Nick Dotson
Bio 504 Final

Examination and Implementation of A Bayesian Hierarchical MCMC Metropolis Analysis

The objective of this paper is to discuss the Journal of Neurophysiology article "Hierarchical Bayesian Modeling and Markov Chain Monte Carlo Sampling for Tuning-Curve Analysis" (B. Cronin et al. 2010) and to provide an example of the analysis in action. The primary focus will be on the components of the hierarchical model, MCMC Metropolis sampling, and the use of credibility intervals to determine changes in individual parameters. The authors make the point that Bayesian analysis are rare in the field of neurophysiology due to difficulties in implementing numerical sampling algorithms. In an effort to change this they provide a Matlab/Java toolbox for tuning-curve analysis. This toolbox was used for the implementation contained in this paper. Also, all data used for the implementation was provided by Charles Gray (Gray Lab, MSU).

The article begins with an explanation of a single cell tuning-curve analysis. The basics of which are cells in primary sensory and motor areas exhibit changes in their firing rates based on the changes in a single parameter of a stimulus. Furthermore, these change in firing rate tend to change smoothly with some stereotyped function. For example, in primary visual cortex the cells have a circular Gaussian function.

The first step in their analysis is to assume a tuning-curve function and a noise model. To determine an appropriate tuning-curve function previous experiments are reviewed. If no tuning-curve function was assumed then it would be more difficult to assess the cells selectivity to stimuli and to compare different cells. The article recommends a circular Gaussian function for orientation selective cells in primary visual cortex based on previous studies. The noise model is established to be Poisson. This is because the probability of a cell firing roughly follows a Poisson process.

This combination of an assumed tuning-curve function and noise model is the hierarchical component of this analysis. Once these components are specified the parameters of interest are estimated.

The parameters are:

1.) Baseline firing rate
2.) Tuning width
3.) Preferred Orientation
4.) Maximum firing rate (Amplitude)

The baseline firing rate is the rate at which the cell fires when it is not responsive to the orientation of the stimulus. Tuning width is the half the width of the distribution at half the distance between the maximum firing rate and the base line. Preferred orientation is the orientation at which the cell fires the most spikes per unit time. This can be found by looking at the orientation associated with the maximum firing rate.

The priors used for these parameters are uniform and realistic. This means that any value in their distributions are equally likely and are within the realm of plausible values. For example, the preferred orientation of a cell might range from 0 to 180 degrees so these are the boundaries of its prior distribution. Another example is the baseline firing rate. A cell typically only fires once every 2 ms due to an absolute refractory period in which the cell can not fire. This limits the maximum rate a cell can fire to 500 Hz so a practical prior distribution would be flat and ranging from 0 to (500 Hz / time interval stimulus is shown). This makes sure that the posterior distribution of these parameters only takes on values that are physiological. Bayes equation is used to estimate the posterior distribution (see article for equation). A hyperparameter is shown in place of individual priors for each parameter. This is because the individual parameters are estimated with information about the tuning curve function and noise model.

To sample the posterior distribution the authors used MCMC Metropolis with 10,000 burn-in samples and then 20,000 samples taking every 400 to avoid serial correlation of samples. The burn-in is used to bring the sampling into stationarity (B. Walsh 2004). This means that it is acurately drawing samples and not just meandering around the edges or features inside. The Metropolis procedure is as follows. Starting with an arbitrary sample value from parameter space a new sample value is chosen at random from a normal proposal distribution (mean = current sample). If the posterior probability of the new sample is greater than the current sample then it is accepted and a the process is repeated. If it is smaller then it will be accepted with a probability (new/current). If it is not accepted then the current value will be recorded again and the procedure will continue.

To implement this analysis a data set from an experiment performed in primary visual cortex on an anesthetized cat was acquired. A component of the experiment that this data set came from was to determine direction selectivity of the cells being recorded from (recordings were done extracellularly). This was done by showing the animal a large gabor on a computer screen for a brief period of time and then changing its orientation. A gabor is an image with alternating black and white stripes at some desired density that moves at some desired speed. As the gabor moves it excites the cell repeatedly. Once all desired orientations were shown a brief period of time was allowed to elapse and the series of orientations were shown again. This was done a total of 5 times which provided 5 data points at each orientation. Figure 1 shows the observations for the cell that will be used for the analysis. It is immediately evident that the cell fires preferentially to a certain orientation and it appears to change smoothly with a Gaussian shape.

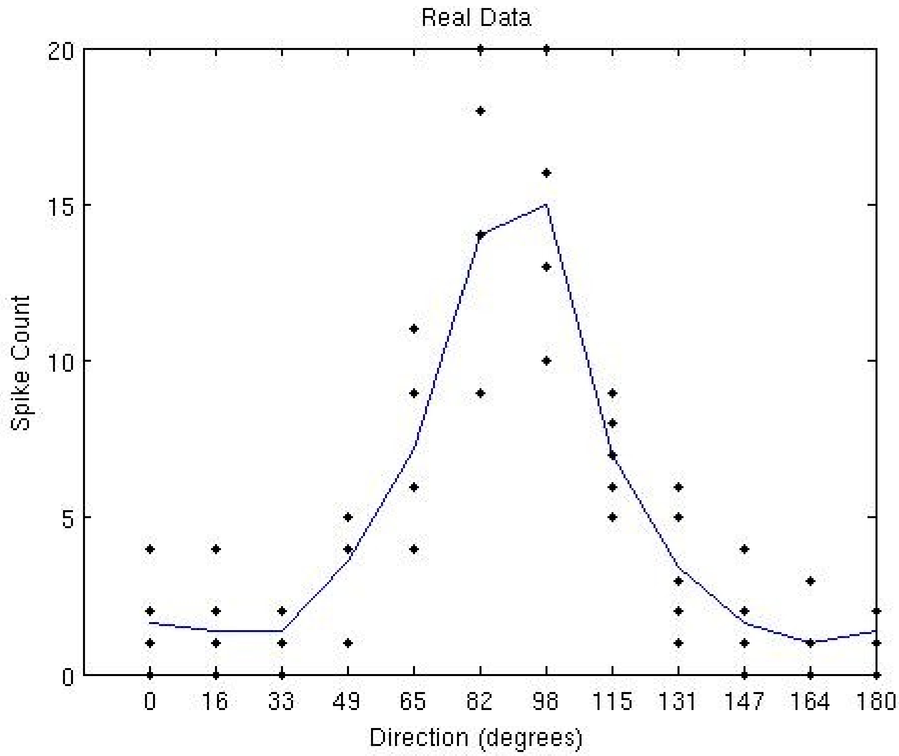

Figure 1: Plot of the spike counts at each orientation for 5 repetitions. The blue line indicates the mean at each orientation. Values are not jittered so some observations overlap.

Some work was required in getting the data into the appropriate format. Once this was accomplished the data was plugged into the Matlab function provided by the authors. As previously described and rationalized by the raw data a circular Gaussian tuning-curve function was used with a Poisson noise model. Samples for each parameter were obtained and plotted in figure 2. Credibility intervals were calculated for each parameter and are reported below. Collectively the posterior distributions provide the user with an actual probability for each value the parameter could take on given the data and the intervals that the true values are most likely to be in. For this analysis the median of each posterior distribution is considered the estimate of the parameter of interest.

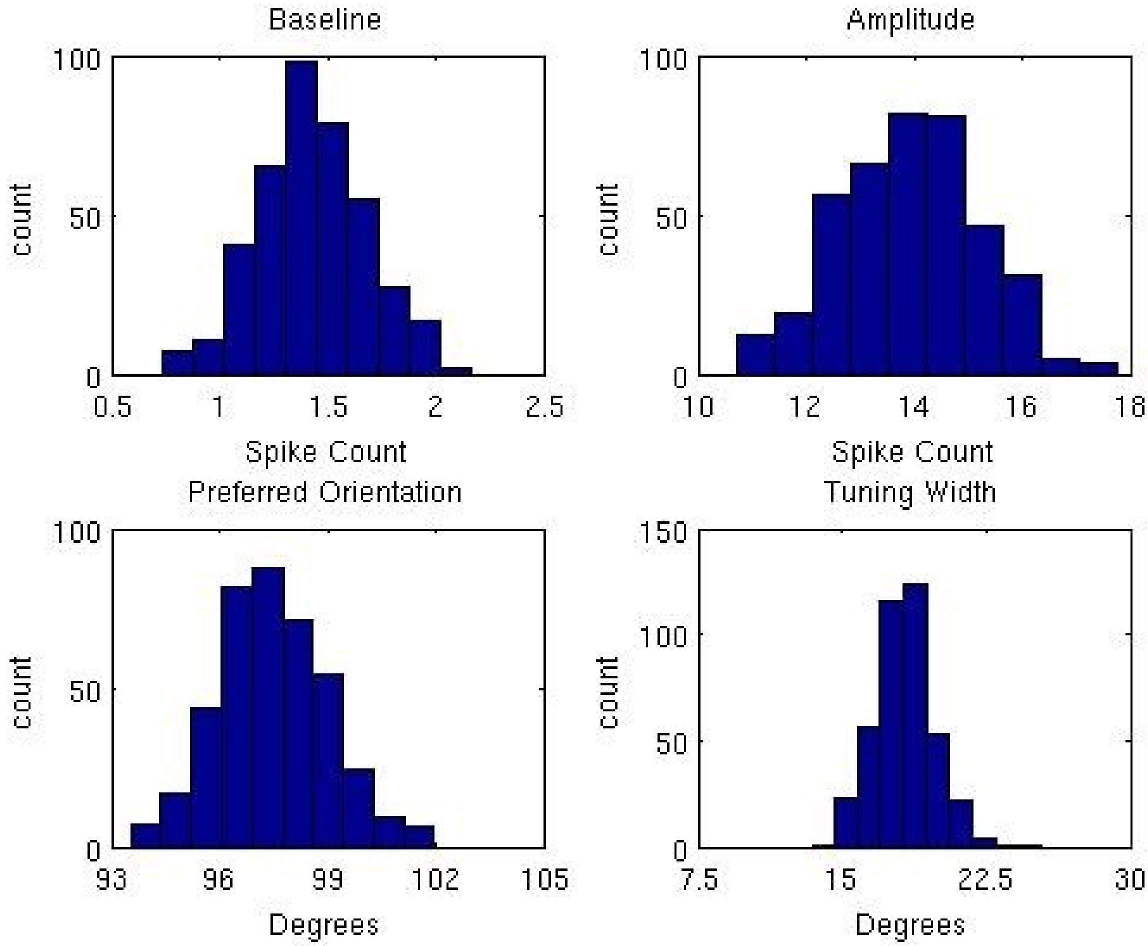

Figure 2: Posterior distributions for the parameters of interest. The y-axis is the count which equates to the probability of each value. The x-axis for baseline and amplitude (maximum firing rate) are in spike counts and preferred orientation and tuning width are in degrees.

Parameter:	**Estimated Value**	**Credibility Interval**
Baseline	1.1	.8 : 1.4
Tuning Width	3.9	2.8 : 5.1
Pref. Orientation	88.6	83.3 : 94.4
Amplitude	19.3	14.0 : 26.8

These posterior distributions can then be collectively summarized on a Gaussian tuning curve function to show the final fit to the data (see figure 3). It can be seen that the median does a very accurate job of fitting the data and that the samples are all fairly tight and uniform around the median.

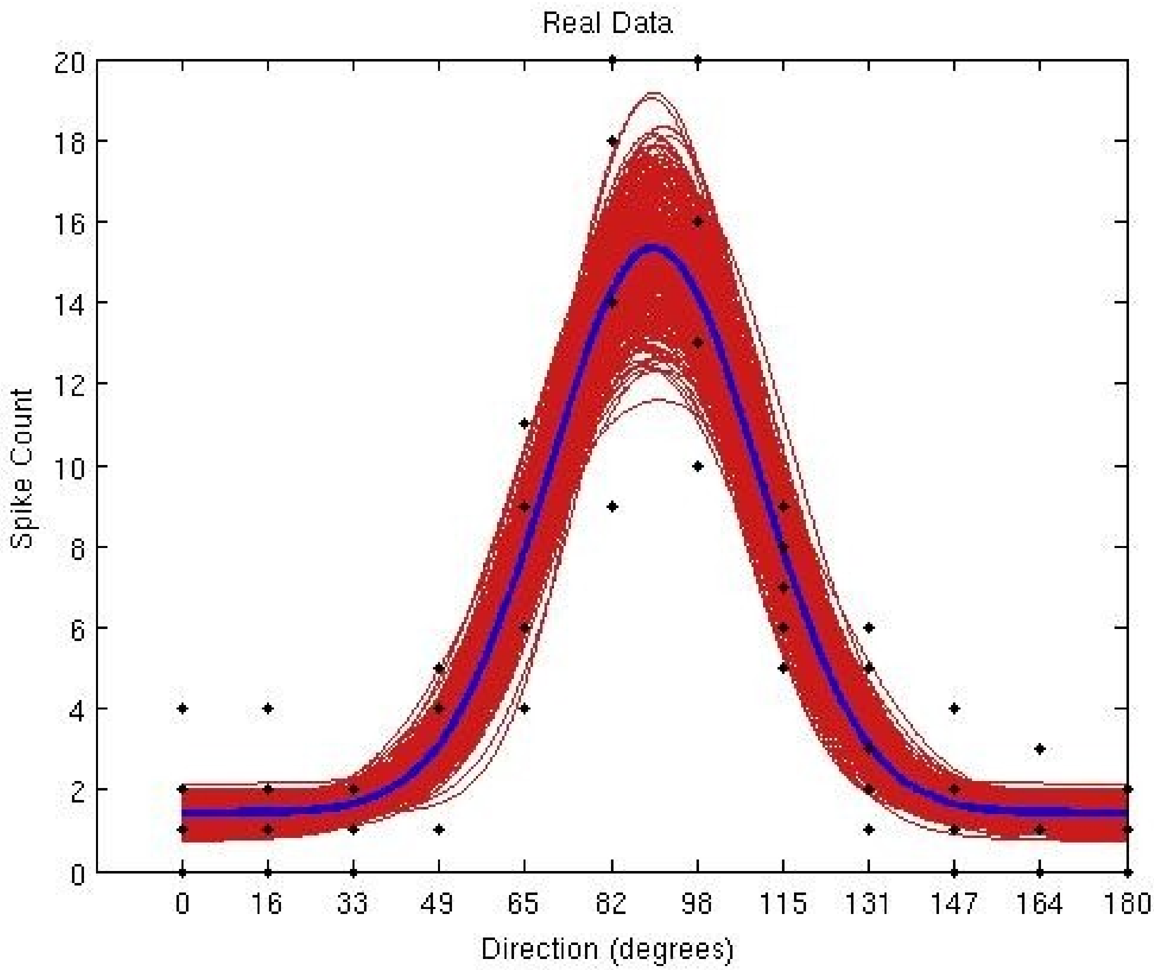

Figure 3: Samples from the posterior distributions from all parameters with a circular Gaussian function. The blue line indicates the median.

Another interesting topic in the article was testing to see if particular parameters change after a perturbation. A typical example would be to test if a cell responds differently to a stimulus when it is first presented than when it is presented at a later time. To illustrate this method using the data set shown above the data was modified for one parameter, baseline firing rate. This was done by subtracting 1.5 from all of the spike counts. This left all other parameter more or less the same. The analysis was completed as described previously with results shown in figures 4, 5.

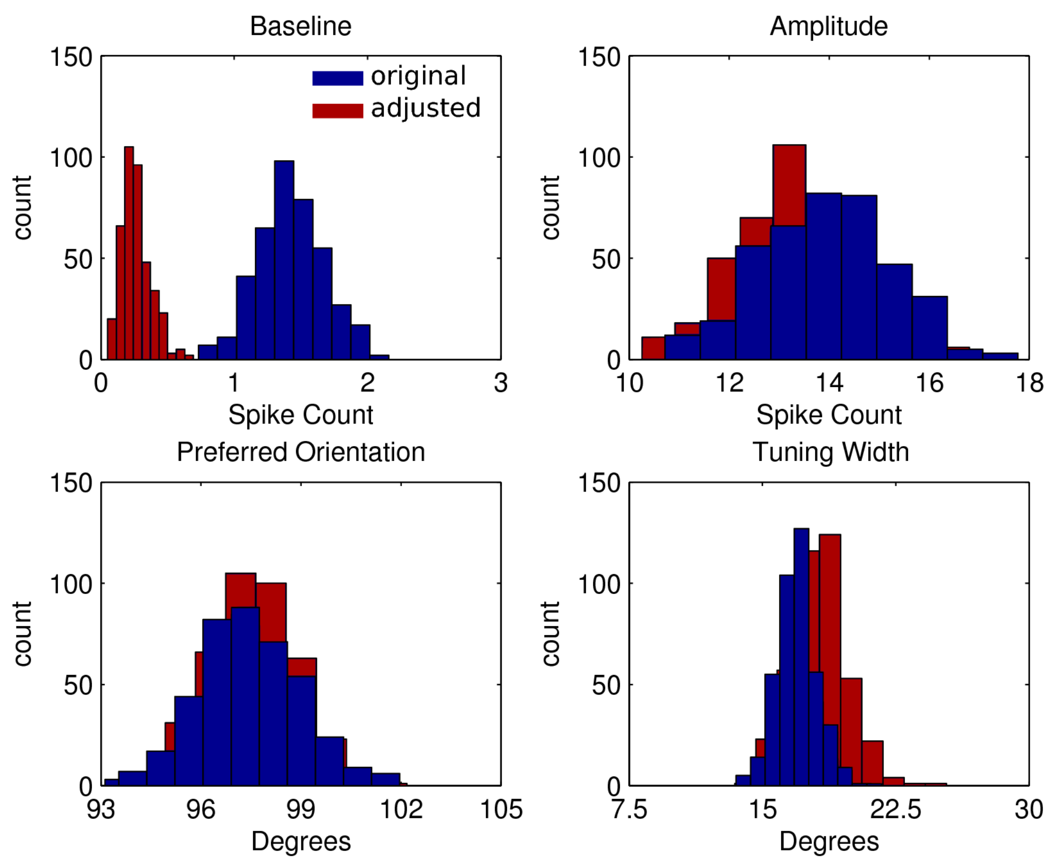

Figure 4: Posterior distributions for the parameters of interest. Blue histograms are from the original data and red histograms are from the baseline augmented data.

The posterior distributions in figure 4 showcase the drastic change in the baseline. Now, using credibility intervals (not shown) the individual parameters can be investigated. In this case a credibility interval is not even necessary for the baseline comparison. The two distributions do not overlap which indicates that the baseline is different between the original and augmented data. This effect can also be seen in figure 5 below. The new baseline now sits near zero.

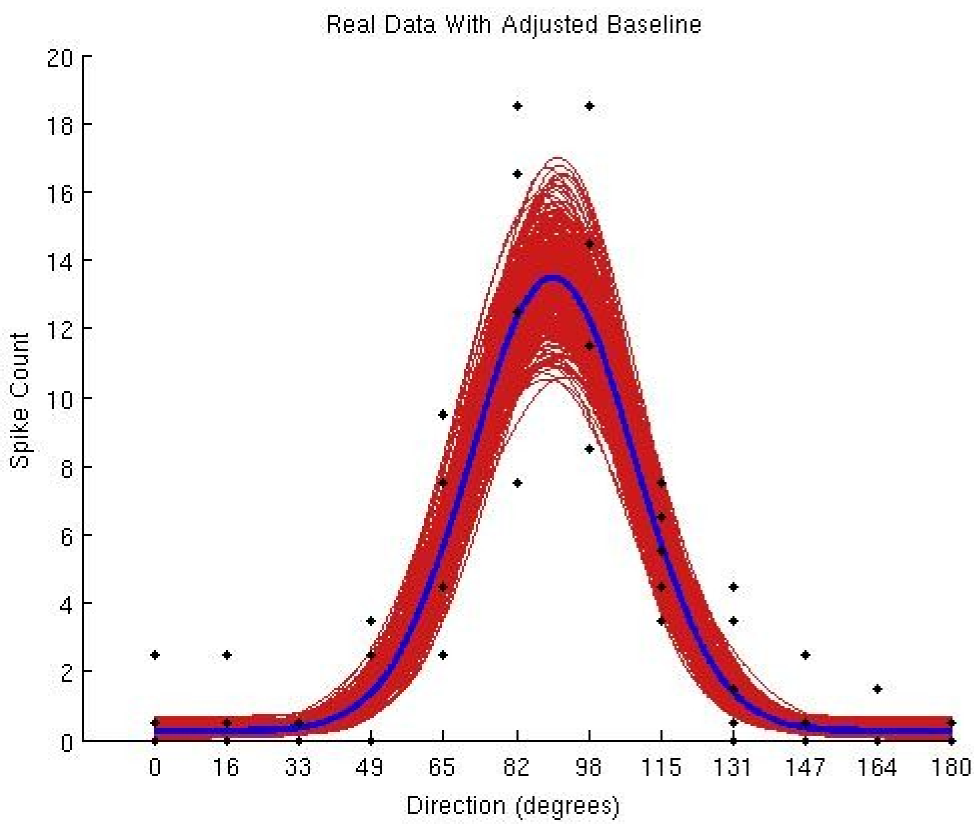

Figure 5: Samples from the posterior distributions from all parameters with a circular Gaussian function for the augmented data. The blue line indicates the median.

From this analysis we can conclude that this particular cell is highly tuned to a specific orientation of a stimulus. It has a fairly steep fall off on either side of its preferred orientation and it is unresponsive outside of a narrow range. The posterior distributions provided by the MCMC Metropolis sampling algorithm were all normal and without large outliers which were easy to work with and to interpret. This does not necessarily mean that the sample space was completely explored and the algorithm could be modified to check for these things. For instance the proposal distribution could be adjusted (W. R. Gilks et al. 1996). The technique was also shown on an augmented data set to illustrate that the credibility intervals could be used to determine if single parameters differ before and after a perturbation.

This concludes the exploration and implementation of the hierarchical MCMC Metropolis analysis. Much of the undiscussed portion of the article being reviewed focused on comparing alternative tuning curve estimation methods. When compared to other methods a highlight of the hierarchical Bayesian method is that it outperforms when the data sets are small. The data set used above is a good example of a small data set and the implementation shows that it provided clear cut and believable results.

References

B. Cronin, I. H. Stevenson, M. Sur, K. P. Kording, Hierarchical Bayesian Modeling and Markov Chain Monte Carlo Sampling for Tuning-Curve Analysis. J. Neurophysiology 103:591-602, 2010

B. Walsh, Markov Chain Monte Carlo and Gibbs Sampling. Lecture Notes for EEB 581, Version 26 April 2004.

W. R. Gilks, Sylvia Richardson, D. J. Spiegelhalter, Markov chain Monte Carlo in practice. Chapman & Hall/CRS 1996