Lecture 9
SNR
Lecture Notes by Assaf Tal

Noise

MRI Noise is "White", And Characterized By Its Standard Deviation

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:

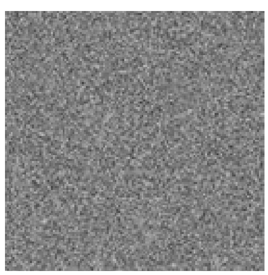

Noise is unavoidable. It comes from resistive elements in our system: the electronics, and even the patient who has some finite resistance. Microscopically speaking, it is because of the thermal fluctuations of our system: in the electronics and of the spins in the patient.

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 **standard deviation** of the noise. They are denoted by $\langle n\rangle$ and $SD \equiv \sqrt{\left\langle \left(n-\langle n\rangle\right)^2\right\rangle}$, 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 <n> is non-zero the signal is said to be **biased**, or have a **DC offset**. 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 SD of our noise is basically its "size". When the SD 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 SD 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 SD "small enough".

Noise Adds Up As The Square Root of the Number of Measurements

Suppose you have several noisy, independent random signals, with the same SD, and you add them together. What will be the SD of the sum? Let's do a little experiment. Plotted below are eight random signals with a standard deviation of unity, before and after addition:

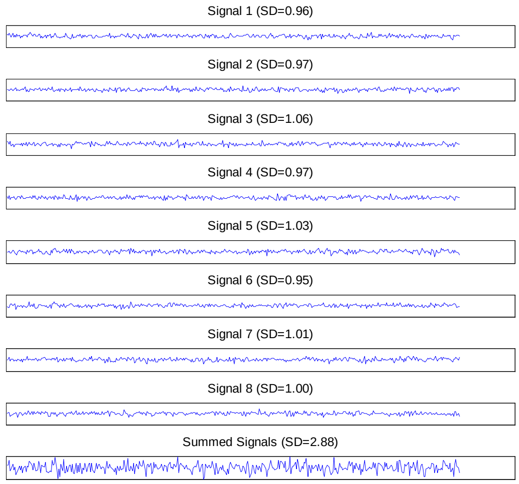

The scale of each plot is exactly the same (±8). You can observe visually that the SD of each individual noisy signal is about 1, and the SD of the sum is about 3. So, while we've added 8 random signals, we didn't increase the SD by a factor of about 3. In fact, we've increased it only by a factor of $\sqrt{8} \approx 2.83$ to be exact. This is a general fact about random signals: Adding N random signals, each having the same SD, X, will yield a random signal with SD $\sqrt{\text{N}}\text{X}$. To see this, suppose we have two noisy signals which are uncorrelated:

$$
\begin{aligned}
n &= \{n_1, n_2, n_3, \ldots\} \\
m &= \{m_1, m_2, m_3, \ldots\}
\end{aligned}
$$

We add them up:

$$n + m = \{n_1 + m_1, n_2 + m_2, n_3 + m_3, \ldots\}.$$

The standard deviation each point is the same as any other point. Looking at the first point, we calculate

$$
\begin{aligned}
\left[SD(n_1 + m_1)\right]^2 &= \left\langle (n_1 + m_1)^2 - \langle (n_1 + m_1) \rangle^2 \right\rangle \\
&= \left\langle n_1^2 - \langle n_1 \rangle^2 \right\rangle + \left\langle m_1^2 - \langle m_1 \rangle^2 \right\rangle \\
&\quad + 2\langle n_1 m_1 \rangle^2 - 2\langle n_1 \rangle \langle m_1 \rangle
\end{aligned}
$$

Because our noise is unbiased, $\langle m_1 \rangle = \langle n_1 \rangle = 0$. Furthermore, because our noise is uncorrelated[1], $\langle n_1 m_1 \rangle = 0$. Therefore

$$\left[SD(n_1 + m_1)\right]^2 = SD(n_1)^2 + SD(m_1)^2.$$

If n=m, we get (after taking the square root):

$$SD(2n_1) = \sqrt{2}SD(n_1),$$

proving the assertion. Intuitively, 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 $\sqrt{\text{N}}$.

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

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

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

Noise in MRI Images Comes From The Patient!

A well known theorem from statistical mechanics states that, the SD in the voltage in an electronic system is given by what's called the **Johnson noise** or the **Nyquist noise** or **thermal noise** of the system:

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

Resistance. Most of the resistance in MRI comes from the patient!

Bandwidth of frequencies being observed.

[1] imagine tossing two coins repeatedly and multiplying the results, where heads=(+1) and tails=(-1): on average you'd expect to get 0, although each experiment will be either +1 or -1

where $k = 1.38\times10^{-23}\ \frac{\text{Joules}}{\text{Kelvin}}$ is Boltzmann's constant, T is the temperature of the system (in Kelvin), R is the its resistance (in Ohms) and $\Delta\nu$ is the range of frequencies we're observing. What is R? There are two sources of resistance in an MRI experiment:

- The RF coils (R_c) and associated electronics.
- 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:

$$R = R_p + R_c \approx R_p .$$

The Patient Loading Increases as The Square of B_0

Here we show that

$$R_p \sim \omega_0^2 = \gamma^2 B_0^2 .$$

You can skip the proof without loss of continuity.

The human body has a certain conductivity σ which is tissue-dependent. A current flowing through a conductor will dissipate into heat because the conductor has some resistance. If we create a time varying flux through the conductor, it will create tiny currents called **eddy currents** which will, in turn, heat up the object – that is, the patient. This principle underlies some of the newer induction heating stoves being sold today.

If we denote by R the resistance of a conductor, by G=1/R its conductance, and by V and I the voltage across and current through the conductor, then the power dissipated in the conductor is simply:

$$P = I^2 R = V^2 G$$

The voltage by a sinusoidal RF field can be calculated via Faraday's law. Let's suppose we have a loop of area A through which we apply a perpendicular RF field of the form $B_1\cos(\omega_0 t)$:

$$V = -\frac{d\phi}{dt} \sim \omega_0 B_1 \cos(\omega_0 t) A$$

The conductance of the loop is

$$G = \sigma\cdot\left(\frac{\text{cross section wire}}{\text{length of wire}}\right) = \sigma\frac{C}{2\pi r}$$

so

$$P = V^2 G \sim \frac{\omega_0^2 B_1^2 AC}{2\pi r} .$$

This result will change for a different geometry, but what will remain the same is the dependence on the RF and B_0 fields, which is both quadratic. If we set I = 1 ampere we get the loop's resistance: $P=I^2R=R$. By reciprocity, when I=1 ampere, B_1 becomes the coil's sensitivity, so we can say

$$R_c \sim \omega_0^2\cdot\left(B_{xy}^{(rec)}\right)^2 .$$

Our conclusion assumes that B_1 is the same everywhere in the body, an assumption that breaks down at higher fields where the near field approximation is no longer valid. However, we'll put aside these issues and simply conclude that

$$R = R_p + R_c \approx R_p \sim \omega_0^2 .$$

The constant of proportionality will depend on the geometry of the body, its conductance, and on the coil's sensitivity pattern, factors we will not trouble ourselves with in this course.

The Noise Increases As The Square Root of the Bandwidth Per Pixel

For a typical 1D MRI experiment, where we acquire in the presence of a gradient G_{read}, $\Delta\nu=\gamma G_{read}FOV$:

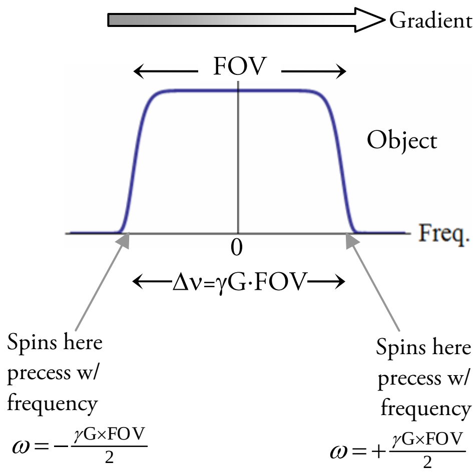

To sum up:

$$\sqrt{\left\langle V^2 \right\rangle} = \sqrt{4kTCB_0^2 \gamma G_{read} FOV_{read}}$$

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).

A Fourier Transform of a Signal With N Points Increases The SNR by √N

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.

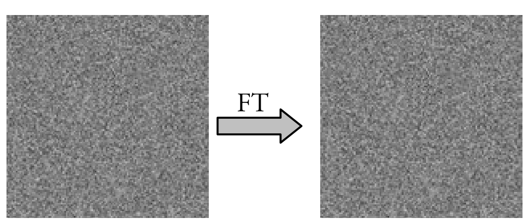

Don't forget our FT is discrete: it's carried over a finite number of points. Because every point in the original (k-space) function affects every point in the Fourier (image) space, this means the noise at some point **r** in our image is added up from all points in **k**-space. If we have a total of N points in k-space, then the SD of the noise in image space will increase as $\sqrt{N}$. This is a

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 SD 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 SD will increase by a factor $\sqrt{N_x N_y}$.

Signal

We've seen the signal in MRI is proportional to

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

The exact form of M_{xy} will depend on the sequence used. Let's assume for simplicity that we have a simple GRE acquisition, so after each excitation:

$$M_{xy}^{(rot)} = M_0 \frac{\left(1-e^{-TR/T_1}\right)\sin(\alpha)}{1-\cos(\alpha)e^{-TR/T_1}} e^{-TE/T_2^*}$$

with

$$M_0 = \frac{PD \cdot (\gamma \hbar)^2 S(S+1)}{3kT} B_0 .$$

The signal in each voxel will be proportional to $M_{xy}^{(rot)}$, assuming the different parameters are slowly varying on the scale of a voxel.

The Discrete Fourier Transform Given A Factor of N to the Signal Amplitude

Let's look at a simple 1D acquisition, in which we acquire N equi-spaced points in k-space:

$$s_j \propto \omega_0 \int_{\text{body}} \overline{B_{xy}^{(rec)}(x)} M_{xy}^{(\text{rot})}(x) e^{-2\pi i k_j x} dx .$$

Here $k_j = -\frac{k_{\max}}{2} + j\Delta k$ (j=0,1,...,N-1). Following a DFT, the signal from the j[th] voxel becomes:

$$\hat{s}_j = \omega_0 \int_{-\infty}^{\infty} \overline{B_{xy}^{(rec)}(x)} M_{xy}^{(rot)}(x) PSF(x_j - x) dx$$

where the PSF was derived in an earlier lecture:

$$PSF(x) = \frac{e^{\pi i \Delta k \cdot x} \sin(\pi k_{\max} x)}{\sin(\pi \cdot \Delta k \cdot x)} .$$

The overall shape of the PSF is such that it is comprised of a main lobe of width approximately given by $\Delta x = \frac{1}{k_{\max}}$, i.e. the voxel size. Furthermore, its height is obtained by taking the limit $x \to 0$ and using $\sin(x) \approx x$, which yields:

$$PSF(x) \xrightarrow{x \to 0} \frac{k_{\max}}{\Delta k} = N .$$

Thus, the **area** of the main lobe is $\Delta x \cdot N$ and we can approximate the signal as coming from x_j, assuming M_{xy} and $\overline{B_{xy}^{(rec)}(x)}$ are constant over the voxel's dimensions:

$$\hat{s}_j \approx \omega_0 \overline{B_{xy}^{(rec)}(x_j)} M_{xy}^{(rot)}(x_j) N \Delta x .$$

This can be immediately generalized to the case of sampling in 2D and 3D k-space:

$$\begin{aligned} \hat{s}_{jm}^{(2D)} &\approx \omega_0 \overline{B_{xy}^{(rec)}(x_j, y_m)} M_{xy}^{(rot)}(x_j, y_m) N_x N_y \Delta x \Delta y \\ \hat{s}_{jmn}^{(3D)} &\approx \omega_0 \overline{B_{xy}^{(rec)}(x_j, y_m, z_n)} M_{xy}^{(rot)}(x_j, y_m, z_n) \\ &\quad \cdot N_x N_y N_z \Delta x \Delta y \Delta z \end{aligned}$$

Non Fourier Transformed Axes Do Not Enjoy The ×N Factor

Multislice 2D imaging, in which slices (say, along z) are excited sequentially and each slice is phase encoded (say, in the k_x-k_y) plane, has a slightly different expression for its signal compared to the 3D case:

$$\begin{aligned} \hat{s}_{jmn}^{(multislice)} &\approx \omega_0 \overline{B_{xy}^{(rec)}(x_j, y_m, z_n)} M_{xy}^{(rot)}(x_j, y_m, z_n) \\ &\quad \cdot N_x N_y \Delta x \Delta y \Delta z \end{aligned}$$

Athough both sequences can cover 3D volumes, they differ by a factor N_z:

$$\hat{s}_{jmn}^{(3D)} = N_z \hat{s}_{jmn}^{(multislice)} .$$

To see why this is so, we go back to the signal equation prior to DFT. The (jmn) data point, which originates from the (k_x^j, k_y^m) point in the k_x-k_y plane in the n^{th} slice, is given by:

$$s_{jmn} \propto \omega_0 \cdot \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \left\{ \begin{matrix} \overline{B_{xy}^{(rec)}(x, y, z)} \\ \cdot M_{xy}^{(\mathrm{rot})}(x, y, z) dz \end{matrix} \right\} e^{-2\pi i \left(k_x^j x + k_y^m y\right)} dy dx$$

The signal is only phase encoded along k_x and k_y, which is why a factor $e^{-2\pi i k_z^n z}$ is lacking. Furthermore, is $M_{xy}^{(rot)}$ and $B_{xy}^{(rec)}$ do not vary inside the slice, we can approximate the integral over z:

$$\begin{aligned} &\int_{-\infty}^{\infty} \overline{B_{xy}^{(rec)}(x, y, z)} M_{xy}^{(\mathrm{rot})}(x, y, z) dz \\ &\qquad \approx \overline{B_{xy}^{(rec)}(x, y, z_n)} M_{xy}^{(\mathrm{rot})}(x, y, z_n) \Delta z \end{aligned}$$

Note the integration is really only carried out within the slice, since M_{xy} is 0 outside the slice (because no magnetization was excited outside the slice by assumption of a slice-selective pulse). Δz is the slice thickness. Using the above, we have

$$s_{jmn} \propto \omega_0 \Delta z \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \overline{B_{xy}^{(rec)}(x, y, z_n)} M_{xy}^{(\mathrm{rot})}(x, y, z_n) e^{-2\pi i \left(k_x^j x + k_y^m y\right)} dy dx$$

which is then DFT-ed along k_x and k_y to yield the signal from the (jmn) voxel:

$$\begin{aligned} \hat{s}_{jmn}^{(multislice)} &\propto \omega_0 \Delta z \overline{B_{xy}^{rec}(x_j, y_m, z_n)} M_{xy}^{(rot)}(x_j, y_m, z_n) \\ &\quad \cdot N_x N_y \Delta x \Delta y \end{aligned}$$

which proves our initial assertion.

Signal-to-Noise

"True" 3D Acquisitions

Focusing on a single voxel where $B_{xy}^{(rec)}$ and $M_{xy}^{(rot)}$ are more-or-less homogeneous, and assuming a true 3D acquisition, we can write down the SNR in the voxel as:

$$SNR^{(3D)}_{voxel} \propto \frac{\omega_0 B^{(rec)}_{xy}\sqrt{N_x N_y N_z}\, M^{(rot)}_{xy}\Delta x \Delta y \Delta z}{\sqrt{4kTC\omega_0^2 \left(B^{(rec)}_{xy}\right)^2 \gamma G_{read} FOV_{read}}}$$

We can simplify if we denote $\Delta V = \Delta x \Delta y \Delta z$, the voxel size, and plug in $M^{(rot)}_{xy}$ for the experiment at hand. Let's assume it's a spoiled GRE:

$$SNR^{(\text{3D GRE})}_{voxel} \propto \frac{PD\cdot(\gamma\hbar)^2 S(S+1) B_0 \Delta V (1-E)\sin(\alpha) e^{-TE/T_2^*}}{3kT\sqrt{4kTC\gamma G_{read}\times FOV_{read}}\,(1-\cos(\alpha)E)}$$

.

where $E \equiv e^{-TR/T_1}$. Some factors, such as $\hbar$, C and kT are constant in MRI and uninteresting from our perspective, so we will drop them and remain with

$$SNR^{(\text{3D GRE})}_{voxel} \propto \frac{PD\cdot\gamma^{3/2}\cdot B_0\sqrt{N_x N_y N_z}\,\Delta V(1-E)\sin(a)e^{-TE/T_2^*}}{\sqrt{G_{read} FOV_{read}}\,(1-\cos(\alpha)E)}$$

2D Multi-slice Acquisitions And The Multiplexing Advantage

We seen that the DFT affects the signal and the noise in the following manner:

$$\text{Signal} \rightarrow N\times\text{Signal}$$
$$\text{Noise} \rightarrow \sqrt{N}\times\text{Noise}$$

Thus, Fourier-transforming along any axis will increase the SNR by $\sqrt{N}$, where N is the number of voxels along that dimension:

$$\text{SNR} \rightarrow \sqrt{N}\times\text{SNR}\ .$$

This is called the **multiplexing advantage** by some authors. The idea is simple, and let's think for a moment about the slice (foot-to-head) direction: in each excitation we acquire the entire volume. Then when we reconstruct our signal, the signal from the j$^{\text{th}}$ slice originates from all N k-space points, which the DFT adds up with varying phases. It is therefore a form of signal averaging. Since there are N slices we end up averaging N "signals", leading to the $\sqrt{N}$ factor (e.g. slice #2 in the following illustration):

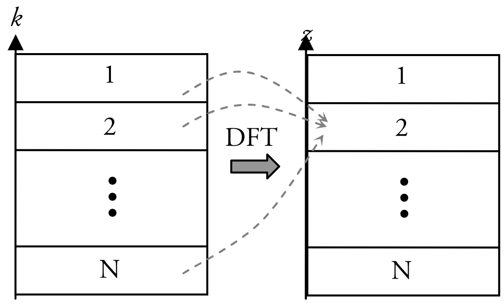

In particular, it should be clear that multislice 2D acquisitions do not enjoy the multiplexing advantage along the slice direction, and the reason should be clear: each excitation excites only a single slice instead of the entire volume, acquiring "less signal". Thus, you will see in some books statements such as:

$$\frac{SNR^{(3D)}}{SNR^{(2D)}} = \sqrt{N_z}\ .$$

This is true but it also omits other important factors. For example, suppose we have N_z slices, and within each slice we phase encode the y-axis and frequency-encode the x-axis. If the total acquisition time T_{acq} is fixed, this would imply that

$$TR = \frac{T_{acq}}{N_y N_z}\ .$$

However, comparing 2D multi-slice and 3D acquisition the **effective TR** for the multislice acquisition is $N_z\cdot TR$. This effective TR_{eff} is defined as the time between sequential excitation of *the same slice*:

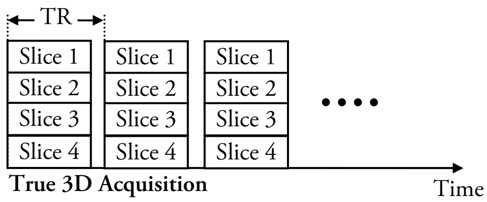

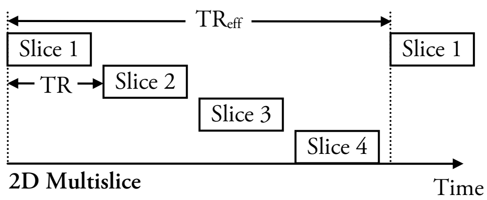

Thus (assuming the total acquisition time is the same),

$$TR_{eff} = N_z \cdot TR .$$

It is the **effective** TR that enters into the dynamic equilibrium factor, since that is the period of time between successive excitations of the same group of spins:

$$\frac{\left(1-e^{-TR_{eff}/T_1}\right)\sin(\alpha)}{1-\cos(\alpha)e^{-TR_{eff}/T_1}} .$$

There is therefore substantially less signal saturation in 2D multislice acquisitions, which can often make them as attractive - or even moreso - than true phase-encoded 3D acquisitions.

In reality, spins with short T_1s are ideal candidates for true 3D acquisitions, since even if we excite them rapidly they still manage to relax quickly back to thermal equilibrium. This isn't always the case for protons, but other spin species - e.g., ^{17}O - have extremely short T_1s (< 1 ms) and are almost impossible to saturate.

Dependence of SNR on Imaging Parameters

Increasing the Bandwidth Per Pixel Decreases SNR but Increases Robustness

A quantity that makes an appearance in many imaging sequences is the bandwidth per pixel: the number of Hz across a single voxel once the readout gradient is turned on. Shortly, if G_{read} is the readout gradient, and Δx is the pixel size (assuming readout is along the x-axis), then the bandwidth per pixel $BW_{1/N}$ is:

$$BW_{1/N} = \not{\gamma} G_{read} \Delta x$$

The BW for the entire FOV is $BW = \not{\gamma} G_{read} FOV_{read} = \not{\gamma} G_{read} N_x \Delta x$, so:

$$BW_{1/N} = \tfrac{BW}{N_x} .$$

This means we can rewrite the above using:

$$SNR_{voxel}^{(\text{3D GRE})} \propto \frac{PD \cdot \not{\gamma}^2 \cdot B_0 \sqrt{N_y N_z}\,\Delta V (1-E)\sin(a)e^{-TE/T_2^*}}{\sqrt{BW_{1/N_x}}\left(1-\cos(\alpha)E\right)} .$$

Thus we see that, along the readout, the SNR is proportional to $1\big/\sqrt{BW_{1/N}}$.

It would seem that increasing $BW_{1/N}$ is detrimental. However, there are also associated benefits: a higher $BW_{1/N}$ minimizes chemical shift displacement and the effects of B_0 inhomogeneity.

Dependence of SNR on Readout Duration

We've just seen that $SNR \propto 1\big/\sqrt{BW_{1/N}}$. This can be slightly rewritten: we know that

$$BW_{1/N} = \not{\gamma} G_{read} \Delta x$$

but also that

$$\Delta x = \frac{1}{k_{\max,x}} = \frac{1}{\not{\gamma} G_{read} T_{read}}$$

where T_{read} is the readout time (i.e. the time during which we acquire a signal while the readout gradient is on). Plugging this back into $BW_{1/N}$:

$$BW_{1/N} = \frac{1}{T_{read}}$$

and so we could equally say:

$$SNR_{voxel}^{(\text{3D GRE})} \propto \frac{PD\cdot\gamma^2\cdot B_0\sqrt{N_y N_z T_{read}}\,\Delta V\left(1-E\right)\sin\left(a\right)e^{-TE/T_2^*}}{\left(1-\cos\left(\alpha\right)E\right)}$$

Dependence on Voxel Size & Number of Voxels (for a Fixed FOV)

If we fix the scan time and number of voxels, our expression for the SNR clearly shows

$$SNR \propto \Delta V \qquad (N_i,\ FOV_i \text{ fixed, } i=x,y,z).$$

However, this is rarely the case in practice. When the voxel volume is halved, the number of voxels is usually doubled because one is often interested in keeping the FOV fixed. One must then also decide whether to keep the BW per pixel fixed (meaning you would have to change either G_{read} or FOV_{read}) or not. If we assume the **total** bandwidth is fixed, so both G_{read} and FOV_{read} remain fixed, our SNR expression tells us that

$$\begin{cases} SNR_{voxel}^{(3D)} \propto \dfrac{1}{\sqrt{N_x N_y N_z}} \propto \sqrt{\Delta V} & FOV,\ BW,\ \text{total acq. time fixed} \\ SNR_{voxel}^{(multislice)} \propto \dfrac{1}{\sqrt{N_x N_y} N_z} \propto \sqrt{\dfrac{\Delta V}{N_z}} & (\text{slice direction: } z) \end{cases}$$

Note the "hidden" assumption: if we keep the same total acquisition time fixed then TR - the time it takes to read out a k-space line in each slice - must be halved, which may or may not be possible.

If the bandwidth per voxel along the readout direction (x) is kept fixed (as opposed to the total BW along the readout direction) - by increasing the gradient and therefore the noise - the above expressions need to be amended by dividing by $\sqrt{N_x}$:

$$\begin{cases} SNR_{voxel}^{(3D)} \propto \dfrac{1}{N_x\sqrt{N_y N_z}} \propto \sqrt{\dfrac{\Delta V}{N_x}} & FOV,\ BW_{1/N},\ \text{total acq. time fixed} \\ SNR_{voxel}^{(m.s.)} \propto \dfrac{1}{N_x\sqrt{N_y} N_z} \propto \sqrt{\dfrac{\Delta V}{N_x N_z}} & (\text{slice direction: } z) \end{cases}$$

Dependence of SNR on Hardware/Sample

Dependence of SNR On Main Field (B_0)

The above expression seems to suggest the SNR increases linearly with B_0. This however does not take into account hidden dependencies of T_1 and T_2 on B_0. Empirical evidence suggests that for biological tissue and at the field strengths encountered in the clinic, $T_1 \propto B_0^a$ with $a \approx 1/3$. For example, $T_1^{WM} \approx 1\text{ sec}$ at 3T, so

$$T_1^{WM} = \left(\frac{B_0}{3T}\right)^{1/3}\cdot\text{sec}\,.$$

For a spin echo experiment (in which T_2^* is swapped by T_2, which has little B_0 dependence), the SNR behaves as

$$SNR_{voxel}^{(\text{3D SE})} \propto B_0\frac{1-e^{-TR/(B_0/3T)^{1/3}}}{1-\cos\left(\alpha\right)e^{-TR/(B_0/3T)^{1/3}}}$$

This dependence is plotted as a function of B_0 (in arbitrary units) for TR=1 sec, α=45°, and compared to the simple linear B_0 dependence in the plot below:

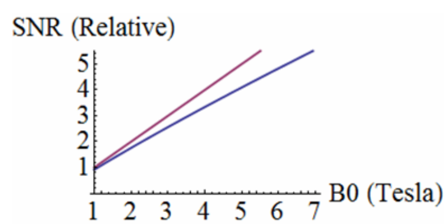

Dependence of SNR On Receiver Coil

An interesting consequence was that both the noise and the signal depend linearly on $B_{xy}^{(rec)}$ in our model, and so cancel out. The final expression of the SNR has no dependence on $B_{xy}^{(rec)}$ and on the sensitivity of the RF coil.

In reality the SNR does show dependence on coil sensitivity. We made some assumptions in our derivation, such as neglecting the coil resistance, which are not 100% correct. In reality, it's quite easy to build a very poor coil which would dominate the noise term.

The quality of a coil is usually characterized by its quality factor, Q. In general, the coil is an inductor with some inductance, L, and it stores magnetic energy. Due to dissipation, some of that energy is lost for every cycle of the RF irradiation. The Q is defined as:

$$Q = \frac{\text{energy stored in coil}}{\text{energy lost per RF cycle}}.$$

The amount of energy lost will depend on what's inside the coil, so an unloaded (empty) coil will have a different and higher Q than a loaded coil with a patient in it, since energy is lost in the patient as well. It is possible to show that

$$\begin{aligned} SNR &\propto \sqrt{Q_{loaded}} \\ Q_{loaded} &= \frac{\omega_0 L}{R_c + R_p} \\ Q_{unloaded} &= \frac{\omega_0 L}{R_c} \end{aligned}$$

If we adhere to the requirement that R_c be minimal, we see that a "good" coil is one for which Q drops dramatically once loaded.

For a good coil, typical values of Q_{loaded} are roughly in the 50-200 range. Unloaded values of Q are in the hundreds.

Dependence of SNR on Gyromagnetic Ratio

Some care must be exercised when expressing the dependence of the SNR on γ (or, equivalently, on $\bar{\gamma}$), since it depends on how precisely we compare two acquisition. If we keep the total acquisition time constant (or, equivalently, the bandwidth per pixel constant), then

$$SNR \propto \bar{\gamma}^2 .$$

It's very difficult to test this prediction in practice, for different reasons:

1. Different nuclei will use different coils.
2. Different nuclei will precess at different resonant frequencies, $\omega_0 = \gamma B_0$. Many physical properties of the tissue being imaged, relaxation constants, conductivities, etc. are all frequency dependent and will affect both the noise and signal.

However, this is one of the main drawbacks of imaging low-γ nuclei, even when at 100% natural abundance (such as ^{31}P), alongside their small concentrations and (sometimes) quadrupolar moments.