Basic MRI Physics and Artifacts

1

Carl Siversson

Magnetic resonance imaging (MRI) is a well-established imaging technique, which is available at most larger hospitals today. Since MRI is sensitive to the chemical surroundings of the atomic nuclei within the body, the images have excellent contrast between different soft-tissues. Due to the combination of this high contrast and the fact that MRI does not emit any hazardous ionizing radiation, MRI is often used for investigation of a large range of pathologies in almost all parts of the body.

A typical MRI scanner consists of a large superconducting electromagnet in which the patient is positioned. Most clinical scanners have a static magnetic field strength of either 1.5 or 3 tesla (T), although efforts are made to increase the field strength further as that would result in even higher image quality.

When the patient is placed in this strong static magnetic field, the tissue in the body becomes temporarily magnetized. By exposing the tissue to a series of radio frequency pulses and magnetic field gradients, the nuclei in the tissue start to emit radio frequency signals of their own. These signals are detected by the MRI scanner and formed into an image. In clinical MRI, signals from hydrogen nuclei are the most widely used, due to the abundance of hydrogen-rich components, such as water and fat, in the human body.

In this chapter an overview of the MRI phenomena is given. Covering the complete theoretical background of MRI within one chapter would not be feasible and is therefore not the purpose of this text. Instead, this text should be considered an introduction for the uninitiated reader who is interested in gaining an intuitive understanding of the relevant concepts. For a deeper theoretical journey through the physics of MRI, the reader is instead referred to any of the excellent textbooks dedicated to this topic [1–3].

Nuclear Spin

Protons and neutrons both have an intrinsic property called spin. In an atomic nucleus the number of protons and neutrons determines whether the nucleus will have an overall spin or not. Of primary interest in clinical MRI is the hydrogen nucleus, which possesses an overall spin since it consists of only a single proton.

In a strict sense, the spin of an individual nucleus should be described from a quantum physics perspective. That said, there are a number of spin models based on classical physics that are not fully comprehensive, but which are still valuable and provide an intuitive understanding of the phenomenon. In one such model, the spin property can be described as each nucleus rotating around its own axis (Fig. 1.1). Since the nucleus is positively

C. Siversson (✉)
Computational Radiology Laboratory,
Boston Children's Hospital, Harvard Medical School,
Boston, MA, USA

Department of Medical Radiation Physics
Lund University, Malmö, Sweden
e-mail: carl.siversson@med.lu.se

Y.-J. Kim (eds.), *Hip Magnetic Resonance Imaging*,
DOI 10.1007/978-1-4614-1668-5_1,

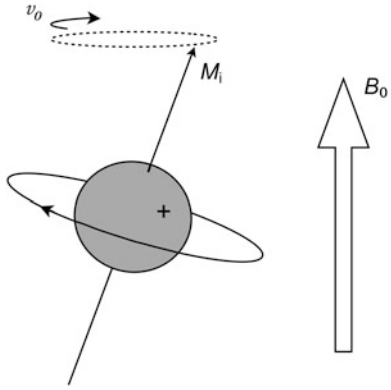

Fig. 1.1 The hydrogen nucleus spins around its own axis, generating an intrinsic magnetic vector M_i, which in turn makes the nucleus precess with the Larmor rotation frequency v_0 about the direction of the external field B_0

charged itself, the rotating charge is equivalent to a rotating current, generating a small magnetic field. Hence, each nucleus has its own intrinsic magnetic vector pointing along its axis of rotation.

Without external influences, the intrinsic magnetic vectors of nuclei in a normal tissue sample are pointing in completely random directions. Hence, from a macroscopic perspective these vectors are averaging themselves out to zero net magnetization.

When placing nuclei in an external magnetic field, B_0, the intrinsic magnetic vector of each nucleus will align along the direction of this field.[1] However, this alignment will not make the vectors become strictly parallel to the B_0 field. Instead, the nuclei will start to precess around the axis of the B_0 field, making the intrinsic magnetic vectors circle at an angle around this direction. This circling will occur with a rotational frequency called the Larmor frequency, v_0 (Fig. 1.1), which varies with the strength of the external B_0 field by

$$v_0 = \gamma \cdot B_0 \tag{1.1}$$

where γ is a nuclei-specific constant called the gyromagnetic ratio ($\gamma = 42.6$ MHz/T for the hydrogen nucleus).

[1] Strictly speaking, there are almost as many hydrogen nuclei aligned against the B_0 field, with only a small surplus of nuclei aligned as described above.

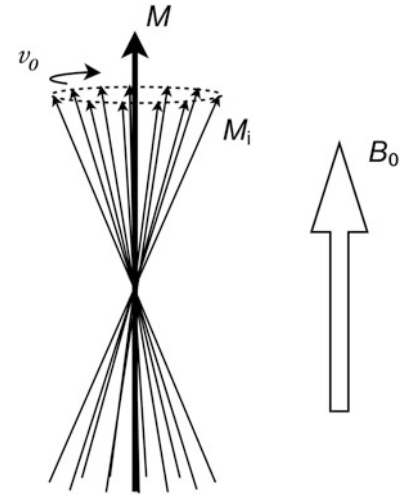

Fig. 1.2 Each individual intrinsic M_i vector has a random angular position. Thus, the sum of all M_i vectors in an isochromat is a new vector, M, pointing in the same direction as the externally applied B_0 field

In any tissue sample, even an extremely small volume contains many millions of hydrogen nuclei. For this reason, it is instructive to consider a large number of nuclei as a single unit, called an isochromat (Fig. 1.2). The intrinsic magnetic vectors of each nucleus in an isochromat are all circling around the B_0 field, but at random angular positions. Thus, adding all intrinsic vectors of the isochromat together results in a bulk magnetization vector, M, pointing in the same direction as the B_0 field.

Coils and RF Pulses

In order to interact with the precessing nuclei, radio frequency (RF) magnetic pulses are used. Such pulses are generated and received by coils, which must be placed in the vicinity of the tissue to be imaged. Coils are available in a variety of types and are usually shaped to come as close as possible to the tissue of interest. For example, for brain imaging a cylindrically shaped coil is often used, efficiently surrounding the entire head, whereas for hip imaging a flat flexible coil that can be directly strapped onto the body surface is more feasible.

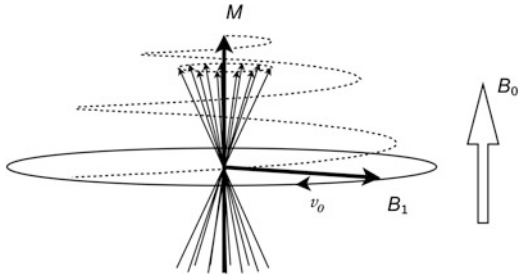

Fig. 1.3 The RF pulse can be thought of as a magnetic B_1 field rotating around the B_0 direction. The nuclei will then precess around both the B_0 and the B_1 fields, which results in the M vector spiraling away from the B_0 direction

In order for the coil to generate an RF pulse, a current oscillating with the Larmor frequency is passed through it. The oscillating magnetic field generated by this current can be divided into separate components. Of particular interest is the component that is perpendicular to the B_0 field, which can be thought of as an additional, weaker magnetic field (called the B_1 field) rotating around the B_0 direction (Fig. 1.3).

From the nuclei's perspective there are now two concurrent magnetic fields around which they will precess. Since one is a strong stationary field (B_0) and the other is a weak rotating field (B_1), the result is that the bulk magnetization vector M will spiral itself away from the B_0 direction at a relatively slow speed. This spiraling will proceed for the duration of the RF pulse. An RF pulse is characterized by the resulting angle between the magnetization vector M and the B_0 direction. This flip angle is primarily determined by the duration (typically 0.5–5 ms) and the amplitude of the pulse.

As soon as the M vector points away from the B_0 direction it will no longer be time and space invariant. Instead, it will rotate around the B_0 direction. Thus, the nuclei themselves generate an oscillating magnetic field, which will induce a signal in the conductors of the receiver coil. At this point the spinning nuclei are said to be excited and the associated RF pulse is often referred to as an excitation pulse. The signal that follows immediately after excitation is denoted free induction decay (FID). However, in clinical MRI other types of signals are more commonly used, which requires a few additional steps to generate. Such signals are referred to as echoes and are described later in this chapter.

For a 90° excitation pulse flip angle the M vector is turned perpendicular to the B_0 field and the generated signal will have its maximum amplitude.

T_1, T_2, and Proton Density

There are three primary parameters required in order to describe the spin-behavior of a set of hydrogen nuclei. These are the T_1, T_2, and proton density (PD) parameters. In most MR images the contrasts origin from the variation of these parameters between tissue types.

The proton density is the most intuitive parameter, as the number of hydrogen nuclei per unit volume is directly proportional to the amplitude of the signal generated by these nuclei. However, in clinical MRI this parameter is often not of primary interest, as it does not provide much information about the chemical composition of the tissue.

The T_1 parameter is a measure of how fast the longitudinal (parallel to B_0) magnetization recovers (Fig. 1.4). When the magnetization vector M is perturbed by an RF pulse, it will take a certain amount of time until it has recovered back to its equilibrium state (i.e., when the isochromats are fully relaxed). This recovery occurs exponentially with time and is characterized by the T_1 parameter, which can be thought of as the time it takes for the nuclei to recover 63 % of the longitudinal magnetization that was lost when the RF pulse was applied.

The T_2 parameter is a measure of how fast the transverse (perpendicular to B_0) magnetization decays (Fig. 1.4). It can be pictured by considering an isochromat immediately after an RF pulse, at which time the M vector has a large rotating transverse component, which can induce a current in an external pickup coil. After a short while, the intrinsic vectors that make up the isochromat will dephase and point in slightly different directions, resulting in a reduction of the rotating component of the M vector. After yet another while the

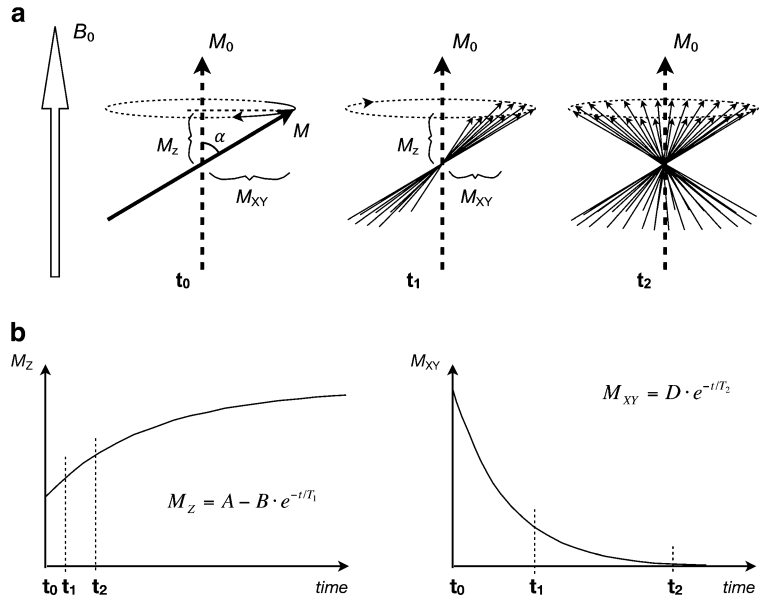

Fig. 1.4 (**a**) Immediately after an RF pulse (t_0) the magnetization vector M is perturbed from its initial M_0 direction by a flip angle α. M_z is the remaining longitudinal magnetization, which will eventually recover to back to M_0. At this point, the transverse magnetization component M_{xy} is time variant, allowing a signal to be induced in an external pickup coil. A short while later (t_1) the individual spinning nuclei that make up the M vector have dephased somewhat, resulting in a smaller time variant component M_{xy} and a decreased amplitude of the induced signal. After yet another while (t_2) the spinning nuclei have dephased totally, resulting in no time variant component at all. At this point, no signal will be induced in the external coil. The dephasing time is characterized by the T_2 parameter. (**b**) The longitudinal magnetization (*left*) recovers exponentially at a rate determined by T_1. The transverse magnetization (*right*) decays exponentially at a rate determined by T_2. The transverse magnetization always decays much faster than the longitudinal magnetization recovers. Time points t_0, t_1, and t_2 refer to illustrated states in (**a**)

intrinsic vectors are completely dephased, leaving no rotating component at all. This decline occurs exponentially with time and is characterized by the T_2 parameter, which can be thought of as the time it takes for the transverse magnetization to decline to 37 % of what it was immediately after the RF pulse.

Any individual nucleus is surrounded by millions of other nuclei. Since all nuclei possess some type of motion (Brownian motion, tumbling motion, etc.), the intrinsic magnetic vectors of all these nearby nuclei will summarize into a slightly fluctuating magnetic field that is superimposed onto the static B_0 field.

The component of this fluctuating magnetic field that is perpendicular to the B_0 field will reveal itself to the nuclei as an additional vector around which they will try to precess. Even though this vector is not strong enough to cause actual precession (as the B_0 and B_1 fields do), it is strong enough to affect the longitudinal relaxation behavior of the nuclei. Thus, this describes one of the primary effects determining the T_1 value.

The other component of the fluctuating magnetic field, parallel to the B_0 field, will add scalar to the B_0 field strength, causing a slight fluctuation of Larmor frequency between individual nuclei. This effect is one of the primary reasons for T_2 relaxation, since a variation in Larmor frequency for nuclei within an isochromat causes its compounding intrinsic magnetic vectors to dephase.

As can be understood, both the T_1 and T_2 relaxation times are determined by the combined characteristics of the fluctuating magnetic fields that are present within a sample. As such, there are no simple rules as to what the T_1 and T_2 values will be for a specific tissue. In general, the T_1 and T_2 values are dependent on what types

of bindings exist between hydrogen nuclei and their surroundings, as well as the mobility of the nuclei. For example, in very low-viscous fluids such as cerebrospinal fluid (CSF) the hydrogen mobility results in very long T_1 values (2,000–3,000 ms), whereas it is very short in fatty tissue (down to 100 ms) where the motion of the nuclei is much more restricted.

In a perfect measurement setup the signal picked up by the receiver coil would strictly follow the decay of the transverse T_2 magnetization, as is described above. However, in practice there are always local inhomogeneities in the applied B_0 field, causing additional variation in Larmor frequency between nuclei and thereby making the nuclei dephase faster than predicted by the T_2 value. These local inhomogeneities are a result of macroscopic effects, such as the technical difficulties of designing an absolutely homogenous magnet in combination with magnetic susceptibility variation between different tissue types. Compared to the mechanism governing the T_2 decay, these local inhomogeneities are static in the sense that they will not vary with time at a specific location. Thus, the dephasing caused by such local inhomogeneities can be reversed using refocusing RF pulses, which is described later in this chapter. The combination of T_2 decay and the decay caused by local inhomogeneities is usually referred to as T_2* decay.

Magnetic Field Gradients

In order to form images it is necessary to be able to differentiate between signals coming from nuclei in different parts of the subject. This is achieved using the gradient coils.

The gradient coils generate magnetic fields that are designed such that they vary in strength along their axis of operation (Fig. 1.5). When a gradient is switched on, its field will superimpose onto the static B_0 field. The total magnetic field strength will then be somewhat higher in one end of the investigated volume and somewhat lower in the other end. Hence, the Larmor frequency of the nuclei in the volume will also vary throughout

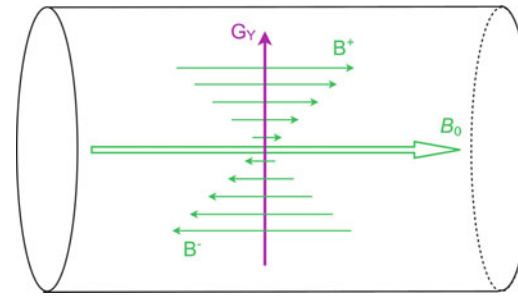

Fig. 1.5 In this figure the *y*-gradient (G_Y) is illustrated inside the MRI scanner, although the principle is the same also for the *x* and *z*-gradients. When the gradient is operated, its field (B^+ trough B^-) will add onto the static B_0 field, so that the total magnetic field is somewhat higher in the upper part of the bore and somewhat lower in the bottom part. Thus, the Larmor frequency of nuclei will vary depending on where they are located. The strength of the gradient can be arbitrarily adjusted by the MRI scanner

the different positions. This is the effect used to form images.

There are three sets of gradient coils in a clinical MRI scanner, generating magnetic field gradients in three orthogonal directions (*x*, *y* and *z* direction). The strength of each gradient can be arbitrarily adjusted by the MRI scanner. By combining these three gradient coils (i.e., running them simultaneously and at different strengths) apparent magnetic field gradients can be generated in any direction. There are typically three different ways to use the gradients for encoding spatial information:

1. Applying a narrow-band RF pulse while a magnetic field gradient is switched on will exclude all nuclei whose Larmor frequencies are outside of the RF pulse bandwidth, meaning that a specific section through the tissue can be selected for excitation. This is referred to as a spatially selective gradient. A slice-selection gradient is a commonly used spatially selective gradient.
2. Applying a gradient for a short while immediately after an RF pulse will create a difference in phase of signals emitted from nuclei at different positions. The phase can be thought of as a measure of how well the emitted signals are synchronized. This is referred to as a phase encoding gradient and the direction in which it

is applied is denoted the phase encoding direction.

3. Applying a gradient during acquisition of the signals will make nuclei at different positions emit signals with different frequencies. Such gradient is referred to as a readout gradient and the direction in which it is applied is denoted the frequency encoding direction.

Thus, by sampling the emitted signals and analyzing their phase and frequency spectrum, information about nuclei from different parts of the subject can be obtained. However, all information required to generate an image cannot be contained within just one sampled signal. Thus, the combination of gradients and RF pulses must be repeated many times, changing the phase encoding gradient slightly for each repetition, in order to obtain the information necessary to generate an image. Such combination of repeated gradients and RF pulses is denoted a pulse sequence and can be configured in a large number of ways in order to enhance the obtained image for different purposes. The execution time of a clinical pulse sequence typically ranges between 1 and 10 min.

When a pulse sequence is executed, the sampled signals together constitute a spatial frequency representation of the object, which can later be Fourier transformed into a real visible image. Such frequency representation is called a *k*-space image.

Gradient Echoes and Spin Echoes

As described earlier, nuclei that are perturbed by an external RF pulse will respond by generating an RF signal of their own (i.e., a FID). In order to use this signal for image formation it is necessary to first subject the nuclei to a series of magnetic field gradients, as signals originating from different parts of the tissue would otherwise not be distinguishable. Since these gradients have a duration in the order of a few milliseconds, the acquisition of signal is usually delayed from the RF pulse, leaving enough time for the gradients to be applied. This type of acquired signals is referred to as echoes, of which the two most commonly used types are the gradient echo and the spin echo.

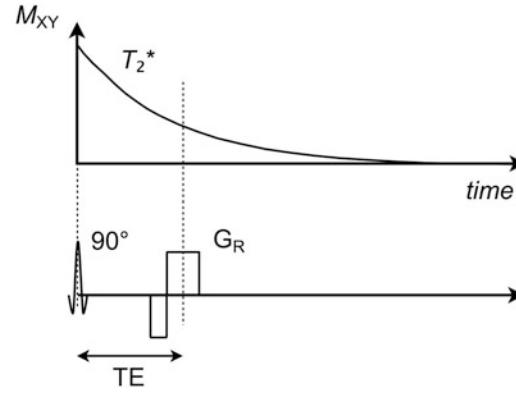

Fig. 1.6 Transverse magnetization (M_{XY}) during a gradient echo. An initial RF pulse generates a FID signal which decays according to the T_2* of the tissue. Immediately after the RF pulse a set of image-encoding gradients is applied, having a final readout gradient (G_R) during which the echo is retrieved. The echo time (TE) can be set arbitrarily and determines the amount of T_2*-weighting within the image

The gradient echo is the simplest type of echo as it basically only involves an excitation RF pulse followed by a set of gradients applied directly on the FID signal. These gradients can be configured in a variety of ways, although they are all characterized by having a final readout gradient during which the actual echo signal is retrieved (Fig. 1.6). The time that arises between the RF pulse and the echo is denoted the echo time (TE), which can be set arbitrarily with respect to the minimum time required for the gradients to apply. However, it should be noted that since the FID signal decays at the relatively fast T_2* speed, the echo time must be kept short (typically 2–20 ms) in order for the echo amplitude not to be unreasonably low. As a consequence, the echo time can be used to adjust how much impact the T_2* of the tissue will have on the echo amplitude and thereby also on the contrast within the resulting image.

A spin echo is characterized by being generated by two separate RF pulses after which an echo arises (Fig. 1.7). The first pulse perturbs the magnetization, resulting in a FID signal which will quickly decay due to T_2* dephasing effects. However, since T_2* is partially an effect of local inhomogeneities that are spatially and temporally static, its impact can be partly reversed by applying

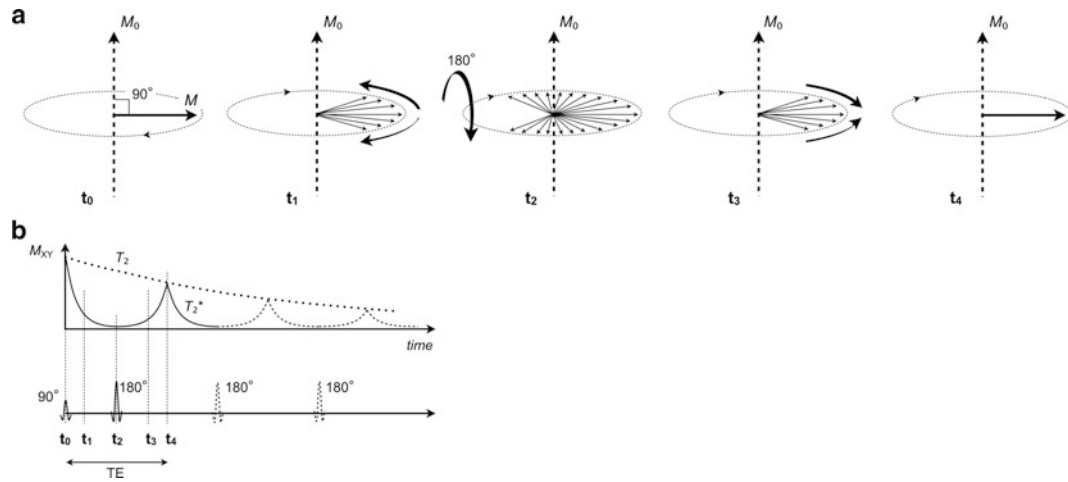

Fig. 1.7 Magnetization during a spin echo. (**a**) Shows the behavior of a set of excited nuclei while (**b**) describes the concurrent RF pulses and the echo signal. At time point t_0 a 90° RF pulse is applied, flipping the M vector perpendicular to its original direction. A FID signal is then emitted, which decays quickly due to T_2* effects (t_1). At time point t_2 a 180° refocusing RF pulse is applied, which reverses the direction by which the individual vectors move relative to each other, meaning that a rephase of the vectors will occur (t_3). At time point t_4 all individual vectors are in phase again and a spin echo occurs. After the spin echo the vectors will once again start to dephase. By adding more 180° RF pulses additional spin echoes can be induced (*dashed line* in (**b**)). The amplitude of successive spin echoes will decline according to the T_2 of the tissue (*dotted line* in (**b**)), since pure T_2 dephasing is not temporally and spatially static and will thus not be rephased

a refocusing RF pulse (typically a 180° pulse). Such a pulse will effectively flip all magnetization vectors upside down, meaning that any dephasing that has occurred due to local inhomogeneities since the first RF pulse will reverse and begin to rephase. At precisely twice the time between the two RF pulses, the magnetization will again be in phase and a spin echo will occur. Since all local inhomogeneity effects are canceled during this echo time, only the decay caused by the actual T_2 of the tissue remains. For this reason the echo time determines how much impact T_2 will have on the echo amplitude and the resulting image contrast. This is often utilized to adjust which tissues and pathologies are to be emphasized in the image.

Immediately after the spin echo, the magnetization will again start to dephase. By adding further RF pulses the magnetization can be refocused again, resulting in additional spin echoes. The amplitude of successive spin echoes will decay according to the T_2 of the tissue (Fig. 1.7b).

Compared to a gradient echo sequence, the echo time for a spin echo sequence can be much longer (since $T_2 > T_2$*) without unacceptable signal loss. As a consequence there are numerous possibilities to incorporate image-encoding gradients, which will be described more in detail later in this chapter.

As to which type of echo is preferred over the other depends on the situation and requested type of image. Spin echo sequences are much more robust and generate stronger signals than gradient echo sequences, since they do not suffer from T_2* losses and are less sensitive to disturbances in the magnetic field. On the other hand, spin echo sequences are often more time-consuming than gradient echo sequences, due to the additional refocusing pulse. For this reason, gradient echoes are usually preferred for three-dimensional high-resolution scans where acquisition time is an issue.

Image Acquisition

There are many different strategies for acquiring an MR image. An MR image can be either a two-dimensional (2D) image, which shows a slice through the subject in an arbitrary direction, or

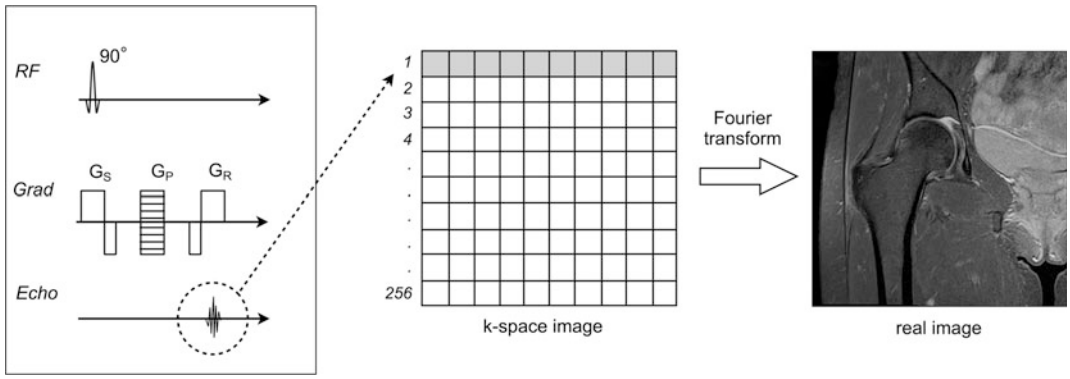

Fig. 1.8 A typical 2D pulse sequence. During the RF pulse a slice selection gradient (G_S) is turned in, which acts to make sure that only one slice of nuclei are excited by the RF pulse. Following the RF pulse a phase encoding gradient (G_P) is applied and during the echo a readout gradient (G_R) is on. The echo is sampled and stored as one line in the *k*-space image. This process is repeated, using different phase encoding, for all lines in *k*-space. When the *k*-space image is fully acquired it is Fourier transformed into a real visible image

it can be a three-dimensional (3D) image where an entire volume of image data is retrieved at once. A 3D sequence is basically constructed the same way as a 2D sequence. The main difference is that instead of just one phase encoding gradient, two such gradients are applied in different directions, in order to account for the spatial encoding of the additional dimension. An alternative to using real 3D imaging is to acquire several 2D images stacked next to each other, creating a so-called multi-slice image. This approach is common when only a few slices are required.

For both 2D and 3D images the number of repeated echoes that are required depends on the desired resolution of the final image. Most pulse sequences collect one full line of *k*-space data for each echo, often using a *k*-space image that is of the same size as the desired final image. For instance, if a 2D image of size 256 × 256 pixels is requested and each sampled echo constitutes one line in the *k*-space image, a total of 256 echoes are required to fill the entire *k*-space (Fig. 1.8). However, if instead a 3D image of size 256 × 256 × 64 is to be acquired a total of 256 × 64 = 16,384 echoes are required.

For this reason, high-resolution 3D images are best suited for use with echoes that can be generated quickly, such as gradient echoes, or the acquisition time might be unfeasibly long. Many techniques have been developed to speed up the acquisition process further, for example by generating multiple echoes following a single excitation pulse or by interleaving the acquisition of 2D slices such that several slices are acquired at once.

Image properties such as pixel size, slice thickness, and field of view can be set almost arbitrarily on an MRI scanner. Technically, these parameters are the result of gradient properties such as strength, duration and increment in combination with the RF pulse bandwidth. However, in any clinical MRI scanner the adjustment of this type of image properties is simplified such that the desired values can be directly typed on the scanner console, thereby letting the user disregard most of the technical details.

A recurring issue in MRI is that of noise in the acquired image. A common measure of noise in an image is the signal-to-noise ratio (SNR), which is high if the amount of noise in the image is low. Typically in MRI, both higher SNR and shorter acquisition time are required, which are parameters that are usually contrary to each other. Thus, the best balance between these parameters always needs to be found.

If the voxel size in an image is decreased fewer echo emitting nuclei will reside within each voxel. The image will then become noisier

and the SNR will be low. This can occur both as an effect of increasing the number of pixels or as an effect of decreasing the field of view. In case the SNR is unacceptably low this can be resolved by acquiring the same image several times, from which a mean image is calculated. This process is commonly known as averaging the image. Other ways of increasing the SNR includes using different types of RF coils or using other types of pulse sequences.

Image Contrast

In any MR image the resulting pixel intensity is a combination of T_1, T_2 (or T_2*), and proton density together with both acquisition parameters and technical parameters, including signal attenuation through the body and signal gain through coils and amplifiers. Thus, the absolute pixel intensities within an image are usually unpredictable and should not be considered when evaluating an image. Instead, it is the relative contrast between different tissues that is of interest in an MR image.

It is generally not possible to generate an MR image in which the image contrast depends on just one of the T_1, T_2, or proton density parameters (unless using a combination of different images, as is described later in this chapter). What is possible, on the other hand, is to adjust which of the parameters are of most impact. This is commonly referred to as weighting. In, for example, a T_1-weighted image, tissue with short T_1 will appear brighter than tissue with long T_1. However, in a T_2-weighted image the relation is opposite, as tissue with long T_2 will appear brighter than tissue with short T_2.

The amount of weighting is determined by the acquisition parameters in combination with the choice of pulse sequence. The two sequence parameters that are of primary interest in this context are the echo time (TE) and the repetition time (TR). As previously described, the echo time is the time between the initial excitation pulse and the generated echo. The repetition time is the time between the excitation pulses.

Table 1.1 The relations between echo time (TE), repetition time (TR), and image weighting

	Short TE	Long TE
Short TR	T_1-weighted	
Long TR	Proton density weighted	T_2 or T_2*-weighted

In a sequence where only one echo is generated per excitation pulse the repetition time is also equal to the time between successive echoes. In short, the relations between these parameters are described in Table 1.1.

The T_1 dependence is mostly determined by the repetition time. At each RF pulse some of the longitudinal magnetization is converted into an echo. The repetition time determines for how long the remaining longitudinal magnetization is allowed to recover before the next RF pulse is applied and the process is repeated. If the repetition time is very long, all longitudinal magnetization will have time to recover and the resulting echo signal will always be high, irrespective of its T_1 value. However, if the repetition time is relatively short, tissue having short T_1 will recover much more during this period than tissue with long T_1 (Fig. 1.9). This is the effect that is utilized to achieve T_1-weighting.

The T_2 dependence is mostly determined by the echo time. The longer the echo time, the larger the impact of T_2 decay (or T_2* decay for a gradient echo sequence) will be on the echo signal. For a short echo time the signal will always be high. However, if the echo time is reasonably long, tissue having short T_2 will experience more signal loss than tissue having long T_2 (Fig. 1.10). This is the effect used to achieve T_2-weighting.

By combining the behavior of the repetition time and the echo time, it is apparent that if both parameters are short there will be more T_1-weighting and less T_2-weighting. Likewise, if both parameters are long there will be more T_2-weighting and less T_1-weighting. Finally, if the repetition time is long and the echo time is short both T_1- and T_2-weighting will be minimized, leaving proton density as the primary contrast contributor.

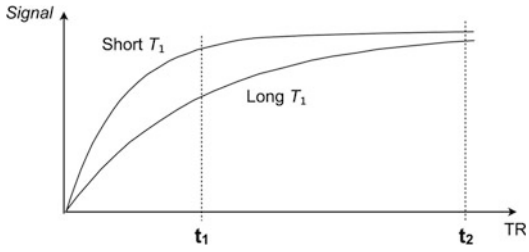

Fig. 1.9 The relation between echo signal and repetition time (TR) for tissues with two different T_1. A long TR (t_2) results in a high signal for both tissues. For a shorter TR (t_1) there is a large signal difference, which depends on the T_1 value. Note that if TR is too short there will be no signal at all

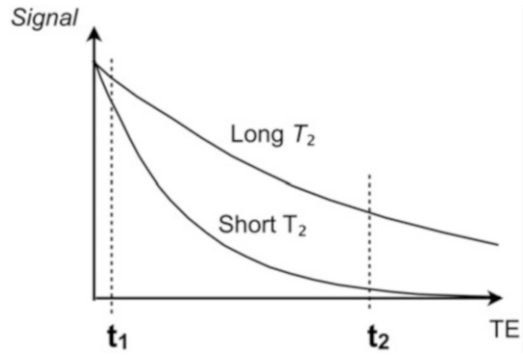

Fig. 1.10 The relation between signal and echo time (TE) for tissues with two different T_2. A short TE (t_1) results in a high signal for both tissues. For a longer TE (t_2) there is a large signal difference, which depends on the T_2 value. Note that if TE is too long all signals will have decayed to zero

Common Imaging Artifacts

An image artifact is any feature which appears in an image that is not physically present in the imaged object. Image artifacts are very common in MRI and can appear for a variety of reasons. Sometimes they are the result of improper operation of the scanner, while other times they are a consequence of natural processes or properties of the human body. It is important to be familiar with the appearance of artifacts since some artifacts can either obscure, or be mistaken for, pathology. As a result of this, image artifacts can lead to both false negative and false positive findings. A few of the most common types of image artifacts are described below.

Phase Wrap Artifacts

Phase wrap artifacts (also known as aliasing artifacts) occur when tissue that should be outside of the field of view suddenly appears within the field of view. The typical example is seen when tissue structures extend outside of the image on one side and then continues into the image again on the opposing side (Fig. 1.11). Depending on how the subject is positioned this may result in structures being imaged on top of each other. Phase wrap artifacts usually only appear in the phase encoding direction of the image.

A simplified explanation of these artifacts is given by picturing the phase encoding gradient when it is operating at a strength such that the phase of the signal is $-180°$ on one side of the field of view and $+180°$ on the other. Signal from nuclei that is positioned slightly outside of the field of view on one side will then have a phase that is slightly outside of this interval, for example $195°$. However, since $195°$ cannot be distinguished from $-165°$ (since $195°–360°=-165°$), such signal will be interpreted as if it was positioned within the field of view but close to the other side.

One solution when this type of artifact appears is to increase the field of view until no tissue extends outside it in the phase encoding direction. However, since this may not always be feasible other solutions include changing the phase encoding direction such that the problem is avoided or using saturation bands to suppress signal outside of the field of view (which is described later in this chapter).

Chemical Shift

Chemical shift misregistration artifacts appear as bright or dark outlines predominantly at the interface between fat and water. This type of artifacts appear since fat might be shifted a few pixels away from water within an MRI image. Due to different chemical surroundings of hydrogen, the Larmor frequencies are slightly different in water

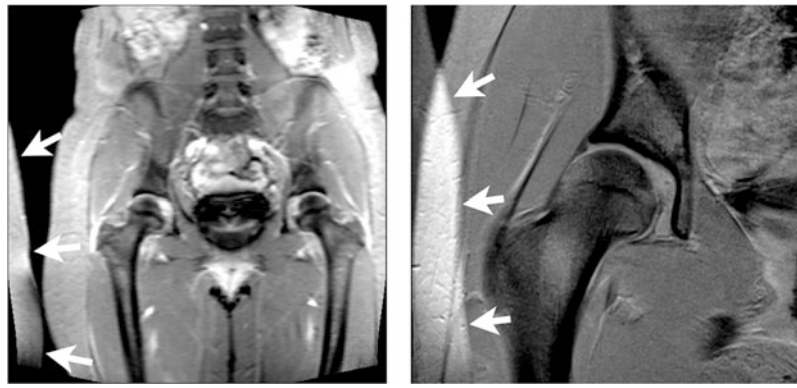

Fig. 1.11 The *white arrows* in the two images point at tissue that is physically located to the *right* of the image, but which appears within the image due to phase wrap artifacts. In the *right* image it can be seen that such artifacts may appear on top of actual structures within the image

and fat (225 Hz difference at 1.5 T). Chemical shift artifacts are an effect of this difference, since echoes having different frequencies will be interpreted as originating from different positions in the frequency encoding direction.

Thus, a dark band appears where fat and water are shifted away from each other (Fig. 1.12), whereas a bright band appears where fat and water are shifted to overlap. Once recognized, this type of artifacts can generally be disregarded when evaluating the images.

Motion Artifacts

Motion artifacts occur due to patient displacement during image acquisition and are often seen as bright noise or repeating densities in the phase encoding direction of the image (Fig. 1.13). Respiratory motion or patient movement are two common reasons for such displacement. Since an MRI scan has a typical duration in the order of minutes, a patient must remain very still in order to avoid motion artifacts. For patients in pain or disease this is sometimes very troublesome.

The artifacts can be explained by considering that many echoes are required to generate an image. Thus, motion that occurs between these echoes will distort the *k*-space information and cause errors in the final image. Motion artifacts are typically distinguished from other types of artifacts in that they extend across the entire field of view.

Flow Artifacts

Movement of body fluid in MRI can produce several types of artifacts. As the most prominent flow in the body is that of blood, image artifacts are possible whenever a large blood vessel appears within the imaging plane (Fig. 1.14).

Some of these artifacts arise from the location change that the flowing hydrogen nuclei experience during the image-encoding gradients. Since the hydrogen nuclei in the blood are moving through these gradient fields, the effect of the gradients on these nuclei will not be the same as on nuclei in stationary tissue. Thus, image encoding of signal from blood may be incorrect. Furthermore, since the pulsation of the blood flow is usually asynchronous to that of the pulse sequence repetition time, these effects may differ between *k*-space lines. The resulting image artifacts are often of similar character as motion artifacts (although more localized and not as severe).

Flow artifacts can be reduced, for example, by changing the phase encoding direction or by applying saturation bands (described later in this chapter) to suppress signal from blood vessels outside of the field of view. Specialized gradients can sometimes also be applied, designed to compensate for blood flow.

Fig. 1.12 Image of the knee joint. Within the bone marrow are large amounts of fat, which is what provides the primary signal within the bone. Thus, a chemical shift artifact can appear between the bone and the surrounding tissue, which is seen in the image as a black band between bone and cartilage on the femoral side

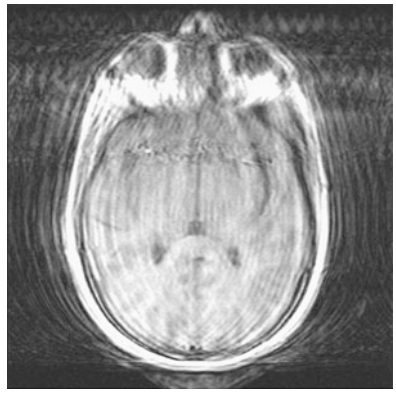

Fig. 1.13 MRI image of the brain. Severe motion artifacts are seen in the horizontal direction of the image

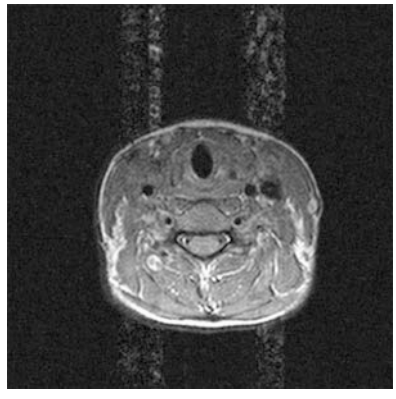

Fig. 1.14 Gradient echo image of the neck with vertical flow artifacts associated with the carotid arteries

Common Imaging Techniques

On a clinical MRI scanner there are typically dozens of different pulse sequences and imaging techniques to choose from, each of which generates images that have different characteristics in terms of acquisition time, signal-to-noise, and image contrast. Some sequences are very similar to each other, while others are more unique. Which pulse sequence should be used depends on the particular application, which is why it is important to understand the different characteristics in order to both request and interpret MRI images. In this section a few of the most frequently used pulse sequences and imaging techniques are described.

Inversion Recovery

Inversion recovery pulse sequences are often used to suppress signals, typically from fat or free fluid, based on their T_1 values. The ability of doing this is valuable in situations when a certain tissue type would otherwise be obscured due to a high fraction of fat or free water in the same local region. In this type of sequence a 180° inversion pulse is applied prior to the excitation pulse. By adjusting the inversion time (TI) between these two pulses, the T_1 values to suppress are decided (Fig. 1.15).

When the inversion pulse is applied all available longitudinal magnetization is inverted, typically from a positive to a negative value. At that time, the longitudinal magnetization will start to recover back towards its positive max value at a speed determined by its T_1 value. Since this recovery starts at a negative value and works its way towards a positive value, it must at some point pass through zero. By setting the inversion time such that the excitation pulse occurs precisely at this time point, any signal from tissue with this specific T_1 value will be effectively nulled.

Provided that a region contains two or more tissue types with different T_1 values, the inversion time can be adjusted such that the signal from one of these tissues is heavily suppressed while the signal from other tissues remains to generate an image. An inversion recovery sequence in which

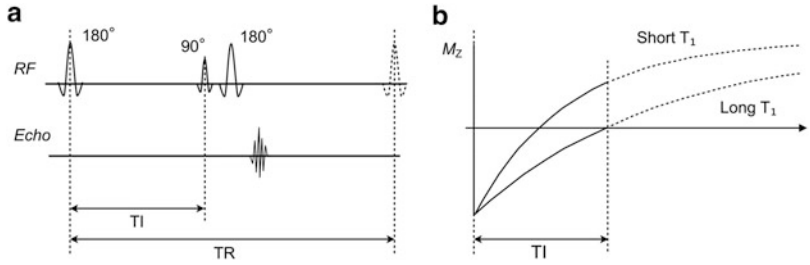

Fig. 1.15 Diagram of an inversion recovery spin echo pulse sequence. Following a 180° inversion pulse, a set of 90° and 180° pulses are applied from which a spin echo is generated (**a**). The inversion pulse and the 90° pulse are separated by an inversion time (TI) during which the longitudinal magnetization recovers (**b**). By adjusting the inversion time appropriately, this recovery can be set such that the signal of tissues having certain T_1 values is effectively nulled, whereas tissues having other T_1 will still yield a usable signal (**b**)

fat is suppressed is often referred to as a STIR sequence (short TI inversion recovery), whereas an inversion recovery sequence where free fluid is suppressed is referred to as a FLAIR sequence (fluid attenuated inversion recovery).

Multiple Spin Echo

A common technique for speeding up the acquisition of a spin echo image is to generate multiple echoes following each single 90° excitation pulse. The technique also goes by several vendor-specific labels, including turbo spin echo (TSE), fast spin echo (FSE), and rapid acquisition with refocusing echoes (RARE).

In a conventional spin echo sequence, a 90° RF pulse is followed by a 180° pulse, after which an echo is generated and encoded to one line in *k*-space. Following the echo a relatively long wait period is applied, until the next 90° pulse. In a multiple spin echo sequence, the echo is instead followed by additional 180° pulses which will refocus the magnetization into additional echoes (cf. Fig. 1.7). Each such echo is generated with a different phase encoding gradient and is assigned to its own line in *k*-space, which can result in a drastically reduced acquisition time (a 256 × 256 scan in typically 30–60 s) compared to the conventional spin echo sequence.

The length of the echo train can be set arbitrarily, although it must be understood that the signal will be decreased between subsequent echoes according to the T_2 value of the tissue. This may result in edge blurring, especially for very long echo trains. In practice, the echo train is usually set at around 5–20 echoes depending on how the image is weighted.

Fast Gradient Echo Sequences

In 3D imaging, gradient echo sequences with short repetition time and low flip angle excitation pulses are often used, as they allow fast acquisition of high-resolution images (a 256 × 256 × 100 scan in typically less than 5 min). This type of images is usually T_1-weighted. Many variations of these sequences exist, which all go by different vendor-specific labels such as fast low-angle shot (FLASH), spoiled gradient echo (SPGR), or fast field echo (FFE).

Since the excitation pulses are typically 5°–30°, only a fraction of the available longitudinal magnetization is lost at each pulse. Due to this, the repetition time can also be kept very short, since not very much magnetization needs to be recovered between pulses. While the advantage of this approach is high acquisition speed, the downside is weaker signal from the echoes and thereby noisier images.

Fast gradient echo imaging is sometimes combined with other techniques, such as a preparatory inversion pule. In such configuration certain features of the inversion recovery sequence can be inherited to the fast gradient echo sequence.

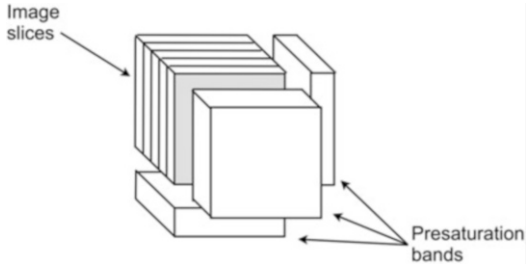

Fig. 1.16 Presaturation bands placed to the *right*, below, and in front of the imaging slices

Presaturation Bands

Presaturation is a method for removing unwanted signal from a specific location. It is particularly useful for reducing phase wrap artifacts and for minimizing artifacts from blood flowing into the image area.

The idea is that a spatially selective 90° pulse is applied which is set to only target a specific region within the body (Fig. 1.16). Immediately after the pulse, the nuclei within this region are said to be saturated since all their longitudinal magnetization is temporarily nulled. By applying such saturation pulses to specific regions immediately before an actual imaging sequence, the lack of longitudinal magnetization will result in no echo signal being generated from these regions.

In situations when it is not feasible to avoid the presence of tissue outside the image in the phase encoding direction, presaturation bands provide a possibility to generate images that are not obscured by phase wrap artifacts. Likewise, when the presence of large blood vessels within an image is unavoidable, presaturation bands can be used for saturating the blood prior to entering the image area, thereby suppressing any signal that would be emitted by the flowing blood.

Quantifying T_1

In certain types of examinations it is of interest to generate images in which each pixel represents the actual T_1 value of the corresponding tissue volume. Such images are generally referred to as T_1 maps. As mentioned previously in this chapter it is not possible to directly generate T_1 maps from an MRI pulse sequence. Instead, T_1 maps can be calculated from a set of two or more images that are acquired using slightly different parameters.

There are several image types that can be used for T_1 quantification, each of which has its own pros and cons in terms of accuracy and acquisition time. In this section two of the most commonly used T_1 quantification methods are briefly described.

2D Inversion Recovery T_1 Quantification

One of the most accurate and stable methods to generate a T_1 map is to use a set of inversion recovery images acquired using different inversion times [4]. This method is generally very stable and usually results in correct T_1 values even if not all parameters are perfectly optimized. The disadvantage is its long acquisition time, which is why this method is only suitable for 2D acquisitions. A 256 × 256 pixel T_1 map is typically acquired in about 5 min.

Commonly, a set of about six inversion recovery images are acquired using inversion times that are spread over an interval larger than the T_1 values of interest (frequently used inversion times are 50, 100, 200, 400, 800, and 1,600 ms). All other parameters, including image position, should be the same for all images.

Consider a specific pixel in each of these six images. As can be understood from Fig. 1.15b the value of this pixel will depend both on its T_1 value and on the inversion time at which it was acquired. The higher the inversion time is, the more magnetization will have had time to recover at the time of the excitation pulse.

The next step is to plot each of these six pixel values in a diagram (Fig. 1.17). As can be seen they will form the shape of a longitudinal recovery curve. Since the equation for this recovery is known, a curve-fitting algorithm can be used to find the T_1 value which generates that particular shape of the curve. Thus, the T_1 value of that specific pixel is calculated. By using computer

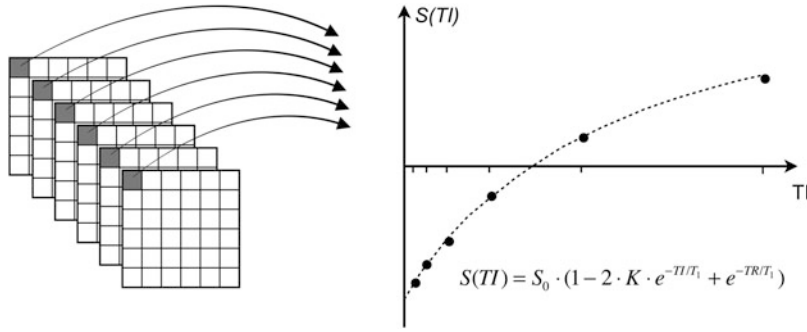

Fig. 1.17 A set of inversion recovery images is acquired using different inversion times (TI). Each same pixel at each TI is then plotted in a diagram where it can be seen that the values follow the longitudinal recovery curve. By fitting these values to the known recovery equation the T_1 value for the pixel is retrieved

software that repeats this process for each pixel, the entire T_1 map is calculated.

3D Dual Flip Angle T_1 Quantification

The dual flip angle T_1 quantification method is common in applications where acquisition of 3D T_1 maps is required. With this method T_1 is quantified using two or more successive fast gradient echo sequences applied with different excitation pulse flip angles. Consequently, 3D dual flip angle T_1 maps can be generated very rapidly using standard 3D gradient echo sequences [5]. A $256 \times 256 \times 100$ voxel T_1 map is typically acquired in less than 10 min.

In a spoiled gradient echo sequence, the longitudinal magnetization is stabilized at a level where exactly as much magnetization is recovered between pulses as is lost at each pulse. This is denoted the longitudinal steady state level and is dependent on a number of parameters, including the T_1 value. Thus, by running the same gradient echo sequence twice, with all parameters equal except the flip angles, a set of equations can be set up from which the T_1 value is solved.

For the dual flip angle method to perform optimally, the flip angles must be selected with regard to the T_1 value to measure. For any T_1 value and repetition time there will be a specific flip angle, denoted the Ernst angle, at which the signal from a gradient echo sequence has a maximum. It has been shown that the optimal flip angle combination for a dual flip angle T_1 measurement is the two flip angles for which the signal is 71 % of that at the Ernst angle (Fig. 1.18) [6]. With this flip angle combination the calculated T_1 values will be least sensitive to any acquired noise.

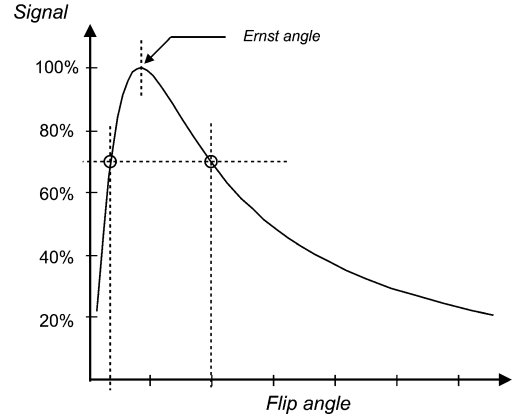

Fig. 1.18 Steady state gradient echo amplitude as a function of the excitation pulse flip angle, for some arbitrary T_1 and TR. The optimal flip angle combination for T_1 mapping is marked with *circles*

The dual flip angle method requires accurate knowledge about the flip angles of the comprised gradient echo sequences in order for the T_1 values to be correctly calculated. The most widely used way of achieving this is to simply assume that the nominal flip angles, as specified in the user interface of the scanner, are sufficiently accurate and

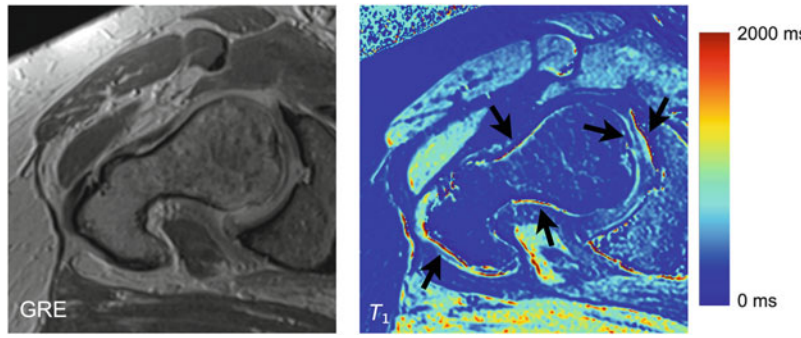

Fig. 1.19 Right image is a T_1 map covering the femoral head and neck. *Left* image is one of the gradient echo images from which the T_1 map is calculated. The *black arrows* in the *right* image point out artifacts due to patient movement, which appears as narrow bands of extremely high or extremely low T_1 values along the bone structures

use them for T_1 calculation. However, there are a number of effects that may be hard to predict which can cause the actual flip angle to deviate severely, resulting in erroneous T_1 calculations. Especially at high magnetic field strength (i.e., 3 T), these effects may cause unacceptable errors [7].

This type of errors can be corrected for if the accuracy of the flip angle is measured at different positions within the field of view. Such measurement is commonly known as a B_1 map, which can be calculated using an additional pulse sequence specifically designed for this purpose [8]. This type of additional B_1 mapping sequence typically will add a few minutes of scan time, which is usually time well spent, considering the gained T_1 mapping accuracy.

Since T_1 maps are calculated from a combination of several images, the accuracy of the T_1 calculation is highly sensitive to patient movement between the different images. This is true for all types of T_1 measurements, although it is of particular importance for dual flip angle T_1 maps since those are calculated from only two images. Errors due to patient movement are sometimes very hard to notice. For a dual flip angle T_1 map such errors are often seen as narrow bands of extremely high (>2,000 ms) or extremely low (<100 ms) T_1 values that appear along the borders of different structures (especially bone structures, Fig. 1.19). If such artifacts are seen, new T_1 maps must be generated. This can be done either by reacquisition of the images or, in some circumstances, by just recalculating the T_1 maps after registering the existing images.

References

1. Mitchell DG, Cohen MS. MRI principles. 2nd ed. Philadelphia, PA: Saunders; 2004.
2. Bernstein MA, King KF, Zhou XJ. Handbook of MRI pulse sequences. London: Academic; 2004.
3. Haacke EM, Brown RW, Thompson MR, Venkatesan R. Magnetic resonance imaging: physical principles and sequence design. New York: Wiley; 1999.
4. Gupta R, Feretti J, Becker E, Weiss G. A modified fast inversion-recovery technique for spin–lattice relaxation measurements. J Magn Reson. 1980;38:447–52.
5. Brookes JA, Redpath TW, Gilbert FJ, Murray AD, Staff RT. Accuracy of T1 measurement in dynamic contrast-enhanced breast MRI using two- and three-dimensional variable flip angle fast low-angle shot. J Magn Reson Imaging. 1999;9(2):163–71.
6. Deoni SC, Rutt BK, Peters TM. Rapid combined T1 and T2 mapping using gradient recalled acquisition in the steady state. Magn Reson Med. 2003;49(3):515–26.
7. Siversson C, Chan J, Tiderius CJ, Mamisch TC, Jellus V, Svensson J, Kim YJ. Effects of B1 inhomogeneity correction for three-dimensional variable flip angle T1 measurements in hip dGEMRIC at 3 T and 1.5 T. Magn Reson Med. 2012;67(6):1776–81.
8. Stollberger R, Wach P. Imaging of the active B1 field in vivo. Magn Reson Med. 1996;35(2):246–51.