

IX SPOILED SPIN AND GRADIENT ECHO IMAGING

Lecture notes by Assaf Tal

We've already met EPI, in which a slice is excited and a signal is consequently acquired in k-space, by varying the gradients. EPI is fast, but suffers from several drawbacks and artifacts, including field inhomogeneity, sub-optimal resolution and high gradients.


In this chapter we set out to devise a robust, fast, all-around imaging sequence with appropriate T_1 and T_2 contrast. What we will come up with are the basic spin & gradient echo imaging sequences. Due to the course's short duration we will only be able to discuss *spoiled* sequences (I'll explain what that means shortly).

1. RAPID PULSING

1.1 DYNAMIC EQUILIBRIUM

Imagine a static bucket with water. The water is said to be in *static equilibrium*, because nothing's happening to it. Next, imaging (i.) poking a hole in the bottom of the bucket, so water start running out, and (ii.) opening a tap just above the bucket, letting water flow in at a constant rate. What will happen? The water may rise or fall, but will eventually reach a new state of equilibrium (remember, the more water there is, the faster it drips out due to pressure; and the less water there is, the slower, so eventually the water flowing in will equilibrate with the water flowing out, even if at first the rates are different). This new equilibrium is termed *dynamic equilibrium*. Something is continuously happening to the system (in and out flow of water), so it's not static anymore, and yet its state doesn't change.

A similar thing happens in MRI when we use rapid pulsing. Consider a train of pulses of flip angle α (that is, $\gamma B_{\text{RF}} t = \alpha$) around, say, the $-y$ axis. Let's call the time between pulses TR (for "Time per Repetition"):



The spins get acted upon by two "forces": the pulses, which repetitively try to take them out of equilibrium, and relaxation, which tries to get them back to equilibrium. There's also precession going on. It can be shown (I won't do it here, but it's not that difficult really) that the spins eventually settle into dynamic equilibrium, regardless of their initial state; that is, after enough pulses have been given, the state of the spins after each pulse is identical. In other words, the magnetization vector at points A, B, C, ... below is the same:



This state will depend on the variables of the system: TR, T_1 , T_2 and α , and also the offset of the spins, $\Delta\omega$. We're going to calculate it, but before we do, we need to make a simplifying assumption in the next section.

1.2 SPOILING


"Spoiling" the magnetization means getting rid of its transverse component; that is, making $M_{xy}=0$. One easy way of doing this is, well, waiting. Waiting enough time ensures T_2 does its job and "eats up" the transverse magnetization, leaving only a longitudinal component.

There are also ways of spoiling the magnetization on purpose. This will turn out to be beneficial in the next section. One such

method is by using very strong field gradients. Let's try and understand why this works in 1D. Imagine you've just excited your spins onto the x axis in your 1D sample, so they're all pointing along the same direction regardless of their position.

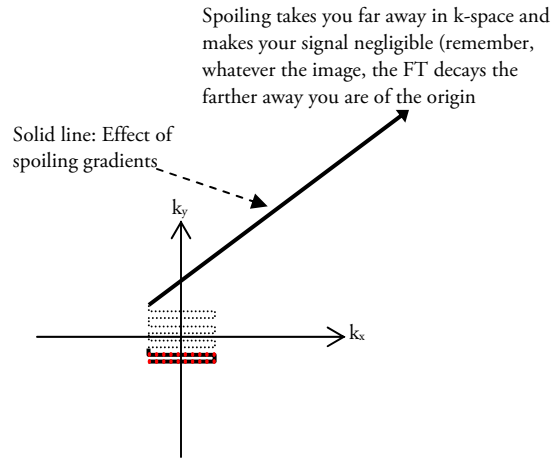


The application of a gradient instills a different precession frequency into each position, $\omega(z) = \Delta\omega_{\text{chemical-shift}} + \gamma Gz$, so after some time the spins will look like this:




Once they do end up pointing in different directions, they will stop contributing to the acquired signal since their vectorial sum will be approximately zero (signal $\sim \int M_{xy} dz \approx 0$).

For gradient-spoiling to work, you need to use very strong gradients for enough time. How strong? Strong enough for subsequent imaging gradients not undo their action and refocus the spins. You can think of this in terms of k-space. What you're doing is "propelling" the spins to a very distant point in k-space, away from the trajectory:



1.3 THE EFFECT OF RAPID PULSING

I'm going to assume that, prior to each pulse in the train, the magnetization is spoiled, either naturally (by waiting) or by applying spoiling gradients. This means that $M_{xy}=0$ just before each pulse. Let's look at three points, A, B and C:




Denote by $\mathbf{M}^{(A)}, \mathbf{M}^{(B)}, \mathbf{M}^{(C)}$ the magnetization vectors at A, B and C, respectively. We're interested in computing $\mathbf{M}^{(B)}$ and, subsequently, the magnetization's evolution between the pulses. Note that:

1. By assumption of dynamic equilibrium, $\mathbf{M}^{(A)} = \mathbf{M}^{(C)}$.
2. Since we're assuming the magnetization is spoiled, $M_{xy}^{(A)} = M_{xy}^{(C)} = 0$.

So:

a



Since $\mathbf{M}^{(B)}$ is a tipped version of $\mathbf{M}^{(A)}$ by an angle α , we have:

$$\begin{aligned} M_z^{(B)} &= M_z^{(A)} \cos(\alpha) \\ M_x^{(B)} &= M_z^{(A)} \sin(\alpha) \\ M_y^{(B)} &= 0 \end{aligned}$$

So $M_{xy}^{(B)} = M_z^{(A)} \sin(\alpha)$. However, we're not interested in the transverse magnetization (yet). We know that the longitudinal component of \mathbf{M} relaxes back to equilibrium with a time-constant T_1 . You've shown in exercise 3 that, following the pulse and after a time TR , M_z will equal

$$M_z^{(C)} = M_z^{(B)} e^{-TR/T_1} + (1 - e^{-TR/T_1}) M_0$$

where M_0 is the thermal equilibrium value of the magnetization (it's in general **not** equal to $M_z^{(A)}$!). Since $M_z^{(C)} = M_z^{(A)}$ and $M_z^{(B)} = M_z^{(A)} \cos(\alpha)$, we can plug these into the above equation,

$$M_z^{(A)} = M_z^{(A)} \cos(\alpha) e^{-TR/T_1} + (1 - e^{-TR/T_1}) M_0$$

and solve for $M_z^{(A)}$:

$$M_z^{(A)} = \frac{1 - e^{-TR/T_1}}{1 - \cos(\alpha) e^{-TR/T_1}} M_0.$$

From this we can compute $M_{xy}^{(B)} = M_z^{(A)} \sin(\alpha)$, and, in fact, deduce

$M_{xy}(t)$ for any time between the pulses ($t=0$ corresponds to the time right after a pulse):

$$\begin{aligned} M_{xy}(t) &= M_{xy}^{(B)} e^{-ik(t) \cdot r} e^{-t/T_2} \\ &= \frac{[1 - e^{-TR/T_1}] \sin(\alpha)}{1 - \cos(\alpha) e^{-TR/T_1}} M_0 e^{-ik(t) \cdot r} e^{-t/T_2} \end{aligned}$$

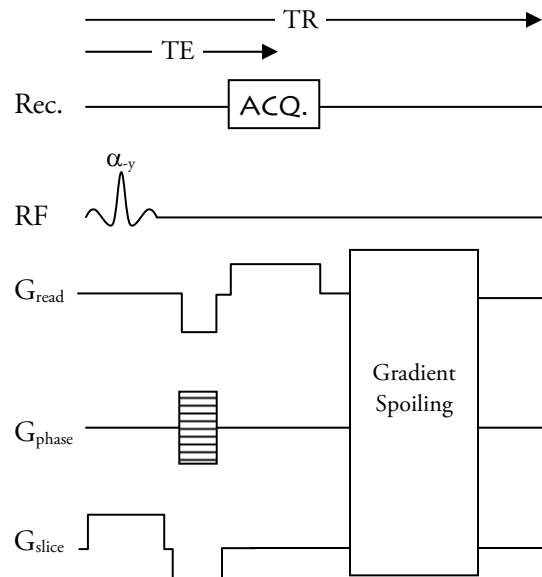
So, the "usual" M_0 is now replaced by

$$M_0 \rightarrow \frac{[1 - e^{-TR/T_1}] \sin(\alpha)}{1 - \cos(\alpha) e^{-TR/T_1}} M_0$$

i.e., by M_0 times some factor which depends on T_1 , T_2 , α and TR . This is the result we were after, and which will be of use to us in a short while.

2. SPOILED GRE IMAGING


The basic spoiled GRE sequence looks like this:



This sequence has several new and interesting elements, which I will now explain. It is repeated **periodically** every TR seconds.

2.1 READ, PHASE, SLICE

The first thing you notice is that I'm no longer referring to gradients as $G_x/G_y/G_z$, but rather as G_{slice} , G_{phase} and G_{read} . Why? Well, had I used



it would imply that I'm selecting a spatial slice perpendicular to the z-axis (a so called "axial slice"). Here slice=z, phase=y and read=x. However, I don't want to commit to a particular direction. What if I wanted to excite a coronal slice? Or some arbitrary slice in the body – even one which is not perpendicular to the xy, yz or xz planes? I want to talk about imaging in a completely general way. The ensuing discussion will only assume slice, read & phase are perpendicular directions. It's going to be up to you, the experimentalist, to decide their assignment. The important thing is that the conclusions we'll derive here will be completely general, regardless of your assignment.

2.2 SEQUENCE OVERVIEW


The sequence begins with a slice-selective pulse along the "slice-direction" (once again, you choose that direction), in the presence of a positive slice-selection gradient. It is

followed by a negative $-G_{\text{slice}}$ to refocus the phase of the spins. The slice-selective pulse (of some duration T) is calibrated such that, on resonance, $\gamma B_{\text{RF}} T = \alpha$. Basically, since this basic unit is going to get repeated many times, we've got a rapid-pulsing scheme here, so we can write for the magnetization right after the pulse:


$$M_z(\mathbf{r}) = \frac{1 - e^{-TR/T_1}}{1 - \cos(\alpha)} \cos(\alpha) M_0(\mathbf{r})$$

$$M_{xy}(\mathbf{r}) = \frac{1 - e^{-TR/T_1(\mathbf{r})}}{1 - \cos(\alpha)} \sin(\alpha) M_0(\mathbf{r})$$


Note how the quantities M_0 , T_1 and T_2 all depend on position, as they usually do. Next come the phase & read gradients. There's a new symbol here for G_{phase} :



This has the following meaning: start from the negative-most G_{phase} shown in the diagram,




in each repetition increase its value, until the maximal value:



where N_p is the total number of repetitions, set by you, the experimentalist. The above drawing is schematic. We'll talk in a moment about how to set it.

In terms of k-space, what the phase & read gradients do is bring you in k-space to a particular point, after which you acquire along a straight line while keeping $G_{\text{phase}}=0$ and $G_{\text{read}}=\text{const}$. Pictorially, this is how it looks like:

Total time \approx 3 minutes



The dashed lines represent the movement in k-space due to the rewinding phase & read gradients prior to acquisition, without acquiring. The solid lines represent lines acquired. There are a total of N_p such lines, since there are N_p scans (each scan having a slightly different G_{phase}).

Now you can also understand how G_{phase} should be varied, and how N_p should be chosen: the imaged object has dimensions $\text{FOV}_{\text{phase}}$ along the phased axis. If you want a voxel size of Δr , then you will choose

$$k_{\text{max,phase}} = \frac{1}{\Delta r}, \quad \Delta k_{\text{phase}} = \frac{1}{\text{FOV}_{\text{phase}}}$$

(with $k_{\text{max,phase}} = N_p \Delta k_{\text{phase}}$.)

Once the basic block has been repeated N_p times **for a particular slice**, acquiring a rectangular grid in the $k_{\text{phase}}-k_{\text{read}}$ plane, a different slice is chosen and imaged (once again, using N_p scans). If you have a total of N_s slices, then the total time for the experiment will be:

$$\text{Time} = \text{TR} \times N_p \times N_s$$


GRE sequences usually have short TRs (for reasons that will become apparent in section 3). A TR of 100 ms is not uncommon. Assuming an image with 16 slices, and 128×128 resolution in each slice (so $N_p = 128$, $N_s = 16$), we have

Note that the resolution along the read direction does **not** affect the imaging time. You get it “for free”, because you can sample as many points as you’d like (limited by your hardware’s minimal dwell time) during TR. Whether you sample 128 or 256 won’t change your TR (it never happens that you sample so many points that they do not fit into your TR. With a dwell time of 100 nanoseconds, you’d have to sample $> 1,000,000$ points for that to happen!).

The important thing to note here is that you don’t have to wait a time $\sim 5T_1$ between experiments for the magnetization to relax back to equilibrium, because of the rapid pulsing. Don’t forget you need to spoil the magnetization before the next scan!

2.3 THE ECHO TIMES

Another parameter which appears in the sequence is TE, the so-called echo time. It is the time between the pulse and the instant you reach the center of the line in k-space, i.e., until you reach the green point in the example scan below::




This includes the time of the selective pulse, rewinding, and (half of) the actual time spend acquiring.

What is special about the “echo time”? The signal in k-space is proportional to the Fourier-transform of the image. To a very good approximation, the signals along the

$k_{\text{phase}}=0$ and $k_{\text{read}}=0$ lines are the largest (with the largest value appearing at $k_{\text{phase}}=k_{\text{read}}=0$). Here is a numerical example: on the one hand, here is a familiar image to all Windows XP users:




and here is the logarithm of the absolute magnitude of the Fourier transform of its red mask, as computed in MATLAB:



This illustrates the above claim. Hence, when you think about your signal, it will peak when you're at the center of the acquired line (at $k_{\text{read}}=0$). This peak is called the echo.

In terms of spins, what's happening is that the initial negative G_{read} gradient dephases the spins, and at $k_{\text{read}}=0$ the spins are all aligned along the same direction, as is illustrated in the following schematic 1D example:



Before acquiring (A), $k_{\text{read}}<0$ and the gradient has “dephased” the spins. At $k_{\text{read}}=0$, $e^{ikx}=1$ (since $k=0$) and the spins all align. At the end of the line, $k_{\text{read}}>0$ and the spins have once again dephased.

2.4 CONTRAST IN GRE

Right after the selective pulse, we have

$$M_{xy}(\mathbf{r}) = \frac{1 - e^{-TR/T_1(\mathbf{r})}}{1 - \cos(\alpha)} \sin(\alpha) M_0(\mathbf{r})$$

This means that, at the echo time,

$$M_{xy}(\mathbf{r}) = \frac{[1 - e^{-TR/T_1(\mathbf{r})}] \sin(\alpha)}{1 - \cos(\alpha)} e^{-TE/T_2^*(\mathbf{r})} M_0(\mathbf{r})$$

Hence, you're not imaging $M_0(\mathbf{r})$, or even $M_0(\mathbf{r})e^{-TE/T_2}$ (note it's T_2^* and not T_2 because we need to take the field inhomogeneities into account). Instead, we're imaging a “pseudo-spin-density” given by the above expression, which is a function of T_2^* , T_1 , α , TE and TR . This is interesting because we can now measure T_1 and T_2^* - albeit not directly. Images for which the image depends on T_1 or T_2 are called “ T_1 -weighted” and “ T_2 -weighted”, respectively. We get a little of both. All images are proton-density weighted because they all have $M_0(\mathbf{r})$, which is

proportional to the density of protons, which is proportional to the density of water/fat.

2.5 FLASH

Acronyms are very important in MRI. They make one feel as if they're doing something more complicated than whatever it is they really are doing. The spoiled-GRE sequence outlined above is sometimes referred to as FLASH on Siemens machines (or T1-FFE on Philips scanners). FLASH stands for Fast Low Angle Shot.

2.6 TURBO-FLASH

When α is very small (~5-10 degrees or less), the GRE sequence outlined above is called turbo-FLASH (on Siemens). Another name it's given is RAGE (Rapid Acquisition by Gradient Echo). In that case, we can approximate (α must be in radians for this to work):

$$\sin(\alpha) \approx \alpha$$

$$\cos(\alpha) \approx 1$$

so

$$\begin{aligned} M_0^{(\text{eff})}(\mathbf{r}) &\approx \frac{[1 - e^{-TR/T_1(\mathbf{r})}]\alpha}{1 - e^{-TR/T_1(\mathbf{r})}} e^{-TE/T_2^*(\mathbf{r})} M_0(\mathbf{r}) \\ &= \alpha e^{-TE/T_2^*(\mathbf{r})} M_0(\mathbf{r}) \end{aligned}$$

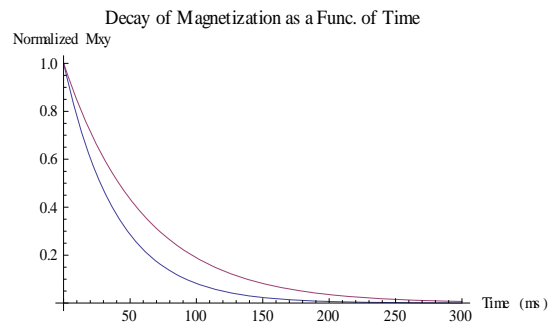
The sequence becomes approximately independent of TR and T_1 , meaning we can select TR to be as short as possible. Typical TR values can range as low as 10ms. The small tip-angle makes it easier for the magnetization to return to equilibrium faster. The drawback? Low signal. Note M_{xy} is proportional to α , which is small by definition. RAGE sequences are often combined with other, preparatory blocks,

such as inversion recovery which was described in chapter 3. I hope we'll get a chance to talk about those in a subsequent lecture.


2.7 GETTING GOOD CONTRAST

Contrast refers to the ability to distinguish between two tissue types having different parameters (T_2 , T_1 , etc). Suppose we're given two tissues which are identical, and only differ in their T_2 values. How would we go about getting the best contrast between them using FLASH?

Consider two tissues having, say, $T_2^* = 40$ ms and $T_2^* = 60$ ms. Assuming all other parameters are equal, here is a plot of the magnitude of M_{xy} , as a function of time, neglecting all common factors except the e^{-t/T_2} factor (see Sec. 2.4):



The purple and blue curves correspond to $T_2^* = 60$ and $T_2^* = 40$, respectively. The **difference** between the two decays is greatest when approximately $t \approx$ average of the two T_2 s. In fact, here is the difference, $e^{-t/T_2^{(1)}} - e^{-t/T_2^{(2)}}$, plotted:



We can find the minimum directly by differentiating $e^{-t/T_2^{(1)}} - e^{-t/T_2^{(2)}}$ and equating the result to zero:

$$\frac{d\left(e^{-t/T_2^{(1)}} - e^{-t/T_2^{(2)}}\right)}{dt} = \frac{e^{-t/T_2^{(2)}}}{T_2^{(2)}} - \frac{e^{-t/T_2^{(1)}}}{T_2^{(1)}} = 0$$

This leads to

$$t_c = \frac{T_2^{(1)}T_2^{(2)}}{T_2^{(1)} - T_2^{(2)}} \ln\left(\frac{T_2^{(1)}}{T_2^{(2)}}\right) \approx \frac{|T_2^{(2)} + T_2^{(1)}|}{2}$$

(the approximation is valid for $T_2^{(1)} \approx T_2^{(2)}$)

This leads to the following rule of thumb:


To maximize T_2^* contrast in spoiled GRE sequences, set

$$TE \approx \frac{|T_2^{(2)*} + T_2^{(1)*}|}{2}$$

Similar reasoning can be applied to maximizing T_1 and $M_0(\mathbf{r})$ (aka proton-density) contrast – some of which you'll apply in your homework exercises.


3. SPOILED SE IMAGING

The basic spoiled SE (Spin Echo) sequence looks a lot like the spoiled GRE one, only now the read gradient doesn't get reversed, and an additional π -pulse is used just prior to acquisition:



The exact same building blocks of the spoiled GRE sequence are used here. A few differences, however, need to be highlighted:

1. Note that now, the π pulse reflects our position in \mathbf{k} space through the origin ($\mathbf{k} \rightarrow -\mathbf{k}$), so in terms of k -space, what we're doing is, e.g., for the first and second scans (shown in black & red, respectively):



We're still scanning k -space, but in a slightly different manner. The advantage? The π -pulse refocuses magnetic field inhomogeneities and

generally produces better quality images.

2. In SE imaging, one usually takes $\alpha=\pi/2$; that is, excites the spins completely onto the xy -plane. This is to ensure the magnetization has no z -component after the slice-selective pulse. Why is that? Suppose we did have such a component. Then the π -pulse would invert it as well before the next slice-selective pulse. If you'd go back to the GRE section, one of the assumptions we've made is that M_z evolves freely between slice-selective pulses; inverting it with a π pulse means our previous reasoning is no longer valid, so we can't use it. While it eventually does settle in some sort of dynamic equilibrium, experience has shown that this state (obtained using flip angles α other than $\pi/2$) yield little to no additional benefits.
3. From (2), we can write down the magnetization at the time of the echo as:

$$M_{xy}(\mathbf{r}) = \left[1 - e^{-TR/T_1(\mathbf{r})}\right] e^{-TE/T_2(\mathbf{r})} M_0(\mathbf{r})$$

4. Note that the decay of the signal in the xy -plane is governed by T_2 , **not** T_2^* , because of the π pulses which refocus the field inhomogeneities, leaving just the intrinsic, microscopic relaxation mechanisms.
5. Because the π pulses take time and because α is very large and the magnetization needs some time before some appreciable z -magnetization is built up, TR tends to be longer in SE sequences compared to GRE sequences. On the other hand, it tends to produce better quality images, because the field

inhomogeneities are refocused by the π pulse. SE imaging was, and still is, considered a "gold standard" as far as image quality goes.

4. UNSPOILED IMAGING

Our entire discussion so far has assumed we've spoiled the magnetization between pulses. It turns out, however, that the magnetization reaches dynamic equilibrium even without spoiling; it's just a lot more complicated. Unfortunately, we don't have the time in this course to discuss unspoiled sequences such as SSFP, bSSFP, FISP, GRASS and others. Those of you who are interested are welcome to ask me for references.