
LECTURE 5
BASIC IMAGING
Lecture Notes by Assaf Tal

The current lecture concerns itself with spatially encoding information: how is it that we're able to image the spatial positions of our spins? We are going to go over the two fundamental approaches for doing so, called **frequency encoding** and **phase encoding**. Once we provide an intuitive understanding of those, we will introduce a mathematical formalism called **k-space** that unifies them. The next lecture will then delve into the mathematical process of reconstructing data acquired in k-space.

FREQUENCY ENCODING

We Cannot Resolve Our Signal Spatially Without a Gradient

Since MRI happens in the near field we have no spatial control over our fields. As we've previously seen, the acquired signal has no dependence on position:

$$s(t) \propto \omega_0 \int_{body} \overline{B_{xy}^{(rec)}(\boldsymbol{r})} M_{xy}^{(rot)}(\boldsymbol{r},t)dV$$

We could get some selectivity by shaping the field of our receiver. For example, in the early days of MR people would acquire a signal by placing a coil close to the object of interest:

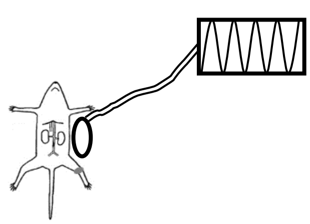

By placing a coil close to the rat's kidney, one can pick up signal mostly from the kidney where $\mathbf{B}_{rec}$ is strongest.

This is highly inefficient; The reception pattern is spatially inhomogeneous (since $\mathbf{B}_{rec}$ is inhomogeneous) and it requires one to mechanically move the coil (or subject) to change their sampling point.

In MRI one takes a different approach: first, the receiver's field is made as homogeneous as possible over the object, and gradient fields are used to spatially resolve our signal through one of two methods: frequency encoding or phase encoding. We describe them in order.

Intuitive Description of Frequency Encoding

Imagine exciting the spins in the object onto the xy-plane. In the absence of a gradient, all the spins see the same field and precess at the same frequency. Imagine being able to "listen" to their frequencies: you would hear one well defined tone.

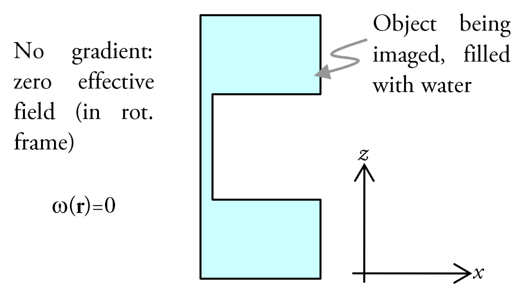

Once you turn on a gradient, say, along the z-axis, we generate a different frequency along each point on the z-axis. You can imagine each position being assigned a different "key" on a "piano":

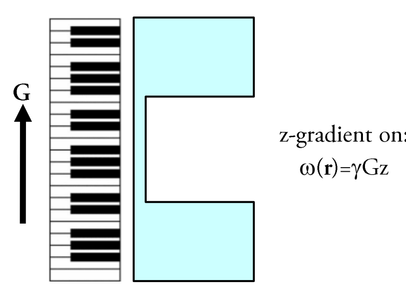

The intensity we "hear" at frequency - at each piano key - is proportional to the number of spins at that position. Therefore, **by "listening" to the signal we can deduce the distribution of spins along the z-axis.** This is the basic idea behind frequency encoding. The image obtained would give us a **projection** of the density of spins on the z-axis:

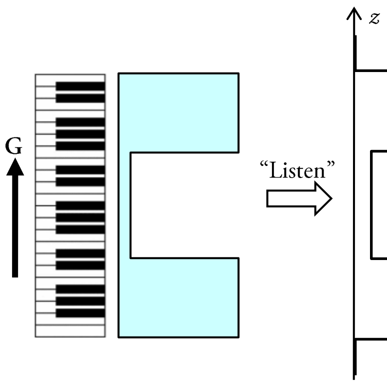

By changing the direction of the gradient we alter the axis of projection, which is parallel to **G**:

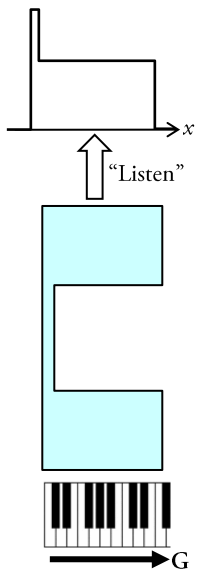

The gradient, **G**, can point in any direction in space, meaning we can create a projection of the 3D spatial distribution of our spins along any spatial axis. For example, **G**=(G_x, G_y, 0) points along an axis in the xy plane which is titled by 45° relative to the x-axis.

The Fourier Transform

We will delve into the math next lecture, but I will just remark that, while our ear can naturally separate a signal into its frequency components, this is not a trivial mathematical operation. For example, here is a signal that is comprised of three frequencies – 10 Hz, 17 Hz and 37 Hz – at relative amplitudes 1:2:3. Few people would be able to deduce this information just by looking at the signal:

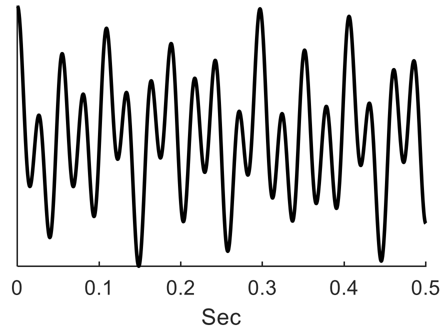

There is, however, a “magical box”, called a **Fourier transform**, which can take a signal such as the one above and spit out a function that describes its frequency content:

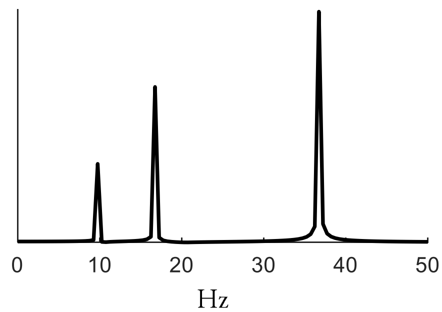

Each peak signifies a separate frequency component. They are located at the respective frequencies (10, 17, 37 Hz) and their intensities are at a 1:2:3 ratio.

After acquiring a signal in the time domain in the presence of a gradient, a Fourier transform (FT) is applied to separate it into its frequency components. The resulting signal in the frequency is a 1D projection of the 3D object, along an axis parallel to the gradient **G**.

The Frequency Encoding Pulse Sequence

The frequency encoding experiment can be described with a **pulse sequence**, which is a diagram indicating the timing and amplitudes of the RF and gradient channels, as well as delays and acquisition blocks:

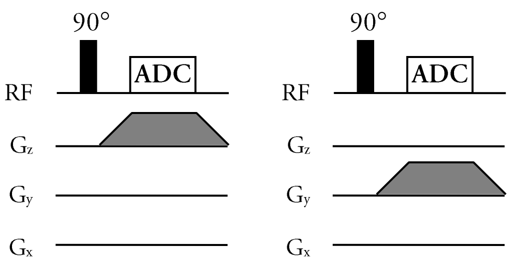

Left: pulse sequence depicting a volume excitation of the entire sample (i.e. not a slice selective excitation), followed by frequency encoding along z. This will create a 1D projection of your object along the z-axis. Right: same sequence, only with frequency encoding along y. This will create a 1D projection of the 3D object along the y-axis.

Multiple Frequency Encoding Scans Can Be Used to Reconstruct an Image Via Projection-Reconstruction

Although it might be unclear how to do this, the reader might feel that, given 1D projections along enough axes, one could infer the 3D spatial distribution of spins in the sample. This is correct, and is known as **projection reconstruction**. This is how Computerized Tomography (CT) scanners work. Some MRI experiments do use this approach, the most famous being the nobel prize winning paper[1] by Paul Lauterbur which first introduced this concept and kickstarted the MRI field, where the first figure shows projections of two test tubes filled with water:

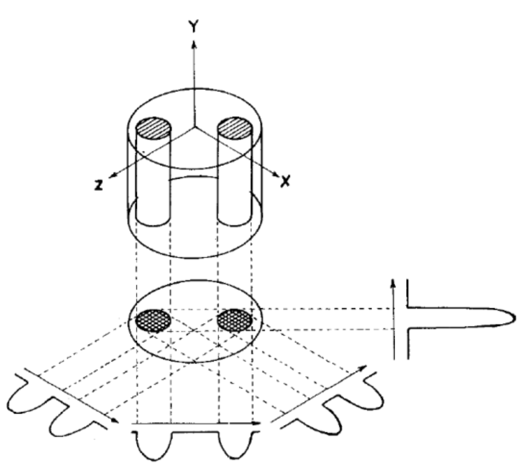

[1] PC Lauterbur, Nature 242:190-191 (1973)

The reader is encouraged to dig up the paper and enjoy this piece of scientific history in the making.

Phase Encoding

A Precessing Spin Has a Phase

The phase of a precessing spin is defined as the phase its projection on the xy-plane makes with the x-axis:

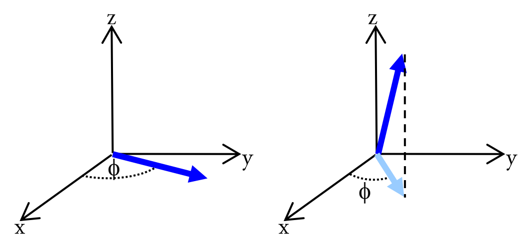

Left: the phase of a spin in the xy-plane is the angle it makes with the x-axis. **Right:** for a spin not in the xy-plane, one examines the phase its projection on the xy-plane makes with the x-axis.

One can think of this mathematically as follows: suppose you are given a magnetization vector

$$\mathbf{M} = \begin{pmatrix} M_x \\ M_y \\ M_z \end{pmatrix}.$$

Its projection on the xy plane is the two dimensional vector

$$\begin{pmatrix} M_x \\ M_y \end{pmatrix}.$$

One can form the complex quantity $M_{xy} = M_x + iM_y = |M_{xy}| e^{i\phi}$. The phase ϕ of this quantity is precisely the phase of the spin.

A spin precessing with a constant angular frequency ω for a time t will accumulate a phase:

$$\phi(t) = \omega t .$$

If ω is time dependent, one must break down the time into small intervals Δt during which ω is approximately constant, and sum up the phase contributions from each:

$$\phi = \omega(t_1)\Delta t + \omega(t_2)\Delta t + \ldots \omega(t_N)\Delta t \,.$$

As $\Delta t \to 0$ this becomes an integral:

$$\phi(t) = \int_0^t \omega(t')\,dt' \,.$$

Precession in the Presence of a Gradient Creates a "Spin-Helix" Along The Gradient Axis

Once a gradient is turned on, the frequency becomes spatially dependent, and so does the phase of the spins. For a z-gradient, $\omega = \gamma G z$ and:

$$\phi(t) = \gamma G z t \,.$$

One should visualize this as a "helical winding" of the spins along the gradient axis:

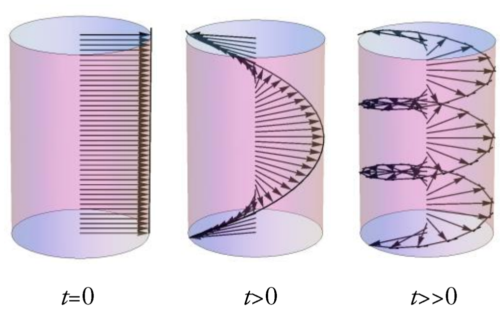

In the presence of a z-gradient, the spins precess with a frequency $\omega=\gamma Gz$ which induces a linear phase $\phi(t)=\gamma Gzt$, imparting a helical shapre to the tips of the magnetization vectors as time progresses.

The Principle of Phase Encoding (PE)

Imagine having a different density of spins at two points, A and B, along the z-axis. Our task is to deduce the density of spins at the two points, M_A and M_B:

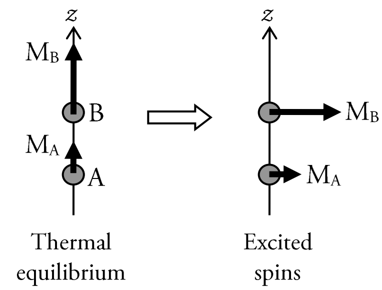

By exciting and acquiring, the signal we would measure would originate from both points equally, and would lack any spatial selectivity:

$$s_1 \propto M_A + M_B \,.$$

We now run a second experiment, in which we apply a gradient just for long enough for the spins to go out of phase (let's assume for simplicity that A corresponds to z=0, so the spin there is stationary):

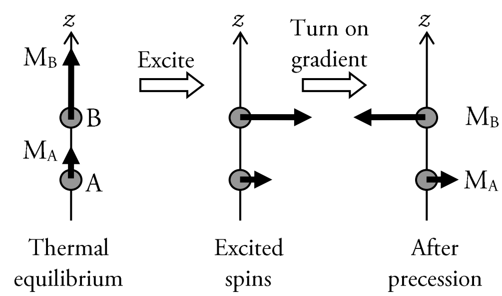

At this point when we acquire a signal, it will be proportional to

$$s_2 \propto M_A - M_B$$

By adding and subtracting the two experiments, one can extract just the signal from A, or just the signal from B:

$$s_1 + s_2 = M_A$$
$$s_1 - s_2 = M_B$$

This idea can be extended to more than just two positions and more than one dimension: **by performing multiple experiments, creating a unique phase distribution in each experiment (using the gradients) and taking linear combinations of those experiments, the signal from multiple points in the sample can be recovered**. That is the principle of phase encoding.

In terms of pulse sequences, the two experiments would look like this:

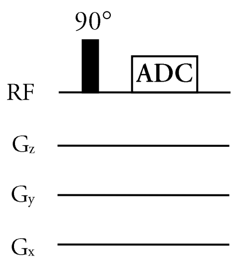

Experiment 1:
First PE step

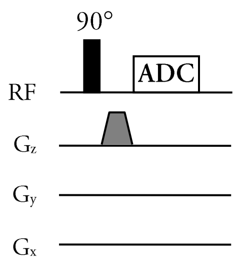

Experiment 2:
Second PE step

These are often notated as:

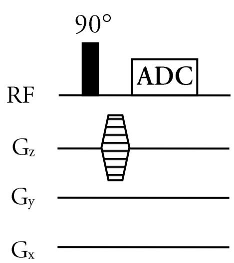

The striped gradient shape indicates that is a phase encoding gradient, which should be incremented by some fixed amount between successive phase encoding steps.

In the above example we needed 2 experiments to phase encode two spatial points, A and B. In general, **one would need N experiments to retrieve N spatial points**. This is true even if the points are in 3D.

Number time. The human head can be enclosed in a box about 20×18×16 cm. If we wanted a spatial resolution of 1 mm^3, we would need $200\times180\times160\approx6\cdot10^6$ voxels, which would require about half a million scans using phase encoding! If we also wait for the spins to return to thermal equilibrium a time ~ T_1 ~ second, this will take ~10^7 seconds to complete, which is ~100 days! As the sole method of imaging, phase encoding is unsuitable for high resolution spatial imaging.

The Origin of the Point Spread Function

Imagine now we were to repeat the above phase encoding experiment, but allow for a continuous distribution of spins, $M_{xy}(z)$, right after excitation. It's obvious that adding and subtracting the two experiments would indeed give us M_A and M_B, but they would be "contaminated" by signals from spins from other positions, which would not cancel out completely. The amount of contamination would depend of course on the exact form of $M_{xy}(z)$.

Phase Encoding is Often Combined with Frequency Encoding and Slice Selection along Orthogonal Axes

To save time, frequency and phase encoding are often combined: a single frequency encoding experiment – which produces a 1D projection – is repeated with phase encoding along an orthogonal axis, all inside a thinly excited slice. In terms of pulse sequences, we might have:

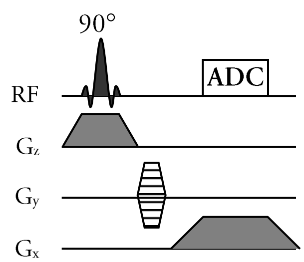

Frequency Encoding is "Fast", Phase Encoding is "Slow"

To produce a 2D image with resolution $N_x\times N_y$, we can use frequency encoding along one axis (say, x) and phase encoding along the other (say, y). We will need a total of N_y scans to acquire the full image, since phase encoding requires one scan for each point it wants to "separate" from all of the others. The N_x points along the frequency encoded axis are for "free"! We'll get a better grasp of this using the mathematical formalism of k-space.

The Imaging Toolbox

Summary: How to Spatially Encode an Image

All methods of spatial encoding fall into one of four categories:

1. **Frequency encoding**, by which positions along a particular axis in space are encoded with different frequencies ($\omega = \gamma \boldsymbol{G} \cdot \boldsymbol{r}$) which are read out in one-shot and then separated using a Fourier transform.
2. **Phase encoding**, by which the phases of the spins are changed in each excitation in a manner which lets us separate the signal from different positions.
3. **Selective excitation**, by which a frequency selective pulse is applied in the presence of a spatial gradient. The pulse, which excites a range of frequencies, will excite a range of positions along the gradient's axis. This effectively excites a 2D "slice". By exciting different slices each time, one can in theory deduce the 3D distribution of spins in the body.
4. **Sensitivity encoding**, by which we use the fact that a small-enough coil only picks up a signal from a localized region in space. By using multiple coils, or perhaps by moving a single coil around and exciting and acquiring signal from different positions, we can uncover the spatial distribution of spins in the body.

The first two approaches are by far the most common. They only require the ability to turn on and off and manipulate gradients, which is readily available to us. They are often combined with the last two approaches; for example, using multiple receiver coils we can accelerate the imaging process, an approach called Parallel Imaging.