HST.583 - Lab 1: fMRI Data Acquisition

In this analysis lab you will examine time-domain signals from fMRI datasets acquired using different spatial resolutions, imaging rates, and RF receive coils. You will use a MATLAB®-based program called *Dview* to examine these signals.

> **Note:** Although not essential, a little experience with and understanding of the basics of MATLAB® will help in these exercises. Sources of MATLAB® help and documentation include Manipulating Matrices, which is part of the MATLAB® 6.x documentation. If you don't understand a command in the MATLAB® code fragments you will see below, you can use the `help` command to get information on it.

The main objectives of the data acquisition and analysis labs are:

- Familiarization of students with a typical functional MRI scanning environment, from data acquisition to offline visualization and analysis.
- Acquisition and examination of image data from a *phantom* (inert test sample) to investigate image intensity non-uniformity, spatial and temporal noise from instrumental sources, and RF receive coil properties.
- Acquisition and examination of *human* data to gain familiarity with 3D anatomic visualization using cross-sectional images and compare physiological and instrumental sources of noise.

Phantom data exercises

First, lets look at data from a phantom (in this case a jug of water doped with a paramagnetic salt). This data was acquired on a Siemens "Sonata" 1.5Tesla whole-body MRI scanner, and should include the following scans:

- A 3D high resolution "anatomic" scan
- An EPI time series (4D dataset) at high spatial resolution
- EPI time series at *high* spatial resolution (2.5mm x 2.5mm x 2.5mm)
- EPI time series at *intermediate* spatial resolution (4mm x 4mm x 4mm)
- EPI time series at *low* spatial resolution (5mm x 5mm x 5mm)

> **Note:** Dview reads MRI data in the Medical Image NetCDF (.mnc) file format. In this format, all dimensions of a scan are contained in a single file. Normally, every file corresponds to a complete dataset (which may be 2D, 3D, 4D, etc.) acquired after one press of the "start" button on the scanner.

Navigation views: Transverse, Sagittal, and Coronal

The three smaller images on the left represent 3 orthogonal cross sections, intersecting at the location of the yellow cursor, of the 3D MR image that you have loaded. You can move to anywhere in the 3D volume by clicking at the desired point in one of the images with the left mouse button. From top to bottom, the views shown are called the *Transverse*, *Sagittal*, and *Coronal* views. With the cylindrical phantom these designations are not too meaningful, but for images of real anatomy they correspond roughly to top, right side, and back views of the head of a subject lying in the scanner. You may be able to recognize the Sagittal (side) view from the air bubble visible at the top of the phantom (on the right side of the image).

Note that the choice of top and back, instead of bottom and front, results in the subject left being on the left side of the image. This is known as the *nuclear medicine* display convention. The opposite scheme (bottom and front views, so that image left = subject right) is known as the *radiological convention*. The current convention is displayed in a status window at the bottom, and can be changed from the **Settings** menu.

Also note the *coordinate system* in which the images are displayed. The default in the Dview program is to use *scanner-based* directions in which X is left-to-right, Y is back-to-front, and Z is foot-to-head (assuming of course that the subject is lying face up in the magnet with their feet out the front):

Head coil and scanner coordinate system (for *Nuclear Medicine* convention)

> **Note:** By default, the Dview program displays images using the **Nuclear Medicine** convention, in which positive *Z* is towards the top of the patient's head. This allows more natural display and plotting from an engineering perspective, but is not used in many medical imaging devices - the *Radiological* convention provides views which correspond to what you would see looking at a patient face-to-face or from the bedside (while maintaining a right-hand coordinate system).

In thinking about data acquisition, it is often more useful to operate in terms of frequency encoding, phase encoding, and slice selection directions. You can change to a *row, column, slice* coordinate system by selecting **Settings->Units->Row,Col,Slice,Frame**. The view corresponding to the *acquisition slice plane* is always indicated by blue axis labels in the Navigation view. Although the MRI datasets are useful as 3D volumes, the concepts of *slice* and *inplane* dimensions are important in determining image characteristics, which brings us to:

Lab question 1: The high resolution anatomic dataset is a "3D" Sagittal acquisition. Describe the differences between the frequency-encoded, phase-encoded, and slice selection directions (row, column, and slice dimensions in this scan). How is this different from later scans? What will happen if the imaged object extends outside the scan boundaries in the different directions?

The three images in the Navigation view are kind of small, but you can move any of the views into the big window by right-clicking on the view you want and selecting **Copy this view to big window**. You can also *zoom* the small Navigation views by dragging the mouse up or down in that view while holding the middle button down. You can switch the action of the middle mouse button from *Zoom* to *Xlate* (translate) and back by pressing the **x** or **z** buttons on the keyboard.

Image intensity non-uniformity

Although the high-resolution T1-weighted scan you have loaded should be of generally good quality, providing clear cross-sectional images of the jug of water, this dataset illustrates the phenomenon of *image intensity non-uniformity*, which is an important concern for intensity-based segmentation of tissue types and cortical surface extraction (topics which will be covered in future lectures). The phantom used for this exercise contains a homogeneous solution of paramagnetic salt in water, so ideally the image intensity values should be the same everywhere. The RF coil (a *head-coil* in this case) used to receive the NMR signal emitted by tissues following excitation does not have uniform sensitivity over space, however, leading to low spatial-frequency variations in the observed image intensity. Typically larger coils offer better uniformity, but at the expense of signal-to-noise ratio

Bonus question: Why is a paramagnetic salt added to the distilled water in the phantom?

With the grayscale colormap, the images will probably appear to be reasonably uniform. Try switching to the "spectral" colormap by selecting **Image->Colormap->spectral** from the top-of-screen menu. The human eye is much better at detecting subtle gradients in color than in luminance, so non-uniformity should be much more readily visible with this mapping (which distributes all visible colors of the spectrum over the image range).

When you are finished, change the colormap back to greyscale in preparation for later exercises.

You can get a better idea of the degree of image intensity non-uniformity by looking at a spatial intensity profile. You can get a profile along the *z* (head-to-foot) axis by right-clicking in the Sagittal (middle) window and selecting **Intensity profiles->Z profile**. You can get profiles along other axes by clicking in any image and selecting the desired profile direction (the path of the profile is indicated by a red line on the images). Clicking on the profile plot axes will move the image navigation cursor to that point along the profile axis, allowing you to see structures in the image that correspond to features on the plot.

Lab question 2: Write a paragraph or two explaining the origin of the observed intensity non-uniformity. What is the percent non-uniformity over a range roughly equal to the size of a human brain? What aspect of data acquisition could be changed to improve the uniformity?

Lab question 3: Place the cursor near the middle of the phantom on the Transverse view and select an intensity profile in the Y direction (right-click). Describe and explain the characteristics of the profile observed at the *edges* of the phantom. Are the edge characteristics the same in the X direction? Why might they be different? (hint: look at the voxel dimensions by selecting **Information->Geometry**)

Spatial noise

A fundamental figure of merit in MRI data assessment is the *signal-to-noise ratio* (SNR). This is especially true in functional MRI, since the effects we wish to detect are a very small fraction of the signal.

The amplitude of thermal noise fluctuations in MRI signals depends primarily on the bandwidth to which the RF receiver is set, electronic characteristics of the coil and preamp circuits, and physical properties of the subject. The amplitude of the *signal* observed in an MR image volume element (or *voxel*) depends on the volume of the voxel, its proton density, the degree of longitudinal and transverse relaxation during excitation and subsequent digitization of the NMR signal, and the coil sensitivity to signals from the voxel's location. Both the noise and signal amplitudes are scaled by a receiver gain factor used to ensure that the maximum received signal does not exceed the dynamic range of the analog-to-digital converter.

Complex numbers and Gaussian, Rayleigh, and Rician noise distributions: If you're not familiar with complex numbers, don't worry - all you need to know about them is that they're a way to describe something that has a length (or *magnitude*) and is pointing somewhere (described by the *phase angle*). They're used in MRI to describe the *transverse magnetization* in voxels and the RF signal induced in the *coil* by precession of the latter in the xy plane (see picture of scanner

coordinate system and coil above). The *real* and *imaginary* components of a complex number refer to Cartesian coordinates which may also be used instead of the polar magnitude and phase.

You are probably familiar with the *Gaussian distribution* from basic probability and statistics. Instrumental noise in electronic systems is very often Gaussian with zero mean, and the real and imaginary channels of the MRI scanner RF receive chain are no exception. Each of these channels is digitized during readout of the post-excitation tissue NMR signal, and the samples merged to form a single complex-valued waveform. As you know, this waveform constitutes one *row* of the 2D (or 3D) Fourier transform *matrix* of the object we want to image. After we've filled the whole matrix, we perform an inverse Fourier transform on it to get the image. Note that the result after inverse Fourier transformation is a *complex valued* image. Although there can be useful information in the phase, we typically discard this part of the complex image data and retain only the *magnitude* image for subsequent processing (all the images you'll look at today are magnitude images).

What does this have to do with noise distributions? Lots. The implications are most obvious if we consider a voxel where the signal *should* be zero (such as outside the phantom). Both the real and imaginary components of the complex image constructed from real data will be the sum of the "true" value of zero, and zero-mean Gaussian noise. Obviously this results in a mean value of zero for both the real and imaginary components, in spite of the additive noise. However when we compute the *magnitude*, which is equal to sqrt(Re^2 + Im^2), we end up with a signal which can not pass through zero to be negative. The resultant noise distribution is called the *Rayleigh distribution* and has a non-zero mean given by **mean = σ x sqrt(π/2)** where **σ** is the standard deviation on the input real and imaginary channels. The apparent *variance* of the Rayleigh-distributed magnitude values, as a function of the standard deviation **σ** of the real and imaginary inputs, is given by **var = (2-π/2)σ^2**.

If we then look at a magnitude signal in a voxel where the "true" value is significantly different from zero (such as within the phantom), the situation is somewhat different. In this case the inability of the magnitude signal to become negative is less relevant, since the noise is riding on a large non-zero "true" component. Although the resultant *Rician* noise distribution is still not exactly Gaussian, it is pretty close. We can therefore treat the apparent mean and standard deviation of a signal *within the phantom* as accurate descriptions of the quasi-Gaussian noise.

But why should you care? Because the mean and standard deviation in background voxels contain useful information about the magnitude of strictly *non-physiological* noise sources, but these numbers need to be

corrected to estimate the instrumental noise contribution in the brain. For example if the apparent standard deviation over time in a background voxel is **σ**, you can estimate the instrumental noise component in the brain as **σ/sqrt(2-π/2)**. If the intra-cerebral standard deviation is higher, the difference can be attributed to physiological sources.

If you need convincing, play around with the following MATLAB® code to generate frequency distribution histograms for real, imaginary, and magnitude signals (remember the `help` command):

```
Re = randn(100000,1);
plot(Re)
[Rn,Rval]=hist(Re,64);
Im = randn(100000,1);
plot(Im,'r')
[In,Ival]=hist(Im,64);
Mag = sqrt(Re.^2+Im.^2);
plot(Mag,'k')
[Mn,Mval]=hist(Mag,64);
plot(Rval,Rn)
hold on
plot(Ival,In,'r')
plot(Mval,Mn,'k','LineWidth',3.0)
xlabel('value');ylabel('frequency')
hold off
legend('real','imaginary','magnitude')
```

In addition to the low spatial-frequency intensity changes in the image profiles you've been examining, you probably noticed higher frequency *noise fluctuations* on the plots of signal as a function of position. In functional MRI, most attention is focused on *temporal* noise (as opposed to spatial noise). Spatial noise is nonetheless important, and can have a significant impact on image segmentation procedures (Lab #4). Also, the uniformity of the phantom used in this part of the experiment allows the spatial noise fluctuations to be seen clearly (the thermal processes contributing to spatial noise are the same as those involved in temporal noise generation, although there are other sources of temporal instability).

Call up an intensity profile in the Z direction and identify the noise component. With the cursor in the phantom, most of the intensity variation you see is due to the low spatial-frequency intensity non-uniformity discussed above. Move the cursor into the background, outside of the image. This will allow you to isolate the thermal noise component in the images, but there are some subtle differences in the statistical properties of noise in background regions.

The following exercise will give you a chance to "play" with image signals in MATLAB® and think about spatial noise characteristics:

Note: In the following exercise you will have to execute a series of MATLAB® commands required to isolate thermal noise in an intensity profile of the phantom. You *could* type these commands at the ">>" prompt in the MATLAB® command window, but if you are viewing this tutorial in the MATLAB® 6 Help Browser, you can execute them simply by highlighting them with the mouse, right-clicking and choosing "Evaluate Selection" from the menu that appears (instead of "cut-and-paste", think "select-and-evaluate"). To avoid typographical errors and save time I recommend using the latter approach. If some of the details of the commands below seem a little mysterious, don't worry (but don't hesitate to ask about them). The commands have been included to provide examples of how to extract and operate on MRI signals in MATLAB® in case anyone feels like experimenting on their own.

1. Obtain an intensity profile along the Z axis, from close to the center of the phantom, by right-clicking on the Sagittal image in Dview and selecting **Intensity profiles->Z profile**. Try to get the longest possible view of the phantom.
2. Save the intensity profile into a MATLAB® workspace variable by right-clicking on the signal and selecting **Save signal->to global variable**. Call the signal *Zprofile*.
3. To make the global variable visible in the MATLAB® workspace, you will have to type (or select-and-evaluate)

   ```
   global Zprofile
   ```

   in the command window (or select the text and evaluate it through the right-click menu).

4. The Zprofile variable is a *structure* with the fields *Signal* and *Xdata*. To get just the Signal part (the .Xdata field contains the positions of the samples) and chop off the ends of the image outside of the phantom, run the following command:

   ```
   Sig = Zprofile.Signal(30:end-30);
   ```

   (you may have to chop off more or less than 30 points, depending on how the phantom was positioned).

5. To get rid of the low spatial-frequency intensity non-uniformity, we will subtract a smoothed copy of the signal from itself. First make the smoothed copy by convolving `Sig` with a 32 point rectangular kernel:

   ```
   SigSmooth = conv2(Sig,ones(1,32)./32,'valid');
   ```

6. You can compare the smoothed and original copies using the **`plot`** command:

```
plot(Sig(16:end-16))
hold on
plot(SigSmooth,'r')
hold off
```

(the 16 points at each end of the signal are not valid after convolution with the 32 point kernel and are therefore not present in `SigSmooth`)

Plots should appear in the big Dview window - they will remain as long as you don't manipulate any Dview controls or click in the images.

7. If this looks good, generate the non-uniformity-corrected version of the signal:

```
SigFlat = Sig(16:end-16) - SigSmooth;
plot(SigFlat);
```

8. Now compute the standard deviation of the noisy residual:

```
std(SigFlat)
```

and look at the result, which is printed in the MATLAB® command window (the semicolons at the end of previous commands suppressed printing of the output).

9. Inspect some Z profiles in *background* regions, and compare the standard deviation (listed above the plot window in Dview) in these areas with the value in the corrected image profile.

Lab question 4: Perform the above exercise for several image profiles, and compare the standard deviations observed in corrected profiles with those seen in background profiles (you don't need to correct for non-uniformity in background). Explain the differences and compare with the expected ratio based on statistical considerations. Is the noise amplitude the same over the entire spatial range of the profile? (try dividing the profile into overlapping 64-point segments and comparing the standard deviations of each - stay away from the edge of the phantom, though).

Temporal noise

Now we'll look at temporal noise, which is of more immediate interest in functional MRI. Open the first EPI scanning run in the list of files for this session (this should be the high spatial resolution EPI scan).

> **Note:** When opening .mnc files, Dview shows the scanning *protocol name* by default. You can change modes to display the *pulse sequence, coil, resolution,* or *none* through the **Settings->'Open file' annotation** menu.

This time, when the file loads, you should see the same column of three navigation views on the left, and a plot of intensity as a function of time in the big window on the right:

> **Note:** Clicking along the time axis in the signal plot will cause the 3D volume (or *frame*) at that time value to be shown in the display windows.

The navigation views show that, unlike the previous 3D anatomic scan, the spatial sampling volume would cover only a relatively thin slab of the head. Normally we try to acquire more, but with the high spatial resolution (*i.e.* thin slices) of this scan, we restricted coverage to avoid having a large number of slices in each of the 128 time points acquired (strictly for convenience in this exercise).

To see the sub-volume covered by this EPI scan, use **Superposition->push file onto superposition stack** *twice* to push first the anatomic scan and then the EPI scan onto the *superposition stack*. This is basically a pile of images which can be displayed semi-transparently so that different types of MRI scan can be merged into a single image. Next, to see the superposition stack, click on **Select->Superposition** and the display should update to show the superimposed volumes. To make the small volume easier to see, you can first **Select** it, then under the **Image** menu change its colormap to **hot**.

> **Note:** Even if a superposition stack has been defined, you can still select the other volumes and view them as usual. Redrawing is slow in superposition mode, so you will generally want to work on the actual volumes except when you actually need to look at the superposition. When you switch between volumes, Dview will try to keep the cursor at the same place in *scanner coordinates* unless you switch this off under **Settings**.

Close the 3D anatomic scan (it takes a lot of memory and we're done with it) by making it current (**Select->*yourfile.mnc***) and then selecting **File->Close current**

file. This should leave you with just the high-resolution EPI scan. Also get rid of the superposition stack by clicking **Superposition->delete superposition stack**, and switch the colormap of the remaining EPI scan back to **gray**.

Spatial vs. temporal noise

When thermal RF noise is dominant, the degree of random fluctuation over space and time should be similar. In the Transverse (top) view, place the cursor in a *background* area well to the bottom of the phantom (we used left-right phase encoding and we don't want to hit any *image ghosts* - try switching to the **spectral** colormap to see and avoid background image artifacts). The plot display will update to show the signal as a function of time at this location, and the mean and standard deviation (and their ratio - the SNR) will be reported above the graph. Note the standard deviation (sigma) value, and then switch the plot display to 'X profile' using the selection box at the lower right corner of the Dview window (don't click on the images, as this will change the position of the cursor). Compare the standard deviation over the background spatial profile with that of the temporal signal (remember the statistical considerations regarding noise distributions in background areas from above). Try this for several background locations (the whole profile must be in background) - you can switch back and forth between profiles and "T signal (raw)" using the lower right selection box.

Spatial resolution and signal-to-noise ratio

In this exercise, we will look at the impact of spatial resolution on SNR. Make a table listing the mean signal value, standard deviation, and SNR values in at least three locations (with approximately similar intensities) from within the phantom from the currently loaded EPI series. By clicking the **Information->Geometry** menu selection, make a note of the voxel dimensions used in the current scan.

Now open the next EPI series, which was acquired at a somewhat lower spatial resolution of $4x4x4mm^3$ (you can choose to annotate file listings with resolution under the **Settings** menu). Also note the mean signal value, standard deviation, and SNR values from several spots in the phantom.

Do the same for the next EPI volume, sampled at $5x5x5mm^3$.

> **Lab question 6:** For each spatial resolution, compute the average SNR value over the different locations you sampled. Plot average SNR as a function of voxel volume and comment on your findings. Describe the signal and noise behavior you would expect as voxel size increases indefinitely (hint - don't forget about relaxation).

Bonus question: In the lower resolution images, there is a lot of empty space around the object. This could be eliminated and a smaller matrix size used (along one direction) by acquiring fewer phase-encoding steps. Would this affect the signal-to-noise ratio?

Ghosting in EPI scans

Another important aspect of image quality in the EPI scans generally used for functional imaging is the degree of *ghosting*. Ghosting arises due to slight phase differences between the even and odd phase encoding lines acquired as the EPI readout gradients take us back and forth through $k_x = 0$ in the frequency domain. This is like taking a perfect image, and adding a complex valued copy with half of the sampling density in the k_y direction and hence half the unaliased field of view. When the magnitude of this complex-valued summation is taken, ghostly aliased copies of the object will be seen wrapping in from the sides (or top and bottom) of the image. The larger the phase error, the stronger these ghost components will be. To correct ghosting, it is necessary to determine the line-to-line phase error by digitizing a single FID in which the same phase-encode line is repeatedly measured (instead of stepping through all the lines required for an image). This information, which is usually acquired at the beginning of each EPI scan, can then be used to correct the k-space data prior to inverse Fourier transformation.

To illustrate how ghosting occurs, we will start with a high quality anatomic image and simulate the effects of EPI acquisition with a phase offset between the odd and even lines:

1. Load the high resolution 3D anatomic dataset (series #2, or `phantom-0-sonata-21006-20010819-134254-2-mri.mnc` in the "File open" list). Save the *Transverse* image by right-clicking on the image and selecting **Save image->as global variable**. You can use the default variable name of *MyImage*.
2. To make the global variable visible in the MATLAB® workspace, you will have to type:

   ```
   global MyImage
   ```

3. The MyImage variable is a *structure* with the fields *Image*, *Xdata*, and *Ydata*. To get just the Image part, run the following command:

   ```
   Image = MyImage.Image;
   ```

4. Generate a phase modulation function that we will apply along the Y direction of the image:

```
PhaseErr = 0.2; % phase offset in radians
PhaseMod = ones(size(Image));
PhaseMod(2:2:end) = PhaseMod(2:2:end).*exp(2.*pi.*PhaseErr);
```

5. Now take the 2D FFT of the image, apply the phase factor, then transform back:

```
ImageFFT = fftshift(fft2(fftshift(Image)));
ImageFFT2 = ImageFFT.*PhaseMod;
Image2 = fftshift(ifft2(fftshift(ImageFFT2)));
imagesc(MyImage.Xdata,MyImage.Ydata,abs(Image2))
axis('image')
```

   and look at the image, which should appear in place of the original in the smaller *Transvserse* viewport of the Dview window.

Lab question 7: Try the above exercise for several values of PhaseErr. Explain briefly how the EPI gradient waveforms have to be modified for estimation of the line-to-line phase error.

This brings us to the end of the phantom data exercises. At this point you should have a feel for

- how to load and view MRI data in the *Dview* software
- how spatial and temporal noise are related within a uniform phantom and in background regions of the image
- how the signal to temporal noise ratio varies with spatial resolution
- what causes ghosting

Human noise data exercises

In the previous exercises, you examined signals from a jug of water. This represents a "best case scenario" for signal-to-noise characteristics for a number of physical reasons:

- the phantom is very close to 100% water by volume, with small amounts of paramagnetic salt in solution. This results in the highest possible number of protons in a voxel of a given size
- the paramagnetic salt shortens T1 enough that the longitudinal magnetization is almost fully relaxed for every EPI volume (multi-slice scan) in a series
- although the paramagnetic solute also shortens T2*, the value in the doped water is still long enough that the signals are very strong.

- because the medium inside the phantom (water) is homogeneous, acceleration of transverse relaxation due to magnetic field inhomogeneity within voxels is minimized (except near the edges).

We're really interested in imaging human brain tissue, though. It is important to recognize that *in vivo* signals may not be quite as strong as those in a phantom, and more importantly are subject to degradation from a number of sources:

- **Motion:**
  Human subjects are generally imaged as they lie comfortably in the scanner after being coached to remain as still as possible. Their efforts to remain still may be assisted using straps, padding, and/or bitebars etc. While it is usually possible to keep motion to well under a centimeter, small shifts of one or two millimeters and rotations of a few degrees are difficult to avoid completely. The result can be large signal changes around high contrast edges in and around the brain.
- **Pulsatile blood flow:**
  In a human subject the heart will be beating at approximately 1Hz, although this can vary considerably with anxiety (subjects may be nervous in the strange and confining scanner environment and/or anxious about correctly following instructions). Even if heart rate is constant, the pulsatile nature of blood flow (and related displacement of tissue) will lead to slight variations in MR signal over the cardiac cycle and contribute a 1Hz (approx.) noise component which will appear *aliased* into fMRI signals when (as is usually the case) lower imaging rates are used (in Lab #3 you will investigate strategies for dealing with this problem).
- **Respiration:**
  Breathing, which may occur at 3-5 breaths per minute (0.05-0.08Hz), also causes small intensity changes whose amplitude can rival that of the activation-related changes we wish to detect. Part of this is due to motion of the head associated with muscular contraction and relaxation required for inhalation and exhalation, but even motion of the chest wall can have an effect by introducing temporal variations in the magnetic field within the volume being imaged. Equally important are changes in the amount of carbon dioxide in the blood that change with respiration rate. CO_2 affects blood flow in the brain, which in turn affects the fMRI signal. Even the oxygen concentration of air in the mouth and sinus cavities can affect the signal slightly, as O_2 is paramagnetic and therefore can affect nearby MRI signals (these topics will be covered in more depth in the Week 5 lecture on physiology and in Lab #2).
- **Uncontrolled changes in cognitive state**
  Of course it's impossible to force someone to maintain a completely constant mental state, so even without any of the other "noise" sources it is likely that there would be signal fluctuations from this source superimposed on (or modulating, or interacting somehow with) responses

to controlled changes in experimental condition (i.e. the stuff we *want* to see).

The above sources of signal fluctuation can be divided into two types:

1. phenomena such as thermal noise, motion, and remote effects from chest wall movement, which cause image intensity values to be "wrong" as estimates of the magnetization at a particular location in the brain
2. others such as CO_2-related changes in brain blood flow and unsolicited changes in cognitive state which constitute accurate measurement of phenomena which are simply uncontrolled experimental variables.

Anyhow, now that you are experts in the noise characteristics of phantom data, it's time to look at some human scans. The scans in this part of the exercise are human datasets, but there were no deliberate changes in experimental condition during data acquisition. The resultant signals are therefore useful as examples of *noise*.

To get to the human noise datasets, cd into the directory `human-0-sonata-21006-20010819-144929/`. You should have the following six files of human data acquired on the same 1.5Tesla scanner as the phantom data:

- A 3D high resolution "anatomic" scan
- An EPI time series (4D dataset) at high spatial resolution
- EPI time series at *high* spatial resolution (2.5mm x 2.5mm x 2.5mm)
- EPI time series at *intermediate* spatial resolution (4mm x 4mm x 4mm)
- EPI time series at *low* spatial resolution (5mm x 5mm x 5mm)
- EPI time series at 4x4x4mm^3 acquired at 2 volumes/s (TR=0.5s) instead of 0.5 volumes/s (TR=2s)

Dview *File->open* menu with human noise data

Human anatomic scan

Load the high resolution 3D anatomic dataset (series #2, or `human-0-sonata-21006-20010819-144929-2-mri.mnc` in the "File open" list). This scan is a *T1-weighted* scan, which is why the grey matter tissue appears darker than the white matter (its T1 is longer). CSF has an even longer T1 value, so it appears essentially black in these images. Using the intensity profile tools, you can see evidence of the coil-related non-uniformity in large areas of white matter.

Physiological vs. thermal noise

In this exercise you should get a feel for the range of SNR values encountered in real human data. You will also look for evidence of physiological noise, which occurs in addition to thermal noise, in human fMRI data.

Lab question 8: As in question 6, make a plot of SNR as a function of voxel volume in human grey matter. Do the same for white matter, and plot the results on the same axes. Compare these plots to the SNR vs. voxel volume plots for phantom data.

Lab question 9: Using the human EPI data, determine the standard deviation over time at three background locations at each spatial resolution. Average the values taken at each voxel size, then compute the corrected estimates of the intra-cerebral thermal noise component (see discussion of Rayleigh vs. Gaussian noise above) as a function of resolution. Repeat this, but get the standard deviations from voxels in grey matter (no correction is required for these). Now plot the *total* noise observed in grey matter as a function of the *estimated* thermal noise component based on the background observations. Do exactly the same thing with phantom data, and plot the relationships (total vs. estimated thermal noise for human and phantom data) on the same axes. Comment on your findings.

Imaging rate and statistical power

Increasing the number of independent observations in an experiment generally results in greater statistical power. In many (but not all) situations, statistical uncertainty will *decrease* with the square root of the number of observations. Based on this rule of thumb, one might suppose that the highest possible imaging rate should be used in an fMRI experiment, so as to acquire the largest possible amount of data and thereby maximize statistical power.

The above supposition would be true if the image properties were independent of the acquisition rate (or TR), but in functional MRI this is not true. At short TR values the image intensity is *attenuated*, due to incomplete recovery of longitudinal magnetization. Thus as the number of observations per unit time increases, the signal-to-noise ratio per unit observation actually decreases. In this exercise you will compare EPI scanning runs that differ only in their imaging rate, to see how these effects compete.

Close all other files and open only the 4x4x4mm^3 EPI scans with 0.5 and 2.0 second TR values (series #4 and #6, with protocol *epi_hires* in the human dataset). Look at signals from different locations in the brain. You will notice that at many locations, the time-domain signal in the short-TR scan exhibits an

exponential-like decay from an initial value down to some lower steady-state value. This is the supression of longitudinal magnetization mentioned above. To keep this signal decay from skewing temporal standard deviation estimates, you should *exclude* the first eight points of the signal. This can be done by *selecting* the high imaging rate volume, then clicking **Tools->Enter frames to exclude->as MATLAB® expression**. When the text-entry window comes up, change it to read `Exclude=[1:8];` (this will highlight the first eight points in the signal in green and exclude them from computations). You could also exclude the first *two* scans in the TR=2s scan, to make the scanning times exactly equal in both cases (the effect of this will be small).

Lab question 10: After exclusion of the first 8 frames, determine the average grey-matter SNR in the short-TR scan. Compare this value with the average grey-matter SNR in the longer-TR scan. If statistical power increases as the square root of the number of observations and linearly with the SNR (as computed in Dview), which dataset would give better results in an fMRI experiment?

Note: an important detail in the superficial discussion of statistical power given above is that it's the number of *independent* observations that determine detection power. Like many types of biomedical data, fMRI signals exhibit *temporal autocorrelation*. This is a fancy way of saying that just because you're imaging really fast does not mean all the observations are fully independent. Although thermal noise fluctuations are pretty independent, a *physiological* noise fluctuation associated with chest wall position at peak inspiratory volume is likely to be followed by other fluctuations associated with exhalation. Estimating the effective number of independent observations (related closely to the number of degrees of freedom), is a statistical problem that will be discussed in future lectures.

Now we are done with the human noise segment, and are ready to move on to looking at human activation data!

Human activation data exercises

By now you hopefully have a pretty good idea of the noise characteristics in human fMRI data. Of course we're not really interested in noise, but rather in our ability to detect activation-related effects that occur against a background of noise.

You have gone through a number of exercises aimed at determine the signal-to-noise ratio of EPI time-series data, where it was defined as the ratio of the average raw MRI signal in a voxel to the standard deviation of that signal over repeated observations. What is really relevant in functional MRI is actually the ratio of the *activation-related signal change* to the temporal noise component in

the signal. This is often referred to as the *contrast-to-noise ratio*, or CNR (although the terms SNR and CNR are often used interchangeably).

The data used in the previous exercises was all acquired on a 1.5Tesla MRI scanner. In this final exercise we will use data acquired on a 3Tesla system, which will allow us to make a limited comparison of the raw SNR observed in brain tissue at the two field strengths. We will do this by comparing unstimulated runs at 4x4x4mm^3 at the two field strengths. Note that this is not the whole story, as part of the benefit from high-field systems derives from enhancement of susceptibility-related signal changes (i.e. increased functional contrast).

To get to the human activation datasets, cd into the directory `human-0-allegra-20006-20010910-173032/`. There should be four files, acquired on a Siemens "allegra" 3Tesla head-only scanner:

- EPI time series at 4x4x4mm3, with no stimulus presentation (noise-only)
- EPI time series at 4x4x4mm3 acquired during periodic presentation of a *low-contrast* visual stimulus
- EPI time series at 4x4x4mm3 acquired during periodic presentation of a *high-contrast* visual stimulus
- EPI time series at 4x4x4mm3 acquired during presentation of a *phase-encoded retinotopic stimulus*.

Dview *File->open* menu with human activation data

Magnetic field strength and SNR

Open the first file in the list (`human-0-allegra-20006-20010910-173032-4-mri.mnc`), which should be the noise-only 4x4x4mm3 EPI series.

> **Lab question 11:** Determine the temporal signal-to-noise ratio in several grey matter voxels, and compute the average grey-matter SNR. Compare this to the average value computed in grey matter at 1.5T for Question #8.

Stimulation-induced signal changes

Now it's time to actually do something to the brain. The idea here is basically to see how much the MRI signal changes when we apply intense stimulation to neuronal tissue. To do this, we will take advantage of the well understood pathway connecting points on the eye's *retina* to a map of the visual field on the brain's occipital cortex.

The first step is to identify cortical areas with direct retinal projections. This can be done by showing the subject a high contrast visual stimulus with a semicircular region comprising 50% of the stimulus area masked out to zero contrast. The polar angle of the masked area is gradually swept through 360 degrees, resulting in periodic modulation of visual input to those cortical locations that have specific retinotopic projections. By acquiring EPI image series during this procedure and then generating an image of the spectral power at the rotation frequency, retinotopic cortex is revealed as a region of enhanced intensity in the spectral image. This type of procedure is known as phase-encoded retinotopic mapping.

Now if we expose the subject to a visual pattern covering the range of visual field used in the previous step, we know that we must be stimulating neurons within the identified retinotopic regions. This allows us to investigate responses to test stimuli in this exercise without having to be able to localize the activated region based solely on the applied stimulus (which in at least one case is a weak stimulus).

So for the next part of the exercise, make sure the noise-only file you had open for the last question is closed, and open the file **`human-0-allegra-20006-20010910-173032-10-mri.mnc`**. This file contains an image series acquired during phase-encoded retinotopic mapping of visual polar angle. Click on the **Tools** menu then select **F statistic for periodic design**. Click **ok** on the window that appears to start generation of the spectral image (this will take about a minute).

At the end of this procedure, you should see a spectral image in which cortex with direct retinal projections exhibits enhanced intensity ("hotter" colors in the default spectral colormap). Due to the intense stimulation and general robustness of this procedure, we will not worry about statistical issues for now.

The file selection list will show a total of three files now - the original EPI series, and two *computed* files with the suffixes "-fftp.mnc" and "-fftm.mnc". The latter is the spectral magnitude image, and will show up by default (fftp is the spectral *phase*). You will use the spectral magnitude image to generate a region of interest (ROI) to guide sampling of tissue in the high and low-contrast visual stimulation runs.

With the spectral magnitude image selected, click on **Tools->Create ROI->by thresholding current volume**. Enter a lower limit of 0.2, and click `OK'. You will see the resultant 3D region of interest labeled in pink on the spectral image.

Now open the file **`human-0-allegra-20006-20010910-173032-8-mri.mnc`**, which contains the EPI series acquired during 3 intervals of high-contrast checkerboard stimulation. Then click **Tools->Apply ROI to current file** and select the ROI you just created. In a second you should see the ROI superimposed on the EPI scan,

and the signal shown in the large window will be the *average* signal within the ROI.

To determine the magnitude of the signal change during stimulation, you need to specify a baseline period and indicate when the stimulus was applied. This can be done by clicking **Tools->Enter design matrix/vector->as MATLAB® expression**. The default values should be appropriate, so you can just click OK at the small window that opens. You should see the *design matrix*, which indicates experimental condition as a function of time, plotted in red over the signal. The title over the plot should now report the average signal change from baseline during stimulation. You may wish to exclude the first few images from computation, due to enhanced longitudinal magnetization in the first couple of scans (**Tools->Enter frames to exclude**.

Lab question 12: Compare the average stimulus-related signal change in visual cortex for high and low-contrast radial checkerboard stimulation. How do these values compare to the fluctuations observed in the noise-only run (series #4)? What are the percentage signal changes?