Dominik Weishaupt
Victor D. Köchli
Borut Marincek

How Does MRI Work?

An Introduction to the Physics and Function of Magnetic Resonance Imaging

Second Edition

constant T2*, which is typically shorter than T2. Most of the inhomogeneities that produce the T2* effect occur at tissue borders, particularly at air/tissue interfaces, or are induced by local magnetic fields (e.g. iron particles). The loss of the MR signal due to T2* effects is called *free induction decay (FID)*. T2* effects can be avoided by using spin echo sequences.

T2 denotes the process of energy transfer between spins, while T2* refers to the effects of additional field inhomogeneities contributing to dephasing.

T1 and T2 relaxation are completely *independent* of each other but occur more or less *simultaneously*! The decrease in the MR signal due to T2 relaxation occurs within the first 100–300 msec, which is long before there has been complete recovery of longitudinal magnetization M_z due to T1 relaxation (0.5–5 sec).

References

1. Gore JC, Kennan RP (1999) Physical principles and physiological basis of magnetic relaxation. In: Stark DD, Bradley WG Jr (eds). Magnetic resonance imaging, 3rd ed. Mosby-Year Book no 33, Mosby, St. Louis
2. Elster AD, Burdette JH (2001) Questions and answers in magnetic resonance imaging, 2nd ed. Mosby, St. Louis
3. Elster AD (1986) Magnetic resonance imaging, a reference guide and atlas. Lippincott, Philadelphia

3 Image Contrast

What determines the contrast of an MR image and how can we influence it?

Having explained the concepts of excitation and relaxation, we can now answer this question. *Three intrinsic features* of a biological tissue contribute to its signal intensity or brightness on an MR image and hence image contrast:

- The *proton density*, i.e. the number of excitable spins per unit volume, determines the maximum signal that can be obtained from a given tissue. Proton density can be emphasized by minimizing the other two parameters, T1 and T2. Such images are called *proton density-weighted* or simply *proton density images*.
- The *T1 time* of a tissue is the time it takes for the excited spins to recover and be available for the next excitation. T1 affects signal intensity indirectly and can be varied at random. Images with contrast that is mainly determined by T1 are called *T1-weighted images (T1w)*.
- The *T2 time* mostly determines how quickly an MR signal fades after excitation. The T2 contrast of an MR image can be controlled by the operator as well. Images with contrast that is mainly determined by T2 are called *T2-weighted images (T2w)*.

Proton density and T1 and T2 times are intrinsic features of biological tissues and may vary widely from one tissue to the next. Depending on which of these parameters is emphasized in an MR sequence, the resulting images differ in their tissue-tissue contrast. This provides the basis for the exquisite soft-tissue discrimination and diagnostic potential of MR imaging: based on their specific differences in terms of these three parameters, tissues that are virtually indistinct on computed tomography (CT) scans can be differentiated by MRI without contrast medium administration.

3.1 Repetition Time (TR) and T1 Weighting

In order to generate an MR image, a slice must be excited and the resulting signal recorded many times. Why this is so will be explained in ► Chapter 4.

***Repetition time (TR)* is the interval between two successive excitations of the same slice.**

Repetition time (TR) is the length of the relaxation period between two excitation pulses and is therefore crucial for T1 contrast. When TR is long, more excited spins rotate back into the z-plane and contribute to the regrowth of longitudinal magnetization. The more longitudinal magnetization can be excited with the next RF pulse, the larger the MR signal that can be collected.

If a *short* repetition time (less than about 600 msec) is selected, image contrast is strongly affected by T1 (TR A in ► Fig. 9). Under this condition, tissues with a short T1 relax quickly and give a large signal after the next RF pulse (and hence appear bright on the image). Tissues with a long T1, on the other hand, undergo only little relaxation between two RF pulses and hence less longitudinal magnetization is available when the next excitation pulse is

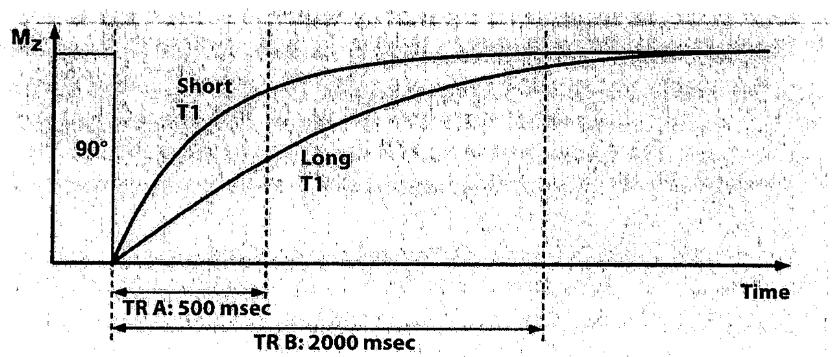

Fig. 9. Relationship between TR and T1 contrast. When TR is short (A), a tissue with a short T1 regains most of its longitudinal magnetization during the TR interval and hence produces a large MR signal after the next excitation pulse whereas a tissue with a long T1 gives only a small signal. When TR is long (B), the signal differences disappear because there is enough time for regrowth of longitudinal magnetization in both tissues

applied. These tissues therefore emit less signal than tissues with a short T1 and appear dark. An image acquired with a short TR is *T1-weighted* because it contains mostly T1 information.

If a fairly *long* repetition time (typically over 1500 msec) is selected, all tissues including those with a long T1 have enough time to return to equilibrium and hence they all give similar signals (TR B in ► Fig. 9). As a result, there is *less T1 weighting* because the effect of T1 on image contrast is only small.

Thus, by selecting the repetition time, we can control the degree of T1 weighting of the resulting MR image:

Short TR → strong T1 weighting
Long TR → low T1 weighting

The relationship between the MR signal of a tissue and its appearance on T1-weighted images is as follows:

Tissues with a *short T1* appear *bright* because they regain most of their longitudinal magnetization during the TR interval and thus produce a stronger MR signal.
Tissues with a *long T1* appear *dark* because they do not regain much of their longitudinal magnetization during the TR interval and thus produce a weaker MR signal.

3.2 Echo Time (TE) and T2 Weighting

What is an echo, anyway?

In ► Chapter 4 we will see that different gradients have to be applied to generate an MR image. For the time being it is sufficient to know that these gradients serve to induce controlled magnetic field inhomogeneities that are needed to encode the spatial origin of the MR signals. However, the gradients also contribute to spin dephasing. These effects must be reversed by applying a refocusing pulse before an adequate MR signal is obtained. The signal induced in the receiver coil after phase coherence has been restored is known as a *spin echo* and can be measured.

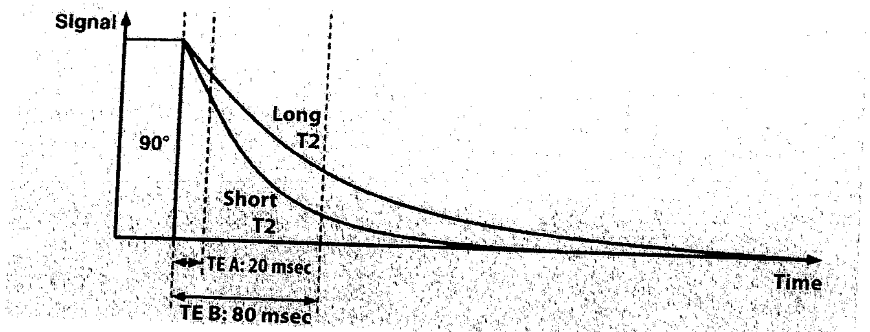

Fig. 10. Relationship between TE and T2 contrast. When TE is very short (A), there is virtually no signal difference between two tissues with different T2 times whereas clear differences become apparent when TE is longer (B): a tissue with a short T2 rapidly loses signal and becomes dark while a tissue with a long T2 retains its brighter signal for a longer time

***Echo time (TE)* is the interval between application of the excitation pulse and collection of the MR signal.**

The echo time determines the influence of T2 on image contrast. T2 is in the range of several hundred milliseconds and therefore much shorter than T1.

If a short echo time is used (less than about 30 msec), the signal differences between tissues are small (TE A in ▸ Fig. 10) because T2 relaxation has only just started and there has only been little signal decay at the time of echo collection. The resulting image has low T2 weighting.

If a longer echo time in the range of the T2 times of tissues (over about 60 msec) is used, the tissues are depicted with different signal intensities on the resulting MR image (TE B in ▸ Fig. 10): tissues with a short T2 having lost most of their signal appear dark on the image while tissues with a long T2 still produce a stronger signal and thus appear bright. This is why, for instance, cerebrospinal fluid (CSF) with its longer T2 (like water) is brighter on T2-weighted images compared with brain tissue.

By selecting an echo time (TE), the operator can control the degree of T2 weighting of the resulting MR image:

Short TE → low T2 weighting
Long TE → strong T2 weighting

▸ Fig. 10 also illustrates the relationship between the T2 value of a tissue and its appearance on T2-weighted images:

Tissues with a *short* T2 appear *dark* on T2-weighted images, tissues with a *long* T2 appear *bright* on T2-weighted images!

The relationships between TR and TE and the resulting image contrast are summarized in ▸ Table 1. ▸ Table 2 lists the signal intensities of different tissues on T1- and T2-weighted images. ▸ Table 3 provides an overview of intrinsic contrast parameters of selected tissues.

A typical T1-weighted spin echo (SE) sequence is acquired with a TR/TE of 340/13 msec. A T2-weighted fast spin echo (FSE) MR image can be acquired with a TR/TE of 3500/120 msec. MR images that combine T1 and T2 effects are known as *proton density-weighted images (PD images)*. PD images with a TE of about 40 msec are also referred to as *intermediate-weighted images*. As a rule, PD images have a higher signal-to-noise ratio (▸ Chapter 5) than comparable T1- and T2-weighted images because the long TR allows recovery of longitudinal magnetization while the short TE minimizes the signal decrease due to the decay of transverse magnetization.

Typical parameters for acquisition of a PD image are for instance a TR/TE of 2000/15 msec for a PD-weighted SE sequence and a TR/TE of 4400/40 msec for a PD-weighted FSE sequence. PD sequences are especially useful for evaluating structures with low signal intensities such as the bones or connective tissue structures such as ligaments and tendons. Proton density weighting is often used for high-resolution imaging. SE sequences are preferred over FSE sequences for PD imaging because SE images are less prone to distortion. In the clinical setting, PD sequences are mainly used for imaging of the brain, spine, and musculoskeletal system.

3.3 Saturation at Short Repetition Times

In the section on repetition time, we already said that there is little time for the regrowth of longitudinal magnetization when TR is very short. The shorter the TR, the smaller the component of longitudinal magnetization that is restored and is available for subsequent excitation. As a consequence, the MR signal decreases as well. When a series of excitation pulses is applied, the MR signal becomes weaker and weaker after each repeat pulse. This process is known as *saturation* (▸ Fig. 11).

Table 1. Image contrast as a function of TR and TE

	TR	TE
T1-weighted	Short	Short
T2-weighted	Long	Long
Proton density weighted (intermediate-weighted)	Long	Short

Table 2. Signal intensities of different tissues on T1- and T2-weighted images

Tissue	T1-weighted image	T2-weighted image
Fat	Bright	Bright
Aqueous liquid	Dark	Bright
Tumor	Dark	Bright
Inflammatory tissue	Dark	Bright
Muscle	Dark	Dark
Connective tissue	Dark	Dark
Hematoma, acute	Dark	Dark
Hematoma, subacute	Bright	Bright
Flowing blood	No signal due to black blood effect (▶ Chapter 7.2)	
Fibrous cartilage	Dark	Dark
Hyaline cartilage	Bright	Bright
Compact bone	Dark	Dark
Air	No signal	No signal

Table 3. Relative proton densities (%) and intrinsic T1 and T2 times (in msec) of different tissues

Tissue	Proton density	T1 (1.5 T)	T2 (1.5 T)
CSF	100	> 4000	> 2000
White matter	70	780	90
Gray matter	85	920	100
Meningioma	90	400	80
Metastasis	85	1800	85
Fat	100	260	80

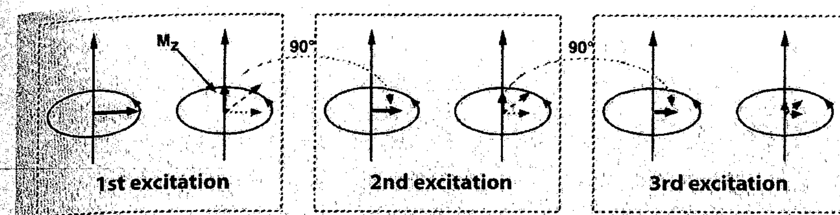

Fig. 11. Mechanism of saturation. With a very short TR, the longitudinal magnetization, M_z, that will recover in the interval and be available for subsequent excitation decreases after each RF pulse. In the example shown, the TR is so short that slightly less than half of the original longitudinal magnetization can regrow before the next excitation pulse is delivered

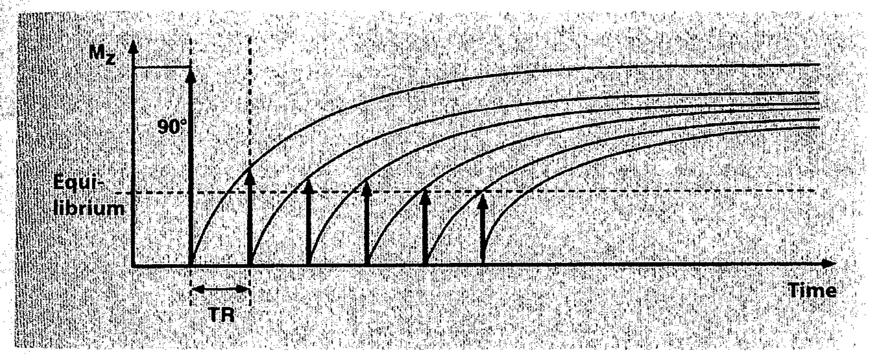

Fig. 12. Longitudinal magnetization at short repetition time. After repeat excitation at very short intervals, the amount of longitudinal magnetization, M_z, restored after each pulse settles at a low level (equilibrium or steady state). In this situation, the individual MR signals that form after each excitation are very weak

Saturation is an important issue when fast or ultrafast MR techniques are used. Here the MR signal may become very weak due to the very short repetition times (▶ Fig. 12). We will return to this phenomenon when we discuss gradient echo sequences.

3.4 Flip Angle (Tip Angle)

Partial flip angle imaging is a technique that can be used to minimize saturation and obtain an adequate MR signal despite a very short repetition time. A smaller flip angle does not deflect the magnetization all the way through

90° but only by some fraction of 90° (e.g. 30°). As a result there is less transverse magnetization and the individual MR signals are smaller while more longitudinal magnetization is available for subsequent excitation even if TR is very short. However, the overall signal is larger than the one obtained with a 90° flip angle. *In general, the shorter the TR, the smaller the flip angle that is needed* to prevent excessive saturation. The flip angle maximizing the signal for a given TR and TE is known as the *Ernst angle.*

3.5 Presaturation

Another option available to modulate image contrast is *presaturation.* This technique employs an initial 90° or 180° inverting pulse that is delivered before the data for image generation is acquired. A presaturation pulse or prepulse can be combined with all basic pulse sequences (SE, FSE, GRE, and EPI sequences). But what is the benefit of this technique?

Fast gradient echo sequences are often limited by poor image contrast because the short repetition times lead to homogeneous saturation of different tissues. As we have seen above, the resultant images are T1-weighted but not very strongly so. Stronger T1 weighting can be achieved by selecting a larger flip angle but the resultant MR signal would be much too weak to obtain a reasonable image quality because saturation would increase as well.

This is why presaturation is used to enhance T1 contrast. A more pronounced T1 effect is achieved with a 180° inverting pulse than with a 90° pulse because a 180° pulse inverts all longitudinal magnetization. As a result, T1 relaxation begins at −1 rather than 0 and twice as much longitudinal magnetization is available. Additionally, the operator can modulate the T1 effect by varying the time interval between the 180° inversion pulse and the excitation pulse (= inversion time, TI). TI can be chosen such that the signal contribution from a specific tissue is eliminated by applying the excitation pulse when the tissue has no magnetization. Thus, a short TI will suppress the signal from fat (► Chapter 7.5) and a long TI the signal from CSF (FLAIR sequence, ► Chapter 7.6). Another practical application is late-enhancement imaging in patients with myocardial infarction (► Chapter 11.8).

3.6 Magnetization Transfer

Without explicitly saying so, we have thus far always referred to free protons (i.e. protons in free water) when talking about protons because only these contribute to the MR signal. In addition to water protons, biological tissues also contain a specific pool of protons bound in macromolecules (usually proteins). These macromolecular protons cannot be directly visualized because of their very short T1. They have a wider range of Larmor frequencies than the water protons. This is why macromolecular protons can also be excited by RF pulses with frequencies slightly different from the Larmor frequency of hydrogen protons. Hence, it is possible to selectively excite a tissue with a large pool of macromolecular protons without directly affecting the protons in free water. Repeated delivery of the magnetization transfer pulse saturates the magnetization of the macromolecular protons from where it is transferred to free protons nearby. This process is associated with a drop in signal that depends in magnitude on the concentration of macromolecules and their interaction with free water and is known as *magnetization transfer* (► Fig. 13). The decrease in signal intensity due to magnetization transfer is large for solid tissues but only small for fluids (as long as their macromolecule content is low) and fatty tissue.

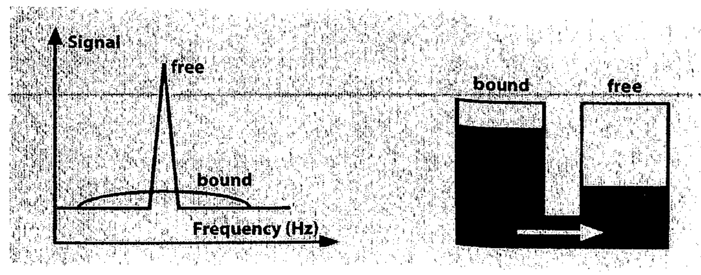

Fig. 13.

The phenomenon of magnetization transfer is exploited to improve image contrast using a technique known as *magnetization transfer imaging.* Magnetization transfer contrast (MTC) is used in cartilage imaging where it improves contrast between synovial fluid and cartilage because synovial fluid contains only few bound protons and thus shows only little magnetization transfer while cartilage contains a large proportion of bound protons and therefore shows pronounced magnetization transfer. In the brain, the MTC technique improves the detection of gadolinium-enhancing lesions.

References

1. Nessaiver M (1996) All you really need to know about MR imaging physics. University of Maryland Press, Baltimore
2. Duerk JL (1997) Relaxation and contrast in MR imaging. In: Riederer SJ, Wood ML (eds) Categorical course in physics: the basic physics of MR imaging. RSNA Publications no 19, Oak Brook
3. Elster AD, Burdette JH (2001) Questions and answers in magnetic resonance imaging, 2nd edn. Mosby, St. Louis

4 Slice Selection and Spatial Encoding

In the preceding sections, we have outlined the MR phenomenon and discussed the role of repetition and echo times. Now, finally, we want to make a picture! As a tomographic technique, MR imaging generates cross-sectional images of the human body. The excitation pulse is therefore delivered only to the slice we want to image and not to the whole body. How is this accomplished and how does the signal provide us with information about its origin within the slice?

For illustration, we consider a transverse (axial) slice or cross-section through the body. The magnetic field generated by most MR scanners is not directed from top to bottom, as in the illustrations we have used so far, but along the body axis of the person being imaged. From now on, this is the direction that will be designated by "z" since, as already said, *z stands for the direction of the main magnetic field*. The magnetic field gradients that now come into play are represented by wedges with the thick side indicating the higher field strength and the tip the lower field strength.

Both the excitation of a specific slice and the identification of the site of origin of a signal within the slice rely on the fact that the *precessional or Larmor frequency is proportional to the magnetic field strength*. In addition, recall that protons are excited only by an RF pulse with a frequency roughly equal to their Larmor frequency (*resonance condition*). If a uniform field of identical strength were generated throughout the body, all protons would have the same Larmor frequency and would be excited simultaneously by a single RF pulse.

To enable selective excitation of a desired slice, the magnetic field is therefore made *inhomogeneous* in a linear fashion along the z-direction by means of a gradient coil. As a result, the magnetic field strength has a smooth *gradient* so that, for example, it is weakest at the patient's head and strongest at the feet. The Larmor frequencies thus change gradually along the z-axis and

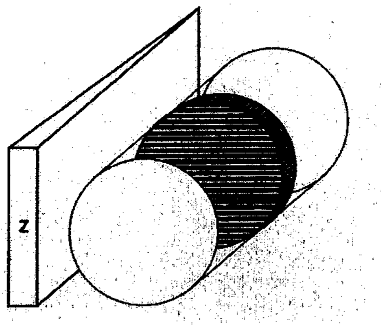

Fig. 14. Slice selection by means of the z-gradient. An RF pulse of a specific frequency excites exactly one slice (hatched) with adjacent slices being unaffected because they have different resonant frequencies

each slice now has its unique frequency. Hence, application of an RF pulse that matches the Larmor frequency of the desired slice will excite only protons within the chosen slice while the rest of the body remains unaffected (► Fig. 14).

Gradients are additional magnetic fields that are generated by gradient coils and add to or subtract from the main magnetic field. Depending on their position along the gradient, protons are temporarily exposed to magnetic fields of different strength and hence differ in their precessional frequencies. A shallow gradient generates a thicker slice while a steep gradient generates a thinner slice (► Fig. 15a). Slice position is defined by changing the center frequency of the RF pulse applied (► Fig. 15b).

Having selected slice position and thickness by application of an appropriate slice-select gradient, we can now proceed to explain how the spatial position of an MR signal is identified. This is accomplished by *spatial encoding*, which is the most difficult task in generating an MR image and requires the application of additional gradients that alter the magnetic field strength along the y- and x-axes. Once we have grasped the concept of spatial encoding, it will be easy to understand the different kinds of artifacts that degrade MR image quality in clinical practice. Spatial encoding comprises two steps, *phase encoding* and *frequency encoding*. These two steps are discussed in their appropriate order, which means that we must first turn to the more difficult technique of phase encoding.

For phase encoding, a gradient in the y-direction (from top to bottom) is switched on after the spins have been excited and precess in the xy-plane. Such a *phase-encoding gradient* alters the Larmor frequencies of the spins according to their location along the gradient. As a result, the excited spins higher up in the scanner experience a stronger magnetic field and thus gain

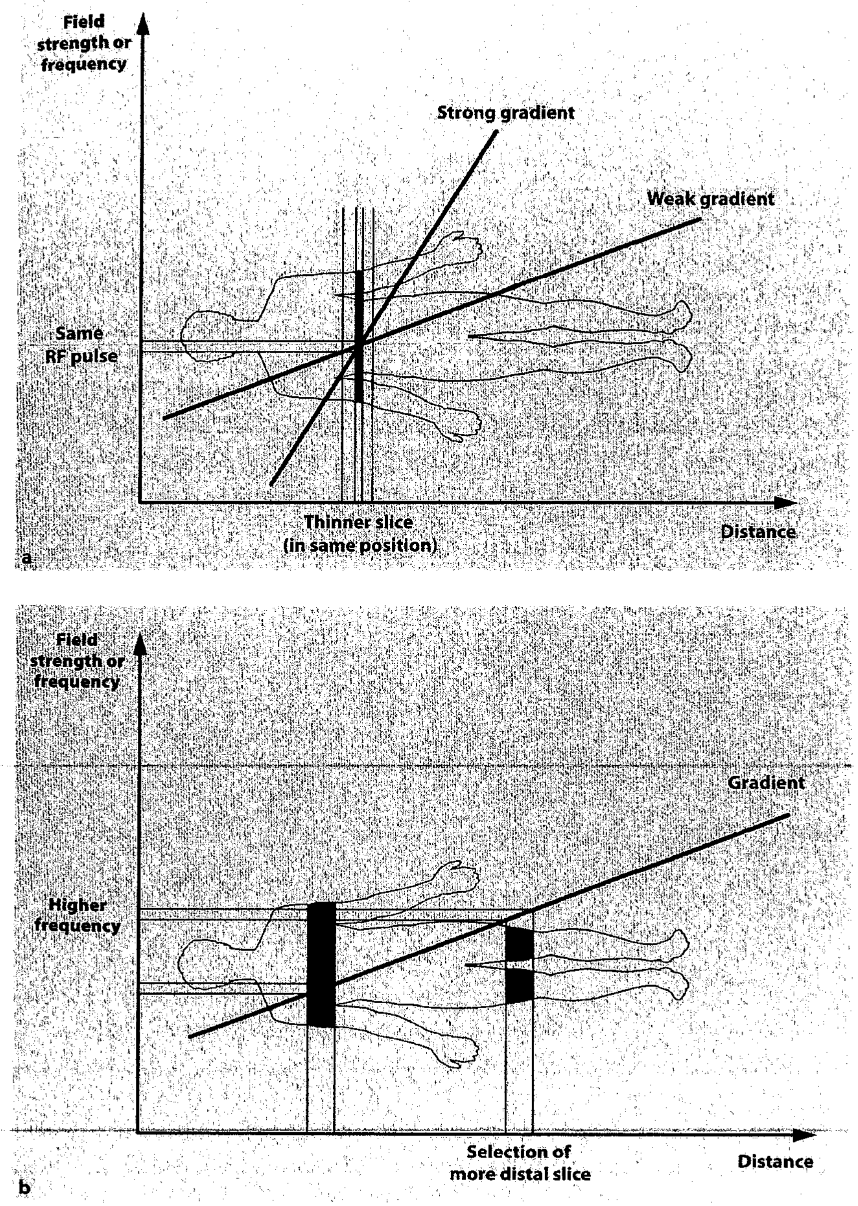

Fig. 15. **a** The strength of the gradient applied defines slice thickness. An RF pulse of a given frequency bandwidth produces a thin slice if the gradient is strong and a thick slice if the gradient is weak. **b** The center frequency of the RF pulse applied determines the location of the slice

phase relative to the somewhat slower spins further down. The result is a *phase shift* of the spins relative to each other (► Fig. 16). The degree of phase shift is determined by the duration and amplitude of the phase-encoding gradient and by the physical location of the precessing nuclei along its length. The phase gain is higher for nuclei closer to the top of the scanner. When the gradient is switched off after some time, all spins return to their initial rate of precession yet are now ahead or behind in phase relative to their previous state. Phase now varies along the y-axis in a linear fashion and each line within the slice can thus be identified by its unique phase.

The second spatial dimension of the MR signal that needs to be identified is encoded by changes in frequency along the x-direction. To this end, a *frequency-encoding gradient* is applied – in our example along the x-axis. This gradient generates a magnetic field that increases in strength from right to left. The corresponding changes in Larmor frequencies make spins on the left side precess slower than the ones on the right side. When we collect the MR signal while the frequency-encoding gradient is switched on, we do not obtain a single frequency but a whole *frequency spectrum* (► Fig. 17) comprising high frequencies from the right edge of the slice and low frequencies from the left edge. Each column of the slice is thus characterized by a specific frequency. Frequency and phase together enable unique spatial identification of each volume element (*voxel*).

The MR signal measured in this way contains two pieces of information. The *frequency* locates the signal along the x-axis. This information can be extracted directly by applying a *Fourier transform* (or frequency analysis) to decompose the signal into its component frequencies along the frequency-encoding direction. This mathematical operation serves to identify the individual frequencies that make up a signal. The *phase distribution* within each frequency provides information on the place of origin of the corresponding signal component along the y-axis. How do we get this second piece of information when we merely have the sum of all spins with the same frequency but different phases? The phases of the individual spins cannot be derived from a single signal but only from a set of signals. In this respect, the MR signal is comparable to a mathematical equation with many unknowns (e.g. 256) of which we only have the result but not the individual unknowns.

To calculate the unknowns, one needs as many *different* equations as there are unknowns. Applied to the MR signal, this means that we must repeat the sequence many times with increasing or decreasing gradient strengths. The set of echoes acquired with different phase encodings allows us to derive the required phase-encoded spatial information by applying a second Fourier

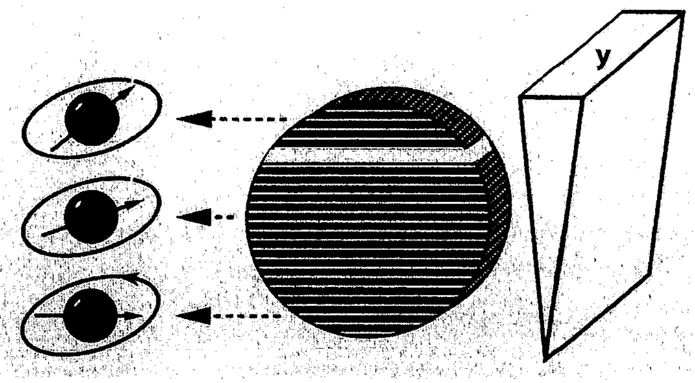

Fig. 16. Phase encoding by means of the y-gradient. Each horizontal line (e.g. the white line in the example) is identified by a unique amount of phase shift

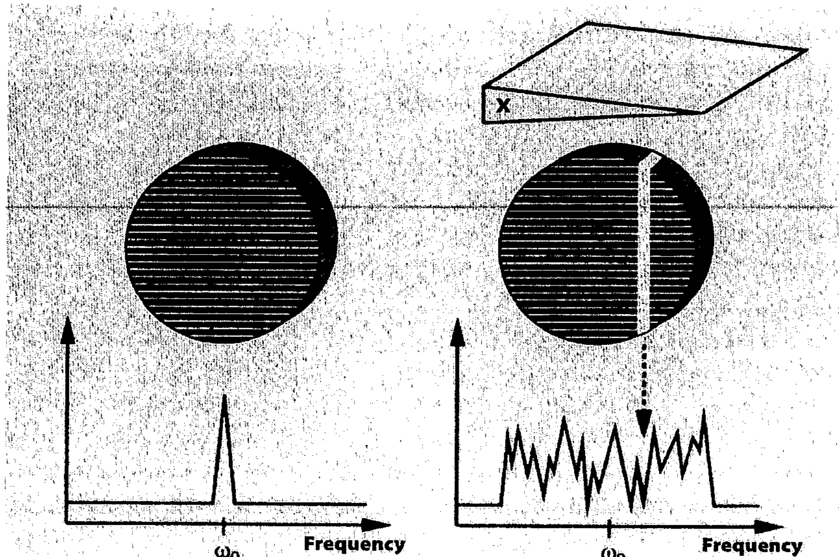

Fig. 17. Frequency encoding by means of the x-gradient. With the gradient switched off (*left*), only a single frequency is received, the Larmor frequency ω_0. With the gradient switched on (*right*), a frequency spectrum is received with each column being identified by its unique frequency

transform, this time along the y-axis. Hence, for spatial encoding in two dimensions, the Fourier transform has to be applied twice, which is why this technique is called two-dimensional Fourier transform (2D-FT). To perform such complex calculations – which corresponds to solving a set of equations with, for example, 256 equations and 256 unknowns – an MR scanner is equipped with a dedicated computer, a so-called array processor.

Repeated measurements are performed with a specific temporal delay, the *repetition time* previously mentioned. The number of phase-encoding steps performed depends on the desired image quality. More phase-encoding steps improve resolution and image quality but also prolong scan time.

4.1 Three-Dimensional Spatial Encoding

It is sometimes desirable to image a whole volume rather than just a number of individual slices, for the following reasons:

- The acquired source data set is to be postprocessed, for example, to generate reconstructions in different planes.
- One wishes to acquire thin slices without drowning the MR signal in noise. Thin slices yield weaker MR signals because fewer spins are excited. This drawback can be overcome by benefiting from the stronger signal generated by an entire volume and extracting the individual slices afterwards.

If we want to excite an entire volume instead of only a single slice, we need an additional step to encode *spatial information in the third direction (z)*. (This is the information provided by the slice-select gradient when a single slice is scanned.)

In volume imaging, the spatial position of a signal along the z-direction is encoded by applying an additional *phase-encoding gradient*, a z-gradient. As with the phase encoding gradient along the y-axis, the number of repetitions performed with different values of the gradient determines image resolution in the z-direction, which corresponds to the slice thickness in 2D imaging. The computation of a volume image is even more time-consuming because a *three-dimensional Fourier transform (3D-FT)* with an additional transform in the z-direction has to be performed. The 3D-FT yields a 3D data set of a volume without interslice gaps from which reconstructions in any plane or projections can be generated with the aid of suitable reconstruction algorithms. These techniques are very useful for MR angiography.

The major drawback of volume imaging is that it may unduly prolong

scan time since spatial encoding in the x- and y-directions must be performed for each phase-encoding step along the z-axis.

4.2 K-Space

Data collected from the signals is stored in a mathematical area known as k-space. K-space has two axes with the horizontal axis (k_x) representing the frequency information and the vertical axis (k_y) the phase information (▸ Fig. 18). It is a graphic matrix of digitized MR data that represents the MR image before Fourier transformation is performed. Each line in k-space corresponds to one measurement and a line is acquired for each phase-encoding step. The center line (0) is filled with the data that is unaffected by the phase-encoding gradient (gradient isocenter).

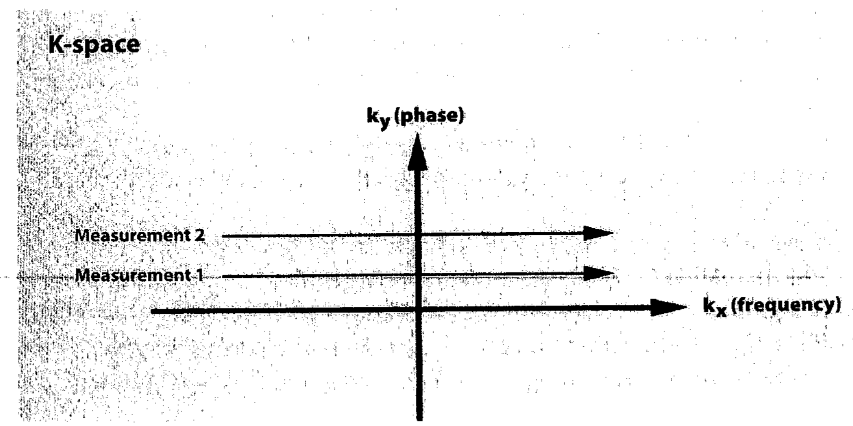

Fig. 18. K-space. k_x is the frequency axis, k_y the phase axis. The data from each measurement fills a different horizontal line

An MR image is created from the raw data by application of 2D-FT after the scan is over and k-space is filled. The lines in k-space *do not* correspond one to one with the lines in the resulting MR image. Rather, data in the *center of k-space* primarily determines *contrast* in the image while the *periphery (the outer lines)* primarily contains *spatial information*. When discussing fast sequences (▸ Chapter 8), we will also learn how we can speed up scanning by filling more than one k-space line with a single acquisition.

References

1. Wehrli FW (1997) Spatial encoding and k-space. In: Riederer SJ, Wood ML (eds). Categorical course in physics: The basic physics of MR imaging. RSNA Publications no 31, Oak Brook
2. Wood ML, Wehrli FW (1999) Principles of magnetic resonance imaging. In: Stark DD, Bradley WG Jr (eds) Magnetic resonance imaging, 3rd edn. Mosby-Year Book no 28, Mosby, St. Louis

5 Factors Affecting the Signal-to-Noise Ratio

In the preceding chapters we have learned how an MR signal is generated and how the collected signal is processed to create an MR image. What we have disregarded so far is that the MR signal can be degraded by noise. Image noise results from a number of different factors:

- Imperfections of the MR system such as magnetic field inhomogeneities, thermal noise from the RF coils, or nonlinearity of signal amplifiers.
- Factors associated with image processing itself.
- Patient-related factors resulting from body movement or respiratory motion.

The relationship between the MR signal and the amount of image noise present is expressed as the *signal-to-noise ratio (SNR)*. Mathematically, the SNR is the quotient of the signal intensity measured in a *region of interest (ROI)* and the standard deviation of the signal intensity in a region outside the anatomy or object being imaged (i.e. a region from which no tissue signal is obtained).

A high SNR is desirable in MRI. The SNR is dependent on the following parameters:

- Slice thickness and receiver bandwidth
- Field of view
- Size of the (image) matrix
- Number of acquisitions
- Scan parameters (TR, TE, flip angle)
- Magnetic field strength
- Selection of the transmit and receive coil (RF coil)

Before we discuss the effects of each of these parameters, it is first necessary to clarify some concepts.

5.1 Pixel, Voxel, Matrix

An MR image is digital and consists of a matrix of *pixels* or picture elements. A *matrix* is a two-dimensional grid of rows and columns. Each square of the grid is a pixel which is assigned a value that corresponds to a signal intensity. Each pixel of an MR image provides information on a corresponding three-dimensional volume element, termed a *voxel* (▸ Fig. 19). The voxel size determines the spatial resolution of an MR image.

The size of a voxel can be calculated from the field of view, the matrix size, and the slice thickness. In general, the resolution of an MR image increases as the voxel size decreases.

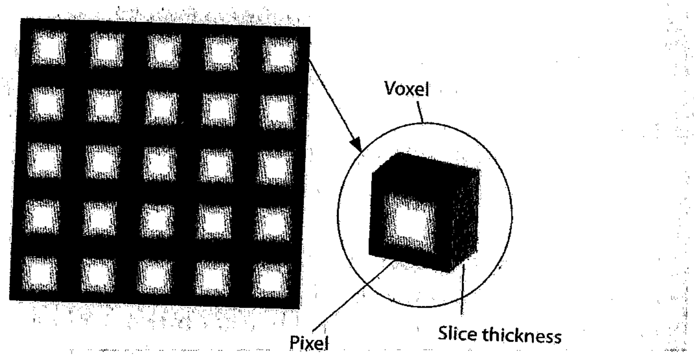

Fig. 19. A voxel is the tissue volume represented by a pixel in the two-dimensional MR image

5.2 Slice Thickness and Receiver Bandwidth

To achieve optimal image resolution, very thin slices with a high SNR are desirable. However, thinner slices are associated with more noise, and so the SNR decreases with the slice thickness. Conversely, thicker slices are associated with other problems such as an increase in partial volume effects.

The poorer SNR of thin slices can be compensated for to some extent by increasing the number of acquisitions or by a longer TR. Yet this is accomplished only at the expense of the overall image acquisition time and reduces the cost efficiency of the MR imaging system.

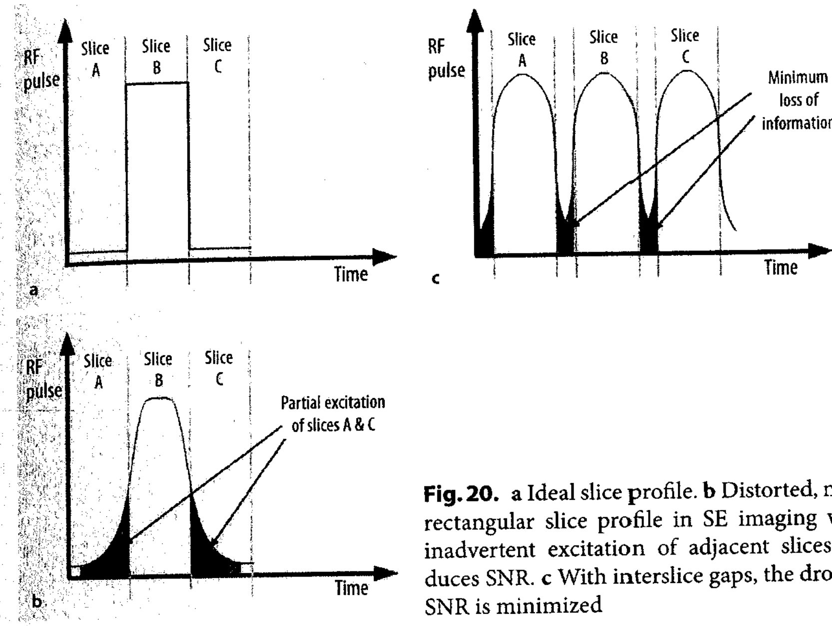

Fig. 20. **a** Ideal slice profile. **b** Distorted, non-rectangular slice profile in SE imaging with inadvertent excitation of adjacent slices reduces SNR. **c** With interslice gaps, the drop in SNR is minimized

The *receiver bandwidth* is the range of frequencies collected by an MR system during frequency encoding. The bandwidth is either set automatically or can be changed by the operator. A wide receiver bandwidth enables faster data acquisition and minimizes chemical shift artifacts (▸ Chapter 13.3) but also reduces SNR as more noise is included. Halving the bandwidth improves SNR by about 30%. With a narrow bandwidth, on the other hand, there will be more chemical shift and motion artifacts and the number of slices that can be acquired for a given TR is limited.

An *interslice gap* is a small space between two adjacent slices. It would be desirable to acquire contiguous slices but interslice gaps are necessary in SE imaging due to imperfections of RF pulses. Because the resultant slice profiles are not perfectly rectangular (▸ Fig. 20), two adjacent slices overlap at their edges when closely spaced. Under these conditions, the RF pulse for one slice also excites protons in adjacent slices. Such interference is known as *cross-talk*.

Cross-talk produces saturation effects and thus reduces SNR (▸ Fig. 20b).

In selecting an appropriate interslice gap one has to find a compromise between an optimal SNR, which requires a large enough gap to completely eliminate cross-talk, and the desire to reduce the amount of information that is missed when the gap is too large. In most practical applications an interslice gap of 25–50% of the slice thickness is used.

Alternatively, the undesired saturation of protons in adjacent slices can be reduced by *multislice imaging*, which will be discussed in ▸ Chapter 7.3. Scan times are somewhat longer unless a shorter TR is used.

Gradient echo (GRE) sequences are different. They do not require a 180° refocusing pulse and thus allow the acquisition of contiguous slices without interslice gaps.

5.3 Field of View and Matrix

There is a close relationship between field of view (FOV) and SNR. When matrix size is held constant, the FOV determines the size of the pixels. *Pixel size in the frequency-encoding direction* is calculated as the FOV in mm divided by the matrix in the frequency-encoding direction and *pixel size in the phase-encoding direction* as the FOV in mm divided by the matrix in the phase-encoding direction.

As illustrated in ▸ Fig. 21, pixel size changes with the FOV. A smaller FOV results in a smaller pixel size as long as the matrix is unchanged. Pixel size is crucial for the spatial resolution of the MR image. With the same FOV, a finer matrix (i.e. a matrix consisting of more pixels) results in an improved spatial resolution (▸ Figs. 22 and 23).

Conversely, a coarser matrix (i.e. one with fewer pixels) results in a poorer spatial resolution when the FOV is held constant (▸ Fig. 23).

From what has been said so far, one might conclude that the matrix should be as large as possible in order to encompass a maximum of picture elements. This is true in terms of image resolution but the minimum pixel size is limited by the fact that, in general, *SNR decreases with the size of the voxel.*

Another limiting factor is image acquisition or scan time, which increases in direct proportion to the matrix size. *Scan time* is the key to the economic efficiency of all MR systems and can be calculated by a simple equation.

Scan time = TR × number of phase-encoding steps × number of signal averages (NSA) [echo train length (ETL)].

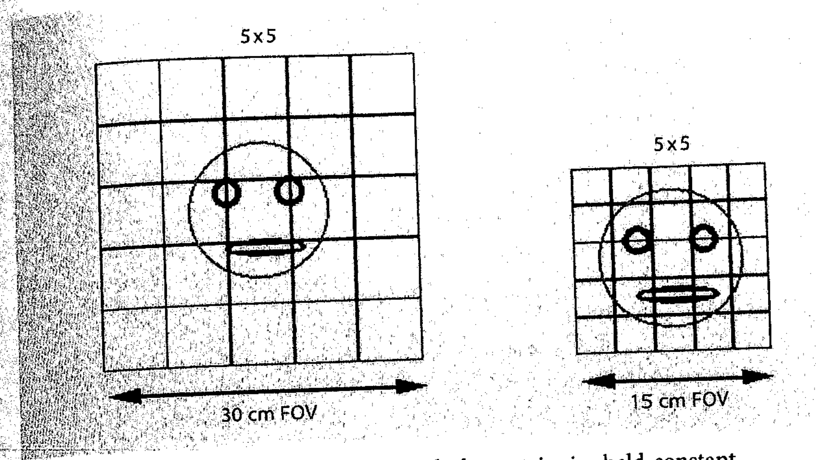

Fig. 21. Effect of the FOV on pixel size with the matrix size held constant

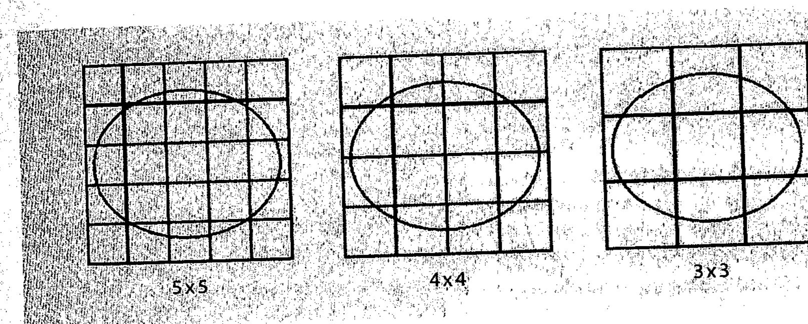

Fig. 22. A smaller matrix size with the FOV held constant results in larger pixels and thus a poorer spatial resolution

A "trick" can be used to achieve a high spatial resolution in a reasonable scan time. This is done by reducing the field of view only in the phase-encoding direction *(rectangular field of view)* and is possible because spatial resolution is determined by the matrix size in the frequency-encoding direction while scan time is determined by the matrix size in the phase-encoding direction. Reduction of the matrix size in the phase-encoding direction therefore does not reduce spatial resolution. Filling only one-half the normal number of phase-encoding lines in k-space reduces imaging time and the FOV by 50%. However, use of a rectangular FOV may be associated with wraparound artifacts when signal outside the FOV in the phase-encoding direction is mapped back into the image at an incorrect location

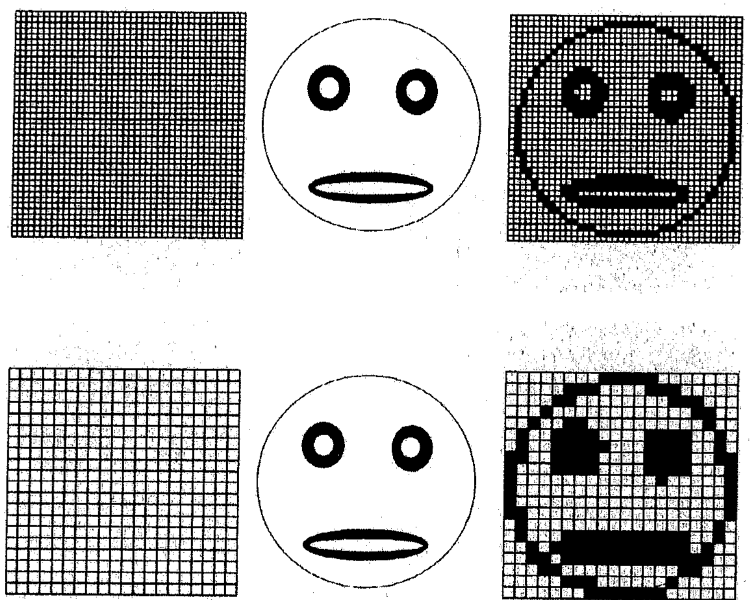

Fig. 23. Effect of matrix size on spatial resolution. Consider we are imaging a smiley face with a fine matrix (*top*) and a coarse matrix (*bottom*). The pixels representing the face are black. The two depictions of the face illustrate the much poorer detail resolution when a coarser matrix (*bottom right*) is used: pupil and eye cannot be distinguished and the open mouth appears to be closed

(► Chapter 13). This kind of foldover can be suppressed by specific anti-aliasing options such as "no phase wrap". Moreover, reduction of the FOV in the phase-encoding direction is associated with a slight drop in SNR. A rectangular FOV is typically used to image the spine and extremities and for MR angiography.

Scan time can be shortened further on state-of-the-art scanners that allow one to use rectangular fields of view in combination with rectangular pixels.

Finally, various *techniques of partial k-space acquisition* (► Figs. 24, 25, and 26) save scan time without one having to change the voxel size. In *partial Fourier imaging*, only half the lines (or slightly more) in the phase-encoding direction are filled (► Fig. 24) while *fractional* or *partial echo imaging* (► Fig. 25) refers to a technique with incomplete filling of the frequency-

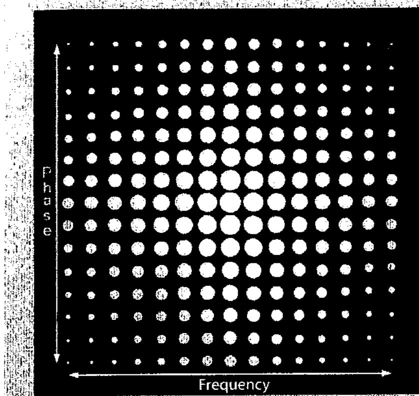

Fig. 24. Complete k-space sampling. Each data point represents one frequency-encoding line and one phase-encoding line

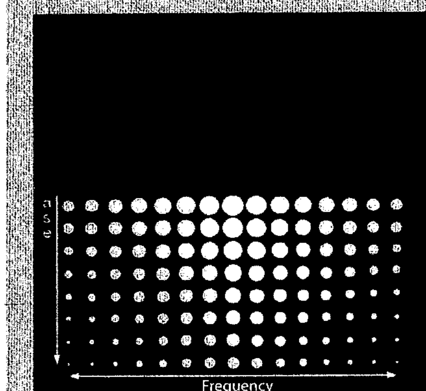

Fig. 25. Partial Fourier imaging. Slightly less than half the k-space lines in the phase-encoding direction are not sampled (*gray dots*). These lines are interpolated

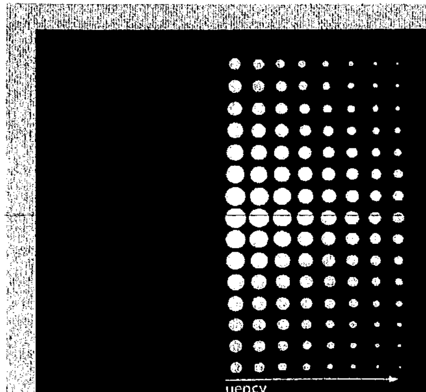

Fig. 26. Fractional echo imaging. Slightly less than half of the k-space lines in the frequency-encoding direction are not filled directly (*gray dots*). The unfilled lines represent the echo portions that have not been sampled. The resulting MR image has a similar resolution but poorer SNR compared with an image generated with complete k-space sampling (► Fig. 24) (as less "true" data is incorporated)

encoding lines by sampling only part of each echo. Both techniques rely on the inherent symmetry of k-space that allows one to interpolate the unfilled lines and to thus reconstruct an MR image when only half or slightly more than half the lines of k-space have been sampled. Both methods *shorten scan time* but this is accomplished *at the expense of SNR*. Partial Fourier and fractional echo imaging are needed for fast imaging techniques (▶ Chapter 8).

In routine 2D Fourier transform or spin-warp imaging, k-space is filled sequentially one line at a time (linear or Cartesian k-space acquisition). More sophisticated sequences use spiral k-space trajectories that fill the lines from the center toward the periphery (*elliptical centric ordering of k-space, CENTRA*). In MR angiography, for instance, this technique is used to fill the center of k-space with the data important for evaluating contrast enhancement patterns.

5.4 Number of Excitations

The *number of excitations (NEX)* or *number of signal averages (NSA)* denotes how many times a signal from a given slice is measured. The SNR, which is proportional to the square root of the NEX, improves as the NEX increases, but scan time also increases linearly with the NEX.

5.5 Imaging Parameters

Other parameters affecting the SNR are the sequence used, echo time (TE), repetition time (TR), and the flip angle. The SNR increases with the TR but the T1 effect is also lost at longer TRs. Conversely, the SNR decreases as the TE increases. With a short TE, the T2 contrast is lost. For this reason, the option of shortening TE to improve SNR is available only for T1-weighted sequences.

5.6 Magnetic Field Strength

Applying a *higher magnetic field strength increases longitudinal magnetization* because more protons align along the main axis of the magnetic field, resulting in an increase in SNR. The improved SNR achieved with high-field systems (▶ Chapter 14) can be utilized to generate images with an improved spatial resolution or to perform fast imaging.

5.7 Coils

An effective means to improve SNR, without increasing voxel size or lengthening scan time, is selecting an appropriate *radiofrequency (RF) coil*. In general, an RF coil should be as close as possible to the anatomy being imaged and surround the target organ. The nearer the coil can be placed to the organ under examination, the better the resulting signal. RF coils can be used either to transmit RF and receive the MR signal or to act as receiver coils only. In the latter case, excitation pulses are delivered by the body coil. The basic coil types that are distinguished are briefly described below.

5.7.1 Volume Coils

Volume coils may be used exclusively as *receive coils* or as *combined transmit/receive coils*. Volume coils completely surround the anatomy to be imaged. Two widely used volume coil configurations are the *saddle coil* and the *birdcage coil*. Volume coils are characterized by a homogeneous signal quality. Another type of volume coil is the *body coil*, which is an integral part of an MR scanner and is usually located within the bore of the magnet itself. Head and extremity coils are further examples of volume coils.

5.7.2 Surface Coils

Most surface coils can only receive the MR signal and rely on the body coil for delivery of RF pulses. Combined transmit/receive surface coils are also available. Surface coils are used for spinal MRI and imaging of small anatomic structures.

5.7.3 Intracavity Coils

Intracavity coils are small *local receive coils* that are inserted into body cavities to improve image quality as a result of the closer vicinity to the target organ. In clinical MRI, endorectal coils are used for imaging of the prostate and the anal sphincter muscle. Experimental applications include endovascular imaging and imaging of hollow organs.

5.7.4 Phased-Array Coils

Phased-array coils serve to *receive* MR signals. A phased-array system consists of several independent coils connected in parallel or series. Each coil feeds into a separate receiver. The information from the individual receivers is combined to create one image. Phased-array coils yield images with a high spatial resolution and allow imaging with a larger field of view as they improve both SNR and signal homogeneity.

► Table 4 summarizes the factors affecting SNR.

► Table 5 summarizes the effects of matrix size, slice thickness, and FOV on spatial resolution.

► Table 6 summarizes the effects of different sequence parameters on scan time.

Table 4. Effects of different imaging and sequence parameters on signal-to-noise ratio (SNR)

Change in parameter	SNR
Increasing slice thickness	Increases
Increasing FOV	Increases
Reducing FOV in phase-encoding direction (rectangular FOV)	Decreases
Increasing TR	Increases
Increasing TE	Decreases
Increasing matrix size in frequency-encoding direction	Decreases
Increasing matrix size in phase-encoding direction	Decreases
Increasing NEX	Increases
Increasing magnetic field strength	Increases
Increasing receiver bandwidth	Decreases
Employing local coils	Increases
Partial Fourier imaging	Decreases
Fractional echo imaging	Decreases

Table 5. Effects of matrix size, slice thickness, and field of view (FOV) on spatial resolution

Change in parameter	Spatial resolution
Increasing matrix size	Increases
Using thicker slices	Decreases
Increasing FOV	Decreases

Table 6. Effects of different sequence parameters on scan time

Change in parameter	Scan time
Using thicker slices	Decreases
Increasing FOV	No direct effect
Using rectangular FOV (in phase-encoding direction)	Decreases
Increasing TR	Increases
Increasing TE	Increases
Increasing matrix size in frequency-encoding direction	Increases
Partial Fourier imaging	Decreases
Fractional echo imaging	Decreases
Increasing NEX	Increases

References

1. Elster AD, Burdette JH (2001) Questions and answers in magnetic resonance imaging, 2nd ed. Mosby, St. Louis
2. Mitchell DG, Cohen MS (2004) MRI principles, 2nd ed. Saunders, Philadelphia
3. Hendrick RE (1999) Image contrast and noise. In: Stark DD, Bradley WG Jr (eds) Magnetic resonance imaging, 3rd ed. Mosby-Year Book no 43. Mosby, St. Louis