The Physics of Magnetic Resonance Imaging (MRI) Lecture 2

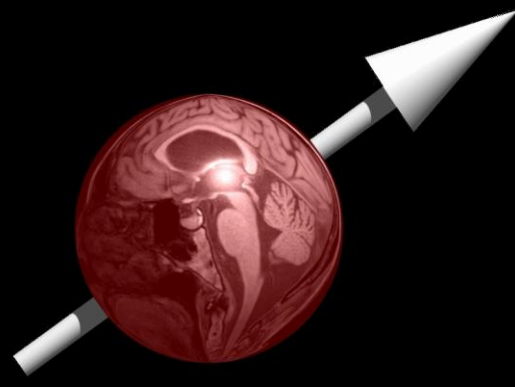

Martin Graves
mjg40@cam.ac.uk

mri-physics.com

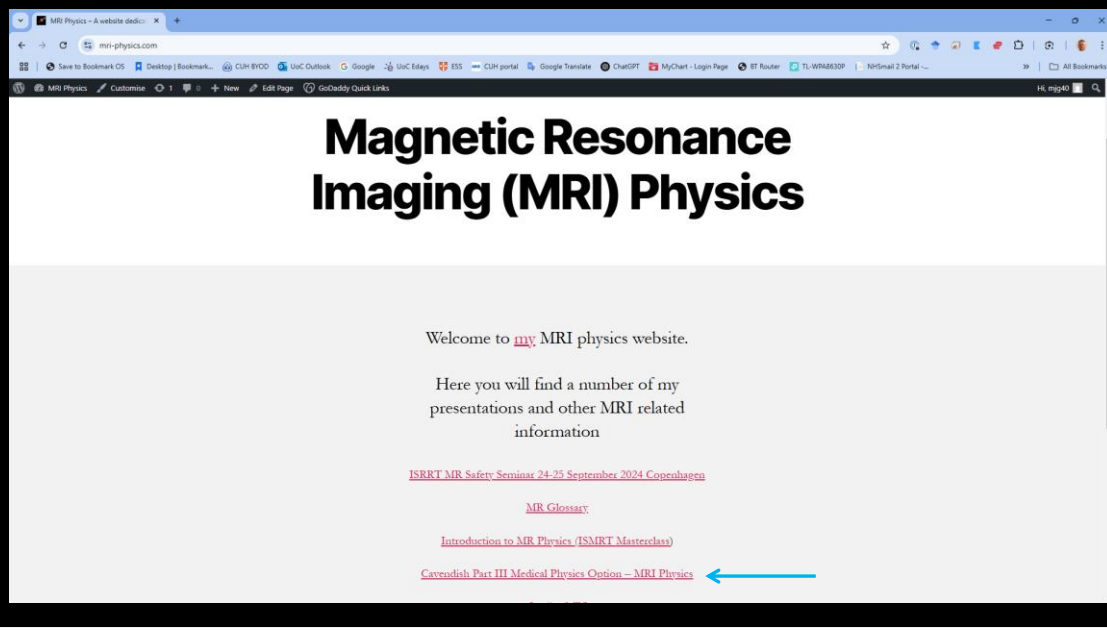

Learning Outcomes

- After these lectures you should be able to:
  - Explain how nuclear spin gives rise to magnetic resonance
  - Understand the principles of T_1, T_2 and T_2^* relaxation
  - Explain the principles of MR image formation
  - Describe the spin echo and gradient echo pulse sequences
  - Outline the basic components of an MRI system
  - Understand the safety issues related to MRI

Nuclear Magnetic Resonance Imaging

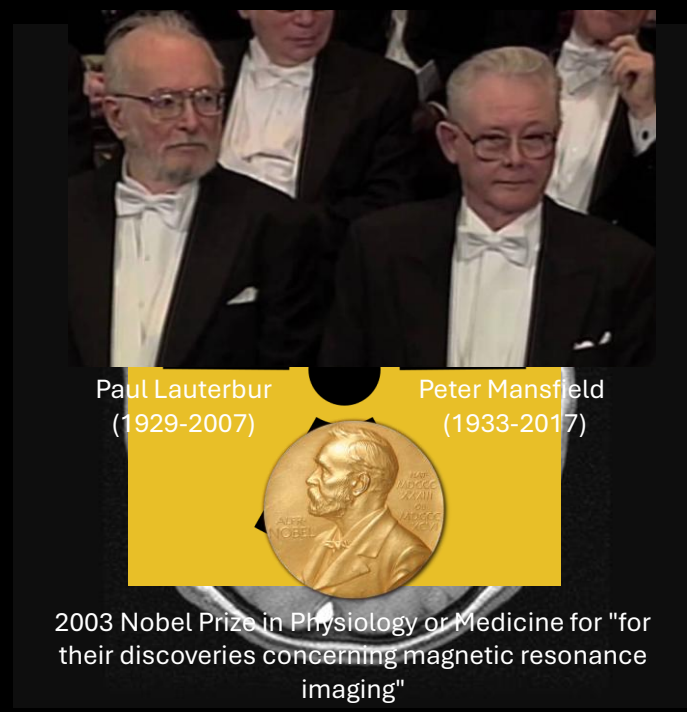

2003 Nobel Prize in Physiology or Medicine for "for their discoveries concerning magnetic resonance imaging"

MRI was original called by the more accurately termed nuclear magnetic resonance imaging but with the negative connotations of the word nuclear, although MRI has nothing to do with radioactivity, that was soon dropped. The Nobel Prize in Physiology or Medicine was awarded in 2003 jointly between Paul Lauterbur of the State University of New York and Professor Sir Peter Mansfield from the University of Nottingham for their discoveries concerning magnetic resonance imaging.

Trapezoidal Gradient Pulses

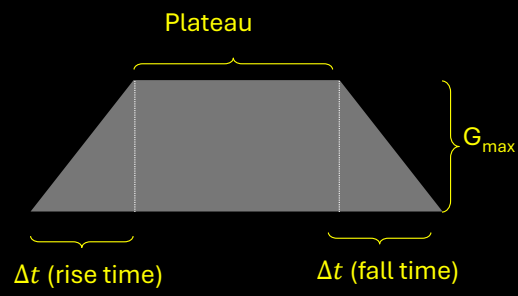

Slew Rate (SR) = $\frac{G_{max}}{\Delta t} = \frac{50 mT/m}{250 \mu s} = 200\ T/m/s$

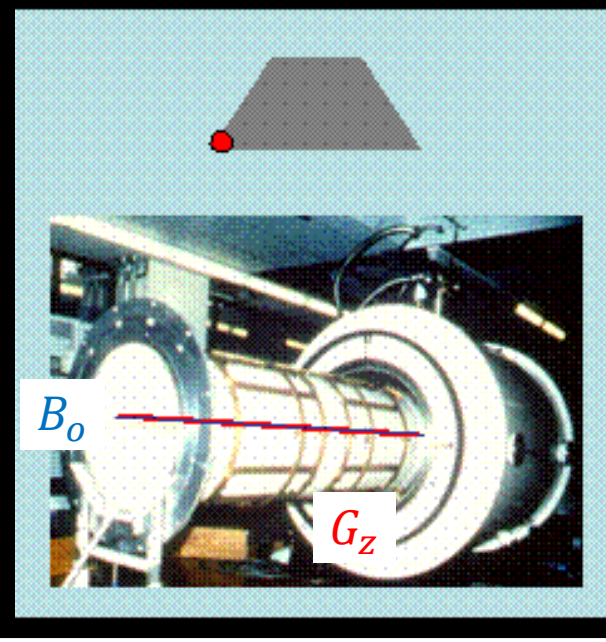

The crucial invention as far as imaging is concerned was made by Lauterbur, who apparently had the idea whilst eating a hamburger. His great idea was to introduce the concept of a magnetic field gradient for spatial localisation. He called his technique "zeugmatography" from the Greek ζεῦγμα "to join". Fortunately, that name did not catch on! The magnetic field gradients are superimposed on B_0. Gradient pulses are typically drawn as trapezoidal shapes with respect to time. We start with the gradient taking a finite time to reach the desired amplitude which is then held for a duration known as the plateau before another period of time for the gradient to switch off. The animation shows a very early gradient coil assembly that has been removed from inside the magnet. I have superimposed a gradient in the z-direction being switch on, held for on a brief period of time and then switched off. The fulcrum of the gradient is in the middle of the gradient coil assembly and that would normally be at the isocentre of the magnet. One metrics used to describe gradient performance is the slew rate which is the maximum possible gradient amplitude G_{max} divided by the minimum time it takes to reach that amplitude. For example, if the maximum gradient amplitude is 50 mT/m and it takes 250 μs to reach that amplitude then gradient will be said to have a slew rate of 200 T/m/s

Imaging & Gradients

- The key to MR imaging is the concept of a magnetic field gradient superimposed on the main static magnetic field

$$G_x = \frac{\partial B_z}{\partial x}, G_y = \frac{\partial B_z}{\partial y}, G_z = \frac{\partial B_z}{\partial z}$$

- The gradient creates a spatially dependent variation of the Larmor frequency, e.g., in the z-direction

$$B(z) = B_0 + G_z z$$

- Gradient amplitudes are typically measured in mT/m or G/cm (10mT/m = 1G/cm)

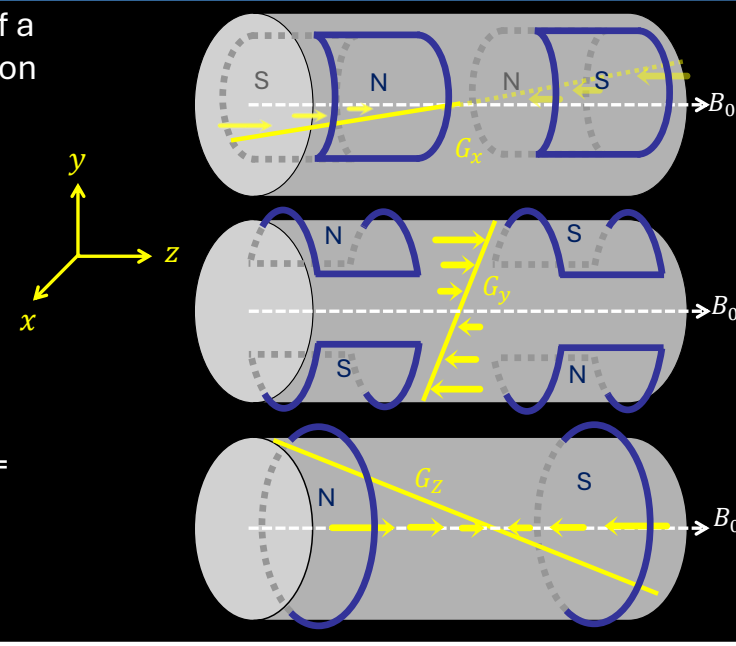

The important concept regarding gradients in MRI is that for MRI to still work the spatial variations in X, Y and Z also require the gradient magnetic field to be always pointing along the Z axis. In the case of the X gradient we have the gradient creating coils which are arranged in what is called a quadripolar or Golay arrangement. The currents are circulating through this gradient coil to make this a north pole and this a south pole. We will then have a strong magnetic field from north to south nearest the surface of the gradient coil assembly and decreasing linearly towards the centre of the coil arrangement. This magnetic field will add to B_0. The gradient will be zero at the isocentre and then because of the opposite polarity arrangement on the other side of the assembly the gradient amplitude will again increase linearly but in the opposite direction, subtracting from B_0. In the Y direction we have the same quadripolar arrangement but just rotated 90° so now the gradient is going from the top to the bottom of the gradient coil. The z-gradient is a little simpler. We have a Maxwell pair in which there are counter-rotating currents. The effect of each gradient is to create a spatially dependent variation of the Larmor frequency. For example, in the z-direction there is a continuous linear decrease in Larmor frequency from the north pole of the gradient to the isocentre, where the Larmor frequency is just given by the static magnetic field, and then a further decrease in Larmor frequency towards the south pole as the gradient is in the opposite direction to B_0.

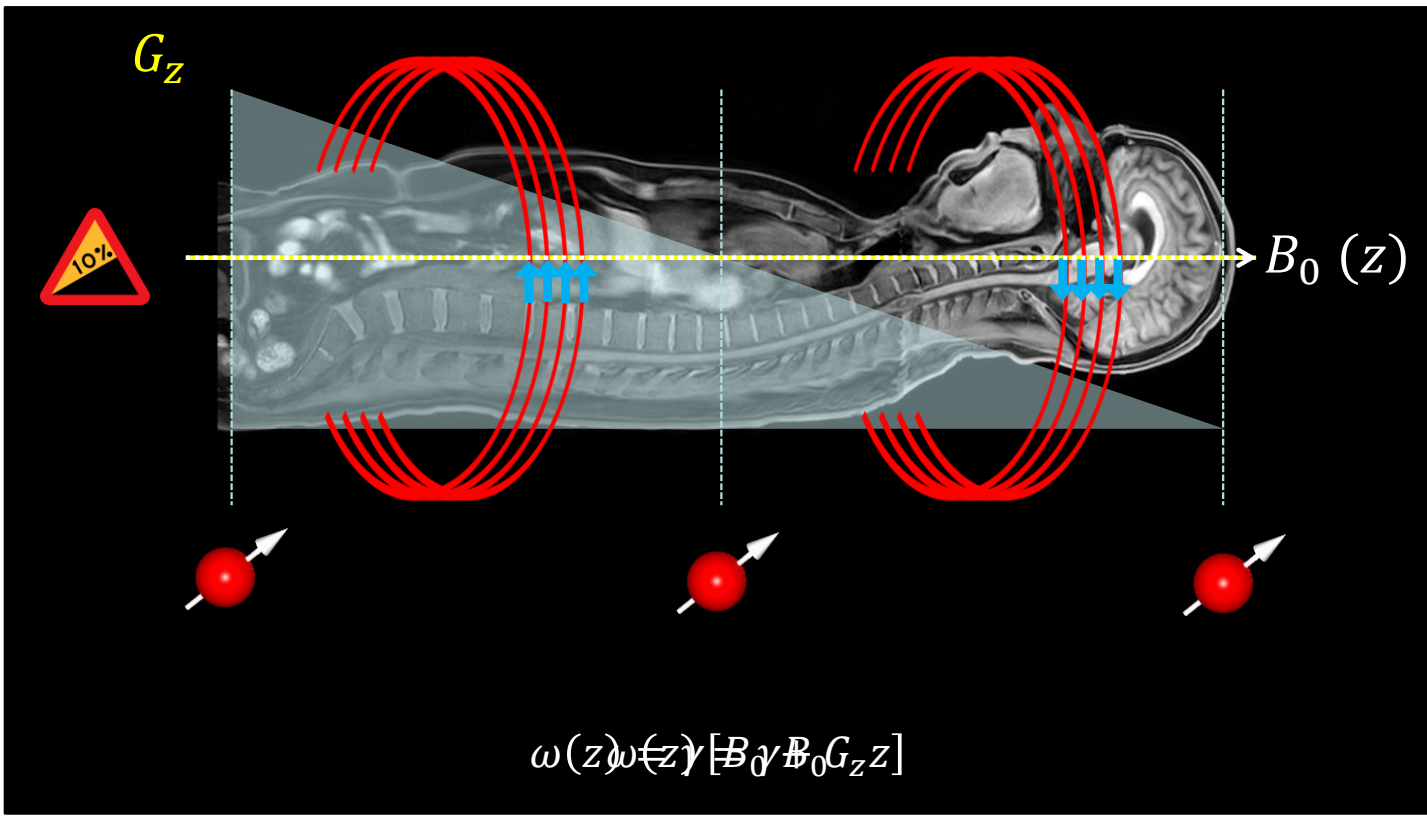

Here is our subject lying inside the magnet with the main magnetic field pointing from their feet to their head. if we look at the protons precessing at three different locations ignoring any non-uniformity in the static magnetic field, we can see that at all three z- locations the protons are all precessing at the same frequency. If we now introduce a z gradient using our Maxwell pair gradient coils, we can see that the proton in the centre is still precessing at gamma times B_0 but the proton to the left is precessing slightly slower and the proton on the right is precessing faster due to the presence of the gradient

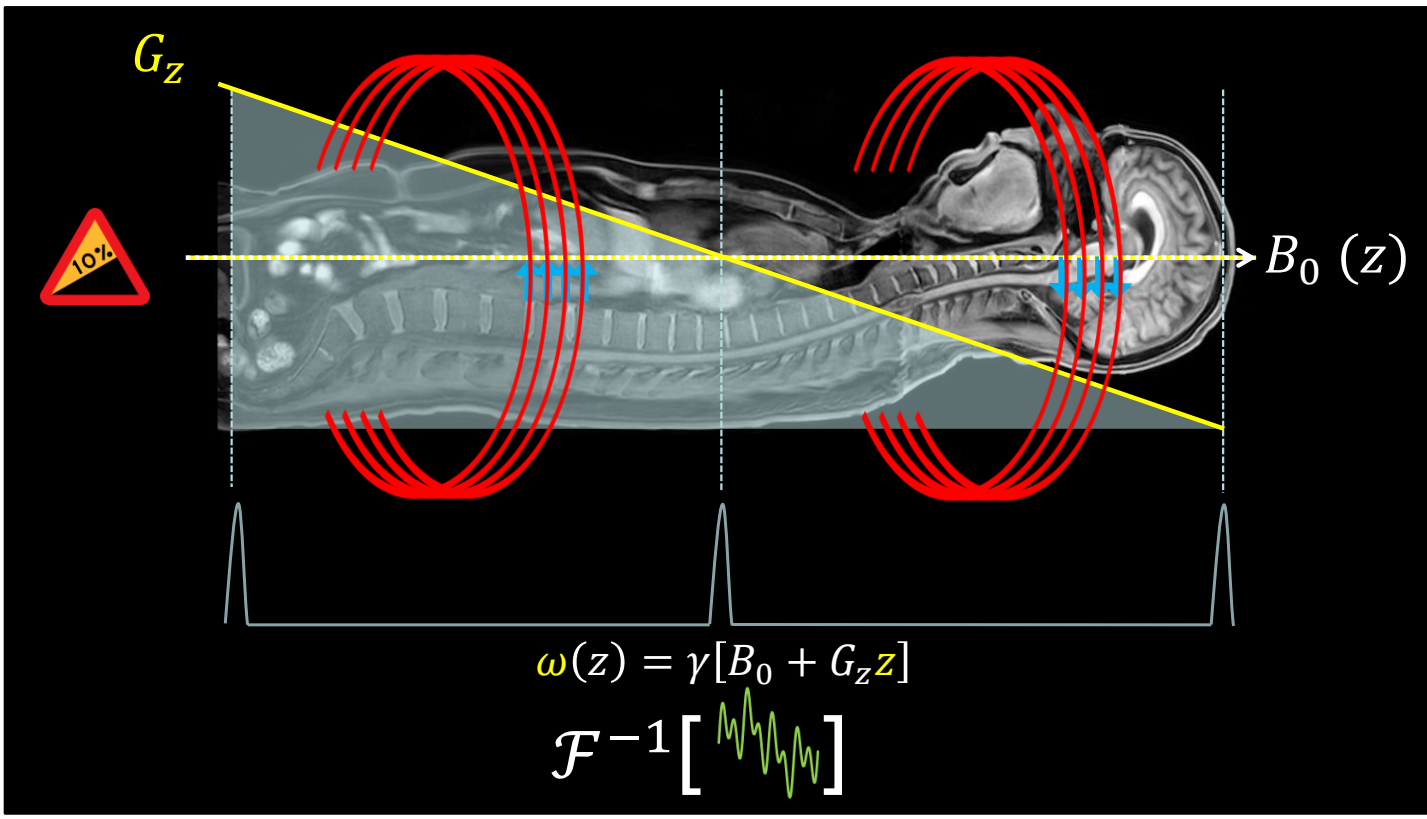

Here is an audio frequency analogy to try and make this a bit clearer. On the left the proton with the highest precessional frequency can be represented as a high frequency note, the proton in the middle has a slightly lower frequency and then the proton on the right has the lowest frequency. However, the signal that we detect from the patient will be a mix of all three frequencies, so we have no positional information. The question is how can we deconstruct this mix of frequencies into its individual components? The answer is, of course, with a Fourier transform. If we Fourier transform the combined signal, we will get three resonant frequencies and if we know the strength of the gradient accurately then we can translate those frequencies to positions. We will have a continuous distribution of frequencies and the Fourier transform will give us a 1- dimensional projection through the patient.

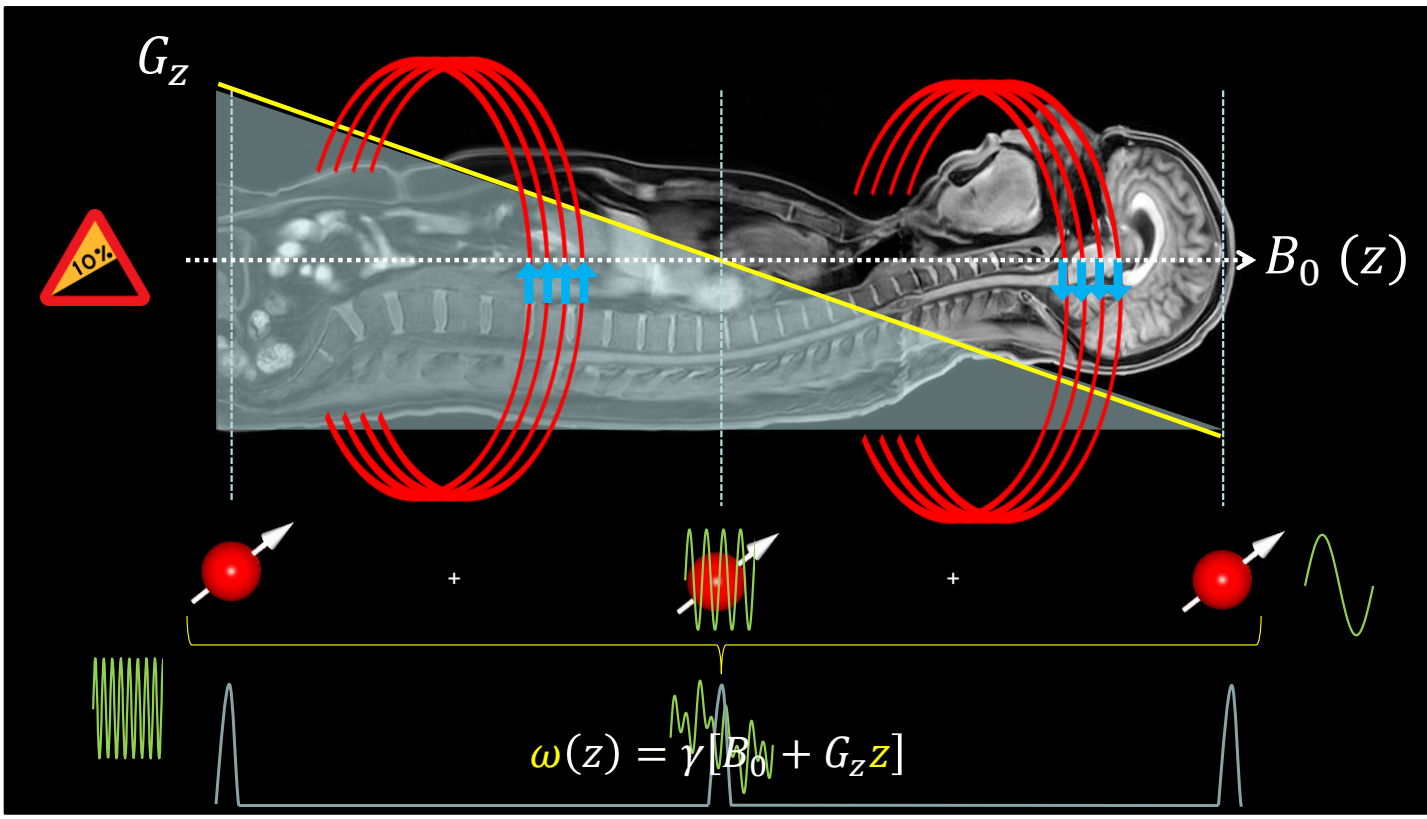

Here is our subject lying inside the magnet with the main magnetic field pointing from their feet to their head. At these three positions the main magnetic field is constant, and the protons are all precessing at the same frequency. If we add yet more coils inside the MRI scanner then by passing current through these coils we can create a gradient, a slope in the magnetic field in addition to the main magnetic field. It is the gradients switching on and off that causes the characteristic knocking noise when imaging. These gradient coils physically flex, not by a lot, but enough to create the acoustic noise. You can see that when the gradient is applied there is now a difference in the precessional frequencies because the protons are now in slightly different magnetic fields due to the gradient.

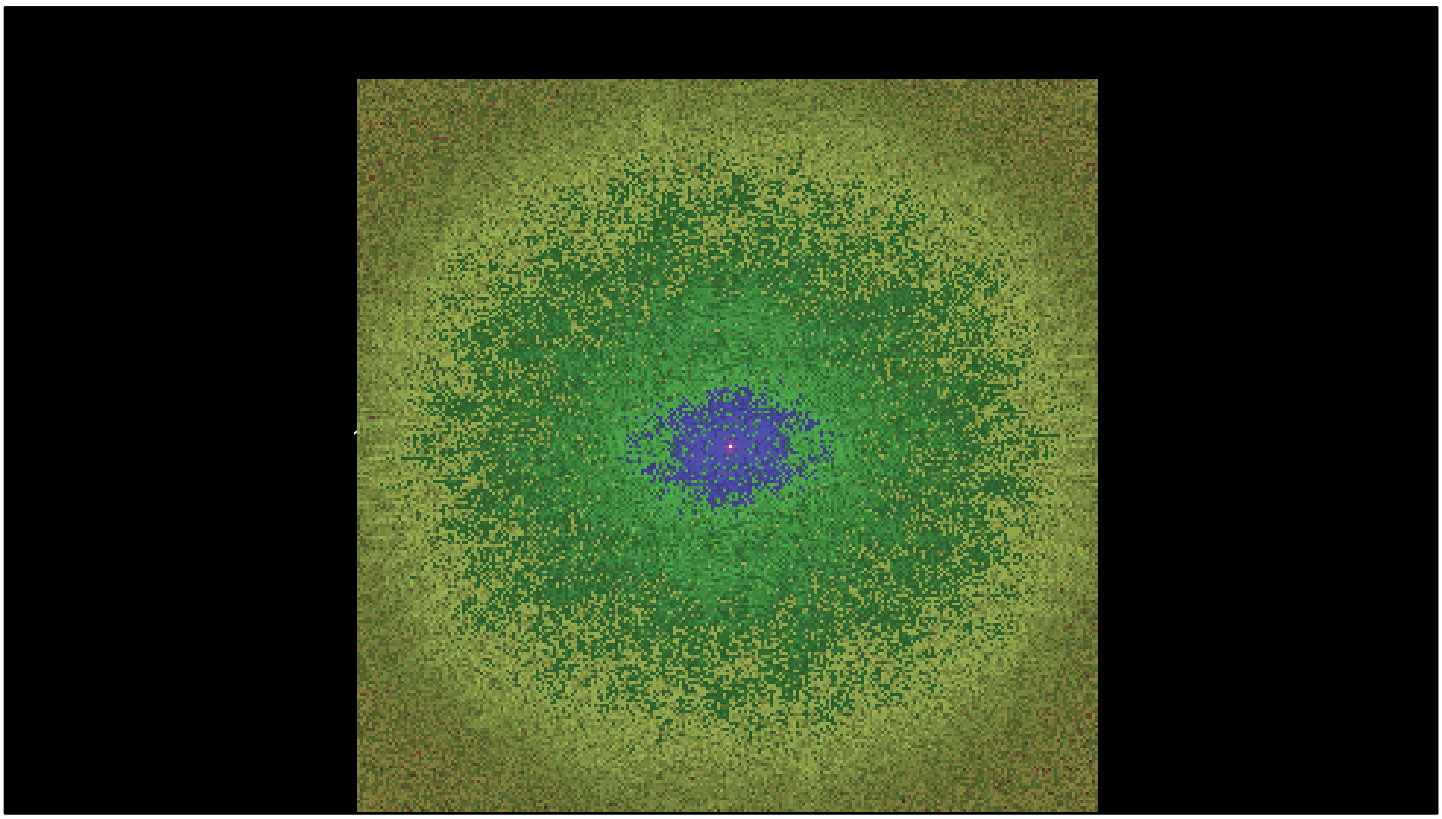

This is a typical echo signal that you would receive from the patient. If we show this as line of raw data in a false, colour it will look like this. If we acquire all the necessary gradient steps to obtain an image, then the full array of raw data will look like this. This raw data matrix is often called *k*-space for reasons we will discuss in a moment.

Slices

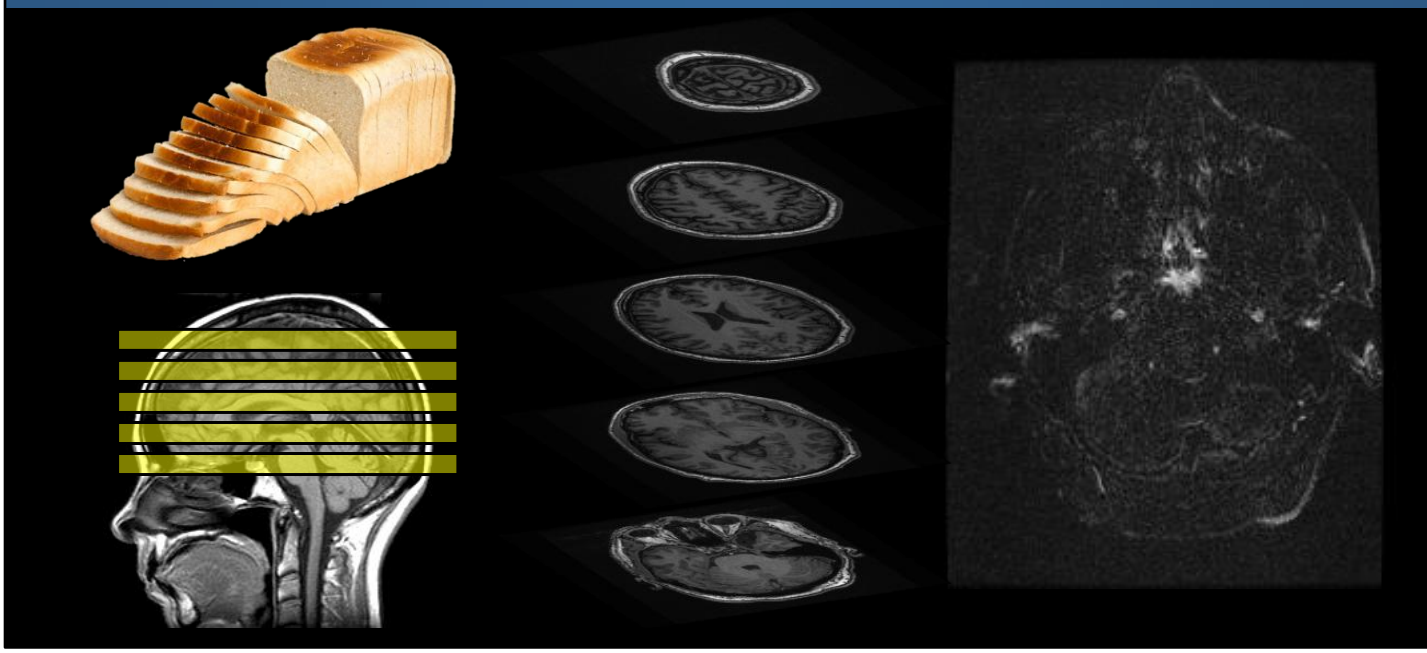

We know that MRI is the greatest invention since slice bread and to maintain the analogy we usually acquire multiple slices through the area of interest. A brain MRI would typically involve the acquisition or probably around thirty 5 mm thick slices.

Spatial Localisation

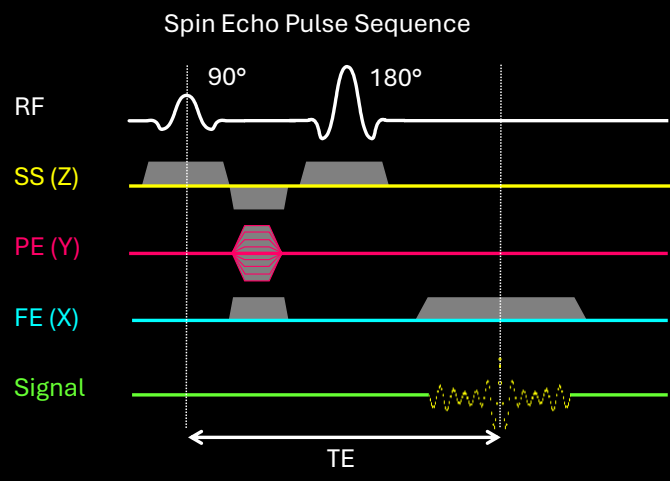

Depending upon the imaging plane the logical gradients (SS, PE, FE) are played out on the relevant physical gradients (X, Y, Z)

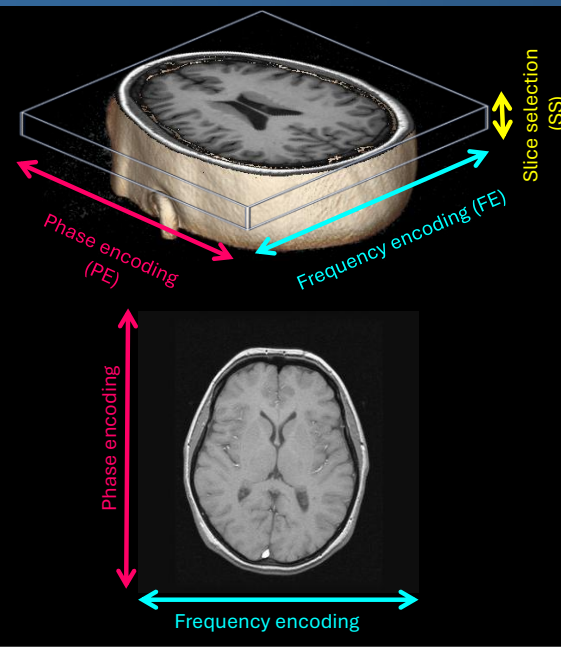

The acquisition of an MR image involves a combination of RF pulses to tip the net magnetization, various time delays to allow the desired contrast to develop, for example T_1 or T_2 weighting, and gradient pulses to spatially encode the MR signal. This combination of RF and gradients is known as a is known as a pulse sequence. This diagram illustrates the standard method for MRI spatial encoding invented at the University of Aberdeen back in the early 1980s. This method calls the two in-plane directions frequency encoding and phase encoding. An MRI system has three orthogonal physical gradients in the X, Y and Z directions, but we will refer to them by their logical names of slice selection, phase encoding and frequency encoding.

Gradient Effects

Gradient pulses induce a change in Larmor frequency with position. They also create a phase accumulation with time

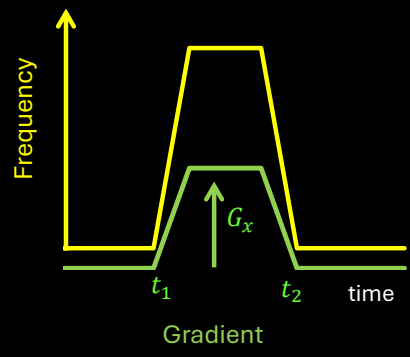

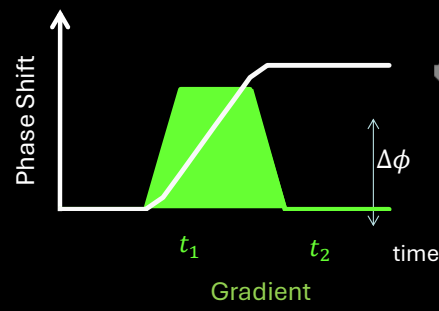

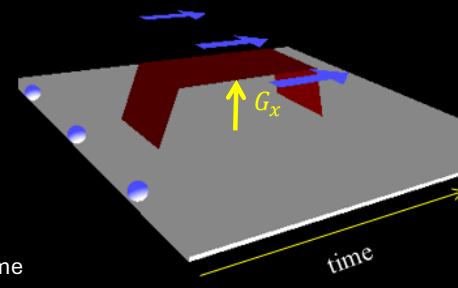

$$\Delta\omega(x) = \gamma\, G_x\, x$$

$$\Delta\phi = \int_{t_1}^{t_2} \Delta\omega(x)\, dt$$

$$\Delta\phi = \gamma \int_{t_1}^{t_2} G_x\, x\, dt$$

A magnetic field gradient will not only cause a transient change in the Larmor frequency but will also induce a relative phase shift in the transverse magnetization proportional to the strength and duration, i.e., the area of the gradient pulse. The diagram shows three spin vectors that will initially all be precessing in phase. I will show the animation in the next slide, but it is important to note how the middle spin will precess faster in the presence of the magnetic field gradient whereas the other two will not. However, at the end of the gradient pulse all three spins will precess at the same frequency but note that the middle spin will be 180° out of phase compared to the other two.

Gradient Effects

Gradient pulses induce a change in Larmor frequency with position. They also create a phase accumulation with time

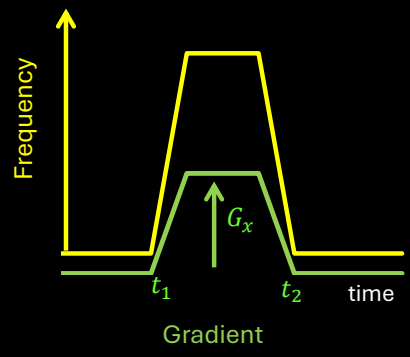

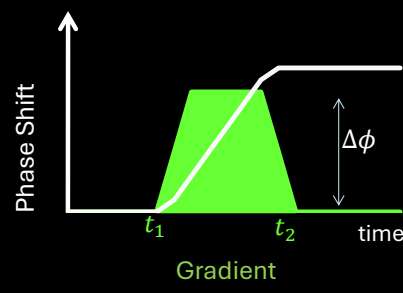

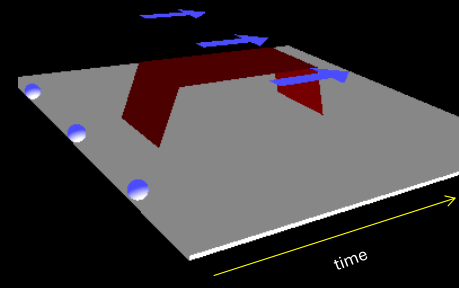

$$\Delta\omega(x) = \gamma\, G_x\, x$$

$$\Delta\phi = \int_{t_1}^{t_2} \Delta\omega(x)\, dt$$

$$\Delta\phi = \gamma \int_{t_1}^{t_2} G_x\, x\, dt$$

Time is represented by the rolling balls. As the middle ball rolls up the gradient its vector will rotate faster. When the ball comes down the other side all three vectors will continue to rotate at the Larmor frequency, but the middle vector will have gained a 180° phase shift with respect to the other two.

Where does '*k*' come from?

- We know that a gradient pulse, e.g., along x, induces a phase shift $\Delta\phi_x$

$$\Delta\phi_x = \gamma \int G_x \, x \, dt$$

- It is usual and convenient to re-write this in terms of a "spatial" frequency ***k*** with units of $radians/cm$, i.e., a frequency that changes with position

$$k_x = \gamma \int G_x \, dt$$

- In the case of gradients with constant amplitude, i.e., trapezoids, we have $k_x = \gamma \cdot G_x \cdot t$ so that $\Delta\phi_x = k_x \cdot x$
- Of which more later …

* Note that for ^{1}H $\gamma = 2.657\text{x}10^8 \; rad/s/\,T$ or $\frac{\gamma}{2\pi} = \text{ɣ} = 42.577 \; MHz/T$

Where does *k* come from when we start talking about *k*-space. We know from our previous discussions that a gradient pulse, for example along the X axis, will induce a phase shift that is equal to the time integral of this gradient of multiplied by its position. We can rewrite this equation in terms of a spatial frequency k in units of rad/cm, assuming gamma is in units of $\text{rad}/\text{s}/\text{T}$. *k* therefore represents a frequency that changes with position. In the case of trapezoidal gradients with constant amplitudes k_x is equal to gamma, times Gx times the duration of the plateau and that the phase shift is simply k_x times the x -position. We will make use of *k* a bit later.

Slice Selection

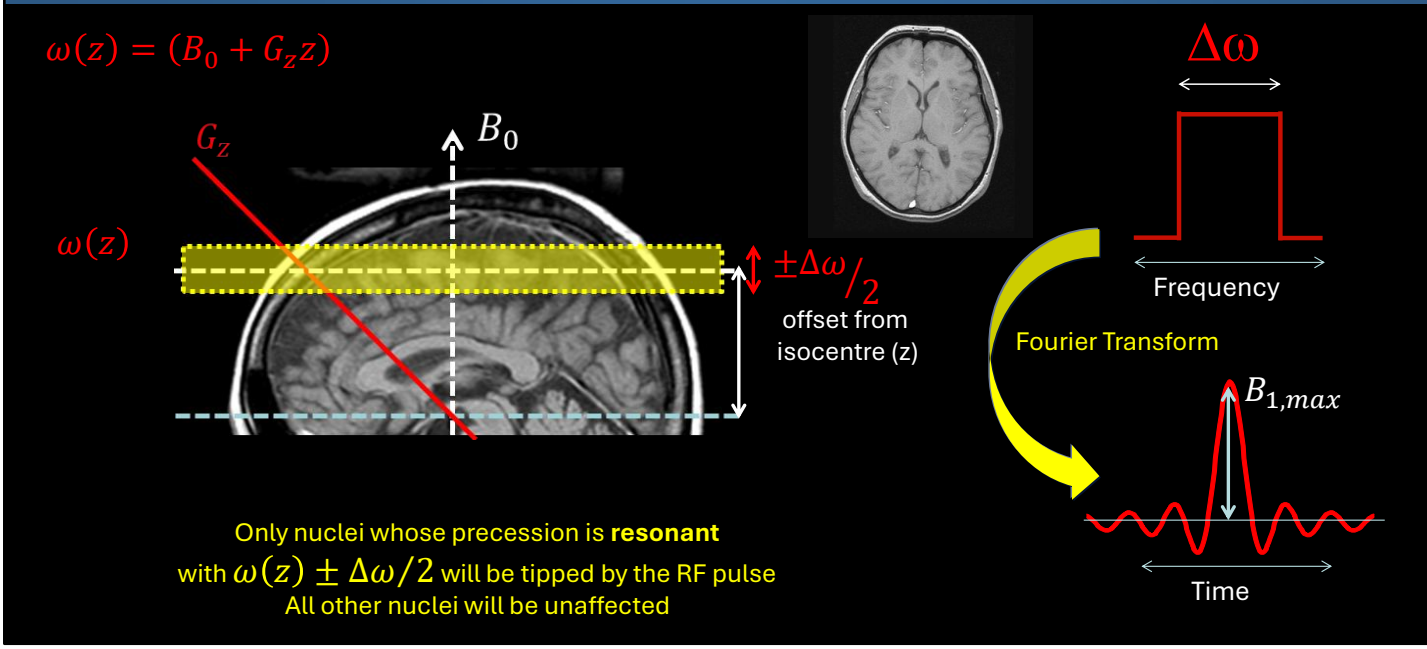

We will start our discussion of spatial localisation in MRI with the principles of slice selection. To acquire an axial slice through the brain we will use the z gradient as the slice select gradient. The frequency at the centre of the slice is given by $\omega(z)$ where z is the distance from isocentre. Because we need to excite a slice of a finite thickness, we need to excite a range of frequencies about this centre frequency. The slice profile in frequency space will look like a top hat function. However, we need to create our slice excitation in the time domain so the conversion from frequency to time is done via the Fourier transformation and the Fourier transform of a top hat function is a sinc function. A windowed sinc pulse (since the sinc is effectively infinite in duration) is then mixed with the $\omega(z)$ carrier frequency and played out in the presence of the slice select gradient Therefore, only nuclei whose precession is resonant with $\omega(z) \pm \Delta\omega/2$ will be tipped, and all other nuclei outside the slice will be unaffected

RF Excitation

$\pm\Delta\omega$

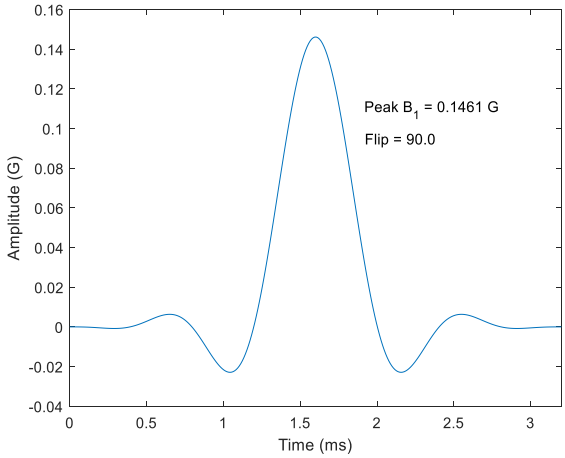

$$\alpha = \gamma \int_0^{t_{RF}} B_1 dt$$

$\omega(z) \pm \Delta\omega$

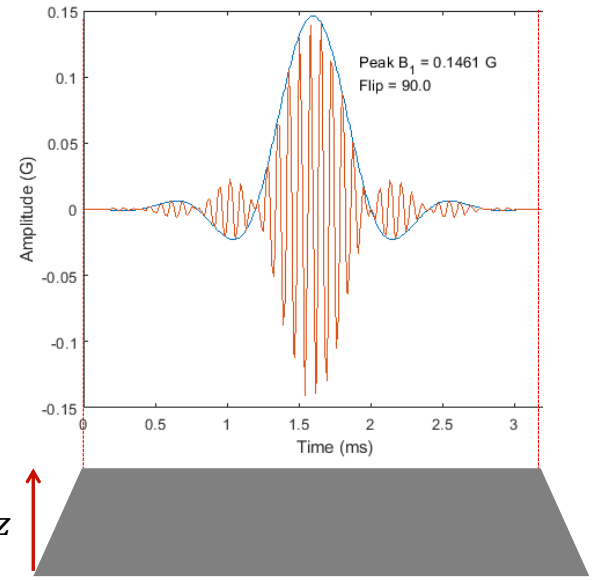

This shows the RF excitation envelope which dictates the slice thickness mixed with the carrier frequency which dictates the slice position. This RF pulse is played out in the presence of the Gz slice select gradient to excite the spins at the desired slice location and thickness. The flip angle of the pulse is given by the time integral of the RF envelope.

RF Shape and Slice Profile

3.2ms Hann windowed SINC (TBW = 8)

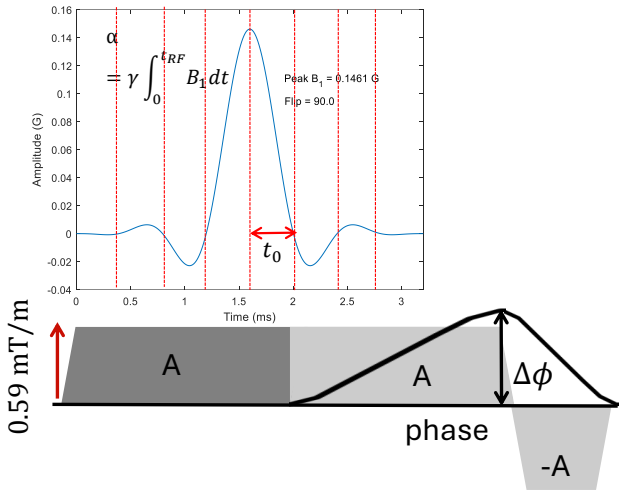

Time-Bandwidth Product (TBW) = no. of zero crossings

$$T\Delta f = N_L + N_R$$

$$\Delta f = \frac{N_L + N_R}{T} = \frac{8}{3.2\ \text{ms}}$$

$$\Delta f = \frac{1}{t_0} = 2.5\ \text{kHz}$$

Slice Profile

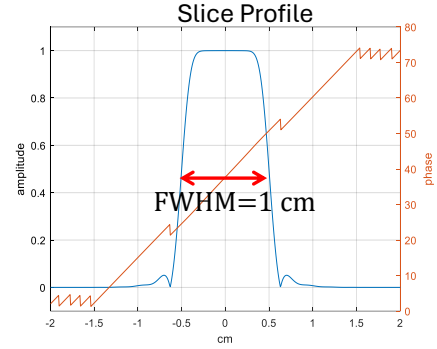

$$G_z = \frac{\Delta f}{\gamma\, \Delta z}$$

The slice select gradient induces a phase across the slice $\Delta\phi$. This is reversed by a subsequent gradient of equal area but opposite polarity

$$G_z = \frac{2.5\ \text{kHz}}{42.57\ \text{MHz/T} \cdot 1\ \text{cm}} = 5.87\ \text{mT/m}$$

Let us look in a little bit more detail at the RF pulse shape and its slice profiles. Since a sinc pulse is effectively infinite in duration it is usually windowed using an appropriate filter to truncate it to an acceptable duration, in this case we use a 3.2 ms Hann windowed sinc. Sinc pulses are designed in terms of their time bandwidth product, which is equal to the number of zero crossings over the total duration of the pulse. In this case we have eight zero crossings. The excitation bandwidth is therefore given by the number of zero crossings divided by the pulse duration which in this case is 2.5 kHz. If we wish to excite a 10 mm thick slice, then we would need to play out a slice select gradient of amplitude 5.9 mT/m. Using the Bloch equations, we can then derive the slice profile for this pulse as you can see on the right. We get a relatively good profile, although the edges are not sharp because we truncated the pulse. The full width half maximum of the profile is equal to the desired slice thickness. As we have previously discussed any gradient induces a phase and in the same way the slice select gradient will induce a phase in the spins across the slice thickness. The start of the pulse is taken as its magnetic centre and therefore to align the spins across the slice profile, we need to play out an additional gradient after the pulse of equal area but opposite polarity.

In-plane Spatial Localisation

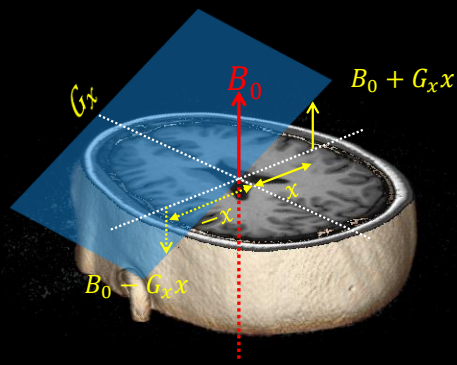

- The B field at a given location x is:
- $B(x) = B_0 + G_x x$
- $\omega(x) = \gamma[B_0 + G_x x]$
- In a rotating frame precessing at ω_0
- $\omega(x) = \gamma G_x x$

- The solution of the Bloch equation in complex notation (ignoring T_2) is
- $M(x,t) = M_0(x)e^{-i\omega(x)t} = M_0(x)e^{-i\gamma G_x x t}$
- The spin density at location x is $\rho(x)$, in which case the signal S between x and dx is given by
- $S(x) = \int_x^{x+dx} \rho(x)e^{-i\gamma G_x x t} dx$

i.e., the signal is the Fourier transform of the spin density

- For a continuous 2D distribution $\rho(x,y)$
- $S(x,y) = \iint \rho(x,y)e^{-i\gamma G_x x t} e^{-i\gamma G_y y t}\, dx\, dy$

Now we have successfully nutated the spins in our desired slice we need to spatially encode the NMR signal in plane. As we have previously seen we do this using magnetic field gradient's and if we express the solution of the Bloch equations for the transverse magnetisation in complex notation and ignoring T_2 relaxation we have an expression for the signal in the x-direction. If instead of the equilibrium magnetisation M_0 we use the proton density at a particular location, x, then the signal is the integral over dx. which is the 1D Fourier transform of the spin density. We can then easily extend this to a continuous two-dimensional distribution of spin density. The proton density distribution can therefore be obtained by a 2D inverse Fourier transformation of the acquired signal data or *k*- space.

1D Projection Imaging

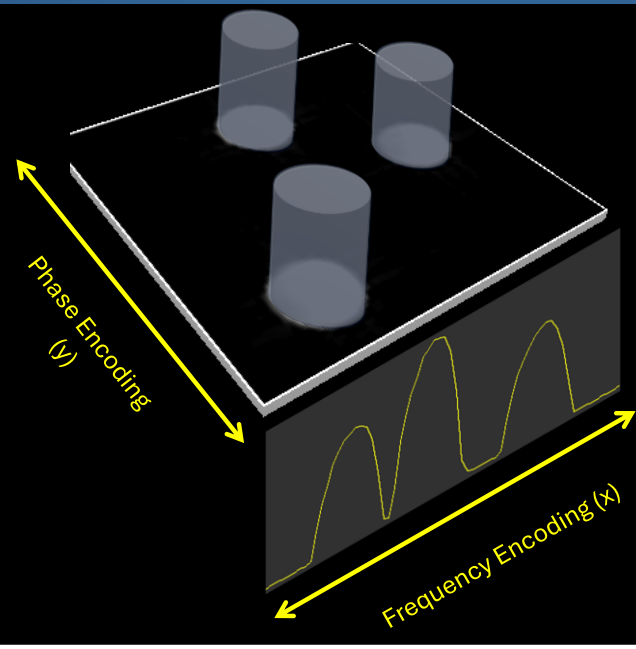

We shall now look at the standard method for MRI in-plane spatial encoding that was called “spin-warp” by the Star Trek inspired inventors at the University of Aberdeen back in the early 1980s. We shall see later the relevance of the “warp”, sadly entirely unrelated to speed. This method names the two in-plane directions frequency encoding and phase encoding. We shall start with frequency encoding and will use a simple arrangement of three tubes of water with slightly different relaxation times to slightly change the signal intensity. Hopefully, you can agree that the signal in the frequency encoding direction represents a one-dimensional projection of the tubes.

Frequency Encoding

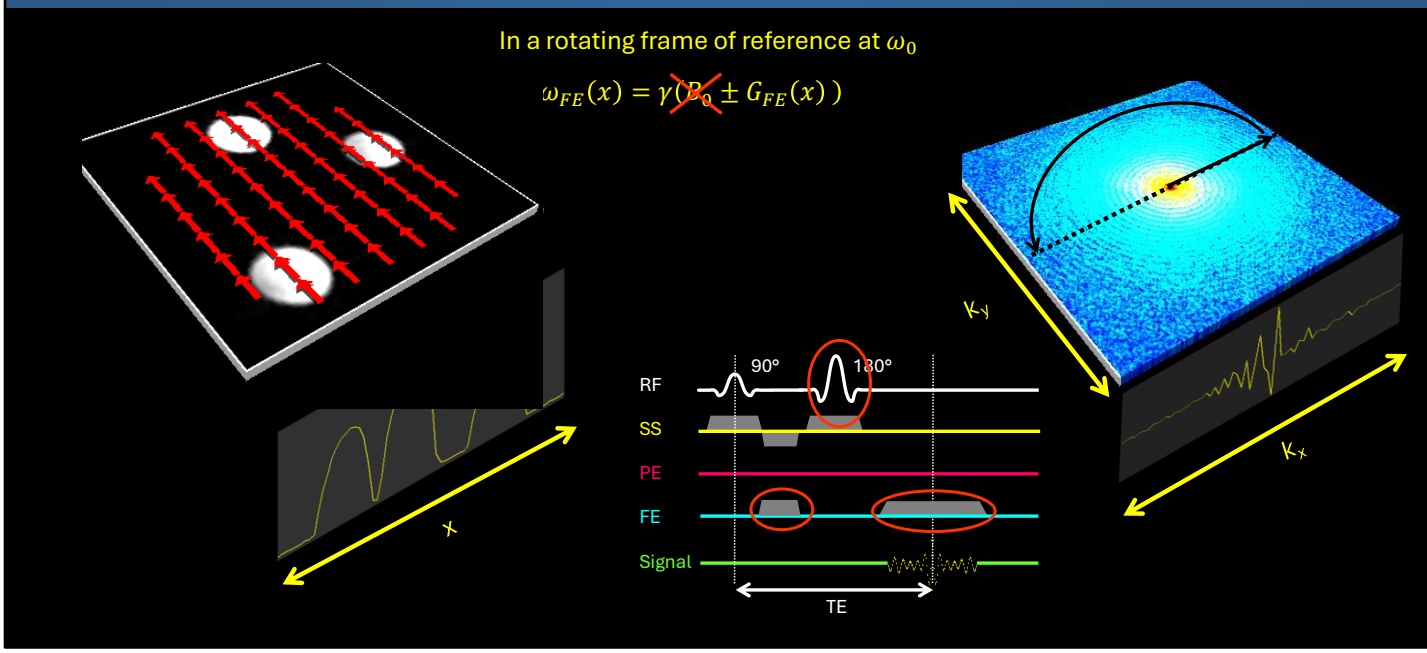

Following slice selection all the spins within the slice will have been tipped into the transverse plane by the RF pulse so will all be precessing at the Larmor frequency. If we work in a frame of reference also rotating at the Larmor frequency, then we can represent the spins at each sample point as a vector that appears stationary. Please ignore the fact that the only signal producing areas would be the bottles and accept that the vectors all have the same length, i.e., we will ignore relaxation effects. Following the excitation pulse the vectors will all have zero phase difference, and we can consider them to be in at the centre of *k*-space. If we first apply the positive polarity frequency prephasing gradient circled the phase of the spins will move to the far-right hand side edge of *k*-space. The 180° pulse will then flip the phases, so the spins are now positioned at the far-left hand side. We will now turn on the frequency encoding gradient represented by the blue plane at the isocentre the vectors still appear stationary since we are in the rotating frame of reference but as we move away from isocentre towards the right the frequency of precession increases in a positive sense due to the increasing strength of the gradient. Similarly, if we move away from isocentre towards the left the frequency of precession increases in a negative sense due to the negative gradient increasing in amplitude. Remember that in the laboratory frame this would just mean that the vectors would be precessing slightly slower than the Larmor frequency. We arrange the gradient timings so that the frequency encoding gradient occurs at the same time as the spin echo is formed and digitised. This digitised echo, after demodulation to remove the Larmor frequency carrier, would then be stored as a line in the centre of *k*-space. The Fourier transform of this line would be the 1D projection of the bottles shown on the left. This is often known as the central slice theorem; however, we do not have any information as to the position of the tubes in the other in-plane direction, which is where phase encoding comes into the process.

Frequency Encoding Gradient Amplitude

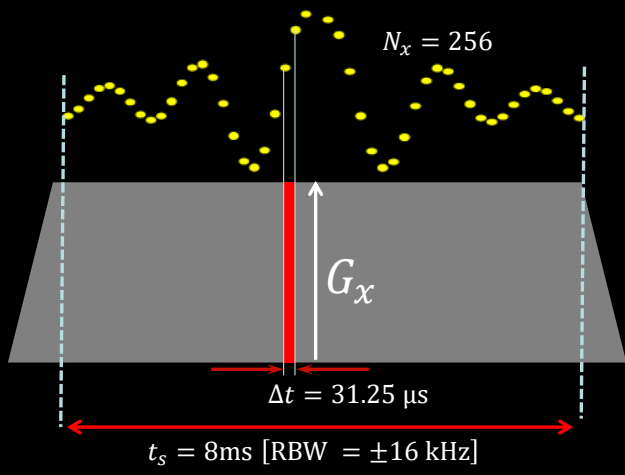

$$\Delta k_x = \gamma G_x \Delta t$$

$$FOV_x = \frac{1}{\Delta k_x}$$

$$t_s = N_x \Delta t$$

$$\Delta k_x = \gamma G_x \frac{t_s}{N_x}$$

$$Receiver\ bandwidth\ (RBW) = \frac{1}{\Delta t} = \frac{N_x}{t_s}$$

$$G_x = \frac{RBW}{\gamma\ FOV_x}$$

Note that the range of frequencies encoded is small cf. the Larmor frequency

$$G_X = \frac{32\ \text{kHz}}{4257\ \text{MHz/T} \cdot 24\ \text{cm}} = 3.13\ \text{mT/m}$$

Just before we talk about phase encoding here is how we calculate the amplitude and duration of the frequency encoding gradient. Each sample in the frequency encoding gradient represents an incremental change in k_x. The operator provides the desired number of samples N_x, and the desired receiver bandwidth and then the pulse sequence calculates the required gradient amplitude and duration (t_s).

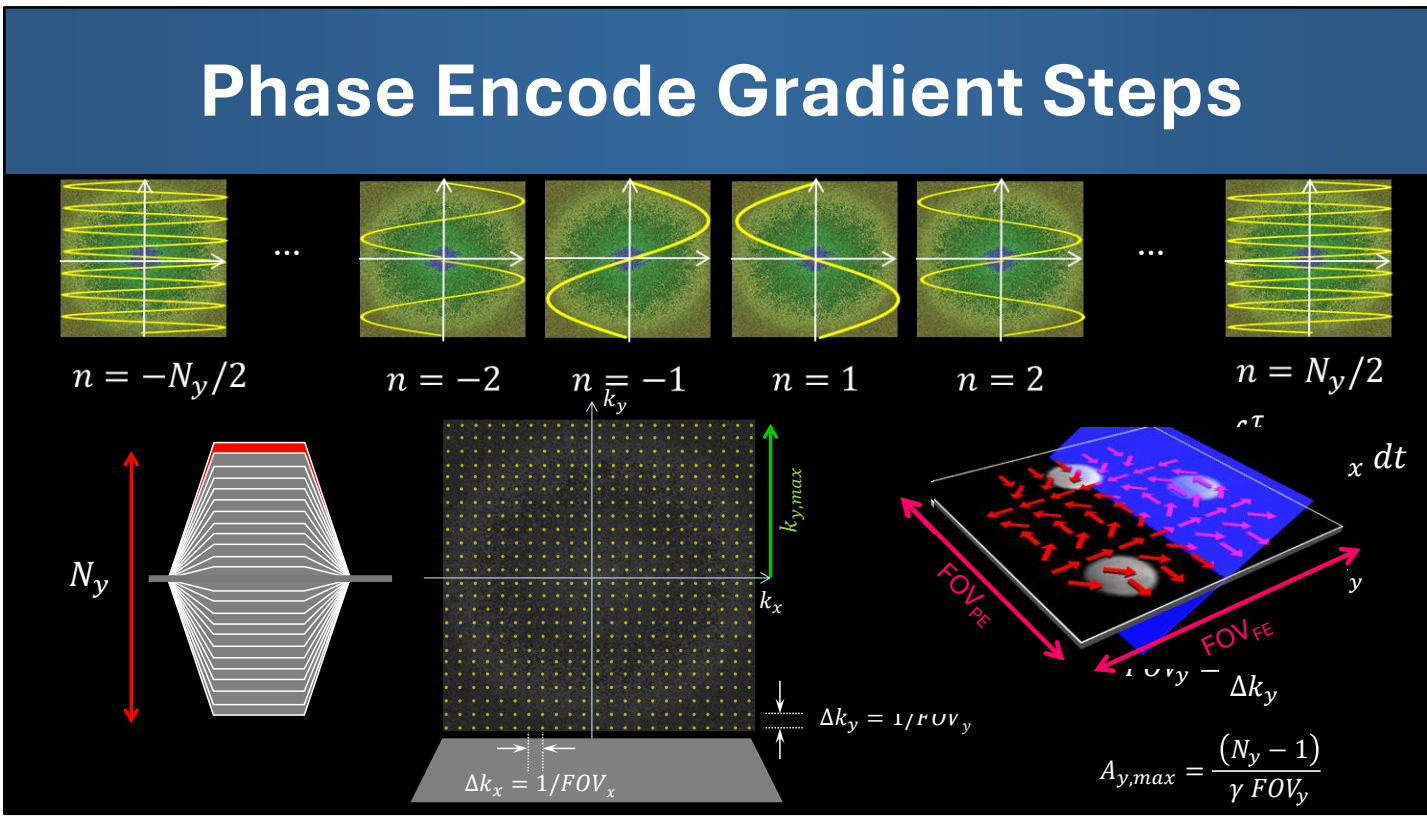

Here you can see the range of phase warps/twists that each phase encode gradient step creates. Unlike the frequency encoding gradient, the important parameter with the phase encoding gradient is the area.

The first animation with the maximum negative phase encoding gradient (n = -N_y/2) step shows the phase encoding gradient area (amplitude multiplied by duration) creating a π phase shift between adjacent voxels in the phase enclose direction. The phase memory is retained when the frequency encoding gradient is applied and you can see, in this example in the rotating frame of reference, that the precessional frequency increases with increasing gradient strength from isocentre and increases in the opposite sense with decreasing gradient strength in the opposite direction. Data acquisition is performed during application of the frequency encoding gradient and the samples are placed into the edge of *k*-space. The process is then repeated with a slightly smaller phase encode gradient area, although importantly the duration of the phase encoding gradient does not change, followed by frequency encoding. The acquired data is then placed in the next line of *k*-space The animation continues until we reach the n= -2 step in which case the phase encode gradient area creates a 4π phase shift across the FOV before the application of the frequency encoding gradient. the next step is n = -1 where the smallest phase encode gradient creates a 2π phase shift across the FOV before the application of the frequency encoding gradient. An n=0 phase encoding gradient is not applied, so the next step is n= 1, resulting in a 2π phase shift but with the opposite phase to n= -1. Similarly, n = 2 creates a 4π phase shift but with the opposite values compared to n = -2. The phase encoding gradient area continues to increase until the maximum positive value n = N_y/2.

The maximum area is calculated as shown here and then that area is divided by the number of phase encoding steps desired. For N_y phase encoding steps, the area covered in *k*-space is (N-1) Δk

Spin Echo *k*-space Trajectory

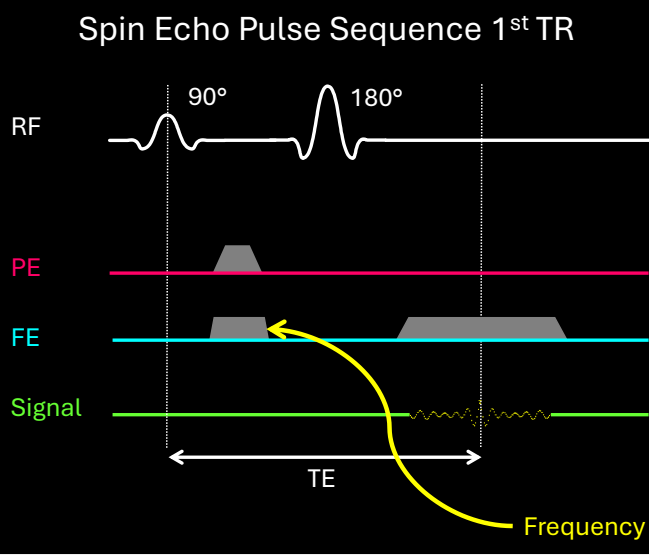

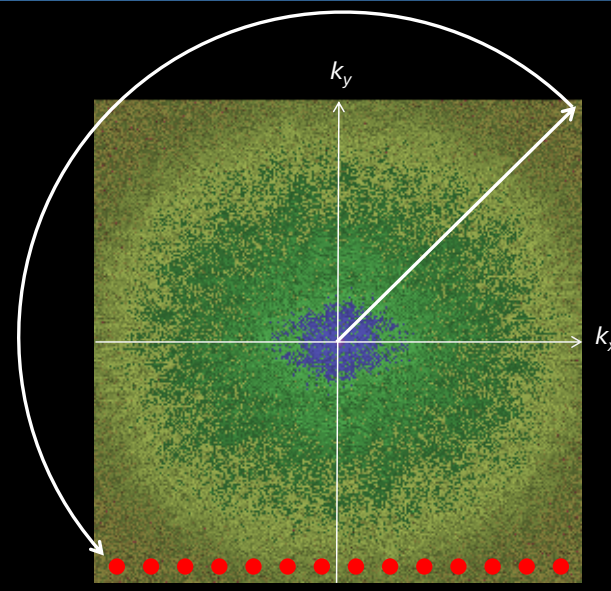

Just to make sure the relationship between the RF and gradient pulses is clear we will go through the *k*-space trajectory again for a spin echo sequence. The pulse sequence is essentially our satnav through *k*-space. The 90° excitation pulse will take us the centre of *k*-space, for the purposes of this illustration we do not need to be concerned with the slice select gradient pulses. We then simultaneously apply the maximum phasing encode gradient and the frequency pre phasing gradient This will take us up to the top right-hand corner of *k*-space. The 180° refocusing pulse will flip the spins to the bottom corner of *k*-space and we then acquire the echo in the presence of the frequency encoding gradient. This digitised echo, after demodulation, will give us our bottom line of raw data

Spin Echo *k*-space Trajectory

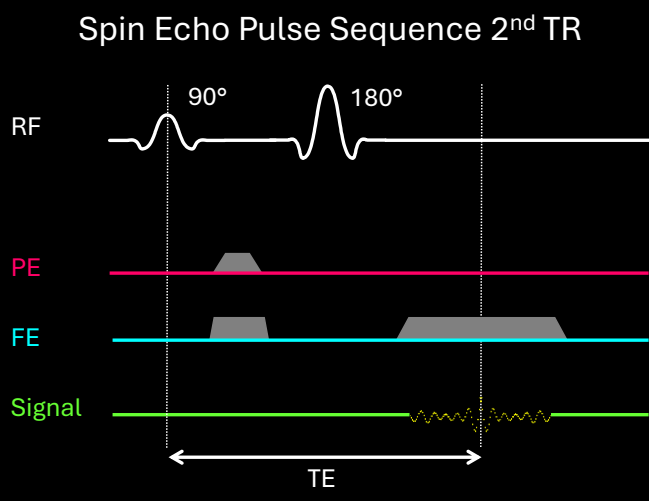

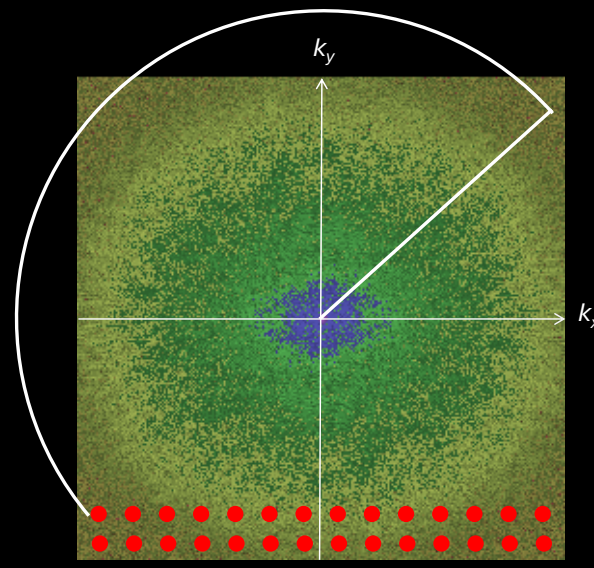

This is the next repetition, again the excitation pulse takes us the centre of *k*-space. The frequency encoding prephasing gradient is the same as the previous acquisition but now we play out the next phase encoding gradient which has a slightly smaller area. This time we do not reach a point on the right hand of *k*-space slightly below the previous acquisition. The 180° pulse then flips us to a position just above the previously acquired line of data and we read out the echo with the same frequency encoding gradient as before, giving us our second line of case space. The whole process is repeated k_y number of times so that the whole of *k*-space is filled. Once filled we can perform a 2D inverse Fourier transform to reconstruct the image.

Spin Echo 'Spin-Warp'

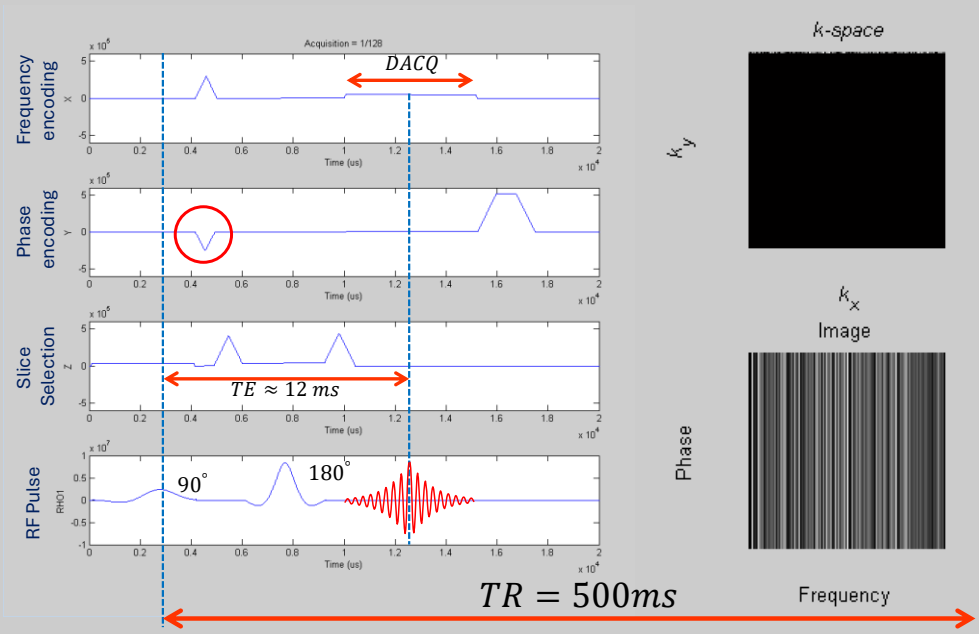

'Spin Warp'

W A 'Bill' Edelstein
1944 - 2014
J M S 'Jim' Hutchinson
1941 - 2018

This movie shows a real spin echo sequence with k_y=128, acquired on one of our clinical MRI systems. Note that all the gradients are the same for every repetition except for the phase encoding gradient (red circle) that is stepped ± 64 times from maximum negative through to maximum positive. Each echo is placed in the appropriate position in *k*-space and in this example, I have performed a Fourier transform of the partially filled k-space as each k_y line is accumulated. Note that in clinical MRI systems the image is not reconstructed until all of *k*-space is filled. This pulse sequence diagram shows the first 20ms of a spin echo sequence. The TE is 12 ms and the TR, to allow for T_1 recovery and hence create a T_1-weighted image, is 500 ms. Therefore, the total acquisition time for this image is 500 ms·128 = 64 s. This is the spin-warp method developed by Jim Hutchison and Bill Edelstein shown here and others at the University of Aberdeen in the 1980s. The team went on to commercialise this system and it was on one of the systems with a field strength of 0.08 mT that I started my career in MRI at St. Bartholomew's Hospital in London in the mid-1980s. It is probably because of this commercialisation that meant that none of the Aberdeen team were also awarded the 2003 Nobel Prize. One person who certainly did think that he also deserved a share of the Nobel Prize was Raymond Damadian who we met in the first lecture, but that is another story altogether!

Spin Echo *k*-space Reconstructions

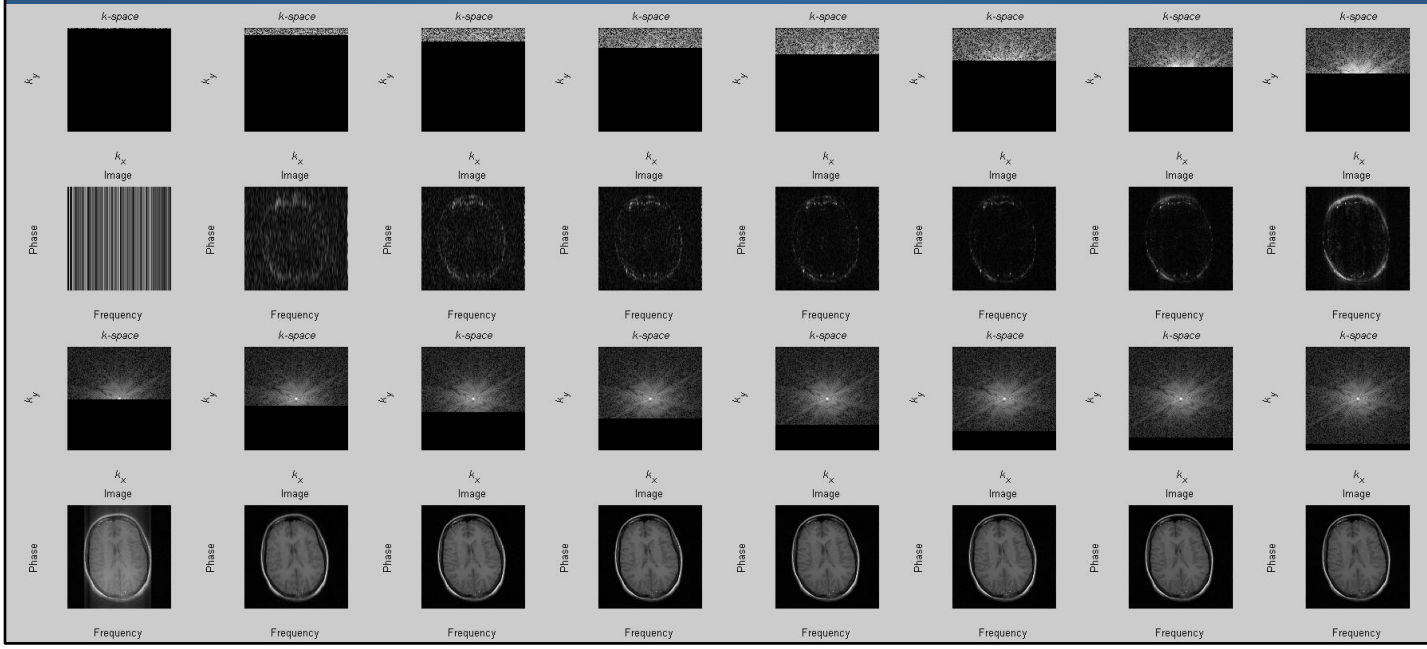

Here is a montage of every 4th *k*-space and reconstructed image from the movie in the previous slide. Top left is the first line of *k*-space with the rest of *k*-space just filled with zeros. Following inverse Fourier transform the image simply shows a 1D projection. As we acquire more k-space we the images start to show the edges of the head, this is because we have so far sampled only the high spatial frequencies in *k*-space. As we pass the centre of *k*-space we start so see the brain as we are now sampling the low spatial frequencies that represent the bulk of the image contrast. Note that as we acquire the bottom half of *k*-space we are improving the spatial resolution.

k-space vs. Image Space

- The spatial frequencies k_x and k_y can be defined as
- $k_x = \gamma \int G_x(t)\, dt$
- $k_y = \gamma \int G_y(t)\, dt$

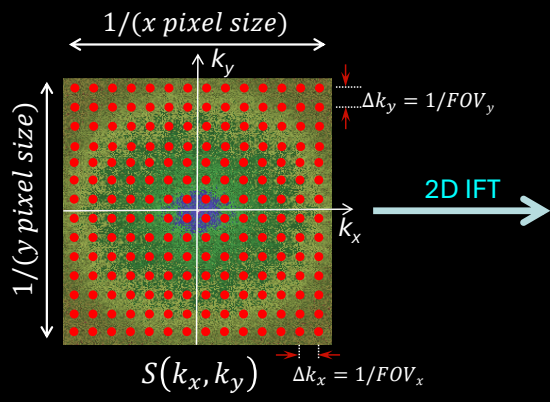

- For simple constant amplitude gradients
- $k_x = \gamma G_x \mathrm{t}$ and $k_y = \gamma G_y \mathrm{t}$
- $S(k_x, k_y) = \iint \rho(x,y) e^{-ik_x x} e^{-ik_y y} dxdy$

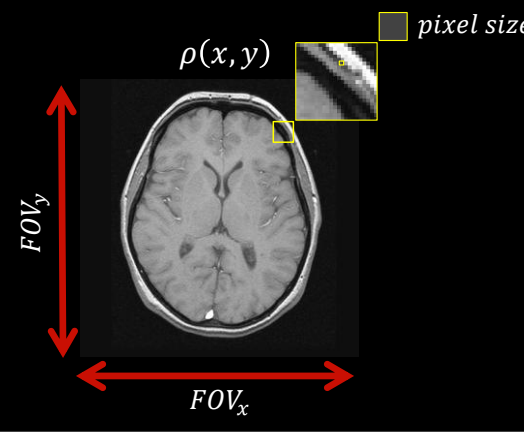

Here are a few observations about the inverse relationship between *k*-space and the reconstructed image space. Note that the extent of *k*-space is inversely related to the pixel size and r spacing between individual samples is inversely related to the field of view. For example, if we wish to increase the field of view but maintain the same number of samples then the *k*-space sampling interval needs to become smaller. This will reduce the extent of *k*-space so the pixel size will therefor become larger. Similarly, if we want to make the pixel size smaller to improve our spatial resolution then we need to increase the extent of *k*-space for the same number of samples. This will increase the *k*-space sampling interval and the field of-view will therefore decrease.

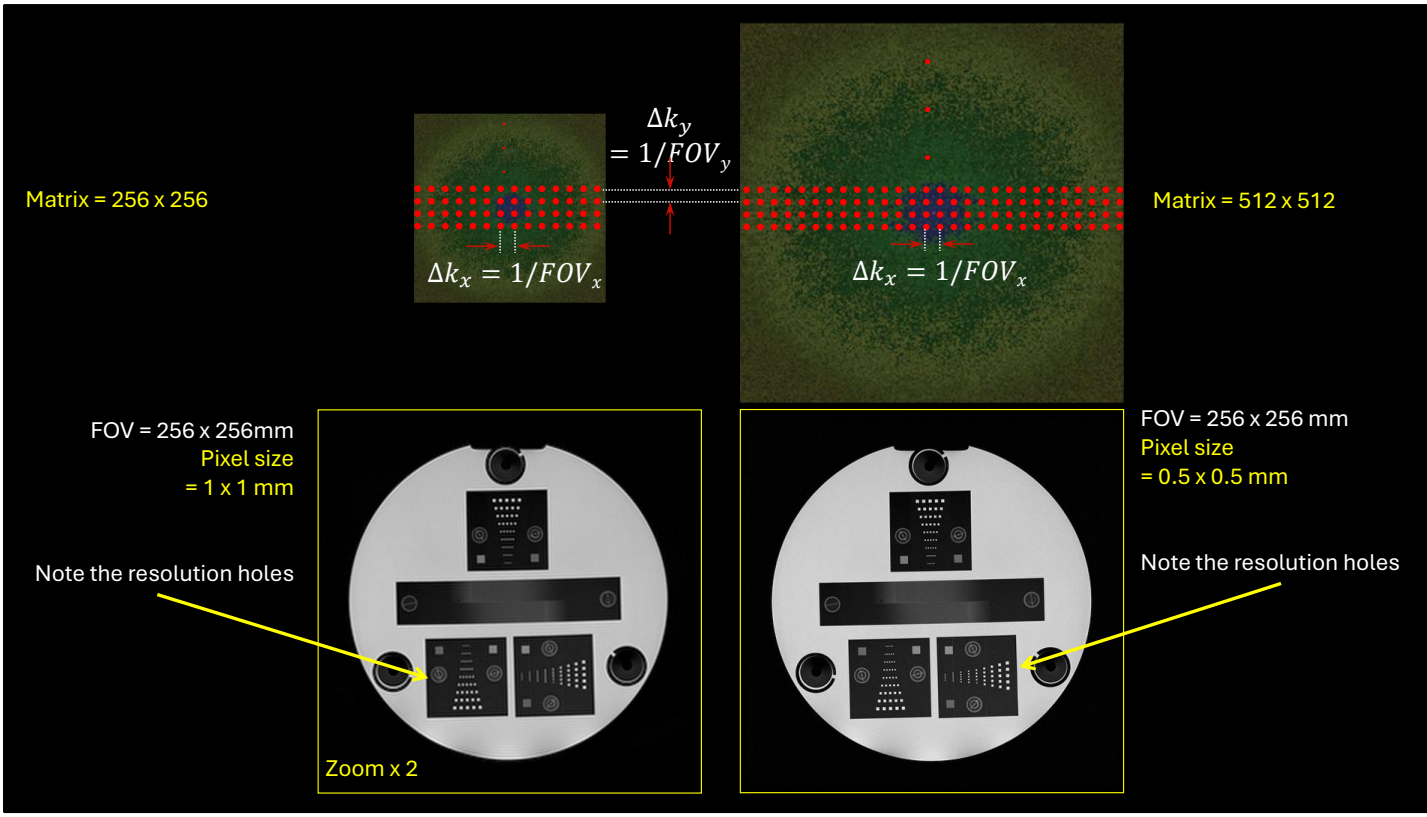

To illustrate this point here increase the matrix from 256^2 to 512^2 so that the extent of k-space is doubled in both directions, but the sampling interval is the same, so the field-of-view is unchanged. The image on the right has a better spatial resolution compared to the one on the left which has been zoomed twice to make the images the same size. Hopefully, you can see the difference in resolution from the spacing of the resolution holes in this image of an MR test object.

k-space Symmetry

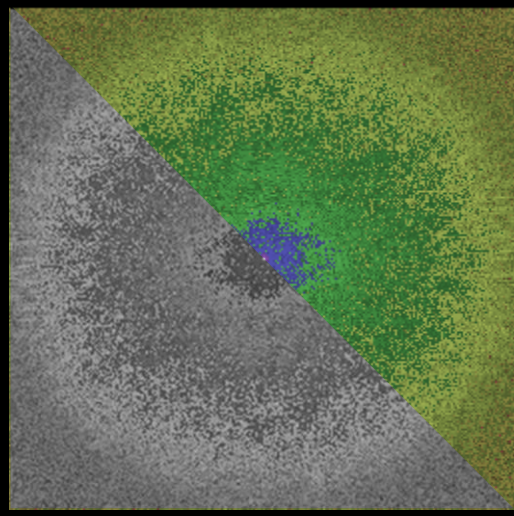

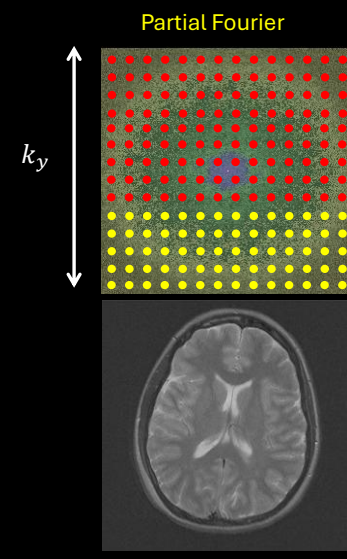

Reduced SNR

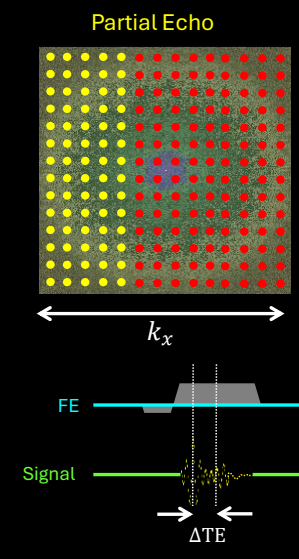

Shorter TE

Another interesting feature of *k*-space is that it has Hermitian symmetry. So, the grey region is the complex conjugate of the coloured region. This means we use techniques such as partial Fourier acquisitions where we only acquire the red samples, thereby reducing the overall acquisition time, but by virtue of this symmetry we can synthesise an approximation to the missing yellow data. The result is an image with the full spatial resolution but a decreased signal-to noise, because we have reduced the number of k_y samples. Alternatively, we can acquire partial echo data in the k_x direction. In this application we reduce the area of the frequency pre-phasing gradient so that the echo does not form in the centre of the frequency encoding gradient but earlier in time. We then sample the red points and synthesise the yellow points. This method sometimes helps if we need to reduce the echo time.

Why are images usually 256 or 512?

- There is special variant of the Fourier transform known as the **Fast** Fourier Transform (FFT) which works with raw data sizes that are a power of two, i.e. $2^6 = 64$, $2^7 = 128$, $2^8 = 256$, $2^9 = 512$, $2^{10} = 1024$
- Raw data that are not a power of 2 are usually zero-padded to the next nearest power of 2, for the FFT, e.g., 288 x288 → 512x512. The displayed image is then shown as 512 x 512.
- The reconstructed image data is **interpolated** (zero filling in *k*-space is equivalent to sinc interpolation in image space. The best possible interpolation for band-limited signals)

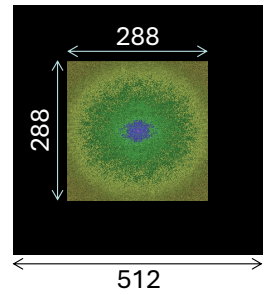

288 x 288 *k*-space is zero-padded to 512 x 512

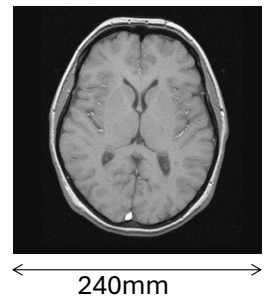

True pixel size = 240/288 = 0.83 mm

Interp. pixel size = 240/512 = 0.47 mm

The dimensions of reconstructed MR images is usually a power of two for example 256 or 512. This is because there is special variant of the Fourier transform which is extremely efficient called the fast Fourier transform which requires the input data to be a power of two. In a typical MRI system if we acquire a non-power of two matrix such as 288 by 288 then the raw data will be zero padded to the next biggest power of two, i.e., 512 by 512. We will then get a 512 by 512 image, but we are sinc interpolating the data. So, the acquired pixel size will be 240mm divided by 288 giving a resolution of 0.83mm but the interpolated resolution will be 240mm divided by 512 or 0.47mm. The use of acquired or interpolated pixel sizes needs to be considered very carefully when making measurements from MR images.

Multiple Slice Interleaving

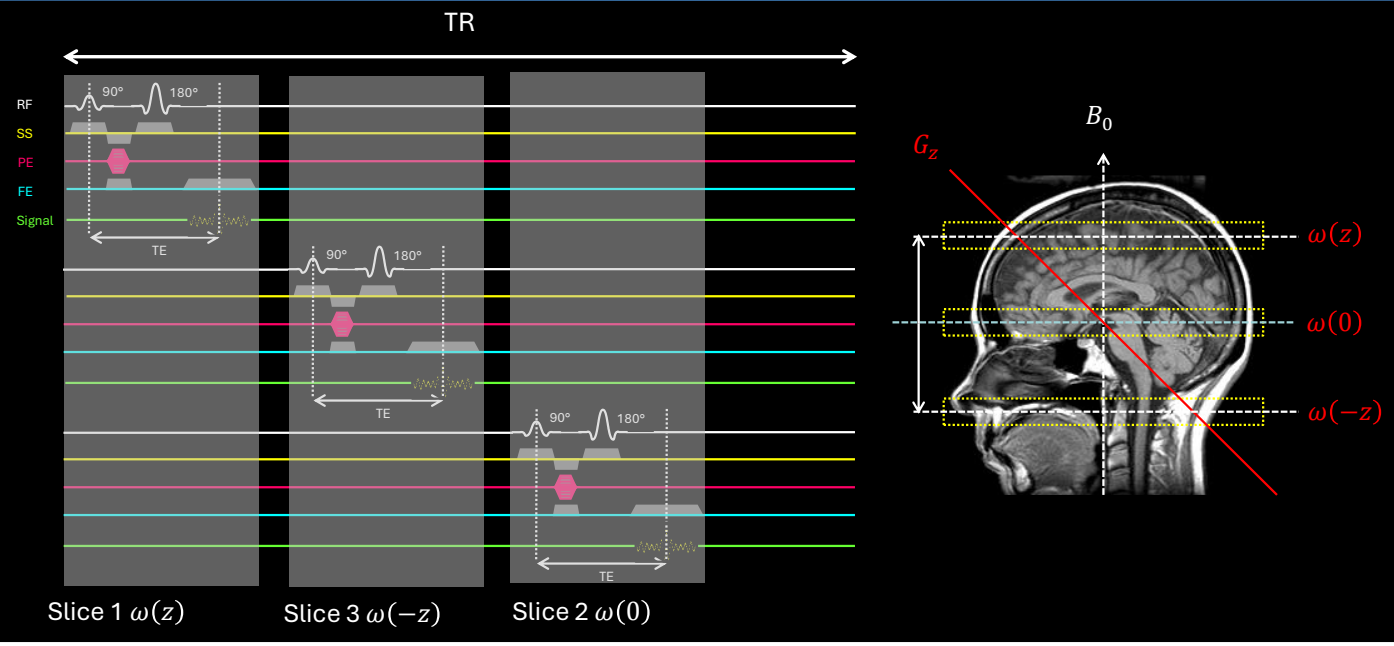

MRI scans can be quite long because of this need for a relatively long TR with generally only one phase encoding step per TR. But this does mean that for most of the TR nothing is happening except T_1 relaxation in order give us the desired contrast. To make MRI more efficient we interleave other slices during the TR period. This simple example shows three slices being interleaved in the TR period. All that is going to change between slices is the carrier frequency for that slice location. The RF pulse envelope will not change since that only dictates the slice thickness and that will be the same for each slice. Since only the carrier frequencies are different only the spins within each slice will be excited and the other slices will recover normally. In a T_1 weighted scan with a TR of around 500ms we could fit approximately 30 slices into the TR. So, although it is still going to take several minutes to acquire the data it does mean that we will acquire 30 slices.

Gradient Echo

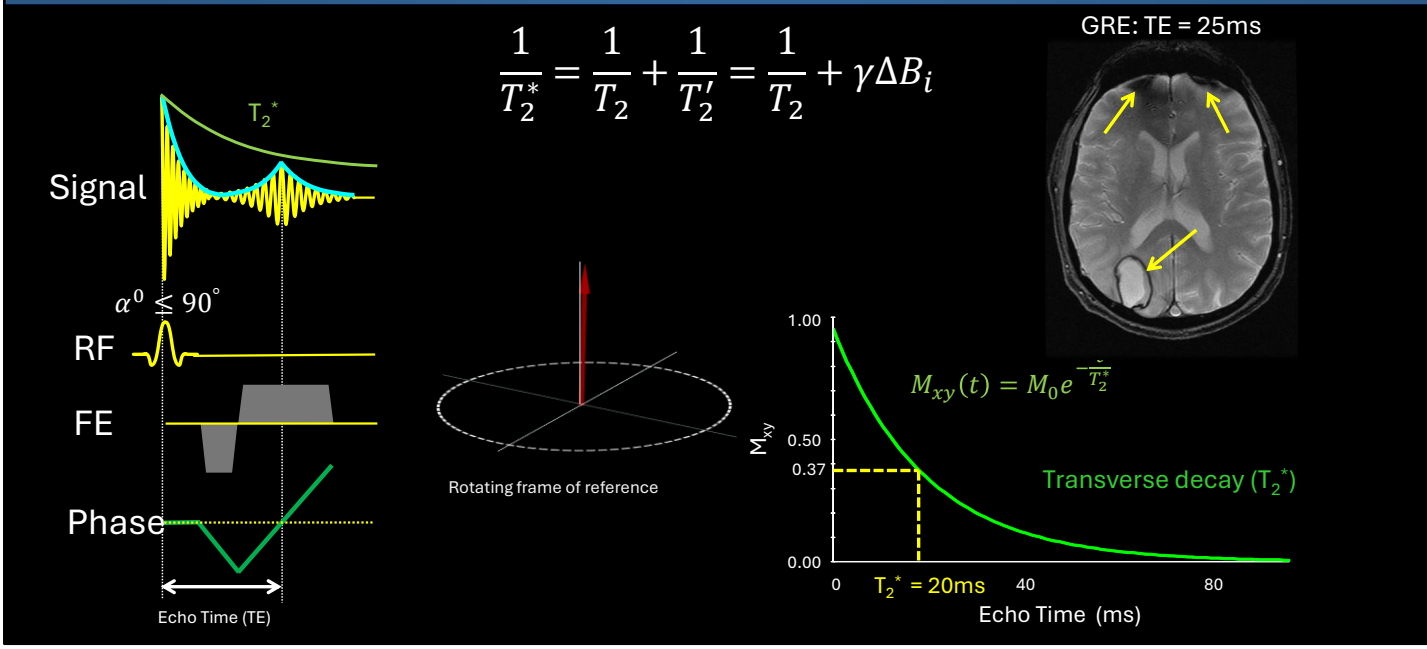

The last pulse sequence I want to briefly explain is the gradient recalled echo sequence. In comparison to the spin echo sequence a gradient echo sequence does not use a 180° refocusing pulse. Instead, the echo is formed by making the frequency encoding prephasing gradient negative which forcibly dephases the magnetisation and then making the frequency encoding gradient the opposite polarity which forcibly rephases the magnetization forming what is known as a gradient recalled echo. Since we are not using a 180° refocusing pulse the gradient echo cannot eliminate the effects of T_2' therefore the image is weighted by T_2^*. The use of T_2^* weighting is often useful in certain pathologies such as in this patient who has had a brain bleed or haemorrhage and the dark outline is due to a rim of haemosiderin, a by-product of the breakdown of red blood cells. The image also shows signal loss at the front of the brain. this is due to the presence of air in the sinuses is below this slice. The difference in magnetic susceptibilities between tissue and air also causes T_2' dephasing resulting in this signal loss. Gradient echo sequences also often use excitation flip angles less than 90°. This means that they can be operated with very short repetition times without substantially saturating the longitudinal magnetisation.

Gradient Echo vs. Spin Echo

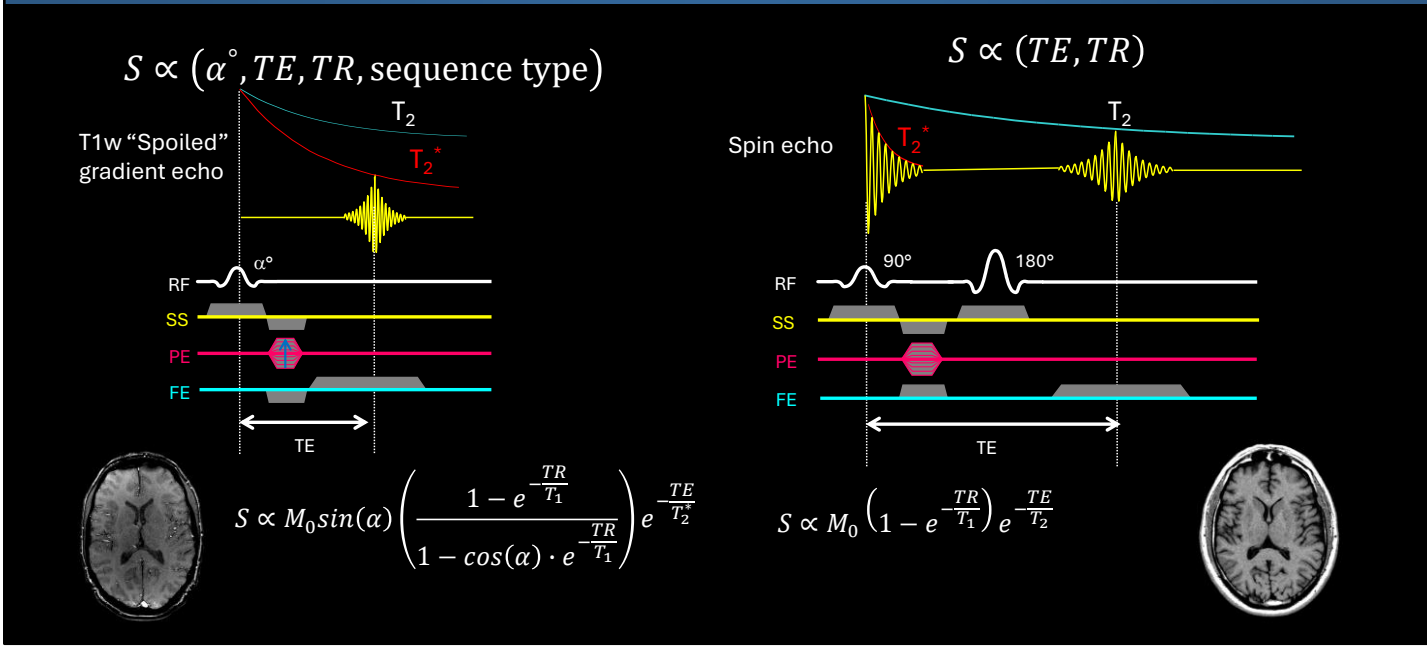

This slide compares the gradient echo pulse sequence with the spin echo. The gradient echo uses a gradient reversal to form the echo whilst the spin echo uses a 180° refocusing pulse. The gradient echo sequence forms under the T_2^* decay curve whereas the spin echo sequence forms under the T_2 decay curve. The gradient echo sequence typically uses a flip angle less than 90°, so the signal equation is somewhat more complicated involving not only TE and TR but also the flip angle (α°). Whereas the spin echo sequence has a relatively straight forward signal equation. Note that there are variants of the gradient echo sequence that have quite different image contrast appearances the image shown here is obtained using a T_1w or spoiled gradient echo sequence. I also show a T_1 w spin echo image for comparison.

T2w Spin Echo vs. T2*w Gradient Echo

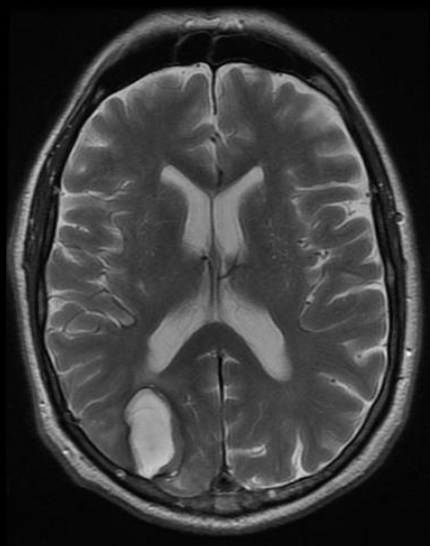

SE: TE = 105ms TR =3906ms

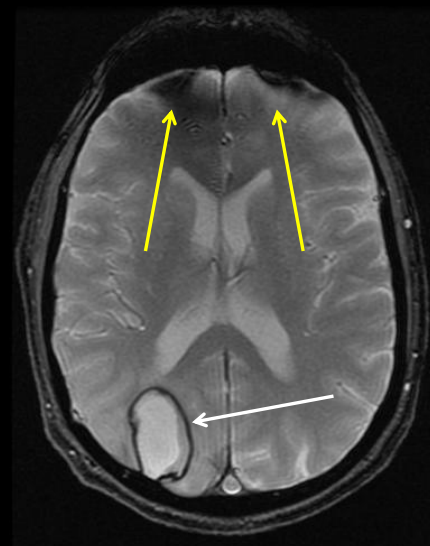

GRE: TE = 25ms TR =800ms

A T_2w spin echo acquisition together with a T_2*w gradient echo acquisition of a patient with a resolving haematoma. Note the intense black outline of a rim of hemosiderin (white arrow) confirming the presence of haemorrhage on the gradient echo image (white arrow). This is due to the susceptibility caused by the iron in this blood breakdown product. The refocusing pulse in the spin echo acquisition reduces this effect. Note also the signal loss at the top of the gradient echo image (yellow arrow). This is due to differences in the magnetic susceptibility of tissue and air in the prefrontal sinuses just below this slice. This artefacts is not seen in the spin echo image.

3D Gradient Echo

Apply a second ‘slab encoding’ in the slice select direction

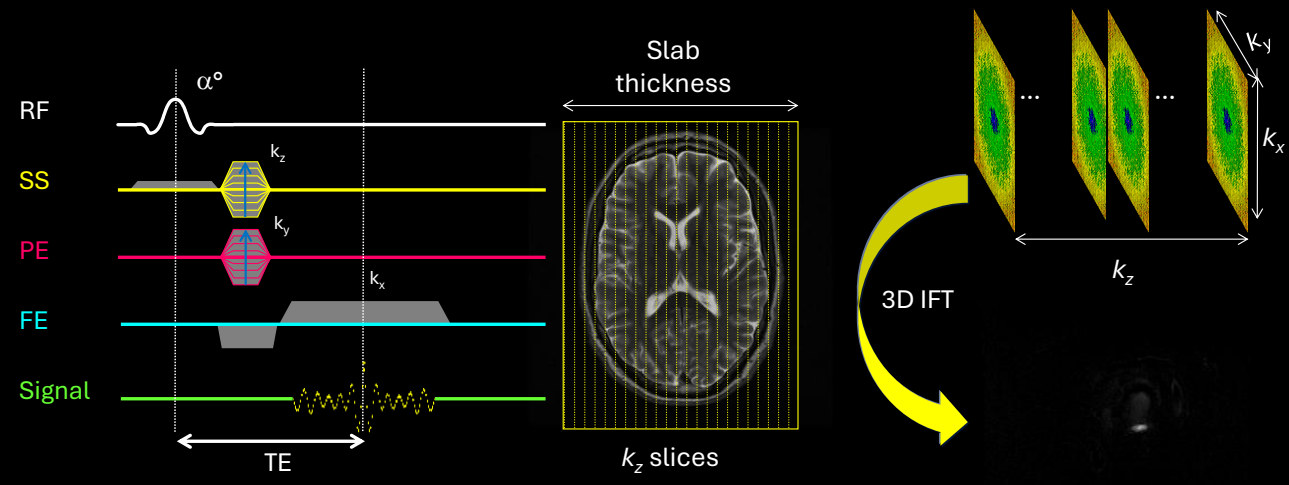

Acq. time = $TR \cdot N_y \cdot N_z \cdot NEX$

One advantage of gradient echo imaging is that with a short TR true 3D imaging can be performed. In 3D imaging a second phase encoding gradient is applied along the slice selection direction. The slice selection gradient excites a thick slab of tissue and the two-phase encoding gradients together with the frequency encoding gradient encodes a 3D *k*-space. The images are reconstructed into conventional slices using a 3D FFT. A short TR is required since the acquisition time is given by TR· . 3D imaging has the advantage that very thin slices can be reconstructed resulting in isotropic voxels that can be reformatted into any plane.

Learning Outcomes

- After these lectures you should be able to:
  - Explain how nuclear spin gives rise to magnetic resonance
  - Understand the principles of T_1, T_2 and T_2^* relaxation
  - Explain the principles of MR image formation
  - Describe the spin echo and gradient echo pulse sequences
  - Outline the basic components of an MRI system
  - Understand the safety issues related to MRI

Non-examinable Material

- Inversion Recovery (IR) pulse sequence
  - Magnitude vs. phase- sensitive reconstruction
  - Short τ (TI) IR (STIR) – suppresses fat
  - FLuid Attenuated IR (FLAIR) – suppresses CSF
- Fat
  - Chemical shift of ^{1}H in water and fat
  - In- and out-of-phase imaging
  - CHEmical Shift Selective (CHESS) fat saturation
  - Comparison of CHESS and STIR

Fat can be a confounder in MRI. The main constituent of body fat in humans are triglycerides, which are relatively large molecules containing lots of CH_2 or methylene groups. Because of their large structure the electrons partially shield the protons from the static magnetic field which means the protons in fat precess at a slightly lower Larmor frequency, which can be seen in this spectrum of fat where the methylene resonance is about 3.4 ppm lower than the water. This is about 220 Hz at 1.5T.

Large molecules like triglycerides that tumble near the Larmor frequency demonstrate more efficient T_1 relaxation and hence have a shorter T_1, so fat appears hyperintense on T_1w imaging.

Inversion Recovery

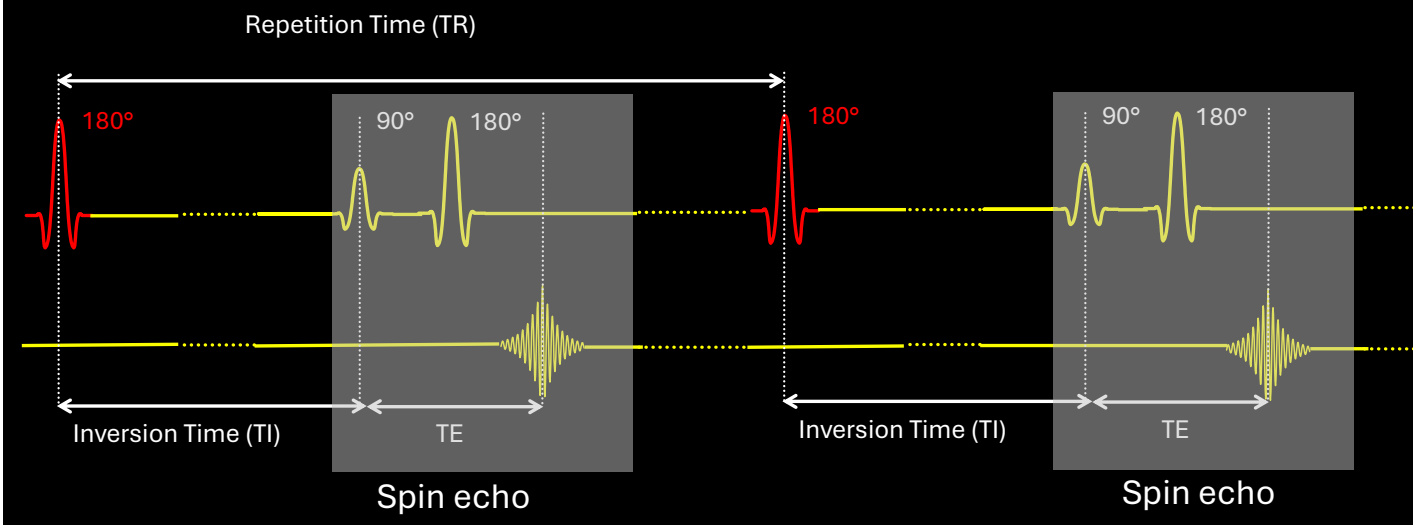

A modification of the spin echo is the inversion recovery sequence. This starts with a 180° pulse which inverts the net magnetization from being along the +z to along −z. We then wait a time from the end of that pulse to the excitation pulse of a spin echo sequence. The time is known as the inversion time during which the magnetisation will recover due to T_1 relaxation. The relevant magnetisation recovery will then be detected in the transverse plane using the spin echo sequence.

Inversion Recovery: Real/Phase Sensitive

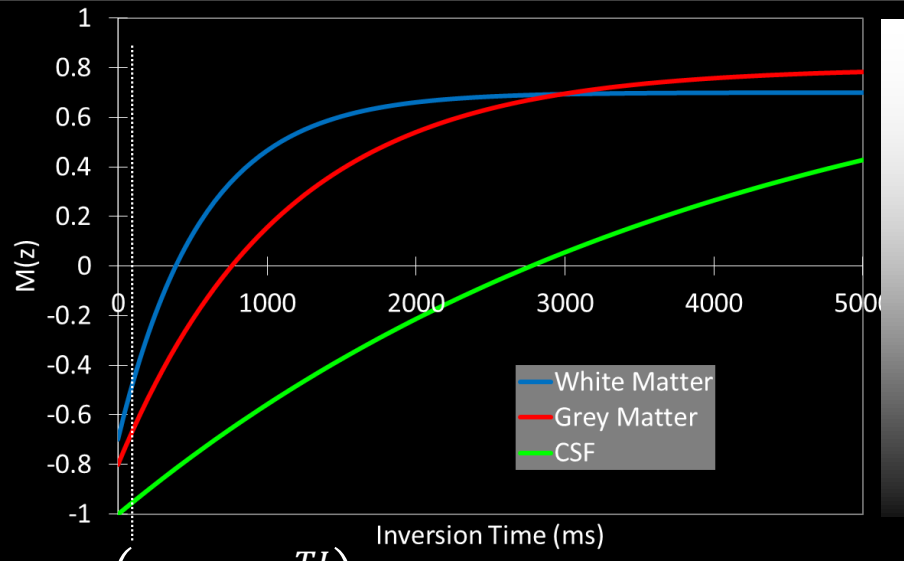

$$S = S_0\left(1 - 2e^{-\frac{TI}{T_1}}\right)$$

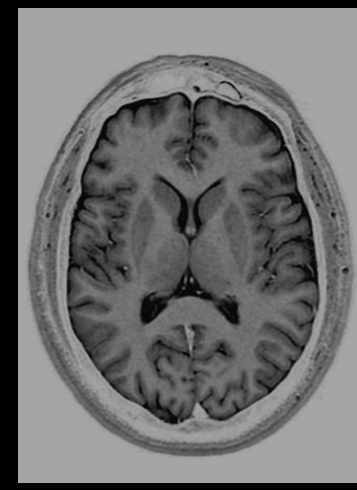

TR = 15000, TE = 5, TI = 160ms

Here we can see the inversion recovery curves for CSF, grey matter, and white matter. In this example we are making use of the real component of the complex data because it encodes the polarity of the magnetisation. The grayscales now go from negative to positive. Real, or sometimes called phase sensitive inversion recovery images, are characterised by the fact that the background air is, with zero signal is displayed as mid-grey. This is an example that has been acquired with an inversion time of about 160 ms so even though the signal is all negative CSF appears darkest with white matter appearing brighter than the grey matter exactly as would be expected from their relative T_1 relaxation times.

Inversion Recovery: Magnitude

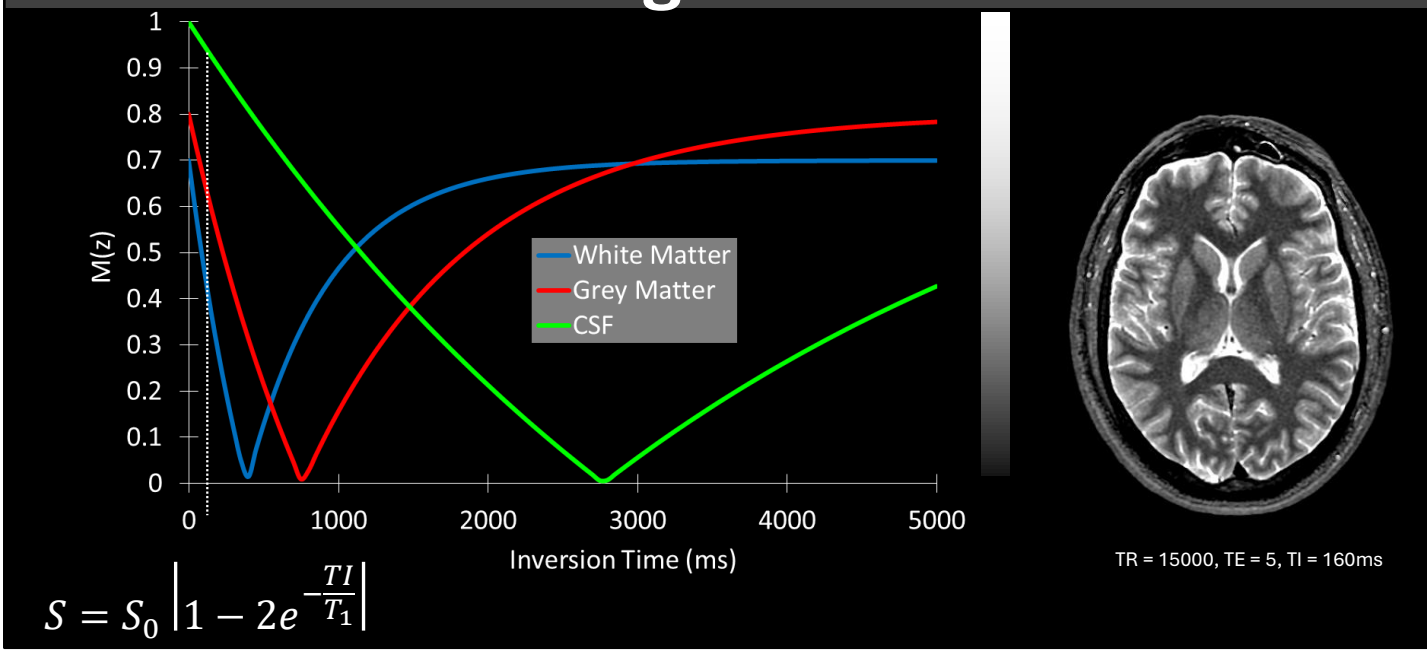

$$S = S_0 \left|1 - 2e^{-\frac{TI}{T_1}}\right|$$

What happens if we just take the magnitude of the data rather than the real component. Everything that was negative simply becomes a positive value and the signal recovery curves look like this. So now for the same 160 ms inversion time we see that CSF is now the brightest tissue followed by grey matter and then white matter. So, does this slightly strange contrast have any useful applications? You can see that for certain inversion times the signal hits zero, so with careful choice of inversion time we could null the signal from any tissue based on a knowledge of its T_1 relaxation time.

Short τ Inversion Recovery (STIR)

$$S = S_0\left(1 - 2e^{-\frac{TI}{T_1}}\right)$$

$$TI_{null} = ln(2) \cdot T_1 = 0.693 \cdot T_1$$
$$0.693 \cdot 260\ \text{ms} = 180\ \text{ms}$$

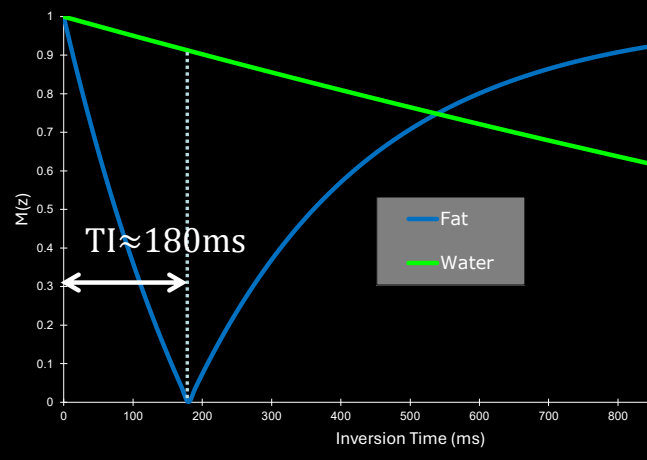

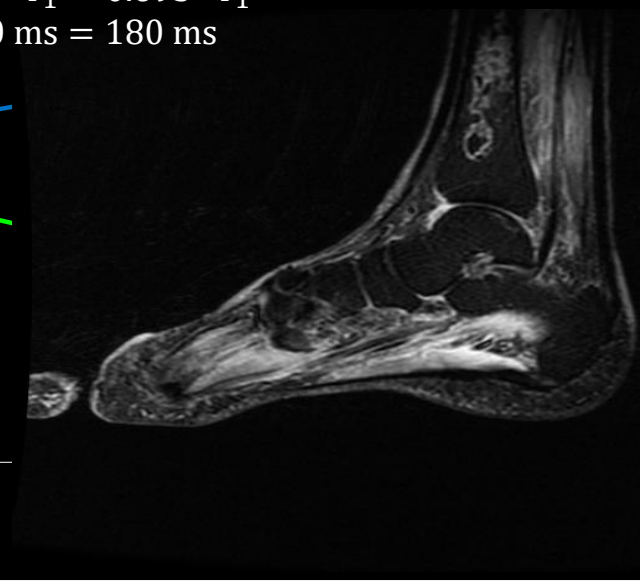

One application of magnitude reconstructed inversion recovery is the so-called short τ inversion recovery or STIR sequence. Fat is often a confounding feature in images because it has a very short T_1 relaxation time and therefore appears very bright in T_1 weighted images. The diagram shows the magnitude recovery curves for fat and water. Note that because of the short T_1 the fat signal reaches zero at an inversion time of approximately 180 ms. Images that are reconstructed using this technique such as this ankle here will demonstrate very good fat suppression of the bone marrow. In this patient there is also some fluid accumulation at the bottom of the foot and because the T_1 of fluid is very long at a TI of 180 ms the fluid appears very bright and may help with the clinical diagnosis. If you know the T_1 of a particular tissue, then the TI for nulling is simply given by the natural log of two multiplied by the T_1

FLuid Attenuated IR (FLAIR)

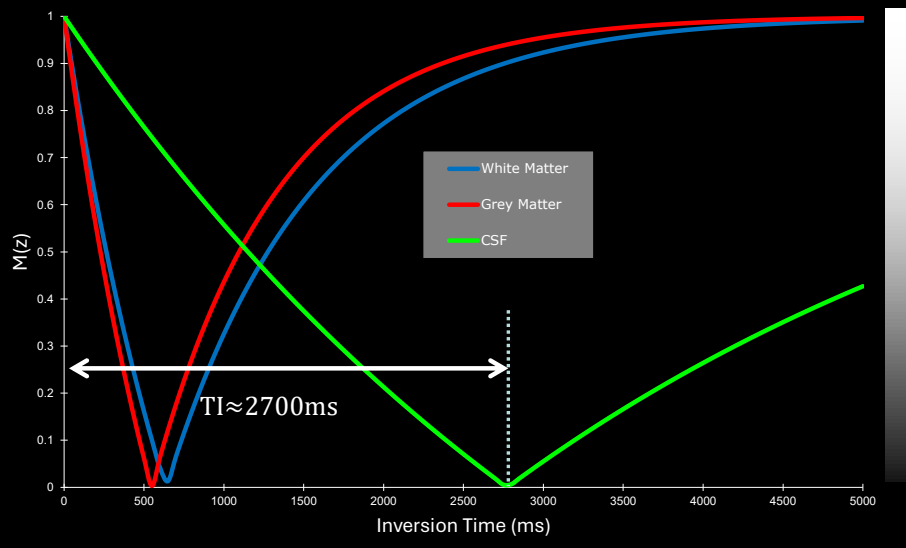

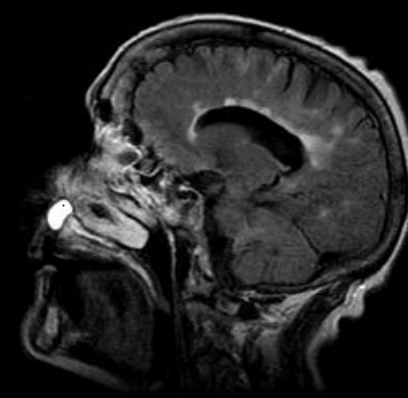

This is an example of nulling the signal from CSF in the brain. This is known as a fluid attenuated inversion recovery or FLAIR sequence where we use a TI of approximately 2700 ms and with magnitude reconstruction, we get this null point for CSF. Since this is a T_2w brain image we would expect the CSF to appear very bright. If the signal from CSF was not suppressed, then it would have been very difficult to see these pathological periventricular white matter signal hyperintensities

T2w vs. FLAIR

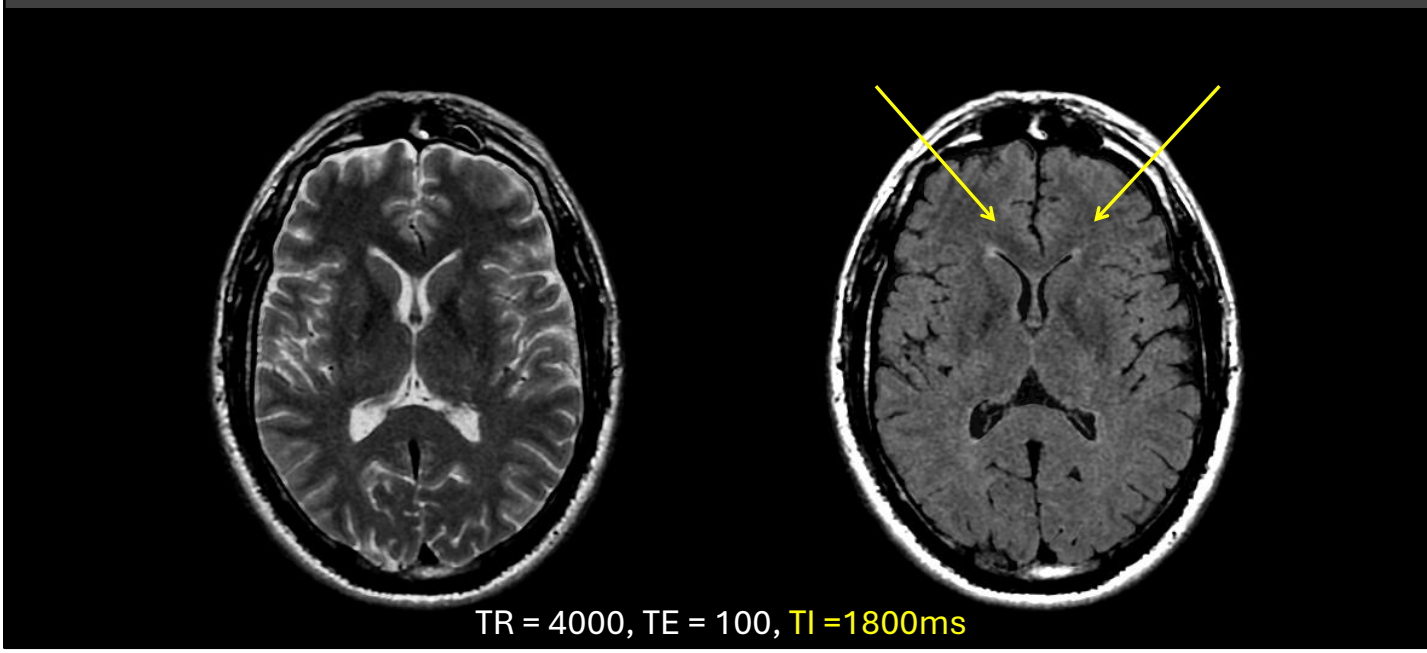

Here is a comparison of an axial T_2w brain image acquired with and without the FLAIR inversion pulse. The small white matter signal hyperintensities arrowed on the FLAIR image are extremely difficult to identify on the conventional T_2 w image.

Fat

- Nuclear shielding gives rise to chemical shifts of 0.9 to 5.3ppm
- Most abundant resonance is methylene (CH_2) at 3.4ppm below water
- Triglycerides are relatively large molecules, so they have a short T_1 of around 270ms at 1.5T
  - Hyperintense on T1w

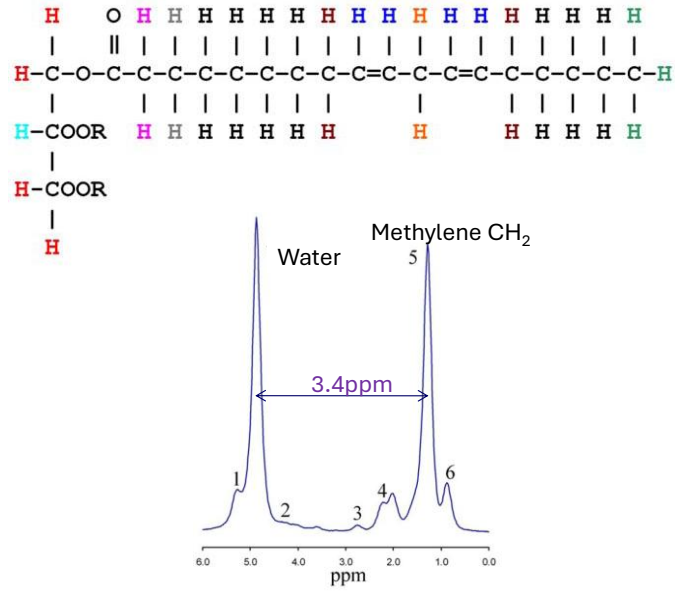

Fat can be a confounder in MRI. The main constituent of body fat in humans are triglycerides, which are relatively large molecules containing lots of CH_2 or methylene groups. Because of their large structure the electrons partially shield the protons from the static magnetic field which means the protons in fat precess at a slightly lower Larmor frequency, which can be seen in this spectrum of fat where the methylene resonance is about 3.4 ppm lower than the water. This is about 220 Hz at 1.5T.

Large molecules like triglycerides that tumble near the Larmor frequency demonstrate more efficient T_1 relaxation and hence have a shorter T_1, so fat appears hyperintense on T_1w imaging.

In- and Out-of-Phase Imaging

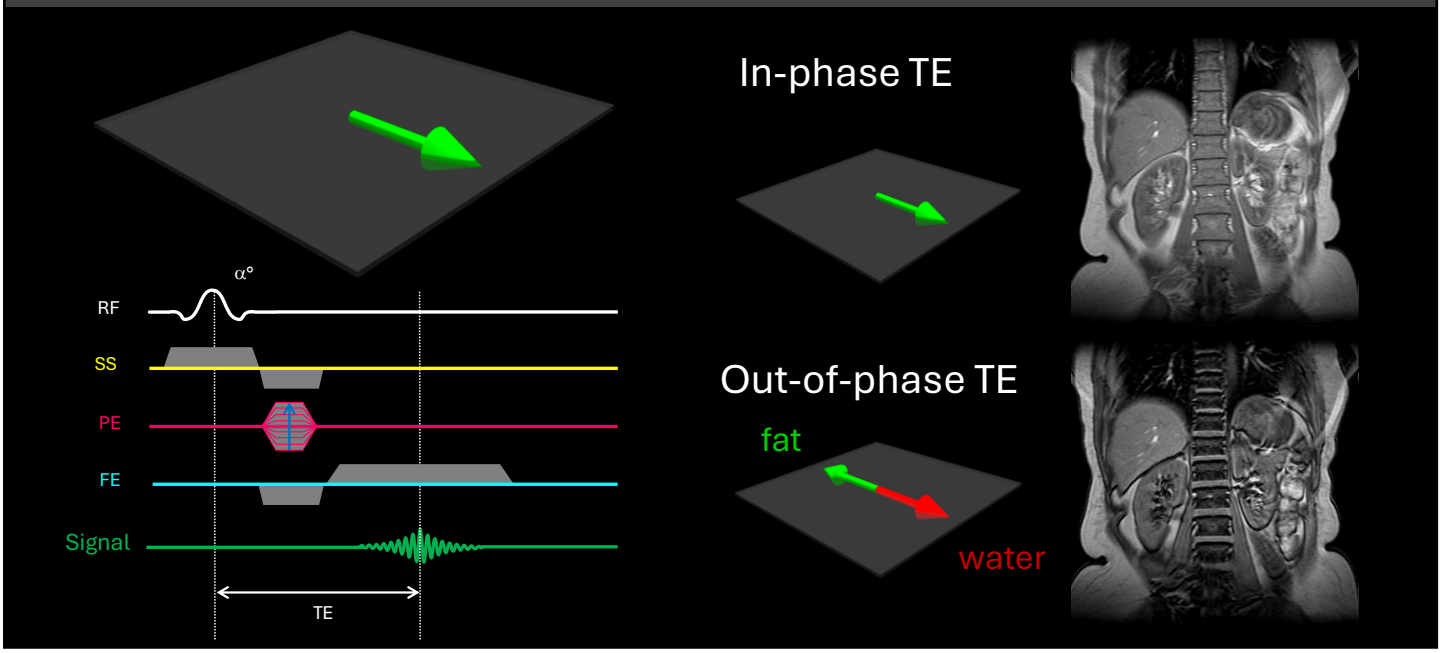

A consequence of this difference in precessional frequency between fat and water means that in gradient echo imaging the fat and water signals will have a different relative phase depending on the echo time (TE). By careful selection of the TE, it is possible to obtain images with fat and water in phase and when fat and water are out of phase. In the out of phase images, you will see a black line surrounding tissues where the voxel contains an equal admixture of water and fat resulting in signal cancellation.

CHEmical Shift Selective (CHESS) Fat Saturation

~90° and ~-220Hz frequency offset at 1.5T
Bandwidth ≈150Hz
Duration ≈ 7ms

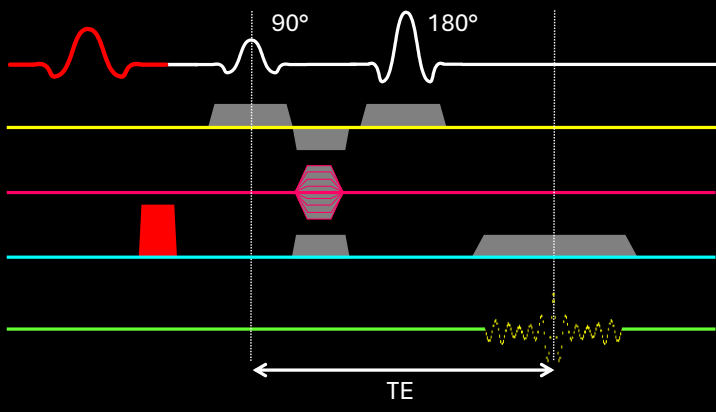

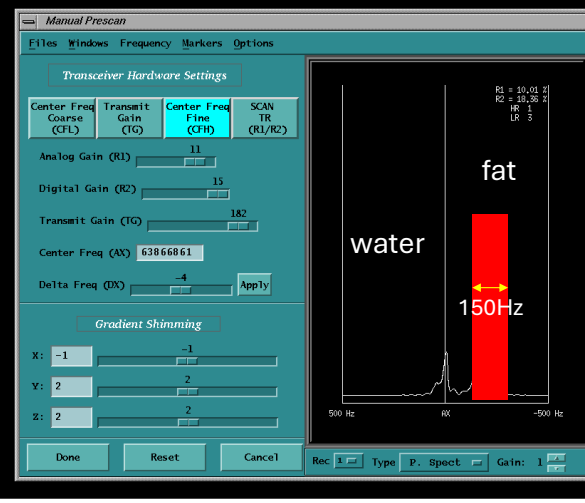

So how can we suppress the signal from fat if we do not want it in our image. One common method is to use a CHESS or chemical shift selective fat saturation pulse prior to the imaging sequence. Here we can see the peaks for water and fat at 1.5 T, with fat approximately 220 Hz lower. If we apply a frequency selective RF pulse centred on the fat frequency, then we can saturate the signal from fat by tipping it into the transverse plane and then using a magnetic field gradient pulse to dephase the signal before we play out our imaging sequence. A simple 7 ms duration RF pulse will excite a relatively narrow bandwidth of approximately 150 Hz. If the pulse is centred on fat, then it will have little impact on the water peak.

CHESS Fat Suppression in Abdomen

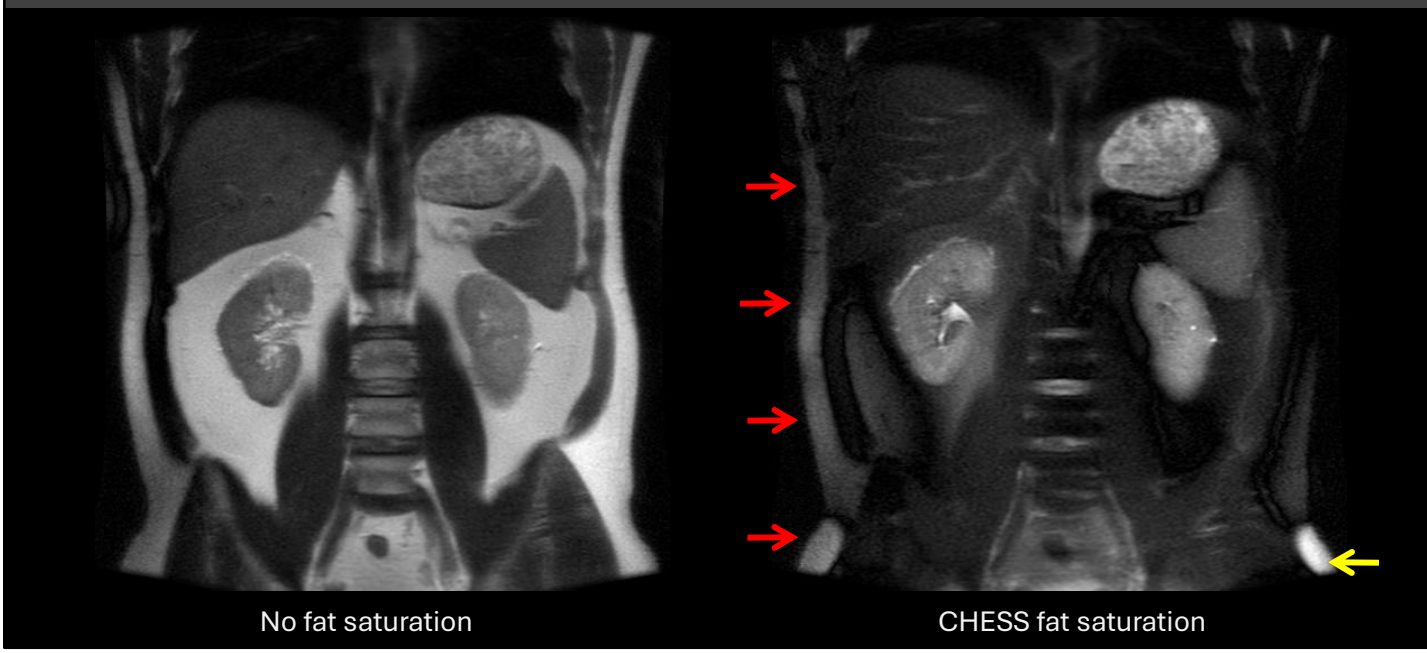

Here are two T_1 weighted images acquired with and without chemical shift selective fat saturation. The fat saturation pulse does a reasonably good job of suppressing the intra-abdominal fat. There is a small focal fat saturation failure at the bottom right of the image (yellow arrow) and generally poor fat suppression on the left of the body (red arrows). This is due to the variation in uniformity of the static magnetic field. In these regions, the Larmor frequency of fat has shifted completely or partially outside of the bandwidth of the fat saturation pulse.

Fat Saturation vs. STIR in the Ankle

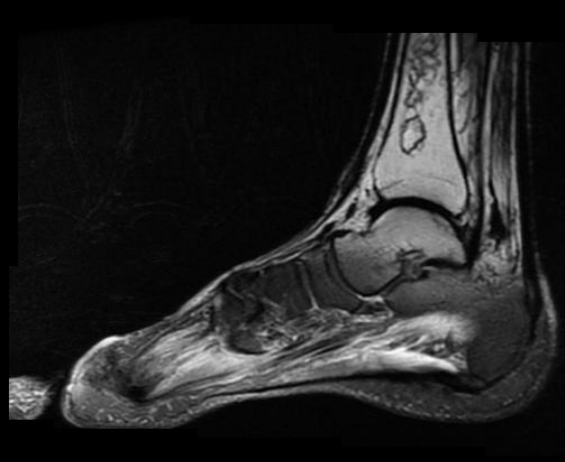

T2w Fat Sat
TE/TR = 80.9/3321 ms

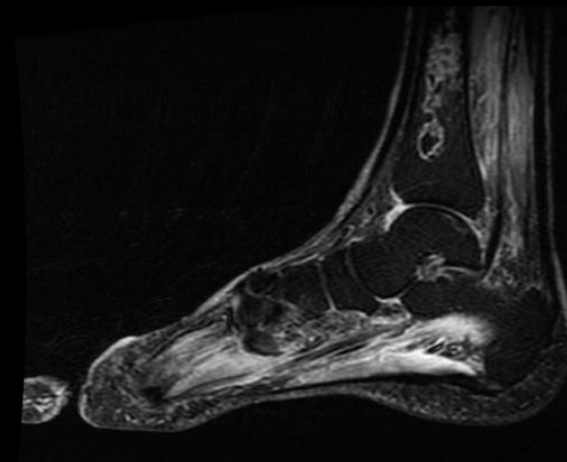

STIR
TE/TR/TI = 37.6/3400/150 ms

We saw this case earlier when talking about STIR, where the fat is suppressed based on its T_1 relaxation time. Here we compared the STIR image with a T_2w image using Chemical shift selective fat saturation. Because the ankle is a difficult area to obtain a good uniform magnetic field, we can see a clear non-uniformity of the chemical shift selective fat saturation. In this case the STIR method does a much better job