


PRACTICAL MR PHYSICS

And Case File of MR Artifacts and Pitfalls



ALEXANDER C. MAMOURIAN

OXFORD

PRACTICAL MR PHYSICS

This page intentionally left blank

PRACTICAL MR PHYSICS

AND CASE FILE OF MR ARTIFACTS AND PITFALLS

Alexander C. Mamourian, MD

Associate Professor of Radiology

Division of Neuroradiology

University of Pennsylvania

Philadelphia, Pennsylvania

OXFORD
UNIVERSITY PRESS

2010

OXFORD
UNIVERSITY PRESS

Oxford University Press, Inc. publishes works that further
Oxford University's objective of excellence
in research, scholarship and education.

Oxford New York
Auckland Cape Town Dar es Salaam Hong Kong Karachi
Kuala Lumpur Madrid Melbourne Mexico City Nairobi
New Delhi Shanghai Taipei Toronto

With offices in
Argentina Austria Brazil Chile Czech Republic France Greece
Guatemala Hungary Italy Japan Poland Portugal Singapore
South Korea Switzerland Thailand Turkey Ukraine Vietnam

Copyright © 2010 by Oxford University Press, Inc.

Published by Oxford University Press Inc.
198 Madison Avenue, New York, New York 10016
www.oup.com

Oxford is a registered trademark of Oxford University Press

First published 2010

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,
without the prior permission of Oxford University Press,

Library of Congress Cataloging-in-Publication Data

Mamourian, Alexander C.

Practical MR physics: And case file of MR artifacts and pitfalls / Alexander C. Mamourian.
p. ; cm.

Includes bibliographical references and index.

ISBN 978-0-19-537281-6

1. Medical physics. 2. Magnetic resonance imaging. 3. Nuclear magnetic resonance.

I. Title. [DNLM: 1. Magnetic Resonance Imaging—Case Reports. 2. Magnetic Resonance
Imaging—Problems and Exercises. 3. Image Interpretation, Computer-Assisted—Case Reports.
4. Image Interpretation, Computer-Assisted—Problems and Exercises. WN 18.2 M265p 2010]

R895.M26 2010


616.07'548—dc22

2009030871

All photographs by Alexander C. Mamourian.

1 3 5 7 9 8 6 4 2

Printed in the United States of America
on acid-free paper



CONTENTS

Preface	vii
Acknowledgments	ix
1 MR PHYSICS	1
2 MR ARTIFACTS	45
3 MR PITFALLS	167
4 TEN PUZZLERS	275
Index	301

This page intentionally left blank



PREFACE

When I filled out my application for a fellowship in MR imaging at the University of Pennsylvania 25 years ago, I was required to write an essay. I included in mine, at the urging of my sister Alicia, a plan to write a book about MR physics. I was never sure of whether it was because of, or in spite of, that essay the department chairman, Dr. Stanley Baum, accepted me for the fellowship but told me during my interview that this was among the most naive ideas he had ever heard from a prospective fellow. There was some truth to that since I was one year out of my residency training. It seems like a better idea now.

While MR physics has never been an easy subject to master, the bar is higher than ever since MR imaging techniques have become more complex. I suspect that many trainees have given up altogether and treat MR with due respect but at arms length, much as I treat my car's engine. Still, it is helpful to have some understanding of MR physics so that you can suggest ways to improve image quality and avoid mistaking common artifacts for disease.

I am not a physicist; this will prove to be a good thing for some readers and disappointing for others. Magnetic resonance imaging has been my professional focus for many years, however, and I hope you will find the explanations and analogies drawn from that experience clear and helpful. This book is not intended to replace the many fine books that focus on MR physics. While the basics of MR physics are included in Chapter 1, my goal is to use case material to illustrate how those principles will help you to identify and understand common artifacts. After reading this book, you will be better prepared to understand more advanced MR techniques as well.

The book is divided into four parts: (1) an overview of MR physics, (2) common MR artifacts, (3) common MR pitfalls, and (4) challenging cases. There are several ways you can use the book. For most readers who want to learn more about MR imaging, cover to cover will do best. You will also find in the book links to five instructional videos on the web that were created to complement this text. I encourage you to view them since the selected topics, such as the motion of precession, are easier to demonstrate than explain. If you are comfortable with the physics, a review of the artifacts and pitfalls should sharpen your imaging skills and will provide a refresher on MR physics. The answer for each case will be defined by a box, usually at the bottom of the page, so avert your gaze until you are ready for it. If you are an experienced imager but happen to encounter what you think is an artifact while you are reviewing clinical MR images, you can use the index in this book to investigate that topic. All readers should try to solve the puzzler cases either before or after reading the book, but the answers will surely come more easily afterward. It is my hope that this book proves to be valuable to you and, in that way, helpful to your patients.

This page intentionally left blank



ACKNOWLEDGMENTS


I am grateful for the time with my grandfather Garabed, who provided the groundwork; my parents, Marcus and Maritza, for helping me find my way; my sister Alicia, especially for drafting that essay; and my children, Ani, Molly, Liz, and Marcus, for giving me all the joyful reasons. To my wife Pam, words cannot express the depth of my gratitude for surely without you this book would have remained an unfulfilled promise.

I also want to thank the many gifted MR technologists I have worked with over the years, with special thanks to Swapan Sen, Chris Harris, Sharon Hurst, Bob Ferranti, Shreve Soule, and Theresa Haron, Theresa for both her competence and her keen eye.

My sincere thanks go to Andrea Seils at Oxford University Press for believing in this project, Josef Debbins, PhD, at the Barrow Institute, for keeping me true to the facts and guiding me to a better understanding of phase encoding, and to Doug Goodwin, MD, at Dartmouth-Hitchcock Medical Center for providing these exquisite musculoskeletal cases.

Finally, my eternal gratitude to Dr. Robert Spetzler and the NICU nursing staff at the Barrow Institute for giving me a second chance and inspiring me to write this book.

This page intentionally left blank



1

MR PHYSICS



The detailed images we can create with a magnetic resonance (MR) scanner using magnets, wires, and electricity are just as magical as X-ray images were for a previous generation. Unlike X-rays, however, MR imaging was not the result of a single moment of discovery. It represents the convergence of discoveries in the fields of physics, chemistry, and mathematics that occurred over the course of nearly 200 years.

In fact, it is hard to establish a single date for the beginning of the story of MR imaging once we acknowledge that the permanent magnets found in compasses have been used by sailors for over 2000 years. Nevertheless, a good starting point is that moment in 1820 when Hans Christian Oersted noticed, during a demonstration of electricity, that the needle of a nearby compass jumped when he attached a wire loop to a battery (Figure 1.1). This was the first reported evidence that these two powerful forces were linked. A decade would go by before Michael Faraday demonstrated the reciprocal effect of magnetic induction when he found that he could generate electrical current in a wire simply by moving a magnet back and forth in its vicinity (YouTube¹ video “Electricity MRphysics PennPhysics”). Throughout this overview of MR physics, you must always keep this interaction between magnetism and electricity in mind to understand how MR images are created.




Figure 1.1 This arrangement of a battery, wire, and compass can be used to demonstrate the interaction of magnetism and electricity.

¹ YouTube video “Electricity MRphysics”: <http://www.youtube.com/watch?v=zVz1MGSTLzQ>.

Magnets

Strong magnets are an essential component of all MR scanners, and in the decades since the introduction of clinical MR, all three types of magnets, i.e., permanent, resistive, and superconducting, have been used for imaging. There have even been reports of imaging using the Earth's magnetic field, but this is closer to a parlor trick. Current medical imaging depends on magnets thousands of times more powerful, and the trends in scanner technology suggest that bigger is better. The reason such powerful magnets are utilized for MR imaging is that they are necessary to align the nuclei of hydrogen, which themselves behave like small magnets, and the stronger the magnet, the larger the effect on these atomic magnets.

Permanent magnets, which can be as humble as iron ore, have a very special property called *ferromagnetism*. Iron will spontaneously form small, discrete regions, called *domains*, where the magnetic poles of iron atoms fall into uniform alignment. The larger the domain, the stronger the magnetic field, with the logical extreme of having all the atoms forming a single domain. If iron is exposed to an extrinsic magnetic field, the small domains can coalesce, increasing their size and therefore their magnetic force, and will then stay that way. This explains why a screwdriver that may have been in use for years will suddenly start to pick up loose screws near its tip.

The formation of domains is temperature dependent, and they disappear above a specific temperature called the *Curie point*, named after Pierre Curie.² Iron, at a temperature above its Curie point, becomes paramagnetic, but as it cools, domains will form spontaneously; this makes iron unique and explains its natural ferromagnetic properties. Other, more exotic elements, like neodymium and boron, can be combined with iron to create powerful permanent magnets for applications as ubiquitous as the speakers inside tiny headphones. In all variations, the principle of innumerable small, atom-sized magnetic fields aligned to form a much larger one remains the same. The term *ferromagnetism* indicates that the magnetic property is intrinsic to the material. In contrast, the term *paramagnetism* is used to describe the transient alignment within an external magnetic field that is quickly lost once that extrinsic magnetic field is removed.

The appeal of these permanent magnets for imaging is obvious since they require no electrical current and provide their own magnetic shielding. The large magnetic fields created by the magnets used for medical imaging may extend far beyond the magnet. These peripheral magnetic fields need to be constrained because of rules and reasonable concerns about unwanted interactions with medical devices, such as pacemakers, when the field extends beyond the scanner room (YouTube³ video “Fringe Field MRphysics”). This *fringe field* may not exceed 5G in hallways and public spaces like stairwells or corridors. Large plates of iron or steel can be used for *shielding* of electromagnets to contain their fringe field within the confines of a small room, but permanent magnets do not require this modification (Figure 1.2).

² Pierre Curie was a professor of physics at the Sorbonne in Paris and he was married to Marie. He and Marie together were awarded one-half of the Nobel Prize in 1903 for their studies of spontaneous radiation, but his own early investigations concerned piezoelectricity and magnetism. The name *Curie point* acknowledges his discovery of this phenomenon as well. At the peak of his career, and just three years after he received the Nobel Prize, Pierre Curie was killed by a horse-drawn carriage as he crossed a street in Paris. This death, from a modern perspective, seems a remarkable intersection of the future and the past.

³ YouTube video “Fringe Field MRphysics”: <http://www.youtube.com/watch?v=8qFeUUegfll>.



Figure 1.2 These large iron blocks are positioned within the housing of a superconducting magnet to make the magnetic field more compact. This allows the MR scanner to be placed in a room that may have formerly housed a CT scanner because the fringe field is now smaller.

The high initial acquisition cost of a permanent magnet for an MR scanner should be considered in light of potential savings in shielding, *cyrogens* (liquid gases), and electricity. While permanent magnets excel for some specialized applications like extremity imaging, they are limited for general medical imaging because of their weight, sensitivity to temperature change, and limits to field strength compared with their electromagnet relatives. For this reason, most manufacturers have turned to electromagnets.

The principle of all electromagnets, both superconducting and resistive, is that electricity flowing in a wire, as noted long ago by Oersted, can create a small magnetic field. This magnetic field can then be amplified by winding the wire around an iron core to form a useful electromagnet.

One essential difference between permanent magnets and electromagnets is that the latter can be switched on or off as needed. This is analogous to the difference between a horse and a car; no matter what the circumstances, only the car can be turned off. Movie fans may recall the large junkyard electromagnet that provided some dramatic effect in the James Bond movie *Goldfinger* when it was used to lift a neatly crushed Lincoln automobile into the air. After the metal cube was suspended for a moment, the electromagnet was turned off, dropping the block onto the bed of a waiting pickup truck. Smaller, and hopefully less dramatic, electromagnets can be found throughout our lives. One that is taken for granted almost daily is the automobile starter. It uses electric current from the battery to energize an electric motor in order to turn the engine over. The ability

to turn electromagnets on and off makes them ideal for their role in MR scanners by creating the transient gradient fields used to modify the strong magnetic field.

While an iron core increases the field strength of an electromagnet, clinical magnets must have an *air core* in order to allow space for the patient in the center. The early scanners used field strengths of about 0.15 tesla (T), which seemed remarkably strong at the time but had only one-tenth the strength of the magnetic field used for routine MR imaging today. These early resistive scanners required a continuous supply of electrical current to create a strong magnetic field and, because of the property of electrical resistance, in the process created a considerable amount of heat. This unwanted by-product (remember that electrical resistance is the property that makes toasters and waffle irons hot) had to be removed in order to provide stability to the magnetic field. This was usually accomplished by cooling the magnet with water using a circulator and a radiator, much like the arrangement in a water-cooled car engine. There was also a tendency for the field to drift with temperature variations in the chilled water. The inherent instability and limited field strength of these early scanners, not to mention the high electricity bills, contributed to their surprisingly rapid replacement with superconducting units.

Superconduction describes the unusual state of electrical current without any loss from resistance. This phenomenon was recognized in 1911 by Heike Kamerlingh Onnes, who discovered that the element mercury, cooled by liquid helium to nearly the temperature at which atomic motion stops (-459°F), loses any resistance to electrical current.⁴ While based on the same fundamental principle as other electromagnets, superconducting magnets can achieve higher magnetic field strengths without the power demands or heating common to all resistive magnets. The wire used in these superconducting electromagnets is made of an alloy of niobium and titanium interwoven in a complex arrangement with a conventional copper conductor and covered by a thin insulator. This highly specialized wire is necessary because the superconducting wire alone is not sufficient to handle the electrical current in the event of a sudden loss of cooling.

The magnet's design and construction also require a considerable degree of sophistication. The wraps of superconducting wire must stay in place as the magnet is energized, and the layers of insulation and copper must be thin enough to allow close proximity of the loops of wire while being, at the same time, sufficient in thickness for their tasks. Finally, the magnetic field must be both strong and uniform, a condition described as *homogeneity*. Some minor imperfections in the field can be corrected using *shim coils*. These are small electromagnets within the larger magnet that can be energized as needed. The wire is finally bathed in liquid helium to create the unique state of superconduction (Figure 1.3).⁵

4 It is no coincidence that Onnes was also the first to generate liquid helium (1908). He received the Nobel Prize for his investigations into this realm of the supercold in 1913.

5 The dewar container seen in Figure 1.3 is named for Sir James Dewar, who studied liquid gases at the turn of the twentieth century. A dewar flask, strictly speaking, is a double-walled glass container that uses an intervening vacuum and a metal coating to provide extreme thermal isolation. While initially designed for the containment of supercold liquid gases, it proved excellent for keeping coffee warm as well and is now produced commercially and sold as the thermos bottle. As the wonderful rhetorical question goes, "It keeps cold things cold and hot things hot. How does it know?" Sir James Dewar, together with Sir Frederick Abel, also invented cordite, which was used as a powerful explosive and propellant for firearms throughout the British Empire.



Figure 1.3 These many Dewars of liquid helium will be used to fill up the superconducting magnet of a 1.5 T MR scanner.

At the time of the introduction of these large superconducting magnets, there were some naysayers who doubted that these technological wonders could be managed in the general medical environment. History shows that their reservations were unfounded. These complex, powerful magnets make up the largest and in some ways the most important component of any MR scanner.

Magnetic fields are measured using the standard unit *gauss* (G), named after the mathematician Carl Friedrich Gauss. The most familiar and ubiquitous magnetic field we experience, the Earth's, measures about 0.5 G, while most refrigerator magnets fall in the 500–1000 G range. Because of their very strong magnetic fields, medical magnets are usually measured using a larger unit, much like pounds and tons, called a *tesla* (T), with 1.0 T equal to 10,000 G.⁶

Recently, clinical scanners using magnets of 3.0 T and research scanners with even stronger magnets have become commonplace. While there are some theoretical and practical limitations to medical imaging at 3.0 T, the motivation behind this migration to higher field strengths is to maximize the signal by recruiting a larger proportion of the available hydrogen nuclei in the body for imaging. The improvement in the signal-to-noise ratio that follows can in principle be used to provide some combination of better resolution, faster imaging, or thinner sections. There is a price to be paid, and not just financially, for these larger magnets. With the increase in magnetic

⁶ While Nicola Tesla was a brilliant inventor in his own right, and a contemporary of Thomas Edison, he never enjoyed the same recognition as Edison. Tesla surpassed Edison in the end, however, with his prediction that alternating current would prove to be better for local power delivery than Edison's choice of direct current.



Figure 1.4 This is a picture of the 1.5 T superconducting magnet that was installed at Hershey-Penn State University Medical Center in 1985. This cylinder contains the superconducting wire and functions like a giant Dewar by providing thermal insulation for the liquid gases. Note the metal plates on the wall and floor that provide RF shielding for the scanner. This isolates the scanner from intruding radio waves that may be detected by the sensitive antenna and projected onto the image.

field strength comes greater power deposition in the body of the patient, decreased T1 contrast, increased risk from metal in the room, and larger chemical shift effects. For some applications, however, like magnetic resonance angiography (MRA) or detailed structural imaging, the trade-off appears to be acceptable.

A superconducting magnet, once energized with electrical power, requires no additional current. For this reason, a superconducting scanner behaves more like a permanent magnet than an electromagnet because it is *always on*. This is also unlike other medical imaging devices that may have an outwardly similar appearance, like computed tomography (CT) scanners, but can be turned off at night.

While the notion of a powerful field with no power costs is very attractive, the energy cost of superconduction comes with the replenishment of liquid helium. This supercold liquid surrounds the wire within a strongly insulated container (Figure 1.4) but boils off continually in spite of its high degree of thermal isolation. Replacing the liquid gas requires access to this container through a special apparatus called a *cold head*, installed at the top of the scanner's magnet (Figure 1.5). The time between helium refills can be extended with the use of a helium pump to regenerate the liquid gas. It is this device that provides that incessant “puckata-puckata” sound that is always audible, even when the scanner is not being used for imaging.




Figure 1.5 Frost forms on the *cold head* at the top of the MR scanner where moisture in the air freezes as it meets the extreme cold due to the liquid helium within. This device provides access to the magnet for the helium refills.

Precession

Magnetic resonance imaging is possible only because hydrogen is intrinsically magnetic and, at the same time, abundant in the human body. The magnetic properties of hydrogen give it polarity, i.e., a north pole and a south pole like those of the Earth or any bar magnet. And, as we have come to expect with a compass needle, these magnetic poles of the hydrogen atom tend to align when they are exposed to the strong external magnetic field of the MR scanner.

The nucleus of a hydrogen atom consists of a single positively charged proton that, as predicted by Pauli in 1924, is constantly spinning. This spin of an electrically charged particle creates a very weak magnetic field that forms, in principle, much like the magnetic field created by moving charges in a wire (see the discussion on page 2). While the composite magnetic field of even hundreds of hydrogen nuclei would be too small to detect, the magnetic field created by the alignment of a small fraction of the trillions of hydrogen nuclei in the human body becomes sufficient to be measured indirectly.

This property of spin not only contributes to the magnetic field of protons, it also gives hydrogen protons the quality of *angular momentum*. Common to moving things, with momentum comes an aversion to change. In the special case of angular momentum, a force on the axis of spin creates torque and movement occurs, improbably, in the direction of the torque that is perpendicular to the displacing force. This results in a special type of movement in a spinning object called *precession*. This is easier to picture with familiar objects than to explain, so consider the toy top or gyroscope (YouTube⁷ video “Precession MRphysics”).

The initial spin gives the top a very high rotational frequency along its long axis. As this axis is pulled down by gravity, the top demonstrates a lazy, large rotational motion around the bottom pin.

⁷ YouTube video “Precession MRphysics”: <http://www.youtube.com/watch?v=V8F-KLhrtTE>.

This second rotational movement is precession. It is important to consider, and is not at all obvious, that the frequency of this precession is directly proportional to the strength of the displacing force. In the case of the top, the displacing force is gravity, so we should expect that the frequency of precession of the top would be lower on the Moon and higher on Jupiter because of their respectively lower and higher gravitational attraction compared with that of the Earth.

We can think of the spinning hydrogen nucleus in the same way as the top, but the displacing force, instead of gravity, is the scanner magnet. And, just like the effect of gravity on the top, the stronger the magnetic field, the higher the frequency of precession. This relationship was reported by Sir Joseph Larmor, who predicted that the exact frequency of nuclear precession could be calculated as the product of the strength of the magnetic field experienced by the hydrogen nuclei and a “gyro-magnetic constant” specific for the nucleus of interest. The symbolic representation of this relationship is called the *Larmor equation*.

$$\omega_0 = \gamma B_0$$

ω_0 = frequency of precession

γ = gyro-magnetic constant specific for the atom

B_0 = the magnetic field

Scanners can be described by their magnetic field strength or by the frequency of precession of the hydrogen nuclei within the magnet. For example, a 1.0 T scanner can be described as a 42 MHz scanner and a 1.5 T scanner as a 63 MHz scanner based on the *Larmor frequency* of hydrogen at each field strength. This tightly linked relationship of field strength and frequency is essential to for your understanding of MR physics and frequency encoding in particular.

Resonance

Resonance: A vibration of large amplitude in a mechanical or electrical system caused by a relatively small periodic stimulus of the same or nearly the same period as the natural vibration period of the system.

Merriam-Webster Online Dictionary

You need look no further than a backyard swing set to find an everyday example of resonance (YouTube⁸ video “Resonance MRphysics”). We intuitively recognize that lifting the passenger on the swing higher and higher requires precise timing of the push. At the same time, we are aware that the incremental force of the push is trivial compared with the energy that would be required to toss the passenger into the air with one push. Resonance describes this process of applying small but optimally timed pushes that match the natural frequency of the swing.

We pick up the story of MR imaging again with the discovery of nuclear magnetic resonance (NMR) by Isidor Isaac Rabi.⁹ Using a device that created a narrow stream of lithium chloride molecules, first developed by Otto Stern and Walter Gerlach, Rabi found a remarkable interaction of

⁸ YouTube video “Resonance MRphysics”: <http://www.youtube.com/watch?v=tpl2skw0TZ4>.

⁹ Rabi is an inspiring example of a somewhat lackluster young student who developed into a brilliant researcher. He is also credited for his work in the years just before receiving the Nobel Prize for the development of radar, which proved to be critical to the Allies’ victory in World War II. While Rabi’s insights laid the groundwork for NMR and MR imaging, strangely he has not received widespread recognition in the medical community.

these charged particles with radio waves and a magnetic field. On the basis of these experiments, he received a Nobel Prize in 1944 for his discovery that the interaction of these charged molecules with the magnetic field and RF energy was due to nuclear resonance.

This principle, established in Rabi's physics laboratory, forms the basis for MR imaging, but four decades would pass before the last pieces of that puzzle would fall into place. In the MR scanner, the magnetic polarity of the many hydrogen nuclei (protons) becomes aligned to some degree by the strong external magnetic field inside the bore of the scanner. However, these nuclei will not just snap into place inside the scanner, all pointing north and south like compass needles. Because of their angular momentum and the displacing force of the magnet, these hydrogen nuclei will precess at a frequency predicted by the Larmor equation. While they are stationary but precessing at the same time, you can think of them like trout swimming in place in a stream. While the trout stay roughly aligned with the current, they are always moving, albeit in one preferred location.

While in practice we use an RF pulse to add additional energy to this natural movement of the hydrogen nuclei, it may be easier to picture the resonant force as a magnet moving on a track around the bore of the magnet. Just like the pushing of a swing, the pace of this smaller magnet must match precisely the natural frequency of the hydrogen precession in order to satisfy the necessary conditions for resonance. Not all but some significant fraction of the many trillions of hydrogen spins will be captured by the field of this moving magnet. As the nuclei become energized by this attraction, their magnetic axis will begin to point away from the long axis of the scanner bore and toward this imaginary magnet on the track until they point perpendicular to the large scanner magnet. At that moment, the individual magnetic poles of all these protons will also be pointing in the same direction, i.e., toward the moving magnet, which maximizes the strength of their composite field. The RF transmitter effectively acts the same way, imparting energy to the ordered but precessing hydrogen nuclei. The hydrogen protons, with their magnetic poles pointing in the same direction but perpendicular to the strong field, can be described as having experienced a 90 degree RF pulse and are now in coherent phase.

I must warn you to avoid unclear thinking due to mixing the image of single hydrogen spins with the composite magnetic field that reflects the sum of many of these spins. At this point, let's change our focus; let's stop thinking about single hydrogen protons and think instead of the net force created by the sum of many individual spins. You can now think of this composite field as a single spinning bar magnet pointing in the direction perpendicular to the long axis of the scanner bore (Figure 1.6), and we know that a moving magnetic field will induce an electrical current in a wire inside the scanner. We can call that wire an antenna and, using it, we can register a rising and falling electrical current as this bar magnet (composed of countless hydrogen nuclei) spins nearby. This waveform will have the same frequency as the hydrogen Larmor frequency and an amplitude determined by the total number of hydrogen nuclei spinning together.

This is precisely how the NMR signal is detected and it represents the state of the art in 1952, the year that Felix Bloch and Edward Purcell received their Nobel Prizes for NMR spectroscopy. Their work brought this phenomenon of NMR out of the laboratory and to the level where it could be used as a laboratory tool for investigation of the composition of solids, liquids, and gases. While MR imaging almost always uses hydrogen nuclei, NMR can be used for a wide range of elements so long as their nuclei have an odd number of charges, such as phosphorus or sodium.

Chemists were entranced by this technique that allowed them to examine nondestructively small chemical samples using only their radio waves. Prior to the discovery of NMR, it was not so easy




Figure 1.6 The net magnetic force of the collectively aligned hydrogen nuclei after a 90 degree RF pulse acts very much like a spinning bar magnet. The reason we can recover signal from a human body comes back to the principle that a moving magnet will induce current in a nearby wire. This spinning magnet will then induce a fluctuating current in a nearby wire (antenna) as it spins towards, and then away from, the antenna. The frequency of the resulting wave can be predicted by the Larmor equation and is based on a gyromagnetic constant and the magnetic field strength the hydrogen spins experience.

to determine the constituents of chemical compounds and it usually involved some mix of torches, solvents, and crystals. What makes NMR so powerful is that the waveform recovered, and then converted to a spectrum, is altered by chemical bonds and other elements in the vicinity of the stimulated atom. To minimize confounding factors, great emphasis was placed on creating a completely homogeneous magnetic field. During the many years that scientists worked exclusively with NMR spectrometers, however, there were no recorded attempts to use it for imaging.

Imaging

The pieces were all in place at this point for the transition from NMR spectra to MR imaging. Diverging from the usual practice of studying chemical compounds with NMR spectrometers, Dr. Raymond Damadian used one to test his hypothesis that the NMR character of liver tumors differed in a predictable fashion from that of normal liver. His experimental proof opened the conceptual door to medical uses for NMR, and he reported the results of his research in the journal *Science* in 1971. Paul Lauterbur in 1973 reported in the journal *Nature* his novel technique for modifying NMR hardware in order to create images,¹⁰ while Peter Mansfield, across the Atlantic, was developing pulse sequences that could be used for imaging. Just 10 years after the appearance of Dr. Lauterbur's paper in *Nature*, clinical MR scanners were being installed throughout the United States, but 20 years more would go by before the Nobel committee decided to award Lauterbur and Mansfield their prizes. It was in these early days of imaging that the terminology changed subtly from *nuclear magnetic resonance* to just *magnetic resonance*, with the thought that the *n* for *nuclear* in the name might be confused with *nuclear medicine* or alarming for some patients.

Dr. Lauterbur's innovation was to modify the strong NMR magnetic field in a uniform and predictable fashion using smaller magnets called *gradient coils*. With this addition, the magnetic field could then be varied from front to back, side to side, and top to bottom. Since we know that the

10 It is of some comfort to researchers everywhere that his paper, later ranked among the most significant scientific papers ever published, was initially rejected.

magnetic field strength predicts the frequency of precession (Larmor equation), and if we know where the field is stronger or weaker, we could in principle establish the location of the returning signal based on its frequency alone. This concept of modifying the magnetic field to establish the spatial location of the returning signal is basic to all MR imaging (Figure 1.7).

This notion of diagnostic MR imaging may seem mundane now, what with MR scanners in trucks and shopping malls, but it was remarkable at the time. It may also come as a surprise that the poor resolution of these early images, compared with the detailed images provided by the more mature technology of CT, led some imagers to downplay the impact of this new imaging tool when the first clinical scanners were introduced (Figure 1.8). Now let's consider exactly how one might use these magnetic gradients to make detailed images of the body. It is easier to explain localization of signal with a simpler object than the human body, so let's start by imaging some wine bottles in a wine rack (Figure 1.9). Let's see if we can use the MR scanner to determine which slots are empty and which slots are filled in a rack with three slots across and three side (3×3).

While imaging can be done by creating a "sweet spot" where resonance occurs, but nowhere else, and moving the rack around it, this approach is not practical. Nevertheless, it is effective and was used to create the first published human body image with MR. In all current clinical MR units, the images are created by using gradient coils that are turned on and off in a sequential pattern. The order and timing of their activation can be most easily depicted using a pulse sequence diagram with a different timeline for each piece of the scanner hardware. This is not much different from a musical score that a composer uses to indicate when the violins come in, the flutes drop out, and all the while the drums play on (Figure 1.10).




Figure 1.7 This is a photograph of a gradient coil assembly prior to installation in an early 1.5 T unit. (The wooden frame supports the coils in transit.) These electromagnets will be placed within the open center of the large superconducting magnet where they will surround the patient. The coils are energized in pairs so that the magnetic field inside the scanner bore may be modified in a roughly linear fashion from top to bottom, side to side, and head to toe.




Figure 1.8 This axial MR scan of the chest was obtained in 1983 on an early 0.15 T resistive MR scanner at the University of Pennsylvania. While not up to the standards of current production MR scanners, it was remarkable at that time and still provided the diagnosis of a lung tumor in this patient. That early clinical scanner, however, provided only one slice at a time, placing a premium on the imager's knowledge of surface anatomy.




Figure 1.9 This drawing indicates the position of three bottles of wine that are arranged in a 9-slot wine rack (three by three) that resides inside the bore of an MR scanner.




Figure 1.10 This is a much simplified pulse diagram for a spin echo MR scan. The different action lines can be thought of just like a music score that indicates when the different instruments come in, for how long, and at what volume. The bottom line indicates when the signal will be detected at the antenna in relation to the other actions.

Slice Selection

Using our wine rack example, we can take the first step toward spatial localization, called *slice selection*. Since most images for diagnosis are displayed like so many slices of bread, it is reasonable to acquire the information in the same fashion. It is worth noting that since this is done entirely with the use of electronic gradients and without any fixed hardware (like an X-ray tube spinning in a CT scanner), any and all planes of acquisition are possible in an MR scanner.

Let's position our wine rack inside the scanner so that the corks of the bottles are pointing away from us as we look inside the scanner bore. We will name the gradient that runs from the top to the bottom of the bottles the *Z gradient* direction. When this gradient is turned on, it modifies the previously uniform strong magnetic field so that it is now higher at the top of the bottles and lower at the bottoms. This leaves a band right in the middle where the hydrogen spins experience no more or less than the native strong magnetic field. It is fairly straightforward, then, to use the radiofrequency (RF) transmitter to impart energy at the frequency that matches the precession frequency of that band in order to move the hydrogen spins into the plane perpendicular to the scanner bore, called a *90 degree pulse*. While all the hydrogen spins inside the scanner will experience this same RF pulse, only those in that narrow slice will resonate with the RF pulse and absorb energy. In other parts of the rack, where the conditions for resonance are not fulfilled, those hydrogen nuclei become effectively invisible. In this fashion, we can get signal back from just one slice through the bottles at a level midway from the neck to the base.

This signal that we recover with the antenna tells us that there are some wine bottles in that slice, but it provides no information about their location. The amplitude of the returning wave tells us indirectly how many bottles are in the rack since more wine bottles mean more hydrogen nuclei and therefore more signal.

The electrical current or signal that registers at the antenna may be displayed as a waveform with a frequency and an amplitude. While it is simple to graph this signal with elapsed time on the horizontal axis and amplitude on the vertical axis, when multiple frequencies are mixed together, it is easier to understand when we *Fourier transform* this into a spectrum, i.e., a graph of amplitude on the vertical axis and frequency on the horizontal axis (Figure 1.11).




Figure 1.11 This illustration demonstrates the aspects of slice selection. By applying a magnetic gradient that varies from the bottom to the top of the bottles, only one slice in the middle will be at the correct frequency for resonance with the RF energy. The pulse diagram shows that just one gradient, the Z or slice select, is used for this step. The waveform we recover with the antenna from the water (actually hydrogen nuclei in the wine) can be displayed as a waveform or, after Fourier transformation, a spectrum with amplitude on the vertical axis and frequency on the bottom. Note that only one frequency is present, at the Larmor frequency of water determined by the magnetic field (1.0 T).

Frequency Encoding

Now at least we have reduced the problem of assigning spatial location to the signal to only two directions, up or down and left or right. We can solve that problem using the two remaining gradient pairs. The left-right location can be determined by turning on the side-to-side gradient (X gradient) during the collection of the signal after the 90 degree pulse. By using a gradient pair to modify the strong magnetic field from side to side, the field becomes stronger on one side than the other. Since the frequency of precession is linked to field strength, and supposing that the field is now higher on the right than on the left, we would then know that any high-frequency signal we recover at the antenna comes from the right side of the scanner and the low-frequency signal comes from the left side. This assignment of spatial location based on frequency alone is called *frequency encoding*. You are already familiar with the concept of spatial location linked to frequency at a more intuitive level involving musical instruments. When you hear a low, rumbling note coming from a piano, you can envision that the player hit a key on the left side of the keyboard without actually seeing the keyboard. You also know, when figuring out how to play a song on the piano, that if the note you strike is too low, you will move to keys on your right, where the frequency of the notes is higher. This predictable spatial–frequency relationship between the keyboard and audible notes is a form of frequency encoding.

This process of frequency encoding brings us to Fourier transformation. This mathematical function is perhaps most easily understood using a biologic example. The function of the cochlea in the ear is to convert the multiple complex frequencies of sound waves that enter the ear canal into discrete neural signals that indicate which frequencies are represented. For example, the low-frequency sound of a tuba stimulates a specific part of the cochlea to provide electrical signals that then register in the cortex of the brain. This function of the cochlea allows us to recognize the unique sound of this and any other musical instrument. The conversion of sound into a signal with an amplitude that also indicates the volume of that sound captures the basic concept of the Fourier transformation.¹¹

Using Fourier transformation, we can convert a complex signal composed of multiple frequencies, detected by the antenna, into its component frequencies, much as the ear and cochlea function together. Fourier transformation proves to be a critical process in MR image reconstruction to the point that the matrix choices we use for imaging (128, 256, 512) are dictated by the requirements of fast Fourier transformation (FFT).

11 This mathematical function is named for Joseph Fourier, and the story of his life may cause you to reconsider any stereotype you might have of a bookish mathematician. The 9th of 12 children who was orphaned at age 9, he nearly went to the guillotine for his role in the French Revolution, became a scientific advisor to Napoleon, and was for a time stranded in Egypt after the British admiral Horatio Nelson eviscerated the French navy. His mathematical work directed at calculating heat transfer proved to be ideal for MR image processing 150 years after he died. He is also credited with the observation that the atmosphere functions as a heat-retaining blanket around the Earth. His contributions are certainly in the forefront in this millennium.

By turning on this second gradient while we recover signal from the one slice through the middle of the bottles, instead of recovering a single frequency from the hydrogen protons at their Larmor frequency, multiple frequencies will be represented in the signal, depending on the location of the bottles. So, at this point, we should be able to figure out which columns (left, center, or right) in the rack have at least one bottle of wine. If there is more than one bottle in any column, we should be able to guess this because the signal at that frequency will have increased amplitude, since there will be more hydrogen protons producing the signal (Figure 1.12). Now we have two dimensions of signal assignment, and we will leave the third coordinate to the technique of *phase encoding*.




Figure 1.12 In this illustration we see the aspects of frequency selection. The gradient field is applied from side to side so that the wine bottles on your right will experience a stronger magnetic field than those on your left. Now, instead of recovering signal at just one frequency at the antenna, we see two different waveforms since the hydrogen nuclei in the two bottles on your right provide signal with both a higher frequency and amplitude since there are more hydrogens with two bottles. After the Fourier transform of the recovered signal, note these two peaks depict those differences in frequency and amplitude of signal.

Phase Encoding Explained

How can we use the phase of the returning signal to indicate if the bottles in any column are at the top, bottom, or middle of that column? And what is phase anyway?

Phase encoding is not a single-step localization process like the first two. It also is the most difficult to explain because it is more of a mathematical process than an intuitive concept. But, rather than say that “something really amazing happens,” let’s examine it carefully since some basic understanding of this principle will help you to better understand the artifacts that follow. It is also helpful to consider at the outset that frequency and phase can be considered interchangeable. When we see a waveform moving up and down along the time axis, this same information can be represented by a measure of its relative phase at specific points in time. This is somewhat analogous to the way music can be stored in an analog form that captures the wave itself (vinyl records) or by sampling that same wave and then storing it in a digital file (CDs). That later process is called *analog-to-digital conversion*, and the number of data points used to measure the sound waves reflects how accurately the sound will be reproduced. This principle of digital sampling of a sound waveform is essentially how music is stored as a digital file in your iPod.

It is worth pointing out again that at the moment the RF pulse is turned off, the energized hydrogen protons are not only rotating perpendicular to the large magnetic field, they are also pointing in exactly the same direction (*coherent phase*). *Phase* is the term used to describe the orientation of the magnetic axis of each hydrogen proton. When the hydrogen protons are considered as a population, and when the magnetic poles of these protons align, their net magnetic force is amplified. However, even when they are energized into the 90 degree plane, the loss of coherent phase among the protons will result in a lower net magnetic force.

When these hydrogen nuclei are considered as a population, there will be both constructive and destructive interactions due to phase on net magnetization. For example, consider a single voxel where there are 1 million energized hydrogen spins all rotating in the transverse plane with coherent phase. In this circumstance, the net magnetic force, which will induce current in the antenna, will be quite large because the small magnetic field of each proton is reinforced by its 999,999 neighbors. This is reminiscent of the property of domains we reviewed in the discussion of permanent magnets. However, if the magnetic axes of these same protons are spread randomly in every direction, even though these protons are spinning in the transverse plane, their net magnetic force will be zero because of cancellation of this force. This happens because protons with opposite phase will cancel each other with respect to their composite or net magnetic field. This is also an illustration of the difference between the dynamics of a single hydrogen proton and that of a whole population.

When we turn the third (Y) gradient on in order to make the magnetic field stronger at the top than at the bottom of the columns, the frequency of precession of the hydrogen protons at the top of the column will become greater than the frequency of precession of the protons at the bottom. This also means that there will be a *phase shift* across the column as the protons speed up or slow down relative to the middle row, where the phase shift is zero. The moment the Y gradient is turned off, the hydrogen nuclei in each column will go back to spinning at the same frequency, but they retain the phase shift acquired while the Y gradient was on. The longer the Y gradient is left on, the larger this phase shift will be from the top to the bottom of the column.

The extent of this phase shift can be described in terms of the relative spread of the phase of individual protons from the top to the bottom, i.e., 90 degrees, 120 degrees, 180 degrees, and so on. If there is no shift at all, this is called the *coherent phase*. Using the gradient, the total spread of available phase shifts is 180 degrees to -180 degrees for a total of 360 degrees. Keep in mind that any phase shifts beyond 360 degrees will become ambiguous since 370 degrees would look the same as 10 degrees to the scanner, just as 0900 hours (military time) looks no different than 2100 on a wall clock.

If at this point we could determine the frequency and absolute phase shift of the signal arising from the protons that we recover at the antenna, we might have the information we need to create an image. Unfortunately, the problem of spatial location using phase encoding proves to be more difficult just because the scanner cannot determine phase directly. Because of the phase interactions of populations of hydrogen protons, however, phase can be measured indirectly, but it requires repeated applications of a varying Y gradient in order to decipher it.

Coming back to the wine rack in the scanner, since we know there are two bottles in the far right column and one in the far left column, we must now determine if they are at the top, middle, or bottom of the columns. To illustrate how this problem can be solved using phase information, let's illustrate the effects of phase change on net magnetization or signal using a belt and two clothespins. Each pin in this example will represent the phase of one of two hydrogen protons. If the belt is placed flat on a table and two pins are placed on the same side of the belt, that arrangement would resemble the state of phase coherence. In that condition, the magnetic fields of these two protons would add together nicely (Figure 1.13).

Now, let's twist the belt buckle one full turn, creating a 180 degree phase shift. If we look at the belt from the single view to replicate the limitations in measuring phase, we would see it as the "end-on" view in Figure 1.14. As you can see, with coherent phase, you cannot tell where the clips reside on the belt. However, with a 180 degree twist (phase shift) in the belt, things become clearer. In this simple example, if the clips point in opposite directions seen from our end-on view, then we could predict that they must be at opposite ends of the belt (Figure 1.15), and if they point in nearly the same direction, they must be very close together indeed.

The scanner, of course, has no end-on view, but it does measure signal amplitude. Consider what might happen to the amplitude of recovered signal at the antenna with these two extremes of phase shift. Coherent phase should produce maximum signal no matter where the bottles reside in the rack. The lowest signal amplitude, or even no signal at all, will be recovered using a 180 degree phase shift if the bottles in the rack are at the top and bottom slots due to cancellation of the opposite phase directions. If, on the other hand, the bottles are in adjacent racks, a small drop in signal will be evident after the 180 degree phase shift. This is how signal intensity alterations after these known phase shifts provide some degree of spatial information. Multiple phase shifts with collection of accompanying signal will be required to determine the exact spatial location of the returning signal, however.

In practice, this requires varying the Y gradient for every row of information in the phase direction of the matrix. Think of expanding this problem of spatial localization from finding three bottles in a wine rack with nine possible slots to an image with a matrix of 512×512 and subtle shadings of signal, and you can better appreciate the calculations and complexities involved in applying this concept of phase encoding to MR imaging (Figure 1.16).




Figure 1.13 This photograph of two clothes pins on a belt are meant to illustrate two individual hydrogen protons with coherent phase, rotating at the same frequency, but at different locations in the phase direction. If we were to briefly apply a magnetic gradient from top to bottom the clips on each belt would experience a slightly different magnetic field strength that would then alter their rotational frequency. The greater the distance between the clips, or hydrogen spins, the larger the difference in their frequencies.




Figure 1.14

180 phase shift






Figure 1.15


Phase Encoding



gradient orientation



before gradient



after gradient application
180° phase shift from top
to bottom of rack

Figure 1.16 In this illustration we see the Y gradient applied so that the magnetic field is briefly higher at the top than at the bottom of the column. This will make the protons at the top precess faster than those at the bottom and, depending on how long or how strongly the gradient is applied, different magnitudes of phase shift can be created. Before the application of the 180 degree phase shift, and after application of the frequency (side to side gradient), we see the two waveforms of different amplitude and frequency shown in figure 12. When the same frequency gradient is applied after a 180 degree shift of the vertical columns of spins, notice how the amplitude of the higher frequency signal is diminished. This is due to phase cancellation that reflects indirectly that there must be a bottle at the top and bottom of that particular frequency column. In this fashion, by varying the phase gradient and recording the changes in the returning signal, the location of the bottles (or really spins) can then be determined in each vertical columns.