Identifying Characteristics of Neural Network Dysfunction in Parkinsonian Mice

Taylor Pospisil
Department of Statistics, Carnegie Mellon University
tpospisi@andrew.cmu.edu

Valérie Ventura
Department of Statistics, Carnegie Mellon University
vventura@stat.cmu.edu

Max G'Sell
Department of Statistics, Carnegie Mellon University
mgsell@stat.cmu.edu

Jordan Rodu
Department of Statistics, Carnegie Mellon University
jrodu@stat.cmu.edu

Aryn Gittis
Department of Biological Sciences, Carnegie Mellon University
agittis@cmu.edu

November 25, 2015

Abstract

Keywords:

1 Introduction

Parkinson's disease (PD) is a disorder of the nervous system characterized by tremors, rigidity, movement dysfunction, and cognitive impairment.

One mechanism of interest in the progression of PD is dopamine loss. In PD patients there is a gradual dying off of dopamine-producing (dopaminergic) cells which leads to a decrease in dopamine levels in the brain. As a neurotransmitter, dopamine is essential to the regulation of brain activity and its loss leads to dysfunctional neuronal activity particularly in the *external globus pallidus* (GPe) of the *basal ganglia*. This region is associated with motor function and it is believed that the motor symptoms of PD originate from this region as normal activity patterns are disrupted by dysfunctional activity. [6]

Previous work has identified some of the characteristics of the network dysfunction caused by low dopamine. These include

β-oscillations in LFP Studies in humans such as [3] and [2] have found evidence of increased power in the β-frequency (15-30 Hz) band of the LFP spectrum. The LFP signal is believed to represent the sum of the electrical activity of all neurons in the vicinity of the probe. Thus an increase in a particular frequency band corresponds to an increase in synchronized oscillation at the corresponding frequency.

Decreased firing rates Initially it was believed that changes in the firing rate of neurons was a cause of PD, however recent studies have failed to detect such changes [8].

Synchronous inhomogenous patterns of firing between neurons An alternative model has been proposed by [10] that a pattern of burst firing may be responsible. These patterns manifest as "on-off" patterns of periods of high firing followed by periods of low firing.

Small time-scale synchrony between neurons In addition to synchrony in firing rates, we also expect short-scale synchrony in which neurons cofire in small time windows [7].

Taken as a whole, these symptoms can be characterized as increased oscillatory activity as waves of activity that dominate the natural patterns present in healthy brains. This rhythmic behavior interferes with the spatial and temporal encoding of network activity

and contributes to the motor dysfunction symptoms. Currently, disrupting this rhythmicity through interventions such as deep brain stimulation have proven effective, however we still lack an understanding of the mechanisms behind this rhythmicity and how it develops as the disease progresses.

Thus for this work we are investigating how previously studied indicators of network dysfunction develop over the course of PD progression. To this end we introduce a novel experimental design in which we gradually deplete dopamine through multiple low-doses of 6-OHDA. In this way we spread the dopamine loss over a span of weeks which more accurately mirrors the course of disease in human patients. Similarly to human patients, motor symptoms do not develop in mice until dopamine levels have been reduced by approximately 80%.

This design allows for the study of the progression of network dysfunction as a function of dopamine loss. We can obtain measurements from mice at various stages of dopamine depletion, ranging from healthy to completely depleted.

However, at the time of this writing we only have data collected from healthy and completely depleted mice with no data from partially depleted mice. As a first step in studying the progression of the disease, we set out to find differenes between the endpoints of healthy and completely depleted subjects.

2 Data

2.1 Experimental Design

Previous experiments studied dopamine depletion with an experimental design in which dopamine is depleted by 80-90% over a short time period through the administration of a single high dose of the toxin 6-OHDA. This toxin selectively kills dopaminergic cells in a process which mimics PD.

However, this short timeline of depletion does not accurately reflect the progression of PD in a clinical setting; dopamine levels of patients can decrease over a period of years before the first motor symptoms appear. Typically dopamine levels have dropped by 60-80% before motor-symptoms present [5]. The period of time before symptoms present is called the *prodromal* phase and the ability to diagnose during this stage could allow for more effective treatment.

2.2 Data Collection

The data we collect is extracellular electro-physiological recordings from awake mice. The area from which we record is the *external globus pallidus* (GPe) of the *basal ganglia*. As mentioned before, this region is associated with motor function and it is believed that the motor symptoms of PD originate from this region.

To obtain the electro-physiological recordings the mice are placed on a free-spinning wheel. The mice are then placed into a stereotaxic frame and a craniotomy is performed in which the piece of skull over the recording site is carefully removed. Then a 16-channel probe is carefully inserted into the target area.

Once inserted, the probe records neuronal activity for a session of approximately 15 minutes. Each channel can record either spike activity or the *local field potential* (LFP) channel.

A channel recording spike activity measures the electrical activity of nearby neurons. When a neuron fires the electrical potential of the cell rapidly rises and then falls. These events can be detected by the characteristic spike in voltage at the recording site. This signal is collected at a 40 kHz sampling rate. For purposes of minimizing data storage, only channels with firing neurons are selected; thus we typically have three or four channels per recording.

A channel recording the LFP measures the sum of action potentials of nearby neurons. Instead of individual spikes, the LFP reflects broader oscillatory patterns in neuronal activity. This signal is collected at a 1 kHz sampling rate.

2.3 Pre-Processing of spike channels

A necessary preprocessing step is to transform the raw channel signal into a *spike-train*, the times of the spikes. This process consists of two main stages: spike detection and spike sorting.

2.3.1 Spike Detection

Spike detection is the process of identifying the spikes occurring in the recording.

A common approach is to use a sub-threshold: if the signal drops below a certain voltage value we record a spike. This value should be low enough to avoid classifying noise as a

spike (false positive) while still high enough to avoid missing low amplitude spikes (false negatives). The difficulty is in deciding which value to set this threshold.

To obviate the need to select a threshold we develop a new method for automatically selecting thresholds that adapts the to noise characteristics of the signal.

For this approach we scan through the data collecting signal segments which pass the following criteria:

- Aligned on a local minimum (to avoid duplication)
- Local minimum below threshold
- Local maximum above threshold
- Local minimum is before local maximum

On this pass through the data we collect spikes along with segements of noise which happen to fit the criteria.

Then we run the procedure on the signal in reverse. Because spikes have their minimum before their maximum, on this pass we don't collect spikes, only noise. This allows us to fit a null distribution for the amplitudes of noise which passes our criteria. We model this distribution as a multivariate normal.

Once we have fit the null distribution the problem reduces to hypothesis testing: for the segments collected in the first pass we test against coming from the null distribution. We select spikes using a FDR of 0.01.

2.3.2 Spike Sorting

Because the activity of several neurons might be recorded at a single site, we need to attribute spikes to individual neurons. To achieve this we apply a technique called *spike sorting*. The spikes of a neuron have a characteristic shape determined by the specific morphology of the neuron's dendrites. We use this characteristic shape to sort spikes into clusters of similar shape which are generated by the same neuron.

Thus we take a short window of time (≈ 2 ms) around each spike, aligned on the minimal voltage of the spike. This represents the entire waveform of the spike. Using these curves we apply a state-of-the-art Bayesian clustering technique for our spike-sorting developed by [4].

3 Exploratory Data Analysis

3.1 Recording Artifacts in LFP channels

One problem encountered during the recording process is the interference of outside signals. The most common source of this interference is power line noise which has a frequency of 60Hz. Because this source is constant it can be removed with a notch filter centered at 60Hz to attenuate this noise.

However, there are other sources of interference which are sporadic. This can be seen in Figure 1; the large spikes in the LFP signal are caused by an extraneous source and are clearly distinct from the LFP signal. These anomalies can be caused in several ways; the contacts in the electrode may be loose or not properly grounded.

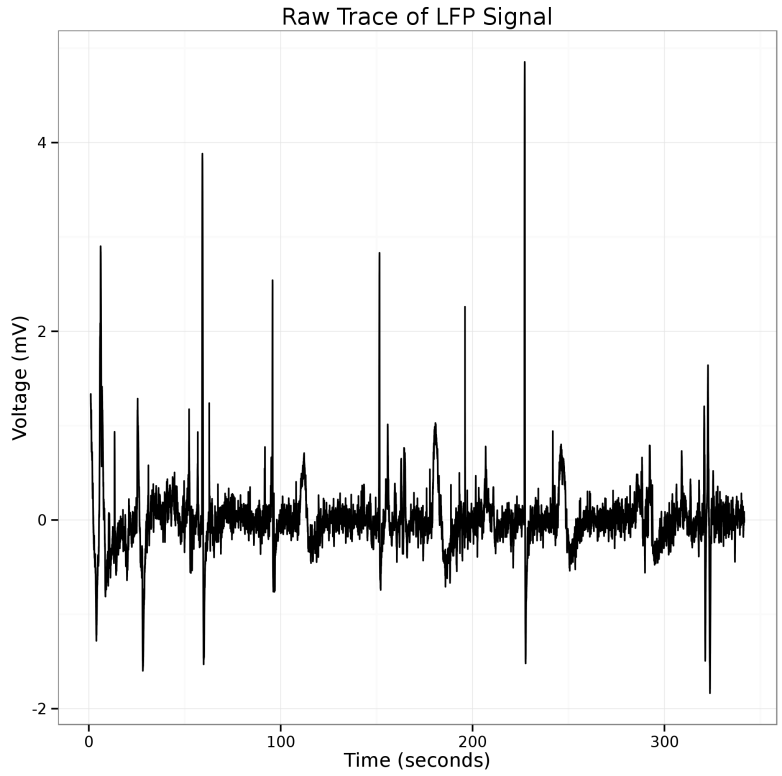

Figure 1: The spikes in the LFP signal are recording artifacts caused by electrical interference

The presence of anomalies can lead to problems estimating the spectrum. We will discuss a solution in Section 4.1.2.

4 Methods

4.1 Spectral Density Estimation

The power spectrum of a signal measure the distribution of power over a range of frequencies. Let $x(t)$ be the recorded LFP signal as a function of time t. Then we have, as a function of frequency f, the Fourier Transform

$$X(f) = \int_{-\infty}^{\infty} x(t) \exp(2\pi i f t) \, \mathrm{d}t \tag{1}$$

.

The power spectrum $p(f)$, again as a function of frequency f, is defined as

$$p(f) = 2|X(f)|^2 \tag{2}$$

.

The power in the β-range (20-40) Hz is then simply the integral

$$\int_{20\mathrm{Hz}}^{40\mathrm{Hz}} p(f) \, \mathrm{d}f \tag{3}$$

4.1.1 Periodogram

The periodogram is a technique to estimate the power spectrum $p(f)$.

This method involves taking the discrete Fourier Transform and defining the function on the discrete grid $f_k = 0, \frac{f_0}{N}, \frac{2f_0}{N}, \ldots, \frac{Nf_0}{2}$,

Because we only have a finite number of samples N, we can approximate the integral by the discrete sum

$$X(f_n) = \int_{-\infty}^{\infty} x(t) \exp(2\pi i f_n t) \, \mathrm{d}t \approx \sum_{k=0}^{N-1} x(t_k) \exp(2\pi i f_n t_k) \Delta_k \tag{4}$$

.

which can be computed quickly using the Fast Fourier Transform (FFT).

Following the definition in 2, the periodogram is again defined as

$$\hat{p}(f_k) = |X(f_k)|^2 \tag{5}$$

.

The sum of these estimates $\sum_{k=0}^{N/2} \hat{p}(f_k)$ is equal to the total power of the signal $\sum_{l=1}^{L} x(l)^2$ so we can interpret this estimator as a binned estimate of the power. That is, $\hat{p}(f_k)$ approximately estimates the binned power in a small neighborhood around f_k.

Expected Value and Variance of the Periodogram The expected value of the periodogram $\hat{p}(f)$ estimate is

$$\mathrm{E}[|\hat{p}(f_k)|^2] = \frac{1}{N} \int_{-\infty}^{\infty} p(s)|W(f_k - s)|^2 ds \tag{6}$$

where $W(s)$ is the spectrum of the window function, here a rectangular window. This is asymptotically unbiased as $N \to \infty$.

The variance of the periodogram is constant regardless of the length of the signal.

$$Var[\hat{p}(f)] \approx p(f)^2$$

It is interesting to note that this implies that $\hat{p}$ is not consistent. The reason is that increasing the size of the sample results in an estimate of $p(f)$ over a finer mesh of frequencies instead of improving the estimate at a fixed frequency.

Asymptotic Distribution of the Periodogram Under the assumption of a stationary Gaussian process, the periodogram has the following asymptotic distribution as $N \to \infty$ [9].

$$\frac{2\hat{p}(f_k)}{p(f_k)} \sim \chi_2^2 \tag{7}$$

.

4.1.2 Bartlett's Method

To obtain a consistent estimate of the power spectral density we use Bartlett's method [1]. This method splits the signal into K segments of length L within which the periodogram is calculated.

We denote the periodogram of these K segments as $\hat{p}_k(f)$. Our estimate then becomes

$$\hat{p}(f) = \frac{1}{K} \sum_{k=1}^{K} \hat{p}_k(f) \tag{8}$$

Because each window is calculating a periodogram, the same bias properties of periodograms hold for Bartlett's method as well. Thus, as long as the segment lengths $L \to \infty$ Welch's method is asymptotically unbiased.

The variance is reduced by a factor $\frac{1}{K}$. Thus as $K \to \infty$ and $L \to \infty$ Welch's method is a consistent estimator.

Correcting for Anomalous Noise Sources As discussed in Section 3.1 we encountered anomalous noise sources in our LFP recordings. These anomalies result in corrupted power spectra due to their high power.

Bartlett's method allows us to correct for these anomalies. When splitting the signal into windows we remove those windows containing an anomaly. This is achieved by a simple filter which removes windows whose power exceeds a threshold. We then use the average of the remaining windows as our estimate. As long as the anomalies are relatively rare we do not filter out too many windows.

Confidence Intervals for Bartlett's Method Using the asymptotic distribution for the periodogram we note that

$$\frac{2\hat{p}_n(f_k)}{p(f_k)} \sim \chi^2_2 \tag{9}$$

When we apply Bartlett's method averaging these variables we obtain the following distribution

$$\frac{2N\hat{p}(f)}{p(f_k)} = \frac{2\sum_{n=1}^{N} \hat{p_n}(f_k)}{Np(f_k)} \sim \chi^2_{2N} \tag{10}$$

For large enough N, the χ^2_{2N} is approximately a Normal distribution with mean $\mu = 2N$ and variance $4N$. Thus we have, approximately

$$\hat{p}(f) \sim N(p(f_k), \frac{p(f_k)^2}{N}) \tag{11}$$

5 Results

5.1 β-oscillations in LFP

To investigate the presence of β-oscillations in the LFP we examine the power spectra of the recordings. In Figure 2 we plot the spectra of two recordings, one with a bump in power in the β-band (15-30 Hz) and the other without.

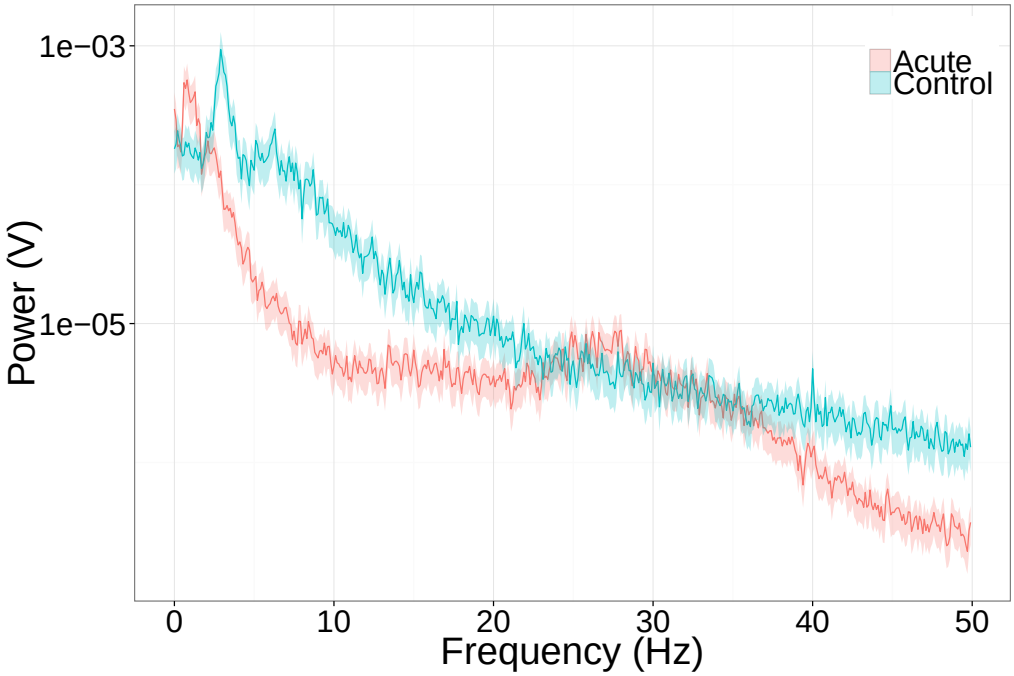

Figure 2: We can see the increased β-oscillations as a bump in the power spectrum of the LFP signals. The Acute (red) mouse has a clear increase in the β (20-40 Hz) range which is not present in the Control (blue) mouse.

To compare across Acute and Control conditions we need to quantify the presence of beta-oscillations. To this end we use the quantity "excess beta" which is the ratio of the actual power over interpolated power. To calculate interpolated power we use linear interpolation of the log power in the β-range. We use a local linear fit to calculate the values at the endpoints to adjust for the variability in estimating the spectrum.

This summary is invariant to differences in power between recordings which occur due to differences in electrode sensitivity.

We plot a histogram of these ratios by conditon in Figure 3. The vertical line at a ratio

value of 0 indicates a spilit between spectra which exhibit a bump above the interpolation and those below. As we can see, the acute distribution clearly has more spectra which exhibit bumps.

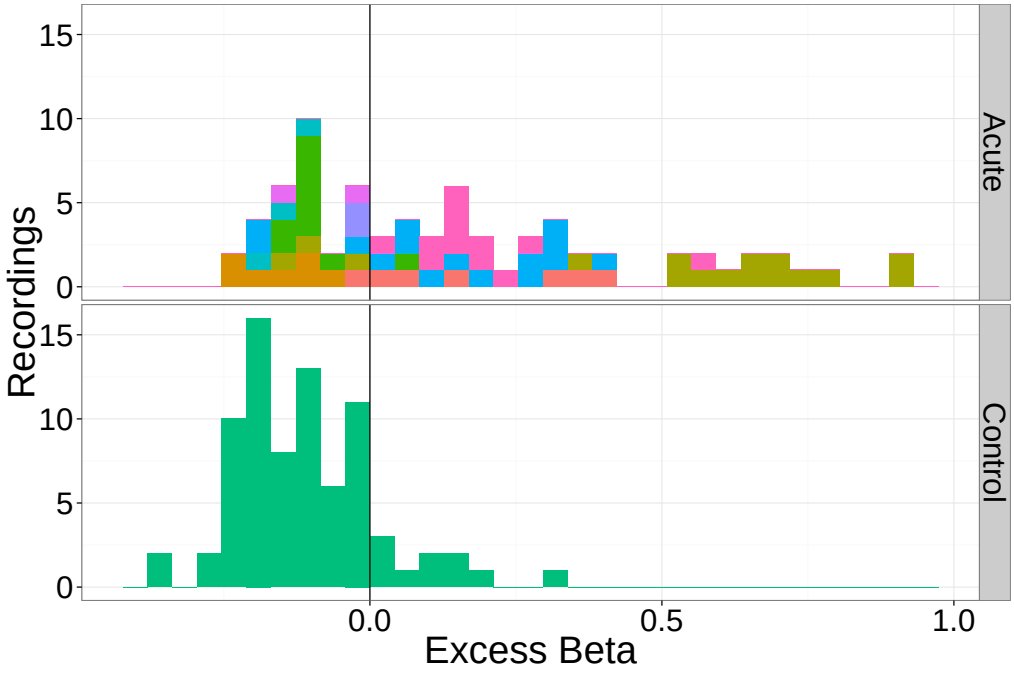

Figure 3:

When we examine a mouse with multiple measurements we see a clear development of the β ratio. Figure 4 shows the LFP at intervals of three, four, and five days post-depletion. We see that the bump is evident by the fifth day. We have 95% point-wise intervals, however because this is only a single mouse we cannot generalize further as we cannot assess between-mouse variability. With more mice and data we hope to be able to quantify this effect as a function of time, a proxy for dopamine depletion.

5.2 Decreased firing rates

We see evidence for decreases in firing rate in Acute mice. Figure 5 shows the distribution of firing rates for all recordings.

As we can see from the plot, the control mice have a longer tail with several neurons exhibiting firing rates about 50Hz while there is only one neuron from an Acute mouse exhibiting such a high firing rate. In general, the distribution for the Acute mice is smaller

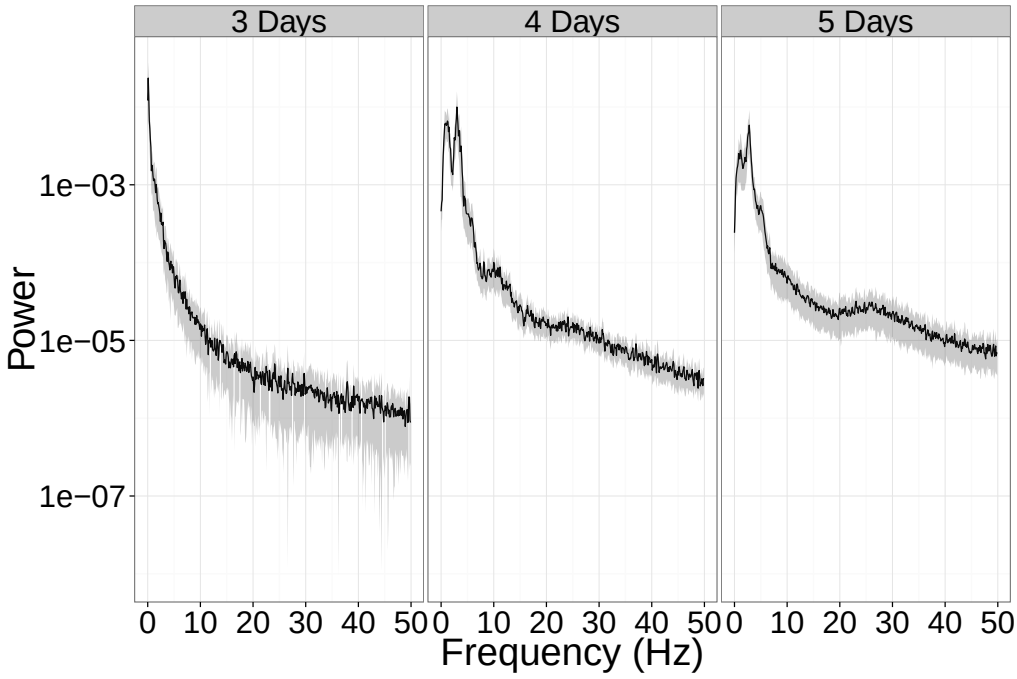

Figure 4: We can see the average spectrum recorded at 3, 4, and 5 post-depletion. We see a bump in the LFP signal develop as time progress and dopamine is depleted.

than the Control mice.

5.3 Synchronous inhomogenous patterns of firing between neurons

Looking at Figure 6 we can see that the firing rates of neurons are not constant. Most notable are that there are patterns of high firing followed by low firing. This inhomogeneous pattern is evidence of "on/off" bursting.

To compare between the acute and control populations we summarize the inhomogeneity using a standard measurement, the coefficient of variation of the interspike-intervals. The coefficient of variation is defined as

$$\text{CV} = \frac{\text{sd}}{\text{mean}}$$

For Poisson firing this ratio is 1. When inhomogenous firing is present, this ratio will be higher.

As seen in Figure 7, we don't see a difference in the Acute and Control distributions.

When we examine the firing rates of several neurons in Figure 8 we do see synchro-

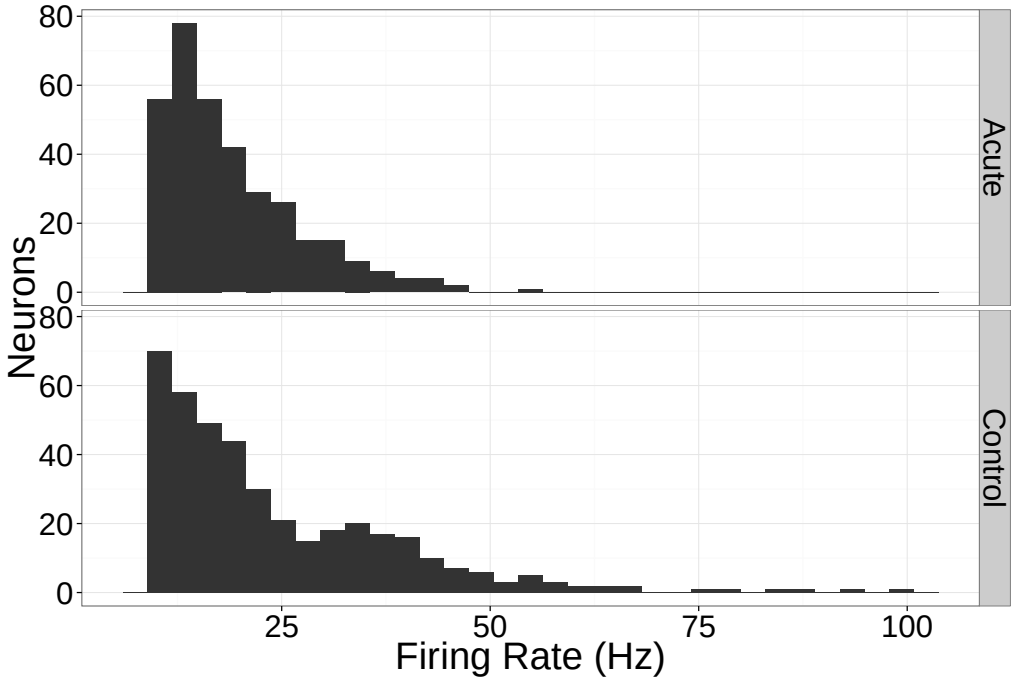

Figure 5: Firing rate for units by condition

nization. However there may be explanations other than synchronization. Plotted on the bottom of Figure 8 is the angular velocity of the mouse. The periods of time which seem to show the most dramatic synchronized changes in firing rate appear to be coincident with times when the mouse is moving. Further work is required to determine whether there is synchrony beyond that explained by movement alone.

5.4 Increased synchrony between neurons

To examine the short-term synchrony between a pair of neurons we can use the crosscorrelgram plot. These plots, seen in Figure 9 measure the firing rate of one neuron (call it A) with respect to the proximity of the other neuron (call it B). We calculate this by taking each spiketime of neuron A and collect the offsets with the spiketimes of neuron B. We do this for every spike of neuron A and plot the histogram of all of the offsets for neuron B.

Examining Figure 9 we see that there is a clear synchronous effect due to the large peak in the histogram. This indicates that the firing rate of the second neuron is much higher around the same time as the first neuron fires.

However, it's possible for this synchrony to be driven by correlations in firing rates as

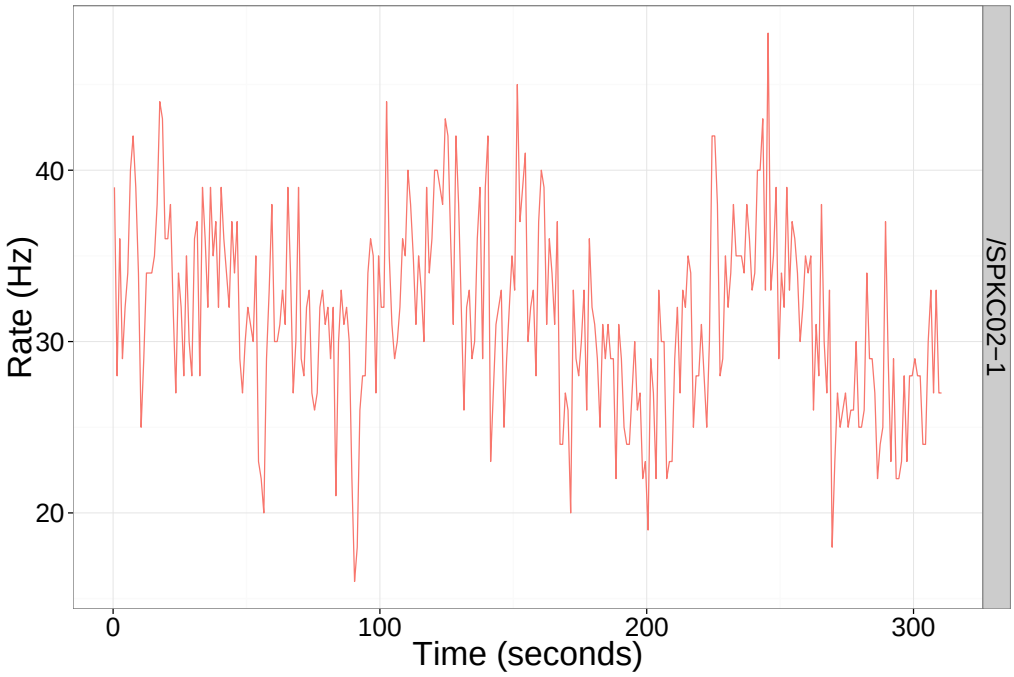

Figure 6:

we saw in the previous section. Thus we need to correct for any effect that's due to firing rates alone. For this purpose we use a jitter test.

The motivation behind the jitter test is to remove short-scale synchrony while keeping firing rates intact. To achieve this, the spike times of each neuron are jittered by sampling from a Normal distribution centered around the spike with a small variance. With these bootstrap samples we calculate the crosscorrelogram again.

To compare between the acute and control populations we need to summarize the synchrony of the neurons. To do this we calculate the ratio between the observed counts in a small window of 5ms and the average bootstrap counts within the same window. A high value of this ratio indicates that there's more synchrony than explained by firing rate.

As we can see in Figure 10, there is no clear difference between the acute and control mice. Most of the values are centered around 1, suggesting that their synchrony if any, is purely explained by firing rate. There aren't enough observations to tell if there are differences in the tails.

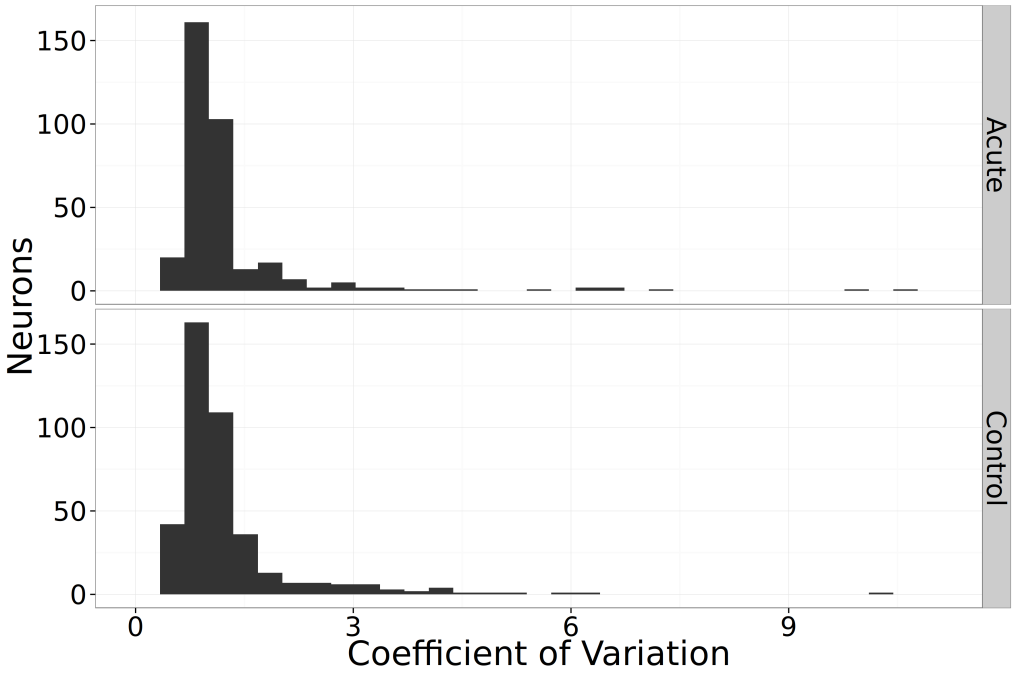

Figure 7: We see little difference in ISI CVs between Acute and Control mice.

6 Discussion

In conclusion, we have mixed results for finding the features of PD mentioned in previous literature. As a summary:

We did find evidence for

- Increased power in the β-frequency (15-40 Hz) band of the Local Field Potential (LFP)
- Decreased firing rate for neurons

We are currently still investigating

- Synchronous inhomogenous patterns of firing between neurons
- Increased synchrony between neurons

Our experiment is still ongoing and we continue to collect data. We are also anticipating receiving data on the dopamine levels of the mice. Because the dopamine depletion process is highly variable, the Acute mice have dopamine levels that can range from completely depleted to having more than half of their dopamine remaining. We plan on analyzing the features above as a function of dopamine rather than group labels.

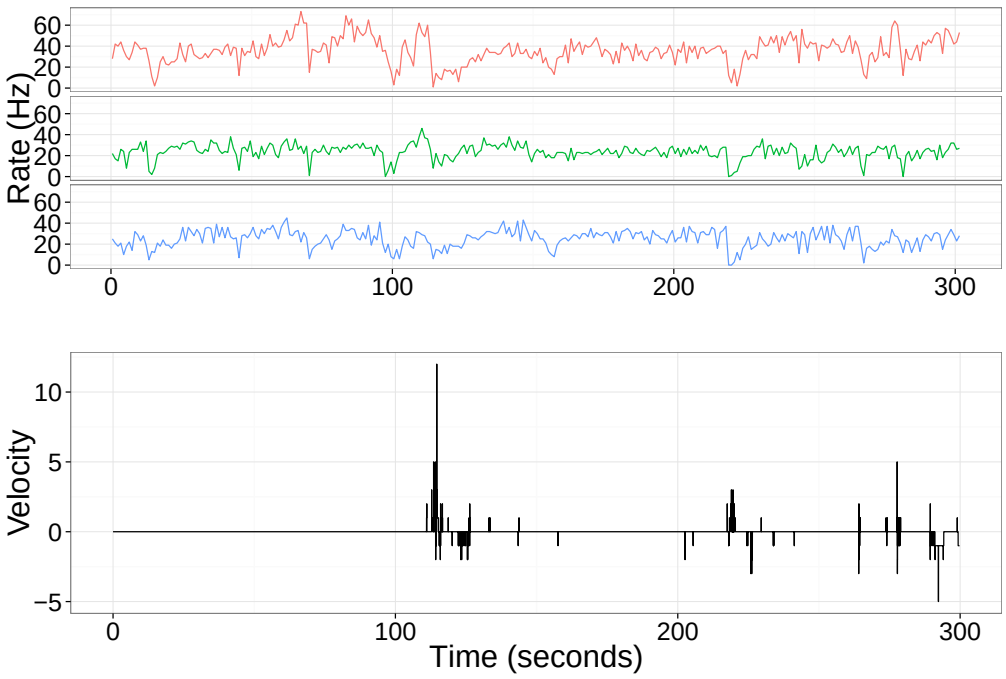

Figure 8: Firing Rate and Movement: Synchonized changes in firing rates appear to be coincident with times when the mouse is moving.

References

[1] MS Bartlett. Smoothing periodograms from time series with continuous spectra. *Nature*, 161(4096):686–687, 1948.

[2] Helen Bronte-Stewart, Crista Barberini, Mandy Miller Koop, Bruce C Hill, Jaimie M Henderson, and Brett Wingeier. The stn beta-band profile in parkinson's disease is stationary and shows prolonged attenuation after deep brain stimulation. *Experimental Neurology*, 215(1):20–28, 2009.

[3] P Brown. Bad oscillations in parkinson's disease. In *Parkinson's Disease and Related Disorders*, pages 27–30. Springer, 2006.

[4] David E Carlson, Joshua T Vogelstein, Qisong Wu, Wenzhao Lian, Mingyuan Zhou, Colin R Stoetzner, Daryl Kipke, Douglas Weber, David B Dunson, and Lawrence Carin. Multichannel electrophysiological spike sorting via joint dictionary learning and mixture modeling. *IEEE Transactions on Biomedical Engineering*, 61(1):41–54, 2014.

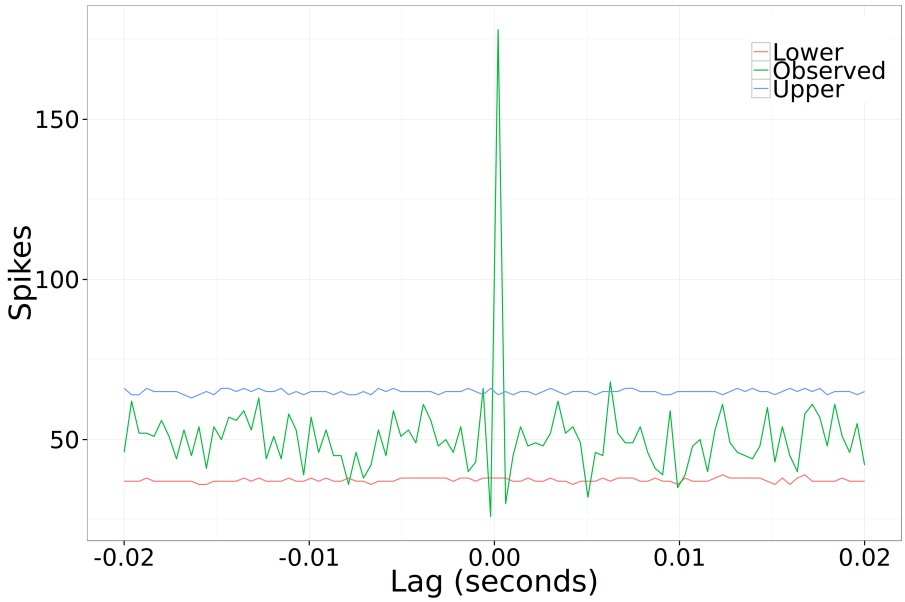

Figure 9: Crosscorrelogram: the lower and upper intervals indicate a 95% bootstrap interval

[5] Hsiao-Chun Cheng, Christina M Ulane, and Robert E Burke. Clinical progression in parkinson disease and the neurobiology of axons. *Annals of Neurology*, 67(6):715–725, 2010.

[6] Mahlon R DeLong. Primate models of movement disorders of basal ganglia origin. *Trends in Neurosciences*, 13(7):281–285, 1990.

[7] Joshua A Goldberg, Thomas Boraud, Sharon Maraton, Suzanne N Haber, Eilon Vaadia, and Hagai Bergman. Enhanced synchrony among primary motor cortex neurons in the 1-methyl-4-phenyl-1, 2, 3, 6-tetrahydropyridine primate model of parkinson's disease. *The Journal of Neuroscience*, 22(11):4639–4653, 2002.

[8] Atsushi Nambu, Yoshihisa Tachibana, and Satomi Chiken. Cause of parkinsonian symptoms: firing rate, firing pattern or dynamic activity changes? *Basal Ganglia*, 5(1):1–6, 2015.

[9] Donald B Percival and Andrew T Walden. *Spectral analysis for physical applications*. Cambridge University Press, 1993.

[10] TA Zirh, FA Lenz, SG Reich, and PM Dougherty. Patterns of bursting occurring in thalamic cells during parkinsonian tremor. *Neuroscience*, 83(1):107–121, 1998.

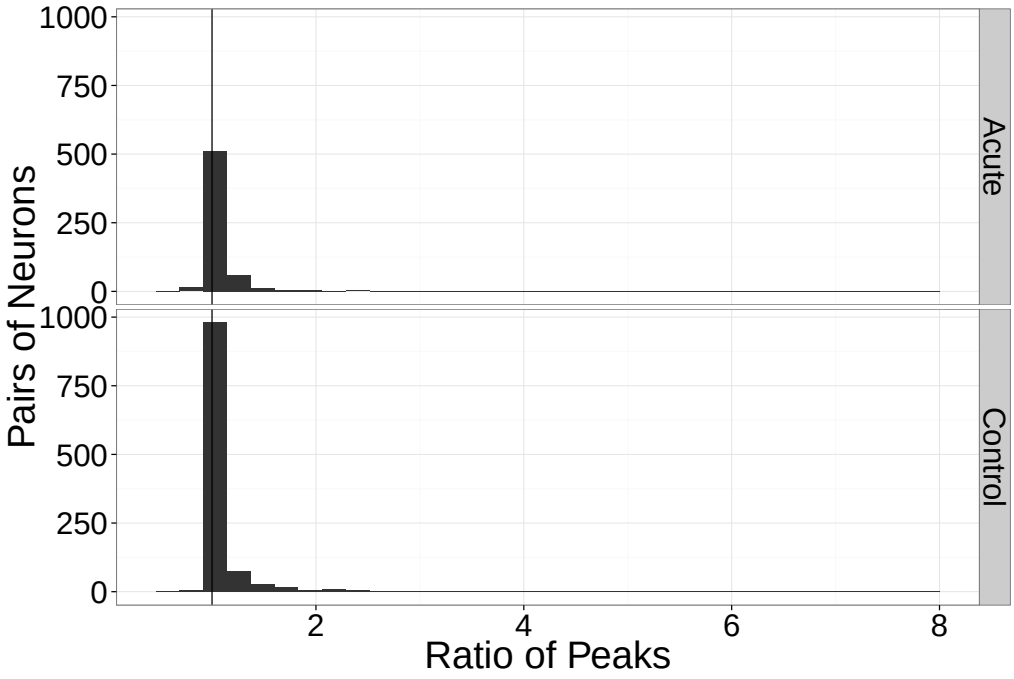

Figure 10: Ratio between observed and bootstrap peaks