



CS/NEUR125 Brains, Minds, and Machines

Lab 8: Using fMRI to Discover Language Areas in the Brain

Due: Wednesday, April 5

In this lab, you will analyze fMRI data from an experiment that was designed to find regions of the brain that are engaged in the processing of language. To perform this analysis, you will learn about the **General Linear Model** that is often used in fMRI studies. These lab activities were adapted from a module developed by Idan Blank for the online MITx course, [Quantitative Biology Workshop](#). Idan also teaches this material in Nancy Kanwisher's course at MIT, *Functional MRI Investigations of the Human Brain*. [Idan's course lectures on Youtube](#) provide an excellent introduction to the General Linear Model and its application to finding language areas of the brain (you do not need to view these videos prior to this lab).

To begin, create a copy of this Google document, modify the title of the copy to include your partner names, and share the copy between partners, as you did in previous labs. The questions for you to answer as you progress through the lab are [shown in blue](#).

[Start MATLAB](#) on your lab Mac and use **Fetch** to download the folder named **fMRI_Lab** from the CS file server to the Desktop on your Mac. You will find this folder inside the cs125/download folder in your individual account on the CS server. For this lab, set the **Current Folder** in MATLAB to the **fMRI_Lab** folder. You will work with two MATLAB programs that allow you to visualize anatomical and functional MRI data, explore the General Linear Model, and apply this analysis technique to the fMRI data. These programs were developed with the help of Isabel D'Alessandro '18.

I. Visualizing Anatomical and Functional MRI Data

In this part of the lab, you will visualize some data collected from an actual fMRI experiment performed by the Kanwisher Lab, aimed at localizing language regions in the brain. In the experiment, a subject's brain was scanned while they performed a reading task. The task was divided into separate trials; on some trials, the subject read a valid English sentence, while on other trials, the subject read a sequence of nonwords that sound like they could be English words, like "blicket" or "florp." Real sentences not only have individual words with meanings, they are also structured according to syntactic rules, in a way that yields meaning for the sentence as a whole.

Recall that fMRI analysis is based on measuring the difference in the blood-oxygen-level dependent (BOLD) signal between two conditions, referred to as a *contrast*. This relies on the fact that a change in cerebral blood flow is related to neuronal activation. An MRI machine can measure the BOLD signal over time, across the entire brain, as different stimuli are presented to the subject. In this lab, you will examine the contrast between the fMRI signal for the two conditions of reading sentences versus reading nonwords.

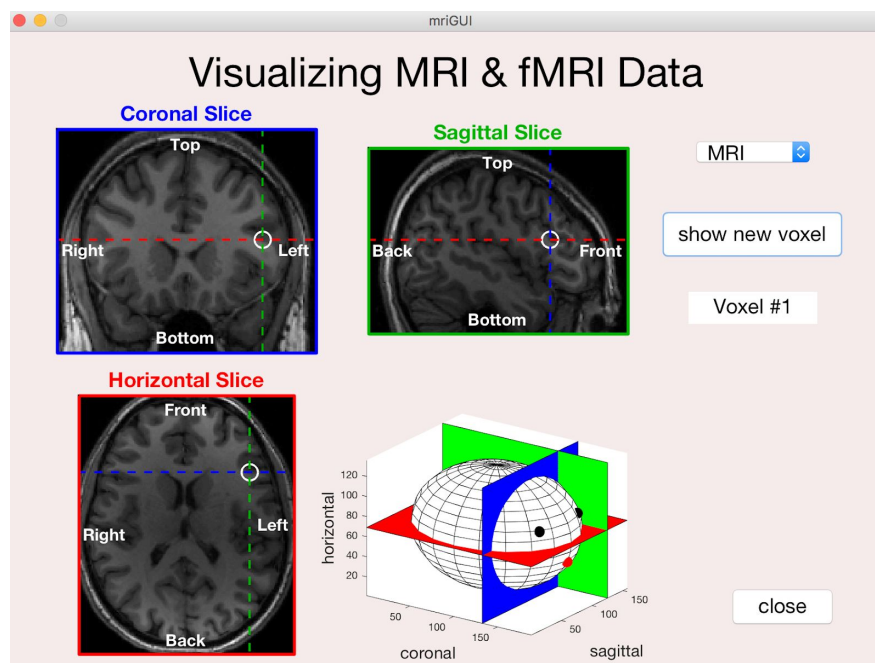
Q1: Why do you think it is important to record fMRI signals when a participant is reading nonwords as well as when they are reading sentences? What brain regions might be activated when reading sentences that are not specifically related to language?

In Assignment 2, you studied a paper about the structural changes observed in the hippocampi of London taxi drivers. That study examined **structural MRI** scans of the brain, also known as **anatomical MRI** scans, which capture the detailed anatomical structure of the brain at a high spatial resolution. **Functional MRI** scans are obtained at a much lower spatial resolution, and are recorded at regular time intervals. In this part of the lab, you will view both anatomical MRI data collected while the subject was resting, and fMRI data collected while the subject was performing the reading task. In the case of fMRI, a “picture” of the entire brain was taken every 2 seconds, with the activity level recorded for each **voxel** (3D pixel) within a 3D volume at every time point. In MATLAB, the volume of anatomical MRI data is stored in a 3D matrix where each entry in the matrix at coordinates (x, y, z) corresponds to a single voxel. The fMRI data is stored in a 4D matrix, where the fourth dimension corresponds to time.

To view the data, enter the following command in the Command Window:


```
>> mriGUI
```

This opens a window that looks like this:



The pull-down menu in the upper right corner allows you to select **MRI** or **fMRI** data. When MRI is selected, anatomical MRI scans are displayed, showing coronal, sagittal, and horizontal slices that intersect at a particular voxel that is circled in white in each slice image. The 3D cartoon in the lower right corner shows the rough position of the slices relative to the head. Note that the eyes and nose are only intended to convey the orientation of the head, and do not accurately depict the location of the subject’s eyes and nose. By dragging your mouse over this graphic,

you can rotate the cartoon in 3D. The numerical values on the three axes of this graphic represent indices of the 3D matrix that stores the MRI data - the head appears elongated in the direction from back to front because there were more slices in the MRI data recorded along this direction (individual voxels in fMRI data are “cubes” of size roughly 1 - 8 mm³). The labels **Left** and **Right** on the slice images refer to the left and right hemispheres of the brain, as shown below. Note that language is listed as a function of the left hemisphere!



<http://brainmadesimple.com/left-and-right-hemispheres.html>

You can view the coronal, sagittal, and horizontal slices that intersect at three different voxels using the button labeled **show new voxel**.

Q2. For each of the three voxels, indicate whether the location of the voxel is in the *anterior* or *posterior* part of the brain, and whether it is in the *left* or *right* hemisphere.

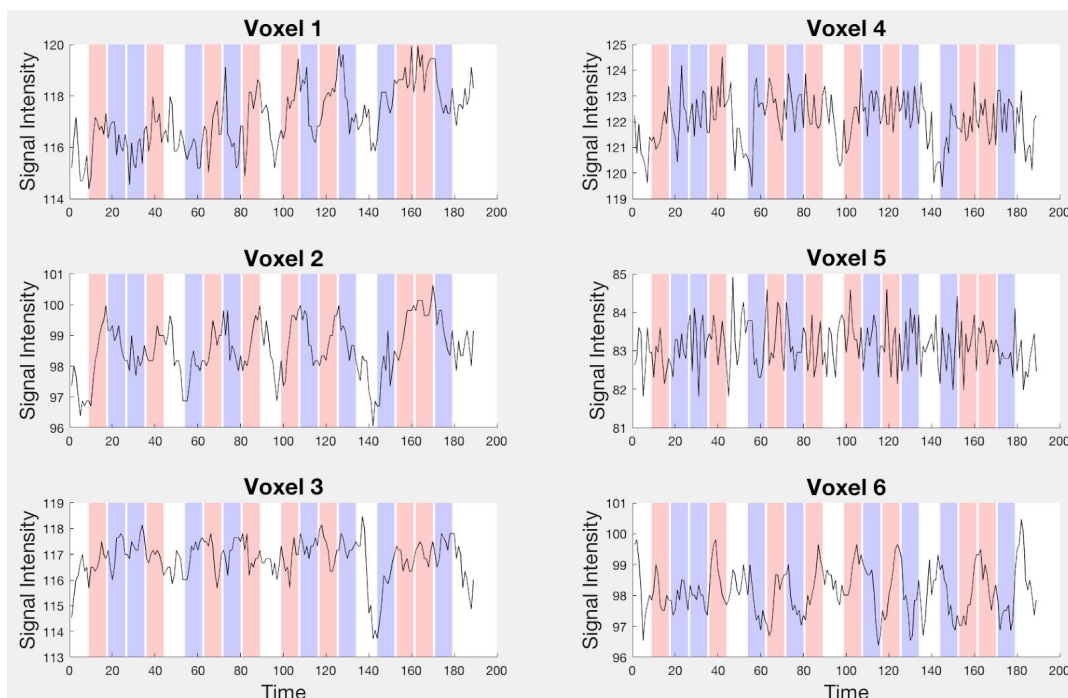
To view the functional data, select the **fMRI** option from the pull-down menu. The label on the **show new voxel** button will change to **show time series**.

Q3. How does the appearance of the fMRI data differ from that of the anatomical data? We noted earlier that the fMRI data has *lower spatial resolution* - *why* is this the case?

Ultimately we want to analyze how the fMRI signal changes over time, as the subject reads sentences versus nonwords over the course of many trials. To see how the activity measured for a few individual voxels varied over time, click the **show time series** button. A new figure window will appear, showing a set of 6 graphs, as shown on the next page. (You may see a warning in the Command Window about the image being too large to fit on the screen - you can ignore this.) Each graph shows the signal intensity (in black) for a different voxel, as a function of time (Voxel #2 in this figure corresponds to the voxel located at the intersection of the three slices of the fMRI data viewed above). The background red and blue bars show the time intervals in which the subject was reading a sentence (red) versus reading a sequence of nonwords (blue). During the intervals with white background, no text appeared.

Q4. For each of the six voxels, indicate which of the following descriptions best characterizes the activity of this voxel, relative to a baseline level of activity that appears during periods of no reading (shown with the white background). Also indicate whether you think this voxel is contained within a brain area that may be engaged during the processing of language.

- (1) signal increases equally in response to both sentences and nonwords
- (2) signal decreases equally in response to both sentences and nonwords
- (3) signal increases in response to sentences, but decreases in response to nonwords
- (4) signal decreases in response to sentences, but increases in response to nonwords
- (5) signal increases strongly in response to sentences, but weakly in response to nonwords
- (6) signal does not appear to change consistently when reading sentences or nonwords, compared to the baseline activity



Q5. The BOLD signal will sometimes show a *linear drift*, which is a very slow overall increase (or decrease) in the signal intensity over time. This slow drift can be the result of instability in the magnetic fields of the MRI scanner, or to slight head movements or changes in basal metabolism in the subject. Consider the 6 voxels shown above. Which voxel shows the strongest linear drift over time, and is this drift a slow *increase* or *decrease* in signal intensity as time increases?

II. The General Linear Model for the Analysis of fMRI Data

In this part of the lab, you will examine the signal time-series for a different set of individual voxels more closely, to determine whether they are located within a brain region that is engaged in language processing. The **General Linear Model** (GLM) is a method used to *predict* the fMRI signal for single voxels based on the known conditions of the experiment and measured


activity within the voxels. The difference between the predicted and actual fMRI signals can be used to assess the relationship between the neural activity within individual voxels and the property being manipulated in the experiment, in this case, the linguistic content of viewed text.

The summary of the General Linear Model below, and the figures, are drawn from Idan Blank's lectures recorded in these [Youtube videos](#).

In the General Linear Model, a predicted fMRI signal is constructed from a **linear combination of predictors**. You can think of the predictors as *mock signals* that each model the contribution to the brain signal we expect from a particular variable in the experiment. In this case, we can approximate the brain signals using three predictors:


- 1) A **constant predictor** that models the baseline response in a voxel - the spontaneous activity, recorded when the stimulus is not present, varies somewhat over time, but this will be modeled with a predictor that remains constant throughout the experiment
- 2) A **sentence predictor** that models the change in a voxel's response when the subject is reading valid sentences
- 3) A **nonwords predictor** that models the change in a voxel's response when the subject is reading a sequence of nonwords

In the figure below, the horizontal dimension represents time. The top pair of signals indicate periods of time (trials) when a sentence is displayed to the subject during the experiment (i.e. times of raised signal in the red plot labeled **S**) and periods of time (trials) when nonwords are displayed (times of raised signal in the blue plot labeled **N**). You learned earlier that the metabolic changes captured by the BOLD signal, which reflect changes in neural activity, take place over a time scale of several seconds. In response to a stimulus change in the experiment, such as the appearance of a sentence on the computer monitor, the change in the BOLD signal will be delayed and will increase over a time period of several seconds. The bottom two plots in the figure below show *predictions* for how the BOLD signal might change over time in response to sentence trials (red plot) and nonwords trials (blue plot).



We will create an overall prediction for the time course of the BOLD signal for a particular voxel by combining predicted responses to the sentence and nonwords stimuli with an estimated level of background activity in the voxel. In particular, we will **approximate the signal time-series** of


a voxel by creating a **linear combination of the three predictors** (also called *regressors*) shown below. Each predictor will be associated with a **weight** that captures how much each component contributes to the overall signal:



Suppose the **baseline**, **sentence**, and **nonwords** predictors shown above are stored in three vectors named **x1**, **x2**, and **x3**, respectively. Let **b1**, **b2**, and **b3** denote the weights associated with each of the three predictors, for a single voxel. The predicted signal (**p**) for this voxel is then given by the following expression:

$$p = b1 * x1 + b2 * x2 + b3 * x3$$

The following figure shows an example of a set of weights associated with the three predictors, and the resulting prediction (purple graph at the bottom) for the true signal computed from the above linear combination, using the specified weights.



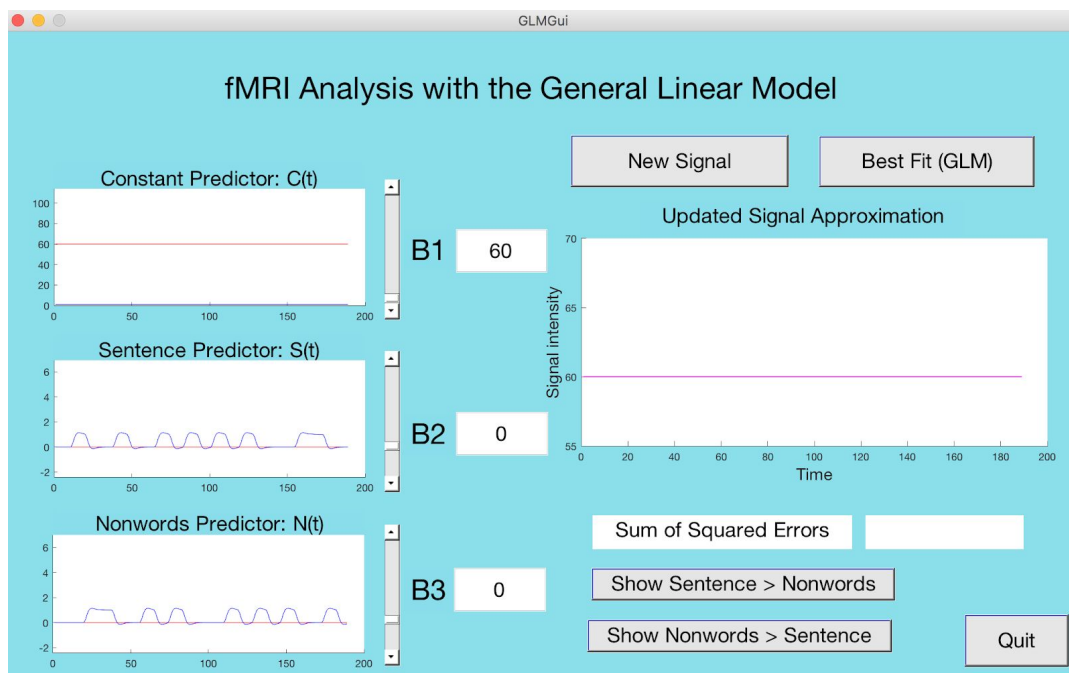
For the first few seconds (on the left), the linear combination has a value of 97.8, the weight associated with the constant predictor ($b1 = 97.8$). The sentence and nonwords predictors have a value of 0 during this time, so they do not contribute to the linear combination. When a response to sentences is expected, the predicted signal increases about 1.2 units above the baseline, because $b2 = 1.2$. When a response to nonwords is expected, the predicted signal is about 0.56 units above baseline, because $b3 = 0.56$.

We will see that it is possible to compute a set of beta weights that yield a predicted signal (i.e. linear combination of predictors) that **best approximates** the true signal for each voxel in the fMRI data. We can then use this analysis to determine which voxels can be reliably associated with language processing regions in the brain.

To explore the General Linear Model, enter the following command in the Command Window:

```
>> GLMGui
```

This opens a window for an interactive program that initially looks like this:



On the left, the three predictors are shown in blue, each with a height of about 1 unit. (In the case of the constant predictor, you'll see a blue line close to the x axis.) The beta weights, B1 (for the constant predictor), B2 (for the sentence predictor), and B3 (for the nonwords predictor) are shown to the right of each graph. You can change the beta weights by adjusting the sliders next to each graph, or by typing a number into the text boxes (**you need to hit the return key after changing the number in the box!**). Each beta weight is restricted to a range of values that are relevant for the experiment here: the value for B1 can be set to any number in the range from 60 to 110, and both B2 and B3 can be set to values in the range from -2 to +6. Initial values for the three weights are shown in the central text boxes. The red plots on the left show the result of multiplying each beta weight by the corresponding predictor.

The graph on the right shows the linear combination of the three predictors (in purple), using the specified beta weights. **Set different values of the weights and observe the changes in the weighted predictors (red curves in the plotting areas on the left) and the predicted signal (shown in purple on the right).** Note that the range of values shown for the signal intensity on the y axis will change as you adjust the weight associated with the constant predictor, B1.

Q6. Set B1 to 90, set B2 to 5, and set B3 to -1.5. Observe the resulting linear combination of predictors shown on the right and answer the following questions:

- At what times does this predicted signal show a response to sentences (list at least 3 sample times)? What is the (approximate) signal intensity during these “sentence response” times?
- At what times does this predicted signal change in response to nonwords (again list at least 3 sample times)? What is the approximate signal intensity during these “nonwords response” times? Does the activity level at these times increase or decrease relative to the baseline level?

Linear combinations of predictors as approximations of real signals

You will now try to construct approximations of real signals drawn from the fMRI data, for a few sample voxels. The program contains four sample signals for you to approximate. To generate each new signal, click the **New Signal** button near the top of the GUI window. A signal (in black) will appear on the graph on the right, the three beta weights will be reset to their initial values, and the linear combination created from these weights will be shown in purple on the right. Below the plot, you will see a text box with a number that shows the **Sum of Squared Errors** between the predicted and actual signals (this number will initially be very large). We have seen this error measure in other contexts. Suppose the actual signal is stored in a MATLAB vector named **s**, and predicted signal is stored in a vector named **p**. The sum of squared errors between these two signals can be computed in MATLAB with the following code statement:

$$\text{SSE} = \text{sum}((\mathbf{s} - \mathbf{p}).^2)$$

Q7. Adjust the beta weights for each of the three predictors to find a good approximation for the first signal (i.e. try to construct a purple trace that matches up fairly well with the black trace on the graph). When you have a good match, record the three beta weights and the resulting error (SSE) in the table below. Then click the **New Signal** button to advance to the next sample signal. Repeat the process of finding a good approximation for each of the four sample signals, recording the weights and errors in the table below. (Note that the signal number will appear above the display area on the right.)

	B1	B2	B3	SSE
Signal #1				
Signal #2				
Signal #3				
Signal #4				

Q8. For each of the four signals, click on the **Best Fit (GLM)** button to see the combinations of beta weights that *best approximate* each signal (i.e. minimize the SSE). Record the best weights and corresponding errors in the table below:

	B1	B2	B3	SSE
Signal #1				
Signal #2				
Signal #3				
Signal #4				

Q9. Based on your own analysis for each of the four voxels, as well as the “best approximation” results above, indicate whether you think the voxel responds more to sentences or to nonwords. What criteria are you using to make this decision?

In this analysis, we only considered three predictors corresponding to baseline activity and responses to sentences and nonwords. There are additional factors that impact the fMRI signals and our ability to predict them with the GLM, including small head movements during the recording session, and variations in the modeled temporal delays in the BOLD responses. Additional predictors for each of these factors can be incorporated into the GLM, leading to improved predictions of the activity of individual voxels. Typical fMRI experiments also include multiple runs to increase the reliability of the results.

Putting it all together

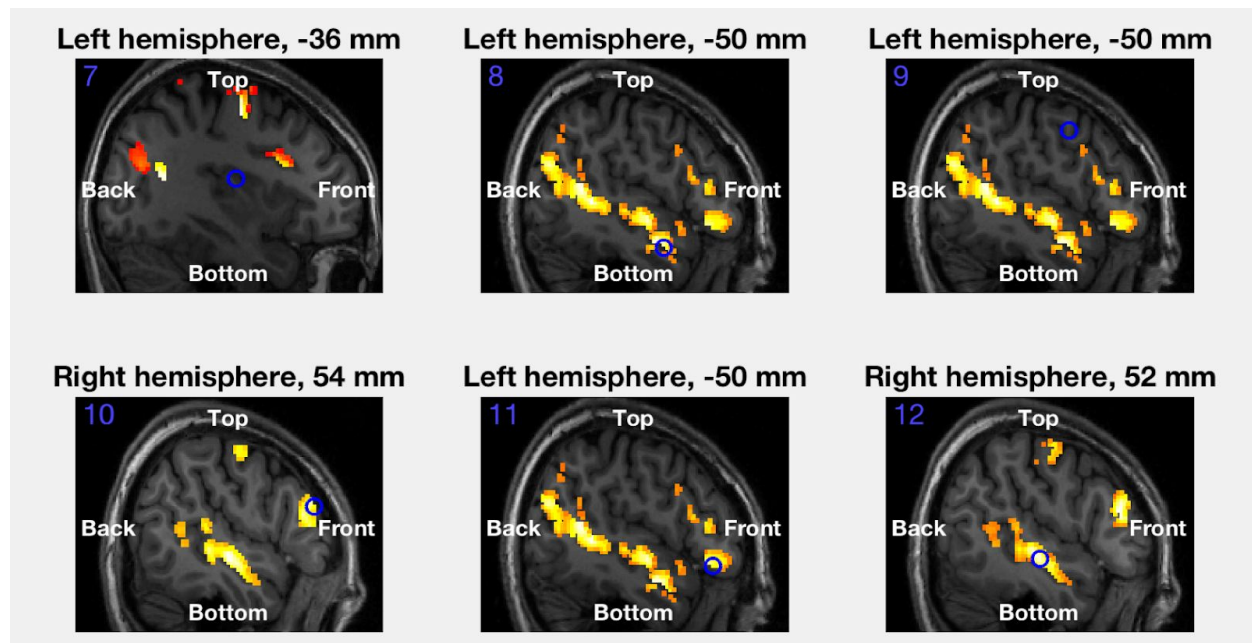
The experiment was designed to localize the language areas in the brain. After the GLM is run for every voxel in the brain, we would like to use the results of this analysis to identify voxels associated with the language areas, i.e. voxels that respond more to sentences than nonwords. This can be done by finding voxels for which the **beta weight for sentences is greater than the beta weight for nonwords**. A comparison between beta weights of this sort is called a **contrast**. By also taking into account the **measure of error (SSE)** in the prediction, we can determine whether a particular difference between two beta weights is **statistically significant**, i.e. large enough to believe that the difference is “real”, reflecting meaningful differences in the underlying neural activity that is captured in the fMRI signals. We will not go into the details of the statistical analysis, but it is important to be aware of this essential step of the process.

In this final exercise, you will view the overall results of the fMRI analysis for the entire brain, and interpret these results. For each voxel in the fMRI data, the following steps were performed using software developed by Idan Blank:

- 1) The full time-series of the fMRI signal was extracted for that voxel
- 2) A GLM was run to find the best beta weights for the predictors based on the signal
- 3) An approximation of the true signal was constructed using a linear approximation
- 4) The errors of the signal approximation were computed

Once the linear approximations are computed, we can define a **sentences>nonwords contrast**, which identifies all of the voxels for which the beta value for the sentences predictor is *significantly greater* than the beta value for the nonwords predictor. We can also define a **nonwords>sentences contrast** for which the opposite is true.

To view some of the results of this process, click on the **Show Sentence > Nonwords** button in the bottom right corner of the GUI window. Two figure windows will appear, each showing images of 6 sagittal slices through the *left hemisphere of the brain*, showing an **activation map** for this contrast. (Again, you may see a warning in the Command Window about the image being too large to fit on the screen - you can ignore this.) The images in one of the figure windows is shown below:



Voxels showing a significant effect are colored on a red-yellow-white axis. The brighter the color (i.e. closer to white), the more significant the contrast (voxels colored in shades of grey do not show a significant effect). **For the Sentences > Nonwords contrast, the activation map will show you the language system in the brain!**

Q10. Each of the 12 images that you see after clicking the **Show Sentence > Nonwords** button shows a specific voxel circled in blue. Referring to the voxel numbers written in blue in the upper left-hand corner of each image, which voxels respond significantly more to sentences than to nonwords? Based on the locations of these voxels, what does this data tell you about where the language system is located in the brain?