

Introduction to Magnetic Resonance Imaging Techniques

Lars G. Hanson, larsh@drcmr.dk
Danish Research Centre for Magnetic Resonance (DRCMR),
Copenhagen University Hospital Hvidovre
Latest document version: <http://www.drcmr.dk/>

Translation to English: Theis Groth

Revised: August, 2009

It is quite possible to acquire images with an MR scanner without understanding the principles behind it, but choosing the best parameters and methods, and interpreting images and artifacts, requires understanding. This text serves as an introduction to magnetic resonance imaging techniques. It is aimed at beginners in possession of only a minimal level of technical expertise, yet it introduces aspects of MR that are typically considered technically challenging. The notes were written in connection with teaching of audiences with mixed backgrounds.

Contents

1	Introduction	2
1.1	Supplemental material	3
1.2	Recommended books	3
2	Magnetic resonance	4
3	The magnetism of the body	8
4	The rotating frame of reference	12
5	Relaxation	12
5.1	Weightings	13
5.2	Causes of relaxation	14
5.3	Inhomogeneity as a source of signal loss, T_2^*	15
6	Sequences	16
7	Signal-to-noise and contrast-to-noise ratios	17

8	Quantum Mechanics and the MR phenomenon	18
8.1	Corrections	19
9	Imaging	20
9.1	Background	21
9.2	Principles	21
9.3	Slice selection	22
9.4	Spatial localization within a slice	22
9.5	Extension to more dimension – k -space	24
9.6	Similarity and image reconstruction	26
9.7	Moving in k -space	28
9.8	Image acquisition and echo-time	31
9.9	Frequency and phase encoding	31
9.10	Spatial distortions and related artifacts	32
9.11	Slice-selective (2D-) versus 3D-sequences	33
9.12	Aliasing and parallel imaging	35
9.13	Finishing remarks on the subject of imaging	35
10	Noise	36
11	Scanning at high field	37
12	MR Safety	38

1 Introduction

This text was initially written as lecture notes in Danish. Quite a few of the students are English speaking, however. Moreover, the approach taken is somewhat different than for most other introductory texts. Hence it was deemed worth the effort to do a translation. The task was taken on by Theis Groth (July 2009 version) whose efforts are much appreciated.

The goal of the text is to explain the technical aspects accurately without the use of mathematics. Practical uses are not covered in the text, but it introduces the prerequisites and, as such, seeks to provide a sturdy base for further studies. While reading, you may find the glossary of help (located towards the end of the notes). It elaborates on some concepts, and introduces others.

Material has been added along the way whenever a need was felt. As a result, the text covers diverse subjects. During the accompanying lectures, software that illustrates important aspects of basic MR physics, including popular sequences, contrast and imaging, was used. It is highly recommended that you download and experiment with these programs, which are available at no charge. Corrections, comments and inspiring questions concerning the text and software are always appreciated.

The text is concerned only with general aspects that are discussed in pretty much any technical textbook on MR. Therefore, very few specific references are given.

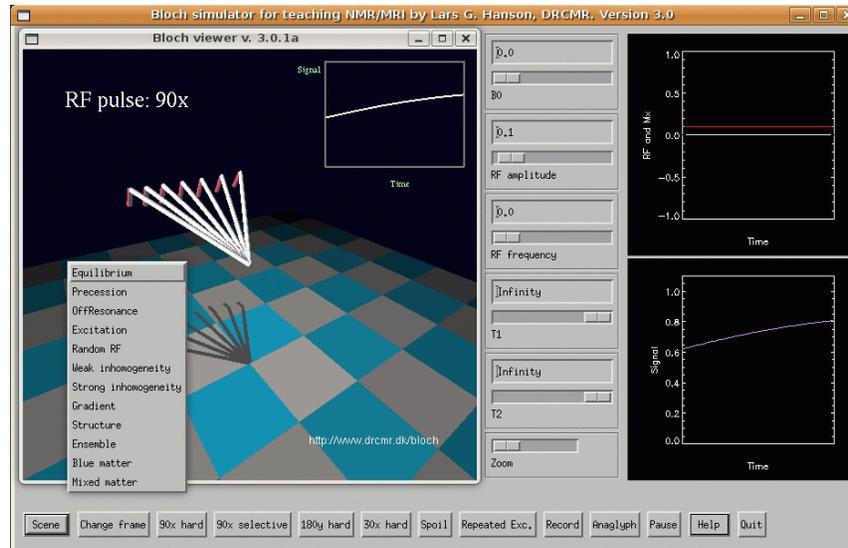


Figure 1: For teaching purposes, a free, interactive simulator, that illustrates a large number of important concepts and techniques, has been developed. The user chooses the initial conditions for the simulation and can thereafter manipulate the components of the magnetization with radio waves and gradients, as is done during a real scanning session. A screenshot from the simulator is shown above. It is described in detail at <http://www.drcmr.dk/bloch>

1.1 Supplemental material

<http://www.drcmr.dk/bloch> Interactive software that can illustrate a broad spectrum of MRI concepts and methods, and that can contribute significantly to the understanding. Animations that illustrate the use of the software and selected MR phenomena are on the homepage.

<http://www.drcmr.dk/MR> Text and animations that briefly describes MR in English. The page was created as a supplement to an article that discusses the connection between the classical and the quantum mechanical approach to the MR phenomenon, and which also describes problems that appear in many basic text books.

<http://www.drcmr.dk/MRbasics.html> Links to this text and an example of accompanying “slides”.

1.2 Recommended books

Even technically oriented readers can benefit from reading a relatively non-technical introduction. Having a basic understanding makes it much easier to understand the formalism in later stages. Once a basic understanding is acquired, the Internet is a good source of more specific texts. Some books deserve special mention:

“Magnetic Resonance Imaging: Physical and Biological Principles” by Stewart C. Bushong. A good introduction that does not demand a technical background from its readers.

“Clinical Magnetic Resonance Imaging” by Edelman, Hesselink and Zlatkin. Three volumes featuring a good mixture of technique and use. Not an intro, but a good follow-up (according to people who have read it. I haven’t).

“Magnetic Resonance Imaging – Physical Principles and Sequence Design” by Haacke, Brown, Thompson and Venkatesan. Broadly oriented textbook with plenty of physics, techniques and sequences. Not an easily read introduction, but suitable for physicists and similar people.

“Principles of Nuclear Magnetic Resonance Microscopy” by Paul T. Callaghan. A classic within the field of probing molecular dynamics. Technically demanding. Should only be read in the company of a grown-up.

“Spin Dynamics: Basics of Nuclear Magnetic Resonance” by Malcolm H. Levitt. Covers theoretical aspects of MR spectroscopy as used in chemical analysis and is thus irrelevant to most who work with MR image analysis. Excels in clear, coherent writing and a description of quantum mechanics devoid of typical misconceptions.

Please note that these books have been published in several editions.

2 Magnetic resonance

Initially it is described how magnetic resonance can be demonstrated with a pair of magnets and a compass. The level of abstraction already increases towards the end of this first section, but despair thee not: As the complexity rises, so shall it fall again. A complete understanding of earlier sections is not a prerequisite for future gain.

If a compass happens to find itself near a powerful magnet, the compass needle will align with the field. In a normal pocket compass, the needle is embedded in liquid to dampen its oscillations. Without liquid, the needle will vibrate through the north direction for a period before coming to rest. The frequency of the oscillations depend on the magnetic field and of the strength of the magnetic needle. The more powerful these are, the faster the vibrations will be.

Radio waves are magnetic fields that change in time (oscillate¹) and as long as the needle vibrates, weak radio waves will be emitted at the same frequency as that of the needle. The frequency is typically reported as the number of oscillations per second, known as Hertz (Hz). If the needle oscillates 3 times per second, radio waves will be emitted at a frequency of 3 Hz, for example. The

¹Throughout this text, the term “radio waves” is used for rapidly oscillating magnetic fields. It has been pointed out by David Hoult in particular, that the wave nature of the oscillating magnetic field is not important for MR except at ultra-high field. The typical choice of wording is therefore unfortunate (“oscillating magnetic field” or “ B_1 field” is preferred over “radio wave field”). I agree, but will nevertheless comply with the somewhat unfortunate standard as it facilitates reading and correctly highlights the oscillatory nature of the B_1 -field. The difference lies in the unimportant electrical component of the field, and in the spatial distribution not felt by individual nuclei.

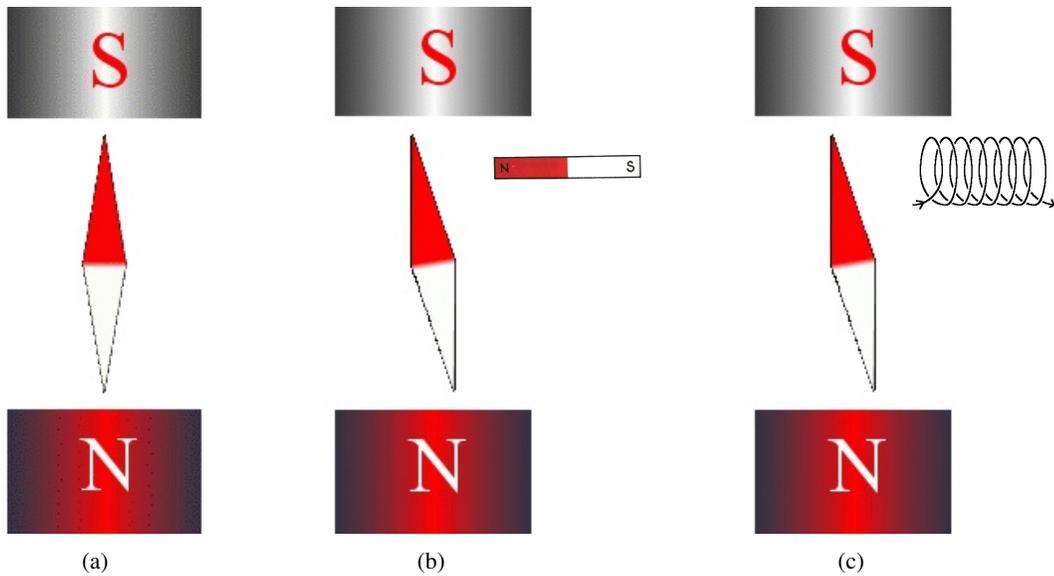


Figure 2: (a) A magnetic needle in equilibrium in a magnetic field. The needle orients itself along the field lines that go from the magnet's northern pole to the southern pole (and yes, the magnetic south pole of the earth is close to the geographical North Pole). (b) A weak perpendicular magnetic field can push the magnet a tad away from equilibrium. If the magnetic field is changed rhythmically in synchrony with the oscillations of the needle, many small pushes can result in considerable oscillation. (c) The same effect can be achieved by replacing the weak magnet with an electromagnet. An alternating current will shift between pushing and pulling the north pole of the magnetic needle (opposite effect on the south pole). Since the field is most powerful inside the coil, the greatest effect is achieved by placing the needle there.

strength of the radio waves (known as the amplitude) is weakened as the oscillations of the needle gradually vanish.

Imagine the following situation: A compass is placed in a magnetic field created by one or more powerful magnets. After a short period of time, the needle has settled and is pointing in the direction of the magnetic field, figure 2(a). If the needle is given a small push perpendicular to the magnetic field (a rotation), it will vibrate through north, but gradually settle again. The oscillations will occur at a frequency that will hereafter be referred to as the resonance frequency. As long as the magnetic needle is oscillating, radio waves with the same frequency as the oscillation will be emitted. The radio waves die out together with the oscillations, and these can in principle be measured with an antenna (coil, figure 2(c)). The measurement can, for example, tell us about the strength of the magnetic needle and the damping rate of its oscillations.

As will be made clear, there are also “magnetic needles” in the human body, as hydrogen nuclei are slightly magnetic. We can also manipulate these so they oscillate and emit radio waves, but it is less clear how we may give the needles the necessary “push.” In order to make an ordinary compass needle swing, we may simply give it a small push perpendicular to the magnetic field with a finger, but this is critically less feasible in the body. Instead, we may take advantage of the fact that

magnets influence other magnets, so that the weak push perpendicular to the magnetic field can be delivered by bringing a weak magnet near the needle, as shown in figure 2(b). With this method, we may push the needle *from a distance* by moving a weak magnet towards the needle and away again.

In an MR scanner, the powerful magnet is extremely powerful for reasons that will be explained below. The magnetic field across is quite weak in comparison. The push that the weak magnet can deliver is therefore too weak to produce radio waves of any significance, if the magnetic needle is only “pushed” once.

If, on the other hand, many pushes are delivered in synchrony with the mentioned oscillations of the magnetic needle, even waving a weak magnet can produce strong oscillations of the magnetic needle. This is achieved if the small magnet is moved back and forth at a frequency identical to the natural oscillation frequency, the resonance frequency, as described above.

What is described here is a classical *resonance phenomenon* where even a small manipulation can have a great effect, if it happens in synchrony with the natural oscillation of the system in question. Pushing a child on a swing provides another example: Even weak pushes can result in considerable motion, if pushes are repeated many times and delivered in synchrony with the natural oscillations of the swing.

Let us focus on what made the needle oscillate: It was the small movements of the magnet, back and forth, or more precisely *the oscillation of a weak magnetic field perpendicular to the powerful stationary magnetic field* caused by the movement of the magnet.

But oscillating magnetic fields is what we understand by “radio waves”, which means that in reality, we could replace the weak magnet with other types of radio wave emitters. This could, for example, be a small coil subject to an alternating current, as shown in figure 2(c). Such a coil will create a magnetic field perpendicular to the magnetic needle.

The field changes direction in synchrony with the oscillation of the alternating current, so if the frequency of the current is adjusted to the resonance frequency of the magnetic needle, the current will set the needle in motion. When the current is then turned off, the needle will continue to swing for some time. As long as this happens, radio waves will be emitted by the compass needle. The coil now functions as an antenna for these radio waves. In other words, a voltage is induced over the coil because of the oscillations of the magnetic field that the vibration of the magnetic needle causes.

In summary, the needle can be set in motion from a distance by either waving a magnet or by applying an alternating current to a coil. In both situations, magnetic resonance is achieved when the magnetic field that motion or alternating currents produce, oscillates at the resonance frequency. When the waving or the alternating current is stopped, the radio waves that are subsequently produced by the oscillating needle, will induce a voltage over the coil, as shown. The voltage oscillates at the resonance frequency and the amplitude decreases over time. A measurement of the voltage will reflect aspects of the oscillating system, e.g. “the relaxation time”, meaning the time it takes for the magnetic needle to come to rest.

The above mentioned experiment is easily demonstrated, as nothing beyond basic school physics is involved. Nobody acquainted with magnets and electromagnetism will be surprised during the experiment. Nonetheless, the experiment reflects the most important aspects of the basic MR phenomenon, as it is used in scanning.

There are, however, differences between MR undertaken with compasses and nuclei. These are due to the fact that the nuclei are not only magnetic, but also rotate around an axis of their own, as shown in figure 3 to the right (a twist to this is discussed in section 8). The rotation, called *spin*, makes the nuclei magnetic along the rotational axis. This is equivalent to a situation where the compass needle described above constantly and quickly rotates around its own length direction (has spin or angular momentum). Such a needle would swing *around* north in a conical motion if the mounting allowed for this, rather than swing in a plane through north. This movement is called *precession* and it is illustrated in figure 4. Think of it as a variant of the normal needle oscillation. Similarly, a pendulum can swing around a vertical axis rather than in a vertical plane, but basically there is no great difference. In the rest of this section, the consequences of spin are elaborated on. The finer nuances are not important for understanding MR scanning.

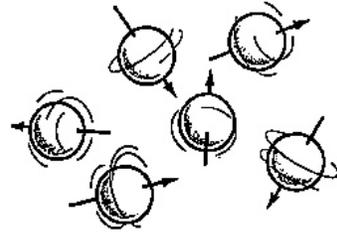


Figure 3: The spin of the nuclei (rotation) makes them magnetic.

Precession around the direction of a field is also known from a spinning top (the gravitational field rather than the magnetic field, in this example): A spinning top that rotates quickly will fall slowly, meaning that it will gradually align with the direction of the gravitational field. Rather than simply turning downwards, it will slowly rotate around the direction of gravity (precess) while falling over. These slow movements are a consequence of the fast spin of the spinning top.

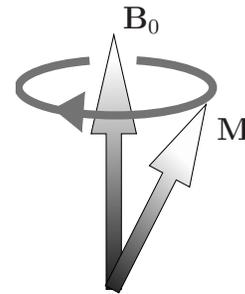


Figure 4: A magnetization \mathbf{M} that precess around the magnetic field \mathbf{B}_0 because of spin (rotation around \mathbf{M}).

The same applies to a hypothetical magnetic needle that rotates quickly around its own length axis, as the atomic nuclei in the body do. In the experiment described above, such a magnetic needle with spin would rotate around north (precess) after having received a push. It would spiral around north until it eventually pointed in that direction, rather than simply vibrate in a plane as a normal compass needle. It would meanwhile emit radio waves with the oscillation frequency, which we may now also call the precession frequency. This frequency is independent of the amplitude, i.e., independent of the size of the oscillations. This consequence of spin is another difference to the situation shown in figure 2 since the natural oscillation frequency of a normal compass is only independent of amplitude for relatively weak oscillations.

Like before, it is possible to push the magnetic needle with a weak, perpendicular magnetic field that oscillates in synchrony with the precession (meaning that it oscillates at the resonance frequency). Just as the magnetic needle precess around the stationary magnetic field, it will be rotated *around* the weak, rotating magnetic field, rather than (as before) moving directly in the direction of it. In practice, this means that a perpendicular magnetic field, that rotates in synchrony with the precession of the needle, will rotate it slowly around the rotating fields direction (meaning a slow precession around the weak rotating magnetic field).

We have now in detail described the magnetic resonance phenomenon, as employed for scanning: The influence of spin on the movement of the magnetic axis can, at first glance, be difficult to

understand: That the force in one direction can result in pushes towards another direction may appear odd (that the pull in a magnet needle *towards* north will make it rotate *around* north, if the needle rotates around its own length axis (has spin)). It does, however, fall in the realms of classical mechanics.

One does not need to know why spinning tops, gyros, nuclei and other rotating objects act queer in order to understand MR, but it is worth bearing in mind that the spin axis rotates *around* the directions of the magnetic fields. The precession concept is thus of importance, and it is also worth remembering that spin and precession are rotations around two different axes.

3 The magnetism of the body

Equipped with a level of understanding of how magnet needles with and without spin are affected by radio waves, we now turn to the “compass needles” in our very own bodies.

- Most frequently, the MR signal is derived from hydrogen nuclei (meaning the atomic nuclei in the hydrogen atoms). Most of the body’s hydrogen is found in the water molecules. Few other nuclei are used for MR.
- Hydrogen nuclei (also called protons) behave as small compass needles that align themselves parallel to the field. This is caused by an intrinsic property called nuclear spin (the nuclei each rotate as shown in figure 3). By the “direction of the nuclear spins” we mean the axis of rotation and hence the direction of the individual “compass needles”.
- The compass needles (the spins) are aligned in the field, but due to movements and nuclear interactions in the soup, the alignment only happens partially, as shown in figure 5 – very little, actually. There is only a weak tendency for the spins to point along the field. The interactions affect the nuclei more than the field we put on, so the nuclear spins are still *largely* pointing randomly, even after the patient has been put in the scanner. An analogy: If you leave a bunch of compasses resting, they will all eventually point towards north. However, if you instead put them into a running tumble dryer, they will point in all directions, and the directions of individual compasses will change often, but there will still be a weak tendency for them to point towards north. In the same manner, the nuclei in the body move among each other and often collide, as expressed by the temperature. At body temperature there is only a weak tendency for the nuclei to point towards the scanner’s north direction.

Together, the many nuclei form a total magnetization (compass needle) called *the net magnetization*. It is found, in principle, by combining all the many contributions to the magnetization, putting arrows one after another. If an equal number of arrows point in all directions, the net magnetization will thus be zero. Since it is generally the sum of many contributions that swing in synchrony as compass needles, the net magnetization itself swings as a compass needle. It is therefore *adequate to keep track of the net magnetization rather than each individual contribution to it*.

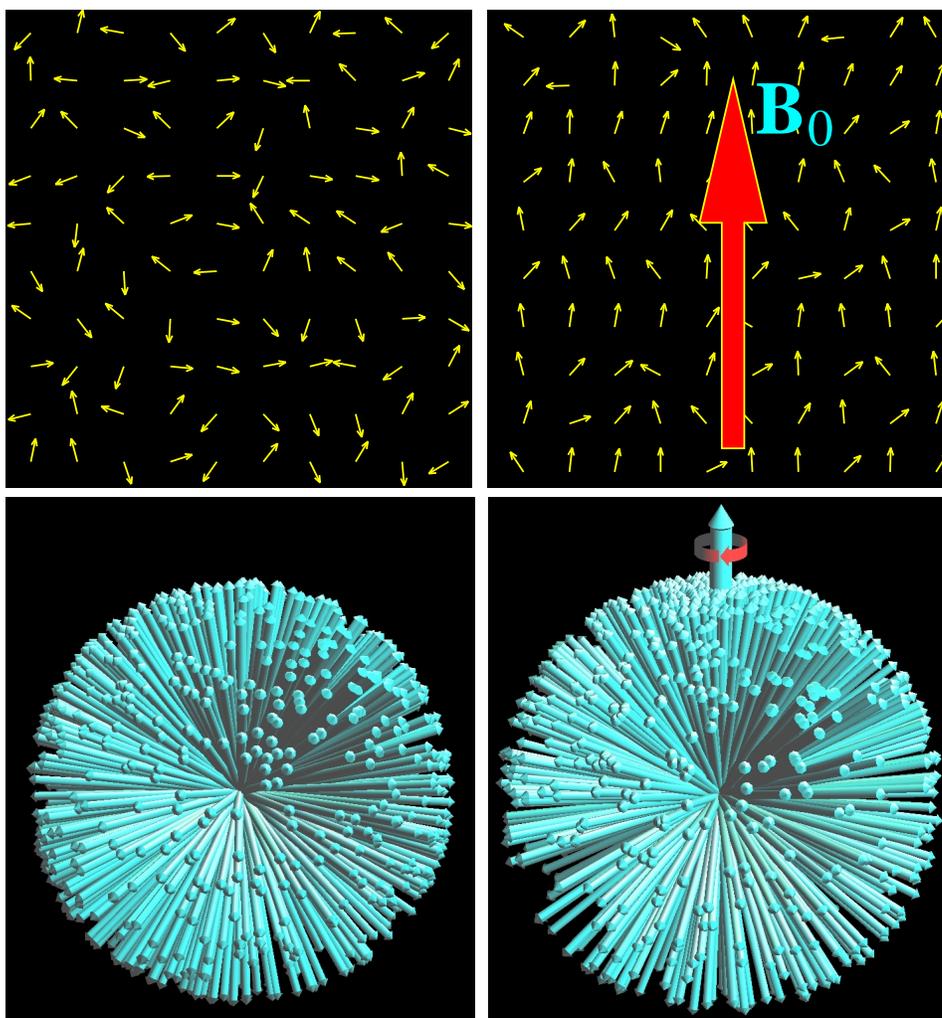


Figure 5: The figure shows the same situations in two and three dimensions (top and bottom, respectively). The nuclear spins are shown as numerous arrows (vectors). In the lower graphs, they are all drawn as beginning in the same point, so that the distribution over directions is clear (implicit coordinate system (M_x, M_y, M_z)). When a patient arrives in the ward, the situation is as shown in the two graphs to the left: The spins are oriented randomly, with a uniform distribution over directions, meaning that there is about an equal number of spins pointing in all directions. The net magnetization is near zero and the nuclei do not precess. When a magnetic field \mathbf{B}_0 is added, as shown in the two figures to the right, a degree of alignment (order) is established. The field is only shown explicitly in the top right figure, but the effect is visible in both: The direction distribution becomes “skewed” so that a majority of the nuclei point along the field. In the lower right figure, the net magnetization (thick vertical arrow) and the precession (the rotation of the entire ball caused by the magnetic field) are shown. The lower figures appear in the article *Is Quantum Mechanics necessary for understanding Magnetic Resonance?* Concepts in Magnetic Resonance Part A, 32A (5), 2008.

As mentioned above, the nuclei in the body move among each other (thermal motion) and the net magnetization in equilibrium is thus temperature dependent. Interaction between neighboring nuclei obviously happens often in liquids, but they are quite weak due to the small magnetization of the nuclei. Depending on the character and frequency of the interaction, the nuclei precess relatively undisturbed over periods of, for example, 100 ms duration. At any time, there is a certain probability that a nucleus takes part in a dramatic clash with other nuclei, and thus will point in a new direction, but this happens rather infrequently.

The net magnetization is equivalent to only around 3 per million nuclear spins oriented along the direction of the field (3 ppm at 1 tesla). This means that the magnetization of a million partially aligned hydrogen nuclei in the scanner equals a total magnetization of only 3 completely aligned nuclei.

With the gigantic number of hydrogen nuclei (about 10^{27}) found in the body, the net magnetization still becomes measurable. It is proportional to the field: A large field produces a high degree of alignment and thus a large magnetization and better signal-to-noise ratio.

- If the net magnetization has been brought away from equilibrium, so it no longer points along the magnetic field, it will subsequently precess around the field with a frequency of 42 million rotations/second at 1 tesla (42 MHz, megahertz). This is illustrated in figure 4. Eventually it will return to equilibrium (relaxation), but it takes a relatively long time on this timescale (e.g. 100 ms as described above). Meanwhile, radio waves at this frequency are emitted from the body. We measure and analyze those. Notice: The position of the nuclei in the body does not change - only their axis of rotation.
- The precession frequency is known as the Larmor frequency in the MR community. The Larmor equation expresses a connection between the resonance frequency and the magnetic field, and it is said to be the most important equation in MR:

$$f = \gamma B_0$$

The equation tells us that the frequency f is proportional to the magnetic field, B_0 . The proportionality factor is 42 MHz/T for protons. It is called “the gyromagnetic ratio” or simply “gamma”. Thus, the resonance frequency for protons in a 1.5 tesla scanner is 63 MHz, for example.

The Larmor equation is mainly important for MR since it expresses the possibility of designing techniques based on the frequency differences observed in inhomogeneous fields. Examples of such techniques are imaging, motion encoding and spectroscopy.

- But how is the magnetization rotated away from its starting point? It happens by applying radio waves at the above mentioned frequency.

Radio waves are magnetic fields that change direction in time. The powerful stationary field pushes the magnetization so that it precesses. Likewise the radio waves push the magnetization around the radio wave field, but since the radio wave field is many thousand times weaker than the static field, the pushes normally amount to nothing.

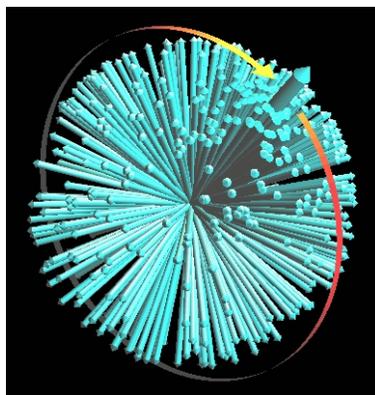


Figure 6: Scene from animation found at <http://www.drcmr.dk/MR> that shows how radio waves affect a collection of nuclear spins precessing around B_0 (vertical axis) at the Larmor frequency. The radio wave field that rotates around the same axis at the same frequency, induces simultaneous rotation around a horizontal axis, as symbolized by the circular arrow. The relative orientation of the nuclei does not change, and it is therefore adequate to realize how the net magnetization (here shown by a thick arrow) is affected by the magnetic field.

Because of this, we will exploit a resonance phenomenon: By affecting a system rhythmically at an appropriate frequency (the systems resonance frequency), a large effect can be achieved even if the force is relatively weak. A well-known example: Pushing a child sitting on a swing. If we push in synchrony with the swing rhythm, we can achieve considerable effect through a series of rather weak pushes. If, on the other hand, we push against the rhythm (too often or too rarely) we achieve very little, even after many pushes.

With radio waves at an appropriate frequency (a resonant radio wave field), we can slowly rotate the magnetization away from equilibrium. “Slowly” here means about one millisecond for a 90 degree turn, which is a relatively long time as the magnetization precesses 42 million turns per second at 1 tesla (the magnetization rotates 42 thousand full turns in the time it takes to carry out a 90 degree turn, i.e., quite a lot faster).

Figure 6 is a single scene from an animation found at <http://www.drcmr.dk/MR> that shows how a collection of nuclei each precessing around both B_0 and a rotating radio wave field as described earlier, together form a net magnetization that likewise moves as described.

The strength of the radio waves that are emitted from the body depends on the size of the net magnetization and on the orientation. The greater the oscillations of the net magnetization, the more powerful the emitted radio waves will be. The signal strength is proportional to the component of the magnetization, that is perpendicular to the magnetic field (the transversal magnetization), while the parallel component does not contribute (known as the longitudinal magnetization). In figure 4, the size of the transversal magnetization is the circle radius.

If the net magnetization points along the magnetic field (as in equilibrium, to give an example) no measurable radio waves are emitted, even if the nuclei do precess individually. This is

because the radio wave signals from the individual nuclei are not in phase, meaning that they do not oscillate in synchrony perpendicular to the field. The contributions thereby cancel in accordance with the net magnetization being stationary along the B_0 -field (there is no transversal magnetization).

- The frequency of the radio waves is in the FM-band so if the door to a scanner room is open, you will see TV and radio communication as artifacts in the images. At lower frequencies we find line frequencies and AM radio. At higher frequencies, we find more TV, mobile phones and (far higher) light, X-ray and gamma radiation. From ultra-violet light and upwards, the radiation becomes “ionizing”, meaning that it has sufficient energy to break molecules into pieces. MR scanning uses radio waves very far from such energies. Heating, however, is unavoidable, but does not surpass what the body is used to.

4 The rotating frame of reference

Confusion would arise if the descriptions to come continued to involve both a magnetization that precesses and radio waves that push this in synchrony with the precession. It is simply not practical to keep track of both the Larmor-precession and the varying radio wave field at the same time, and this is necessary to establish the direction of the pushes, and thus how the magnetization changes. Instead, we will now change our perspective.

An analogy: We describe how the horse moves on a merry-go-round in action. Seen by an observer next to the merry-go-round, the horse moves in a relatively complicated pattern. However, if the merry-go-round is mounted, the movement of the horse appears limited to a linear up-and-down movement. We say that we have changed from the stationary frame of reference to a rotating frame of reference.

In the same vein, we can simplify the description of the resonance phenomenon by moving to a rotating frame of reference and getting rid of the Larmor-precession. We mount the merry-go-round that rotates with the Larmor frequency and recognize that in this frame the magnetization is stationary until we employ radio waves. The effect of B_0 therefore appears to be gone, as shown in figure 7.

5 Relaxation

Interactions happening at near-collisions between nuclei give rise to the magnetization constantly approaching the equilibrium size. This is called relaxation. The speed at which relaxation occurs depends on the protons interactions with their neighbors, which in turn depends on the firmness of the substance (the consistency). It is the difference in consistency and the presence of large molecules that limit the waters free movement, which causes most of the contrast we see in MR images.

The relaxation occurs on two different time scales: The magnetization perpendicular to the magnetic field (the transversal magnetization) often decreases relatively rapidly, while it can take considerably longer to recover the magnetization along the field (the longitudinal magnetization).

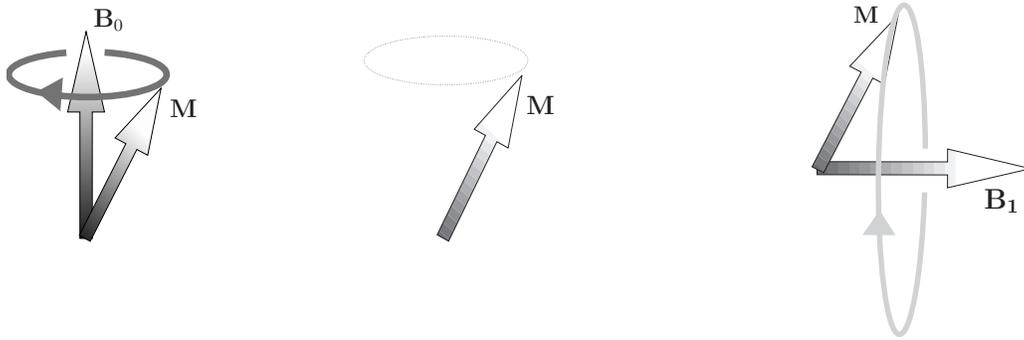


Figure 7: In the stationary frame of reference, a rapid precession occurs (left). In the rotating frame of reference (middle) the magnetization usually only moves slowly due to relaxation and smaller differences in frequency. When resonant radio waves are transmitted, the effect of these is to induce precession around an axis perpendicular to the B_0 field (right). Descriptions in the rotating and stationary frames of reference are equally precise, but explanations are almost always given in the rotating frame of reference, since the movement of the magnetization is more simple there. Furthermore, the scanner always presents measurements as if they were obtained in the rotating frame of reference. This happens because the scanner removes the rapid variation in the signals before they are stored in the computer (the signals are modulated down from the Larmor frequency to lower frequencies).

- The transversal magnetization (M_{xy}) decreases exponentially on a timescale T_2 (e.g. around 100 ms for brain tissue. Several seconds for pure water).
- The longitudinal magnetization (M_z) approaches equilibrium M_0 on a timescale T_1 (for example approximately 1 s for brain tissue. Several seconds for pure water).

The relaxation times depend on the mobility of the molecules and the strength of the magnetic field, as discussed in section 5.2.

5.1 Weightings

The contrast in an MR-image is controlled by the choice of measuring method (sequence and sequence parameters, which will be discussed later). For example, we call an image T_2 -weighted if the acquisition parameters are chosen so the image contrast mainly reflect T_2 -variations. One must understand, however, that even in a heavily T_2 -weighted image, the contrast will often reflect more than just T_2 -variation. To provide an example, variation in water content always results in some contrast.

The echo time, TE, is the period from we rotate the magnetization into the transversal plane until we decide to measure the radio waves (a more precise definition will follow later). Meanwhile, a loss of magnetization and signal will occur due to T_2 -relaxation. The echo time is thus the period within the measurement which gives T_2 -weighting in the images. A long TE compared to T_2 will thus result in considerable T_2 -contrast, but only little signal. The greatest sensitivity to T_2 -variation will be achieved when $TE \simeq T_2$.

Often, we will repeat similar measurements several times, e.g. once per line in an image. The

repetition-time, TR, is the time between these repetitions. Every time we make such a measurement, we (partially) use the longitudinal magnetization present (the magnetization is rotated into the transversal plane which results in emission of radio waves while the transversal component gradually disappears). If we use the magnetization often (short TR), every repeat will therefore only produce a small signal. If we, on the other hand, wait longer between repetitions (long TR), the magnetization will nearly recover to equilibrium between repetitions. What is meant by short and long TR? It is relative to T_1 that is the time scale on which the longitudinal magnetization is rebuilt.

If the magnetization is completely rebuilt between measurements for all types of tissue in the scanner, meaning if TR is significantly longer than the maximum T_1 , the T_1 -contrast will disappear. In this case, the transversal magnetization immediately following the rotation of the nuclei reflects the equilibrium magnetization. The radio waves do so, as well. The equilibrium magnetization is governed by the hydrogen concentration, also known as the proton density (PD).

Thus, we may conclude that using a long TR results in limited T_1 -weighting but a strong signal. If we apply a shorter TR, the signal is reduced for all types of tissue, but the signal becomes more T_1 -weighted, meaning that the images will be less intense, but with a relatively greater signal variation between tissues with differing T_1 .

Finally, we can minimize both T_1 - and T_2 -contrast, which results in a PD-weighted image. In such an image, variation in the water content is the primary source of contrast, since the proton density is the density of MR-visible hydrogen that is largely proportional to the water content.

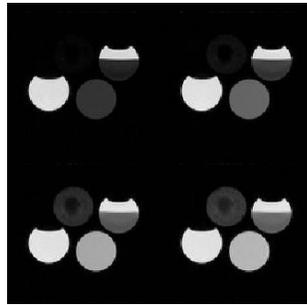
Summarized, the following apply to weightings in a simple excitation-wait-measure-wait-repeat sequence.

- T_1 -weighted images are made by applying short TR and short TE, since T_1 contrast is thereby maximized and T_2 contrast is minimized.
- T_2 -weighted images are made by applying long TR and long TE, since T_1 contrast is thereby minimized and T_2 contrast maximized.
- PD-weighted images are made by applying long TR and short TE, since both T_1 - and T_2 -contrast are thereby minimized. Meanwhile, the signal is maximized, but it is of no use if the contrast disappears, since only small variations exist in the water content of the tissue.

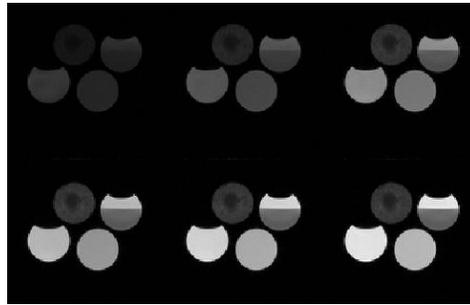
The “lacking” combination of long TE and short TR results in a weak signal and mixed T_1 - and T_2 contrast. Images of bottles with jello in picture 8 illustrates how the contrast and signal in an image varies with TE, TR and consistency.

5.2 Causes of relaxation

The difference between T_1 and T_2 is due to a difference in the causes of relaxation. For the protons in firm matter, the spins will rapidly be dephased after excitation, meaning that they will point in all directions perpendicular to the field (have all kinds of *phases*). This is caused by the small local contributions to the magnetic field that the individual spins are making, and that make the neighboring nuclei precess at altered frequencies. In firm matter, these interactions are constant in time, while they vary in fluids, since nuclei are constantly experiencing new neighbors. Thus, the spins can remain in phase for relatively long (seconds) in fluids, while they lose their common



(a) Four images, all obtained with a common TR=5 seconds and TE=90, 50, 20, 15 ms (shown in reading order).



(b) Six images obtained with a common TE=15 ms and TR=500, 1000, 2000, 3000, 4000, 5000 ms (shown in reading order).

Figure 8: Phantom data which illustrates signal intensity and contrast for bottles filled with jello of varying consistency. Where is T_1 long/short? How long, how short? The same for T_2 ? Which bottles might be pure water? Which jello is most firm? What pictures are the most T_1 -, T_2 - and PD-weighted?

orientation in a matter of milliseconds or less in firm matter. We can therefore conclude that T_2 is short in firm matter.

The described process affects the individual spins Larmor-frequency, but does not give rise to a change of the longitudinal magnetization, since the interaction of just two nuclei cannot alter the combined energy, which is proportional to the longitudinal magnetization. Therefore this type of nuclear interaction does not contribute T_1 -relaxation – that requires more drastic nuclear interactions that involve an exchange of energy with the surroundings. All processes that result in T_1 -relaxation also result in T_2 -relaxation, ensuring that T_1 is never shorter than T_2 .

Generally T_2 becomes shorter with increasing firmness of the matter, but this does not apply to T_1 , which is long in very firm matter and very fluid matter (e.g., several seconds), but is short for semi-firm matter.

The shortest T_1 is achieved precisely when random interactions with the neighbors fulfill the resonance condition, meaning, for example, that at 1 tesla, T_1 is shortest for matter where a proton meets about 42 million other nuclei in 1 second (the Larmor frequency is 42 MHz). This is understandable – even random pushes to a swing can result in considerable oscillations if the frequency is approximately correct. If we instead push way too frequently or rarely, very little will be accomplished regardless of whether the pushes are random or not.

5.3 Inhomogeneity as a source of signal loss, T_2^*

Interactions between nuclei in the ever changing environment of molecules is the cause of radio signal loss on a timescale T_2 and longitudinal magnetization recovery on a longer timescale T_1 . However, a loss of transversal magnetization is also observed because of inhomogeneity in the field, meaning variation in B_0 . As expressed through the Larmor-equation, the nuclei precess at a frequency depending on the magnetic field. If this varies over the investigated area, the nuclei

will after some time, point in all directions transversally, even if they had some degree of alignment after excitation. This process is called *dephasing*. Since the measured signal is proportional to the transversal net magnetization, inhomogeneity gives rise to a loss of signal. The larger the field inhomogeneity, the faster the dephasing. How quickly this happens is denoted by the time constant T_2^* (pronounced “ T_2 -star”). The degree of inhomogeneity depends partly on the scanner's ability to deliver a uniform field (called a good *shim*). T_1 and T_2 are tissue-dependent parameters for which normal values are published. This does not apply to T_2^* values, since they depend, for example, on the size of the voxel since the inhomogeneity increases with this.

Loss of signal due to interactions of nuclei is irreversible, but through a sly trick, signal lost due to inhomogeneity can be recovered. The inverse process of dephasing is called refocusing, and it involves gradual re-alignment of nuclei in the transversal plane. Refocusing is triggered by repeated use of radio waves as described below. The recovered signal is known as an “echo”.

6 Sequences

A “sequence” is an MR measurement method and there exists an extraordinary number of these. A measurement is characterized by the sequence of events that play out during the measurements. In a classic “spin-echo-sequence” as shown in figure 9, the magnetization is first rotated 90 degrees away from equilibrium through the use of radio waves. A short waiting period follows with dephasing caused by inhomogeneity. Subsequently the magnetization is turned an additional 180 degrees with radio waves. After a little more waiting time and refocusing, the signal is measured. This course is usually repeated. In the sequence description above, a couple of sequence parameters were involved, e.g. the waiting time from excitation to measurement. A measurement method can thus be described by a sequence name and the associated sequence parameters. Together they control the measurement's sensitivity to different tissue parameters, and the choice is therefore made on the background of the actual clinical situation. Sequences include elements such as

excitation Turning of the magnetization away from equilibrium.

dephasing Field inhomogeneity causes the nuclei to precess at different speeds, so that the alignment of the nuclei – and thus the signal – is lost.

refocusing pulse After excitation and dephasing, part of the signal that has been lost due to inhomogeneity can be recovered. This is achieved by sending a 180-degree radio wave pulse (in this context known as a refocusing pulse) that turns the magnetization 180 degrees around the radio wave field. The effect is that those spins which precess most rapidly during dephasing are put the furthest back in evolution and vice versa (mirroring around the radio wave field causes the sign of the phase angle to be reversed). In the following “refocusing period” where the nuclei are still experiencing the same inhomogeneity, they will therefore gradually align again (come into phase). The lost signal is recovered, so that an “echo” is measured.

readout Measurement of MR signal from the body.

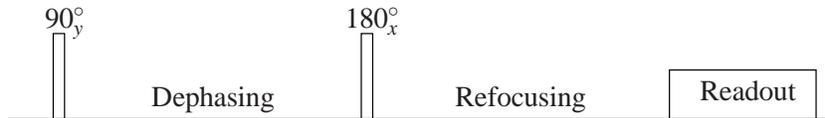


Figure 9: Sequence diagrams illustrate with an implicit time axis from left to right, the events during a measurement. The first blocks in this spin-echo sequence illustrate radio wave pulses, meaning periods wherein the radio wave sender is turned on. The diagram is read as follows: Initially radio waves are sent for the duration needed to turn the magnetization by 90 degrees around the y -axis in the rotating frame of reference. After a short waiting period wherein the spins will dephase, radio waves are again sent for a period appropriate for turning the magnetization 180 degrees around the x -axis. This gives rise to gradual refocusing. The signal is read out (measured) when this is complete, which happens after a brief waiting period equal to the duration of the dephasing. Sequence diagrams most often contain several events. An example with gradients is found in figure 17

waiting times Periods wherein the relaxation contributes to the desired weighting.

Further sequence elements (building blocks) that can be mentioned include inversion, bipolar gradients, spoiling and imaging gradients. Some of these are described later in the main text or in the glossary.

7 Signal-to-noise and contrast-to-noise ratios

In section 5.1, it was hinted that PD-weighted images are rarely the best at reflecting pathology. This is in spite of the fact that the signal-to-noise ratio is often good in such images, because imaging with fully relaxed magnetization and short echo-time increases the signal strength. This exemplifies that the contrast-to-noise ratio *always* is more important for diagnostics than the signal-to-noise ratio, meaning that a good image is characterized by the signal difference between conditions or tissues that we wish to distinguish (the contrast) being large compared to the relevant noise (that is, the signal variations that make the distinction difficult, whether these are random or not). If the contrast-to-noise ratio is small, on the other hand, the relevant signal differences drown in noise.

An image can be quite noisy, meaning that it has a low signal-to-noise ratio, but still be good at reflecting a given pathology. One example is diffusion-weighted images in connection with ischemia, since diffusion weighting lowers the signal for all types of tissue, but improves contrast between normal and ischemic tissue. Something similar is often true for images acquired with long echo-time.

In rough terms, the signal-to-noise ratio is an expression for how nice-looking an image is, while the contrast-to-noise ratio is a measure of the usefulness *insofar* as the contrast is defined meaningfully from all the parameters that can help answering the clinically relevant questions. It is thus not trivial to define the contrast for a given clinical situation (a metric), but it is *in principle* simple to optimize it. In practice the optimization can easily be the most difficult part, however, as it may require many measurements to be performed with different sets of parameters.

8 Quantum Mechanics and the MR phenomenon

This section is non-essential for understanding the rest. Here, some common misconceptions found in many introducing textbooks (even good ones) are pointed out. These misconceptions have been recognized by a minority for a long time, but they are sadly still appearing in many new books, as authors often quote each other without a deep understanding of MR phenomena. The problem derives mainly from the fact that few authors of MR textbooks have a good background in quantum mechanics, which is what the misconceptions center around. Such a background is needed, if authors present explanations based on quantum mechanics, but luckily such explanations are not needed.

The fundamental problem is that quantum mechanics is often presented as necessary in order to understand MR, where a classical description is fully adequate. If the two are combined incorrectly, one easily ends up with something that is contrary to both common sense and experimental evidence.

First of all, a little background: In physics, there exists a distinction between classical phenomena and quantum phenomena. The first are those that can be described by classical physics, such as Newton's laws. The vast majority of the experiences we gain from everyday life belong in the classical category. In certain situations, however, the classical description is inadequate, so that predictions derived from it turn out to be inconsistent with observations. In such cases, one needs quantum mechanics that is believed to be able to describe all of the currently observed physical phenomena, but which is difficult to understand, since we make very few unclassical observations in our everyday life. One example of a quantum phenomenon is superconductivity (resistance free electrical conductance), that we use when creating powerful magnetic fields in the scanner. Superconductivity would not be possible if classical physics were accurate on an atomic level.

Quantum mechanics, on the other hand, describes superconductivity, meaning that if one takes the quantum mechanical equivalents to Newton's fundamental laws as a basis, it can be shown that a broad range of surprising phenomena are predictable, including superconductivity. All phenomena that can be described using classical mechanics can also be described with quantum mechanics, although the complexity grows tremendously, since explanations require familiarity with the quantum mechanical terminology rather than simply experience from everyday life.

Magnetic resonance is not a quantum phenomenon, although it is often presented this way. It *can* be described with quantum mechanics, but the approach necessitates considerable background knowledge. MR can alternatively be described classically, since it can be shown that the classical description of magnetic dipoles behavior in the magnetic field is a direct consequence of quantum mechanics.

The existence of the proton spin and magnetic moment is, on the contrary, a quantum mechanical phenomenon, and the classical figure 3 should therefore not be taken too literally (it has its limitations, although these are not evident in MR experiments). If one accepts the existence of the magnetic needle quality of nuclei, the rest follows from classical physics.

For a number of introductions to MR, the classical approach is not chosen. On the other hand, the authors typically do not possess the necessary background in order to use the heavy (and in this case superfluous) quantum mechanical formalism in a meaningful way. Instead, we are often presented with a semi-classical description with elements of classical physics and elements of quantum physics.

Such a description *can* be precise and useful, but easily gets to include elements of nonsense, which is the case for much literature on the subject. Examples are presented below.

Although quantum mechanics is not needed to understand the magnetic resonance phenomenon or pretty much anything else going on in scanners, physicists and chemists can still benefit from learning about it. Quantum mechanics is, for example, necessary in order to determine the *size* of T_1 and T_2 correctly from theoretical assumptions alone (without actual measurements), while the *existence* of T_1 and T_2 is purely classical (they can be calculated classically, but the results will be wrong). In the same vein, quantum mechanics is practical in connection with certain types of spectroscopy.

There are consequently good reasons to use quantum mechanics in descriptions aimed at chemists and physicists with the necessary quantum background, while an incorrect pseudo-quantum introduction is never helpful. Finally, quantum mechanics is necessary to describe measurements carried out on single protons. Such measurements are never performed, however, as the signal always comes from a large collection of protons, a so-called statistical ensemble. For such an ensemble of non-interacting protons, it is a relatively easy to show, quantum mechanically, that the classical description is exactly right.

Quantum mechanics is most often employed in the following contexts in introductions to magnetic resonance:

- Quantum mechanics is very convenient for “explaining” the resonance phenomenon. This is often done by drawing two lines, calling them energy levels and say that quantum mechanics tell us that the photon energy must match the energy difference between the energy levels. A problem with this approach is that it is non-intuitive and raises more questions than it answers, unless the reader is already schooled in quantum mechanics (why are almost half of the nuclei pointing *opposite* the field? Why doesn’t continuous use of radio waves level out the population differences? What happens if the photon-energy isn’t quite precise? Alternatively, it can be intuitively argued (as attempted above) that the radio waves must push the magnetization in synchrony with the precession in order to rotate the magnetization away from equilibrium.
- Quantum mechanics can conveniently be used to calculate the size of the equilibrium magnetization. In equilibrium, the coherence terms average to zero and the partition function consequently is reduced to a sum of just two terms. Classically, more states contribute, and the calculation is therefore more complicated. The results of the two calculations are equal, however.

8.1 Corrections

After this tirade, time has come to comment on some of the myths and incorrect figures often found in literature. It is more important to recognize and ignore the erroneous figures than it is to understand why they are wrong. The details can be found in the article *Is Quantum Mechanics necessary for understanding Magnetic Resonance?* Concepts in Magnetic Resonance Part A, 32A (5), 2008, that contains detailed accounts.

MR is a quantum mechanical phenomenon No, the resonance phenomenon is classical. See above. The existence of nuclear spin and aspects of certain experiments are, however, quantum mechanical.

Protons can only be in two conditions: Near-parallel or anti-parallel with the field.

This situation, as illustrated in figure 10, does not occur. There are two conditions that are characterized quantum mechanically by having a well-defined energy. As expected, however, the spins can point in all directions both classically and quantum mechanically.

The radio waves align the spins on the two cones The illustration is the result of the erroneous belief that the protons can only be in one of two conditions, combined with the fact that we can rotate the magnetization. The result is pure hocus-pocus: Why are the individual spins not rotated the same way by the radio waves? How do the protons “know” where they are each supposed to be pointing? According to this view, it appears possible to change the magnetization with radio waves. Why don't we see this happening in practise? The questions are numerous, simply because the drawing is wrong. It is relatively simple to show (classically as well as quantum mechanically) that the locally homogeneous fields that magnet and RF coils generate, can never change the relative orientation of spins locally. Only inhomogeneous fields such as those the nuclei affect each other with, can do that. They are the source of relaxation.

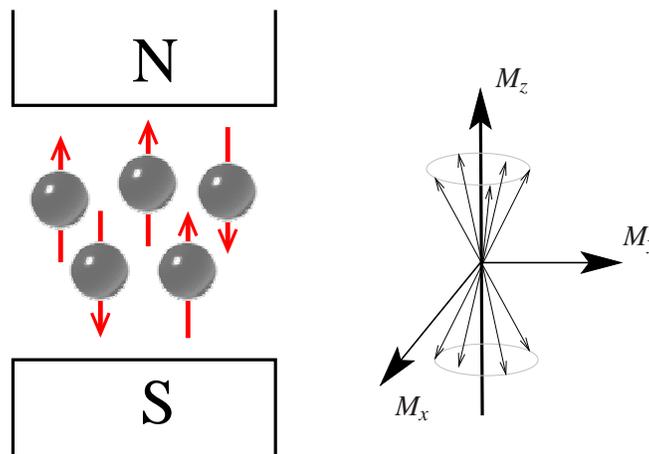


Figure 10: These figures are seen often, but they are misleading. They reflect a factually wrong perception of spins as only being capable of being in spin-up or spin-down conditions.

9 Imaging

So far, it has been discussed how the body can be brought to emit radio waves, but it has not been mentioned how radio waves from one position in the body can be distinguished from radio waves

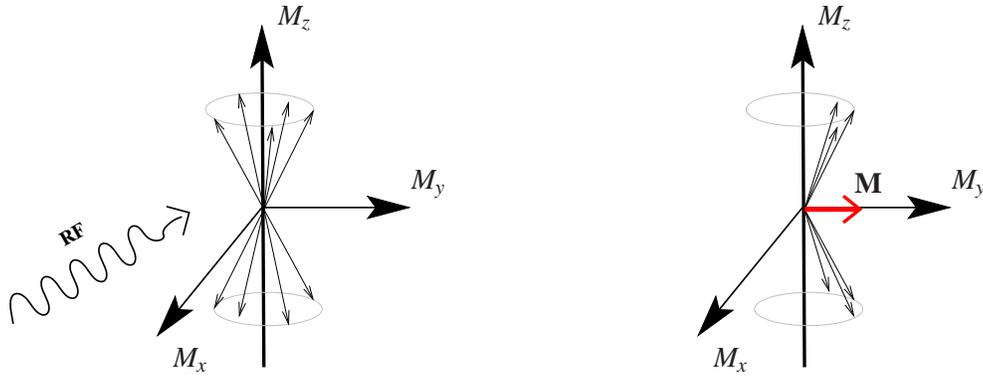


Figure 11: The spins are not aligned on two cones with a net magnetization in the transversal plane, as this figure claims. Homogeneous radio wave fields *cannot* change the relative orientation of spins.

emitted at another position, which is rather necessary for imaging. In this section, techniques for spatial localization and imaging are described. Initially, it is explained why these techniques are fundamentally different from more easily understood techniques known from everyday life.

9.1 Background

The most obvious methods for MR imaging could be imagined to be projection or the usage of antennas that can detect where in the body the radio waves are emitted. X-ray and normal microscopy are examples of such “optical” techniques, and it would appear obvious to extend this type of imaging to MR. Optical techniques are, however, “wavelength-limited”, which means that they cannot be used to acquire images more detailed than approximately one wavelength.

In other words: Due to fundamental causes, one cannot localize the source of radio waves more precisely than about one wavelength when using lenses or other direction-sensitive antennas. The radio waves used in MR scanning are typically several meters long, so with optical techniques we can hardly determine whether the patient is in or outside the scanner (this argument is really only valid in the far field, which is the background for parallel imaging. More on this later).

Optical techniques as we know them from binoculars, eyesight, CT, X-ray, ultrasound and microscopes, are thus practically useless for MR-imaging, and a fundamentally different principle is necessary. This principle was introduced by Paul Lauterbur in 1973, and it resulted in the Nobel Prize in Medicine in 2003. Basically, Lauterbur made the protons give their own locations away by making the frequency of the emitted radio waves reflect the position. Lauterbur shared the prize with Sir Peter Mansfield, who also contributed greatly to the development of techniques used in magnetic resonance imaging (MRI).

9.2 Principles

A requirement for MR imaging is that the scanner is equipped with extra electromagnets called “gradient coils” that cause linear field variations. Direction and strength can be altered as desired.

The spatial localization occurs according to different principles, of which the most simple is slice selection. Other forms of coding involve the so-called k -space, which will be introduced in a later section.

9.3 Slice selection

By using gradient coils, the magnetic field strength can be controlled so that it, for example, increases from left to right ear, while the direction is the same everywhere (along the body). This is called a field gradient from left to right. By making the field inhomogeneous in this way, the resonance frequency varies in the direction of the field gradient. If we then push the protons with radio waves at a certain frequency, the resonance condition will be fulfilled in a plane perpendicular to the gradient as shown in figure 12. The spins in the plane have thus been rotated significantly, while spins in other positions simply vibrate slightly. Thus we have achieved slice selective excitation of the protons and a sagittal slice has been chosen.

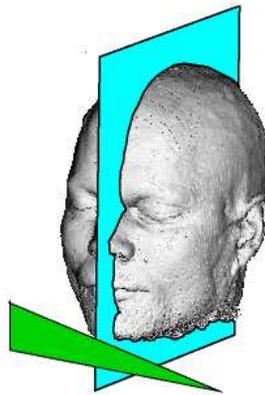


Figure 12: Spin is influenced selectively in a sagittal slice, if a gradient from left to right is applied while radio waves are transmitted.

9.4 Spatial localization within a slice

After the protons in a slice are excited, they will all emit radio waves. In order to create images of the slice, we must introduce a way to distinguish the signals from different positions within the slice. The fundamental principle can appear a bit foreign, but will be explained in detail at a later stage. Briefly told, different patterns in the magnetization throughout the slice are created with the help of gradients. The strength of the radio signals that are returned tell us how much the object in the scanner “looks like” the applied pattern. By combining patterns according to their measured similarity to the object, the well-know MR-images are created.

What is meant by “patterns” is first illustrated in one dimension, meaning that we consider spins placed on a line (e.g. between the ears) and watch their location and direction immediately after excitation.



As illustrated, immediately after excitation the spins all point in the same direction perpendicular to the magnetic field, which points out of the paper. They will thereafter precess around the magnetic field, that is, they will rotate in the plane of the paper at a frequency that is dependent on the magnetic field. Insofar as the field is made to increase from left to right by applying a field gradient briefly, the spins will each turn an angle that depends linearly on the nucleus’ position:



This so-called “phase roll” is an example of the above mentioned spin patterns, that can be created in the patient by use of gradients. The word “phase” expresses the direction that the spins point in. It is seen that the neighboring spins point in nearly the same direction, but throughout the object, the magnetization has rotated several turns. The longer time a gradient is turned on, and the greater the field variation that arises, the more “phase roll” is accumulated (more rotations per unit length).

We have through use of gradients, made the spins point in all directions in a controlled fashion and have thus simultaneously lost the signal. This is seen by comparing the two situations above, since the measured magnetization is the sum of all the contributions from the individual spins. When the spins are in phase (that is, pointing in the same direction), they jointly create a considerable magnetization that gives rise to radio waves being emitted. When the spins point in all directions as when a gradient has been applied, their sum is quite small. As a result, comparably weak radio waves are emitted.

The gain from using the gradient can thus appear quite small: We have simply lost the signal. That does not, however, have to be the case. Look, for example, at the situation illustrated below where there are not (as above) protons uniformly distributed from left to right, but where there is a regular variation in water content instead.

Before gradient:



After gradient:



In the figure we see how the spins point before and after applying the same gradient as above. The only difference between the two earlier figures are “holes” that indicate a lack of protons (locally water-free environment – bone, for example).

We now consider the total magnetization and thus the signal in this new situation. As long as the spins all point in the same direction, i.e. before the gradient is applied, we receive less signal than before, as fewer spins are contributing to the magnetization. The signal before the gradient is applied is thus a measure of the total water content.

After the gradient is applied, on the other hand, we receive *more* signal than in the comparable situation above (homogeneous water distribution). This may appear odd, since there are *fewer* protons that contribute to the magnetization. The remaining protons all largely point in the same direction, however, with the result that their combined magnetization is relatively large.

That the object with structure (varying water content) emits radio waves even after the gradient has been applied, is caused by the *match* of the phase roll and the structure in the object, understood in the sense that the “wave length” of the phase roll is the same as the distance between the water puddles in the example.

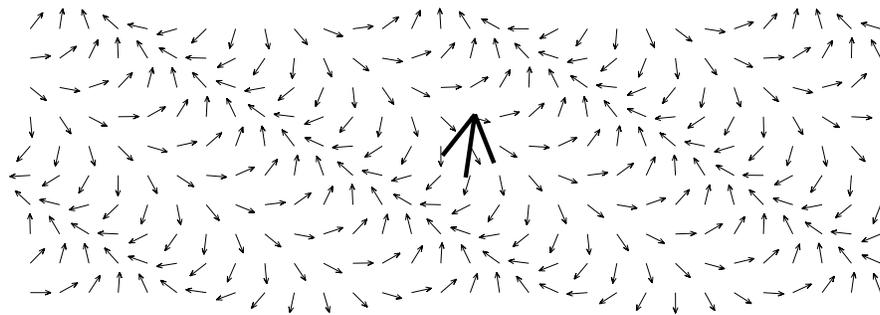
If this is not the case, there will be tendency for the different contributions to the magnetization to cancel each other. The signal is thus a measure of “similarity” between the phase roll pattern and the structure in the object: After a phase roll has been introduced in the object by applying a gradient, the signal we receive is a measure of whether there is structure in the object matching the phase roll.

Thus we now have a basis for understanding MR-imaging: Different phase roll patterns in the body are drawn one after the other. The resulting radio wave signals are recorded for each of the patterns. The size of the signal for a particular phase roll pattern tells us whether there is similarity to the structure of the body.

In few cases there is an obvious similarity between structure and phase roll pattern. Most often, the similarity is very limited. Generally, the images are obtained by combining patterns weighted by the appropriate measured signals. This principle is actually known from the game MasterMind, which was popular long ago (figure 13).

9.5 Extension to more dimension – k -space

So far, we have considered a situation with hydrogen nuclei placed as pearls on a string. We will now extend to two dimension and consider the situation in a slice of tissue with uniform water content after a gradient has been applied:



The figure show a phase roll in the gradient direction (upwards and right). The “thick arrow” on the middle of the drawing shows the net magnetization, that is, the sum of all the small magneti-

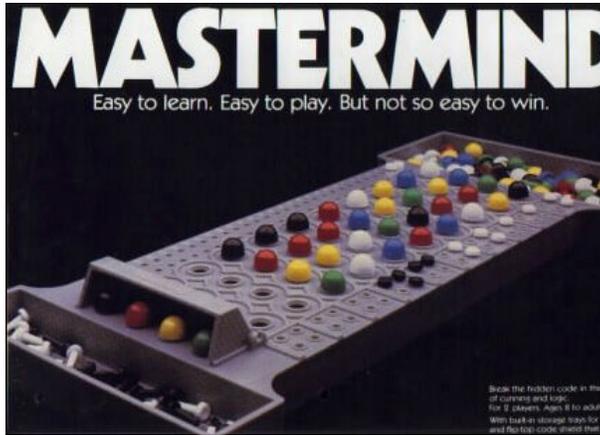
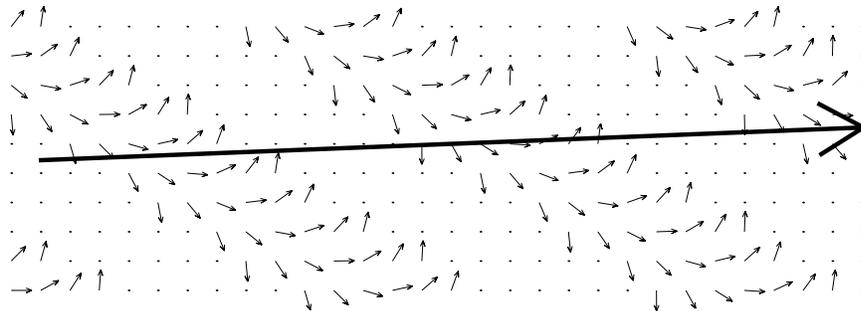


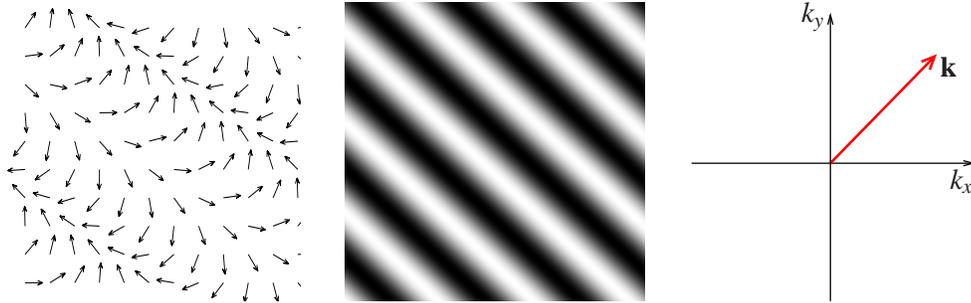
Figure 13: The game Mastermind. One player places a row of colored pieces hidden behind a screen at one end of the board. The opposing player must then guess the combination of pieces. After each guess (pieces placed in the opposite end) the first player informs to which degree the guess was correct, that is, how many pieces were placed correctly (though not which). The guessing player will after a few rounds be able to predict the hidden combination based on responses to earlier guesses. The same principle is in action in MR imaging: The inner parts of the patient are hidden, but through gradients we may “draw” spin patterns in the patient. In return, we receive radio waves whose strength reveals the similarity between patient structure and the drawn pattern.

zation vectors (a vector is an arrow). It is seen that the sum is relatively small, since there is an approximately equally many arrows pointing in all directions.

As in the one-dimensional case, there may still be a considerable net magnetization even after a gradient has been applied. This happens if there, as shown in the following figure, is structure in the object that matches the phase roll (the dots mark places where there are no protons).



For the sake of clarity, the spin patterns in the following will be shown in two different ways. Rather than showing arrows pointing in different directions, phase rolls will be shown as an intensity variation as in the figure in the middle, where the color reflects the direction of the arrows²:



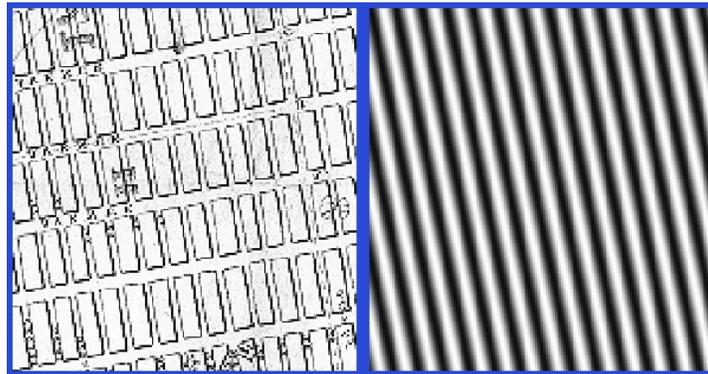
For each such pattern, a vector \mathbf{k} is assigned, whose direction indicates the direction of change and whose length indicates the density of stripes as shown to the right.

By applying gradients, we “draw” patterns in the patient, one by one, in a predetermined order specific for the chosen sequence. All directions of stripes and all densities of stripes are drawn up until a certain limit. As said, the returned radio wave signal depends on the similarity between object and pattern and it is registered as a function of \mathbf{k} , as shown in figure 14. Example patterns are shown, but for each single point in k -space there is a corresponding pattern.

9.6 Similarity and image reconstruction

Above, it has been described how the similarity between object and phase roll is measured, and it has been said that the image is calculated by adding patterns together that are weighted by the measured MR-signal. This section elaborates on what this actually means.

The concept of “similarity” is illustrated in this case with an example.



To the left is shown part of a city map (Manhattan). To the right is shown the unique stripe pattern that looks most like the map in an MR sense. The similarity may seem small, but it is evident when

²The figure is not brilliant: the left situation corresponds to a stripe pattern with only 2 stripes rather than the 6 shown in the middle. The direction, however, is spot on in all three illustrations.

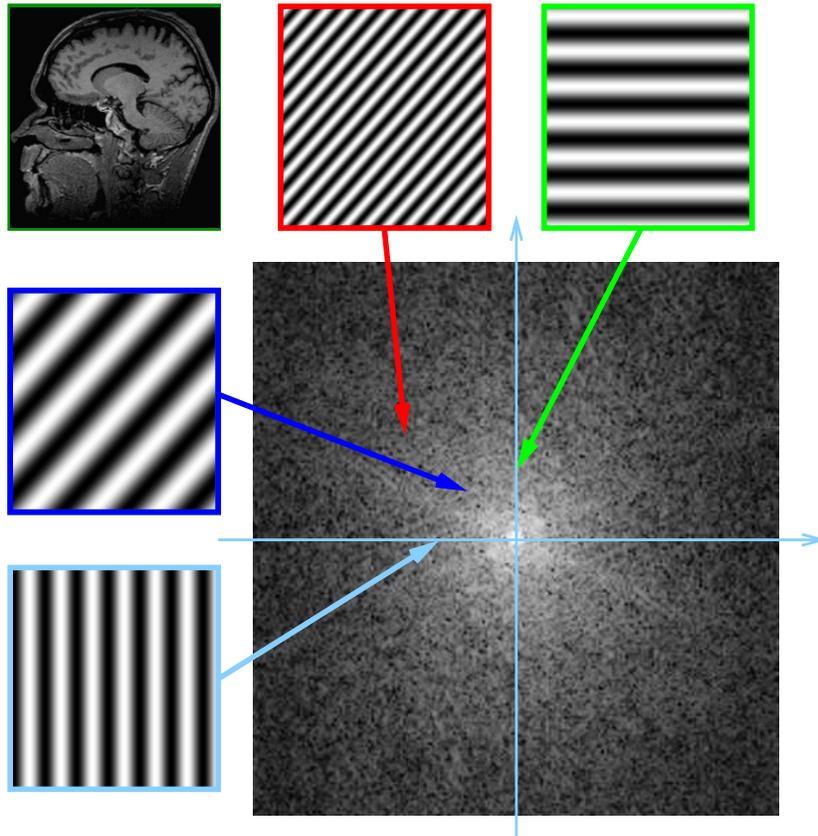


Figure 14: The structure of k -space. The signal is measured as a function of the \mathbf{k} -vector (k_x, k_y) , which changes when field gradients are applied. The lighter region in the middle of k -space tells us that the object in the scanner (a head) has greater similarity to slowly varying patterns, or expressed differently, the most powerful radio waves are emitted from the body when low-density phase roll patterns are created.

it is realized that the stripes are going in the same directions as the north-south roads, that they have about the same distance as these and that the pattern is shifted, so that it is bright where the map is bright.

We find almost equally high similarity with a similar east-west oriented pattern. If the two patterns are put on top of each other (stacked as two overheads), one would almost have a usable road map. This is how image reconstruction is done, with the exception that it usually takes a lot of patterns to create a useful image. This is caused by the fact that the similarity between object and each individual pattern is typically extremely small, but not zero.

It may appear surprising that very complex pictures can be created by simply adding different patterns. There is also an important difference from the overhead example, which would provide darker and darker pictures the more patterns are added: The difference lies in the fact that one must also be able to subtract patterns in order to produce usable images, meaning that the sign of the MR-signal (the phase) is important. Figure 15 shows how patterns can be added to form brain images.

9.7 Moving in k -space

The phase roll is controlled by applying gradients with different directions, strengths and durations. As long as a gradient is active, and the field thus varies linearly with position, the phase roll will change, but it will *always* have the shape of a stripe pattern as shown earlier. Meanwhile, there will continuously be emitted radio waves from the body which tell us how well the pattern matches the structure of the object.

Example: If a gradient going from ear to ear (the x -direction) is applied, the result is a phase roll in this direction corresponding to the \mathbf{k} -vector growing along the k_x -axis. As long as the gradient is “on”, the stripes in the pattern are getting closer as the phase rolls accumulate (the wavelength, defined as the distance between stripes, becomes shorter). If another gradient in the direction from scalp to foot is applied, a phase roll in this direction will accumulate and the \mathbf{k} -vector grows along the k_y -axis.

If both gradients are applied, a tilted phase roll is achieved. This is true regardless of whether the gradients are applied one at a time or simultaneously: By applying them one at a time, the \mathbf{k} -vector moves along the k_x -axis first and then along k_y . By applying them simultaneously, it moves directly to the same point and the final result is the same.

One example of the path through k -space is shown for the spin-echo sequence in figure 16. The gradient, which is itself described by a vector, determines at all times where the measurement is headed in k -space and is thus the speed of the \mathbf{k} -vectors movement in k -space.

That the emitted signal depends entirely of the position in k -space and not on the path there, is important for the success of the k -space description of MR-imaging. Many sequences differ by the path through k -space (the order in which patterns are created), but image reconstruction is fundamentally identical.

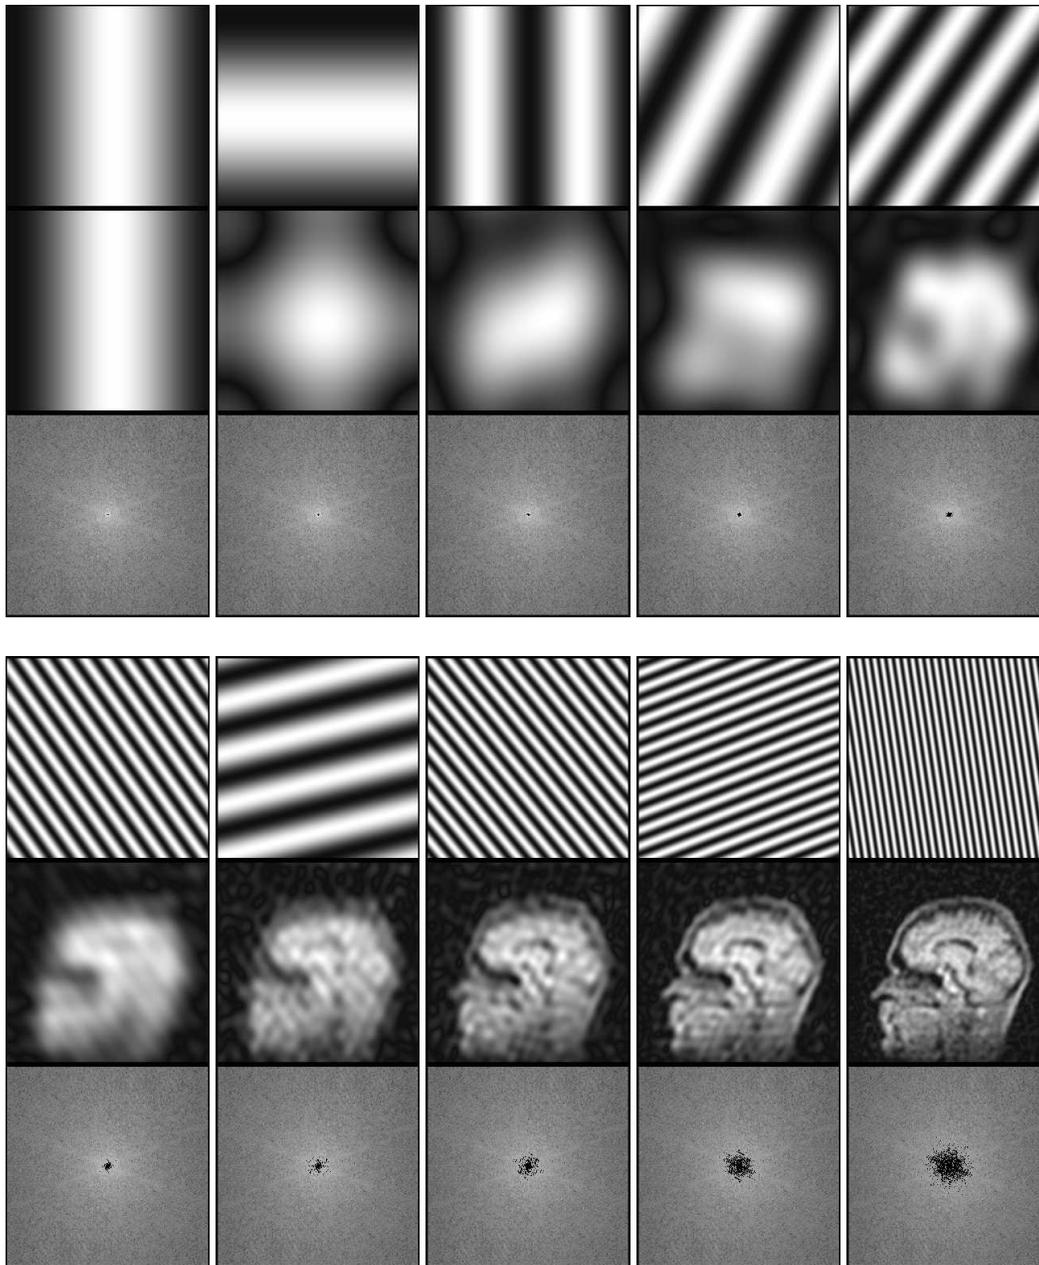


Figure 15: Image reconstruction. The figure shows how simple patterns (line 1) can be summed up to complex images (line 2). The reconstructed images are, in this case, created from the patterns that have the greatest similarity with the subject in the scanner, meaning the areas of k -space where the strongest radio waves were measured (line 3). Still more patterns are included in the image reconstruction (indicated by the dark areas on the k -space images and the reconstructions are likewise becoming more and more detailed (the number of included patterns is doubled in each step). The final reconstructed image (lower right) has been created on the background of the thousand patterns appearing most similar to the excited slice. The top line shows the “most recently added” pattern. It is seen that the slow signal variations (intensities) are measured in the middle of k -space, while the sharpness of the edges are derived from the measurements further out in the k -space.

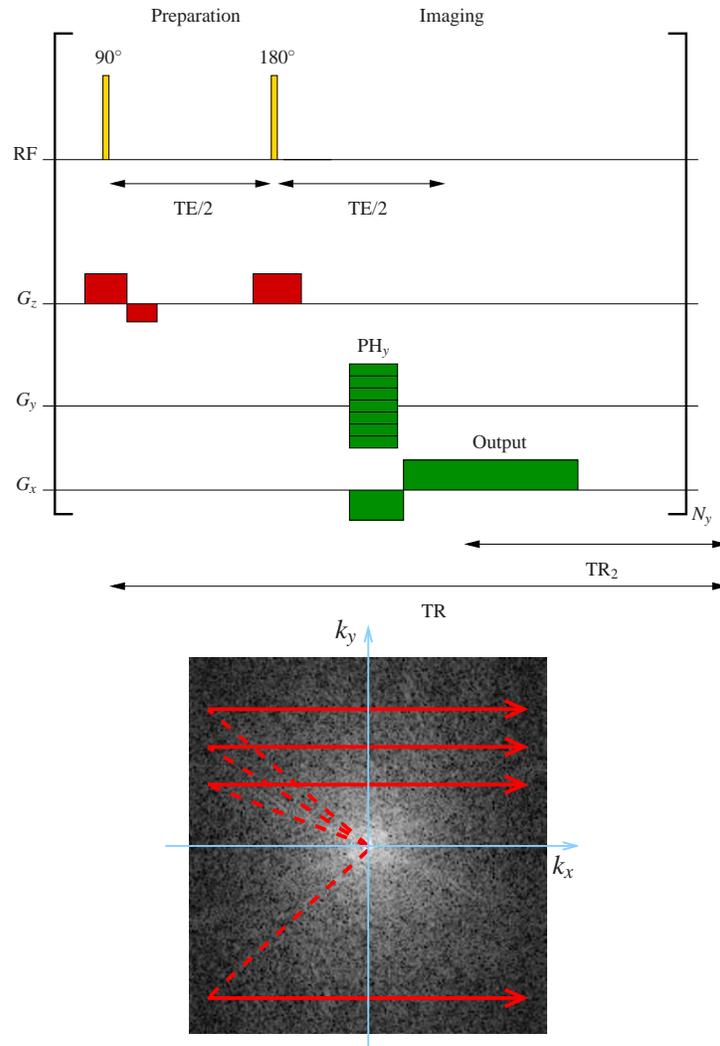


Figure 16: Here, the spin-echo sequence illustrates how gradients can be used to maneuver in k -space. As in figure 9, the top line indicates the use of radio wave pulses, while the bottom three show gradients in perpendicular directions. Initially, the magnetization is rotated 90 degrees into the transversal plane. Midway between excitation and output, an RF refocusing pulse is applied so that the signal at time TE after excitation is unaffected by field inhomogeneity, if such is present. A gradient along the z -direction is applied simultaneously with the two radio wave pulses, so that these become slice-selective as shown in figure 12. The gradient direction is inverted shortly after the first RF pulse in order to remove the phase roll that is accumulated in the period immediately following excitation. There is consequently no phase roll immediately prior to application of the imaging gradients in the x - and y -directions: The local net magnetization point in the same direction everywhere and we are at this point located in the middle of k -space. The field is now made inhomogeneous in both the x - and y -directions (positive G_y , negative G_x) so that a phase roll is accumulated and (k_x, k_y) is likewise changed toward the upper left corner of k -space. When the y -gradient is turned off and the x -gradient inverted, readout of a line in k -space is begun. The process is repeated after a period TR with a new value of the y -gradient so that a new line of k -space is read out. This is indicated in the sequence diagram with a gradient table (PH_y) and an outer bracket that indicates repetition for each line in k -space.

9.8 Image acquisition and echo-time

Earlier in the text, the echo time was described as the duration from excitation until the point where radio waves are measured. Since the transversal magnetization in this period is lost on a timescale of T_2 (or T_2^* , if a refocusing pulse is not used), the echo time TE is thus determining the corresponding relaxation time weighting of the measurement.

After the introduction of imaging, we now need to reconsider the definition of echo time, since several positions in k -space are typically being measured after each excitation. A typical approach to imaging is, for example, to measure individual points along a line in k -space one by one after each excitation.

The single points (corresponding to different stripe patterns) are consequently not measured at the same time after excitation, and the echo-time definition has therefore become blurry. For echo-planar imaging (EPI), this problem is extreme, since the entire image is measured after a single excitation, meaning that some points in k -space are measured milliseconds after excitation, while others are measured, for example, 100 ms later and thus with a completely different T_2 weighting.

Is the possibility to characterize T_2 -weighting by a single parameter therefore lost? No! It turns out that the echo-time definition can be adjusted, so that it still can be interpreted as above. A surprising characteristic comes into play here: Even though parts of k -space are acquired shortly after excitation and other parts a long time after, the reconstructed image looks (contrast-wise) as if it has been acquired at a very certain time after excitation, that being the time where the middle of k -space has been acquired. As such, it makes good sense to define the echo-time as the duration from excitation until the time where the middle of k -space is measured (see an example in figure 16).

9.9 Frequency and phase encoding

Introductions to MR imaging methodology often takes their outset in “frequency encoding” that is readout of signal in presence of a gradient so that the frequency of the radio waves reflect the position. It is intuitively easy to understand how a frequency analysis of the measured signal reveals how the body water is distributed along the direction of the gradient and how frequency encoding thereby offers simple 1D imaging. Even though frequency encoding is easier to understand, a direct k -space approach was chosen above. The reason is partly that frequency encoding leaves a need to explain “phase encoding” which is a technique necessary to obtain information about the water distribution in the direction perpendicular to the readout and slice selection gradient (figure 16). Phase encoding is most easily understood from phase roll and similarity considerations (nuclei are *not* given a unique phase along the phase encoding direction though attempts to explain the technique that way are often made – look at figure 14 for confirmation). Another reason to choose the direct k -space approach is that the frequency and phase encoding approaches are inadequate for understanding sequences such as spiral EPI and to some extent also normal “blipped” EPI (even the necessity of the pre-winding gradient for normal frequency encoding appears non-obvious). In contrast, the k -space approach gives a far more general understanding of MR imaging (see glossary). A related advantage is that frequency and phase encoding in k -space appears as simple variations on a general theme. Even sequences such as spiral EPI appear conceptually similar to other sequences.

It is important to know of frequency and phase encoding for several reasons, however. Firstly, the corresponding directions are often referred to in articles and in scanner software as the scan duration, for example, often depends on the number of phase encoding directions and steps. Secondly, the artifacts in the phase and frequency encoding directions are very different. Examples are given in the following sections, and it will be advantageous to have frequency encoding in mind when reading them. Whether readout gradients were used or not, it is often fruitful to consider in terms of frequency encoding how a particular artifact would appear, had k -space along a particular direction been filled using a readout gradient respecting the timing of the actual filling.

9.10 Spatial distortions and related artifacts

As described above in section 9.8, MR images appear contrast-wise to have been captured at a certain time after excitation (the echo time) although measurements are far from always performed that way. The cause of this is that the central parts of k -space primarily contains information about the image contrast whereas the outer regions contain information about the finer structural details (edges). That the acquisition of the images occurs during a longer period is therefore not reflected in the contrast that is determined in the relatively short time it takes to pass the middle of k -space. The duration is instead primarily reflected in spatial distortions (displacements) and other oddities.

The so-called “chemical shift” artifact and the EPI distortion artifact are both such consequences of different parts of k -space being acquired at different times after excitation. The artifacts appear as relative displacements of nuclei in the image plane, happening if these precess at different frequencies, even in absence of imaging gradients. It is a fundamental assumption in imaging that all frequency differences are caused by linear imaging gradients, and violations cause artifacts.

For the distortion artifacts, it is differences between phase rather than intensity that cause the displacements. Intensity differences also have an effect, however, but it is hard to see at first glance: Because the outer parts of k -space contain information about finer structural details, and because variation in signal over k -space is T_2 -dependent, the spatial resolution of a measurement is dependent on T_2 .

Among other things, this has the surprising effect that a blood vessel under a bolus-passage of contrast agent will not only change intensity on EPI images, but will also appear to be “smeared” over a larger area of the image. Because of the phase differences that induced field changes give rise to, the vessel will even appear displaced under the bolus passage.

Displacements in the image plane of two substances that oscillate at different frequencies (fat and water, for example) is proportional to the frequency difference. In order to calculate the displacement in pixels, the frequency difference must be divided by the “bandwidth per pixel” that can often be examined or even controlled from the scanner software. In accordance with the above, the bandwidth is inversely proportional to the time it takes to traverse k -space in a given direction. The displacement in a given direction is thus simply the product of the frequency difference and the time it takes to pass k -space in this direction. The clock is “reset” at each new excitation.

9.11 Slice-selective (2D-) versus 3D-sequences

A distinction is made between multi-slice and 3D-sequences. This is confusing since the result is a 3D-dataset in both cases (most often represented as a series of slices). The somewhat unfortunate choice of words does not refer to the result but rather to the method of acquisition: For slice-selective sequences, protons in one slice at a time are affected as described above (section 9.3), while in 3D sequences all spins in the coil are affected identically and simultaneously. Subsequently the k -space methods described in chapter 9.4 are used to distinguish signals from different positions (3D k -space is similar to 2D k -space). A frequently used 3D sequence is MPRAGE (Magnetization Prepared Rapid Acquisition Gradient Echo sequence). This is shown in figure 17 and is described in appendix A.

3D sequences are almost always applied with a short repetition time TR, since the measurements will otherwise become very time consuming. Therefore, they are practically always T_1 -weighted, but can furthermore contain significant T_2 -weighting, as is the case for the CISS sequence, for example, that has mixed contrast. Sometimes, the T_1 weighting is isolated in the edges of k -space, so that the image appears T_2 weighted. Because of the short TR, there is much less room for contrast variation compared to multi-slice sequences.

High spatial resolution in all directions simultaneously can be achieved for 3D measurements, e.g. isotropic (1 mm)³ resolution (=1 microliter). For regular 2D sequences, the case is different since slice thicknesses under 3 mm are rarely employed. The resolution in the image plane is often set high for 2D sequences, which results in good-looking images but quite oblong voxels. This may cause problems when distinguishing small or flat structures such as grey matter. In summary:

3D sequences Allows for high spatial resolution and isotropic voxels, but most often only T_1 weighting and mixed weighting turn out well for short measurement time.

Multi slice sequences High “in plane” spatial resolution, but rarely under 3 mm slice-thickness. Flexible contrast.

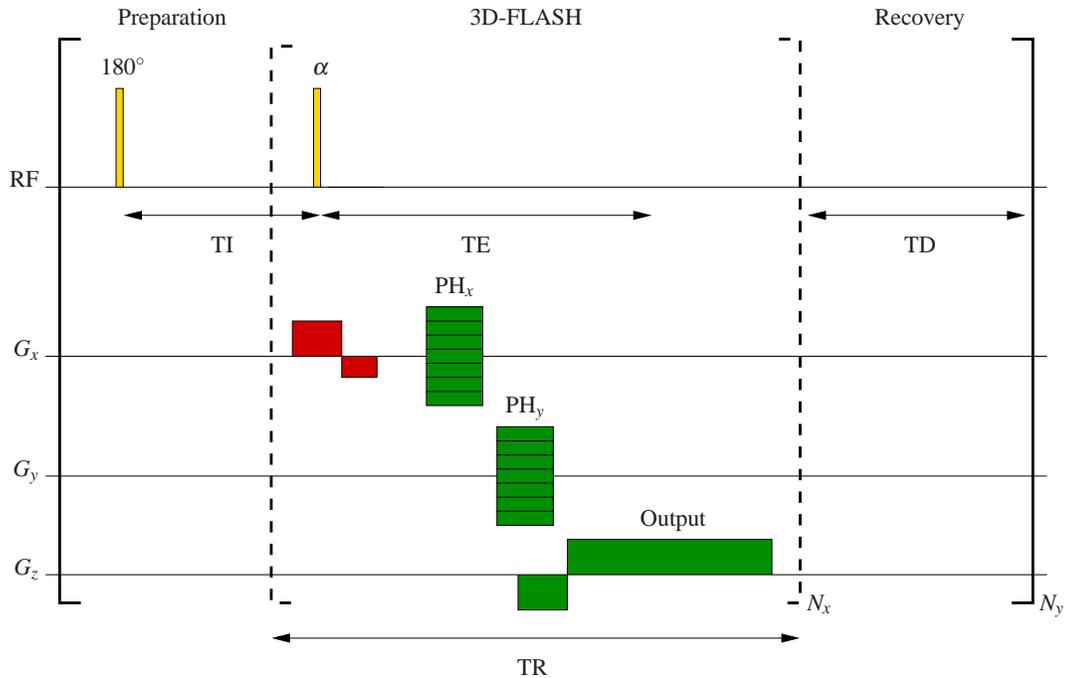


Figure 17: The MPRAGE sequence is an example of a 3D sequence where excitation is non-selective or weakly selective, and where the 3D k -space must be traversed in order to achieve spatial resolution in all directions. Initially, the magnetization is rotated by radio waves so that it points opposite the magnetic field in the following period TI (the “inversion time”). This is done to increase the T_1 -weighting of the magnetization and thus also of the image. In order to detect the magnetization, it is tilted slightly into the transversal plane by an RF pulse with a small flip-angle α . The signal is coded spatially by the green gradients and it is measured in the readout period. The inner, stippled bracket indicates that the enclosed part of the sequence is repeated once for each k_x value. By applying small flip-angles, the longitudinal magnetization is only affected slightly each time, so that many lines along the k_z axis can be measured before the magnetization must once again be inverted. This happens after a recovery period TD. The entire sequence is repeated for each k_y value, so that all of k -space is covered. The gradient applied simultaneously with the α pulse is not strictly necessary, but can be used to choose a thick slice (slab) within which phase encoding can be carried out. Thus, aliasing may be avoided, even if only relatively few phase encoding steps are employed. Further information regarding MPRAGE and FLASH can be found in the appendix.

9.12 Aliasing and parallel imaging

That k -space imaging is a Fourier technique (meaning that it is based on the addition of phase rolls) is reflected in the artifacts that may appear in images, if the underlying assumptions are not met. If we do not have knowledge about the subject in the scanner *a priori*, the similarity with all possible phase roll patterns must in principle be measured in order to reconstruct a perfect MR image. It does, however, take time to create patterns, and in practise only a central, rectangular part of k -space is covered with finite resolution, as illustrated in figure 15. Usually, it is the number of measured lines that determines the duration of a measurement. A limited density in k -space of these is therefore desired.

If the chosen density is too low, it results in a characteristic “aliasing artifact”: If sampling in k -space is not dense enough to catch signal variations (i.e., if the Nyquist-criteria is not met), errors will appear in the reconstructed images. The errors consist of signals being misplaced and overlapping, as seen to the right.

In 1997, however, Sodickson and Manning brought attention to the fact that if several antennas with different sensitivity profiles are measuring the same k -space, the “lacking” k -space samples can be calculated, thus avoiding aliasing. Simultaneously applying two antennas for signal reception can in principle therefore accelerate the measurement by a factor of two, although in practice it requires several more antennas to get reliable images. Pruessmann *et al* soon proved that sets of reconstructed, aliased images acquired with different antennas can intuitively be “un-aliased”, which is principally equivalent. The techniques are today known under the common name “parallel imaging”.

That parallel imaging is an independent technique, even though it was first proposed as a modification of the k -space technique, is most easily understood through a thought experiment: In the extreme limit “each-pixel-has-its-own-antenna” (e.g. many tiny antennas spread equidistantly through the imaged object), the gradients are not necessary to measure detailed images. It is simultaneously realized that parallel imaging is not a Fourier technique, meaning that reconstruction is not based on the summation of stripe patterns.

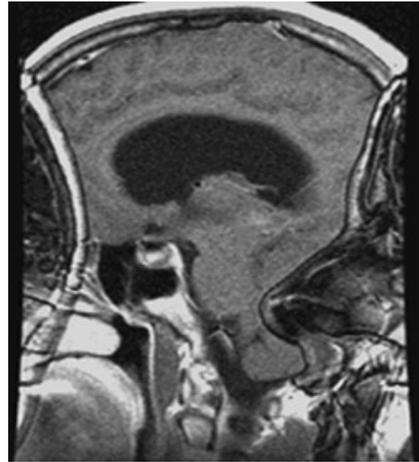


Figure 18: Aliasing: The image displays a brain in profile, but since the sampling density in k -space is too low, the nose ends up in the neck and vice versa. Thus, pathology can be hidden.

9.13 Finishing remarks on the subject of imaging

Typically, relaxation times are relatively long which is responsible for the enormous flexibility in MR scanning. In 100 ms, for example, series of RF and gradient pulses can be applied to prepare the magnetization in a way that allows its size or direction to reflect physiologically interesting parameters, in other words manipulating the contrast in the images. Additionally, long T_2 's allows for the generation of a large number of phase roll patterns and for measurement of the corresponding

radio wave signals before the system's memory of the preparation has been erased. Actually, the whole k -space can be covered after a single excitation (echo-planar imaging, EPI).

It is interesting to note how the wavelength limitation known from optical techniques was avoided. In the case of k -space techniques, this happened by using frequency differences rather than far field properties, for spatial source localization of radio waves. The traditional wavelength limitation is thus replaced with a requirement that the neighboring nuclei's frequencies must be distinguishable during the measurement period, if the positions are to be distinguished. This can be expressed alternatively: No structural details considerably less than the shortest wavelength of the phase roll can be observed. Thus, a new wavelength limitation is introduced, which is intuitively understandable from figure 15. Since it takes time to move in k -space, the ultimate limitations for the spatial resolution are determined by gradient strengths, relaxation times and movement of the water molecules, including diffusion.

In parallel imaging, the limited and different sensitivities of the coil elements are used to avoid the wavelength limitation. Since the sensitivity area's dimensions are more or less equal to those of the coil elements, the spatial resolution is determined by the coil elements' size for "pure" parallel imaging. This resolution is usually inadequate for practical purposes.

Parallel imaging is, however, completely compatible with k -space imaging, in that the two techniques can be combined and images acquired faster than normally. In some situations, speed is more important than spatial resolution, and therefore experimentation is now done with images acquired without applying gradients. Such images are acquired with sub-millisecond time resolution of up to hundreds of antennas simultaneously. These images are of interest for imaging of substances with short relaxation times, for example.

10 Noise

The limiting factor for many MR examinations is noise. We can, for example, not directly detect substances *in vivo* in concentrations below a few millimolars on a reasonable timescale, because the signal is drowned by noise.

The noise can be physiological (pulse, respiration, movement), but even if the patient is lying completely still, there exists an upper limit for the image quality achievable in a given period. In the absence of physiological noise and under a couple of other assumptions (which are rarely completely fulfilled), the signal to noise ratio is proportional to the voxel size and the square root of the time it has taken to acquire the image.

It is essential to realize that the signal-to-noise ratio is not dependent on the number of voxels in the image. If you double the matrix (points along one side of the image) and the field-of-view it will typically imply that the signal-to-noise ratio is increased by a factor of $\sqrt{2}$ because the measuring time is thereby typically doubled while the voxel size remains unaltered. That it may only be extra air that is included by the expanded field-of-view is irrelevant, and the noise level is not affected by the fact that 4 times as many voxels are being measured.

Insofar as the scanner is well functioning, the electronics is not the primary source of noise. Instead, that is the random motion of charged particles (ions) in the patient. When charged particles

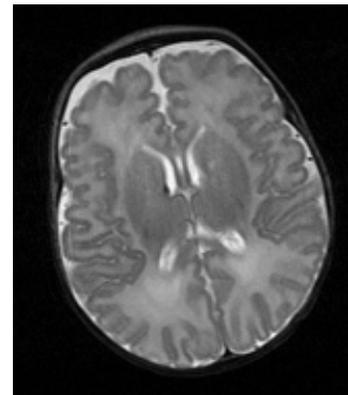
diffuse, they emit random radio waves as they change their direction of motion. The higher the temperature and conductivity of a material, the more noise it emits.

Thermal noise is evenly distributed despite the fact that it is emitted almost only from within the patient, who may only fill part of the image area. This is caused by the noise not being an MR signal, and the gradient-induced spatial coding of the signal is therefore not affecting the noise. Instead, the noise is received in a steady stream during the entire measurement and it is therefore evenly spread over k -space and consequently also evenly over the MR image.

The noise from patients cannot be avoided, but we can, to some extent, avoid measuring it. The idea is to use a small coil that only detects noise from a small area of the patient (a surface coil, for example). It is a common (but non-essential) misconception that using a small coil primarily enhances the signal-to-noise-ratio through its improved sensitivity to signal. Improved signal sensitivity, however, also increases sensitivity to noise that is generated in the same regions of the body. Improving the sensitivity therefore does not by itself improve the signal-to-noise ratio. Instead, the surface coil limits the noise by being sensitive to a smaller part of the body, of size similar to that of the surface coil. A small surface coil only detects the noise (and signal) from a small part of the body, and this noise appears evenly distributed over the entire image. The surface coil therefore improves the signal-to-noise ratio by reducing the sensitivity to distant noise sources. These effects are well described by Redpath in *The British Journal of Radiology*, **71**:704-7, 1998, for example.

11 Scanning at high field

Three tesla is the largest, clinically available field strength at present. The high field is especially beneficial for imaging with a high spatial resolution, as well as for performing measures that normally suffer under a low signal-to-noise-ratio, including perfusion measurements and spectroscopy. The image to the right depicts a newborn child and demonstrates high resolution and fine contrast after only 3 minutes acquisition at 3T. A small periventricular bleeding is visible in the right hemisphere (left on the image, according to radiological convention). The relaxation times are changed at high field (typically longer T_1 and shorter T_2). Sequences and parameters must therefore be adjusted to the field strength in order to obtain good images.



For functional imaging, fMRI, the high field strength provides enhanced contrast upon brain activation, so it can be measured in which areas activation occurs in a given situation (see fMRI in the vocabulary). The methods can, for example, be used clinically for planning of surgery preserving essential functions such as language. There are, however, several pitfalls doing fMRI and the quality of the results must thus be evaluated carefully before conclusions can be drawn.

12 MR Safety

Is the MR technique harmless, and just how certain are we of this? It is a natural question, since we do use radio waves that have been suspected of being carcinogenic in connection with high voltage wiring and mobile telephones. Furthermore, magnetic fields of quite significant strength are employed, from 25.000 to 75.000 times the average field of the earth. By nature, we cannot know with total certainty whether MR is completely harmless. At this time, no harmful effects have been documented, although the question has been explored in depth.

It is known that possible harmful effects are rare or have very limited effect. This knowledge is partially derived from animal experiments with long term exposure. Population studies have been conducted of radiographers and others working close to magnets or radio wave fields (chemists and physicists have worked in high magnetic fields since the 40's and many people working with radar have been exposed to monstrous amounts of radio waves).

The area is regulated by international standards and the scanners are equipped with monitoring equipment that ensures that no radio waves more powerful than allowed are emitted, even by accident.

The most significant risks associated with MR scanning are probably resulting from misinterpretations of the images, as well as purely practical issues concerning handling. Reported injuries include hearing damage, burns as a result of wire loops (in jewelry, for example), as well as serious incidents with metal objects such as the flying chair shown to the right. There are scary examples of oxygen canisters, floor waxers and other things going astray, found in anecdotes at http://www.simplyphysics.com/flying_objects.html



Even though the scanning itself is harmless, it has been shown that MR contrast agents injected intravenously in connection with certain types of scanning, can be dangerous to patients with kidney problems. Because of differences in the contents and chemical stability of the contrast agents used, the risk varies. In the case of patients with serious kidney problems, the use of gadolinium based contrast agents should be avoided when possible, in accordance with the guidelines found on the websites of the European and International MR Societies.

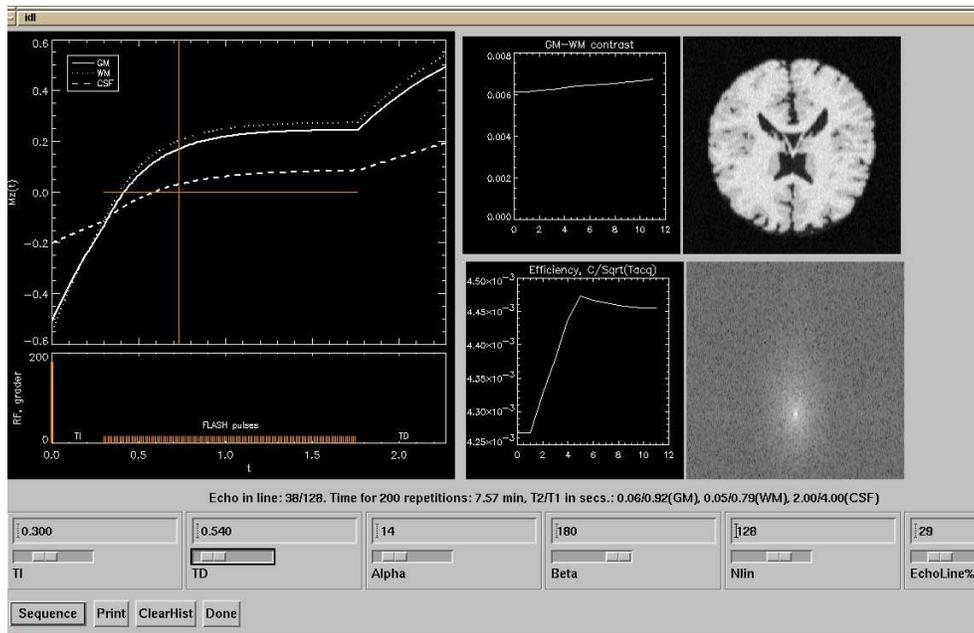
<http://www.esmrm.org/> and <http://www.ismrm.org/>

Usually, the contrast agents are excreted from the body through the kidneys shortly after scanning, but in the case of kidney patients, it can be trapped in the body. It must be emphasized that the problem only occurs in very special cases (exclusively when combining a high dose of contrast agent, a heavily impaired kidney function and the use of certain substances). At Hvidovre Hospitals MR department, there are no suspected cases, although contrast have been used for 20 years for thousands of patients. The same is true of most other departments with MR scanners.

Appendices

Appendix A: The educational software mprage

This appendix is the online help for the software “mprage”. It also describes the design and important qualities of this often used 3D-sequence. The program is unfortunately not in a form where it can be widely spread, but it can be started in Hvidovre from the linux command line by writing `mprage`. It can also be distributed on demand.



Introduction

This is the online manual for a simulation program for the sequences MPRAGE and FLASH. The program was developed in connection with a research project (inhomogeneity correction) and has since been adapted for sequence optimization and educational use.

The MPRAGE sequence

The MPRAGE sequence as shown in figure 17, is an often used T1-weighted 3D sequence that includes two repetitions (loops): An inner loop repeating an excitation pulse with a small flip angle, and an outer loop that begins with a 180 degree inversion pulse.

The sequence is a 3D FLASH with phase encoding in two directions, meaning that the magnetization is rotated away from the direction of the magnetic field by the use of many RF pulses with small flip angles in quick succession (typically over a hundred with about 10ms interval). In contrast to the spin echo sequence

where the magnetization is used up completely at each excitation, a typical FLASH flip angle is small, e.g. 10 degrees.

Each pulse only affects the magnetization slightly, but when many pulses are employed, the longitudinal magnetization approach a level that depends on T_1 and the strength and frequency of the RF pulses.

A FLASH sequence usually results in T_1 weighted images, since the degree of saturation is T_1 dependent. Said in another way: The many RF pulses eventually drive the magnetization to a new equilibrium level (“driven equilibrium”), that reflects how quickly the magnetization is rebuilt. For the purpose of providing extra T_1 -weighting, the MPRAGE sequence has a built-in inversion pulse, employed before each FLASH module.

After each excitation, the fresh transversal magnetization is affected by gradients that creates a phase roll corresponding to a new position in k -space. A readout gradient causes movement through k -space along a straight line while signal is acquired for a few milliseconds. When the inner and outer loops are finished, the entire 3D k -space is covered and an image of the whole volume can be reconstructed. Each line in k -space is weighted differently since the preceding FLASH pulses influence the magnetization.

T_2 does not play a significant role concerning the contrast in FLASH and MPRAGE, since the echo time typically is a few milliseconds and thus much shorter than T_2 .

Using the program

The parameters from MPRAGE can be varied in the program, and the equivalent sequence and evolution of the longitudinal magnetization, M_z , can be followed (left graphs). The parameters occurring are

TI The inversion time. The period from the inversion pulse until the first FLASH pulse.

Beta The flip angle of the inversion pulse (180° for MPRAGE, 0° for FLASH, 90° for saturation-recovery-FLASH).

Alpha The FLASH flip-angle.

TD Pause from last FLASH pulse until the next inversion pulse.

Nlin The number of lines that are acquired in one of the k -space phase encoding directions. The other directions (phase encoding and readout) are not so interesting in this context since they do not influence contrast or voxel shape. The uninteresting phase encoding direction is here chosen perpendicular to the image plane.

EchoLin% A percent rate that indicates when in the FLASH period the middle of k -space is passed. The contrast is primarily determined here. As an example, a choice of 50% indicates that the center of k -space is measured after half of the FLASH pulses are played out. Often, the center of k -space is measured asymmetrically to increase the effect of the inversion pulse.

The naming of the parameters are in accordance with older Siemens scanners. On more recent scanners, the parameters are defined and named differently. Ask.

On the graph seen to the left, a single iteration of the outer loop is illustrated. The evolution of the magnetization is shown at the top, and the RF pulse timings are shown at the bottom. The horizontal red line in the top graph shows the timing of the FLASH module and the vertical line illustrates when the middle of k -space is reached.

The outer loop is typically repeated many times (e.g. 200) and after a few iterations, the same curve is achieved every time (the one shown). The magnetizations shown in the right and left side of the graph are therefore connected in the sense that only the sign differs between the first and last point for beta equal to 180 degrees, for example (an inversion pulse inverts the magnetization, hence M_z changes sign).

Along with the contrast evolution, simulated k -space images (raw data) and reconstructed images are shown. The raw data are generated by adding k -space images for the single tissue types, but weighted line-wise with the relevant M_z during the FLASH period. The reconstructed images are produced by Fourier-transforming the k -space images.

When the parameters are changed, the corresponding changes in contrast and efficiency are shown. The latter is a contrast measure corrected for sequence duration (achieving a better contrast is easy, if the duration of the measuring is increased). The contrast is calculated as the signal difference between grey and white matter (GM, WM) at the time when the middle of k -space is passed. This contrast measure is reasonable, if the voxel shape is not bizarre for any tissue type, i.e., if the k -space weighting is reasonably uniform. Efficiency is defined as the contrast per $\sqrt{\text{duration}}$ (so that it is independent of the number of repeated measures (signal averages)).

Important points

The program illustrates a number of important aspects:

- There are many parameters to adjust in a sequence like MPRAGE. The contrast and image quality depends strongly on these.
- It is not possible to provide easy rules for how the contrast changes when a given parameter is adjusted – it depends on the other parameters. Simulation tools like this one are thus required. The software also contains optimization tools. Ask, if this interests you.
- Contrast and spatial resolution (matrix) are not independent.
- During the pauses, M_z approaches the equilibrium magnetization (proportional to proton density) on a timescale of T_1 .
- During the FLASH period, when pulsing is frequent, the magnetization approaches another level dictated by T_1 , pulse interval and strength. The greater the flip angle alpha, the faster “driven equilibrium” is achieved, and the less the influence of the inversion pulse.
- The signal from one type of tissue can be “nulled” in the middle of k -space through appropriate parameter choices. The signal from the tissue type in question is thereby suppressed. The tissue still contributes with signal in other regions in k -space, however, and bright edges may therefore result.
- The voxel shape and spatial resolution depend on the type of tissue, since the k -space weighting (meaning the signal variation over k -space) depends on the type of tissue. This is most obvious when a null-crossing appears close to the k -space center (often producing ugly images).
- Although the GM/WM contrast, as defined above, can be good for a given set of parameters, the image may be rubbish. This occurs if the k -space weighting produces an unfortunate voxel shape for one or more types of tissue (see above). Thus it is difficult to automatically optimize the sequence in a meaningful way (tradeoff between voxel shape and contrast).
- The importance of the inversion pulse for the contrast, is dependent on the duration and number of FLASH pulses that occur between inversion and passage of the k -space center.

Precautions

The simulation is not more precise than the parameters entering the simulation, including

- the assumed relaxation times (shown in the window - these depend on the field strength),
- the assumed proton densities,

- sequence parameters, including FLASH-TR, which at the moment can not be varied without altering the program slightly,
- the assumed noise level,
- the choice of phase encoding directions,
- the lack of influence from RF-inhomogeneity that is not simulated. The program has “hidden” graphics, that illustrate this sensitivity (press Print). Ask, if your interest has been aroused.

Appendix B: Glossary

This glossary can be used for repetition and for learning more about individual topics.

The BOLD effect This effect is the basis for most fMRI. The rate at which the MR signal dies out is weakly dependent on the concentration of oxygen in the blood, since oxygenation changes the magnetic qualities of the blood. Deoxyhemoglobin is a paramagnetic contrast agent contrary to oxygenated hemoglobin that is diamagnetic like the tissue.

The signal intensity is usually increased by neuronal activation, since the blood supply (perfusion) increases more than the oxygen consumption. This is the so-called BOLD-effect (blood oxygenation level dependent effect) that results in signal variation of a few percent in gradient-echo sequences. The greatest sensitivity occurs at $TE \simeq T_2^*$ (consider why this is).

Coil This word is used for several concepts in MR. It is typically used for the antennas that send and receive radio waves. The coil is usually fitted to the body part that is being scanned, so it is near the area of investigation. A head coil surrounds the entire head, for example, and not much else. The choice of coil involves a trade-off of sensitivity, homogeneity and coverage.

In contrast, gradient coils, which are electromagnets, are not used for transmitting and receiving radiowaves, but for creating linear field gradients (see Gradients). Shim coils are similarly used to eliminate non-linear, background field gradients.

Contrast agents Substances that affect the contrast in images. Contrary to X-ray contrast agents, for example, the MR contrast agents are usually not seen directly. Rather, we see their effect on the protons (increased relaxation – see gadolinium). Contrast agents are, for example, used for detecting defects in the blood-brain barrier: Gd-DTPA usually stays in the blood stream, but is temporarily trapped in the tissue in the case of blood-brain barrier breakdown. Magnetically tagged blood can also be used as a contrast agent (spin labelling).

Diamagnetism Magnetic fields deform the electronic orbits so that the magnetic field is weakened slightly. This phenomenon is called diamagnetism and it is exhibited by all substances. However, there are substances where the diamagnetic quality is completely overshadowed by far more powerful influences (paramagnetic and ferromagnetic substances). The so-called magnetic susceptibility expresses how much the field is strengthened or weakened by a material.

Diffusion Diffusion is measured by employing a powerful gradient that introduces a short-wave phase roll over the object. When the gradient is then inverted (or equivalently, if a refocusing pulse is applied), the spins will be rotated back in phase, forming an echo *insofar as the spins have not moved meanwhile*. If the spins have diffused, meaning that they moved randomly among each other, the signal will only be partially restored. The signal intensity therefore reflects the mobility of the protons (the diffusion coefficient).

Echo-planar imaging, EPI A rapid imaging method where k -space is traversed using oscillating gradients after a single excitation. The acquisition time for a single EPI image is typically

below 100 ms. EPI often exhibits low signal-to-noise ratio and pronounced artifacts, including spatial distortions and signal voids near air-filled spaces. It is used frequently, however, when time resolution is crucial.

Eddy current See induction.

Electrons Negatively charged particles that orbit the positive atomic nuclei. The electron cloud is responsible for the chemical properties of atoms and molecules. Its appearance is weakly reflected in the MR signal, since the electron cloud partly screens off the magnetic field. This is exploited in spectroscopy. Electrons also have spin, and ESR is the term used for Electron Spin Resonance. Unfortunately, the short relaxation times of electrons make ESR useless for medical imaging.

Excitation The magnetization cannot be measured before it is rotated and has acquired a component perpendicular to the magnetic field (a transversal component). The rotation is caused by transmitted resonant radio waves, a process called excitation.

Exponential function When “something” decreases exponentially (the transversal magnetization, for example), it means that a specific fraction is lost when we wait for a certain period of time: When we say that T_2 is 100 ms it means that about 60% of the signal is lost in 100 ms. The same fraction of the remaining signal is lost again if we wait another 100 ms. The T_1 relaxation too involves exponential functions, since the longitudinal magnetization approaches equilibrium exponentially after excitation (the difference decreases exponentially).

Ferromagnetism Diamagnetism provides a minimal weakening of the field, and paramagnetism typically provides a more powerful increase of the field. Such changes are in the range of per mille or less. There are, however, a few agents that have far more powerful magnetic attributes, among them iron, nickel and cobalt that are ferromagnetic. These materials may not be introduced in the scanner room without special precautions. Be aware that only metallic iron is ferromagnetic – ionic iron is paramagnetic.

Flow Velocity and direction of uniform movement can be measured with a technique similar to that described above for measuring diffusion. Yet, the employed gradients are typically weaker and the measurements are interpreted differently, since it is the net rotation rather than the signal loss that reflects the speed. After employing the bipolar gradients, the spins have achieved a total phase change that is proportional to the velocity along the gradient. The technique is therefore called phase contrast MRI and the velocity is determined by comparing phase images.

fMRI Functional MRI. Mapping of the areas of the brain that are activated in a given situation. Usually based on the BOLD-effect.

Frequency encoding If signal is read out while a gradient coil is active, the frequency content will directly reflect the distribution of nuclei along the gradient direction. A one-dimensional image (a projection) can consequently be acquired by doing a frequency analysis of the

measured signal. Since the gradient is active during signal readout, its direction is called the readout direction. Even though frequency encoding is easy to understand, it has proven fruitful to think instead of how the phase roll that the gradient is creating, is causing the signal to reflect sample structure. Frequency encoding is consequently a special case of k -space imaging. See also “phase encoding”.

Gadolinium Gadolinium (Gd) is one of the less well-known elements, a heavy metal. It is characterized by having unpaired electrons in inner shells, which results in the atom having magnetic qualities that are far stronger than that of protons. The field near a gadolinium atom is strongly inhomogeneous. Protons that pass near gadolinium therefore receive a powerful magnetic influence which gives rise to increased relaxation. Free gadolinium is toxic and it is therefore placed in a larger organic structure – DTPA is typically used. Contrary to X-ray contrast agents, among others, it is not the actual agent that is seen with MR, but rather the protons that come near gadolinium (or those that do not, for T_2 -weighted images). Thus, there exists long range effects of gadolinium (e.g. over the blood-brain barrier, which water can pass but Gd-DTPA cannot).

Gradients Gradients are variations in the magnetic field. The scanner is equipped with gradient coils that give linear field variations whose direction and strength can be altered as needed. This is essential for slice selection, imaging and more. Field alterations result in mechanical forces on the scanner, vibrations, and thus acoustic noise.

Non-linear gradients also appear due to imperfect shimming. On a microscopic scale, the field is inhomogeneous, e.g., near paramagnetic particles such as gadolinium and deoxyhemoglobin. For this reason, T_2^* is always shorter than T_2 .

Gradient echo A gradient creates a phase roll, i.e. a position-dependent rotation of the spins. If the gradient is subsequently inverted, the phase roll will gradually rewind until the spins are back in phase. Thus the magnetization is recovered (in fact, it has only been hidden). Signal recovered in this manner is called a gradient-echo, and it is employed for imaging and for flow- and diffusion measurements. The related spin-echo removes the effect of inhomogeneities caused by variation in the magnetic properties of the tissue. This does not apply to the gradient echo, since we can only actively invert the gradient direction for the gradients we control, i.e. the ones that are created with electromagnets.

Induction If the magnetic field through a wire loop is changed, a current in the wire will be induced (an eddy current). This may, for example, happen to electrodes that form a noose in the scanner, which may result in burns. Radio waves are typically the most significant source of heating. Eddy currents also appear in the scanner metal. Magnetic fields formed in this way can cause distortions of the MR-signals.

Inversion Often, the equilibrium magnetization is inverted for a period by use of a 180-degree pulse before it is flipped into the transversal plane to be measured. The inversion time, TI, is this period with T_1 -weighting, i.e., the time from inversion until the excitation pulse. Inversion is

typically introduced to add T_1 -weighting (inversion recovery) or to get rid of signal derived from fat or CSF (STIR, FLAIR).

Magnetic dipole Spin makes nuclei act as small compass needles. More precisely, they act as magnetic dipoles, meaning that they have a south pole and a north pole. Furthermore the angular momentum (rotation) makes them precess in the magnetic fields, meaning that the north/south axis rotate around the direction of the field.

Metabolites Substances that play a part in the metabolism. The most easily detectable metabolites *in vivo* are NAA, choline, creatine and lipids. Others are lactate, inositols, GABA, NAAG, taurine, glutamine, glutamate and alanine. Some of these are painstakingly difficult to measure reliably. There are many other agents in the body, of course, but only small, mobile molecules in high concentrations are detectable in a clinical setting.

MR Magnetic Resonance.

MRI Magnetic Resonance Imaging.

NMR “Nuclear Magnetic Resonance”. The same as MR. The reference to the nuclei is said to be removed to avoid associations with radioactivity, nuclear power and atomic weapons. The chemists have maintained the N and they use the technique for identifying agents and mapping protein structures by spectroscopy, for example.

Paramagnetism All substances have diamagnetic properties, but some are also paramagnetic, which makes them alter the local field significantly more than the diamagnetic contributions do. Gadolinium and deoxygenated blood are examples. In pure paramagnetic agents, the field is typically increased in the range of a few per mille. See Diamagnetism above.

Perfusion Contrary to bulk flow where a large amount of blood is flowing in the same direction, the blood flow in the smallest vessels (capillaries) is far from uniform within a voxel (there are millions of capillaries oriented randomly within a voxel of tissue). The blood flow in these is called perfusion and it is responsible for oxygen supply to the tissue. It is typically measured by following the passage of a contrast agent through the vasculature on a series of images. The subsequent analysis is demanding, especially for quantitative measurements.

Phase The word “phase” is used for the direction of the magnetization perpendicular to the B_0 field (meaning the direction in the transversal plane). This is affected by the magnetic field that causes precession. Immediately after excitation, the individual contributions to the net magnetization point along the same direction perpendicular to the field – the spins are said to be in phase (or having the same phase). Subsequently, a gradual dephasing and signal loss follows. This reflects that the common phase is lost, resulting in a loss of net magnetization. This loss can be caused by random interactions between nuclei (in which case the loss is irreversible), or it can be caused by field gradients (and is then typically reversible, since the signal can be recovered after using a refocusing pulse).

Phase images Usual MR images show the size (amplitude) of the magnetization in the transversal plane measured at the echo time, but they do not tell us anything about the direction (the phase). This must, however, be known to measure flow as the velocity is often reflected in the phase. For most sequences, one can therefore choose to acquire phase images in addition to regular amplitude images. They are most frequently zebra-striped, which indicates the presence of a phase roll: The phase angle moves from 0 to 360 degree, jumps to 0 and increases steadily again. This is equivalent to the spins being rotated one or more turns.

Phase encoding Both slice selection and frequency encoding are relatively easy to understand without considering phase rolls, but these techniques are inadequate to achieve spatial resolution in the third dimension. Phase encoding is often the solution: The sequence is repeated with different strengths of a gradient inserted immediately after the excitation. The direction of this is called the phase encoding direction. Phase encoding is most easily explained in term of *k*-space imaging: Different phase rolls are induced along the gradient direction, and the resulting signal strength reflects the degree to which the patient structure “matches” the phase roll. This signal is subsequently frequency encoded and measured. Even simple imaging becomes relatively complicated when phase and frequency encoding are combined, but described separately. They both appear on equal footing, however, in the *k*-space description of imaging, which has proven much more viable and general than phase encoding and frequency encoding separately. There are still good reasons, however, to get acquainted with frequency and phase encoding as described in the main text.

Proton Nuclear particle. All atomic nuclei are built of protons and neutrons. The hydrogen nucleus is the smallest and consists of only one proton.

Precession Spin, which is the rotation of the nuclei around their own axis, gives rise to the nuclear magnetic property as shown in figure 3. This rotation should not be confused with “precession”, which is the rotation of the magnetization in a magnetic field as shown in figure 4), i.e., the rotation of the nuclear north-south axis.

Quantum Mechanics A fascinating subject that you should be wise not to ask me about, unless you have plenty of time on your hands. Irrelevant for understanding most aspects of MR, including nearly everything that scanners are used for. Introductions to MRI often contain misunderstood QM.

Radio waves, RF Radio waves in the relevant frequency range are often simply called 'RF' (for radio frequency). The employed frequencies overlap with the ones used in radio communication. MR must therefore be carried out in an RF shielded room (also known as an RF cabin or a Faraday cage). The shielding can be tackled a number of ways. In animal experimental scanners, the “plugs” in the ends of the scanner act as boundaries for the shielded volume.

Refocusing pulse A pulse of radio waves transmitted with the intention of creating a spin echo, is often called a refocusing pulse. The pulse typically has a flip angle of 180 degrees.

Relaxation times T_1 and T_2 are time constants that describe how quickly the magnetization is approaching equilibrium. The magnetization is longitudinal in equilibrium, meaning that it is pointed along the magnetic field ($M_{xy} = 0$, $M_z = M_0$). Away from equilibrium, the transversal magnetization M_{xy} decreases on a timescale T_2 , while the longitudinal magnetization M_z approaches M_0 on a timescale T_1 .

Resonance The magnetization will only be rotated in a useful manner if it is subject to radio waves at the right frequency. This is a resonance phenomenon as known from the example of a swing: The amplitude of the oscillations will only change if the swing is pushed in synchrony with its natural oscillation frequency (i.e., “on resonance”).

Saturation Most often, the purpose of excitation is to create a signal, but not always so. Saturation pulses are used to get rid of signal in a given position, from a given type of tissue or from a particular metabolite. Examples include saturation slices and lipid suppression.

Saturation more generally refers to the attenuation of signal caused by reduced magnetization, e.g. resulting from incomplete recovery of longitudinal magnetization after prior excitation pulses.

Slice selection When excitation is performed while a gradient is active, the protons in some positions will precess in synchrony with the radio waves, and the nuclei will hence be rotated. At other positions, the resonance condition will not be fulfilled and the nuclei are hardly influenced. Specifically, a slice perpendicular to the direction of the gradient will be excited. When radio wave transmission and the gradient is subsequently turned off, a signal from the nuclei in the slice can not immediately be detected, however. Perpendicular to the slice, the last part of the slice selection gradient will have induced an unwanted phase roll that needs to be unwind before a signal from the slice can be measured. Slice selection will therefore typically be followed by a gradient directed opposite to the slice direction gradient (figure 16). This brings the nuclei in the slice into phase.

Spectroscopy The common term for techniques based on measuring differences in frequency content (color sight, variance in tone, NIRS. . .). Pretty much any MR technique is spectroscopic (imaging, for example), but the term MR-spectroscopy is usually used to refer to techniques aimed at distinguishing metabolite signals.

Spin Spin is a quantum mechanical quality possessed by protons and neutrons. In brief, this may be perceived as rotation of individual nuclei around an axis of their own. The magnetic characteristics of the nuclei (the dipole quality) is caused by spin and the dipole is along the direction of the spin. In nuclei, protons pair up in entities that seemingly have no spin. The same is true of neutrons. There are consequently only few nuclei that are suitable for MR (an unequal number of protons and/or neutrons is required).

Spin-echo If the magnetic field is inhomogeneous, the spins will loose alignment after excitation since some spins precess faster than others – signal is lost. Using pulses of radio waves, the dephasing can be reversed so that the lost signal is recovered: A 180 degree pulse inverts

the phase of the spins, so that those that have precessed the most are set most back in the process. Waiting a while longer, the signal can be recovered in the form of a spin-echo since the inhomogeneity persists.

Tesla The unit for magnetic field strength. The earth magnetic field is on average about 0.04 mT, while the magnetic fields used in MR scanning are typically 1-3 T.

Voxel Volume element. Typically refers to the volume that a given point in an image corresponds to. The word pixel (picture element) might be better known, but it does not reflect the depth-aspect of any MR measurement (2D rather than 3D). The voxel size is the same as the spatial resolution. The voxel dimensions are the edge-lengths of a voxel.