incorrect. It is the magnetic moments of hydrogen nuclei that align with B_0 not hydrogen nuclei themselves. The hydrogen nucleus does not change direction but merely spins on its axis.

Quantum theory uses the *energy level* of the spins (or hydrogen nuclei) to illustrate alignment. Protons of hydrogen nuclei couple with the external magnetic field B_0 (termed **Zeeman interaction**) and cause a discrete number of energy states. For hydrogen nuclei, there are only two possible energy states (Figure 1.4):

- **Low-energy nuclei** do not have enough energy to oppose the main B_0 field (shown as a white arrow on Figure 1.4). These are nuclei that align their magnetic moments parallel or spin-up to the main B_0 field in the classical description (shown in blue in Figure 1.4).
- **High-energy nuclei** do have enough energy to oppose the main B_0 field. These are nuclei that align their magnetic moments antiparallel or spin-down to the main B_0 field in the classical description (shown in red in Figure 1.4).

Quantum theory explains why hydrogen nuclei only possess two energy states – high or low (Equation (1.1)). This means that the magnetic moments of hydrogen spins only align in the parallel or antiparallel directions. They cannot orientate themselves in any other direction. The number of spins in each energy level is predicted by the **Boltzmann equation** (Equation (1.2)). The difference in energy between these two states is proportional to the strength of the external magnetic field (B_0) (ΔE in the Boltzmann equation). As B_0 increases, the difference in energy between the two energy states increases, and nuclei therefore require more energy to align their magnetic moments in opposition to the main field. Boltzmann's equation also shows that the patient's temperature is an important factor that determines whether a spin is in the high- or low-energy population. In clinical imaging, however, thermal effects are largely discounted, as the patient's temperature is usually similar inside and outside the magnetic field. This is called **thermal equilibrium**.

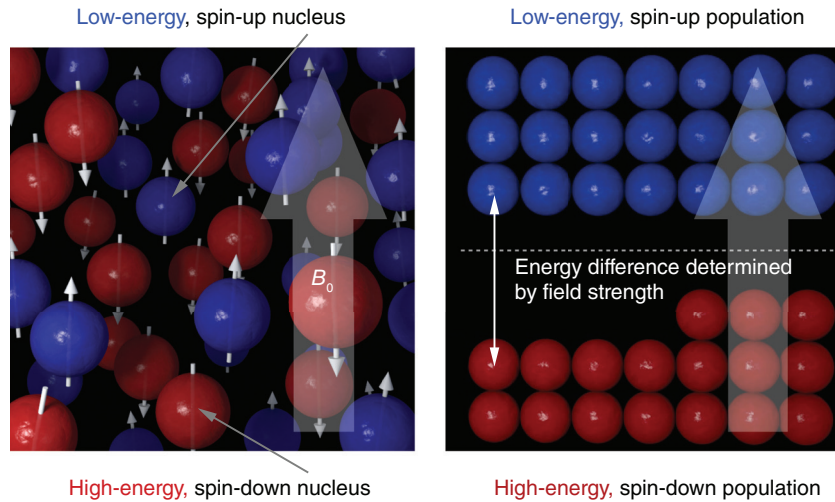


Figure 1.4 Alignment – quantum theory.

Equation 1.1

Number of energy states = $2S + 1$

S is the spin quantum number. The value of S for hydrogen is $\frac{1}{2}$

This equation explains why hydrogen can only possess two energy states. If $S = \frac{1}{2}$, then the number of energy states is $2 \times \frac{1}{2} + 1 = 2$

Learning tip: What is B_0 ?

B_0 refers to the large magnetic field of the MRI scanner. This static magnetic field is measured in teslas (T) using the Systeme Internationale (SI). B is the universally accepted notation for magnetic flux density, and the zero annotation indicates that this is primary magnetic field of the scanner. Other magnetic fields are also used in MRI. These include graded or sloped magnetic fields (called gradients, used to produce images) and an oscillating magnetic field that causes a phenomenon called resonance. This oscillating field is termed B_1 . It has a magnitude several orders lower than B_0 (milliteslas as opposed to teslas).

NET MAGNETIC VECTOR (NMV)

Magnetic moments of hydrogen spins are constantly changing their orientation because, due to Zeeman interaction, they are always moving between high- and low-energy states. Spins gain and lose energy, and their magnetic moments therefore constantly alter their alignment relative to B_0 . In thermal equilibrium, at any moment in time, there are a greater proportion of spins with their magnetic moments aligned in the same direction as B_0 than against it. As there is a larger number aligned parallel, there is always a small excess in this direction that produces a net magnetic moment (Figure 1.5). This is called the NMV and reflects the relative balance between spin-up and spin-down nuclei. It is the sum of all magnetic moments of excess spin-up nuclei and is measurable (in the order of microteslas) [3]. It aligns in the same direction as the main magnetic field in the longitudinal plane or z-axis.




Figure 1.5 The net magnetic vector.

The number of spins that constitute this small excess depends on the number of molecules per gram of tissue and the strength of B_0 . According to Avogadro's law, there are about 6×10^{23} molecules per gram of tissue, and the number of excess spins is in the order of 6×10^{17} per gram of tissue [4]. In thermal equilibrium, the strength of the external field also determines the relative quantities of spin-up to spin-down nuclei because this also affects the difference in energy levels between the two energy states (see Equation (1.2)). As the magnitude of the external magnetic field increases, more magnetic moments line up in the parallel direction because the amount of energy spins must possess to align their magnetic moments in opposition to the stronger field (and line up in the antiparallel direction) increases. As the field strength increases, fewer spins possess enough energy to align their magnetic moments in opposition to the larger B_0 field. As a result, the low-energy population increases in size, the high-energy population decreases in size, and therefore the number of excess number of spins also increases. At 1.5 T, the number in excess is about 4.5 for every million protons; at 3 T, this increases to about 10 per million [5]. Consequently, the NMV also increases in size and is one of the reasons why the signal-to-noise ratio (SNR) increases at higher field strengths (see Chapter 7).

Equation 1.2

$$N^+ / N^- = e^{-\Delta E / kT}$$

N^+ and N^- are the number of spins in the high- and low-energy populations, respectively.

ΔE is the energy difference between the high- and low-energy populations in Joules (J)

k is Boltzmann's constant (1.381×10^{-23} J/K)

T is the temperature of the tissue in Kelvin (K)

This equation enables prediction of the number of spins in the high- and low-energy populations and how this is dependent on temperature. In MRI, thermal equilibrium is presumed in that there are no significant changes in body temperature in the scan room

Table 1.3 Things to remember – alignment.

10

When placed in an external magnetic field, the magnetic moments of hydrogen align in a spin-up, low-energy or spin-down, high-energy orientation
At thermal equilibrium, there are more spin-up, low-energy than spin-down, high-energy spins so the net magnetic vector (NMV) of the patient is orientated in the same direction as B_0
The difference in energy between these populations is mainly determined by the strength of B_0
As B_0 increases the energy difference between the two populations also increases as the number of spin-up, low-energy spins increases relative to the number of spin-down, high-energy spins
The signal-to-noise ratio (SNR) increases at higher values of B_0 (see Chapter 7)

PRECESSION AND PRECESSIONAL (LARMOR) FREQUENCY

Each hydrogen nucleus spins on its axis as in Figure 1.6. The influence of B_0 produces an additional spin or wobble of the magnetic moments of hydrogen around B_0 . This secondary spin is called **precession** and causes the magnetic moments to circle around B_0 . The course they take is called the **precessional path**, and the speed at which they precess around B_0 is called the **precessional frequency**. The precessional frequency is often called the **Larmor frequency** because it is determined by the **Larmor equation** (Equation (1.3)). The unit of precessional frequency is hertz (Hz) where 1 Hz is one cycle or rotation per second (s), and 1 megahertz (MHz) is one million cycles or rotations per second. The magnetic moments of all spin-up and spin-down nuclei precess around B_0 on a precessional path at a Larmor frequency determined by B_0 (Figure 1.7).

Equation 1.3

$$\omega_0 = \gamma B_0 / 2\pi$$

simplified to

$$\omega_0 = \gamma B_0$$

ω_0 is the precessional or Larmor frequency (MHz)
 γ is the gyromagnetic ratio (MHz/T)
 B_0 is the strength of the external magnetic field (T)

This is the Larmor equation. The 2π function enables the conversion of ω_0 from angular to cyclical frequency. As γ is a constant, for a given MR-active nucleus ω_0 is proportional to B_0

The **gyromagnetic ratio** expresses the relationship between angular momentum and the magnetic moment of each MR-active nucleus. It is constant and is expressed as the precessional frequency of the magnetic moment of a specific MR-active nucleus at 1 T. The unit of the gyromagnetic ratio is therefore MHz/T. The gyromagnetic ratio of hydrogen is 42.58 MHz/T. Other MR-active nuclei have different gyromagnetic ratios, so their magnetic moments have different precessional frequencies at the same field strength (Table 1.4).




Figure 1.6 Precession.




Figure 1.7 Precession of the spin-up and spin-down populations.

Table 1.4 Magnetic characteristics of common elements.

Element	Nuclear spin	Gyromagnetic ratio (MHz/T)	Larmor frequency at 1.5 T (MHz)
¹ H (hydrogen)	1/2	42.5774	63.8646
¹³ C (carbon)	1/2	10.7084	16.0621
¹⁵ N (nitrogen)	1/2	4.3173	6.4759
¹⁷ O (oxygen)	5/2	5.7743	8.6614

In addition, magnetic moments of MR-active nuclei have different precessional frequencies at different field strengths. For hydrogen, for example:

- At 1.5 T, the precessional frequency is 63.87 MHz (42.58 MHz × 1.5 T).
- At 1.0 T, the precessional frequency is 42.57 MHz (42.58 MHz × 1.0 T).
- At 0.5 T, the precessional frequency is 21.29 MHz (42.58 MHz × 0.5 T).

These frequencies fall into the **radiofrequency (RF)** band of the electromagnetic spectrum (Figure 1.8).




Figure 1.8 The electromagnetic spectrum.

Learning tip: What does the Larmor equation tell us?

All MR-active nuclei have their own unique gyromagnetic constant or ratio so that when they are exposed to the same field strength, their magnetic moments precess at different frequencies, i.e. magnetic moments of hydrogen precess at a different frequency to magnetic moments of either fluorine or carbon, which are also MR-active nuclei. This allows specific imaging of hydrogen. Other MR-active nuclei are ignored because the precessional or Larmor frequency of their magnetic moments is different to that of hydrogen (we explore how this is done later). In addition, as the gyromagnetic ratio is a constant of proportionality, B_0 is proportional to the Larmor frequency. Therefore, if B_0 increases, the Larmor frequency increases proportionally and vice versa.

PRECESSIONAL PHASE

Phase refers to the position of magnetic moments on their precessional path at any moment in time. The unit of phase is a **radian**. A magnetic moment travels through 360 rad or 360° during one rotation. In this context, **frequency** is the rate of change phase of magnetic moments, i.e. it is a measure of how quickly the phase position of a magnetic moment changes over time. In MRI, the relative phase positions of all magnetic moments of hydrogen are important.

- **Out of phase** or **incoherent** means that magnetic moments of hydrogen are at *different* places on the precessional path at a moment in time.
- **In phase** or **coherent** means that magnetic moments of hydrogen are at the *same* place on the precessional path at a moment in time.

When the only influence is B_0 , the magnetic moments of the nuclei are out of phase with each other, and therefore the NMV does not precess.

Table 1.5 Things to remember – precession.

Magnetic moments of all the spins precess around B_0 at the Larmor frequency that is proportional to B_0 for a given MR-active nucleus. Frequency therefore refers to how fast magnetic moments of spins are precessing and is measured in MHz in MRI

For field strengths used in clinical imaging, the Larmor frequency of hydrogen is in the radio-frequency (RF) band of the electromagnetic spectrum

Phase refers to the position of a magnetic moment of a spin on its precessional path at any moment in time

At equilibrium, the magnetic moments of the spins are out of phase with each other

RESONANCE

Resonance is a phenomenon that occurs when an object is exposed to an oscillating perturbation that has a frequency close to its own natural frequency of oscillation. When a nucleus is exposed to an external force that has an oscillation similar to the natural frequency of its magnetic

moment (its Larmor frequency), the nucleus gains energy from the external source. If energy is delivered at a different frequency to that of the Larmor frequency, resonance does not occur, and the nucleus does not gain energy. As magnetic moments of hydrogen nuclei precess in the RF band of the electromagnetic spectrum, for resonance of hydrogen to occur, an RF pulse of energy is applied at the Larmor frequency of hydrogen. Other MR-active nuclei, whose magnetic moments are aligned with B_0 , do not resonate, because the precessional frequencies of these magnetic moments are different to that of hydrogen. This is because their gyromagnetic ratios are different.

Resonance is achieved by transmitting an RF pulse called an **RF excitation pulse**. This is produced by a transmit coil (see Chapter 9). As with any type of electromagnetic radiation, it consists of an electric and magnetic field that propagate in waves at 90° to each other. These waves have a frequency that resides in the RF band of the electromagnetic spectrum. The RF excitation pulse is derived from the magnetic component only (the electric field produces heat), and unlike the B_0 field, which is stationary, the RF excitation pulse produces an oscillating magnetic field, termed B_1 . The B_1 field is applied at 90° to B_0 at a narrow range or bandwidth of frequencies centered around a central frequency (termed the transmit bandwidth; see Chapters 5 and 6). The magnetic field associated with the RF excitation pulse B_1 is very weak compared with that of the main external field B_0 [6].

The results of resonance – classical theory

From the classical theory perspective, application of the B_1 field in a plane at 90° to B_0 , termed the **transverse plane** or **x-y-axis**, causes magnetic moments of the spins to precess around this axis rather than about the longitudinal plane or z-axis. As we have just learned, the Larmor equation determines that precessional frequency is proportional to the field strength. As the B_1 magnetic field associated with the RF excitation pulse is weak, the magnetic moments of spins precess at a much lower frequency than they do when they are aligned in the longitudinal plane and experience the much larger B_0 field. The transition results in a spiral motion downward of the NMV from the longitudinal to the transverse plane. This spiral motion is called **nutation** and is caused by two precessional motions that happen simultaneously; precession around B_0 and a much slower precession around B_1 [7].

Another consequence of the RF excitation pulse is that the magnetic moments of the spin-up and spin-down nuclei move into phase with each other. Magnetic moments that are in phase (or coherent) are in the same place on their precessional path at any given time. When resonance occurs, all magnetic moments move to the same position on the precessional path and are then in phase (Figure 1.9).

The results of resonance – quantum theory

Application of an RF pulse that causes resonance is termed **excitation**, which means it is “energy-giving.” The RF excitation pulse gives energy to hydrogen nuclei and causes a net increase in the number of high-energy, spin-down nuclei (Figure 1.10). This is because the spin-up, low-energy hydrogen nuclei absorb energy from the RF excitation pulse and move into the high-energy population. At the same time, the spin-down, high-energy nuclei are stimulated to release energy and return to the low-energy state. However, because there are more low-energy spins, the net effect is of energy absorption [8].

Learning tip: **B_0 vs B_1**

15

The RF excitation pulse is characterized by its amplitude (B_1) and its frequency. For resonance to occur, the frequency of the RF excitation pulse must equal the Larmor frequency of magnetic moments of the hydrogen nuclei. If this match occurs, B_1 causes magnetic moments of the hydrogen nuclei to precess in the transverse plane. How fast they precess in the transverse plane is derived from the Larmor equation, which states that precessional frequency is proportional to the field strength (see Equation (1.3)). As B_1 is much smaller than B_0 , magnetic moments of the hydrogen nuclei precess at a much lower frequency than they do before resonance when affected only by B_0 . Before resonance, not only do they precess faster but their magnetic moments are out of phase, and they therefore have no net transverse component. However, when the B_1 field is applied in the transverse plane, magnetic moments align with this field and, in doing so, gain phase coherence. This causes an increase in transverse magnetization. The combination of development of phase coherence and nutation results in coherent magnetization that precesses in the transverse plane. During the RF excitation pulse, the transverse magnetization precesses at a frequency dependent on the amplitude of the B_1 field [4].




Figure 1.9 In phase (coherent) and out of phase (incoherent).

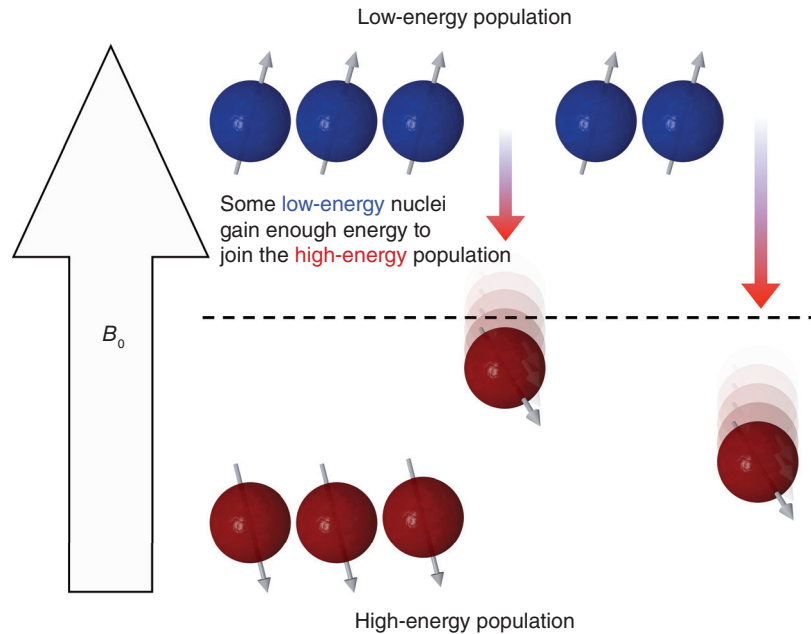


Figure 1.10 Energy transfer during excitation.

If just the right amount of energy is absorbed, the NMV lies in the transverse plane at 90° to B_0 . When it does, it has moved through a **flip** or **tip angle** of 90° (Figure 1.10). The energy and frequency of electromagnetic radiation (including RF) are related to each other, and, consequently, the frequency required to cause resonance is related to the difference in energy between the high-energy and low-energy populations and thus the strength of B_0 (Equation (1.4)). As the field strength increases, the energy difference between the two populations increases so that more energy (higher frequencies) is required to produce resonance.

Equation 1.4

$$E = h\omega_0$$

E is the energy of a photon (Joules, J)
 h is Planck's constant (6.626×10^{-34} J/s)
 ω_0 is the frequency of an electromagnetic wave (Hz)

Planck's constant relates the energy of a photon of electromagnetic radiation to its frequency. Photons are particles that possess energy and at the same time behave like waves that have frequency (wave-particle duality)

$$\Delta E = h\omega_0 = h\gamma B_0$$

ΔE is the energy difference between the spin-up and spin-down populations
 h is Planck's constant (6.626×10^{-34} J/s)
 ω_0 is the precessional or Larmor frequency (MHz)
 γ is the gyromagnetic ratio (MHz/T)

This equation shows that when the energy of the photon matches the energy difference between the spin-up and spin-down populations, energy absorption occurs. This is proportional to the magnetic field strength B_0

Learning tip: The flip angle

The magnitude of the flip angle depends on the amplitude and duration of the RF excitation pulse. Usually, the flip angle is 90° , i.e. the NMV is given enough energy by the RF excitation pulse to move through 90° relative to B_0 . However, as the NMV is a vector, even if flip angles other than 90° are used, there is always a component of magnetization in a plane perpendicular to B_0 . With a flip angle of 90° , the nuclei are given sufficient energy so that the longitudinal NMV is completely transferred into a transverse NMV. When flip angles less than or more than 90° are used, only a portion of the NMV is transferred to the transverse plane. The flip angle depends on the strength of the B_1 field and for how long it is applied (Equation (1.5)).

It can be seen from Equation (1.5) that a flip angle of 180° is caused when the RF excitation pulse is twice the magnitude of that used to produce a 90° flip angle [8]. In quantum mechanics, a 180° RF pulse produces an inversion of the spin populations, i.e. all the low-energy spins have enough energy to locate in the high-energy population and all the high-energy spins have been stimulated to give up their energy and locate in the low-energy population. This is called **saturation** and is caused when the spins are unable to absorb more energy or to be stimulated and release more energy. The amount of RF needed to produce a 90° flip angle is half of this value and relates to equalizing the high- and low-energy spins [6].

Equation 1.5

$$\theta = \omega_1 \tau$$

Therefore from the Larmor equation

$$\theta = \gamma B_1 \tau$$

$$90^\circ = \pi/2 = \gamma B_1 \tau$$

$$180^\circ = \pi = \gamma B_1 \tau$$

θ is the flip angle ($^\circ$)

ω_1 is the precessional frequency of B_1 (μT)

B_1 is the magnetic field associated with the RF excitation pulse (mT)

τ is the duration of the RF excitation pulse (ms)

This equation shows that the flip angle is determined by the strength of the B_1 field and the duration of the pulse. In trigonometry, a factor of 2π represents 360° . A flip angle of 90° can therefore be written as $\pi/2$; a flip angle of 180° is π . Replacing θ with these values shows that an RF pulse producing a flip angle of 90° has either half the power or half the duration of an 180° RF pulse [9].

Analogy: The watch analogy



The terms frequency and phase are used many times in this book, and it is important we understand the difference between them and how they relate to each other. The easiest analogy is the hour hand on an analog watch. Frequency is the time it takes the hour hand to make one revolution of the watch face, i.e. 12 h. The unit of frequency is Hz, where 1 Hz is one cycle or rotation per second. Using the watch analogy, the frequency of the hour hand is $1/4320 \text{ s} = 0.000\,023\,1 \text{ Hz}$ as it moves around the watch face once every 12 h.

The phase of the hour hand, measured in degrees or radians, is the time on the watch, e.g. 1 o'clock, 2 o'clock, which corresponds to its position around the watch face when you look to see what time it is (Figure 1.11). The phase of the hour hand depends on its frequency (speed). If the frequency is correct, then the hour hand always tells the correct time. If the watch goes fast or slow, i.e. the frequency either increases or decreases, then the watch tells an incorrect time.

Imagine a room full of people, with watches that tell perfect time, who are asked to synchronize their watches at 12 noon. One hour later, all their watches say 1 o'clock because they have kept perfect time. They are in phase or coherent because they all tell the same time, and their hour hands are all at the same place on the watch face at the same time. However, if after synchronization the watches on the left-hand side of the room go fast for 1 h, and the watches on the right-hand side of the room go slow for 1 h, then at 1 o'clock they tell different times. The watches on the left-hand side of the room tell a time greater than 1 o'clock, e.g. 1.15 p.m., and those on the right-hand side of the room tell a time less than 1 o'clock, e.g. 12.45 p.m. Therefore, the watches are out of phase or incoherent because they tell different times, and their hour hands are not at the same place on the watch face at the same time.

The phase difference between them depends on their relative frequencies between the time 12 noon and 1 o'clock. If the difference in frequencies is large, then the difference in phase is greater than if the frequency difference is small. Phase and frequency are therefore connected. In this context, the frequency of the hour hand is related to its change of phase over time. In other contexts, used later in this book, frequency is a change of phase over distance. The watch analogy is used many times in this book. Look out for the watch symbol in the margin.

Learning tip: Stationary vs rotating frame of reference

The **stationary frame of reference** refers to the observer (i.e. you) viewing something moving. You and the room you are situated in are stationary, and what you are observing moves. You are an outsider looking in.

The **rotating frame of reference** refers to the observer viewing this from a different perspective. Imagine you are "the thing" that moves: what would the room look like? You are stationary, and the room would appear to move around you because you are now part of the rotating system.

A good example of this is to imagine what happens during an RF excitation pulse. If you were to observe this from the stationary frame of reference, then you would observe nutation of the NMV around B_0 and simultaneously around B_1 . As B_0 is larger than B_1 , the outside observer sees fast precession around B_0 and a much slower spiraling down onto the transverse plane around B_1 . If, however, you were to observe this from inside the rotating frame of reference, then you would see something different. Imagine that you are riding along with the NMV inside the rotating system at the frequency associated with B_0 . You would then only observe the slow precession of the NMV from the z-axis onto the x-y-axis caused by B_1 [4].

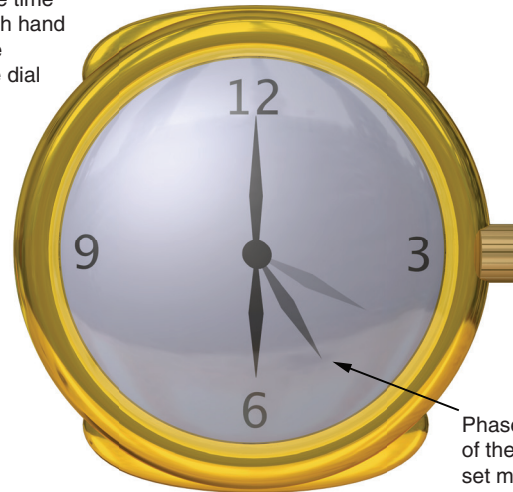


Refer to animation 1.1 on the supporting companion website for this book:
www.wiley.com/go/westbrook/mriinpractice

MR SIGNAL

Because of resonance, in-phase or coherent magnetization precesses in the transverse plane. This changing magnetic field generates an electric current. **Faraday's law** of induction explains this phenomenon. The change of magnetic flux through a closed circuit induces an **electromotive**

Frequency is the time taken for a watch hand to complete one revolution of the dial



Phase is the position of the hand at any set moment in time

Figure 1.11 Phase and frequency (the watch analogy).

force (emf) in the circuit. The emf is defined as the energy available from a unit of charge traveling once around a loop of wire. The emf drives a current in the circuit and is the result of a changing magnetic field inducing an electric field.

The laws of electromagnetic induction (Faraday) state that the induced emf:

- is proportional to the rate of change of magnetic field and the area of the circuit
- is proportional to the number of turns in a coil of wire (Equation (1.6))
- is in a direction so that it opposes the change in magnetic field that causes it (**Lenz's law**).

According to Faraday's law, a changing magnetic field causes movement of charged particles, i.e. electrons. This flow of electrons is a current, and if a receiver coil or any conductive loop is placed in a moving magnetic field, i.e. the magnetization precessing in the transverse plane, a voltage generated by this current is induced in the receiver coil. This voltage is called **signal** and is produced when coherent (in phase) magnetization cuts across the coil. The frequency of signal depends on the frequency of rotation of the magnetic field – the magnitude of signal depends on the amount of coherent magnetization present in the transverse plane (Figure 1.12).

Equation 1.6

$$\varepsilon = -N d\Phi / dt$$

ε is the emf in volts (V)

N is the number of turns in a coil

$d\Phi$ is changing magnetic flux in a single loop (Vs)

dt is changing time (s)

This equation shows that the amount of induced current in a coil is related to the rate of change of magnetic flux (how fast the magnetic lines of flux are crossed) and the number of turns in a coil




Figure 1.12 Generation of the signal.



Refer to animations 1.2 and 1.3 on the supporting companion website for this book: www.wiley.com/go/westbrook/mriinpractice

THE FREE INDUCTION DECAY (FID) SIGNAL

When the RF excitation pulse is switched off, the NMV is influenced only by B_0 , and it tries to realign with it. To do so, the hydrogen nuclei lose energy given to them by the RF excitation pulse. The process by which hydrogen loses this energy is called **relaxation**. As relaxation occurs, the NMV returns to realign with B_0 because some of the high-energy nuclei return to the low-energy population and therefore align their magnetic moments in the spin-up direction. At the same time, but independently, the magnetic moments of hydrogen lose coherency due to dephasing. This occurs because of inhomogeneities in the B_0 field and due to interactions between spins in the patient's tissue (see Chapter 2). As the magnitude of transverse coherent magnetization decreases, so does the magnitude of the voltage induced in the receiver coil. The induction of decaying voltage is called the **free induction decay (FID)** signal. This is because spins *freely* precess influenced only by B_0 , signal *decays* with time, and magnetic moments of the spins *induce* a current in the receiver coil.

The magnitude and timing of the RF pulses form part of **pulse sequences**, which are the basis of contrast generation in MRI (see Chapters 2–4).

Learning tip: Vectors

21

The NMV is a vector quantity. It is created by two components at 90° to each other. These two components are magnetization in the longitudinal plane and magnetization in the transverse plane (Figure 1.13). Before resonance, there is full longitudinal magnetization parallel to B_0 . After the application of the RF pulse and assuming a flip angle of 90° , the NMV is flipped fully into the transverse plane. There is now full transverse magnetization and zero longitudinal magnetization.

Once the RF excitation pulse is removed, the NMV recovers. As this occurs, the longitudinal component of magnetization grows again while the transverse component decreases (shown later in Figure 2.5). As the received signal amplitude is related to the magnitude of the coherent transverse component, signal in the coil decays as relaxation occurs.




Figure 1.13 Longitudinal and transverse magnetization.

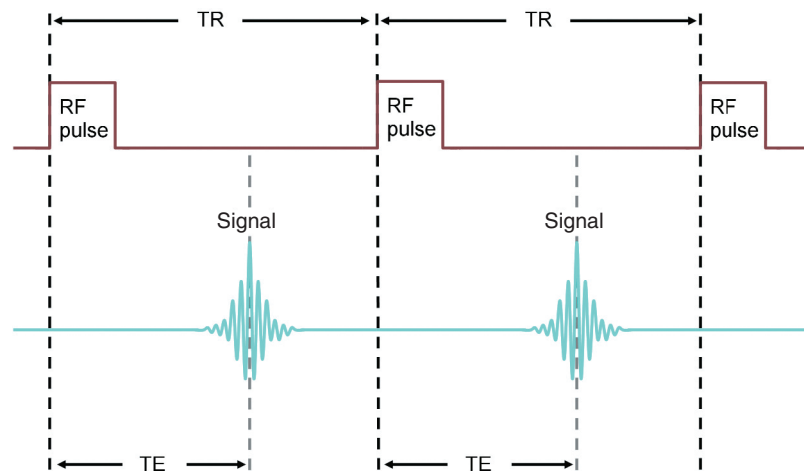


Figure 1.14 A basic pulse sequence.

Table 1.6 Things to remember – excitation and signal generation.

Application of RF energy at the Larmor frequency causes a net absorption of energy (excitation) and changes the balance between the number of spins in the low- and high-energy populations
The orientation of the NMV to B_0 depends on this balance. If there are a similar number of spins in each population, the NMV lies in a plane at 90° to B_0 (transverse plane)
Resonance also causes the magnetic moments of all spins to precess in phase. The result is coherent transverse magnetization that precesses in the transverse plane
If a receiver coil (conductor) is placed in the transverse plane, the movement of the rotating coherent transverse magnetization causes a voltage in the coil
When the RF excitation pulse is removed, the magnetic moments of all spins dephase and produce a FID

PULSE TIMING PARAMETERS

A very simplified pulse sequence is a combination of RF pulses, signals, and intervening periods of relaxation (Figure 1.14). It is important to note that a pulse sequence as shown diagrammatically in Figure 1.14 merely shows in simple terms the separate timing parameters used in more complicated sequences, i.e. repetition time (TR) and echo time (TE).

A pulse sequence consists of several time periods. The main ones are outlined below.

- The **TR** is the time from the application of one RF excitation pulse to the application of the next RF excitation pulse for each slice and is measured in millisecond. The TR determines the amount of longitudinal relaxation that occurs between the end of one RF excitation pulse and application of the next. The TR thus determines the amount of T1 relaxation that has occurred when signal is read (see Chapter 2).
- The **TE** is the time from the application of the RF excitation pulse to the peak of signal induced in the receiver coil and is also measured in millisecond. The TE determines how much decay of transverse magnetization occurs. TE thus controls the amount of T2 relaxation that has occurred when signal is read (see Chapter 2).

In this chapter, we explored and describe the basic principles behind signal creation. The application of RF pulses and the receiving of signals at predefined times produce contrast in MRI images. In the next chapter, we look at these concepts in detail.



For questions and answers on this topic, please visit the supporting companion website for this book: www.wiley.com/go/westbrook/mriinpractice

References

1. Cox, B. and Forshaw, J. (2012). *The Quantum Universe: Everything that Can Happen Does Happen*, 16. Penguin, London.
2. Odaibo, S.G. (2012). *Quantum Mechanics and the MRI Machine*, 5. Arlington, VA: Symmetry Seed Books.
3. McRobbie, D.W., Moore, E.A., Graves, M.J. et al. (2017). *From Picture to Proton*, 3, 127. Cambridge: Cambridge University Press.
4. Hashemi, R.H., Bradley Jr, W.G., and Lisanti, C.J. (2010). *MRI: The Basics*, 3, 24. Philadelphia, PA: Lippincott Williams and Wilkins.
5. Elmaoglu, M. and Celik, A. (2012). *MRI Handbook, MR Physics, Patient Positioning and Protocols*, 9. New York: Springer.
6. McRobbie, D.W., Moore, E.A., Graves, M.J. et al. (2017). *From Picture to Proton*, 129. Cambridge: Cambridge University Press.
7. Hashemi, R.H., Bradley Jr, W.G., and Lisanti, C.J. (2010). *MRI: The Basics*, 3, 24. Philadelphia, PA: Lippincott Williams and Wilkins.
8. Dale, B.M., Brown, M.A., and Semelka, R.C. (2015). *MRI Basic Principles and Applications*, 5, 11. Wiley.
9. Hashemi, R.H., Bradley Jr, W.G., and Lisanti, C.J. (2010). *MRI: The Basics*, 3, 37. Philadelphia, PA: Lippincott Williams and Wilkins.

2

Image weighting and contrast

Introduction	24	Relaxation in different tissues	32
Image contrast	25	T1 contrast	36
Relaxation	25	T2 contrast	40
T1 recovery	26	Proton density contrast	41
T2 decay	27	Weighting	42
Contrast mechanisms	31	Other contrast mechanisms	51

After reading this chapter, you will be able to:

- *Differentiate between intrinsic and extrinsic contrast parameters.*
- *Explain T1 recovery and T2 decay.*
- *Understand how contrast is generated in different tissues.*
- *Apply what you have learned to create images of different weighting.*
- *Describe techniques that affect image contrast.*

INTRODUCTION

All clinical diagnostic images must demonstrate contrast between normal anatomical features and between anatomy and pathology. If there is no contrast, it is impossible to identify anatomy or detect abnormalities within the body. One of the main advantages of MRI compared to other imaging modalities is its excellent soft tissue discrimination. The contrast characteristics of each image depend on many variables, and it is important that the mechanisms that affect image contrast in MRI are clearly understood.

IMAGE CONTRAST

The factors that affect image contrast in diagnostic imaging are usually divided into two categories.

- **Intrinsic contrast parameters** are those that *cannot* be changed because they are inherent to the body's tissues.
- **Extrinsic contrast parameters** are those that *can* be changed because they are under our control.

For example, in radiography, intrinsic contrast parameters include the density of structures through which the X-ray beam passes and is attenuated by, while extrinsic contrast parameters include the exposure factors. These determine image contrast in a radiograph, but exposure factors can be changed, whereas tissue density cannot. In MRI, there are several parameters in each group.

Intrinsic contrast parameters are as follows:

- T1 recovery time
- T2 decay time
- Proton density (PD)
- Flow
- Apparent diffusion coefficient (ADC).

All these are inherent to the body's tissues and cannot be changed. T1 recovery time, T2 decay time, proton density, and ADC are discussed in this chapter. Flow is discussed in Chapter 8.

Extrinsic contrast parameters are as follows:

- TR
- TE
- Flip angle
- TI
- Turbo factor/echo train length
- b value.

These are all selected in the scan protocol. Some of these parameters depend on the **pulse sequence** we choose (see Table 4.18). TR, TE, and the b value are discussed later in this chapter, and the other parameters are described in Chapters 3 and 4 in the relevant pulse sequence sections. A list of acronyms of the five main system manufacturers is provided at the beginning of the book. This includes the contrast parameters and pulse sequences described in these chapters.

RELAXATION

At the end of Chapter 1, we explored the consequences of switching off the radio frequency (RF) excitation pulse. To recap, as soon as the B_1 field is removed, hydrogen nuclei are only under the influence of B_0 . One of the principles of thermodynamics is that a system always seeks its lowest possible energy level. This occurs in MRI when the RF excitation pulse is switched off, and therefore hydrogen nuclei return to their low-energy state and their magnetic moments dephase [1]. The process by which this occurs is called **relaxation**.

During relaxation, hydrogen nuclei give up absorbed RF energy, and the net magnetic vector (NMV) returns to B_0 . At the same time but independently, magnetic moments of hydrogen

Learning tip:**T1 recovery vs T2 decay**

T1 recovery and T2 decay occur at two different rates. T2 decay occurs 5–10 times faster than T1 recovery. An important learning step is to understand the timing of different components of a pulse sequence (see Chapter 3).

nuclei lose phase coherence. Relaxation therefore results in recovery of magnetization in the longitudinal plane and decay of coherent magnetization in the transverse plane.

- The recovery of longitudinal magnetization is caused by a process termed **T1 recovery**.
- The decay of coherent transverse magnetization is caused by a process termed **T2 decay**.

T1 RECOVERY

T1 recovery is caused by hydrogen nuclei giving up their energy to the surrounding environment or molecular lattice. The term *recovery* refers to the recovery of longitudinal magnetization, and *T1* relates to the fact that it is the primary relaxation process. (It is not, however, the first process that occurs. The learning tip above shows that T1 recovery takes 5–10 times longer than T2 decay.) This type of relaxation is called **spin–lattice energy transfer**. Energy released by spins to the surrounding molecular lattice causes magnetic moments of hydrogen nuclei to recover their longitudinal magnetization. According to quantum theory, the number of high-energy spins decreases, and the number of low-energy spins increases as energy is lost by high-energy spins during the relaxation process. According to classical theory, the NMV gradually realigns itself in the longitudinal plane as the proportion of spin-up and spin-down hydrogen nuclei changes.

The rate of T1 recovery is an exponential process and occurs at different rates in different tissues (Table 2.1). As illustrated in Figure 2.1, longitudinal magnetization is related exponentially to recovery time. This means that most longitudinal recovery happens at the beginning of the time frame. As time progresses, gradually less and less longitudinal recovery occurs until the longitudinal magnetization is fully recovered. There is a time constant associated with this exponential relationship. This is called the **T1 recovery time** and is the time it takes for 63% of the longitudinal magnetization to recover in a tissue (Equation (2.1)) (Figure 2.1). The T1 recovery time of a tissue is an intrinsic contrast parameter that is inherent to the tissue. The time during which T1 recovery occurs is the time between one RF excitation pulse and the next. This is the repetition time (TR) (see Chapter 1). The TR therefore determines how much T1 recovery occurs in a tissue. It is therefore the variable shown on the horizontal axis of Figure 2.1.

Table 2.1 Typical T1 recovery times of brain tissue at 1 T.

Tissue	T1 recovery time (ms)
Water	2500
Fat	200
CSF	2000
White matter	500

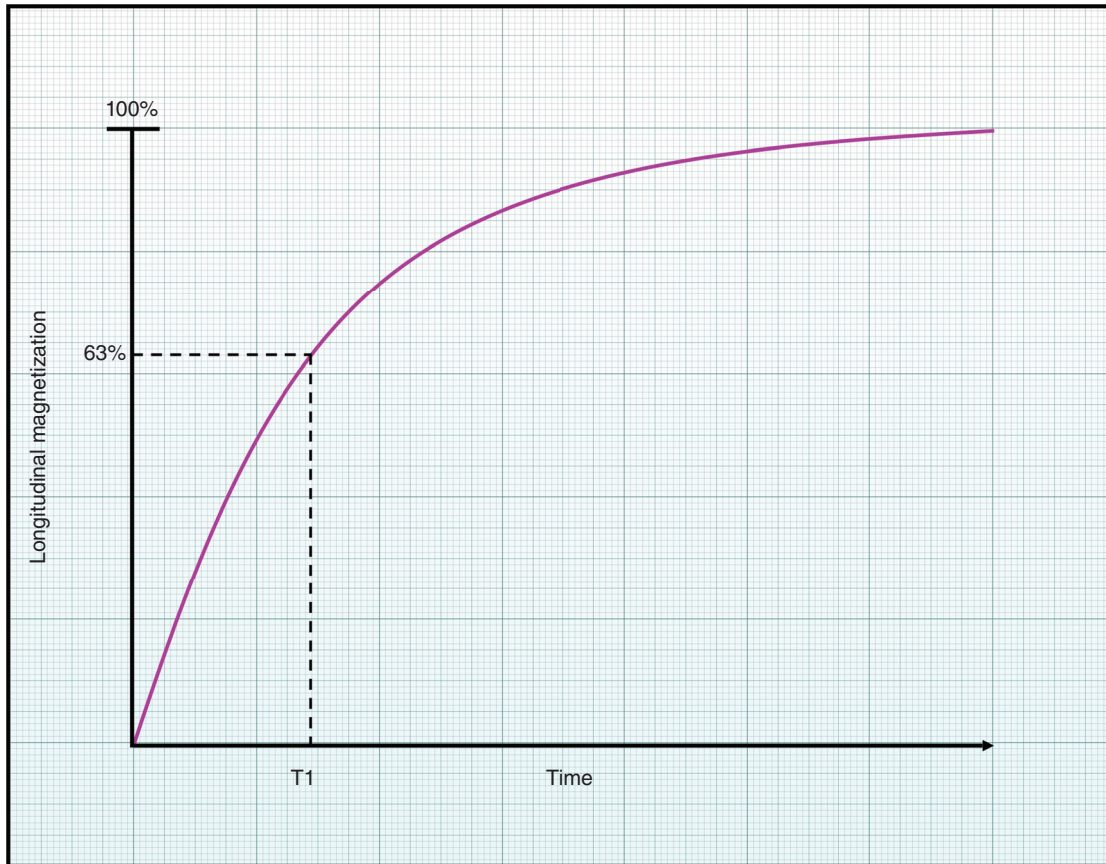


Figure 2.1 The T1 recovery curve.

Equation 2.1

$$Mz_t = Mz (1 - e^{-t/T1})$$

Therefore

$$SI = (1 - e^{-t/T1})$$

Mz_t is the amount of longitudinal magnetization at time t (ms) after the removal of the excitation pulse

Mz is full longitudinal magnetization

$T1$ is the T1 recovery time (ms) and is the time taken to increase the longitudinal magnetization by a factor of e .

SI is the signal intensity in a tissue

This equation plots the size of the recovering NMV as a function of time after the removal of the excitation pulse and the T1 recovery time. When $t = T1$, 63% of the longitudinal magnetization recovers. When $t = 2 \times T1$, 86% recovers and when $t = 3 \times T1$, 95% recovers. It usually takes between 3 and 5 T1 recovery times for full recovery to occur

T2 DECAY

T2 decay is caused by the magnetic fields of neighboring hydrogen nuclei interacting with each other. The term *decay* refers to the loss of coherent transverse magnetization, and T2 relates to the fact that it is the secondary relaxation process. This type of relaxation is termed **spin-spin relaxation** and causes dephasing of magnetic moments of the spins. Spin-spin relaxation is caused by one spin transferring energy to another spin rather than into the lattice. It occurs because

hydrogen nuclei are in the same environment and experiencing the same B_0 field [2]. Magnetic moments of all the hydrogen nuclei (spin-up and spin-down) lose phase coherence in this way.

Imagine two spins close to each other, one with its magnetic moment aligned in the same direction as B_0 and the other in the opposite direction. The spin whose magnetic moment is aligned in the same direction as B_0 creates a slightly larger magnetic field than is experienced by the neighboring spin. As a result, the precessional frequency of the magnetic moment of this spin increases. Conversely, the spin whose magnetic moment is aligned in the opposite direction to B_0 causes a slightly lower magnetic field than is experienced by the other spin, and its precessional frequency decreases [3]. These small changes in frequency are sufficient to cause dephasing of magnetic moments of the spins.

Spin-spin interaction is inherent to the tissue, but dephasing is also caused by inhomogeneities in the B_0 field. Inhomogeneities are areas within the magnetic field that do not exactly match the external magnetic field strength. Some areas have a magnetic field strength slightly less than the main magnetic field (shown in purple in Figure 2.2), while other areas have a magnetic field strength slightly higher than the main magnetic field (shown in red in Figure 2.2).

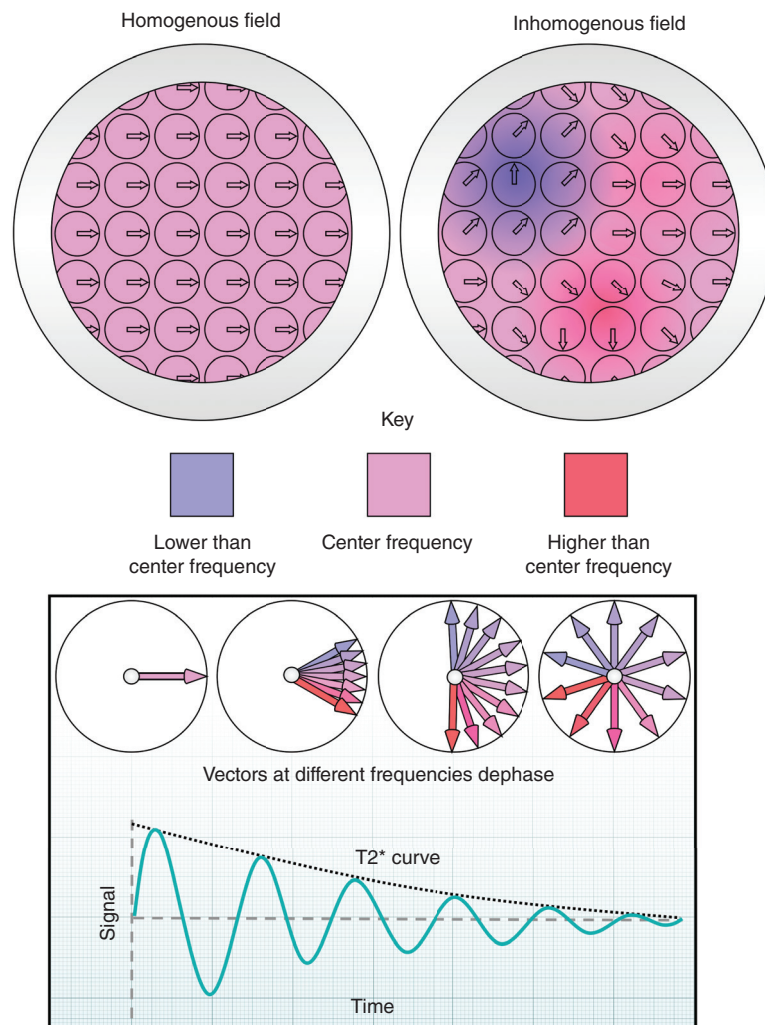


Figure 2.2 T2* decay and field inhomogeneities.

According to the Larmor equation, the Larmor frequency of an MR-active nucleus is proportional to the magnetic field strength it experiences. If a hydrogen nucleus lies in an area of inhomogeneity with higher field strength, the precessional frequency of its magnetic moment increases, i.e. it speeds up. However, if a hydrogen nucleus lies in an area of inhomogeneity with lower field strength, the precessional frequency of its magnetic moment decreases, i.e. it slows down. This relative acceleration and deceleration of magnetic moments due to magnetic field inhomogeneities, and differences in the precessional frequency in certain tissues, causes immediate dephasing of the magnetic moments of the spins and produces a free induction decay (FID) as shown in Figures 2.2 and 2.3.

Learning tip: Inhomogeneities



Do you remember the watch analogy in Chapter 1? The change of phase of magnetic moments caused by inhomogeneities in the field is the same as several watches telling different times because the frequencies of their hands are different.

The rate of T2 decay is an exponential process and occurs at different rates in different tissues (Table 2.2). As Figure 2.4 illustrates, coherent transverse magnetization is related exponentially to decay time. This means that there is more coherent transverse magnetization at the beginning of the time-frame and, as time progresses, there is less coherent transverse magnetization until all the magnetic moments dephase. There is a time constant associated with this exponential relationship. It is called the **T2 decay time** and is the time it takes for 63% of the transverse magnetization to dephase (37% is left in phase) in a tissue (Equation (2.2)) (Figure 2.4).

The T2 decay time of a tissue is an intrinsic contrast parameter that is inherent to the tissue. The time during which this occurs is the time between an RF excitation pulse and when signal is collected in the receiver coil (see Chapter 1). The echo time (TE) therefore determines how





Figure 2.3 Dephasing and free induction decay.

Table 2.2 Typical T2 decay times of brain tissue at 1T.

Tissue	T2 decay time (ms)
Water	2500
Fat	100
CSF	300
White matter	100

**Figure 2.4** The T2 decay curve.

much T2 decay occurs in a tissue when signal is collected. It is therefore the variable shown on the horizontal axis of Figure 2.4.

Dephasing caused by inhomogeneities in the B_0 field produces its own decay curve. This is differentiated from T2 decay by using the term $T2^*$. When the RF excitation pulse is switched off, magnetic moments of the hydrogen nuclei dephase very quickly (within about 10 ms), and this is caused by T2* decay. T2 decay, from inherent tissue dephasing, takes longer than T2* decay (Equation (2.3)). The purpose of pulse sequences is to refocus or rephase magnetic moments of the hydrogen nuclei so that inherent tissue dephasing is measured at time TE and images of different contrast can be acquired (see Chapter 3).

Equation 2.3

$$1/T2^* = 1/T2 + 1/2\gamma\Delta\beta_0$$

T2 and T2* are the tissues T2 and T2* relaxation times (ms)
 γ is the gyromagnetic ratio (MHz/T)
 ΔB_0 is the variation in magnetic field (parts per million, ppm)

This equation shows how T2 and T2* are related to each other. Poor field inhomogeneities result in T2* being much shorter than T2, and a fast decaying signal

Table 2.3 Things to remember – relaxation.

Relaxation is a general term that refers to a loss of energy. In MRI, this is energy that is delivered to the spins via the RF excitation pulse and then lost once it is switched off

Spin lattice energy transfer is a relaxation process where spins give up the energy absorbed through RF excitation to the surrounding molecular lattice of the tissue. It causes T1 recovery

T2 decay is an irreversible loss of phase coherence due to spin–spin interactions on an atomic and molecular level. It is one of the causes of T2 decay

Pulse sequences are mechanisms that permit refocusing of spins so that images can be acquired with different types of contrast

CONTRAST MECHANISMS

An MR image has contrast if there are areas of high signal (hyperintensity – white in the image) and areas of low signal (hypointensity – black in the image). Some areas have an intermediate signal (shades of gray in between white and black). The NMV is separated into individual vectors of the tissues such as fat, cerebrospinal fluid (CSF), and muscle.

A tissue has a high signal if it has a large transverse component of coherent magnetization at time TE. If there is a large component of coherent transverse magnetization, the amplitude of signal received by the coil is large, resulting in a hyperintense area on the image. A tissue returns a low signal if it has a small or no transverse component of coherent magnetization at time TE. If there is a small or no component of transverse coherent magnetization, the amplitude of signal received by the coil is small, resulting in a hypointense area on the image.

Images obtain contrast mainly through the mechanisms of T1 recovery, T2 decay, and proton or spin density. The **proton density (PD)** of a tissue is the number of mobile hydrogen protons per unit volume of that tissue. The higher the proton density of a tissue, the more signal available from that tissue. T1 and T2 relaxation depend on two factors:

- *If the molecular tumbling rate matches the Larmor frequency of hydrogen.* If there is a good match between the rate of molecular tumbling and the Larmor frequency of magnetic moments of hydrogen, energy exchange between hydrogen nuclei and the molecular lattice is efficient. When there is a bad match, energy exchange is not as efficient. This is important in both T1 recovery and T2 decay processes.
- *If the molecules are closely packed together.* In tissues where molecules are closely spaced, there is more efficient interaction between the magnetic fields of neighboring hydrogen nuclei. The reverse is true when molecules are spaced apart. This is especially important in

T2 decay processes, which rely on the efficiency of interactions between the magnetic fields of neighboring hydrogen nuclei (spin–spin relaxation).

RELAXATION IN DIFFERENT TISSUES

As discussed earlier, T1 recovery and T2 decay are exponential processes with time constants T1 recovery time and T2 decay time, which represent the time it takes for 63% of the total magnetization to recover in the longitudinal plane due to spin–lattice energy transfer (T1 recovery time), or lost in the transverse plane via spin–spin relaxation (T2 decay time). Generally, the two extremes of contrast in MRI are fat and water (Figure 2.5). (In this book, fat vectors are drawn in yellow and water vectors in blue.) Let's explore how contrast is generated in these tissues.

Fat and water

Fat molecules contain atoms of hydrogen arranged with carbon and oxygen. They consist of large molecules called lipids that are closely packed together and whose molecular motion or tumbling rate is relatively slow. Water molecules contain two hydrogen atoms arranged with one oxygen atom (H_2O). Its molecules are spaced apart, and their molecular tumbling rate is relatively fast. The oxygen in water tends to steal the electrons away from around the hydrogen nucleus. This renders it more available to the effects of the main magnetic field. In fat, the carbon does not




Figure 2.5 The magnitude of transverse magnetization vs amplitude of signal.




Figure 2.6 T1 recovery in fat.

take the electrons from around the hydrogen nucleus. They remain in an electron cloud, protecting the nucleus from the effects of the main field. Therefore, hydrogen in fat recovers more rapidly along the longitudinal axis than water and loses transverse magnetization faster than water. Subsequently, fat and water appear differently in MR images.

T1 recovery in fat

T1 recovery occurs due to hydrogen nuclei giving up their energy to the surrounding molecular lattice. Fat has a low inherent energy and easily absorbs energy into its lattice from hydrogen nuclei. The slow molecular tumbling in fat allows the T1 recovery process to be relatively rapid because the molecular tumbling rate matches the Larmor frequency. Consequently, there is efficient energy exchange from hydrogen nuclei to the surrounding molecular lattice. This means that magnetic moments of fat hydrogen nuclei quickly relax and regain their longitudinal magnetization. The NMV of fat realigns rapidly with B_0 , so the T1 recovery time of fat is short (Figure 2.6).

T1 recovery in water

T1 recovery occurs due to hydrogen nuclei giving up energy acquired from the RF excitation pulse to the surrounding lattice. Water has a high inherent energy and does not easily absorb energy into its lattice from hydrogen nuclei. In water, molecular mobility is high, resulting in less efficient T1 recovery because the molecular tumbling rate does not match the Larmor frequency and does not allow efficient energy exchange from hydrogen nuclei to the surrounding molecular lattice. Magnetic moments of water hydrogen nuclei take longer to relax and regain their longitudinal magnetization. The NMV of water takes longer to realign with B_0 , and so the T1 recovery time of water is long (Figure 2.7).

Learning tip:**T1 recovery and B_0**

34

T1 recovery is affected by the strength of B_0 . The precessional frequency of magnetic moments of the hydrogen nuclei within a tissue varies slightly, but efficient energy exchange, due to molecular motion, is optimal at the Larmor frequency. The Larmor frequency is proportional to B_0 , and therefore T1 recovery takes longer as B_0 increases because there are fewer molecules moving at relaxation-causing frequencies [4].




Figure 2.7 T1 recovery in water.

T2 decay in fat

T2 decay occurs because the magnetic fields of hydrogen nuclei interact with each other. This process is efficient in hydrogen in fat, as the molecules are packed closely together, and therefore spin–spin interactions are more likely to occur. It also occurs because magnetic moments of hydrogen nuclei in fat precess at a similar frequency to molecular tumbling. As a result, magnetic moments dephase quickly, and there is a rapid loss of coherent transverse magnetization. The T2 decay time of fat is therefore short (Figure 2.8).

T2 decay in water

T2 decay in water is less efficient than in fat, as the molecules are spaced apart, and spin–spin interactions are less likely to occur. In addition, magnetic moments of hydrogen nuclei in water precess much faster than molecular tumbling. As a result, magnetic moments of hydrogen nuclei dephase slowly, and there is a gradual, rather than rapid, loss of coherent transverse magnetization. The T2 decay time of water is therefore long (Figure 2.9).

Learning tip:**T2 decay and B_0**

T2 decay is affected only to a small degree by the strength of B_0 . Spin-spin relaxation is optimal when molecular motion occurs at the Larmor frequency. The Larmor frequency is proportional to B_0 and therefore T2 decay takes marginally longer as B_0 increases because there are fewer molecules moving at relaxation-causing frequencies [1].

35




Figure 2.8 T2 decay in fat.




Figure 2.9 T2 decay in water.

T1 CONTRAST

36

The term **T1 contrast** means that image contrast is derived from differences in the T1 recovery times of the tissues rather than any other mechanism. T1 contrast is likely to occur if vectors do not fully recover their longitudinal magnetization between each RF excitation pulse. It therefore increases if the TR is short. If the TR is longer than the relaxation times of the tissues, full recovery occurs in all tissues, and, therefore, it is not possible to produce an image that demonstrates contrast based on the differences in their T1 recovery times.

The T1 recovery time of fat is much shorter than that of water, so the fat vector realigns with B_0 faster than the water vector. If the TR is shorter than the total recovery times of the tissues, the longitudinal component of magnetization of fat is larger than that of water. When the next RF excitation pulse is applied, it flips the longitudinal components of magnetization of both fat and water into the transverse plane (assuming a 90° RF excitation pulse is applied) as in Figure 2.10. As there is more longitudinal magnetization in fat before the RF excitation pulse, there is more transverse magnetization in fat after the RF excitation pulse. Fat therefore has a high signal and is hyperintense. As there is less longitudinal magnetization in water before the RF excitation pulse, there is less transverse magnetization in water after the RF excitation pulse. Water therefore has a low signal and appears relatively hypointense.

Learning tip:

Saturation

Whenever the NMV is pushed beyond 90° , it is said to be **partially saturated**. When the NMV is pushed to a full 180° , it is said to be **fully saturated**. If partial saturation of the fat and water vectors occurs, T1 contrast is maximized.

Look at Figure 2.11. Before application of the first RF excitation pulse, the fat and water vectors are aligned with B_0 . When the first 90° RF excitation pulse is applied, the fat and water vectors are flipped into the transverse plane. The RF excitation pulse is then removed, and the vectors begin to relax and return to B_0 . Fat has a shorter T1 recovery time than water and so returns to B_0 faster than water. If the TR is shorter than the T1 recovery times of the tissues, the next RF excitation pulse flips the vectors beyond 90° and into partial saturation because their recovery was incomplete. The fat and water vectors are saturated to different degrees because they were at different points of recovery before the 90° RF excitation pulse was applied. The transverse component of magnetization for each vector is therefore different. The transverse component of fat is greater than that of water because its longitudinal component recovers to a greater degree before the next RF pulse is applied, and so more longitudinal magnetization is available to be flipped into the transverse plane. The fat vector therefore generates a higher signal than does water. Fat is hyperintense, and water is relatively hypointense.

Now look at Figure 2.12. If the TR is longer than the T1 recovery times of the tissues, both fat and water fully recover before the next RF excitation pulse is applied. Both vectors are flipped directly into the transverse plane and are not saturated. The magnitude of the transverse component of magnetization for fat and water depends only on their individual proton densities rather than the rate of recovery of their longitudinal components. The flip angle has a significant impact on saturation effects. This is discussed in more detail in Chapter 4.

Learning tip: Achieving the steady-state

It is clear from the previous learning tip that if the TR is shorter than the T1 relaxation times of the tissues, the first few RF excitation pulses result in different amounts of transverse magnetization. This is because they have recovered different amounts of longitudinal magnetization before the RF excitation pulse was applied. It takes a few TR periods for things to settle down into what is called the **steady-state**. When the steady-state is achieved, vectors recover to the same point and achieve the same amount of longitudinal magnetization during the TR period and they are always flipped to the same point by the 90° RF excitation pulse. They therefore create the same amount of transverse magnetization every TR.

Once the steady-state is achieved, signals (or echoes as they are usually called) are detected by the receiver coil. Before then, signals are not detected because they are different every TR. This is because the amount of transverse magnetization that is created is different. The first few RF excitation pulses are known as **preparatory** or **dummy pulses** because the signals they produce are ignored. Once the longitudinal and transverse vectors of magnetization have settled down and are steady, then these signals are detected and used to create the image. The time it takes to achieve the steady-state depends on B_0 , proton density, flip angle, T1 relaxation time, and the duration of the RF excitation pulse [5]. The steady-state is discussed further in Chapter 4 in relation to gradient-echo sequences.




Figure 2.10 T1 contrast generation.

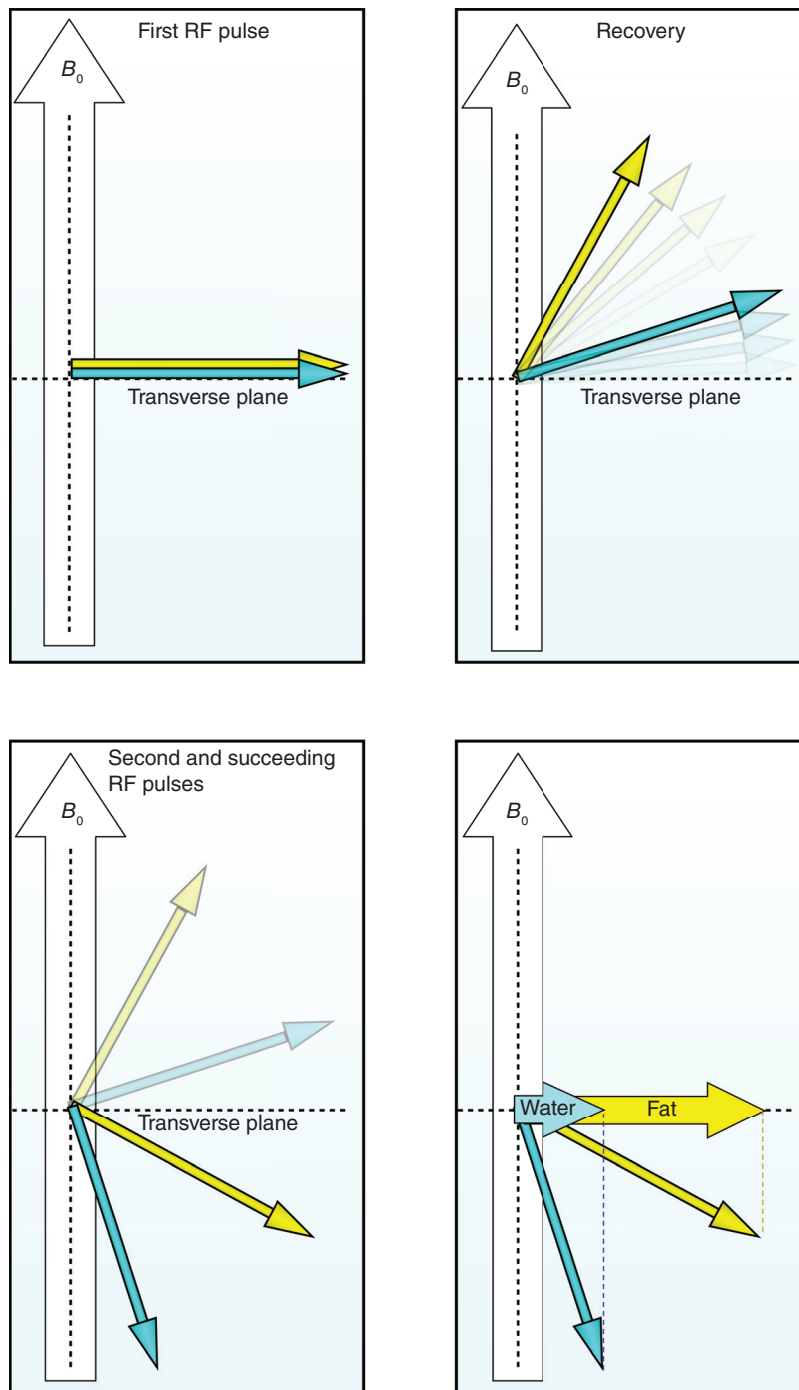


Figure 2.11 Saturation with a short TR.




Figure 2.12 No Saturation with a long TR.

Table 2.4 Things to remember – T1 recovery.

Fat has a short T1 recovery time
Water has a long T1 recovery time
T1 recovery is caused by spin–lattice energy transfer. The efficiency of this process depends on the inherent energy of the tissue and how well the rate of molecular tumbling matches the Larmor frequency
T1 recovery times are dependent upon magnetic field strength. As field strength increases, tissues take longer to relax
T1 contrast is controlled by the TR. For good T1 contrast, the TR must be short

T2 CONTRAST

The term **T2 contrast** means that image contrast is derived from differences in the T2 decay times of the tissues rather than any other mechanism. T2 contrast is likely to occur if vectors dephase and there is a difference in coherent transverse magnetization in each tissue. It therefore increases if the TE is long. Magnetic moments of the hydrogen nuclei dephase at different rates, so if the TE is long, it is possible to produce an image that demonstrates differences in their T2 decay times. If the TE is short, then little dephasing occurs, and therefore it is not possible to produce images that demonstrate differences in T2 decay times of the tissues.

The T2 decay time of fat is shorter than that of water so the transverse component of magnetization in fat decays faster than the transverse component in water. There is more coherent transverse magnetization in water than in fat. Water, therefore, has a high signal and is hyperintense, and fat therefore has low signal and is relatively hypointense on a T2 contrast image (Figure 2.13).

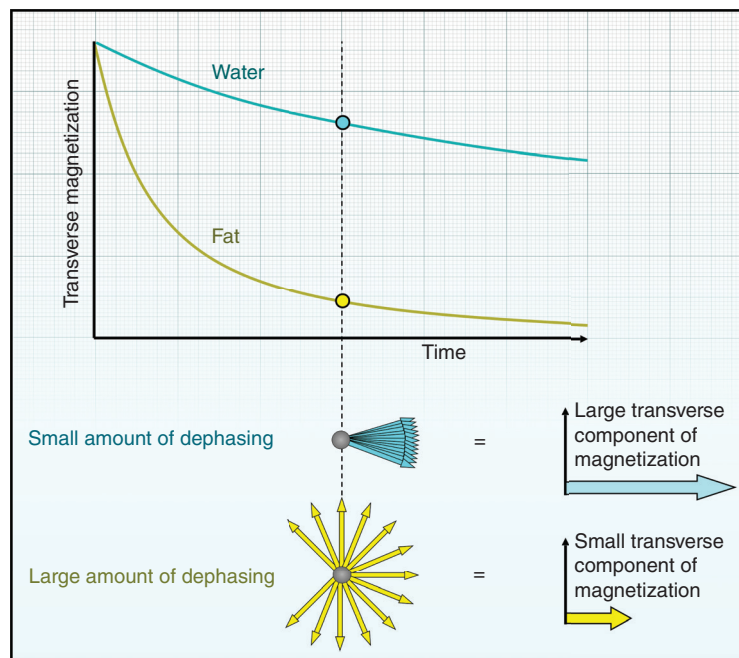
**Figure 2.13** T2 contrast generation.

Table 2.5 Things to remember – T2 decay.

Fat has a short T2 decay time
Water has a long T2 decay time
T2 decay is caused by spin–spin relaxation. The efficiency of this process depends on how closely the molecules are packed together
T2 decay times are dependent upon magnetic field strength. As field strength increases, tissues take longer to dephase
T2 contrast is controlled by the TE. For good T2 contrast, the TE must be long

PROTON DENSITY CONTRAST

Proton density contrast refers to differences in signal intensity between tissues that are a consequence of their relative number of mobile hydrogen protons per unit volume. To produce contrast due to differences in the proton densities between the tissues, the transverse component of magnetization must reflect these differences. Tissues with a high proton density have a large transverse component of magnetization (and therefore a high signal) and are hyperintense. Tissues with a low proton density have a small transverse component of magnetization (and therefore a low signal) and are relatively hypointense (see Figure 2.18). Proton density contrast is always present and depends on the patient and the area under examination. The signal intensity of a tissue therefore depends on its intrinsic contrast properties. Three of these are:

- proton density (PD)
- T1 recovery time
- T2 decay time.

Signal intensity also depends on extrinsic contrast parameters. Two of these are:

- TR
- TE (Equation (2.4)).

Now let's explore how image contrast is created by altering the TR and TE.

Equation 2.4

$$SI = PD e^{-TE/T2} (1 - e^{-TR/T1})$$

SI is the signal intensity in a tissue
 PD is proton density
 TE is the echo time (ms)
 T2 is the T2 relaxation time of the tissue (ms)
 TR is the repetition time (ms)
 T1 is the T1 relaxation time in the tissue (ms)

This equation shows why the signal intensity from a tissue depends on intrinsic and extrinsic contrast parameters. Equation (4.1) shows how this equation is modified in gradient-echo pulse sequences

WEIGHTING

42

All the intrinsic contrast parameters listed at the beginning of this chapter simultaneously affect image contrast, and therefore, it is possible to obtain images of mixed appearance. This means that when looking at an image it is very difficult to determine the relative contribution of each intrinsic contrast parameter to the contrast observed. To minimize this, extrinsic contrast parameters are selected to *weight* image contrast toward one of the intrinsic contrast parameters and away from the others. This is achieved by applying our knowledge of how extrinsic contrast parameters control the relative contribution from each intrinsic contrast parameter. To demonstrate T1, T2, or proton density weighting, specific values of TR and TE are selected. The appropriate selection of these parameters weights an image so that one contrast mechanism *dominates* the other two.

T1 weighting

A **T1-weighted image** is one where contrast depends predominantly on the differences in the T1 recovery times between fat and water (and all the tissues with intermediate T1 recovery times). The TR controls how far each vector recovers before the slice is excited by the next RF excitation pulse. To achieve T1 weighting, the TR must be short enough so that neither the vector in fat nor the vector in water has sufficient time to fully return to B_0 . If the TR is too long, both the vectors in fat and water return to B_0 and fully recover their longitudinal magnetization. When this occurs, T1 recovery is complete in both tissues, and the differences in their T1 recovery times are not demonstrated (Figure 2.14). T1-weighted images are used to show anatomy (Figure 2.15) and pathology after administration of a contrast agent (Table 2.6).

- TR controls the amount of T1 contrast.
- For T1 weighting, the TR must be short and the TE must also be short (Equation (2.5)).

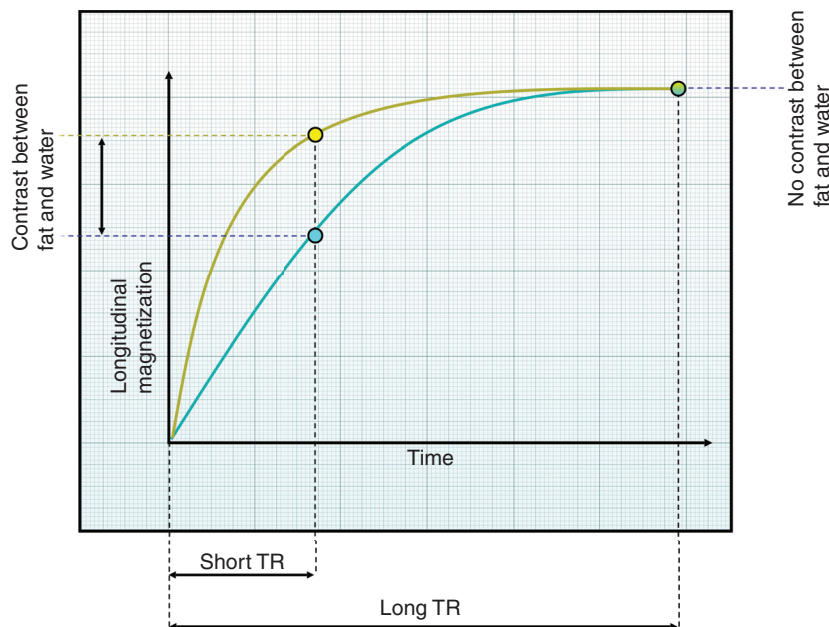


Figure 2.14 The difference in T1 recovery between fat and water.



Figure 2.15 Coronal T1-weighted image of the knee.

Table 2.6 T1 contrast examples.

High signal	<ul style="list-style-type: none"> Fat Hemangioma Intraosseous lipoma Radiation change Degeneration Fatty deposition Methemoglobin Cysts with proteinaceous fluid Paramagnetic contrast agents Slow-flowing blood
Low signal	<ul style="list-style-type: none"> Cortical bone Avascular necrosis Infarction Infection Tumors Sclerosis Cysts Calcification
No signal	<ul style="list-style-type: none"> Air Fast-flowing blood Tendons Cortical bone Scar tissue Calcification

Equation 2.5

$$SI = PD e^{-TE/T2} (1 - e^{-TR/T1})$$

SI is the signal intensity in a tissue
Referring to Equation (2.4):
 $e^{-TE/T2}$ is the T2 component
 $(1 - e^{-TR/T1})$ is the T1 component

This equation shows that if the TE is infinitely short then $e^{-TE/T2} = 1$. Therefore, T2 contrast is minimized and signal intensity depends mainly on PD and T1 contrast

Table 2.7 Things to remember – T1 weighting.

TR controls T1 contrast. TE controls T2 contrast

To produce a T1-weighted image, it is necessary to create contrast in which the differences in the T1 recovery times of the tissues dominate image contrast

A short TR combined with a short TE maximizes T1 and minimizes T2 contrast, respectively

T1-weighted images are used for anatomy and pathology postcontrast enhancement

T2 weighting

A **T2-weighted image** is one where contrast predominantly depends on the differences in the T2 decay times between fat and water (and all the tissues with intermediate T2 decay times). The TE controls the amount of T2 decay that occurs before signal is received. To achieve T2 weighting, the TE must be long enough to give the vectors in both fat and water time to dephase. If the TE is too short, neither the vector in fat nor the vector in water has had time to dephase, and, therefore, the differences in their T2 decay times are not demonstrated (Figure 2.16). T2-weighted images are used to image pathology because most pathology has a high water content and is therefore relatively hyperintense on T2-weighted images (Figure 2.17 and Table 2.8).

- TE controls the amount of T2 contrast.
- For T2 weighting, the TE must be long and the TR must also be long (Equation (2.6)).

Equation 2.6

$$SI = PD e^{-TE/T2} (1 - e^{-TR/T1})$$

SI is the signal intensity in a tissue
Referring to Equation (2.4):
 $e^{-TE/T2}$ is the T2 component
 $(1 - e^{-TR/T1})$ is the T1 component

This equation shows that if the TR is infinitely long, then $(1 - e^{-TR/T1}) = 1$. Therefore, T1 contrast is minimized, and signal intensity mainly depends on PD and T2 contrast

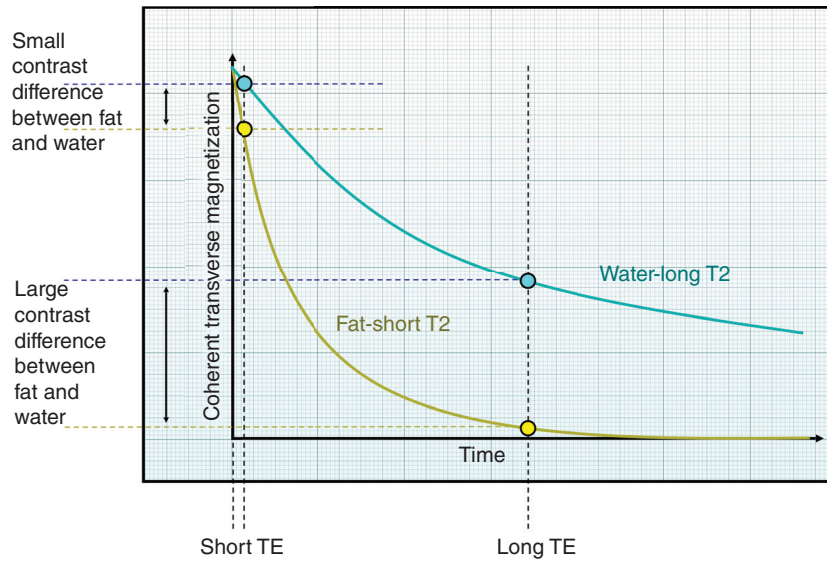


Figure 2.16 The difference in T2 decay between fat and water.




Figure 2.17 Axial T2-weighted image of the wrist.

Table 2.8 T2 contrast examples.

High signal	Water Synovial fluid Hemangioma Infection Inflammation Edema Some tumors Hemorrhage Slow-flowing blood Cysts
Low signal	Cortical bone Bone islands Deoxyhemoglobin Hemosiderin Calcification T2 paramagnetic agents
No signal	Air Fast-flowing blood Tendons Cortical bone Scar tissue Calcification

Table 2.9 Things to remember – T2 weighting.

TR controls T1 contrast. TE controls T2 contrast
To produce a T2-weighted image, it is necessary to create contrast in which the differences in the T2 decay times of the tissues dominate image contrast
A long TR combined with a long TE minimizes T1 and maximizes T2 contrast, respectively
T2-weighted images are used for pathology



2 Refer to animation 2.1 on the supporting companion website for this book: www.wiley.com/go/westbrook/mriinpractice

Proton density weighting

A **PD-weighted** image is one where differences in the number of mobile hydrogen nuclei per unit volume of tissue are the main determining factor in forming image contrast Table (2.10). PD weighting is always present to some extent. To achieve PD weighting, the effects of T1 and T2 contrast are diminished so that proton density contrast dominates. A long TR allows the vectors in both fat and water to fully recover their longitudinal magnetization and so diminishes T1 contrast. A short TE does not give the vectors in fat or water time to dephase and so diminishes T2 contrast (Equation (2.7)). PD-weighted images are used to image anatomy and pathology Figure (2.18).

Equation 2.7

$$SI = PD e^{-TE/T2} (1 - e^{-TR/T1})$$

SI is the signal intensity in a tissue
Referring to Equation (2.4):
 $e^{-TE/T2}$ is the T2 component
 $(1 - e^{-TR/T1})$ is the T1 component

This equation shows that if the TR is infinitely long, then $(1 - e^{-TR/T1}) = 1$ and if the TE is infinitely short then $e^{-TE/T2} = 1$. Therefore, T1 and T2 contrast are minimized and signal intensity mainly depends on PD

47

Table 2.10 Proton density contrast examples.

High signal	CSF Synovial fluid Slow-flowing blood Infection Inflammation Edema Cysts Fat
Low or no signal	Air Fast-flowing blood Tendons Cortical bone Scar tissue Calcification




Figure 2.18 Sagittal proton density weighted image of the ankle.

Table 2.11 Things to remember – proton density weighting.

TR controls T1 contrast; TE controls T2 contrast

To produce a PD-weighted image, it is necessary to create contrast in which the differences in the proton densities of the tissues dominate image contrast

A long TR combined with a short TE minimizes T1 and T2 contrast, respectively, so that proton density can dominate
PD-weighted images are used for anatomy and pathology

In any image, contrast due to the inherent proton density together with T1 and T2 mechanisms occurs simultaneously and contributes to image contrast (see Equation (2.4)). To weight an image so that one process is dominant, the influence of the other processes is minimized.

Analogy: The heat analogy



The mechanisms of weighting are well described using an analogy of a gas stove that has two knobs labeled TR and TE. The TR knob controls the amount of T1 contrast; the TE knob controls the amount of T2 contrast. The TR knob turns the heat up or down on T1 contrast. The TE knob turns the heat up or down on T2 contrast.

- Turning the TR knob down turns the heat up on T1 contrast, i.e. T1 contrast increases.
- Turning the TE knob up turns the heat up on T2 contrast, i.e. T2 contrast increases.

To weight an image, it is necessary to turn the heat up on one intrinsic contrast parameter and the heat down on the others. For example, for T1 weighting turn the heat up on T1 contrast and the heat down on T2 contrast so the image is weighted toward T1 contrast and away from T2 contrast (proton density depends on the relative number of protons and cannot be changed for a given area).

- To turn the heat up on T1 contrast, the TR is short (TR knob down).
- To turn the heat down on T2 contrast, the TE is short (TE knob down) (Figure 2.19).

For T2 weighting, turn the heat up on T2 contrast and the heat down on T1 contrast. The image is weighted toward T2 contrast and away from T1 contrast (proton density depends on the relative number of protons and cannot be changed for a given area).

- To turn the heat up on T2 contrast, the TE is long (TE knob up).
- To turn the heat down on T1 contrast, the TR is long (TR knob up) (Figure 2.20).

For proton density weighting, turn the heat down on T1 contrast and the heat down on T2 contrast. In this way, proton density contrast predominates.

- To turn the heat down on T1 contrast, the TR is long (TR knob up).
- To turn the heat down on T2 contrast, the TE is short (TE knob down) (Figure 2.21).

The heat analogy is used elsewhere in this book. Look out for the flame symbol in the margin.




Figure 2.19 T1 weighting and the heat analogy.




Figure 2.20 T2 weighting and the heat analogy.




Figure 2.21 Proton density weighting and the heat analogy.

Scan tip: Understanding weighting

A good first step in learning weighting is to look for the water content in the image, and, if it has a high signal, the image is likely to be T2-weighted and acquired with a long TE. If water has a low signal, it is likely to be T1-weighted and acquired with a short TR. However, depending on the area of the body, some PD-weighted images also have dark water. Fat is an unreliable marker, as it is bright in many types of weighting depending on the pulse sequence.

To demonstrate variables in image contrast, look at Figure 2.22. It was acquired using a standard spin-echo sequence and is a T1-weighted image so the contrast is predominantly due to differences in the T1 recovery times of the tissues. It has contrast expected of an image acquired with a short TR and TE, i.e. fat in the scalp and bony marrow of the clivus is hyperintense, and water in the CSF is hypointense. However, looking more closely, it is clear that not all areas of high signal are fat, and not all areas of low signal are water. For example, the area labeled A, which has a high signal, is not fat but slow-flowing blood in the superior sagittal sinus. The area labeled B, which has a low signal, is not water but air in the sphenoid sinus. Although this image is predominantly T1-weighted, there are also flow and proton density effects contributing to image contrast.



Figure 2.22 Midline sagittal T1-weighted image of the brain.

Learning tip:

Relaxation definitions

It is very important to understand the differences between T1 recovery, T1 recovery time, T1 contrast and T1 weighting, and the equivalent distinctions in terms of T2 relaxation processes. Table 2.12 lists the various definitions of these terms.

OTHER CONTRAST MECHANISMS

We have explored the main image contrast mechanisms. These are the types of contrast you are likely to encounter on a day-to-day basis. However, there are other techniques that are used to generate very specific image contrast:

- Diffusion-weighted imaging (DWI)
- Functional MRI (fMRI)

- Magnetization transfer contrast (MTC)
- Susceptibility-weighted imaging (SWI)
- Contrast agents.

Table 2.12 Image contrast definitions.

T1 recovery	The recovery of longitudinal magnetization due to spin–lattice relaxation after the RF excitation pulse is switched off
T1 recovery time	The time it takes for 63% of the longitudinal magnetization to recover in a tissue
T1 contrast	An image where fat is hyperintense and water is relatively hypointense because the TR is short enough to not allow full recovery of the vectors
T1 weighting	An image whose contrast is predominantly due to the differences in the T1 recovery times of the tissues
T2 decay	The decay of coherent transverse magnetization due to spin–spin relaxation after the RF excitation pulse is switched off
T2 decay time	The time it takes for 63% of the coherent transverse magnetization to decay in a tissue
T2 contrast	An image where fat is hypointense and water is relatively hyperintense because the TE is long enough to allow full dephasing
T2 weighting	An image whose contrast is predominantly due to the differences in the T2 decay times of the tissues

Diffusion-weighted imaging (DWI)

Diffusion is a term used to describe the movement of molecules in the extracellular space due to random thermal motion. This motion is restricted by boundaries such as ligaments, membranes, and macromolecules Figure (2.23). Sometimes restrictions in diffusion are directional, depending on the structure of the tissues, and diffusion is also restricted in pathology. The net displacement of molecules diffusing across an area of tissue per second is called the **apparent diffusion coefficient (ADC)**, and this is one of the intrinsic contrast parameters listed at the beginning of this chapter (Table 2.13). It is therefore a parameter that affects image contrast, but it is intrinsic to the tissue and not therefore under our control. In areas of restricted diffusion, the ADC is low because the extracellular space is small. Examples of this type of tissue are ligaments, and many types of pathology. In areas of free diffusion, the ADC is high because the extracellular space is large. Examples of this type of tissue are normal gray matter and normal liver tissue.

Diffusion Weighted Images (DWI) denotes those whose contrast is determined by the ADC. This is achieved by using gradients (see Chapter 5). In this technique, differences in the ADC are revealed by applying two gradients. The first gradient dephases magnetic moments of hydrogen nuclei, and the second gradient attempts to rephase them. In tissues where the ADC is low, the molecules (and therefore the hydrogen nuclei that make them up) are essentially stationary because their diffusion is limited. Magnetic moments of these spins acquire no net phase change after the gradients are applied. This is because they do not move between each gradient application. The first gradient dephases the magnetic moments of the hydrogen nuclei, but then the second gradient rephases them. As a result, a high signal is obtained from tissues with a low ADC, as the magnetic moments of spins within them are coherent and produce a large component of transverse magnetization.

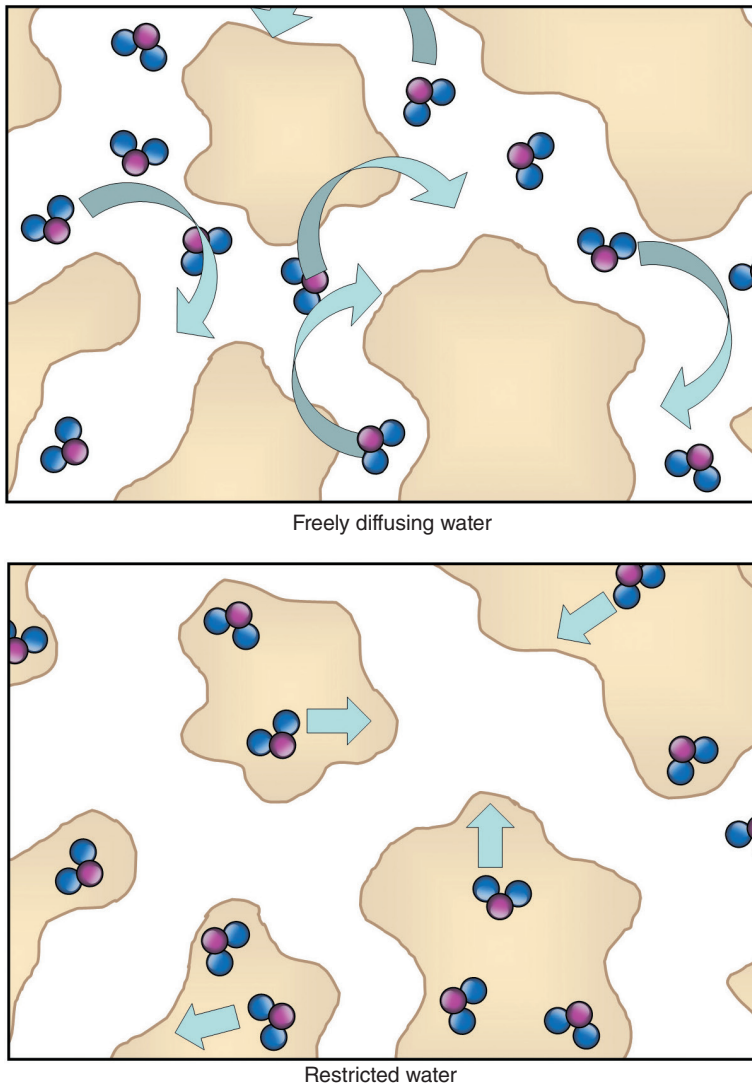


Figure 2.23 Free and restricted diffusion in water.

Table 2.13 Typical ADC values in the brain.

	ADC ($\times 10^{-3}$ mm ² /s)	Relative signal when $b=1000$
Cerebral spinal fluid	2.94	0.05
Gray matter	0.76	0.47
White matter	0.45	0.63

Magnetic moments of moving hydrogen nuclei, however, acquire phase change, and this results in a signal loss. This is because the molecules (and therefore the hydrogen nuclei that make them up) diffuse and therefore move between the application of each gradient. The first gradient dephases the magnetic moments of the hydrogen nuclei, but then the second gradient cannot rephase them because they move in the meantime. In diffusion imaging, normal tissue that exhibits a high ADC has a lower signal intensity than abnormal tissue that has a low ADC, as the molecules within it are free to move. Diffusion becomes restricted when pathology is present, and so the signal intensity is higher. Signal change depends on the ADC of the tissue and the strength, duration, and interval of the gradients (collectively known as the ***b* factor/value** expressed in units of s/mm^2) (Equation (2.8)). This is one of the extrinsic contrast parameters discussed earlier in this chapter. In DWI, an extrinsic contrast parameter (*b* factor) controls how much a tissue's intrinsic contrast parameter (ADC) contributes toward image weighting. As the *b* factor increases, so does diffusion weighting because the contribution from differences between the ADC of different tissues to image weighting also increases.

Equation 2.8

$$b = \gamma^2 \times G^2 \times \delta^2 \times (\Delta - \delta/3)$$

b is the *b* value or *b* factor (s/mm^2)
 γ is the gyromagnetic ratio (MHz/T)
 G is the gradient amplitude (mT/m)
 δ is the gradient duration (ms)
 Δ is the time between two gradient pulses (ms)

The *b* value or *b* factor is a function of the amplitude, duration, and interval of the gradients in the Stejskal–Tanner scheme



Refer to animation 2.2 on the supporting companion website for this book:
www.wiley.com/go/westbrook/mriinpractice

Functional MRI

Functional MRI (fMRI) is a rapid MRI technique that acquires images of the brain during activity or stimulus and then at rest. Image contrast depends on a physiological process called **blood oxygenation level dependent (BOLD)**. BOLD exploits differences in the magnetic susceptibility of oxyhemoglobin and deoxyhemoglobin because of increased cerebral blood flow and little or no increase in local oxygen consumption that occurs during stimulation. Because deoxyhemoglobin is paramagnetic, vessels containing a significant amount of this molecule create local field inhomogeneities causing dephasing and therefore signal loss. During activity, blood flow to the cortex increases, causing a reduction of deoxyhemoglobin, which results in a decrease in dephasing and a corresponding increase in signal intensity. These effects are very short-lived and therefore require extremely rapid sequences (see Chapter 4). To exploit the rapid dephasing effects, BOLD images are usually acquired with a long TE while the task is modulated on and off. The “off” images are then subtracted from the “on” images, and a sophisticated statistical analysis is performed.

Regions that were activated above some threshold level are overlaid on anatomical images Figure (2.24). It is these regions that reflect brain activity.

Magnetization transfer contrast

Magnetisation Transfer Contrast (MTC) is another mechanism that generates images with a certain contrast. It involves the fast exchange of energy between bound and free nuclei. Bound nuclei are those that are restricted and have a very short T2 decay time. Their T2 decay times are so short they cannot normally be imaged. However, bound nuclei reduce the signal intensity of the free nuclei. Free nuclei are observable because they have longer T2 decay times [6]. The magnetic moments of the bound nuclei have a much broader precessional frequency range than the magnetic moments of free nuclei and are therefore excited by an RF excitation pulse that is several kilohertz away from the frequency of the free nuclei. The energy absorbed by the bound pool of nuclei causes saturation, and magnetization is transferred to the free pool of nuclei. This causes a reduction in their signal intensity.

Susceptibility weighting (SWI)

SWI uses the magnetic susceptibility differences between tissues to generate image contrast. Gradient-echo sequences, in conjunction with a long TE, are commonly used, as they enhance the differences in magnetic susceptibility between tissues (see Chapter 4).




Figure 2.24 BOLD images of the brain. Functional areas are shown in red.

Table 2.14 Things to remember – image contrast special techniques.

DWI is a technique that sensitizes a spin-echo-type sequence to diffusion motion by using strong gradients
The ADC is an intrinsic contrast parameter and signifies the net displacement of molecules in the extracellular space per second
The <i>b</i> value is an extrinsic contrast parameter that controls how much the intrinsic ADC influences image contrast – hence the term diffusion-weighted imaging
Functional imaging techniques are used to image the function or physiology of a system rather than its anatomy
fMRI relies on a process called BOLD to produce a signal in areas of the brain where there is increased activity after performing a function (such as finger tapping)

Contrast agents

Pathology often has a large number of water nuclei, and therefore T2-weighted images display a good intrinsic contrast between pathology and normal tissue. To increase conspicuity, contrast enhancement agents may be used that selectively change the relaxation times of certain tissues. This effect is known as **relaxivity** and enables better visualization of tissues that are affected by the agent. Contrast agents are not imaged themselves but are seen because they indirectly affect the relaxation times of water nuclei.

Contrast agents are usually characterized by whether they affect T1 or T2 relaxation times. Those that shorten T1 recovery times are called **T1 agents**, and those that shorten T2 decay times are called **T2 agents**. The degree of shortening depends on the concentration of the agent. Most contrast agents used in MRI are T1 agents. The most commonly used agent is gadolinium. Gadolinium (Gd) is a rare-earth metal and, in its natural form, is highly toxic. It cannot be excreted by the body and would cause long-term side effects as it binds to membranes. It is made safe by binding or chelating the gadolinium to other molecules such as diethylene triaminepentaacetic acid (**DTPA**) (a ligand), which is safely excreted.

As we learned earlier in this chapter, molecules that tumble with a frequency at or near the Larmor frequency have shorter T1 recovery times than other molecules. Gadolinium is a paramagnetic agent and therefore has a large magnetic moment. It is a trivalent lanthanide element that has seven unpaired electrons and an ability to allow rapid exchange of bulk water to minimize the space between itself and water within the body. The chelated molecules of gadolinium are relatively large and are bonded to metal ions. Fat molecules are too large to get close enough to the gadolinium to affect their relaxivity. Water molecules, however, are able to diffuse close to the tumbling molecules of gadolinium. These fluctuations occur near the Larmor frequency and decrease the T1 recovery times of nearby hydrogen nuclei. Enhancing tissues are hyperintense on a T1-weighted image.

T2 agents are typically superparamagnetic macromolecules of iron. They distort the local magnetic field so that the magnetic moments of neighboring hydrogen nuclei dephase. This decreases the T2 decay times of enhancing tissue, and they are therefore hypointense on T2-weighted images.

In this chapter, we explored the different methods we use to generate image contrast in MRI. In most examinations, this involves creating images that are T1, T2, PD, and diffusion-weighted. Most modern systems have a facility to acquire all these weightings from a single

acquisition. Images are calculated using a variety of different TRs and TEs from a single specialized scan. It is also possible to calculate the parameters necessary to null certain tissues [7]. In the next two chapters, the fundamental principles of image contrast are used to describe how pulse sequences generate images of a certain type of weighting.



For questions and answers on this topic please visit the supporting companion website for this book: www.wiley.com/go/westbrook/mriinpractice

References

1. Hashemi, R.H., Bradley Jr, W.G., and Lisanti, C.J. (2010). *MRI: The Basics*, 3, 40. Philadelphia, PA: Lippincott Williams and Wilkins.
2. Dale, B.M., Brown, M.A., and Semelka, R.C. (2015). *MRI: Basic Principles and Applications*, 5, 21. Wiley.
3. Hashemi, R.H., Bradley Jr, W.G., and Lisanti, C.J. (2010). *MRI: The Basics*, 3, 42. Philadelphia, PA: Lippincott Williams and Wilkins.
4. McRobbie, D.W., Moore, E.A., Graves, M.J. et al. (2017). *From Picture to Proton*, 3, 138. Cambridge: Cambridge University Press.
5. Dale, B.M., Brown, M.A., and Semelka, R.C. (2015). *MRI: Basic Principles and Applications*, 5, 19. Wiley.
6. McRobbie, D.W., Moore, E.A., Graves, M.J. et al. (2017). *From Picture to Proton*, 3, 140. Cambridge: Cambridge University Press.
7. Hagiwara, A., Warntjes, M., and Hori, M. (2017). SyMRI of the brain: rapid quantification of relaxation rates and proton density with synthetic MRI, automatic brain segmentation and myelin measurement. *Investigative Radiology* 52(10): 647–657.

3

Spin-echo pulse sequences

Introduction	58	Inversion recovery (IR)	78
RF rephasing	59	Short tau inversion recovery (STIR)	82
Conventional spin-echo	65	Fluid attenuated inversion recovery (FLAIR)	84
Fast or turbo spin-echo (FSE/TSE)	68		

After reading this chapter, you will be able to:

- *Explain the purpose of pulse sequences.*
- *Describe how spin-echoes are created.*
- *Understand the mechanisms of common spin-echo pulse sequences.*
- *Apply what you have learned to understand how images of different weighting are created using spin-echo pulse sequences.*

INTRODUCTION

Pulse sequences enable control of the way in which the system applies RF pulses and gradients. They are used to determine image weighting. Dephasing, caused by magnetic field inhomogeneities, produces a rapid loss of coherent transverse magnetization (and therefore signal) so that it reaches zero before most tissues have had time to attain their T1 or T2 relaxation times. As we explored in Chapter 2, the FID decays within about 10 ms, which is too fast to measure any significant relaxation. Pulse sequences are methods used by the MR system to rephase the magnetic moments of hydrogen nuclei at a later point in time. This rephasing produces a signal called an **echo**. As data are collected from the echo later in the sequence, image contrast relies on the differences in the T1 recovery times, T2 decay times, or proton density between tissues.

There are two ways of rephasing the magnetic moments of hydrogen nuclei to produce an echo – by using an additional 180° RF pulse or by using gradients. Sequences that use a 180° RF rephasing pulse to generate an echo are called **spin-echo pulse sequences**; those that use a gradient are called **gradient-echo pulse sequences** (see Chapter 4).

Learning tip: What is a pulse sequence diagram?

Pulse sequences are a timed sequence of events. These events are RF pulses and gradients. Different pulse sequences are illustrated using schematic diagrams that represent what each of the hardware components of the system is doing at different points in time in the sequence. The time that elapses during the sequence is represented on a horizontal axis. There are usually five horizontal lines spaced vertically from each other. Three lines represent each of the three gradients (slice select, phase encoding, and frequency encoding – see Chapter 5); the other two lines represent RF pulses and signal respectively. Gradients are drawn as shapes above and below the horizontal line. If the shape is above the line, this indicates a positive polarity gradient; below the line, a negative polarity gradient. The amount of deviation from the horizontal indicates the amplitude of the gradient (see Chapter 5).

Analogy: The dance analogy

The definition of a pulse sequence is a series of RF pulses, gradient applications, and intervening time periods. RF pulses are applied for excitation or rephasing purposes. Gradients are applied to spatially encode signal (see Chapter 5) and sometimes to rephase and dephase magnetic moments of hydrogen nuclei to produce an echo (see Chapter 4). The intervening time periods refer to the time intervals between these various functions, some of which are extrinsic contrast parameters that we select in the scan protocol (see Chapters 2 and 7). Therefore, a pulse sequence is a carefully coordinated and timed sequence of events that generates a certain type of image weighting.

Pulse sequences may be thought of like dances. All dances involve the movement of arms and legs, just as all pulse sequences involve RF pulses and gradients. However, just as the timing and coordination of arms and legs determine the type of dance, e.g. tango, foxtrot, etc., so the timing and coordination of RF pulses and gradients determine image weighting.

There are many different pulse sequences, and each is designed for a specific purpose. This chapter discusses the mechanisms, uses, and parameters for each of the common spin-echo pulse sequences and their advantages and disadvantages. Each manufacturer uses different acronyms to distinguish between individual pulse sequences. A table is included that compares the common acronyms for spin-echo pulse sequences for the main manufacturers (Table 3.1). A more comprehensive table is also provided at the beginning of the book. These tables are only a guide; they are not meant to compare the performance or specification of each system. The parameters included in this chapter depend on field strength and the nuances of individual systems. However, they should be suitable for most field strengths used in clinical imaging.

RF REPHASING

All spin-echo pulse sequences are characterized by RF rephasing. The spin-echo pulse sequence commonly uses a 90° RF excitation pulse to flip the NMV fully into the transverse plane. The NMV precesses in the transverse plane inducing a voltage in the receiver coil. A FID occurs when the 90° RF excitation pulse switches off (see Chapter 1). $T2^*$ dephasing from inhomogeneities


Table 3.1 Spin-echo pulse sequences and their common acronyms.

Generic	GE	Philips	Siemens	Toshiba	Hitachi
Conventional spin-echo	SE	SE	SE	SE	SE
Fast or turbo spin-echo	FSE	TSE	TSE	FSE	FSE
Inversion recovery	IR	IR	IR	IR	IR
STIR	STIR	STIR	STIR	Fast STIR	STIR
FLAIR	FLAIR	FLAIR	FLAIR	Fast FLAIR	FLAIR
DRIVE	FR-FSE	DRIVE	RESTORE	T2 PULS FSE	Driven equilibrium FSE

Abbreviations used in Table 3.1

SE	spin-echo	STIR	short tau inversion recovery
FSE	fast spin-echo	FLAIR	fluid attenuated inversion recovery
TSE	turbo spin-echo	DRIVE	driven equilibrium
IR	inversion recovery		

in the B_0 field occurs almost immediately, and signal decays to zero. After a time called **tau**, another RF pulse is used to compensate for this dephasing and refocus or rephase the magnetic moments of hydrogen nuclei (Figure 3.1). It commonly has a magnitude of 180° and is called the **180° RF rephasing pulse**.

**Figure 3.1** T2* dephasing.

T_2^* dephasing causes the magnetic moments of hydrogen nuclei to dephase or “fan out” in the transverse plane. The magnetic moments are now out of phase with each other on the transverse plane, i.e. they are at different positions on the precessional path at any given time. The magnetic moments that slow down form the trailing edge of the fan (shown in blue in Figures 3.2 and 3.3). The magnetic moments that speed up form the leading edge of the fan (shown in red in Figures 3.2 and 3.3). The 180° RF rephasing pulse is then applied. It has sufficient energy to move the NMV through 180° . As the NMV is still in the transverse plane, it remains in this plane but at a physically opposite position to that before the RF rephasing pulse is applied. The 180° RF rephasing pulse flips the individual magnetic moments through 180° (rather like flipping a pancake). They are still in the transverse plane, but magnetic moments that formed the trailing edge before the 180° RF rephasing pulse now form the leading edge. Conversely, magnetic moments that formed the leading edge before the 180° RF rephasing pulse now form the trailing edge (as shown in the bottom half of Figure 3.2).

The direction of precession remains the same, and so the trailing edge begins to catch up with the leading edge because, due to magnetic field inhomogeneities, these magnetic moments precess faster than the trailing edge. At a specific time (TE), the two edges are superimposed. The magnetic moments of hydrogen nuclei are momentarily in phase because they are all at the same place on the precessional path. At this instant, there is in-phase transverse magnetization, and so maximum signal is induced in the receiver coil. This signal is called a **spin-echo**.

In spin-echo pulse sequences, T_2^* dephasing is eliminated by the 180° RF rephasing pulse because magnetic field inhomogeneities are largely predictable. T_2 decay is not affected by the

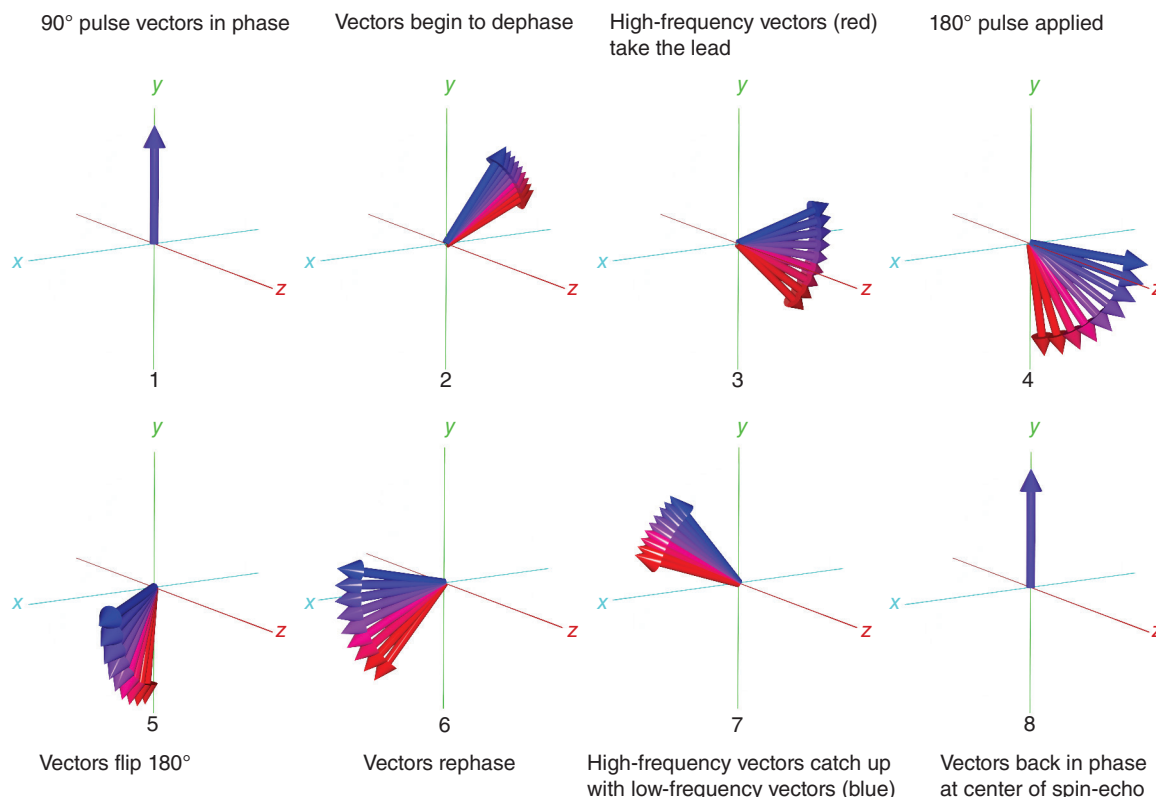


Figure 3.2 180° RF rephasing.

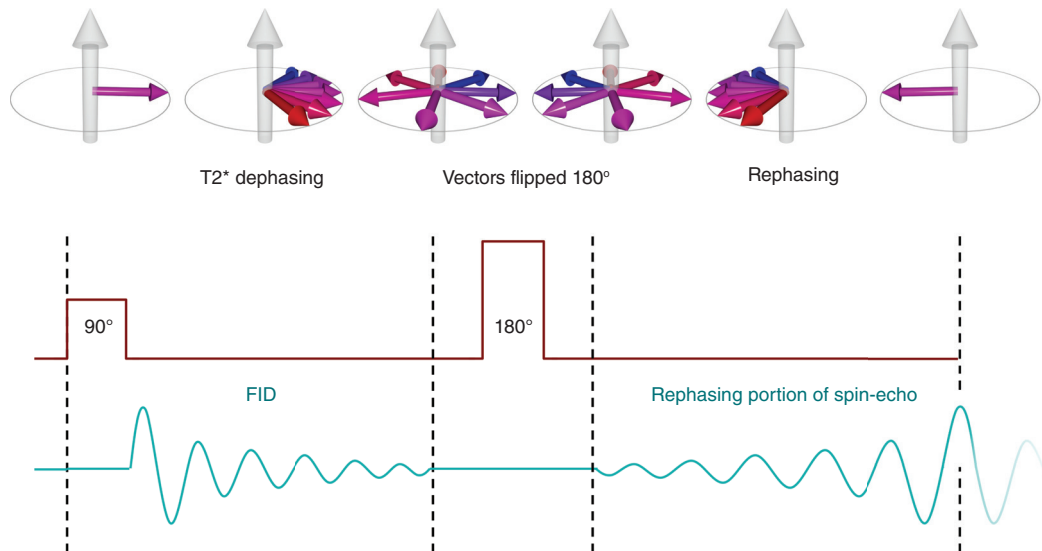


Figure 3.3 A basic spin-echo pulse sequence.

180° RF rephasing pulse because this is caused by spin–spin interactions, which randomly fluctuate [1]. In addition, by rephasing the spin-echo at a later point in time in the pulse sequence, time is allowed for tissues to reach their T1 and T2 relaxation times, and therefore a different image weighting is obtained (Figure 3.3).



Refer to animation 3.1 on the supporting companion website for this book:
www.wiley.com/go/westbrook/mriinpractice

Analogy: The Larmor Grand Prix

An easy way to understand 180° RF rephasing is to imagine three cars on a circular racetrack. The cars relate to three magnetic moments of three different hydrogen nuclei and the circular racetrack to the precessional path of the magnetic moments. The cars have varying speeds; one is a racing car, one a family saloon, and one a tractor (Figure 3.4). At the sound of the start gun, the cars set off around the track. Very shortly, the racing car pulls ahead of the family car, which in turn sprints ahead of the tractor. They are now out of phase with each other, as they are in a different place on the track to each other at a given time. The longer the race runs, the more dephasing between the vehicles.

The start gun is fired again. The start gun now refers to the 180° RF rephasing pulse. On hearing the gun, the cars turn around through 180° and head back towards the start line again. The racing car is now at the back because it traveled furthest at the beginning of the race. The tractor is at the front because it traveled slower at the beginning of the race. The family saloon is somewhere in-between. Assuming the cars travel back to the start line at the same speed as they traveled out at the beginning of the race, the racing car and family saloon catch up with the tractor and are at the same place at the same time when they get back to the start line. They are therefore back in phase and, if they were magnetic moments of hydrogen nuclei, they would generate a spin-echo at this point. The time taken for the cars to complete the whole race (from the start line to the point where they turn around and back to the start line again) corresponds to the TE.

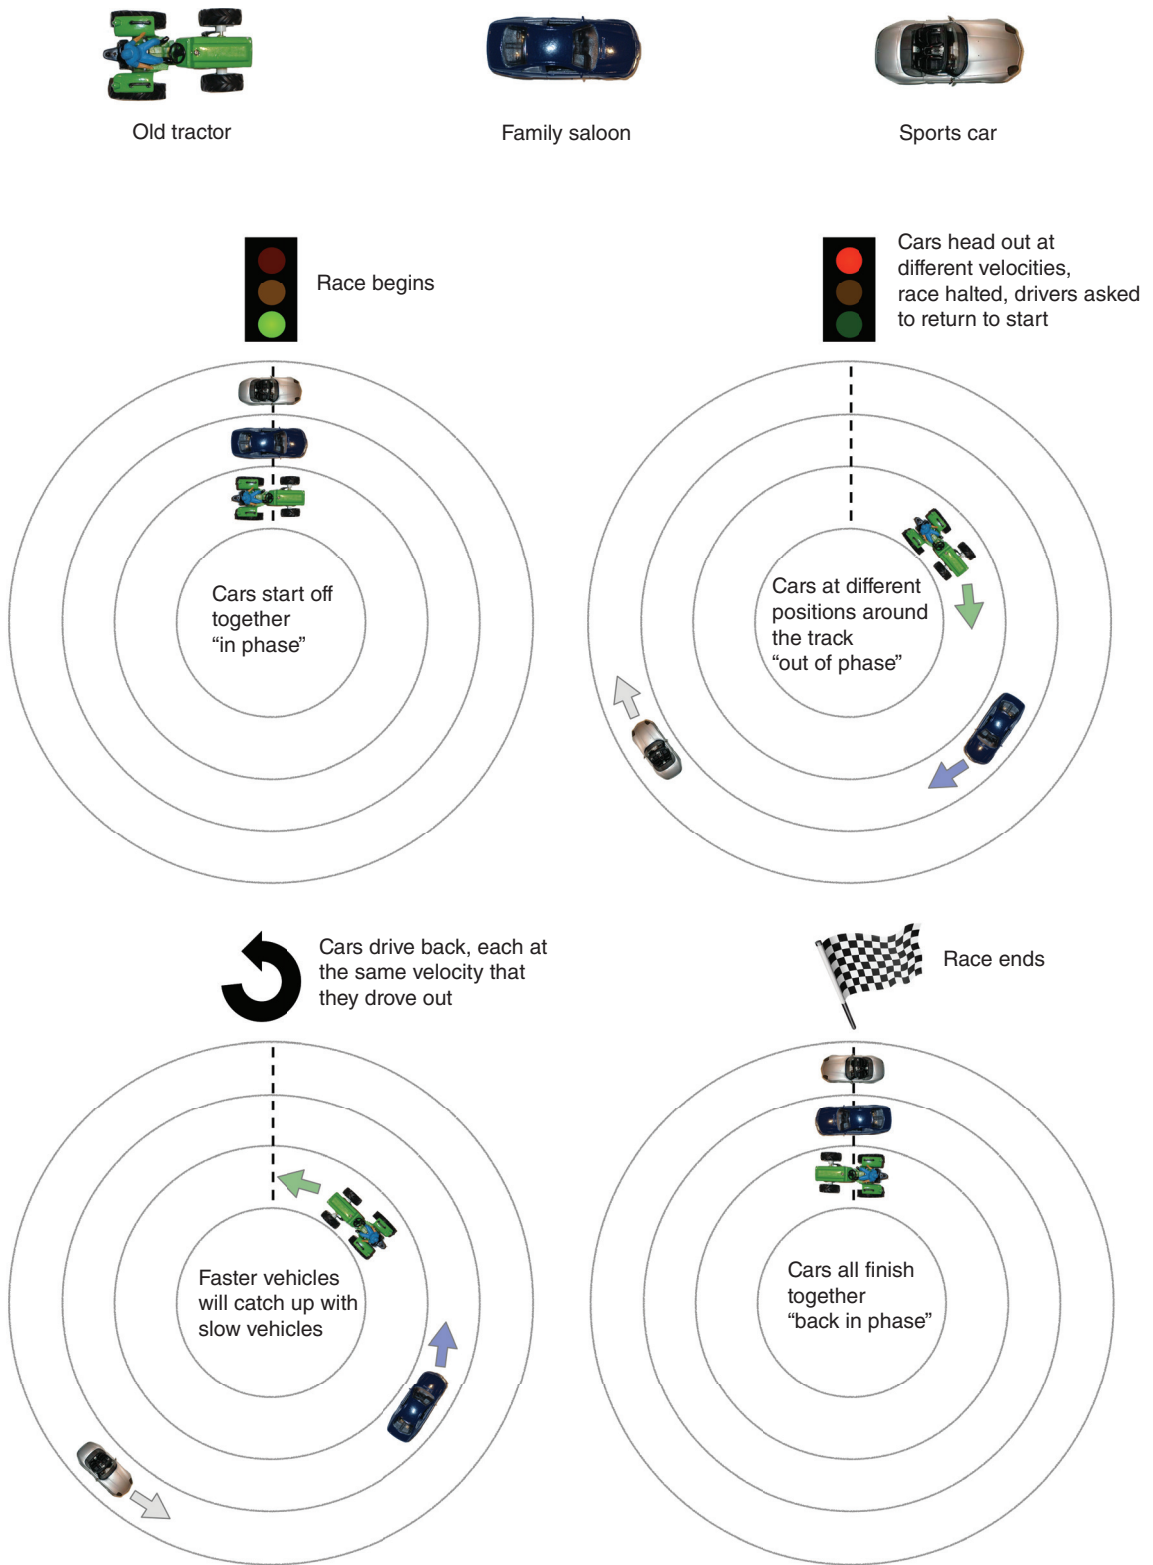


Figure 3.4 The Larmor Grand Prix.



Refer to animation 3.2 on the supporting companion website for this book: www.wiley.com/go/mriinpractice

The TR is the time between each 90° RF excitation pulse for each slice. The TE is the time between the 90° RF excitation pulse and the peak of the spin-echo (Figure 3.5). The time taken to rephase after the 180° RF rephasing pulse equals the time to dephase when the 90° RF excitation pulse is withdrawn. This time is called tau. The TE is therefore twice tau and the system times the 180° RF rephasing pulse by halving the TE selected in the scan protocol.

Look at Figure 3.5 and note the symmetry of the spin-echo. As the magnetic moments of hydrogen nuclei gradually come into phase, signal gradually builds, reaching a peak at the TE when all magnetic moments are in phase. However, the magnetic moments that are precessing rapidly soon overtake those that are precessing slowly, and dephasing occurs again. This results in a gradual loss of signal, which mirrors the gradual growth before the peak of the echo. This accounts for the symmetry of the spin-echo.

Having described the fundamental principles of rephasing, it is time to explore the variety of pulse sequences in the spin-echo family. These are generically called;

- conventional spin-echo
- fast or turbo spin-echo (FSE/TSE)
- inversion recovery, which includes STIR and FLAIR.

The mechanisms of pulse sequences and their appropriate timing parameters are important. In the following section, scan tips link the theory of spin-echo pulse sequences to practice. Theory is related to what is going on “behind the scenes” when we select a timing parameter in the scan protocol.




Figure 3.5 Tau.

CONVENTIONAL SPIN-ECHO

Mechanism

65

Conventional spin-echo uses a 90° RF excitation pulse followed by one or more 180° RF rephasing pulses to generate one or more spin-echoes. Each 180° RF rephasing pulse generates a separate spin-echo that is received by a coil and used to create an image. Although any number of echoes may be created, spin-echo sequences typically generate either one or two echoes.

Contrast is mainly determined by the spin-echo, but there is also a contribution from the fact that the magnetic moments of hydrogen nuclei are rephased by negative polarity applications of the slice select and frequency-encoding gradients [2]. In addition, spoiler gradients are applied at the end of each TR period to ensure that there is no coherent transverse magnetization at the beginning of the next repetition (see Chapter 4).

Spin-echo using one echo

This pulse sequence is used to produce T1-weighted images by selecting a short TR and a short TE. One 180° RF rephasing pulse is applied after the 90° RF excitation pulse. The single 180° RF rephasing pulse generates a single spin-echo. Timing parameters are usually selected to produce a single T1-weighted image. A short TE ensures that the 180° RF rephasing pulse and subsequent spin-echo occur early so that only a little T2 decay occurs. Differences in T2 decay times of the tissues are minimized and do not, therefore, dominate the spin-echo and its contrast. A short TR, however, ensures that fat and water vectors do not fully recover, and so the differences in their T1 recovery times dominate the spin-echo and its contrast (Figure 3.6). A single T1-weighted image is therefore obtained for every slice location.




Figure 3.6 Spin-echo with one echo.

Spin-echo using two echoes

This is used to produce both a proton density and a T2-weighted image in the TR period. The first spin-echo is generated early by selecting a short TE. Only a little T2 decay occurs, and so T2 decay time differences between the tissues are minimized in this echo. The second spin-echo




Figure 3.7 Spin-echo with two echoes.

is generated much later by selecting a long TE. A significant amount of T2 decay occurs, and so differences in the T2 decay times of the tissues are maximized in this echo. The TR is long so that T1 recovery differences between the tissues are minimized in each spin-echo. The first spin-echo therefore has a short TE and a long TR, and is PD-weighted. The second spin-echo has a long TE and a long TR, and is T2-weighted (Figure 3.7). Two images are therefore produced for every slice location. One is PD-weighted, and the other is T2-weighted.

Learning tip: The first echo is “free”

In dual echo spin-echo sequences (or indeed in any spin-echo pulse sequence where more than one 180° RF rephasing pulse is applied to generate more than one spin-echo), the first spin-echo is “free.” This means it does not “cost” anything in terms of scan time to acquire it. The scan time is not reduced by not acquiring the first spin-echo in a dual echo spin-echo sequence because it is necessary to wait for the second spin-echo to occur anyway. While waiting to collect data from the second spin-echo, data are collected from the first spin-echo during the waiting period. If multiple spin-echoes are acquired, e.g. four, then the first three spin-echoes are free. This is not the case, however, in FSE/TSE, which is discussed later.

Uses

Spin-echo pulse sequences are considered the gold standard in that the contrast they produce is understood and predictable. They produce T1-, T2-, and PD-weighted images of good quality and in most parts of the body (Table 3.2). However, due to relatively long scan times, PD- and T2-weighted images are often acquired using FSE/TSE (see next section).

Table 3.2 Advantages and disadvantages of spin-echo.

Advantages	Disadvantage
Good image quality	Long scan times
Very versatile	
True T2 weighting	
Available on all systems	
Gold standard for image contrast and weighting	

Suggested parameters

Single echo (for T1 weighting):

- TR 300–700 ms
- TE 10–30 ms.

Dual echo (for PD/T2 weighting):

- TR 2000 ms+
- TE1 20 ms
- TE2 80 ms.

Table 3.3 Things to remember – conventional spin-echo.

Spin-echo sequences are characterized by 180° RF rephasing pulses that refocus the magnetic moments of spins to produce an echo

T1, T2, and PD weighting are all achievable using conventional spin-echo

Conventional spin-echo is traditionally used to acquire one or two echoes to achieve T1, T2, or proton density weighting

Although they are old sequences, they are still considered the gold standard and can be used to image anatomy and pathology in all body areas

Scan tip:

Parameter selection in spin-echo – what's going on behind the scenes?

When we alter extrinsic contrast parameters in spin-echo pulse sequences, behind the scenes, we determine image weighting (see Chapter 2). When we select the TR in the scan protocol, we control how much T1 recovery is permitted between each RF excitation pulse. We therefore control to what extent T1 contrast influences image weighting. We also determine the SNR, the scan time and the slice number (see Chapter 7) but these factors are not usually as important as weighting.

When we select the TE in the scan protocol we control how much T2 decay is permitted between the RF excitation pulse and the peak of the spin-echo. We therefore control to what extent T2 contrast influences image weighting. We also determine the SNR (see Chapter 7) but this is not usually as important as weighting.

FAST OR TURBO SPIN-ECHO (FSE/TSE)

Mechanism

68

As the name suggests, FSE or TSE is a spin-echo pulse sequence but with scan times that are much shorter than conventional spin-echo. It is also known as **RARE (Rapid Acquisition with Relaxation Enhancement)** [3].

Learning tip: TSE, gradients, and k-space

A good understanding of spatial encoding and k-space is needed for the following section, so it might be a good idea to read Chapters 5 and 6 before learning this sequence.

The scan time is a function of the TR, the number of signal averages (NSA) and the phase matrix (see Chapter 7 and Equation (6.7)). The scan time is reduced by decreasing one or more of these parameters. In TSE, the scan time is decreased by modifying the phase matrix component of this equation. The number of phase-encoding steps is maintained so that the phase matrix is unchanged; however, in TSE, the number of phase-encoding steps *per TR* is increased. As a result, *k*-space is filled more efficiently, and the scan time decreases.

In conventional spin-echo, one phase-encoding step is applied per TR on each slice, and therefore only one line of *k*-space is filled per TR (Figure 3.8). In TSE, the scan time is reduced by performing more than one phase-encoding step and subsequently filling more than one line of *k*-space per TR. This is achieved by using several 180° RF rephasing pulses to produce several spin-echoes to form an **echo train** (Figure 3.9). Each rephasing produces a spin-echo, and a different phase-encoding step is performed on this echo. In conventional spin-echo, raw image data from each spin-echo are stored in *k*-space, and each spin-echo is used to produce a *separate image* (usually PD- and T2-weighted; see earlier in this chapter). In TSE, data from each spin-echo are placed into *one* image. The number of 180° RF rephasing pulses performed every TR corresponds to the number of spin-echoes produced in the echo train and the number of lines of *k*-space filled with data from these echoes. This number is called the **turbo factor** or the **echo train length (ETL)**. The higher the turbo factor, the shorter the scan time, as more phase-encoding steps are performed per TR. For example:

- In conventional spin-echo, if a 256 phase matrix is selected, 256 phase encodings are performed. Assuming 1 NSA is also selected, 256 TR periods must elapse to complete the scan.
- In TSE, using the same parameters but selecting a turbo factor of 16, 16 phase-encoding steps are performed every TR. Therefore $256 \div 16$ (i.e. 16) TR periods must elapse to complete the scan. The scan time is therefore reduced to 1/16 of the original (Equation (3.1)).

At each 180° RF pulse/phase encoding combination, a different amplitude of phase-encoding gradient slope is applied to fill a different line of *k*-space. In TSE, several lines corresponding to the turbo factor are filled every TR (Figure 3.9). Therefore, *k*-space is filled more rapidly, and the scan time decreases.




Figure 3.8 Spatial encoding in conventional spin-echo.




Figure 3.9 The echo train.

Equation 3.1

$$ST = \frac{TR \times M(p) \times NSA}{ETL}$$

ST is the scan time (s)
 TR is the repetition time (ms)
 M(p) is the phase matrix
 NSA is the number of signal averages
 ETL is the echo train length or turbo factor

This equation enables the scanner to calculate the scan time in TSE. The longer the echo train, the shorter the scan time, but this may result in fewer slices per TR

Learning tip: The chest of drawers and TSE



In Chapter 6, we sometimes use the chest of drawers analogy to explain *k*-space. Using this analogy, in conventional spin-echo one drawer is opened per TR to fill one line of *k*-space with data points. In TSE, all the drawers are still filled, but more than one drawer is opened per TR to fill *k*-space more quickly, so the scan time decreases. This is achieved by performing more than one application of the phase-encoding gradient per TR, each one to a different slope and/or polarity to open a different drawer.

For example, if 10 drawers are opened per TR, then the phase-encoding gradient is applied 10 different times to 10 different amplitudes and/polarity per TR to open 10 different drawers. Once the drawers are opened, there must be data to put into them. This requires producing 10 echoes, one for each drawer. To do this, 10 different 180° RF rephasing pulses are applied. The number of 180° RF rephasing pulses corresponds to the number of echoes and the number of drawers opened per TR. This is called the ETL or turbo factor and determines how much faster the scan is compared with conventional spin-echo, i.e. a turbo factor of 10 indicates 10 drawers are opened per TR, and the scan time is 10 times faster than conventional spin-echo.

Weighting in TSE

70

Spin-echoes are generated at different TEs, and therefore data collected from them have a different contrast. All these data are stored in k -space and ultimately used to create a single image. So how is it possible to correctly weight a TSE sequence? The answer is that the selected TE is only an **effective TE**. In other words, it is the TE at which we wish to *effectively* weight the image. To achieve this weighting, the system organizes where to place data from each spin-echo in k -space. As described in Chapter 6, each phase-encoding step applies a different slope of gradient and/or polarity to phase encode each spin-echo differently. The data from these spin-echoes are therefore placed in different lines in k -space (Figure 3.10).

Very steep phase encoding slopes result in spin-echoes with a low amplitude, and the data from them are placed in the outer lines of k -space. Shallow phase-encoding slopes result in spin-echoes that have a high amplitude, and the data from them are placed in the central lines of k -space (Figure 3.11) (see Chapter 6). In TSE, the system orders the phase-encoding steps so that instead of filling k -space lines in a linear fashion (either from top to bottom or from bottom to top), the lines are filled depending on how close the TE of the spin-echo from which the data are collected matches the effective TE. Shallow slopes, which allow maximum signal, are centered on spin-echoes that have a TE at or near the effective TE. Steep slopes, which allow less signal, are placed on spin-echoes that have a TE away from the effective TE. The resultant image contains data from all spin-echoes in the echo train, but data from spin-echoes collected around the effective TE have more impact on image weighting, as they fill the central lines of k -space that contribute signal and contrast to the image. Data from spin-echoes collected at the “wrong” weighting (other TEs) have much less of an effect, as they fill the outer lines of k -space and therefore make less contribution to the signal and contrast of the image (Figure 3.12).

For example, if T2 weighting is required, an effective TE of 100 ms, a TR of 4000 ms, and a turbo factor of 16 might be used. During the sequence, the shallowest phase encodings are performed on spin-echoes occurring around 100 ms. Data acquired from these echoes have a TE at or close to 100 ms. Phase encodings performed at the very beginning and end of the echo train are steep, and the signal amplitude of these echoes is therefore small. These echoes contain either PD- or very heavily T2-weighted data. Although data from these early and late spin-echoes contribute to image contrast, their impact is significantly less than from spin-echoes whose TE is at or around the effective TE. Consequently, one image is obtained for each slice location that is mainly T2-weighted.




Figure 3.10 Phase-encoding gradient slopes.




Figure 3.11 Phase encoding vs signal amplitude.

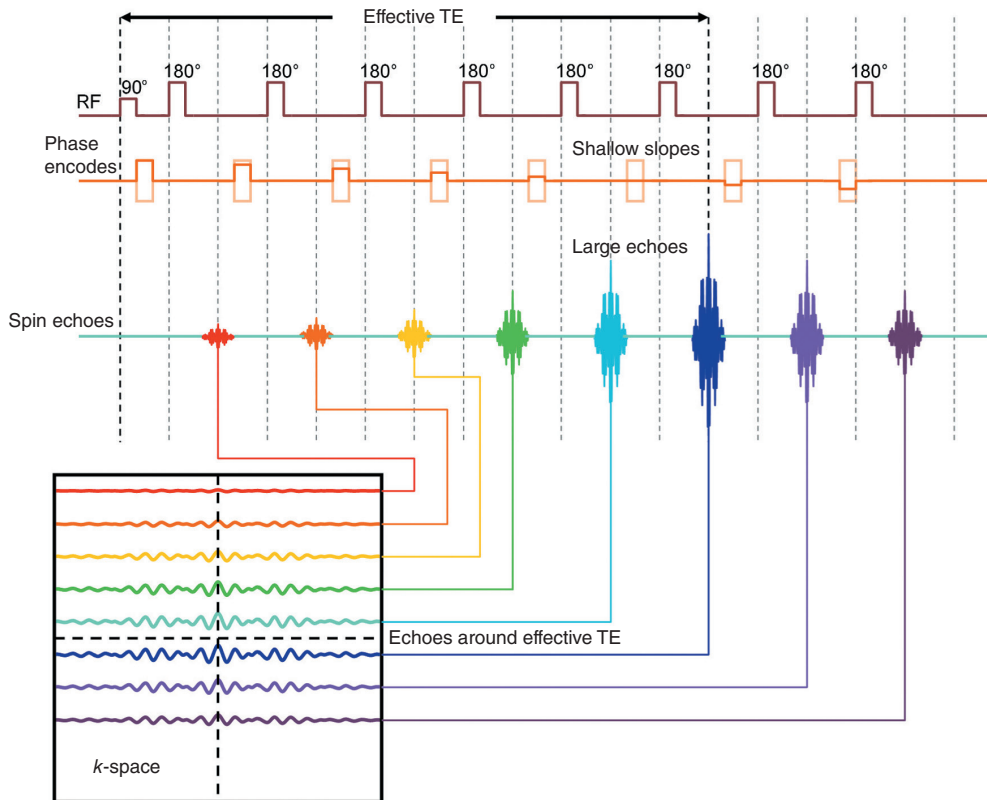


Figure 3.12 *k*-Space filling and phase reordering.

Learning tip: Phase-reordering

72

The amplitude and polarity of the phase-encoding gradient are varied by a process called **phase reordering**. In TSE, this is performed to match the effective TE with the required amplitude of phase-encoding gradient. The resultant images rely on a weighted average of all the data collected in *k*-space. This means that in terms of image contrast, more emphasis is placed on data acquired at or around the effective TE than on data acquired at other TEs [3,4]. Phase reordering is also carried out in some techniques used to reduce motion artifacts (see Chapter 8).

Uses

Contrast in TSE images is similar to spin-echo, and therefore these sequences are useful in most clinical applications. There are, however, two contrast differences between conventional spin-echo and TSE, both of which are due to the repeated, closely spaced 180° RF rephasing pulses in the echo train. Firstly, fat remains bright on T2-weighted images due to the multiple 180° RF rephasing pulses that reduce the effects of spin-spin interactions in fat (**J coupling**) (Figure 3.13). However, fat saturation techniques can be used to compensate for this (see Chapter 7). Secondly, repeated 180° RF rephasing pulses might increase magnetization transfer effects so that muscle, for example, appears darker on TSE images than in conventional spin-echo (see Chapter 2). In addition, multiple 180° RF rephasing pulses reduce magnetic susceptibility effects (see Chapter 8), which can be detrimental when looking for small hemorrhages.



Figure 3.13 Sagittal TSE T2-weighted image of a female pelvis. Source: Westbrook 2014 [5]. Reproduced with permission of John Wiley & Sons.

Table 3.4 Advantages and disadvantages of TSE.

Advantages	Disadvantages
Short scan times	Some flow artifacts increased
High-resolution imaging	Incompatible with some imaging options
Increased T2 weighting	Some contrast interpretation problems
Magnetic susceptibility decreases ^a	Image blurring possible

^aAlso a disadvantage, e.g. hemorrhage not detected/delineated (see Chapter 8).

Image blurring may occur in TSE images at the edges of tissues with different T2 decay times. This is because each line of k -space that is filled during the echo train contains data from echoes with a different TE. When using long echo trains, late echoes that have a low signal amplitude contribute to the resolution of k -space. If these echoes are negligible, then resolution is lost, and blurring may occur. This, however, is reduced by decreasing the spacing between echoes (known as the **echo-spacing**) and/or the turbo factor. In addition, artifact from metal implants is significantly reduced when using TSE because repeated 180° RF rephasing pulses compensate for magnetic field inhomogeneity (Table 3.4) (see Chapter 8).

Suggested parameters

These are similar to conventional spin-echo. However, the turbo factor now plays an important role in image weighting. The higher the turbo factor, the shorter the scan time, but the resultant image has more of a mixture of weighting because there are more data collected at the “wrong” TE. This is not as important in T2-weighted scans as proton density data are offset somewhat by heavily T2-weighted data. In T1 and PD weighting, on the other hand, longer turbo factors place too much T2 weighting in the image, and hence shorter turbo factors are used. Scan time savings in T1-weighted imaging are therefore not as great as they are in T2 weighting.

If proton density and T2 weighting are required together in the same sequence, there are three possible ways of achieving this.

- *Full echo train.* The full echo train is used to acquire data for the PD-weighted image, and then the whole scan is repeated using the full echo train to acquire the T2-weighted image.
- *Split echo train.* The echo train is split, with data from the first half used to acquire the PD-weighted image and the second half the T2-weighted image.
- *Shared echo train.* Data from the first spin-echoes in the echo train are used to acquire the PD-weighted image, and data from the last spin-echoes for the T2-weighted image. Data from the other spin-echoes are shared for both images. This technique permits shorter turbo factors and therefore more slices per TR than the other two techniques (see below).

Learning tip: TSE – the first echoes are not “free”

In an earlier learning tip, we saw that in multiple-echo conventional spin-echo sequences, the first echoes are free. It does not cost anything in terms of scan time to acquire them. In TSE, these echoes are not free because data from every single echo fill a line of k -space, and each image has its own k -space area. Therefore, if a dual echo sequence is required, data for the PD-weighted image fills one k -space area, and a completely different k -space area is filled for the T2-weighted image. This means that the scan time doubles.

For T1 weighting (Figure 3.14):

- TR 300–700 ms
- Effective TE minimum
- Turbo factor 2–8.

For proton density weighting (Figure 3.15):

- TR 3000–10 000 ms (depending on required slice number)
- Effective TE minimum
- Turbo factor 4–12.

For T2 weighting (Figure 3.13):

- TR 3000–10 000 ms (depending on required slice number)
- Effective TE 80–140 ms
- Turbo factor 12–30.



Figure 3.14 Axial T1-weighted TSE image of a male pelvis. Source: Westbrook 2015 [6]. Reproduced with permission of John Wiley & Sons.




Figure 3.15 Coronal TSE PD-weighted image of a knee. Source: Westbrook 2014 [5]. Reproduced with permission of John Wiley & Sons.

Scan tip:

Parameter selection in turbo spin-echo – what’s going on behind the scenes?

In TSE, selection of the TR and TE follows the same principles as conventional spin-echo. TR controls T1 contrast, and TE controls T2 contrast. When we select the turbo factor or ETL in the scan protocol, behind the scenes, we determine how many lines of k -space are filled every TR. We therefore determine how many times the RF rephasing pulse is applied every TR and how many times the phase-encoding gradient is switched to a different amplitude every TR. Selection of a long turbo factor means many of lines of k -space are filled every TR so the scan time decreases. Long turbo factors also lead to more T2 contrast. The opposite is true of short turbo factors.

Scan tip: TR selection in TSE

The TR used in TSE is often much longer than that used in conventional spin-echo. The 180° RF rephasing pulses take time to perform, and so fewer slices are available for a given TR. The system cannot go on to excite the next slice until it has acquired data from the current slice. The longer it spends “sitting” at each slice waiting to collect data, the fewer slices it can acquire in a TR period. In conventional spin-echo, the time the system spends waiting to collect data is the TE. Longer TEs result in fewer slices per TR. In TSE, this time depends firstly on how many echoes there are (ETL or turbo factor). The longer the turbo factor, the longer the system spends at each slice collecting data and so fewer slices are permitted in a TR period. Echo spacing is also important. This is the time between each echo and contributes to how long it takes to complete the entire echo train. If the echo spacing is long then an echo train of a given value takes longer than if the echo spacing is short.

If we select a long turbo factor and/or echo spacing in the scan protocol, the TR is extended to acquire a given number of slices, and this increases the scan time. For example, to acquire 20, T2-weighted slices with a long TE in conventional spin-echo, a TR of about 2000 ms might be needed. In TSE, a TR of 4000 ms or longer might be necessary to acquire the same 20 slices (Equation (3.2)). The longer TRs associated with TSE are relevant in T2 weighting but are far less significant than the huge scan time savings produced by a long echo train. In T1 weighting, however, increasing the TR is not a good strategy, as this decreases T1 contrast. Therefore, in T1 weighting, the TR is kept short, and several acquisitions are often necessary to cover anatomy.

Equation 3.2

$$N \text{ slices} = \frac{\text{TR}}{\text{ETL} \times \text{E}_s}$$

N slices is the number of slices allowed per TR
TR is the repetition time (ms)
ETL is the echo train length or turbo factor
ES is the echo spacing (ms)

This equation shows how many slices are allowed in TSE and is less than in conventional spin-echo

Table 3.5 Things to remember – turbo spin-echo.

Turbo or fast spin-echo sequences involve applying the phase-encoding gradient multiple times in a TR period to varying amplitudes and polarity
This means that multiple lines of k -space are selected per TR. The number is equal to the echo train length (ETL) or turbo factor
Multiple echoes are produced by multiple applications of an RF rephasing pulse, and data from each echo are placed in a different line of k -space
Scan times are reduced by a factor equal to the turbo factor or ETL
Image weighting is controlled by phase reordering so that data collected from echoes at or around the effective TE are placed in the signal and contrast areas of k -space

Single-shot turbo spin-echo (SS-TSE)

It is possible to acquire TSE images in even shorter scan times by using a technique known as **single-shot fast** or **turbo spin-echo (SS-TSE)**. In this technique, all the lines of k -space are acquired at once. SS-TSE combines a partial Fourier technique with TSE (see Chapter 6). Half the lines of k -space are acquired in one TR, and the other half are transposed. This technique yields a reduction in imaging time, as all the image data are acquired at once.

77

Driven equilibrium

In a modification of TSE called **driven equilibrium (DRIVE)**, a reverse flip angle excitation RF pulse is applied at the end of the echo train. This drives any transverse magnetization into the longitudinal plane so that it is available for excitation at the beginning of the next TR period. Therefore, it is not necessary to wait long periods for T1 relaxation to occur. Some manufacturers rephase the transverse magnetization with a 180° RF rephasing pulse before the restoration 90° pulse is applied. As water has the longest T1 and T2 relaxation times, most of this magnetization is composed of water, and therefore this has a higher signal intensity. This sequence produces increased signal intensity in fluid-based structures such as CSF when using a shorter TR than normal in TSE (Figures 3.16 and 3.17).




Figure 3.16 The DRIVE pulse sequence.




Figure 3.17 Axial DRIVE image through the right internal auditory meatus. Note high signal intensity in the CSF.

Learning tip: Different rephasing angles in TSE

In Chapter 1, we learned that an RF pulse is a pulse of electromagnetic radiation that consists of a magnetic and electric field oscillating in waves at 90° to each other. The oscillating magnetic field causes resonance. The oscillating electric field produces heat and causes a rise in the specific absorption rate (SAR; see Chapter 10). In TSE, the SAR increases significantly because multiple RF rephasing pulses are applied in quick succession. This usually manifests itself in a reduction in the number of permissible slices and it can therefore be difficult to get the required coverage in a single acquisition. To reduce the SAR, the rephasing angle of the RF rephasing pulses may be reduced. For example, instead of 180° , a rephasing angle of 150° or 120° is used. Be careful: some manufacturers refer to this as the flip angle; however, it is the rephasing pulse angle rather than the excitation pulse angle). In Chapter 4, the characteristics of RF pulses are discussed in more detail.

INVERSION RECOVERY (IR)

Mechanism

Inversion recovery (IR) is a spin-echo pulse sequence that uses an RF inverting pulse to suppress signal from certain tissues, although it is also sometimes used to generate heavy T1 contrast. The IR pulse sequence begins with a 180° RF pulse. This is applied at the beginning of the TR

period when the NMV is aligned in the same direction as B_0 in the longitudinal plane (termed $+z$). The RF pulse inverts the NMV through 180° (Figure 3.18), which means that after the pulse, the NMV still lies in the longitudinal plane but in the opposite direction to B_0 (termed $-z$). When the RF inverting pulse is removed, the NMV relaxes back to B_0 because of T_1 recovery processes. At a certain time-point during this recovery, a 90° RF excitation pulse is applied and then switched off. The resultant FID is then rephased by another 180° RF rephasing pulse to produce a spin-echo at time TE (Figure 3.19).




Figure 3.18 The 180° inverting pulse in an IR pulse sequence.




Figure 3.19 The inversion recovery sequence.

The time from the 180° RF inverting pulse to the 90° RF excitation pulse is known as the **TI (time from inversion)**. Image contrast depends primarily on the TI. If the 90° RF excitation pulse is applied after the NMV has relaxed back through the transverse plane, image contrast depends on the amount of longitudinal recovery of each vector (as in spin-echo). The resultant image is T1-weighted (Figure 3.20). If the 90° RF excitation pulse is not applied until the NMV has reached full recovery, a PD-weighted image is produced, as both fat and water are fully relaxed (Figure 3.21).




Figure 3.20 T1 weighting in inversion recovery.




Figure 3.21 PD weighting in inversion recovery.

Uses

IR is conventionally used to produce heavily T1-weighted images to demonstrate anatomy (Figure 3.22). The 180° RF inverting pulse produces a large contrast difference between fat and water because full saturation of the fat and water vectors is achieved at the beginning of each TR. Tissues begin their recovery from full inversion as opposed to from the transverse plane as in conventional spin-echo. This allows more time for differences in the T1 recovery times between tissues to become apparent, and therefore IR pulse sequences produce heavier T1 weighting than conventional spin-echo. In addition, as gadolinium primarily shortens the T1 recovery times of certain tissues (see Chapter 2) IR pulse sequences increase signal from contrast-enhanced structures.

81

Suggested parameters

The TI is the most potent controller of contrast in the IR sequence. Medium TI values produce T1 weighting but, as the TI is lengthened, the image becomes more PD-weighted. The TR should always be long enough to allow full recovery of the NMV before application of the next




Figure 3.22 Axial T1-weighted inversion recovery sequence of the brain. AT1 of 700 ms was used.

RF inverting pulse. If this is not so, individual vectors recover to different degrees, and weighting is affected. For example, to achieve full recovery of the NMV at 1 T, the TR should be longer than 3000 ms. Most systems use IR in conjunction with TSE (see below).

When IR is used to produce predominantly heavily T1-weighted images, the TE controls the amount of T2 decay, and so it is usually kept short to minimize T2 contrast. However, it can be lengthened to give tissues with a long T2 decay time a high signal. This is called **pathology weighting** and produces an image that is predominantly T1-weighted but where pathology is bright.



Refer to animation 3.3 on the supporting companion website for this book: www.wiley.com/go/westbrook/mriinpractice

T1 weighting:

- Medium TI 400–800 ms (varies at different field strengths)
- Short TE 10–20 ms
- Long TR 3000 ms+ .

Proton density weighting:

- Long TI 1800 ms
- Short TE 10–20 ms
- Long TR 3000 ms+ .

Pathology weighting:

- Medium TI 400–800 ms
- Long TE 70 ms+
- Long TR 3000 ms+ .

Fast inversion recovery

In this sequence, the 180° RF inverting pulse is followed at time TI by the 90° RF excitation pulse and a train of 180° RF rephasing pulses to fill out multiple lines of *k*-space as in TSE. This reduces the scan time compared to conventional IR. However, instead of T1-weighted images, fast IR is usually used to suppress signal from certain tissues in conjunction with T2 weighting so that water and pathology return a high signal. The two main sequences in this category are **STIR** and **FLAIR**.

SHORT TAU INVERSION RECOVERY (STIR)

Mechanism

STIR is an IR pulse sequence that uses a TI that corresponds to the time it takes the fat vector to recover from full inversion to the transverse plane so that there is no longitudinal magnetization corresponding to fat. This is called the **null point** (Figure 3.23). As there is no longitudinal




Figure 3.23 The STIR sequence.

component of fat when the 90° RF excitation pulse is applied, there is no transverse component after excitation, and signal from fat is nulled. A TI of 100–175 ms usually achieves fat suppression, although this value varies slightly at different field strengths. The TI required to null signal from a tissue is always 0.69 times its T1 relaxation time (Equation (3.3)).

Equation 3.3

$$SI = 1 - 2e^{-t/T1}$$

SI is the signal intensity in a tissue
t is the TI (ms)
T1 is the T1 relaxation time in the tissue (ms)

This equation represents T1 recovery when the NMV is fully inverted. The SI = zero when $t = 0.693 \times T1$. This is the TI required to null signal from a tissue [7]

Uses

STIR is an extremely important sequence in musculoskeletal imaging because normal bone, which contains fatty marrow, is suppressed, and lesions within bone such as bone bruising and tumors are seen more clearly (Figure 3.24). It is also a very useful sequence for suppressing fat in general MR imaging (see Chapter 7).

Scan tip: When not to use STIR

STIR should not be used in conjunction with contrast enhancement, which shortens the T1 recovery times of enhancing tissues, making them relatively hyperintense (see Chapter 2). The T1 recovery times of these structures are shortened by the contrast agent so that they approach the T1 recovery time of fat. In a STIR sequence, therefore, enhancing tissue may also be nulled.




Figure 3.24 Coronal STIR of the knee. Source: Westbrook 2014 [5]. Reproduced with permission of John Wiley & Sons.

Suggested parameters

- Short TI (tau) 150–175 ms (to suppress fat depending on field strength)
- Long TE 50 ms+ (to enhance signal from pathology)
- Long TR 4000 ms+ (to allow full longitudinal recovery)
- Long turbo factor 16–20 (to enhance signal from pathology).

FLUID ATTENUATED INVERSION RECOVERY (FLAIR)

Mechanism

FLAIR is another variation of the IR sequence. In FLAIR, a TI corresponding to the recovery of the vector in CSF from full inversion to the transverse plane is selected. This TI nulls signal from CSF because there is no longitudinal magnetization present in CSF. As there is no longitudinal component of CSF when the 90° RF excitation pulse is applied, there is no transverse component after excitation, and signal from CSF is nulled. FLAIR is used to suppress high CSF signal in T2 weighted images so that the pathology adjacent to CSF is seen more clearly (Figure 3.25). A TI of 1700–2200 ms usually achieves CSF suppression (although this varies slightly at different field strengths and is calculated by multiplying the T1 relaxation time of CSF by 0.69).




Figure 3.25 Axial T2-weighted FLAIR of the brain. Source: Westbrook 2014 [5]. Reproduced with permission of John Wiley & Sons.

Uses

FLAIR is used in brain and spine imaging to see periventricular and cord lesions more clearly because high signal from CSF that lies adjacent is nulled. It is especially useful in visualizing multiple sclerosis plaques, acute subarachnoid hemorrhage, and meningitis. Another modification of this sequence in brain imaging is selecting a TI that corresponds to the null point of white matter. This TI value nulls signal from normal white matter so that lesions within it appear much brighter by comparison. This sequence (which requires a TI of about 300 ms) is very useful for white matter lesions such as periventricular leukomalacia and for congenital gray/white matter abnormalities [7] (Figure 3.26). The advantages and disadvantages of IR sequences are summarized in Table 3.6.

Table 3.6 Advantages and disadvantages of inversion recovery.

Advantages	Disadvantage
Versatile	Long scan times (conventional IR)
Good image quality	
Sensitive to pathology	




Figure 3.26 Coronal IR-TSE T2-weighted image with a TI selected to null white matter. Source: Westbrook 2014 [5]. Reproduced with permission of John Wiley & Sons.

Learning tip: FLAIR and gadolinium

Sometimes gadolinium is given to enhance pathology in the FLAIR sequence. This oddity (gadolinium enhancement in T2-weighted images) may be due to the long echo trains used in FLAIR sequences that cause fat to remain bright on T2-weighted images because of J-coupling. As gadolinium reduces the T1 recovery time of enhancing tissue so that it is similar to fat, enhancing tissue may appear brighter than when gadolinium is not given.

Suggested parameters

- Long TI 1700–2200 ms (to suppress CSF depending on field strength)
- Long TE 70 ms+ (to enhance signal from pathology)
- Long TR 6000 ms+ (to allow full longitudinal recovery)
- Long turbo factor 16–20 (to enhance signal from pathology).

Scan tip: Parameter selection in IR – what's going on behind the scenes?

When we select the TE in the scan protocol, the same principles as in conventional spin-echo are applicable, i.e. TE controls T2 contrast. The TR does not control T1 contrast in IR because it is always long. A long TR ensures that there is full longitudinal recovery of all the vectors in-between RF inverting pulses. When we select the TI in IR sequences in the scan protocol, behind the scenes, we determine how long after the inverting pulse the 90° RF excitation pulse is applied. We therefore control which tissues are nulled.

IR prep sequences

There are two further modifications of fast IR that were specifically developed to null blood in cardiac imaging. **Double IR prep** begins with two 180° RF inverting pulses. One is non-slice selective and inverts all spins in the imaging volume, and the other is slice selective and re-inverts spins within a slice. A TI corresponding to the null point of blood (about 800 ms) completely nulls signal from blood in the slice so that black blood imaging results. This is useful when looking at the morphology of the heart and great vessels. **Triple IR prep** adds a further inverting pulse at the TI of fat (about 150 ms) to null fat and blood together. This is useful when determining fatty infiltration of the heart walls.

Learning tip: Image reconstruction in IR

Images can be reconstructed based on either the magnitude of signal or its phase. Signal intensity in magnitude images depends only on signal amplitude. Tissues with either very long or very short T1 recovery times are assigned high pixel values. Phase images also include magnitude data but, in addition, incorporate the contribution made by the inverting pulse. If the TI is short, then tissues with a long T1 recovery time are fully inverted by the 180° RF inverting pulse. A low pixel intensity is assigned to these tissues [8]. The inclusion of these phase-sensitive data extends the contrast range of the image [9].

Table 3.7 Things to remember – inversion recovery.

IR sequences are spin-echo sequences with an initial 180° RF inversion pulse that saturates the longitudinal magnetization
The TI determines the time between this and the 90° RF excitation pulse that follows. This parameter controls weighting in conventional and fast versions of this sequence
When combined with an echo train as in TSE, scan times are reduced, and this sequence is typically used to suppress the signal from either CSF (FLAIR) or fat (STIR)
A TI equivalent to the null point of either CSF or fat is used for this purpose. The ETL determines by how much the scan time is reduced

Table 3.8 summarizes what is going on behind the scenes in spin-echo pulse sequences.

Table 3.8 Spin-echo pulse sequences – summary of what is going on behind the scenes.

Parameter	Behind the scenes
TR	Controls the amount of T1 recovery and therefore T1 contrast
TE	Controls the amount of T2 decay and therefore T2 contrast. The 180° RF rephasing pulse is applied at half the selected TE (τ)
TI	Controls the point at which longitudinal magnetization is nulled
ETL/TF	Determines the number of k -space lines filled every TR. It therefore controls how many RF rephasing pulses are applied every TR and how many times the phase-encoding gradient is applied to different amplitudes every TR

Spin-echo pulse sequences are very commonly used in clinical imaging, but there are circumstances where their use might not be optimal. In the next chapter, we explore the principles behind the other family of pulse sequences – gradient-echo.



For questions and answers on this topic please visit the supporting companion website for this book: www.wiley.com/go/westbrook/mriinpractice

References

1. Hashemi, R.H., Bradley Jr, W.G., and Lisanti, C.J. (2010). *MRI: The Basics*, 3, 83. Philadelphia, PA: Lippincott Williams and Wilkins.
2. Dale, B.M., Brown, M.A., and Semelka, R.C. (2015). *MRI: Basic Principles and Applications*, 5, 67. Wiley.
3. Henning, J., Naureth, A., and Friedburg, H. (1986). RARE imaging: a fat imaging method for clinical MR. *Magnetic Resonance in Medicine* 3: 828.
4. Hashemi, R.H., Bradley Jr, W.G., and Lisanti, C.J. (2010). *MRI: The Basics*, 3, 221. Philadelphia, PA: Lippincott Williams and Wilkins.
5. Westbrook, C. (2014). *Handbook of MRI Technique*, 4, 2014. Wiley Blackwell.
6. Westbrook, C. (2015). *MRI at a Glance*, 3. Wiley Blackwell.
7. Westbrook, C. (2014). *Handbook of MRI Technique*, 4, 68–69. Wiley Blackwell.
8. Dale, B.M., Brown, M.A., and Semelka, R.C. (2015). *MRI: Basic Principles and Applications*, 5, 80. Wiley.
9. Liney, G. (2010). *MRI from A to Z*, 2, 171. London: Springer.

4

Gradient-echo pulse sequences

Introduction	89	Incoherent or spoiled gradient-echo	109
Variable flip angle	90	Reverse-echo gradient-echo	113
Gradient rephasing	91	Balanced gradient-echo	119
Weighting in gradient-echo pulse sequences	94	Fast gradient-echo	122
Coherent or rewind gradient-echo	106	Echo planar imaging	122

After reading this chapter, you will be able to:

- *Explain how gradient-echo sequences differ from spin-echo.*
- *Describe how gradient-echoes are created.*
- *Analyze the steady state and why it is important in gradient-echo pulse sequences.*
- *Understand the mechanisms of common gradient-echo pulse sequences.*
- *Apply what you have learned to understand how images of different weighting are created using gradient-echo pulse sequences.*

INTRODUCTION

This chapter discusses the mechanisms, uses, and parameters for each of the common gradient-echo pulse sequences and their advantages and disadvantages. A table is included that compares the common acronyms for gradient-echo pulse sequences for the main manufacturers (Table 4.1). A more comprehensive table is also provided at the beginning of the book. As with Table 3.1 and the acronym table that compares acronyms for spin-echo sequences, these tables are only a guide; they are not meant to compare the performance or specification of each system. The parameters included in this chapter depend on field strength and the nuances of individual systems. However, they should be suitable for most field strengths used in clinical imaging.

Table 4.1 Gradient-echo pulse sequences and their common acronyms.

Generic	GE	Philips	Siemens	Toshiba	Hitachi
Coherent or rewind gradient-echo	GRASS	FFE	FISP	SSFP	Rephased SARGE
Incoherent or spoiled gradient-echo	SPGR	T1-FFE	FLASH	Fast FE	RF spoiled SARGE
Reverse-echo gradient-echo	SSFP	T2-FFE	PSIF	No sequence	Time-reversed SARGE
Balanced gradient-echo	FIESTA	B-FFE	True FISP	True SSFP	Balanced SARGE
Fast gradient-echo	Fast GRASS or SPGR	TFE	Turbo FLASH	Fast FE	RGE
Echo planar imaging	EPI	EPI	EPI	EPI	EPI

Abbreviations used in Table 4.1

GRASS	gradient recalled acquisition in the steady state	FLASH	fast low angled shot
SPGR	spoiled GRASS	PSIF	reverse FISP
SSFP	steady state free precession	EPI	echo planar imaging
FIESTA	free induction echo stimulated acquisition	RGE	rapid gradient-echo
FFE	fast field echo	SARGE	steady acquisition rewind gradient-echo
FISP	fast imaging with steady precession		

Gradient-echo pulse sequences differ from spin-echo pulse sequences in two ways:

- They use *variable RF excitation pulse flip angles* as opposed to 90° RF excitation pulse flip angles that are common in spin-echo pulse sequences.
- They use *gradients* rather than RF pulses to rephase the magnetic moments of hydrogen nuclei to form an echo.

The main purpose of these two mechanisms is to enable shorter TRs and therefore scan times than are common with spin-echo pulse sequences. Let's explore these strategies in more detail.

VARIABLE FLIP ANGLE

A gradient-echo pulse sequence uses an RF excitation pulse that is variable and therefore flips the NMV through any angle (not just 90°). Typically, a flip angle of less than 90° is used. This means that the NMV is flipped through a lower angle than it is in spin-echo sequences when a larger 90° flip angle is usually applied. As the NMV is moved through a smaller angle in the excitation phase of the pulse sequence, it does not take as long for the NMV to achieve full relaxation once the RF excitation pulse is removed. Therefore, full T1 recovery is achieved in a much shorter TR than in spin-echo pulse sequences. As the TR is a scan time parameter (see Equations (6.7)–(6.9)), this leads to shorter scan times.

GRADIENT REPHASING

After the RF excitation pulse is withdrawn, the FID immediately occurs due to inhomogeneities in the magnetic field and T_2^* decay. In spin-echo pulse sequences, the magnetic moments of hydrogen nuclei are rephased by an RF pulse. As a relatively large flip angle is used in spin-echo pulse sequences, most of the magnetization is still in the transverse plane when the 180° RF rephasing pulse is applied. Consequently, this pulse rephases this transverse magnetization to create a spin-echo. In gradient-echo pulse sequences, an RF pulse cannot rephase transverse magnetization to create an echo. The low flip angles used in gradient-echo pulse sequences result in a large component of magnetization remaining in the longitudinal plane after the RF excitation pulse is switched off. The 180° RF pulse would therefore largely invert this magnetization into the $-z$ direction (the direction that is opposite to B_0) rather than rephase the transverse magnetization [1]. Therefore in gradient-echo pulse sequences, a **gradient** is used to rephase transverse magnetization instead.

Gradients perform many tasks, which are explored fully in Chapter 5. In this chapter, they are discussed specifically in relation to how they are used to rephase or dephase the magnetic moments of hydrogen nuclei.

How gradients dephase

Look at Figure 4.1. With no gradient applied, all the magnetic moments of hydrogen nuclei precess at the same frequency, as they experience the same field strength (in reality they do not because of magnetic field inhomogeneities, but these changes are relatively small compared with those imposed by a gradient). A gradient is applied to coherent (in phase) magnetization (all the magnetic moments are in the same place at the same time). The gradient alters the magnetic field strength experienced by the coherent magnetization. Some of the magnetic moments speed up, and some slow down, depending on their position along the gradient axis. Thus, the magnetic moments fan out or dephase because their frequencies are changed by the gradient (see the watch analogy in Chapter 1).

The trailing edge of the fan (shown in purple) consists of nuclei whose magnetic moments slow down because they are situated on the gradient axis that has a lower magnetic field strength relative to isocenter. The leading edge of the fan (shown in red) consists of nuclei whose magnetic moments speed up because they are situated on the gradient axis that has a higher magnetic field strength relative to isocenter. The magnetic moments of nuclei are therefore no longer in the same place at the same time, and so magnetization is dephased by the gradient. Gradients that dephase in this way are called **spoilers**, and the process of dephasing magnetic moments with gradients is called **gradient spoiling**.

How gradients rephase

Look at Figure 4.2. A gradient is applied to incoherent (out of phase) magnetization to rephase it. The magnetic moments initially fan out due to T_2^* decay, and the fan has a trailing edge consisting of nuclei with slowly precessing magnetic moments (shown in purple) and a leading edge consisting of nuclei with faster precessing magnetic moments (shown in red). A gradient is then applied so that the magnetic field strength is altered in a linear fashion along the axis of the gradient. The direction of this altered field strength is such that the slowly precessing magnetic moments in the trailing edge of the fan experience an increased magnetic field strength




Figure 4.1 How gradients dephase.




Figure 4.2 How gradients rephase.

and speed up. In Figure 4.2, these are the purple spins that experience the red “high end” of the gradient. At the same time, the faster precessing magnetic moments in the leading edge of the fan experience a decreased magnetic field strength and slow down. In Figure 4.2, these are the red magnetic moments that experience the purple “low end” of the gradient. After a short period of time, the slow magnetic moments speed up sufficiently to meet the faster ones that are slowing down. At this point, all the magnetic moments are in the same place at the same time and are therefore rephased by the gradient. A maximum signal is induced in the receiver coil, and this signal is called a **gradient-echo**. Gradients that rephase in this way are called **rewinders**.

Whether a gradient field adds or subtracts from the main magnetic field depends on the direction of current that passes through the gradient coils. This is called the **polarity** of the gradient. Gradient-echoes are created by a **bipolar gradient**. This means that it consists of two lobes, one negative and one positive. The frequency-encoding gradient is used for this purpose (see Chapter 5). It is initially applied negatively, which increases dephasing and eliminates the FID. Its polarity is then reversed, which rephases only those magnetic moments that were dephased by the negative lobe. It is only these nuclei (those whose magnetic moments are dephased by the negative lobe of the gradient and are then rephased by the positive lobe) that create the gradient-echo at time TE (Figure 4.3). The area under the negative lobe of the gradient is half that of the area under the positive lobe [2].




Figure 4.3 A basic gradient-echo sequence showing how a bipolar application of the frequency-encoding gradient produces a gradient-echo.

Learning tip: The good and bad of gradient-echo pulse sequences

Gradients rephase the magnetic moments of hydrogen nuclei much faster than RF pulses, and therefore echoes are generated faster than in spin-echo pulse sequences. The TEs are therefore shorter than in spin-echo. The TE is not part of the scan time equation, but, as we saw in Chapter 3, the TE determines how long we sit at each slice waiting for an echo. When the TE is short, a given number of slices are acquired in a short TR, and, therefore, the scan time is shorter than that of the spin-echo pulse sequences.

However, in gradient-echo sequences, there is no compensation for magnetic field inhomogeneities. Gradient rephasing does not remove the contribution made by T2* decay processes. This is because the rephasing lobe of the bipolar gradient only affects the magnetic moments that are dephased by the dephasing lobe of the gradient. Magnetic moments dephased due to magnetic field inhomogeneities are not affected. Gradient-echo sequences are therefore very susceptible to certain artifacts that rely on magnetic field inhomogeneities such as magnetic susceptibility (see Chapter 8). They are also heavily reliant on T2* relaxation processes (see Chapter 2). As a result, in gradient-echo pulse sequences, T2 weighting is termed **T2* weighting**, and T2 decay is termed T2* decay to reflect the contribution made by magnetic field inhomogeneities to image contrast.

WEIGHTING IN GRADIENT-ECHO PULSE SEQUENCES

The weighting mechanisms in gradient-echo pulse sequences are quite complex, and this is one of the many reasons why they are hard to understand (and explain!). There are essentially three different processes that affect weighting in gradient-echo pulse sequences, and sometimes all three overlay each other in the image. These are as follows:

- Extrinsic parameters (TR, TE, and flip angle)
- The steady state
- Residual transverse magnetization.

Let's explore each of these processes in detail.

Table 4.2 Things to remember – gradient-echo pulse sequences.

Gradient-echo sequences use gradients to rephase the magnetic moments of hydrogen nuclei and usually flip angles less than 90°. Both of these strategies permit a shorter TE and TR than in spin-echo pulse sequences

Low flip angles mean that, as less longitudinal magnetization is converted to transverse magnetization during the excitation phase of the sequence, less time is required for relaxation. This is why a short TR can be used

The speed of rephasing is increased using a gradient. A bipolar application of the frequency-encoding gradient enables magnetic moments of the hydrogen nuclei to rephase faster than when using an RF rephasing pulse. This permits a short TE, which means that a shorter TR can be used for a given number of slices than in spin-echo

Although faster than RF rephasing, inhomogeneities are not compensated for in this type of sequence. Magnetic susceptibility artifacts therefore increase

Weighting mechanism 1 – extrinsic contrast parameters

The influence of the TR and TE on image weighting was explored in Chapter 2. The discussion assumed that the RF excitation pulse flip angle is 90° (as in spin-echo pulse sequences). Under these circumstances, TE controls T2 contrast, and T2 contrast increases as the TE increases. The same is true in gradient-echo pulse sequences except that T2 is termed T2* to reflect the fact that magnetic field inhomogeneities are not compensated for by gradient rephasing. In spin-echo pulse sequences, TR controls T1 contrast, and T1 contrast increases as the TR decreases. This is because short TRs do not permit complete recovery of the vectors, and, therefore, subsequent 90° RF excitation pulses cause saturation. In gradient-echo pulse sequences, the TR *and* the flip angle control the amount of T1 relaxation and saturation that occurs (Equation (4.1)).

95

Equation 4.1

$$SI = PD e^{-TE/T2^*} (1 - e^{-TR/T1})$$

$$[\sin \theta / (1 - \cos \theta e^{-TR/T1})]$$

SI is the signal intensity in a tissue
 PD is the proton density
 TE is the echo time (ms)
 T2* is the T2* relaxation time of the tissue (ms)
 TR is the repetition time (ms)
 T1 is the T1 relaxation time in the tissue (ms)
 θ is the flip angle
 $[\sin \theta / (1 - \cos \theta e^{-TR/T1})]$ is the flip angle function

This equation shows why the signal intensity from a tissue depends on intrinsic and extrinsic contrast parameters. Compare this equation with Equation (2.4). The flip angle function is added, and T2 becomes T2*. The flip angle function shows how the flip angle, TR, and T1 relaxation time all determine whether a tissue is saturated. If $\alpha=0^\circ$ or 90° , then $\sin \alpha=1$ and $\cos \alpha=0$. This equation is then identical to Equation (2.4) [8]

Apart from the added variable of the flip angle, weighting rules in gradient-echo are the same as in spin-echo (see the heat analogy in Chapter 2).

The trick is to imagine how far the vectors are flipped by the RF excitation pulse (flip angle) and then how long they are given to recover their longitudinal magnetization (TR).

- If the combination of flip angle and TR causes saturation of the vectors (i.e. they never fully recover their longitudinal magnetization during the TR period), then T1 contrast is maximized.
- If the combination of flip angle and TR *does not* cause saturation of the vectors (i.e. they recover most, or all, of their longitudinal magnetization during the TR period), then T1 contrast is minimized.

These rules, along with those of how the TE controls T2* contrast, are used to weight images in gradient-echo pulse sequences.

Using extrinsic contrast parameters in gradient-echo – T1 weighting

To obtain a T1-weighted image, differences in the T1 recovery times of the tissues are maximized, and differences in the T2* decay times of the tissues are minimized. To maximize differences in T1 recovery times, neither fat nor water vectors are given time to recover full longitudinal

magnetization before the next RF excitation pulse is applied. To avoid full recovery of their longitudinal magnetization, the flip angle is large and the TR short so that the fat and water vectors are still in the process of recovering when the next RF excitation pulse is applied. To minimize differences in $T2^*$ decay times, the TE is short so that neither fat nor water has time to decay (Figure 4.4).

Using extrinsic contrast parameters in gradient-echo – $T2^*$ weighting

To obtain a $T2^*$ -weighted image, differences in the $T2^*$ decay times of the tissues are maximized, and differences in the $T1$ recovery times are minimized. To maximize differences in $T2^*$ decay times, the TE is long so that fat and water vectors have had time to dephase. To minimize differences in $T1$ recovery times, the flip angle is small and the TR long enough to permit full recovery of the fat and water vectors before the next RF excitation pulse is applied. In practice, small flip angles produce such little transverse magnetization that full longitudinal recovery occurs even if the TR is short (Figure 4.5).

Using extrinsic contrast parameters in gradient-echo – PD weighting

To obtain a PD-weighted image, both $T1$ and $T2^*$ processes are minimized so that the differences in proton density of the tissues are demonstrated. To minimize $T2^*$ decay, the TE is short so that neither the fat nor the water vectors have had time to decay. To minimize $T1$ recovery, the flip angle is small and the TR long enough to permit full recovery of longitudinal magnetization before the next RF excitation pulse is applied.

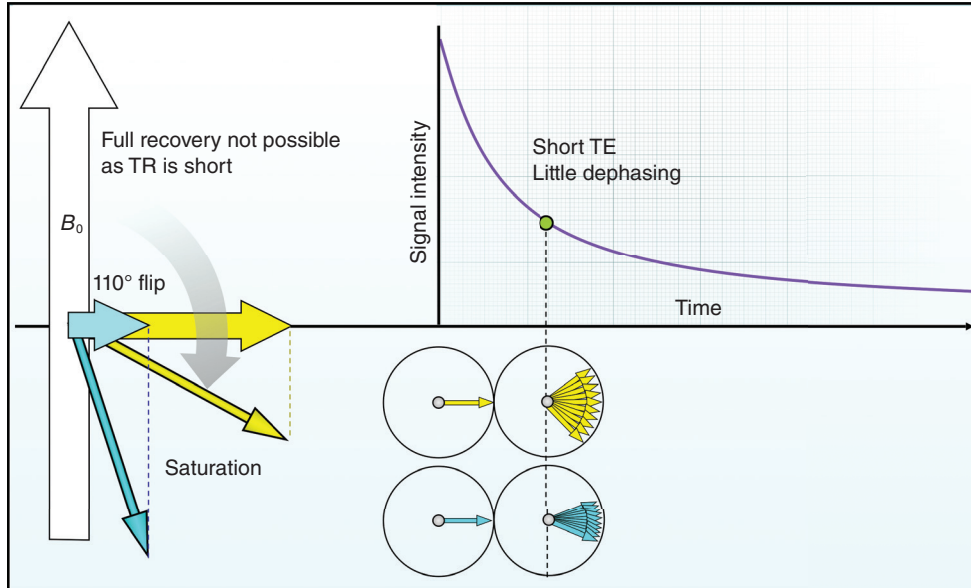


Figure 4.4 T1 contrast in gradient-echo.

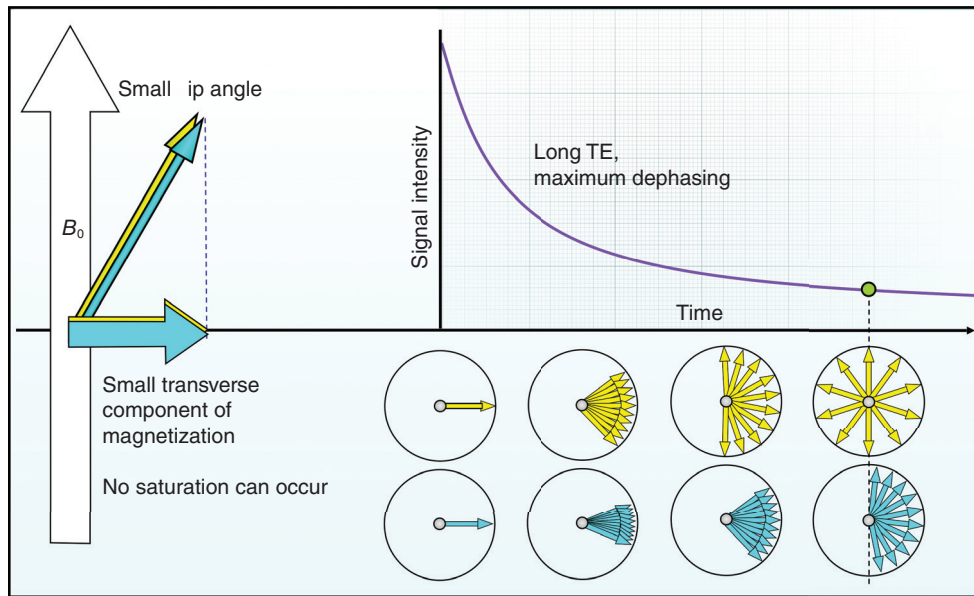


Figure 4.5 T2* contrast in gradient-echo.

Learning tip:

Weighting in gradient-echo using the heat analogy



Look at Figures 4.6–4.8. These are similar to Figures 2.19–2.21. The difference is the addition of a flip angle knob on the stove. This reflects the fact that the flip angle is an extrinsic contrast parameter in gradient-echo pulse sequences. In spin-echo pulse sequences, the flip angle is usually the same (90°) and is therefore not included in the analogy in Chapter 2.

For T1 weighting, turn the heat up on T1 contrast and the heat down on T2* contrast.

- To turn the heat up on T1 contrast, the TR is short (TR knob down), and the flip angle is high (flip angle knob up).
- To turn the heat down on T2* contrast, the TE is short (TE knob down) (Figure 4.6).

For T2* weighting, turn the heat up on T2* contrast and the heat down on T1 contrast.

- To turn the heat up on T2* contrast, the TE is long (TE knob up).
- To turn the heat down on T1 contrast, the TR is long (TR knob up), and the flip angle is low (flip angle knob down) (Figure 4.7).




Figure 4.6 T1 weighting in gradient-echo and the heat analogy.




Figure 4.7 T2* weighting in gradient-echo and the heat analogy.




Figure 4.8 PD weighting in gradient-echo and the heat analogy.

For proton density weighting, turn the heat down on T1 contrast and the heat down on T2* contrast. Proton density contrast therefore predominates.

- To turn the heat down on T1 contrast, the TR is long (TR knob up) and the flip angle low (flip angle knob down).
- To turn the heat down on T2* contrast, the TE is short (TE knob down) (Figure 4.8).

Table 4.3 Comparison of extrinsic parameters – spin-echo and gradient-echo.

Sequence	TR	TE	Flip angle
Spin-echo	Long 2000 ms+	Long 70 ms+	90°
	Short 300–700 ms+	Short 10–30 ms+	90°
Gradient-echo	Long 100 ms+	Long 15–25 ms	Small 5°–20°
	Short less than 50 ms	Short less than 5 ms	Medium 30°–45°
			Large 70°+

Table 4.4 Things to remember – weighting mechanism gradient-echo pulse sequence.

TR and flip angle control whether the NMV is saturated. Saturation is required for T1 weighting only
TE controls T2* weighting
For a T1-weighted gradient-echo, the flip angle and TR combination ensures that saturation occurs. The flip angle is large and the TR short to achieve this. In addition, the TE is short to minimize T2*
For T2*-weighted gradient-echo, the flip angle and TR combination prevents saturation. The flip angle is small and the TR long to achieve this. In addition, the TE is long to maximize T2*
For PD-weighted gradient-echo, the flip angle and TR combination prevents saturation. The flip angle is small and the TR long to achieve this. In addition, the TE is short to minimize T2*

Weighting mechanism 2 – the steady state

The steady state is a term referred to in Chapter 2 but is commonly associated with gradient-echo sequences. It has a significant impact on image weighting in these pulse sequences. The steady state is generically defined as a stable condition that does not change over time. For example, if a pot of water is placed on a stove, the heating element of the stove gradually heats it up. Heat energy is lost through processes of conduction, convection, and radiation. If the amount of heat energy gained from the heating element of the stove equals the amount of heat energy lost by convection, conduction, and radiation, the temperature of the pot and water remains constant and stable. This is an example of the steady state because the energy “in” equals the energy “out,” and, therefore, the temperature of the whole system remains unchanged for a time.

This analogy works well in MRI. The RF excitation pulse gives energy to hydrogen nuclei, and the amount of energy applied is determined by the flip angle. Energy is lost by hydrogen nuclei through spin-lattice energy transfer, and the amount of lost energy is determined by the TR. Therefore, by selecting a certain combination of TR and flip angle, the overall energy of the system remains constant, as the energy “in” as determined by the flip angle equals the energy “out” as determined by the TR (Figure 4.9). As RF has a low frequency and hence low energy, for most values of flip angle very short TRs are required to achieve the steady state. In fact, the required TRs are shorter than the T1 and T2 relaxation times of tissues. Therefore, unlike spin-echo, where even with short TRs some transverse magnetization decays, in gradient-echo there is no time for transverse magnetization to decay before the pulse sequence is repeated. This magnetization influences weighting as the receiver coil is positioned in the transverse plane.

The number of TR periods needed to reach the steady state depends on the TR, flip angle, and the relaxation times of tissues [3]. However, in gradient-echo sequences, a short TR is deliberately used to minimize the scan time. As the TR is so short, magnetization in tissues does not have time to reach its T1 recovery or T2 decay times before the next RF excitation pulse is applied. Therefore, in the steady state, image contrast is not due to differences in the T1 recovery and T2 decay times of tissues but rather due to the ratio of T1 recovery time to T2 decay time. In tissues where T1 recovery and T2 decay times are similar, signal intensity is high, and where they are dissimilar, signal intensity is low. In the human body, fat and water have this parity (fat, very short T1 recovery and T2 decay times; water, very long T1 recovery and T2 decay times); therefore these tissues return high signal intensity in steady state sequences (Table 4.5). Tissues such as muscle do not have this parity (very short T2 decay time and very long T1 recovery time), so they return a low signal in steady state sequences.




Figure 4.9 The steady state.

Table 4.5 T1 and T2 relaxation times and signal intensity of brain tissue in the steady state at 1 T.

Tissue	T1 time (ms)	T2 time (ms)	T1/T2	Signal intensity
Water	2500	2500	1	↑
Fat	200	100	0.5	↑
Cerebral spinal fluid	2000	300	0.15	↓
White matter	500	200	0.2	↓

Typically, TRs less than 50 ms are considered appropriate to maintain the steady state. The optimum flip angle is determined by the Ernst angle equation (Equation (4.2)). The **Ernst angle** is the flip angle that provides optimum signal intensity for a tissue with a given T1 recovery time scanned using a given TR. Figure 4.10 illustrates typical Ernst angles for three tissues in the brain using a TR of 30 ms. The optimum signal intensity in all three tissues is about 12°, but to obtain good contrast between them, larger flip angles between 30° and 45° are required. As contrast generation is important, medium flip angles in the range shown in Figure 4.10 are common.

Equation 4.2

$$\text{Ernst} = \cos^{-1}[e^{(-TR/T1)}]$$

Ernst is the Ernst angle in degrees
 TR is the repetition time (ms)
 T1 is the T1 relaxation time of a tissue (ms)

This equation determines the maximum signal intensity for a tissue with a certain T1 relaxation time at different TR values. When the flip angle is larger than the Ernst angle, saturation and therefore T1 contrast increase. When the flip angle is less than the Ernst angle, contrast relies more on PD




Figure 4.10 Ernst angle graphs in the brain using a TR of 30 ms.

Weighting mechanism 3 – residual transverse magnetization

In the steady state, there is coexistence of both longitudinal and transverse magnetization. The transverse component of magnetization does not have time to decay and builds up over successive TRs. This transverse magnetization is produced because of previous RF excitation pulses but remains over several TR periods in the transverse plane (see learning tip below). It is called **residual transverse magnetization**, and it affects image contrast, as it induces a voltage in the receiver coil. Tissues with long T2 decay times (i.e. water) are the main component of this residual transverse magnetization and enhance T2 contrast.

Figure 4.11 demonstrates a typical image acquired using a gradient-echo sequence in the steady state. The extrinsic contrast parameters (TR, TE, and flip angle) are selected to generate the steady state and to enhance T2* contrast. However, the influence of the other two weighting mechanisms is also evident. The effect of residual transverse magnetization is seen from the high signal from water in the stomach. Water is also hyperintense in this image because water has a good parity between its T1 recovery and T2 decay times. Fat is also bright on this image for the same reason. Muscle is hypointense because it does not have a parity between its T1 recovery and T2 decay times.




Figure 4.11 Axial steady state image.

Learning tip: **Echo formation in the steady state**

The steady state involves repeatedly applying RF excitation pulses using very short TRs. As the TR is shorter than either the T1 or T2 relaxation times of the tissues, there is a build-up of residual transverse magnetization over successive TR periods. This is because there is not enough time in-between RF excitation pulses for the transverse magnetization to either dephase or realign with B_0 . In the steady state, this residual transverse magnetization is rephased by subsequent RF excitation pulses in the sequence and produces echoes. This might seem odd as, so far, we have assumed that RF excitation pulses excite, and RF rephasing pulses rephase. RF excitation pulses usually have a magnitude of 90° and rephasing pulses a magnitude of 180° . In fact, any pair of RF pulses may produce an echo. Their magnitude is irrelevant – they can all excite and rephase. Let's see how.

RF pulses are pulses of electromagnetic radiation and consist of an oscillating magnetic field. If the frequency of the oscillating magnetic field matches the precessional frequency of the magnetic moments of hydrogen nuclei (Larmor frequency), resonance occurs. The magnitude of the magnetic field (B_1) determines the energy of the pulse and therefore the flip angle (see Chapter 1). An RF excitation pulse resulting in a 90° flip angle is common in spin-echo pulse sequences. However, we have already learned that RF pulses with flip angles other than 90° are still excitation pulses (see earlier in this chapter). The same is true of RF rephasing pulses – they still have the capability

to rephase magnetic moments even if they do not have a magnitude of 180° (see Chapter 3 where RF rephasing angles of less than 180° are discussed in relation to TSE). It is always necessary for the frequency of an RF pulse (whether it be an excitation or rephasing pulse) to oscillate at the Larmor frequency; otherwise excitation and rephasing does not occur. However, the amplitude or magnitude of that pulse is not important. Any RF pulse, regardless of its amplitude, has the potential to excite and rephase.

In spin-echo pulse sequences, this phenomenon is not usually apparent because the TRs and TEs are long enough for there to be little transverse magnetization present at the end of each TR period. Spoiling gradients are often applied at the end of each TR to eliminate any transverse magnetization should there be any present. In addition, crusher pulses are sometimes used around the 90° and 180° RF pulses that destroy the unwanted effects of each pulse. Therefore, in spin-echo pulse sequences, not only is there nothing in the transverse plane at the end of the TR period to rephase, but the RF excitation pulses only excite, and RF rephasing pulses only rephase. Crusher pulses make sure that they only have one function. In gradient-echo pulse sequences in the steady state, these strategies are not used. There is therefore residual transverse magnetization present at the end of each TR period. In addition, every RF pulse is an excitation and rephasing pulse, and rephases the residual transverse magnetization to produce an echo [4].

Look at Figures 4.12 and 4.13. You will see in these diagrams that a train of RF pulses generates multiple echoes. Two “signals” are created at each RF pulse:

- A FID
- An echo (called a spin-echo in Fig. 4.12 and 4.13. These echoes are also termed Hahn or stimulated echoes).

In Figure 4.12, the first RF pulse (RF pulse 1, shown in red) is an excitation pulse and, therefore, produces a FID when it is switched off (also shown in red). This happens regardless of the flip angle because it is applied at the Larmor frequency. As discussed previously, flip angles between 30° and 45° are typically used in the steady state.

The second RF pulse (RF pulse 2, shown in orange) is applied a short TR period later (less than 50 ms as previously described). This is another RF excitation pulse so it also produces a FID (also shown in orange). However, because the TR between RF pulses 1 and 2 is shorter than the




Figure 4.12 Echo formation in steady state I.




Figure 4.13 Echo formation in steady state II.

relaxation times of the tissues, residual transverse magnetization is still present when the RF pulse 2 is applied. Therefore, RF pulse 2 produces a FID but also rephases the residual transverse magnetization still present from the first RF pulse. An echo is therefore produced. This occurs at the same time as the third RF excitation pulse (RF pulse 3, shown in blue) because the time for rephasing this residual transverse magnetization after RF pulse 2 is the same time it took to dephase before RF pulse 2 (Figure 4.13).

These echoes are termed **Hahn** or **stimulated echoes** depending on the amplitude of the RF pulses. Any two 90° RF pulses produce a Hahn echo (after Erwin Hahn who discovered them). Any two RF pulses with varying amplitude, i.e. with flip angles other than 90° produce stimulated echoes. This latter type of echo is used in steady state gradient-echo sequences as variable flip angles are common. In practice, echo production is so rapid that the tails of FID signals merge with stimulated echoes resulting in a continuous signal of varying amplitude.

Think of this as listening to a musical note, let's say from a synthesizer. You always hear the note, but the loudness of the note varies over time as its amplitude changes. It increases, then decreases, and then increases again, but it is always heard. This is analogous to the FID decaying (musical note starts loud and then becomes quieter), followed immediately by the stimulated echo (musical note gets louder again and then quieter as it decays), followed immediately by the FID in the next TR (musical note gets louder again) and so on. In the interest of simplicity, the diagrams in this chapter show the FID and the stimulated echo separately.

Table 4.6 Things to remember – the steady state.

The steady state is created when the TR is shorter than the T1 and T2 relaxation times of tissues. Residual transverse magnetization therefore builds up over time

The residual transverse magnetization is rephased by subsequent RF pulses to form stimulated echoes

The resultant image contrast is therefore determined by the ratio of T1 and T2 in a tissue and whether the FID or the stimulated echo is sampled

The weighting of various steady state gradient-echo sequences depends on whether the stimulated echo, or the FID, or both are used to generate the gradient-echo. Their contrast is determined by which of these are used to create the gradient-echo. This is the echo that is read by the MR system to produce an image.

- The stimulated echo contains mainly $T2^*/T2$ weighted information because it is generated from the residual transverse magnetization. As water has the longest $T2$ decay time, water is a large component of the residual transverse magnetization and therefore the stimulated echo. Water is likely to be hyperintense when the stimulated echo is used to create the gradient-echo.
- The FID tends to create contrast that relies on $T1$ and proton density effects. This is because it does not contain residual transverse magnetization. Water is likely to be hypointense when the FID is used to create the gradient-echo.
- When both the stimulated echo and the FID are used to create the gradient-echo, $T1$, proton density, and $T2^*$ weighting are achievable.

Gradient-echo sequences are classified according to which of these signals they use. They are generically referred to as follows:

- Coherent or rewind gradient-echo
- Incoherent or spoiled gradient-echo
- Reverse-echo gradient-echo
- Balanced gradient-echo
- Fast gradient-echo.

Pulse sequences in the gradient-echo family are now discussed in terms of their mechanism, examples of when they might be used, and suggested parameters. In the following section, scan tips link the theory of gradient-echo pulse sequences to practice. Theory is related to what is going on “behind the scenes” when a timing parameter is selected in the scan protocol.

COHERENT OR REWOUND GRADIENT-ECHO

Mechanism

Coherent or rewind gradient-echo pulse sequences use a variable flip angle RF excitation pulse followed by gradient rephasing to produce a gradient-echo. The steady state is maintained by selecting a TR shorter than the $T1$ and $T2$ relaxation times of tissues. There is therefore residual transverse magnetization left over when the next RF excitation pulse is applied. These sequences maintain the coherency of this residual magnetization by rewinding. This is achieved by reversing the slope of the phase-encoding gradient after readout [5] (Figure 4.14). Rewinding rephases all transverse magnetization regardless of when it is created so that it is in phase or coherent at the beginning of the next TR period. Therefore, the resultant gradient-echo contains information from the FID and the stimulated echo. These sequences may therefore be used to achieve $T1$ -, PD -, or $T2^*$ -weighted images, although traditionally they are used in conjunction with a long TE to produce $T2^*$ weighting.




Figure 4.14 The coherent gradient-echo sequence.

Uses

Coherent or rewind gradient-echo pulse sequences are generally used to create T2*-weighted images in a very short scan time (Figures 4.15 and 4.16). As water is hyperintense, they are often said to have an angiographic, myelographic, or arthrographic effect. They can be used to determine whether a vessel is patent or whether an area contains fluid. They may be acquired slice by slice or in a 3D volume acquisition. As the TR is short, slices can be acquired in a single breath-hold.

Suggested parameters

To maintain the steady state:

- Flip angle 30°–45°
- TR 20–50 ms.

To maximize T2*:

- Long TE 10–15 ms.




Figure 4.15 Axial coherent gradient-echo sequence of the abdomen. Source: Westbrook 2014 [6]. Reproduced with permission of John Wiley & Sons.



Figure 4.16 Sagittal coherent gradient-echo sequence of the knee. Source: Westbrook 2014 [6]. Reproduced with permission of John Wiley & Sons.

Table 4.7 Advantages and disadvantages of coherent or rewind gradient-echo.

Advantages	Disadvantages
Very fast scans	Reduced SNR in 2D acquisitions
Very sensitive to flow so useful for angiography	Magnetic susceptibility increases
Can be acquired in a volume acquisition	Loud gradient noise

In addition:

- use gradient moment rephasing to accentuate T2* and reduce flow artifact (see Chapter 8)
- average scan time – seconds for single slice, minutes for volumes.

Scan tip: Something interesting about coherent gradient-echo

Coherent gradient-echo is a sequence that was first used in the 1980s before other fast sequences such as turbo or fast spin-echo were developed. As both the FID and stimulated echo are sampled in this sequence, it is possible to obtain PD-, T1-, and T2*-weighted images. It is also normally possible to use this sequence out of the steady state with a TR similar to that used in T1-weighted spin-echo (e.g. 400 ms). This enables acquisition of multiple slices in the TR period and reduces some artifacts. It is worth noting that if any of the other gradient-echo sequences yield poor image quality it might be worth trying coherent gradient-echo.

The following parameters are a good place to start:

- T1 weighting TR 400 ms/TE 5 ms/flip angle 90°
- PD weighting TR 400 ms/TE 5 ms/flip angle 20°
- T2* weighting TR 400 ms/TE 15 ms/flip angle 20°.

Table 4.8 Things to remember – coherent or rewind gradient-echo.

Coherent gradient-echo is a steady state sequence that utilizes a short TR and medium flip angle
A reversal of the phase-encoding gradient rewinds all transverse magnetization so that its coherency is maintained
Both the FID and the stimulated echo are sampled so that T1, T2*, and PD weighting are possible
This sequence is usually used with T2* weighting with a long TE to image water

INCOHERENT OR SPOILED GRADIENT-ECHO

Mechanism

Incoherent or **spoiled gradient-echo** pulse sequences begin with a variable flip angle RF excitation pulse and use gradient rephasing to produce a gradient-echo. The steady state is maintained so that residual transverse magnetization is left over from previous TR periods. These sequences

dephase or spoil this magnetization so that its effect on image contrast is minimal. Only transverse magnetization from the previous excitation is used, i.e. the FID, enabling T1 and proton density contrast to dominate. There are two spoiling methods:

- *RF spoiling.* RF excitation pulses are transmitted not only at a certain frequency to excite each slice but also at a specific phase. Every TR, the phase angle of the transverse magnetization is changed [6]. A **phase-locked circuit** is used, which means that the receiver coil discriminates between transverse magnetization that has just been created by the most recent RF excitation pulse and residual transverse magnetization created by previous RF excitation pulses. This is possible because the phase angle of the residual transverse magnetization is different from that of the newly created transverse magnetization.



Look at Figure 4.17 and use the watch analogy from Chapter 1 to see how this sequence works. Imagine you are in the rotating frame of reference or “riding along” with the vectors at the Larmor frequency. This means that you disregard the precessional rotation of transverse magnetization. The first RF excitation pulse has a phase angle of 3 o’clock. This means that the resultant transverse magnetization is created at 3 o’clock in the transverse plane. The magnetic moments dephase and are rephased by a gradient to produce a gradient-echo. The receiver coil, situated in the transverse plane, detects this gradient-echo. A short TR period




Figure 4.17 RF spoiling in the incoherent gradient-echo sequence.

later the process is repeated, but this time the RF excitation pulse creates transverse magnetization at a different phase angle such as 6 o'clock. The magnetic moments dephase and are rephased by a gradient to produce a second gradient-echo. The receiver coil detects this gradient-echo. However, as the TR is so short, residual transverse magnetization created at 3 o'clock is still present, as it has not had time to decay. As this residual transverse magnetization has a different phase angle than the transverse magnetization that was just created, it is not detected by the receiver coil and therefore does not impact image contrast. This is **RF spoiling** and enables only gradient-echoes produced from the most recently created transverse magnetization to affect image contrast.

- *Gradient spoiling.* Gradients are used to dephase and rephase residual magnetization. Gradient spoiling is the opposite of rewinding. In gradient spoiling, the slice-select, phase encoding, and frequency-encoding gradients are used to dephase residual magnetization so that it is incoherent at the beginning of the next TR period. T2* effects are therefore reduced. The uses and parameters involved in these pulse sequences are similar to those used in RF spoiling. However, most manufacturers use RF spoiling in incoherent or spoiled gradient-echo pulse sequences.

Uses

As the stimulated echo contains mainly T2* and T2 information, and this is spoiled, these pulse sequences produce T1- or PD-weighted images. This is because image contrast is mainly influenced by the FID that contributes T1 and proton density contrast. However, flowing water (blood and CSF) may have a rather high signal due to gradient rephasing (see Chapter 8). These sequences are used for 2D and volume acquisitions, and, as the TR is short, 2D acquisitions are used to obtain T1-weighted breath-hold images. Incoherent or spoiled gradient-echo sequences also demonstrate good T1 anatomy and pathology after gadolinium contrast enhancement (Figures 4.18 and 4.19).

Suggested parameters

To maintain the steady state:

- Flip angle 30°–45°
- TR 20–50 ms.

To maximize T1:

- Short TE 5–10 ms.

In addition:

- Average scan time – several seconds for single slice, minutes for volumes.

Figure 4.18 Sagittal incoherent gradient-echo sequence of the ankle. Source: Westbrook 2014 [6]. Reproduced with permission of John Wiley & Sons.



Figure 4.19 Coronal incoherent gradient-echo after contrast enhancement. Source: Westbrook 2014 [6]. Reproduced with permission of John Wiley & Sons.



Table 4.9 Advantages and disadvantages of incoherent or spoiled gradient-echo.

Advantages	Disadvantages
Shorter scan times	Reduced SNR in 2D acquisitions
Can be used after gadolinium injection	Magnetic susceptibility increases
Can be acquired in a volume acquisition	Loud gradient noise
Good SNR and anatomical detail in 3D	

Table 4.10 Things to remember – incoherent or spoiled gradient-echo.

Incoherent gradient-echo is a steady state sequence that utilizes a short TR and medium flip angle

RF spoiling ensures that residual transverse magnetization is not sampled. This is achieved by altering the phase angle of each RF excitation pulse every TR and locking this to the receiver coil

Only the FID is sampled so that T1 weighting predominates

Scan tip: T2* vs true T2

The difference between the terms T2 and T2* is well demonstrated in imaging of the cervical spine. If, for example, the suspected pathology is a herniated disk causing cervical myelopathy using a T2* gradient-echo sequence such as coherent or rewind, gradient-echo is a good choice. The disk is demonstrated as a low-signal-intensity disk bulge herniating into a high-signal-intensity CSF-filled thecal sac. If, however, the pathology is subtle, for example, a small multiple sclerosis plaque within the cord, then this might be missed in gradient-echo sequences. As the TE is not long enough to measure the differences in actual T2 decay times of the tissues, subtle pathologies that do not produce any changes around them become less visible.

To see these pathologies, it is important to use pulse sequences that use long TEs. They are likely to produce images where the differences in the T2 decay times of the tissues are observable because there is enough time for these processes to occur before the echo is generated. Conventional spin-echo and TSE are a good choice, but there are several disadvantages with these sequences (see Chapter 3). It is difficult to use a long TE in gradient-echo sequences because they are designed to be used with a short TR to achieve short scan times. Reverse-echo gradient-echo pulse sequences allow the combination of a short TR and a long TE so that true T2 weighting is achieved at the same time as a short scan time. How is this possible? Read on!

REVERSE-ECHO GRADIENT-ECHO

Mechanism

In gradient-echo sequences, the TE is not long enough to measure the T2 decay time of tissues, as a TE of at least 70 ms is required for this. In addition, gradient rephasing is inefficient so that gradient-echoes are dominated by T2* effects. True T2 weighting is difficult to achieve. The **reverse-echo gradient-echo** overcomes this problem in obtaining images that have a sufficiently long TE and less T2* than in other steady state sequences.

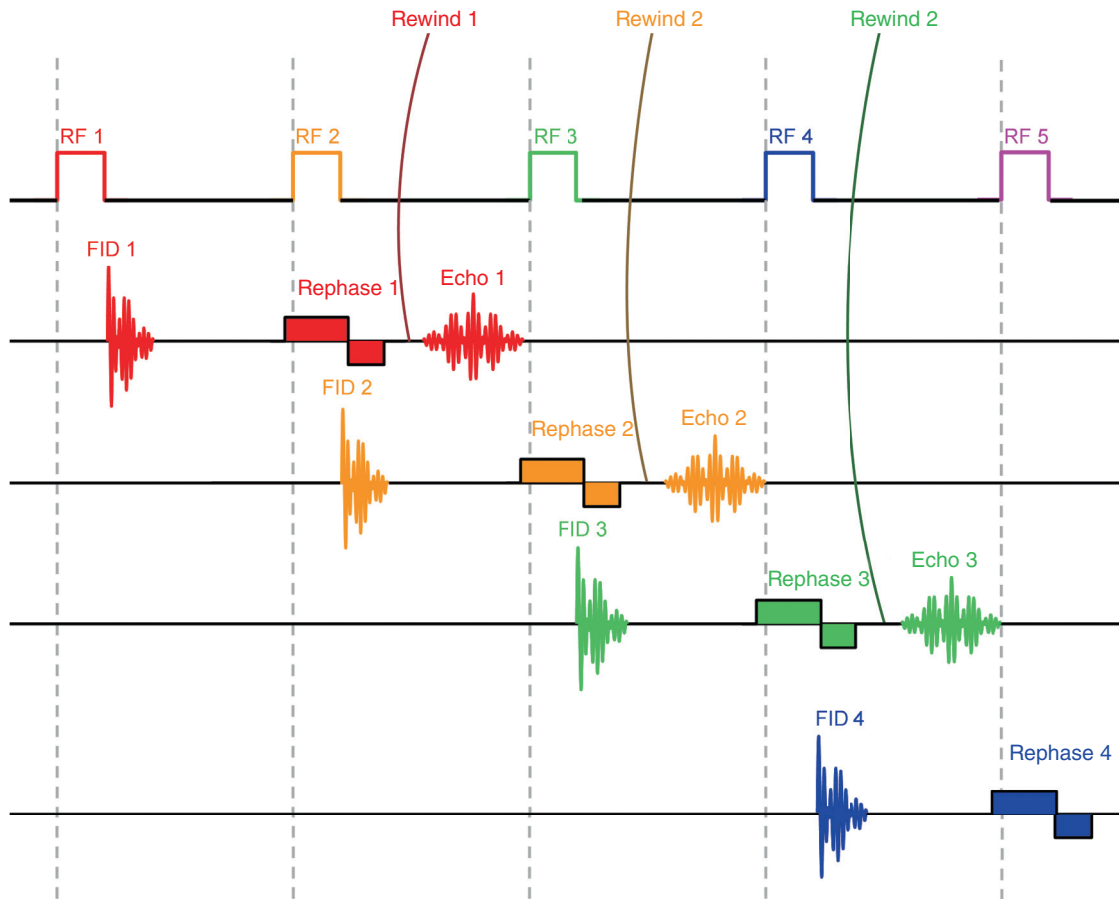


Figure 4.20 The reverse-echo sequence.

As previously described, every RF pulse regardless of its net magnitude contains energies that have sufficient magnitude to rephase spins and produce a stimulated echo. In this sequence, only the stimulated echo is read. To do this, the stimulated echo is repositioned so that it does not occur at the same time as the subsequent RF excitation pulse. A rewinder gradient accelerates rephasing so that the stimulated echo occurs sooner than it normally does (Figure 4.20). Rewinding is achieved by applying the positive lobe of the frequency-encoding gradient [7]. The resultant gradient-echo is dominated by the stimulated echo and therefore demonstrates better T2 weighting than conventional gradient-echo sequences. This is because the TE (called the effective TE) is longer than the TR. In this sequence, there are usually two TEs:

- The **actual TE** is the time between the peak of the gradient-echo and the *next* RF excitation pulse. This is the TE selected in the scan protocol in these sequences, but it is not the TE that determines T2 contrast.
- The **effective TE** is the time from the peak of the gradient-echo to a *previous* RF excitation pulse (i.e. the RF pulse that created its FID). This is the TE that determines T2 contrast, as this is the time allowed for T2 decay in the gradient-echo.




Figure 4.21 Axial reverse-echo gradient-echo in the brain.

Equation (4.3) shows how the effective TE is calculated, and, as it is longer than the TR, T2 weighting is possible using a short TR.

Equation 4.3

$$TE_{\text{eff}} = 2 \times TR - TE_{\text{act}}$$

TE_{eff} is the effective TE in ms
 TE_{act} is the TE set at the console in ms
 TR is the repetition time in ms

This equation shows that the effective TE is longer than the TR so T2 weighting increases. It also shows that shorter actual TEs increase T2 weighting

Uses

Reverse-echo gradient-echo pulses sequences were used to acquire images that demonstrate true T2 weighting (Figure 4.21). They were especially useful in the brain and joints with both 2D and 3D volumetric acquisitions. TSE has now largely replaced this sequence, as it produces better T2 weighting in short scan times. However, the process of shifting the stimulated echo is used in sequences where rapid data acquisition and long TEs are required. An example of this is in perfusion imaging (Figure 4.22).




Figure 4.22 Perfusion imaging using an echo shifting sequence. Source: Westbrook 2015 [9]. Reproduced with permission of John Wiley & Sons.

Table 4.11 Advantages and disadvantages reverse-echo gradient-echo.

Advantages	Disadvantages
Fast shorter scan times	Reduced SNR in 2D acquisitions
Truer T2 than in conventional gradient-echo	Loud gradient noise
Can be acquired in a volume acquisition	Susceptible to artifacts
Good SNR and anatomical detail in 3D	Image quality can be poor

Suggested parameters

To maintain the steady state:

- Flip angle: 30°–45°
- TR: 20–50 ms
- The actual TE affects the effective TE. The longer the actual TE, the shorter the effective TE. The actual TE should therefore be as short as possible to enhance T2 contrast.

Table 4.12 Things to remember – reverse-echo gradient-echo.

Reverse-echo gradient-echo is a steady state sequence that utilizes a short TR and medium flip angle
Rephasing of the stimulated echo is initiated with an RF pulse but the echo is repositioned by a rephasing gradient
Only the stimulated echo is sampled, and due to its repositioning, the TE of this echo is long enough to include T2 rather than T2* contrast

In addition:

- Average scan time – seconds for slice-by-slice acquisitions to several minutes for volumes.

Learning tip: How to differentiate common steady state sequences

The steady state produces two signals:

- A FID made up of transverse magnetization that has just been created by switching off an RF excitation pulse
- A stimulated echo (labeled spin-echo in Figs 4.23–4.25) made up of the residual transverse magnetization component that builds up over time.

Coherent gradient-echo, incoherent gradient-echo, and reverse-echo gradient-echo pulse sequences are differentiated according to whether they use one or both of these signals.

- *Coherent* gradient-echo samples both the FID and the stimulated echo to produce either T1-/PD- or T2*-weighted images depending on the TE (Figure 4.23).
- Incoherent gradient-echo samples the FID only to produce images that are mainly T1/PD weighted (Figure 4.24).
- Reverse-echo gradient-echo samples the stimulated echo only to produce images that are T2 weighted (Figure 4.25).




Figure 4.23 Echo formation in coherent gradient-echo.




Figure 4.24 Echo formation in incoherent gradient-echo.




Figure 4.25 Echo formation in reverse-echo gradient-echo.

Scan tip:**Parameter selection in gradient-echo – what's going on behind the scenes?**

When we alter extrinsic contrast parameters in gradient-echo pulse sequences, behind the scenes, we determine image weighting (see Chapter 2). When we select the TR and the flip angle in the scan protocol, we control how much T1 recovery is permitted between each RF excitation pulse and how far the vectors are moved by the RF excitation pulse. We therefore control the extent to which T1 contrast influences image weighting. We also determine the SNR, the scan time, and the slice number (see Chapter 7), but these factors are not usually as important as weighting.

In most gradient-echo pulse sequences, the TR and the flip angle are selected to maintain the steady state rather than to control T1 contrast. When we select the TE in the scan protocol, we control how much T2* decay is permitted between the RF excitation pulse and the peak of the gradient-echo. We therefore control the extent to which T2* contrast influences image weighting. We also determine the SNR (see Chapter 7), but this is not usually as important as weighting.

119

BALANCED GRADIENT-ECHO

Mechanism

The **balanced gradient-echo** pulse sequence is a modification of the coherent gradient-echo sequence. It uses a balanced gradient scheme to correct for phase errors in flowing blood and CSF, and an alternating RF excitation scheme to enhance steady state effects. The balanced gradient system is shown in Figure 4.26. As the area of the gradient under the line equals that above the line, moving spins accumulate a zero-phase change as they pass along the gradients. As a result, the magnetic moments of flowing spins are coherent and have a high signal intensity. This gradient scheme is the same as flow compensation or gradient moment rephasing (see Chapter 8). In balanced gradient-echo, these gradients are applied in the slice and frequency axes.




Figure 4.26 Balanced gradient system in balanced gradient-echo.




Figure 4.27 Maintenance of the steady state in balanced gradient-echo.

In addition, higher flip angles and shorter TRs are used than in coherent gradient-echo producing a higher SNR and shorter scan times. Normally, this combination of flip angle and TR results in saturation and therefore enhanced T1 contrast. However, saturation is avoided by changing the phase of the RF excitation pulse every TR. This is achieved by selecting a flip angle of 90°, for example, but in the first TR period only applying half of this, i.e. 45°. In successive TRs, the full flip angle is applied but with alternating phase angles so that the resultant transverse magnetization is created at a different phase every TR (Figure 4.27). Consequently, saturation is avoided, and fat and water, which have T1/T2 relaxation times approaching parity, return a higher signal than tissues that do not. The resultant images display a high SNR; good CNR between fat, water, and surrounding tissues; and fewer flow voids. Also, these images are obtained in a very short scan time.

Uses

Balanced gradient-echo was developed initially for imaging the heart and great vessels, but is also used in spinal imaging (Figure 4.28), especially the cervical spine and internal auditory meatus, as CSF flow is reduced. It is also sometimes used in joint and abdominal imaging.

121




Figure 4.28 Axial balanced gradient-echo image of the lumbar spine.

Table 4.13 Advantages and disadvantages of balanced gradient-echo.

Advantages	Disadvantages
Fast shorter scan times	Reduced SNR in 2D acquisitions
Reduced artifact from flow	Loud gradient noise
Good SNR and anatomical detail in 3D	Susceptible to artifacts
Images demonstrate good contrast	Requires high performance gradients

Table 4.14 Things to remember – balanced gradient-echo.

Balanced gradient-echo is a steady state sequence in which longitudinal magnetization is maintained during the acquisition, thereby preventing saturation

This is achieved by altering the phase angle of each RF excitation pulse every TR

A balanced gradient scheme is used to correct for flow artifacts

Suggested parameters

- Flip angle variable (larger flip angles increase signal)
- Short TR less than 10 ms (reduces scan time and flow artifact)
- Long TE 5–10 ms.

FAST GRADIENT-ECHO

Very fast versions of some gradient-echo pulse sequences acquire a volume in a single breath-hold. These usually employ coherent or incoherent gradient-echo sequences, but the TE is significantly reduced. This is achieved by applying only a portion of the RF excitation pulse so that it takes much less time to apply and switch off. Only a proportion of the echo is read (partial echo) and the receive bandwidth is widened (see Chapter 6). In addition, a technique called **ramped sampling** is used. Sampling begins before the frequency-encoding gradient reaches its maximum amplitude. These measures ensure that the TE is kept to a minimum so that the TR and therefore the scan time are reduced.

Many fast sequences use extra pulses, applied before the pulse sequence begins, to premagnetize the tissue. This premagnetization is achieved by applying a 180° RF inverting pulse before the rest of the pulse sequence begins. This inverts the NMV into full saturation, and at a specified delay time, the pulse sequence itself begins. This enhances T1 contrast and may also null signal from certain organs and tissues as in inversion recovery pulse sequences (see Chapter 3).

Fast gradient systems permit multislice gradient-echo sequences with very short TEs. Multiple images are therefore acquired in a single breath-hold and are free from respiratory motion artifacts. In addition, fast gradient-echo acquisitions are useful when temporal resolution is required. This is especially important after the administration of contrast agents when the selection of fast gradient-echo permits dynamic imaging of an enhancing lesion.

ECHO PLANAR IMAGING

Echo planar imaging (EPI) is a rapid acquisition technique that begins with a sequence of one or more RF pulses and is followed by a series of gradient-echoes. These gradient-echoes are typically generated by oscillation of the readout gradient (see Chapter 5). A different image contrast is achieved by beginning the sequence either with a variable RF excitation pulse termed **gradient-echo EPI (GE-EPI)** or with 90° and 180° RF pulses termed **spin-echo EPI (SE-EPI)**. GE-EPI begins with an RF excitation pulse of any flip angle and is followed by EPI readout of gradient-echoes (Figure 4.29). In this scenario, images are acquired in one TR period.




Figure 4.29 GE-EPI sequence.

SE-EPI begins with a 90° RF excitation pulse followed by a 180° RF rephasing pulse followed by an EPI readout of gradient-echoes (Figure 4.30). Application of the rephasing pulse helps to “clean up” some of the artifacts caused by magnetic field inhomogeneities and chemical shift (see Chapter 8). SE-EPI has longer scan times but generally a better image quality than GE-EPI, but the extra RF pulses increase RF deposition to the patient (see Chapter 10). EPI sequences can begin with any type of RF pulse. An example is EPI-FLAIR ($180^\circ/90^\circ/180^\circ$ followed by EPI readout) where CSF is nulled, but the sequence is significantly faster than in conventional FLAIR sequencing (Table 4.15).




Figure 4.30 SE-EPI sequence.

Table 4.15 Comparison of single-shot and multishot techniques.

Sequence	Mechanism	Readout	Time
SS or MS-FSE	90/180 echo train	Spin-echo	min/s
SS or MS SE-EPI	90/180	Gradient-echo	s/sub-s
SS or MS GE-EPI	Variable flip	Gradient-echo	s/sub-s
SS or MS IR-EPI	189/90/180	Gradient-echo	s/sub-s

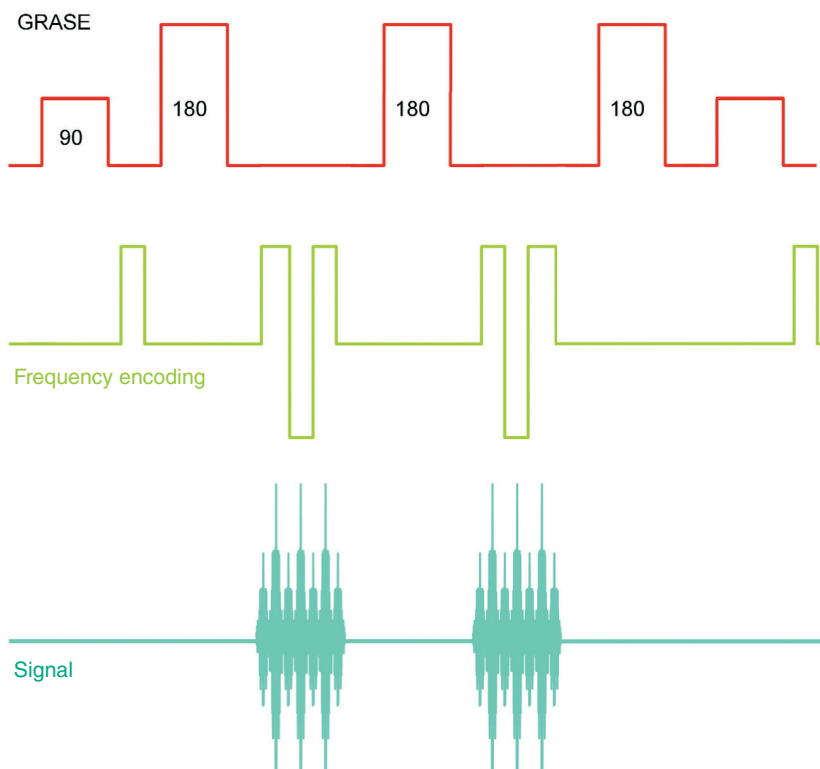


Figure 4.31 GRASE sequence.

EPI sequences are often run in conjunction with single-shot imaging (see Chapter 6). In all single-shot techniques, all k -space is filled at once so the recovery rates of vectors in individual tissues are not critical. For this reason, the TR is said to equal infinity (because it is infinitely long). Either PD or T_2^* weighting is achieved by selecting either a short or long effective TE, which corresponds to the time interval between the RF excitation pulse and when the center of k -space is filled. T_1 weighting is produced by applying an RF inverting pulse before the RF excitation pulse to produce saturation.

Hybrid sequences combine gradient- and spin-echoes, such as **GRASE** (gradient- and spin-echo). Typically, a series of gradient rephasing is followed by an RF rephasing pulse (Figure 4.31). The hybrid sequence uses the benefits of both types of rephasing methods; i.e. the speed of gradient rephasing and the ability of the RF pulse to compensate for T_2^* effects.

Uses and limitations

The rapid scan times of the EPI pulse sequence decreases physiological motion in MR images, which is advantageous when imaging the heart and coronary vessels and when performing interventional techniques (Figure 4.32). Rapid imaging also enables visualization of physiology such as perfusion and blood oxygenation. As this pulse sequence requires rapid switching of gradient



Figure 4.32 Axial SE-EPI of the abdomen. Source: Westbrook 2015 [9]. Reproduced with permission of John Wiley & Sons.

polarity, particularly in the frequency encoding axis, nerve stimulation sometimes occurs. In addition, gradient noise is high so acoustic insulation and ear protection are essential (see Chapter 10).

In addition, many artifacts are seen in EPI including distortion and chemical shift (see Chapter 8). As each gradient-echo is acquired rapidly, there is relatively little chemical shift in the frequency direction. However, there is a larger chemical shift along the phase axis. This phase directional chemical shift artifact does not appear in conventional spin or gradient-echo acquisitions since echoes with different phase-encoding gradient applications are acquired at the same time after RF excitation. In single-shot imaging, however, the length of time required to perform a train of phase-encoding gradient applications means that phase data are encoded at different times after RF excitation. This results in chemical shift that is larger than in spin-echo imaging.

Other artifacts seen in single-shot imaging include blurring and ghosting. **Blurring** occurs as a result of T_2^* decay during the acquisition. If the train of gradient-echoes takes a similar time to decay, signal from the end of the acquisition is reduced, resulting in a loss of resolution and blurring. In EPI acquisitions, **half FOV ghosts** occur as a result of small errors in the timing and shape of readout gradients. This causes differences between echoes acquired with positive and negative readout gradients. These errors cause a ghost of the real image that appears shifted in the phase direction by one half of the FOV. Since it is difficult to eliminate these errors, a correction is usually performed during image reconstruction using information acquired during a reference scan.

There are a few other specialist sequences that are unique to certain manufacturers. For example, Double Echo Steady State (DESS) is a steady state sequence that generates two echoes. One is a coherent gradient and the other a reverse-echo gradient-echo. Both are combined in the final image. The coherent gradient-echo provides resolution and the reverse-echo gradient-echo T_2 contrast.

Table 4.16 Advantages and disadvantages of EPI.

Advantages	Disadvantages
Very fast shorter scan times	Chemical shift artifact is common
Reduced artifact from respiratory and cardiac motion	Peripheral nerve stimulation due to fast switching of gradients
All three types of weighting can be achieved	Susceptible to artifacts
Functional information acquired	
Scan time savings can be used to improve phase resolution	

Table 4.17 Things to remember – fast imaging techniques.

Fast or turbo versions of the traditional gradient-echo sequences use strategies such as ramped sampling and fractional echo to reduce scan times
EPI is a method of filling k -space in a single or multiple shot by oscillating the frequency-encoding gradient and reading the resultant gradient-echoes
Ultrafast sequences are commonly used to acquire functional rather than anatomical information

The last three chapters describe various mechanisms that determine image contrast and demonstrate how we control these mechanisms via pulse sequences. We must select the most appropriate pulse sequence in the scan protocol and then accurate extrinsic contrast parameters within that pulse sequence. This is how images that best demonstrate anatomy and pathology are obtained. Pulse sequences are like vehicles. All vehicles get you from A to B but the controls that are available to us differ depending on the type of vehicle we use. These controls are the user-selectable extrinsic contrast parameters.

Table 4.18 lists the different pulse sequences (vehicles) and their extrinsic parameters (controls).

Table 4.18 Extrinsic contrast parameters selected in each sequence.

		TR	TE	TI	FA	TF	b value
Spin-echo	CSE	✓	✓				
	TSE	✓	✓			✓	
	IR	✓	✓	✓			
	IR-TSE	✓	✓	✓		✓	
	STIR-TSE	✓	✓	✓		✓	
	FLAIR-TSE	✓	✓	✓		✓	
Gradient-echo	Coherent GE	✓	✓		✓		
	Incoherent GE	✓	✓		✓		
	Reverse echo GE	✓	✓		✓		
	Fast GE	✓	✓		✓		
	MP-GE	✓	✓	✓	✓		
	EPI		✓	✓	✓		
Spin-echo or gradient-echo	DWI	✓	✓				✓

Table 4.19 Gradient-echo pulse sequences – summary of what’s going on behind the scenes.

Parameter	Behind the scenes
TR	Controls the amount of T1 recovery and therefore T1 contrast. In practice, selected to maintain the steady state
TE	Controls the amount of T2* decay and therefore T2 contrast
Flip angle	Controls the amount of saturation and therefore T1 contrast. In practice, selected to maintain the steady state
<i>b</i> value	Determines how much phase shift there is across an area of tissue per s in DWI

Table 4.19 summarizes what is going on behind the scenes in gradient-echo pulse sequences. This table also includes the *b* value, which is used in DWI (see Chapter 2).

Now we have a thorough understanding of pulse sequences, let’s discover how slices are acquired and how they are located.



For questions and answers on this topic please visit the supporting companion website for this book: www.wiley.com/go/westbrook/mriinpractice

References

1. Hashemi, R.H., Bradley Jr, W.G., and Lisanti, C.J. (2010). *MRI: The Basics*, 3, 235. Philadelphia, PA: Lippincott Williams and Wilkins.
2. Hashemi, R.H., Bradley Jr, W.G., and Lisanti, C.J. (2010). *MRI: The Basics*, 3, 236. Philadelphia, PA: Lippincott Williams and Wilkins.
3. McRobbie, D.W., Moore, E.A., Graves, M.J. et al. (2017). *From Picture to Proton*, 3, 135. Cambridge: Cambridge University Press.
4. McRobbie, D.W., Moore, E.A., Graves, M.J. et al. (2017). *From Picture to Proton*, 3, 208. Cambridge: Cambridge University Press.
5. McRobbie, D.W., Moore, E.A., Graves, M.J. et al. (2017). *From Picture to Proton*, 3, 214. Cambridge: Cambridge University Press.
6. Hashemi R.H., Bradley Jr, W.G. and Lisanti, C.J. (2010). *MRI: The Basics*, 3, 248. Philadelphia, PA: Lippincott Williams and Wilkins.
7. McRobbie, D.W., Moore, E.A., Graves, M.J. et al. (2017). *From Picture to Proton*, 3, 217. Cambridge: Cambridge University Press.
8. Hashemi, R.H., Bradley Jr, W.G., and Lisanti, C.J. (2010). *MRI: The Basics*, 3, 242. Philadelphia, PA: Lippincott Williams and Wilkins.

5

Spatial encoding

Introduction	128	Frequency encoding	142
Mechanism of gradients	129	Phase encoding	145
Gradient axes	134	Bringing it all together – pulse sequence timing	152
Slice-selection	135		

After reading this chapter, you will be able to:

- *Describe gradients and how they work.*
- *Explain slice-selection.*
- *Understand how gradients spatially locate signal in a slice.*
- *Apply what you have learned to explore how gradients are used in common pulse sequences.*

INTRODUCTION

In Chapter 1, we learned that a radio frequency (RF) excitation pulse is applied at 90° to B_0 at the precessional frequency of the magnetic moments of hydrogen nuclei to cause them to resonate. The RF excitation pulse gives energy to the hydrogen nuclei. This creates magnetization in the transverse plane and puts individual magnetic moments of hydrogen nuclei into phase. The resultant coherent transverse magnetization precesses at the Larmor frequency of hydrogen in the transverse plane.

A voltage or signal is therefore induced in the receiver coil positioned in the transverse plane. It is caused by oscillation of coherent transverse magnetization relative to the receiver coil. This signal is an alternating voltage that has a frequency equal to the Larmor frequency regardless of the origin of signal. As all magnetic moments precess at the same frequency, all

signals oscillate at the same frequency so the system cannot spatially locate it. In other words, the MRI system has no idea where individual signals are coming from because they all have the same frequency.

To produce an image, the MRI system must calculate how much signal is coming from each three-dimensional location in the patient. This location is called a **voxel**. The simplest way to do this is to first locate a slice and then locate signal at each two-dimensional location within it. This location is called a **pixel**. The process is called **spatial encoding**, and it is performed by **gradients**. In this chapter, the mechanisms of gradients are discussed in relation to spatial encoding. Scan tips are also used to link the theory of spatial encoding to practice.

MECHANISM OF GRADIENTS

The concept of gradients was first introduced in Chapter 4 and is further discussed in Chapter 9. When no gradient is applied, all magnetic moments of the hydrogen nuclei precess at the same frequency as they experience the same field strength (in fact inhomogeneities in the field cause magnetic moments to precess at slightly different frequencies, but these changes are relatively small compared with those imposed by a gradient). To locate individual signals, the main magnetic field is altered so that it slopes in a linear and, therefore, predictable way. Graded or sloped magnetic fields are generated by cylindrical electromagnets situated in the warm bore of the cryostat. These coils are called **gradient coils**, and, at certain time points within a pulse sequence, current is passed through each of these coils. According to Faraday's law of electromagnetic induction, when current is passed through a gradient coil, a magnetic field is induced around it (see Chapter 1). This magnetic field is superimposed onto the main magnetic field (B_0) in such a way that the magnetic field strength along the axis of the gradient coil is sloped.

Look at Figures 5.1 and 5.2. In Figure 5.1, a gradient is applied that increases the magnetic field strength toward the right-hand side of the magnet (shown in red) and decreases it toward the left-hand side (blue). Figure 5.2 shows that the gradient coil has three terminals, one at each end of the coil and one in the middle. Current is passed through these terminals into the gradient coil. These terminals enable control of the direction of the current flowing through the gradient coil. This, in turn, determines gradient **polarity**. The polarity of a gradient depends on which end of the gradient magnetic field is higher than B_0 and which is lower. If current flows clockwise through the coil, then the magnetic field induced around the coil adds to B_0 . This increases the magnetic field strength relative to B_0 . If the current flows anticlockwise through the coil, then the magnetic field induced around the coil subtracts from B_0 and decreases the magnetic field strength relative to B_0 . The middle of the axis of the gradient remains at the field strength of the main magnetic field even when the gradient is switched on. This is called the **magnetic isocenter**.

Therefore, to achieve the gradient polarity shown in Figure 5.1, current is applied clockwise through the gradient coil on the right-hand side of Figure 5.2 from the center to the right-hand terminal. The current is also applied anticlockwise through the gradient coil on the left-hand side of Figure 5.2 from the center to the left-hand terminal. The combination of these two currents produces a linear alteration in (or sloped) magnetic field strength from a low magnetic field on the left that gradually increases to the right. In all gradient diagrams in this book, gradient fields higher than the magnetic isocenter are shown in red and those lower in blue.

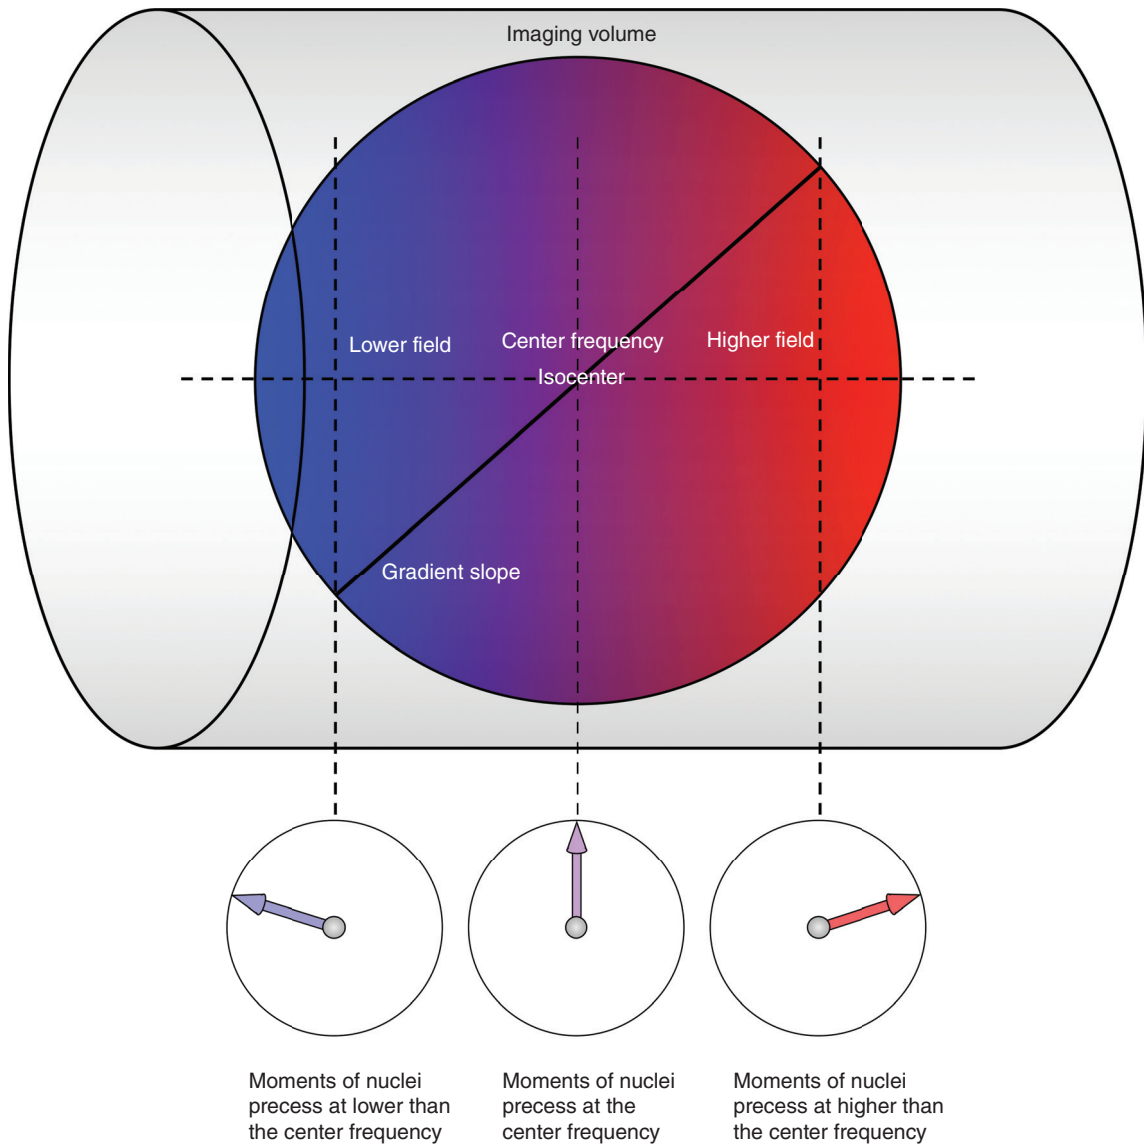


Figure 5.1 How gradients change field strength and precessional frequency.

The Larmor equation states that the precessional frequency of magnetic moments of hydrogen nuclei increases or decreases depending on the magnetic field strength they experience at different points along the gradient (see Figure 5.1). The precessional frequency increases when the magnetic field increases, and decreases when the magnetic field decreases. Magnetic moments of hydrogen nuclei experiencing an increased field strength due to the gradient speed up,




Figure 5.2 Three-terminal gradient coil.

i.e. their precessional frequency increases. Magnetic moments of hydrogen nuclei experiencing a decreased magnetic field strength slow down, i.e. their precessional frequency decreases. Therefore, the position of a spin located along a gradient is identified according to the precessional frequency of its magnetic moment (Table 5.1) (Equation (5.1)). The changes in field strength and therefore frequency imposed by gradients are quite small. They usually vary the magnetic field and therefore the frequency of magnetic moments of spins located along them by less than 1% [1].

Equation 5.1

$$B_{p1} = B_0 + Gp_1$$

B_{p1} is the magnetic field strength at point 1 along the gradient (T)
 B_0 is the main magnetic field strength (T)
 G is the total amplitude of the gradient (mT/m) at p_1 position p_1

This equation shows how the field strength experienced by a spin at any point along the gradient depends on its position along it and the amplitude of the gradient

Table 5.1 Frequency changes along a linear gradient that alters the magnetic field strength by 1 G/cm at a main field strength of 1 T.

Position along gradient	Field strength (G)	Larmor frequency (MHz)
Isocenter	10 000	42.580 0
2 cm positive from isocenter	10 002	42.588 5
1 cm positive from isocenter	10 001	42.584 2
1 cm negative from isocenter	9 999	42.575 7
2 cm negative from isocenter	9 998	42.571 4
10 cm negative from isocenter	9 990	42.537 4

Learning tip: Units of magnetic flux density

The tesla (T) is commonly used to denote magnetic field strength or magnetic flux density using the SI system. This unit is appropriate when quantifying a relatively large magnetic field such as that of the MRI scanner. However, sometimes another unit, used by the centimeter–gram–second (CGS) system, the gauss (G), is applicable instead. This is used to quantify magnetic fields on a smaller scale (1 T is equivalent to 10 000 G). The earth's magnetic field, for example, averages 0.46 G at its surface [2]. As this is such a small magnetic field, the gauss is more useful than the tesla. In MRI, the gauss is often used when referring to the small changes in magnetic field imposed by a gradient. These are the units used in Table 5.1.

Gradients also cause magnetic moments of the hydrogen nuclei to change phase. This is because as the frequency of magnetic moments increases they, gain phase relative to the magnetic moments of other hydrogen nuclei. Likewise, as the frequency of magnetic moments decreases, they lose phase relative to the magnetic moments of other hydrogen nuclei.



The watch analogy discussed in Chapter 1 comes in handy when learning about gradients and understanding why they alter frequency and phase. Imagine that the magnetic moments of hydrogen nuclei sitting along a gradient are watches. When no gradient is switched on, the precessional frequency of the hands of these watches is the same because they all experience the same field strength (B_0).

Imagine that the gradient illustrated in Figure 5.1 is switched on. At the magnetic isocenter, the hands of the watch at this location continue to precess at the same frequency as they were precessing when there was no gradient. This is because there is no change in field strength at the isocenter even when gradients are applied. However, the watches situated to the right-hand side of the magnetic isocenter experience a progressively increased magnetic field strength. As precessional frequency is proportional to the magnetic field strength, the hands of these watches precess increasingly faster and therefore tell a time ahead of the watch located at the magnetic isocenter. This means that they *gain phase* relative to the watch at the magnetic isocenter.

Similarly, the watches situated to the left-hand side of the magnetic isocenter experience a progressively decreased magnetic field strength. As precessional frequency is proportional to the




Figure 5.3 (a) Steep and (b) shallow gradient slopes.

magnetic field strength, the hands of these watches precess increasingly slower, and therefore they tell a time behind that of the watch located at the magnetic isocenter. This means that they *lose phase* relative to the watch at the magnetic isocenter. The farther away from the magnetic isocenter a watch is located, the greater the difference in time (either ahead or behind) between them and the watch at the magnetic isocenter. The degree of difference depends on the amplitude or steepness of the gradient.

Gradient amplitude determines the rate of change of the magnetic field strength along the gradient axis. Steep gradient slopes alter the magnetic field strength between two points more than shallow gradient slopes. Steep gradient slopes therefore alter the precessional frequency of magnetic moments of hydrogen nuclei between two points more than shallow gradient slopes (Figure 5.3). Another way of saying this is that steep gradient slopes cause a large difference in precessional frequency and therefore phase between the magnetic moments of hydrogen nuclei situated along the gradient. Shallow gradient slopes cause a small difference in precessional frequency and therefore phase between the same two points.

Table 5.2 Things to remember – gradient mechanisms.

When a moving current is passed through a conductor, a magnetic field is induced around it
Gradient coils are conductors that cause a linear change in the magnetic field strength along their axes when a current is passed through them
The amount of current passing through the coil determines the amplitude, strength, or slope of the gradient
The direction of the current passing through the coil determines its polarity
When a gradient is switched on, it causes a linear change in the magnetic field strength and therefore precessional frequency and phase of the magnetic moments of hydrogen nuclei that lie along it

Learning tip: Gradients in pulse sequence diagrams

In pulse sequence diagrams, gradients are drawn as shapes above and below a central line. These are called **lobes**. The degree to which a gradient alters the magnetic field strength depends on its amplitude and/or duration. A large area inside the lobe indicates a steep or high-amplitude gradient and/or a long duration; a small area, a low-amplitude gradient and/or a short duration. The polarity of the gradient is illustrated by the position of the gradient lobe relative to a central line. A gradient lobe drawn above the line indicates a positive polarity gradient; a gradient lobe drawn below the line indicates a negative polarity gradient. Each gradient is usually allocated a different line in the pulse sequence schematic.

GRADIENT AXES

There are three gradient coils situated within the bore of the magnet, and these are named according to the axis along which they act when they are switched on. Figure 5.4 shows these directions in a typical superconducting magnet. However, some manufacturers may use a different system, so it is important to check the convention on your scanner.

- The *z gradient* alters the magnetic field strength along the *z*-(*long*) axis of the magnet (from head to foot of the patient).
- The *y gradient* alters the magnetic field strength along the *y*-(*vertical*) axis of the magnet (from the back to the front of the patient).
- The *x gradient* alters the magnetic field strength along the *x*-(*horizontal*) axis of the magnet (from right to left of the patient) (Table 5.3).




Figure 5.4 Gradient axes in a typical superconducting system.

Table 5.3 Gradient labeling.

	Slice-selection	Phase encoding	Frequency encoding
Sagittal	X	Y	Z
Axial (body)	Z	Y	X
Axial (head)	Z	X	Y
Coronal	Y	X	Z

X is across the bore of the magnet from right to left.

The magnetic isocenter is the center point of the axis of all three gradients and the bore of the magnet. The magnetic field strength and therefore the precessional frequency and phase remain unaltered here even when the gradients are applied. Permanent magnets (see Chapter 9) have different axes. The z-axis is vertical, not horizontal, as shown in Figure 5.4.

Gradients perform many important tasks during a pulse sequence as described in Chapters 3 and 4. Can you remember what these are? Gradients are used to either dephase or rephase the magnetic moments of nuclei. However, gradients also perform the following three main tasks. Their purpose is to spatially locate or encode signal depending on where it is located along these three gradients.

- **Slice-selection.** Locating a slice within the selected scan plane.
- Spatially locating (encoding) signal along the long axis of the slice – this is called **frequency encoding**.
- Spatially locating (encoding) signal along the short axis of the slice – this is called **phase encoding**.

Now let's explore each of these processes in turn.

SLICE-SELECTION

How does it work?

When a gradient coil is switched on, the magnetic field strength and therefore the precessional frequency of magnetic moments of hydrogen nuclei located along its axis are altered in a linear and predictable way. Therefore, magnetic moments of hydrogen nuclei located at any point along the axis of the gradient have a specific precessional frequency, and this is dependent on the amplitude of the gradient and the position of the nucleus along the gradient (Equation (5.1)).

A slice is selectively excited by transmitting an RF excitation pulse that is oscillating at the same or similar frequency to the precessional frequency of the magnetic moments of hydrogen nuclei in the slice. Consequently, resonance occurs in these nuclei. Nuclei situated in other slices along the gradient do not resonate because their precessional frequencies are different due to the presence of the gradient.

Analogy: The tuning fork analogy and slice-selection

Look at Figure 5.5 in which tuning forks are used to illustrate how slice-selection is performed. In Figure 5.5a, a gradient is applied to change the magnetic field strength from low (blue) to high (red) along the z-axis of the magnet. Imagine we are trying to select slice A. With the amplitude of gradient used in Figure 5.5, magnetic moments of hydrogen nuclei in this slice have a precessional frequency of 41.20 MHz when the gradient is applied. Magnetic moments of hydrogen nuclei on either side of this slice have a different precessional frequency because they experience a different magnetic field strength. Without the gradient, all magnetic moments precess at the same frequency, and it is not possible to differentiate them. As the gradient is on in this diagram, magnetic moments of hydrogen nuclei in different slices precess at different frequencies. Therefore, it is possible to differentiate them because slices are selectively excited by transmitting an RF excitation pulse that has the same frequency as the magnetic moments of hydrogen nuclei within each individual slice.

This is analogous to tuning forks tuned to different frequencies located at different locations along the gradient. To produce resonance and excite spins in slice A, an RF excitation pulse is applied that matches the precessional frequency of magnetic moments in slice A, i.e. 41.20 MHz. Doing so causes resonance just in nuclei in slice A; nuclei in other slices do not resonate because their magnetic moments are precessing at different frequencies to the transmitted RF. To produce the same effect in slice B (b), an RF excitation pulse is applied with a frequency of 43.80 MHz to produce resonance in nuclei in slice B. In this example, axial slices are excited (assuming the patient is lying either supine or prone on the scan table). In practice, the location of each slice is changed by altering the center frequency of the RF excitation pulse and keeping the amplitude of the slice-select gradient the same [3].

The selected scan plane determines which of the three gradients performs slice-selection during the RF excitation pulse (Figure 5.6). Typically, they are as follows (although some manufacturers may vary, so please check the convention on your scanner).

- The *z gradient* alters the field strength and precessional frequency along the z-axis of the magnet and therefore selects axial slices (head to foot of the patient).
- The *x gradient* alters the field strength and the precessional frequency along the x-axis of the magnet and therefore selects sagittal slices (left to right of the patient).
- The *y gradient* alters the field strength and the precessional frequency along the y-axis of the magnet and therefore selects coronal slices (back to front of the patient).
- Oblique slices are selected using two or three gradients in combination.

When does slice-selection occur?

In spin-echo pulse sequences, the slice-select gradient is switched on during the application of the 90° RF excitation pulse and during the 180° RF rephasing pulse to selectively excite and rephase each slice (Figure 5.7). In gradient-echo pulse sequences, the slice-select gradient is switched on only during the RF excitation pulse. The significance of this is explored in Chapter 8.

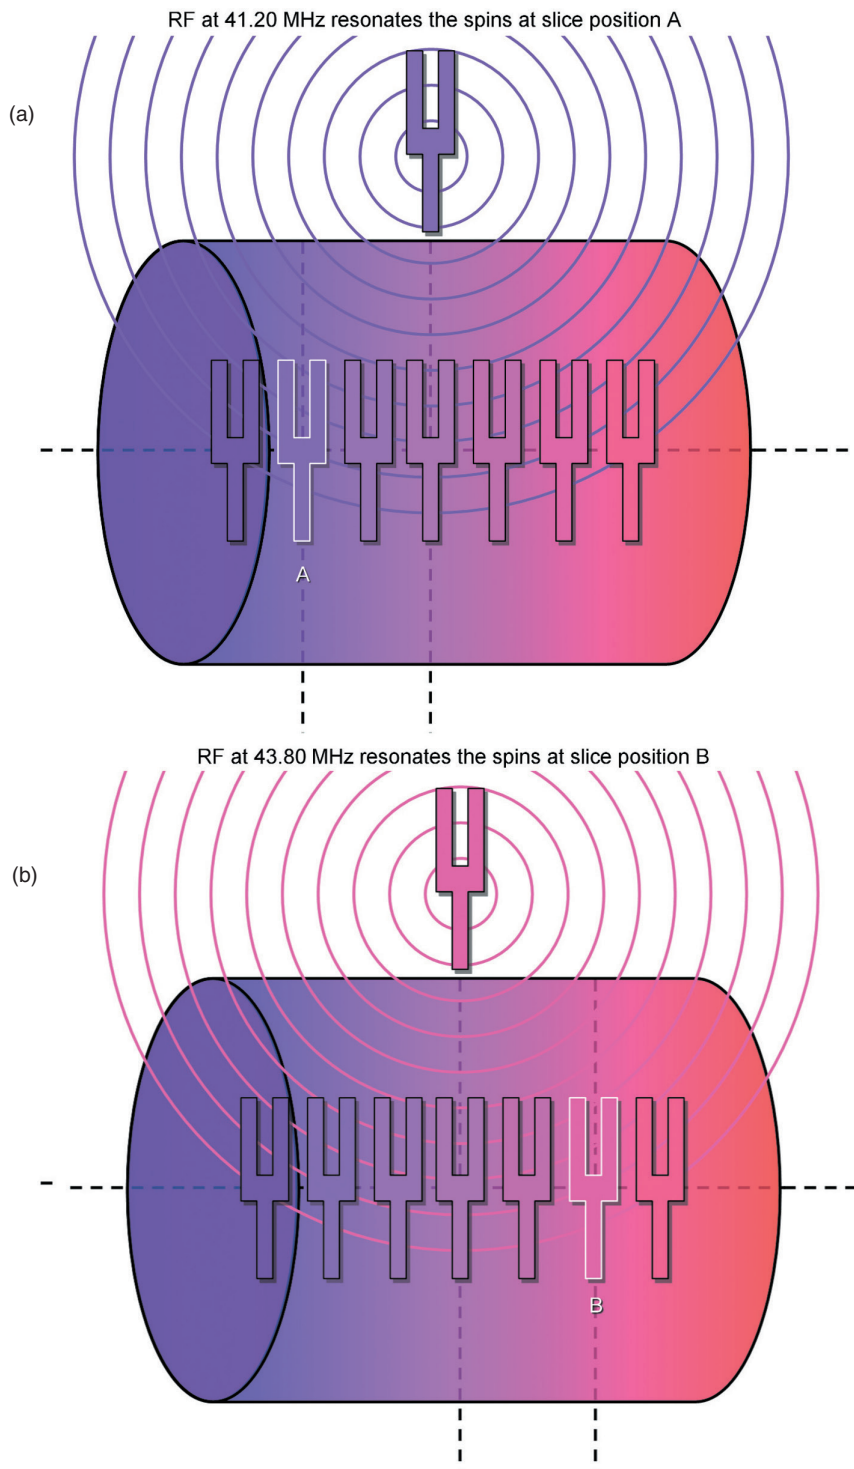


Figure 5.5 Slice-selection and the tuning fork analogy.




Figure 5.6 X, Y, and Z as slice selectors.

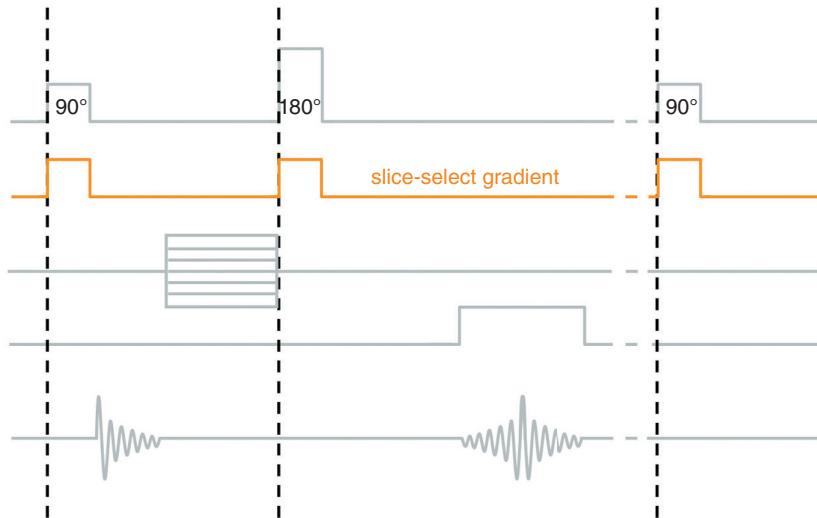


Figure 5.7 Timing of the slice-select gradient in a spin-echo pulse sequence.

Scan tip:

Slice-selection – what's going on behind the scenes?

The MRI system knows which gradient is allocated for slice-selection by the plane of the scan. When we select the scan plane in the scan protocol (axial, coronal, sagittal, or oblique), behind the scenes, we determine which gradient (x, y, or z, or a combination of these) is applied during the RF excitation pulse (and during the RF rephasing pulse in spin-echo pulse sequences).

Slice thickness and slice-selection

To give each slice a thickness, a “band” of nuclei are excited by the RF excitation pulse (and rephased by the RF rephasing pulse in spin-echo pulse sequences). The slope or amplitude of the slice-select gradient determines the difference in precessional frequency between two points on the gradient. Steep gradient slopes result in a large difference in precessional frequency between two points on the gradient, while shallow gradient slopes result in a small difference in precessional frequency between the same two points (see Figure 5.3).

The slice-select gradient is applied at the same time as the RF excitation pulse. The slope of the gradient determines the range of frequencies present between the edges of the slice. The RF excitation pulse contains the same range of frequencies to match the difference in precessional frequency between these edges. This frequency range is called the **bandwidth**, and, as the RF is transmitted at this point, it is specifically called the **transmit bandwidth**. The slice thickness is proportional to the transmit bandwidth (Equation (5.2)) [4].

We previously discovered that the amplitude of any gradient is determined by how much current passes through it when it is switched on. In relation to the slice-select gradient, the amplitude determines the slice thickness.

- To achieve *thin slices*, a steep slice-select slope and/or narrow transmit bandwidth is applied.
- To achieve *thick slices*, a shallow slice-select slope and/or broad transmit bandwidth is applied (Equation (5.2)).

Equation 5.2

$$SI_t = \frac{TBW}{\gamma G}$$

SI_t is the slice thickness

TBW is the transmit bandwidth (KHz)

γ is the gyromagnetic ratio derived from the Larmor equation (MHz/T)

G is the total amplitude of the gradient (mT/m)

This equation shows how the slice thickness is controlled by either changing the gradient amplitude or the transmit bandwidth

In practice, the system automatically applies the appropriate gradient slope and transmit bandwidth according to the required slice thickness. The slice is excited by transmitting RF at the center frequency corresponding to the precessional frequency of magnetic moments of hydrogen nuclei in the middle of the slice, and the transmit bandwidth and gradient slope determine the range of nuclei that resonate on either side of the center (Figure 5.8). The center frequency is called the **carrier frequency** (see Chapter 6). This is a relatively high frequency because it is near to the Larmor frequency (which is on the MHz scale). The transmit bandwidth is a small range of frequencies on either side of the carrier frequency that typically has a range 1 to 2 KHz (which is a much smaller scale) [5]. The gap between the slices is determined by the gradient slope and by the thickness of the slice. The size of the gap is important in reducing image artifact (see Chapter 8).




Figure 5.8 Transmit bandwidth, gradient slope, and slice thickness.

Scan tip: Slice thickness – what's going on behind the scenes?

When we select the slice thickness and gap between the slices in the scan protocol, behind the scenes, we determine how much current is applied through the slice-select gradient during the pulse sequence. Thin slices are produced by passing a large current through the slice-select gradient and thick slices by passing less current. For example, if we select axial slices with a thin, 3 mm slice thickness and 1 mm slice gap, behind the scenes, the z gradient is switched on during RF excitation (and RF rephasing in spin-echo pulse sequences) and a large current is passed through it when it is applied. A large current passing through the slice-select gradient coil produces a high-amplitude or steep gradient. This means that there is a large difference in frequency between two points that are located along the slice-select gradient. If a narrow transmit bandwidth is used in combination with this steep gradient slope, then resonance occurs in a narrow range of frequencies, and a thin slice is excited.

Conversely if we select a thick, 10 mm slice thickness and 5 mm slice gap, behind the scenes, a small current is passed through the slice-select gradient when it is applied. Low current passing through the slice-select gradient coil produces a low-amplitude or shallow gradient. This means that there is a small difference in frequency between two points that are located along the slice-select gradient. If a broad transmit bandwidth is used in combination with this shallow gradient slope, then resonance occurs in a broad range of frequencies, and a thick slice is excited.

Some systems allow us to select a low, high, or medium transmit bandwidth. A low transmit bandwidth produces a good slice profile and is useful to reduce the SAR (see Chapter 10), but images are susceptible to distortion, and the minimum TE increases. A high transmit bandwidth has the opposite effect and is sometimes used in fast gradient-echo sequences to reduce the TE. In practice, the slice thickness is determined by the amplitude of the slice-select gradient and transmit bandwidth. Once the slice thickness is determined, these variables remain fixed during the pulse sequence. The position of each slice is altered by changing the center or carrier frequency [6].

Table 5.4 Things to remember – slice-selection.

Slices are selected by applying a gradient at the same time as the RF excitation pulse (and the RF rephasing pulse in spin-echo pulse sequences)
The slice-select gradient changes the magnetic field strength and therefore the precessional frequency of the magnetic moments of hydrogen nuclei that lie along it
An RF excitation pulse at the specific frequency of magnetic moments of hydrogen in a particular slice on the gradient causes resonance of the slice
RF is transmitted with a bandwidth or range of frequencies on either side of the center frequency of the slice
Slice thickness is altered by changing the slope of the slice-select gradient and the transmit bandwidth
Thin slices require a steep slice-select gradient slope and a narrow transmit bandwidth
Thick slices require a shallow slice-select gradient slope and a broad transmit bandwidth

FREQUENCY ENCODING

How does it work?

142

Once a slice is selected, signal coming from it is located along both axes of the image. Signal is usually located along the long axis of the anatomy by a process known as **frequency encoding**. When the frequency-encoding gradient is switched on, the magnetic field strength, and therefore precessional frequency of magnetic moments of hydrogen nuclei located along the axis of the gradient are altered in a linear, predictable way. The gradient therefore produces a frequency difference along its axis. Nuclei (and signal they create) can now be located along the axis of the gradient according to the frequency of their magnetic moments. Nuclei located at the high end of the gradient field experience a large magnetic field strength, and therefore, their magnetic moments precess quickly. Their precessional frequency is higher than those located at the low end of the gradient. These nuclei experience a low magnetic field strength, and therefore their magnetic moments precess more slowly. Nuclei at the magnetic isocenter experience no difference in the magnetic field strength when the gradient is switched on, so the precessional frequency of their magnetic moments remains unchanged.

The system can now differentiate between nuclei positioned at different locations along the frequency-encoding gradient because their magnetic moments have a different precessional frequency.

Signals produced by these nuclei are used by the computer of the MRI system to place them in different positions in the image (see Chapter 6). The different frequencies created along the frequency-encoding gradient are similar to the different notes made on a piano keyboard. Each note sounds different because the sound waves that each key makes when it is pressed are different. High notes are those that have a high frequency, low notes a low frequency (Figure 5.9).

We usually select the direction of frequency encoding so that it encodes signal along the longest axis of the anatomy. It may help to refer to the images in Chapter 2 to work out which gradient was used for each spatial encoding function. Remember the patient is usually lying supine along the z -axis while on the table (in a superconducting system). Using this standard, it is easy to work out the long and short axis of anatomy.

- In *coronal* and *sagittal* images, the long axis of the anatomy lies along the z -axis of the magnet (head to foot of the patient), and, therefore, the z *gradient* performs frequency encoding.
- In *axial* images, the long axis of the anatomy usually lies along the horizontal axis of the magnet (left to right of the patient), and, therefore, the x *gradient* performs frequency encoding. However, in imaging of the head, the long axis of the anatomy usually lies along the anteroposterior axis of the magnet (back to front of the patient), so in this case the y *gradient* performs frequency encoding.

When does frequency encoding occur?

The frequency-encoding gradient is switched on when signal is received (i.e. when the echo occurs) and is often called the **readout** or **measurement gradient**. The echo is usually centered in the middle of the frequency-encoding gradient so that the gradient is switched on during the rephasing and dephasing part of the echo and its peak (Figure 5.10). The frequency-encoding gradient is switched on for a very specific time. This is called the **sampling time** or **sampling window** (see Chapter 6). For example, if the typical sampling time is 8 ms, the frequency-encoding gradient is switched on for 8 ms during 4 ms of rephasing, the peak, and then 4 ms of dephasing of the echo.

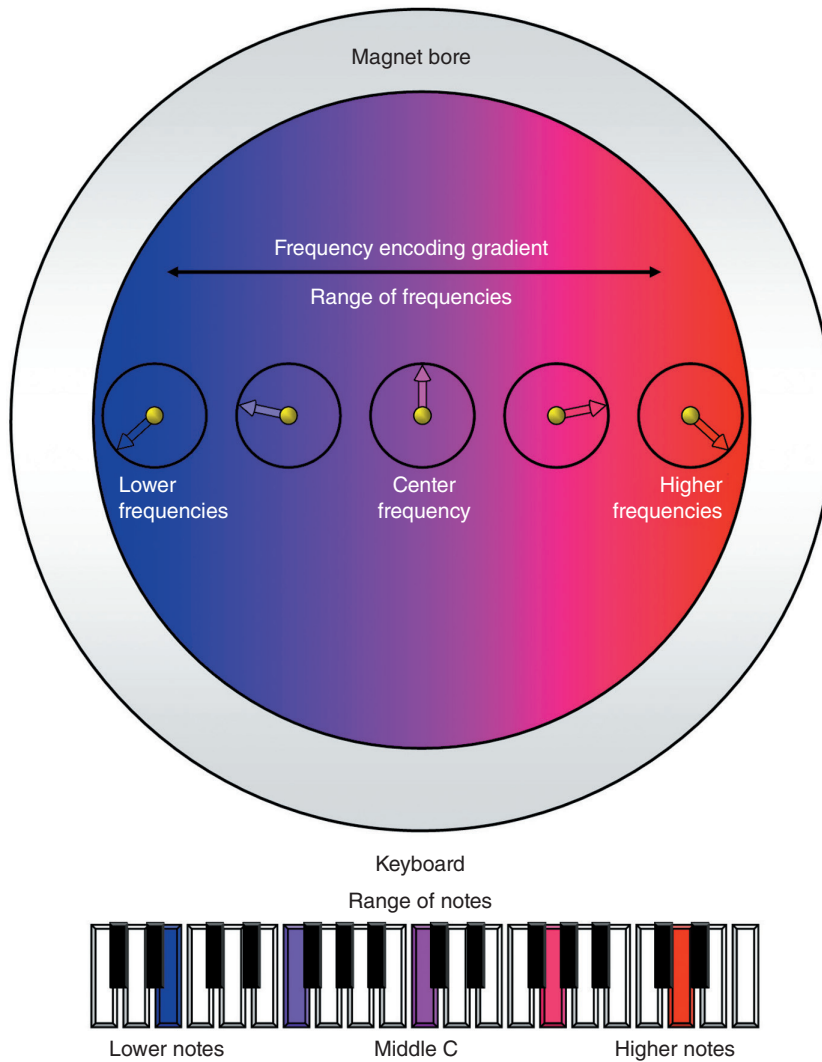


Figure 5.9 Frequency encoding and the keyboard analogy.

Scan tip:

Frequency encoding – what's going on behind the scenes?

The MRI system knows which gradient is allocated for frequency encoding by the direction of frequency encoding. When we select the frequency-encoding direction (head to foot, anterior to posterior, right to left), behind the scenes, we determine which gradient (x, y, or z) is switched on during the echo.

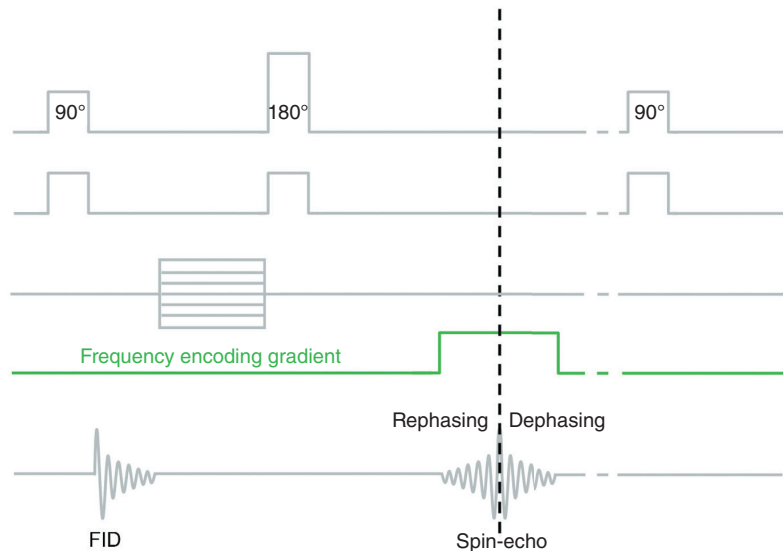


Figure 5.10 Timing of the frequency-encoding gradient in a spin-echo pulse sequence.

FOV and frequency encoding

As previously discussed, the amplitude of any gradient is determined by how much current passes through it when it is applied. In relation to the frequency-encoding gradient, the amplitude is one of the factors that determines the FOV in the frequency direction of the image. The steepness of the slope of the frequency-encoding gradient partly determines pixel resolution (see Chapter 6).

- To achieve a *small FOV* in the frequency direction, a steep frequency-encoding gradient is applied.
- To achieve a *large FOV* in the frequency direction, a shallow frequency-encoding gradient is applied.

A range of frequencies (called the **receive bandwidth**, see Chapter 6) is mapped across the FOV in the frequency direction. The frequency range within each pixel (called the bandwidth per pixel) is determined by the receive bandwidth for the whole FOV divided by the number of pixels in the frequency direction (called the **frequency matrix**, see Chapter 6).

Scan tip:

FOV(f) – what's going on behind the scenes?

There are a few parameters that alter the FOV in the frequency direction of the image. One of these is the amplitude of the frequency-encoding gradient. When we select the frequency FOV in the scan protocol, behind the scenes, we determine how much current is applied through the frequency

encoding gradient during the pulse sequence. A small FOV is produced by passing a large current through the frequency-encoding gradient coil and a large FOV by passing less current. For example, in the sagittal plane, the long axis of the anatomy lies head to foot in the z-axis. Therefore, behind the scenes, the z gradient is switched on during the echo to perform frequency encoding. If we select a small FOV, a large current is passed through the frequency-encoding gradient when it is applied. The large current produces a high-amplitude or steep gradient, and this creates a small FOV dimension in the frequency axis. Conversely, if we select a large FOV, behind the scenes, less current is passed through the frequency-encoding gradient when it is applied. Low current passing through the frequency-encoding gradient coil produces a low-amplitude or shallow gradient, and this creates a large FOV dimension in the frequency axis. This concept and other parameters that determine FOV and pixel size are explained in more detail in Chapter 6.

Table 5.5 Things to remember – frequency encoding.

Slices are frequency-encoded by applying a gradient along one axis of the two-dimensional image (usually the longest axis)
The frequency-encoding gradient is switched on during the echo. Typically, the peak of the echo occurs in the middle of the application of this gradient (it is often called the read or measurement gradient because the MRI system is reading or measuring the echo when it is switched on)
The frequency-encoding gradient changes the magnetic field strength and therefore the precessional frequency of the magnetic moments of hydrogen nuclei that lie along it
The change of frequency is measured and enables the system to spatially encode signal in the frequency-encoding direction of the image
The amplitude of the frequency-encoding gradient determines the size of the FOV in the frequency encoding axis of the image (see Chapter 6)
A steep frequency-encoding gradient produces a small FOV dimension in the frequency axes of the image
A shallow frequency-encoding gradient produces a large FOV dimension in the frequency axes of the image

PHASE ENCODING

How does it work?

Signal is located along the remaining short axis of the anatomy, and this localization of signal is called **phase encoding**. When the phase-encoding gradient is switched on, the magnetic field strength and, therefore, precessional frequency of magnetic moments of hydrogen nuclei located along the axis of the gradient are altered. As the speed of precession of the magnetic moments of hydrogen nuclei changes, so does their accumulated phase along their precessional path. Nuclei located at the high end of the phase-encoding gradient (shown in pink/red in Figure 5.11) experience a high magnetic field strength. The precessional frequency of these magnetic moments increases due to the presence of the gradient. This causes them to move further around their precessional path (or gain phase) than they would if the gradient had not been applied. Nuclei located at the low end of the phase-encoding gradient (shown in blue/




Figure 5.11 Phase encoding.

purple in Figure 5.11) experience a low magnetic field strength. The precessional frequency of these magnetic moments decreases due to the presence of the gradient. This causes them to lose phase relative to their phase position if the gradient had not been applied. Nuclei at magnetic isocenter experience no difference in the magnetic field strength when the gradient is switched on, so the precessional frequency and phase of their magnetic moments remain unchanged.

The system can now differentiate between nuclei positioned at different locations along the phase-encoding gradient because their magnetic moments have a different precessional phase. Signals produced by these nuclei are therefore different, and the computer of the MRI system uses these data to place these signals in different positions in the image (see Chapter 6).

Analogy: Phase encoding and the watch analogy



The watch analogy referred to in Chapter 1 is a very easy way of understanding how phase encoding works. Imagine a watch telling the time of 12 o'clock. The hour and minute hand are both located over the number 12. Assume that the position of the hour hand at this point is equivalent to the phase of a magnetic moment of a nucleus experiencing B_0 . When the phase-encoding gradient is switched on, the magnetic field strength, precessional frequency, and phase of the magnetic moments of nuclei change according to their position along the phase-encoding gradient. Magnetic moments of nuclei experiencing a higher magnetic field strength gain phase, i.e. move further around the watch to say 4 o'clock, because they travel faster while the gradient is switched on. Magnetic moments of nuclei experiencing a lower field strength lose phase, i.e. move back around the watch to say 8 o'clock, because they travel slower while the gradient is switched on. Magnetic moments of nuclei at magnetic isocenter do not experience a changed field strength, and their phase remains unchanged, i.e. 12 o'clock.

There is now a phase difference or shift between magnetic moments of nuclei positioned along the axis of the phase-encoding gradient. When the phase-encoding gradient is switched off, the magnetic field strength experienced by nuclei returns to the main field strength B_0 and therefore the precessional frequency of all the magnetic moments of hydrogen nuclei returns to the Larmor frequency. However, the phase difference between the nuclei remains. The nuclei travel at the same speed (frequency) around their precessional paths, but their phases or positions on the watch are different because the phase-encoding gradient was applied. Another way of saying this is that the phase difference between the magnetic moments of nuclei located along the gradient is remembered when the phase-encoding gradient is switched off. This difference in phase is used to determine the position of nuclei along the phase-encoding gradient.

147

We usually select the direction of phase encoding so that it encodes signal along the shortest axis of the anatomy. The purpose is to avoid aliasing (see Chapter 8). It may help to refer to the images in Chapter 2 to work out which gradient was used for each spatial encoding function. Always remember the patient is usually lying supine along the z-axis while on the table (in a superconducting system). Using this standard, it is easy to work out the long and short axis of anatomy.

- In *coronal* images, the short axis of the anatomy lies along the x-axis of the magnet (right to left of the patient), and, therefore, the *x gradient* performs phase encoding.
- In *sagittal* images, the short axis of the anatomy lies along the y-axis of the magnet (back to front of the patient), and, therefore, the *y gradient* performs phase encoding.
- In *axial* images, the short axis of the anatomy usually lies along the vertical axis of the magnet (back to front of the patient), and therefore, the *y gradient* performs phase encoding. However, in imaging of the head, the short axis of the anatomy usually lies along the right to left axis of the magnet (right to left of the patient), so in this case the *x gradient* performs phase encoding.

Scan tip:**Phase encoding – what’s going on behind the scenes?**

The MRI system knows which gradient is allocated for phase encoding by the direction of phase encoding. When we select the phase encoding direction (head to foot, anterior to posterior, right to left), behind the scenes, we determine which gradient (x, y, or z) is switched on during phase encoding.

148

When does phase encoding occur?

The phase-encoding gradient is usually switched on after the application of the RF excitation pulse (Figure 5.12). The purpose of the phase-encoding gradient is to impose a phase shift or difference across an axis of the patient. This phase shift occurs because the phase-encoding gradient causes a change in frequency of magnetic moments of spins located along the gradient, and this causes a change in their phase. This is true of any gradient but, in the case of phase encoding, it is the phase shift that is important. Once phase shift occurs, the phase-encoding gradient is switched off. The magnetic moments of hydrogen nuclei return to precessing at the Larmor frequency (because they all experience B_0 when the gradient is removed), but the phase shift remains. It remains until all signal is lost or another gradient is applied, so the phase-encoding gradient does not have to be applied at a specific time. In most pulse sequences, the phase-encoding gradient is applied as soon as possible after the RF excitation pulse to ensure that it is completed before other tasks in the sequence.

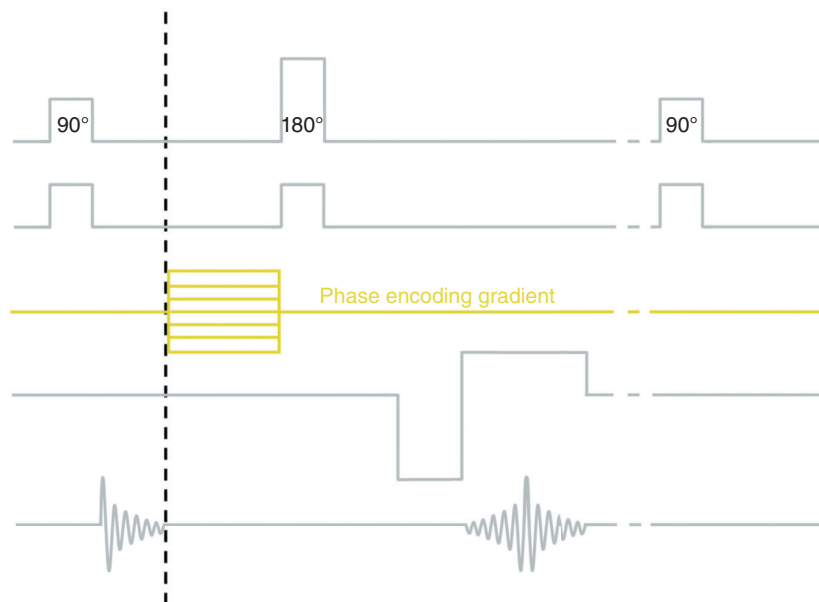


Figure 5.12 Timing of the phase-encoding gradient in a spin-echo pulse sequence.

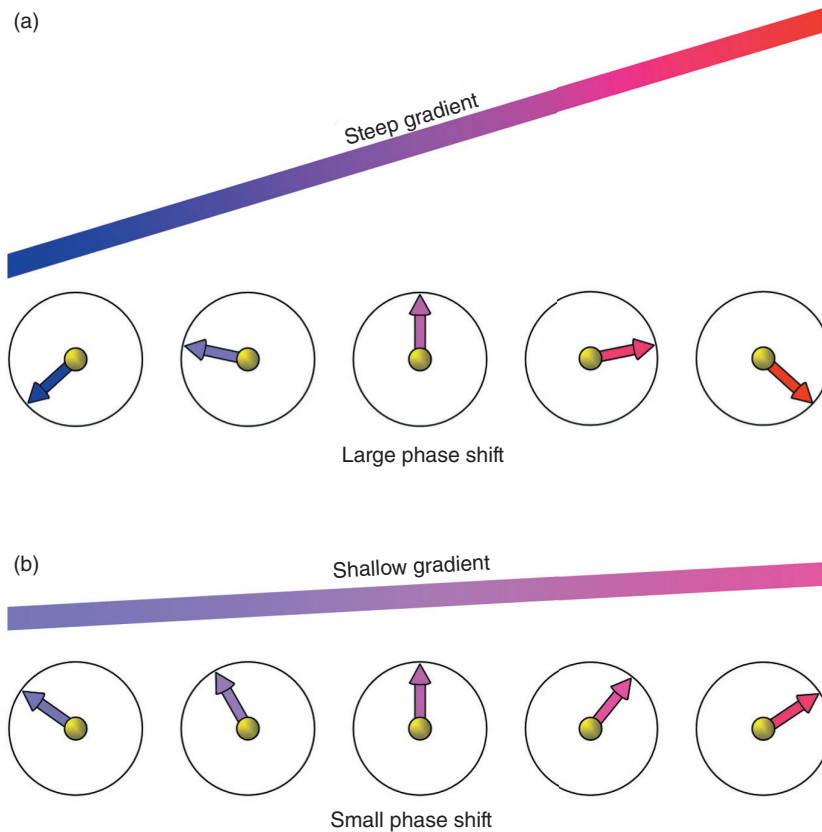


Figure 5.13 (a) Steep and (b) shallow phase-encoding gradients.

Phase matrix, phase resolution, and phase encoding

The amplitude of any gradient is determined by how much current is passed through it when it is switched on. The amplitude of the phase-encoding gradient determines the degree of phase shift between two points along the gradient (Figure 5.13). A steep phase-encoding gradient causes a large phase shift between two points along the gradient, for example, 8 o'clock and 4 o'clock, while a shallow phase-encoding gradient causes a smaller phase shift between the same two points along the gradient, for example, 10 o'clock and 2 o'clock, as shown in Figure 5.13.

Unlike the frequency-encoding gradient, which is switched on to the same amplitude each time it is applied, the amplitude of the phase-encoding gradient is changed throughout the pulse sequence. Its polarity is also changed and is often illustrated as a ladder in pulse sequence diagrams (see Figure 5.12). Different rungs or steps in the ladder are achieved by changing the amplitude and polarity of the phase-encoding gradient (see Chapter 6). The number of steps in the ladder or how many times the phase-encoding gradient is switched on to a

different amplitude or polarity during the sequence is determined by the **phase matrix** we select in the scan protocol.

- To achieve a *high phase matrix*, the phase-encoding gradient is applied many times during the pulse sequence.
- To achieve a *low phase matrix*, the phase-encoding gradient is applied fewer times during the pulse sequence.

The steepest slope of phase-encoding gradient, either positive or negative, determines pixel size in the phase direction of the image. This is the phase encoding slope that corresponds to either the highest or lowest step in the ladder [7].

- To achieve a *small pixel*, the highest or lowest step in the ladder is selected using a steep phase-encoding gradient.
- To achieve a *large pixel*, the highest or lowest step in the ladder is selected using a shallow phase-encoding gradient.

Scan tip: Phase matrix and resolution – what’s going on behind the scenes?

When we select the phase matrix in the scan protocol, behind the scenes, we determine how many times the phase-encoding gradient is switched on to a certain slope and polarity during the pulse sequence. If the phase matrix is high or fine, for example, 512, then, behind the scenes, the phase-encoding gradient is switched on 512 times to a different slope and/or polarity during the sequence. If it is low or coarser, for example, 128, then it is switched on fewer (128) times during the sequence.

The amplitude of the steepest application of the phase-encoding gradient determines pixel size in the phase encoding axis of the slice. How much current passes through the phase-encoding gradient when it is at its steepest point either negatively or positively determines pixel size. A small pixel is produced by passing a large current through the phase-encoding gradient and a large pixel by less current.

So, for example, in the coronal plane, the short axis of the anatomy lies right to left on the x-axis. Therefore, behind the scenes, the x gradient is switched on to perform phase encoding. If we select a small pixel, a large current is passed through the phase-encoding gradient when it is applied to its steepest point both negatively and positively. If a larger pixel is required, less current is passed through the phase-encoding gradient when it is applied to its steepest point both negatively and positively. The “behind the scenes” story so far is summarized in Table 5.7.

Table 5.6 Things to remember – phase encoding.

Slices are phase encoded by applying a gradient along one axis of the two-dimensional image (usually the shortest axis)
The phase-encoding gradient can be switched on at any time in a pulse sequence, but it is usually applied as soon as possible after the RF excitation pulse has been switched off
The phase-encoding gradient changes the magnetic field strength and therefore the precessional frequency and phase of the magnetic moments of hydrogen nuclei that lie along it
Once this change of phase has occurred, the phase-encoding gradient is switched off so that the magnetic moments of hydrogen nuclei precess at the Larmor frequency again but their phase difference remains
The magnetic moment of each spin therefore has a slightly different phase position to its neighbor along the gradient
The phase-encoding gradient is altered to a different amplitude and polarity during the pulse sequence. The number of times it is applied to a different amplitude determines the phase matrix
Its steepest application determines pixel resolution in the phase axes of the image (see Chapter 6)

Table 5.7 Image acquisition – summary of what is going on behind the scenes (so far).

Parameter	Behind the scenes
Slice thickness	Amplitude of the slice-select gradient and transmit bandwidth
Slice gap	Amplitude of the slice-select gradient and transmit bandwidth
Frequency FOV	Amplitude of the frequency-encoding gradient (also see Chapter 6)
Phase FOV	See Chapter 6
Frequency matrix	See Chapter 6
Phase matrix	Number of different phase-encoding steps
Phase resolution (pixel size)	Amplitude of the steepest phase-encoding gradient positively and negatively
Frequency resolution (pixel size)	See Chapter 6
NSA or NEX	See Chapter 6

Learning tip: Compensatory gradients

All gradients produce a change of frequency and a change of phase of magnetic moments of spins located at different positions along them. A change of phase is a natural consequence of changing frequency. If the speed of a watch is changed, then it naturally causes the hands of the watch to change position and tell a different time. In phase encoding, this phase change is needed to encode signal along the phase encoding axis of the image. However, this information is not required for slice-selection or frequency encoding. In these processes only the frequency change information is required. Therefore, in a pulse sequence diagram, extra compensatory gradients are usually illustrated on the slice-select and frequency-encoding gradients. These rephase the magnetic moments of the spins so that once the slice-select or frequency encoding process is finished there is no net phase change.

To summarize, every spatial location (pixel) within each slice is identified according to its phase and frequency. Its frequency is determined by its position along the frequency-encoding gradient. Its phase is determined by its position along the phase-encoding gradient.

Analogy: Using the watch analogy to understand spatial encoding



The watch analogy is a good way of remembering how all gradients encode. Imagine two people wearing watches that are synchronized and tell perfect time. They walk into the MRI scan room for 15 min. The magnetic field of the scanner affects the timekeeping of the watches because it magnetizes the hands of the watches. The person standing nearest to the magnet is affected the most because the magnetic field here is strongest. The person standing farthest away is affected to a lesser degree because the magnetic field here is less strong. If they then walk out of the room so they are no longer affected by the magnetic field, a stranger will be able to tell which person was standing nearer to the magnet and which was standing farther away simply by looking at their watches. This is because the hands of the watch of the person standing closer to the magnet are more out of phase from the synchronized time than the watch of the person standing farther away. In other words, the stranger has used the frequency and phase shift of the hands of the watch, produced because of applying a magnetic field to the watches, to spatially encode the relative positions of each person while they were in the room.

BRINGING IT ALL TOGETHER – PULSE SEQUENCE TIMING

In Chapters 3 and 4, we learned that pulse sequences consist of RF pulses, gradients, and intervening time periods. Some of these time periods are user-dependent (repetition time (TR) and echo time (TE) for example). It is very important to get to grips with what happens at different time points in the pulse sequence. Look at Figure 5.14, which illustrates gradient timing in a conventional spin-echo pulse sequence. Start on the left side of this diagram and work across to the right as you read the following description.

Every pulse sequence begins with an RF excitation pulse. This is always switched on at the same time as the slice-select gradient to selectively excite each slice. The plane of the scan (we select in the scan protocol) determines which gradient is used to select slices. The amplitude of this gradient in combination with the transmit bandwidth determines the slice thickness. In gradient-echo pulse sequences, this is the only time the slice-select gradient is applied. In spin-echo pulse sequences, it is switched on again each time an RF rephasing pulse is applied to ensure that spins in each slice are also rephased.

As soon as resonance occurs in the selected slice, the RF excitation pulse and the slice-select gradient are switched off. Magnetic moments of hydrogen nuclei within the selected slice return to precessing at the Larmor frequency of their resting state (the state before the RF excitation pulse was applied). They are also in phase. This is because the RF




Figure 5.14 Gradient timing in a spin-echo pulse sequence.

excitation pulse causes resonance of spins in the slice. Within about 10 ms, the magnetic moments dephase due to inhomogeneities in the magnetic field of the scanner producing a free induction decay signal (FID). Magnetization is still in the transverse plane, but it is not coherent and therefore returns no signal.

Usually, the phase-encoding gradient is the next gradient to be switched on in the sequence and is commonly applied along the shortest axis of the slice. The shape of anatomy within the slice indicates which is the shortest axis and therefore which gradient is chosen to perform phase encoding. This gradient produces a change of frequency and a change of phase of magnetic moments of hydrogen nuclei along its axis. The phase-encoding gradient is then switched off. All magnetic moments return to precessing at the Larmor frequency of their resting state, but the phase difference imposed by the gradient remains.

The amplitude and polarity of the phase-encoding gradient change throughout the scan. The simplest way of doing this is to change it every TR. Every TR, the current passed through the phase-encoding gradient changes, and this alters its amplitude. This means that a different phase shift is produced every TR. Data from these phase shifts are used later to encode signal along the phase encoding axis of the slice. Magnetization is still in the transverse plane at this point but it is out of phase. Some of the dephasing is caused by inhomogeneities, some by spin–spin interactions, and some by the phase-encoding gradient itself.

The next event in the pulse sequence is the rephasing mechanism. This is needed to rephase magnetic moments of at least some of the hydrogen nuclei in the slice to produce an echo. In spin-echo pulse sequences, this is achieved by a 180° RF rephasing pulse (see Chapter 3). In gradient-echo pulse sequences, a bipolar application of the frequency-encoding gradient is used for this purpose (see Chapter 4). The magnetic moments of hydrogen nuclei within the slice rephase and therefore produce an echo. These nuclei undergo T2 decay during the time from the RF excitation pulse to the peak of the echo. This time interval is called the TE and is one of the parameters we select in the scan protocol. Longer TEs allow more T2 decay to occur than short TEs when the echo is read (see Chapter 2).

The frequency-encoding gradient is then switched on during the echo. This is usually applied along the longest axis of the slice. The shape of anatomy within the slice indicates which is the longest axis and therefore which gradient is chosen to perform frequency encoding. This gradient produces a change of frequency along its axis. Data from this frequency change are used to encode signal along the frequency encoding axis of the slice. The peak of the echo is usually

located in the middle of the application time of the frequency-encoding gradient. This means that it is switched on when the magnetic moments of hydrogen nuclei are rephasing, when they are all in phase and reach a peak, and when they are dephasing again. While the echo occurs and the frequency-encoding gradient is switched on, the MR system reads the frequencies present in the echo. The echo is a waveform that consists of many different frequencies that reflect what happened during phase and frequency encoding in the slice in each TR period. Once the magnetic moments of the spins dephase, coherent transverse magnetization is lost, and the frequency-encoding gradient is switched off.

The system then usually repeats the above process for the next slice (excite, slice-select, phase encode, rephase, frequency encode, read echo) and then does the same to all the other slices sequentially after that. Once all the slices are encoded and their echoes are read, the whole sequence of events begins again by returning to the first slice. The time between RF excitations applied to a slice is the TR period. In the simplest variety of spin-echo or gradient-echo pulse sequence, every TR, each slice is encoded, and data from each echo from each slice are collected every TR. The TR is repeated until there are enough data to create an image.

Scan tip: What is the TR?

This might seem like an obvious question with an obvious answer! The TR is the time between RF excitation pulses but it is not the time between exciting a slice and then exciting the next slice in the imaging stack. It is the time between exciting slice 1, for example, then following that by exciting the next and every slice after that and then coming back to slice 1. In other words

- It is not the time between applying successive RF excitation pulses to the patient.
- It is the time between applying successive RF excitation pulses to each slice.

This explains why the TR determines how many slices are possible in a 2D acquisition. Longer TRs mean that there is more time to individually excite, rephase, and encode more slices. If the TR is short, there is less time to do this so fewer slices are possible.

Learning tip: The dead time and T1 recovery

Let's imagine for a moment that we excite, rephase, encode, and read the echo at time TE from slice 1 in the slice stack. We then do the same to all the other slices and clearly it takes some time to do this and return to slice 1. During the time the system is attending to all the other slices in the stack (called the **dead time**), T1 recovery occurs in slice 1. If you remember from Chapter 2, T1 recovery is much longer than T2 decay. T1 recovery occurs over 100 s and sometimes 1000 s of ms, and it is during the dead time, while the system is encoding and collecting data from all the other slices in the stack, that T1 recovery occurs in slice 1. How much T1 recovery there is depends on how long the system waits to apply the next RF excitation pulse to that slice. This is the TR. Longer TRs allow more T1 recovery to occur than short TRs. This explains why the TR determines not only how many slices are possible but also T1 contrast (see Chapter 2).




Figure 5.15 How a waveform is produced from a change of phase over distance.


For reasons discussed in the next chapter, the phase shift imposed by the phase-encoding gradient must be converted into a frequency. This is not as difficult as it sounds. The watch analogy explains how frequency is a change of phase over *time*. If we plot the change of phase of a hand on the watch as it precesses over time we generate a waveform. The phase, slice-select, and frequency-encoding gradients also cause a change of phase but this is over *distance* along a gradient. With phase encoding, this change of phase is extrapolated as a frequency by creating a waveform by connecting all the phase values associated with a certain phase shift along the gradient (Figure 5.15). This demonstrates a change of phase over distance rather than over time and is therefore called a **spatial frequency**.

The waveform has a spatial frequency that depends on the degree of phase shift produced by the gradient. Steep phase-encoding gradients produce large phase shifts across a given distance along the gradient and are able to encode high spatial frequencies, while low-amplitude phase-encoding gradients produce small phase shifts across the same distance and are able to encode low spatial frequencies (Figure 5.16). The frequency change imposed by the frequency-encoding gradient also produces spatial frequency information because the frequency of magnetic moments of spins depends on their spatial position along the gradient. There are also other frequencies that reflect the fact that each slice is excited with a bandwidth or range of frequencies. The echo therefore contains many spatial frequencies from the encoding process.

The multitude of different frequencies in the echo is similar to different musical notes produced by a piano keyboard. Each key is tuned to produce a certain note when pressed. Different notes are characterized by the fact that they resonate a piano wire at different frequencies so that note A, for example, has a different frequency from note C. Each note has a different position or spatial location on the keyboard. If all the keys are pressed at once, then a chord (of sorts) is produced. This chord is made up of many different frequencies. In addition, some of the keys might be pressed hard to produce a loud note and some more softly to produce a quiet note. Experienced pianists, on hearing this chord, would be able to figure out:

- which notes make up the chord (their frequencies);
- where on the keyboard each key (or frequency) is located (their spatial position);
- how hard each key was pressed (the amplitude of each note or frequency).

- (a) Steep phase encoding gradient, capable of determining high spatial frequencies



- (b) Shallow phase encoding gradient, capable of determining low spatial frequencies




Figure 5.16 Spatial frequency vs amplitude of phase-encoding gradient.

The experienced pianist who converts these frequencies to something the computer uses to create an MR image of each slice is called **Fast Fourier Transform (FFT)**. FFT is the math that gets us from frequencies in the echo to the image. To achieve this, spatial frequencies in the echo are digitized and stored until FFT is performed. This next step is discussed in detail in Chapter 6.



For questions and answers on this topic please visit the supporting companion website for this book: www.wiley.com/go/westbrook/mriinpractice

References

1. Dale, B.M., Brown, M.A., and Semelka, R.C. (2015). *MRI: Basic Principles and Applications*, 5, 26. Wiley.
2. Odaibo, S.G. (2012). *Quantum Mechanics and the MRI Machine*, 26. Arlington, VA: Symmetry Seed Books.
3. McRobbie, D.W., Moore, E.A., Graves, M.J. et al. (2017). *From Picture to Proton*, 3, 108. Cambridge: Cambridge University Press.

4. Odaibo, S.G. (2012). *Quantum Mechanics and the MRI Machine*. Arlington, VA: Symmetry Seed Books, 24.
5. Dale, B.M., Brown, M.A., and Semelka, R.C. (2015). *MRI: Basic Principles and Applications*, 5, 28. Wiley.
6. Dale, B.M., Brown, M.A., and Semelka, R.C. (2015). *MRI: Basic Principles and Applications*, 5, 29. Wiley.
7. McRobbie, D.W., Moore, E.A., Graves, M.J. et al. (2017). *From Picture to Proton*, 3, 117. Cambridge: Cambridge University Press.

6

k-Space

Introduction	158	Part 3: Some important facts about <i>k</i> -space!	184
Part 1: What is <i>k</i> -space?	159	Part 4: How do pulse sequences fill <i>k</i> -space?	197
Part 2: How are data acquired and how are images created from these data?	165	Part 5: Options that fill <i>k</i> -space	199

After reading this chapter, you will be able to:

- Describe the characteristics of *k*-space.
- Explain different ways in which *k*-space is filled.
- Understand how pulse sequences determine how and when *k*-space is filled with data.
- Apply this understanding when altering parameters in the scan protocol.

INTRODUCTION

In Chapter 5, we discovered that spatial encoding selects an individual slice and produces a frequency shift of magnetic moments of spins along one axis of the slice and a phase shift along the other. The system now has a way of locating an individual signal within the image by measuring the number of times the magnetic moments of spins cross the receiver coil (frequency) and their position around their precessional path (phase).

The phase shift caused by the phase-encoding gradient creates a spatial frequency. This is because a waveform is derived from plotting the change of phase of magnetic moments of hydrogen nuclei located at different locations along the gradient. Spatial frequencies are also obtained from slice-selection and frequency encoding because they too cause frequencies that are dependent on spatial position. Data from these spatial frequencies are used by Fast Fourier

Transform (FFT) mathematics to produce an image. During the scan, data are acquired and stored in **k-space**. *k*-Space is a spatial frequency domain, i.e. where information is stored about the frequency of a signal and where it comes from.

Learning tip: What does the “*k*” stand for?

Why is *k*-space called that? Why not “*z*-space” or “*w*-space”? The reason is that *k* denotes spatial frequency (strictly speaking it refers to the angular wavenumber, which is a function of spatial frequency and wavelength, but for our purposes it’s acceptable to just use the term spatial frequency). *k*-Space, therefore, is a storage space for spatial frequencies. The units of *k*-space are rad/m (or cm). This is because spatial frequency is represented as a phase change over distance along a gradient. The units of phase are **radians** (units of degrees in a circle) and the units of distance are m (or cm). This is different from the standard unit of frequency, rad/s or Hz. These units are appropriate when looking at the magnetic moment of a single spin because, in that context, the phase change of the magnetic moment is measured as it precesses over time. Spatial frequency is different because it measures the change of phase between the magnetic moments of a row of spins along the gradient.

159

Every parameter we select in the scan protocol changes how *k*-space is filled with data. *k*-Space is therefore a very important concept. This chapter is broken down into the following five parts, and scan tips are used to link the theory of *k*-space to practice.

1. What is *k*-space?
2. How are data acquired and how are images created from these data?
3. Some important facts about *k*-space!
4. How do pulse sequences fill *k*-space?
5. What are the different ways in which *k*-space is filled?

The following discussion is a jigsaw puzzle made up of these five parts. You need the whole puzzle to get the whole picture!

PART 1: WHAT IS *k*-SPACE?

k-Space is a storage device. It stores digitized data produced from spatial frequencies created from spatial encoding (see Chapter 5). Figure 6.1 illustrates *k*-space for *one slice*. *k*-Space is rectangular and has two axes perpendicular to each other. The frequency axis of *k*-space is horizontal and is centered in the middle of several horizontal lines. Data from frequency encoding are positioned in *k*-space along this axis. The phase axis of *k*-space is vertical and is centered in the middle of *k*-space perpendicular to the frequency axis. Data from phase encoding are positioned in *k*-space along this axis. In the simplest method of *k*-space filling, data are stored in horizontal lines that are parallel to the frequency axis and perpendicular to the phase axis of *k*-space.





Figure 6.1 The axes of *k*-space.


**Analogy:
The chest of drawers**




k-Space is analogous to a chest of drawers. Look at Figure 6.2, in which *k*-space with its horizontal lines is illustrated. These lines look like drawers in a chest of drawers, which, like *k*-space, is a storage device. The number of drawers corresponds to the number of lines of *k*-space that are filled with data to complete the scan. The chest of drawers analogy is referred to many times in this book. Look out for the chest of drawers symbol in the margin.



Diagrammatic



Data



The chest of drawers

Figure 6.2 *k*-space – the chest of drawers.

Let's for a moment imagine a room with a very high ceiling. Inside this room there are 20 chests of drawers, that extend from floor to ceiling. In each chest of drawers, there are 100 drawers and in the middle of the room is a large pile of clothes of different types – all jumbled up together. Imagine that you are tasked with sorting these clothes into different drawers and that the aim is to ensure that each drawer, in each of the chests of drawers, contains the same type and quantity of clothes. For example, the top drawer of every chest of drawers in the room contains 50 pairs of green socks. The next drawer down of every chest of drawers contains 50 red shirts, and the next, 50 yellow pairs of trousers, and so on. The task is complete when the pile of clothes in the center of the room is sorted into all the correct drawers of every chest of drawers in the room.

The information that you would probably need to complete this task (now called the clothes-sorting exercise) is as follows:

1. The total number of chests of drawers in the room
2. The total number of drawers in each chest of drawers
3. Numbered drawers so that each is quickly identifiable
4. The most efficient system to fill the drawers with the required quantity and type of clothing.

The MRI computer system needs the same information to fill k -space as you would need to complete the clothes sorting exercise. A single chest of drawers is analogous to k -space for a single slice. All the chests of drawers represent all the selected slices. The drawers inside a chest of drawers are equivalent to the lines of k -space filled with data from echoes from that slice. The large pile of jumbled-up clothes represents all the spatial frequencies, from all the echoes, from all the slices produced from spatial encoding during the scan. The task is to ensure that the data from these spatial frequencies are sorted out and located correctly in each drawer or line of k -space. So, what are the answers to the questions above?

1. The number of chests of drawers in total equals the number of k -space areas, and this equals the total number of selected slices. This is a parameter selected in the scan protocol (20 in the example above).
2. The number of drawers or lines of data equals the phase matrix selected in the scan protocol (100 in the example above). If a phase matrix of 256 is selected, then 256 lines or drawers are filled with data to complete the scan. If a phase matrix of 128 is selected, then 128 lines or drawers are filled with data to complete the scan (Figure 6.3).
3. The lines of k -space are numbered so that the system always knows where it is in k -space. The following convention is used. The lines are numbered with the lowest number near to the central axis (e.g. lines $\pm 1, 2, 3, \dots$) and the highest numbers toward the outer edges (e.g. $\pm 128, 127, 126, \dots$) (Figure 6.4). The lines in the top half of k -space are called positive lines, and those in the bottom half are called negative lines. This is because the line filled with data is determined by the polarity of the phase-encoding gradient. Positive polarity phase-encoding gradients are associated with lines in the top half of k -space, whereas negative polarity phase-encoding gradients are associated with lines in the bottom half of k -space.

4. There are many different k -space filling methods. The computer uses different systems depending on parameters we select in the scan protocol. The simplest method is called **Cartesian filling**. In this method, k -space is filled in a linear manner from top to bottom or bottom to top.

162




Figure 6.3 k -Space – phase matrix and the number of drawers.




Figure 6.4 *k*-Space – labeling.

Learning tip: How the phase-encoding gradient fills *k*-space

In Chapter 5, we discovered that the slope and sometimes polarity of the phase-encoding gradient are altered throughout the scan. This is necessary to fill different lines of *k*-space with data. The phase-encoding gradient picks which line of *k*-space, or which drawer is filled with data, in a certain repetition time (TR) period. Positive polarity phase-encoding gradients pick lines in the top half of *k*-space; negative polarity phase-encoding gradients pick lines in the bottom half. In addition, the slope of the phase-encoding gradient determines which line is selected. Steep gradients, both positive and negative, select the most outer lines, while shallow gradients select the central lines. If the slope of the phase-encoding gradient decreases from its steepest to shallowest amplitude, the system steps down through lines of *k*-space from the most outer lines to more central lines (see Figure 6.4). The opposite is true if the slope of the phase-encoding gradient increases from its shallowest to steepest amplitude.

If the phase-encoding gradient is not changed, then the same line is filled every TR, i.e. only one line is filled with data. As the number of lines filled determines the phase matrix, leaving the phase-encoding gradient unchanged throughout the scan results in an image with a phase matrix of one. Therefore, both the polarity and the slope of the phase gradient are altered every TR to produce multiple pixels in the phase direction of the image.

Using the Cartesian method and the chest of drawers analogy, look at Figure 6.5, which illustrates a conventional spin-echo pulse sequence. The top half of the diagram illustrates when gradients are applied during the pulse sequence. The bottom half shows *k*-space for a single slice, drawn as a chest of drawers. It may also help at this point to refer to the end of Chapter 5 when we explored the timing of different elements within a pulse sequence.

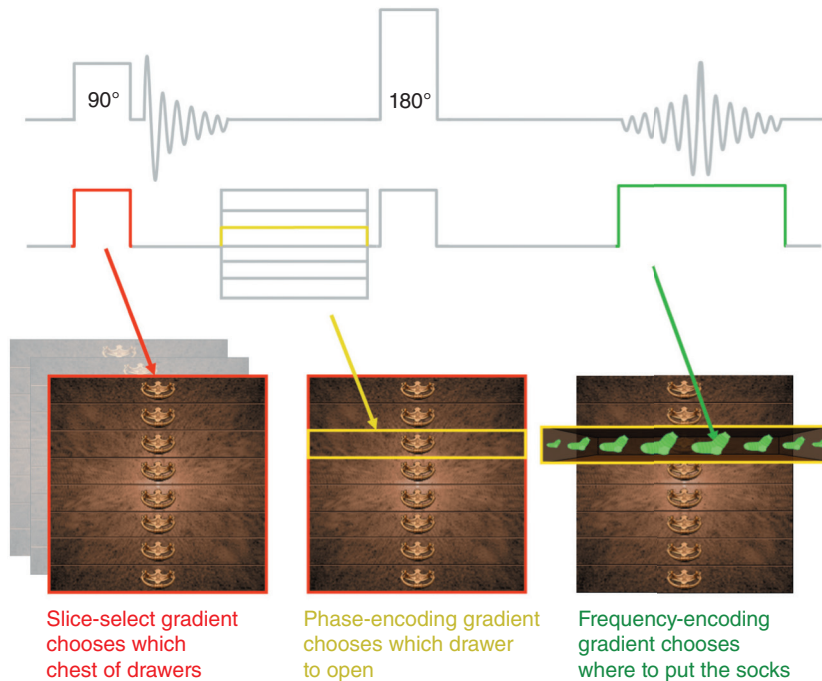


Figure 6.5 *k*-Space filling in a spin-echo sequence.

The slice-select gradient is applied during the RF excitation and rephasing pulses to selectively excite and rephase a slice. The slice-select gradient determines which slice is excited or which chest of drawers is selected. Each slice has its own area of *k*-space or chest of drawers. Note that although three chests of drawers are shown in Figure 6.5, they do not represent *k*-space for three separate slices in this diagram. In Figure 6.5, each chest of drawers represents the *same* slice at three different times in the sequence when each of the three gradients is switched on.

The phase-encoding gradient is then applied. This determines which line or drawer is filled with data. In the example shown in Figure 6.5, the third drawer down in the chest of drawers is opened. Let us say that this is equivalent to line +126 in *k*-space. To do this, the phase-encoding gradient is applied positively and steeply. Application of this gradient selects line +126 in *k*-space. The frequency-encoding gradient is then switched on. During its application, frequencies in the echo are digitized, and data are acquired and located in line +126 in *k*-space. These data are arranged in **data points**, and they are equivalent to pairs of green socks in the clothes sorting exercise (or red shirts, or yellow pairs of trousers!). The number of data points is determined by the frequency matrix (more on this later). After data points are collected and located in line +126, the frequency-encoding gradient switches off, and the slice-select gradient is applied again to excite and rephase the next slice (slice 2). This is equivalent to walking up to another chest of drawers (not shown in Figure 6.5). The phase-encoding gradient is applied again to the same polarity and amplitude as for the first slice, filling line +126 in *k*-space for slice 2. The process is repeated for slice 3 and all the other slices, with the same line filled for each *k*-space area.

Once this line is filled in every chest of drawers, the TR is repeated. The slice gradient selects slice 1 (or chest of drawers 1) again, but this time a different line of *k*-space is filled to that

Table 6.1 Things to remember – *k*-space.

<i>k</i> -Space stores information about where frequencies within the slice are located
Data points acquired over time are laid out in <i>k</i> -space during the scan and mathematically converted into information related to amplitude via FFT
<i>k</i> -Space is analogous to a chest of drawers where the number of lines filled is the number of drawers in the chest of drawers
Each gradient determines when and how the chest of drawers is filled
In a standard sequence, the same drawer is filled for each of the chest of drawers in a single TR period
The number of drawers equals the phase matrix
The number of socks or data points in each drawer equals the frequency matrix

in the previous TR. If the Cartesian *k*-space filling method is used, either the line above or the line below is filled. In our example, let's assume this is the line below that is labeled +125. To do this, the phase-encoding gradient is switched on positively but to a shallower slope than in the previous TR. This opens the next drawer down in the chest of drawers, or selects line +125. When the frequency-encoding gradient is switched on, data points are laid out in this line. When this is complete, the slice-select gradient is applied again to select slice 2. The same amplitude and polarity of phase-encoding gradient are applied to open the same drawer or pick the same line (+125) in *k*-space for slice 2.

This process is repeated for all the slices. As the pulse sequence continues, every TR the phase encoding amplitude is gradually decreased to step down through the lines of *k*-space until the center line is filled. To fill the bottom lines, the phase-encoding gradient is switched negatively with gradually increased amplitude every TR to progressively fill the outer lines. The central line (line 0) is always filled with data. To do this, the phase-encoding gradient is not switched on. In Cartesian *k*-space filling, if, for example, a 256 phase matrix is selected, the system fills 128 lines in the top half of *k*-space, line 0 and 127 lines in the bottom half of *k*-space (total 256). This is written as (+128, 0, -127). The system can, however, just as easily start at the bottom and work its way up through *k*-space, in which case it is written as (-128, 0, +127). This is the most common *k*-space filling method, although there are many others, and some of these are discussed in Part 5.



**Refer to animation 6.1 on the supporting companion website for this book:
www.wiley.com/go/westbrook/mriinpractice**

PART 2: HOW ARE DATA ACQUIRED AND HOW ARE IMAGES CREATED FROM THESE DATA?

This section specifically explores how spatial frequencies are converted into data points stored in *k*-space and how they create images. What are the socks in our clothes-sorting exercise? This is a difficult subject and one you may need some time to learn. However, it is important, as it affects several parameters we always select in the scan protocol.

As we learned earlier, a waveform is created by plotting the change of phase of magnetic moments either over time or over distance. Figure 6.6 shows three waveforms. These waveforms represent the change of phase over time of the magnetic moments of three spins precessing at three different frequencies (1, 2, and 4 Hz). The echo received by the receiver coil at time TE contains hundreds of different frequencies, and, unlike those shown in Figure 6.6, they also have different amplitudes. These frequencies are dependent on where signal is coming from in the slice during the scan. This is complicated because there are hundreds of different frequencies and many different amplitudes.

166

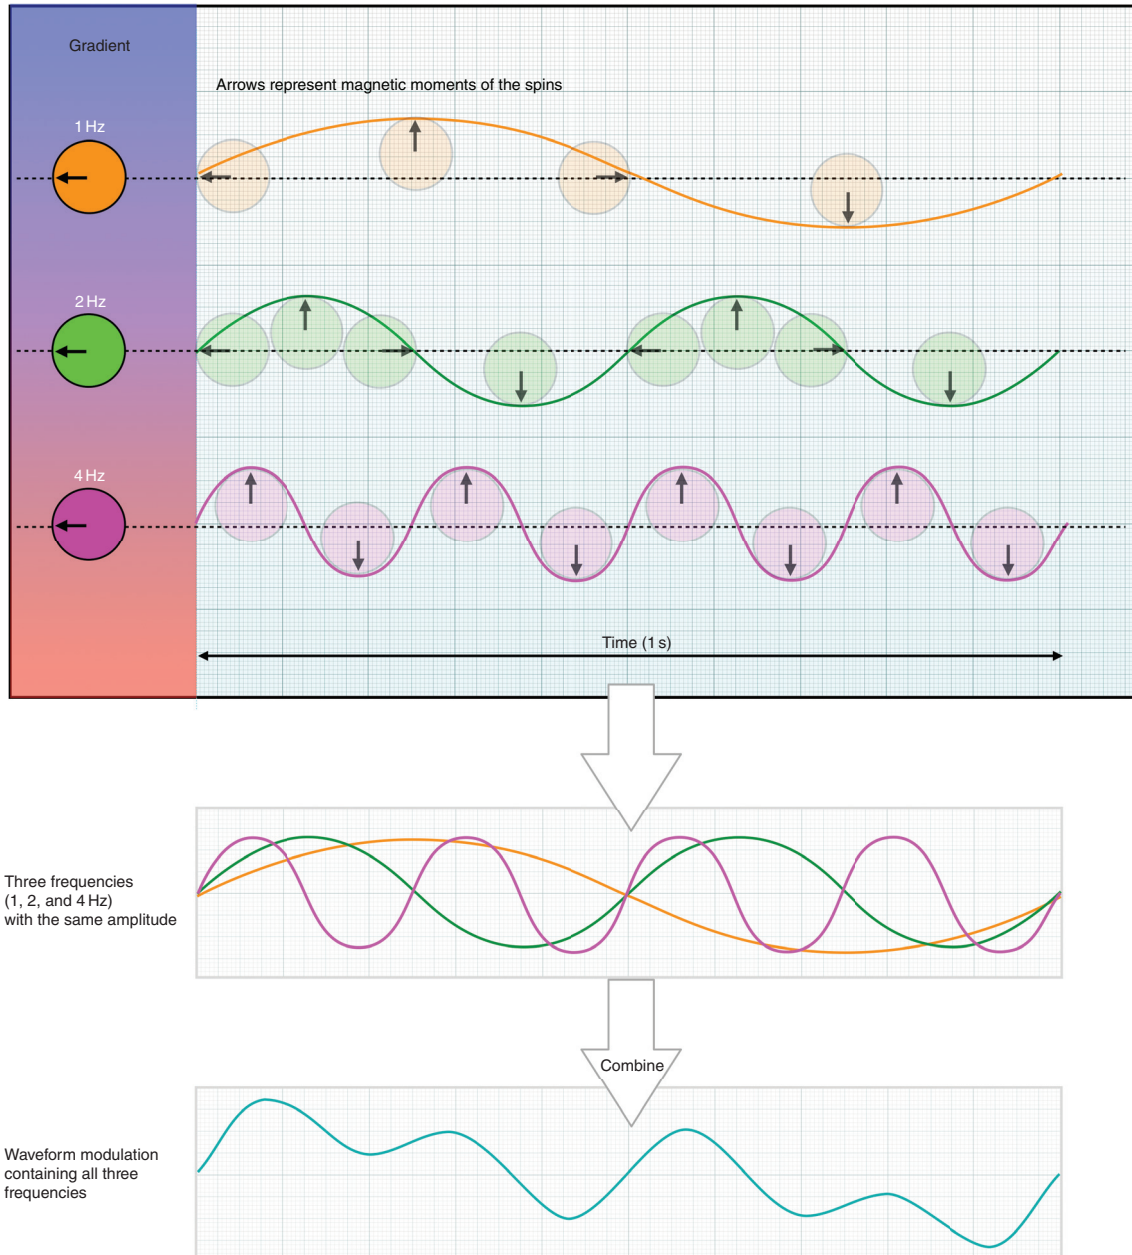


Figure 6.6 Waveforms from three different magnet moments precessing at three different frequencies and their amplitude modulation.

A few preliminary steps are needed to ensure that the data are in a format required for FFT mathematics to create an image of each slice. The first step is to simplify the frequencies and amplitudes present in the echo. This step is called **frequency** and **amplitude modulation**. The multiple waveforms in the echo are simplified into one waveform that represents the average amplitude of all the different frequencies in the echo at different time points (see bottom of Figure 6.6). This modulation is still a waveform (although it does not perhaps look like one). The echo that occurs at time TE is this modulation. The echo is a symmetrical representation of signal, and the peak of the echo usually occurs in the middle of the application of the frequency-encoding gradient (see Chapter 5). The next step is digitizing the modulation (called the echo from now on). This process is called **analog-to-digital conversion (ADC)**. Analog is a term used for information about a variable illustrated as a waveform. Digitization displays the same information but in binary numbers. On modern scanners, digitization of analog signal is done inside either the receiver coil assembly or the body of the MRI scanner (see Chapter 9).

Analogy: Using the watch analogy to understand ADC



Watches display the variable of time in two different ways. Watches that display the time by moving hands around a watch face are called analog watches. A waveform from any hand on a watch is created by plotting the change of phase of the hand as it moves around the watch face over time. A watch that displays time as a series of numbers is called a digital watch. The measured variable is still time, but digital watches display it as numbers. These numbers are acquired by breaking down time into bite-size chunks of data.

The frequency-encoding gradient is switched on while the system reads the echo and digitizes it. It is therefore sometimes called the **readout** or **measurement gradient**. The duration of the frequency-encoding gradient is called the sampling time, sampling window, or **acquisition window** (called the sampling window from now on). During the sampling window, the system samples or measures the echo at certain time points. Every time a sample is taken, this is stored as a data point in *k*-space. These are the pairs of socks used in the clothes-sorting exercise.

During the sampling window, several data points are acquired as the system samples the echo, and, in the simplest method, these are laid out in a line of *k*-space from left to right. The number of data points is determined by the frequency matrix. If the frequency matrix is 256, then 256 data points are acquired during the sampling window. If the frequency matrix is 512, then 512 data points are acquired. Every TR, a different line of *k*-space is filled with the same number of data points determined by the frequency matrix. This continues until all the lines are full of data points when the scan ends.

The process of data acquisition therefore results in a grid of data points:

- The number of data points horizontally in each line equals the frequency matrix, e.g. 512, 256, 1024, etc.
- The number of data points vertically in each column corresponds to the phase matrix, e.g. 128, 256, 384, 512, etc. (Figure 6.7).




Figure 6.7 Data points in k -space. The number in each column is the phase matrix. The number in each row is the frequency matrix.

Scan tip:

Phase and frequency matrix – what’s going on behind the scenes?

The number of k -space areas needed to complete the scan depends on the number of slices in the slice stack. When we prescribe slices in the scan protocol, behind the scenes, we determine the number of different k -space areas in the array processor. This is the total number of chests of drawers in the clothes-sorting exercise.

The number of lines required to complete the scan depends on the phase matrix (i.e. the number of pixels displayed in the phase direction of the image). When we select the phase matrix in the scan protocol, behind the scenes, we determine how many lines are filled with data in each k -space area. This is the total number of drawers in each chest of drawers in the clothes-sorting exercise.

The number of data points in each line of k -space depends on the frequency matrix (i.e. the number of pixels displayed in the frequency direction of the image). When we select the frequency matrix in the scan protocol, behind the scenes, we determine how many data points are laid out in each line of k -space. This is the number of pairs of socks in each drawer of the chest of drawers in the clothes-sorting exercise.

Sampling

During the sampling window, the system samples or measures spatial frequencies in the echo. Every sample results in a data point that is placed in a line of k -space. Each data point contains information about spatial frequencies in the echo at different time points during the sampling window.

Learning tip: Data points vs photographs

169

If it's easier, think of data points as photographs. Photographs are a record of what is happening to the objects within the photograph at a certain moment in time. If we take several photographs of the same moving object, each photograph contains information about the object at different time points as it moves. This is like different frames in a video. Similarly, data points are measurements of spatial frequencies in the echo at different time points during the sampling window.

The rate at which sampling occurs is called the **digital sampling rate** or **digital sampling frequency**. It is the rate at which data points are acquired per second during the sampling window and has the unit Hz. If one data point is acquired per second, the digital sampling frequency is 1 Hz, 100 data points per second is 100 Hz, 1000 per second is 1000 Hz or 1 KHz, and so on. The digital sampling frequency is not directly selected in the scan protocol, but it affects several other parameters that are selected, so read on!

The digital sampling frequency determines the time interval between each data point. This time interval is called the **sampling interval** and is calculated by dividing the digital sampling frequency by 1 (Equation (6.1)). This relationship means that if the digital sampling frequency increases, then the sampling interval decreases and vice versa. For example:

- If the digital sampling frequency is 32 000 Hz (32 KHz), the sampling interval is 0.031 ms ($1 \div 32\ 000$).
- If the digital sampling frequency halves to 16 000 Hz (16 KHz), the sampling interval doubles to 0.062 ms ($1 \div 16\ 000$).

Equation 6.1

$$\omega_{\text{sampling}} = 1/\Delta T_s$$

ω_{sampling} is the digital sampling frequency (KHz)
 ΔT_s is the interval between each data point or sampling interval (ms)

This equation shows that if the digital sampling frequency increases, the sampling interval decreases

The digital sampling frequency is important. If it is too low, then there may not be enough data points in k -space to create an accurate image. If it is too high, then the resultant files might be large and unmanageable, and sampling might take too long. In addition, some of the frequencies sampled in the echo are unwanted noise frequencies. As the digital sampling frequency increases, more noise data are acquired, and this affects image quality (see Chapter 7).




Figure 6.8 The Nyquist theorem.

The most optimum digital sampling frequency is determined by the **Nyquist theorem**. (The full name is Whittaker–Kotelnikov–Shannon–Raabe–Someya–Nyquist theorem [1]. Thankfully, it is abbreviated to the Nyquist theorem!) This calculates the minimum digital sampling frequency needed to acquire enough data points to create an accurate image. The Nyquist theorem states that when digitizing a modulation of several frequencies, the highest frequency present in the modulation must be sampled at least twice as frequently to accurately digitize or represent it.

Look at Figure 6.8. Sampling once per cycle, or at the same frequency as the frequency we are trying to digitize, results in a representation of a straight line or an absent frequency in the data (middle diagram). Sampling at less than once per cycle represents a completely incorrect frequency that leads to an artifact called **aliasing** (bottom diagram) (see Chapter 8). Sampling twice per cycle, or at twice the frequency we are trying to digitize, results in the correct representation of that frequency in the data (top diagram).

As long as the highest frequency present is sampled twice, it is represented correctly in the data. Lower frequencies are sampled more often at the same digital sampling frequency and are also represented accurately in the data. Digital sampling frequencies higher than this produce more data and therefore a more accurate representation of the original analog frequencies. However, due to time constraints, the digital sampling frequency is sometimes limited. High digital sampling frequencies also result in more noise data. The digital sampling frequency is therefore usually kept at just twice the highest frequency in the modulation

(termed the **Nyquist frequency**) to avoid aliasing while still sampling in the most time-efficient manner. The digital sampling frequency is usually limited, therefore, to twice the Nyquist frequency (Equations (6.2)–(6.4)).

Equations 6.2–6.4: Sampling equations.

$$\omega_{\text{sampling}} = 2 \times \omega_{\text{Nyquist}}$$

ω_{sampling} is the digital sampling frequency (KHz)
 ω_{Nyquist} is the Nyquist frequency (KHz) – the highest frequency in the echo that can be sampled

If Nyquist is obeyed, the highest frequency is sampled twice as fast as the Nyquist frequency, and this determines the digital sampling frequency

$$\text{RBW} = 2 \times \omega_{\text{Nyquist}}$$

therefore

RBW is the receive bandwidth (KHz)
 ω_{Nyquist} is the Nyquist frequency (KHz)

The receive bandwidth is the range of frequencies sampled on either side of center frequency. The RBW is therefore twice the highest frequency sampled

$$\omega_{\text{sampling}} = \text{RBW}$$

ω_{sampling} is the digital sampling frequency (KHz)
 RBW is the receive bandwidth (KHz)

Combining the first two equations shows that when Nyquist is obeyed, the receive bandwidth has the same numerical value as the digital sampling frequency



Refer to animation 6.2 on the supporting companion website for this book:
www.wiley.com/go/westbrook/mriinpractice

Analogy: Using the sprinter analogy to understand sampling

This difficult concept is perhaps better understood by using the following analogy. Imagine you are asked to take several photographs of a sprinter during a race. You would need the following information before you could start:

- How many photographs must you take in total?
- How many photographs per second can your camera take?
- How long is the race?

Each photograph is equivalent to a data point, as it is effectively a sample or measurement that shows the position of the runner's arms and legs at certain time points in the race.

- The total number of photographs is equivalent to the frequency matrix.
- The time you have available to take the photographs or the length of the race is equivalent to the sampling window.
- How many photographs per second your camera can take is equivalent to the digital sampling frequency.

For example, if the digital sampling frequency is one photograph per second and the race is 10 s long, 10 photographs can be taken, so using this example, a frequency matrix of 10 is achieved, thus:

- The digital sampling frequency is 1/s (1 Hz).
- The number of samples taken is 10.
- The sampling interval is 1 s.
- The sampling window is 10 s.

172

If the digital sampling frequency increases, then the time interval between each photograph is shorter so you can take 10 photographs in a much shorter time. For example, if you use a camera that can take two photographs per second, then the sampling interval is 0.5 second and you acquire 10 photographs in 5 s instead of 10 s.

The digital sampling frequency is not a parameter directly selected in the scan protocol. However, when Nyquist principles are obeyed, there is a selectable parameter that has the same *numerical value* as the digital sampling frequency. This is called the **receive bandwidth**.

Receive bandwidth

The bandwidth (or range of frequencies) used in the RF excitation pulse to excite a slice is called the transmit bandwidth (see Chapter 5). The transmit bandwidth permits slices of a certain thickness. The *receive* bandwidth is the range of frequencies accurately sampled or digitized during the sampling window. The receive bandwidth is determined by applying a filter to the frequencies encoded by the frequency-encoding gradient. This is achieved by selecting the center frequency of the echo and defining the upper and lower limits of frequencies that are accurately digitized on either side of this center frequency.

Learning tip: Larmor frequency vs bandwidth frequency

There are 55 radio stations in the St. Augustine area of Florida. Examples are Radio WSOS, which transmits at a frequency of 103.9 MHz, Radio WAYL at 91.9 MHz, and Radio WFCF at 88.5 MHz. These are their frequency modulations (FM). However, we cannot hear these frequencies because they are too high. Audible frequencies range from 1000 to 2000 Hz (1–2 KHz) [2]. Unfortunately, radio stations cannot transmit in the audible range because the bandwidth is too narrow. There is insufficient range of different frequencies to separate out signals from the different radio stations. If radio stations transmitted in kilohertz, there would be interference between the radio stations WSOS, WAYL, and WFCF.

To overcome this problem, the narrow audible range of frequencies from 1 to 2 KHz is modulated onto a center frequency, which is much higher (on the megahertz scale). Modulation means

that a low-frequency “piggy backs” onto a higher frequency. This means that the lower audible frequency can be transmitted over long distances, and signals from different radio stations are easily separated because a broader range of frequencies are permitted on the megahertz scale. The center high frequency is called a carrier frequency, and the lower audible frequency is modulated as a narrow band of frequencies onto this. In the case of radio stations, a bandwidth of audible low frequencies, from music or the voice of the radio presenter, is modulated onto a much higher carrier frequency, e.g. 103.9 MHz for Radio WSOS. The signal is then transmitted and, in the radio of the listener, demodulated back to audible low frequencies so that music and voices are heard.

The same principles apply in MRI. The Larmor frequency, which is on the megahertz scale in clinical MRI, can be too high for some systems to either transmit RF or receive signal. To overcome this problem, a range of low frequencies are centered on the much higher carrier frequency of spins in the slice. RF is transmitted at the center frequency of the slice, which is on the megahertz scale, but with a transmit bandwidth of just a few kilohertz on either side of the carrier frequency (see Chapter 5). When it is time to receive frequencies in the echo, a similar range of low frequencies are received or sampled. The carrier frequency is the center frequency of the echo and on the megahertz scale. Although most modern systems can now accurately sample at very high frequencies, the MR signal is still usually demodulated down from the higher carrier frequency to a lower frequency [3]. This is commonly in the order of several kilohertz on either side of the carrier frequency and is the receive bandwidth. The highest frequency present in the modulation (the Nyquist frequency) is therefore the highest frequency in the receive bandwidth, not the much higher carrier frequency.

The receive bandwidth is equivalent to representing frequencies of half the receive bandwidth above the center frequency to half the receive bandwidth below the center frequency. For example, if the receive bandwidth is 32 KHz across the whole echo, then this represents 16 KHz above the center frequency of the echo to 16 KHz below. These frequencies are mapped across the field of view (FOV) after FFT.

Using the example above, the highest frequency in the echo is 16 KHz above the center frequency of the echo. This is the Nyquist frequency. The receive bandwidth is therefore twice the Nyquist frequency ($2 \times 16 = 32$ KHz). As we have just learned, the digital sampling frequency is also twice the Nyquist frequency. Therefore, although the receive bandwidth and the digital sampling frequency are different parameters, they have the same numerical value (Equations (6.2)–(6.4)) when Nyquist principles are applied.

As the receive bandwidth is a selectable parameter in the scan protocol, it is used to determine the digital sampling frequency. When the receive bandwidth increases, the highest frequency sampled in the echo also increases. To accurately sample this higher frequency, the digital sampling frequency also increases (aliasing results if this does not occur). If the receive bandwidth increases from 32 to 64 KHz, for example, the Nyquist frequency is 32 KHz and the sampling frequency is twice this, i.e. 64 KHz which is the same as the receive bandwidth. Using a digital sampling frequency of 64 KHz, means that 64 000 data points are acquired per second so that the sampling interval halves ($1 \div 64\ 000$ is half $1 \div 32\ 000$). Hence, the required number of data points (as determined by the frequency matrix) is acquired in half the sampling window. The opposite is true if the receive bandwidth decreases.

Scan tip: **Receive bandwidth – what's going on behind the scenes?**

When we select the receive bandwidth in the scan protocol, behind the scenes, we determine how frequently the computer samples the echo per second during the sampling window. This is the digital sampling frequency and has the same numerical value as the receive bandwidth. If a receive bandwidth of 64 KHz is selected, then the echo is sampled 64 000 times per second during the sampling window. If it is 16 KHz, the echo is sampled 16 000 times per second and so on. A high or wide receive bandwidth means that the echo is sampled more frequently than when the bandwidth is low. Data are therefore collected more quickly, and, consequently, the sampling window is short. In other words, all the required data points are collected quickly, and the time needed to collect them (sampling window or length of the race in the sprinter analogy) is short. The opposite applies when the receive bandwidth is low or narrow. Data are collected more slowly, and the sampling window is long.

Sampling window (sampling time)

The sampling window is not a user-selectable parameter in the scan protocol. However, as the echo is usually centered in the middle of this time (i.e. the peak of the echo corresponds to the middle of the application of the frequency-encoding gradient), the duration of the sampling window indirectly affects the TE (which *is* selectable in the scan protocol). For example, if the frequency-encoding gradient is switched on for 8 ms (i.e. the sampling window is 8 ms), then the peak of the echo occurs after 4 ms. If the sampling window increases, the frequency-encoding gradient switches on for longer. Hence, the peak of the echo occurs later, increasing the time from the peak of the echo to the RF excitation pulse that created it (i.e. TE increases). The opposite is true if the sampling window decreases.

Learning tip: The relationship between TE, receive bandwidth, and frequency matrix

So why bother learning about sampling? The reason is that the receive bandwidth, frequency matrix, and minimum TE are related to one another and have a significant impact on data acquisition. These three parameters are selected in every scan protocol, and it's important that they work together effectively. To understand this more clearly, let us recap the following:

- The receive bandwidth determines the range of frequencies accurately digitized during the sampling window. It has the same numerical value as the digital sampling frequency when Nyquist principles are applied.

- The digital sampling frequency determines the number of data points acquired per second through ADC.
- The frequency matrix determines the number of data points collected during the sampling window.
- The minimum TE is affected by the duration of the sampling window because the echo is usually centered in the middle of this time window.

Let's return to the sprinter analogy for a moment. Suppose we wish to take 10 photographs of our sprinter but with a camera that only takes a photograph every 2 s instead of every 1 s. This is equivalent to halving the digital sampling frequency and the receive bandwidth. We still require 10 photographs of the sprinter to work out exactly how she was running during the race. This is equivalent to wanting to maintain the selected frequency matrix. One of the ways to achieve this is to make the race twice as long, i.e. the race takes 20 s instead of 10 s. This is equivalent to doubling the sampling window. We ask the sprinter to run the race but, instead of slowing down and doing a victory lap, continue running at the same pace for another 10 s. That way, using a camera that takes a picture every 2 s, 10 photographs are still taken. It just takes longer to acquire these 10 photographs than when we used our original camera that takes pictures at a higher frequency (every 1 s). As the race takes twice as long, this is equivalent to increasing the TE because the echo is repositioned to occur in the middle of the longer race.

The same is true if we need 20 photographs instead of 10. This is equivalent to doubling the frequency matrix. Let's assume we take one photograph per second. To take 20 photographs at this rate, we double the length of the race from 10 to 20 s. Both these scenarios require a longer race – either halving the number of photos taken per second (which is equivalent to halving the digital sampling frequency and receive bandwidth) – or doubling the number of photographs (which is equivalent to doubling the frequency matrix). Doubling the length of the race is equivalent to doubling the sampling window, and this increases the TE, as the echo is located in the middle of the longer sampling window.

For example, if a frequency matrix of 256 is required, 256 data points are acquired during the sampling window. If a receive bandwidth of 32 KHz is selected, the digital sampling frequency is also 32 KHz. This means that 32 000 data points are collected per second (Equation (6.5)). As the sampling interval is $1 \div$ digital sampling frequency, a data point is acquired every 0.031 ms. Therefore, to acquire 256 data points, the sampling window is 256×0.031 ms or 8 ms (Equation (6.6)). Hence, the frequency-encoding gradient is switched on for 8 ms to allow enough time for 256 data points to be acquired when sampling once every 0.031 ms at a sampling frequency of 32 KHz. The sampling window includes 4 ms as the magnetic moments of hydrogen nuclei rephase and reach a peak at time TE, and then 4 ms as the magnetic moments dephase. The peak of the echo therefore occurs half way into the sampling window after 4 ms.

Equation 6.5

$$RBW = 1/\Delta T_s$$

RBW is the receive bandwidth (KHz)
 ΔT_s is the interval between each data point (ms)

The interval between each data point or sampling interval is determined by the receive bandwidth. If the sampling interval is short, the receive bandwidth increases and vice versa

Equation 6.6

$$W_s = \Delta T_s \times M(f)$$

W_s is the sampling window (ms)
 ΔT_s is the sampling interval (ms)
 $M(f)$ is the frequency matrix

The frequency matrix determines how many data points are collected in each line of k -space. Therefore, the sampling window is this number multiplied by the sampling interval

176

If the receive bandwidth halves to 16 KHz, the digital sampling frequency also halves to 16 KHz, and so 16 000 data points are acquired per second. If the sampling window is still 8 ms, only 128 data points can be collected in this time instead of the required 256. To collect 256 data points at that receive bandwidth, the sampling window is doubled to 16 ms and results in a 4 ms increase in the minimum permissible TE, i.e. the peak of the echo moves to occur in the middle of the longer sampling window.

For example, if the minimum TE is 10 ms using a bandwidth of 32 KHz and a frequency matrix of 256, halving the receive bandwidth to 16 KHz increases the minimum TE to 14 ms (Figure 6.9). This is because the sampling window is 8 ms longer than before, and, therefore, the peak of the echo occurs 8 ms into the sampling window rather than 4 ms. Remember that the time between the RF excitation pulse and the beginning of the sampling window does not change. It is 6 ms in both these scenarios. What increases is the sampling window. The TE increases because the peak of the echo is in the middle of this longer window.

- A receive bandwidth of 32 KHz and a frequency matrix of 256 result in a minimum TE of $6 + 4 = 10$ ms.
- A receive bandwidth of 16 KHz and a frequency matrix of 256 result in a minimum TE of $6 + 8 = 14$ ms.

Increasing the frequency matrix has the same effect as reducing the receive bandwidth. Using the example above, if the frequency matrix doubles from 256 to 512, then 512 data points are required, and frequencies must be sampled 512 times during the sampling window. If the receive bandwidth is maintained at 32 KHz, then the sampling window and, therefore, the minimum TE increase to obtain the required number of data points.

Table 6.2 outlines this more clearly. The default is shown in the top line where a sampling window of 8 ms is used with a 32 KHz bandwidth with a frequency matrix of 256. If the bandwidth halves, not enough data points are acquired (128 instead of the required 256). To solve this, the sampling window doubles to 16 ms, which increases the TE by 4 ms (as the peak of the echo is situated in the middle of the sampling window as shown in Figure 6.10). The same occurs if a frequency matrix of 512 is required. The sampling window doubles to acquire 512 data points. This also increases the TE by 4 ms.

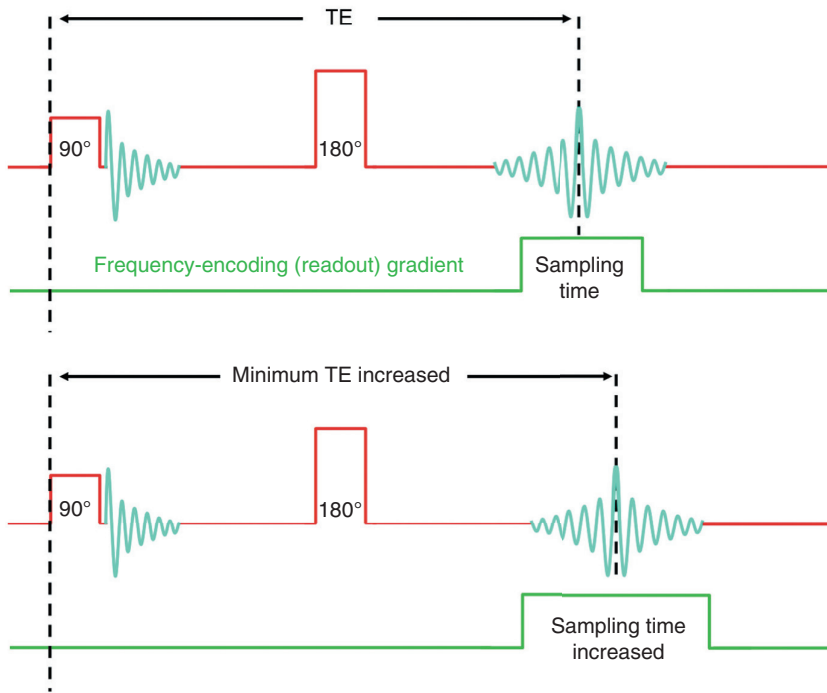


Figure 6.9 Sampling time (acquisition window) and the TE.

Table 6.2 Receive bandwidth, sampling window, and frequency matrix.

Frequency matrix	Receive bandwidth (KHz)	Sampling window (ms)
256	32	8
128	16	8
256	16	16
512	32	16




Figure 6.10 FFT.

Table 6.3 Things to remember – sampling.

The sampling window is how long the system has to acquire the data. It is the time during which the frequency-encoding gradient is switched on
The digital sampling frequency is how often the system samples frequencies per s during the sampling time. When Nyquist is obeyed, this has the same numerical value as that of the receive bandwidth
The frequency matrix determines how many data points are collected during the sampling time
The digital sampling frequency and the sampling window determine how many data points can be collected and therefore the frequency matrix
If either the receive bandwidth or the frequency matrix is altered, the sampling window is changed and this impacts on the TE as the peak of the echo is positioned in the middle of the new sampling window

Scan tip: How different manufacturers display the receive bandwidth

Different manufacturers do not use the same terms for the receive bandwidth. Some display the receive bandwidth across the whole FOV. If this is the case, then the receive bandwidth typically is on a scale of KHz or 1000 Hz. Others choose to display it as the receive bandwidth per pixel in the frequency direction of the image. To do this, the receive bandwidth for the whole FOV is divided by the frequency matrix. If this is the case, then the receive bandwidth typically is on a scale of Hz. Some display this parameter as the number of pixels that fat and water are separated by in the image. The magnetic moments of fat and water precess at different frequencies at the same field strength. For example, there is a 220 Hz difference in their precessional frequencies at 1.5T. The receive bandwidth determines how many pixels apart signals from fat and water are mapped into the FOV, and, therefore, the receive bandwidth is expressed as the number of pixels they are shifted apart in the image.

For example, at 1.5T using a receive bandwidth of 32 KHz and a frequency matrix of 256, the receive bandwidth might be displayed in the following formats (see the table at the beginning of the book for acronyms for this parameter).

- As bandwidth over the whole FOV ± 16 KHz (32 KHz)
- As a bandwidth per pixel $32\,000 \div 256 = 125$ Hz/pixel
- As fat water shift of $220 \div 125 = 0.176$ pixels.

Scan tip: Slice number and the TE

The TE is not a scan time parameter but, under certain circumstances, it can influence it. Imagine a scenario where we set up a scan protocol with the maximum number of slices allowed for the selected TR. Just before beginning the scan, we decide to double the frequency matrix from 256 to 512. This means that the system collects 512 data points during the sampling window as opposed to 256. If the receive bandwidth (and therefore the digital sampling frequency) remains

unchanged, to collect the required number of data points, the sampling window doubles. As the echo is situated in the middle of the sampling window, the peak of the echo is repositioned to occur in the middle of this longer window. This increases the TE.

As the system now “sits” at each slice for longer waiting for the echo to occur, and it cannot excite the next slice until it has read the echo from the current slice, it cannot acquire all the slices in the selected TR. The system either increases the TR to fit the slices in, or acquires half of the slices in one acquisition and repeats the whole scan again to acquire the other half. Both strategies cause an increase in scan time because the TR increases.

Fast Fourier transform (FFT)

MR images consist of a matrix of pixels, the number of which is determined by the number of lines filled in k -space (phase matrix) and the number of data points in each line (frequency matrix). After FFT, each pixel is allocated a color on the grayscale. This corresponds to the amplitude of specific frequencies coming from the same spatial location as represented by that pixel.

Each data point contains phase and frequency information from the *whole* slice at a certain time during the sampling window. In other words, frequency amplitudes are represented in the time domain. FFT converts this into frequency amplitudes in the *frequency* domain. This is necessary because gradients spatially locate signal according to their frequency, not their time. Instead of looking at the amplitude of a waveform as it plays out over time, FFT looks at its amplitude dependent on its frequency. A waveform with a certain frequency and amplitude becomes a spike at one frequency after FFT. Two waveforms, with two different frequencies, become two spikes after FFT. The height of these spikes depends on the amplitude of the original waveform (Figure 6.10). FFT is basically the math that gets us from the data points in k -space to the image. If you are a clinical MRI practitioner and not a mathematician, it is perfectly acceptable to leave it at that! A more detailed, but still nonmathematical, explanation now follows.

The frequency-encoding gradient is applied during the echo. It causes a change of frequency of the magnetic moments of hydrogen nuclei according to their spatial location along the gradient. During application of this gradient, these different frequencies are sampled by the computer. This information is stored in each data point laid out in a line of k -space during the brief sampling window of just a few millisecond. Phase encoding is more complicated. The phase-encoding gradient produces a change of phase of magnetic moments of hydrogen nuclei located along the gradient. Once this occurs, the gradient switches off, and the magnetic moments of spins located along it return to precessing at the Larmor frequency. Their phase shift, however, remains. A waveform is derived from plotting the change of phase of the magnetic moments of spins imposed by this gradient over distance. The spatial frequency of this waveform depends on the slope or amplitude of the phase-encoding gradient (see Chapter 5). If the slope of the phase encoding changes every TR, then the resultant phase shift, and, therefore spatial frequency, is different every TR.

Spins located at a certain position in the slice experience a different phase-encoding gradient every TR, and therefore the spatial frequency obtained for those spins is also different every TR. Unlike frequency encoding, which relies on what happened during the short time-frame of the sampling window, the phase encoding process relies on what happened during the *entire scan*. Spins located at a unique position in the slice experience a unique pattern of different slopes

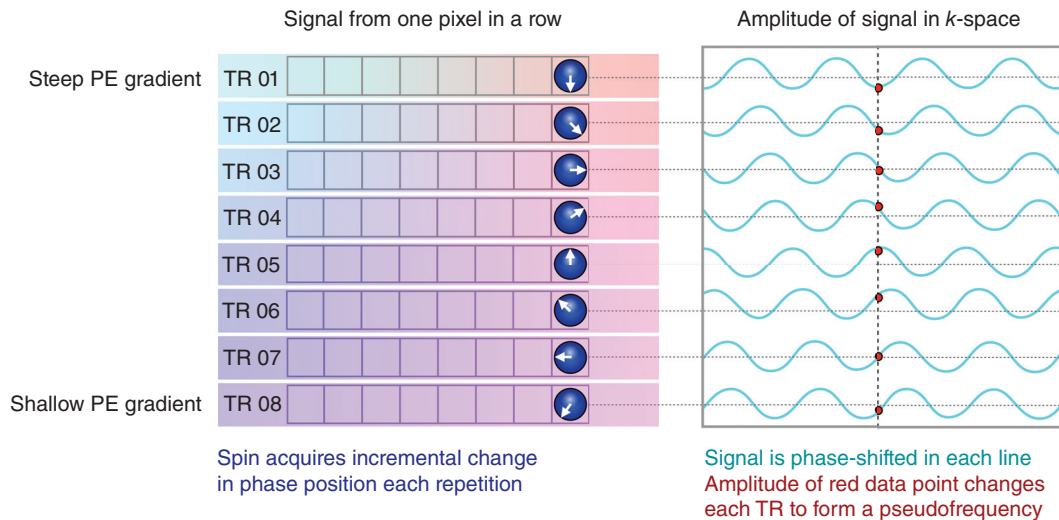


Figure 6.11 Signals in k -space from one voxel.

of phase-encoding gradient throughout the scan, and, therefore, a unique pattern of spatial frequencies are obtained. FFT calculates the exact position of spins located along the phase encoding axis according to this unique pattern.

To illustrate this more clearly, look at Figure 6.11. This shows how data from spins located at the edge of the FOV are mapped into k -space. Spins located here experience a large variation in phase encoding slope throughout the scan. As they are located at the edge of the FOV, in some TRs spins located here experience very steep applications of the phase-encoding gradient, both positively and negatively. In some TRs, however, they experience very shallow applications of the phase-encoding gradient. The point is that there is a great *variation* in the slope of phase-encoding gradient throughout the scan at this location. Spins are therefore phase-shifted to varying degrees during the scan, and therefore the spatial frequencies derived from these phase changes are also very varied (from low to high spatial frequencies) [4]. This variation from one TR to the next is mapped into k -space (see Figure 6.11). The waveforms on each k -space line are created in different TRs by different slopes of phase-encoding gradient, and so their phase information is different. The red dots in Figure 6.11 represent the data point corresponding to the middle of each line, and their position on the waveforms is different. The change of location of the red dots on each waveform represents the change in phase of spins at this location because the slope of the phase-encoding gradient is changed every TR.

Now look at Figure 6.12. k -Space is rotated to show its phase axis. The green dots in Figure 6.12 are the red dots shown in Figure 6.11. When seen from the side (as in the bottom diagrams of Figure 6.12), you can again see how their position changes every TR. If the green dots are joined together, another waveform emerges. This is called a **pseudo-frequency** and is shown in the bottom right-hand diagram of Figure 6.12. It represents the change of phase of spins at this location, at the edge of the FOV, throughout the *entire scan*, from the first TR in the scan, to the last. Spins located elsewhere in the slice experience a different pattern of changing phase-encoding gradient slope throughout the scan and therefore have a different pseudo-frequency than the one illustrated in Figures 6.11 and 6.12.

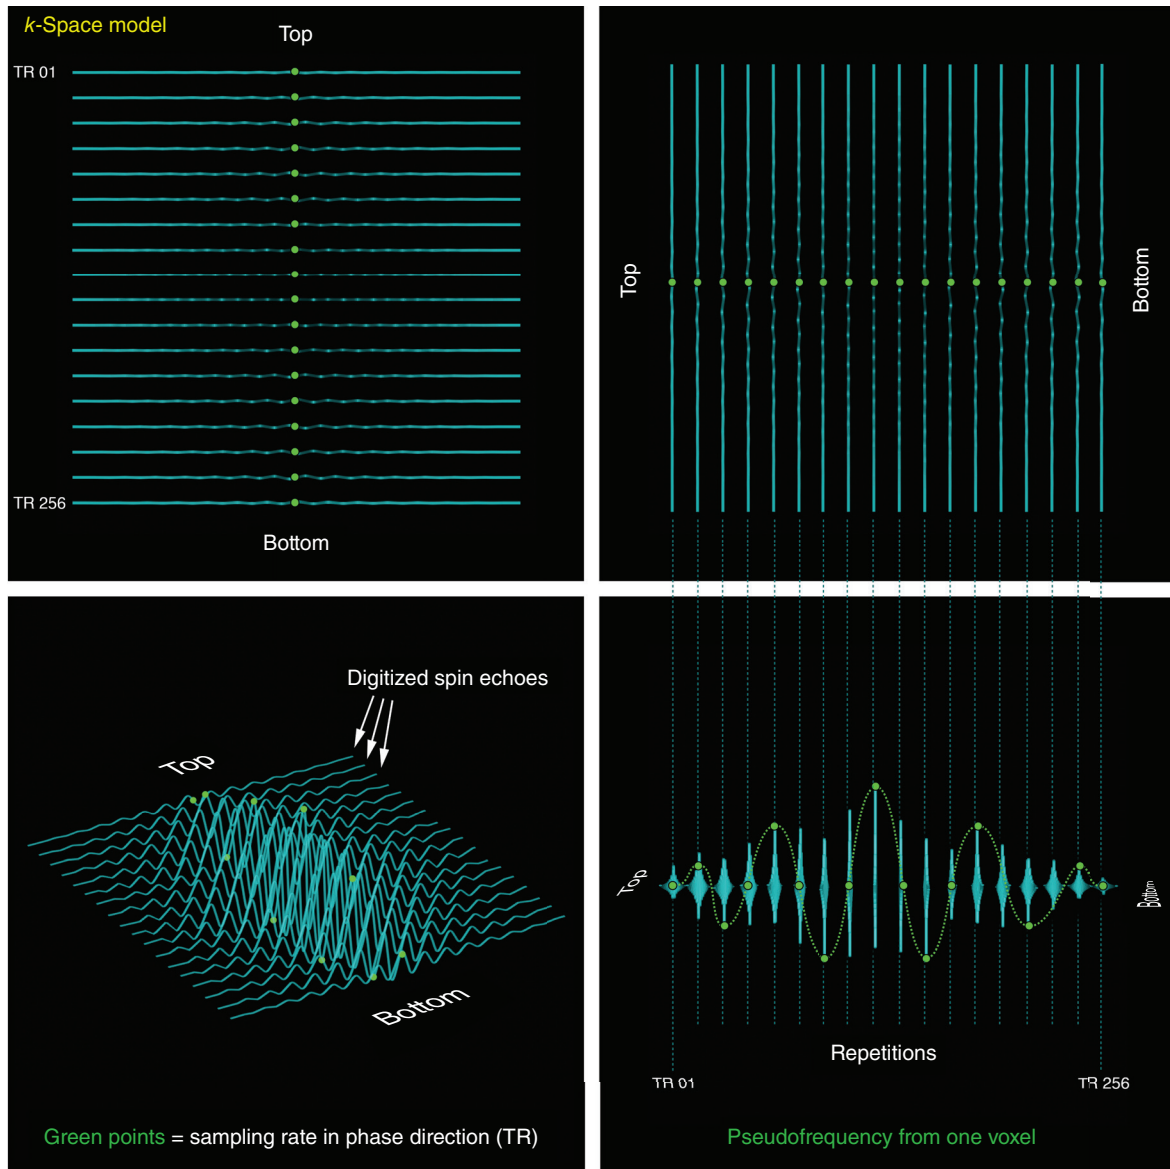


Figure 6.12 Generation of pseudo-frequency from one voxel.

Each data point in *k*-space therefore has information from the entire slice “locked” inside it. This information includes:

- spatial frequencies obtained from frequency encoding;
- spatial frequencies (called pseudo-frequencies) obtained from phase encoding.

FFT is required to mathematically “unlock” these data so that the system can locate signal at each pixel location in the slice and assign a signal intensity to it. This signal intensity depends on the amplitude of signal at each pixel location.

Perhaps the easiest example is to figure out how a pixel located at magnetic isocenter, right in the middle of the slice in both frequency and phase directions, is encoded. The magnetic moments of hydrogen nuclei located here experience no change of frequency or phase during the scan because the gradients do not change the magnetic field strength at this location. In terms of frequency encoding, magnetic moments of hydrogen located here always precess at the Larmor frequency regardless of the amplitude of the frequency-encoding gradient. In terms of phase encoding, the magnetic moments of these spins never undergo a change of phase because their frequencies are unaltered [5]. Therefore, when data from the phase encoding process are mapped into k -space, the position of the green dot from one k -space line to the next remains unchanged throughout the scan. If these green dots are joined together in the phase axis of k -space, the result is a straight line. In other words, there is no pseudo-frequency or spatial frequency from the phase encoding process for spins located at magnetic isocenter. FFT maps this zero change of frequency and zero change of phase in a pixel at magnetic isocenter. Another way of putting this is that any signal that demonstrates a zero change in frequency and phase *must* be coming from this point in the slice – it cannot be coming from anywhere else because it is only at magnetic isocenter that these conditions apply. If there is a lot of signal that behaves in this way (high signal amplitude), then the pixel is allocated a white tone on the grayscale. If there is little or no signal that behaves in this way (low signal amplitude), then the pixel is allocated a black tone on the grayscale.

Table 6.4 Things to remember – phase data in k -space.

The slope of the phase-encoding gradient is changed every TR (assuming 1 NSA). This changes the degree of phase shift across a certain distance in the patient
Steep phase-encoding gradient slopes produce more phase shift than shallow ones
The phase position of magnetic moments in voxels across the whole slice creates a waveform across the slice. The spatial frequency of this waveform depends on the amount of phase shift over distance along the gradient
The phase position of magnetic moments of spins within each voxel is altered differently when the slope of the phase-encoding gradient is altered, and this is reflected by the production of a pseudo-frequency. This is obtained by observing the change of phase shift of magnetic moments in a voxel during the <i>entire scan</i>
Each pixel in the slice therefore has its own unique pseudo-frequency so hundreds of pseudo-frequencies are obtained throughout the scan
These pseudo-frequencies are sampled along with the frequencies obtained by frequency encoding, and information from both processes is contained within a data point
Each data point therefore has information about what happened during the spatial encoding process for the entire slice
FFT is required to “unlock” each data point and calculate the signal intensity for each pixel position in the slice

This process is repeated for every pixel location in the slice. Each pixel location has its own unique pattern of spatial frequency, and, therefore, its own pseudo-frequency, and there are differing amplitudes of signal at these different pixel locations. FFT works out what these unique patterns are and the spatial location of each unique pattern (pixel). It then allocates a tone on the grayscale (from black to white and hundreds of shades of gray in-between) to each pixel depending on the amplitude of signal with that unique pattern.

Learning tip: Why does the phase gradient have to change?

Remember that we need to change the amplitude of the phase-encoding gradient to fill different lines of k -space and therefore create multiple pixels in the phase direction of the image. Another way of looking at this is that by changing the slope of the phase-encoding gradient the data “look different” than in the previous TR. This is how the system knows to place these data points in a new line of k -space. If the data look the same every TR, then the system places the data in the same line every TR, and the resultant image has a resolution of only 1 pixel in the phase direction of the image. To create a phase matrix, the data must be different every TR.

The frequency data cannot be changed from one TR to the next because, to do this, the slope of the frequency-encoding gradient changes every TR. This, in turn, changes the size of pixels and the FOV in the frequency direction of the image every TR (see later), which is obviously undesirable. The slice-selection data also cannot be changed every TR because they alter the slope of the slice-select gradient applied to a slice every TR. This, in turn, changes the slice thickness every TR, which is again undesirable. The only gradient slope that can be changed without having any serious undesirable effect is the phase-encoding gradient, and, by doing so, this alters the phase data. This is what the system needs to place these “different” data in a different line of k -space and thus provide multiple pixels in the phase direction of the image.

PART 3: SOME IMPORTANT FACTS ABOUT k -SPACE!

The difficult part of k -space is over and it's time to explore several important points. These are useful pieces in the k -space jigsaw puzzle.

Fact 1: k -Space is not the image

It is tempting to think that a single data point in k -space becomes a single pixel in the image. Images have a grid or matrix of pixels, and k -space has a grid or matrix of data points of the same number.

- The number of pixels in the frequency encoding axis of the image is the frequency matrix, and this is the same number of data points in each row of k -space.
- The number of pixels in the phase encoding axis of the image is the phase matrix, and this is the same number of data points in each column of k -space.

However, the grid of data points in k -space does *not* correspond to the grid of pixels in the image. One data point does *not* equal 1 pixel. If it did, we wouldn't need FFT! As we have seen, each data point contains information for the *whole slice*, as spatial frequencies come from the whole echo, and the echo comes from the whole slice. This explains why we need FFT to “unlock” each data point in k -space.




Figure 6.13 *k*-Space symmetry – phase.

Fact 2: Data are symmetrical in *k*-space

Data in the top half of *k*-space are identical to those in the bottom half. This is because the slope of phase-encoding gradient required to select a certain line in one half of *k*-space is identical to that required to select the same line in the opposite side of *k*-space. Although the polarity of gradient is different, the slope is the same so the spatial information in each line is also the same (Figure 6.13).

In addition, data on the left side of *k*-space are identical to data on the right. This is because data points are laid out in a line sequentially from left to right as the echo is rephasing, reaching its peak and dephasing. The peak of the echo corresponds to the central vertical axis of *k*-space. As echoes are symmetrical features, frequency data digitized from the echo are the same on one side as they are on the other (Figure 6.14). The symmetry of *k*-space is called **conjugate symmetry**. As *k*-space is symmetrical top to bottom and left to right, each data point in *k*-space is the same as one diagonally opposite to it. Conjugate symmetry is used to reduce scan times in some imaging options (see Part 5).

Fact 3: Data acquired in the central lines contribute signal and contrast, while data acquired in the outer lines contribute resolution

The central lines of *k*-space are filled using shallow phase-encoding gradient slopes, and the outer lines are filled using steep phase-encoding gradient slopes. Shallow slopes result in low spatial frequencies because of small changes of phase over distance along the phase-encoding gradient. To produce echoes with a large signal amplitude, the magnetic moments of hydrogen nuclei must be coherent or in phase. By minimizing phase change with a shallow phase-encoding




Figure 6.14 k-Space symmetry – frequency.

gradient slope, the resultant echo has a high signal amplitude and contributes largely to signal and contrast in the image.

Steep phase-encoding gradient slopes result in high spatial frequencies because of large changes of phase over distance along the phase-encoding gradient. The resultant echo therefore has relatively low signal amplitude and does not contribute signal and contrast in the image (Figure 6.15). However, large changes of phase along the phase-encoding gradient mean that two points close together in the patient are likely to have a phase difference and are differentiated from each other. They therefore provide resolution information.

To summarize:

- The *central* portion of *k-space* contains data that have *high signal* amplitude and *low resolution*.
- The *outer* portion of *k-space* contains data that have *low signal* amplitude and *high resolution*.

Signal and resolution are important image quality factors and are discussed in Chapter 7. If all *k-space* is filled during an acquisition, then both signal and resolution are obtained and displayed in the image. However, there are many different permutations of *k-space* filling where the relative proportion of central to outer lines filled is altered. Under these circumstances, image quality is significantly affected. For example, this fact is used in TSE to weight images (see Chapter 3). In this pulse sequence, *k-space* lines are reordered so that central lines, that contribute signal and contrast, are filled during echoes that have a similar TE to the effective TE we select in the




Figure 6.15 Phase gradient amplitude vs signal amplitude.

scan protocol. Their influence on image weighting is maximized. Outer, resolution k -space lines are filled during echoes with TEs that are dissimilar to the effective TE. These echoes contribute resolution but, because their amplitude is low, they do not contribute signal and contrast. Their influence on image weighting is minimized.

It is also worth noting that when the phase matrix decreases the outer lines are dropped and only the central lines of k -space are filled with data. For example, if the phase matrix decreases from 512 to 128, then rather than just fill lines (+128 to 0), lines (+64, 0, -63) are filled. These are primarily the signal and contrast lines of k -space (Figure 6.16). This convention is used because signal and contrast are usually more important than resolution in MR images (see Chapter 7). When resolution is also required, this is achieved by increasing the proportion of outer lines that contain resolution data.




Figure 6.16 k -Space – signal and resolution.

Learning tip: *k*-Space, resolution, and signal

Figure 6.17 shows an image acquired using all *k*-space. Both resolution and signal are seen. Figure 6.18 illustrates what happens if an image is created only from data at the outer edges of *k*-space. This image has good resolution in that the detail of the hair and eyes is well demonstrated but there is very little signal. Figure 6.19 shows what happens if an image is created only from data in the center of *k*-space. The resultant image has excellent signal but poor resolution. This example also demonstrates that *k*-space is not the image (see Fact 1). If it were, the image in Figure 6.18 would lose its nose, and Figure 6.19 would show only the nose. Both images, however, show all the face, even though only a small percentage of the total number of data points in *k*-space are used in their creation. This is because every data point in *k*-space contains information from the entire slice (see Fact 1).

187




Figure 6.17 *k*-Space using all data.




Figure 6.18 *k*-Space using resolution data only.




Figure 6.19 *k*-Space using signal data only.

Table 6.5 Things to remember – signal, contrast, and resolution in *k*-space.

The outer lines of *k*-space contain data with high spatial resolution as they are filled by steep phase-encoding gradient slopes

The central lines of *k*-space contain data with low spatial resolution as they are filled by shallow phase-encoding gradient slopes

The central portion of *k*-space contains data that has high signal amplitude and low spatial resolution

The outer portion of *k*-space contains data that has high spatial resolution and low signal amplitude

Fact 4: The scan time is the time to fill *k*-space

When the system calculates the scan time, it usually multiplies the following three parameters:

- TR
- Phase matrix
- **Number of signal averages (NSA).**

There are additional scan time parameters in TSE and volume imaging (Equations (6.7)–(6.9)), but the scan time is always based on the parameters listed above. Let's see why.

Equations 6.7–6.9: Scan time equations.

In 2D imaging:
 $ST = TR \times M(p) \times NSA$

ST is the scan time (s)
 TR is the time to repeat (ms)
 M(p) is the phase matrix
 NSA is the number of signal averages

This equation shows how the scan time is calculated by the system

Equations 6.7–6.9: (Continued)

<p>In 2D TSE or FSE: $ST = \frac{TR \times M(p) \times NSA}{ETL}$</p>	<p>ST is the scan time (s) TR is the time to repeat (ms) M(p) is the phase matrix NSA is the number of signal averages ETL is the echo train length</p>	<p>The ETL determines how many lines of <i>k</i>-space are filled per TR. The higher the ETL, the shorter the scan time</p>
<p>In 3D imaging: $ST = TR \times M(p) \times NSA \times N_s$</p>	<p>ST is the scan time (s) TR is the time to repeat (ms) M(p) is the phase matrix NSA is the number of signal averages N_s is the number of slice locations</p>	<p>The number of slices in 3D imaging is equivalent to a slice matrix. This is why the scan time is multiplied by this number as it is similar to the phase matrix</p>

Repetition time (TR)

Every TR, each slice is selected, phase encoded, and frequency encoded. Its echo is sampled, and frequencies are digitized. Therefore, every TR a line of data is laid out in *k*-space for each slice. Once this occurs, the system repeats this process for every other slice in the slice stack. In the next TR, the first slice is excited again, a different line of *k*-space is filled with data, and this is repeated for every slice. Long TRs mean that there is a long time between filling one *k*-space line and the next for each area of *k*-space. The opposite is true of short TRs. It therefore takes longer to fill all *k*-space areas and therefore complete the scan when the TR is long compared to when it is short.

Scan tip: TR and slice number

Slices are not usually selected together but sequentially, i.e. slice 1 is selected and encoded, and frequencies from its echo are digitized. Then the next slice is selected, encoded, and digitized, and so on. This explains why the maximum number of available slices depends on the TR (Figure 6.20). Longer TRs allow more slices to be selected, encoded, and digitized than short TRs. A TR of say 500 ms may allow 30 slices, while a TR of 2000 ms allows 40 slices. This point was also explored in Chapter 5.

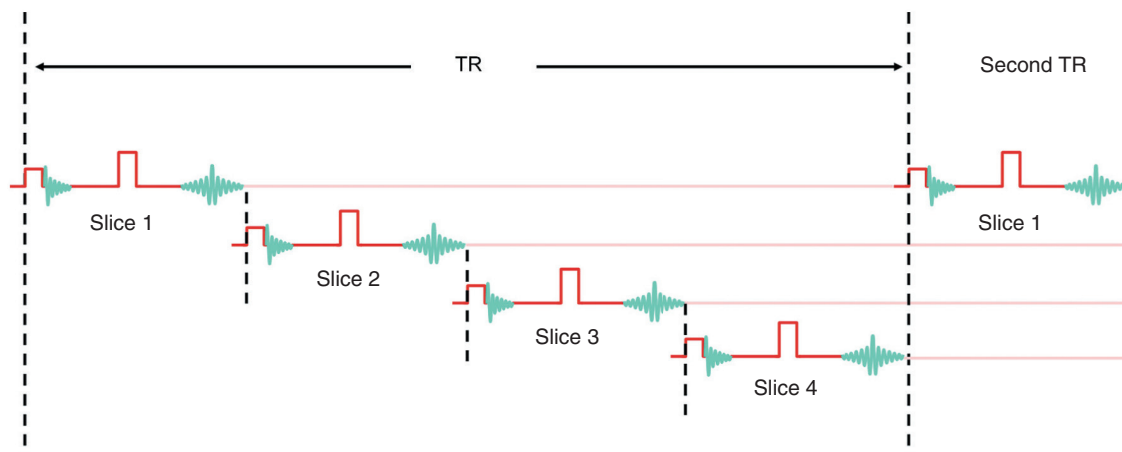


Figure 6.20 TR vs slice number.

Phase matrix

This determines the number of lines of k -space filled to complete the scan. If one line is filled per TR, then if:

- a phase matrix of 128 is selected, 128 lines are filled, and 128 TRs are completed to finish the scan;
- a phase matrix of 256 is selected, 256 lines are filled, and 256 TRs are completed to finish the scan.

As the phase matrix increases, the scan time also increases because it takes longer to fill k -space than when using a low or coarse phase matrix.

NSA (also known as number of excitations or NEX)

This is the number of times each line is filled with data. The echo can be sampled more than once by maintaining the same slope of phase gradient over several TRs instead of changing it every TR. The same line of k -space is filled several times so that each line contains more data than if data are placed there only once. As there are more data in each line, the resultant image has a higher signal-to-noise ratio (SNR) (see Chapter 7), but the scan time is proportionally longer.

For example:

- TR 1000 ms, phase matrix 256, 1 NSA, scan time = 256 s
- TR 1000 ms, phase matrix 256, 2 NSA, scan time = 512 s.

Usually, to fill each line more than once, the same slope of phase-encoding gradient is used over two or more successive TRs, rather than filling all the lines once and then returning to repeat the process again. In TSE pulse sequences, the scan time is calculated using an additional parameter, the turbo factor or ETL (see Chapter 3). It is also modified in volume and parallel imaging (see Part 5).

Analogy:

Using the chest of drawers analogy to explain k -space and scan time



- The TR is the time between filling the top drawer of the first chest of drawers and filling the next drawer down in the same chest of drawers. During that time, the top drawer in all the other chests of drawers is filled sequentially one after the other.
- The phase matrix is the number of drawers in each of the chests of drawers.
- The NSA is the number of times each drawer is filled, e.g. once, twice, thrice, etc.

The scan is over when all the drawers, in all the chests of drawers, are filled.

The number of data points in each line of k -space does not increase when the NSA increases, i.e. when two signal averages are used (2 NSA), the second signal average does not double the number of data points in each line of k -space. The number of data points in each line of k -space is the frequency matrix, so it is important to ensure that the number of data points in each line remains the same. When the NSA increases, there are the same number of data points in each line but each data point contains more information. If data points were pairs of socks, then the socks get larger but the number of socks in each line remains the same.

Scan tip: NSA – What’s going on behind the scenes?

When we select the NSA in the scan protocol, behind the scenes, we determine for how many successive TRs the phase-encoding gradient slope is applied to the same slope and polarity to fill the same line of k -space with data.

Learning tip: Scan time and k -space dimensions

Although data points in k -space contain spatial frequency information, the dimensions of k -space are in the time domain. The width of the chest of drawers depends on the sampling window measured in millisecond. The height of the chest of drawers is determined by how many drawers there are of a certain depth (see later). These drawers are filled over successive TR periods measured in millisecond. A large area of k -space usually takes longer to fill than a smaller area and explains why the scan time is the time to fill k -space.

Fact 5: The incremental step between each data point in k -space determines the FOV

Data points in k -space are separated from each other. Those in each column represent the process of phase encoding and are separated into lines. The data in each line are acquired after application of a different slope of phase-encoding gradient applied in each TR period. The degree of phase change between the data points in each column is therefore different. The data points in each row represent the process of frequency encoding. They are laid out in k -space, and the separation between each one is the sampling interval. Fact 5 tells us that the difference between each data point in each column determines the size of the FOV in the phase direction of the image, and the difference between data points in each row determines the FOV in the frequency direction of the image. Let’s see how.

In the phase axis of k -space, the difference in phase shift between each data point in each column is inversely proportional to the size of the phase FOV (Equation (6.10)). The difference in phase shift depends on the incremental step between each line of k -space, and there is an inverse relationship between these two variables. This means that if the incremental step between each line of k -space doubles, then the phase FOV halves and vice versa.

Equation 6.10

$$\text{Size of pixel } (p) = \frac{1}{G(p)M(p)\phi}$$

$G(p)$ is the max amplitude of the phase-encoding gradient (mT/m)

$M(p)$ is the phase matrix
 ϕ is the incremental step between each line

This equation shows that a pixel size in the phase axis is determined by the maximum amplitude of the phase-encoding gradient, the incremental step, and the total number of lines or phase matrix (i.e. the height of the chest of drawers)

192

Analogy: Using the chest of drawers to understand Fact 5



The depth of each drawer in our chest of drawers (i.e. its dimension from the bottom to the top of the drawer) determines the size of the FOV in the phase direction of the image. This is because this dimension is analogous to the incremental step between each line of k -space. Deep drawers produce a smaller phase FOV, and shallow drawers a larger phase FOV. So, for example, let's assume that we no longer want to store pairs of socks in each drawer but want to store something larger like thick, woolly sweaters. To fit these sweaters into the drawers, each drawer needs to be deeper from top to bottom. To keep the math simple, let's make them two times deeper from top to bottom than when socks are stored in the chest of drawers. As each drawer is twice as deep, this equates to a phase FOV that is halved in size compared to when socks are stored. If we do the opposite and store something much thinner than socks in each drawer (for example, butterflies), then this means that the drawers are shallower from top to bottom. To keep the math simple, let's make them half as deep from top to bottom than when socks are stored in the chest of drawers. As each drawer is half as deep, this equates to a phase FOV that is doubled in size compared to when socks are stored.

In the frequency axis of k -space, the sampling interval between each data point in each row is inversely proportional to the size of the frequency FOV (Equation (6.11)). This means that if the sampling interval between each data point halves, then the frequency FOV doubles and vice versa. We learned earlier that the sampling interval is determined by the receive bandwidth (sampling interval is $1 \div$ receive bandwidth). Therefore, if the sampling interval halves, the receive bandwidth doubles. Doubling the receive bandwidth doubles the frequency FOV and vice versa (Equation (6.12)). More on this later. The dimension of the FOV in the phase direction is usually expressed as a percentage of the FOV in the frequency direction. If the phase and frequency FOV are the same, then the phase FOV is 100%. If the phase FOV is half that of the frequency FOV, then this is expressed as a 50% phase FOV.

Equation 6.11

$$\text{Size of pixel (f)} = \frac{1}{G(f)M(f)\Delta T_s}$$

$G(f)$ is the amplitude of the frequency-encoding gradient (mT/m)
 $M(f)$ is the frequency matrix
 ΔT_s is the sampling interval (ms)

This equation shows that pixel dimension in the frequency axis is determined by the amplitude of the frequency-encoding gradient, the number of data points (frequency matrix) and the sampling interval between them (the width of the chest of drawers)

193

Equation 6.12

$$\text{FOV(f)} = \frac{\text{RBW}}{G(f)}$$

FOV(f) is the frequency FOV (cm)
 RBW is the receive bandwidth (KHz)
 $G(f)$ is the amplitude of the frequency-encoding gradient (mT/m)

This equation shows that a small FOV in the frequency direction is obtained by either increasing the amplitude of the frequency-encoding gradient or decreasing the RBW

The phase FOV dimension is changed in several imaging options including

- rectangular or asymmetric FOV (see Chapter 7)
- antialiasing (see Chapter 8)
- parallel imaging (see later in Part 5).

Fact 6: The dimensions of *k*-space determine pixel size

When we describe resolution in MR images, we refer to the size of each pixel, and this determines **spatial resolution**. Pixel size governs how well two points, close together in the patient, are resolved in the image. In the image, pixel size is determined by the FOV and the image matrix. The size of the pixel in the phase direction of the image depends on the phase FOV and the phase matrix. In the frequency direction of the image, it depends on the frequency FOV and the frequency matrix. In the simplest terms, pixel size depends on the size of the FOV and how many boxes (pixels) divide it up. Spatial resolution increases by either keeping the FOV the same and increasing the image matrix or keeping the matrix the same and decreasing the FOV (see Chapter 7). In *k*-space (which, remember, is not the image), resolution is referred to as **frequency resolution** [6] because *k*-space is a spatial frequency domain. So how does this play out in *k*-space? What's going on behind the scenes to change resolution? Let's explore how *k*-space controls the matrix, FOV, and pixel size. What follows is quite complicated, so you may need to read it a few times!

Image matrix

This is determined by the number of data points in *k*-space.

- The phase matrix depends on how many lines of *k*-space are filled with data points. It is therefore the total number of data points in each column of data in *k*-space.
- The frequency matrix is the number of data points in each row of *k*-space.

FOV

According to Fact 5, the FOV is determined by the difference between each data point in k -space in either direction. It is inversely proportional to the size of the FOV in that direction [7].

- In the phase axis of k -space, this is the incremental step between each line of k -space, or each data point in a column of data.
- In the frequency axis of k -space, this is the distance between each data point in k -space or sampling interval between each data point in each row of data.

194

Now let's combine these facts to determine the pixel size in the phase and frequency direction of the image and how this is related to the matrix and FOV. It might help to draw your own chest of drawers to see how the scenarios described below play out. At this point, it is probably worth emphasizing again that k -space is *not* the image. Individual data points in k -space do not in any way relate to individual pixels in the image, and the dimensions of k -space do not relate to the FOV.

Pixel size (phase)

If we combine the fact that the incremental step between each line of k -space determines the phase FOV with the fact that the phase matrix equals the number of data points in each column of k -space, pixel size in the phase direction of the image is determined by the amplitude of the phase-encoding gradient both positively and negatively [8]. In other words, it corresponds to the height of the chest of drawers. Another way of saying this is that for a given depth of drawer (phase FOV) and a given number of drawers (phase matrix), the overall height of the chest of drawers determines how big these pixels can be. If the depth of each drawer doubles but the *height* of the chest of drawers remains the same, fewer drawers of this depth are possible (half the number in fact) and therefore the phase matrix halves. For example, instead of 256 pixels of a certain size making up the phase FOV, only 128 pixels of that size make up the phase FOV. This explains why the phase FOV halves when the incremental step between each line doubles.

Pixel size (frequency)

If we combine the fact that the sampling interval between each data point in each row of k -space determines the frequency FOV with the fact that the frequency matrix equals the number of data points in each row, pixel size in the frequency direction of the image is determined by the sampling window [8]. In other words, it corresponds to the *width* of the chest of drawers. Another way of saying this is that for a given sampling interval (frequency FOV) and a given number of data points (frequency matrix), the overall width of the chest of drawers determines how big these pixels can be. If the sampling interval doubles (which, remember, is equivalent to halving the receive bandwidth) but the width of the chest of drawers remains the same, fewer data points are possible (half the number in fact) and therefore the frequency matrix halves. For example, instead of 256 pixels of a certain size making up the frequency FOV, only 128 pixels of that size make up the frequency FOV. This explains why the frequency FOV halves when the sampling interval doubles.

Scan tip:**What's going on behind the scenes? – k -space and FOV**

Now we have these facts, let's use them to figure out how the phase and frequency FOV are changed and what's going on behind the scenes.

FOV (phase). This is changed by adjusting the degree of phase-encoding gradient amplitude change from one application to the next [4]. When we decrease the size of the FOV in the phase direction in the scan protocol, behind the scenes, we increase the degree of phase-encoding gradient amplitude change from one line to the next. The opposite applies when we increase the size of the phase FOV.

FOV (frequency). This is changed by adjusting the sampling interval or the time between data points in each row of k -space. If the sampling interval increases, the frequency FOV decreases and vice versa. We have also just learned that there is a relationship between the receive bandwidth and the frequency FOV. This is because the sampling interval is inversely related to the receive bandwidth (see Equation (6.5)). Increasing the frequency FOV increases the receive bandwidth and the opposite applies if the frequency FOV decreases (Equation (6.12)). In practice, the FOV in the frequency direction is usually altered by changing the slope of the frequency-encoding gradient rather than changing the receive bandwidth. Therefore, when we decrease the size of the FOV in the frequency direction in the scan protocol, behind the scenes, a large current is applied through frequency-encoding gradient coil. This produces a steep gradient and results in a small-frequency FOV because the range of frequencies along the gradient is squeezed into a small distance in the patient. When we increase the size of the FOV in the frequency direction in the scan protocol, behind the scenes, a small current is applied through frequency-encoding gradient coil. This produces a shallow gradient and results in a large-frequency FOV because the range of frequencies along the gradient is spread apart over a large distance in the patient.

Scan tip:**What's going on behind the scenes? – k -space and pixel size**

When we change pixel size in the scan protocol, behind the scenes, we determine the vertical and horizontal dimensions of k -space. Pixel size in the phase axis of the image depends on the vertical dimension of k -space. Pixel size in the frequency axis of k -space depends on the horizontal dimension of k -space. Each of these dimensions is affected by the FOV and the image matrix, and explains why these two parameters determine pixel size and therefore spatial resolution.

Table 6.6 provides the complete picture of what's going on behind the scenes and relates parameters we change in the scan protocol to k -space and the chest of drawers analogy.

Table 6.6 Image acquisition – summary of what’s going on behind the scenes.

Parameter	Behind the scenes
Slice thickness	Amplitude of the slice-select gradient and transmit bandwidth
Slice gap	Amplitude of the slice-select gradient and transmit bandwidth
Frequency FOV	Amplitude of the frequency-encoding gradient and incremental step between data points in each row of k -space (sampling interval)
Phase FOV	Incremental step between data points each column of k -space. Degree of phase change from one line to the next
Frequency matrix	Number of data points in each line of k -space
Phase matrix	Number of data points in each column of k -space
Phase resolution (pixel size)	Amplitude of the steepest phase-encoding gradient positively and negatively (height of the chest of drawers)
Frequency resolution (pixel size)	Sampling window (width of the chest of drawers)
NSA or NEX	Number of times the same line of k -space is filled with data. Achieved by maintaining the slope and polarity of the phase-encoding gradients over successive TR periods

Learning tip: Why is there an inverse relationship between FOV and distance between each data point?



Before we leave this tricky section, let’s clarify why there is an inverse relationship between the FOV and the distance between data points. It’s really just math, but the chest of drawers analogy is also really helpful here. The inverse relationship comes about because the size of the FOV is calculated differently in the image than it is in k -space.

- In the image, the FOV is calculated by multiplying the pixel size by the number of pixels.
- In k -space, the FOV is calculated by dividing the dimensions of k -space by the distance between each data point.

Note that in the image, a multiplication calculation is made, whereas in k -space, a division calculation is made. How does this play out using the chest of drawers analogy?

Imagine our chest of drawers has a wooden frame into which we slot a certain number of drawers and into which we place a certain number of socks. Remember that the:

- height of the wooden frame = pixel size in the phase direction of the image
- phase matrix = the number of drawers we fit into the wooden frame
- width of each drawer = pixel size in the frequency direction of the image
- frequency matrix = the number of socks we fit into each drawer.

In the phase axis of the image, if the number of pixels of a certain size increases, then the phase FOV increases, but in k -space, if there are more drawers, and the height of the wooden frame remains the same, then the step between each drawer decreases (so that we can fit them all in). The opposite is true if the number of pixels of a certain size decreases. The phase FOV of the image decreases, but in k -space, if there are fewer drawers, and the height of the wooden frame remains the same, then the step between each drawer can increase, as there are fewer drawers to fit in. Do you see why either an increase or decrease in the phase FOV is caused by the opposite change in the step between each line or drawer?

In the frequency axis of the image, if the number of pixels of a certain size increases, then the frequency FOV increases, but in k -space, if there are more socks and the width of each drawer remains the same, then the distance between each sock decreases (so that we can fit them all in). The opposite is true if the number of pixels of a certain size decreases. The frequency FOV of the image decreases, but in k -space, if there are fewer socks, and the width of each drawer remains the same, then the distance between each sock can increase, as there are fewer socks to fit in. Do you see why either an increase or decrease in the frequency FOV is caused by the opposite change in the distance between each data point?

PART 4: HOW DO PULSE SEQUENCES FILL k -SPACE?

The way in which k -space is traversed and filled depends on a combination of the polarity and amplitude of both frequency and phase-encoding gradients and the RF pulses. You may remember in Chapter 3, pulse sequences were defined as a series of RF pulses, gradient applications, and intervening time periods. Pulse sequences were likened to dances. All dances involve doing two things; moving arms and moving legs. The combination and timing determines the type of dance.

Using this analogy, imagine k -space as a dance floor. RF pulses and gradients are the two components of the pulse sequence dance, and they move the system around the k -space dance floor so that it is filled in different ways. The rules of the dance are as follows:

- The RF excitation pulse centers the system into the middle of k -space where the phase and frequency axes of k -space cross.
- The amplitude of the phase-encoding gradient determines how far up and down a line of k -space is filled. The steepest phase gradient slope in the acquisition determines the outermost lines filled both positively and negatively, and this, in turn, determines the size of the pixels in the phase direction of the image.
- The polarity of the phase-encoding gradient determines whether a line in the top or bottom half of k -space is filled.
 - Phase-encoding gradient *positive* – fills *top* half of k -space.
 - Phase-encoding gradient *negative* – fills *bottom* half of k -space.
- The amplitude of the frequency-encoding gradient (among other things) determines how far to the left and right k -space is traversed, and this, in turn, determines the size of the pixels in the frequency direction of the image.
- The polarity of the frequency-encoding gradient determines whether k -space is traversed from left to right or right to left.
 - Frequency-encoding gradient *positive* – k -space traversed from *left to right*.
 - Frequency-encoding gradient *negative* – k -space traversed from *right to left*.




Figure 6.21 How gradients traverse k -space in a gradient-echo sequence.

These processes are best described using an illustration of a typical gradient-echo sequence (Figure 6.21). The starting point is at the center of k -space when the pulse sequence begins with an RF excitation pulse. k -Space is then traversed from the center to the left to a distance (A) that depends on the amplitude of the negative lobe of the frequency-encoding gradient. The phase encode in this example is positive, and therefore a line in the top half of k -space is filled. The amplitude of this gradient determines the distance traveled (B). The larger the amplitude of the phase-encoding gradient, the higher up in k -space is the line that is filled with data from the echo in that TR. Therefore, the combination of the phase-encoding gradient and the negative lobe of the frequency-encoding gradient determine at what point in k -space data storage begins.

The frequency-encoding gradient is then switched positively, and during its application data are read from the echo. As the frequency-encoding gradient is positive, data are placed in a line of k -space from left to right. The distance traveled depends on the amplitude of the positive lobe of the gradient and determines the pixel size in the frequency direction of the image. This is only one example of how k -space is filled. If the phase-encoding gradient is negative, then a line in the bottom half of k -space is filled in the same manner as above.

k -Space traversal in spin-echo sequences is more complex. The frequency-encoding gradient is typically applied positively at the same time as the phase-encoding gradient. This moves the point in k -space to the right side of a k -space line (which line depends on the slope and polarity of the phase-encoding gradient in that TR). The 180° RF rephasing pulse is then applied. This

Table 6.7 Things to remember – *k*-space traversal (the *k*-space dance).

The combination of RF pulses and gradients governs the type of pulse sequence, which in turn determines how <i>k</i> -space is traversed
The slice-select gradient determines which area of <i>k</i> -space is being traversed
RF excitation pulses result in the system centering itself in the center of <i>k</i> -space
RF rephasing pulses result in a point in <i>k</i> -space flipped to the mirror point on the opposite side of <i>k</i> -space
The polarity of the phase-encoding gradient governs whether a line in the top or bottom half of <i>k</i> -space is filled. Its amplitude controls which line is filled. In the image, this determines the pixel size in the phase direction
The polarity of the frequency-encoding gradient governs whether <i>k</i> -space is traversed from right to left or left to right. Its amplitude controls how far <i>k</i> -space is traversed in this axis. In the image, this determines the pixel size in the frequency direction and indirectly the frequency FOV

twists the *k*-space point from the right to the left and to the equivalent *k*-space line in the opposite portion of *k*-space, i.e. if before the 180° RF rephasing pulse line +120 is selected by the phase-encoding gradient, and the right side of that line is selected by the frequency encoding, then after the 180° RF rephasing pulse, that point in *k*-space is flipped to the left side of line -120. The frequency-encoding gradient is then applied again positively to fill the line with data points from left to right.

Now that we understand the rules of the *k*-space dance, we can use them to work out how any pulse sequence traverses *k*-space!

PART 5: OPTIONS THAT FILL *k*-SPACE

The way in which *k*-space is filled depends on how the data are acquired and manipulated to suit the circumstances of the scan. This is especially helpful to reduce the scan time. *k*-Space filling is manipulated in many imaging options, sequences, and types of acquisition. These include the following and are discussed in the indicated chapters:

- Rectangular FOV (see Chapter 7)
- Antialiasing (see Chapter 8)
- Turbo spin-echo (TSE) sequences (see Chapter 3)
- Respiratory compensation (see Chapter 8).

They also include these imaging options that are discussed in this chapter:

- Partial or fractional averaging or half Fourier
- Partial echo
- Parallel imaging
- Single-shot
- Spiral
- Propeller or radial
- Sequential and 3D acquisition.

A list of acronyms of the five main system manufacturers is provided at the beginning of the book. This includes some of the scan parameters and imaging options that alter how *k*-space is filled.


Table 6.8 *k*-Space filling options.

Option	Resolution	SNR	Scan time	Purpose
Partial averaging	Same	Less	Less	Reduces time when SNR is good
Partial echo	Same	Same	Same	Automatic for a short TE
Rectangular FOV (phase short axis)	Same	Less	Less	Reduces scan time when anatomy is rectangular
Antialiasing (GE/older Philips)	Same	Same	Same	Eliminates aliasing
Antialiasing (Siemens/newer Philips)	Same	More	More	Eliminates aliasing
Turbo spin-echo (for scan time)	Same	Same	Less	Reduces scan time
Respiratory compensation	Same	Same	Same	Reduces respiratory artifact
Parallel imaging (for scan time)	Same	Less	Less	Reduces scan time
Parallel imaging (for resolution)	More	Less	Same	Increases resolution

Partial, fractional averaging, or half Fourier

The negative and positive halves of *k*-space on each side of the phase axis are symmetrical and a mirror image of each other. As long as over half of the lines of *k*-space are filled during the acquisition, the system has enough data to create an image. For example, if only 75% of *k*-space is filled, only 75% of the phase-encoding steps are needed to complete the scan (Figure 6.22). The scan time therefore decreases, but fewer data are acquired so there is less signal. It is possible to extrapolate the missing data and end up with full *k*-space but, because the vertical phase axis of *k*-space is more likely to display motion artifacts (see Chapter 8), this is not usually done. Zeros are placed in the empty lines of *k*-space, and there are fewer data than when all the lines are filled.

Partial averaging is used where a reduction in scan time is necessary and where the resultant signal loss is not of paramount importance. A good example is in volume imaging where the scan times are quite long but where there is inherently good signal. Partial averaging allows a reduction in scan time but, as there are fewer data in *k*-space, truncation artifact is more likely (see Chapter 8).

**Figure 6.22** Partial Fourier.

Partial echo

Partial echo is performed when only part of the echo is read during application of the frequency-encoding gradient. The peak of the echo is usually centered in the middle of the sampling window. For example, if the frequency-encoding gradient is switched on for 8 ms, frequencies are digitized during 4 ms of rephasing and 4 ms of dephasing. This signal is mapped relative to the frequency axis of k -space, and the left half of the frequency area of k -space is the mirror image of the right half (see Fact 3). When a very short TE is required, the echo rephases sooner than with a long TE. This normally means that the frequency-encoding gradient is applied sooner. However, gradient limitations might mean that this is not possible and limits the minimum achievable TE. This problem is overcome by selecting partial or fractional echo. This technique switches on the frequency-encoding gradient as soon as it is possible to do so, but also moves the peak of the echo so that it is no longer centered in the middle of the sampling window and occurs sooner. This means that only the peak and the dephasing part of the echo are sampled, and, therefore, initially only half of the frequency area of k -space is filled (the right side of k -space). However, due to the right-to-left symmetry of k -space, the system can extrapolate the data in the right side and place it also in the left side. Therefore, although initially only the right side of k -space is filled with data, after extrapolation, both sides contain data, and overall no data are lost. Partial echo imaging is routinely used when a very short TE is selected in the scan protocol. This maximizes T1 and PD weighting and the slice number achievable for a given TR (Figure 6.23).

201

Parallel imaging

Parallel imaging or **sensitivity encoding** is a technique that fills k -space more efficiently than conventional imaging. It does this by filling multiple lines of k -space per TR (as in TSE). Unlike TSE, however, these lines are acquired by assigning them to certain coils that are coupled together in an array to enable them to acquire data simultaneously (see Chapter 10). Coils that are specifically designed for this purpose are needed. The coils are constructed with multiple coils or with multiple coil elements each having their own channel for carrying data back to the host computer. Multiple channels are selected, and the maximum number depends on the level and sophistication of the hardware and software of the system. In this example, a four-channel configuration is described.

Look at Figure 6.24. There are four coil elements (coils) shown, each collecting data that is sent to the host computer via a separate channel. Each channel is allocated a k -space line as follows:

- Channel 1 acquires line 1 and every fourth line thereafter.
- Channel 2 acquires line 2 and every fourth line thereafter.
- Channel 3 acquires line 3 and every fourth line thereafter.
- Channel 4 acquires line 4 and every fourth line thereafter.

Hence every TR, four lines of k -space are acquired. In the first TR period:

- Channel 1 acquires line 1.
- Channel 2 acquires line 2.
- Channel 3 acquires line 3.
- Channel 4 acquires line 4.




Figure 6.23 Partial echo.

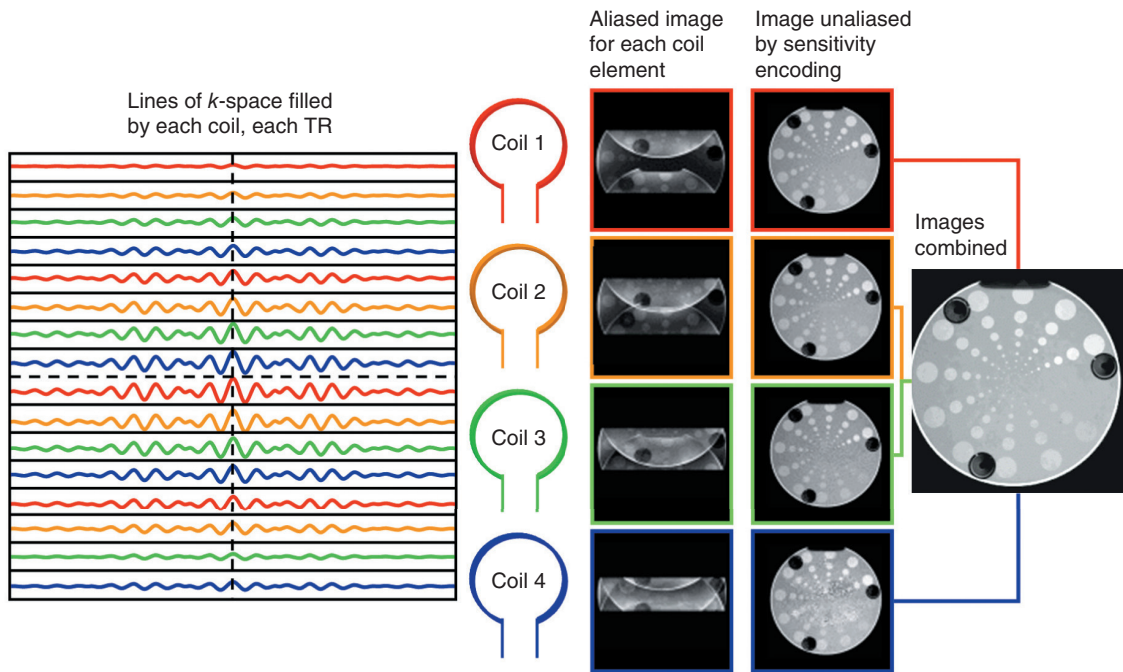


Figure 6.24 Parallel imaging.

In TR period 2:

- Channel 1 acquires line 5.
- Channel 2 acquires line 6.
- Channel 3 acquires line 7.
- Channel 4 acquires line 8, and so on.

The process is repeated until all k -space lines are filled. As four lines are acquired per TR in this example, the scan time decreases by a factor of 4. This is called the **reduction** or **acceleration factor** and is similar to the turbo factor in TSE (see Chapter 3). The reduction factor equals the number of channels in the configuration. The coil configuration can also be used to increase spatial resolution, e.g. achieve a phase matrix of 512 in the time of a 128. In addition, it is possible to combine the decreased scan time and improved spatial resolution. For example, two coils or channels are used to halve the scan time, and two are used to double the phase resolution achieved for a given scan time.

Now let's look at the lines acquired by each coil. You can see from Figure 6.16 that each channel acquires every fourth line of k -space and that, as a result, the gap between each line is four times greater than if k -space is filled in the conventional way. Using the chest of drawers analogy, this means that the depth of each drawer quadruples, and, as this dimension is inversely proportional to the size of the FOV in the phase direction, the phase FOV decreases to a quarter of its original size (see Part 5). As a result, aliasing of tissue outside the FOV in the phase direction occurs, and therefore each coil produces an aliased image (see Chapter 8).

The system rectifies this either in the image or in k -space. Strategies that rectify aliasing in the image use the sensitivity profile of each channel to calculate where signal is coming from relative to the coil so that it can correctly map it onto the image. This profile determines the position of signal relative to the coil based on its amplitude. Signal coming from near to the coil has a higher amplitude than that farthest away. Consequently, the image is unwrapped, and, using algorithms, the unwrapped data from each coil are combined to produce a single image.

Strategies that rectify aliasing in k -space fill extra central k -space lines during the scan. The data from these lines are used to produce low-resolution, unaliased images for each channel. These data and algorithms are used to unwrap the full-resolution images. The scan time savings of this method are not as good as the image rectification method because extra time is needed to fill the additional central k -space lines [9]. Parallel imaging can be used with all pulse sequences. Although it has obvious benefits in terms of scan time and/or resolution, it results in a loss of SNR. In addition, chemical shift may increase as different resonant frequencies are mapped across each coil. Patient movement also causes misalignment between undersampled data and reference scans.

Learning tip: What is an algorithm?

An **algorithm** is like a recipe. It is a set of steps that a computer completes to perform a task. This task might be a calculation or solution to a problem. Algorithms are frequently used in MRI. For example, in parallel imaging, the problem of how to unwrap aliased images is solved by a set of methodological steps (algorithm) by the system computer. Protocols are another example of algorithms. The problem (e.g. create a T1-weighted sagittal image of the brain) is solved by a set of programmed parameters, each of which is designed to produce the required images (see Chapter 7).

Single-shot

204

The scan time is significantly reduced by filling more than one line of k -space per TR. A good example is TSE (see Chapter 3). Taking this concept to the limits, the fastest possible scan time is where all the lines are filled at once. This is termed **single-shot (SS) imaging** and this method collects all the data required to fill lines of k -space from a single echo train. The echo train may consist of spin-echoes (generated by a train of RF rephasing pulses) termed single-shot fast/turbo spin-echo (SS-FSE or SS-TSE) or a train of gradient-echoes termed echo planar imaging (EPI) (see Chapter 4). To achieve this, multiple echoes are generated, and each is phase encoded by a different slope of gradient to fill all the required lines of k -space in a single-shot. For example, if a phase matrix of 128 is required, then an echo train of 128 echoes is produced and individually phase encoded to fill 128 lines of k -space. To fill k -space in this way, the read-out and phase-encoding gradients are rapidly switched on and off, and change direction. The frequency-encoding gradient switches from positive to negative; positively to fill a line of k -space from left to right and negatively to fill a line from right to left. This rapid change in gradient polarity also rephases the FID produced after the RF excitation pulse to generate gradient-echoes. As the frequency-encoding gradient switches its polarity so rapidly, it is said to oscillate.

The phase-encoding gradient also switches on and off rapidly, but its polarity does not need to change in this type of k -space traversal. Look at Figure 6.25. The first application of the phase-encoding gradient is maximum positive to fill the top line. The next application (to encode the next echo in the echo train) is still positive, but its amplitude is slightly less so that the next line down is filled. This process is repeated until the center of k -space is reached when the phase-encoding gradient switches negatively to fill the bottom lines. The amplitude gradually increases until it reaches maximum negative polarity, and the bottom line of k -space is filled. This type of gradient switching is called **blipping**, and single-shot imaging is its simplest form in that although all lines are filled in one shot, lines are filled linearly. There is no TR in single-shot imaging because there is only one RF excitation pulse at the beginning of the pulse sequence. A second RF excitation pulse does not happen. The TR is therefore infinity.




Figure 6.25 k -Space filling in EPI.




Figure 6.26 Spiral k -space filling.

Spiral k -space filling

A more complex type of k -space traversal is shown in Figure 6.26. In this example, both frequency and the phase-encoding gradients switch their polarity rapidly and oscillate. In this spiral form of k -space traversal, the frequency-encoding gradient oscillates to fill lines from left to right and then right to left, and k -space filling begins at the center. The phase-encoding gradient must also oscillate to fill a line in the top half followed by a line in the bottom half. To understand this more clearly, place a pen at the center of k -space on the diagram and work out the amplitude and polarity of each gradient as you move your pen along. In this example, the pen is never removed from the paper, indicating that there is no TR; all k -space is filled in one shot.

Propeller or radial k -space filling

In this imaging option, data points are acquired in strips, and these are rotated about the central axis of k -space. As the strips are rotated, the central portion of k -space is acquired every TR, and therefore signal and contrast increase. In addition, because the central portion of k -space is sampled every TR, this is equivalent to using multiple NSA and results in a reduction in motion artifact due to motion averaging (see Chapter 8). However, difficult math is required in propeller imaging. Algorithms are needed so that data are placed in the correct location in k -space as the strips of data rotate around the central axis. In the process, a lot of data are thrown out, and, therefore, scan times increase as it takes longer to fill k -space.

Sequential and 3D (volume) acquisition

There are three ways of acquiring data:

- Sequential
- Two-dimensional volumetric
- Three-dimensional volumetric.

206

Sequential acquisitions

Acquire all the data from slice 1 and then go on to acquire all the data from slice 2 (all k -space lines are filled for slice 1, and then all are filled for slice 2, etc.). The slices are therefore displayed as they are acquired.

Two-dimensional (2D) volumetric acquisitions

Fills one line of k -space for slice 1 and then goes on to fill the same line of k -space for slice 2, etc. When this line is filled for all slices, the next line of k -space is filled for slices 1, 2, 3, etc. This is the most common type of data acquisition and the version used in explanations earlier on in this chapter.

Analogy: Acquisition type and the chest of drawers



Let's go back to the chest of drawers analogy to explain the different types of acquisition. Imagine three chests of drawers representing three slices in the acquisition. Sequential acquisition is one in which all the drawers for a chest of drawers are filled before going onto the next chest of drawers. This type of acquisition might be used for breath-holding techniques. In this technique, images are displayed while the scan is still in progress. Once an entire chest of drawers (or k -space area) is filled, an image of the slice is displayed. Two-dimensional volumetric acquisition is one in which the top drawer in each of the three chests of drawers is filled in one TR, and then in the next TR, the next drawer down in each of the three chests of drawers is filled. This is the most typical type of acquisition and the one assumed for many explanations in this chapter (Figure 6.27).



**Refer to animation 6.3 on the supporting companion website for this book:
www.wiley.com/go/westbrook/mriinpractice**

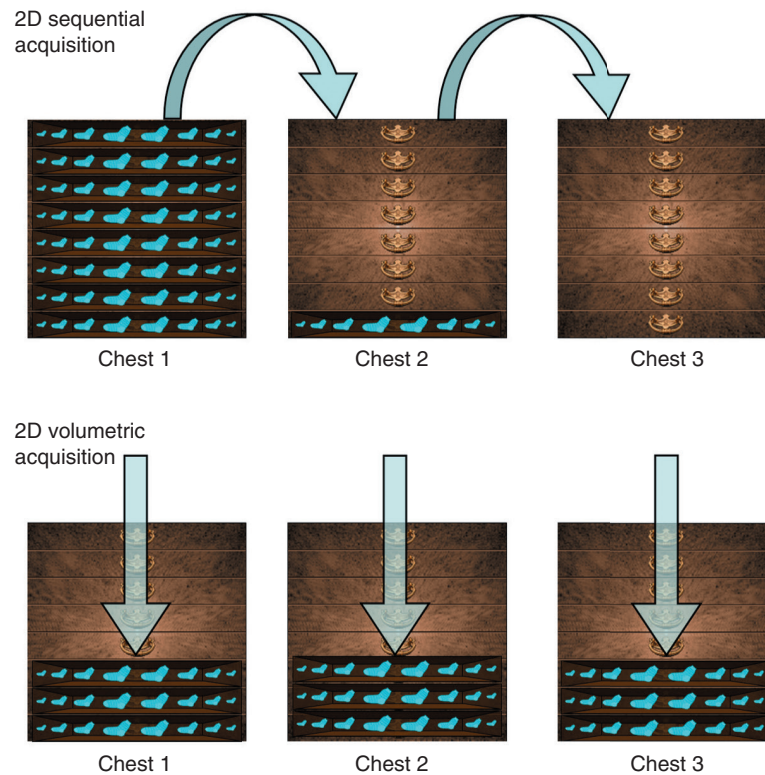


Figure 6.27 Data acquisition methods.

Three-dimensional (3D) volumetric acquisition (volume imaging)

In this type of acquisition, data are acquired from an entire volume of tissue rather than in separate slices. The RF excitation pulse is not slice-selective, and the whole prescribed imaging volume is excited. At the end of the acquisition, the volume or slab is divided into discrete locations or partitions by the slice-select gradient that, when switched on, separates slices according to their phase value along the gradient. This process is now called **slice encoding**. Many slices are possible (typically 128–256) without a slice gap. In other words, the slices are contiguous. The advantages of volume imaging are discussed in more detail in Chapter 7.

Well done for making it to the end of this difficult chapter! Hopefully you now understand why every parameter we select in the scan protocol is just changing how k -space is filled with data. In summary:

- The height of the chest of drawers determines the pixel size/resolution in the phase axis of the image.
- The width of the chest of drawers determines the pixel size/resolution in the frequency axis of the image.
- The number of drawers equals the phase matrix.

- The number of socks in each drawer equals the frequency matrix.
- The phase FOV is inversely proportional to the depth of each drawer.
- The frequency FOV is inversely proportional to the sampling interval.

It is not necessary to get too immersed in the math of k -space, but getting a good grip of the basic concepts is important because it makes everything else in MRI physics a lot easier to understand. In the next chapter, we start to bring all the physics together and apply it to practice, specifically how to optimize scan protocols.



For questions and answers on this topic please visit the supporting companion website for this book: www.wiley.com/go/westbrook/mriinpractice

References

1. Odaibo, S.G. (2012). *Quantum Mechanics and the MRI Machine*, 83. Arlington, VA: Symmetry Seed Books.
2. Hashemi, R.H., Bradley Jr, W.G., and Lisanti, C.J. (2010). *MRI: The Basics*, 3, 105. Philadelphia, PA: Lippincott Williams and Wilkins.
3. McRobbie, D.W., Moore, E.A., Graves, M.J. et al. (2017). *From Picture to Proton*, 161. Cambridge: Cambridge University Press.
4. Dale, B.M., Brown, M.A., and Semelka, R.C. (2015). *MRI: Basic Principles and Applications*, 5, 34. Wiley.
5. Hashemi, R.H., Bradley Jr, W.G., and Lisanti, C.J. (2010). *MRI: The Basics*, 3, 117. Philadelphia, PA: Lippincott Williams and Wilkins.
6. Dale, B.M., Brown, M.A., and Semelka, R.C. (2015). *MRI: Basic Principles and Applications*, 5, 32. Wiley.
7. Dale, B.M., Brown, M.A., and Semelka, R.C. (2015). *MRI: Basic Principles and Applications*, 5, 50. Wiley.
8. McRobbie, D.W., Moore, E.A., Graves, M.J. et al. (2017). *From Picture to Proton*, 117. Cambridge: Cambridge University Press.
9. McRobbie, D.W., Moore, E.A., Graves, M.J. et al. (2017). *From Picture to Proton*, 229. Cambridge: Cambridge University Press.

7

Protocol optimization

Introduction	209	Scan time	237
Signal-to-noise ratio (SNR)	210	Trade-offs	238
Contrast-to-noise ratio (CNR)	226	Protocol development and modification	238
Spatial resolution	232		

After reading this chapter, you will be able to:

- *Discuss factors that affect protocol optimization.*
- *Analyze parameters that influence each of these characteristics.*
- *Use what you have learned to effectively modify scan protocols.*

INTRODUCTION

In this chapter, we explore how to develop and modify scan protocols. We discuss common protocol parameters and how image characteristics and the acquisition are affected when these parameters are changed. Protocol optimization enables us to maximize image quality and acquire diagnostic images in the shortest scan time. These skills are a major part of clinical MRI practice.

Although it is common to view a protocol as a way of examining a certain area or pathology, e.g. brain protocol, tumor protocol, protocols must be considered in a much broader context than this. A **protocol** is defined as a “set of rules,” and in MRI these rules are a variety of different parameters that we select at the imaging console. They include extrinsic contrast parameters, geometry parameters, and a variety of imaging options and data acquisition methods. Protocols are judged by how well they show anatomy and pathology, and this is based on producing images that demonstrate the following four characteristics:

- High **signal-to-noise ratio (SNR)**
- Good **contrast-to-noise ratio (CNR)**

- High spatial resolution
- Short **scan time**.

In an ideal world, all four of these characteristics are achieved in every image. However, due to a variety of constraints, this is not usually possible. Optimizing parameters in favor of one of the aforementioned characteristics usually means compromising another. The skill lies in making informed decisions about which is most important for each patient and pathology, and using knowledge of underpinning physics to appropriately balance protocol parameters.

Let's investigate the four main characteristics that determine protocol optimization. Each is defined and then the factors that affect them are explored. This chapter also explains how changing any of these parameters affects another. These are called **trade-offs**. A list of acronyms of the five main system manufacturers is provided at the beginning of the book. This includes some of the scan parameters and imaging options described in this chapter. Scan tips are used to apply the theory of protocol optimization to practice.

SIGNAL-TO-NOISE RATIO (SNR)

The SNR is defined as the ratio of the amplitude of signal received to the average amplitude of the background noise.

- Signal is the voltage induced in the receiver coil by the precession of coherent magnetization in the transverse plane at, or about, time TE.
- **Noise** represents frequencies that exist randomly in space and time.

Signal is cumulative and predictable. It occurs at, or near, time TE and at specific frequencies at, or near, the Larmor frequency. It depends on many factors and can be altered. Noise, on the other hand, is not predictable and is detected by the whole coil volume [1]. It occurs at all frequencies and is also random in time and space. It is equivalent to the hiss on a radio when the station is not tuned in properly, and some of it is energy left over from the “Big Bang.” In the context of MRI, the main source of noise is from thermal motion in the patient [2] but it is also generated by background electrical noise of the system. Noise is constant for every patient and depends on the build of the patient, the area under examination, and the inherent noise of the system. The purpose of optimizing SNR is to make the contribution from signal larger than that from noise. As signal is predictable and noise is not, this usually means using measures that increase signal relative to noise, rather than reducing noise relative to signal.

Therefore, any factor that affects signal amplitude affects the SNR. These are as follows:

- Magnetic field strength of the system
- Proton density of the area under examination
- Coil type and position
- TR, TE, and flip angle
- Number of signal averages (NSA)
- Receive bandwidth
- Voxel volume (Equation (7.1)).

Equation 7.1

$\text{SNR} \propto (\text{voxel volume})$ $\text{SNR} \propto \sqrt{M(p) M(f) \text{NSA} / \text{RBW}}$	<p>M(p) is the phase matrix</p> <p>M(f) is the frequency matrix</p> <p>NSA is the number of signal averages</p> <p>RBW is the receive bandwidth (KHz)</p>	<p>This equation shows some of the parameters related to SNR. The parameters relate to the SNR by a square root of the total sampling time of the slice. Therefore, another way of expressing this is that the SNR is proportional to the voxel volume and to the square root of the total sampling time</p>
--	---	--

Magnetic field strength

The magnetic field strength plays an important part in determining SNR. As we discovered in Chapter 1, as field strength increases, so does the energy gap between high- and low-energy nuclei. As the energy gap increases, fewer nuclei have enough energy to align their magnetic moments in opposition to B_0 . Therefore, the number of spin-up nuclei increases relative to the number of spin-down nuclei. The NMV increases at higher field strengths, and there is more available magnetization to image the patient. SNR therefore increases. Although the magnetic field strength cannot be altered, when imaging with low-field systems, SNR may be compromised, and it might be necessary to alter protocol parameters that boost the SNR. This usually manifests itself in longer scan times.

Proton density

The number of protons in the area under examination determines the amplitude of received signal. Areas of low proton density in terms of those that are MR-active (such as the lungs) have low signal and therefore low SNR, whereas areas with a high proton density (such as the pelvis) have high signal and therefore high SNR. The proton density is inherent to the tissue and cannot be changed (that is why it is an intrinsic contrast parameter (see Chapter 2)). However, as the SNR is likely to be compromised when imaging areas of low proton density, steps may have to be taken to boost the SNR that are not necessary when scanning areas with a high proton density. In addition, when measures are taken to null or saturate signal from a tissue, SNR decreases because signal contribution from that tissue is removed (see Presaturation later in this chapter).

Type of coil

The type of coil affects the amount of received signal and therefore the SNR. Larger coils receive more noise in proportion to signal than smaller coils because noise is received from the entire receiving volume of the coil [1]. Coil types are discussed in Chapter 9. Quadrature coils increase SNR because several coils are used to receive signal. Phased array coils increase SNR as the data from several coils are added together. Surface coils placed close to the area under examination also increase SNR. The use of the appropriate receiver coil plays an extremely important role in optimizing SNR. In general, the size of the receiver coil should be chosen such that the volume of tissue optimally fills the sensitive volume of the coil. Large coils, however, increase the likelihood of aliasing, because tissue outside the FOV is more likely to produce signal (see Chapter 8). The position of the coil is also very important for maximizing SNR. To induce maximum signal, the coil must be positioned in the transverse plane perpendicular to B_0 . Angling the coil, as sometimes happens when using surface coils, results in a reduction of SNR (Figure 7.1).

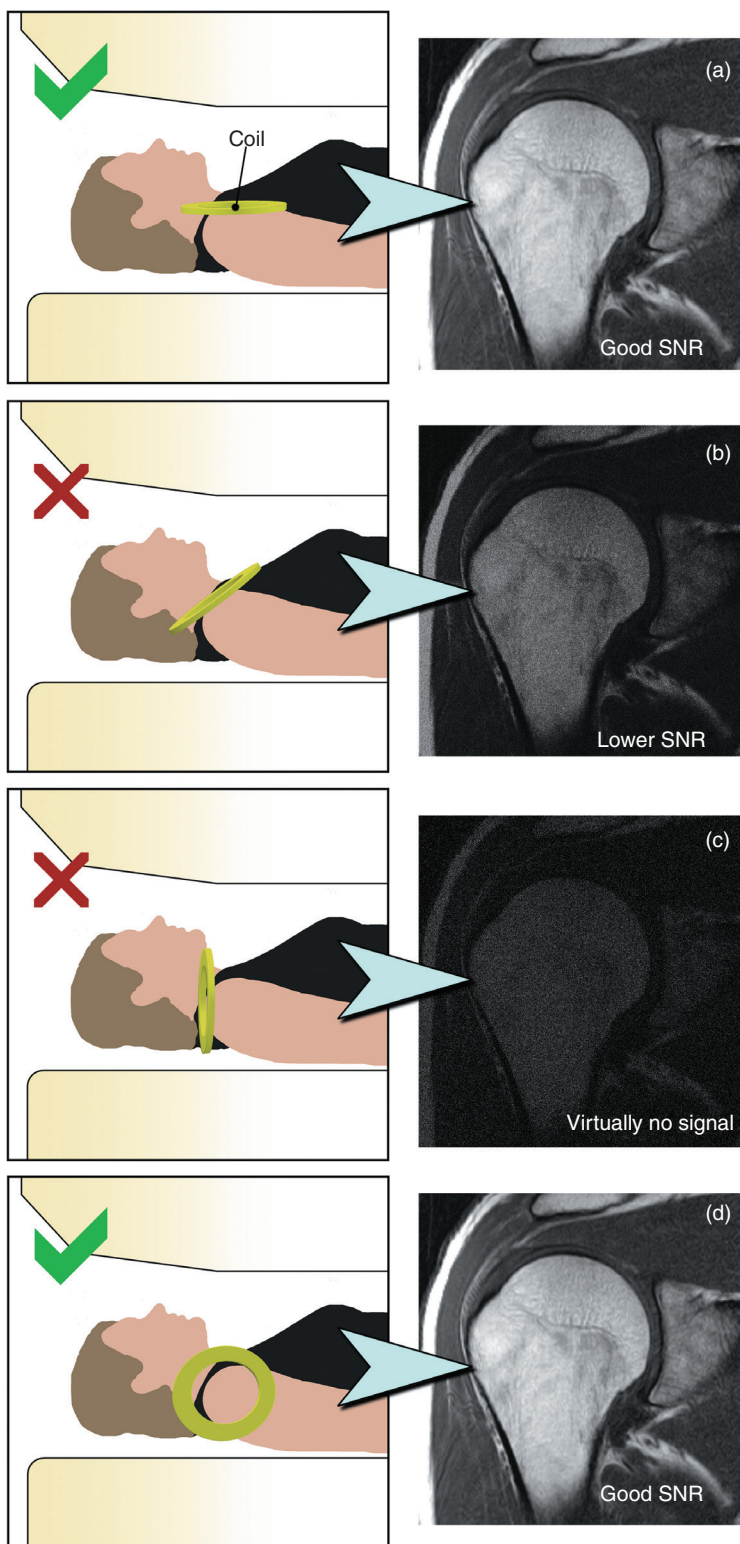


Figure 7.1 Coil position vs SNR.

TR, TE, and flip angle

Although TR, TE, and flip angle are usually considered parameters that influence image contrast (see Chapters 2–4); they also affect SNR and therefore image quality.

- The *TR* controls the amount of longitudinal magnetization that recovers before the next RF excitation pulse is applied. A long TR allows full recovery of longitudinal magnetization so that more is available to be flipped into the transverse plane in the next TR. A short TR does not allow full recovery of longitudinal magnetization so less is available to be flipped. Look at Figure 7.2 where the TR increases from 140 to 700 ms (at 1.5 T). The SNR improves as the TR increases. This is also seen in Figure 7.3 in images acquired at 3 T as the TR increases from 100 to 750 ms. However, as the TR is one of the factors that affect scan time (see Chapter 6), increasing the TR also increases scan time and the chance of patient movement.
- The *flip angle* controls the amount of transverse magnetization created by the RF excitation pulse, which induces a signal in the receiver coil (Figure 7.4). If the TR is long, maximum signal amplitude is created with flip angles of 90° because full recovery of longitudinal magnetization occurs with a long TR, and this is fully converted into transverse magnetization by a 90° flip angle. Look at Figures 7.5 and 7.6 in which the flip angle changes from 10° to 90° . SNR significantly decreases in the lower flip angle image. If the TR is short, the flip angle required to generate maximum transverse magnetization and therefore signal is less than 90° and is governed by the Ernst angle equation (see Equation (4.2)).
- The *TE* controls the amount of coherent transverse magnetization that decays before an echo is collected. A long TE allows considerable decay of coherent transverse magnetization before the echo is collected, while a short TE does not (Figure 7.7). Look at Figure 7.8 where the TE increases from 11 to 80 ms (at 1.5T). SNR decreases as the TE increases because there is less transverse magnetization available to rephase and produce an echo. This explains why T2-weighted sequences, which use a long TE, usually have a lower SNR than T1- or PD-weighted sequences that use a short TE. This is also seen in Figure 7.9 in images acquired at 3 T as the TE increases from 76 to 150 ms.

Number of signal averages (NSA or NEX)

The NSA controls the amount of data stored in each line of *k*-space (see Chapter 6). It is the number of times data are collected with the same amplitude of phase encoding slope and therefore how many times a line of *k*-space is filled with data. Doubling the NSA therefore doubles the amount of data that are stored in each line of *k*-space, while halving the NSA halves this amount. These data contain both signal and noise. Signal is additive over each signal average but the noise is not. It therefore increases by the factor of a square root [3]. For example, doubling the NSA only increases the SNR by $\sqrt{2}$ (≈ 1.4) (Figure 7.10). To double the SNR, the NSA and therefore the scan time are increased by a factor of 4 [4]. To triple the SNR requires a ninefold increase in NSA and scan time. Increased scan time increases the chances of patient movement.




Figure 72 (a) TR 700 ms, (b) TR 500 ms, (c) TR 300 ms, (d) TR 140 ms.




Figure 73 Changing TR at 3 T.




Figure 7.4 Flip angle vs SNR.




Figure 7.5 Axial gradient-echo image of the brain using a flip angle of 10° at 3 T.




Figure 7.6 Axial gradient-echo image of the brain using a flip angle of 90° at 3 T.




Figure 7.7 SNR vs TE.




Figure 7.8 (a) TE 11 ms, (b) TE 20 ms, (c) TE 40 ms, (d) TE 80 ms.




Figure 7.9 Changing TE at 3 T.




Figure 7.10 SNR vs NSA.

Look at Figures 7.11 and 7.12 where the NSA increases from 1 to 4. The SNR is undoubtedly greater in Figure 7.12 (exactly twice) but took four times longer to acquire than in Figure 7.11. Therefore, increasing the NSA is not necessarily the best way of increasing SNR because it results in a disproportionate increase in scan time. However, increasing the NSA also reduces motion artifact (see Chapter 8).




Figure 7.11 Sagittal brain using 1 NSA.




Figure 7.12 Sagittal brain using 4 NSA.




Figure 7.13 SNR vs receive bandwidth.

Receive bandwidth

This is the range of frequencies that are accurately sampled during the sampling window (see Chapter 6). Reducing the receive bandwidth results in less noise sampled relative to signal. By applying a filter, noise frequencies much higher and lower than signal frequencies are filtered out. Look at Figure 7.13. The areas shaded green and red represent the ratio of signal to noise, respectively (the squares are shaded orange where signal frequencies are the same as noise frequencies). In the left-hand diagram (which represents a broad receive bandwidth), there are 15 green signal squares and 7 red noise squares. Therefore, the SNR is approximately 2 : 1. In the right-hand diagram (which represents a narrow receive bandwidth), there are still 15 green signal squares but only 5 red noise squares. Therefore, the SNR increases to 3 : 1. Although the height of the signal curve is lower in the left-hand diagram, the area under each curve is the same (i.e. 15 green squares). The height of the signal curve in the left-hand diagram is lower because signal frequencies are spread over a wider frequency range. Therefore, as the receive bandwidth decreases, SNR increases as less noise is sampled as a proportion of signal. Halving the bandwidth increases the SNR by about 40%, but increases the sampling window [5]. As a result, reducing the bandwidth increases the minimum TE (see Chapter 6). Reducing the bandwidth also increases chemical shift artifact (see Chapter 8).

Scan tip: When to use a reduced receive bandwidth

Although these restrictions apply, there are some clinical situations where a narrow receive bandwidth is advantageous. Lengthening TEs are not important when a long TE is required for T2 weighting. In addition, chemical shift artifact only occurs when water and fat coexist in the same voxel. Therefore, reducing the receive bandwidth is a useful way to significantly improve SNR when performing T2-weighted images in conjunction with chemical saturation techniques, which remove

signal from either fat or water (see Chapter 8). Alternatively, broadening the receive bandwidth is often necessary when very short TEs are required. In fast gradient-echo pulse sequences where a very short TE is required, this is obtained by increasing the receive bandwidth. This decreases SNR because more noise frequencies are sampled relative to signal frequencies but, as the sampling window is short, the TE is also short (see Chapters 4 and 6). In TSE, broadening the receive bandwidth, which shortens the TE and sampling window, means that each echo in the echo train is sampled more time-efficiently. The consequence is that echo spacing decreases, and longer echo trains are permitted for a given TR. This reduces the scan time. Finally, shorter TEs obtained from broadening the receive bandwidth reduce magnetic susceptibility artifact (see Chapter 8).

Voxel volume

The building unit of a digital image is a **pixel**. The brightness of the pixel represents the strength of the MRI signal generated by a unit volume of patient tissue (**voxel**) (see Chapter 6). Voxel dimensions are determined by the pixel area and the slice thickness (Figure 7.14). The pixel area is determined by the **field of view (FOV)** and the number of pixels in the FOV or **image matrix**.

Large voxels contain more spins or nuclei than small voxels and therefore have more nuclei to contribute toward signal. Large voxels consequently have a higher SNR than small voxels (Figure 7.15). SNR is therefore proportional to the voxel volume, and any parameter that alters the size of the voxel changes the SNR. Any selection that decreases the size of the voxel decreases the SNR and vice versa. There are three ways to achieve this:

- *Changing the slice thickness.* In Figure 7.16, voxel size is altered by halving the slice thickness from 10 to 5 mm. Doing so halves the voxel volume from 1000 to 500 mm³ and hence halves the SNR.
- *Changing the image matrix.* The image matrix is the number of pixels in the image. It is identified by two numbers: one denotes the number of pixels in the frequency direction (usually the long axis of the image), the other one the number of phase pixels (usually the short axis of the image). Look at Figures 7.17 and 7.18 where the phase matrix increases from 128 (Figure 7.17) to 256 (Figure 7.18). As the FOV remains unchanged, there are smaller pixels and therefore voxels in Figure 7.18 than in Figure 7.17. Therefore, as the voxel volume halves, the SNR also halves.
- *Changing the FOV.* Look at Figures 7.19–7.21. The FOV halves, which halves the pixel dimension along both axes. Therefore, the voxel volume and SNR decrease to one quarter of the original value (from 1000 to 250 mm³). Comparing Figure 7.20 with Figure 7.21, it is evident that SNR significantly decreases in Figure 7.21. Depending on the area under investigation and the receiver coil, it is sometimes necessary to take steps to increase SNR when using a small FOV especially in conjunction with a large coil.

Learning tip: Volume imaging and SNR

In volume imaging, the entire volume of tissue is excited, and the volume contains no gaps, so the SNR increases. Slices are located by a technique known as slice encoding. This is another series of phase-encoding steps along the slice-select axis. Therefore, just as the number of phase-encoding steps increases the scan time in conventional spin-echo, the number of slices affects the scan time in volume imaging. The greater the number of slices, the longer the scan time (see Equation (6.9)). However, this is offset by the fact that the greater the number of slices, the greater the SNR, and so the NSA can be reduced [6]. In volume imaging, the number of slices is an added factor in Equation (7.1) so it is easy to see why SNR is higher in volume imaging compared to two-dimensional imaging [7].




Figure 7.14 The voxel. The large blue square is the FOV.




Figure 7.15 Voxel volume and SNR (spin numbers are not representative).




Figure 7.16 SNR vs slice thickness.




Figure 7.17 Sagittal brain using 128 phase matrix.




Figure 7.18 Sagittal brain using 256 phase matrix.




Figure 7.19 SNR vs FOV.




Figure 7.20 Sagittal brain using a square FOV of 240 mm.




Figure 7.21 Sagittal brain using a square FOV of 120 mm.

Table 7.1 SNR relationships.

$\text{SNR} \propto \text{FOV}$ $\text{SNR} \propto 1/\text{matrix}$ $\text{SNR} \propto S_t$ $\text{SNR} \propto \sqrt{\text{NSA}}$ $\text{SNR} \propto \sqrt{1/\text{RBW}}$	S_t is slice thickness (mm) RBW is the receive bandwidth in KHz	The proportional sign is used in these relationships as there are many other factors that affect SNR such as TR, TE, flip angle, and proton density
---	--	---

Scan tip: To optimize the SNR select the following parameters in the scan protocol

- Long TR and short TE
- 90° flip angle in spin-echo or the Ernst angle in gradient-echo pulse sequences
- The correct coil (ensure that it is well tuned and correctly positioned)
- A low or coarse image matrix
- A large FOV
- Thick slices
- High NSA
- Narrow receive bandwidth.

CONTRAST-TO-NOISE RATIO (CNR)

CNR is defined as the difference in the SNR between two adjacent areas (Equation (7.2)) [8]. It is controlled by the same factors that affect SNR. CNR is probably the most critical factor affecting image quality, as it directly determines the eye's ability to distinguish areas of high signal from areas of low signal. As image contrast depends on both intrinsic and extrinsic parameters (see Chapter 2), these factors also affect CNR. From a practical point of view, the CNR is increased in the following ways:

Table 7.2 Things to remember – SNR.

Signal amplitude is altered in several ways including using a long TR, a short TE, a large flip angle, and a good coil
Noise is random and largely not alterable although when using a narrow receive bandwidth, fewer noise frequencies are sampled
The SNR is therefore usually improved by increasing signal relative to noise rather than the other way around
The trade-offs associated with improving SNR are summarized in Tables 7.5 and 7.6

Equation 7.2

$$\text{CNR} = (SI_1 - SI_2) / \sigma$$

SI is the mean signal intensity from two regions 1 and 2
 σ is the standard deviation of the background signal or noise

This equation enables quantitative measurement of the relative CNR between two regions

- *T2 weighting.* Although T2-weighted images often have a lower SNR than T1-weighted images (due to the longer TE), the ability to distinguish pathology from normal tissue is often much greater because of the high signal of pathology compared with low signal of surrounding anatomy, i.e. the CNR is higher (Figure 7.22).
- *Contrast agents.* The purpose of administering contrast agents is to increase CNR between pathology (which enhances) and normal anatomy (which does not) (see Chapter 2).
- *Magnetization transfer contrast (MTC).* There is always a transfer of magnetization between the bound and the free protons, which causes a change in the T1 recovery times of free spins. This is exploited by selectively saturating the bound spins, which reduces the intensity of signal from the free spins due to MTC (see Chapter 2).
- *Flow-related techniques.* There are techniques that are specifically designed to produce signal only from nuclei with certain characteristics. Nuclei that do not possess such characteristics do not produce signal, and so there is a good CNR between them and those that do. For example, phase-contrast MR angiography is a technique that only images flowing spins. Stationary spins produce no signal, and so there is a good CNR between vessels and the tissue around them (see Chapter 8).




Figure 7.22 A heavily T2-weighted image of the buttock. A very long TE was used in this image. The CNR is optimized showing pathology clearly.

Presaturation

In addition to the strategies listed here, the CNR is optimized by using presaturation pulses. By saturating normal anatomy (which often contains fat), pathology (which is mainly water) is often seen more clearly especially if it has high signal intensity. Hydrogen exists in different chemical environments in the body, mainly fat and water (see Chapter 2). The precessional frequency of magnetic moments of fat nuclei is therefore slightly different from that of water. This is called **chemical shift**. As the magnetic field strength increases, this frequency difference also increases. For example, at 1.5 T, the precessional frequency between fat and water is approximately 220 Hz, so magnetic moments of fat nuclei precess 220 Hz lower than those of water nuclei. At 1.0 T, this frequency difference is reduced to 147 Hz. To saturate either fat or water, the precessional difference between the magnetic moments of both types of nuclei must be sufficiently large so that they can be isolated from each other. **Fat or water saturation** is therefore most effectively achieved on high-field systems.

To saturate fat signal, a 90° presaturation RF pulse is applied at the precessional frequency of fat to the whole FOV (Figure 7.23). The RF excitation pulse is then applied to the slices, and the magnetic moments of fat nuclei are flipped into saturation. If they are flipped to 180°, they do not have a component of transverse magnetization and produce no signal. The water nuclei, however, are excited, and their magnetic moments are rephased and produce signal (Figure 7.24). Water signal is also saturated in the same way as fat. The presaturation pulse is applied at the precessional frequency of water to the whole FOV (Figure 7.25) to specifically saturate these spins. Only the fat nuclei produce signal.

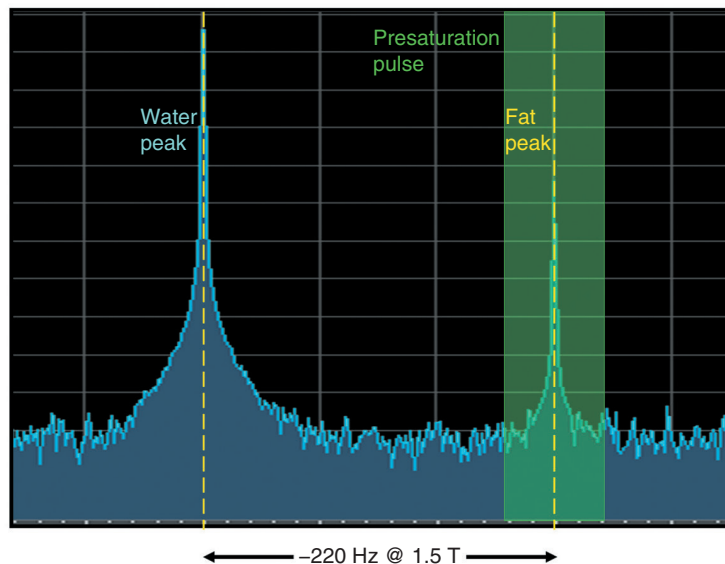


Figure 7.23 Fat saturation.




Figure 7.24 Sagittal ankle with fat saturation. Source: Westbrook 2014 [9]. Reproduced with permission of John Wiley & Sons.




Figure 7.25 Water saturation.

Effective presaturation requires an even distribution of fat or water throughout the FOV. Presaturation RF is transmitted at the same frequency and evenly to the whole FOV so that a particularly dense area of fat receives the same presaturation energy as an area with very little fat. Under these circumstances, fat saturation is less effective. In addition, the gradients applied for spatial encoding vary the frequency across each slice. Therefore, presaturation sometimes appears nonuniform across the slice or imaging volume. Optimal saturation occurs at the center of a slice or in the central portion of the imaging volume.

Scan tip: What is the SAT TR?

230

Presaturation pulses deliver extra RF and therefore reduce the number of slices available for a given TR because the SAR increases (see Chapter 10). The presaturation pulses are delivered to the FOV before the excitation of each slice. The interval between the presaturation pulses is called the **SAT TR** and is equal to the scan TR divided by the number of slices. If the SAT TR is longer than the T1 recovery times of fat or water, the magnetic moments of fat or water nuclei may not be saturated, as they have had time to recover before each presaturation pulse is delivered. To prevent this, prescribe the maximum number of slices available for a given TR so that the SAT TR is minimized.

Spatial inversion recovery (SPIR) is a variation of presaturation. An RF pulse at the precessional frequency of magnetic moments of fat nuclei is applied to the imaging volume but, unlike presaturation, this has a magnitude of 180° . The magnetic moments of fat nuclei are therefore totally inverted into the $-z$ direction. After time TI, which corresponds to the null point of fat, the 90° RF excitation pulse is applied. As fat has no longitudinal magnetization at this point, the RF excitation pulse produces no transverse magnetization in fat. Therefore, the fat signal is nulled (Figures 7.26 and 7.27).




Figure 7.26 Sagittal T2-weighted image of the pelvis without fat saturation.




Figure 7.27 Sagittal T2-weighted image of the pelvis with fat saturation.

Scan tip: SPIR vs STIR

SPIR combines fat saturation and inverting mechanisms used in the STIR pulse sequence to eliminate fat signal. It has several advantages. Standard presaturation is very dependent on homogeneity of B_0 as it requires the precessional frequency of magnetic moments of fat nuclei to be the same throughout the imaging volume. SPIR is much less susceptible to this because nulling also occurs by selecting an inversion time corresponding to the null point of fat. This depends on the T1 recovery time of fat rather than its precessional frequency, and relaxation times are not affected by small changes in homogeneity. However, as STIR relies on T1 recovery times to null signal rather than precessional frequencies, they are less likely to be affected by inhomogeneity than fat saturated methods such as SPIR or fat saturation. However, in STIR, gadolinium may be nulled along with fat. Gadolinium shortens the T1 recovery time of enhancing tissues so that it is close to the T1 relaxation time of fat (see Chapter 3). Therefore, STIR sequences should not be used after giving gadolinium. However, in SPIR sequences this does not occur because fat is selectively inverted and nulled, leaving gadolinium untouched. Therefore, SPIR may be used to null signal from fat in contrast-enhanced sequences.

It is possible to null signal from many types of tissues. This is achieved by applying a saturation pulse at the specific frequency of the tissue we wish to null before applying the RF excitation pulse. Tissue is also nulled by using an inverting pulse followed by an RF excitation pulse at a delay time equivalent to the null point of the tissue (see Chapter 3). Liver and spleen may

Table 7.3 Things to remember – CNR.

CNR is the difference in SNR between two adjacent areas
It is important to maximize the CNR so that pathology is clearly seen as distinct from normal anatomy or so that one structure is clearly seen next to another
CNR is improved by increasing the signal from pathology or structures that are important to see (e.g. positive contrast agents, T2 weighting, flow techniques)
CNR is improved by decreasing signal from normal structures (e.g. chemical suppression, MTC)

be specifically nulled, as can materials such as silicone. Presaturation is also useful to reduce artifacts such as phase mismapping and aliasing (see Chapter 8).

Out-of-phase imaging (**Dixon technique**) is used in gradient-echo sequences to null signal from voxels in which fat and water nuclei coexist. This is achieved by selecting a TE when the magnetic moments of nuclei in fat and water are out of phase with each other. As they are incoherent, no signal is received from the voxel (see more on the phase difference between fat and water in Chapter 8).

Scan tip:

To optimize the CNR select the following parameters in the scan protocol

- Very long TE and TR
- Saturation techniques to null certain tissues
- Techniques that remove signal from certain nuclei (e.g. phase contrast angiography).

SPATIAL RESOLUTION

Spatial resolution is the ability to distinguish between two points as separate and distinct, and is controlled by the voxel size. Small voxels result in high spatial resolution, as small structures are easily differentiated. Large voxels, on the other hand, result in low spatial resolution, as small structures are not resolved so well. In large voxels, individual signal intensities are averaged together and not represented as distinct within the voxel. This is called **partial voluming**. The voxel size is affected by:

- slice thickness
- FOV
- number of pixels or image matrix

Thin slices increase spatial resolution in the slice-select plane, while thick slices reduce spatial resolution. However, thinner slices result in smaller voxels, and, because there are fewer spins in small voxels, SNR decreases (see Figure 7.15).

The image matrix determines the number of pixels in the FOV. For a given FOV, a high or **fine matrix** divides the FOV by more pixels than a low or **coarse matrix**. Therefore, each pixel is smaller. Small pixels increase spatial resolution, as they increase the ability to distinguish between two structures close together in the patient. Increasing the image matrix therefore increases the spatial resolution. However, high image matrices result in smaller voxels, and therefore SNR decreases (see Figure 7.18).

The FOV also determines pixel dimension. For a given matrix, a large FOV results in large pixels, while a small FOV produces small pixels. Increasing the FOV therefore decreases spatial resolution. However, a small FOV results in smaller voxels, and therefore SNR decreases (see Figures 7.19 and 7.21).

Scan tip: Pixel dimension – what’s going on behind the scenes

MRI systems employ a variety of different methods to alter voxel geometry. We are able to individually select FOV dimensions, image matrix, and slice thickness. Alternatively, we may just be permitted to control voxel size overall. If this is the case, then, behind the scenes, the system alters the FOV, image matrix, and slice thickness to produce voxels of the required dimension. Whichever method is used, the principles are the same. If the FOV is fixed and square, and the phase matrix is less than the frequency matrix, pixels are longer in the phase direction than in the frequency direction. Spatial resolution is therefore reduced along the phase axis. If the FOV is rectangular, the pixels are square if the matrix produces a pixel with the same dimensions along the phase as well as frequency. Square pixels always provide better spatial resolution than rectangular pixels, as the image is equally resolved along both the frequency and phase axes.

The terms resolution and matrix are often confused. Although the image matrix is one of the parameters that determine resolution, they are not the same thing. Resolution refers to the *size* of each pixel, whereas matrix is how *many* pixels there are of that size. The size of each pixel multiplied by the image matrix determines the FOV dimension. As we discovered in Chapter 6, spatial resolution parameters depend on certain *k*-space characteristics. Let’s recap:



- *Pixel size* is determined by the distance traveled in *k*-space. Pixel size in the phase axis of the image is determined by the steepest phase-encoding gradient, both positively and negatively (height of the chest of drawers). Pixel size in the frequency axis of the image is determined by the sampling window (width of the chest of drawers).
- *Image matrix* depends on the number of data points in *k*-space. The phase matrix is the number of *k*-space lines (or drawers if you prefer). It is the number of data points in each column of *k*-space. The frequency matrix is the number of data points in each line of *k*-space (or pairs of socks if you prefer). It is the number of data points in each row of *k*-space.
- *FOV* depends on the distance between each data point in *k*-space. The size of the FOV in the phase axis of the image is inversely proportional to the distance between data points in each column of *k*-space. This is the incremental step between each line of *k*-space (or the depth of each drawer). The size of the FOV in the frequency axis of the image is inversely proportional to the distance between data points in each row of *k*-space. This is the sampling interval between each data point.

Learning tip: How resolution affects the minimum TE

Resolution is controlled by voxel size. To achieve a small voxel and, therefore good resolution, thin slices, a small FOV and a fine matrix are required.

- The slice thickness is determined by the slope of the slice-select gradient. Therefore, to achieve thin slices, the slice-select gradient slope is steep.
- The size of the frequency FOV is determined by the sampling interval (and therefore the receive bandwidth, see Chapter 6) but also by the slope of the frequency-encoding gradient. To achieve a small FOV, the frequency-encoding gradient slope is steep.
- The phase matrix size is determined by the number of phase encodings steps. To achieve a fine phase matrix, a high proportion of the phase-encoding gradient slopes are steep both positively and negatively. The frequency matrix is determined by the number data points in each line of *k*-space. To achieve a fine frequency matrix, for a given digital sampling frequency (as determined by the receive bandwidth) and sampling window, the amplitude of the frequency-encoding gradient is steep.

If gradient slopes are steep during a pulse sequence because we select thin slices, fine matrices, or a small FOV, their rise times are greater. The **rise time** of a gradient is the time required for it to achieve the correct slope or amplitude (see Chapter 9). Steep gradient slopes result in a higher rise time for the gradient than shallow gradient slopes. Steep gradient slopes therefore stress the gradient coils more than do shallow gradient slopes. This, therefore, increases the minimum TE, as the system cannot sample the echo until all the gradient functions are complete (the frequency-encoding gradient application is not complete by the time the echo is read, but most of it is completed by this time). A small FOV, thin slices, and fine matrices increase the minimum TE and may result in fewer slices per TR. If the TE increases, it takes longer to select and encode each slice, and therefore fewer slices are permissible in a given TR.

Rectangular FOV

A rectangular FOV may be desired when scanning anatomy that has a smaller dimension in the phase axis than frequency. A **rectangular** or **asymmetric FOV** maintains spatial resolution because pixel size is unchanged, but reduces the scan time, as only a portion of the total number of phase-encoding steps are performed. The dimension of the FOV in the phase direction is reduced compared to that in the frequency direction and so should be used when imaging anatomy that fits into a rectangle, for example, a sagittal lumbar spine.

Scan tip: Rectangular FOV – what's going on behind the scenes



Rectangular FOV does not affect pixel size. There are, however, fewer pixels of this size in the phase direction of the image. This explains why the phase FOV decreases relative to the frequency FOV and a rectangular FOV is obtained. To understand this concept more clearly, imagine that our chest of drawers is made up of a wooden

frame, and drawers are inserted inside this frame. Pixel size in the phase axis of the image is determined by the height of the wooden frame. In Chapter 6, we referred to this dimension as the height of the chest of drawers. The number of drawers inserted into the frame is the number of pixels of that size, i.e. the phase matrix.

In Chapter 6, we explored the analogy of changing what we store in k -space from pairs of socks to thick, woolly sweaters. This increases the depth of each drawer because thick sweaters require more room than do pairs of socks. The depth of each drawer is inversely proportional to the phase FOV, so as the depth of each drawer increases, the phase FOV decreases. However, if each drawer is deeper, and the height of the wooden frame of the chest of drawers is unchanged, the number of drawers must decrease to fit them inside the frame. Fewer drawers mean fewer lines of k -space, and therefore the scan time decreases.

For example, let's assume that we select the following parameters in the scan protocol:

- Frequency FOV 256 mm
- Phase FOV (50%)
- Image matrix 256×256 .

The required 50% phase FOV means that in the phase direction of the image, the FOV measures 128 mm, and, as the frequency FOV is 256 mm, a rectangular FOV is obtained. However, pixel size remains the same. In this example, each pixel measures $1 \text{ mm} \times 1 \text{ mm}$ ($\text{FOV} = 256 \text{ mm}/256 \text{ pixels}$). In the frequency axis, 256 pixels of this size are displayed, whereas in the phase direction only 128 pixels of this size are displayed.

To achieve this, behind the scenes, the system maintains the height of the wooden frame of the chest of drawers (thereby maintaining the $1 \text{ mm} \times 1 \text{ mm}$ pixel size), but inserts 128 deep sweater drawers inside the frame instead of 256 shallower sock drawers. As the depth of each drawer doubles, the phase FOV halves (50%). As only 128 drawers instead of 256 drawers are filled, the scan time halves. Using this analogy, it is easy to see that some signal may be lost because fewer data are acquired (fewer lines or drawers are filled with data) (Figure 7.28).

Scan tip:

To optimize the spatial resolution select the following parameters in the scan protocol

- Thin slices
- Small FOV
- High/fine image matrix.

Learning tip:

Volume imaging and resolution

In a volume acquisition, equal resolution in every plane and at every angle of obliquity is obtained by ensuring that each voxel is symmetrical (*isotropic*). The voxel has equal dimensions

in every plane. If this is not so, the volume has poorer resolution in the planes other than the one in which it was acquired. For example, if an FOV of 240 mm and matrix of 256×256 are selected, each pixel has a dimension of 0.9 mm (FOV/matrix). If the slice thickness is 3 mm, resolution is worse when the voxel is viewed from the side. Under these conditions, the voxel is **anisotropic**.

236

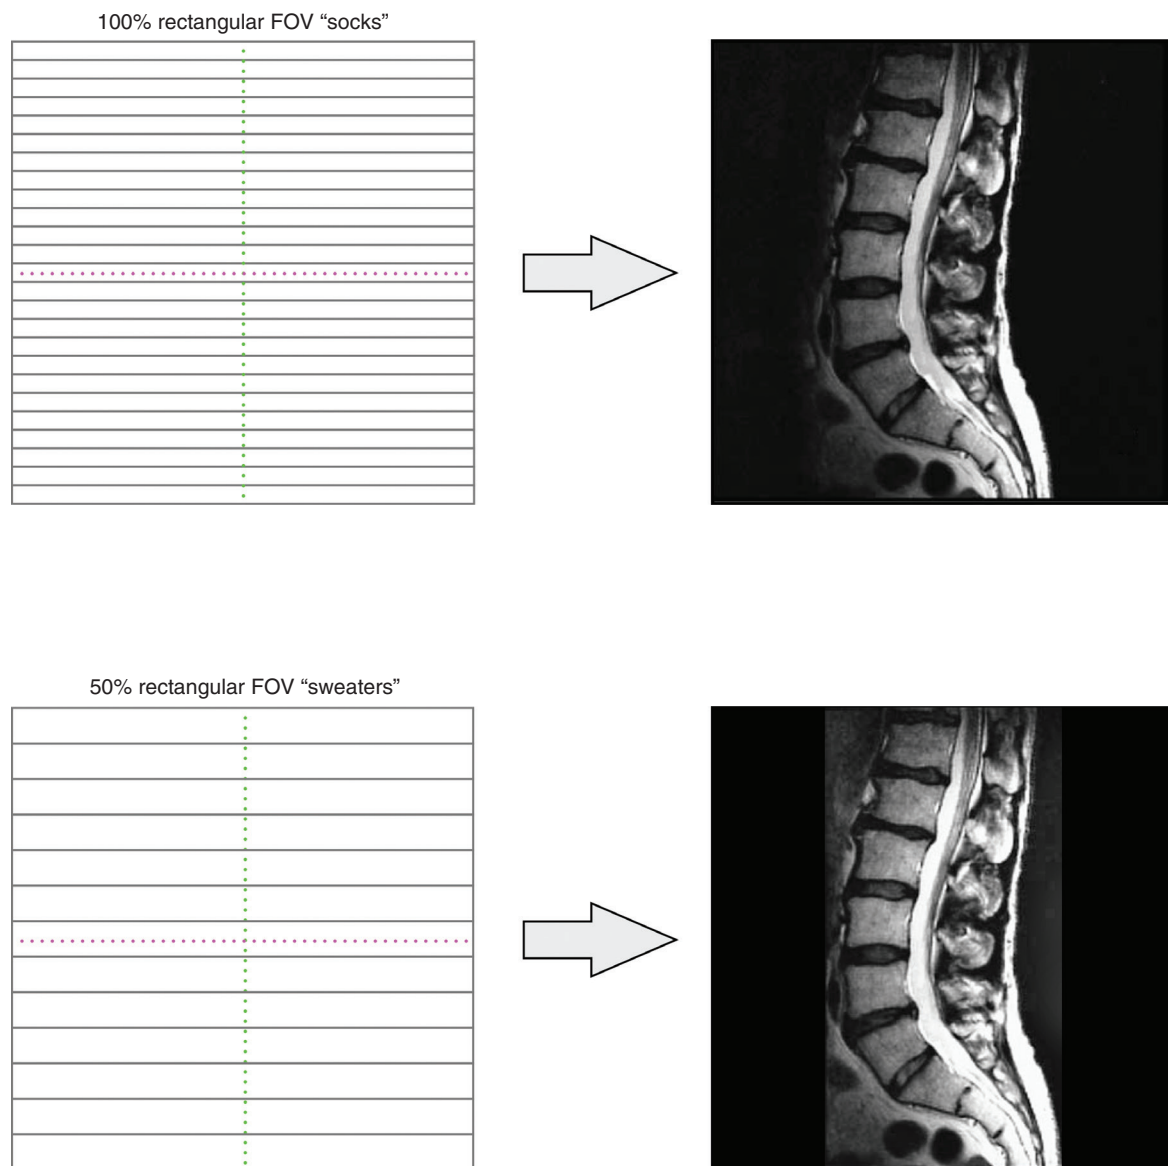


Figure 7.28 Square and rectangular FOV and the chest of drawers.

Table 7.4 Things to remember – spatial resolution.

Resolution is the ability to visualize two points that are close together in the patient as separate points in the image
Resolution depends on voxels being small and is therefore achieved by using a small FOV, a thin slice, and a high matrix
Small voxels result in poor SNR
In volume imaging, isotropic voxels give equal resolution in every plane
Increasing the slice number increases the SNR, but also increases the scan time
The trade-offs associated with improving resolution are summarized in Tables 7.5 and 7.6

SCAN TIME

The scan time is the time to complete data acquisition or the time to fill k -space (see Chapter 6). It is not a protocol parameter per se but determined by other factors that are selected in the scan protocol. As we discovered in Chapter 6, the scan time is proportional to the following:

- *TR*. The time of each repetition or the time between filling consecutive drawers or lines. Increasing the TR increases the scan time and vice versa.
- *Phase matrix*. The number of phase-encoding steps determines the number of lines of k -space or the number of drawers filled with data to complete the scan. Increasing the phase matrix increases the scan time and vice versa.
- *NSA*. The number of times data are collected with the same slope of phase-encoding gradient or the number of times each drawer or line is filled with data. Increasing the NSA increases the scan time and vice versa (see Equation (6.7)).

Scan time optimization is important, as long scan times give the patient more chance to move during the acquisition. Any movement of the patient is likely to degrade images. As multiple slices are selected during two- and three-dimensional volumetric acquisitions, movement during these types of scan affects all the slices. During a sequential acquisition, movement of the patient only affects those slices acquired while the patient is moving.

Scan tip:

To optimize the scan time select the following in the scan protocol

- Short TR
- Low phase matrix
- Low NSA
- High turbo factor in TSE
- Use imaging options that reduce scan time, e.g. rectangular FOV or partial Fourier (see Chapter 6).

Scan tip: **How to improve resolution but not increase scan time**

Usually, improving resolution requires a change in the phase matrix, which increases scan time. Sometimes, however, resolution can be increased without a corresponding increase in scan time. This is done by the following:

- Changing the frequency matrix only. The frequency matrix does not affect scan time but, if increased, increases resolution.
- Using asymmetric FOV. This maintains the size of the FOV along the frequency axis but reduces the FOV in the phase direction. Therefore, the resolution of a square FOV is maintained, but the scan time decreases in proportion to the reduction in the size of the FOV in the phase direction. This option is useful when anatomy fits into a rectangle.

TRADE-OFFS

It is evident that there are many trade-offs involved in protocol modification. Ideally, an image has high SNR, has good spatial resolution, and is acquired in a very short scan time. However, this is rarely achievable, as changing one parameter inevitably affects another. It is vital that we have a good understanding of all parameters that affect each image quality characteristic and their trade-offs. Table 7.5 lists the results of optimizing image quality. Table 7.6 provides common parameters and their associated trade-offs.

PROTOCOL DEVELOPMENT AND MODIFICATION

Protocol modification depends on the area under examination and the condition and co-operation of the patient. There are really no rules in MRI. This can be very frustrating when trying to learn, but also makes the subject interesting and challenging. Every facility has protocols established with the co-operation of the manufacturer and clinicians. However, never assume that your protocols are optimized. MRI is developing quickly so it is good practice to regularly review even the most well-established protocols. Here are a few tips on how to do this:

- Always choose the correct coil and position it correctly. This often makes the difference between a good, or bad-quality examination.
- Make sure that the patient is comfortable. This is very important, as a patient is more likely to move if they are uncomfortable. Immobilize the patient as much as possible to reduce the likelihood of movement.
- Try to ascertain from the radiologist exactly what protocols are required before the scan. This saves a lot of time, as radiologists are often difficult to track down!
- Spatial resolution is not generally as important in MR images as it is in X-radiation modalities. SNR is the most important image quality factor. There is no point in having an image with good spatial resolution if the SNR is poor. Sometimes, however, good spatial resolution is vital, but if the SNR is low, images are likely to be of poor quality, and the benefit of good spatial resolution is lost. It is worth remembering that, for example, a 1 mm difference in slice thickness makes all the difference in improving SNR without noticeably reducing the spatial resolution. Also remember that as the FOV size decreases, the dimensions of the pixel

along both axes are reduced (using a square FOV). Under these circumstances, the FOV is the most potent controller of SNR. Using a 160 mm FOV instead of an 80 mm FOV may be important in maintaining SNR.

- If the area under examination has inherently good signal (for example, the brain), and the correct coil is selected, it is usually possible for a fine matrix and fewer NSA to be used to achieve good-quality images in terms of SNR and spatial resolution. However, when examining an area with inherently low signal (for example, the lungs), it may be necessary to select more NSA and a coarser matrix.
- It is very important to keep the scan time as short as possible. There is no point in having an image with great SNR and spatial resolution if it took so long to acquire that the patient moves during the scan. Any patient can move – not just a restless one. The longer the patient is expected to lie on the table, the more likely that they move.

Table 7.5 Results of optimizing image quality.

To optimize image	Adjusted parameter	Consequence
Maximize SNR	↑ NSA	↑ Scan time
	↓ Image matrix (fixed FOV)	↓ Scan time (pMatrix)
	—	↓ Resolution
	↑ Slice thickness	↓ Resolution
	↓ Receive bandwidth	↑ Minimum TE
	—	↑ Chemical shift
	↑ FOV (fixed matrix)	↓ Resolution
	↑ TR	↓ T1 contrast (during incomplete recovery)
	—	↑ Number of slices
	↓ TE	↓ T2 contrast
Maximize resolution (assuming a square FOV)	↓ Slice thickness	↓ SNR
	↑ Image matrix (fixed FOV)	↓ SNR
	—	↑ Scan time (pMatrix)
	↓ FOV (fixed matrix)	↓ SNR
Minimize scan time	↓ TR	↑ T1 contrast (during incomplete recovery)
	—	↓ SNR (until incomplete recovery)
	—	↓ Number of slices
	↓ Phase matrix (fixed FOV)	↓ Resolution
	—	↑ SNR
	↓ NSA	↓ SNR
	—	↑ Movement artifact
	↓ Slice number in volume imaging	↓ SNR

Table 7.6 Parameters and their trade-offs.

Parameters	Benefits	Limitations
TR ↑	↑ SNR (until full recovery)	↑ Scan time
	↑ Number of slices	↓ T1 contrast (until full recovery)
TR ↓	↓ Scan time	↓ SNR (during incomplete recovery)
	↑ T1 contrast (during incomplete recovery)	↓ Number of slices
TE ↑	↑ T2 contrast	↓ SNR
TE ↓	↑ SNR	↓ T2 contrast
NSA ↑	↑ SNR	↑ Scan time
	↑ Signal averaging	
NSA ↓	↓ Scan time	↓ SNR
		↓ Signal averaging
Slice thickness ↑	↑ SNR	↓ Slice resolution
	↑ Coverage	↑ Partial voluming
Slice thickness ↓	↑ Slice resolution	↓ SNR
	↓ Partial voluming	↓ Coverage
FOV ↑ (fixed matrix)	↑ SNR	↓ Resolution
	↑ Coverage	
	↓ Aliasing (pFOV)	
FOV ↓ fixed matrix)	↑ Resolution	↓ SNR
		↓ Coverage
		↑ Aliasing (pFOV)
pMatrix ↑ (fixed FOV)	↑ Resolution	↑ Scan time
		↓ SNR if pixel small
pMatrix ↓ (fixed FOV)	↓ Scan time	↓ Resolution
	↑ SNR if pixel large	
Receive bandwidth ↑	↓ Chemical shift	↓ SNR
	↓ Minimum TE	
Receive bandwidth ↓	↑ SNR	↑ Chemical shift
		↑ Minimum TE
Large coil	↑ Area of received signal	↓ SNR
		Sensitive to artifacts
		Aliasing with small FOV
Small coil	↑ SNR	↓ Area of received signal
	Less sensitive to artifacts	
	Less prone to aliasing with small FOV	

Manipulating SNR, image contrast, spatial resolution, and scan time is a real art and takes some time and experience. Even after many years we all occasionally get things wrong! However, perseverance is important and eventually results in optimized protocols. Increased automation threatens to reduce these skills and may impact on the need for healthcare practitioners, especially in specialist practice [10, 11]. New MRI practitioners often learn how to operate software rather than MRI and risk demotion to mere “button-pushers.” However, educated professionals are able to make more precise and independent decisions than button-pushers, and this should be a way of investing in efficiency in the future. This is one of the many reasons why it is so important to understand protocol optimization.

In the next chapter, we explore image quality in more detail by evaluating the source of image artifacts and how to reduce or eliminate them.



For questions and answers on this topic please visit the supporting companion website for this book: www.wiley.com/go/westbrook/mriinpractice

References

1. Liney, G. (2010). *MRI from A to Z*, 2, 318. London: Springer.
2. Liney, G. (2010). *MRI from A to Z*, 2, 220. London: Springer.
3. Hashemi, R.H., Bradley Jr, W.G., and Lisanti, C.J. (2010). *MRI: The Basics*, 3, 177. Philadelphia, PA: Lippincott Williams and Wilkins.
4. Elmaoglu, M. and Celik, A. (2012). *MRI Handbook, MR Physics, Patient Positioning and Protocols*, 73. New York: Springer.
5. Hashemi, R.H., Bradley Jr, W.G., and Lisanti, C.J. (2010). *MRI: The Basics*, 3, 178. Philadelphia, PA: Lippincott Williams and Wilkins.
6. Westbrook, C. (2014). *Handbook of MRI Technique*, 4, 20. Wiley.
7. Hashemi, R.H., Bradley Jr. W.G., and Lisanti, C.J. (2010). *MRI: The Basics*, 3. Philadelphia, PA: Lippincott Williams and Wilkins, 179.
8. Liney, G. (2010). *MRI from A to Z*, 2, 62. London: Springer.
9. Westbrook, C. (2014). *Handbook of MRI Technique*, 4. Wiley Blackwell.
10. Frey, C.B. and Osbourne, M.A. (2013). The future of employment: How susceptible are jobs to computerization? http://www.oxfordmartin.ox.ac.uk/downloads/academic/The_Future_of_Employment.pdf (accessed 4 September 2017).
11. Harari, Y.N. (2017). *Homo Deus – A Brief History of Tomorrow*, 375. London: Vintage Books.

8

Artifacts

Introduction	242	Shading artifact	276
Phase mismatching	243	Moiré artifact	277
Aliasing	253	Magic angle	279
Chemical shift artifact	261	Equipment faults	280
Out-of-phase signal cancellation	265	Flow artifacts	280
Magnetic susceptibility artifact	269	Flow-dependent (non-contrast-enhanced) angiography	298
Truncation artifact	272	Black-blood imaging	303
Cross-excitation/cross-talk	273	Phase-contrast MRA	304
Zipper artifact	275		

After reading this chapter, you will be able to:

- *Recognize common MRI artifacts.*
- *Understand why these artifacts occur.*
- *Explain remedies for common MRI artifacts.*
- *Analyze the mechanisms of flow and how they are used to image vessels.*

INTRODUCTION

All MRI images have artifacts. Some artifacts degrade the image and may mask or even mimic pathology. It is therefore very important to understand their causes and how to compensate for them. Other artifacts are beneficial, and we deliberately create them to demonstrate flow, visualize pathology, or characterize lesions. Some artifacts are irreversible and are only reduced rather than eliminated. Others can be avoided altogether.

In this chapter, we discuss the appearances, causes, and remedy of the most common artifacts encountered in MRI. We also explore flow phenomena and how the artifacts they cause

are used to image flowing vessels. A list of acronyms of the five main system manufacturers is provided at the beginning of this book. This includes some of the parameters used to compensate for artifacts and angiography techniques described in this chapter. As with other chapters, scan tips link the theory of artifacts and flow to practice.

PHASE MISMAPPING

Appearance

Phase mismapping or **ghosting artifact** produces replications of moving anatomy across the image in the phase encoding axis. Phase mismapping usually originates from anatomy that moves periodically throughout the scan such as the anterior abdominal wall during respiration (Figure 8.1), pulsation of vessels and CSF, swallowing, and eye movement. Ghosting is less obvious with distance from the source of motion. The physical separation between each ghost depends on the scan time parameters (see Chapters 6 and 7) and period of motion (Equation (8.1)).

243

Equation 8.1

$$S_p = \frac{TR \times M(p) \times NSA}{T_m}$$

therefore

$$S_p = \frac{\text{scan time}}{T_m}$$

S_p is the separation between ghosts in pixels

TR is the repetition time in millisecond

M(p) is the phase matrix

NSA is the number of signal averages

T_m is the period of motion of the moving object

T_m is calculated thus: heart rate = 60 beats/min or 1 beat/s $T_m = 1$ s. This equation calculates the separation in pixels and if multiplied by the pixel size determines the actual distance between ghosts



Figure 8.1 Axial image through a breathing abdomen showing phase mismapping.

Cause

Phase mismatching is produced by anatomy moving along the phase-encoding gradient during the pulse sequence. Unlike the frequency encoding and slice-select gradients that are applied at the same amplitude every TR, the phase-encoding gradient is applied at a different amplitude (see Chapter 5). As anatomy moves during the scan, its reconstructed signal is misplaced in the phase encoding direction as the gradient amplitude changes. Imagine the anterior abdominal wall moving during the scan as shown in Figure 8.2. Anatomy is located at a position along the phase-encoding gradient during a particular TR period but may move to another position during

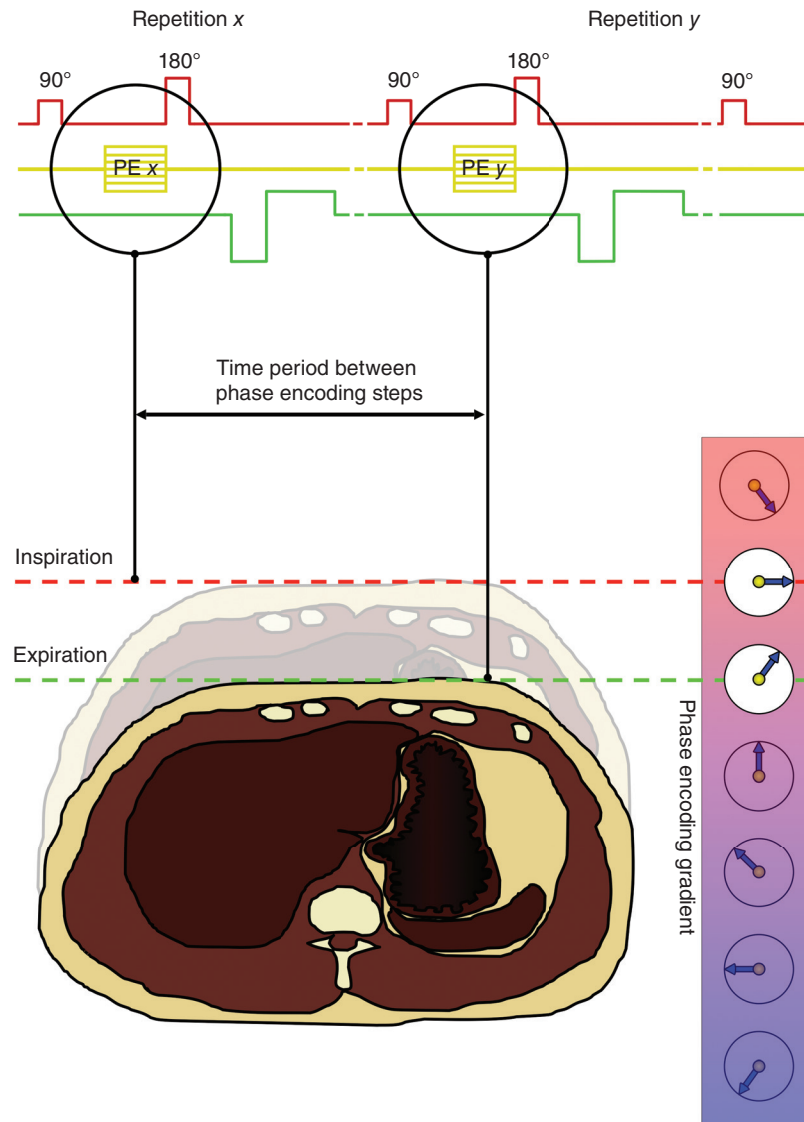


Figure 8.2 One of the causes of phase mismatching.

the next phase-encoding step. Signal from the abdominal wall acquires different phase values depending on its position along the gradient. Accurate reconstruction of the image relies on each line of data having a discrete incremental change in phase. Repetitive motion results in periodic perturbations of the data collected in k -space. After FFT, it is these perturbations that result in “ghosts” of moving anatomy being mismapped into incorrect spatial locations across the image. This is also sometimes described as **inter-view** or **view-to-view** mismapping.

Secondly, there is usually a time delay between phase encoding and readout (sometimes called intra-view mismapping), so anatomy may move between phase encoding and the echo. This time factor is very short, typically having a value in millisecond, so little motion is likely during this period. In addition, mismapping does not usually occur along the frequency axis of the image, as frequency encoding is performed as the frequencies in the echo are read and digitized.

Learning tip:

Phase mismapping and k -space

There is another reason why ghosting artifact only occurs in the phase encoding direction. The time interval between data points in each column of k -space is much longer than that between data points in each row. Each data point in k -space reflects the influence of the frequency and phase-encoding gradients at different time points during the scan. The sampling interval is the time between these time points. In the frequency direction of k -space (left to right), data are collected over a short time frame (in the order of a few millisecond – this is the sampling window). The sampling interval between each data point in a line of k -space is also short and is calculated by dividing the number 1 by the receive bandwidth. For example, if the receive bandwidth is 32 KHz, the sampling interval is $1 \div 32\,000$ or 0.031 ms (see Chapter 6).

In the phase axis of k -space (top to bottom), data are collected over the entire scan so the time frame is much longer than for data collection in each row (usually in the order of several minutes – this is the scan time). The sampling interval between each data point in a column of k -space is a TR period. This is because a new line of k -space is filled every TR, and a column of data points represents what happens from one TR period to the next as different lines are filled. The TR is usually several hundred or even thousands of millisecond.

Therefore, the sampling interval between a data point in a line of k -space is much shorter than the sampling interval between a data point in each column. There is, therefore, much more chance of something moving in the phase axis of k -space than in the frequency axis [1]. Can you think of anything that moves in the human body in 0.031 ms? However, there is quite a good chance that something moves in several hundreds or thousands of millisecond. This explains why motion artifact occurs in the phase encoding axis of the image, as this represents data in the phase axis of k -space.

Remedy

There are several ways to reduce phase mismapping. Equation (8.1) shows that the amount of separation between ghosts depends on the TR, phase matrix, NSA, and the period of motion. Therefore, if the TR, phase matrix, and NSA increase, the separation between ghosts increases




Figure 8.3 Axial T1-weighted image of the chest. Phase is anterior to posterior. Source: Westbrook 2014 [3]. Reproduced with permission of John Wiley & Sons.

and there are therefore fewer ghosts in the image. The same applies when imaging structures that pulsate rapidly because this reduces the period of motion [2]. There are other remedies, however, so let's discuss these in more detail.

Swapping phase and frequency

As ghosting only occurs along the phase encoding axis of the image, its direction can be changed so that artifact does not interfere with the area of interest. For example, imagine we select an axial stack of slices through the chest of a patient with a suspected coarctation of the aorta. Frequency encoding is usually performed right to left across the patient, as this is the longest axis of anatomy in the axial plane (Figure 8.3). The phase encoding direction is therefore performed anterior to posterior. Pulsation of the heart and great vessels along the phase encoding axis produces ghosting in this direction of the image and may obscure the aorta. Swapping phase and frequency so that frequency encoding occurs anterior to posterior, and phase encoding occurs right to left, places the artifact away from the area of interest (Figure 8.4). This strategy does not eliminate or reduce movement artifact. It repositions it so that it no longer obscures the area under investigation.




Figure 8.4 Axial T1-weighted image of the chest. Phase is right to left. Source: Westbrook 2014 [3]. Reproduced with permission of John Wiley & Sons.

Scan tip: Swapping the phase and frequency encoding axes

The system defaults to apply the frequency-encoding gradient along the longest axis of the anatomy and phase encoding along the shortest axis.

- Do you know why?
- When would you choose to swap these axes?
- What strategy might you have to use if you do?

There are circumstances where swapping phase and frequency is a good idea. Examples include sagittal knee or sagittal spine imaging where a major vessel runs from head to foot (popliteal and aorta, respectively). In sagittal imaging, the shortest axis is usually anterior to posterior but, if phase encoding occurs in this axis, then pulsation artifact from these major vessels runs from anterior to posterior right across the area under investigation. It is sometimes a good strategy to swap phase and frequency so that frequency is anterior to posterior, and phase is head to foot. This results in phase mismatching artifact passing behind the joint (knee imaging) and in front of the spine rather than straight over the region of interest [3]. If we swap the phase and frequency directions in the scan protocol, then it is likely that antialiasing options are needed because if phase is head to foot in sagittal imaging, there is anatomy outside the FOV in the phase direction, and this wraps into the image. Depending on your system, there may be penalties such as increased scan time or less signal averaging (see later).

Presaturation pulses

Presaturation nulls signal from specified areas. If we place presaturation pulses over the source of motion, signal from this motion decreases or is nulled altogether. For example, in sagittal imaging of the cervical spine, swallowing produces ghosting artifact along the phase axis (anterior to posterior) and obscures the spinal cord. A presaturation pulse placed over the throat reduces this artifact. Saturation pulses may also be positioned outside the FOV to saturate the signal from inflowing blood. This reduces phase mismapping from pulsatile flow.

Respiratory compensation techniques

248

In chest and abdominal imaging, respiratory motion along the phase encoding axis is likely to obscure important structures. There are several techniques that either reduce or eliminate this artifact. In fast sequences, it is usually possible for the patient to hold their breath during the scan. In longer sequences, a method known as **respiratory compensation** reduces motion artifact from respiration. Manufacturers may offer two types of respiratory sensor. The first kind is a physical device such as respiratory bellows positioned around the patient's chest, or a pad-shaped sensor that can be placed in proximity to moving anatomy. The second type uses a software-based strategy that tracks the position of the diaphragm using a navigator echo. In the first instance, a strap is placed around the patient's chest. The strap features a corrugated hollow tube (bellows) that expands and contracts as the patient breathes (Figure 8.5). This expansion and contraction causes a change in volume and air pressure within the bellows that are connected by hollow rubber tubing to a transducer that sends data to a slave computer responsible for physiological monitoring. A transducer is a device that converts the change in pressure within the bellows to an electrical signal. The system analyzes this signal, the amplitude of which corresponds to the maximum and minimum motion of the chest wall during respiration. It then reorders the lines of k -space so that when all the lines are filled, the data in k -space do not reflect the periodic motion of breathing throughout the scan. Modern systems also provide the option to use **respiratory navigator echoes** that are built into the pulse sequence. These RF pulses are designed to excite a narrow strip of spins and are typically positioned over the right hemidiaphragm. A sample is taken approximately five times per second and can track the position of the diaphragm over time. The amplitude of this signal corresponds to the maximum and minimum excursions of the diaphragm during respiration. The slave computer

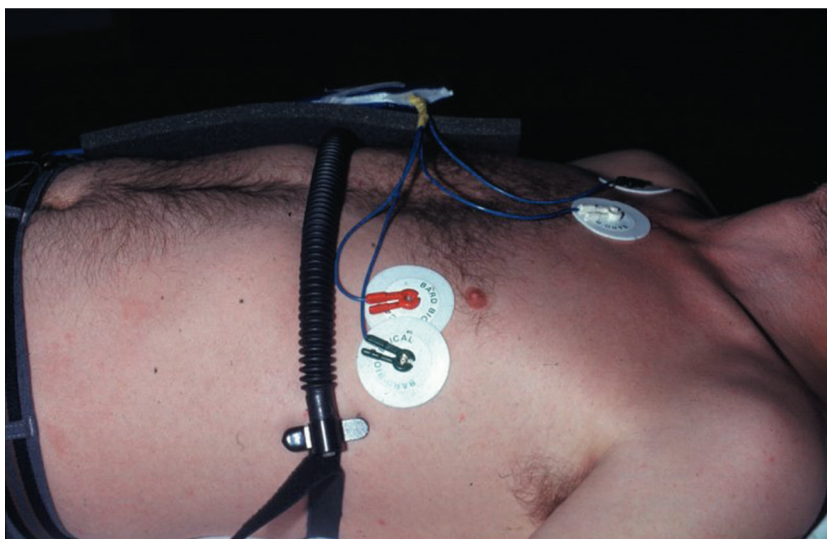


Figure 8.5 Placement of respiratory compensation and cardiac gating leads.

then reorders the lines of k -space so that when all the lines are filled, the data in k -space do not reflect the periodic motion of breathing throughout the scan.

Scan tip:

What's going on behind the scenes – respiratory compensation



As we learned in Chapter 6, the central lines of k -space are filled after shallow phase-encoding gradient slopes (that result in good signal and contrast), while the outer lines are filled after steep phase-encoding gradient slopes that result in high spatial resolution. The system reads the electrical signal from the measurement device or navigator and matches the order in which k -space lines are filled with the respiratory cycle [4].

When we select respiratory compensation in the scan protocol, behind the scenes, the system reorders the phase-encoding steps. The shallow phase-encoding gradient slopes that fill the central lines of k -space are filled when the chest or abdominal wall movement is minimal. Most of the data that provide signal and contrast are therefore acquired when chest wall motion is low. The steep phase encoding slopes that fill the outer lines are reserved for when the chest wall movement is maximum (Figure 8.6). Data positioned in k -space in this manner look nonperiodic, and therefore ghosting artifact from respiratory motion decreases. Look at Figures 8.7 and 8.8. Phase mismapping shown in Figure 8.7 is reduced by using respiratory compensation in Figure 8.8.




Figure 8.6 Respiratory compensation and k -space.




Figure 8.7 Image showing respiratory motion.




Figure 8.8 Image without respiratory motion.

Respiratory gating and triggering

These techniques utilize signal collected by the transducer or navigator echoes. In this context, the term **triggering** refers to an action that is triggered by an event. This could be the application of an RF pulse that is triggered by the inspiratory movement of the chest wall. **Gating** is slightly different in that we select a threshold (gate) that only allows data through when a certain condition is met. This could be used to cease data acquisition at a predefined threshold of the respiratory cycle and only

allow data to be collected when the diaphragm is relatively motionless at full expiration. In respiratory triggering, the RF excitation pulse is triggered by each respiratory excursion. Each echo is therefore obtained at the same phase of respiration. However, this method has two main drawbacks. Firstly, the periodicity of the patient's respiratory cycle becomes the TR and therefore determines the image contrast. This is likely to be around 15 breaths/min and limits the TR to around 4000 ms (or multiple of such). Secondly, the patient's rate of respiration may change during the acquisition leading to mixed contrast. Respiratory gating does not affect the TR because this technique permits us to select the normal user-defined parameters. The gate threshold is selected to only allow data through that have been collected when the diaphragm is relatively motionless and to reject data that are acquired at the extremes of the respiratory cycle. The image therefore only includes data collected within an acceptance range when the diaphragm is at a comparatively stationary neutral position between breaths. The main trade-off is that up to 60% of the data may be rejected, and this means that the acquisition time increases significantly in order to fill every line of data (Figure 8.9).

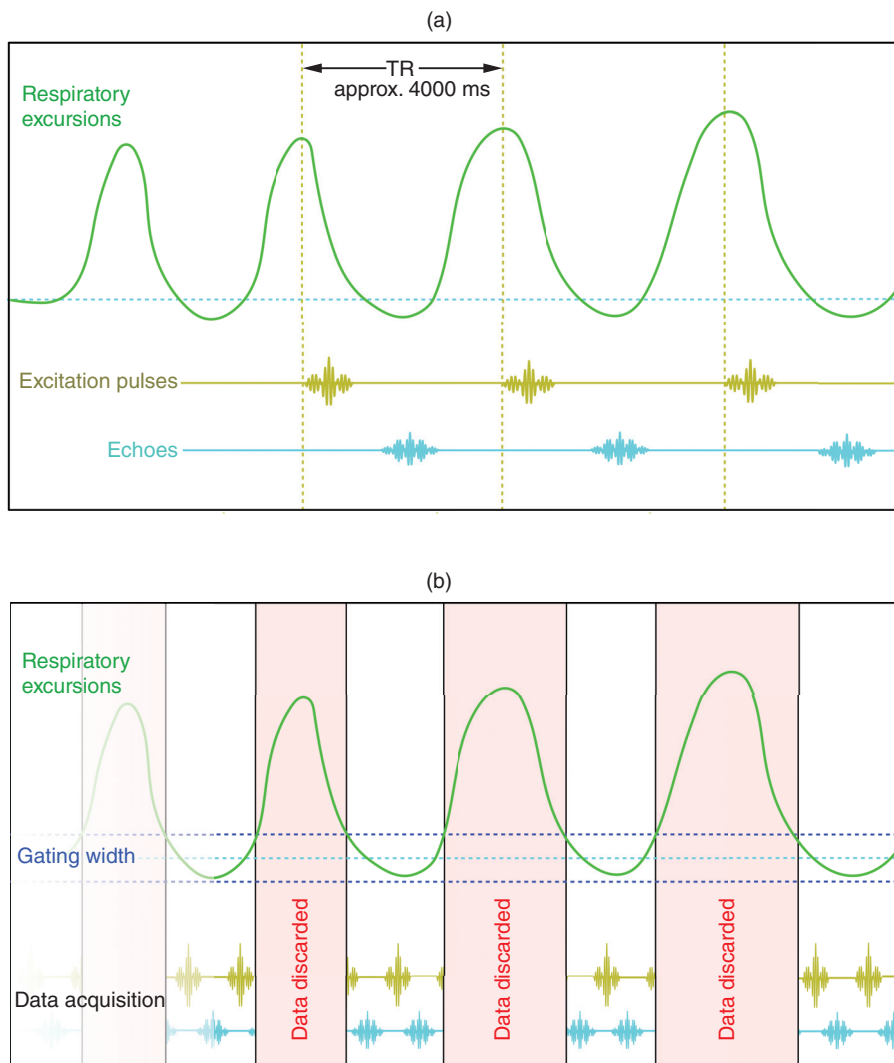


Figure 8.9 (a) Respiratory triggering and (b) gating.

Multiple NSA

Increasing the NSA may be a useful strategy to reduce motion artifact. Correctly encoded signal appears at the same position and is reinforced by every signal average. The mismatched signal that creates the ghosts is somewhat randomly located on each signal average and becomes less apparent because it averages out of the image with increasing NSA. The trade-off is that scan time is increased by an entire acquisition period for every additional signal average (see Chapters 6 and 7). This trade-off is compensated by using radial k -space filling in which the central region of k -space where most of the signal is stored is filled every TR, but the edges are only filled once. As a result, the image benefits from multiple signal averaging, but with a smaller trade-off in scan time (see Chapter 6). Voluntary motion is also reduced by making the patient as comfortable as possible and immobilizing them with pads and straps. A nervous patient always benefits from a thoughtful explanation of the procedure and a constant reminder to keep still over the system intercom. A relative or friend in the room can also help in some circumstances. In extreme cases, sedation of the patient may be required. Involuntary physiological motion, such as bowel peristalsis, is controlled by administering antispasmodic agents.

Cardiac triggering

Triggering may also be used in cardiac studies. Cardiac triggering monitors cardiac motion by coordinating the RF excitation pulse with the R wave of cardiac systole. This is achieved by using an electrical signal generated by cardiac motion to trigger each RF excitation pulse. There are two forms of triggering used:

- Electrocardiogram (ECG, EKG) gating uses electrodes and lead wires that are attached to the patient's chest to produce an EKG (see Figure 8.9). This is used to determine the timing of each RF excitation pulse. Each slice is acquired at the same phase of the cardiac cycle, and therefore phase mismapping from cardiac motion decreases. EKG triggering should be used when imaging the chest, heart, and great vessels.
- Peripheral pulse triggering uses a light sensor attached to the patient's finger to detect the pulsation of blood through the capillaries. The pulsation is used to trigger RF excitation pulses so that each slice is acquired at the same phase of the cardiac cycle. Peripheral triggering is not as accurate as EKG triggering, so it is not very useful when imaging the heart itself. However, it is effective at reducing phase mismapping when imaging small vessels or the spinal cord where CSF flow may degrade the image.

Learning tip: Gating in cardiac imaging

The term gating is also used in cardiac imaging to describe pulse sequences where data are collected continuously and retrospectively reordered into a stack of images that show the heart at different positions along the cardiac cycle. Although the term is not used in the same sense as respiratory gating, it serves to identify that data acquisition is not triggered by the R wave. Strictly speaking, this is not a remedy for phase mismapping but rather a data-sorting mechanism.

Gradient moment rephasing

This reduces ghosting caused by flowing nuclei moving along gradients (see later).

Table 8.1 Things to remember – phase mismapping.

Phase mismapping, ghosting, or motion artifact is caused by periodic motion mainly as a result of spins moving between each phase encode

It mainly originates from breathing and pulsatile motion of vessels and CSF

Respiratory compensation, gating, presaturation, and gradient moment nulling are the main techniques used to reduce this artifact

Artifacts and their remedies are summarized in Table 8.8

ALIASING

Appearance

Aliasing or **wrap** is an artifact where anatomy that exists outside the FOV is folded onto the top of anatomy inside the FOV. In Figure 8.10, the FOV in the phase direction is smaller than the




Figure 8.10 Sagittal image of the brain showing aliasing or wrap around.

anterior-to-posterior dimensions of the head. Therefore, signal outside the FOV in the phase direction is wrapped into the image.

Cause

Anatomy outside the FOV still experiences the effects of the gradients and produces a signal if it is within the receiving volume of the receiver coil. The signal from this anatomy has frequencies that are higher or lower than those within the FOV because nuclei are positioned on parts of the gradient that extend beyond the FOV. If the frequency exceeds the Nyquist frequency, it is not accurately digitized and is represented as a lower frequency [5].

254

Frequency wrap

Aliasing along the frequency encoding axis is known as **frequency wrap**. When the FOV is smaller than the anatomy in the frequency direction of the image, frequencies outside the FOV are higher than the Nyquist frequency and are mapped to a lower frequency. This is called **high-frequency aliasing** [6] (Figure 8.11, bottom image).




Figure 8.11 Aliasing and undersampling.

Phase wrap

Aliasing along the phase axis of the image is known as **phase wrap**. This is caused by undersampling of data along the phase axis of the image. Signal originating outside the FOV in the phase direction is allocated a phase value and therefore a pseudo-frequency that has already been given to signal originating from inside the FOV.

Look at Figure 8.12 where the FOV in the right-to-left phase encoding axis of the image is smaller than the dimensions of the abdomen. The phase-encoding gradient is applied in this direction and produces a change of phase across the bore of the magnet. This gradient is applied numerous times at different amplitudes during the acquisition and incrementally changes the phase position of the magnetic moments of each spin every TR (see Chapter 6). This produces a certain pseudo-frequency (i.e. the magnetic moments of each spin appear to have completed a number of cycles over the duration of the scan), and, using the gradient slope shown in Figure 8.12, this is the same as those inside the FOV (red and blue areas in the diagram). As they have the same pseudo-frequency, these red and blue areas are wrapped inside the FOV.

Learning tip: Calculating aliasing

The degree of aliasing is determined by Equation (8.2). The frequency that the system perceives is determined by comparing its actual frequency with the Nyquist frequency. For example, if the receive bandwidth is 64 KHz, the maximum frequency is +32 KHz, and the minimum is -32 KHz relative to the center frequency. Let's assume that some anatomy that exists outside the FOV has an actual frequency of 35 KHz because it is positioned further along toward the high end of the gradient. Using Equation (8.2), the perceived frequency is -29 KHz ($35 - [2 \times 32]$), which is much lower than the actual frequency of +35 KHz. Signal from this anatomy is therefore mapped to the opposite, low-frequency side of the FOV.


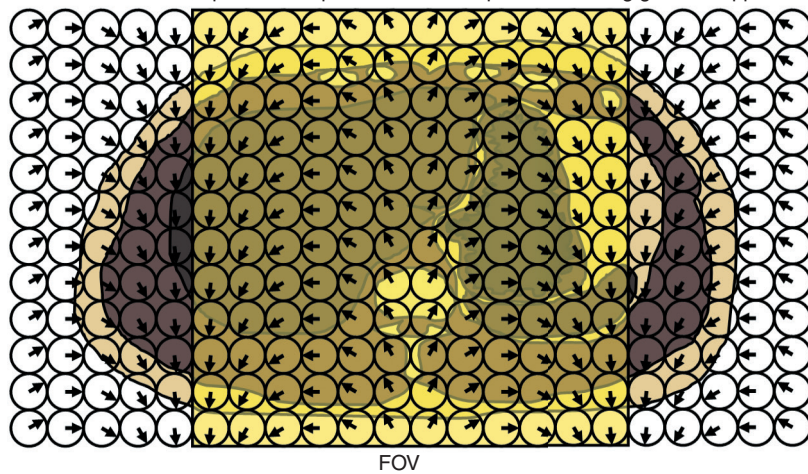
Equation 8.2

$f_p = f_t - 2$ (Nyquist frequency)	f_p is the perceived frequency in KHz f_t is the actual frequency in KHz	This equation calculates the degree of aliasing
-------------------------------------	---	---

Remedy

Aliasing along both the frequency and phase axes can degrade an image and should be compensated for. This is achieved by enlarging the FOV to incorporate all signal-producing anatomy. However, in the case of the phase FOV, this strategy increases the scan time. Another option is to use presaturation bands on areas outside the FOV that may wrap into the image. These can sometimes null signal from these areas and reduce aliasing. There are, however, two antialiasing software methods that compensate for wrap.

Axial abdomen slice, spins exhibit phase curve after phase encoding gradient application



Spins outside the field of view having same phase value as those outside




Figure 8.12 Phase wrap.




Figure 8.13 Antialiasing along the frequency axis.

Antialiasing along the frequency axis

Aliasing in the frequency direction is eliminated by increasing the digital sampling rate so that all frequencies are sufficiently digitized. This is achieved by decreasing the sampling interval while maintaining the same sampling window. The number of data points increases, but only those related to the central frequencies are displayed [6]. This is called a low-pass frequency filter. It essentially eliminates the frequencies that exceed the bandwidth [7] (Figure 8.13) and is like filtering out the bass and treble on a music system with a graphic equalizer.

Antialiasing along the phase axis

This is termed **no phase wrap**, **phase oversampling**, or **fold-over suppression**. Antialiasing software oversamples by increasing the number of phase encodings steps. This increases the number of k -space lines, and more data are stored so that there is no duplication of spatial frequencies. However, increasing the number of phase encoding increases the scan time. This is automatically compensated for by some manufacturers by reducing the NSA.

Scan tip:

Antialiasing what's going on behind the scenes



Antialiasing does not affect pixel size. There are, however, more pixels of this size in the phase direction of the image. This explains why the phase FOV increases relative to the frequency FOV, and aliasing is reduced or eliminated. To understand this concept more clearly, imagine that our chest of drawers is made up of a wooden frame, and

drawers are inserted inside this frame. Pixel size in the phase axis of the image is determined by the height of the wooden frame. In Chapter 6, we referred to this dimension as the height of the chest of drawers. The number of drawers inserted into the frame is the number of pixels of that size, i.e. the phase matrix.

In Chapter 6, we explored the analogy of changing what we store in k -space from pairs of socks to butterflies. This decreases the depth of each drawer because butterflies do not require as much room as pairs of socks. The depth of each drawer is inversely proportional to the phase FOV, so as the depth of each drawer decreases, the phase FOV increases. However, if each drawer is shallower, and the height of the wooden frame of the chest of drawers is unchanged, more drawers can be fitted inside the frame. More drawers mean more lines of k -space, and therefore the scan time increases (Figure 8.14).

For example, let's assume that we select the following parameters in the scan protocol:

- Frequency FOV 256 mm
- Antialiasing where the phase FOV is 200%
- Image matrix 256 × 256.

The required 200% phase FOV means that in the phase direction of the image, the FOV measures 512 mm, and, as it is twice as large as the frequency FOV, anatomy is likely to fit inside the larger FOV so that aliasing does not occur. The size of each pixel remains the same. In this example, each pixel measures 1 mm × 1 mm (FOV = 256 mm ÷ 256 pixels). In the frequency axis, 256 pixels of this size are displayed, whereas in the phase direction, 512 pixels of this size are displayed.

To achieve this, behind the scenes, the system maintains the height of the wooden frame (thereby maintaining the 1 mm × 1 mm pixel size), but inserts 512 shallow butterfly drawers inside the frame instead of 256 deeper sock drawers. If the depth of each drawer halves, the phase FOV doubles (200%). As 512 drawers are filled instead of 256, the scan time doubles. Using this analogy, it is easy to see that more data are acquired (more lines or drawers are filled with data) so the data are oversampled. Different systems have different ways in which antialiasing affects the acquisition. All techniques oversample the data so that the phase FOV increases, and, as there are more data in k -space, the system can more accurately locate signal inside and outside the FOV in the phase direction. However, there are two possible consequences of using antialiasing depending on which one your system uses:

- Option 1. The phase FOV increases to incorporate anatomy in the phase FOV, and the scan time increases proportionally. We decide whether to compensate for this by changing other parameters in the scan protocol.
- Option 2. The phase FOV automatically doubles, and the extended portions are discarded, leaving the original FOV. The scan time doubles, but this is automatically compensated for by the system by halving the NSA.

Option 2 means that phase mismatching artifacts might be more evident than expected because although 4 NSA may be displayed, for example, only 2 NSA are used. The SNR, however, remains largely unchanged. This is because oversampling of data compensates for the reduction in NSA (Figures 8.15 and 8.16).

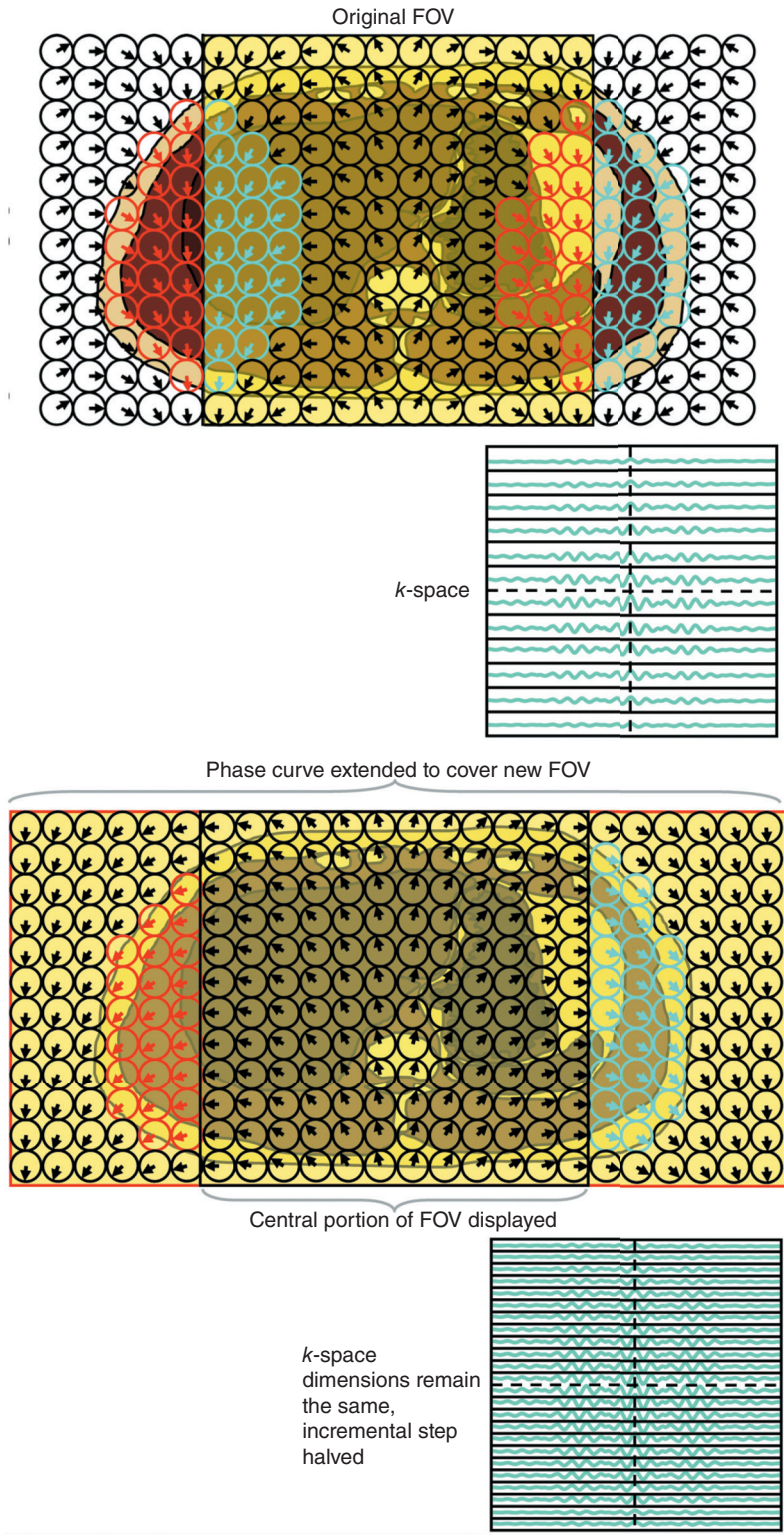


Figure 8.14 Antialiasing along the phase axis.




Figure 8.15 With wrap.




Figure 8.16 Without wrap.

Table 8.2 Things to remember – aliasing.

Aliasing is caused by undersampling of frequencies
If frequencies are not sampled often enough, the system cannot accurately represent those frequencies in the image
In the frequency-encoding direction of the image, this is avoided by ensuring that the digital sampling frequency is at least twice that of the highest frequency present, and additional low-pass filters are used
In the phase encoding direction of the image, aliasing is remedied by using antialiasing software. This increases the phase FOV, thereby oversampling during the phase encoding process. As a result, there is less likelihood that anatomy exists outside the larger FOV
Artifacts and their remedies are summarized in Table 8.8

Aliasing can also occur in the slice-select direction in volume imaging. This is because the slice-select gradient acts like a second phase-encoding gradient and is used to locate each slice within the imaging volume (see Chapter 6). Aliasing occurs because of undersampling during this process. The effect is reduced by ensuring the volume includes all of the anatomy lying outside the region under investigation. For example, signal from the patient's nose may otherwise wrap into the occipital area of the brain on the image.

CHEMICAL SHIFT ARTIFACT

Appearance

Chemical shift artifact causes the misplacement of signal from fat in the image (Figure 8.17). It can also create both signal voids and signal superimposition (high signal) in areas where




Figure 8.17 *Chemical shift.* In this image, the signal from the fatty bone marrow in the talus has been shifted inferiorly in the frequency-encoding direction. This is more apparent on the 16 KHz image where the signal void mimics a thickening of cortical bone.

fat and water interface. The renal area is a good example of this, as fluid-filled kidneys are surrounded by perirenal fat. The signal loss represents a boundary where the signal from fat has shifted by a certain number of voxels, leaving an area devoid of signal. The hyperintense area on the opposite side of the kidney is caused by the shifted fat signal being superimposed upon the signal from the underlying anatomy. The overall appearance gives an embossed effect where the anatomical structures appear to have been “lit” from one direction, casting a “shadow” in the opposite aspect. For this reason, the artifact is often referred to as the **bas-relief artifact**.

Cause

262

Chemical shift artifact is caused by different chemical environments of fat and water. Although fat and water both contain hydrogen atoms, fat consists of hydrogen arranged with a chain of carbon atoms, while in water, hydrogen is arranged just with a single oxygen atom (see Chapter 2). Fat is therefore a much larger molecule than water and the hydrogen atoms within each molecule are surrounded by many other atoms. Triglyceride fat molecules are sometimes described as a “self-shielding” because the electron clouds shield the hydrogen atoms from the static field B_0 . This self-shielding is much more pronounced in fat than in water and results in a lowering of the Larmor frequency of the magnetic moments of hydrogen nuclei in fat [8].

The difference in the precessional frequencies between magnetic moments in fat and water is called chemical shift, or sometimes fat/water shift. The amount of chemical shift is often expressed in arbitrary units known as parts per million (ppm) of the main magnetic field strength. Its value is always independent of the main field strength and equals 3.5 ppm. From this, the chemical shift between fat and water is calculated at different field strengths (Equation (8.3)). For example, at 1.5 T the difference in precessional frequency is approximately 220 Hz. The magnetic moments of fat nuclei precess 220 Hz lower than the magnetic moments of water nuclei. At 1.0 T, this difference is 147 Hz, and at lower field strengths (0.5 T or less), it is even lower and usually insignificant.

Equation 8.3

$$\omega_{\text{csf}} = \omega_0 \times C_s$$

ω_{csf} is the chemical shift frequency difference between fat and water (Hz)

ω_0 is the precessional frequency (Hz)

C_s is the chemical shift (3.5 ppm or 3.5×10^{-6})

At 1.5 T, for example, the precessional frequency is 63.86 MHz (63.86×10^6). Therefore, the chemical shift frequency difference between fat and water at 1.5 T is 220 Hz

Chemical shift artifact occurs because the difference in frequency between the magnetic moments of fat and water causes them to be placed into different pixels in the image. The receive bandwidth determines the range of frequencies that are accurately sampled and displayed across the FOV in the frequency direction of the image. The receive bandwidth and the number of frequency samples (or data points in each line of k -space) determine the bandwidth of each pixel. For example, if the receive bandwidth is ± 16 KHz, 32 000 Hz is mapped across the FOV. If 256 data points are collected, the FOV is divided into 256 frequency pixels. Each pixel therefore has an individual frequency range of 125 Hz/pixel ($32\ 000 \div 256$ Hz) (Figure 8.18). At a field strength of 1.5 T, the precessional frequency difference between the magnetic moments of fat and water nuclei is 220 Hz, and therefore using the above example, fat and water protons that exist adjacent to one another in the patient are mapped 1.76 pixels apart ($220 \div 125$) (Figure 8.18, middle diagram). The actual dimensions of this artifact depend on the size of the FOV, as this determines the size of each pixel.

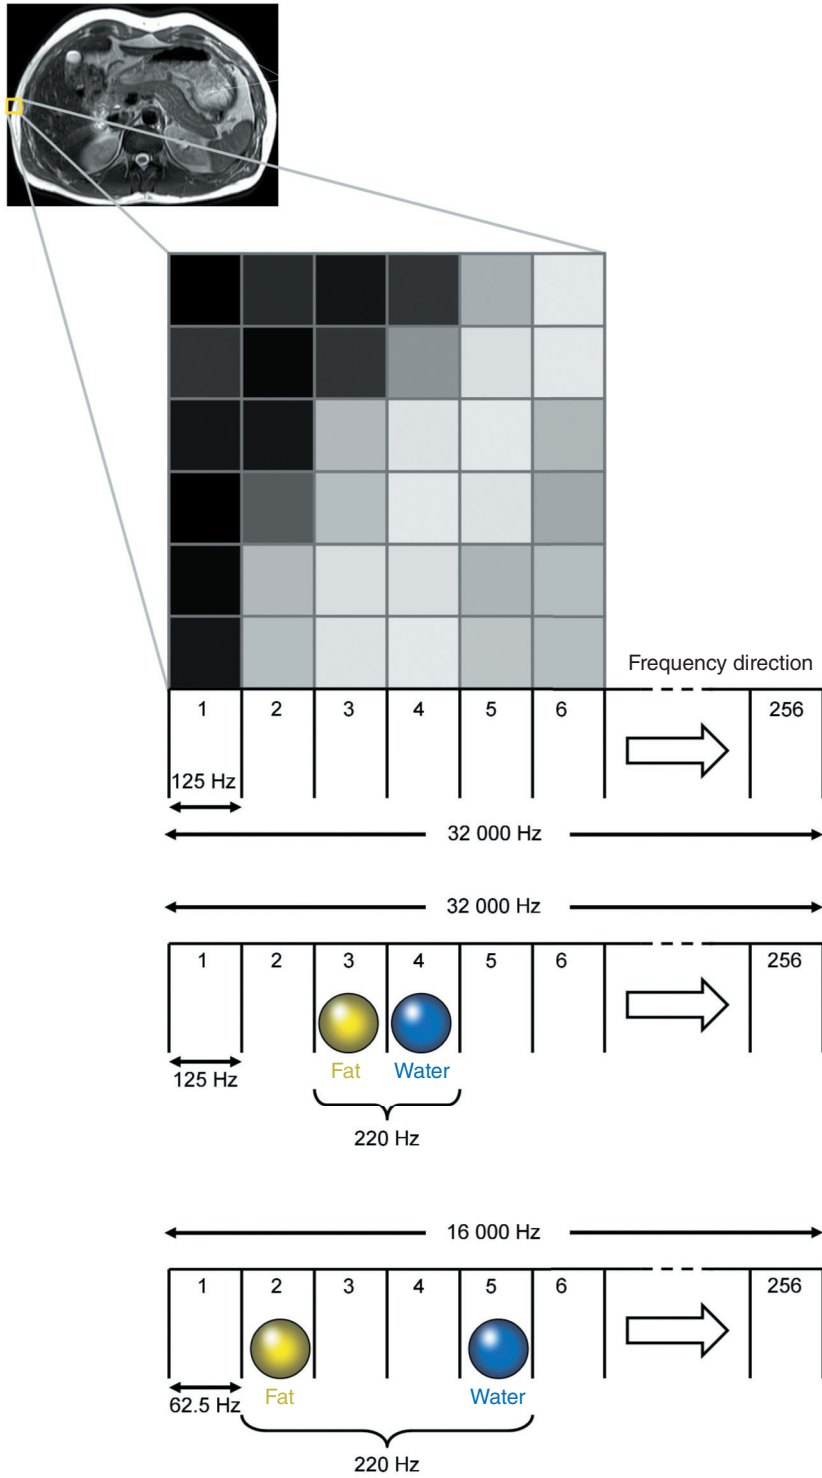


Figure 8.18 Chemical shift and pixel shift.

For example, an FOV of 240 mm and 256 frequency columns result in a pixel size of 0.93 mm. A pixel shift of 1.76 therefore results in an actual chemical shift between fat and water of 1.63 mm (0.93×1.76 mm) (Equation (8.4)).

Equation 8.4

$$CS_p = \frac{C_s \gamma B_0 \text{FOV}}{\text{RBW}/M(f)}$$

CS_p is the pixel shift (mm)

C_s is the chemical shift (3.5 ppm or 3.5×10^{-6})

γ is the gyromagnetic ratio (MHz/T)

B_0 is the main magnetic field strength (T)

FOV is the field of view (cm)

RBW is the receive bandwidth (Hz)

$M(f)$ is the frequency matrix

This equation calculates the pixel shift in millimeters caused by the chemical shift between fat and water. To calculate the actual number of pixels that fat and water are shifted, remove the FOV function from this equation

Remedy

Chemical shift artifact is reduced by scanning at a lower field strength and by minimizing the FOV. At high field strengths, chemical shift is limited by increasing the receive bandwidth. As the receive bandwidth increases, a larger frequency range is mapped across a certain number of frequency pixels (depending on the frequency matrix). The individual frequency range of each pixel (its bandwidth) therefore increases, so the 220 Hz difference in precessional frequency between fat and water is translated into a small pixel shift. For example, if the receive bandwidth is ± 32 KHz, 64 000 Hz is mapped across 256 frequency pixels. Each pixel has a bandwidth of 250 Hz ($64\,000 \div 256$ Hz). The 220 Hz precessional frequency difference between the adjacent fat and water protons at 1.5 T is translated into a pixel shift of less than 1 ($220 \div 250$). If the receive bandwidth decreases to ± 8 KHz, 16 000 Hz is mapped across 256 frequency pixels. Each pixel has a bandwidth of 62.5 Hz ($16\,000 \div 256$ Hz). The 220 Hz precessional frequency difference between adjacent fat and water protons at 1.5 T is now translated into a pixel shift of 3.52 pixels ($220 \div 62.5$) (Figure 8.18, lower diagram).

Scan tip:

How to reduce the receive bandwidth and avoid chemical shift artifact

In Chapter 7, we learned that SNR increases when using a narrow receive bandwidth. However, this strategy increases chemical shift artifact. One way to avoid this is to saturate signal from either fat or water (see Chapter 7). By doing so, there is nothing for one tissue to shift against, and therefore chemical shift artifact is eliminated.

Table 8.3 Things to remember – chemical shift artifact.

Fat and water precess at different frequencies. This is 3.5 ppm and is called chemical shift
Chemical shift causes a displacement of fat and water signals in the frequency direction of the image. This is dependent upon the field strength, the receive bandwidth, and the FOV
Artifacts and their remedies are summarized in Table 8.8

OUT-OF-PHASE SIGNAL CANCELLATION

265

Appearance

Out-of-phase signal cancellation produces a ring of dark signal around certain organs where fat and water interfaces occur within the same voxel (Figure 8.19). It degrades the image in gradient-echo pulse sequences because gradient rephasing does not compensate for this artifact.

Learning tip: Out-of-phase signal cancellation is different to chemical shift

Out-of-phase signal cancellation is sometimes referred to as “chemical shift of the second kind”; however, it must not be confused with the artifact described above. In fact, the term is a misnomer because no shift in signal occurs. Chemical shift artifact is caused by the difference in precessional frequency between the magnetic moments of hydrogen nuclei in fat and water. Fat signal is, therefore, shifted along the frequency encoding axis of the image. In out-of-phase signal cancellation, the artifact is created by the difference in phase position and occurs in any pixel that contains both fat and water. It can occur in any direction and location in the image. For example, in fatty infiltration of the liver, the entire organ may demonstrate a reduction in signal in the out-of-phase image.

Cause

Out-of-phase signal cancellation artifact is caused by the fact that the magnetic moments of hydrogen precess at different frequencies, and the differences in their phase positions are exhibited by fat and water vectors at discrete points in time. The magnetic moments of fat and water nuclei are in phase at certain times and out of phase at others. As the magnetic moments travel at different frequencies around their precessional paths, they are usually at different positions on the path, but periodically they are at the same position and, therefore, in phase. Sometimes, they are opposed by 180° and completely out of phase with each other.




Figure 8.19 Out-of-phase signal cancellation show as a black line around the abdominal organs border at the boundaries between fat and muscle.

Learning tip: Out-of-phase signal cancellation and the watch analogy



This artifact is analogous to the hour and minute hands of a watch. Both hands travel at different speeds around the watch: the hour hand moves through 360° in 12 h, while the minute hand moves the same distance in 1 h. However, at certain times of the day, the hands are superimposed or in phase, i.e. approximately at 12 noon, 1.05 a.m., 2.10 a.m., 3.15 a.m., etc. At other times, they are out of phase, i.e. approximately at 12.33, 1.38, 2.49, 3.50, etc. (Figure 8.20).




Figure 8.20 Out-of-phase signal cancellation and the watch analogy.



Refer to animation 8.1 on the supporting companion website for this book: www.wiley.com/go/westbrook/mriinpractice

Remedy

268

Select a TE that matches the periodicity of fat and water. The periodicity depends on the field strength (Figures 8.21 and 8.22). At 1.5 T, for example, a TE that is a multiple of 4.2 ms reduces out-of-phase signal cancellation, while at 0.5 T, this increases to 7 ms because the periodicity between fat and water is longer. In addition, use spin-echo sequences rather than gradient-echo pulse sequences, as 180° RF rephasing pulses are very effective at compensating for differences in phase between fat and water. Gradient rephasing does not have the same effect.

Although defined as an artifact, out-of-phase signal cancellation is clinically very useful. Any pathology that contains both fat and water exhibits reduced signal on an out-of-phase image. Adrenal tumors are a good example. Furthermore, out-of-phase images provide a method of background suppression in inflow angiography when examining the cerebral circulation. There are computer algorithms that can manipulate data from in-phase and out-of-phase images to provide fat-only and water-only images. This is known as the **Dixon Technique**.




Figure 8.21 The periodicity of fat and water.




Figure 8.22 Axial gradient-echo images when the magnetic moments of fat and water nuclei are out of phase (below) and in phase (above). Source: Westbrook 2014 [3]. Reproduced with permission of John Wiley & Sons.

MAGNETIC SUSCEPTIBILITY ARTIFACT

Appearance

This artifact produces distortion of the image together with large signal voids. Figure 8.23 shows magnetic susceptibility artifact from a dental implant within the image volume.

Cause

Magnetic susceptibility is the ability of a substance to become magnetized. Different tissues magnetize to different degrees, which results in a difference in precessional frequency and phase of magnetic moments of nuclei within the tissue. This causes dephasing at the interface of these tissues and a loss of signal. The geometric accuracy of an MRI image relies on a homogenous




Figure 8.23 Magnetic susceptibility from a dental implant causing massive distortion of the image.

imaging volume and approximately linear gradient applications. The distortion caused by magnetic susceptibility artifact is due to the lines of magnetic flux becoming attracted to a ferromagnetic item and causing a corresponding nonlinearity of the gradients used in spatial encoding. In practice, the main causes of this artifact are metals within the imaging volume, although it can also be seen from naturally occurring hemorrhage or iron deposition, as they magnetize to a much greater degree than the surrounding tissue. Magnetic susceptibility artifact is more prominent in gradient-echo sequences, as the gradient reversal cannot compensate for the phase difference at the interface.

Remedy

This artifact can, under some circumstances, aid in diagnosis. Small hemorrhages are sometimes only seen because they produce a magnetic susceptibility effect. The use of fruit juices such as pineapple as negative “contrast agents” in the bowel also relies on susceptibility-related dephasing. However, in general, this artifact is undesirable and degrades the image. There are several remedies.

- *Remove all metal objects.* Always ensure that the patient has removed all metal objects before the scan. Always check whether the patient has any aneurysm clips or metal implants. Most implants can be scanned but may cause local heating effects (see Chapter 10).
- *Use spin-echo sequences instead of gradient-echo.* The 180° RF rephasing pulse used in spin-echo sequences is very efficient at compensating for phase differences caused by inhomogeneity, while gradient-echo sequences do not employ this pulse and are less effective. In




Figure 8.24 Sagittal gradient-echo images of the knee with pins in the tibia. Magnetic susceptibility has produced a large distortion of the image.

Figures 8.24 and 8.25, gradient-echo and spin-echo sequences are used to acquire images of the knee in which there are tibial screws. Metal artifact produces magnetic susceptibility artifact on both images but is significantly less in the spin-echo sequence. The same effect is also seen in SS-TSE. Additional 180° RF rephasing pulses in a long echo train increase rephasing and reduce this artifact.

- *Decrease the TE.* Longer TEs allow more dephasing between tissues with susceptibility differences; therefore, using a short TE reduces this artifact. This is achieved by selecting a wide receive bandwidth in the scan protocol (see Chapter 6) but decreases SNR (see Chapter 7).
- *Use a wide receive bandwidth when scanning a metal implant.* The distortion that occurs with susceptibility artifact is caused by the metal affecting the homogeneity of the field. In turn, this creates a range of anomalous frequencies around the implant that create distortion in the frequency-encoding direction. Just as in chemical shift artifact, this results in a change in frequency over a given number of pixels. Although the artifact causes a shift in precessional frequency, a wide receive bandwidth restricts this to a smaller number of pixels than a narrow bandwidth. This is because a wide bandwidth results in more frequencies distributed across each pixel. For example, if the shift caused by the artifact is 200 Hz, and a wide bandwidth corresponds to 125 Hz/pixel, the distortion only extends for two pixels in the frequency direction.




Figure 8.25 Sagittal spin-echo images of the same patient as shown in Figure 8.24. The artifact is reduced.

TRUNCATION ARTIFACT

Appearance

Truncation artifact produces a banding artifact at the interfaces of high and low signal (Figure 8.26). It creates a low-intensity band running through a high-intensity area.

Cause

This artifact results from undersampling of data (too few k -space lines are filled) so that interfaces of high and low signal are incorrectly represented on the image. It is most common when tissue is still producing a high signal at the end of data collection or when the peak of the echo is not centered in the middle of the sampling window. The latter is common when using a very short TE.




Figure 8.26 Truncation artifact.

Remedy

Undersampling of data should be avoided. One way to do this is to increase the number of phase-encoding steps. For example, use a 512×512 matrix instead of 512×128 . In addition, avoid using techniques that only partially fill k -space (e.g. partial Fourier; see Chapter 6). Fat suppression techniques used in T1-weighted imaging might reduce this artifact by nulling high signal from fat at the beginning and end of the sampling window. Alternatively, filters can be used that force the signal amplitude to zero at the end of the sampling window [9].

CROSS-EXCITATION/CROSS-TALK

Appearance

This artifact causes a reduction in SNR in adjacent slices in the slice stack (Figure 8.27).

Cause

Ideally, the profile of a slice should be square, or rather rectangular, when viewed from the edge. In practice, an RF excitation pulse is not able to achieve this. The width of the pulse should be half its amplitude, but this normally varies by up to 10%. Adjacent slices receive energy from the RF excitation pulse of their neighbors where the overlap occurs (Figure 8.28). This energy




Figure 8.27 Contrast changes between slices as a result of cross-excitation.




Figure 8.28 Cross-excitation.

pushes the magnetic moments of nuclei toward the transverse plane, so that they may become saturated when they are excited by their intended RF pulse. This effect is called **cross-excitation**, and the resulting saturation reduces SNR.

The literature sometimes discriminates between cross-excitation and cross-talk, stating that **cross-talk** occurs in multiangle imaging when slices physically intersect. The physical principles are, however, the same. The RF intended for one slice is absorbed by another. Where slices physically intersect, the artifactual appearance creates dark bands of signal loss corresponding to the position of the overlapping slices.

Remedy

275

Cross-excitation can be reduced by ensuring that there is at least a 30% gap between the slices. This is 30% of the slice thickness itself and reduces the likelihood of RF exciting adjacent slices.

Scan tip: Avoiding cross-excitation in practice

One way of avoiding this artefact is to select a large slice gap. For example, if the slice thickness selected is 5 mm, use a skip or gap of 2 mm (40% of 5 mm) rather than a 1 mm gap (20% of 5 mm). In addition, most systems excite alternate slices during the acquisition so that there is some time for cross-excitation in adjacent slices to decay before it is their turn to be excited. For example, the excitation order of slices is 1, 3, 5, 7, 2, 4, 6, 8. Slices 1–7 have time to decay their cross-excitation, while slices 2–8 are excited (approximately half the TR). A process known as **interleaving** or **concatenation** extends this time even further. When interleaving slices, alternate slices are excited and divided into two acquisitions. Cross-excitation created in adjacent slices has the time of a whole acquisition to decay before it is its turn to be excited. For example, the excitation order of slices is 1, 3, 5, 7 in the first acquisition and 2, 4, 6, 8 in the second. Slices 1–7 have the time of a whole acquisition (several minutes) to decay while slices 2–8 are excited. When using interleaving, no gap is required between the slices.

Some systems use software to square off the RF pulses so that the adjacent nuclei are less likely to become excited. This reduces cross-excitation, and, therefore, much smaller gaps between slices are possible. However, there may be some loss of signal as a proportion of the RF pulse is lost in the squaring-off process. It is still wise to use a small gap of 10% when employing this software. If appropriate, a volume acquisition offers the advantages of thin slices (partitions) with no gap required. Cross-excitation in multislice/multiangle can be avoided by ensuring the angle between slice blocks is not too acute. If this is not feasible (for example, due to the angulation of the intervertebral discs), the localizer may be split in multiple acquisitions, one for each block.

ZIPPER ARTIFACT

Appearance

Zipper artifact appears as a dense broken line across the image perpendicular to the frequency-encoding direction (Figure 8.29).




Figure 8.29 *Zipper artifact.* In both examples shown, there is a line of high signal running perpendicular to the frequency-encoding direction. This represents external interference at a discrete frequency.

Cause

This is caused by extraneous RF entering the room at a frequency that matches the frequencies expected in the echo. Frequencies that are close to the Larmor frequency can be generated by broadcast radio, by wireless computer networks, or from “unintentional emitters” (see Chapter 9). Unintentional emitters are devices having worn electric motors that emit extraneous interference. Common examples include air-conditioning fans, floor-polishing machines, and rotary dimmer switches. The frequency of the interference tends to be of a much higher amplitude compared to the spin-echo and therefore appears in the image as a high-intensity line representing that particular frequency after FFT. This contamination of signal is usually caused by operating the scanner with the magnet room door open or by a breach in the RF cage.

Remedy

Always close the magnet room door during data acquisition. Call the engineer to locate any breaches in the RF shielding and repair them.

SHADING ARTIFACT

Appearance

Shading produces a difference in signal intensity across the imaging volume.

Cause

The main cause of shading artifact is uneven excitation of nuclei within the patient due to RF excitation pulses applied at flip angles other than 90° and 180° . Shading is also caused by abnormal loading on the coil or by coupling of the coil at one point. This may occur with a large patient who touches one side of the body coil and couples it at that point. Shading is also caused by inhomogeneities in the main magnetic field, which is improved by shimming (see Chapter 9).

Remedy

Always ensure that the coil is loaded correctly, i.e. the correct size of coil is used for the anatomy under examination and that the patient is not touching the coil at any point. The use of foam pads or water bags between the coil and the patient usually suffices. In addition, ensure that appropriate prescan parameters are obtained, as these determine the correct excitation frequency and amplitude of the applied RF pulses.

277

MOIRÉ ARTIFACT

Appearance

The term **moiré** pattern refers to a pattern generated when there is interference between two other patterns. You can sometimes see this effect when looking through two sheets of net curtain. There are shifting stripes generated where the weave of the netting either superimposes or does not. The same appearance is used in fabrics such as moiré silk. There are two image aberrations that can be described as moiré artifacts because they result in a stripy appearance. The first is caused by multiple noise spikes, the second one by field inhomogeneity.

Multiple noise spikes. The receiver coil and the components that transfer the analog signal to the analog-to-digital-converter can be sensitive to interference. Occasionally, a single burst of noise from, for example, a static discharge in the magnet room can cause a spike in the data at just one single point in k -space. After FFT, this is reconstructed in the image as a grid of parallel lines. The lines may be close together (high spatial frequency area of k -space) or farther apart (low spatial frequency area of k -space). This depends on whether the spike occurs after a steep or shallow phase-encoding step. The angle of the lines also depends on whether the spike occurs during the rephasing or dephasing part of the spin-echo. In the event that two noise spikes occur during the same acquisition, the two sets of lines may interfere on the image leading to various appearances that all look like types of woven fabric, ranging from silk to canvas depending on the spacing of the lines (Figure 8.30). As noise spikes are usually caused by static discharges in the magnet room, to remedy this artifact, ensure that the humidity of the room is within normal limits. Noise spikes may also be caused by arcing if the plug on a receive coil cable is not firmly sited in its socket.

Field inhomogeneity. There is also a type of moiré effect seen in gradient-echo sequences that is caused by a combination of aliasing and field inhomogeneity. Aliasing occurs because signal-producing anatomy exists outside the FOV. Field inhomogeneity causes the superimposed signal to be in phase and out of phase with the signal in the voxels onto which it is wrapped. This creates a black-and-white banding appearance around the edges of the FOV (Figure 8.31). This artifact often resembles zebra stripes because the edge of the homogenous imaging volume is typically

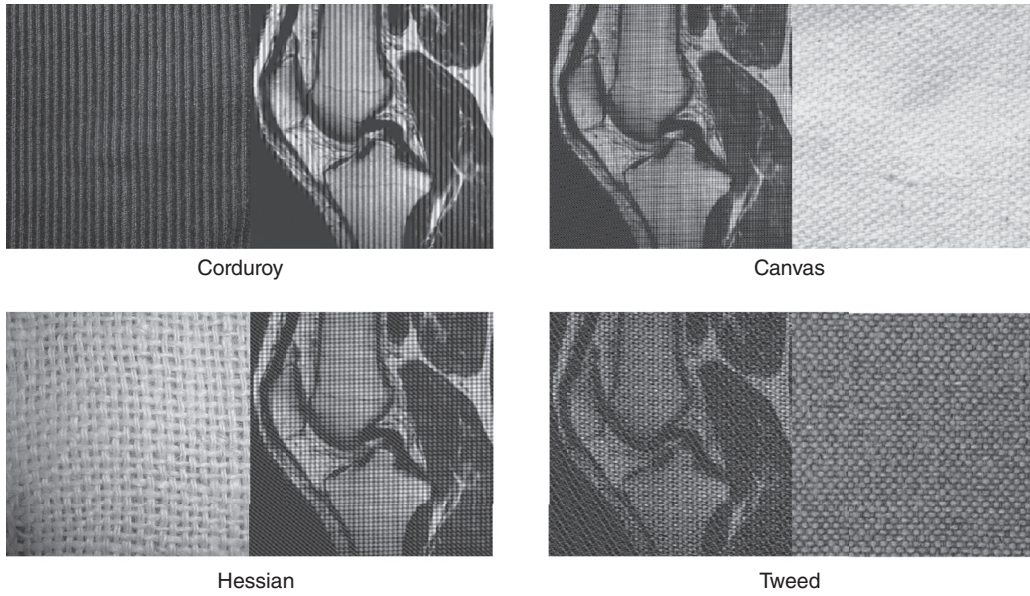


Figure 8.30 Various appearances of multiple noise spikes.



Figure 8.31 Moiré artifact seen as zebra lines on the edge of the FOV.

distorted because a segmented solenoid is used to generate B_0 . This gives the edge of the imaging volume a somewhat spiky profile. To remedy this artifact, use spin-echo sequences rather than gradient-echo, and use antialiasing remedies (see previous section).

MAGIC ANGLE

Appearance

Magic angle artifact produces abnormally high signal intensity in tissues that contain collagen (such as tendons). In Figure 8.32, this is seen in the patellar tendon and may mimic pathology.

279

Cause

Although collagen is rich in hydrogen, it normally exhibits a very low signal on MRI images. This is because water molecules bound to the collagen fibers experience very rapid T2 decay. Magic




Figure 8.32 Magic angle artifact shown with a high signal intensity at the lower border of the patellar tendon.

Table 8.4 Things to remember – other artifacts.

Magnetic susceptibility is caused by different tissues being magnetized differently. It is reduced by using spin-echo sequences. Sometimes it is a good artifact in that it increases the conspicuity of hemorrhage when using gradient-echo sequences

Truncation is caused by undersampling of phase data. It is remedied by filling more of k -space

Cross-excitation is caused by the nonrectangular profile of RF pulses and their Fourier transforms. One of the ways to remedy this is to introduce a gap or skip between slices

Artifacts and their remedies are summarized in Table 8.8

angle artifact occurs when collagen structures lie at an angle of 55° to the main field. The anisotropic shape of the molecules in collagen causes reduction to zero of spin–spin interactions. The T2 decay time increases in collagen structures when they lie at this angle to B_0 . This causes increased signal intensity in the structure when a short TE is used.

Remedy

Alter the angle of the patient's anatomy relative to the B_0 field or increase the TE. If the signal remains hyperintense, it is likely to be fluid (ligament tear) rather than magic angle artifact.

EQUIPMENT FAULTS

There are some other artifacts caused by major equipment malfunction. The loss of a gradient, for example, causes image distortion, and eddy currents induced in the gradient coils can cause phase artifacts as they create additional unwanted phase shifts. In addition, data-acquisition errors cause a variety of different artifacts, most of which are characterized by a geometric appearance in the image, such as well-defined bands of missing signal. These may disappear if the scan is repeated but, if persistent, may require an engineer to replace image-processing hardware such as array processor.

FLOW ARTIFACTS

Flow artifacts may result in significantly different degrees of signal from the lumen of blood vessels and CSF. There are some negative aspects to this phenomenon, for example when pulsatile high signal from arterial flow produces phase mismapping, or saturation in slow-moving flow gives a misleading impression of vessel size. There are, however, some benefits to flow artifacts. For example, if the pulse sequence parameters are manipulated to saturate background signal from anatomy while maintaining high signal in vessels, angiographic images may be acquired. To fully understand the benefits and trade-offs of flow phenomena, it is necessary to examine the mechanism behind the artifactual appearances seen on the image. There are two main flow phenomena described in the scientific literature that have a significant effect on signal intensity: **entry-slice phenomenon (ESP)** and **time-of-flight (TOF) phenomenon**. ESP tends to create **flow-related enhancement** of signal, whereby the lumen of a vessel appears hyperintense on the image. Time of flight tends to cause **high-velocity signal loss** whereby the lumen of the vessel appears less hyperintense on the image as the velocity of flow increases.

Entry-slice phenomenon (ESP)

ESP is related to the excitation history of nuclei. Stationary nuclei that receive repeated RF excitation pulses during an acquisition with a short TR are saturated because their magnetic moments are more likely to be orientated in the spin-down direction (see Chapter 2). This is because the TR is not long enough for full longitudinal recovery of magnetization in the tissues in which the nuclei reside. Inflowing nuclei do not receive these repeated RF excitation pulses because until they enter the slice, they have not been at the appropriate resonant frequency to be affected by the slice-selective RF excitation pulse. These inflowing spins are therefore unsaturated as their magnetic moments are mainly orientated in the spin-up direction. This means that their NMV can be flipped into the transverse plane by the RF excitation pulse that they experience on first entering the slice. As a result, the inflowing blood returns a much higher signal than the saturated stationary background spins. This is called flow-related enhancement (Figure 8.33). This effect is seen when nuclei flow is perpendicular to the slice and is called ESP (or inflow effect) as it is most prominent in the first slice of a stack. The effect decreases as the spins move further into the imaging volume because they receive multiple RF excitation pulses as they pass through each slice, especially if




Figure 8.33 Inflow and flow-related enhancement.

they are traveling in the same direction as slice excitation. This causes flowing spins to become saturated, in some cases more so than the background tissue. The signal intensity of flowing nuclei in these middle slices depends on the TR, flip angle, slice (volume) thickness, flow velocity, and the direction of flow compared to the direction of slice excitation.

- *TR*. The TR is the time between each RF excitation pulse. A short TR increases the rate at which RF is delivered. A short TR decreases the time between successive RF excitation pulses. A short TR may cause inflowing spins to receive multiple RF pulses as they pass through the slice. This causes a degree of saturation and therefore reduces the ESP.
- *Flip angle*. For a given TR, the higher the flip angle, the more saturation occurs. Higher flip angles therefore reduce the ESP.
- *Slice thickness*. Flowing nuclei with a constant velocity take longer to travel through thick slices than through thin slices. Nuclei traveling through thick slices are likely to receive more RF excitation pulses than nuclei traveling through thin slices. ESP is therefore more pronounced in thin slices than in thick slices.
- *Velocity of flow*. The velocity of flow also affects the number of RF excitation pulses received by flowing spins as they traverse a slice. Fast-flowing nuclei only take a short amount of time to traverse a slice or volume and therefore only receive one or two RF excitation pulses. This is not sufficient to cause significant saturation. ESP therefore becomes more pronounced when there is high velocity flow.
- *Direction of flow*. The direction of flow is probably the most important factor in determining the degree to which ESP penetrates into a stack of slices. Nuclei that are flowing in the same direction as slice excitation (sometimes termed **co-current flow**) are more likely to receive repeated RF excitation pulses as they move from one slice to the next. They therefore become saturated relatively quickly, so ESP decreases rapidly. When flow is in the opposite direction to slice excitation (sometimes termed **counter-current flow**), flowing nuclei stay relatively unsaturated, as, when they enter a slice, they are less likely to have received previous RF excitation pulses. High signal from the vessel lumen therefore penetrates more deeply into the stack of slices (Figure 8.34).




Figure 8.34 Entry-slice phenomenon and direction of flow.

Scan tip: ESP in practice

Look at Figures 8.35–8.38. These are four axial slices of the abdomen prescribed and excited from the most inferior position to the most superior position, i.e. Figure 8.35 is slice 1, Figure 8.36 is slice 2, Figure 8.37 is slice 3, and Figure 8.38 is slice 4 in the stack of slices. Slice 1 is acquired first and slice 4 last in the acquisition. Look at the signal intensity of the aorta and inferior vena cava (IVC) in these images. Although they both contain blood and should have the same signal intensity in all slices, this is not the case. In slice 1, the IVC has a high signal intensity and the aorta low signal intensity. In slice 4, the contrast is opposite, i.e. the IVC is hypointense, and the aorta is hyperintense. In addition, the IVC has a lower signal intensity in slice 4 than the aorta has in slice 1.

These appearances are due to entry-slice phenomena. In slice 1, nuclei in the IVC are unsaturated because they have traveled from an anatomical region that is outside of the stack of slices. Excitation can only occur when they reach a point along the slice-select gradient that matches the frequency of the transmitted RF excitation pulse. Therefore, in slice 1, these nuclei receive their first RF excitation pulse and produce a high signal as their magnetic moments are mainly in the spin-up direction and may, therefore, be flipped into the transverse plane. Nuclei in the aorta, however, are saturated and produce a low signal because they are excited by RF excitation pulses as they travel down through the stack of slices during the acquisition. Their magnetic moments are primarily orientated in the spin-down direction.

In slice 4, the effect is opposite to that in slice 1. Nuclei in the aorta are unsaturated as they travel into the stack from above, having received no previous RF excitation pulses. Therefore, in slice 4, these nuclei receive their first RF pulse and produce a high signal as their magnetic moments are mainly orientated in the spin-up direction and may be flipped into the transverse plane. Nuclei in the IVC, however, are saturated by repeated RF excitation pulses as they travel through the stack of slices during the acquisition, and their magnetic moments are primarily orientated in the spin-down direction. In slices 2 and 3, the inflow effect decreases as nuclei in both vessels receive RF excitation pulses. The IVC has a lower signal intensity in slice 4 than the aorta has in slice 1 because flow in the IVC is in the same direction as slice excitation while flow in the aorta is in the opposite direction. Therefore, nuclei in the IVC receive more RF excitation pulses than those in the aorta. This effect is rarely seen in clinical imaging because flow compensation techniques such as spatial presaturation eliminate it. This is discussed later.




Figure 8.35 Entry-slice phenomenon: slice 1 (most inferior).



Figure 8.36 Entry-slice phenomenon: slice 2 (middle inferior).



Figure 8.37 Entry-slice phenomenon: slice 3 (middle superior).




Figure 8.38 Entry-slice phenomenon: slice 4 (most superior).



Refer to animation 8.2 on the supporting companion website for this book:
www.wiley.com/go/westbrook/mriinpractice

Time-of-flight (TOF) phenomenon

Nuclei must receive both RF excitation and rephasing pulses to produce signal. Stationary nuclei always receive both pulses, but flowing nuclei present in the slice for the RF excitation pulse may have exited the slice before the RF rephasing pulse is applied. Alternatively, they may receive the RF rephasing pulse but were not present in the slice for the RF excitation pulse. In both cases, no signal is received from these nuclei. This is called TOF phenomenon (Figure 8.39). Its effects depend on the type of pulse sequence.

Table 8.5 Things to remember – entry-slice phenomenon.

- Entry-slice phenomenon increases
 - at the first slice in the stack
 - when using a long TR in thin slices
 - with fast flow
 - when flow is in the opposite direction to slice excitation




Figure 8.39 Time-of-flight phenomenon.

TOF phenomenon in spin-echo pulse sequences. In a spin-echo pulse sequence, a 90° RF excitation pulse and a 180° RF rephasing pulse are applied to each slice. Every slice is therefore selectively excited and rephased. Stationary nuclei within the slice receive both the 90° and the 180° RF pulses and produce signal. Nuclei flowing perpendicular to the slice may be present within the slice during the 90° RF excitation pulse but may have exited the slice before the 180° RF rephasing pulse is delivered. These nuclei are excited but not rephased and do not therefore produce signal. Alternatively, nuclei not present in the slice during the RF excitation pulse may be present during rephasing. These nuclei have not previously been excited and do not therefore

produce signal. TOF phenomena result in a signal void from flowing nuclei, and vessels are hypointense. TOF effects depend on the following (Equations (8.5) and (8.6)):

Equation 8.5

$$SI \propto 1 - v(1/2TE)/SI_t$$

SI is the signal intensity

v is the velocity of flow (mm/s)

TE is the echo time (ms)

SI_t is the slice thickness (mm)

This equation shows how the time-of-flight effect depends on TE, slice thickness, and velocity of flow

287

Equation 8.6

$$v = SI_t/TR$$

v is the velocity of flow (cm/s)

SI_t is the slice thickness (mm)

TR is the repetition time (ms)

This equation calculates the velocity of flow required to replace all the saturated spins in a slice with fresh unsaturated spins

- *Velocity of flow.* As the velocity of flow increases, a smaller proportion of flowing nuclei are present in the slice for both 90° and the 180° RF pulses. As the velocity of flow increases, the TOF phenomenon increases. This is called high-velocity signal loss. As the velocity of flow decreases, a higher proportion of flowing nuclei are present in the slice for both 90° and 180° RF pulses. Therefore, as the velocity of flow decreases, the TOF phenomenon decreases.
- *TE.* As the TE increases, a higher proportion of flowing nuclei exit the slice between the 90° RF excitation pulse and the 180° RF rephasing pulse. Therefore, with a long TE, more nuclei receive only one pulse, and the signal decreases (Figure 8.40).
- *Slice thickness.* For a given constant velocity, nuclei take longer to travel through a thick slice than through a thin slice. Therefore, nuclei are more likely to receive both 90° and 180° RF pulses in thick slices. As the thickness of the slice decreases, nuclei are more likely to receive only one pulse, and the signal decreases.



Refer to animation 8.3 on the supporting companion website for this book: www.wiley.com/go/westbrook/mriinpractice

TOF phenomenon in gradient-echo pulse sequences. In gradient-echo pulse sequences, an RF excitation pulse with a variable flip angle is followed by gradient rephasing (see Chapter 4). Because a 180° RF rephasing pulse is not used, the TE is typically much shorter than that used in a spin-echo sequence. This means that most of the spins do not exit the slice in which they were excited. In addition, RF excitation is slice-selective, but gradient rephasing affects the entire imaging volume. Therefore, a flowing nucleus that receives an RF excitation pulse is likely to be rephased regardless of its location and more likely to produce signal. Gradient-echo pulse sequences are therefore often said to be flow-sensitive.

In addition to ESP and TOF, there are other factors that affect the appearance of flowing blood. These are determined by the nature of flow inside a vessel.




Figure 8.40 Time of flight vs TE.

Table 8.6 Things to remember – time-of-flight flow phenomenon.

Time-of-flight phenomenon produces flow-related enhancement or high-velocity signal loss

- Flow-related enhancement increases as the
 - velocity of flow decreases
 - TE decreases
 - slice thickness increases
- High-velocity signal void increases as the
 - velocity of flow increases
 - TE increases
 - slice thickness decreases

There are four principal types of flow (Figure 8.41):

- *Laminar flow (also known as parabolic flow)* is flow that is at different but consistent velocities across a vessel. The flow at the center of the lumen of the vessel is faster than at the vessel wall where resistance slows down the flow. However, the velocity difference across the vessel is constant (Equation (8.7)). Laminar flow causes two problems. Firstly, the magnetic

Equation 8.7

$$Re = d\rho v/\eta$$

Re is the Reynolds number
 d is the density (g/cm³)
 v is the velocity of flow (cm/s)
 m is the diameter of the vessel (cm)
 η is the viscosity of blood (g/cm s)

If Re is less than 2100, then flow, is laminar. If Re is more than 2100 then flow is turbulent




Figure 8.41 The different types of flow.

moments of nuclei moving along a gradient either gain or lose phase according to the direction of the gradient and the distance moved along it. If the flow is at different velocities across the lumen of the vessel, the spins in the center move further along the gradient than those at the edge. This may lead to intravoxel dephasing and loss of signal. Secondly, relatively slow-flowing blood near the vessel wall may experience more RF pulses as it travels through the slice (or volume) and saturate more quickly than the higher velocity flow at the center of the vessel. This may cause vessels to look narrower than their true size on the reconstructed image. Diseases such as sickle-cell anemia may exacerbate this effect, as the blood viscosity and vascular patency are affected.

- *Spiral flow (helical flow)* is where the velocities (described earlier) coil around each other in a corkscrew formation. This is typically seen in large vessels such as the thoracic aorta. The multidirectional nature of the flow and the anatomy may necessitate the use of phase contrast angiography (see later).
- *Vortex flow* is, by definition, a whirlpool. Blood having passed through a stricture or flowing into an aneurysm may begin to exhibit this continuously circling type of flow. The issue with vortex flow is that it may cause spins to remain in a particular location for sufficient time to become saturated by multiple RF pulses and is also likely to cause a degree of intravoxel dephasing.
- *Turbulent flow* is where flow velocity difference changes erratically across the vessel. The difficulty in imaging this type of flow also relates to phase dispersion within a voxel.

Learning tip: Intravoxel dephasing

Gradients alter the magnetic field strength, precessional frequency, and phase of magnetic moments of hydrogen nuclei (see Chapter 5). The precessional frequency of magnetic moments of nuclei flowing along a gradient rapidly increases or decreases depending on the direction of flow and gradient application. The magnetic moments of flowing nuclei therefore either gain phase (if they are accelerated) or lose phase (if they are decelerated) (refer to the watch analogy in Chapter 1). Across the lumen of a vessel, there are different velocities of flow because blood is viscous and experiences friction from the vessel wall. Faster-flowing spins in the center of the vessel travel further along an applied gradient, causing their magnetic moments to gain or lose phase compared to those spins having a lower velocity. Therefore, magnetic moments of nuclei within the same voxel are out of phase with each other, decreasing signal amplitude from the voxel. This is called **intravoxel dephasing** (Figure 8.42). The magnitude of intravoxel dephasing also depends on the degree of turbulence. In turbulent flow, intravoxel dephasing effects are irreversible. In laminar flow, intravoxel dephasing can be compensated for using a technique known as gradient moment rephasing. This is most effective when the velocity and direction of flow are constant. This type of flow is sometimes defined as **first-order motion**. Accelerating flow (second-order motion) and jerky flow (third-order motion) are not compensated by using gradient moment rephasing. Intravoxel dephasing can also be reduced by the use of a short TE (allowing less time for dephasing to occur) and by using smaller voxels. The use of smaller voxels to increase signal is somewhat counterintuitive but, as shown in Figure 8.42, a large voxel is more likely to incorporate a wider number of spin isochromats. If these have opposing phase positions, it results in cancellation of signal. A smaller voxel is more likely to contain spins having broadly similar phase positions.




Figure 8.42 Intravoxel dephasing.

Flow artifact remedies

Signal returned from a blood vessel can be artificially manipulated, and, as such, the appearance may be defined as an artifact. However, from an image quality viewpoint, there is no particular disadvantage to visualizing high or low signal from a vessel. There is, however, an issue when anatomy returning high signal moves between phase-encoding steps because this results in phase mismatching artifact as previously described. This can be compensated for by applying a presaturation pulse outside of the imaging volume to saturate inflowing spins. There may, however, be occasions when a high signal is required from the vasculature, such as inflow angiographic studies. In these cases, flow compensation gradients called **gradient moment rephasing** (also known as **gradient moment nulling** or **flow compensation**) can be applied.

Gradient moment rephasing. This technique compensates for the altered phase values of magnetic moments of nuclei flowing along a gradient. It uses additional gradients to correct the altered phases back to their original values and follows the same principles as the balanced gradient system used in balanced gradient-echo sequences (see Chapter 4). The magnetic moments of flowing nuclei do not gain or lose phase due to the presence of the main gradient. The phase coherence of these magnetic moments is maintained, and their nuclei produce signal. Gradient moment rephasing is performed by the slice-select gradient and/or the frequency-encoding gradient. Gradient polarity changes from positive to double negative and then back to positive again. A flowing nucleus traveling along this gradient experiences different magnetic field strengths, and the phase of its magnetic moment changes accordingly. This is shown in Figure 8.43 where the magnetic moment of a flowing nucleus gains 90° of phase as it passes along the first positive lobe of the gradient and then loses 180° of phase as it passes through the double negative lobe of the gradient. The net phase change of the magnetic moment at this stage is that it has lost 90° of phase. As the flowing nucleus passes through the last positive lobe of the gradient, this is corrected so that the net phase change of its magnetic moment




Figure 8.43 Gradient moment rephasing (nulling).




Figure 8.44 Without GMN.

is zero. Gradient moment rephasing predominantly reduces intravoxel dephasing. As the phase shifts of the magnetic moments of flowing nuclei are corrected, flow motion artifacts decrease (Figures 8.44 and 8.45). Gradient moment rephasing assumes a constant velocity and direction along the gradient. It is most effective on slow laminar flow and is therefore often termed **first-order motion compensation**. Pulsatile flow is not strictly constant, so gradient moment rephasing is often more effective on venous rather than arterial flow. It is also less effective on turbulent fast flow perpendicular to the slice. As gradient moment rephasing uses extra gradients, it causes an increase in the minimum TE. If the system performs extra gradient tasks, more time must elapse before it is ready to read an echo. As a result, fewer slices may be available for a given TR, or the TR and therefore the scan time may automatically increase to acquire the selected slices.



Refer to animation 8.4 on the supporting companion website for this book:
www.wiley.com/go/westbrook/mriinpractice

Spatial presaturation

Spatial presaturation pulses nullify signal from flowing nuclei so that the effects of entry-slice and TOF phenomena are minimized. It works on a similar principle to saturation pulses that specifically null signal from certain tissues to improve CNR (see Chapter 7). Spatial presaturation delivers a 90° RF pulse to a volume of tissue outside the FOV. A flowing nucleus within the volume receives this pulse. When it then enters the slice stack, it receives the 90° RF excitation pulse and is saturated. If it is fully saturated, it has no transverse component of magnetization and produces no signal (Figure 8.46).




Figure 8.45 With GMN.




Figure 8.46 Spatial presaturation.

Scan tip: Using presaturation pulses in practice

Spatial presaturation pulses should be placed between the origin of flow and the imaging volume so that signal is nulled from flowing nuclei that enter the FOV. In sagittal and axial imaging, presaturation pulses are usually placed above and below the FOV so that arterial flow from above and venous flow from below are saturated. Right and left presaturation pulses are sometimes useful in coronal imaging (especially in the chest) to saturate flow from the subclavian vessels [10]. Spatial presaturation pulses can also be brought into the FOV itself and placed over artifact-producing areas (such as the aorta). The presaturation pulse saturates spins and reduces phase mismatching artifact that originates from these structures.

Presaturation pulses increase the amount of RF delivered to the patient that may increase heating effects and decrease the number of slices available (see Chapter 10). They are only useful if they are applied to tissue so they need to be positioned correctly and only used when they are likely to be effective. Presaturation pulses are applied around each slice just before the RF excitation pulse. The TR and the number of slices therefore govern the interval between the delivery of each presaturation pulse. To optimize presaturation, use all the slices permitted for a given TR. As presaturation nulls signal, it is usually used in T1- and PD-weighted images (Figures 8.47 and 8.48).




Figure 8.47 Without presaturation.




Figure 8.48 With presaturation.

Table 8.7 Things to remember – flow compensation techniques.

Even-echo rephasing uses balanced echoes in which even echoes demonstrate less dephasing than odd echoes. It reduces intravoxel dephasing and is mainly used in T2-weighted sequences

Gradient moment rephasing uses additional gradients to correct altered phase values. It reduces artifact from intravoxel dephasing and is most effective on slow, laminar flow within the slice

Presaturation uses additional RF pulses to nullify signal from flowing nuclei. It reduces artifact due to time-of-flight and entry-slice phenomena. It is effective on fast and slow flow, and increases the RF deposition to the patient

Table 8.8 Artifacts and their remedies.

Artifacts	Axis	Remedy	Penalty
Flow motion	Phase	Swap phase and frequency axes	May need antialiasing
		Gating/triggering	Variable TR
			Variable image contrast
			Increased scan time
		Presaturation	May lose a slice
Gradient moment rephasing	Increases minimum TE		

Artifacts	Axis	Remedy	Penalty
Chemical shift	Frequency	Increase receive bandwidth	Decreases minimum TE available
		Reduce FOV	Decreases SNR
			Increases resolution
		Use chemical saturation	Reduces SNR
			May lose slices
Out-of-phase signal cancellation	Frequency and phase	Select a TE at periodicity of fat and water	May lose a slice if TE is significantly increased
Aliasing	Frequency and phase	No frequency wrap	None
		Antialiasing/no-phase wrap	May increase scan time depending on manufacturer
			May increase motion artifact depending on manufacturer
		Enlarge FOV(p)	May increase scan time depending on manufacturer
Presaturation bands	Increase in SAR		
Zipper	Frequency	Call engineer	Irate engineer!
Magnetic susceptibility	Frequency and phase	Use SE, TSE/FSE, or SS-TSE or metal artifact reduction techniques	Blood product may be missed
		Remove metal	None
Shading	Frequency and phase	Check shim	None
		Load coil correctly	None
Patient motion	Phase	Use antispasmodics	Invasive
		Immobilize patient	None
		Counseling of patient	None
		All remedies for flow motion	See above
		Sedation	Possible side effects
Invasive			
Costly			
Cross-excitation	Slice select	Interleaving	Doubles the scan time
		Squaring-off RF pulses	Reduces SNR
Moiré	Frequency and phase	Use SE, TSE/FSE	None
Magic angle	Frequency and phase	Change TE slightly	None
		Alter position of anatomy	None

FLOW-DEPENDENT (NON-CONTRAST-ENHANCED) ANGIOGRAPHY

Although artifactual in nature, flow phenomena may be used to increase the diagnostic value of an examination by using **magnetic resonance angiography (MRA)**. This may utilize “bright-blood” imaging to demonstrate high-signal flow within the vasculature, or “black-blood” imaging in the cases of arterial dissection to distinguish between low-signal intraluminal flow and stationary blood trapped in the intima (lining) of an artery. The main benefit of flow-dependent techniques is that no exogenous contrast medium is required. This makes the procedure very fast and, more importantly, very safe and noninvasive for the patient. Traditionally, flow-dependent angiograms were only suitable for small regions of interest, such as the circle of Willis or the carotid bifurcation, because imaging larger structures can be very time-consuming. Scan time is protracted because contrast on inflow studies relies on the ESP, for which flow must ideally be perpendicular to the vessel. Covering the entire abdominal aorta would, therefore, require a large number of slices, and because these are typically acquired sequentially, the scan time is unacceptably long. For this reason, gadolinium contrast-enhanced magnetic resonance angiograms (CEMRA) became the technique of choice for large vessels. Evidence that suggests gadolinium is the primary agent in the cause of nephrogenic systemic fibrosis, and more recent concerns involving accumulation of gadolinium in patients with no previous renal impairment, has resulted in a reappraisal of noncontrast techniques. Flow-dependent angiography can be broadly divided into the following main categories.

Inflow MRA

The aim of an inflow angiogram is to produce an image with high contrast between high-signal vascular flow and saturated, low-signal stationary background anatomy (Figure 8.49). The technique

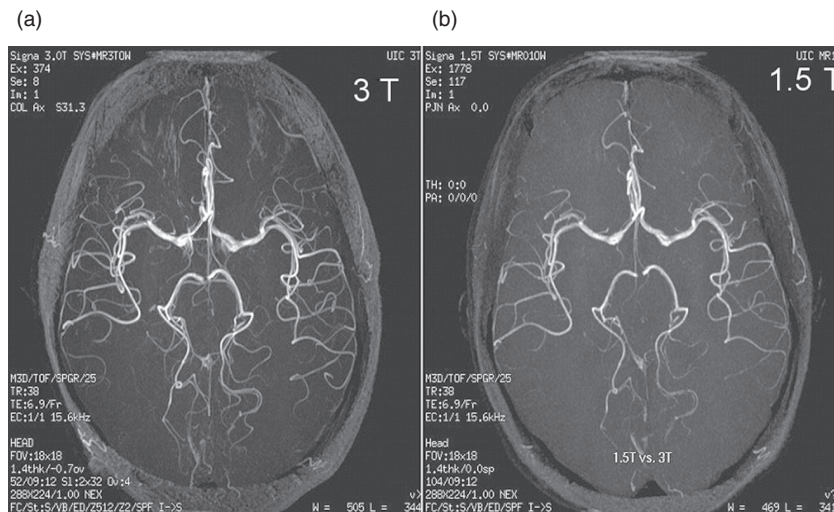


Figure 8.49 Axial 3D inflow MRA images of the brain to evaluate vasculature in the circle of Willis. These images were acquired at 3 T (a) and 1.5 T (b). Note the improvement in vascular contrast due to greater SNR and CNR in the 3 T image.




Figure 8.50 *In-plane flow.* Tortuous vessels or vessels in which the flow is along the plane of the slice may not be demonstrated on a MIP. The spins become saturated by numerous excitation pulses.

relies on the ESP to create high signal in the vessel lumen. However, the TOF effect must also be taken into consideration as long TEs otherwise cause the RF-excited bolus to leave the slice before the echo is sampled. This causes a signal void in fast-flowing arterial blood. Signal from the stationary background tissue is saturated by the use of a short TR.

Inflow angiograms are acquired using a flow-compensated gradient-echo sequence and may be performed as a sequential 2D acquisition or a 3D volumetric acquisition. Flow compensation by gradient moment rephasing is required to prevent phase mismapping of the high-signal pulsatile flow. A saturation band may also be employed outside of the imaging volume to selectively saturate arterial or venous flow if the vessels are not required to be visualized on the images. In 2D sequential acquisitions, the saturation band follows the position of each slice as it is acquired (sometimes called a **traveling SAT band** or travel-sat). Critically, inflow angiograms are only suitable for imaging blood that flows perpendicular to the slice. In-plane flow spends a longer amount of time in the slice and receives multiple RF pulses causing saturation and a loss of signal (Figure 8.50).

2D sequential inflow angiograms

The benefit of using a 2D sequential acquisition is that each slice is acquired separately. This means, in theory, the FOV and slice coverage are only limited by the size of the receive coil. In addition, there is less risk of spins becoming partially saturated by the RF excitation pulses of other slices, as is often the case in a 2D volumetric acquisition. This effect is further reduced by acquiring individual slices in the opposite direction to the flow. Another advantage offered by a 2D sequential acquisition is that even slow-flowing spins experience a short transit through a thin slice. This means that they are only likely to receive one or two RF excitation pulses and are not saturated inside the slice. As a result, signal is maximized throughout the data set. The ability to image slow flow means that both arteriography and venography can be performed using this technique. There are also some disadvantages, however. SNR tends to be poorer

Table 8.9 Advantages and disadvantages of inflow MRA.

Advantages	Disadvantages
<i>2D inflow MRA</i>	
Large area of coverage	Lower resolution
Sensitive to slow flow	Saturation of in-plane flow
Sensitive to T1 effects	Venetian blind artifact
<i>3D inflow MRA</i>	
High resolution for small vessels	Saturation of in-plane flow
Sensitive to T1 effects	Small area of coverage

Table 8.10 Overcoming the advantages of inflow MRA.

Susceptibility artifacts	Use short TEs and small voxel volumes
Poor background suppression	Use a TE that acquires data when fat and water are out of phase
	Implement magnetization transfer techniques
Venetian blind artifacts	Use breath-hold techniques
Limited coverage (3D)	Acquire images in another plane
	Use MOTSA (multiple overlapping thin-section angiography)
Suppression of in-plane signal	Use ramped RF pulses
	Administer contrast media
Pulsation artifacts	Time acquisition to the cardiac cycle

than the 3D volumetric option, and the slice thickness is likely to be greater than a 3D partition (Tables 8.9 and 8.10).

This is important because the resulting image undergoes reconstruction using an algorithm (known as **maximum intensity projection** or MIP) that allows the anatomy to be viewed from any angle. Thick slices when viewed from the edge may give a stepped appearance on the reconstructed image. As the slices are acquired sequentially, there is a potential for very conspicuous image anomalies if the patient shifts position during the procedure. Rather than the typical ghosting artifact, the reconstructed image appears to demonstrate gaps, breaks, and noncontinuity of the vasculature structures. Like any heavily T1-weighted sequence, the base data may continue to return signal from fat or any tissue that has a short T1 recovery, such as methemoglobin. Fat can be eliminated from the MIP by ensuring that only the vessels are included in the region to be reconstructed. Early subacute hemorrhage may be difficult to distinguish from the vessels in instances such as a brain hemorrhage and may obscure the area of interest. In some cases, it may mimic the appearance of an arteriovenous malformation.

Learning tip: MIP

Vascular contrast on flow-dependent angiograms is created by motion. Strictly speaking, the resulting images do not show the true anatomy of the vessels so much as the ability of blood to move through them. Of course, this could be said of any angiogram; however, a CEMRA shows the morphology of the vessels in their entirety as the contrast passes through them. Inflow MRA, on the other hand, does not. The base-data consist of slices that are acquired at 90° to the direction of flow and only reveal the vessel lumen in cross-section. This is not a particularly useful way to view vascular structures. To create a more anatomically correct image, a computer algorithm known as the MIP is used. This technique does exactly as its name suggests. Each pixel in a slice is given a numerical value according to its intensity. The background is largely saturated and therefore has low pixel-intensity values. The inflowing spins inside the vessel lumen have very high values. Figure 8.51 shows how these maximum intensity values can be projected onto two compiled images representing an anteroposterior view and a lateral view of the vasculature. Each row and each column within a slice have a maximum intensity pixel at a particular location. This maximum intensity is projected onto a pixel in the compiled image. In the example shown, this has been from two projections at 90° to each other; however, in practice it is usual to include projections at (for example) every 5° of rotation between 0° and 90°. The resulting set of compiled images gives a more realistic impression of the vascular structures and may be presented as a rotating cine loop. There is little value in compiling images over a greater range than 90° because normal MIP images are orthographic in nature. This means that there are no depth cues or perspectives to the compiled images. Reconstructions that are rotated more than 90° to each other therefore contain the same diagnostic information but appear as “mirror-images.”




Figure 8.51 MIP reformatting. The maximum intensity projection ascertains the maximum intensity in each row or column of pixels and assigns this value to a pixel in a projected plane. In this diagram, there are two such planes representing an anterior and lateral projection of the data.



Figure 8.52 Image reconstructions from MOTSA acquisition (multiple overlapping thin slabs). Signal saturation has occurred at the edge of a slab resulting in an apparent (and highly suspicious) bilateral “stenosis.”

3D volumetric inflow angiograms

The benefits of using a 3D volumetric acquisition are that the partitions are typically thinner than 2D slices. When used with a high matrix, this permits the use of small isotropic (cubic) voxels. Thin partitions and cubic voxels are beneficial because they ensure that the MIP reconstructed image is of high spatial resolution and that the voxel size is the same from every angle of rotation. Small voxels also reduce the amount of intravoxel dephasing and improve SNR. The main trade-off when using a 3D volume over 2D sequential acquisition is that spins spend longer in transit through a thick slab compared to a thin slice. They therefore experience the RF excitation pulse for long enough to become saturated. This is unlikely to cause problems with small regions of interest such as the circle of Willis, but makes a larger coverage impossible without multiple overlapping thin slab acquisitions (**MOTSA**). Although MOTSA increase the examination time, this technique offers the benefits of 3D volumetric acquisition coupled with the coverage of a 2D acquisition. It is important to overlap the slabs by a certain percentage to avoid a striped “**venetian blind**” effect on the compiled image. This occurs because partially saturated anatomy at the top of each slab tends to exhibit a lower SNR than the entry-slice (partition) of the adjacent slab. In addition, this can cause the appearance of stenosis in vessels that coincidentally run along the direction of the uppermost partitions in a slab. Another factor when choosing between 2D and 3D inflow sequences is that slow flow also saturates before signal has penetrated into a 3D volume, making this technique unsuited to venography (Figure 8.52).

Scan tip: Ramped RF

Poor signal penetration into the slab may be offset by the use of spatially varying RF pulses, also known as **ramped RF**. This works by reducing the flip angle at the edge of the slab where the inflowing spins enter the volume and increasing the flip angle at the other edge of the slab. This creates a linear range of flip angles throughout the slab in the direction of flow and improves signal penetration by increasing the contrast-to-noise ratio between flowing and stationary spins in the distal partitions of the slab. Any flow in the opposite direction to the ramped RF saturates more quickly, which is an advantage when the intention is to reduce the signal from venous return.

Table 8.11 Things to remember – inflow MRA.

Inflow MRA is a technique that produces images where unsaturated spins coming into the slice produce a higher signal intensity than the stationary spins within the slice

Saturation pulses are also used to null signal from unwanted flow (venous)

3D inflow-MRA is useful in high-velocity flow regions. 2D inflow MRA is useful in slower flow regions

BLACK-BLOOD IMAGING

In contrast to inflow studies, the objective of black-blood imaging sequences is to completely nullify the signal from flowing blood, leaving the vessel lumen black. The simplest way to achieve this is to use a T1-weighted image and prevent ESP by saturating inflowing spins with a saturation slab outside of the slice stack (Figure 8.53). Reducing signal from

303




Figure 8.53 Spatial presaturation to produce black blood. Note that the magnetic moments of nuclei in the vessel (at the top of the illustration) are aligned with the magnetic field (B_0) along the z-axis. As the blood within the vessel flows down into the saturation volume, they receive a 90° RF pulse and their vectors enter the transverse plane. As blood continues to flow down into the slice, nuclei receive another 90° RF pulse and are aligned 180° from their original position at the top of the diagram. At this point (and with no time to recover), nuclei with the blood are saturated and are dark on the image.




Figure 8.54 Black-blood imaging. Normal flow may return a similar signal intensity to the vessel lumen. In the event of an arterial dissection, stationary blood trapped in the intima may also exhibit a similar intensity. By saturating inflowing blood, the signal is removed resulting in a more accurate and conspicuous representation of vessel patency.

inflowing spins may also be achieved using inversion prepulses and applying an RF excitation pulse at the null point of blood. Black-blood images are typically viewed individually rather than using projection algorithms, although a minimum intensity projection is theoretically possible. The main advantage of removing signal from the lumen of a vessel is that the tunicae (vessel walls) become more clearly delineated, revealing the extent of arterial plaque or trapped blood in arterial dissection. In cardiac imaging, the ventricles and atria are similarly well distinguished. The conspicuity of vessels, however, is reduced compared to bright-blood techniques, particularly in anatomical areas having low-inherent background signal such as the base of the skull or the lungs. Figure 8.54 demonstrates how black-blood imaging increases the conspicuity of an arterial dissection. In bilaterally symmetrical vessels such as the vertebral arteries, this is even more apparent, as a comparison can be made between the two sides.

PHASE-CONTRAST MRA

Phase-contrast angiography takes advantage of the fact that the magnetic moments of spins moving along a bipolar gradient acquire a phase shift, whereas those of stationary spins do not (Figure 8.55). The degree and direction of the shift can be manipulated by the use of **velocity-encoding** (VENC) gradients built into the pulse sequence. Stationary spins experience the positive and negative lobes of the gradient in equal measure, and, as a result, their magnetic moments are restored to their original phase position. The magnetic moments of moving spins become either phase-advanced or phase-retarded according to their flow velocity and direction.

Importantly, the VENC is a user-definable parameter that must be set prior to the acquisition. The amplitude of the gradient is modified to suit the speed of flow in centimeters per






Figure 8.55 PCA gradient.

(cm/s) (Figure 8.56). Venous flow is likely to be around 15 cm/s, but arterial flow may exceed 200 cm/s in vessels such as the thoracic aorta. The importance of setting this factor correctly is that the acquired phase shift must not exceed 180° , otherwise aliasing may occur (Figure 8.57) or blood may appear to be flowing in the opposite direction to that being measured. This is because a phase shift of $+185^\circ$ cannot be distinguished from a phase shift of -175° , which would be expected from flow in the opposite direction. Having acquired phase-shifted data in three orthogonal directions, a subtraction is performed using a flow-compensated mask (see gradient moment rephasing later in this chapter). As a result, only the moving spins are visualized on the reconstructed image. Because subtraction is used, high signal from early subacute hemorrhage and fat is eliminated from the image. This is one of the main advantages of the phase-contrast technique. It also requires the patient to keep very still and may require physiological motion correction to prevent breathing and bowel peristalsis from misregistering in the subtracted image. The scan time in phase-contrast studies may be considerably longer than in inflow angiograms because at least four separate acquisitions must be made in order to quantify flow in three orthogonal planes and also create the mask.

Phase-contrast angiograms offer a unique advantage over other flow-dependent methods in that they can provide two different types of information when reconstructed as magnitude



Low VENC



High VENC




Figure 8.56 VENC.




Figure 8.57 VENC aliasing.

images or phase images. **Magnitude images** offer the normal anatomical-looking appearance with high-signal vessels against a black background. **Phase images**, on the other hand, demonstrate the direction of flow as being white (when in the same direction as the VENC) or black (when in the opposite direction to the VENC) (Table 8.12).

Table 8.12 Advantages and disadvantages of phase-contrast MRA (PC-MRA).

Advantages	Disadvantages
Sensitive to a variety of vascular velocities	Long imaging times with 3D
Sensitive to flow within the FOV	More sensitive to turbulence
Reduced intravoxel dephasing	
Increased background suppression	
Magnitude and phase images	

Table 8.13 Things to remember – phase-contrast MRA.

PC-MRA uses gradients to sensitize the sequence to flow. Flowing spins have a higher signal than stationary spins
The amplitude of the sensitizing gradients is controlled by the VENC. If the VENC is too low, aliasing occurs. If the VENC is high, vessel wall delineation may be compromised
3D PC-MRA produces images with a better SNR and spatial resolution than 2D but the scan times are long

ECG-triggered 3D FSE (flow-spoiled fresh blood imaging)

The main disadvantage of traditional inflow methods is that slices must be acquired sequentially at 90° to the vessel of interest, prolonging acquisition time. This technique typically uses an ECG-gated, 3D half-Fourier, TSE pulse sequence using read-out spoiler gradients to acquire a 3D volume acquisition. This permits a large FOV as required for imaging the great vessels, but within an acceptable scan time (usually under 4 min). The technique derives contrast by exploiting the fact that the T2 decay time of blood is longer than that of the surrounding background tissues. The background signal may also be digitally subtracted from the image, or may be reduced by the application of an inversion pulse as in short-tau inversion recovery (STIR). The technique may be described as flow-dependent, but only in the context of distinguishing between slow-flowing (venous) flow and high-velocity (arterial) flow. To achieve this, EKG data are acquired during the procedure, and two sets of images are acquired at systole and diastole (Figure 8.58).




Figure 8.58 EKG-triggered subtraction imaging. This technique exploits the T2 contrast difference between blood (a fluid) and the background tissue. Data acquisition is synchronized to the EKG permitting the visualization of arteries or veins. The mechanism is related to the speed of flow.

Signal from arterial flow varies greatly according to its velocity. Unlike venous flow, arterial flow is pulsatile. At systole, the flow has a high velocity, and there is a corresponding loss of signal due to both TOF effect and phase shift along a gradient. At diastole, arterial flow has a comparatively low velocity between heartbeats. Venous flow has a low velocity at both systole, and diastole because flow is not pulsatile. The aim of the technique is to acquire images at both systole and diastole and then digitally subtract one from the other. Depending on the type of system used, it may be possible to acquire both systolic and diastolic images in the same acquisition. This is advantageous, as it reduces the risk of patient movement between acquisitions and is therefore more likely to provide subtraction without misregistration. The systolic data only demonstrates slow venous flow and serve as a venogram if required. The diastolic data demonstrate both arteries and veins, but the systolic data may be used as a mask to digitally subtract the signal from veins to provide an arteriogram. The use of subtraction also offers the advantage of eliminating any signal from the stationary background tissues. A STIR prepulse may still be employed to improve the conspicuity of the vessels when viewing the nonsubtracted images.

In summary, this technique offers many of the advantages of CEMRA and CT angiography but without the need for exogenous contrast media or ionizing radiation. On a cautionary note, remember that the appearance of vasculature on the arterial image is actually a flow-difference image and as such is susceptible to some flow-related artifacts as described earlier. For example, slow arterial flow that is sometimes found proximal to a stenosis may still be visible on the systolic image and, as a result, may not subtract out. In the cases where narrowing is apparent on the image, evaluation of the unsubtracted images helps to rule out any overestimation of the stenosis. A final point to note is that although a single acquisition is approximately 4 min in duration, if peripheral imaging is required, this increases the overall examination time considerably.

MRA summary

In summary, flow-dependent MRA capitalizes on the artifactual appearances caused by ESP, the TOF phenomenon, and the phase shifts seen in the magnetic moments of spins moving along a gradient. Pulse sequence parameters such as the TR, TI, and TE may be adjusted to maximize the effects of these phenomena to provide high contrast between flowing blood and stationary background tissue. The contrast is usually manipulated to represent high-signal flow against a low-signal background. The flow-related nature of the technique is more analogous to a Doppler study than a conventional angiogram; however, postprocessing algorithms are able to present morphological-looking images that are capable of demonstrating stenoses, aneurysms, and other vascular pathology without the necessity for exogenous contrast media.

In the preceding chapters, we learned how images are acquired and how their contrast and quality are manipulated and modified. In the next chapter, the MRI system and its various components are discussed.



For questions and answers on this topic please visit the supporting companion website for this book: www.wiley.com/go/westbrook/mriinpractice

References

1. Hashemi, R.H., Bradley Jr, W.G., and Lisanti, C.J. (2010). *MRI: The Basics*, 3, 198. Philadelphia, PA: Lippincott Williams and Wilkins.
2. Hashemi, R.H., Bradley Jr, W.G., and Lisanti, C.J. (2010). *MRI: The Basics*, 3, 199. Philadelphia, PA: Lippincott Williams and Wilkins.

3. Westbrook, C. (2014). *Handbook of MRI Technique*, 4, 336. Wiley Blackwell.
4. McRobbie, D.W., Moore, E.A., Graves, M.J. et al. (2017). *From Picture to Proton*, 82. Cambridge: Cambridge University Press.
5. Hashemi, R.H., Bradley Jr, W.G., and Lisanti, C.J. (2010). *MRI: The Basics*, 3, 186. Philadelphia, PA: Lippincott Williams and Wilkins.
6. Dale, B.M., Brown, M.A., and Semelka, R.C. (2015). *MRI: Basic Principles and Applications*, 5, 107. Wiley.
7. Odaibo, S.G. (2012). *Quantum Mechanics and the MRI Machine*, 84. Arlington, VA: Symmetry Seed Books.
8. McRobbie, D.W., Moore, E.A., Graves, M.J. et al. (2017). *From Picture to Proton*, 88. Cambridge: Cambridge University Press.
9. Dale, B.M., Brown, M.A., and Semelka, R.C. (2015). *MRI: Basic Principles and Applications*, 5, 112. Wiley.
10. Westbrook, C. (2014). *Handbook of MRI Technique*, 4, 218. Wiley Blackwell.

9

Instrumentation

Introduction	311	Shim system	328
Magnetism	313	Gradient system	330
Scanner configurations	315	RF system	337
Magnet system	318	Patient transport system	343
Magnet shielding	326	Computer system and graphical user interface	344

After reading this chapter, you will be able to:

- *Differentiate between the different types of magnetism.*
- *Understand the differences in MRI scanner design and form-factor.*
- *Explain the function of the technical components found inside an MRI scanner.*

INTRODUCTION

This chapter critically evaluates equipment design in MRI and examines the components that make up a modern MRI system. As stated in previous chapters, MRI scanning requires a homogeneous, powerful magnetic field and a system to transmit and receive pulses of electromagnetic radiation in the RF spectrum (see Chapters 1 and 5). In addition, spatial encoding requires sophisticated manipulation of the static field in three orthogonal planes (see Chapter 5). All MRI scanners must therefore incorporate the following:

- A powerful magnet to create magnetic field over a 40–50 cm spherical volume
- A shim system to improve the homogeneity of the magnetic field
- A gradient system to create linear slopes in field strength in any direction
- An RF transmission system to generate and transmit pulses of electromagnetic radiation

- A set of RF receiver coils to detect signal from the patient
- A computer system to allow input of parameters and displaying images
- A computer subsystem capable of coordinating the application of RF pulses and gradients and reconstructing the acquired data into images and storing them.

The configuration of the scanner is largely dictated by the clinical or research requirements that it is required to address. Figure 9.1 shows a schematic diagram of the main components of a closed-bore MRI system. The components are arranged in layers, forming concentric cylinders around the magnet bore. Each of these components is evaluated in this chapter, but first let's examine the meaning of magnetism.

312

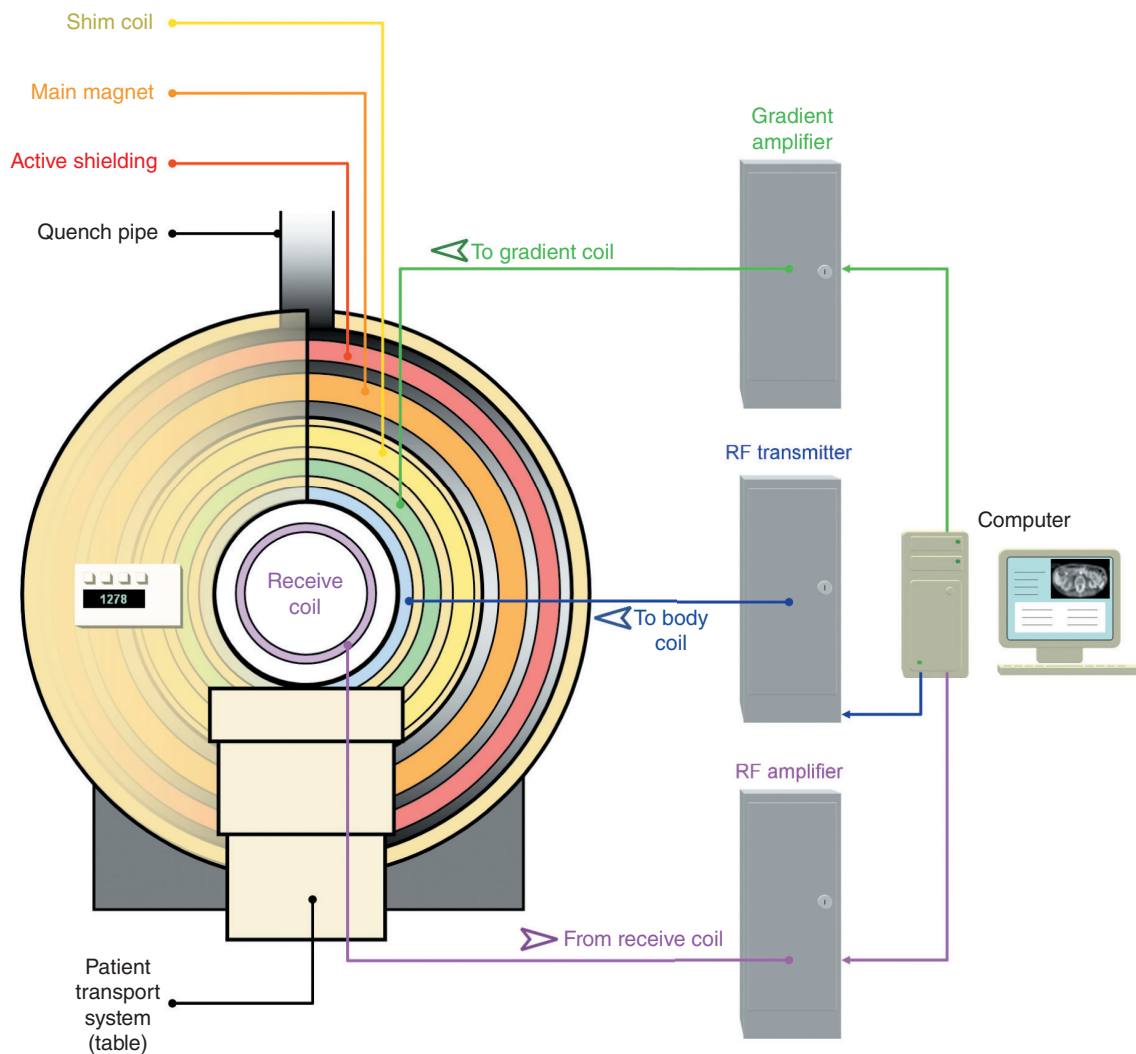


Figure 9.1 Closed-bore MRI scanner in axial cross-section revealing the principal components to be arranged in concentric circles, most of them being cylindrical electromagnets.

MAGNETISM

Magnetism is the second most powerful fundamental force of nature [1]. However, it is difficult to determine when humans were first aware of this phenomenon. The prehistoric iron age occurred over 4000 years ago, but there are naturally occurring magnetic minerals such as magnetite that are likely to have been known to earlier civilizations. The writings of Aristotle suggest that Thales of Miletus (600 BCE) was one of the first Greek philosophers to examine ferromagnetism from a scientific viewpoint. An understanding of electromagnetism is first attributed to Hans Christian Ørsted in 1820. Ørsted accidentally discovered that a compass-needle aligned to an electrical conductor in his laboratory, anticipating Michael Faraday's work on electromagnetic induction in the 1830s. In the modern world, we know that many chemical elements exhibit magnetic properties, and these are categorized according to a phenomenon known as magnetic susceptibility (see Table 9.1). Magnetic susceptibility, as the name suggests, refers to how susceptible (responsive) a material is to an applied magnetic field, to what extent it is magnetized, and whether it is attracted or repelled by the external field. Differences occur due to the atomic or molecular structure of the material in question, specifically the number of electrons surrounding the atomic nucleus and how they move and spin (Equation (9.1)). The types of magnetism are defined in the following sections [2]:

313

Equation 9.1

$$B_o = H_o (1 + x)$$

B_o is the magnetic field in teslas (T)

H_o is the magnetic intensity in amperes/m

This equation shows the apparent magnetization of an atom. A substance is diamagnetic when $x < 0$. A substance is paramagnetic when $x > 0$

Diamagnetism

Diamagnetic compounds are characterized by the fact that they exhibit a weak repulsion to an external magnetic field. This is known as having a small negative magnetic susceptibility. Diamagnetic elements have atoms in which all electrons are evenly paired. As a fast-spinning negatively charged particle, a single unpaired electron induces a powerful magnetic moment. However, when the spins are paired, their magnetic fields cancel each other. This is because the electrons spin in opposite directions. As a result, a diamagnetic material does not retain any net magnetism when removed from the external field and does not possess any magnetic moment of its own. When placed in an external field, the lines of magnetic flux diverge around a diamagnetic material as shown in Figure 9.2. A total of 31 elements are identified as being diamagnetic in the periodic table, including hydrogen and helium, and quite a few metals such as gold, silver, and lead.

Paramagnetism

Paramagnetic compounds are characterized by the fact that they exhibit a weak attraction to an external magnetic field. This is known as having a small positive magnetic susceptibility. Paramagnetic elements increase the strength of an external magnetic field into which they are introduced. The effect is due to the presence of unpaired electrons, which, as stated above, generates a net magnetic moment. However, on removal from an external field, the electron paths lose alignment; the paramagnetic material does not retain any net magnetism, and does not possess any magnetic moment of its own. When placed in an external field, the lines of magnetic




Figure 9.2 Effect of a diamagnetic substance on a homogenous magnetic field.




Figure 9.3 Effect of a paramagnetic substance on a homogenous magnetic field.

flux converge toward a paramagnetic object as shown in Figure 9.3. The periodic table lists 29 paramagnetic elements including calcium, oxygen, and many metals including aluminum, titanium, and platinum.

Ferromagnetism

Ferromagnetic compounds are said to have a large positive magnetic susceptibility and are powerfully attracted to an external magnetic field. This is a safety concern in MRI from a projectile hazard perspective (see Chapter 10). Magnetic domains (whereby the atomic magnetic moments are aligned parallel by an external magnetic field) cause ferromagnetic elements to retain their magnetic moments when removed from an external field. As can be seen in Figure 9.4, the flux lines of an external magnetic field are powerfully distorted by a ferromagnetic object, and this causes geometric distortion of images in MRI patients with ferromagnetic implants. There are just four naturally occurring elements that are ferromagnetic at a normal ambient temperature; iron, nickel, cobalt, and gadolinium.




Figure 9.4 Ferromagnetic substance in a homogenous magnetic field.

Table 9.1 Things to remember – magnetism.

Paramagnetic substances add to (increase) the applied magnetic field
Super-paramagnetic substances have a magnetic susceptibility that is greater than that of paramagnetic substances but less than that of ferromagnetic materials
Diamagnetic substances slightly oppose (decrease) the applied magnetic field
Diamagnetic effects appear in all substances. However, in materials that possess both diamagnetic and paramagnetic properties, the positive paramagnetic effect is greater than the negative diamagnetic effect, and so the substance appears paramagnetic
Ferromagnetic substances are strongly attracted to, and align with the applied magnetic field. They are permanently magnetized even when the applied field is removed
Moving a conductor through a magnetic field induces an electrical charge in it
Moving electrical charge in a conductor induces a magnetic field around it

Learning tip:**How can gadolinium be ferromagnetic, yet be classed as a paramagnetic contrast medium?**

In the field of MRI contrast agents, gadolinium is always considered to be paramagnetic, but in its refined state, gadolinium is a silver metal that has unpaired electrons and magnetic domains, which are the characteristics of ferromagnetism. Ferromagnetic elements have what is known as a “Curie temperature,” above which they cease to exhibit ferromagnetic properties. In the case of gadolinium, this is 20 °C. Below this temperature, gadolinium exhibits a magnetic moment of its own, which does not rely on the presence of an external magnetic field. At human body temperature (or any temperature greater than 20 °C), gadolinium is paramagnetic, hence its description as such when used as a contrast agent. Where a compound contains atoms of different magnetic susceptibility, the net magnetic susceptibility is dictated by the number of each type of atom in the compound and their electron configuration. For example, in water, oxygen is paramagnetic; however, water has a surfeit of hydrogen atoms and therefore exhibits a net diamagnetic effect. This is one of the factors that cause an adverse effect on field homogeneity when a patient is placed into an MRI scanner (see next section).

SCANNER CONFIGURATIONS

There are currently three main types of scanner configuration in clinical use [3]:

- Closed-bore systems
- Open systems
- Extremity systems.

Closed-bore systems

Closed-bore systems are the most popular type of MRI scanner worldwide. They feature the familiar tunnel-shaped magnet bore and resemble, in shape, a larger version of a computed tomography (CT) scanner. Longitudinal table movement allows the patient to be positioned with the region of interest lying at the center of the magnet bore. This encloses the patient to the

front, back, and sides, but still allows limited access. Patients having lower extremity scans may be positioned feet first, which permits most of the body to remain outside of the bore. Closed-bore scanners generate the main magnetic field using toroidal superconducting **solenoid electromagnets** positioned in circumference to the cylindrical bore. This type of scanner can generate very high magnetic field strengths, typically between 1 and 3 T for clinical use and up to 8 T (and above) for research studies. The highest field currently generated by this type of magnet for live-animal research is 21 T.

Open systems

316

Open systems have a different design, whereby the patient is positioned on a wider imaging table that is maneuvered between two magnetic poles that are located above and below the imaging volume. This only encloses the patient above and below, leaving a relatively unobstructed view from all sides. This is advantageous when scanning large animals, humans having a large habitus (broad or obese), and nervous/claustraphobic patients (such as small children), who may find the open access more tolerable. The design also facilitates easy side access to the patient by clinicians when undertaking interventions such as biopsies. Importantly, these scanners also permit a degree of sideways table movement. This is very useful when imaging lateral body structures such as the shoulder or elbow, as it allows the region of interest to be positioned closer to the isocenter of the magnet rather than at the edge of the imaging volume where there may be poorer field homogeneity. Flexion and extension views of the spine are also possible as patients have the space to adopt positions that are not possible in the confines of a closed-bore scanner. At least one manufacturer offers upright open MRI systems that permit weight-bearing examinations. Open scanners use large **permanent magnets** or superconducting solenoids to generate the main magnetic field. The maximum currently available field strength for an open superconducting MRI system is 1.2 T.



Refer to animation 9.1 on the supporting companion website for this book: www.wiley.com/go/westbrook/mriinpractice

Extremity systems

Extremity scanners, as the name suggests, are designed to scan limbs and are smaller in size than their whole-body counterparts. The typical design is approximately the size and shape of a domestic washing machine having a narrow aperture in the center that is large enough to accommodate an arm or leg. Slightly larger models are the size of a fluoroscopy unit and may be angled to allow weight-bearing views of the spine, hips, and knees. The magnetic field is typically generated by permanent magnets and is therefore restricted to below 1 T. This has certain negative trade-offs in terms of scan time and image quality but offers advantages too. The small physical size of the scanner and reduced magnetic fringe field means that they can be located in small rooms and offices. They are also cheaper to purchase, and running costs may be lower, as permanent magnets do not require electrical power or liquid helium fills to maintain the magnetic field.

The following sections of this chapter examine the components of closed-bore superconducting MRI scanners, but be aware that many of the same principles apply to open systems, particularly those using solenoid electromagnets. The primary difference is that in open systems the solenoids are positioned horizontally above and below the patient (anteroposteriorly) rather than inferosuperiorly (see Figure 9.5).



Figure 9.5 Differences in solenoid configuration in (a) closed-bore and (b) open MRI scanners.

MAGNET SYSTEM

Creating the magnetic field required for high-quality anatomical imaging is a demanding task. There are six main requirements, each of which presents technological challenges:

- The field strength (flux density) must be high, typically between 1.0 and 8.0 T.
- The **fringe field** having a strength of 0.5 mT (5 G) or greater must not extend outside of safety Zones III and IV, and should ideally be contained within the magnet room (see Chapter 10).
- The field must be spatially homogenous to a very high degree.
- The homogeneity must extend over a large spherical imaging volume (40 cm) to accommodate the required anatomical FOVs.
- The field must be temporally stable. This means that the flux density must not vary over time (for example, if the environmental temperature fluctuates).
- The weight and bulk of the magnet must be kept at a level that does not pose any problems with installation in a normal imaging department.

Fortunately, modern MRI systems have features that address all these requirements. They are now covered in more detail [4].

Permanent magnets

Permanent magnet MRI scanners do not employ electromagnets. Instead, they are equipped with large discs of a ferromagnetic alloy such as neodymium, boron, and iron, or aluminum, nickel, and cobalt (**alnico**). Neodymium magnets are also known as rare-earth magnets (despite neodymium being neither “rare” nor “earth”) and are some of the most powerful permanent magnets. The ferromagnetic discs are known as **pole shoes** and are typically mounted on a yoke that positions them directly above and below the imaging volume (Figure 9.6). The magnetic field is created by the inherent ferromagnetism of the alloy, namely the combined forces of unpaired electrons in the atoms of the metal that create a macroscopic magnetic field. The advantages of this type of magnet are that it does not require electrical power or cryogenic cooling. These advantages are somewhat offset by the fact that these magnets are unable to generate a flux density of more than 0.5 T, are typically very heavy (17 US tons), and cannot be switched off in an emergency. Furthermore, the flux density of a permanent magnet is not stable and may change with the environmental temperature of the magnet room.

Resistive electromagnets

The phenomenon of electromagnetism was first discovered by Hans Christian Ørsted in 1820. Ørsted observed that a direct current flowing through a conductive wire caused a magnetic field to be induced around the conductor. Together with Michael Faraday’s law of electromagnetic induction, Ørsted’s law provides an explanation for the mechanism of electromagnets such as those used in MRI scanners. The direction of the induced lines of magnetic flux is visualized by using what is known as the “right-hand grip rule” as shown in Figure 9.7. This analogy supposes that a conductor, such as the length of a wire, is gripped in the right hand. The thumb indicates the conventional direction of current flow along the wire (+ to –). The direction of the fingers as they curl around the wire indicates the direction of the induced magnetic field. This model can also be adapted for solenoid magnets. In this case, the fingers represent the direction of the current as it flows through the windings of the solenoid, and the direction of the thumb indicates




Figure 9.6 Typical design of a permanent-magnet open scanner. The flux lines of the static field run vertically in this type of scanner.




Figure 9.7 Right-hand grip rule.

the direction of the induced magnetic field. If we assume that current flows around the solenoid windings in a clockwise direction, and we stand at the front end of a closed-bore MRI scanner, the flux lines run parallel to the bore, with the north pole of the magnet at the far end. The flux density of these magnets is determined by the number of windings in the solenoids and the magnitude of the current flowing through them (see Equation (1.6)).

Resistive MRI scanners employ copper-wound solenoids that operate just below normal room temperature. The principal advantage of this type of system is that the field strength can be adjusted and the magnet switched off safely after use. Industrial resistive magnets can achieve an ultra-high field strength; however, they typically feature very narrow magnet bores. To attain a maximum flux density of around 0.4 T, in a solenoid of a size required for human scanning, a current of over 10 kilowatts (kW) is required. Like an electric bar fire, the resistivity in the windings produces significant heat, and water cooling is required to prevent damage to the system (which would otherwise become incandescently hot). This is achieved by siting the solenoid magnets inside a water-filled vessel through which chilled water is constantly circulated. **Superconducting magnets** were introduced to avoid resistivity issues. These devices use **cryogenics** (coolants) to reduce the temperature of the windings to within 4° of absolute zero (4 kelvin (K)). This enables a substantially higher flux density using a solenoid large enough to fit a patient inside.

Superconducting electromagnets

Superconducting electromagnets create a magnetic field in the same way as a resistive magnet; however, the windings of the solenoid are spun from a type of metal alloy that is superconductive [5] (typically niobium/titanium). This means that the resistivity of the metal decreases to zero when the metal is cooled below a certain critical temperature (known as the transition temperature). To understand how this affects the design of the MRI scanner, it is first necessary to describe how the cooling system works.

Cryostat

The term cryostat is derived from the Greek words meaning “cold” and “stable.” The cryostat is a somewhat larger version of a thermal vacuum flask that you might use to keep your wine chilled. The cryostat contains the cryogen liquid helium, which has a boiling point of just 4.2 K (−268.9 °C). The primary function of the cryostat is to prevent heat transfer from the adjacent system components (particularly the gradient coils) to the cryogen. This thermal insulation reduces the rate at which the liquid helium boils off to the atmosphere.

The physical construction of a cryostat is shown in Figure 9.8. The outer structure consists of a hollow cylindrical steel tank. This is almost entirely seal-welded except for an aperture through which the helium is filled and through which also passes an atmospheric exhaust vent (quench pipe). The entire outer tank is evacuated of air, which largely reduces heat transfer by thermal convection. On top of the outer shell of the cryostat is a refrigeration unit that chills the metal superstructure of the cryostat, helping to prevent heat transfer by conduction. The area inside the cylinder of the cryostat is known as the **warm bore**. This contains not only the patient bore but also the components of the MRI system that operate at room temperature.

Inside this outer tank is a similarly shaped secondary cryogen chamber constructed from aluminum. In Figure 9.8, the cryogen chamber is shown half removed from the steel cryostat for clarity. The outside wall of the cryogen chamber is swathed in layers of aluminized polyester sheeting with insulating spacers. This highly reflective insulating material is familiar to anyone who has seen a space-blanket used to protect patients (or marathon runners) from hypothermia. The term “space-blanket” refers to the fact that the material was originally developed by NASA




Figure 9.8 Construction of an MRI cryostat.

(National Aeronautics Space Administration) to provide an insulation layer in space-suits. This was necessary to protect the astronauts from extremes in temperature during the moon landings. Its highly reflective surface forms a very efficient heat shield that prevents heat transfer by thermal radiation. This combination of features considerably reduces heat transfer, thereby the helium boil-off rate. Many modern scanners also feature a helium recondensing or recycling system that further reduces helium loss to a negligible amount [6]. Such scanners are unlikely to require a helium refill during their operational lifetime.

Cryogenics – liquid helium

Liquid helium is the cryogen of choice for superconducting magnets because of its extremely low temperature. Achieving a temperature lower than 4 K is very difficult outside a specialized plasma laboratory. Helium is readily available because in some geographical areas, it forms a (small) percentage of natural gas. However, as a finite and diminishing resource, it is important not to waste it. There are currently large extraction plants operating in several countries including the United States of America (USA) and Qatar, where there are sizable underground reserves of

fossil fuels. Due to the likely diminishing availability of helium and the inherent risks of quench, current research is exploring superconducting magnets that can operate at higher temperatures (20 K). These use alloys such as yttrium barium copper oxide rather than niobium/titanium [7].

Solenoid magnets

As previously discussed, the main magnetic field is induced using a solenoid magnet. In practice, several separate solenoid segments are employed. The reason for this is that a single, evenly wound, cylindrical solenoid of a suitable length for an MRI scanner is not capable of creating a large enough homogenous imaging volume to achieve the required FOV. To achieve a 40 cm spherical imaging volume, it is necessary to segment the solenoid into sections, each having a certain number of windings (also known as turns). Figure 9.9 shows a model representation of a segmented solenoid magnet of the type used in a closed-bore superconducting MRI scanner. Solenoid magnets are typically wound onto a former (or bobbin). In the model, there are two large circular solenoids at each end of the former with four smaller solenoids positioned along the length of the structure. The end solenoids are responsible for generating the bulk of the main magnetic field, while the four evenly spaced solenoids ensure that the homogenous imaging volume is large enough to cover the requisite 40 cm sphere.

322

Learning tip: Superconductivity

To achieve superconductivity, it is necessary to use an appropriate metal alloy for the magnet windings. Alloys such as niobium/titanium exhibit superconductivity below a certain critical temperature (10 K). The windings of a superconducting MRI solenoid are formed from strands of this alloy in a matrix of copper. Copper is normally considered an efficient conductor of electricity, but at a temperature of 4 K, electrons show a preference to pass through the niobium/titanium alloy. Copper thus becomes the insulator for the wire and protects the windings in the event of a quench by helping to spread the thermal load when resistivity returns to the solenoid. Below the critical temperature, the thermodynamic motion of niobium–titanium atoms slows to a degree that permits electrons to pass through the molecular lattice unimpeded. This means that the current continues to persist indefinitely, even when the external power supply is detached. Although on first thoughts this phenomenon seems like a type of perpetual motion, or “free” electricity, remember energy is still applied to the system to reduce cryogen loss. Without this, the cryogens would eventually boil off to the atmosphere, the windings would return to a resistive state at room temperature, and the current flow would be impeded.

Ramping a magnet

Bringing an MRI magnet up to the required field strength involves a simple parallel circuit featuring what is known as a “persistent switch.” The parallel winding is a superconductor with a heating coil positioned around the outside. When heated, the terminals of the solenoid are connected to an external power supply to power up the magnet. When the required field strength




Figure 9.9 Bobbin used to create a segmented solenoid electromagnet.

is achieved, the heater is turned off, and the persistent switch becomes superconductive. This creates a closed loop of superconductive wire and effectively bypasses the external power supply because the flowing electrons show a preference for the nonresistive circuit (Figure 9.10). The process of energizing an MRI solenoid is called **ramping**, reflecting the fact that the current is gradually increased during the process. To ramp down the magnet when decommissioning the system, the heater is reactivated on the persistent switch. This diverts the current from the solenoids through a resistor to dissipate the energy. If a large ferromagnetic object becomes lodged inside the MRI scanner, this technique is used to safely remove it without a quench of the system. However, this is only done in a non-emergency event where there is no danger to life or limb (see Chapter 10).

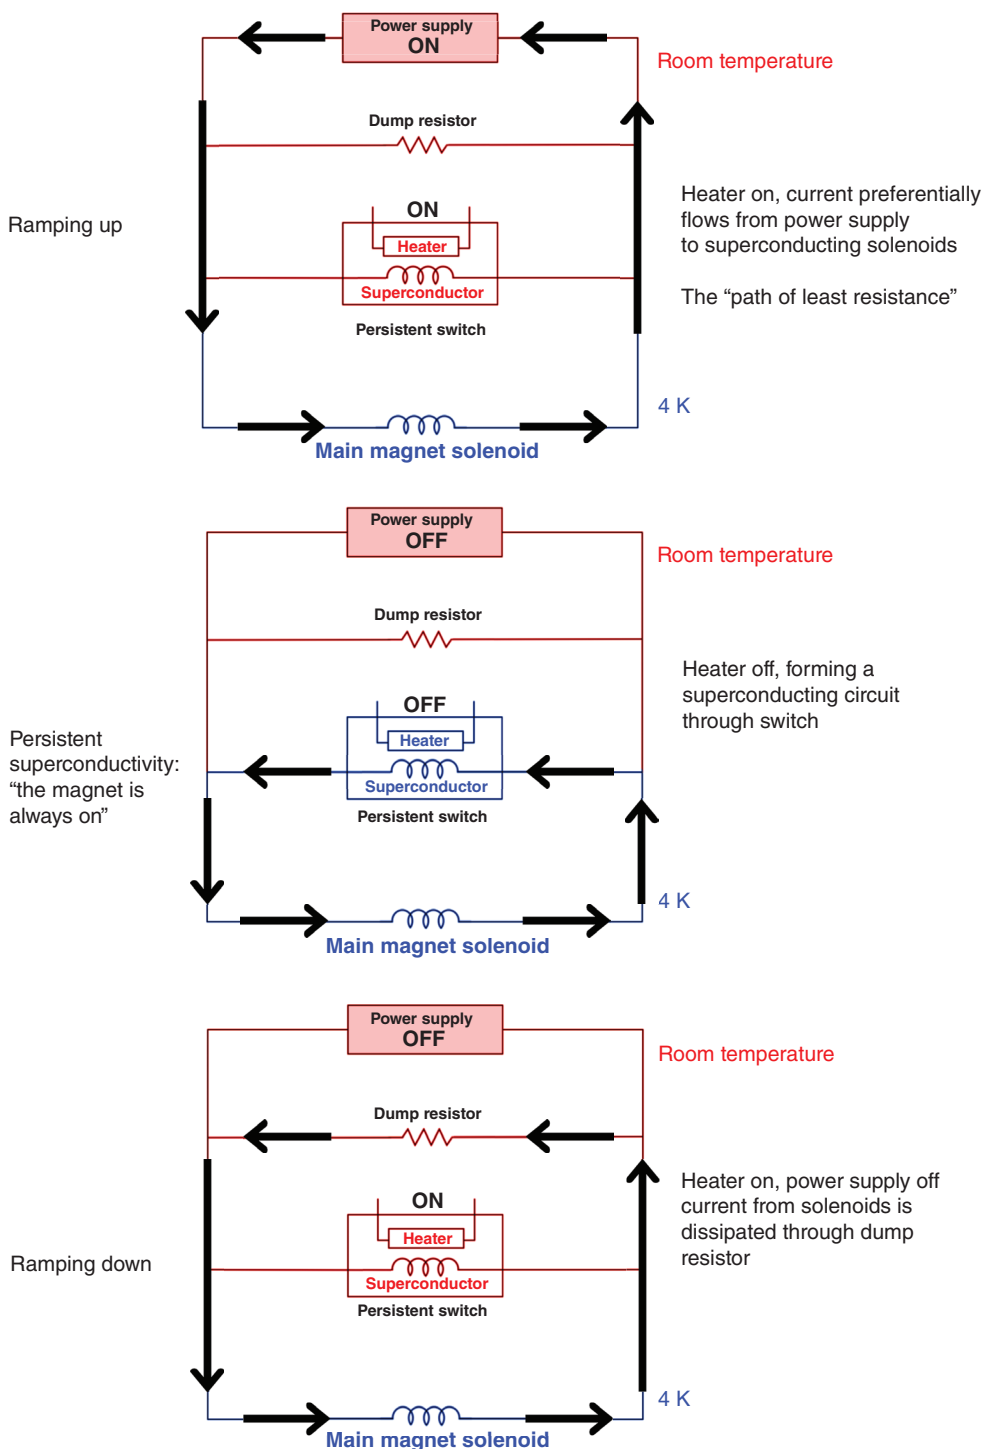


Figure 9.10 Parallel circuit and persistent switch used for ramping a magnet.

Field strength (flux density)

Magnetic fields are measured in two main ways. Strictly speaking, magnetic field strength is measured in amperes (A) per m and is given the symbol H . A magnetic field is also visualized as having lines of flux that are seen when iron-filings are sprinkled around a bar magnet. Flux density describes the number of flux lines passing through a given area. In the context of MRI, this is the spherical imaging volume at the isocenter of the bore. Flux density is affected by the magnetic permeability of the medium through which the lines pass. This is seen in passive shielding where the flux lines preferentially pass through a steel plate rather than air (see next section). Although the terms “field strength” and “flux density” are not strictly interchangeable, for the purposes of MRI, they can be considered to have a close relationship. This is because the permeability of the patient and the air in the magnet bore remain constant (see Equation (9.1)). Flux density is measured in teslas and is given the symbol B , and for this reason the main magnetic field of the scanner is known as B_0 . The field strength (flux density) of an MRI magnet varies according to the scanner design and typically ranges between 0.15 and 8 T. In SI units, the tesla is used to quantify flux density. In the older CGS system, the gauss is used (where 1 T equals 10 000 G – see Chapter 1). Although CGS units are generally being superseded by SI units, it is useful to keep the gauss unit, because it is sometimes considered a more meaningful way to measure lower magnetic field strengths. The earth’s magnetic field is usually around 0.5 G depending on one’s position in relation to the equator (this is also stated as 50 μ T or 0.000 05 T).

Learning tip: Field strength and SNR

The principal advantage of a high-field MRI magnet is an inherently higher SNR. Signal and noise both increase with field strength. Signal increases in proportion to the square of the external field strength, and noise increases in a linear relationship. This results in an approximately linear relationship when combined (see Chapter 7). Ignoring all other factors, doubling the field strength approximately doubles the SNR (see Chapter 1). This is certainly the case when scanning a water-filled phantom; however, it may not hold true when scanning a patient. At ultra-high field strengths, magnetic susceptibility effects are accentuated, which may result in more rapid dephasing of signal (see Chapter 8). In addition, chemical shift increases, necessitating an increase in receive bandwidth to reduce the appearance of the artifact on anatomical images. This has a negative trade-off with SNR (see Chapter 7).

A potential disadvantage of an ultra-high field strength is the corresponding increase in energy deposition to the patient. Much of the RF energy used to flip the NMV is also dissipated as heat in the patient’s body tissues, and this increases with the square of the field strength (it is the electric field component of an RF pulse that causes heating; see Chapter 1). The corresponding increase in specific absorption rate (SAR) may have consequences for the TR in spin-echo pulse sequences and the flip angle in gradient-echo pulse sequences (see Chapter 10).

These issues aside, the increase in SNR permits techniques such as functional MRI (fMRI) where signal changes in the brain cortex might otherwise be difficult to distinguish from background noise at a lower field strength. Blood-oxygen-level-dependent (BOLD) techniques such as fMRI rely on magnetic susceptibility, so the effect becomes more pronounced at high field

Table 9.2 Things to remember – superconductive magnet systems.

Superconducting magnets are the commonest type of magnet used in clinical MRI
The resistivity of the electromagnetic coils is zero because they are immersed in a cryogen that reduces their temperature to almost absolute zero
Current therefore flows, and the magnetic field is retained as long as the electromagnetic coils have no resistance
High magnetic strengths permit fast imaging and produce images with a high SNR. However, some artifacts (e.g. chemical shift and magnetic susceptibility artifact) are more evident at higher field strengths

(see Chapter 2). The increase in SNR also permits the use of smaller voxels, which, in turn, increases the spatial resolution of the image (see Chapter 7). Although an increase in chemical shift may be detrimental to image quality, at ultra-high field strengths, it permits the use of MR spectroscopy. The various metabolites within a voxel have very similar precessional frequencies at 1 T, but these become more distinguishable at ultra-high field strengths. In a similar vein, spectral fat suppression techniques become more effective because the fat peak is isolated from the water peak and is more readily targeted by an off-resonant RF pulse without affecting water (see Chapter 7).

MAGNET SHIELDING

When siting an MRI scanner, it is desirable that the main magnetic field is not permitted to extend into areas other than the magnet room. This is because a powerful magnetic field can adversely affect nearby equipment (in imaging departments) and creates a potential safety hazard if members of the public gain access to a field strength of 5 G or greater (see Chapter 10). Shielding is achieved in one of two ways: **passive shielding**, which requires the scanner to be surrounded by large steel plates, or **active shielding**, which uses additional solenoid magnets.

Passive shielding

Before the introduction of active shielding in 1987, all MRI scanners were either passively shielded or not magnetically shielded at all [8]. For a 1.5 T scanner, it was possible for the 5 G threshold to be located as far as 10 m from the isocenter of the magnet. To maintain safety, it was necessary to site unshielded scanners in buildings with a large footprint, away from other hospital departments, thereby preventing the fringe field from extending into adjacent public areas. The fringe field was further reduced by passive shielding, which requires large steel plates incorporated either around the scanner or in the walls of the magnet room. Lines of magnetic flux travel through ferromagnetic media in preference to air. Passive shielding reroutes the fringe field away from the outside environment and back toward the scanner (Figure 9.11). Passive shielding has several major disadvantages. The iron cladding can weigh over 20 tons, it is very expensive, and the proximity of ferromagnetic metal can adversely affect the homogeneity of the scanner that it is intended to shield. For these reasons, passive shielding has now been replaced by active shielding in most clinical scanners and in some ultra-high-field research systems.




Figure 9.11 Passive shielding causes the lines of flux to pass through steel cladding in preference to air.

Active shielding

In addition to the main magnet solenoids, Figure 9.9 shows two larger diameter solenoids positioned at each end of the bobbin. These are colloquially known as **bucking coils** in that they oppose (buck) the effect of the main magnet windings. Their function is to actively shield the local environment by constraining the 5 G footprint of the fringe field to within a short distance from the scanner. To do this, the bucking coils carry a current flowing in the opposite direction to the main magnet windings, reversing the flux. This design was formerly suited only to magnets with a flux density up to 3 T, because in ultra-high-field research magnets, the induced mechanical stresses due to the Lorentz force were likely to have a self-destructive effect on the bobbin. In modern ultra-high-field-strength systems, active shielding is permitted, even at 8 T. Remember that although active shielding allows convenient siting of an MRI scanner, it does not remove the safety risk from projectiles. If anything, we should be more cautious about projectile safety because active shielding causes a very steep static field gradient. By the time we are aware that an item is ferromagnetic, there is little we can do to prevent it becoming a projectile as it suddenly enters the powerful part of the field (see Chapter 10).



Refer to animation 9.2 on the supporting companion website for this book:
www.wiley.com/go/westbrook/mriinpractice

Learning tip:**What is the Lorentz force?**

Faraday's law of electromagnetic induction states that charge, movement, and magnetism are inextricably linked. One cannot have any two of these factors present without generating or inducing the third. The Lorentz force is defined as "the force exerted on a charged particle moving with velocity through an electric and magnetic field" (Equation (9.2)). If a current-carrying conductor is placed perpendicular to an external magnetic field, the Lorentz force acts on the moving charges in the wire and causes physical deflection of the conductor. This explains the challenges in using active shielding at ultra-high field, and the acoustic noise emitted by the gradient coil during activation.

328

Equation 9.2

$$F = qE + qv \times B$$

F is the Lorentz force. It is the total emf in volts (V)
 q is the charge of a particle in coulombs (C)
 v is the velocity in cm/s
 E is an electric field vector
 B is the magnetic field vector

This equation shows that if a current-carrying conductor is placed perpendicular to an external magnetic field, the Lorentz force acts on the moving charges in the wire and causes physical deflection of the conductor

SHIM SYSTEM

The first component to be encountered within the warm bore of the cryostat is the shim system. This system operates at room temperature and is located around the circumference of the inner wall of the cryostat. The function of the shim system is to ensure homogeneity of the magnetic field within the imaging volume. One of the first tasks after the installation of a new MRI scanner is to assess the homogeneity of B_0 with the scanner in situ. A homogenous magnetic field is desirable for two main reasons. Firstly, any distortions in the magnetic field lead to geometric distortion of the images. Secondly, excitation of the hydrogen nuclei is frequency-dependent. To create resonance in a group of hydrogen nuclei, the transmitted RF pulse must match the precessional frequency of their magnetic moments. Inhomogeneity of the field results in a change in precessional frequency of magnetic moments of hydrogen nuclei within the imaging volume, and this has implications for the SNR and techniques that require spectral fat suppression. Such sequences employ an additional RF pulse each TR, tuned to the frequency of the magnetic moments of hydrogen nuclei in fat at the field strength in use (see Chapter 7). Any variations in frequency due to the harmonics of the magnetic field prevent the saturation pulse from achieving its purpose. Fat remains relatively hyperintense in areas where the frequency of the saturation pulse does not match the frequency of the magnetic moments of hydrogen nuclei in fat.

Homogeneity

Magnetic field homogeneity in MRI is typically described in ppm over the 40 cm (diameter) spherical imaging volume. Despite design improvements to the segmented solenoid over the

last three decades, the homogeneity of the bare magnetic field is still likely to vary by around 100 ppm on delivery. This may be further compromised by the local magnetic environment of the magnet room, for example, by steel in the framework of the building and by the ferromagnetic components of the scanner itself. To achieve high image quality, imaging at 1.5–3 T requires a homogeneity of at least 10 ppm or better.

A procedure known as **shimming** the magnet is employed to achieve this. The term “shim” is used in carpentry and engineering, and refers to a narrow strip of material (often tapered) that might be used to level a shelf or make minor adjustments to components. In the context of MRI, shims are typically small ferromagnetic plates that are positioned around the inner circumference of the cryostat. The task is not to physically adjust the magnet solenoids, but rather to compensate for inhomogeneity. Three main types of shimming are described in the literature, not all of which use ferromagnetic shims. These are now discussed in more detail.

329

Passive shimming

Passive shimming uses shims to adjust for large changes in magnetic field homogeneity. The inner circumference of the warm bore of the cryostat is lined with several long plastic trays that fit along the full length of the bore (Figure 9.12). These shim trays have around 14–16 spaced compartments designed to accommodate the small ferromagnetic shims described in the section above. The shims must be strategically positioned to balance out any inhomogeneity present in the bare field by manipulating the lines of flux. Remember, ferromagnetic objects powerfully attract the flux lines. To determine the optimum placement of the shims, an MRI engineer must first measure the flux density of the magnet over the entire 40 cm imaging volume. This is done




Figure 9.12 Passive shim system.

with a piece of equipment known as a plotting-rig. The rig is mounted inside the bore and has an adjustable probe with a small sample surrounded by a receiver coil. The sample is positioned at various spherical coordinates inside the imaging volume, and frequency measurements are obtained at each angular distance (azimuth) from the isocenter. A computer algorithm is then used to determine the correct placement for the shims. Following the shimming procedure, the homogeneity of the magnetic field is tested by comparing the phase evolution of the magnetic moments of the spins between the first and second echoes acquired using a dual-echo, gradient-echo pulse sequence. In a homogenous field, the magnetic moments of all spins should evolve to the same degree. Inhomogeneity causes an uneven phase shift across the FOV, and this is shown on a magnetic field map that is created using data comparing both echoes.

Active shimming

Active shimming uses electromagnets instead of ferromagnetic shims [9] and is used in addition to passive shimming. There have been various designs over the years using both resistive (room temperature) coils and superconducting solenoids. Resistive shimming coils were often positioned close to the gradient coils. The advantage of using resistive shim coils is that the shim can be manipulated at any time by adjusting the current flowing through the windings. In modern scanners, active shimming is usually performed by additional superconducting solenoids inside the cryostat. Their advantage is that no additional electrical power is required. Like passive shimming, this procedure is usually carried out once, at the time of installation, as manipulating the current flow in superconducting coils is time-consuming (see section on ramping a magnet).

Gradient offset (dynamic) shimming

The final method for shimming uses the gradient set, which is another electromagnet that is designed to manipulate the magnetic field during image acquisition. The primary functions of the gradient system are discussed in the following section, but in terms of shimming it is possible to apply a current to the gradient coils that offsets any minor inhomogeneity in the main magnetic field. As stated earlier, it is desirable to achieve a homogeneity of better than 10 ppm for good image quality. By introducing a large patient into the magnet bore, the homogeneity can be changed by up to 9 ppm by diamagnetic repulsion. Dynamic shimming can help to correct for this. The primary purpose of the gradient system, however, is to create linear slopes along the main magnetic field.

GRADIENT SYSTEM

Moving toward the center of the warm bore, the next component to be found is the gradient system. As discussed in Chapter 5, the system must localize the origin of the MRI signal. The accuracy of the location of data samples in k -space relies on the use of high-precision magnetic field gradients (see Chapter 5). This task is achieved by the gradient system by creating linear magnetic field gradients (slopes) along the three orthogonal axes of the imaging volume. These mutually perpendicular axes are labeled x , y , and z according to Cartesian geometry. In a closed-bore magnet system, with a patient being positioned in the supine position, the y -axis lies vertically (posteroanteriorly), the x -axis is orientated horizontally from left to right, and the

z-axis lies horizontally along the length of the magnet bore (inferosuperiorly or head to foot). In an open MRI system, the flux lines of B_0 run vertically between the poles of the magnet, and the z and y gradient directions are therefore reversed in comparison to the closed-bore configuration described above.

Learning tip: Remembering gradient directions

The easy way to remember gradient directions is that the z gradient is always in the direction of the main magnetic field, and the x gradient always lies across the patient left to right (as an aide memoire, the shape of the letter x is also a cross – across). The y gradient therefore lies in the one remaining plane.

331

Gradient coil

The gradient coil is a cylindrical solenoid electromagnet; however, its design differs to those previously discussed. Rather than using conductive windings, the modern gradient coil is a copper-plated cylinder with the conductive elements etched into the surface of the metal plating. This allows a more complex configuration in a relatively compact and light form-factor. This style is known as a “fingerprint” design, reflecting the fact that the etched pattern looks like a human fingerprint. The copper etching is supported by a shell to prevent physical vibration or distortion of the coil by Lorentz forces during use. Importantly, the elements of the gradient coil are supplied by three separate power sources (called **gradient amplifiers**) allowing manipulation of the main magnetic field in any direction. As the gradient set operates at room temperature, application of electrical current causes heating due to resistivity. This usually necessitates the use of a water cooling circuit, especially in high-power gradient sets. Another consideration in gradient coil design is that the gradient magnetic field should be applied across the imaging volume, but should not extend further into the cryostat. For this reason, the outer layer of the gradient coil provides active shielding. The principle is similar to that of the active shielding of the static field whereby an equal but opposite current is applied to the shielding coil. The main difference is that active gradient shielding operates at room temperature. The reason that the gradient field should not be permitted to extend into the main magnet solenoids is that the resulting interaction causes unwanted currents (known as eddy currents) that can affect the quality of the main magnetic field.

The basic principle behind the mechanism of a gradient coil is illustrated in Figure 9.13. This diagram shows a representation of the gradient coil located inside the solenoids that generate B_0 (shown in cross-section). In this diagram, current flow in the gradient coil is represented by red and blue. Power is supplied by the z gradient amplifier, and, as a result, the current flows in different directions at each end of the coil in the z direction. The red elements of the gradient coil carry current in the same direction as the windings of the main magnet solenoids. This has the same effect as increasing the number of windings in this solenoid and correspondingly increases the magnetic flux density at the +z end of the imaging volume. The blue elements of the gradient coil carry current in the opposite direction to the windings of the main magnet solenoids. This has the same effect as reducing the number of windings in this solenoid and subtracts from the main magnetic field at the –z end of the imaging volume. Between these two




Figure 9.13 Mechanism of a gradient coil.

extremes, there is a fairly linear slope along the main magnetic field in the z-axis as shown in Figure 9.14. The diagram shows that the linearity begins to decrease at the limits of the imaging volume. If uncorrected, this causes distortion of the anatomy at the edges of the FOV. This effect is typically reduced by computer manipulation of the reconstructed image. A similar technique is also used to improve the geometric accuracy of images acquired on open MRI systems having poor static field homogeneity. The polarity of a gradient may be changed by reversing the current flow in the gradient coil elements. In open magnet systems, the gradient coils are flattened into disc-like structures that are positioned above and below the patient between the magnet shoes (or magnet solenoids) and the RF transmitter. Although they are structurally different from the cylindrical coils used in closed-bore systems, they function in a similar way.

Gradient amplifiers

A gradient coil requires a power supply. The power is generated in three gradient amplifiers, one for each orthogonal gradient direction. The gradient amplifiers are typically housed in a cabinet remote from the scanner. Their function is to supply the current required by the gradient coils during image acquisition. Older models of MRI scanner used linear analog amplifiers; more recent models are equipped with a different type of amplifier having what is known as a pulse width modulated (PWM) design [10]. Linear amplifiers are similar in operation to a music




Figure 9.14 How gradients change field strength.

amplifier. The waveform created by the output is identical to the input waveform, only larger (i.e. amplified). For example, a disc jockey (DJ) deck tracks the fluctuations of a needle in the groove of a record. It creates a tiny oscillating waveform that must be amplified to drive the speakers to fill a large room with sound. In MRI, a gradient amplifier receives a comparatively low-amplitude signal from a pulse generator that must be amplified to a level that is used to drive the coil. The applied current may be as high as 900 A, and the voltage ranges between 1000 and 2000 volts (V). The main drawback of this linear method of amplification is that it causes power losses (heating) in the gradient coil, but also in other parts of the circuit, including amplifiers [10]. Efficient gradient operation is problematic in linear amplifiers because of losses in the components that perform this task (semiconductors). Put simply, the higher the percentage of time that a gradient spends at maximum amplitude (known as the **duty cycle**), the more heating occurs. PWM amplifiers help reduce this burden by applying voltage to the coil in short closely spaced discrete bursts. This technique improves the efficiency of the gradient coil, and, if the pulses are of sufficiently high frequency (closely spaced in time), the coil interprets the supply as smooth and continuous.

Gradient characteristics

During data acquisition, the gradient coils are required to ensure that the transmitted RF is targeted at the volume of tissue required in each individual slice (see next section). Additionally, the gradient set is also responsible for determining the spatial origin of signal returning from each

slice along the two remaining axes of anatomy, for example, from left to right and from anterior to posterior in an axial slice (see Chapter 5). During the acquisition, each of the elements of the gradient coil is activated thousands of times. Every time the coil is activated, power is applied to the conductive coil elements until the gradient slope reaches the maximum required amplitude in a particular direction. The element is then briefly deactivated before the next stage of the acquisition. Some gradient functions (such as phase encoding) require gradients to be switched on at different amplitudes and polarities (directions) throughout the acquisition (see Chapters 5 and 6). Importantly, while a gradient coil is activated, the magnetic moments of the hydrogen nuclei within the imaging volume experience slightly different flux densities along the length of the slope of the induced gradient. The Larmor frequencies of these magnetic moments therefore vary in a linear way along the axis of the gradient slope. When the gradient coil is deactivated, magnetic moments of the hydrogen nuclei return to the center frequency, but exhibit different phase positions (Figure 9.14). The vectors that momentarily experienced a higher flux density during the gradient application are phase-advanced compared to those that were at a lower flux density at a different location along the gradient slope.

The ideal gradient set should be powerful, capable of high-amplitude gradient slopes, and rapid to reduce scan time. These attributes are described by the terms gradient amplitude, rise time, **slew rate**, and duty cycle. These features are shown in Figure 9.15.

Gradient amplitude

Gradient amplitude defines the power of the gradient, specifically how steep a gradient slope can be achieved when the gradient coil element is activated. Because the gradient causes changes

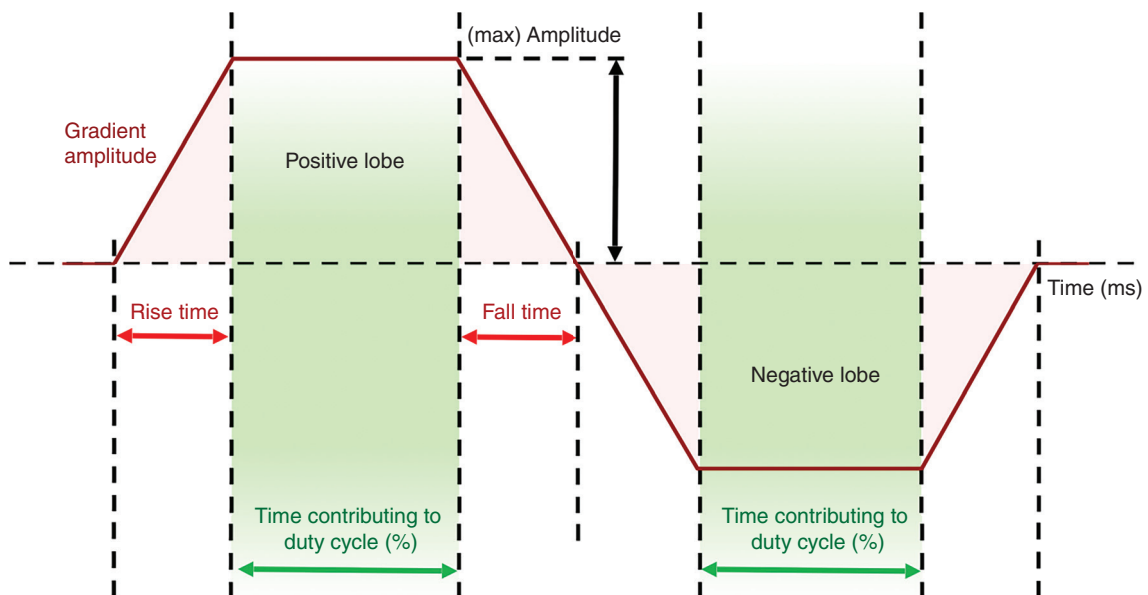


Figure 9.15 Characteristics of a magnetic field gradient.

in magnetic flux density over a certain distance (for example, 40 cm from left to right across the magnet bore), the units used to measure gradient amplitude are mT/m or G/cm (remember the units of gauss are 10 000 times smaller than teslas). On modern scanners, typical values are below 40 mT/m.

Gradient rise time

Gradient rise time is defined as the time taken for the magnetic field gradient to reach the required amplitude. It is measured in microseconds (μs) and values are typically below 1000 μs . Techniques such as diffusion imaging may use gradient rise times as short as 130 μs . Rise time cannot be used as an indicator of overall gradient performance *per se*, because a low-amplitude gradient is likely to have a shorter rise time than a high-amplitude gradient.

335

Gradient slew rate

Gradient **slew rate** combines the two previously mentioned factors and describes how quickly the gradient magnetic field can be applied at a given amplitude and over a given distance measured in meters. For this reason, the units used for measurement are T/m/s. Slew rate is determined by dividing the amplitude of the gradient by the rise time. For example, if the required gradient amplitude is 30 mT/m and the achievable rise time is 0.3 ms, the slew rate is 100 mT/m/ms (or 100 T/ms). For safety reasons, values are typically kept below 200 T/m/s.

Gradient power duty cycle

The gradient power duty cycle is, perhaps, a more meaningful measure of gradient performance. This is defined in different ways by different manufacturers and different sources but put simply, the duty cycle can be thought of as gradient “on-time.” This is the time that the gradient coil is operating at the required maximum amplitude and is expressed as a percentage of the total acquisition time. An efficient duty cycle ideally requires gradient amplifiers employing high-amplitude gradient pulses with very short spacing between pulses. In practice, this is not always feasible, not only because of the rise time of the gradient but also because of heating concerns.

These factors are important in practice because they determine several key aspects of an MRI acquisition and the features of the resulting image. Gradient slew rate and duty cycle contribute to the overall acquisition time. Acquisition (scan) time is also affected by the TR and TE (see Chapters 6 and 7); however, in pulse sequences that require a high temporal resolution, such as perfusion techniques or fMRI, a rapid slew rate and high duty cycle are needed. There are, however, technological and physiological concerns that may require limitation of these factors. In terms of the MRI hardware, rapid gradient switching can cause overheating of the gradient coils. Some implanted devices may also be adversely affected, and manufacturers place conditions on scanning related to gradient slew rate. The physiological limitation to gradient slew rate relates to the fact that a time-changing magnetic field causes an induced voltage in a nearby conductor. This includes the patient’s nervous system (see Chapter 10). In terms of the geometry of the MRI image, the gradient amplitude is determined by the slice thickness, phase resolution, and FOV (see Chapter 7).

Acoustic gradient noise

Another potential issue related to the gradient system is the acoustic noise that has become a trademark of MRI examinations. Earlier in this chapter, we discussed that a Lorentz force is exerted when a charge moves through a magnetic field. In the gradient system, rapid pulsing of electric current through the conductive elements of the coil results in rapidly oscillating mechanical forces on the coil structure. Provision is made for the firm attachment of the gradient coil within the warm bore of the cryostat to prevent gross physical motion during image acquisition. This does not, however, prevent the coil from vibrating, and the resulting acoustic noise can be considerable. Pulse sequences that require rapid gradient switching (such as EPI) tend to cause the highest levels of acoustic noise with some studies reporting intensities of up to 115 decibels (dB) or higher (see Chapter 10).

336

Balanced gradients

Figure 9.15 shows a plot of gradient amplitude over time in millisecond. The horizontal axis represents time, and the vertical axis represents gradient amplitude. The gradient is applied using two lobes. The first lobe shows a gradient applied positively, with the maximum amplitude above the horizontal axis of the graph. The second lobe shows the gradient applied with an equal but opposite polarity whereby the maximum amplitude lies below the horizontal axis. In practice, this is achieved by reversing the direction of current flow through the gradient coil (see Chapter 5). The diagram also shows that the gradient is balanced, in that the area under each lobe is the same size. If the gradient is only applied as a single lobe, there is dephasing of the magnetic moments of spins along the slope of the gradient. This causes a loss of signal and is therefore undesirable for most of the gradient functions, for example, when selecting a slice, or during the sampling of signal.

To correct this, the second lobe of the gradient is applied with opposite polarity. The mechanism relies on the fact that a gradient magnetic field either adds to, or subtracts from, the main magnetic field. As previously discussed, the magnetic moments of spins at the high end of the gradient exhibit a higher precessional frequency while the gradient is activated. The magnetic moments of spins at the lower field end of the gradient exhibit a lower precessional frequency. By applying the second lobe of the gradient with negative polarity, the direction of the slope is reversed. The spins that previously experienced a higher field during the application of the first lobe experience an equal but opposite negative field and vice versa. The result is that the magnetic moments of all the spins are put back into phase. To use the examples provided earlier, during slice-selection the RF excitation pulse is only applied during the first lobe of the gradient, ensuring that RF only excites the spins at a matching frequency along the gradient slope. The undesirable dephasing of the magnetic moments of spins caused during this process is then corrected by the second lobe. In simple frequency encoding, the gradient is activated during sampling. The timing is therefore determined by the required digital sampling rate and the number of samples to be acquired (see Chapter 6). These factors are influenced by the receive bandwidth and the required spatial resolution of the image in the frequency direction.

If a corrective lobe were to be applied at the same amplitude and timing as the lobe applied during readout, the resulting TR may have to be lengthened. This is undesirable, as it may reduce the maximum number of slices that can be acquired. To avoid this, the corrective lobe is applied at a higher amplitude but for a shorter time, the area under each lobe remaining equal but opposite. In spin-echo pulse sequences, the corrective lobe is applied first, and sampling takes place during the second lobe of the frequency-encoding gradient. This means that the first

Table 9.3 Things to remember – gradient coils.

There are three gradients x, y, and z that perform several functions during a sequence
The amplitude of a gradient is determined by the amount of current passing through the coil
The polarity of a gradient is determined by the direction of the current flowing through the coil
The amplitude of a gradient determines the spatial resolution
The slew rate determines how fast data can be acquired

lobe may be applied at any convenient time before the spin-echo. The only caveat is that if the first lobe is applied before the 180° RF rephasing pulse, it must be at the *same* polarity as the second lobe, because the 180° RF rephasing pulse essentially reverses the effect of the gradient (called the Stejskal–Tanner scheme). This bipolar gradient application is commonly used in MRI pulse sequences unless the aim of the gradient is to create a net dephasing of the magnetic moments spins, which is the case in phase encoding (see Chapter 5).

In summary, the gradient system is an essential component of the MRI system, determining scan time, slice location, spatial encoding, and (indirectly) image contrast in studies using gradient-echo pulse sequences, perfusion, and diffusion imaging.

RF SYSTEM

The next major component of the MRI scanner is the RF system. The purpose of the RF system is to transmit and receive electromagnetic radiation during image acquisition. In classical physics, electromagnetic radiation describes waves that propagate through a vacuum at the speed of light and have both electrical and magnetic properties (see Chapter 1). For this reason, RF used in MRI is often referred to as B_1 being a secondary magnetic field (in addition to B_0 the main magnetic field) (see Chapter 1). The purpose of the transmitted RF is to target a volume (slice) of tissue within the region to be examined and increase the energy level of a proportion of the hydrogen nuclei within that slice. This energy transfer is achieved by the phenomenon of nuclear magnetic resonance whereby the transmitted RF is applied at the Larmor frequency of magnetic moments of the hydrogen nuclei (see Chapter 1). The Larmor equation states that the precessional frequency of the magnetic moment of an atomic nucleus is proportional to the external magnetic field (see Equation (1.3)). Hydrogen has a gyromagnetic ratio of 42.58 MHz/T. This means that in an external magnetic field of flux density 1 T, the magnetic moments of hydrogen nuclei precess at a rate of 42.58 million times per second (MHz). At 1.5 T, the rate increases to 63.87 MHz, and at 3 T the frequency increases to 127.74 MHz (see Chapter 1). All these frequencies lie in the RF portion of the electromagnetic spectrum, specifically the VHF (very high frequency) band, which is shared with many other applications such as nondigital FM radio and television (TV) broadcasts, maritime communication, and air traffic control. For this reason, the magnet room is shielded from external radio sources.

RF shielding

The amplitude of signal generated within the tissues of the patient during image acquisition is vanishingly small compared to similar frequencies transmitted for radio messaging and broadcasting. If these external frequencies were to be detected by the system, they would create

artifactual appearances on the resulting images. In addition to this, the RF produced during the scanning process may cause interference with other nearby electrical equipment. To circumvent this problem, it is necessary to shield the entire scan room using what is known as a Faraday cage. Unlike traditional X-ray rooms, which may require lead shielding, RF shielding is accomplished by lining the walls, floor, ceiling, and door with any conductive metal. Copper was the original material of choice, due to its excellent conductivity, but it is comparatively expensive and heavy. Aluminum panels are now favored by many manufacturers, as they are easier to handle during equipment upgrades and are suitable for mobile installations where weight is a prime concern. Wherever there needs to be a breach in the metal shield, such as the observation window or air conditioning system, it is necessary to employ a wire mesh or honeycomb filter. This type of filter has an array of small cells (hence the name) that are designed to absorb and reflect radio waves, preventing ingress into the magnet room. *Ad hoc* access to the magnet room (for example, piped gases or fluids) is provided by a penetration plate. This panel is usually located near the floor of the room and is pierced by metal cylinders known as wave guides. The design of the wave guide is such that small items can pass through into the magnet room, but radio waves of the frequencies used in MRI cannot. The penetration panel also usually features RF filtered multipin connectors that are used to supply electricity to any devices inside the magnet room.

RF transmit system

Having isolated the magnet room from outside sources of radio interference, RF pulses are transmitted and received during image acquisition. The transmitted RF is an electromagnetic wave, which contains a narrow range of frequencies, centered around the Larmor frequency of the magnetic moments of hydrogen nuclei at the external field strength of the system. This narrow range of frequencies is known as the transmit bandwidth and, in turn, determines the duration of the transmitted pulse (see Chapter 5). To excite a rectangular slice profile, the pulse must be rectangular in the frequency domain. This means that the pulse should ideally only contain the narrow range of frequencies imposed by the slice-select gradient across the thickness of the slice and that these frequencies should be at the same amplitude. In the field of signal processing, the shape of the RF pulse used in MRI is known as a sinc pulse (sine cardinal) [11]. After FFT this type of pulse creates a slice profile with well-defined boundaries and an approximately rectangular profile when viewed from the edge.

As shown in the top half of Figure 9.16, the sinc pulse is generated by a digital waveform generator and is immediately converted into an analog waveform. This is then processed by a synthesizer that ensures that the frequencies in the pulse are centered around the Larmor frequency of the magnetic moments of the spins. In addition, the synthesizer determines the phase of the




Figure 9.16 RF transmit and receive chain.

waveform, as it may be desirable in some sequences to adjust the phase of the transmitted RF pulses (see Chapter 4). The waveform is then passed through an **RF amplifier** to provide enough energy to couple with the nuclei in the region of interest. This is a rack-mounted solid-state device typically operating at 18 kW or greater in 1.5 T systems and up to 40 kW in 3 T systems. The RF pulses thus generated by the coil are transmitted into the patient repeatedly throughout data acquisition. The pulses tend to be described according to their function (excitation pulses, rephasing pulses, inversion pulses) and are timed to occur at certain time periods throughout each repetition (e.g. TR, tau, and TI) (see Chapters 2–4).

RF transmit coils

The **RF transmit coils** used in MRI are known as resonator coils [12]. When this type of coil is caused to resonate at the required frequency, large oscillations in voltage and current are produced in the coil. These, in turn, generate radio waves at the desired frequency, which are transmitted into the anatomical area under investigation. The aim of the transmit coil is to transfer energy to the hydrogen nuclei. The most efficient way to do this is to generate a field that not only oscillates at the Larmor frequency but also rotates in the same orientation as the spins. To achieve this, the RF transmitter is circularly polarized (also known as a quadrature design). This means that the coil is connected to two power supplies coming from an RF amplifier. The supply from the second terminal is delayed by 90° with respect to the first. This results in two RF fields that are perpendicular to each other, maximizing energy coupling between the coil and the spins. The radiofrequencies are generated either by a large RF transceiver that is situated immediately around the inner circumference of the magnet bore or, in some systems, by smaller detachable coils such as those used for scanning the head. The term **transceiver** is used to denote the fact that this coil transmits RF and, if needed, also receives RF.

In closed-bore MRI systems, the main RF transceiver is colloquially known as the body coil and typically features what is known as a **birdcage design**. This name reflects the symmetrical arrangement of several electrically conductive elements situated inside a cylindrical plexiglass structure broadly resembling a birdcage. As shown in Figure 9.17, there are two large circular conductive loops (known as end-rings) located at each end of the structure, and these are connected along the length of the magnet bore by an even number of straight copper strips. The end-rings must be larger than the circumference of the patient bore, greater than 70 cm in diameter, and the length of the coil is typically around 50 cm in length. This is desirable to achieve a large FOV, such as is required in abdominal imaging. A similar design is used in a transmit head coil, but scaled down in size.

With any MRI coil, it is necessary to ensure good RF homogeneity and signal quality. Signal quality is largely dictated by a feature known as the **filling factor** of a coil. This is determined by how well the shape and size of the coil are matched to the anatomical region under investigation. The RF homogeneity of the coil (i.e. the ability of the coil to provide RF at the same amplitude across the entire imaging volume) is dictated by the configuration of the conductive elements of the coil. RF homogeneity tends to improve as the coil size increases; however, there is then a trade-off with filling factor. Generally, therefore, RF transmit coils are positioned close to the patient, must enclose the entire region of interest, and must have evenly spaced, rotationally symmetrical elements surrounding the anatomical area.

RF receive system

As the name suggests, the aim of the RF receive system is to receive signal returning from the patient at time TE. This process relies on the fact that the oscillating electromagnetic field




Figure 9.17 The RF transmit coil. This is a birdcage resonator coil consisting of two end rings linked by a symmetrical array of straight conductors.

generated by the magnetic moments of the hydrogen nuclei induces a corresponding oscillating voltage in a closely placed conductor (coil). This returning signal is typically of VHF, but low amplitude. High-frequency, low-amplitude signal can be difficult to process, due to the limitations of the components used for this. Older scanners required that signal be demodulated to a lower frequency before processing (see Chapter 6). With recent advances in technology, signal can now be sampled directly.

In addition to useful signal, the receiver coil also detects a certain amount of undesirable noise. Noise in MRI is mainly caused by the thermodynamic motion of electrons both in the coil and inside the body tissues (see Chapter 7). Noise occurs at all frequencies and is therefore manifested over the entire image unless filtered out or smoothed. Amplification of signal and filtering of noise are achieved by a series of components found in the receive chain highlighted in red in the bottom half of Figure 9.16. The digitized signal is then routed back to the computer for processing. On modern scanners, the amplification, filtering, and digitization of analog signal are done inside the receiver coil assembly or the body of the MRI scanner. This permits the digital signal to be sent to the computer via fiber-optic cables, which are immune to signal losses, external interference, and noise spikes that can affect the older coaxial wire conductors.

RF receive coils

All MRI scanners are provided with several receiver coils. These are designed to receive signal from any of the anatomical areas that may be examined using the MRI scanner. As such, to achieve a satisfactory filling factor, the structure and appearance of the coils vary greatly to match the size and shape of the anatomy under investigation. Despite the range of size and shape, there are essentially only two types of receiver coil, **surface coils** and **volume coils**. As the

name suggests, surface coils are used to image surface anatomy, and volume coils are required when a large volume of tissue is examined.



Refer to animation 9.3 on the supporting companion website for this book: www.wiley.com/go/westbrook/mriinpractice

Surface RF receive coils

Surface coils are not typically used as RF transmitters because of their poor RF homogeneity. They are described as receive-only coils. Their design, as shown in Figure 9.18, usually consists of a single conductive loop or an array of loops that are either connected or overlapped and operated independently via separate channels. The latter configuration is known as a phased array. In normal use, surface coils are positioned flat, in proximity to the anatomy under examination. This provides an optimum filling factor whereby the entire sensitive volume of the coil contains tissue. This may not be the case with volume coils, which may contain significantly more patient tissue than the area that is being scanned as well as empty space. Surface coils therefore receive a high signal from the surface anatomy, offering a two to fivefold increase in sensitivity compared to volume coils. An additional factor that contributes to the high SNR from this type of coil is that they only receive noise from a small region of the anatomy. A volume coil (especially the integrated RF transceiver) typically contains areas of the body other than the region of interest, and although signal only returns from a specific thin volume of tissue (slice), noise may be detected from all other anatomy inside the coil.

The trade-off for this high SNR is that the sensitivity decreases with the distance from the coil. Useful signal is only returned from a tissue depth equal to the diameter of the conductive loop. This makes the coils unsuitable for imaging anatomy deep inside the patient. For example, a small 10 cm coil offers an FOV of 10 cm and only detects useful signal to a depth of around 5–7 cm. For these reasons, single-loop surface coils are falling out of favor in modern-day MRI scanning. They are, however, ideally suited to MRI microscopy with dedicated small coils available for high-resolution examinations of areas such as the finger joints or the orbits.

Phased array surface coils are now offered as standard and are ubiquitous. They offer the SNR advantages of single-loop coils, but with a greatly extended FOV and many different configurations. A good example is the linear phased array used in spine imaging. In this device, there is an array of small surface coils (elements), each having its own channel through which the detected signal is carried to the signal processing components of the system. As stated earlier,


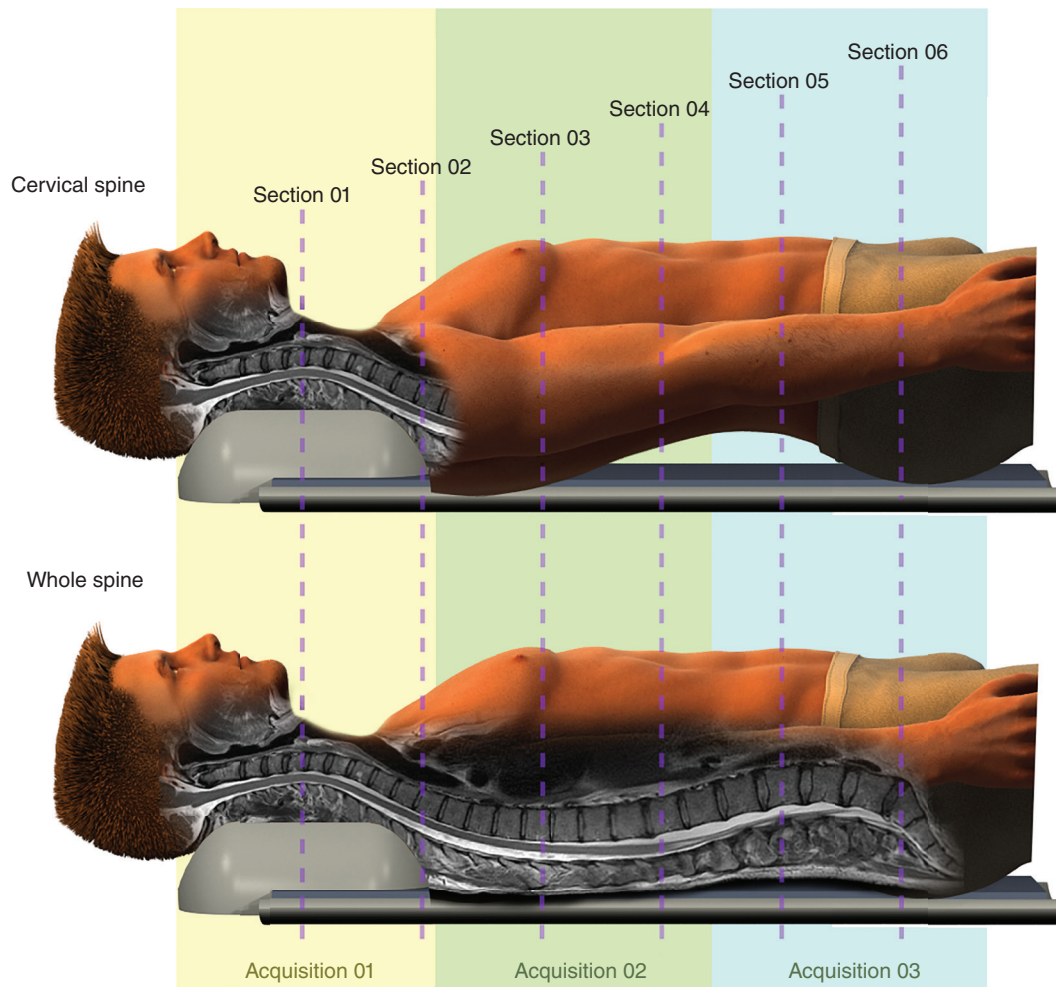


Figure 9.18 Schematic diagram showing surface coil configurations.



342

Figure 9.19 Phased array spine coil.

these components may now be incorporated into the receive coil. In Figure 9.19, there are six sections to the coil, each of which contains one or more overlapping coil elements. The dotted lines represent the center of the sensitive volume for each section. These are usually marked on the chassis of the coil to aid positioning. The coil elements contained in section 01 are incorporated into a neck former and may be used in isolation to acquire a cervical spine study. Should a larger FOV be required, we can select more coil elements and the patient transport system adjusts the position of the patient to the appropriate isocenter. By performing several acquisitions, the entire length of the spine is imaged. The number of coil elements and channels available in an MRI coil varies between 2 and 128 depending on the purpose of the coil and the design offered by the manufacturer. In the spine array, the coil elements are located along the full length of the anatomy, allowing a maximum FOV. A secondary advantage is that each channel contributes not only its own signal but also its own proportion of background noise. Although seemingly counterintuitive, there is an improvement in SNR because thermal noise is

Table 9.4 Things to remember – receiver coils.

Receiver coils are a critical part of image optimization. The selection of the appropriate receiver coil for the area under examination is very important

Large coils provide large coverage but a relatively poor SNR. This is because the anatomical area may not fill the entire sensitive area of the coil. The filling factor is low.

Small coils provide small coverage but relatively good SNR. This is because the sensitive volume of the coil will be filled with tissue. The filling factor is high. However, aliasing is more likely

Phased array coils of the linear and volume type are the best option, as they combine the benefits of using small coils with those of using large ones

343

random and the contributions to noise made by each channel tend to average out in comparison to signal. A final advantage of phased array coils is that they permit parallel data acquisition (parallel imaging). This is facilitated connecting the multiple channels with miniaturized data processing components inside the coil. Each coil element simultaneously contributes data to k -space, reducing scan time (see Chapter 6).

Volume RF receive coils

There are several types of conductor with volume coil configurations, the most common being the birdcage design. Other configurations include the saddle-shaped design and the transverse electromagnetic resonator. As we saw earlier in the chapter, volume coils are typically employed to transmit RF, but there are also receive-only volume coils. The primary advantage of a volume coil is that, unlike surface coils, they offer a comparatively homogenous RF field over a large imaging volume. This means that volume coils can be used to image anatomy at any location within the region of interest, at any depth. Signal from any anatomical structure enclosed by the coil is received with equal sensitivity. Volume coils are now offered with phased array capability. Head coils have up to 48 channels, and some modern scanners offer up to 128 channel capability for certain coils such as advanced cardiac and body coils. The flexibility to arrange multiple receivers around an anatomical area has therefore blurred the distinction between volume and surface coils. The only potential disadvantage to using a volume coil is that the filling factor tends to be lower than that of a surface coil. The coil is not totally filled by the anatomy and contains empty space, reducing the SNR. In addition, areas of anatomy that *are* contained within the coil contribute to the thermal noise, but may not contribute to signal. This SNR issue is largely compensated for by the fact that these coils are multichanneled. Importantly, manufacturers ensure that volume coils are designed to fit closely around the desired anatomical area.

PATIENT TRANSPORT SYSTEM

The patient transport system is typically a large nonferromagnetic patient couch (table) that is raised and lowered to facilitate patient access and driven horizontally into the magnet bore. This movement is controlled from a panel situated on the front cone of the scanner cover. From a geometric perspective, the aim of the couch movement is to position the center of the region of interest at the true isocenter of the imaging volume where the **magnetic homogeneity** is greatest. The front cone of the scanner is equipped with laser positioning lights. The region to be scanned is centered to a crosshair marker formed by the lasers, and the position is recorded by pushing

a button on the front panel. The table is then activated to automatically move the anatomical region into the isocenter of the bore. In open MRI systems, the couch top may also be adjusted sideways, which is of great benefit when imaging lateral structures such as the shoulders or elbows. The patient couch may also include posterior receive coil elements and sockets into which surface coils can be mounted. There is typically a light source inside the magnet bore, and a patient cooling fan may also be included. Manufacturers also offer patients a prismatic or reflective headset that allows them to watch videos, or pleasant immersive scenery during the procedure. This may alleviate nervousness and provide information to the patient about the duration of each acquisition (see Chapter 10). There is also a call-button and a patient microphone to allow two-way communication between the patient and staff at the imaging console.

344 COMPUTER SYSTEM AND GRAPHICAL USER INTERFACE

The entire process of an MRI acquisition is orchestrated by the host computer. The graphical user interface (GUI) of the system is used to identify the patient, usually via a network link to the radiology information system. The age, weight, gender, and physical orientation of the patient may also be recorded. This ensures that the system is prepared to calculate the amplitude of the RF pulses and label the physical orientation of anatomy shown on the images. The interface also usually contains pages or tabs relating to the various parameters that are required for the scan. These may be grouped by function, for example, a geometry page that allows adjustments to FOV, image matrix, and slice thickness. The parameters required for an MRI scan are numerous, and it is undesirable (from a patient-throughput perspective) for us to manually enter these factors in every acquisition. The computer is therefore preprogrammed by the manufacturer with commonly used MRI protocols. These can be adjusted, optimized, and resaved as required (see Chapter 7). Some scanners permit sharing and downloading of protocols from online databases. The host computer also displays patient images and permits their manipulation, changing window settings and performing various postprocessing functions such as MIP (see Chapter 8). Like a modern airliner, the latest MRI scanners are also directly connected (via a wide area network) to the manufacturer allowing every aspect of the system to be remotely monitored and faults detected.

In addition to the host computer, the system uses dedicated modules to control various aspects of the MRI process. The most important of these is a **pulse generator module (PGM)**, an independent computer that is responsible for sending instructions relating to the timing, amplitude, and shape of the transmitted RF pulses and the timing and duration of the sampling window. The PGM is also connected to the gradient amplifiers and generates pulses used to determine the shape and duration of the gradients used during the scan. For these reasons, the PGM also needs to be connected to any physiological sensors that are monitoring the breathing rate or pulse rate of the patient. Another function of this module is to ensure that the patient table is adjusted to ensure that the area under investigation is always positioned at the isocenter of the imaging volume. The MRI system may also employ a dedicated image-processing module. This has become increasingly important since the introduction of multiple channel imaging and parallel imaging where image reconstruction is more complex. A dedicated image processor speeds up this task. This, in turn, increases patient throughput, because an examination cannot be concluded until the images have been reconstructed and checked for quality. A fast acquisition offers no great advantage if the image reconstruction time is long.

Finally, the host computer must also be connected to a server for image archiving. This contains a redundant array of independent drives for storing patient images in DICOM format (digital

imaging and communications in medicine). The redundant array is used to protect from data loss. Data from the archive are also accessed via the PACS (picture archiving and communications system) access network permitting teleradiology, and access from other hospital departments, such as the operating room.

In this chapter, the different components of the MRI system were analyzed in detail. In the next and final chapter, issues associated with the safe use of this equipment are described along with other safety considerations in MRI.



For questions and answers on this topic please visit the supporting companion website for this book: www.wiley.com/go/westbrook/mriinpractice

References

1. Cox, B. and Cohen, A. (2017). *Forces of Nature*, 25. London: Collins.
2. National Research Council (2013). *High Magnetic Field Science and Its Application in the United States: Current Status and Future Directions*, 185–195. Washington, DC: National Academies Press.
3. Overweg, J. (2006). MRI main field magnets. *Proceedings of the 14th ISMRM*, Seattle, WA (6–12 May 2006), 1. ISMRM.
4. Overweg, J. (2006). MRI main field magnets. *Proceedings of the 14th ISMRM*, Seattle, WA (6–12 May 2006), 5. ISMRM.
5. Charifouline, Z. (2006). Residual resistivity ratio (RRR) measurements of, LHC superconducting NbTi cable strands. *IEEE Transactions on Applied Superconductivity* 16 (2): 1188–1191.
6. Cosmus, T.C. and Parizh, M. (2011). Advances in whole-body MRI magnets. *IEEE Transactions on Applied Superconductivity* 21, 2108.
7. Slade, R.A., Parkinson, B.J., and Walsh, R.M. (2014). Test results for a 1.5 T MRI system utilizing a cryogen-free, YBCO magnet. *IEEE Transactions on Applied Superconductivity* 24(3).
8. Hawksworth, D.G., McDougall, I.L., Bird, J.M. et al. (1987). Considerations in the design of MRI magnets with reduced stray fields. *IEEE Transactions on Magnetism* 23 (2): 1309.
9. Poole, M.S. and Shah, N.J. (2014). Convex optimisation of gradient and shim coil winding patterns. *Journal of Magnetic Resonance* 244: 36–45.
10. Schmitt, F. (2013). The gradient system. *Proceedings of the 21st ISMRM*, Seattle, WA (20–26 April 2013), 5. ISMRM.
11. Bernstein, M.A. and King, K.F. (2004). *Handbook of MRI Pulse Sequences*, 37. Cambridge, MA: Academic Press.
12. Vaughan, J.T. and Griffiths, J.R. (2012). *RF Coils for MRI*. Oxford: Wiley Blackwell.

10

MRI safety

Introduction (and disclaimer)	346	Time-varying gradient magnetic fields	363
Definitions used in MRI safety	347	Cryogenics	365
Psychological effects	350	Safety tips	367
The spatially varying static field	351	Additional resources	368
Electromagnetic (radiofrequency) fields	357		

After reading this chapter, you will be able to:

- *Recognize the main hazards associated with MRI scanning.*
- *Understand the mechanisms behind MRI-related injuries.*
- *Identify and ameliorate potential risks to the patient undergoing an MRI procedure.*

INTRODUCTION (AND DISCLAIMER)

MRI safety is a convoluted subject. Over the last 20 years, the field of MRI has increased greatly in complexity. Scanners now operate at a wide range of field strengths and manufacturers offer different machine configurations, some of which operate with powerful and rapidly switching gradient magnetic fields. Coil design, magnet bore width, and other factors also vary across different platforms. In addition to variations in hardware, patients present with a diverse variation in body habitus and an increasing variety of implanted devices and cosmetic body jewelry. Implanted devices may be conditionally safe below a certain field strength, spatial gradient magnetic field, or **specific absorption rate (SAR)**. As a result, there are many combinations of factors that could potentially contribute to a safety hazard. This chapter is therefore not intended to be a definitive safety guide for the use of MRI but rather a general introduction to the topic of MRI safety. By the same token, pharmacological safety issues with gadolinium-based contrast agents are not within the brief of this book. Instead, this chapter aims to identify the generic safety factors relating to the MRI equipment and associated magnetic fields and explains the mechanisms behind the common types of injury that may occur during an MRI procedure.

The overarching philosophy behind MRI safety recognizes that every case is different and should be considered on its own merits. For patients with implanted devices, this is likely to require a careful risk/benefit analysis. On the one hand, patients should not be needlessly excluded from an MRI procedure that is necessary for their ongoing care. However, any potential risks need to be carefully considered. The primary responsibility for this task lies with the referring physician and the radiology lead. Crucially, the MRI practitioner should also be permitted to decide whether it is safe to proceed with an examination, or whether it should be postponed until sufficient evidence has been gathered related to the safety of the request. This is important because additional safety concerns (such as body piercings) may come to light when the patient attends for the examination. Such factors may not have been appreciated at the time of referral. Advice may also be sought from medical physicists and other professionals to gain a more informed and balanced view of the inherent risks in certain cases. As new devices are introduced to the market on a regular basis, it is vital that we keep up to date with the latest developments in the field of MRI safety and the latest safety research. There are recommended resources listed at the end of this chapter.

DEFINITIONS USED IN MRI SAFETY

The following definitions relate to the layout of the MRI department, the personnel who work in this environment, and the devices that are permitted to be taken into the MRI magnet room or subjected to MRI scanning.

Safety zones

The American College of Radiology (ACR) has published a guidance document on MRI safety [1] that makes recommendations related to policy and practice in the field. One of the key recommendations is that the MRI facility should be zoned according to risk. The zones are represented as a floor plan in Figure 10.1. The aim of using zones is to prevent unauthorized access to areas where the high magnetic field may cause injury or death. The ACR defines the zones as follows:

- *Zone I.* "...all areas that are freely accessible to the general public. This area is typically outside the MR environment itself and is the area through which patients, healthcare personnel, and other employees of the MR site access the MR environment."
- *Zone II.* "...the interface between the publicly accessible, uncontrolled Zone I and the strictly controlled Zones III and IV. Typically, patients are greeted in Zone II and are not free to move throughout Zone II at will, but are rather under the supervision of MR personnel. It is in Zone II that the answers to MR screening questions, patient histories, medical insurance questions, etc. are typically obtained."
- *Zone III.* "...the region in which free access by unscreened non-MR personnel or ferromagnetic objects or equipment can result in serious injury or death as a result of interactions between the individuals or equipment and the MR scanner's particular environment. These interactions include, but are not limited to, those involving the MR scanner's static and time-varying magnetic fields. Zone III regions should be physically restricted from general public access by, for example, key locks, passkey locking systems, or any other reliable, physically restricting method that can differentiate between MR personnel and non-MR personnel."




Figure 10.1 MRI safety zones as recommended by the ACS Guidance Document on MR-Safe Practices 2013.

- *Zone IV.* “...the physical confines of the room within which the MR scanner is located. Zone IV should also be demarcated and clearly marked as being potentially hazardous due to the presence of very strong magnetic fields. Zone IV should be clearly marked with a red light and lighted sign stating: *The Magnet is On*. Except for resistive systems, this light and sign should be illuminated at all times and should be provided with a backup energy source to continue to remain illuminated for at least 24 h in the event of a loss of power to the site.”

Full definitions are found in the ACR Guidance Document on MR Safe Practices: 2013 (see References).

Personnel

The ACR guidance document on MRI safety identifies three levels of personnel:

- *Non-MRI personnel* include patients, visitors, or facility staff who have not undergone formal safety training (within the last 12 months) as designated by the MRI safety director.
- *Level 1 personnel* include office staff and patient aides who have passed minimal safety education to ensure their own safety as they work within Zone III.
- *Level 2 personnel* include MRI technologists, radiologists, and nursing staff who have been extensively trained in MRI safety, including issues relating to thermal loading, burns, and neuromuscular excitation from rapidly changing gradients.

Device safety

In 2005, the American Society for Testing and Materials International reviewed the terminology used to describe the MRI safety of implants and other medical devices that may pose risks in the MRI environment. These may include electromagnetic field interactions leading to heating or malfunction. Devices are now divided into three main categories: MRI safe, MRI conditional, and MRI unsafe (Figure 10.2). These are defined as follows:

- *MR safe*. “An item that poses no known hazards in all MR imaging environments. With this terminology, MR safe items are nonconducting, nonmetallic, and nonmagnetic items, such as a plastic Petri dish. An item may be determined to be MR safe by providing a scientifically based rationale rather than test data.”
- *MR conditional*. “An item that has been demonstrated to pose no known hazards in a specified MR environment with specified conditions of use. Field conditions that define the MR environment include static magnetic field strength, spatial gradient, time rate of change of the magnetic field (dB/dt), RF fields, and SAR. Additional conditions, including specific configurations of the item (e.g. the routing of leads used for a neurostimulation system), may be required. For MR conditional items, the item labeling includes results of testing sufficient to characterize the behavior of the item in the MR environment. In particular, testing for items that may be placed in the MR environment should address magnetically induced displacement force and torque, and RF heating. Other possible safety issues include but are not limited to: thermal injury, induced currents/voltages, electromagnetic compatibility, neurostimulation, acoustic noise, interaction among devices, the safe functioning of the item, and the safe operation of the MR system. Any parameter that affects the safety of the item should be listed and any condition that is known to produce an unsafe condition must be described.”
- *MR unsafe*. “An item that is known to pose hazards in all MR environments. MR unsafe items include magnetic items such as a pair of ferromagnetic scissors.”

349

Assessing the safety of MRI conditional devices can be a complex procedure. Importantly, devices that are conditionally safe at 1 T may not be safe at higher field strengths (see “Additional resources” at the end of this chapter for further information on this topic).




Figure 10.2 Device-labeling icons developed by the American Society for Testing and Materials International and recognized by the FDA.

An MRI examination might subject the patient to adverse psychological and biological effects. Some of these effects are transient and have no long-term safety implications, and others may cause serious injury or death. The five main factors that are considered to have an impact on patient safety are as follows:

- Psychological effects
- The spatially varying static magnetic field
- Electromagnetic (radiofrequency) fields
- Time-varying gradient magnetic fields
- Cryogenics.

This chapter analyzes each of these areas in turn and provides a broad overview to assist you in making an informed decision on whether it is safe to scan.

PSYCHOLOGICAL EFFECTS

The design of MRI equipment, particularly closed-bore scanners, may increase levels of anxiety and emotional distress in patients. Early closed-bore scanners required a very long magnet bore to ensure a large imaging volume of appropriate homogeneity. One undesirable trade-off with this design was that patients were required to be almost entirely confined within the narrow bore. A study performed in 1998 revealed that 14.3% of patients required sedation or anesthesia to overcome feelings of anxiety or claustrophobia during the scan procedure [2]. Notably, a majority of these patients were scheduled for a brain examination. This requires the patient's head to be positioned at the isocenter of the bore with the roof of the bore in proximity to the patient's face. The use of a head coil also adds to the feeling of enclosure. More recent research from 2008 seemed to indicate that the confined space of an MRI scanner was still an issue. In addition to claustrophobia, patients expressed fears of suffocation and fear of fainting, and exhibited symptoms of panic attack. The factors that contribute to patient anxiety are complex. The narrow bore plays a significant role. However, patients also report feelings of isolation and other factors such as scan duration, acoustic noise, and fear of a negative diagnosis are also thought to contribute to emotional distress.

On first thoughts, patient noncompliance may appear to be nothing more than a minor inconvenience; however, there may be negative consequences to diagnostic outcomes. Firstly, a nervous patient may find it difficult to keep still or comply with instructions related to breath-holding. This is likely to degrade image quality due to phase mismatching artifacts (see Chapter 8). Furthermore, a patient suffering a panic attack may terminate the procedure before any diagnostic images are acquired. This may require the patient to be referred for alternative tests such as CT. A recent comparison revealed that CT was the preferred modality for 42% of patients undergoing cardiac imaging compared with just 12% who preferred MRI. The use of other modalities such as CT or isotope studies to yield a satisfactory diagnosis results in an unnecessary dose of ionizing radiation. Examinations prolonged due to patient intolerance may place a burden on staff, decreasing patient throughput. Failed examinations also waste scanner time and other valuable resources, incurring financial penalties. Finally, rescheduling may delay diagnosis, adversely affect waiting list times, and increase the burden on administrative staff. For these reasons, it is important to have a strategy for anxious or claustrophobic patients.

There are various approaches employed to reduce emotional distress in MRI patients. Aromatherapy has been used with varying degrees of success, but it is difficult to assess whether this offers a statistically significant benefit. Scanner design has improved in recent years; magnet

bores are now considerably shorter and wider than earlier models, offering a brighter and more spacious environment for the patient. This has reduced the incidence of claustrophobia three-fold [3]. The bore is often equipped with a fan to reduce the feeling of enclosure. Open scanners also appear to be better tolerated and allow relatives to sit close to the patient for comfort during the examination. Electrophysiology tests indicate that patients are most anxious at the very beginning of the procedure when they are moved into the bore and that the level of anxiety typically decreases throughout the procedure. This suggests that any strategies for combating anxiety should focus on this stage of the process. Equipment manufacturers have recently started to place more emphasis on the patient experience as a critical factor in the success of an MRI examination. Ambient lighting, wall and ceiling panels displaying relaxing scenery, and immersive in-bore virtual reality style video imagery are employed to shift the patient's focus away from the mechanics of the procedure itself. Patients may even be invited to select a theme for their scan before entering the magnet room distracting them from any anxiety at this critical time. The in-bore video presentation may also include a count-down timer that keeps the patient informed about how long they are required to keep still and the remaining examination time. This may provide an additional motivation to complete the procedure successfully. For very young children, or patients who are very claustrophobic, sedation or anesthesia may be required. These patients require additional care and monitoring throughout the scan. Special care should be taken to ensure that monitor leads and other conductors are not left in contact with the skin of anesthetized or sedated patients.

THE SPATIALLY VARYING STATIC FIELD

The static field of the MRI scanner presents four main implications for patient safety:

- Transient biological effects
- Projectile hazards
- Torque on implanted devices
- Foreign bodies in the static field.

Transient biological effects

The main magnetic field B_0 is responsible for certain transient biological effects experienced by patients, staff, and research volunteers, particularly when ultra-high-field systems are used. The sensitivity of living organisms to external magnetic fields is well known. Around 50 different animal species appear to be sensitive to the relatively insignificant magnetic field of the planet Earth (0.5 G). Migrating birds such as pigeons are thought to have a neural substrate that acts as a magnetoreceptor for navigational purposes. There are also species of bacteria, Gram-negative prokaryotes, that are known to be magnetotactic, passively aligning themselves and actively moving along the flux lines of the magnetic north pole.

As the external magnetic field increases to a magnitude that is 160 000 times more powerful than that of the earth (8 T), biological effects become more pronounced. Patients and staff who are required to enter magnetic fields of 3–7 T report transient symptoms including a metallic taste in the mouth and vertigo. Research indicates that these effects tend to occur when workers are required to move through the flux lines of the stray static magnetic field, and when patients are moved to the isocenter of the magnet. Vertigo, causing dizziness, nausea, nystagmus (involuntary eye movement), and postural instability, is predominantly elicited at fields of 7 T and

above. Patients report a sensation of moving along a curved trajectory, despite the fact that the table movement occurs in a straight path along the z -direction of the magnet bore. The effect appears to be due to a temporary overcompensation by the vestibular system of the inner ear and can persist for a short time post examination. The mechanism of this effect has been theorized to be related to Lorentz force on the ions contained in the fluid within the semicircular canals. Any flow can cause vertigo to persist even when the patient is immobilized and stationary. The reason that this can be considered a safety hazard is that some patients appear to be more sensitive to the ultra-high field than others and, in some cases, may be advised to avoid driving or operating machinery until any symptoms subside. Normal function usually returns within 15 min from being removed from the magnetic field. Some studies have shown that the severity of induced vertigo increases in relation to the velocity at which the patient is introduced into the magnet bore. Using a slow table movement is therefore desirable when positioning the patient in an ultra-high-field system. The effects of vertigo do not seem to be influenced by RF or rapidly switching gradients.

352

In addition to the Lorentz force, there is Faraday's law of electromagnetic induction to consider. This phenomenon causes the induction of an electrical current through a conductor that is moved through an external magnetic field (see Chapter 1). This is understood to be the mechanism behind phosphenes, whereby the optic nerve acts as a conductor and stimulation of the nerve causes the appearance of flashes and optical disturbances in the patient's visual field. Metallic taste sensations are also reported by patients as they move along the static field gradient. Although the exact mechanism is still under investigation, it is likely that these sensations are caused by either nerve stimulation of the taste-buds or possibly electrolysis in the saliva. In terms of safety, the threshold for these effects is very high. Most patients, even those scanned on ultra-high-field equipment, do not report any perception of taste changes.

In summary, transient biological effects are not of great concern from a safety perspective. The International Committee on Non-Ionizing Radiation states that there is no evidence of serious adverse health effects from whole-body exposure up to 8 T. In some cases, some of the effects described by patients may be psychological in origin. Research participants assigned to a control group occasionally report symptoms when positioned inside a realistic dummy scanner having no active magnetic field.

Projectile hazards

Unlike the benign transient biological effects associated with the static field, ferromagnetic projectiles have the potential to be very dangerous. When we consider gradient magnetic fields in MRI, we often think about the time-varying gradients created by the gradient coils for spatial encoding. However, the static magnetic field also has a gradient. The static field is homogenous over the imaging volume at the isocenter of the magnet bore; however, the flux density changes dramatically over increasing distance from the end of the magnet bore. This creates a spatially varying gradient, rather than a time-varying gradient. Active shielding by the bucking coils considerably increases this effect because it is desirable to maintain the 5 G threshold as close to the magnet as possible (see Chapter 9). As a result, the magnetic field may vary from 3 T to 0.5 mT (5 G), for example, over just a few meters with the steepest part of the gradient at the end of the bore. The force required to accelerate a ferromagnetic item, such as a steel oxygen cylinder, toward the isocenter of the magnet is proportional to the product of the flux density of the magnet (field strength) and the spatial static-field gradient. In older, poorly shielded systems and some modern ultra-high-field systems, the static field gradient extends to such a distance that large ferromagnetic items are attracted from up to several meters away. This results in a greater capacity for acceleration along the gradient and a high terminal velocity.




Figure 10.3 Static magnetic field gradient of an actively shielded closed-bore MRI scanner. The values shown will vary with scanner model.

However, modern actively shielded systems contain the fringe field very efficiently. The 0.5 mT (5 G) threshold may extend little further than the end of the patient transport assembly, which is 3 m from the end of the bore or less (Figure 10.3). From a projectile-safety viewpoint, this could be considered a double-edged sword. There is less risk of items being dragged from a long distance and accelerated along a powerful magnetic field. However, there is a much smaller margin of error if a ferromagnetic wheelchair or other such item is accidentally taken into the magnet room. On unshielded systems, we might notice a moderate attractive force exerted on a ferromagnetic item (for example, coins in pockets) while still some distance from the scanner and can therefore take corrective action. On actively shielded systems, by the time it becomes apparent that the item is being attracted to the magnet, there is nothing that can be done to prevent it becoming a projectile. Figure 10.4 shows an example of this, where a patient was taken almost to the end of the magnet bore in a ferromagnetic wheelchair. The wheelchair only became attracted to the magnet when the patient stood upright, pushing the chair into an intensely higher magnetic field. The same effect applies to smaller items such as pens, scissors, coins, hairgrips, pagers, phones, and other such items that may be inadvertently introduced into the magnet room if screening is not undertaken thoroughly. To date, there have been at least two cases of very serious injury caused by oxygen cylinder projectile incidents and two known fatalities. In addition to the human cost, a high-velocity projectile can cause severe damage to the scanner hardware resulting in days or weeks of lost imaging time.



Figure 10.4 Ferromagnetic wheelchair taken into an actively shielded magnet room. No attraction was appreciated until the chair was moved close to the end of the bore.

In the event of a large projectile incident, such as those involving oxygen cylinders or wheelchairs, it is possible that the superconducting magnet requires **quenching** to remove the projectile. This is an absolute necessity if there is an immediate danger to life or limb. There have been at least two reported cases where patients or staff have been trapped by oxygen cylinders and very seriously injured. On one occasion, the patient sustained facial fractures. In the other incident, a ward assistant sustained a fractured forearm, and a radiographer a fractured pelvis. In both these cases, the patients were physically trapped by the oxygen cylinder and were unable to be released because the quench circuit failed. The patient with facial fractures had to be extracted from the rear end of the magnet bore. The ward assistant and radiographer were trapped for approximately 4 h until an engineer managed to bypass the quench mechanism and free them. In this case, the radiographer had lost consciousness due to internal bleeding. Such cases make it very clear that all new staff members should be made aware of the location of the quench button. Quench circuitry should also be regularly checked by the service engineer.

In addition to the projectile effect, there is another potential safety concern when small ferromagnetic items are trapped inside the scanner. Such items may have a deleterious effect on image quality or image geometry/distortion. For example, some types of patient immobilization devices (“sand” bags) may be filled with steel-shot. In addition to a potential projectile hazard, they are counter-effective as immobilization devices, as they are attracted to the magnet. There have been at least two incidences where such bags have ruptured and shed their contents into

the warm bore of the cryostat, permanently affecting image quality. Any small ferromagnetic items such as coins, hair grips, and safety pins that accumulate inside the MRI scanner may also adversely affect the homogeneity of the imaging volume. This causes image-quality concerns but may also become a safety issue when the examination is required to make accurate measurements of anatomy, for example, when planning image-guided surgery, RF ablation, radiotherapy, or pelvimetry. In all these examples, inaccurate measurement may have seriously unfavorable consequences for the patients, if their treatment plan is incorrect.

In summary, projectile hazards can be avoided by vigilance, the use of physical barriers, unambiguous warning signs, and a robust staff training program. If metallic items (such as halovests or external fixators) have to be taken into the magnet room, a powerful hand-held magnet helps to determine safety. Be aware that large devices or monitoring equipment may have a plastic cover, but a ferromagnetic chassis. For any patient monitoring, it is strongly recommended that only purpose-manufactured MRI certified devices are used.

Learning tip: Further thoughts on the dangers of projectiles

History shows that projectile incidents tend to share a common theme, whereby a member of staff who does not work in the MRI facility gains access to the magnet room without being observed by the MRI staff. These have included cleaners, nurses, anesthesiologists, porters, anesthetic machine engineers, patient escorts, a respiratory therapist, and workmen such as plumbers. At least seven oxygen cylinder incidents have occurred this way, in addition to scores of other less serious incidents, particularly involving cleaning equipment such as floor-polishing machines. Projectile hazards in the United States have become further complicated by police officers taking firearms into MRI magnet rooms, at least one of which was discharged on entering the magnet bore, firing a bullet into the rear wall of the magnet room. These incidents confirm the need for protective zoning of the MRI suite, a locked physical barrier between any public areas and the magnet room, and the training of all hospital staff on the dangers of MRI (see Figure 10.1). Unless a resistive MRI system is used, hospital staff must be made aware of the fact that a superconducting magnet is always activated (always on) and is equally dangerous when scanning is not in progress. It is also advisable to train nonhospital workers such as firefighters who may attempt to gain access to the magnet room out-of-hours. When an unauthorized person gains access to the magnet room, the blame usually lies with the MRI department.

Torque on implanted devices

The static field may affect implanted medical devices. These include stents, clips, neurostimulators, and cardiac pacemakers. In addition to ferromagnetic attraction, such devices may also experience **torque**. This term describes a rotational force that causes the alignment of an object to the lines of magnetic flux. Any ferromagnetic implant is likely to experience this turning force until it becomes aligned with the static field. For devices such as aneurysm clips, this could be for the entire duration of the MRI procedure. There are various factors that affect the level of risk associated with torque effects. Some implanted devices, such as orthopedic intramedullary nails or hip replacements, are anchored to the bone with screws or cement. These would be unlikely to experience torque-related movement, even if they were ferromagnetic.




Figure 10.5 An implanted device or clip may experience torque. This is a turning force that causes the long axis of a device to become aligned with the static field.

Some devices become immobilized due to fibrosis (scar tissue) over the weeks following implantation. On the other hand, some body areas, such as the brain, are composed of very soft tissue where fibrosis does not occur and where there is no bony anchor. The shape of the implant is also an important factor. In Figure 10.5, a long, narrow device (such as a surgical clip) lying at an angle to the lines of flux experiences a rotational force. It essentially behaves like a lever. The fulcrum (center of rotation) is likely to be where the clip is anchored to the blood vessel. As the clip aligns with the external field, there is a risk of damage to the neighboring anatomy due to the twisting rotational forces. To date, there has been at least one fatality, attributed to torque on an implanted aneurysm clip. In this case, the patient was aware of the clip but was allegedly misinformed about its MRI safety by the hospital that performed the clipping procedure. This highlights an important lesson for practitioners. Unless the pedigree of an implanted clip or device can be ascertained with certainty, the procedure should be canceled or postponed until further investigations have shown beyond a reasonable doubt that the examination is safe to go ahead. Documentary evidence such as the patient notes must be taken into consideration, and an examination should never be initiated on verbal evidence alone.

Other implanted devices such as cardiac pacemakers and implantable cardiac defibrillators may contain ferromagnetic components and batteries that result in an interaction with an external magnetic field. This translational attraction and rotational force may cause the device to move when introduced into a high spatial magnetic field gradient. Some modern implanted cardiac devices are less ferromagnetic, have advanced circuitry, and may be conditionally safe to scan. Under these circumstances, the device must be remotely programmed into an MRI safe mode. The examination should be conducted under the jurisdiction of a cardiology team and under strictly controlled conditions.

Nonferrous metals may also experience torque when exposed to an external static magnetic field gradient. This is due to the Lenz effect. In Chapter 1, we learned that Lenz's law states

that “the direction of an induced current is always such as to oppose the change in the circuit or the magnetic field that produces it.” When certain nonferrous metals are moved through a spatially varying magnetic field gradient, currents are generated within the metal that oppose the external magnetic field. In vitro research on older types of metal cardiac valve prosthesis have shown that this effect may impede valve function in some cases as the device is moved along the B_0 static field gradient [4].

Foreign bodies in the static field

Ferromagnetic foreign bodies are likely to be contraindicated for the same reasons as implanted devices, namely, the risk of movement due to translational and rotational forces. The risk is particularly prevalent in areas where movement of the item has the potential to cause hemorrhage or damage to delicate anatomical structures. For this reason, special attention is given to screening patients for intra-orbital foreign bodies. High-velocity ferromagnetic fragments have been known to penetrate the globe of the eye, and in at least two cases have interacted with the static field resulting in unilateral blindness. In one case, the damage was irreversible due to a vitreous hemorrhage. In the other case, the injury was treatable by the implantation of a prosthetic lens. The procedure for screening such patients is subject to a degree of controversy. Radiographs deliver a lens dose, and the combined cost of negative examinations is very costly. In the case of permanent blinding mentioned above, the patient was unaware of the presence of the metal fragment. In the second case, the patient knew of the fragment, however radiographs were obtained, but were reported as normal. A decision was made to progress with the scan, and the resulting injury pulled a 1 mm steel fragment from the lens of the eye into the anterior chamber of the globe. The resulting trauma resulted in the formation of a lenticular cataract over the months following the scan. Patients should, therefore, be screened for intraocular foreign bodies, and their verbal history should be respected as much as their tick-box responses on the screening form. All patients who report having suffered a penetrating eye injury should be considered for plain radiography if no recent imaging is available for review.

357

ELECTROMAGNETIC (RADIOFREQUENCY) FIELDS

A commonly cited misconception is the statement that “MRI does not use radiation.” It is more accurate to say that MRI uses *non-ionizing* radiation. In conventional radiography, adverse effects center around the use of ionizing radiation and the risk of deterministic and stochastic effects from increasing dose. Stochastic effects are somewhat random in nature and are defined as those that increase the probability of the patient developing a radiation-induced cancer at some time in the future. Deterministic effects are those that can be shown to have a nonrandom cause-and-effect with increasing absorbed dose. These effects include erythema or hair loss. The degree to which electromagnetic radiation causes damage to living tissue is determined by its inherent energy. All electromagnetic radiation travels at the speed of light through a vacuum, but the energy contained in the beam is determined by the wavelength. X-rays have a wavelength of fewer than 10 nanometers (nm) so no more than 20 atoms fit between the peaks of the waveform. Radiation at this frequency is energetic enough to cause damage by ionizing atoms within the patient’s tissues, causing the discharge of an electron. An MRI scanner employs RF at a frequency to match the Larmor frequency of the magnetic moments of hydrogen nuclei inside the imaging volume. This is determined by the field strength of the magnet (see Chapter 1). At 1 T, the center frequency is 42.58 MHz, and an electromagnetic wave at the corresponding

frequency has a wavelength of approximately 7 m. At 1.5 T, it is 63.87 MHz (4.69 m). The energy contained in a photon of electromagnetic radiation is directly proportional to its frequency and inversely proportional to its wavelength. As a result, the RF employed in MRI has significantly lower energy than is found in conventional radiography or CT. The energy is too low to disrupt chemical bonds or cause mutations in the DNA. As a result, the RF used in MRI is non-ionizing, but this does not mean that it is completely safe. There are three principal safety hazards associated with RF: patient heating/burns, the antenna effect, and adverse effects on implant function due to induced currents.

Heating

358

As we learned in Chapter 1, resonance involves the transfer of energy from an RF pulse to the hydrogen nuclei inside the imaging volume. Much of the energy deployed by the RF pulse is absorbed by the patient's body tissues and causes an increase in thermodynamic temperature. The mechanism behind this heating effect relates to the fact that human tissues are conductive, and eddy currents are formed in accordance with Faraday's law of electromagnetic induction. Resistivity losses cause heating. For dosimetric purposes, it is necessary to determine the rate at which the energy is absorbed. Energy is measured in joules (J), the rate at which it is absorbed requires a measurement over time in seconds (s) into a mass of tissue in kilograms (kg). The dosimetric units for the SAR are, therefore, watts per kilogram (W/kg). For this reason, it is necessary to weigh the patient accurately before the scan procedure and enter the weight into the scan computer. The system is then able to estimate the SAR for a given pulse sequence. The prescan calibration applies short bursts of RF that are needed to create a 90° and 180° flip angle. The amount of energy required for all the pulses in a pulse sequence is summed, divided by the TR, and then divided by the mass of the patient to calculate the whole-body SAR.

Some anatomical areas are more susceptible to heating compared to other areas. The lens of the eye is acutely sensitive because its lack of blood supply impedes the dispersion of heat. The germ cells of the male reproductive system can also be adversely affected by heating. Animal studies have shown transient infertility at SAR levels of around 6 W/kg. This is the level required to raise the testicular temperature by around 3.5 °C. For these reasons, different SAR thresholds are recommended for the trunk, the head, and the extremities.

The International Electrotechnical Commission (IEC) has defined three modes of operation for MRI scanners with increasing RF deposition [5]:

- *Normal operating mode.* As the name suggests, this is the mode used in routine scanning. It is defined as “the mode of operation in which none of the outputs have a value that would be anticipated to cause physiological stress to subjects.”
- *First-level controlled operating mode.* This is “the mode of operation in which one or more outputs reach a value that may cause physiological stress to subjects, which needs to be controlled by medical supervision.”
- *Second level controlled operating mode.* This is “the mode of operation of the MR equipment in which one or more outputs reach a value that can produce significant risk for subjects, for which explicit ethical approval is required according to local requirements.”

To ensure patient safety by preventing heat stress and local tissue damage, the IEC and the British Standards Institution (BSI) have published a set of standards for SAR limits. These are summarized in Table 10.1. Note that different countries apply different standards, for example in the USA the FDA have guidelines that vary slightly from the European standard.

Table 10.1 SAR limits as defined by the BSI 60601-2-33:2010+A2:2016.

Body region	SAR limits (W/kg) averaged over 6 min The SAR limits over any 10 s period shall not exceed two times the stated values					
	Volume transmit coils			Local transmit coils		
	Whole	Partial		Head	Trunk	Extremities
	Whole body	Exposed part	Head			
Normal operating mode	2	2–10	3.2	10	10	20
First-level controlled operating mode	4	4–10	3.2	20	20	40
Second-level controlled operating mode	>4	>(4–10)	>3.2	>20	>20	>40

The standard also sets limits for patient/research participant body temperature:

- *Normal mode* shall only increase core body temperature by 0.5 °C, raising the core or local body temperature to no greater than 39 °C.
- *First-level controlled mode* shall only increase core body temperature by 1 °C, raising the core or local body temperature to no greater than 40 °C.
- *Second-level controlled mode* may increase the core body temperature by more than 1 °C, raising the core or local body temperature to greater than 40 °C.

In addition, when using a local RF transmit coil, the temperature of the orbits should be limited to 1 °C.

On occasions, the MRI system notifies you that the planned acquisition cannot be initiated because the SAR limit would be exceeded. In these cases, you may decide to seek medical advice as to whether it would be safe to switch from normal mode to first-level controlled mode. Second-level controlled operating mode is a research-only mode and should never be used in a clinical examination. The SAR limits for such research projects fall under the responsibility of the local board that has approved the research. Software locks are employed to ensure that second-level control mode cannot be selected accidentally by nonauthorized personnel.

When deciding on the use of first-level controlled mode, various factors need to be considered. Some conditions may adversely affect the ability of the patient to dissipate heat. These include age (geriatric and neonate), body habitus (obesity), a pre-existing elevated body temperature (fever), or medications that alter the thermoregulatory response of the patient. Because effective monitoring is required, patients who are sedated or unable to communicate (due to age or language barriers) may also be excluded from first-level controlled mode scanning. In cases where the risk of changing to second-level mode is thought to exceed the potential benefit to the patient, it is necessary to make adjustments to the scan protocol in order to reduce RF deposition and reduce the SAR to within the limits of the normal operating mode. The interventions that may be used to reduce the RF deposition include the following:

- *Using a lower-field-strength scanner.* This is not usually a viable or practicable option. The energy deposition decreases because the center-frequency of the system is lower, and therefore the frequency of the transmitted RF is reduced. The trade-off is a reduction in SNR if all other factors remain equal. This may require a higher NSA, and this at least doubles the examination time.

- *Removing or reducing presaturation pulses.* Remove any additional RF pulses that are being used to saturate signal in structures inside or outside the FOV. The trade-off may be an increase in phase mismatching artifact that the saturation pulses were intended to reduce (see Chapter 8).
- *Increasing the TR.* Reduces the number of RF excitations per unit time. The trade-off is a change in weighting, especially in a T1-weighted sequence (see Chapter 2).
- *Decreasing the number of slices.* Reduces the number of RF pulses required per repetition and therefore reduces the number of RF pulses required per unit time. The trade-off is a reduced coverage of the anatomical area unless extra acquisitions are used, and this increases the scan time.
- *Using a different receiver coil.* A quadrature coil, for example, is more efficient than a linear coil. There is no trade-off if an appropriate coil is selected at the outset.
- *Avoiding TSE or FSE or reducing the ETL/TF in these pulse sequences.* Energy deposition increases with the square of the flip angle. 180° RF pulses deploy four times more energy than 90° RF pulses.
- *Reduce the flip angle of the RF rephasing pulses.* Some manufacturers allow us to select the flip angle of the RF rephasing pulses in spin-echo pulse sequences. This reduces the RF deposition for the same reason mentioned above (see Chapter 3). The trade-off may be a slight reduction in image quality but is unlikely to render the images nondiagnostic.

A final point to note about RF and patient heating is that burns may occur where naturally occurring loops are formed by the patient's anatomical position. If the hands are in contact with the sides of the legs, or if the heels are resting against each other, there may be a small surface area of skin in contact. This creates a "biological circuit" through which induced current may flow [6]. Because skin contact may only be over a few square centimeters, the amount of induced current is concentrated in this small area, and intense heating may result (Figure 10.6). Depending on the posture and body habitus, the same effect may occur between the knees or the calves. For this reason, MRI equipment manufacturers make recommendations about patient positioning and provide padding to insulate these body areas by breaking any naturally occurring conductive loops.

The antenna effect

A second safety concern associated with RF is known as the **antenna effect**. Any conductor acts as an antenna (radio aerial) if it happens to be of a certain critical length. Dipole antennae, such as the ones used to receive commercial radio broadcasts, are designed to be exactly half as long as the wavelength that they are intended to receive. This type of antenna is known as a **resonator** because the applied RF creates a standing wave in the dipole. This phenomenon, first described by Michael Faraday, occurs when waves are reflected back and forth inside a medium resulting in a stationary wave that oscillates over time. The wave has points where the amplitude is always at a minimum (called nodes) and points where the amplitude reaches its greatest value (antinodes). Remember that resonance is a transfer of energy (see Chapter 1). If the conductor is at the resonant frequency of the wave, the antinodes are found at the ends of the wire, and because this is the point of greatest energy, heating may occur. In vitro testing of pacemaker wires in an MRI scanner revealed a temperature increase at the tip of the wire, raising its temperature to a maximum of 80 °C. This temperature is sufficiently high to cause burns or tissue damage [7]. Dipole heating of conductive wires is shown in Figure 10.7.

To prevent burns, it is of key importance to ensure that conductive wires are kept away from the patient's skin surface, particularly if they are sedated, unconscious, or anesthetized. Such




Figure 10.6 Anatomy may form “biological circuits” through which an induced current may flow. Contact areas such as the hands or heels may concentrate the current flow through a small area causing burns.

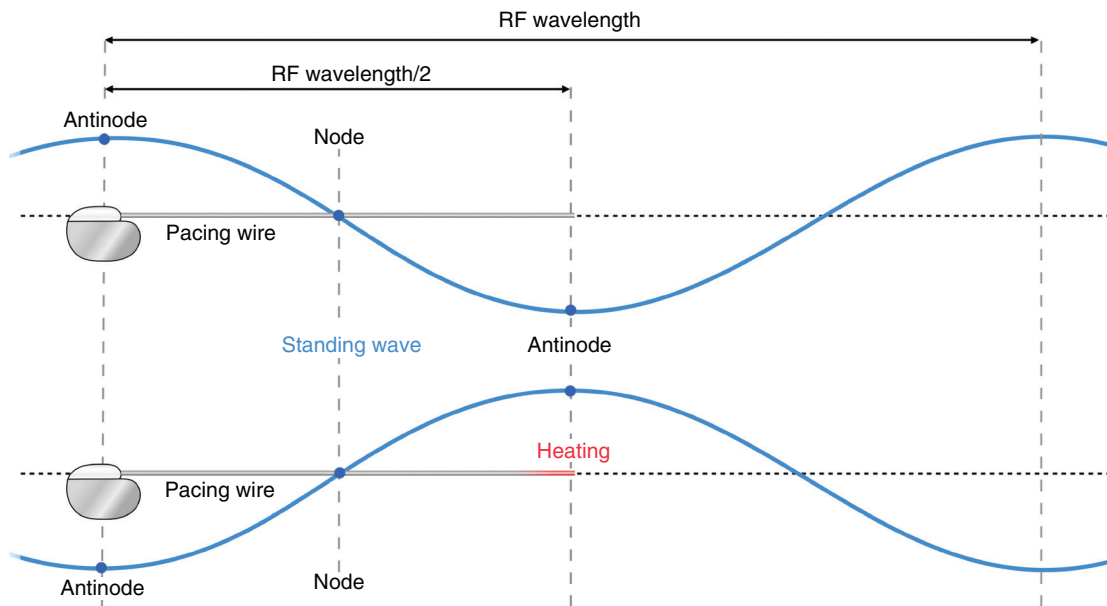


Figure 10.7 Dipole heating in a conductive wire such as a pacemaker lead.

patients are unable to report burning sensations, and there have been tragic consequences where serious burns have occurred because they went unnoticed until the end of the examination. In one case, an MRI unsafe pulse oximeter wristband experienced inductive heating during an MRI examination under anesthesia, resulting in fourth-degree burns to a five-week-old baby girl. The injury resulted in nonviability of the remaining tissue leading to amputation of the entire hand and forearm [8]. This case highlights the fact that seemingly innocuous devices may be extremely dangerous in the MRI environment.

Learning tip: Wavelength and the safety of conductors

362

Sharp-eyed readers may be wondering how a wavelength of up to 7 m could affect conductors such as pacing wires, which are unlikely to be at the required dipole length of 3.5 m. The reason is that the wavelength of electromagnetic radiation is usually stated as it travels at the speed of light through a vacuum. This wavelength changes if the medium changes, for example, when passing through biological tissue. This is because human tissue is both a conductor and an insulator. It contains charges that move but also fixed molecules with a magnetic dipole. In the presence of an external electromagnetic field, these dipoles align in opposition to the external field and, in doing so, reduce its effect. The degree to which different molecules (tissues) reduce the RF field is defined as the **relative permittivity** of the medium in question. The relative permittivity of a perfect vacuum is exactly 1. Various biological tissues have different relative permittivities, all of which are greater than 1. At 63 MHz water has a value of 81 and blood of 67. The blood value is the important one because pacing wires are necessarily immersed in this medium within the vessel. A relative permittivity of 67 reduces the wavelength of the RF to approximately 58 cm, and therefore the critical dipole length of a conductor is approximately 29 cm.

Induced currents in implanted devices

In addition to heating effects in the leads of devices such as cardiac pacemakers, there are other effects reported in the literature including temporary or permanent changes to the function of such devices. The literature shows that complications associated with the inadvertent scanning of pacemakers are quite rare. Unfortunately, on the occasions that complications do occur, the result can be life-threatening. To date, there have been at least 13 reported fatalities attributed to pacemaker malfunction, making this, by far, the most common cause of death related to the MRI procedure. The physiological mechanism behind pacemaker fatality is usually asystole or ventricular fibrillation. The technical reasons relate to the effect of RF and/or magnetic field gradients on the internal sensing circuitry of the device resulting in inappropriate cardiac stimulation. In vivo studies conducted on animals and humans show that overly rapid pacing causes a significant drop in arterial blood pressure. Other reported effects include disruption of the internal power supply and activation of the magnetic reed switch resulting in asynchronous pacing. In recent years, pacemaker manufacturers have addressed many of the MRI incompatibility concerns associated with these devices:

- Leads have been modified to reduce tip-heating.
- Internal sensing circuitry has been improved to reduce inappropriate cardiac stimulation.
- The number of ferromagnetic components has been reduced.

- Filters have been employed to reduce disruption to the internal power supply.
- The reed switch has been replaced with a more complex analog transducer.

This has resulted in a range of devices that are conditionally safe to scan.

TIME-VARYING GRADIENT MAGNETIC FIELDS

Nerve stimulation

The final type of magnetic field that may cause safety issues in MRI is the time-varying gradient magnetic field that is applied during the acquisition. Faraday's law of electromagnetic induction predicts that an external magnetic field will interact with a conductor to produce an emf (see Chapter 1). One of the earliest electrical generators, invented by Antoine Hippolyte Pixii in 1832, capitalized on this phenomenon featuring a hand-cranked magnet positioned in proximity to a wire solenoid. This device is known as an alternator. The alternation of the applied magnetic field induces an alternating current in the conductor. This is a good analogy for the action of magnetic field gradients applied during an MRI acquisition. The gradients are applied in all three orthogonal directions and the polarity of the gradient field may be rapidly alternated many times per second. This creates the potential (no pun intended) for the induction of a current in the body tissues. The minimum current that is required to depolarize the cell membrane of a neuron or muscle cell is known as the **rheobase**. The parameter that determines the excitability of biological tissue is known as the **chronaxie value**. This is defined as the minimum pulse duration over which an electric current at twice the value of the rheobase needs to be applied to stimulate a muscle or neuron. Put simply, if the induced current exceeds a certain threshold, over a certain time interval nerve stimulation occurs. Research volunteers have reported this effect as causing tapping, tingling, and throbbing sensations in the extremities.

From a safety viewpoint, nerve stimulation could pose a risk if the cardiac muscle or the diaphragm were affected. The chronaxie value is important, as this is the value used in clinical cardiac pacing to determine the pulse duration needed to depolarize myocardial muscle. In this context, the term “pulse sequence” takes on a whole new meaning, as it may start acting as a cardiac pacemaker if the induced current and chronaxie value reach the critical level. When using a standard diagnostic MRI scanner, inappropriate cardiac pacing is unlikely, as the threshold for cardiac stimulation is considerably greater than that required for peripheral nerve stimulation [9]. The patient would be pressing the alarm button somewhat vigorously under these circumstances! Current safety guidelines state that in normal operating mode, the gradient system should operate at a level that does not exceed 80% of the mean threshold for peripheral nerve stimulation and should never exceed 100%. Research modes may increase the threshold to 120%, a level above which discomfort may occur. Canine research studies show that respiratory stimulation would require a threshold of 300% and cardiac stimulation a threshold of 900% of the mean threshold for peripheral nerve stimulation. This is unlikely using a diagnostic MRI scanner, as the slew rate of the gradient coil is restricted by the manufacturer to safe levels (typically up to 200 T/m/s).

Effect of time-varying gradients on implanted devices

Time-varying gradients may have an adverse effect on implanted electronic medical devices due to interference with the circuitry. They induce a current in metallic objects such as cervical fixation devices. The induced current may, in turn, induce a magnetic moment in the metalwork of the device causing an interaction with the static field. This is reported to create uncomfortably high-frequency vibrations experienced by the patient when the gradients are active.

Acoustic noise

Another safety concern arising from the gradient coil is the acoustic noise generated during the MRI examination. Lorentz forces (see Chapter 9) are responsible for this noise, whereby the current flowing through the gradient coil interacts with the static magnetic field. Although the gradient coil is securely anchored in place, the stresses imposed by the Lorentz forces are propagated through the structure of the MRI scanner as vibrations and are transferred to the surrounding air as longitudinal sound waves. Sound waves are pressure waves, the wavelength of which is determined by compressions and rarefactions of the medium through which it travels. These fluctuations in air pressure are detected by the tympanic membrane of the ear and perceived as sound. Sound intensity is measured on a decibel (dB) scale. The logarithmic nature of this scale means that a 3 dB increase in sound pressure level represents a doubling of the loudness (intensity) of the sound. Table 10.2 gives some real-world examples of loudness and the thresholds above which hearing damage may become a concern.

Noise-induced hearing loss is second only to age-related hearing loss in human beings. As a rule, damage to hearing increases with both amplitude and duration of the sound. A high-amplitude sound of short duration such as an explosion or a longer duration sound such as a rock concert may inflict the same amount of damage. The Occupational Safety and Health Administration (OSHA) recommends that exposure to noise should be maintained below a level equivalent to 85 dB over an 8 h period to minimize the risk of hearing loss. At 100 dB, this drops to just 15 min/day. Table 10.2 shows that MRI scans may reach a level of 100 dB or higher. However, even if this acoustic level were maintained for longer than 15 min, the sound is not continuous (there are brief pauses between acquisitions), and the patient does not undergo the procedure

Table 10.2 Examples of acoustic noise (approximate loudness in dB).

Sound	Upper threshold loudness (dB)
Gunshot/explosions	170
Jet aircraft engine (takeoff)	140
Threshold above which sound may become painful or cause tympanic rupture (140 dB)	
Rock concert (front row)	125
3 T MRI scanner bore (echo-planar sequence)	115
Night club (dancefloor)	110
1.5 T MRI scanner bore, aircraft cabin, orchestra	100
MP3 player at high volume (headphones)	95
Motorcycle engine	90
Threshold above which chronic sound exposure may cause hearing loss (85 dB)	
Motor car	70
Normal conversation	60
Whisper	20
Threshold below which sounds become barely audible to humans	
Breathing	10

on a regular basis. For this reason, it is unlikely that permanent hearing damage is caused by a single MRI procedure. The use of hearing protection is advisable, however, because levels above 75 dB can become annoying and unpleasant for many people. The International Commission on Non-Ionizing Radiation Protection (ICNIRP) recommends the use of ear protection for noise levels above 80 dB.

In summary, the time-varying gradients employed during an MRI acquisition may induce currents in biological tissue, leading to peripheral nerve stimulation above a critical threshold. However, this is unlikely to occur at the gradient amplitudes and slew rates used in routine scanning.

Scan tip:

Practical tips for protecting your patient's hearing during a scan

365

Occasionally, a patient may express concern about the level of noise, for example, if their ear defenders become loose or fall out during the procedure. Under these circumstances, reassure them that the sound levels to which they have been exposed are largely equivalent to those that they would have experienced on a jet aircraft flight or even a loud orchestral performance. If a tight head-coil design prevents the use of over-the-ear defenders, ear plugs should be used. When using ear plugs, it is important that these are positioned to fill the entire ear canal. It is worth taking some time to ensure that the patient is taught how to fit these correctly. Some manufacturers are now offering sequences that are "silent" in that they do not appreciably increase the normal ambient noise level of the magnet room. This reduces gradient stresses by the use of a spiral radial k -space trajectory (see Chapter 6). This technique employs gradient steps that are relatively small compared to traditional imaging sequences. If your scanner is equipped with such sequences, they are recommended for pediatric patients, geriatric patients, and anyone of a nervous disposition who may become noncooperative if alarmed by loud noise.

CRYOGENS

As we discussed in Chapter 9, the cryogen used to facilitate superconductivity in modern MRI systems is liquid helium. There are four primary safety concerns associated with its use:

- Thermal sensitivity
- Asphyxia
- Quench
- Explosion.

Thermal sensitivity

Liquid helium has a boiling point of just 4 K above absolute zero (see Chapter 9). This extremely low temperature poses a hazard when handling liquid helium as splashes may cause damage to sensitive biological tissues, especially the cornea of the eye. Serious skin damage has been reported in at least one case where liquid helium entered the gloves of an MRI engineer requiring skin grafts. For this reason, any on-site helium should be handled and stored correctly and only in dewars designed

for the purpose. A helium **dewar** is a highly pressurized double container having an inner chamber separated from the outer container by an evacuated annulus. To maintain safe storage, dewars should be kept upright. These containers are fitted with a sophisticated pressure control system and have several release valves. These should only be operated by trained service engineers in the event of a helium fill, and then only when wearing appropriate personal protective equipment and full-face protection. Liquid helium must never be handled by nontrained staff and should never be poured out of a container. Leaking vessels may be identified by a plume of water vapor and should be periodically checked for frost formation. Such signs must be reported immediately to the manufacturer and, in the event of a leak, the vessel moved to an outdoor location remote from public areas. If a patient or member of staff should come into contact with liquid helium, urgent medical advice must be sought. If there is injury to the hands, they should be placed in a bowl of water at body temperature as a first-aid measure until medical help arrives.

Asphyxia

Liquid helium is seemingly innocuous. It is odorless, tasteless, nontoxic, inert, and noncorrosive; however, it displaces oxygen and can cause asphyxiation. Normal atmospheric air has an oxygen content of approximately 21%, the remainder mostly consisting of nitrogen. The OSHA defines an optimal oxygen level for normal activity to lie in the range between 19.5% and 23.5%. There could be a danger to personnel if helium gas reduces the oxygen concentration to a level below 19.5%. Magnet rooms have an oxygen sensor that is calibrated to sound an alarm if the oxygen level drops below this threshold. Oxygen concentrations below 19.5% cause impaired perception and judgment, which, environmental safety experts say, are responsible for preventing individuals from self-rescue. If the oxygen concentration falls below 8%, an individual in the magnet room is likely to lose consciousness. At concentrations less than 6%, loss of consciousness occurs in as little as 10 s and heart failure and death within 2–4 min. Death caused by asphyxia from inhaling the helium from party balloons is well recognized. The National Inhalant Prevention Coalition reported nine such deaths in 1 year in the state of Florida. The last case in the UK involved the death of a 13-year-old child in 2010 who suffered a cardiac arrest after inhaling the contents of a helium balloon. An MRI engineer was also killed in 2000 by a leaking nitrogen pipeline during an equipment installation. Nitrogen is no longer used as an MRI cryogen but exhibits the same dangers as helium in that it displaces oxygen in a confined space.

Quench

To avoid the risk of asphyxiation, a wide-bore cryogen vent (quench pipe) is designed to evacuate any gaseous helium from the building in the event of a quench. A fully filled MRI cryostat may contain over 1600 liters (L) of liquid helium. The expansion ratio of helium is 1 : 754. This means that every liter of liquid evaporates very rapidly into 754 l of gas. The cryogen vent is therefore required to conduct over 1 000 000 l of gas in a short space of time. Occasionally, the quench pipe may rupture under the pressure, and helium gas may be vented into the magnet room or the MRI building. This is thought to be caused by blockages in the vent, or rainwater collecting inside the pipe and forming a solid plug of ice when hit by the freezing helium gas. Under such circumstances, the building must be completely and quickly evacuated. Remember that helium gas is lighter than air and may find its way to the highest rooms in a multistory building. If the patient requires evacuation from the magnet room, MRI staff must be mindful of the fact that low oxygen levels may cause rapid loss of consciousness both for the patient and

the rescuer. Extreme caution should be exercised, particularly when working alone. In the nitrogen asphyxia case described above, several people who ran into the trailer to attempt to rescue the victim were also temporarily hospitalized due to the effects of asphyxia. Remember that the golden rule of first aid is to ensure your own safety first.

One final complication that may arise during a quench pipe failure is a sudden increase in atmospheric pressure in the magnet room. A relatively minor increase in air pressure exerts a significant amount of force on an inwardly opening magnet room door. This may prevent access to the magnet room until the pressure difference is equalized on both sides of the door [10]. In an emergency, this is achieved more quickly by breaking the control room window. Some magnet rooms may have pressure release panels built into the wall. Even when the quench pipe operates correctly, it is necessary to site the outside discharge point at least 8 m away from anywhere with public access to avoid accidental exposure to the freezing effects of the escaping gas.

Explosion

The cryostat is a pressurized container. As we saw in Chapter 9, the helium boil-off is controlled by the use of vacuum insulation. The resulting boil-off rate is therefore very slow. During a quench, a valve (known as a burst disk) opens to allow the helium gas to be vented as described in the previous section. A potentially dangerous issue may arise if this valve fails, particularly if room temperature air is allowed into the vacuum chamber, increasing the helium boil-off rate. Although it is likely that some events go unreported, there have been several occasions on which MRI cryostats have exploded due to a build-up of pressure inside the cryogen chamber that failed to activate the bursting disk.

Probably the most well-known example was caught on film by a local news team when covering a story at a regional medical center in the United States. The report was intended to be about the delivery of a new MRI unit, but the old cryostat was sitting on a trailer waiting for collection in full camera shot when it exploded. Other such incidents have occurred including a decommissioned unit exploding in Kennesaw, Atlanta, and more recently at a veterinary hospital in New Jersey. The Kennesaw incident was thought to be due to a faulty gauge that failed to indicate that 500 l of liquid nitrogen remained inside the cryostat. The explosion occurred as the machine was removed from the magnet room. The entire end of the cryostat was blown off causing shrapnel injuries to two workers and reportedly knocking a 3 m hole through the wall of the building. In the New Jersey accident, three engineers were injured under similar circumstances, one very seriously. Although the building was left structurally intact, over one million dollars worth of damage was reported. Such explosions are unlikely to occur in the clinical setting, but these events highlight the need for regular service inspections and meticulous decommissioning procedures.

In summary, cryogenics tend to remain out-of-sight and out-of-mind, and although helium appears to be chemically innocuous, there are serious risks when the liquid or the gas comes into contact with humans. There should be a policy in place for any adverse events such as a quench, particularly if there is a failure of the cryogen venting system.

SAFETY TIPS

- Never assume that any member of staff has prior knowledge about the dangers of entering a magnet room. Screen them with the same rigor as used for the patients, even if they have been before.
- Avoid using a digital code lock for Zone III. Codes can easily be shared with staff from other departments who have no knowledge of MRI safety.

- Never assume that a patient who has been scanned elsewhere yesterday is safe to scan at your facility today. Your field strength may be greater; your gradients may be stronger and faster; the patient may have discovered that they are pregnant this morning.
- Ensure that patients are asked about contraindications before their appointment. This prevents wasted time and aggravation when patients wait a long time for their appointment, only to find that they cannot be scanned.
- Provide plenty of information to the patient in advance of the procedure. This may help to alleviate any fears caused by scare-stories and misinformation from other sources.
- If a patient appears to be a “poor historian” giving tangential answers to screening questions, take extra care, check history with relatives, and get them fully changed into a gown.
- If there is recent imaging available, such as chest and skull X-rays, these might help rule out the presence of pacemakers or aneurysm clips.
- Place MRI unsafe stickers on any large ferromagnetic objects kept in the MRI department. Always use a nonferrous alternative if available. For example, if you keep a set of steps to allow patients to climb onto trolleys in the waiting room, make sure that you choose nonferrous steps. Even if they are not intended to go into the magnet room, they may accidentally be taken in. Choose nonferrous fire extinguishers for the same reason.
- Always verbally check for contraindications in addition to the written screening form. Patients occasionally identify potential hazards when asked verbally that they did not indicate on the form.
- Never allow yourself to be bullied or persuaded into scanning a patient that you have concerns about. If in doubt, postpone the examination.

In summary, MRI safety is a very important topic. The noninvasive nature of MRI and the benefits of nonionizing radiation must be balanced against the potential for serious injury due to projectile hazards, implant malfunction, RF-related burns, and cryogen hazards. The number of reported MRI-related fatalities is thankfully very low, fewer than one per year. It is important to remember, however, that this is due to the constant vigilance of the MRI department staff and responsible referring clinicians.



For questions and answers on this topic please visit the supporting companion website for this book: www.wiley.com/go/westbrook/mriinpractice

ADDITIONAL RESOURCES

- Shellock, F.G., 2018. *MRI Safety*, www.mrisafety.com
- Shellock, F.G., 2017. *Reference Manual for Magnetic Resonance Safety, Implants and Devices*. Los Angeles: Biomedical Research Publishing Group.
- Shellock, F.G. and Crues, J.V. 2013, *MRI Bioeffects Safety and Patient Management*. Los Angeles: Biomedical Research Publishing Group.

For more detailed information on safety legislation dosimetry and statistics related to equipment testing, publications are available from the following MRI safety organizations:

- American Society for Testing and Materials (ASTM). <https://www.astm.org>, particularly the ASTM document *Standard Practice for Marking Medical Devices and Other Items for Safety in the Magnetic Resonance Environment* (ASTM F2503 available from <http://www.astm.org/cgi-bin/resolver.cgi?F2503>)

- American College of Radiology (ACR). <https://www.acr.org>
- Food and Drug Administration (FDA). www.fda.gov
- Institute for MRI Safety, Education and Research. www.imrser.org
- International Electrotechnical Commission (IEC). www.iec.ch
- National Electrical Equipment Manufacturer's Association (NEMA). www.nema.org.

References

1. Kanal, E., Barkovich, A.J., Bell, C. et al. (2013). ACR guidance document on, MR safe practices. *Journal of Magnetic Resonance Imaging* 37 (3): 501–530.
2. Murphy, K.J. and Brunberg, J.A. (1997). Adult claustrophobia, anxiety and sedation in MRI. *Magnetic Resonance Imaging* 15 (1): 51–54.
3. Dewey, M., Schink, T., and Dewey, C.F. (2007). Claustrophobia during magnetic resonance imaging: cohort study in over 55,000 patients, *Journal of Magnetic Resonance Imaging* 26: 1322–1327.
4. Edwards, M.B., Mclean, J., Solomondis, S. et al. (2015). In vitro assessment of the Lenz effect on heart valve prostheses at 1.5 T. *Journal of Magnetic Resonance Imaging* 41 (1): 74–82.
5. Keevil, S. (2016). Safety in magnetic resonance imaging. *Medical Physics International Journal* 4 (1): 28.
6. Calamante, F., Faulkner, W.H., Itterman, B. et al. (2015). MR system operator: recommended minimum requirements for performing MRI in human subjects in a research setting. *Journal of Magnetic Resonance Imaging* 41 (4): 899–902.
7. Dempsey, M.F., Condon, B., and Hadley, D.M. (2001). Investigation of the factors responsible for burns during MRI. *Journal of Magnetic Resonance Imaging* 13 (4): 627–631.
8. Haik, J., Daniel, S., Tessone, A. et al. (2009). MRI induced fourth-degree burn in an extremity, leading to amputation. *Burns* 35 (2): 294–296.
9. Bourland, J.D., Nyenhuis, J.A., and Schaefer, D.J. (1999). Physiologic effects of intense MRI gradient fields. *Neuroimaging Clinics of North America* 9: 363–377.
10. Gilk, T. and Kanal, M. (2006). Integrating MRI facility design and risk management. *Journal of Healthcare Risk Management* 26 (3): 29–37.

Glossary

A

Acceleration factor	see reduction factor .
Acquisition window	see sampling time .
Active shielding	uses additional superconducting coils located at each end of the main magnet inside the cryostat to shield the system.
Active shimming	additional solenoid magnets to adjust field homogeneity.
Actual TE	time between the echo and the next RF pulse in reverse-echo gradient-echo pulse sequences.
Algorithm	a set of steps used by a computer to solve a problem.
Aliasing	artifact produced when data are undersampled.
Alignment	when the magnetic moments of hydrogen nuclei have the same orientation as the main static magnetic field.
Alnico	alloy used to make permanent magnets.
Amplitude modulation	modulation or simplification of waveforms according to their amplitude.
Analog-to-digital conversion	changing a waveform into binary numbers via digitization.
Angular momentum	spin of MR active nuclei, which depends on the balance between the number of protons and neutrons in the nucleus.
Anisotropic	voxels that are not the same dimension in all three planes.
Antenna effect	causes heating in conductive wires of a critical length.
Antialiasing	also called no phase wrap and fold-over suppression . Oversamples along the phase encoding axis by increasing the number of phase-encoding steps.

Antiparallel alignment	alignment of magnetic moments in the opposite direction to B_0 .
Apparent diffusion coefficient	the net displacement of molecules across an area of tissue per second.
Asymmetric FOV	see rectangular FOV .
Atom	a tiny element that is the basis for all things.
Atomic number	sum of protons in the nucleus – this number gives an atom its chemical identity.
Atomic weight	see mass number .
B	
B_0	main magnetic field measured in teslas.
b factor/value	strength and duration of the gradients in diffusion-weighted imaging.
Balanced gradient-echo	pulse sequence characterized by a balanced gradient scheme that corrects for flow. Alternating RF pulses with varying phase angles prevent saturation and maintain the steady state.
Bandwidth	range of frequencies. See also transmit bandwidth and receive bandwidth .
Bas-relief artifact	see chemical shift .
Bipolar gradient	gradient with two lobes, positive and negative. Often used to dephase and then rephase magnetic moments to produce a gradient-echo.
Birdcage design	symmetrical birdcage arrangement of several electrically conductive elements in a transceiver coil.
Black-blood imaging	acquisitions in which blood vessels are hypointense.
Blipping	used in single shot to step down through phase-encoding steps.
Blood oxygen level dependent	a functional MRI technique that uses the differences in magnetic susceptibility between oxyhemoglobin and deoxyhemoglobin to image areas of activated cerebral cortex.
Blurring	result of $T2^*$ decay during the course of the EPI acquisition.
Boltzmann equation	equation that predicts the number of spins in each energy level.
Bright blood imaging	acquisitions in which blood vessels are hyperintense.
Bucking coils	used in active shielding to constrain the 5 G footprint of the fringe field to within a short distance from the scanner.
C	
Carrier frequency	frequency of spins located in the middle of a slice. RF is transmitted at this frequency with a small range on either side – see transmit bandwidth .

Cartesian filling	linear filling of k -space.
Central lines	area of k -space filled with the shallowest phase encoding slopes.
Chemical shift	artifact along the frequency axis caused by the frequency difference between fat and water.
Chronaxie value	parameter that determines the excitability of biological tissue.
Classical theory	Newtonian physics that uses variables such as mass, velocity and force to explain how the universe works.
Coarse matrix	a matrix that results in a low number of pixels in the FOV.
Cocurrent flow	flow in the same direction as slice excitation.
Coherent	see in phase .
Coherent gradient-echo	pulse sequence characterized by a rewinder gradient. All transverse magnetization (FID and stimulated echo) is rephased so all types of weighting are possible.
Concatenation	see interleaving .
Conjugate symmetry	symmetry of data in k -space.
Contrast-to-noise ratio (CNR)	difference in SNR between two points.
Counter-current flow	flow in the opposite direction to slice excitation.
Cross-excitation	energy given to nuclei in adjacent slices by the RF pulse.
Cross-talk	occurs in multiangle imaging when slices physically intersect. Is caused by the same principles as cross-excitation .
Cryogenics	substances used to supercool the coils of wire in a superconducting magnet.
D	
Data point	point in k -space that contains digitized information from spatial encoding.
Dead time	time in a pulse sequence where T1 recovery occurs in spins within a slice while waiting for the next RF excitation pulse.
Dewar	a highly pressurized double container for storing and transporting helium.
Diamagnetic compounds	exhibit a weak repulsion to an external magnetic field.
Diffusion	movement of molecules due to random thermal motion.
Diffusion weighted imaging	technique that produces images whose contrast is due to the differences in ADC between tissues.
Digital sampling rate or frequency	rate at which a waveform is sampled per second.
Dixon technique	uses the fact that magnetic moments are usually out of phase because of their chemical shift.

Double IR prep	sequence in which two 180° pulses are used to saturate blood in black-blood imaging.
Driven equilibrium	pulse sequence that achieves a very high signal intensity from water even when using short TRs.
DTPA	diethylene triaminepentaacetic acid, a gadolinium chelate.
Dummy pulses	see preparatory pulses .
Duty cycle	percentage of time that a gradient spends at maximum amplitude.
E	
Echo	signal detected in the receiver coil at time TE. It is made up of rephased magnetic moments in the transverse plane.
Echo-planar imaging	single-shot or multishot acquisition that fills <i>k</i> -space with data from gradient-echoes.
Echo-spacing	time between each echo in the echo train in TSE or FSE.
Echo time (TE)	time in millisecond from the application of the RF pulse to the peak of the signal induced in the coil. TE determines how much decay of transverse magnetization is allowed to occur.
Echo train	series of 180° RF rephasing pulses and echoes in a fast or turbo spin-echo pulse sequence.
Echo train length	number of 180° RF rephasing pulses/echoes/phase encodings per TR in fast or turbo spin-echo.
Effective TE	TE selected in TSE. Used to weight the image as accurately as possible given that echoes with different TEs are used to determine image weighting. (Note: Effective TE is also used in reverse-echo gradient-echo where it is the time between the gradient-echo and the previous RF pulse.)
Electro-motive force (emf)	energy available from a unit of charge traveling once around a loop of wire.
Electrons	particles that spin around the nucleus.
Entry-slice phenomena	contrast difference between flowing nuclei relative to the stationary nuclei because they enter a slice fresh. More likely to happen in the first slice in a stack of slices.
Ernst angle	flip angle that generates the highest signal intensity in a tissue with a given T1 recovery time and in a given TR.
Even-echo rephasing	technique that uses two echoes of a multiple of the same TE to reduce flow artifact.
Excitation	application of an RF pulse that causes resonance – energy giving.
Extrinsic contrast parameters	the parameters that can be changed by the operator.

F

Faraday's law	law that describes the connection of a magnetic and electric field with motion.
Fast Fourier transform (FFT)	mathematical conversion of frequency/time domain to frequency/amplitude.
Fast spin-echo (FSE)	see turbo spin-echo (TSE) .
Fat saturation	technique that nulls signal from fat by applying an RF pulse at the frequency of fat to the imaging volume before slice excitation.
Ferromagnetic compounds	have a large positive magnetic susceptibility and are powerfully attracted to an external magnetic field.
Field of view (FOV)	area of anatomy covered in an image.
Filling factor	determines how well the shape and size of the coil are matched to the anatomical region under investigation. Needed to ensure good signal quality from a coil.
Fine matrix	matrix that results in a large number of pixels in the FOV.
First-order motion compensation	gradient moment nulling.
FLAIR	fluid-attenuated inversion recovery. A pulse sequence that suppresses signal from CSF.
Flip angle	angle of the NMV to B_0 .
Flow compensation	see gradient moment nulling .
Flow phenomena	artifacts produced by flowing nuclei.
Flow-related enhancement	decrease entry-slice phenomenon due to the decrease in velocity of flow.
Fold-over suppression	see antialiasing .
Fractional averaging	see partial averaging .
Fractional echo	see partial echo .
Free induction decay (FID)	loss of signal due to relaxation.
Frequency	rate of change of phase – the speed of rotation.
Frequency encoding	locating a signal according to its frequency.
Frequency matrix	number of pixels in the frequency axis of the image.
Frequency modulation	modulation or simplification of waveforms according to their frequencies.
Frequency resolution	refers to the resolution of spatial frequencies in k -space.
Frequency wrap	aliasing along the frequency encoding axis.
Fringe field	stray magnetic field outside the bore of the magnet.
Fully saturated	when the NMV is pushed to a full 180° .

Functional MRI (fMRI)	allows MRI to be used to assess function and physiology based on the level of oxyhemoglobin in cerebral flow.
G	
Gating	technique of reducing phase mismapping from periodic motion caused by respiration, cardiac, and pulsatile flow motion.
Gauss	unit of magnetic flux density – usually for small changes in field strength. See also teslas .
Ghosting artifact	motion artifact in the phase axis.
Gibbs artifact	line of low signal in the cervical cord due to truncation.
Gradient	a linear change of the magnetic field strength.
Gradient coils	coils of wire that alter the magnetic field strength in a linear fashion when a current is passed through them.
Gradient amplifiers	supply power to the gradient coils.
Gradient amplitude	steepness of a gradient. Determined by how much current passes through it. How steep a gradient slope can be achieved when the gradient coil element is activated.
Gradient-echo	echo produced as a result of gradient rephasing.
Gradient-echo-EPI (GE-EPI)	gradient-echo sequence with EPI readout.
Gradient-echo pulse sequence	sequence that uses a gradient to regenerate an echo.
Gradient moment nulling	system of gradients that compensates for intravoxel dephasing (also called gradient moment rephasing or flow compensation).
Gradient moment rephasing	see gradient moment nulling .
Gradient spoiling	use of gradients to dephase magnetic moments – the opposite of rewinding.
GRASE	gradient-echo and spin-echo.
Gyromagnetic ratio	the precessional frequency of an element at 1.0 T.
H	
Hahn echoes	echoes formed when any two 90° RF pulses are used in steady state sequences.
Half Fourier	see partial averaging .
Half FOV ghosts	seen in EPI where a ghost of the real image appears shifted in the phase direction by one half of the FOV.
High-energy nuclei	nuclei that have enough energy to align their magnetic moments in opposition to the main B_0 field.
High-frequency aliasing	frequencies outside the FOV are higher than the Nyquist frequency and are mapped to a lower frequency causing aliasing.

High-velocity signal loss	increase in time-of-flight phenomenon due to an increase in the velocity of flow.
Hybrid sequences	combination of fast spin-echo and EPI sequences where a series of gradient-echoes are interspersed with spin-echoes.
Hydrogen	the most abundant atom in the body.
I	
Image matrix	number of pixels in the frequency and phase axes of the image.
Incoherent	see out of phase .
Incoherent gradient-echo	pulse sequence that spoils the residual transverse magnetization so that the FID makes the most contribution to image contrast. Used for T1 weighting.
Inflow effect	see entry-slice phenomenon .
Inflow MRA	see time-of-flight angiography
Inhomogeneities	areas where the magnetic field strength is not exactly the same as the main field strength – magnetic field unevenness.
In phase	magnetic moments that are in the same place on the precessional path around B_0 at any given time.
Interleaving	a method of acquiring data from alternate slices and dividing the sequence into two acquisitions – no slice gap is required.
Inter-view mismapping	see phase mismapping .
Intravoxel dephasing	phase difference between flow and stationary nuclei in a voxel.
Intrinsic contrast parameters	parameters that cannot be changed because they are inherent to the body's tissues.
Ions	atoms with excess or deficit of electrons.
Ionization	the process of knocking out electrons from the atom, thereby causing an electrical imbalance.
Inversion recovery	pulse sequence that begins with a 180° inversion pulse to suppress the signal from certain tissues.
Isocenter	see magnetic isocenter .
Isotopes	atoms with an odd mass number.
Isotropic	voxels that are the same dimension in all three planes.
J	
J coupling	causes an increase in the T2 decay time of fat when multiple RF pulses are applied in TSE or FSE.

K

***k*-space** an area in the array processor where data on spatial frequencies are stored.

L

Larmor equation equation that determines the precessional frequency of an MR active nucleus.

Larmor frequency see **precessional frequency**.

Lenz's law law that states that induced emf is in a direction that opposes the change in the magnetic field that causes it.

Lobes denote gradient amplitude, duration, and polarity in a pulse sequence diagram.

Longitudinal plane axis parallel to B_0 .

Low-energy nuclei nuclei that do not have enough energy to align their magnetic moments in same direction as the main B_0 field.

M

Magic angle artifact seen in collagen structures at certain TEs and when anatomy lies at 55° to B_0 .

Magnetic homogeneity evenness of the magnetic field.

Magnetic isocenter center of the bore of the magnet in all planes.

Magnetic moment denotes the direction of the north/south axis of a magnet and the amplitude of the magnetic field.

Magnetic susceptibility ability of a substance to become magnetized and whether it is attracted or repelled by the external field.

Magnetization transfer contrast technique used to suppress background tissue and increase the CNR.

Magnitude image unsubtracted image combination of flow sensitized data.

Mass number sum of neutrons and protons in the nucleus.

Matrix see **image matrix**.

Maximum intensity projection technique that uses a ray passed through an imaging volume to assign signal intensity according to their proximity to the observer.

Measurement gradient see **frequency-encoding gradient**.

Molecules where two or more atoms are arranged together.

Moiré artifact caused by a combination of aliasing and inhomogeneities in gradient-echo pulse sequences.

MOTSA	method combining a number of high-resolution 3D acquisitions to produce an image that has good resolution and a large area of coverage.
MR active nuclei	nuclei that possess an odd mass number.
MR angiography	method of visualizing vessels that contain flowing nuclei by producing a contrast between them and the stationary nuclei.
Multishot	where k -space is divided into segments, and one segment is acquired per TR.
N	
Net magnetic vector	magnetic vector produced as a result of the alignment of excess hydrogen nuclei with B_0 .
Neutron	neutrally charged element in an atomic nucleus.
NEX	see NSA .
No phase wrap	see antialiasing .
Noise	frequencies that exist randomly in time and space.
Nucleons	particles in the nucleus.
Number of excitations (NEX)	see number of signal averages (NSA) .
Number of signal averages (NSA)	number of signal averages. Number of times a line of k -space is filled with data.
Null point	the point at which there is no longitudinal magnetization in a tissue in an inversion recovery sequence.
Nutation	a spiral motion downward of the NMV from the longitudinal to the transverse plane.
Nyquist frequency	highest frequency that is accurately sampled. Occurs if the digital sampling frequency is at least twice the frequency being sampled.
Nyquist theorem	states that a frequency must be sampled at least twice in order to reproduce it reliably.
0	
Outer lines	area of k -space filled with the steepest phase-encoding gradient slopes.
Out of phase	when magnetic moments are not in the same place on the precessional path.
Out-of-phase signal cancelation	artifact caused by incoherence between magnetic moments of fat and water nuclei due to their chemical shift.

P

Parallel alignment	alignment of magnetic moments in the same direction as B_0 .
Parallel imaging	technique that uses multiple coils to fill segments of k -space.
Paramagnetic compounds	exhibit a weak attraction to an external magnetic field.
Partial averaging	filling only a proportion of k -space with data and putting zeros in the remainder.
Partial echo	sampling only part of the echo and extrapolating the remainder in k -space.
Partially saturated	occurs when the NMV is flipped beyond 90° (91° – 179°).
Partial voluming	loss of spatial resolution when large voxels are used.
Passive shielding	shielding accomplished by surrounding the magnet with steel plates.
Passive shimming	uses shims to adjust for large changes in field homogeneity.
Pathology weighting	achieved in IR pulse sequence with a long TE – pathology appears bright even though the image is T1 weighted.
Permanent magnets	magnets that retain their magnetism.
Phase	position of a magnetic moment on its precessional path at any given time.
Phase-contrast angiography	technique that generates vascular contrast using the phase difference between stationary and flowing spins.
Phase encoding	locating a signal according to its phase.
Phase image	subtracted image combination of flow-sensitized data.
Phase-locked circuit	used in incoherent or spoiled gradient-echo sequences to enable the receiver coil to lock onto the most recently transmitted transverse magnetization. Residual magnetization that has a different phase angle is ignored.
Phase matrix	the number of pixels in the phase axis of the image.
Phase mismatching	see ghosting artifact .
Phase oversampling	technique used to reduce aliasing – see antialiasing .
Phase reordering	process by which the amplitude and polarity of the phase-encoding steps are varied rather than applied linearly.
Phase wrap	aliasing along the phase encoding axis.
Pixel	two-dimensional location within a slice.
Polarity	direction of a gradient, i.e. which end is greater than B_0 and which is lower than B_0 . Depends on the direction of the current through the gradient coil.
Pole shoes	disks of a ferromagnetic alloy such as neodymium, boron, and iron or aluminum, nickel, and cobalt (alnico) used in permanent magnets.

Precession	secondary spin of magnetic moments around B_0 .
Precessional frequency	speed of precession.
Precessional path	circular pathway of magnetic moments as they precess around B_0 .
Preparatory pulses	first few RF excitation pulses in a sequence that are used to obtain the steady state.
Protocol	a set of rules. An algorithm.
Proton	positively charged element of an atomic nucleus.
Proton density	the number of protons per unit volume of that tissue.
Proton density contrast	image contrast derived from the differences in the proton densities of the tissues rather than any other mechanism.
Proton density weighting	image that demonstrates the differences in the proton densities of the tissues.
Pseudo-frequency	frequency that is indirectly derived from a pattern of phase change at a spatial location throughout the scan.
Pulse generator module (PGM)	responsible (among other things) for sending instructions related to transmitted RF pulses and the timing and duration of gradients.
Pulse sequence	a series of RF pulses, gradient applications, and intervening time periods.
Q	
Quantum theory	based on a subatomic understanding of how the universe works.
Quenching	sudden loss of the superconductivity of the magnet coils so that the magnet becomes resistive.
R	
Radians	another unit for degrees (in a circle).
Radio frequency (RF)	low-energy, low-frequency electromagnetic radiation. Used to excite hydrogen nuclei in MRI.
Ramped RF	spatially varying RF pulses that improve signal penetration in inflow angiograms.
Ramped sampling	where sampling begins before the frequency-encoding gradient reaches its maximum amplitude.
Ramping	the process of energizing an MRI solenoid.
RARE	rapid acquisition with relaxation enhancement. Another term for fast or turbo spin-echo .
Readout gradient	see frequency-encoding gradient .
Receive bandwidth	range of frequencies that are sampled during readout. Also called receiver bandwidth.

Rectangular FOV	also known as asymmetric FOV – uses an FOV in the phase direction that is different from that in the frequency direction of the image.
Reduction factor	see acceleration factor .
Region of interest (ROI)	a small area or volume of tissue that the system can measure. Used to calculate the SNR for example and also to trigger data acquisition.
Relative permittivity	degree to which different tissues reduce the RF field.
Relaxation	process by which spins lose energy.
Relaxivity	effect of a substance on relaxation rate.
Repetition time (TR)	time between each excitation pulse.
Residual transverse magnetization	transverse magnetization left over from previous RF pulses in steady state conditions.
Resistive MRI scanners	employ copper-wound solenoids that operate just below normal room temperature.
Resonance	phenomenon that occurs when an oscillating object is exposed to a frequency having the same or similar frequency to that of the object.
Resonator	a dipole antenna is a conductive rod that is designed to be half as long as the maximum RF wavelength to be transmitted/received.
Respiratory compensation	uses the mechanical motion of air in bellows attached on the patient's chest to order <i>k</i> -space filling and reduce respiratory motion artifact. Also known as respiratory ordered phase encoding (ROPE).
Respiratory gating/triggering	gates the sequences to chest wall movements to reduce respiratory motion artifacts.
Reverse-echo gradient-echo	pulse sequence characterized by a rewinder gradient that repositions the stimulated echo so that it is read by the system. True T2 makes a significant contribution to image contrast because the residual transverse magnetization contains mainly T2 contrast data.
Rewinders	gradients that rephase.
Rewound gradient-echo	see coherent gradient-echo .
RF amplifier	supplies power to the RF transmitter coils.
RF excitation pulse	short burst of RF energy that excites nuclei into a high-energy state.
RF homogeneity	the ability of a coil to provide RF at the same amplitude across the entire imaging volume.

RF rephasing pulse	additional RF pulse (usually with a magnitude of 180°) that rephases magnetic moments in the transverse plane to produce a spin-echo.
RF spoiling	use of a digitized RF phase-lock circuit to transmit and receive at a different phase every TR. Residual transverse magnetization is differentiated from new transverse magnetization by its phase value.
RF transmit coil	coil that transmits RF at the resonant frequency of hydrogen to excite nuclei and move them into a high-energy state.
Rheobase	minimum current required to depolarize the cell membrane of a neuron or muscle cell.
Rise time	time it takes a gradient to switch on, achieve the required gradient slope, and switch off again.
Rotating frame of reference	refers to the observer being part of a rotating system.
S	
Sampling interval	time between each data point in a line of k -space.
Sampling rate or frequency	see digital sampling rate or frequency .
Sampling time	time that the readout gradient is switched on for.
Sampling window	see sampling time .
SAR	specific absorption rate – a way of measuring the USA Food and Drug Administration limit for RF exposure.
SAT TR	time between each presaturation pulse.
Saturation	occurs when the NMV is flipped to a full 180° .
Scan time	time of the total acquisition. The time to fill k -space.
Sensitivity encoding	see parallel imaging .
Sequential acquisition	acquisition where all the data from each slice are acquired before going on to the next.
Shimming	process whereby the evenness of the magnetic field is optimized.
Signal	voltage induced in the receiver coil.
Signal-to-noise ratio (SNR)	ratio of signal relative to noise.
Single-shot FSE (SS-FSE)	fast spin-echo sequence where all the lines of k -space are acquired during a single TR period.
Single-shot imaging	technique that fills k -space in one shot.
Slew rate	strength of the gradient over distance.
Slice encoding	separation of individual slice locations by phase in volume acquisitions.

Slice-selection	selecting a slice using a gradient.
Solenoid electromagnets	magnets that use current passed through the coils of a wire to generate a magnetic field.
Spatial encoding	encoding or locating signal in spatial three dimensions of the imaging volume.
Spatial frequency	change of phase over distance in the patient caused by a gradient.
Spatial inversion recovery (SPIR)	method used to saturate fat.
Spatial resolution	refers to resolution in the image. The ability to distinguish two points in the image as separate.
Specific absorption rate	see SAR .
Spin-down	population of high-energy hydrogen nuclei that align their magnetic moments antiparallel to B_0 .
Spin-echo	echo produced as a result of a 180° RF rephasing pulse.
Spin-echo–EPI (SE-EPI)	spin-echo sequence with EPI readout.
Spin-echo pulse sequence	sequence that uses a 180° rephasing pulse to generate an echo.
Spin–lattice energy transfer	process by which energy is given up by spins to the surrounding molecular lattice.
Spin–spin relaxation	process by which interactions between the magnetic fields of adjacent nuclei cause dephasing.
Spin-up	population of low-energy hydrogen nuclei that align their magnetic moments parallel to B_0 .
Spoilers	gradients that dephase.
Spoiled gradient-echo	see incoherent gradient-echo .
Stationary frame of reference	refers to the observer viewing something moving. You are an outsider looking in.
Steady state	condition where the TR is less than T1 and T2 relaxation times of tissues. Also defined generically as a stable condition that does not change over time.
Stimulated echoes	echoes formed when any two RF pulses are used in steady state sequences.
STIR	short tau inversion recovery. A pulse sequence that suppresses signal from fat.
Superconducting magnets	solenoid electromagnets that use supercooled coils of a wire so that there is no inherent resistance in the system. The current flows and, therefore, the magnetism are generated without a driving voltage.
Surface coils	receive-only coils used to image surface anatomy.
Susceptibility weighting	technique that uses differences in magnetic susceptibility between tissues to generate image contrast.

T

T1 agent	agents that shorten T1 relaxation in tissues that take up the agent.
T1 contrast	image contrast derived from the differences in the T1 recovery times of the tissues rather than any other mechanism.
T1 recovery	growth of longitudinal magnetization as a result of spin–lattice relaxation.
T1 recovery time	time taken for 63% of the longitudinal magnetization to recover.
T1-weighted image	image that demonstrates the differences in the T1 recovery times of tissues.
T2*	dephasing due to magnetic field inhomogeneities.
T2* weighting	used in gradient-echo sequences to indicate T2 weighting. Recognizes that T2* decay makes a large contribution to image contrast.
T2 agents	agents that shorten T2 relaxation times in tissues that take up the agent.
T2 contrast	image contrast derived from the differences in the T2 decay times of the tissues rather than any other mechanism.
T2 decay	loss of transverse magnetization as a result of spin–spin relaxation.
T2 decay time	time taken for 63% of the transverse magnetization to decay.
T2-weighted image	image that demonstrates the differences in the T2 decay times of tissues.
Tau	time between the RF excitation pulse and the 180° RF rephasing pulse and the time between this and the echo. Sometimes used in STIR sequences as an alternative to the TI.
Teslas	unit of magnetic flux density usually for large magnetic fields. See also gauss .
Thermal equilibrium	assumes patient's temperature is constant and therefore does not influence the thermal energy of hydrogen during the MR experiment.
Three-dimensional volumetric	acquisition where the whole imaging volume is excited so that the images can be viewed in any plane.
Time-of-flight angiography	technique that generates vascular contrast by using the inflow effect.
Time-of-flight phenomenon	rate of flow in a given time – causes flowing nuclei to receive only one RF pulse (excitation or rephase but not both) and therefore produce no signal.
Time from inversion (TI)	time from 180° RF inverting pulse to 90° RF excitation pulse in inversion recovery pulse sequences.
Time to echo (TE)	see echo time .

Tip angle	see flip angle .
Torque	a rotational force that causes the alignment of an object to the lines of magnetic flux.
TR	see repetition time .
Trade-offs	consequences of changing protocol parameters.
Transceiver	coil that both transmits RF and receives the MR signal.
Transmit bandwidth	range of frequencies transmitted in an RF excitation pulse.
Transverse plane	axis perpendicular to B_0 .
Traveling SAT band	a saturation band follows the position of each slice as it is acquired.
Triggering	an action that triggers an event such as respiratory movement triggering an RF excitation pulse.
Triple IR prep	adds a further inverting pulse to a double IR prep sequence to null fat and blood together.
Truncation artifact	artifact caused by undersampling so that edges of high and low signal are not properly mapped into the image.
Turbo factor	see echo train length .
Turbo spin-echo (TSE)	pulse sequence that uses multiple applications of the phase-encoding gradient to encode multilines of k -space in a TR. Also called RARE and fast spin-echo .
Two-dimensional volumetric	acquisition where a line of data are acquired from each slice before repeating the TR.
V	
Velocity encoding	used in phase-contrast angiography to sensitize the pulse sequence to flow.
Venetian blind	artifact seen in MOTSA as partially saturated anatomy at the top of each slab exhibits a lower SNR than the entry-slice of the adjacent slab.
View–view mismapping	see phase mismapping .
Volume coil	coil that transmits and receives signal over a large volume of the patient.
Voxel	three-dimensional location in the patient.
Voxel volume	volume of tissue inside one voxel.

W**Warm bore**

area inside the cylinder of the cryostat. Contains the patient bore and components of the MRI system that operate at room temperature.

Water saturation

technique that nulls signal from water by applying an RF pulse at the frequency of water to the imaging volume before slice excitation.

Wrap

see **aliasing**.

X**x-y-axis**

see **transverse plane**.

Z**z-axis**

see **longitudinal plane**.

Zeeman interaction

coupling between hydrogen nuclei and the external magnetic field B_0 .

Zipper artifact

caused by extraneous RF entering the scan room.

Index

Note: Page number followed by italics are for figures and bold are for tables, respectively.

- abdominal imaging
 - aliasing along the phase axis of abdomen, *256*
 - entry-slice phenomenon (ESP) in, *283, 283–5*
 - out-of-phase signal cancellation in image of abdomen, *266*
 - phase mismatching in, *243*
 - using balanced gradient-echo, *121*
- acceleration factor, *203*
- acoustic gradient noise, *336*
- active shielding, *327*
- active shimming, *330*
- ADC
 - in brain, **53**
 - in diffusion imaging, *54*
 - extracellular space and, *52*
- adrenal tumors, out-of-phase imaging of, *268*
- algorithms, *203*
- aliasing, *170–1, 173, 232, 261, 277*
 - appearance, *253–4*
 - calculating, *255*
 - causes, *254*
 - FOV in, *253–4*
 - frequency wrap, *254*
 - in *k*-space, *203*
 - phase wrap, *255*
 - remedy, *255–61*
 - velocity-encoding (VENC) gradients, *307*
- alignment of nuclei, *6–7, 8*
 - antiparallel, *7*
 - classical theory explanation, *6, 6–7*
 - parallel, *7*
 - quantum theory explanation, *7, 8*
- American College of Radiology (ACR)
 - levels of personnel, *348*
 - safety zones, definition, *347–8, 348*
- American Society for Testing and Materials International, device-labeling icons, *349, 349*
- amplitude modulation, *167*
- analog, *167*
- analog-to-digital conversion (ADC), *167*
 - watch analogy and, *167*
- analog watches, *167*
- aneurysm clips, *270, 355, 368*
- angiograms
 - 2D sequential inflow, *299–300*
 - 3D volumetric inflow, *300*
- angular momentum, *4*
- anisotropic voxel, *236*
- antialiasing, *257–8*
 - along frequency axis, *257, 257*
 - along phase axis, *257, 260*
- anti-foldover (no phase wrap), *257*
- antiparallel alignment, *7*
- aorta
 - phase-contrast MRA of, *305*
 - spatial presaturation pulses for imaging, *295*
 - swapping phase and frequency in imaging of, *246*
- array processor, *280*
- artifacts
 - aliasing, *253–61*
 - bas-relief, *262*
 - chemical shift, *261–4, 265*
 - cross-excitation/cross-talk, *273–5*
 - from data acquisition errors, *280*
 - from equipment faults, *280*
 - flow, *280–3*
 - magic angle, *279–80*
 - magnetic susceptibility, *269–72*
 - Moiré, *277–9*
 - out-of-phase signal cancellation, *265–9*
 - phase mismatching or ghosting, *243–53*
 - remedies for, **296–7**
 - shading, *276–7*
 - truncation, *272–3*
 - zipper, *275–6*
- asphyxiation, *366*
- asymmetric FOV, *234–5*
- atomic number, *2*
- atomic weight, *2*
- atoms, *2, 6*
 - electrical charge of, *2*
 - motion in, *2–4, 3*
 - structure of, *2*
- Avogadro's law, *9*
- axial images
 - frequency-encoding gradients, *142*
 - phase-encoding gradients, *147*
- balanced gradient-echo, **122**
 - advantages and disadvantages, *121*
 - balanced gradient system in, *119, 119*
 - of lumbar spine, *121*
 - maintenance of steady state in, *120*
 - mechanism, *119–20*
 - parameters, *122*
 - steady state in, *120*
 - uses, *121*
- balanced gradients, *336*
- bandwidth, *139*
 - per pixel, *144*
- bas-relief artifact, *262*
- bellows, *248, 249*
- B_0 (external magnetic field), *8, 25, 313, 325*
 - in Boltzmann's equation, *7*
 - B_1 vs, *15*
 - Larmor frequency vs, *10*
 - number of spins and, *9*
- b* factor/value, *54, 56*
 - in DWI, *54*
- B_1 field, *8, 25, 337*
 - B_0 vs, *15*
- bipolar gradient, *93*
- birdcage design, *339*
- black-blood imaging sequences, *303–4, 304*
 - advantages, *304*
 - in cardiac imaging, *304*
- blipping, *204*
- blood oxygenation level dependent (BOLD) technique, *54, 325*

- blurring, 125
- Boltzmann equation, 7
- bone
 - STIR, 83
 - T1 contrast of, 43
- brain, 356
 - ADC values in, **53**
 - axial 3D inflow MRA images of, 298
 - axial gradient-echo image of, 215–16
 - axial reverse-echo gradient-echo in, 115, 115
 - axial T2-weighted FLAIR of, 84, 85
 - axial T1-weighted inversion recovery sequence of, 81
 - BOLD images of, 55
 - 3D inflow MRA images of, 298
 - Ernst angle graphs of, 102
 - fMRI of, 54, 325
 - hemorrhage, 300
 - midline sagittal T1-weighted image of, 51
 - number of signal averages (NSA) of sagittal, 219
 - sagittal image of, 219, 223–5, 253
 - showing aliasing or wrap around, 253
 - T1 and T2 relaxation times and signal intensity of, **101**
 - T2 decay times of, **30**
 - T1 recovery times of, **26**
- breath-hold images, T1-weighted, 111
 - brightness of pixel, 221
- bucking coils, 327
- burns
 - prevention during MRI, 360
 - RF-related, 358, 368
- carbon (C¹³), **5**
- cardiac pacemakers, risk of damage/ malfunction during scanning, 353, 355–6, 362–3
- cardiac triggering, 252
- carrier frequency, 139
- Cartesian filling, 162
- Cartesian *k*-space filling method, 165
- cerebrospinal fluid (CSF), 31
 - ADC values, **53**
 - T1 and T2 relaxation times and signal intensity of, **101**
 - T2 decay times of, **30**
 - T1 recovery times of, **26**
- cervical spine, 342
 - balanced gradient-echo imaging of, 121
 - sagittal imaging of, 248
 - T2* vs true T2 in images of, 113
- chemical shift, 228
- chemical shift artifact, 261, **265**
 - appearance, 261–2
 - causes, 262–4
 - pixel shift vs, 263, 264
 - reducing receive bandwidth for eliminating, 264
 - remedy for, 264
- chest of drawers analogy, 69, 160, 160–1, 164, 195, 203
- data acquisition type and, 206
- explanation of *k*-space and scan time, 190
- rectangular FOV and, 234–5, 236
- size of FOV and, 192
- children
 - open system scanner for, 316
 - psychological effects during scanning, 351
- chronaxie value, 363
- classical theory, 1, 4–5
 - direction of magnetic moments of spins, 7
- closed-bore systems, 316
- clothes sorting exercise, 161
- co-current flow, 282
- coherent magnetization, 15, 18–20, 26, 31, 91, 210
- coherent or rewound gradient-echo, **110**
 - of abdomen, 108
 - advantages and disadvantages, 109
 - echo formation in, 117
 - of knee, 108
 - mechanism, 106
 - parameters, 107–9
 - uses, 107
- coils
 - bucking, 327
 - coil position vs signal-to-noise ratio (SNR), 212
 - filling factor of, 339
 - gradient, 129, 131
 - head, 343
 - phased array, 211
 - quadrature, 211
 - RF homogeneity of, 339
 - RF receive, 340–3
 - RF transmit, 339
 - surface, 211
 - surface RF receive, 341–3
 - volume RF receive, 343
- compensatory gradients, 151
- computer system and graphical user interface, 344–5
- conjugate symmetry, 184
- contrast agents, 56–7
 - contrast-to-noise ratio (CNR), 227
- contrast-enhanced magnetic resonance angiograms (CEMRA), 298, 309
- contrast-to-noise ratio (CNR), 209
 - contrast agents, 227
 - definition, 226
 - factors affecting, 226–7
 - flow-related techniques, 227–8
 - magnetization transfer contrast (MTC), 227
 - optimization of, 232
 - T2-weighted images in, 227
 - using presaturation pulses, 228–32
- conventional spin-echo, 64, **67**
 - advantages and disadvantages, **67**
 - mechanism, 65
 - parameters, 67
 - spatial encoding in, 68
 - uses, 66
 - using one echo, 65, 65
 - using two echoes, 65–6, 66
- counter-current flow, 282
- cross-excitation/cross-talk, 273–5, 274
 - appearance, 273
 - causes-275, 273
 - remedy for, 275
- cryogenics (coolants), 320, 322
 - asphyxiation risk, 366
 - explosion risk, 367
 - quenching risk, 366–7
 - thermal sensitivity from, 365–6
- cryostat, 320–1, 321
 - physical construction of, 321, 321
 - primary function of, 321
- Curie temperature, 315
- dance analogy, 59
- data points, 164
- dead time and T1 recovery, 154
- dephasing, 28, 29
- diamagnetic compounds, 313
- diamagnetism, 313–14, 314
- diethylene triaminepentaacetic acid (DTPA), 56
- diffusion, 52, 53
- diffusion weighted imaging (DWI), 52–4, **56**
- digital sampling frequency, 169–71, 179
 - in terms of Nyquist frequency, 173
- digital sampling rate, 169
- digital watch, 167
- digitization of analog signal, 167
- discs, ferromagnetic, 318
- Dixon technique. *see* out-of-phase imaging (Dixon technique)
- Double Echo Steady State (DESS), 125
- double IR prep, 87
- driven equilibrium (DRIVE), 77, 77
- 2D sequential inflow angiograms, 299–300
- dual echo spin-echo sequences, 66
- duty cycle, 333
- 3D volumetric inflow angiograms, 300

- ECG-triggered 3D FSE, 308–9
- echo planar imaging (EPI), 122–4
 advantages and disadvantages, **126**
 contrast parameters, **126**
 uses and limitations, 124–7
- echo-spacing, 73
- echo time (TE), 22, 29
 magic angle artifact and, 280
 relationship between receive bandwidth, frequency matrix, and, 174–6, **177**, **177**
 resolution, effects of, 234
 slice number and, 178–9
 susceptibility artifact and, 271
 truncation artifact and, 272
- echo train, 68, 69
 full, 73
 shared, 73
 split, 73
- echo train length (ETL), 68
- eddy currents, 331
- EKG-triggered subtraction imaging, 308–9
- electrocardiogram (ECG, EKG) gating, 252
- electromagnetic (radiofrequency) fields,
 hazard effects of
 antenna effect, 360–2
 deterministic effects, 357
 patient heating/burns, 358–60
 temporary or permanent changes to implanted devices, 362–3
- electromagnetic spectrum, 12
- electromagnetism, 313
- electromotive force (emf), 18–19
- entry-slice phenomenon (ESP), 280–3,
 281–2, 284–5, **285**
 direction of flow, 282, 283
 flip angle, 282
 flow velocity, 282
 slice thickness and, 282
 TR in, 282
- Ernst angle, 101, 102
- even-echo rephasing, **296**
- extremity scanners, 317
- extrinsic contrast parameters, 25, 226
 T2* contrast and, 102
- Faraday's law of electromagnetic induction, 4–5, 18–19, 129, 313, 320, 352
- Fast Fourier Transform (FFT), 156, 158–9, 178, 179–83
- fast gradient-echo, 122
- fast or turbo spin-echo (FSE/TSE), 64, **76**
 advantages and disadvantages, **73**
 different rephasing angles in, 78
- filling of *k*-space, 69
 mechanism, 68–9
 parameters, 73–5
 PD-weighted image of a knee, 75 ?
 TR selection in TSE, 76
 T2-weighted image of a female pelvis, 72
 T1-weighted image of a male pelvis, 74
 uses, 72–3
 weighting in TSE, 70–1
- fat
 chemical shift artifact of, 262
 hyperintensity of, 328
 magnetic moments of, 33, 178, 228, 230–1, 262, 265, 269
 molecules, 32
 out of phase magnetic moments of, 269
 partial saturation of, 36
 periodicity of, 268
 saturation, 228, 228–9
 suppression, 83, 273, 326
 T2 decay in, 34
 T1 recovery in, 33
 T1 recovery time of, 36
- ferromagnetic compounds, 314
- ferromagnetism, 314–15, 315
- field of view (FOV)
 chest of drawers analogy and size of, 192
 distance between each data point and, 196–7
 frequency, 144–5, **151**, 194–5
 frequency encoding and, 144–5
 half FOV ghosts, 125
k-space, 191–5
 phase, 192–5
 pixel dimension and, 233
 receive bandwidth for whole, 178
 rectangular, 234–5
 resolution and, 234
 size of, 191–2, 194
 voxel dimensions and, 221, 223
- field strength (flux density), 325
- filling factor of a coil, 339
- filling options for *k*-space, 199–206, **200**
 in EPI, 204
 NSA controls, 190–1, 215
 parallel imaging, 201–3, 202
 partial echo, 201, 202
 partial or fractional averaging or half Fourier, **200**, 200, 200–1
 propeller, 205
 radial, 205
 single-shot (SS) imaging, 204–5
 spiral, 205, 205
- fine matrix, 233
- first-order motion, 290
 compensation, 293
- first spin-echo, 66
- flip, 16
- flip angle
 entry-slice phenomenon (ESP), 282
 gradient-echo pulse sequences, 90
 signal-to-noise ratio (SNR) vs, 214
 variable, 90
- flow artifacts, 280–3
 entry-slice phenomenon (ESP), 280–3, 284–5
 flow-related enhancement of signal, 280
 high-velocity signal loss, 280
 time-of-flight (TOF) phenomenon, 280, 285–90
- flow compensation techniques, **296**
 even-echo rephasing, **296**
 gradient moment rephasing (gradient moment nulling), 291–3, 292
 spatial presaturation pulses, 293–5, 294, 296
- flow-dependent (non-contrast-enhanced) angiography
 benefit of, 298
 2D sequential inflow angiograms, 299–300
 3D volumetric inflow angiograms, 300
 inflow MRA, 298–9, **303**
 scan time, 298
- flow-related enhancement, 281, 281
- flow-related techniques, 227–8
- flow-spoiled fresh blood imaging, 308–9
- flow velocity
 in entry-slice phenomenon (ESP), 282
 in time-of-flight (TOF) phenomenon, 287
- fluid attenuated inversion recovery (FLAIR), **60**, 64
 of brain, 85
 EPI-FLAIR, 123
 gadolinium and, 86
 mechanism, 84
 parameters, 86–7
 uses, 85
- flux density, 325
- fold-over suppression (no phase wrap), 257
- foreign bodies, risk from, 357
- free induction decay (FID) signal, 21, 29, 29, 153
- frequency encoding, 135, **145**, 151
 FOV and, 144–5
 keyboard analogy and, 143
 procedure, 142
 rephasing mechanism in, 142–3
 timing of, 142, 144

- frequency FOV, 144–5, **151**, 194–5
 frequency matrix, 144, **151**, 164, **177**,
 178–9, 184, 191, 193–4, **196**,
 208, 233–4
 of 10, 172
k-space and, 167–8
 relationship between TE, receive
 bandwidth, and, 174–6, 177, 177
 scan time and, 238
 sprinter analogy and, 171–2
 frequency modulation, 167
 frequency resolution
 in *k*-space, 193
 frequency wrap, 254
 fringe field, 318
 functional MRI (fMRI), 54–5, 55, **56**,
 325–6
- gadolinium (Gd), 56, 315
 contrast-enhanced magnetic resonance
 angiograms (CEMRA), 298
 ferromagnetic nature of, 315
 FLAIR and, 86
 incoherent or spoiled gradient-echo
 sequences, effect in, 111
 in IR pulse sequences, 81
 SPIR vs STIR sequences, 231
 gating, 250
 in cardiac imaging, 252
 electrocardiogram (ECG, EKG), 252
 gauss (unit), 132, 325, 335
 ghosting artifact. *see* phase mismatching
 gradient amplifiers, 332–3
 gradient amplitude, 133
 gradient axes, 134–5
x gradient, 134
y gradient, 134
z gradient, 134–5
 gradient coils, 129, 131, 331–2, 332
 gradient-echo, 93
 gradient-echo EPI (GE-EPI), 122, 123
 gradient-echo pulse sequences, 58, 89,
 93, **94**, **127**
 classification of, 106
 coherent or rewound gradient-echo,
 106–9
 common acronyms, **90**
 comparison of single-shot and
 multishot techniques, **123**
 extrinsic parameters, **99**
 gradient rephasing, 91–3, 92, 94
 parameter selection in, 119
 variable flip angle, 90
 vs spin-echo pulse sequences, 90
 weighting mechanisms in, 94–106
 gradient labeling, **135**
 gradient moment rephasing (gradient
 moment nulling), 253, 291–3, 292
- gradient off-set (dynamic) shimming, 330
 gradient polarity, 129, 130
 gradient power duty cycle, 335
 gradients, 8, 129
 field strength and precessional
 frequency, changes in, 130, **132**
 mechanism of, 129–33, **133**
 in pulse sequence diagrams, 134
 steep and shallow gradient slopes,
 133, 133
 watch analogy and, 132
 gradient spoiling, 91, 111
 gradient system, 330–7
 acoustic gradient noise, 336
 balanced gradients, 336
 characteristics of a magnetic field
 gradient, 334
 field strength, 333
 gradient amplifiers, 332–3
 gradient amplitude, 334–5
 gradient coil, 331–2, 332
 gradient power duty cycle, 335
 gradient rise time, 335
 gradient slew rate, 335
 graphical user interface, 344–5
 GRASE (gradient- and spinecho), 124,
 124
 gray matter, ADC values of, 52, **53**
 gyromagnetic ratio, 10
- Hahn echoes, 105
 half Fourier, 199–201
 half FOV ghosts, 125
 head coils, 343
 headphones, **364**
 heart
 balanced gradient-echo imaging of, 121
 gating in cardiac imaging, 252
 IR prep sequences of, 87
 phase encoding axis, 246
 heat analogy, 48, 49–50
 weighting in gradient-echo using, 97,
 98–9
 Heisenberg's Uncertainty Principle, 2
 helium dewar, 366
 high-energy nuclei, 7
 homogenous magnetic field, 328
 human body, common elements in, **5**
 hybrid sequences, 124
 hydrogen, 2
 energy states of, 7
 in fat, 33
 gyromagnetic ratio of, 10
 magnetic characteristics of, **12**
 nucleus/nuclei, 5, 7
 value of *S* for, 8
 hydrogen nuclei, magnetic moments of,
 5, 6–7, 64, 148, 152, 182, 328,
 338
- Boltzmann's equation, 7
 direction of, in classical and quantum
 theory, 7
 precessional frequency of, 130–1
- image blurring, 73, 125
 image contrast in diagnostic imaging,
 25, **56**
 contrast agents, 56–7
 contrast mechanism, 31–2
 extrinsic contrast parameters, 25
 functional MRI (fMRI), 54–5, 55,
56
 intrinsic contrast parameters, 25
 magnetization transfer contrast
 (MTC), 55
 proton density contrast of tissues, 41
 relaxation, 25–6, 32–5
 signal intensity of a tissue, 41
 susceptibility weighting (SWI), 55
 T1 contrast, 36
 T2 contrast, 40–1, **41–2**
 T2 decay, 26–31, 28
 T1 recovery, 26–7, 27
 weight image contrast, 42–8
 image matrix, 221
 in *k*-space, 233
 number of pixels in the FOV and, 233
 voxel size and, 233
 voxel volume and, 221
 incoherent or spoiled gradient-echo,
 112, **113**
 advantages and disadvantages, **113**
 echo formation in, 118
 mechanism, 109–11
 parameters, 111
 spoiling methods, 110–11
 uses, 111
 inflow MRA, 298–9, **303**
 advantages and disadvantages, **300**
 overcoming advantages of, **300**
 in phase (coherent), 13, 15
 inter-view mismatching, 245
 intravoxel dephasing, 290, 291
 intrinsic contrast parameters, 25, 226
 inversion recovery (IR), 64, **87**
 advantages and disadvantages, **85**
 fast, 82
 image reconstruction in, 87
 180° inverting pulse in, 79
 mechanism, 78–80
 parameters, 81–2
 PD weighting in, 80
 prep sequences, 87
 sequence, 79
 T1 weighted, of brain, 81
 T1 weighting in, 80
 uses, 81

- ionization, 2
- ions, 2
- isotopes, 2
- isotropic voxel, 235

- J-coupling, 72
- joint imaging
 - using balanced gradient-echo, 121

- k*-space, 68–70, 73
 - central lines of, 184, 186
 - chest of drawers analogy and, 69, 160, 160–1
 - contrast lines of, 186
 - data acquisition and image creation in, 165–83
 - data arrangement in, 164
 - data points vs photographs, 169
 - definition, 159
 - digital sampling frequency, 169–70
 - filling and phase reordering, 71
 - filling methods, 159, 161–2, 164
 - filling options, 199–206, **200**
 - FOV dimensions, 233
 - frequency axis of, 159, 160, 192, 194
 - frequency FOV, 192–5
 - frequency matrix, 168, 193
 - frequency resolution, 193
 - grid of data points in, 167, 168, 184
 - image matrix in, 233
 - labeling, 163
 - lines of, 161
 - meaning of *k*, 159
 - multiple noise spikes, 277
 - negative and positive halves of, 200
 - phase axis of, 159, 160, 182, 192, 194
 - phase-encoding gradient and, 163–5
 - phase FOV, 192–5
 - phase gradient amplitude vs signal amplitude, 186
 - phase matrix, 168, 190, 193, 233
 - phase mismapping and, 245
 - pixel size, , 193–5, 233
 - positive polarity phase-encoding gradients, 161
 - pulse sequences, 197–9
 - respiratory compensation and, 248, 249–50
 - sampling window, 167, 169–72
 - scan time to fill, 188–91
 - sequential and 3D (volume)
 - acquisition of data in, 206–8
 - signal and resolution of image, 185–7, 186–7
 - signals from one voxel, 180
 - in a single shot, 204
 - size of FOV, 191–2, 194
 - slice-select gradient application, 164–5
 - spatial resolution, 193
 - symmetrical features, 182, 184, 185
 - units of, 159
 - using signal data only, 188
 - vs spatial frequency, 159
 - waveforms, 166180

- laminar flow (parabolic flow), 288–90
- Larmor equation, 10
- Larmor Grand Prix analogy, 62, 63
- Lenz's law, 19
- level 1 personnel, 348
- level 2 personnel, 348
- lipids, 32
- liquid helium, 322
 - risks from, 365–7
- liver
 - ADC values, 52
 - nulling of signals in, 231–2
- lobes, 134
- longitudinal and transverse magnetization, 21
- longitudinal plane, 7
- Lorentz force, 328, 336, 352
- low-energy nuclei, 7

- magic angle artifact, 279–80
 - appearance, 279
 - causes, 279–80
 - in collagen rich tissues, 279–80
 - echo time and, 280
 - of patellar tendon, 279
 - remedy for, 280
- magnetic flux density, units of, 132
- magnetic homogeneity, 343
 - in MRI, 328–9
- magnetic isocenter, 129, 132–3
 - location of pixel at, 182
- magnetic moment, 4, 9, 13–14
- magnetic moments of fat nuclei
 - inversion of, 230
 - precessional frequency of, 228, 231
 - saturation of, 228
- magnetic resonance angiography (MRA)
 - flow-dependent, 309
 - inflow, 298–9, **303**
 - pulse sequence parameters, 309
- magnetic resonance imaging (MRI)
 - classical theory in, 1
 - components of, 312
 - computer system and graphical user interface, 344–5
 - extrinsic contrast parameters, 25
 - fat and water, 32, 32–3
 - frequency encoding in, 143
 - fundamentals of, 1
 - geometric accuracy of, 269–70
 - image contrast in, 24
 - intrinsic contrast parameters, 25
 - k*-space filling methods in, 161–2
 - magnetic field homogeneity in, 328–9
 - patient transport system, 343–4
 - principles of, 2
 - scanners, 311–12, 312
 - slice-selection in, 138
 - thermal equilibrium in, 9
 - voxel geometry and pixel dimension, 233
- magnetic susceptibility, 313
- magnetic susceptibility artifact, 269–71
 - appearance, 269
 - causes, 269–70
 - of a dental implant, 270
 - echo time and, 271
 - of knee, 271
 - receive bandwidth and, 271
 - remedy, 270–2
 - spin-echo sequences and, 270, 272
- magnetism, **313**, 313–15
 - diamagnetism, 313–14
 - ferromagnetism, 314–15
 - paramagnetism, 314
- magnetization transfer contrast (MTC), 55, 227
- magnet shielding
 - active, 327
 - passive, 326–7, 327
- magnet system, requirements of
 - field strength, 318
 - fringe field, 318
 - homogeneity aspect, 318
 - weight and bulk of magnet, 318
- magnitude images, 307
- mass number, 2, 4
- maximum intensity projection (MIP), 300–1, 301
- Moiré artifact, 277–9
 - appearance, 277–9
 - field inhomogeneity, 277–9
 - multiple noise spikes, 277, 278
 - as zebra lines, 277, 278
- molecules, 2
- MR active nuclei, 4–5, 10
 - magnetic moments of, 12
- multiple overlapping thin slab acquisitions (MOTSA), 302, 302
- muscle, 31
 - cardiac, 363
 - images in steady-state sequences, 100
 - magnetization transfer effects of, 72
 - rheobase, 363
 - T1 recovery and T2 decay times, 102

- negative lines, 161
- neodymium magnets, 318
- net magnetic vector (NMV), 7–9, 9, 21, 26, 31, 61, 79–80, 82, 90
 - of fat, 33
 - fully saturated, 36
 - partially saturated, 36
 - signal-to-noise ratio (SNR), effect on, 9
 - of water, 33
- neurostimulators, risk of damage/ malfunction during scanning, 355
- neutrons, 2
- nitrogen (N^{15}), **5**
 - magnetic characteristics of, **12**
- noise, 210
- non-MRI personnel, 348
- nucleons, 2
- null point, 82
- number of signal averages (NSA), 190–1, **196**
 - multiple, 252
 - phase mismapping and, 258
 - of sagittal brain, 219
 - scan time and, 237
- nutration, 14
- Nyquist frequency, 171
- Nyquist theorem, 170, 170

- Ørsted, Hans Christian, 313, 319
- open systems, 316
- optimal saturation, 230
- optimizing image quality, **239**
- out-of-phase imaging (Dixon technique), 232, 268
- out of phase (incoherent), 13, 15
- out-of-phase signal cancellation, 265–9
 - appearance, 265
 - benefits of, 268
 - causes, 265
 - watch analogy and, 266, 267
- oxygen (O^{17}), **5**
 - magnetic characteristics of, **12**

- PACS (picture archiving and communications system) access network permitting teleradiology, 345
- parallel alignment, 7
- parallel imaging, 201–3, 202
- paramagnetic compounds, 314
- paramagnetism, 314, 314
- partial averaging, 200–1
- partial echo, 201, 202
- partial voluming, 232
- passive shielding, 326–7, 327
- passive shimming, 329–30, 330
- pathology weighting, 82

- patient transport system, 343–4
- PD weighting, in coherent gradient-echo, 109
- peripheral pulse triggering, 252
- permanent magnet MRI scanners, 318–19, 319
- permanent magnets, 316
- phase, 13
- phase-contrast angiography, 304–7, **307**
 - advantages and disadvantages, 305–7, **307**
 - gradient, 305
- phase-contrast MR angiography, 228
- phase encoding, 135, 146, **151**, 179
 - amplitude and polarity changes in, 149–50, 153
 - of axial images, 147
 - of coronal images, 147
 - long and short axis of anatomy, 147
 - phase matrix and phase resolution in, 149–50
 - phase shift along the gradient, 155, 156
 - procedure, 145–7
 - of sagittal images, 147
 - selection of direction, 147
 - spatial frequencies in, 155–6, 156, 158
 - steep and shallow gradients, 149, 149–50
 - timing of, 148, 148–9
 - watch analogy and, 147
- phase FOV, 194–5
- phase images, 307
- phase-locked circuit, 110
- phase matrix, 149–50, **151**
 - scan time and, 237
- phase mismapping, 232, 243–53, **253**
 - appearance, 243
 - cause, 244, 244–5
 - equation, 243
 - gradient moment rephasing for reducing, 253
 - image through a breathing abdomen, 243
 - k -space and, 245
 - multiple NSA for reducing, 252
 - phase-encoding gradient, 244
 - presaturation pulses for reducing, 248
 - remedy for, 245–53
 - respiratory compensation for reducing, 248–9
 - respiratory gating and triggering for reducing, 250–1, 251
 - swapping phase and frequency for reducing, 246
- phase oversampling (no phase wrap), 257
- phase reordering, 72
- phase resolution, 149–50, **151**

- phase wrap, 255
- pixel, 129, 221
- pixel dimension, 233
- pixel size
 - frequency, 194, **196**
 - in k -space, 233
 - phase, 194–5, **196**
- polarity of gradient, 93
- pole shoes, 318
- precession, 10, 11
 - spin-up and spin-down populations, 11
- precessional (Larmor) frequency, 10, 13–14, 33, 55, 128, 139, 148
 - between fat and water, 228
 - gyromagnetic ratio, 10
 - of hydrogen, 31
 - Larmor equation, 14
 - magnetic field strength and, 28–9, 130, 130, 135
 - of magnetic moments of fat nuclei, 228
 - of magnetic moments of hydrogen nuclei, 34, 103, 130–3, 135–6, 139, 142, 145–7, 228
 - shallow gradient slopes, effect of, 133
 - slice-select gradient and, 136, 139
- precessional path, 10
- presaturation pulses, 228–32, 293–5, 294, 296
 - for reducing phase mismapping, 248
- projectile hazards, 352–5, 353–4
- propeller k -space filling, 205
- protium, 5, **5**
- protocol (s)
 - characteristics, 209–10
 - definition, 209
 - modification, 238–41
 - optimization of, 210
- proton density contrast of tissues, 41
 - examples, **47**
- proton density (PD) of a tissue, 31
- proton density weighting, 46, **48**
 - of ankle, 47
 - heat analogy, 50
- protons, 2
- pseudo-frequency, 180
- psychological effects of MRI, 350–1
- pulsatile flow, 293
- pulse generator module (PGM), 344
- pulse sequence, 22, 22, 25
- pulse sequences, 21, 58
 - definition, 59
 - diagram, 59
- pulse sequence timing, 152–4
- pulse timing parameters, 22
- pulse width modulated (PWM) design, 332

- quantum theory, 1, 7
 quenching, 354, 366–7
- radial *k*-space filling, 205
 radian (s), 13, 159
 radio frequency (RF), 12
 radio frequency (RF) excitation pulse,
 14–15, 29, 128
 dummy pulses, 37
 magnetization in water, effect on, 36
 preparatory, 37
 ramped RF, 302
 ramped sampling, 122
 ramping a magnet, 323, 324
 rare-earth magnets, 318
 RARE (Rapid Acquisition with
 Relaxation Enhancement), 68
 readout or measurement gradient, 142,
 167
 receive bandwidth, 144, 172–4
 parameters of different
 manufacturers, 178
 reducing for eliminating chemical
 shift artifact, 264
 reducing of, 220–1
 relationship between TE, frequency
 matrix, and, 174–6, **177**, 177
 signal-to-noise ratio (SNR) vs, 220,
 220
 susceptibility artifact and, 271
 for the whole FOV, 178
 rectangular FOV, 234–5
 reduction factor, 203
 relative permittivity, 362
 relaxation, 20, 25–6, **31**
 in different tissues, 32–5
 relaxivity, 56
 renal area, chemical shift artifact of, 262
 repetition time (TR) period, 22, 154
 acceleration factor, 203
 in balanced gradient echo, 120
 in coherent gradient-echo, 120
 entry-slice phenomenon (ESP), 282
 equal infinity, 124
 frequency-encoding gradient in, 183
k-space filling and, 163–5, 189, 201,
 203
 longitudinal magnetization during, 95
 reduction factor, 203
 respiratory gating and triggering,
 effects of, 251
 SAT TR, 230
 scan time and, 237
 signal to noise ratio vs, 213–15
 T1 recovery and, 154
 vs slice number, 189, 189
 resistive electromagnets, 318–20
 resistive MRI scanners, 320
- resonance, 8, 13–16
 resonator, 360
 respiratory compensation techniques,
 248–9, 248–50
 respiratory gating and triggering, 250–1,
 251
 respiratory navigator echoes, 248
 respiratory sensor, 248
 reverse-echo gradient-echo, **114**, **117**
 actual TE, 114
 advantages and disadvantages, **116**
 in brain, **115**
 echo formation in, **118**
 effective TE, 114
 mechanism, 113–14
 parameters, 116–17
 steady-state sequences in, 117–18
 uses, 115
 rewinders, 93
 RF amplifier, 339
 RF homogeneity of the coil, 339
 RF receive system, 339–40
 RF rephasing, 59–62, 64
 180° RF rephasing pulse, 60–1, **61**
 RF rephasing in spin-echo pulse
 sequences, 139, 141
 frequency encoding and, 142–3
 in hybrid sequence, 124
 in spin-echo EPI (SE-EPI), 123
 RF shielding, 337–8
 RF spoiling, **110**, 110–11
 RF system, 337–43
 RF receive coils, 340–3
 RF receive system, 339–40
 RF shielding, 337–8
 RF transmit coils, 339
 surface RF receive coils, 341–3
 transmit system, 338–9
 volume RF receive coils, 343
 rheobase, 363
 right-hand grip rule, 319, **320**
 rise time of gradient, 335
 rotating frame of reference, 18
- safety (MRI)
 of devices, 349–50
 MRI conditional, definition, 349
 MRI safe, definition, 349
 MRI unsafe, definition, 349
 of patient, 350
 philosophy behind, 347
 protection of patient's hearing during
 scan, 365
 tips, 367–8
 zones, 347–8
 sagittal brain, using a square FOV, **225**
 sampling
 digital sampling frequency, 169–71
 digital sampling rate, 169
- equations, 171
 at less than once per cycle, 170
 once per cycle, 170
 sprinter analogy and, 171–2
 sampling interval, 169
 sampling time, 142
 sampling window, 142, 173–4, **178**, **179**
 frequency-encoding gradient in,
 174–5
 impact on TE, 174
 sampling window (acquisition window),
 167
 SAT TR, 230
 saturation, 17
 saturation of fat signal, 228
 saturation pulses, 248
 scanner configurations
 closed-bore systems, 316
 extremity systems, 317
 open systems, 316
 scan time, 68–9, 210
 in flow-dependent (non-contrast-
 enhanced) angiography, 298
 frequency matrix and, 238
 for improving resolution, 238
 NSA and, 237
 optimization, 237
 phase matrix and, 237
 TR and, 237
 sensitivity encoding. *see* parallel
 imaging
 sequential acquisitions of data, 206,
 207
 sequential and 3D (volume) acquisition
 of data in *k*-space, 206–8
 shading artifacts, 276–7
 appearance, 276
 causes, 276–7
 remedy for, 277
 shimming, 328–30
 active, 330
 gradient off-set (dynamic), 330
 passive, 329–30, **330**
 shim system, 328–30
 short tau inversion recovery (STIR), **60**,
 64, 308–9
 of knee, **84**
 mechanism, 82–3
 parameters, 84
 pulse sequence, 231
 uses, 83
 sickle-cell anemia, 290
 signal intensity of a tissue, 41
 signal-to-noise ratio (SNR), 209, **226**,
 233
 coil position vs, **212**
 definition, 210
 equations, 211
 field strength and, 325–6

- signal-to-noise ratio (*continued*)
 flip angle vs, 214, 214
 magnetic field strength and, 211
 net magnetic vector (NMV), effect of, 9
 number of signal averages vs, 215–18, 218–19
 optimization of, 226
 phased array coils and, 211
 proton density and, 211
 quadrature coils and, 211
 receive bandwidth vs, 220, 220
 relationships between protocol parameters and, **226**
 signal amplitude affecting, 210
 slice thickness vs, 222
 spatial resolution and, 239, 241
 surface coils and, 211
 TE controls vs, 215, 216–18
 trade-offs in, 238, 240
 TR controls vs, 213, 213–15
 T2-weighted images in, 227
 type of coil and, 211
 volume imaging and, 223
 voxel volume vs, 221–3
- single shot fast or turbo spin-echo (SS-TSE), 77, 204
- single-shot (SS) imaging, 204–5
- slew rate, 334–5
- slice encoding, 207
- slice gap, **151, 196**
- slice-selection, 135–9, **141, 151**
 gradients performing, 136, 138
 in MRI system, 138
 procedure, 135–6
 slice thickness and, 139–41, 140
 timing of, 136, 138
 tuning fork analogy and, 136, 137
- slice thickness, 139–41, **151, 196**
 entry-slice phenomenon (ESP), 282
 time-of-flight (TOF) phenomenon and, 287
- solenoid electromagnets, 316–17, 317–18, 323
- solenoid magnets, 319, 322
- space-blanket, 321
- spatial encoding, 129, 161, 230, 270, 352
 in conventional spin echo, 68
 gradient axes, 134–135
 mechanisms of gradients, 129–133
 slice-selection, 135–139
 watch analogy and, 152
- spatial frequency, 155
 musical notes on a piano keyboard, 155
 from slice-selection, 158
 in steep phase-encoding gradients, 155
- vs amplitude of phase-encoding gradient, 155–6, 156
- spatial inversion recovery (SPIR), 230
 STIR vs, 231
- spatial presaturation pulses, 293–5, 294, 296
 for black-blood imaging sequences, 303
- spatial resolution, 210, 232–4, **237**
 in *k*-space, 193
 in MR, 239
 voxel size and, 232
- specific absorption rate (SAR), 78, 346
- spinal imaging
 using balanced gradient-echo, 121, 121
- spine
 flexion and extension views of, 316
 spin-echo, 61
 spin-echo EPI (SE-EPI), 122–3, 123
 of abdomen, 125
 spin-echo pulse sequences, 58, 62, **88**
 common acronyms, **60**
 elimination of T2* dephasing, 61
 gradient timing in, 152–4, 153
k-space filling in, 164
- spin–lattice energy transfer, 26
- spin–spin interaction, 28
- spin–spin relaxation, 27–8
 optimality of, 35
- spiral flow (helical flow), 290
- spiral *k*-space filling, 205, 205
- spleen, nulling of signals in, 231–2
- spoilers, 91
- sprinter analogy, 171–2
- static field of MRI scanner, implications for safety
 foreign bodies, 357
 projectile hazards, 352–5
 torque on implanted devices, 355–7
 transient biological effects, 351–2
- stationary frame of reference, 18
- steady-state, 37, **105**
 echo formation in, 103–6, 104–5
 weighting of steady-state gradient-echo sequences, 100–1, 101, **102**
- stents, risk of damage/malfunction during scanning, 355
- superconducting MRI scanner, 322
- superconductive magnet systems, **326**
 electromagnets, 320–2
- superconductivity, 322
- suppression
 CSF, 84
 fat, 83, 273, 326
 fold-over, 257
 in inflow angiography, 268, **300**
- surface RF receive coils, 341–2, 341–3
- susceptibility weighting (SWI), 55
- T1 agents, 56
- T2 agents, 56
- tau, 60, 64, 64
- T1 contrast, 36, **52**
 achieving steady-state, 37
 examples, **43**
 impact on saturation effects, 36, 38–9
 longitudinal magnetization in fat and water, 36, 37
- T2 contrast, 40–1, **41–2, 52**
 examples, **46**
- T2* contrast, 97
 extrinsic contrast parameters and effect on, 102
 in gradient-echo, 97
 TE controls and effect on, 95
- T2 decay, 26–31, 28, **41, 44, 52, 61, 65–6, 94, 102, 153**
 of brain tissue, **30**
 in fat, 34
 between fat and water, 45
 strength of B_0 and, 35
 in water, 34
- T2* decay, 94
- T2 decay time, 29, **52, 55**
 of fat and water, 40
- T2* dephasing, in spin-echo pulse sequences, 59–61, 60, 61, 62
 in magnetic moments of hydrogen nuclei, 60, 61
- teslas (T), 8, 132, 325
- thermal equilibrium, 7
 spin-up to spin-down nuclei in, 9
- thermal sensitivity, cryogen, 365–6
- three-dimensional (3D) volumetric acquisition (volume imaging), 207–8
- time-of-flight (TOF) phenomenon, 280, 285–90, 286, **288**
 flow velocity, 287
 in gradient-echo pulse sequences, 287
 slice thickness and, 287
 in spin-echo pulse sequences, 286–7
 TE vs, 288
 type of pulse sequence and, 285
 types of flows, 288–90, 289
- time-varying gradient magnetic fields, hazard effects of
 adverse effect on implanted devices, 363
 nerve stimulation, 363
 noise-induced hearing loss, **364, 364–5**
- tip angle, 16
- tissues
 aliasing of, 203
 contrast agents and visualization of, 56
 dephasing between, 271

- diffusion imaging of, 54
- EPI sequences on, 124
- gradient-echo imaging, 95–6
- IR pulse sequences, effect of, 81–2
- magnetization in, 100, 122, 269, 281
- nulling of signals in, 231–2, 293
- proton density of, 211
- RF rephasing effect on, 62, 65
- susceptibility weighting (SWI) of, 55–6
- T1 and T2 relaxation times of, 100, **101**, 103–6
- T2 decay time of, 65–6, 73, 102, 113
- turbo spin-echo imaging, 73
- T2-weighted images of, 56
- TI (time from inversion), 80
- torque on implanted devices, effect on patient, 355–7, *356*
- trade-offs, 210
- transceiver, 339
- transient biological effects, 351–2
- transmit bandwidth, 139
- transverse plane, 14
- traveling SAT band, 299
- T1 recovery, 26–7, 27, 42, **52**
 - dead time and, 154
 - in fat, 33
 - between fat and water, 42, 42
 - in water, 33, 34
- T1 recovery times, 26, **52**, 81
 - of brain tissue, **26**
 - of fat, 36
- triggering, 250
 - cardiac, 252
 - EKG, 252
 - peripheral pulse, 252
- triglyceride fat molecules, 262
- triple IR prep, 87
- truncation artifacts, 272–3, 273
 - appearance, 272
 - causes, 272
 - echo time and, 272
 - remedy, 273
- turbo factor, 68–9
- T2-weighted image
 - of buttock, 227
 - of pelvis, 230–1
- T1-weighted images, **44**, **52**, 300, 303
 - of brain, 51
 - breath-hold images, 111
 - chest, 246–7
 - fat suppression techniques used in, 273
 - heat analogy, 49
 - of knee, 43
- T2-weighted images, 44, **46**, **52**
 - contrast-to-noise ratio (CNR) in, 227
 - heat analogy, 49
 - signal-to-noise ratio (SNR) in, 227
 - of wrist, 45
- T1 weighting
 - in coherent gradient-echo, 109
- T2* weighting, 94, 124
 - in coherent gradient-echo, 106, 109
 - in gradient-echo and heat analogy, 97, 98
- two-dimensional (2D) volumetric acquisitions of data, 206, 207
- ultra-high-field research magnets, 327
- ultra-high-field-strength systems, 327, 351–2
- unpaired electrons, 56, 314–15, 319
- user interface, 344–5
- variable flip angle, 90
- vascular patency, 290
- vectors, 21, 303, 334
 - at different frequencies dephase, 28
 - extrinsic contrast parameters, effect of, 95
 - fat, 35, 36, 42, 44, 46, 65, 81, 95–6, 265
 - in 180° RF rephasing, 61
 - RF spoiling effect on, 110
 - water, 35, 36, 42, 44, 46, 65, 81, 95–6, 265
- velocity-encoding (VENC) gradients, 304, 306–7
- venetian blind effect, 302
- view-to-view mismapping, 245
- volume imaging, 188, 200
 - aliasing in, 261
 - resolution and, 235–6
 - SNR and, 223
 - three-dimensional (3D) volumetric acquisition, 207–8
- volume RF receive coils, 343
- vortex flow, 290
- voxel, 129
- voxel size
 - FOV vs, 233
 - number of pixels or image matrix and, 233
 - slice thickness vs, 232
- voxel volume
 - changes in FOV and, 223
 - image matrix and, 221
 - signal-to-noise ratio (SNR) vs, 221–3, 222
 - voxel size and slice thickness, 221
- warm bore, 321, 321
- warning signs, 355
- watch analogy, 17–18, 132
 - change of phase over time, 155
 - out-of-phase signal cancellation and, 266, 267
 - phase encoding and, 147
 - spatial encoding and, 152
- wrap. *see* aliasing
- water
 - chemical shift artifact of, 262
 - magnetic moments of, 178, 262, 265, 269
 - molecules, 32
 - net magnetization vector (NMV) of, 33
 - out of phase magnetic moments of, 269
 - periodicity of, 268
 - saturation, 228, 229
 - T2 contrast image of, 40, 40
 - T2 decay in, 34
 - T1 recovery in, 33, 34
- weight image contrast, 42–8
 - proton density weighting, 46
 - T1-weighted image, 42. *see also* T1-weighted images
 - T2-weighted image, 44. *see also* T2-weighted images
- weighting mechanisms in gradient-echo pulse sequences, 94–106, **100**
 - extrinsic contrast parameters, 95–6
 - heat analogy and, 97, 98–9
 - PD weighting, 96, 99
 - residual transverse magnetization, 102
 - steady state, 100–1, 101, 102, 106
 - T1 weighting, 95–6, 96
 - T2* weighting, 96, 97
- white matter
 - ADC values, **53**
 - FLAIR imaging of, 85
 - T1 and T2 relaxation times and signal intensity of, **101**
 - T2 decay times of, **30**
 - T1 recovery times of, **26**
- x gradient, 134, 136, 142, 147, 150, 331
- x–y-axis, 14
- y gradient, 134, 136, 142, 147, 331
- z-axis, 7
- Zeeman interaction, 7–8
- z gradient, 134, 136, 141–2, 145, 331
- z gradient amplifier, 331
- zipper artifact, 275–6
 - phase mismapping, 275, 276
 - causes, 276
 - remedy for, 276

WILEY END USER LICENSE AGREEMENT

Go to www.wiley.com/go/eula to access Wiley's ebook
EULA.