Section 1 Body MRI overview

Chapter 1

Introduction

Introduction

Body MRI is a dynamic, exciting modality. If read carefully, you will find in this small book the essentials to protocol, understand, and interpret abdominal/pelvic MRI. This is not a tome. It is light on physics. This sentence contains the only mention of *k*-space.

The utility of MRI for evaluation of the chest, abdomen, and pelvis has improved dramatically in the past decade due to more powerful scanners, better pulse sequences, and improved coils. It is the test of choice for evaluation of focal and geographic liver disease and the biliary tree. It is also commonly used to evaluate lesions of the kidneys, adrenal glands, and, more recently, the small bowel. With unparalleled soft tissue resolution, MRI has become the gold-standard imaging exam for evaluation of the female pelvis and staging of pelvic malignancies in either gender.

Fundamentally, MRI remains a problem-solving modality. Exams are/should be targeted to a specific diagnostic problem. The goal of MRI interpretation is to put a diagnostic issue to rest with one final test. It is critical that all previous imaging studies, as well as laboratory data and patient history, be thoroughly reviewed before protocoling and interpreting MRI studies. If the purpose of an exam cannot be determined from all available data, the referring clinician should be contacted.

When we attempt to interpret studies in a vacuum, everyone loses. Many clinicians hold fast to the belief that an ultrasound is better than a plain film, a CT is better than an ultrasound, and an MRI is the best there is. Therefore MRI is typically the last imaging stop, the final frontier – nowhere to go from here except biopsy. We need to be as definitive as possible.

Remember, the geniuses are the ones who thought it might be possible to create an image of the inside of the body using a really big magnet – and then made it happen. Imagine if they were asked to write a guide to creating and building an MRI scanner. It would be a heck of a lot more complicated than this book.

The interpretation is the easy part.

How MRI works in a few paragraphs (and T1 and T2 in a few more)

The MRI scanner works by placing the patient in a strong magnetic field, sending multiple radiofrequency (RF) signals into the patient, and then waiting, measuring, and localizing what comes out. The RF signal entering the patient comes from RF coils built into the magnet. The antennae or receiver coil can be the same "body coil" which is also built into the magnet, but we try whenever possible to use a "surface coil" – a coil placed on the surface (usually the belly) of the patient. (We call them coils, really they look more like foam pillows.)

Now let's tackle the basic MRI terms.

TE stands for echo time. It is the time we wait after sending the RF signal into the patient before we "listen" to the signal coming out. It is measured in milliseconds.

TR stands for repetition time. It is the time between sending the first RF signal into the patient and the next one. It is also measured in milliseconds.

That's it! *TE and TR are the two basic parameters that we change at the scanner console to determine what the image will look like.*

What about T1 and T2? These are *not* parameters we vary on the magnet, but *are innate physical properties of matter* that describe how tissues behave in a magnetic field. T1 is the time for tissues to recover longitudinal magnetization in a magnetic field and T2 is the time to lose phase coherence. Please now feel free to forget the last sentence (if you haven't already).

But, *do not forget that T1 and T2 are innate physical properties of matter*. We are born with them. MRI generates valuable images because there is tremendous

natural variation in the T1 and T2 of different tissues; the result is magnificent soft-tissue contrast.

There is much more variation in the T1 and T2 times of various tissues than there is in density, which is what we measure and "see" on CT. That's why a CT of the knee shows the bone as bright, the fat as black, and all the muscles and tendons as a gray mush, whereas MRI can separate each soft tissue from its neighbor.

By placing the patient in the magnet and varying the TR and TE to specific parameters (figured out by MRI physicists, not you and me) we can generate an image that emphasizes the T1 properties of the tissue. This is called a T1-"weighted" image. By changing the TR and TE, we can instead emphasize the T2 characteristics of the tissue. This image is then called T2-"weighted". By learning what different tissues and pathologies look like on T1-weighted and T2-weighted images, we can not only detect abnormalities, but also characterize the lesions that we detect.

It is because MRI has more variables than CT that we can be more specific. (CT has essentially only two variables – attenuation and enhancement.)

On T1-weighted images, fluid is dark. On T2-weighted images, fluid is bright. Thus, fluid in the bladder or CSF space is generally dark on T1-weighted images and bright on T2-weighted images. If you are looking at a set of images and don't know if they are T1- or T2-weighted, first find a fluid containing structure in the image (bladder or CSF space.) If it is dark, you are looking at a T1-weighted image, if it is bright, you are looking at a T2-weighted image (Figure 1.1).

Of course, "weighting" is not all or nothing. Most modern pulse sequences have a combination of T1- and T2-weighting. The beginning reader should think of T2-weighted sequences as fluid-sensitive – that is, *if fluid is bright, consider it a T2-weighted image*. Don't lose the forest for the trees.

So. Fluid is bright on T2-weighted images. What is bright on T1-weighted images?

The list of things that are bright on T1-weighted images is one of the great lists in radiology – great, because it is short. This is not the differential diagnosis of white-matter disease in the brain or, my favorite, a lecture I (twice) attended entitled "The twenty-one mimickers of pneumonia on chest x-ray." There are only four things:

(1) *Blood* Blood is very important. Don't miss blood.
One nice thing about body imaging is that the signal of blood products within the abdomen/pelvis does not necessarily follow the complex, mnemonic-ridden sequence it does in the brain. Don't try to "age" blood products in the abdomen/pelvis. Just be happy you detected them. Note also that flowing blood within vessels may be variably T1-bright as well.
(2) *Fat* Fat is bright on T1-weighted images. That is, as long as there is no fat saturation.
(3) *Protein* Various proteins are bright on T1. This includes melanin. Metastatic melanoma can be bright on T1.
(4) Enhancement with gadolinium-based contrast agents.

Fat is inherently bright on T1-weighted images and may be variably bright on T2-weighted images. Many sequences may be performed with or without fat saturation. Fat saturation is a parameter modified on the scanner to make fat lose its bright signal. Chemically selective fat saturation causes *macroscopic* fat to lose signal and become black. This is considered proof beyond any reasonable doubt that the tissue contains fat.

One final topic of unnecessary confusion. Pulse sequences are proprietary, they are developed and owned by the vendors, and therefore the vendor gets to name them. Commercially available scanners, from whichever vendor, all have essentially the same basic array of sequences for clinical use, but will name them differently. (This is analogous to generic and brand-name drugs. The generic name is most important, it describes what the pulse sequence actually is. The brand name is designed to sound good.)

For example, all vendors have a T1-weighted fat-saturated sequence used before and after gadolinium administration. Siemens calls it VIBE, General Electric liked the sound of LAVA, Philips named theirs THRIVE. In this section, we will list the sequences by "generic" names with "brand" names in parentheses.

Each sequence in each exam should be performed for a specific reason. *When looking at a series of images, ask yourself specifically what that sequence adds to the overall exam*. If you don't know why a sequence is performed or what it adds to the exam – ask someone else. *If no one knows, it should be eliminated*.

So, without further ado, commonly used sequences. . .

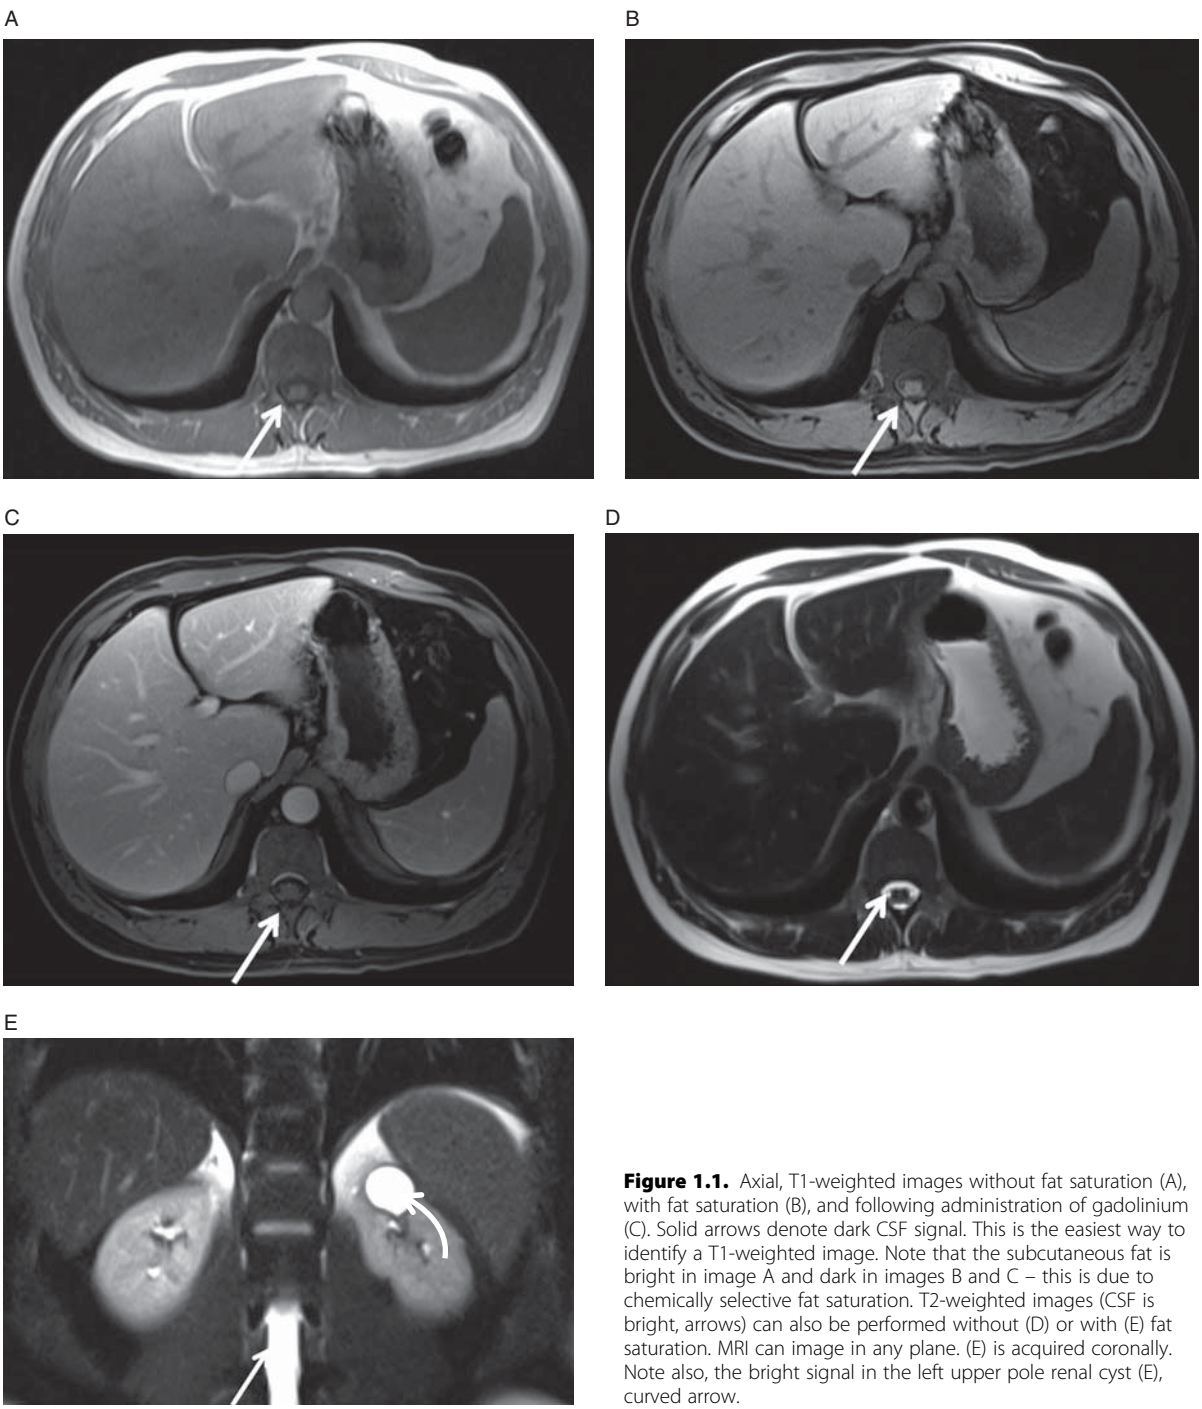

Figure 1.1. Axial, T1-weighted images without fat saturation (A), with fat saturation (B), and following administration of gadolinium (C). Solid arrows denote dark CSF signal. This is the easiest way to identify a T1-weighted image. Note that the subcutaneous fat is bright in image A and dark in images B and C – this is due to chemically selective fat saturation. T2-weighted images (CSF is bright, arrows) can also be performed without (D) or with (E) fat saturation. MRI can image in any plane. (E) is acquired coronally. Note also, the bright signal in the left upper pole renal cyst (E), curved arrow.

Commonly used pulse sequences

Many are intimidated by the seemingly never-ending and often-changing array of pulse sequences. Do not fear. In body imaging there is really only T1, T2, and diffusion. That's it.

There is no T3.

Below are the pulse sequences we commonly use as well as a brief description of what they are used for. Since there's only T1, T2, and diffusion, first we will discuss T1-weighted sequences, then T2-weighted sequences and then diffusion. Then we'll stop.

Nearly all the sequences used in body MRI are performed while the patient holds their breath. (There are rare exceptions which will be described later.) Good breath-holding is absolutely critical to generating quality images, particularly post-contrast images which are often the most time-consuming and therefore motion sensitive. It is critical that the MRI technologist coach and coax the patient, offer encouragement (and breaks when necessary) in order to achieve the best possible breath-holds. When the images are clear, the interpretation is easy.

T1-weighted sequences

The original T1-weighted images were spin echo sequences. This is the classic sequence described in radiology physics classes with the 90° pulses, etc. . . . Forget it. Takes too long. Not relevant to body imaging (unless the patient can hold their breath for 12 minutes).

The T1-weighted images we perform now are gradient echo. The physics of this is not particularly important for our needs. But, you should know that they are gradient echoes because you'll look stupid if you ever tell someone they are spin echoes. One of the reasons that gradient echo images are so fast is that we can use very small flip angles – less than 90°. This is the last time we will mention flip angles. In clinical practice, you will never change the flip angle.

We don't just get a T1-weighted gradient echo image. What we actually get is called a dual-echo or chemical-shift image.

In- and out-of-phase/chemical-shift images (General Electric = GRE, Philips = FFE, Siemens T1-weighted images, no fat saturation)

Chemical-shift imaging is also known is in/out-of phase imaging (IP/OOP). These are T1-weighted images without fat saturation acquired with two different TEs. You get both sets of images during the same acquisition. The difference is the TE, i.e. when we "listen" for the echo. Both are T1-weighted, but one set is in phase and the other set is out of phase.

Remember high-school physics? Did you ever make waves with jump ropes in a hallway or in little pools of water?

When two waves collide and their amplitudes are in the same direction (up or down) the amplitudes are additive. If the waves collide when one is up and one is down, they cancel.

The signals coming back to the scanner from the patient are also a wave. The signal is generated by protons, but the frequency of that signal changes based upon the microenvironment of each proton.

A proton bound to water emits a signal with a different frequency than a proton bound to fat. People smarter than you, me, or anyone either of us know figured out the precise frequencies of these signals. Of course, they vary with magnetic field strength, but at 1.5 T waves from protons bound to fat and protons bound to water will be out of phase at a TE of approximately 2.2 ms and in phase at 4.4 ms.

It is important to remember that we are talking about microscopic fat and the signal cancellation is on a per voxel basis.[1] That is, when protons bound to water and protons bound to fat are in the same voxel, their signals will be out of phase if imaged at a TE of 2.2 ms and the scanner will actually receive little signal. This voxel will be dark on the out-of-phase images and bright on the in-phase images (waves additive).

Our solid organs contain a great deal of water and are surrounded by fat. Thus, at the borders of solid organs there are protons bound to water (within the organ) and protons bound to fat (outside the organ). On an out-of-phase image, these "border" voxels have both protons bound to water and protons bound to fat and will therefore lose signal (turn black). This creates "etching" or "India-ink" artifact surrounding the solid organs.

The presence of microscopic fat is extremely useful for characterization of both focal and diffuse disease. If an adrenal mass has microscopic fat, if its signal drops significantly on the out-of-phase images compared to the in-phase images, it is an adenoma (Figure 1.2). If geographic areas of the liver lose signal on out-of-phase images this is diagnostic of hepatic steatosis.

[1] A voxel is a three-dimensional pixel.

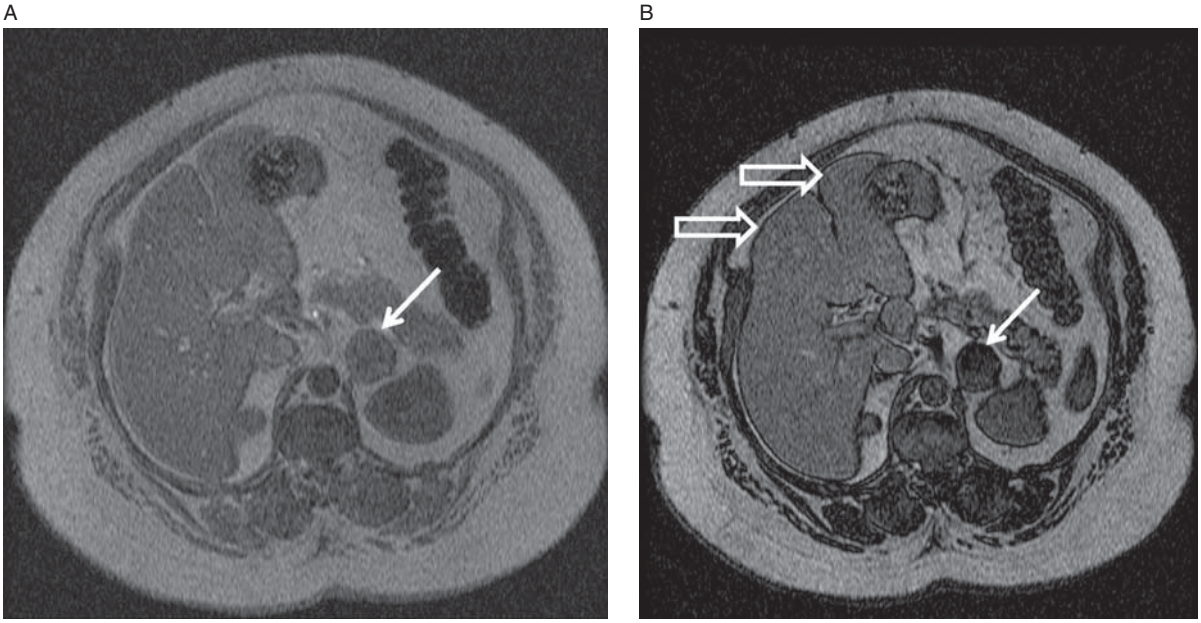

Figure 1.2. Adrenal adenoma. Solid arrows demonstrate marked signal loss in the left adrenal mass on the out-of-phase image (B) compared to the in-phase image (A). This proves that it contains microscopic fat, diagnostic of an adrenal adenoma. Open arrows demonstrate the appearance of "etching" artifact at the interface of the liver (which contains lots of water) and the peritoneal fat on the out-of-phase image.

Now, for the advanced student . . . those striving for mediocrity can proceed to the next section (you know who you are). Tissues that contain microscopic fat lose signal on out-of-phase images. What might lose signal on in-phase images?

Remember that the TE of the in-phase image (the time we wait before we listen for the signal back from the patient) is twice as long as that of the out-of-phase image. Iron deposition causes inhomogeneity in the magnetic field and inhomogeneity causes signal loss. The longer you wait to image, the more signal loss. Therefore, in patients with iron overload, the sites of iron deposition will lose signal as the TE lengthens (Figure 1.3). The degree of signal loss is directly related to the amount of iron in mg/dl.

Remember, in-/out-of-phase images assess microscopic fat. Signal loss occurs on the out-of-phase image when water and fat share the same voxel. To evaluate macroscopic fat, look for a sequence which utilizes a chemically selective fat-saturation pulse (typically, the pre-contrast T1-weighted images – see below). Signal drop on a chemically selective fat-saturated image is considered histologic proof of macroscopic fat, such as in an ovarian dermoid.

Volume-interpolated gradient echo (General Electric = LAVA, Philips = THRIVE, Siemens = VIBE) T1-weighted, fat-saturated images

These are our pre- and post-contrast images.

In the United States, all currently approved MRI contrast agents are chelates of gadolinium. These agents function by increasing the T1-relaxivity (shortening the T1 time) of blood and tissues.

Therefore, contrast-enhanced images are always T1-weighted.

Gadolinium is bright on T1-weighted images.

You can run a T2-weighted sequence post-contrast – it will look remarkably similar to a T2-weighted sequence run pre-contrast. The gadolinium (with few exceptions) has no noticeable effect on a T2-weighted image.

As we've discussed, fat is inherently bright on T1-weighted images. Pre- and post-contrast images are performed with chemically selective fat saturation so that the enhancement stands out prominently against a dark background. On modern scanners, this is accomplished by pressing a button on the scanner which squashes the spectral peak of fat so that we are essentially imaging only water. Macroscopic fat, which is bright on the in-/out-of-phase images, becomes black.

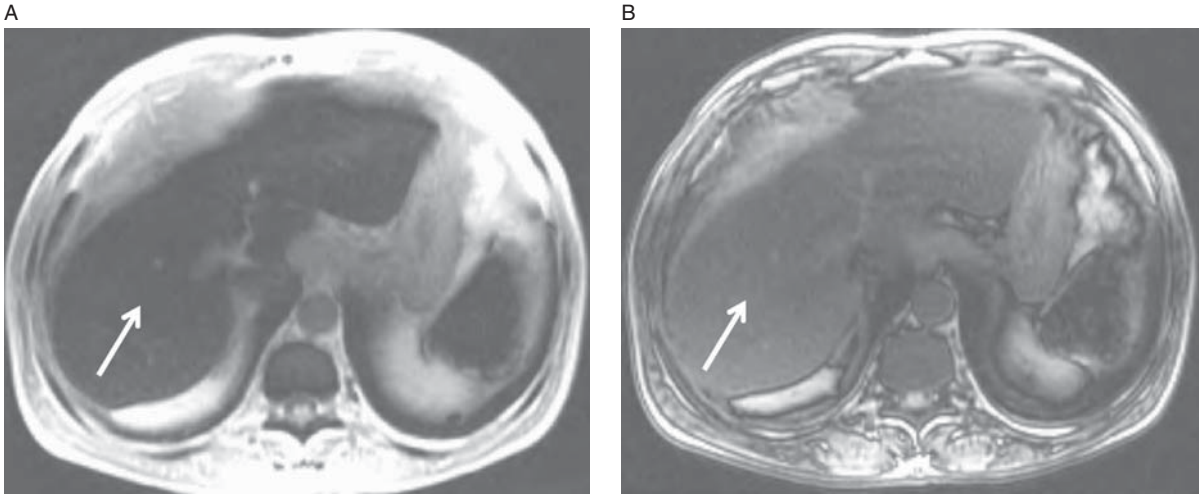

Figure 1.3. Iron overload, chemical shift imaging. Arrows demonstrate dark signal in the liver, spleen, and bone marrow on the axial out-of-phase image (B) which drops significantly on the in-phase image (A). This is due to the longer echo time of the in-phase image. The longer you wait to listen, the more signal you lose due to the inhomogeneity introduced by the iron in the magnetic field.

Whenever giving contrast, pre-contrast images should *always* be obtained. Imaging parameters must remain identical for pre- and post-contrast images so that subtraction images can be generated. Pre-contrast fat saturated images are performed to:

(1) Assure adequate coverage of intended imaging area.
(2) Assess (and correct) artifacts.
(3) Identify lesions that are T1-bright before contrast so that we *don't mistake them for enhancing lesions!* (This is a significant advantage of MRI over CT. You wouldn't routinely obtain pre-contrast imaging with CT because of the additional radiation dose. With only one phase, it can be difficult to tell enhancement from calcification, etc.)
(4) Identify macroscopic fat by comparing to in-phase images which are also T1-weighted but are performed without chemically selective fat saturation.

Subtraction images

Based on all the variables we have discussed, the scanner assigns a numeric value to each voxel in an image which is then displayed as a dot of variable brightness.[2] Subtraction images are generated by subtracting the value of each voxel of the pre-contrast images from the corresponding voxel on the post-contrast images and then displaying this "new" value.

For example, let's suppose a renal lesion is bright on T1-weighted images before contrast with a signal intensity of 200. It will also be bright on post-contrast images, but did it enhance? The subtraction image takes the post-contrast signal intensity and subtracts the initial value of 200. So, if the signal after contrast is 200, the subtracted pixel value will be 200 (post) – 200 (pre) = 0. The lesion will appear black. If, post-contrast, the signal intensity is greater than 200, it will appear bright.

Subtractions allow confident assessment of the presence or absence of true enhancement of lesions which are T1-bright before contrast administration

[2] This is a critical distinction between the way CT and MRI work. Hounsfield units are a measure of density. They are a physical property of matter based on the ability of a substance to attenuate x-rays. The density of water is defined as zero Hounsfield units. We can then "window" the image, alter the parameters by which the image is displayed, to accentuate voxels of different values. For example, lung windows accentuate voxels of very low number. However, changing the windows does not change the voxel value in Hounsfield units. Signal intensity values on MRI are not uniform, standard, or necessarily consistent. Fat will always be bright on a T1-weighted image, but its numeric value can change. This is because the scanner, in an effort to do us a favor, automatically scales each image for us before displaying it to create the "optimal" image. The scale typically changes only when a new pulse sequence is run but it can also change suddenly, without warning or explanation. Better not to draw ROI's unless you are prepared to also draw them on an internal standard and perform a more complicated calculation.

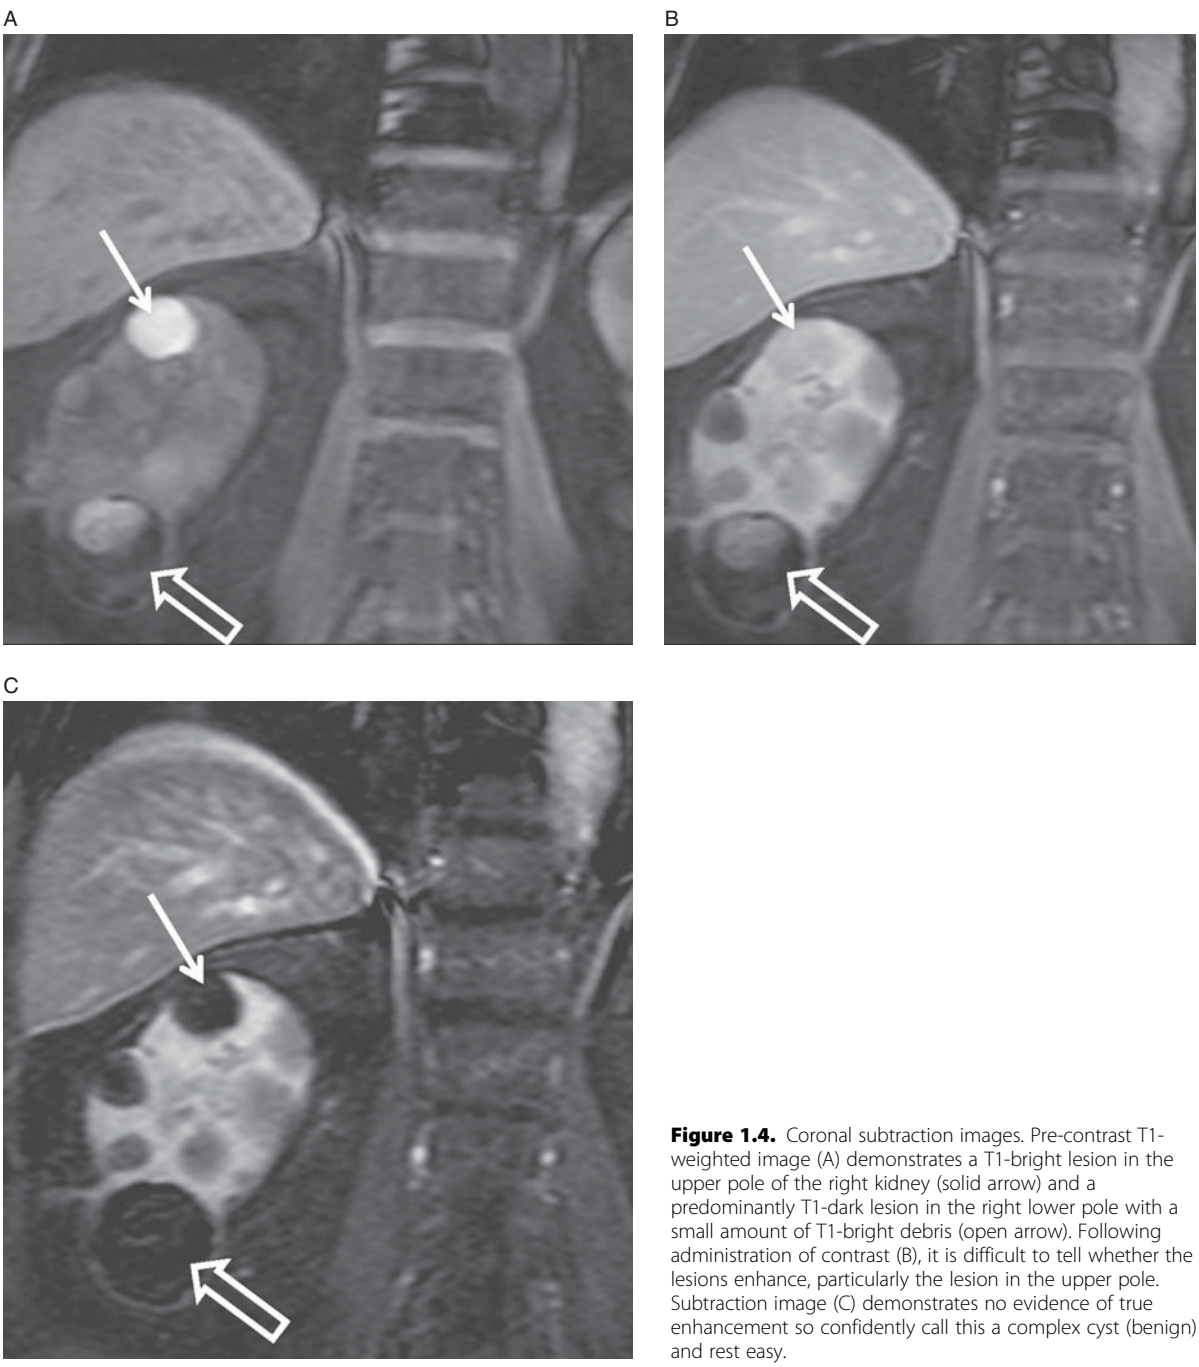

Figure 1.4. Coronal subtraction images. Pre-contrast T1-weighted image (A) demonstrates a T1-bright lesion in the upper pole of the right kidney (solid arrow) and a predominantly T1-dark lesion in the right lower pole with a small amount of T1-bright debris (open arrow). Following administration of contrast (B), it is difficult to tell whether the lesions enhance, particularly the lesion in the upper pole. Subtraction image (C) demonstrates no evidence of true enhancement so confidently call this a complex cyst (benign) and rest easy.

(Figure 1.4). They may also accentuate subtle hypervascularity which can be difficult to appreciate on conventional contrast-enhanced images. Subtraction images can only be created if all imaging parameters are *identical* on the pre- and post-contrast sequences.

T2-weighted sequences

Just as with T1-weighted images, the original T2-weighted images were spin echo. Forget about them. They take too long. We don't use them in body imaging.

Fast-spin echo (General Electric = FSE, Philips = TSE, Siemens = TSE) T2-weighted images; can be performed with or without fat saturation

This was the first, practical T2-weighted sequence for body imaging. It has the potential to generate excellent resolution with unparalleled T2-weighting. At many institutions, including possibly your own, they are a part of every body MRI exam. Unfortunately, for my taste, they still take too long, which makes them prone to motion artifacts when the patient breathes. We typically only use them in the pelvis, which, at least while in the MRI scanner, is relatively motionless.

Single-shot fast-spin echo (General Electric = SSFSE, Philips = SSTSE, Siemens = HASTE) T2-weighted images; can be performed with or without fat saturation

Now we're talking fast.

SSFSE is fast because only 1 RF pulse is used. We then "listen" for numerous echoes to create the image.

SSFSE images are very fast, T2-weighted images which may be obtained with or without fat saturation. Because the acquisition time is short, they are almost always free of motion. Initially they were (and are) the backbone of MRI cholangiopancreatography (MRCP) but they have now replaced TSE/FSE at many institutions (including our own). There's no free lunch here.

The pro: they are fast and therefore motion free.

The con: they are not as heavily T2-weighted as TSE/FSE. (In later chapters we will discuss why we are willing to sacrifice the T2-weighting for the sake of speed.)

As discussed above, T2-weighted images are "fluid-sensitive" – i.e., fluid is bright. But, beware, not all that is bright represents fluid. Classic mistakes include mucinous and neuroendocrine tumors which can be bright on T2. Close inspection will show that these lesions are not quite *as* bright as fluid – correlation with other sequences is critical.

Diffusion-weighted images

Diffusion-weighted images and ADC (General Electric = to be announced, Philips = DWIBS, Siemens = Reveal)

We like diffusion-weighted images (DWI), but you should know that this sentiment is not shared by everyone. They are not strictly necessary for good body MRI and therefore not everyone uses them. I think they make our job easier and I like things that make our job easier.

DWI is an exciting, relatively new area of body MRI which attempts to bridge anatomic and functional imaging. DWI measures and displays the ability of water molecules to diffuse through tissue. The ADC (actual diffusion coefficient) map is a plot of the actual diffusion coefficient.

Normal tissues consist predominantly of interstitial or extracellular space through which water can easily diffuse. Tumors are hypercellular when compared with normal tissues and this increased number and concentration of cells decreases the amount of free interstitial space for water to diffuse through easily, replacing it with intracellular space which creates boundaries and obstacles to diffusion. The result is restricted diffusion which is displayed as bright on DWI and dark on the ADC map (Figure 1.5).

Fundamentally, conventionally acquired DWI are T2-weighted – it is therefore critical to correlate DWI findings with the ADC map. Lesions that are bright on DWI but also bright on ADC do *not* represent true restricted diffusion, but "T2 shine-through."

Currently, the best use for DWI is lesion detection. DWI is now widely considered to be the most sensitive sequence for the detection of liver lesions (this may simply be due to their extreme T2-weighting and one reason we are willing to sacrifice the heavily T2-weighted FSE in favor of the single-shot).

It is tempting to use DWI to characterize masses; however, do so with caution. Despite the early excitement of many investigators, benign and malignant lesions can demonstrate significant overlap in ADC. DWI should therefore be viewed as one tool in our armamentarium – excellent for lesion detection, and one of many tools which should be evaluated together for lesion characterization.

The sensitivity of DWI to water motion can be varied depending on the gradient amplitude, duration of applied gradient, and time interval between paired gradients. The "b-value" can be changed on modern MRI scanners to increase sensitivity to diffusion, which is proportional to the above three factors, with emphasis on the gradient amplitude.

Technically speaking, the ADC is a curve, and the more b-values one uses, the more accurately that curve will reflect the true ADC. Some institutions will therefore use ever-increasing numbers of b-values to more accurately define the true ADC. The drawback to this is that more b-values typically mean increased acquisition time.

But, the actual ADC value doesn't much matter. As we've discussed, you can't use it as a strict predictor of histology. So if the specific number doesn't exactly matter, why measure it exactly?

We currently use b-values of 50 s/mm^2 and 500 s/mm^2. Obviously, with only two points we get

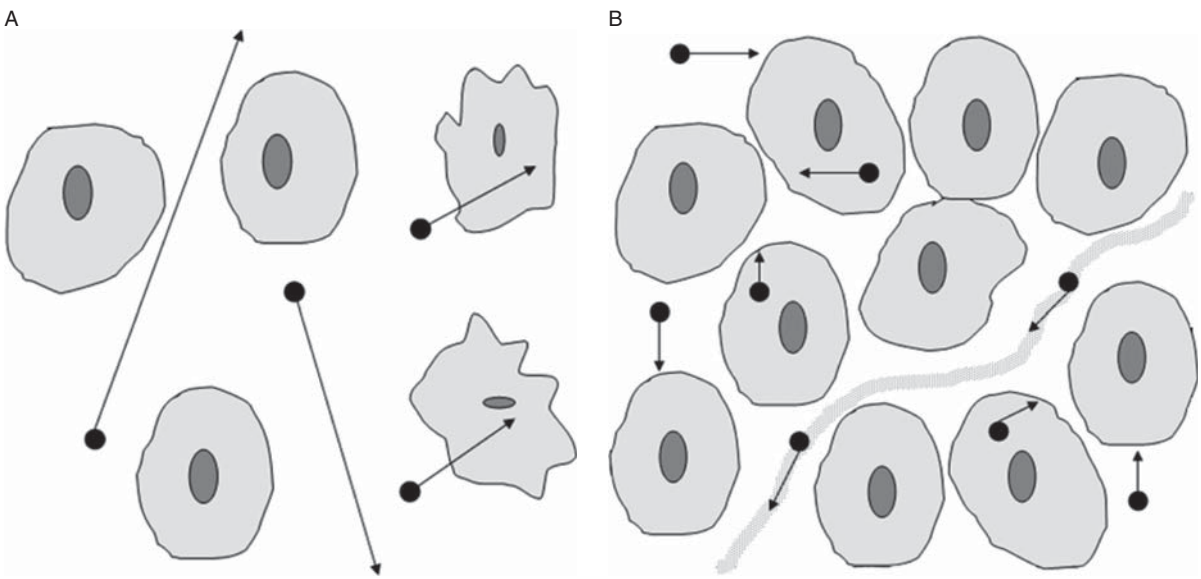

Figure 1.5. Diffusion-weighted imaging. Diagram of normal cells (A) demonstrating extensive extracellular space through which water can easily diffuse. Diagram of abnormal tissue (B) showing decreased extracellular space due to cellular crowding. Water cannot easily diffuse in the space around these cells. This is restricted diffusion. (Reprinted with permission from *AJR*, 6/2007 **188**: 1622–1635.)

a line, not a true curve. But, what's bright is bright and since all we really care about is detection (not characterization) it is accurate enough.

Water molecules with a large degree of motion or great diffusion distance will be hyperintense at low b-values. Alternatively, larger b-values (e.g. 500 s/mm^2), are required to perceive slow-moving water molecules or small diffusion distances.

Steady-State Free-Precession (General Electric = FIESTA, Philips = Balanced-FFE, Siemens = TruFISP)

Consider this a T2 weighted-sequence (fluid is bright).

This sequence features a very short acquisition time and can be used to generate cine images for cardiac or bowel imaging. It also depicts flowing blood within vessels as high signal intensity ("bright-blood" MRI). In the current era of nephrogenic systemic fibrosis (NSF) hysteria, there is renewed interest in these sequences as non-contrast MRAs.

And . . . well, that's it. Don't be intimidated by MRI pulse sequences and protocols. Fundamentally, there just isn't that much to it. You've got T1, T2, DWI, and +/− fat saturation. Each of these provides one more little piece of information which will lead us to our diagnosis.

RF coils used in body MRI

The coil is the antenna used to receive MRI signals from inside the body. The MRI signal which we are trying to capture is quite weak so the closer we can place our antenna to the body part we wish to image, the more signal we obtain (i.e., the better our signal-to-noise ratio).

Body coil

The MRI scanner arrives equipped with a send/receive coil within it – the "body" coil. While always used to send RF pulses into the body, it can also be used as an antenna to receive signal back. Obviously, this coil is as far from the body part we wish to image as a coil can be (while still being in the scanner) and thus should only be used when a surface coil will not fit between the patient and the bore of the magnet (i.e. obese patients).

Phased-array coils

Whenever possible, a phased-array surface coil is used. The term "phased-array" means that the surface coil is actually made of multiple (currently up to 32) individual coils which can each receive signal independently (these are receive-only coils). Having lots of individual coils within the larger coil allows us to perform parallel imaging. The physics of this is beyond the scope of this book, but parallel imaging decreases the acquisition time of pulse sequences by factors of 2 and beyond.

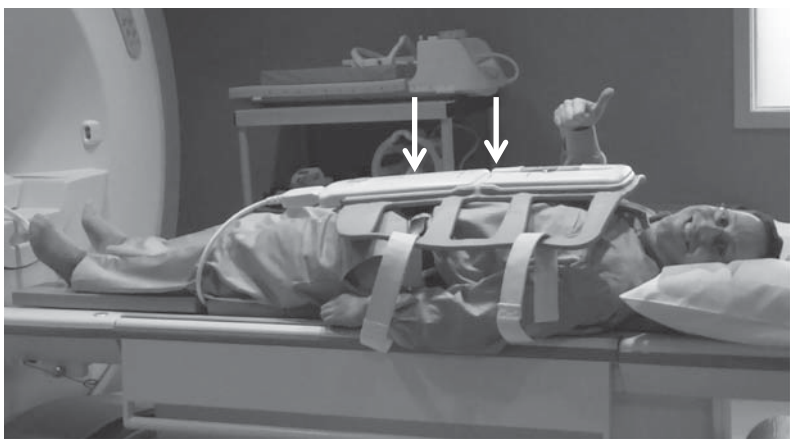

Figure 1.6. Surface coil placement. Arrows indicate the phased-array surface coil on the ventral surface of the patient (author). Depending on the manufacturer, either the patient lies directly on the posterior surface coil, or it is embedded in the table.

The surface coil is placed on the ventral surface of the patient. We call them "coils" but they look like foam pads. The coils are hidden on the inside (Figure 1.6). Modern scanners combine this information with signal information from the phased-array spine coil which either the patient lies directly on or is embedded within the table. Careful observation of images acquired particularly from larger patients will show brighter signal at the ventral and dorsal surfaces of the patient with a decrease in signal toward the center of the patient as the coils each become more distant.

Endocavitary coils

Everyone's favorite topic. Knowledge of these coils is more useful to disgust your friends than for the actual practice of radiology. But they do currently have a few clinical indications.

Remember that the signal we are listening for is very weak and diminishes rapidly with distance. The closer we can place our coil to the area of interest, the better the image. Hence, endocavitary coils placed within the body.

There are a variety of these coils, but the most commonly used is the endorectal coil which is used to image the prostate. Suffice it so, you can't get much closer to the prostate than this, but it often doesn't get there without a fight. Endocervical and even endourethral coils are available but are not widely used outside of large academic centers.

Contrast agents used in body MRI

Intravenous contrast agents

All intravenous MRI contrast agents available in the United States are chelates of gadolinium. They function by increasing the T1 relaxivity of blood and tissue, rendering them bright on T1-weighted images. They are typically power-injected and imaging may be performed dynamically, i.e. in multiple phases. Practically, there are two categories of gadolinium agents.

Extracellular or non-specific

These are the most commonly used agents. They are used for all types of MRI. Prior to the advent of Nephrogenic Systemic Fibrosis (NSF), they were essentially interchangeable.

Gadolinium-based contrast agents are extremely safe because the gadolinium ion is bound to something else. Free gadolinium is bad. The strength with which the gadolinium moiety is bound to its chelate varies between contrast agents, and the working theory of NSF is that it occurs when the gadolinium ion breaks free from its chelate and wreaks havoc. Theoretically, the tighter the binding, the safer the agent.

Gadopentetate dimeglumine (Magnevist)
Gadoteridol (ProHance)
Gadoversetamide (OptiMARK)

Liver specific

There are currently two available agents in the United States that are considered liver specific.

Gd-BOPTA (MultiHance) and gadoxetate disodium (Eovist) are taken up by functioning hepatocytes in the liver and have variable biliary excretion, as opposed to the renal only excretion of the non-specific agents.

Perhaps the main value of these agents is in distinguishing focal nodular hyperplasia (FNH) from hepatic adenoma. The imaging features of these two lesions overlap when using non-specific gadolinium