

XI SIGNAL-TO-NOISE (SNR)

Lecture notes by Assaf Tal




1. NOISE

1.1 CHARACTERIZING NOISE

Noise is a random signal that gets added to all of our measurements. In 1D it looks like this:



while in 2D it looks like the “snow” on your TV screen:



Noise is unavoidable. It comes from resistive elements in our system: the electronics, and even the patient who has some finite resistance. Microscopically speaking, it is because of the thermal fluctuations of our system: in the electronics and of the spins in the patient.


A noise signal, $n(t)$, cannot be represented by an analytical function. To characterize noise, we need to speak in statistical terms. The two most important characteristics are the mean (also known as the average) and the variance of the noise. They are denoted by $\langle n \rangle$ and $\langle n^2 \rangle$, respectively, and are illustrated below:

In “well behaved” systems, the mean of the noise is 0; it sometimes gets added and sometimes gets subtracted from our signal at random. When $\langle n \rangle$ is non-zero the signal is said to be “biased”. Bias is easy to detect and remove, so we won’t focus on it here, and will assume the mean of our noise is 0.

The variance of our noise is basically its “size”. When the variance of the noise becomes as large as the signal being measured, it becomes extremely hard to discern the two. Ideally, we would like to make the variance as small as possible. In practice, we often settle for making it “small enough”; that is, small enough *with respect to the signal we’re looking at*, so as to make the features that interest us discernable. This chapter will mostly be devoted to ways of making the noise’s variance “small enough”.

1.2 SIGNAL TO NOISE RATIO (SNR)


As noted before, noise gets added to every signal we measure.




The signal to noise ratio of the noisy signal is defined as:

$$\text{SNR} = \frac{\text{magnitude of signal}}{\text{variance of noise}}$$

If the signal changes considerably from point to point, so will the SNR. In other words, the SNR is a function of position (if we're dealing with images) or time (if we're dealing with time signals). For example, take a look at this function:



Now we add some random noise to it:







The SNR of the first few "steps" is much lower than the SNR of the last few "steps".

A low SNR means the noise is "big" and will make it hard for us to determine the magnitude of the signal accurately, or even see the signal. A high SNR means the noise is "small".


Measuring the SNR in practice can be tricky. The best thing to do is to find some region in your signal/image where you know there's only noise, and to use it to estimate your noise's variance.

1.3 ADDING NOISE

Suppose you have several noisy, independent random signals, with the same variance, and you add them together. What will be the variance of the sum? Let's do a little experiment. Here are four random signals:



and here is their sum:



You can observe visually that the variance of each individual noisy signal is about 0.5, and the variance of the sum is about 1. So, while we've added 4 random signals, we didn't increase the variance by a factor of 4. In fact, we've increased it only by a factor of $\sqrt{4} = 2$. This is a general fact about random signals:

Adding N random signals, each having the same variance, X, will yield a random signal with variance $\sqrt{N}X$.

I will not prove this fact in these notes. Intuitively, though, expecting an increase by a factor of N is unreasonable, because the noise sometimes adds constructively and sometimes destructively. This leads to a corollary:

Adding N measurements improves the SNR by a factor of \sqrt{N} .

This is because the signal multiplies by N, the noise's variance by \sqrt{N} , and their ratio by \sqrt{N} :

$$\text{SNR} = \frac{\text{signal}}{\text{noise}} \xrightarrow{N \text{ scans}} \frac{N \times \text{signal}}{\sqrt{N} \times \text{noise}} = \sqrt{N} \times \text{SNR}$$

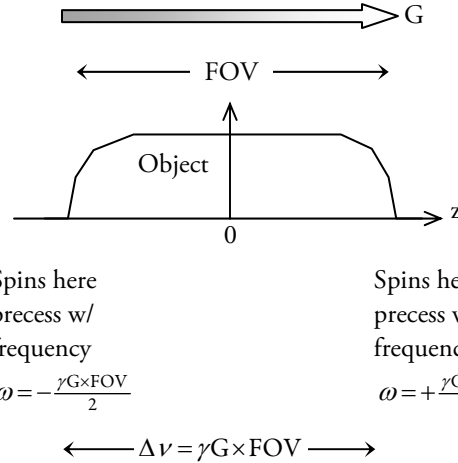
This is sometimes also called **signal averaging**.

1.4 NOISE IN MRI

A well known theorem from statistical mechanics (the so-called Nyquist theorem) states that, the variance in the voltage in an electronic system is given by

$$\langle V^2 \rangle = \sqrt{4k \times T \times R \times \Delta \nu}$$

where $k = 1.38 \times 10^{-23} \frac{\text{Joules}}{\text{Kelvin}}$ is Boltzmann's constant, T is the temperature of the system (in Kelvin), R is its resistance (in Ohms) and $\Delta \nu$ is the range of frequencies we're observing. For a typical 1D MRI experiment, where we acquire in the presence of a gradient G_{read} , $\Delta \nu = \gamma G_{\text{read}} \text{FOV}$:



What is R? There are two sources of resistance in an MRI experiment:

- The RF coils (R_c).
- The patient (R_p).

Both the coils and the patient are conductors, to a degree. When a magnetic field infringes upon a conductor it dissipates partially as heat. We are basically made out of water, which is a conductor. When a magnetic field tries to penetrate a conductor it creates "eddy currents" as it dissipates slowly. This is known as the **skin effect**. The currents induced in the patient then induce currents in the coils that are picked up as noise. This is called **patient loading**. It turns out that for high fields (~1 Tesla and above in practice), patient loading is more important than the intrinsic hardware noise.

Calculating a patient's resistance is difficult and usually not very constructive, so we will simply treat it as a constant. The only important fact is that it is (approximately) proportional to the square of the main field: B_0^2 . So

$$R = R_p + R_c \approx R_p \approx C B_0^2.$$


where C is some constant, whose value will be assumed unknown. It can be calculated, but it will not serve our purpose (which is inferring how changing the experimental parameters will affect the SNR). To sum up:

$$\langle V^2 \rangle = \sqrt{4kTCB_0^2 \gamma G_{\text{read}} \times \text{FOV}_{\text{read}}}$$

Note I've added the "read" subscript, to emphasize that the range of frequencies we observe during acquisition is determined by the **read** gradient (and not, say, the slice selection or phase-encoding gradients).

1.5 FOURIER TRANSFORMING NOISE

The MRI signal is measured in k-space and consequently Fourier transformed to yield an image. The Fourier transform of noise is just ... more noise.



Don't forget our FT is discrete: it's carried over a finite number of points. Because every point in the original (k-space) function affects every point in the Fourier (image) space, this means the noise at some point \mathbf{r} in our image is added up from all points in k-space. If we have a total of N points in k-space, then the variance of the noise in image space will increase as \sqrt{N} . However, the discrete Fourier transform also has a factor of $1/N$ in its definition. Without going into the technical details, here is the bottom line that's relevant for us:

Fourier transforming noise over a discrete set having N points decreases its variance by \sqrt{N} .

This works in 2D and 3D as well. For a 2D grid having N_x points along the k_x axis and N_y points along the k_y axis, the noise's variance will decrease by a factor $\sqrt{N_x N_y}$.

2. SNR IN MRI

2.1 RELATIVE SNR

Being able to calculate the absolute SNR might be interesting theoretically but quite formidable,

because the noise will ultimately depend on the hardware, patient, and image. Rather, one looks at how the SNR changes as we change the imaging parameters: the resolution, voxel size, FOV, acquisition time, gradients, B_0 (not really a parameter, but still interesting), etc. This is precisely what this section is all about.

2.2 SIGNAL


We've spent the last section talking about noise, but haven't said anything about signals. We've remarked how the finite sampling in k-space (k_{max}) causes blurring in the signal, so the signal at point \mathbf{r} is actually an average of the signal in a voxel around it:

$$I(\mathbf{r}) = \omega_0 \int_{\text{voxel centered around } \mathbf{r}} M_0(\mathbf{r}') d\mathbf{r}' \approx \omega_0 M_0(\mathbf{r}) \Delta V$$

where ΔV is the size of the voxel. Remember the ω_0 in front is there because of Faraday's law (the signal is proportional to the time derivative of the magnetization, which is proportional to ω_0 . See chapter 2, section 3.2).

2.2 SNR IN 3D IMAGING

Here is a 3D GRE sequence:



Suppose we've:

1. Collected N_x , N_y and N_z points along the k_x , k_y , k_z axes.
2. Have a voxel size $\Delta V = \Delta x \Delta y \Delta z$ (note the voxel doesn't have to be a cube, i.e., Δx isn't necessarily equal to Δy or Δz , etc).
3. Because we're reading along the x-axis, we have a bandwidth of $\Delta \nu = \gamma G_x \text{FOV}_x$.
4. The readout time along x is T_s (see sequence above).
5. We acquire the same image N_{acq} times (for signal averaging).

Then

$$\text{SNR}(\mathbf{r}) = \frac{\text{signal}(\mathbf{r})}{\text{noise}} \propto \frac{M_0(\mathbf{r}) \Delta V \sqrt{N_{\text{acq}}}}{\sqrt{\frac{\Delta \nu}{N_x N_y N_z}}}$$

Using

$$T_s = N_x \delta t = \frac{N_x \delta k}{\gamma G_x} = \frac{N_x \delta k \text{FOV}_x}{\gamma G_x \text{FOV}_x} = \frac{N_x}{\Delta \nu}$$

we get


$$\boxed{\text{SNR}(\mathbf{r}) \propto M_0(\mathbf{r}) \Delta V \sqrt{N_z N_y N_{\text{acq}} T_s}}$$

The SNR of a 3D scan (assuming read is along the x-axis)

This is also valid for 3D spin-echo imaging.

2.3 SNR IN 2D IMAGING

2D imaging is just like 3D imaging, with one omission: you don't sample along the 3rd dimension. Rather, you use a slice-selective gradient. Let's take the slice-selective direction to be z, with x & y being the read & phase axes, respectively. This means that we need to omit N_z from the above expression, because we're not Fourier-transforming over that direction. Furthermore, note that ΔV , the voxel's volume, has a thickness equal to the slice's thickness along the z-axis (let's call it naturally Δz):



Thus:

$$\boxed{\text{SNR}_{2D}(\mathbf{r}) \propto M_0(\mathbf{r}) \Delta V \sqrt{N_y N_{\text{acq}} T_s} = \frac{\text{SNR}_{3D}}{\sqrt{N_z}}}$$

In terms of SNR, 3D imaging is superior to 2D imaging. Should we always use 3D? That depends. 3D often requires more scans to achieve the same slice thickness, to avoid aliasing (you need smaller Δk 's to cover the entire FOV along the 3rd axis; in 2D imaging, you don't care about aliasing because you're selectively exciting a slice and imaging all of it). Another problem is "ringing" artifacts having to do with the Fourier reconstruction. As a rule of thumb, for "thick" slices (a few mm and above), you should go for 2D imaging. For "thin" ones ($-mm$ and thinner), go for 3D.

There are also other, phenomenon-specific reasons for going 2D; in MR "time-of-flight" angiography sequences, for example, slowly flowing blood yields better contrast in slice-selective, rather than 3D, imaging.

2.4 SNR DEPENDENCE ON B_0

One thing the above equations don't show is the dependence of the SNR on B_0 . Recall that, for high fields (which interest us),

$$\text{noise} \sim \sqrt{\langle V^2 \rangle} \sim \sqrt{B_0^2} = B_0$$

while

$$\text{signal} \sim \omega_0^2$$

(one ω_0 comes from M_0 , the other comes from the Faraday's law: the signal we acquire is proportional to the time derivative of $M_{xy} \sim e^{i\omega_0 t}$.) Hence:

$$\text{SNR}(B_0) \sim \frac{B_0^2}{B_0} = B_0$$

The SNR should increase linearly with the field. In practice, this is only approximate.