Magnetic Resonance Imaging basics and fMRI

Mark Cohen

UCLA Center for Cognitive Neuroscience

Departments of Psychiatry, Neurology, Radiology, Psychology, Biomedical Physics. Biomedical Engineering

Nuclear Spin

- "Spin" is a property of many particles. It is a type of angular momentum.
- Angular momentum is a vector quantity: has *magnitude* & *direction**.

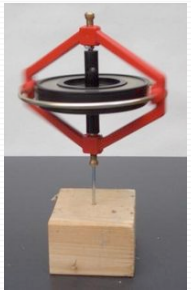

$$\mathbf{S} = \hbar\sqrt{s(s+1)}, \text{ where } s = \{0, \tfrac{1}{2}, 1, \tfrac{3}{2}, 2, \ldots\}$$

$$\hbar = 1.0546 \text{ X } 10^{-34} \textbf{J-s}$$

* Quantum effects prohibit knowing both magnitude *and* three-dimensional orientation. We *can* know both the z-component and magnitude.

UCLA Medical Center

©2008 Mark Cohen, all rights reserved

www.brainmapping.org

Center for Cognitive Neuroscience

Nuclear spin is a quantized property. The quantum angular momentum behaves like the classical spin of a moving top. Unpaired nuclear spin results in the creation of a magnetic moment, oriented along the axis of spin.

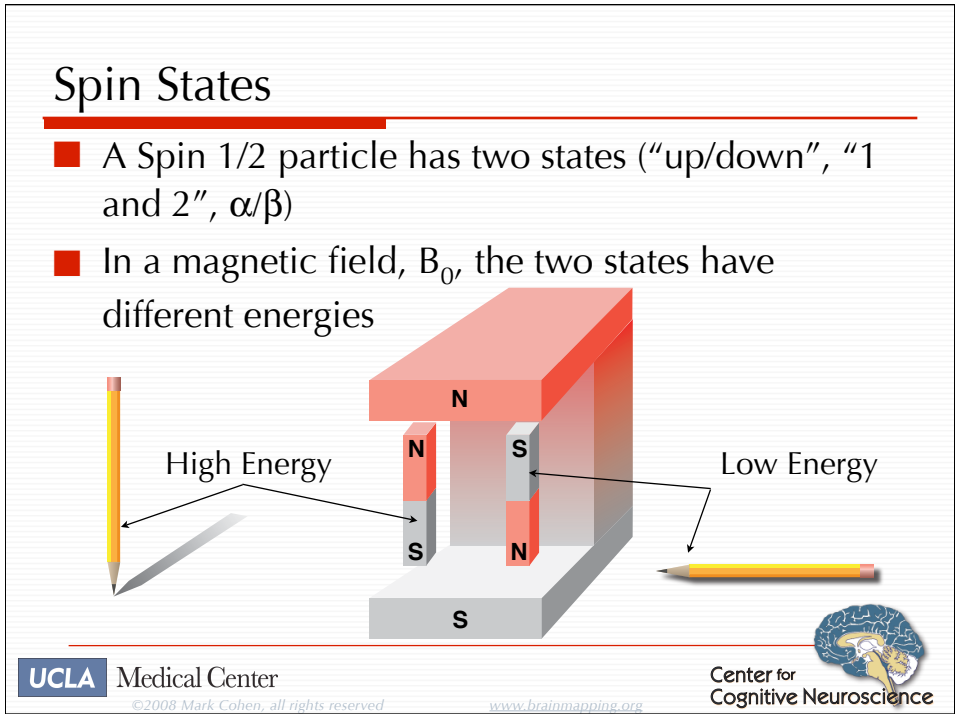

When the nuclear magnet is placed into a magnetic field it can adopt only one of two states - aligned with or aligned against - the external field. Then the individual nuclear spins are aligned AGAINST the external field they are in a lower energy state. Moving a system to its lowest energy state is called "relaxation" and amounts to a loss of stored heat energy.

Proton Responses to Magnetic Field

- Spin Alignment Along Net Applied Field
  spins align parallel or anti-parallel to the applied field
- Precession About the Magnetic Field
  at a precession frequency of: $\gamma \times B$, known as the Larmor frequency
- Spin Alignment Occurs at the Rate, T1

©2008 Mark Cohen, all rights reserved

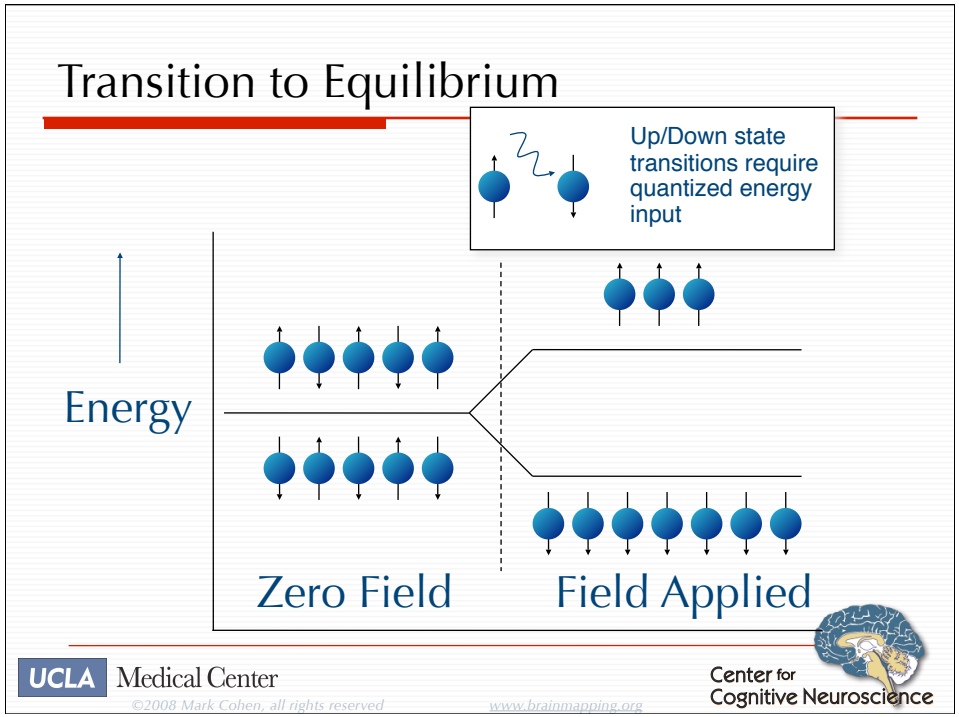

Absent an external field, there will be an equal number of spins in the two allowed states. When an external field is applied, some of the spins change state to oppose it. The local effects on individual spins is much larger than the effect of an external field; therefore only a small number of spins change state (1 in a million or so). The state changes actually require the capture of energy from the environment, which occurs rarely. Therefore, the relaxation to magnetic equilibrium can take some time.

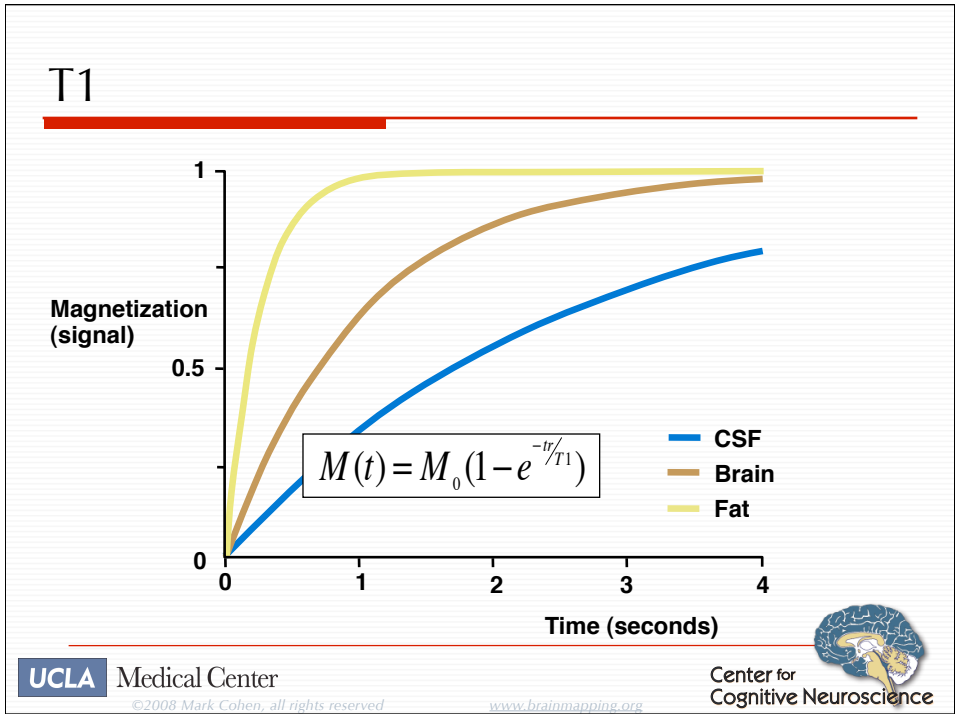

The approach to magnetic equilibrium is exponential. When the system is far from equilibrium, most interactions with the environment will move individual spins to the lower energy state. As the system reaches equilibrium, however, the state transitions become nearly equally likely to be either up or down.

The exponential process is governed by a single time constant, T1, which is characteristic of the tissue type. The T1 of water (CSF) is up to several seconds, while the T1 of lipid (fat) is a few tenths of seconds. Most body tissues lie between. The difference in T1 forms a major contrast mechanism in MRI.

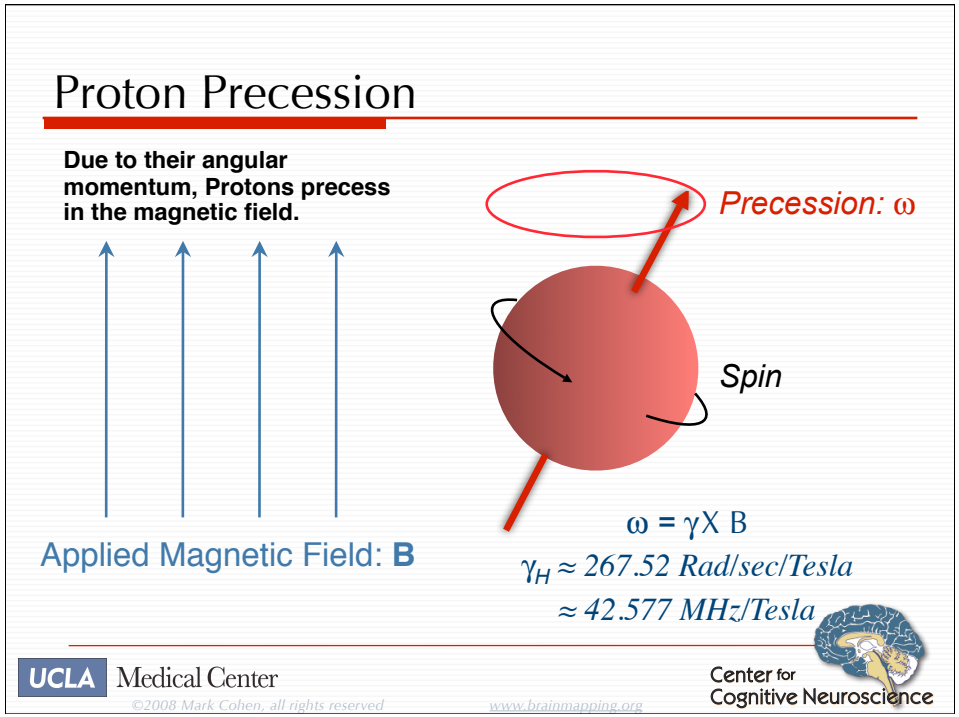

The angular moment of the spins is *conserved*. This means that when an external field is applied, the individual spins will *precess* about it. The rate of precession is known as the Larmor frequency and is proportional to the strength of the external field. The "Larmor constant" or "gyromagnetic ratio" is about 42.577 million cycles/second per Tesla.

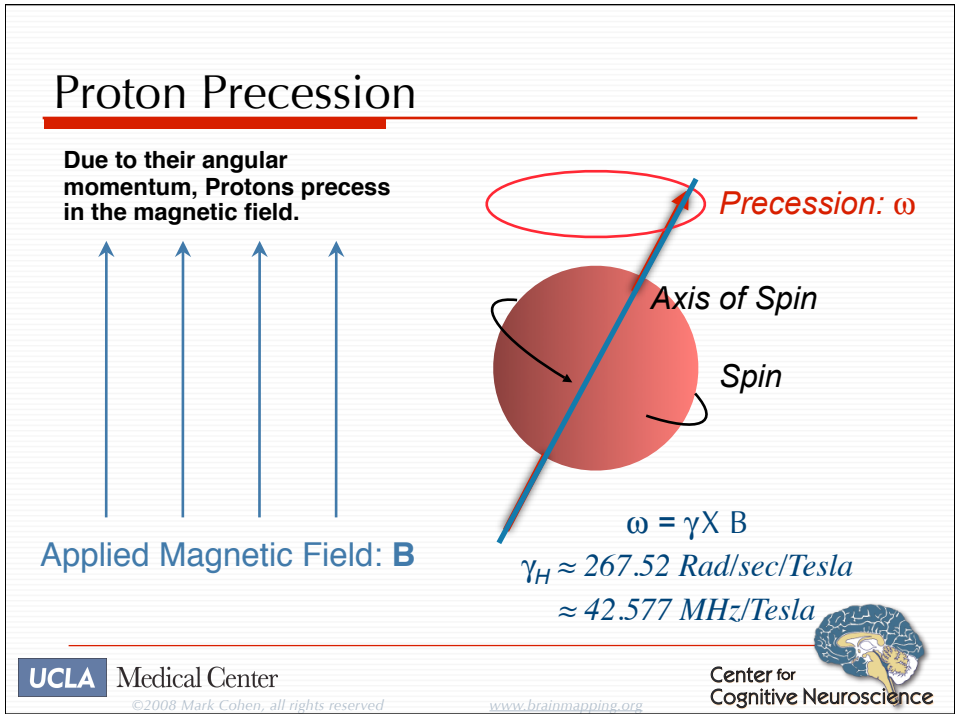

The angular moment of the spins is *conserved*. This means that when an external field is applied, the individual spins will *precess* about it. The rate of precession is known as the Larmor frequency and is proportional to the strength of the external field. The "Larmor constant" or "gyromagnetic ratio" is about 42.577 million cycles/second per Tesla.

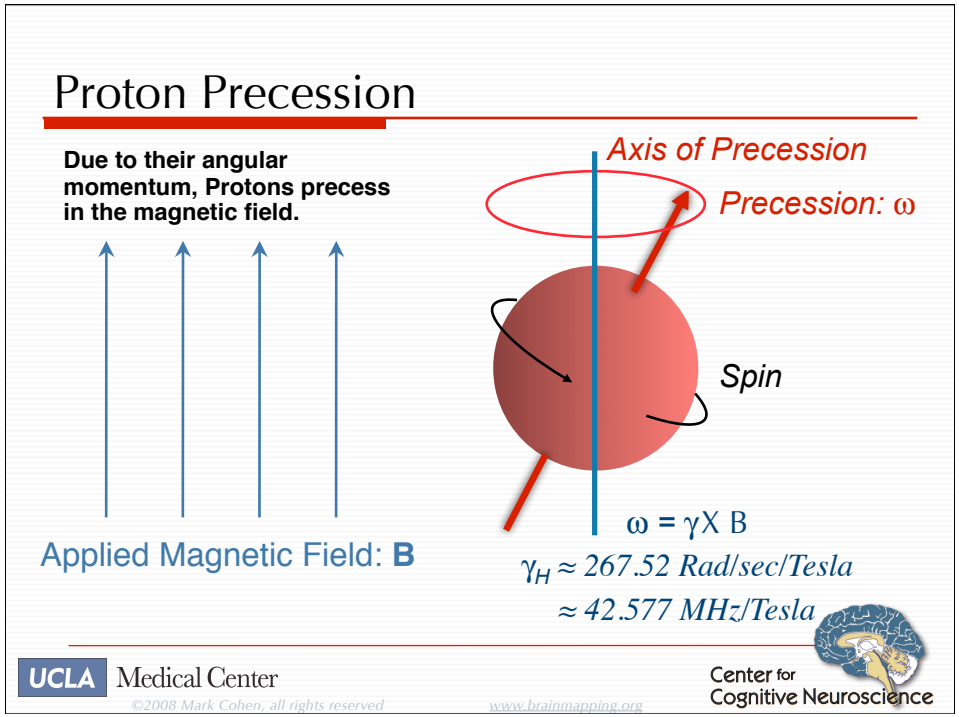

The angular moment of the spins is *conserved*. This means that when an external field is applied, the individual spins will *precess* about it. The rate of precession is known as the Larmor frequency and is proportional to the strength of the external field. The "Larmor constant" or "gyromagnetic ratio" is about 42.577 million cycles/second per Tesla.

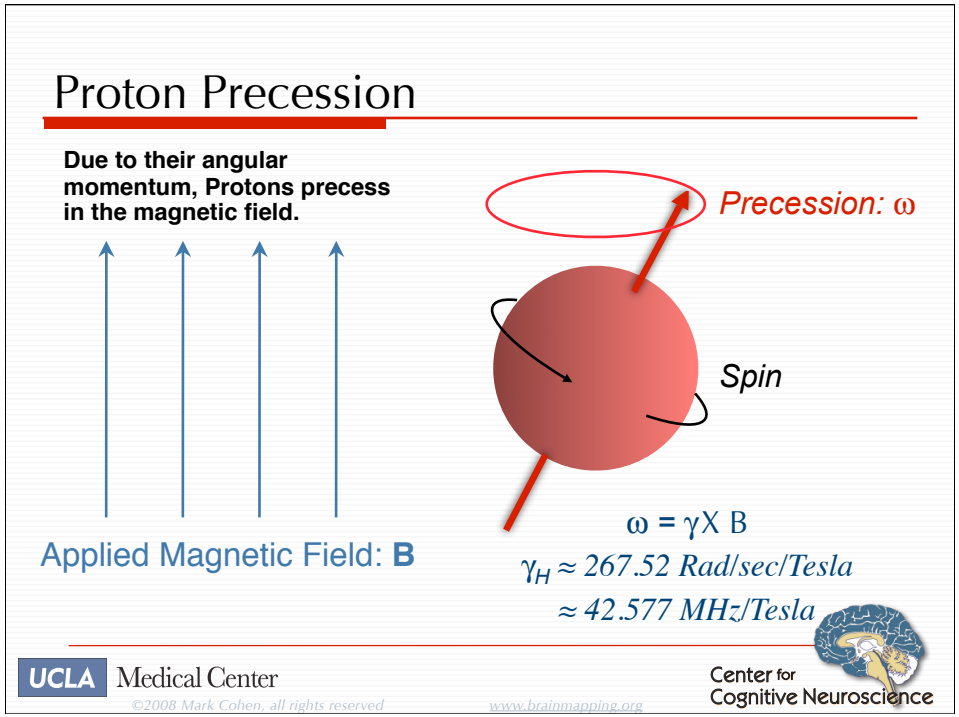

The angular moment of the spins is *conserved*. This means that when an external field is applied, the individual spins will *precess* about it. The rate of precession is known as the Larmor frequency and is proportional to the strength of the external field. The "Larmor constant" or "gyromagnetic ratio" is about 42.577 million cycles/second per Tesla.

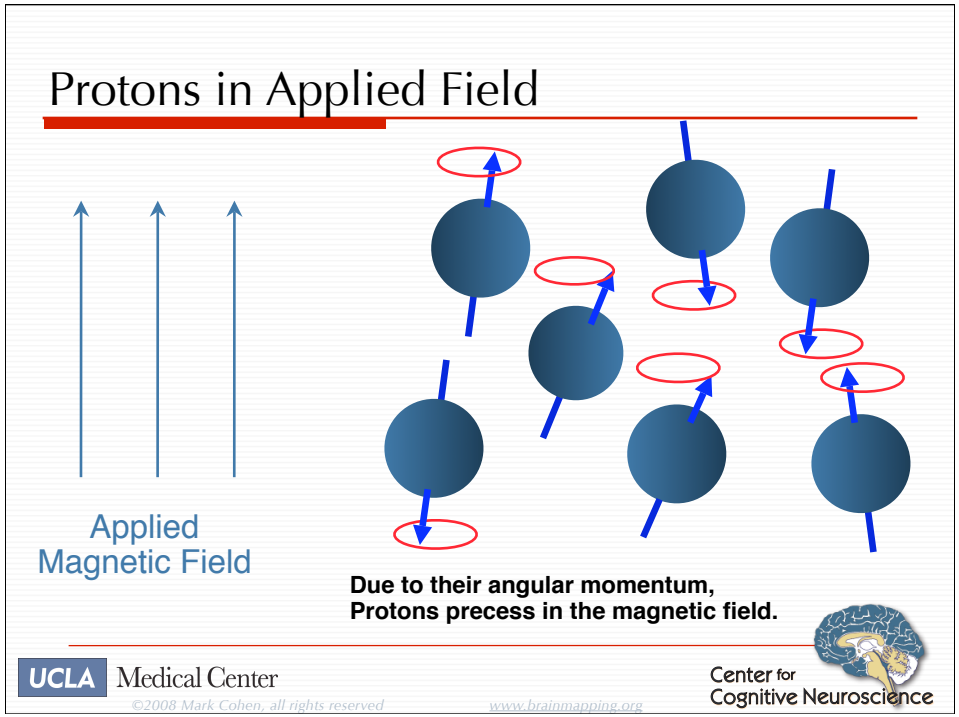

When an ensemble of protons is placed into a magnetic field each of the individual spins adopts the spin up or spin down condition. The precession of the spins also implies that each has a rotational phase. In general, the phases of the individual spins can be expected to be random with respect to one another.

Each spin has a Longitudinal vector component about which it precesses and a Transverse component that rotates about the applied field.

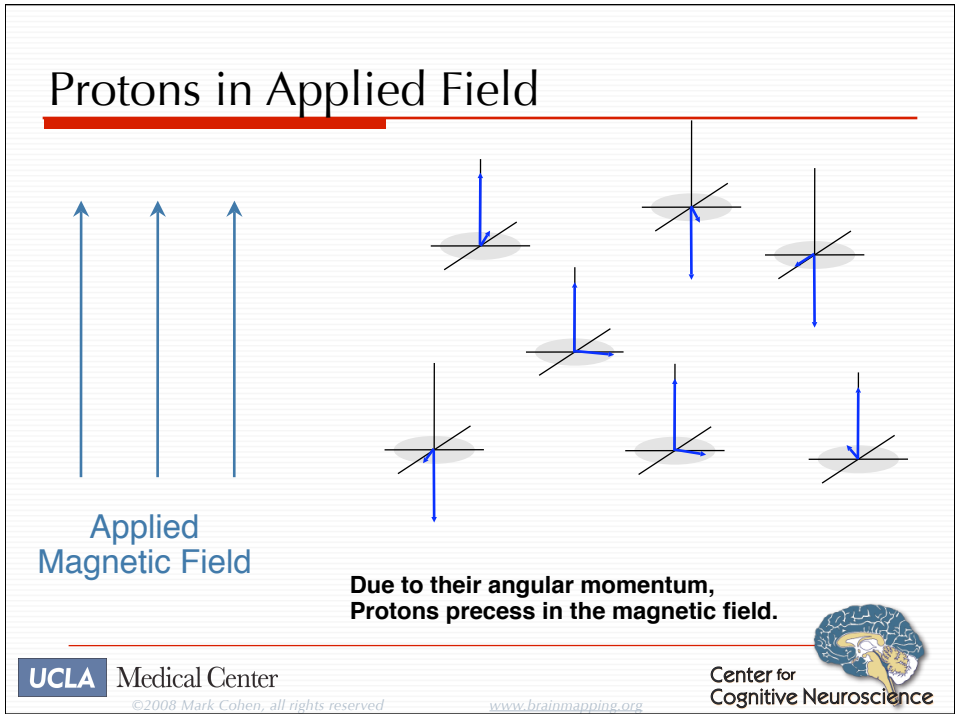

When an ensemble of protons is placed into a magnetic field each of the individual spins adopts the spin up or spin down condition. The precession of the spins also implies that each has a rotational phase. In general, the phases of the individual spins can be expected to be random with respect to one another.

Each spin has a Longitudinal vector component about which it precesses and a Transverse component that rotates about the applied field.

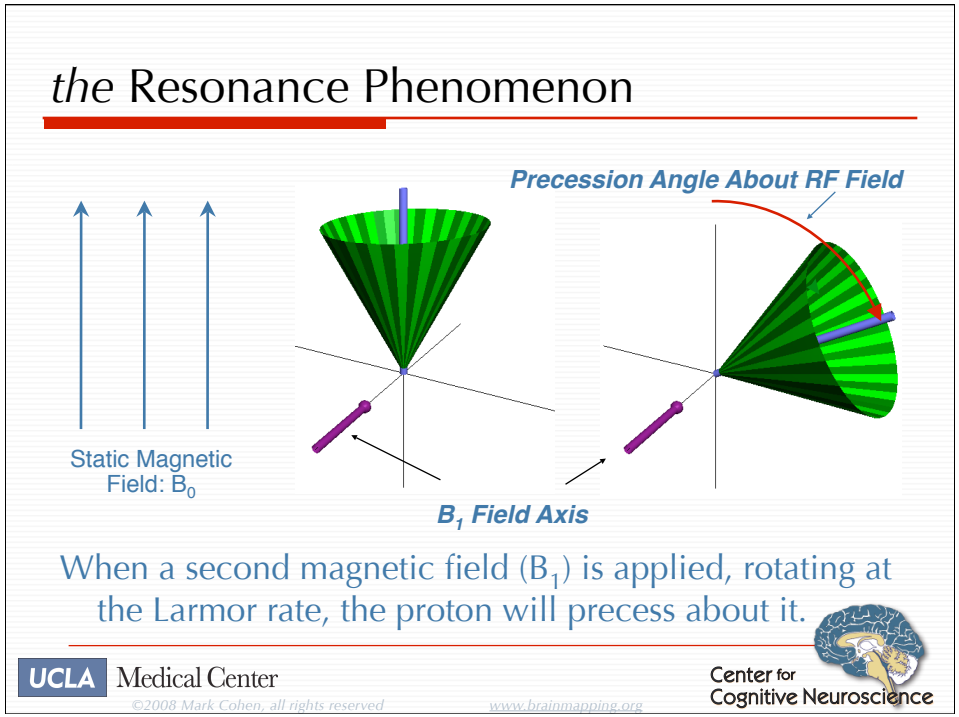

If an additional magnetic field is applied the nuclei will precess about their vector sum. In the special case that the second field, B1, is made to rotate at the same (Larmor) rate as the proton nuclei, the B1 field will seem stationary with respect to any invidividual proton. They will thus precess about B1 in a simple manner. This view of the interaction of B1 and B0 (the static field) is called the rotating frame. In the rotating frame there is no apparent precession about B0.

Because the Larmor rate at reasonable magnetic fields is typically in the tens to hundreds of MegaHertz, and because B1 is applied just long enough to produce the desired precession away from the longitudinal axis, the B1 field is typically called a Radio Frequency, or RF pulse.

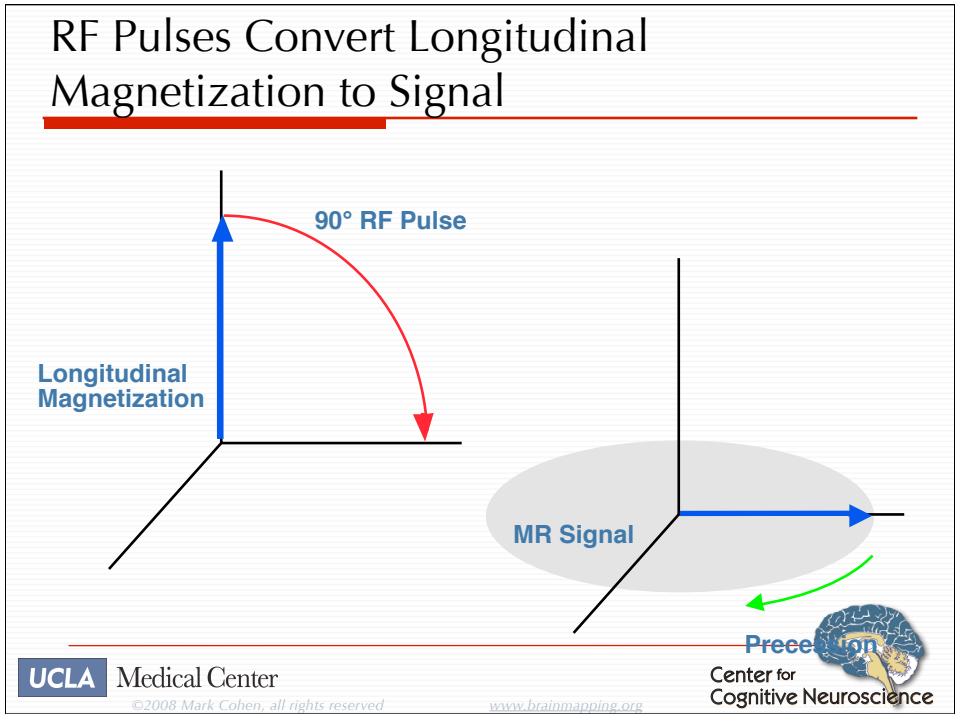

In simple terms, then, the effect of a 90° RF or B1 pulse is to convert any longitudinal magnetization into detectable signal.

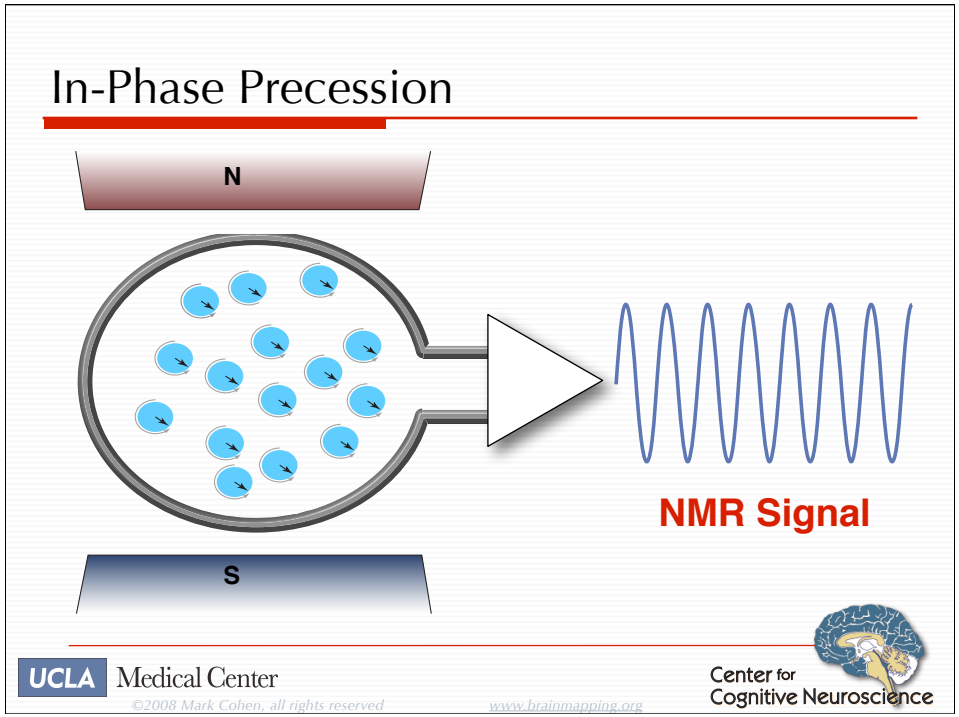

After an RF pulse, the phases of the individual spins will instantaneously be the same - that is, the spins will be precessing “in-phase” and their transverse components will add. This gives rise to a rotating magnetic field that can be detected easily because a time-varying magnetic field creates an equivalent electrical field. If a conductor (antenna) is placed in the vicinity, the electrical field will create a current that can be amplified. Generally, the signal will be a sinusoid at the Larmor frequency.

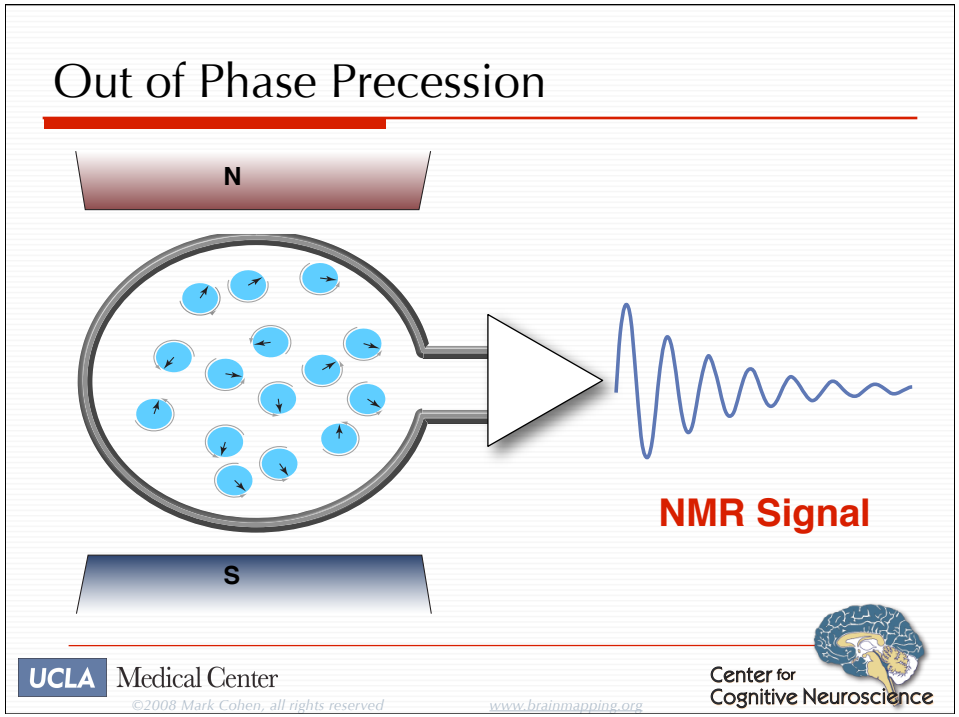

©2008 Mark Cohen, all rights reserved
www.brainmapping.org

The MR signal, once formed, decays quickly. This is because the spins will precess in phase only to the extent that they precess at precisely the same frequency. Even tiny variations in field cause the dephasing to take place over a couple of tenths of a second. These field variations arise from many factor. For example, we cannot manufacture a magnet field that is homogeneous to more than a few parts per million. Further, the individual spins affect one another, causing local inhomogeneities, and different tissues magnetize to different degrees.

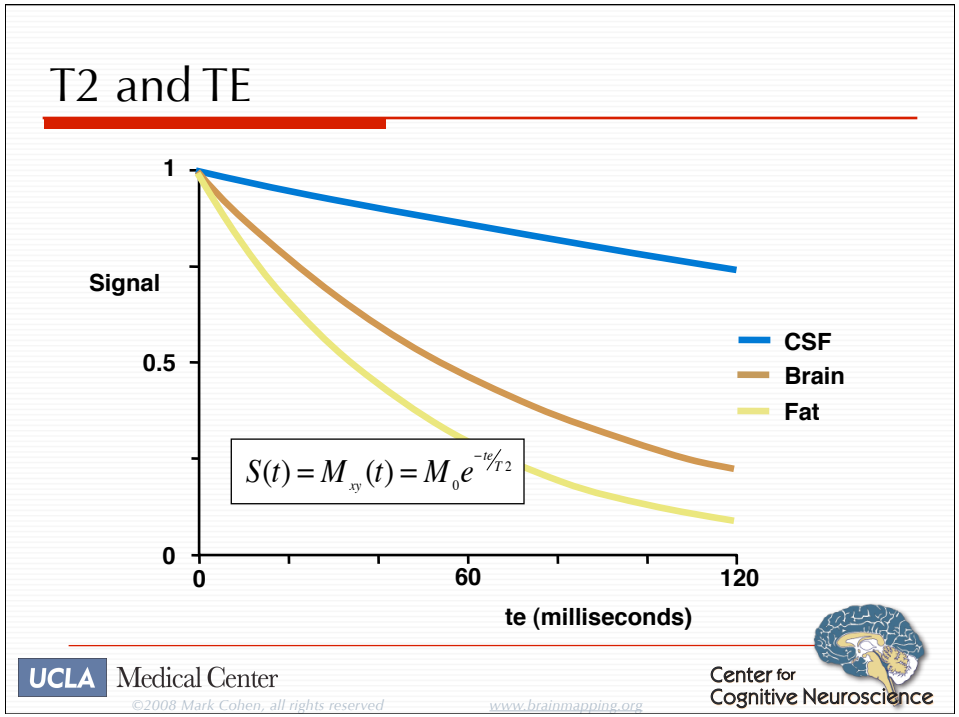

The MR signal decays exponentially as a result of spin dephasing. The decay is governed by a time constant, T2, that is tissue specific. In the body, the T2 ranges from a few hundred microseconds in bone to hundreds of milliseconds in water (CSF, urine). The adjustable parameter, te, specifies the time after signal excitation at which the data are collected. Longer te's result in less signal overall, but also in increased contrast between tissues with differing T2

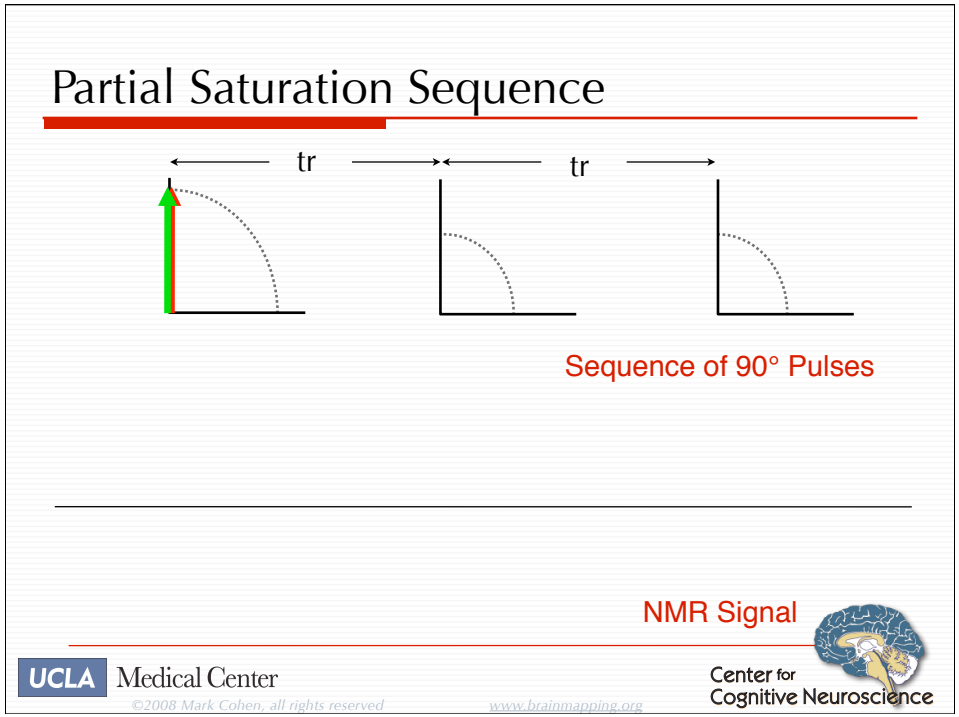

Repeating the excitation and data collection at a fixed interval, tr, introduces contrast based on T1 differences. This is because if the tr is short compared to T1, the spins do not regain their equilibrium state. Each time that a 90° pulse is applied, all of the longitudinal magnetization is lost, and starts again from zero. If a tissue has a short T1, more magnetization is regained, and more signal is created with each 90° pulse. The signal from longer T1 tissues is weaker, however.

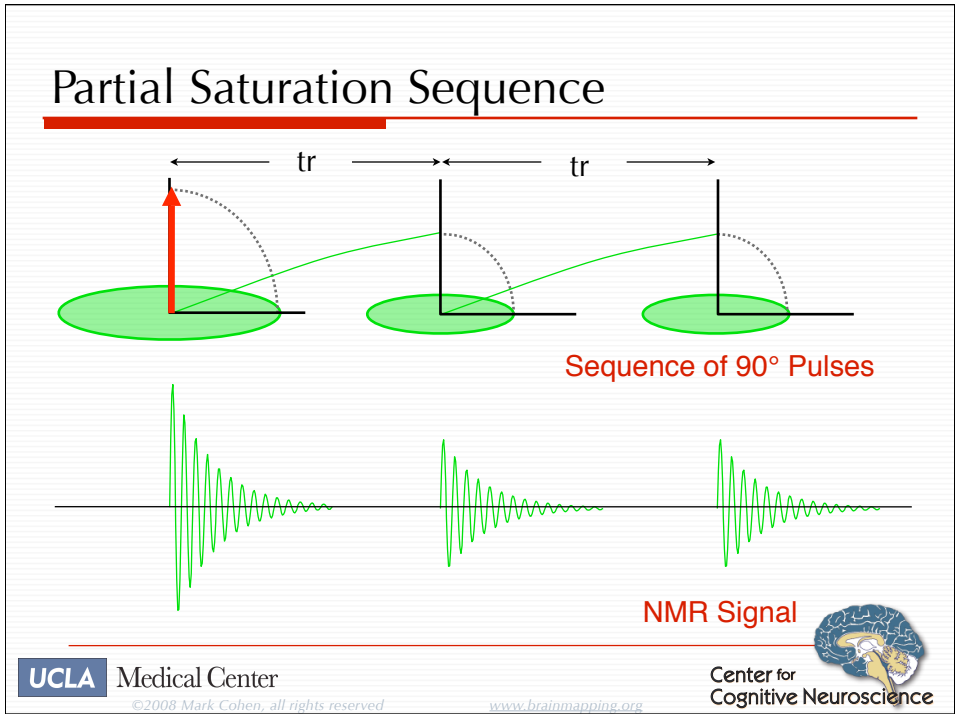

Repeating the excitation and data collection at a fixed interval, tr, introduces contrast based on T1 differences. This is because if the tr is short compared to T1, the spins do not regain their equilibrium state. Each time that a 90° pulse is applied, all of the longitudinal magnetization is lost, and starts again from zero. If a tissue has a short T1, more magnetization is regained, and more signal is created with each 90° pulse. The signal from longer T1 tissues is weaker, however.

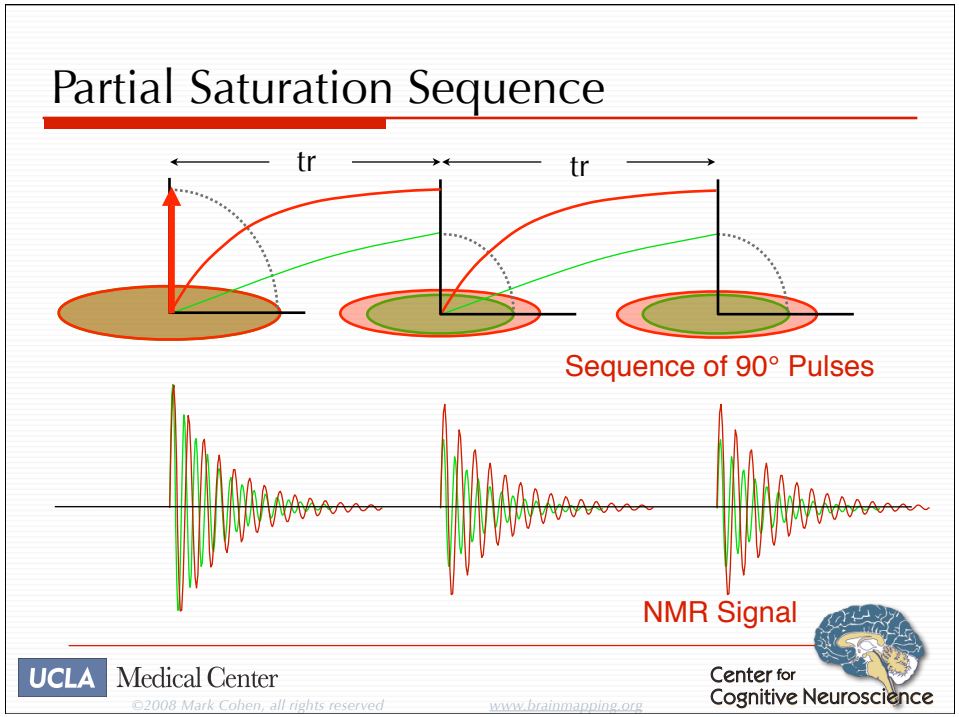

Repeating the excitation and data collection at a fixed interval, tr, introduces contrast based on T1 differences. This is because if the tr is short compared to T1, the spins do not regain their equilibrium state. Each time that a 90° pulse is applied, all of the longitudinal magnetization is lost, and starts again from zero. If a tissue has a short T1, more magnetization is regained, and more signal is created with each 90° pulse. The signal from longer T1 tissues is weaker, however.

Effects of TE at long TR

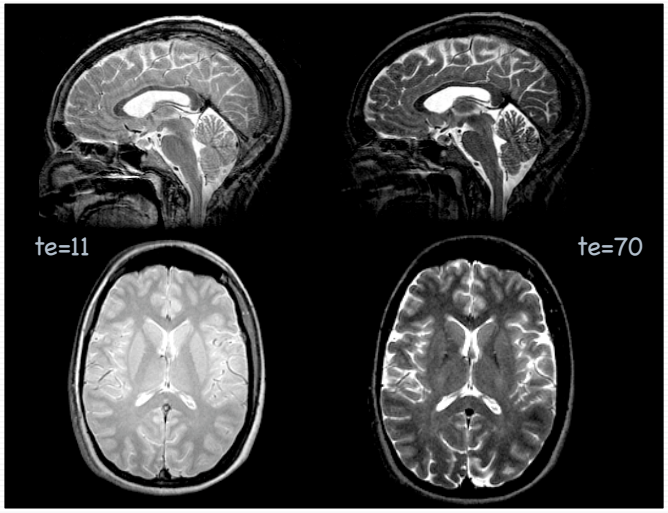

UCLA Medical Center

©2008 Mark Cohen, all rights reserved

www.brainmapping.org

Center for Cognitive Neuroscience

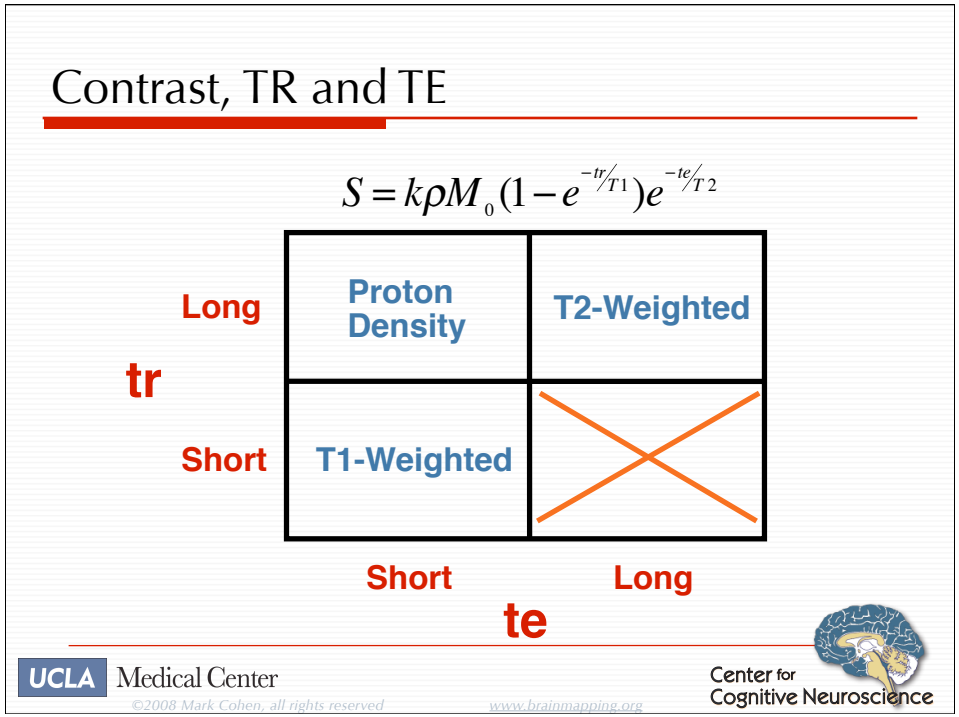

This chart suggests the four principle variations in contrast used in MR imaging. If we consider long tr to be tr >> T1, the images will have little T1 contrast, as all tissues will magnetize almost fully. If tr << T1 the images will have T1 contrast. Likewise, te << T2 minimizes the contrast differences from T2 and te >> T2 maximizes T2 contrast. The equation shows this in analytic form. When tr is long, and te is short, the major contrast determinant would be the density of protons per unit volume. Images with mixed T1 and T2 contrast are typically poor, as shown in the next slide.

Contrast, TR and TE

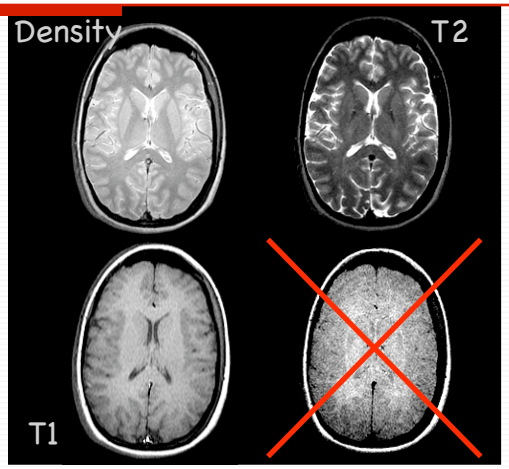

Long

TR

Short

Short TE Long

UCLA Medical Center
©2008 Mark Cohen, all rights reserved
www.brainmapping.org
Center for Cognitive Neuroscience

the Observed Transverse Relaxation Rate, T2*, is the sum of several components:

$$\frac{1}{T2^*} = \frac{1}{T2} + \frac{1}{T2'} + \frac{1}{T2_D} + \ldots$$

Molecular

Field Inhomogeneity

Diffusion

Many physical effects result in dephasing. The process that dephases the spins most quickly dominates the contrast in the final images, thus the net observed T2 (T2*, “T2-star”) is shorter than any of the individual T2’s. Most MR imaging is concerned chiefly with three dephasing effects: molecular interactions (T2), local variations in field from, for example, tissue boundaries (T2’) and dephasing from molecular motion from diffusion (T2D).

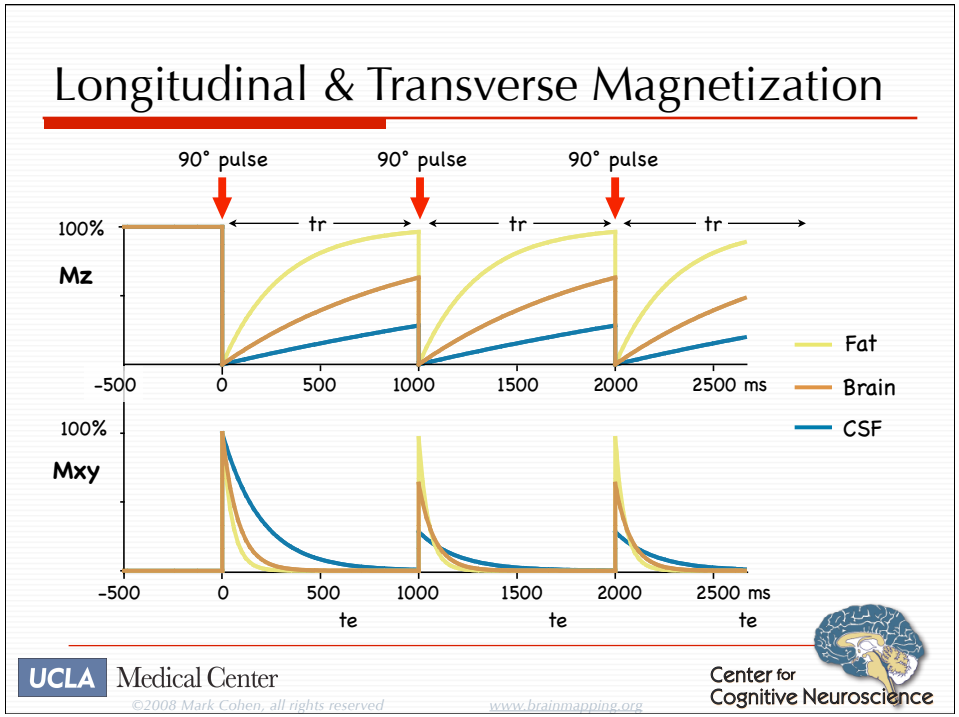

This graph shows the combined effects of T1 and T2 on the longitudinal and transverse magnetization during repeated excitation pulses at fixed tr. This model was made with a tr of 1 second, and with typical T1 and T2 of head tissues. Notice that the signals for brain, CSF and fat all cross over with a te of about 70 msec at this tr. The result is that there is no contrast between these tissues. If te is made very short the T1 effects dominate. If te is made very long, the T2 effects dominate, but the signal will be VERY weak.

Summary Animation

UCLA Medical Center

©2008 Mark Cohen, all rights reserved

www.brainmapping.org

Center for Cognitive Neuroscience

Summary Animation

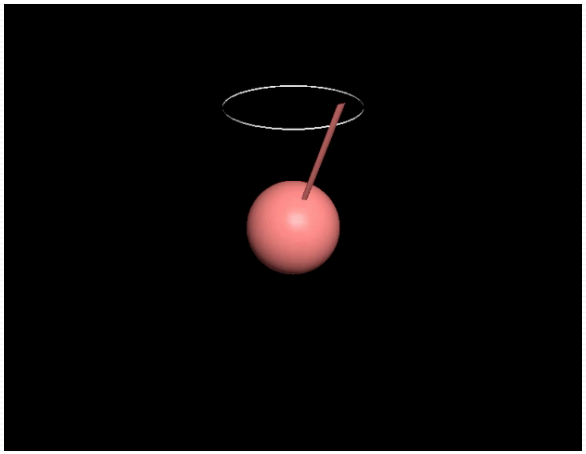

UCLA Medical Center

©2008 Mark Cohen, all rights reserved

www.brainmapping.org

Center for Cognitive Neuroscience

Sample Data Set (normal)

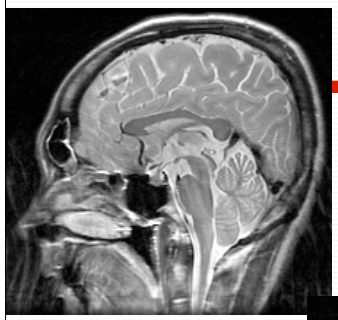

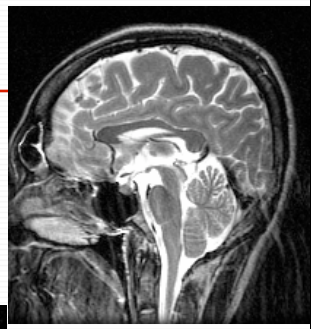

Fast Spin Echo
3 mm Slices

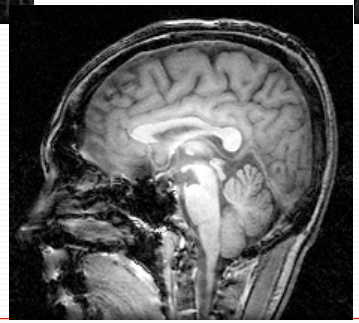

3D IR-SPGR
TE = 3.2, TI = 700

Sample Data Set (normal)

UCLA Medical Center

©2008 Mark Cohen, all rights reserved

www.brainmapping.org

Center for Cognitive Neuroscience

Contrast to Noise Ratio

(Gray-White)

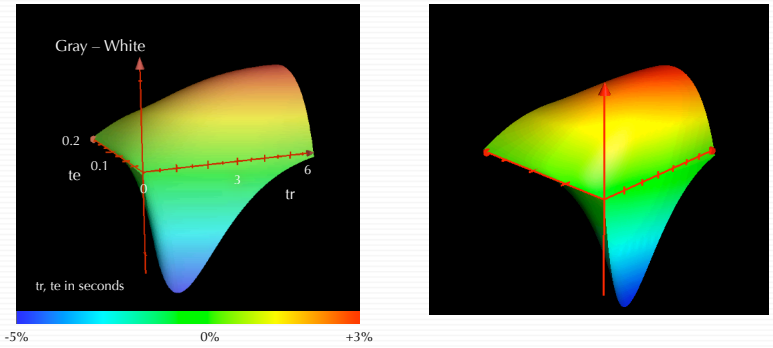

$$Contrast = [(1 - e^{-tr/1.2})e^{-te/.08}],\ gray\ matter$$
$$-[(1 - e^{-tr/1.0})e^{-te/.07}],\ white\ matter$$

UCLA Medical Center

©2008 Mark Cohen, all rights reserved

www.brainmapping.org

Center for Cognitive Neuroscience

The graphs above plot the difference in signal (the “contrast”) between grey and white matter as a function of tr and te. Notice that there is a broad range of combined tr and te where the signals will be isointense (no contrast). The blue regions of the graphs are regions of high T1 contrast, and the red regions have high T2 contrast.

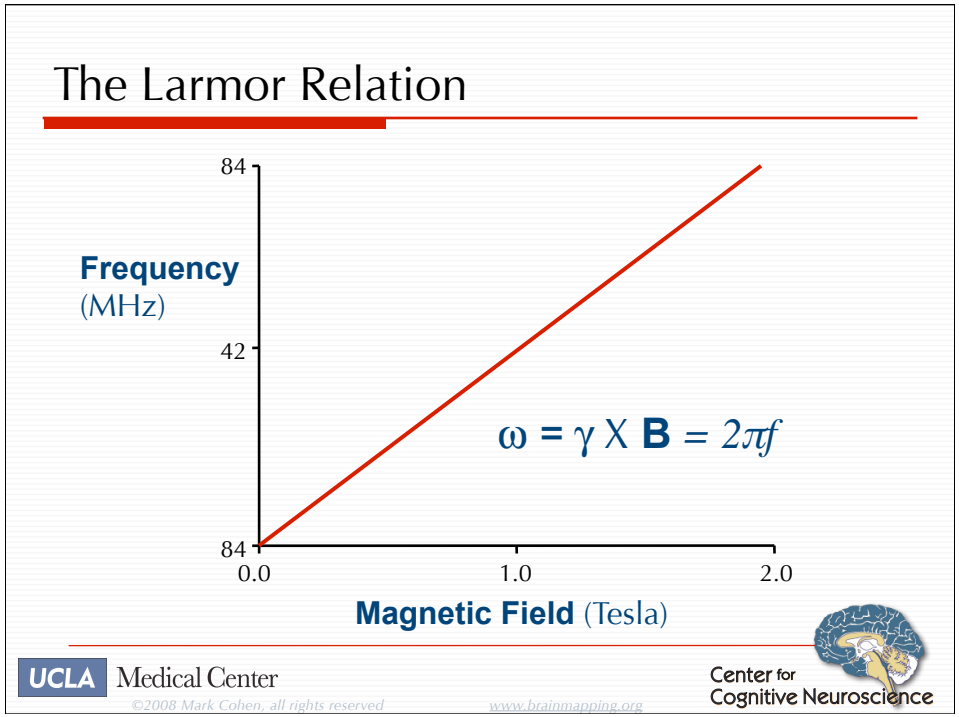

Given what we know about signal intensity and contrast, our next step is to make an image. This implies finding the location of the signal in 3 dimensional space. To do this we rely on the Larmor relation that the precession frequency is proportional to magnetic field. The trick will be to cause the magnetic field to vary with position.

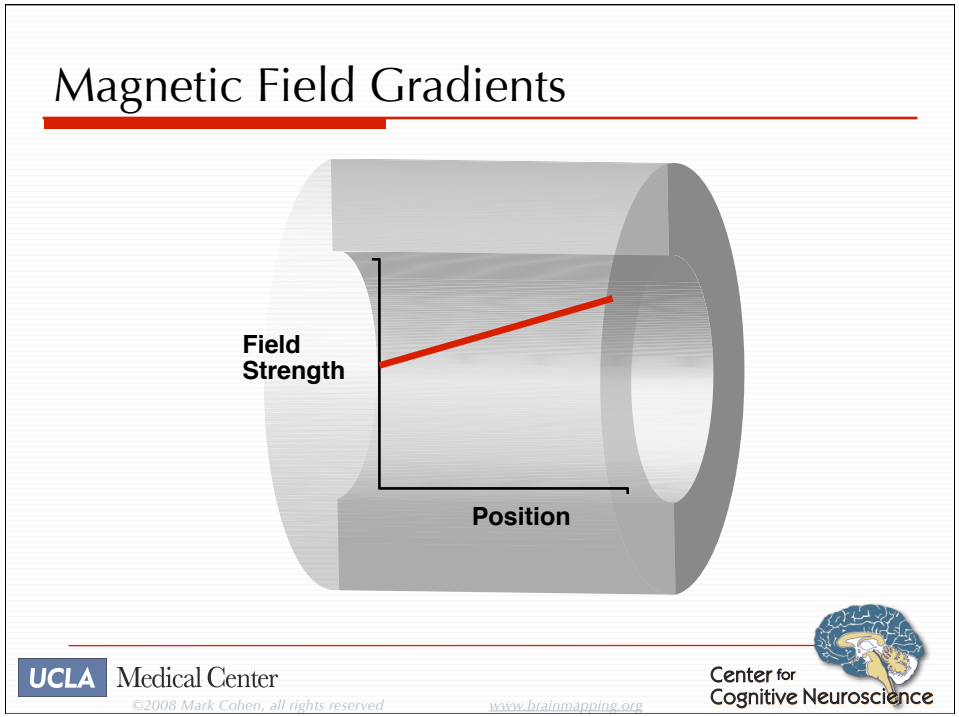

We create spatially varying (“gradient”) magnetic fields within the bore of the imaging magnet through the use of physical “gradient” coils. The spatial variations are typically small compared to the strong imaging magnetic field. These cause the frequency.

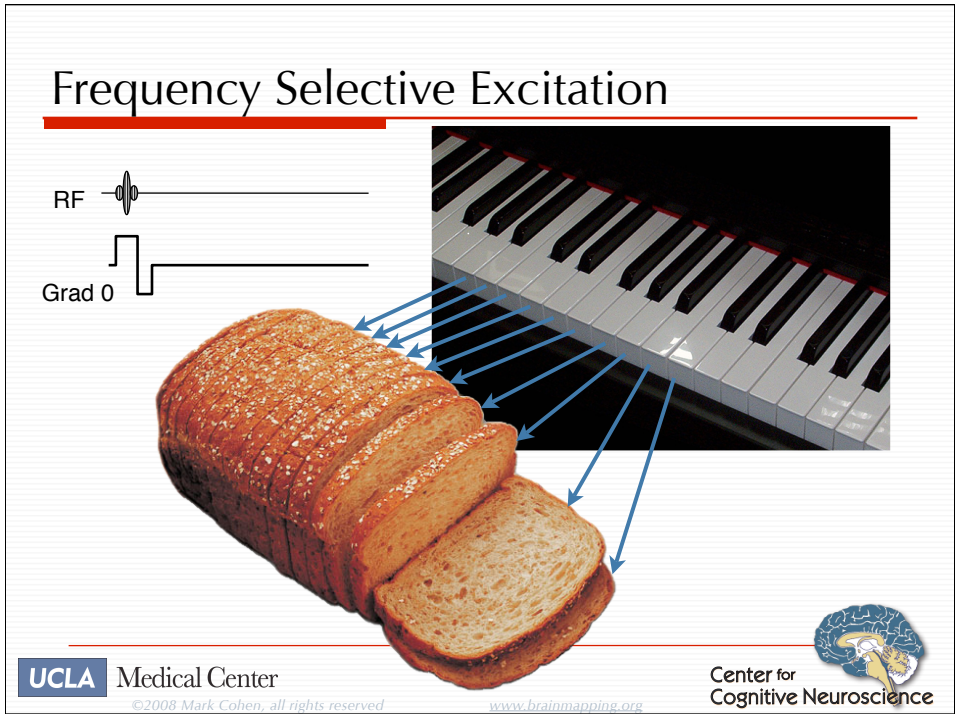

Slice selective excitation is an invention of Nobel-laureate, Sir Peter Mansfield. He noted that if an excitation RF pulse was applied at a fixed frequency in the presence of a gradient field, only spins at the corresponding location in the field would be affected by the pulse and become rotated into the transverse plane (excited). The position corresponds to frequency and can be selected easily.

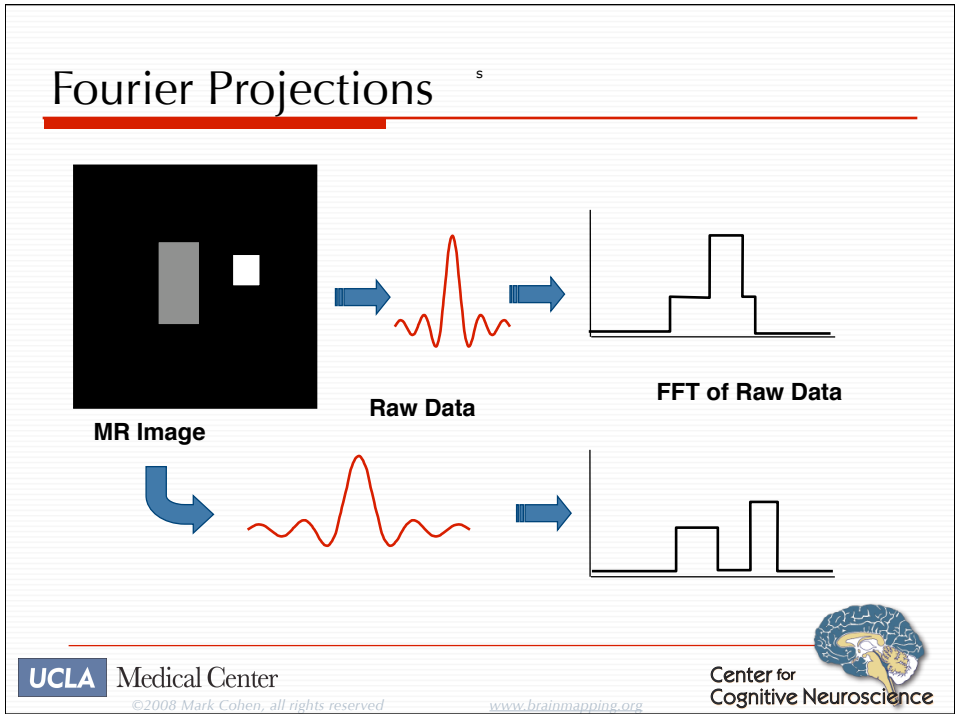

Afer a slice has been excited, we can apply a field gradient in a different direction, causing the spin frequencies to vary along this second axis. The Fourier transform of the signal represents calculates the intensity of the signal at each frequency. As frequency depends on position, we can interpret this as the intensity as a function of location along the same axis as the gradient field was applied. We can readily apply gradient fields along any axis, and thus collect these “projections” in any orientation.

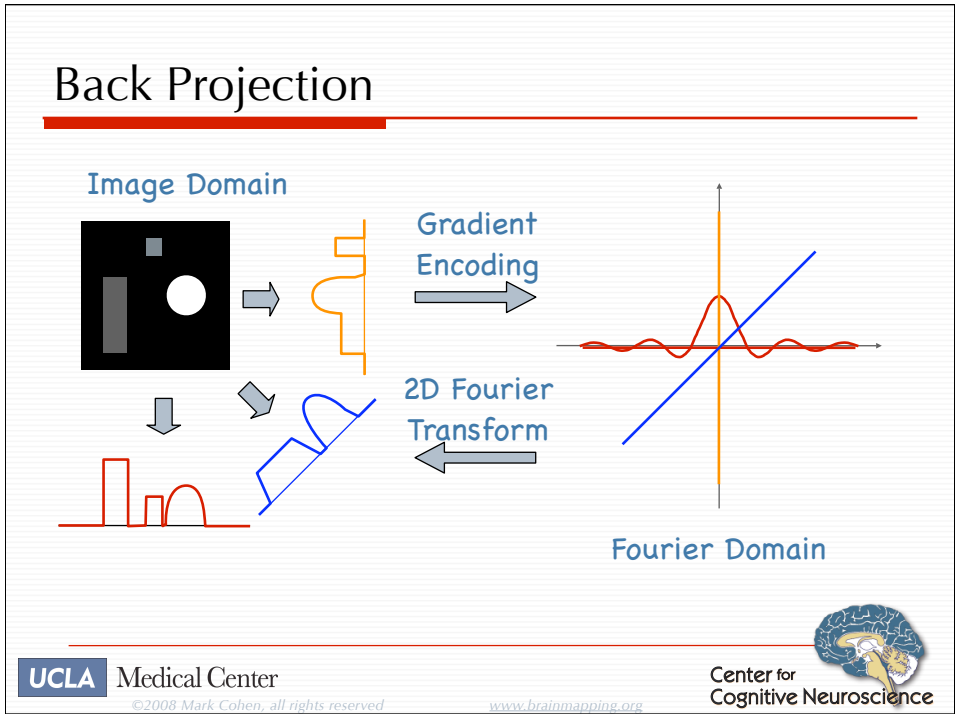

Here we show the projected and transformed signal along three different axes. We can project the ‘untransformed’ data lines into a raw data space according to the direction of the gradient field. The combination of all of these radial projections can looks a bit like rings of waves in a puddle. The 2 dimensional Fourier transform of this combined raw data is the image. A 2D Fourier transform amounts simply to calculating the Fourier transform of the data set line by line, then calculating the Fourier trasform of each column in this intermediate data set.

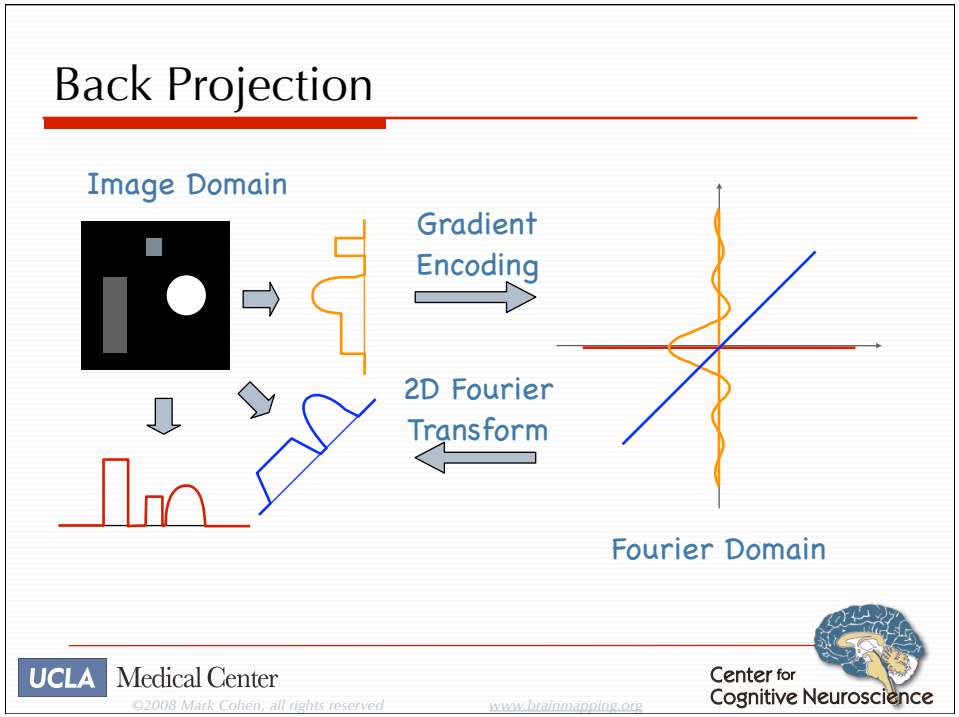

Here we show the projected and transformed signal along three different axes. We can project the ‘untransformed’ data lines into a raw data space according to the direction of the gradient field. The combination of all of these radial projections can looks a bit like rings of waves in a puddle. The 2 dimensional Fourier transform of this combined raw data is the image. A 2D Fourier transform amounts simply to calculating the Fourier transform of the data set line by line, then calculating the Fourier trasform of each column in this intermediate data set.

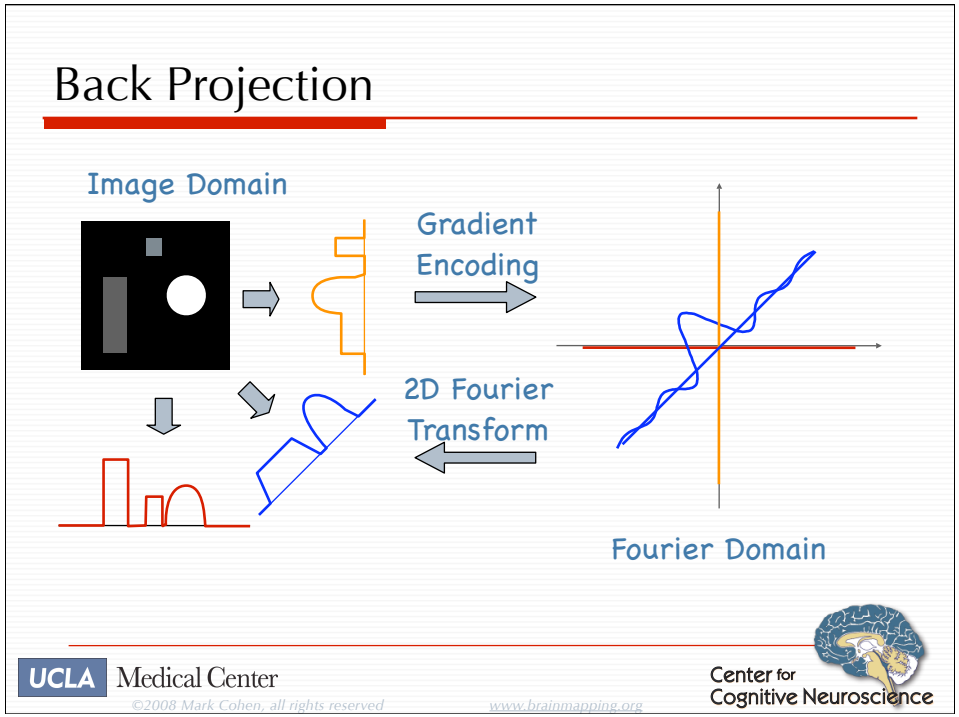

Here we show the projected and transformed signal along three different axes. We can project the ‘untransformed’ data lines into a raw data space according to the direction of the gradient field. The combination of all of these radial projections can looks a bit like rings of waves in a puddle. The 2 dimensional Fourier transform of this combined raw data is the image. A 2D Fourier transform amounts simply to calculating the Fourier transform of the data set line by line, then calculating the Fourier trasform of each column in this intermediate data set.

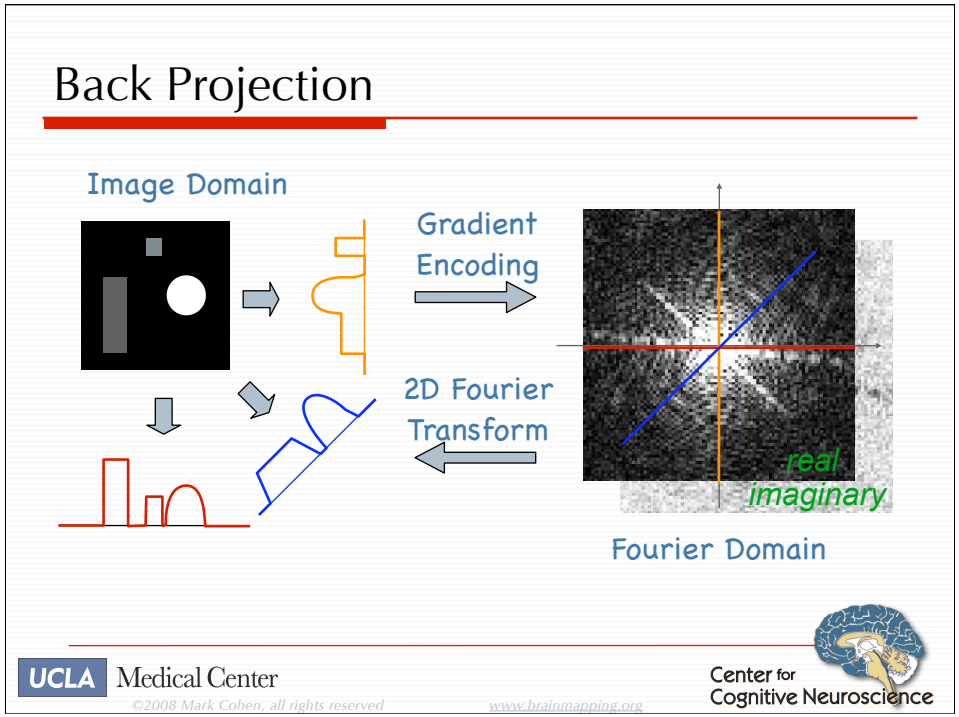

Here we show the projected and transformed signal along three different axes. We can project the ‘untransformed’ data lines into a raw data space according to the direction of the gradient field. The combination of all of these radial projections can looks a bit like rings of waves in a puddle. The 2 dimensional Fourier transform of this combined raw data is the image. A 2D Fourier transform amounts simply to calculating the Fourier transform of the data set line by line, then calculating the Fourier trasform of each column in this intermediate data set.

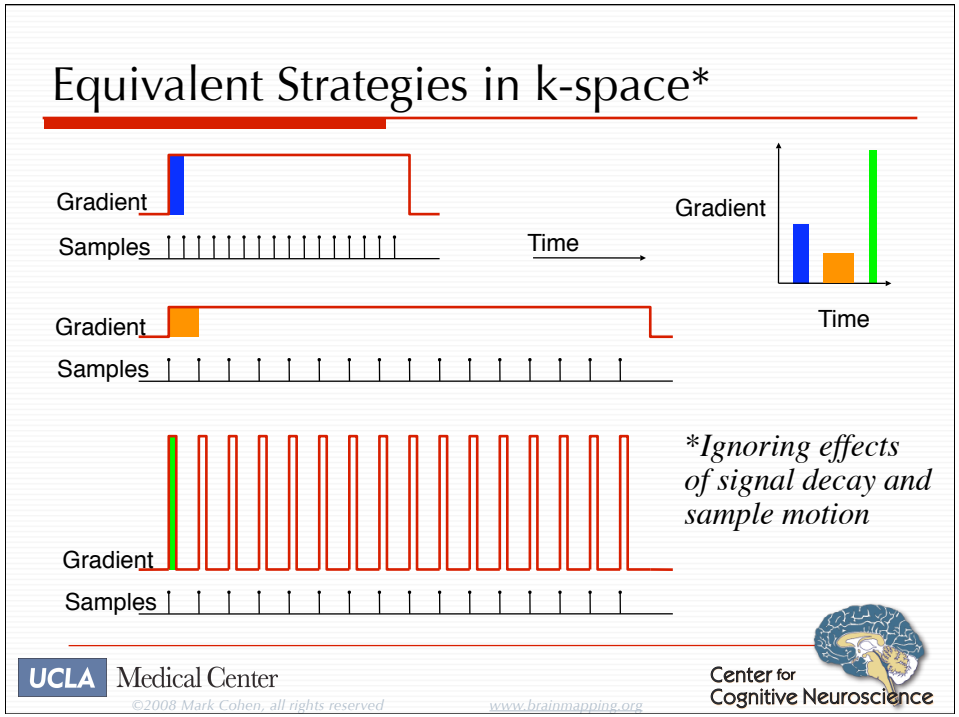

Crucially, if we ignore all of the contrast effects considered above, the only thing that causes the intensity of the MR signal to change is the effect of the applied magnetic field gradients. While the signal evolves in a continuous manner, in practise, it is sampled discretely and digitized. In the figure, the circles on stems represent sampling time points. In the time between samples (incidentally, this is called the “dwell time”) the differences in precession frequency of spins in different locations causes a phase difference to accumulate as a function of the product of the frequency difference between these locations and the dwell time. Thus, leaving the gradients on at a high amplitude and short duration is equivalent to leaving the gradients on for a longer time at a shorter amplitude. From the perspective of the encoding difference between each sample time point there is no difference. In fact, we don’t have to leave the gradients on for the duration of the dwell interval and instead, we may pulse the gradients at very high amplitude and get the same effect.

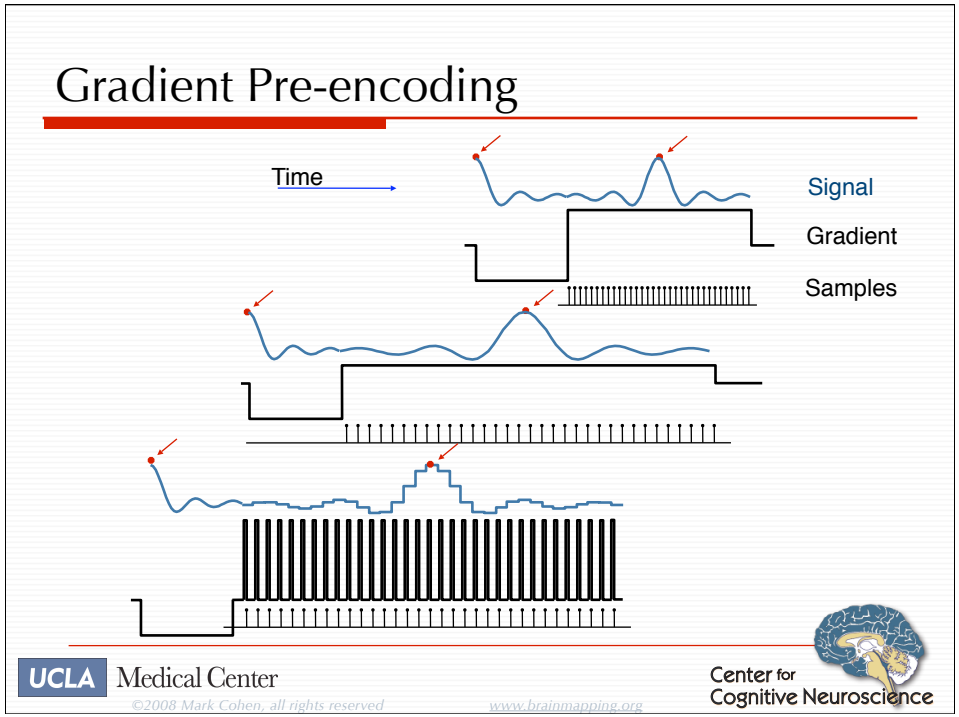

The pulsed gradient method (bottom) causes the signal to make discrete jumps in intensity. If the gradients are left on continuously, but the signal is sampled discretely, we collect identical data. Note in these figures that the spins are typically “pre-encoded” by applying the field gradients in their opposite sense, before collected the signal. This ensures that all of the spins will be in-phase at the center of the readout period.

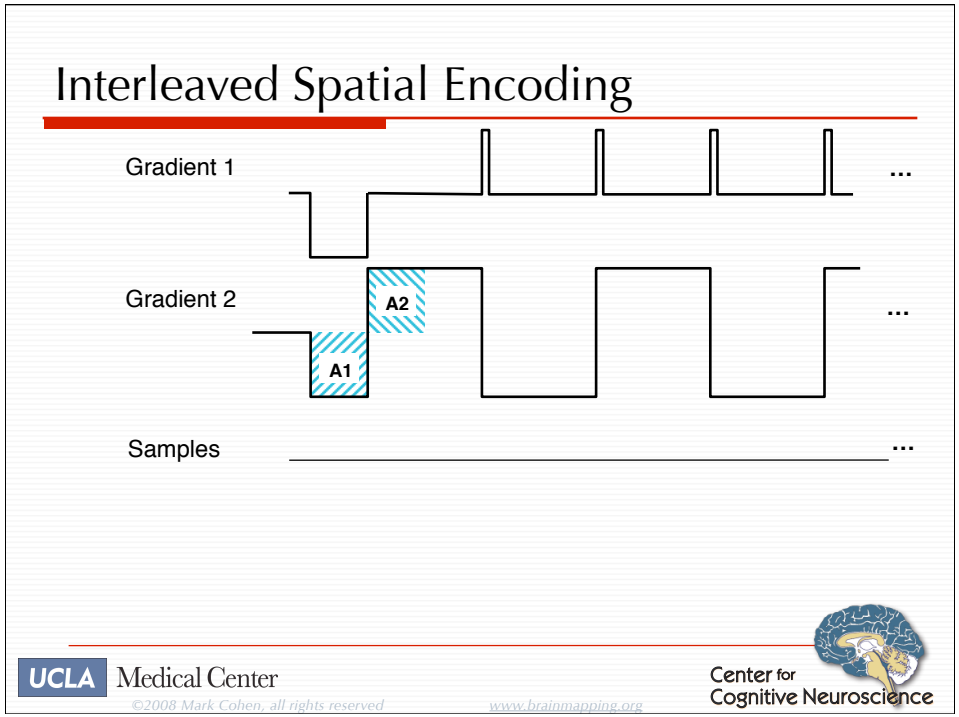

As a gradient is applied for a time period (A1) the spins in different location go out of phase with respect to one another. When the gradient is applied for the same duration in the opposite sense - reversing which side of the instrument is at higher field and which is lower - the spins go back into phase.

Considering only Gradient 2 above, at the time points indicated in black, the spins are always in phase. Alternating the sign of Gradient 2 makes it invisible at these time points. However, if an orthogonal gradient, Gradient 1 is applied between these oscillations, it will cause a phase accumulation along the Gradient 1 direction.

For example, if Gradient 1 is directed along the vertical (Y) axis of the magnet, the time points indicated in balck represent a line with spatial encoding along this axis only. The Fourier transform of this, of course, is a signal intensity projection along the Y axis. The points indicated in blue are affected by both gradients and fill in the other points in the 2D raw data space.

In this way, we can interleave the spatial encoding in the X and Y axes creating a 2D image all at once. This scheme is known as Echo Planar Imaging or EPI and is the fastest practical method to form MR images at this time.

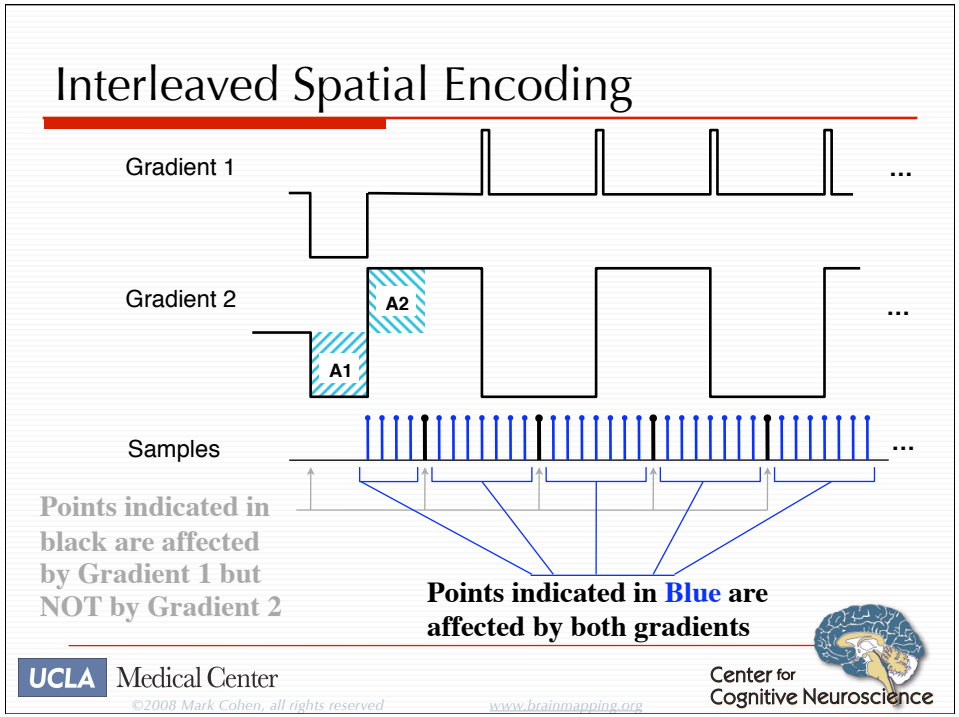

As a gradient is applied for a time period (A1) the spins in different location go out of phase with respect to one another. When the gradient is applied for the same duration in the opposite sense - reversing which side of the instrument is at higher field and which is lower - the spins go back into phase.

Considering only Gradient 2 above, at the time points indicated in black, the spins are always in phase. Alternating the sign of Gradient 2 makes it invisible at these time points. However, if an orthogonal gradient, Gradient 1 is applied between these oscillations, it will cause a phase accumulation along the Gradient 1 direction.

For example, if Gradient 1 is directed along the vertical (Y) axis of the magnet, the time points indicated in balck represent a line with spatial encoding along this axis only. The Fourier transform of this, of course, is a signal intensity projection along the Y axis. The points indicated in blue are affected by both gradients and fill in the other points in the 2D raw data space.

In this way, we can interleave the spatial encoding in the X and Y axes creating a 2D image all at once. This scheme is known as Echo Planar Imaging or EPI and is the fastest practical method to form MR images at this time.

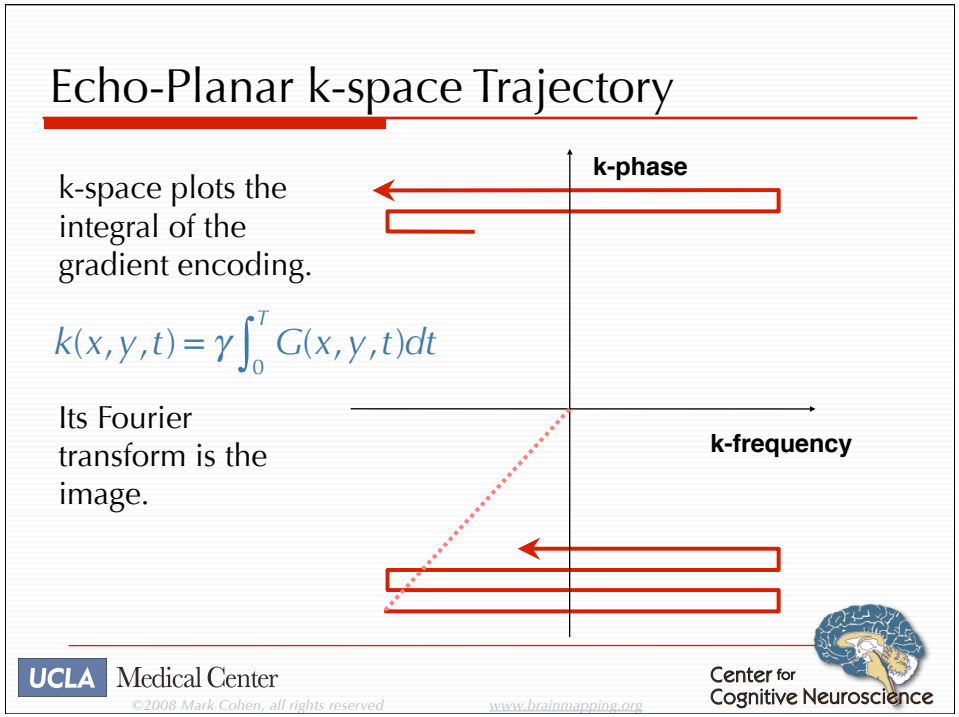

If we examine the locations of the raw data points that are collected in EPI we see that they are collected line-by-line in alternating directions. The brief pulses of Gradient 1 move the locations from one line to the next. The net effect is to make a raster pattern in the raw data space.

This raw data space is known as “k-space” and is the 2D Fourier transform of the image. As long as there are no other things changing the signal, the order in which the data are placed in k-space makes no difference.

Since the amount of spatial encoding between spins is determined by the product of the gradient amplitude and duration we generalize this as the integral of the gradient-time product. This product becomes the phase difference between spin locations; k-space is sometimes referred to as phase space for this reason.

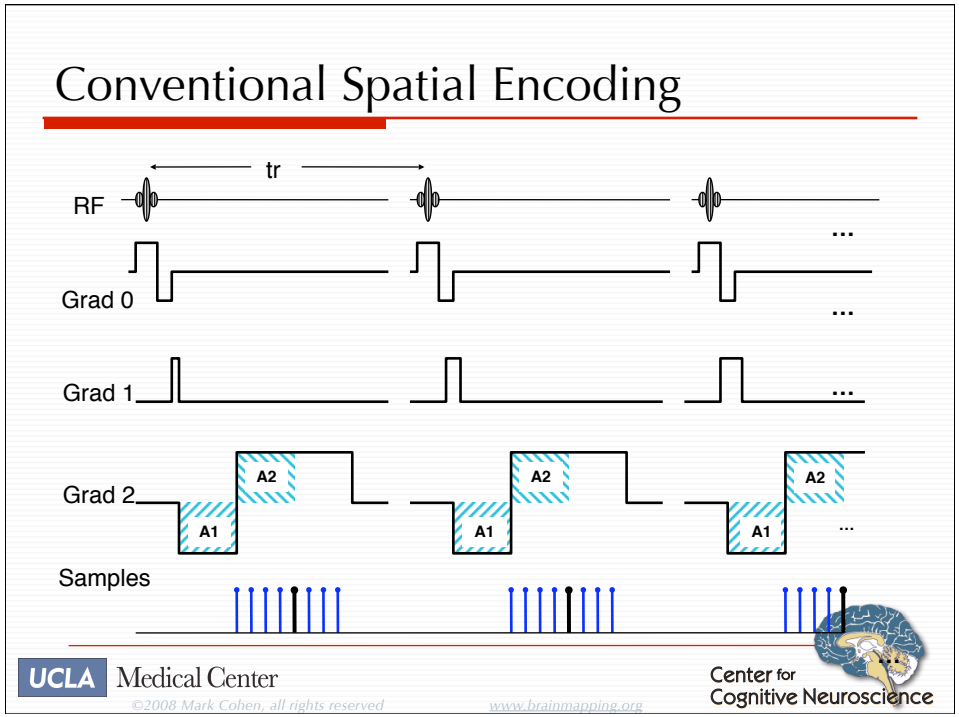

Before echo-planar imaging became possible, a different encoding scheme was used that collected only one line of k-space with each tr. The position along the y-axis of k-space is established by applying a brief pulse of the Y-gradient (Grad 1) before collecting the data in the presence of the X-gradient. With each tr, the area of the Y-gradient pulse is increased. The diagram above also shows the selective excitation pulse, wherein an RF pulse is transmitted in the presence of a gradient directed along the z-axis (Grad 0). Patterns of this kind that indicate the timing of events in MRI are known as pulse sequences.

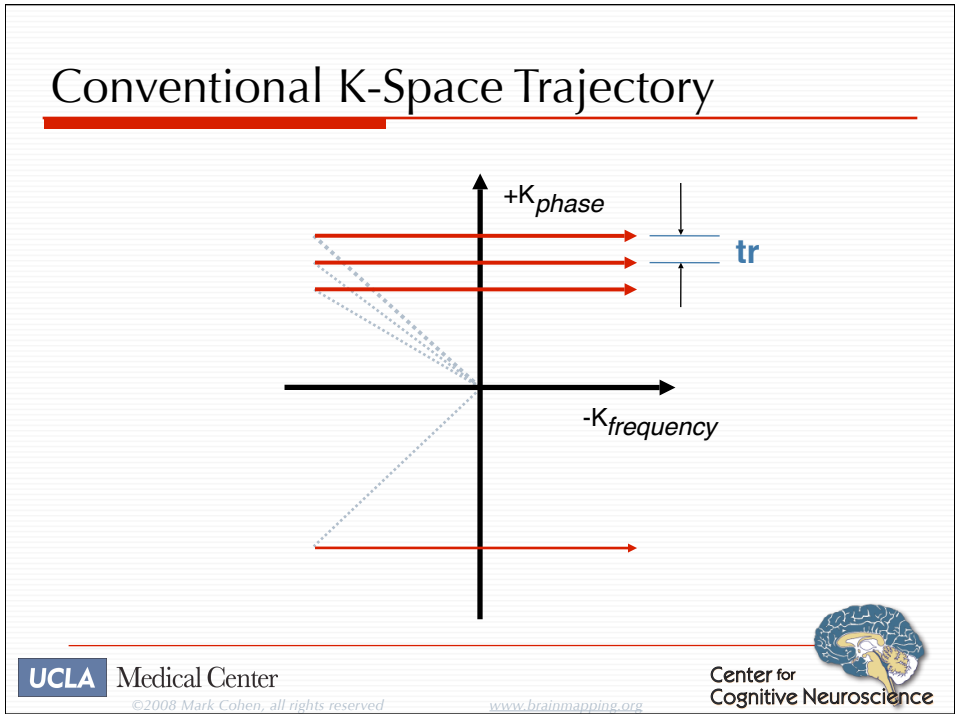

The traditional line-by-line manner of forming MR images results in the k-space filling pattern shown above. Note that time does not appear explicitly in the k-space diagram. In the conventional imaging case, the total imaging time is determined by the number of lines of resolution in the Y or phase encoding axis and by the time between lines, which is tr.

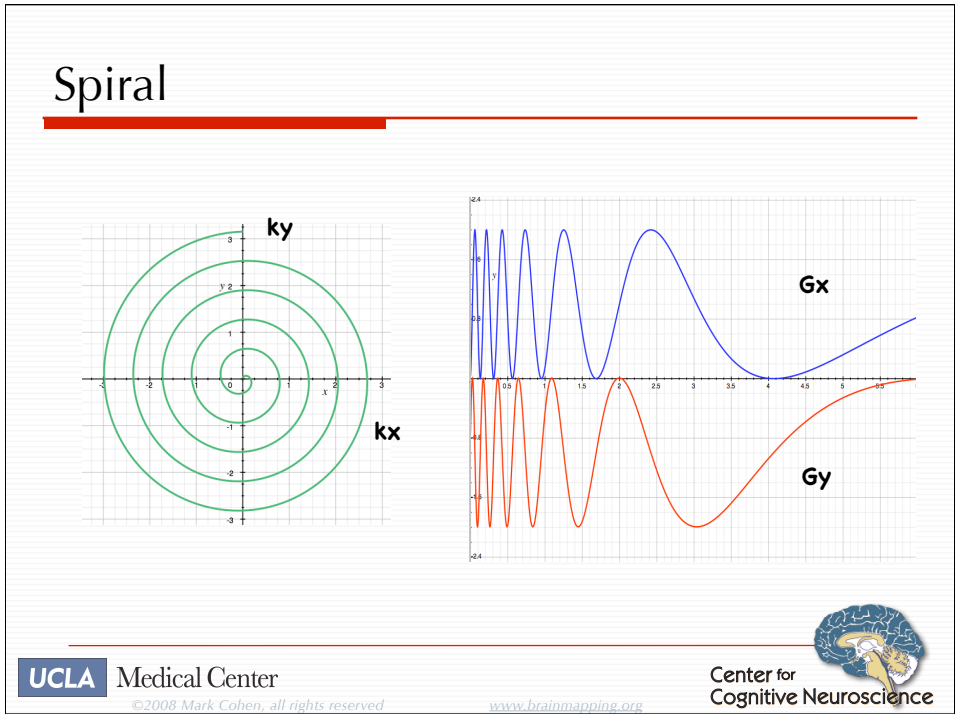

The order in which k-space is populated is somewhat arbitrary. One filling pattern, “Spiral Scanning” has significant popularity. The idea is to make a spiral trajectory in k-space. It is easy to see that this requires a gradient pattern in X and Y that is a slowing sinusoid.

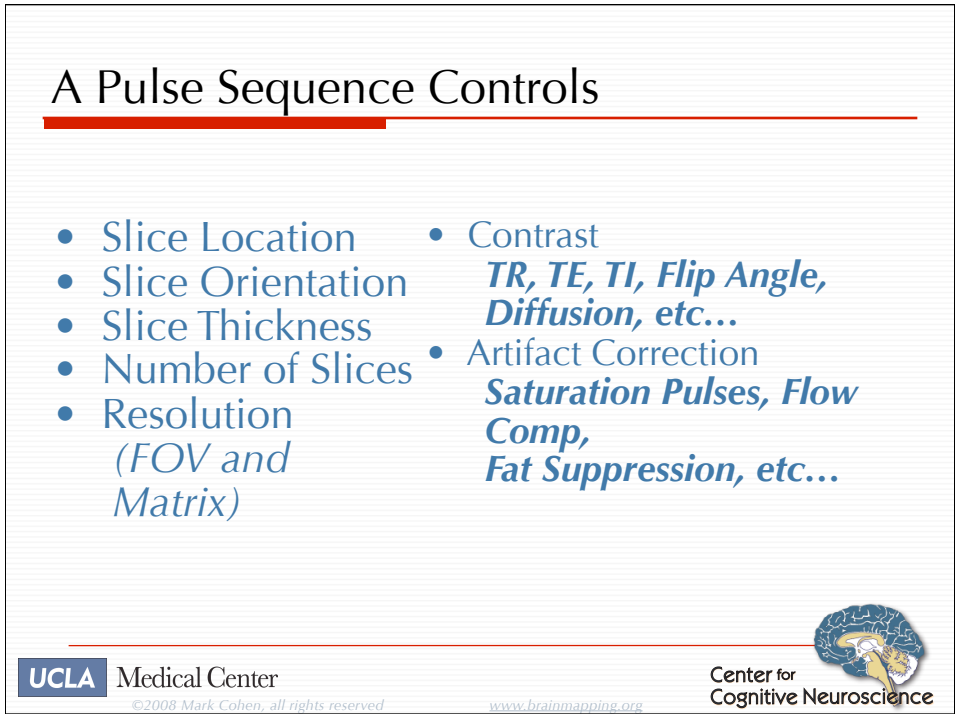

The pulse sequence ultimately controls a host of factors. The user seldom interacts directly with the gradient and RF pulse waveform controls, but instead enters parameters for acquisition such as the slice locations resolution features and contrast controls.

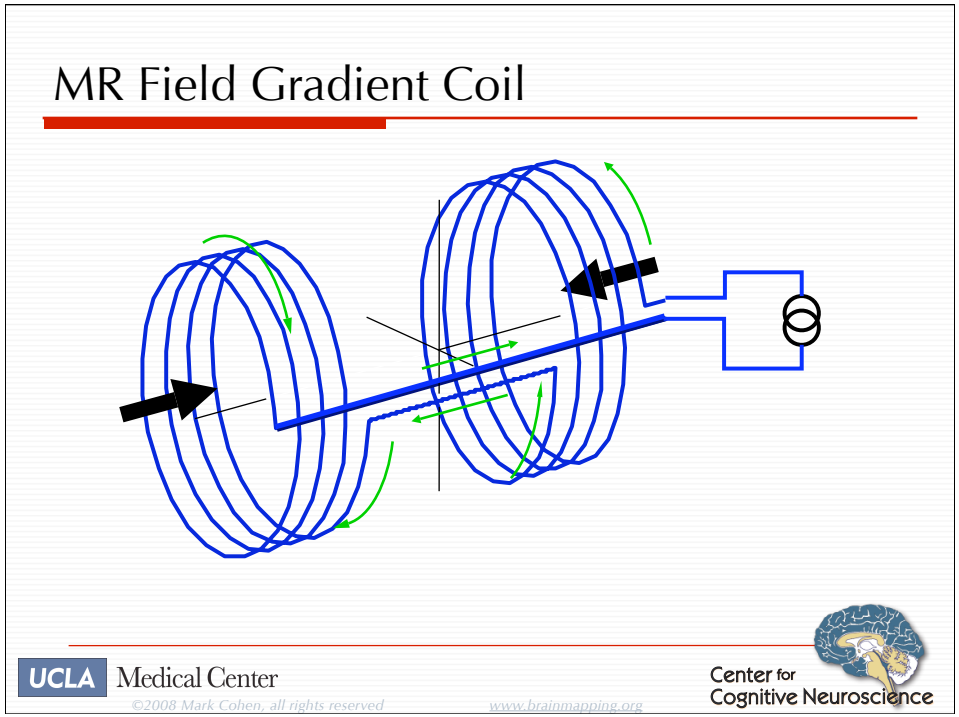

Creating the field gradients requires passing electrical current through large coils (Gradient Coils). In MR imaging systems several hundred Ampere's of current are required to create large enough gradients. The pulse sequences generally require that the gradients are turned on and off as rapidly as possible. Doing so requires very high Voltages because of the coil inductance.

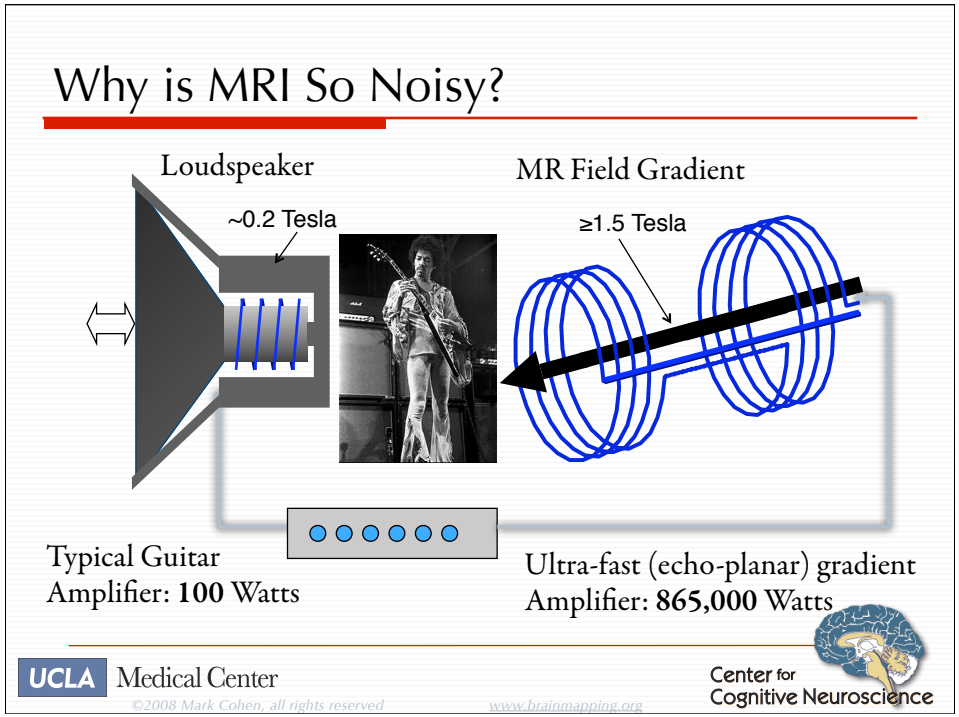

The construction of a gradient coil is very similar to the construction of a loudspeaker. Both utilize a coil of wire to produce a time-varying magnetic field in the presence of a large stationary field. The magnetic field in the coil experiences forces against the stationary field that tend to make it move. However, the magnitudes of all of the forces are much greater in MRI. Despite great efforts to reduce it, MRI is necessarily extremely noisy.

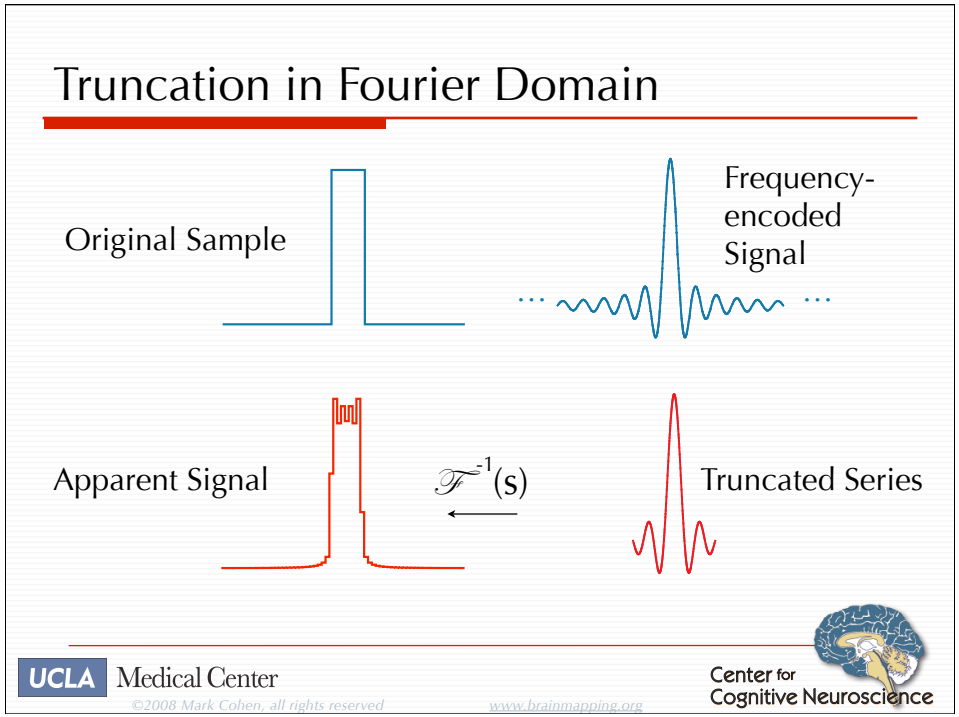

Remembering that the MRI raw data encode position by frequency, it is useful to consider factors that control the fidelity of that encoding. The Fourier transform represents the amplitudes (and phases) of the frequencies in a signal or waveform. Consider a simple rectangular signal as shown at top left. Its representation in frequency is shown at top right. Notably, the transform has infinite tails because the waveform contains energy at all frequencies.

If the waveform at upper right were the MR signal for a rectangular object, we would necessarily be able to acquire only part of it as doing otherwise wold require infinite time. When we then take the Fourier transform of the truncated signal, we find a distorted representation of the original waveform. Notably, the signal now contains ringing at the edges.

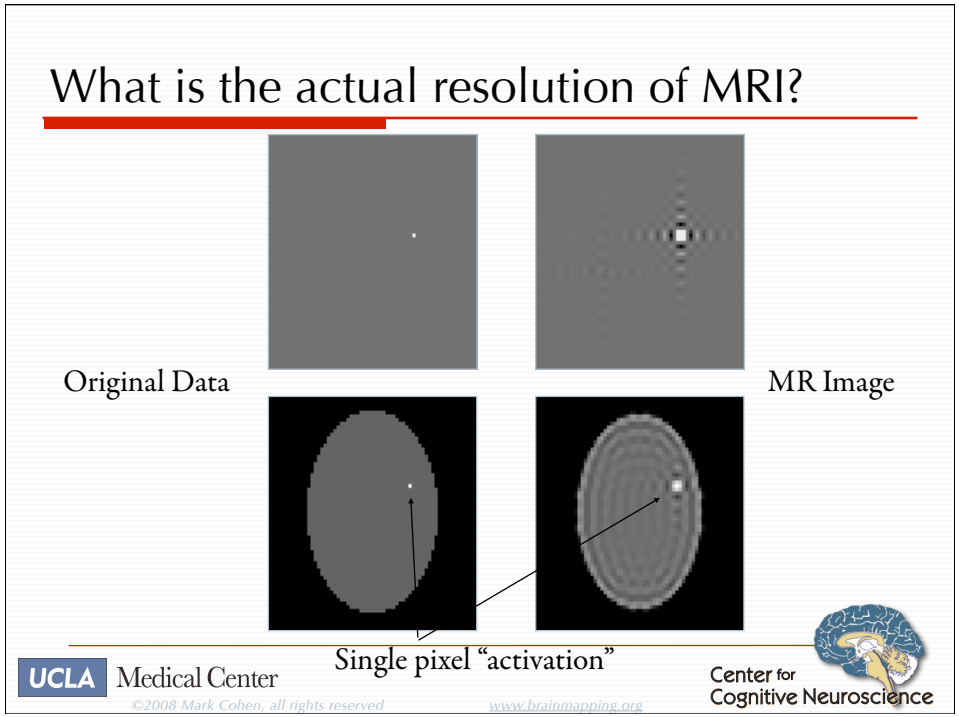

Because actual MR imaging is always based on a truncated signal collection, the resulting images are always distorted. In this simulation, we consider the images that would result from a small area of bright signal. The image representation will spread far away from the signal source as a sequence of dark and light bands. Further, a distinct, though low intensity, focus of signal will appear in the opposite X and Y locations on the image. If the bright location were a focus of brain activation that caused a BOLD signal increase, the apparent location of the activity would be represented throughout the image. An accurate statistical map based on the images would include signal decreases in the pixels next to the center of the activation.

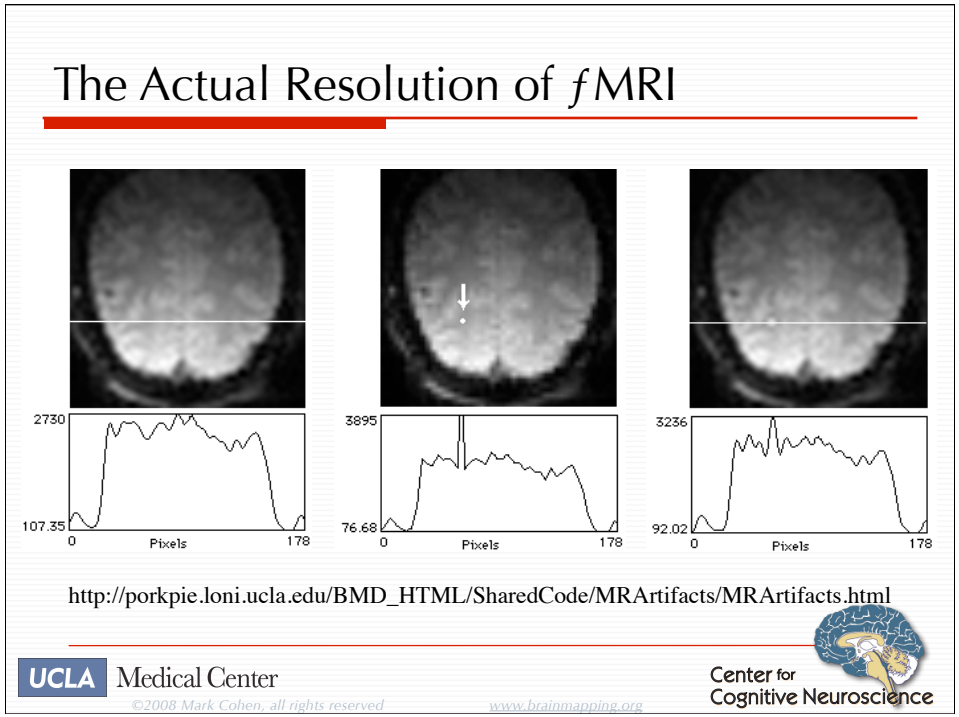

The apparent representation on an MR image of a signal point of true signal increase

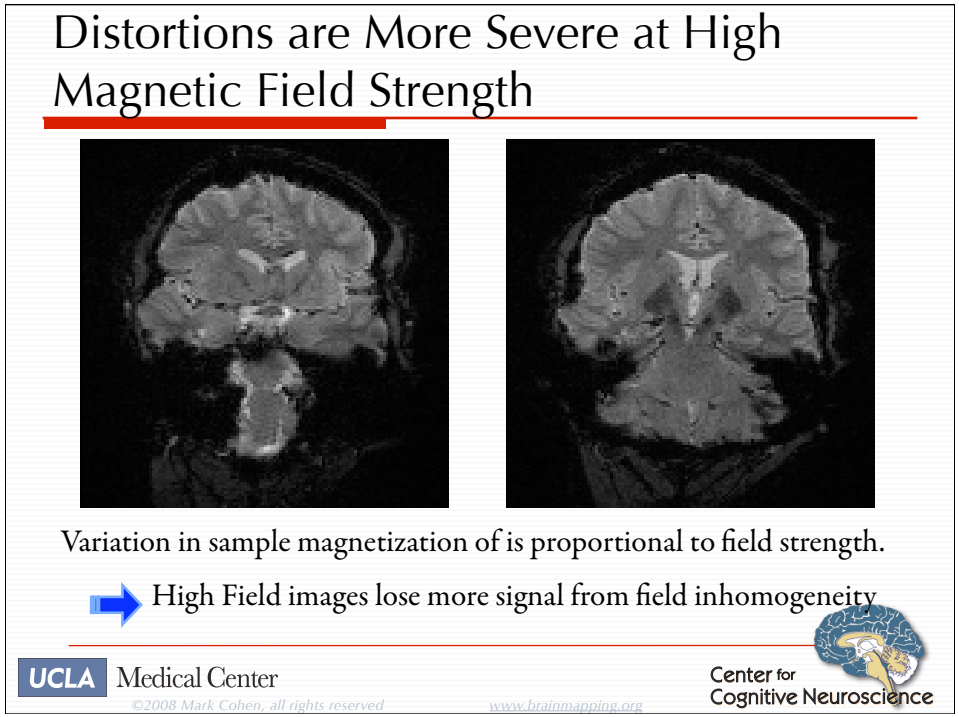

Location is detected by the frequency of spin precession, which in turn is determined by the local strength of the magnetic field. If the field is distorted, the images will be as well. Unfortunately, the subject him or herself distorts the primary field because body tissues magnetize unequally. The images become distorted by this. Various pulse sequence factors either tend to mitigate or amplify this effect, but the combination of parameters used to create fMRI images tends to amplify these shape distortions.

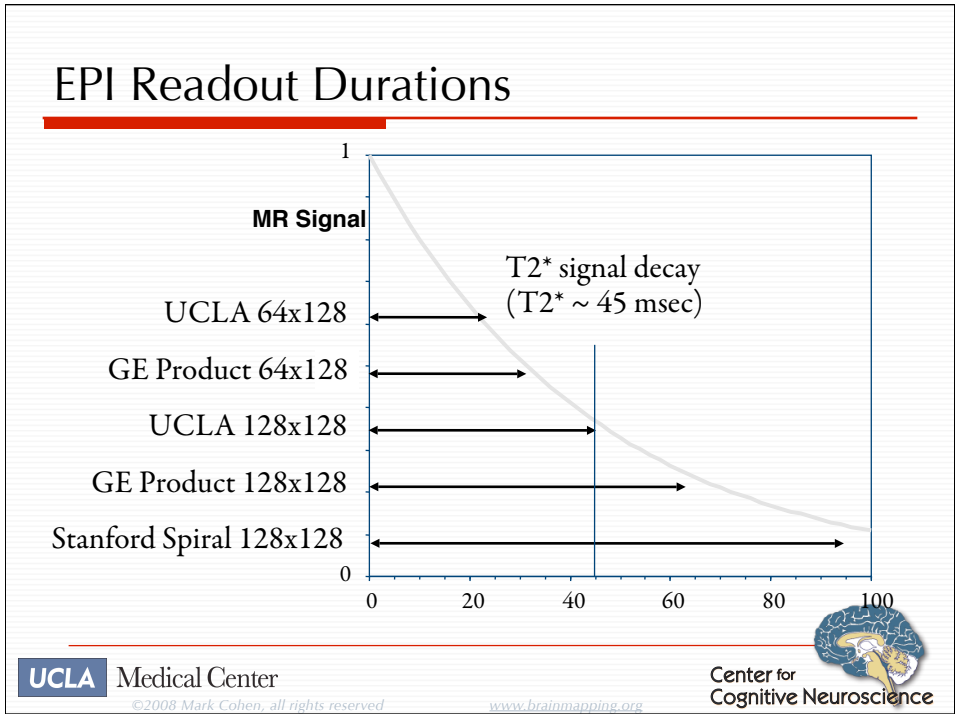

Henceforth all of the discussion about gradient encoding and image distortion pointedly neglected T1 and T2 effects on the signal. The actual effects on the spatial localization are not insignificant, however. During the time that we are encoding the location, the signal is changing. For example, during the readout of an echo-planar or spiral scan, T2* effects cause the signal to decrease. In long readout sequences, such as spiral scans the signal may decrease by 80% during the actual readout.

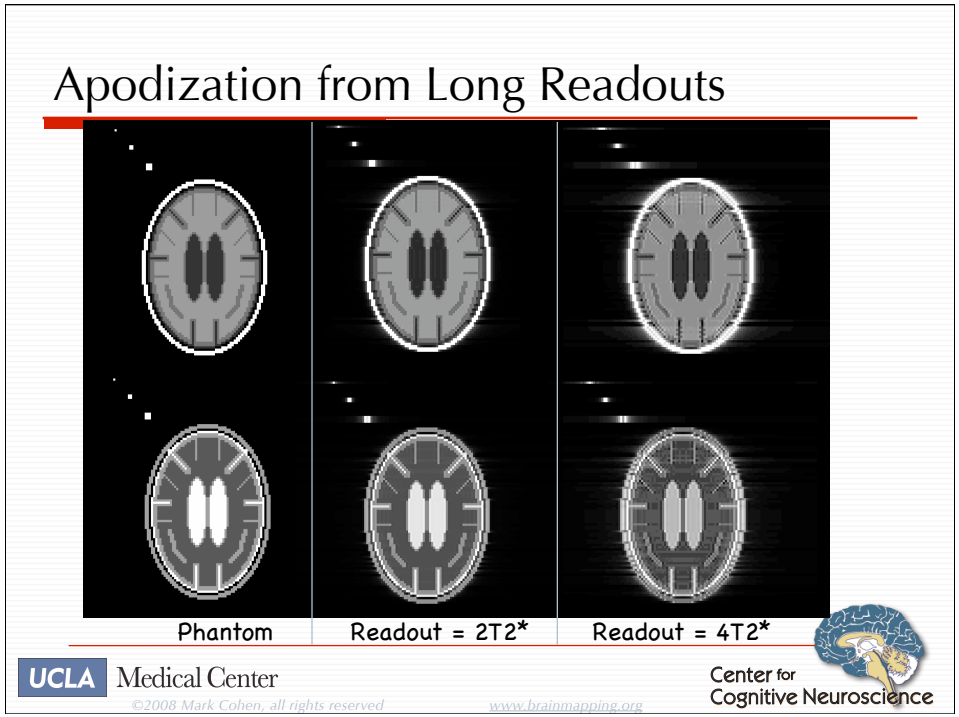

If the readout period is long compared to T2*, the images tend to become blurred or apodized. This simulation shows the blurring that would occur with various readout durations. T2* in the human brain is typically around 40 msec (but depends strongly on the scanner itself) and readout periods of 50 to 80 msec are not uncommon. Thus the images are detectably blurred. This simulation considers only rectilinear k-space traversals, such as EPI. In spiral scans the blurring that results is more circular and much more complex.

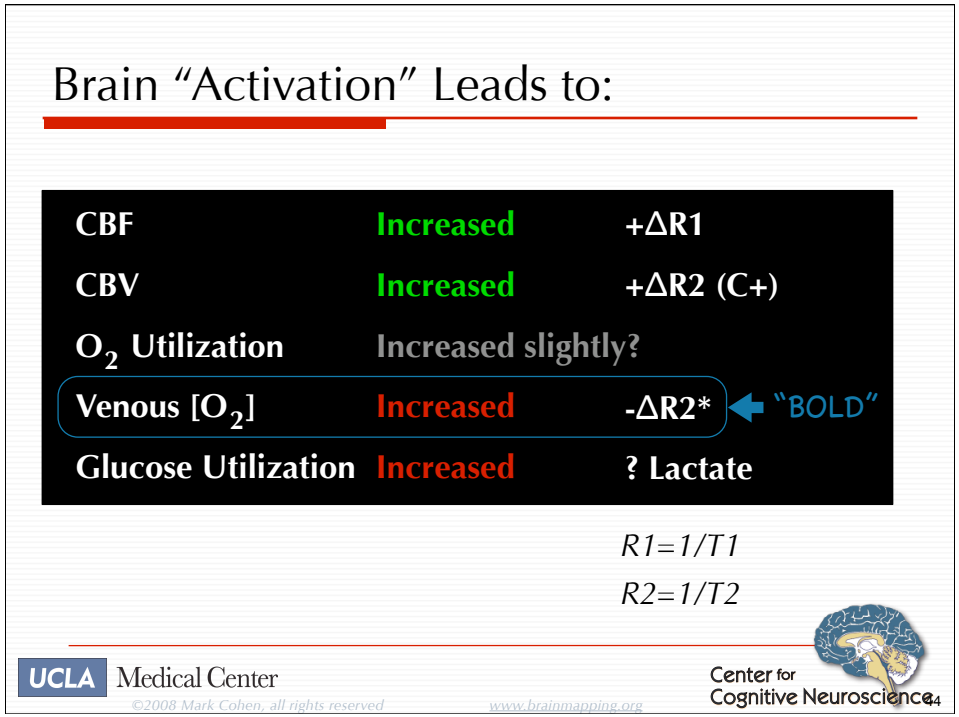

There are many biophysical effects associated with increased neuronal electrical activity. At least five coupling events are known to be observable during fMRI, though most studies at present rely on detecting the change in venous oxygen concentration which follows increases in synaptic activity.

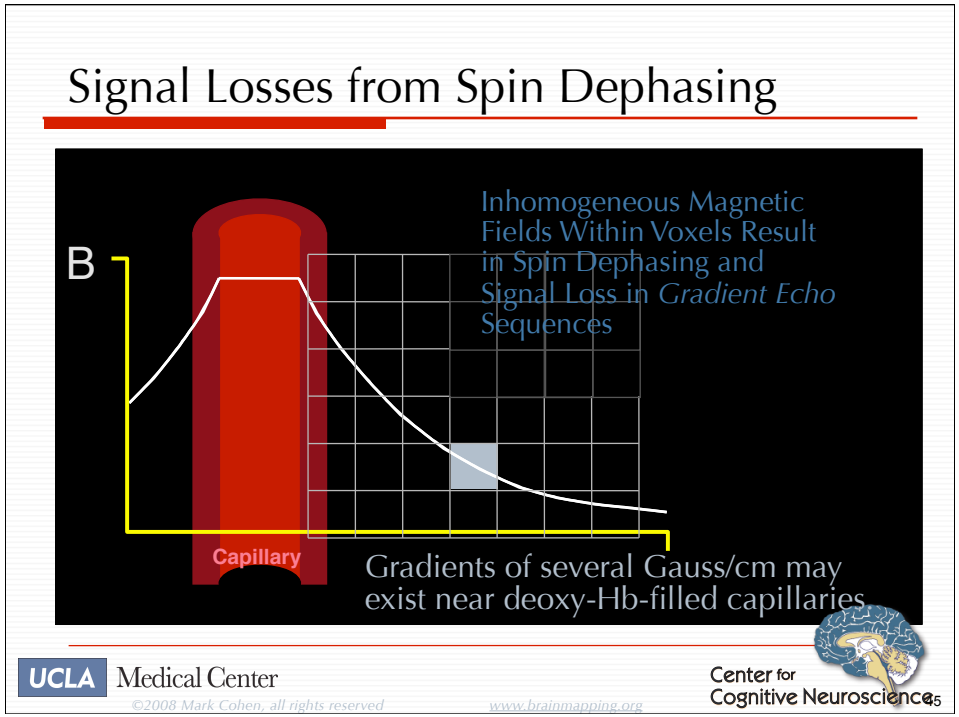

MRI signal depends on the coherence of individual proton spins. In the presence of magnetic field distortions, the MR signal is reduced.

Deoxygenated blood is paramagnetic and results in such field distortions in the immediate vicinity of the capillaries.

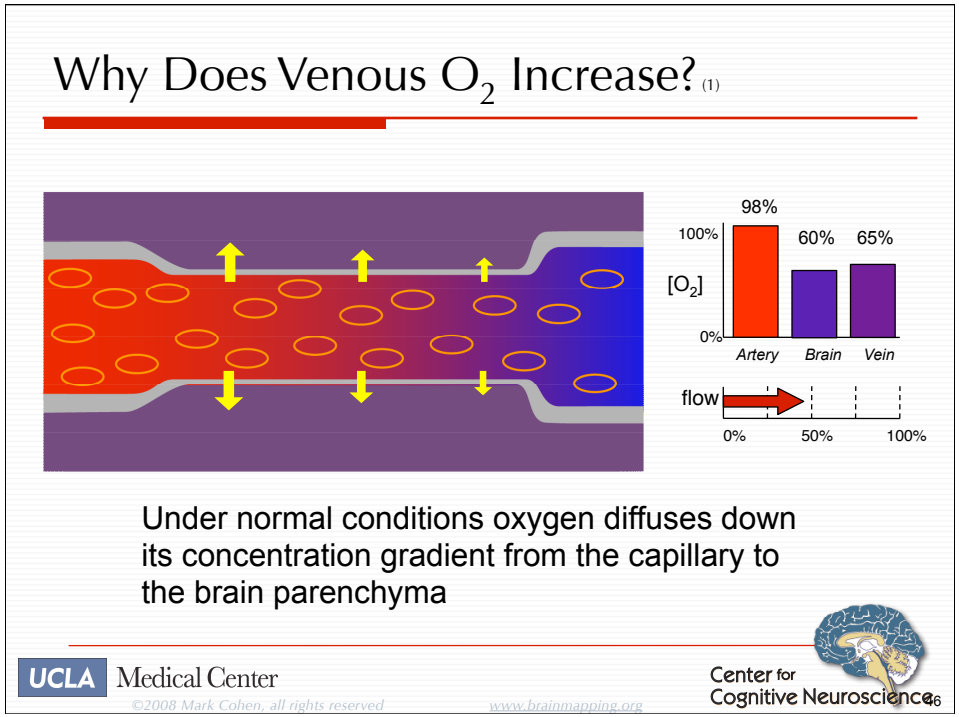

This series of slides outlines the dominant hypothesis as to why the BOLD signal is increased when brain metabolic demand is increased. Some believe, however, that this demand is driven not by oxygen but by glucose. The effects would be the same, however.

Why Does Venous O_2 Increase? (4)

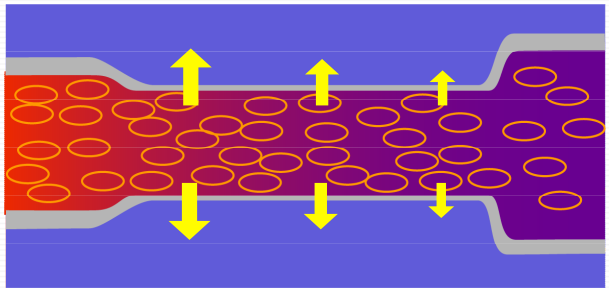

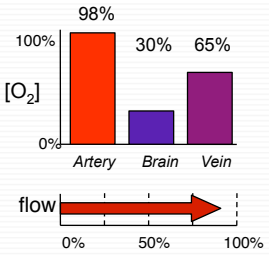

The vascular system responds by increasing blood flow so that more oxygenated blood is available throughout the capillary

©2008 Mark Cohen, all rights reserved

Why Does Venous O_2 Increase? (5)

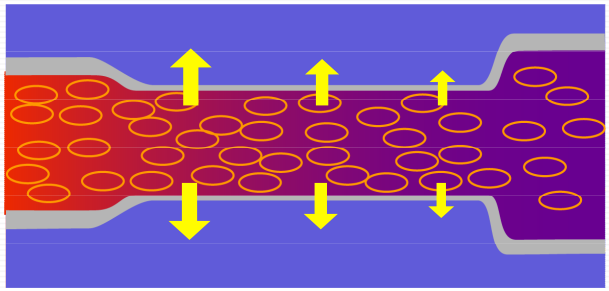

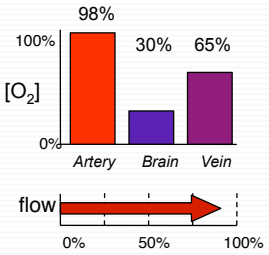

Because the blood flow is increased more oxygenated blood passes into the venous end of the capillary

©2008 Mark Cohen, all rights reserved

Why Does Venous O_2 Increase? (6)

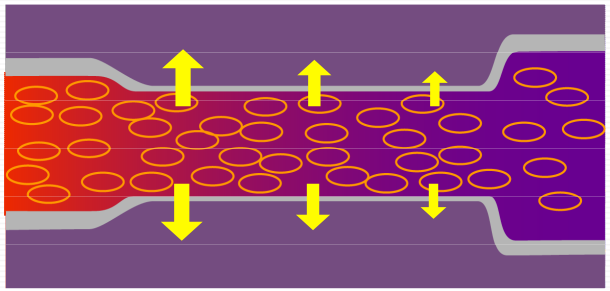

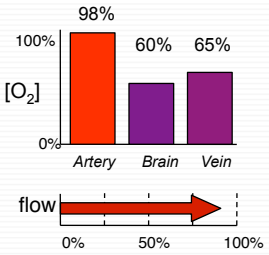

Because the blood flow is increased more oxygenated blood passes into the venous end of the capillary

©2008 Mark Cohen, all rights reserved

BOLD Contrast & Field Strength

- BOLD Contrast arises from susceptibiity differences
- The absolute field distortion (from BOLD) is proportional to the magnetic field strength
- The absolute change in MRI signal is proportional to *both* the field distortion and the signal strength.

$$\text{BOLD } \textit{should} \text{ go as } kB_0^2$$

©2008 Mark Cohen, all rights reserved

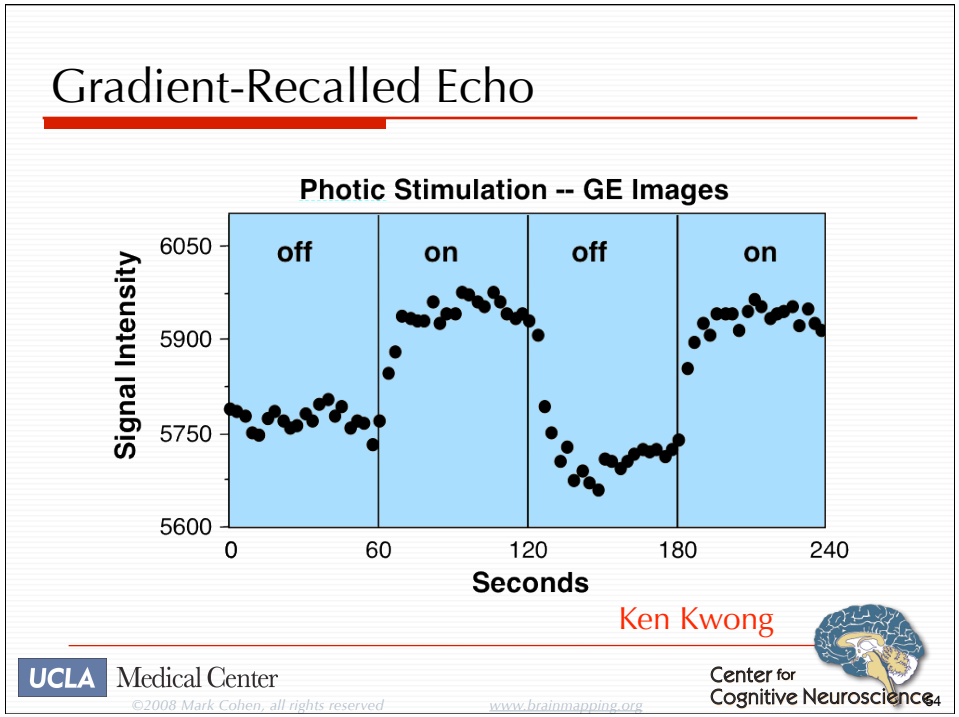

fMRI Signal increases relatively slowly (several seconds), and returns to baseline slowly.
No exogenous contrast needed

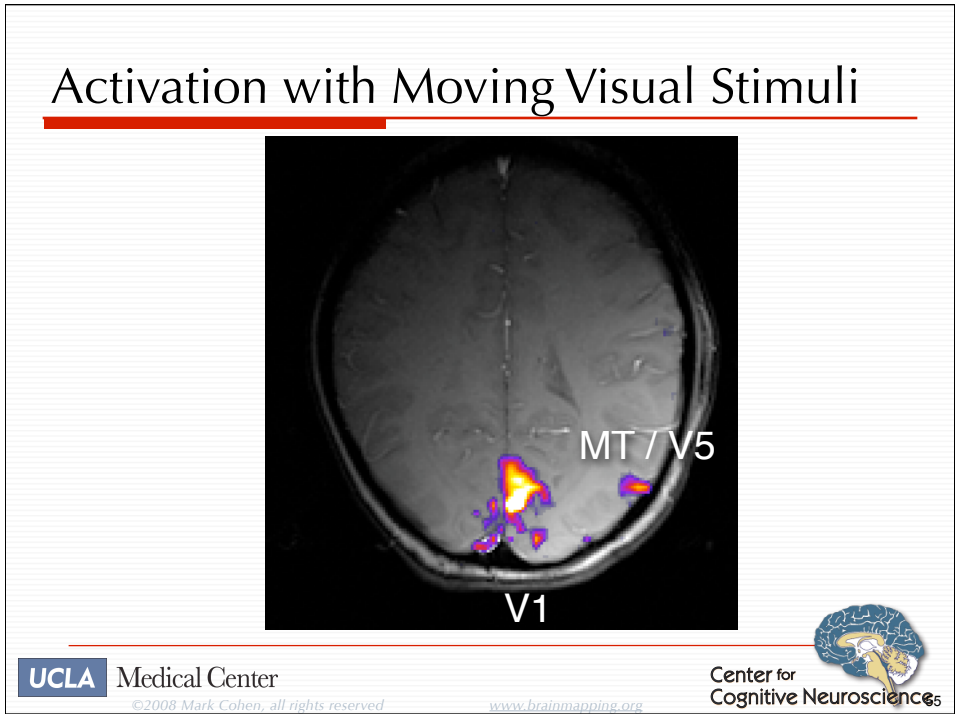

One of the more interesting and important features of fMRI is its ability to collect multiple experiments on the same subject to tease apart the responsiveness of various brain regions to complex stimuli. In this study, Roger Tootell looked at the fMRI signal change when a subject observed a visual stimulus containing moving vertical bars. Multiple brain areas showed signal increases, but his hypothesis was that the detection of the bars and of their motion were represented separately in the brain.

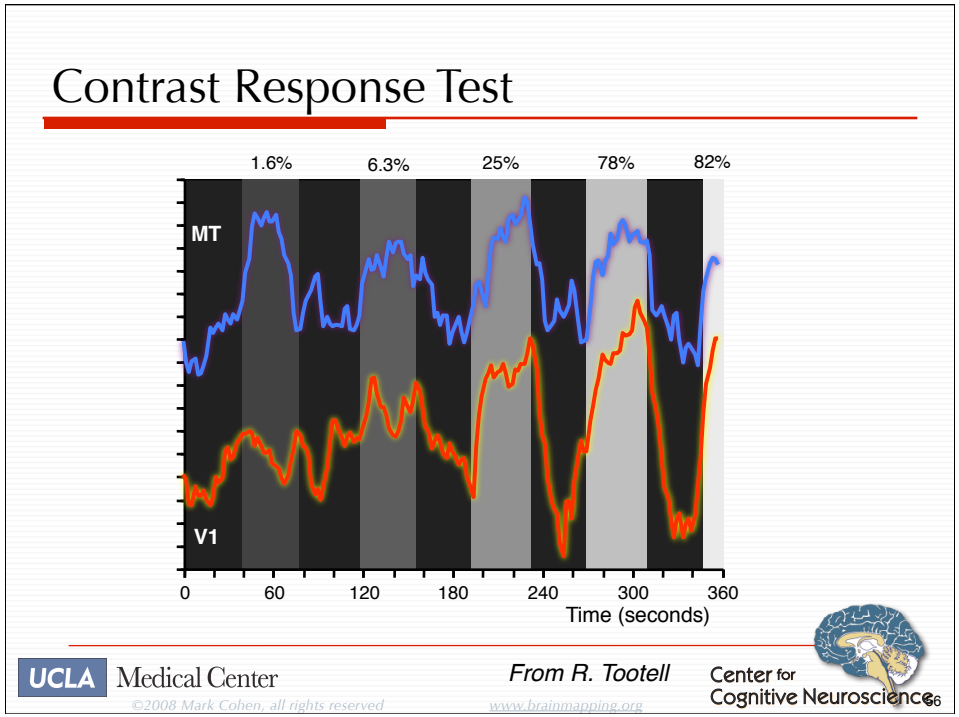

To test this hypothesis, Tootell looked at the fMRI signal as a function of contrast between the moving bars and their background. As the contrast was increased, there was a parametric increase in the signal response within area V1 (the central blob in the prior slide) whereas the more lateral extrastriate blob responded similarly regardless of contrast.

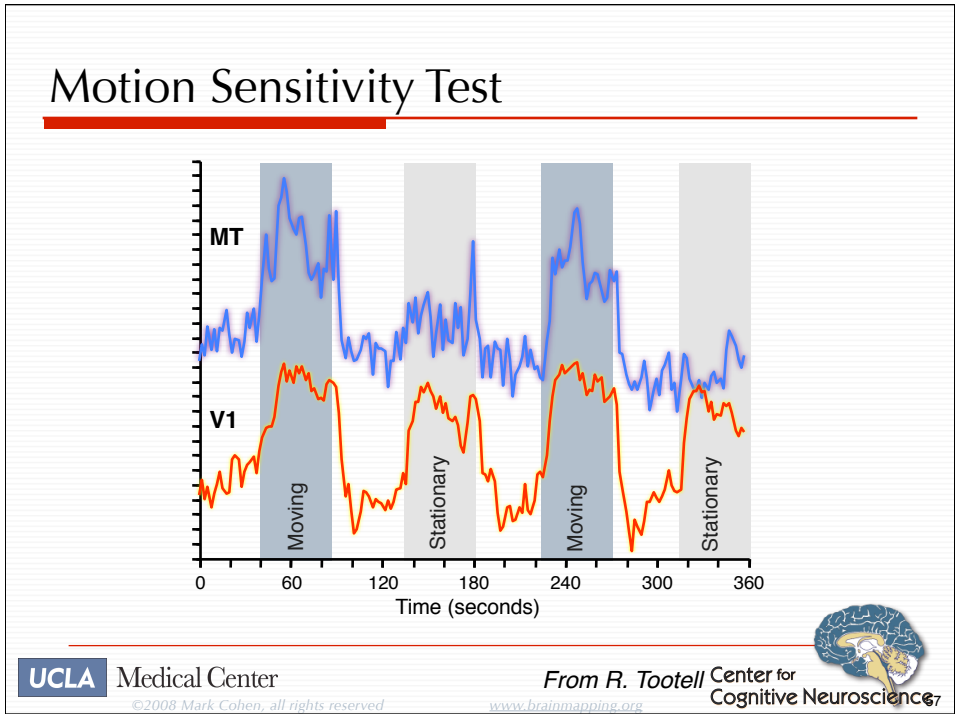

When Tootell instead observed the responses to lines that were either moving or stationary, the signal in V1 was of essentially equal amplitude whether or not the bars were moving, whereas the signal changes were much greater in the extrastriate area, MT, when the bars were moving. This is a double dissociation experiment that was able to functionally identify area MT.

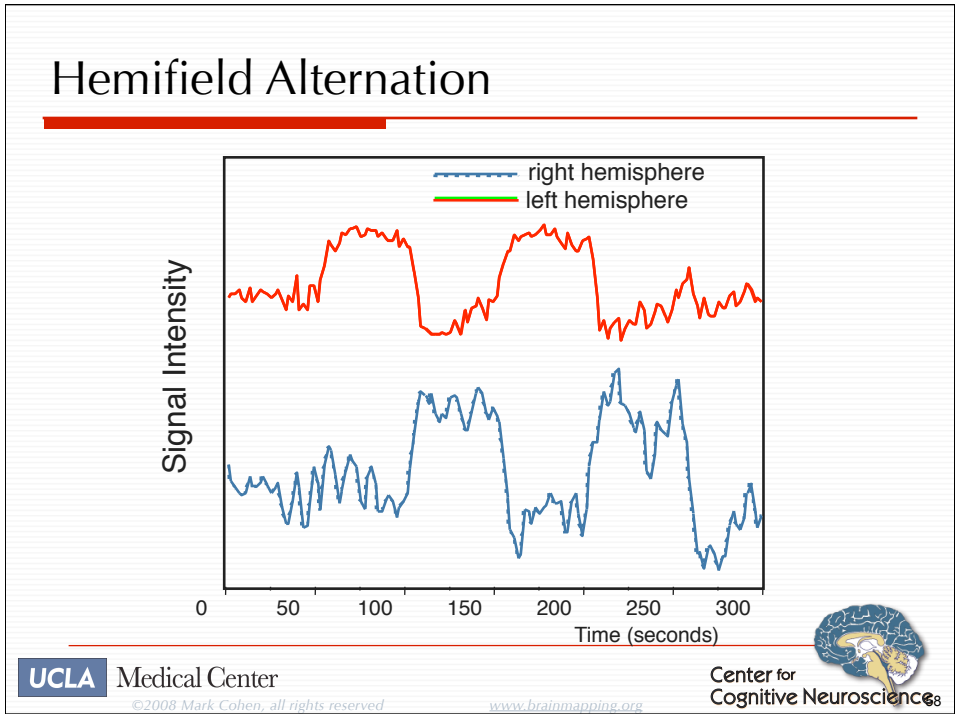

If stimuli are modulated relatively slowly, the BOLD fMRI response follows the stimulus timing rather faithfully.

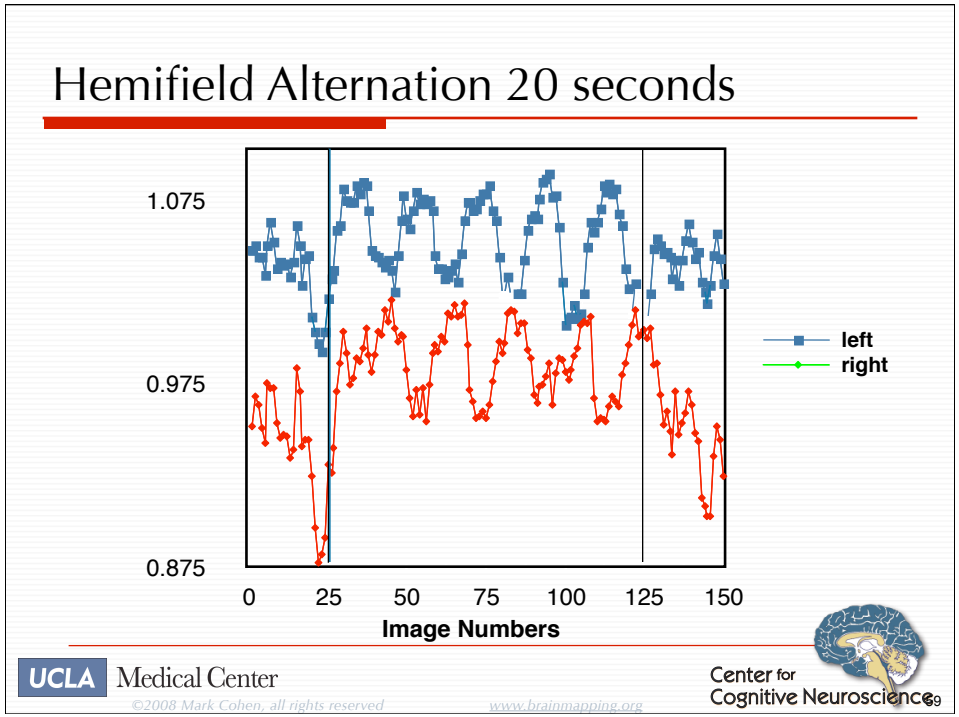

When stimuli are varied quickly, the BOLD signal changes are much reduced and do not look a great deal like the stimulus timing.

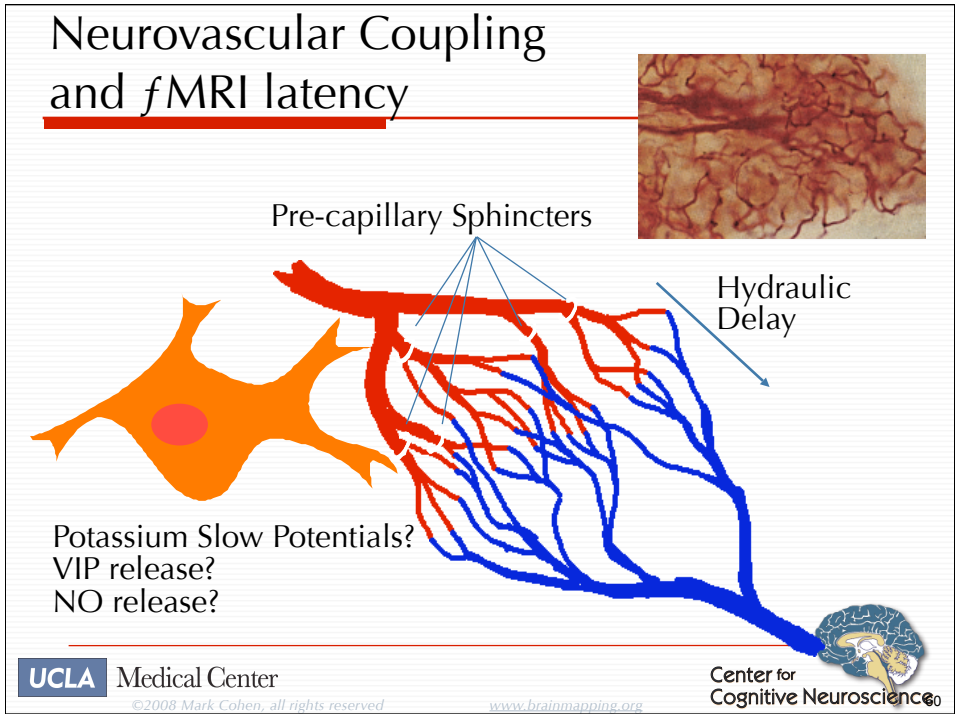

The fMRI response takes some time to occur. Much of this delay is thought to be related to the signaling from neurons to the vascular system. At this time it is not entirely clear how this signaling takes place, though there are many candidate signaling molecules and mechanisms that are probably all involved to varying degrees.

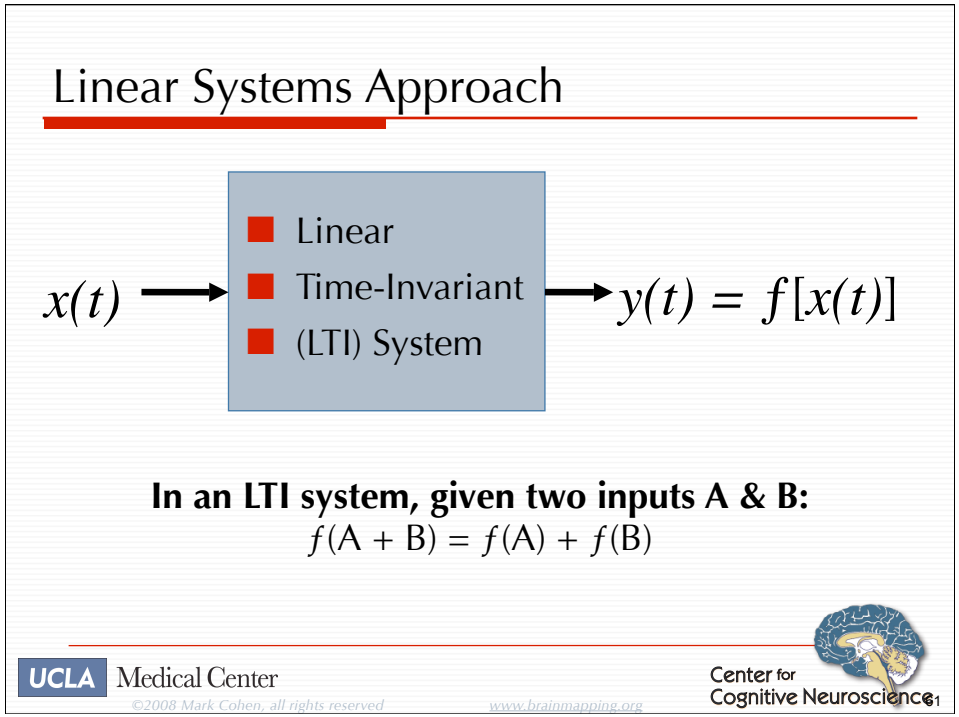

Almost all fMRI analysis is now performed under the rubric of linear systems analysis. This is because the properties of linear systems are very well-worked out and the toolsets available to analyze them are mature and relatively simple. A linear, time-invariant system is defined as one which obeys the property that the response of the system to two inputs presented together is the same as the sum of the responses to those two inputs presently individually.

One of the most important properties of a linear system is that by measuring its response to a standard input - usually an instantaneous transient input, we can calculate and predict its response to any input through the use convolution. Specifically, the predicted output, Y, to any stimulus, X is the convolution of X with the systems response, h, to an impulse: Y = X * h, where * is used to indicate convolution.

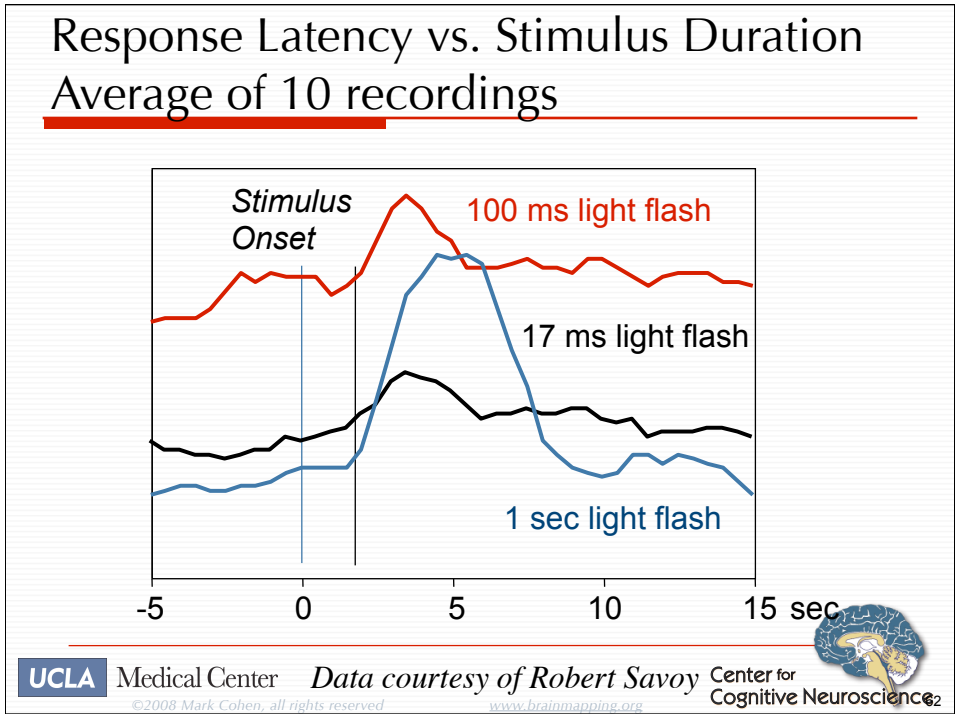

To capture the power of linear systems analysis for fMRI data it is useful to estimate the form of the brain's response to an impulse stimulus. These data of Robert Savoy show the response in V1 to brief flashes of light, which is a reasonable approximation to an impulse stimulus.

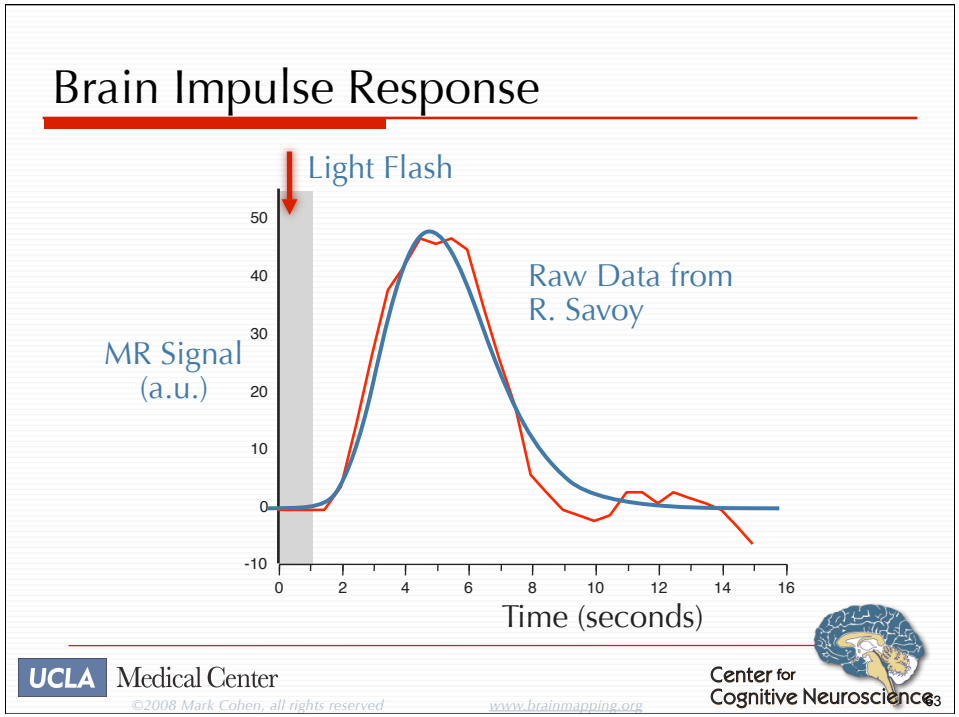

Using the observed V1 impulse response one popularly used model of the brain impulse response is the Gamma variate waveshape, shown here in blue. The Gamma function is a relatively function whose shape is very similar to the observed brain impulse data.

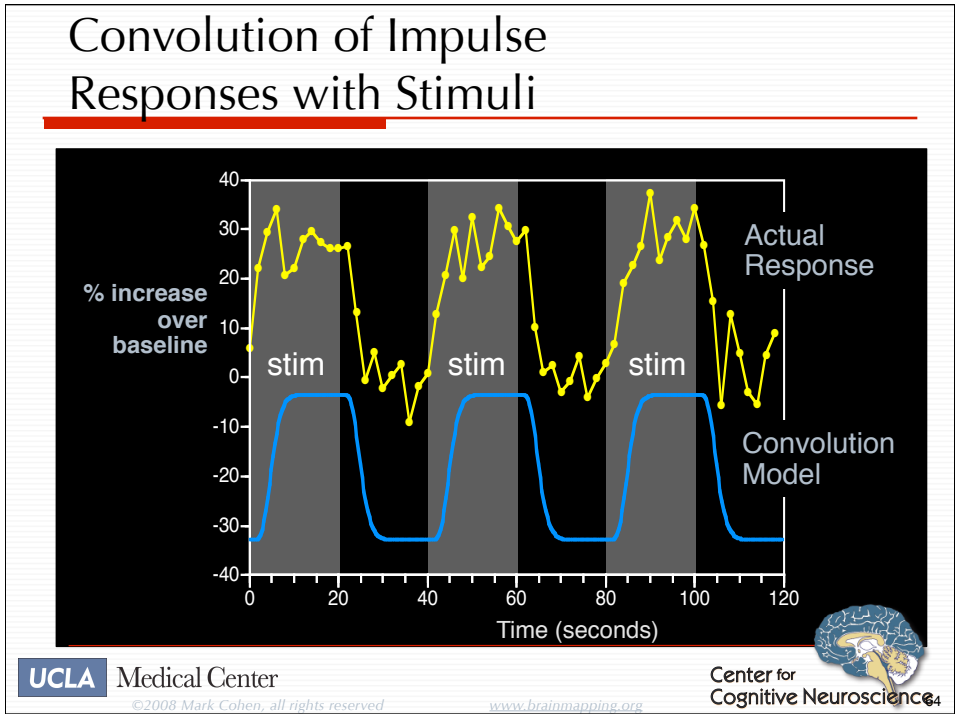

The yellow points indicate the signal intensity observed in V1 in response to an intense flashing light stimulus that was turned on and off every twenty seconds. Convolving the time course of the stimulus with the Gamma function discussed above results in the estimate for brain activity shown in blue. Visual comparison of the two curves suggests that the convolution model is a reasonably good approximation to the observed response. The canonical means of detecting brain activation by fMRI is to model the brain response to the stimulus and to search for brain regions that behave similarly to the model. Thus, the convolution approach based on linear systems theory is a powerful tool for this work.

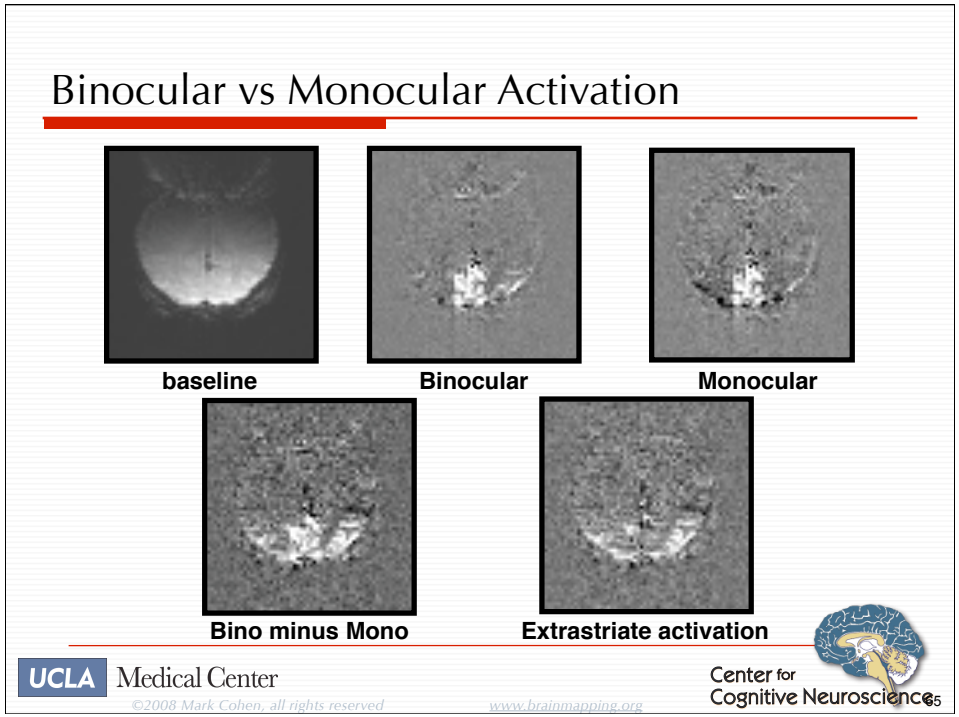

Given the importance of the linear systems analysis, it is instructive to test whether the brain fMRI signal response is linear at all. These data from Ken Kwong show the brain responses to visual stimulation to both eyes simultaneously and to one eye alone. Because the primary visual cortex is known to be independent for each eye, we would expect the difference between the binocular and monocular stimulation to be the same as the response to the other eye alone. The data however show that this is simply not the case. Thus our assumption of linearity is strongly violated.

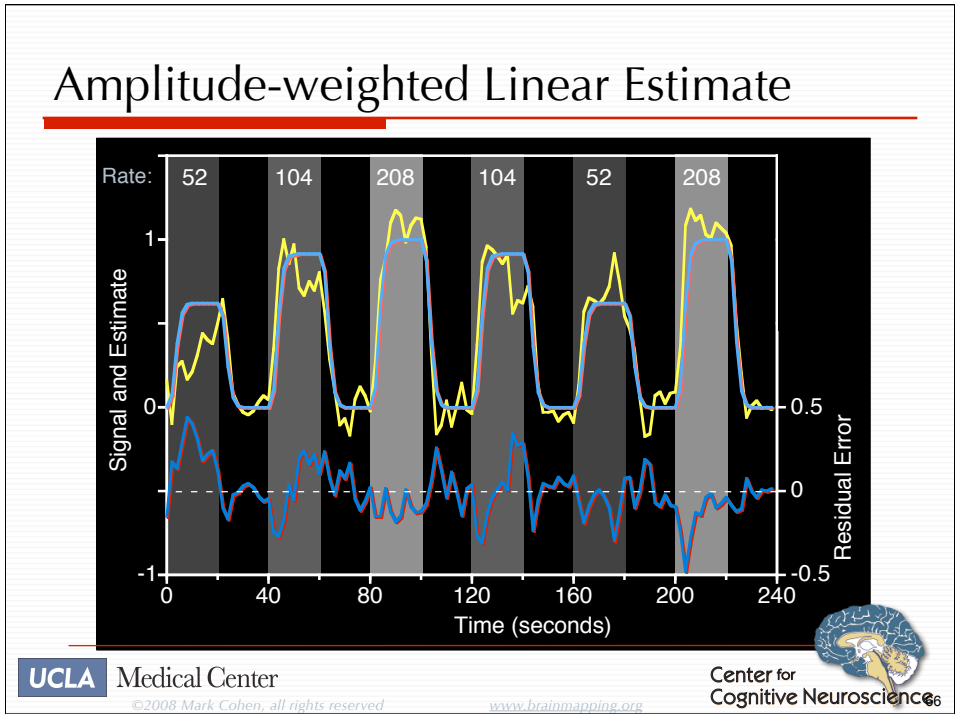

Various degrees of non-linearity can be incorporated into the linear systems analysis in a principled manner. Here, the estimated BOLD response was modeled as increasing in amplitude with the log of the rate at which subjects performed a finger motion task. Notice that the response estimates closely fit the data.

People using fMRI for their data collections should be cognizant of the assumptions made in the analysis and thoughtful of how they relate to the underlying physiology of the brain. In many cases, the assumptions of linearity will produce grossly distorted activation maps.

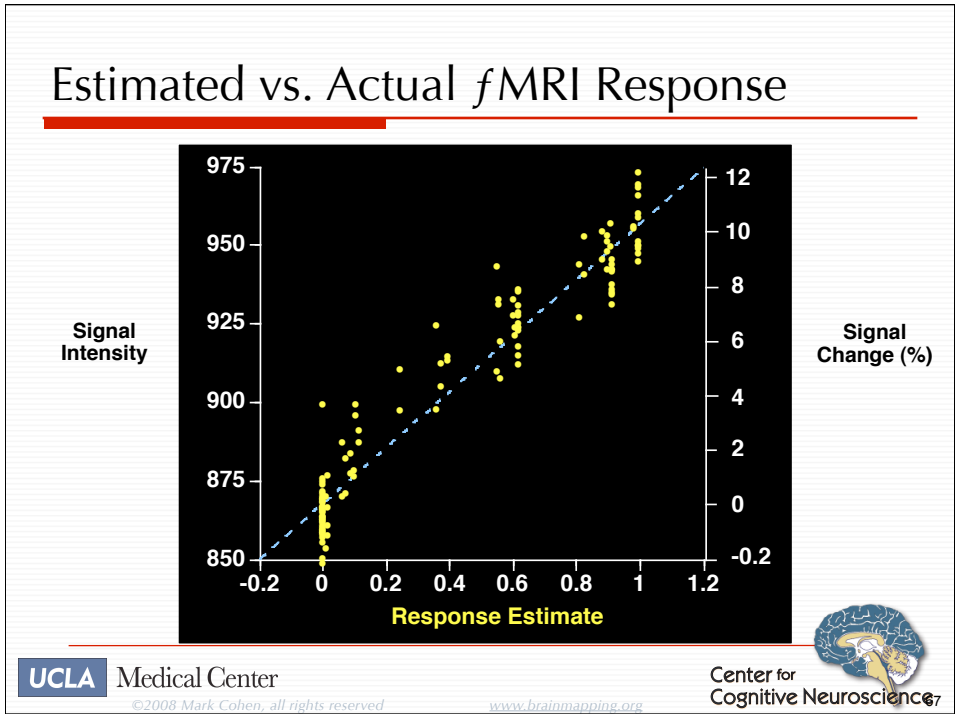

These are the data from the prior slide plotted instead as the observed signal intensity as against the modeled intensity and show that the log transform model closely fits the actual data.

So long, and thanks for all the fish*.

*…nod to Douglas Adams

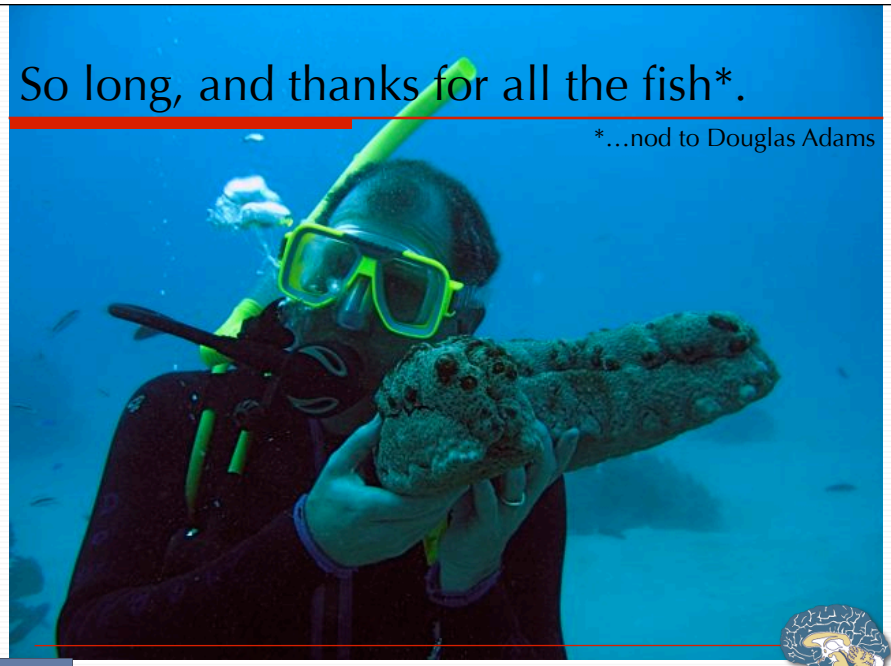

UCLA Medical Center

Center for Cognitive Neuroscience

©2008 Mark Cohen, all rights reserved