Noise Correlations in Primary Motor Cortex are Modulated by Reward during a Grip Force Task

A Thesis

Presented to

the Faculty of the Department of Biomedical Engineering

University of Houston

In Partial Fulfillment

of the Requirements for the Degree

Master of Science

in Biomedical Engineering

By

Sheng Khang

May 2019

Noise Correlations in Primary Motor Cortex are Modulated by Reward during a Grip Force Task

Sheng Khang

Approved:

Chair of the Committee
Dr. Joseph T. Francis
Professor,
Dept of Biomedical Engineering
University of Houston

Committee Members:

Dr. Nuri Ince
Associate Professor,
Dept of Biomedical Engineering
University of Houston

Dr. Yingchun Zhang
Associate Professor,
Dept of Biomedical Engineering
University of Houston

Dr. Suresh K. Khator
Associate Dean,
Cullen College of Engineering
University of Houston

Dr. Metin Akay
Founding Chair and Professor
Dept of Biomedical Engineering
University of Houston

Acknowledgments

I would first like to thank my thesis advisor, Dr. Joseph T. Francis, for providing me an environment to learn and for the opportunity to have first-hand experience in performing research in the field of neuroscience. I was always in awe at the way that Dr. Francis could break down complicated subjects into little bite size pieces we lab members could comprehend. I am greatly appreciative for the humor he brought into the lab, especially if it was dark or morbid.

I would also like to thank the committee members for their advice, guidance, and excitement for student research.

I would like to thank my lab members, past and current, who have helped me along the way. Specifically, Dr. Jack Hessburg and Yao Zhao as they were the members of the lab who worked with the subjects and collected the data used in this thesis. I would like to thank my Houston lab mates who allowed me to ask them questions. So, so many questions, no matter how basic the query.

I would also like to thank my family. Their unconditional support during my time as a Master student has been invaluable and I could not have accomplished this without them.

And lastly I would like to thank my friends, new and old. Specifically, Leni Tran for her unending support and constant reminders of self-care. I would also like to thank the friend I have made during my time here in Houston, especially those in my cohort.

Noise Correlations in Primary Motor Cortex are Modulated by Reward during a Grip Force Task

An Abstract

of a

Thesis

Presented to

the Faculty of the Department of Biomedical Engineering

University of Houston

In Partial Fulfillment

of the Requirements for the Degree

Master of Science

in Biomedical engineering

by

Sheng Khang

May 2019

Abstract

Reinforcement-learning brain machine interfaces could reduce the number of times a neural prosthesis has to be updated daily for individual use. This can be achieved by using reward signals that are present in cortical neural activity but firing rates have been shown to be varied across trials, even if the same stimulus is being presented. This trial-to-trial variability, also known as noise correlation, is shared across a neuronal population and has been shown to be modulated by attention, learning, and behavior. To investigate this the current study performs a noise correlation analysis using data recorded from the primary motor cortex (M1) of two rhesus monkeys performing a grip force task (GFT). It is seen that the noise correlation generally increases with the presentation of a rewarding cue. The results also suggest a stimulus dependence as well as independence to changes in firing rates.

Table of Contents

List of Figures

Chapter 1 Introduction

One of the challenges in neuroscience is understanding the communication between a population of neurons, and how their interplay can lead to a variety of cognitive and physiological processes. One major area that has great interest is understanding how these populations respond to desirable experiences, tasks, or outcomes. A diverse body of literature suggests that midbrain structures, known as the reward system, control responses to a rewarding stimulus and are thought to use the neurotransmitter dopamine to communicate between neurons (Schultz et al., 1997; Yager et al., 2015). This system has been based on the theory of incentive salience and has been linked to learning tasks or habits and behavioral choices (Cohen and Frank, 2009). These reward-linked cognitive processes have been found in the cortical regions of the brain as well (Kapogiannis et al., 2008; Molina-Luna et al., 2009; Hosp et al., 2011; McNiel et al., 2016; Ramakrishnana et al., 2017) Since reward is an essential part of development and growth, studying this system and how it encodes for reward is valuable.

Many studies have focused on the relevant similarities or differences in the reaction to reward across species, while others aim their attention at how the excitability of neurons differs for different regions of the brain. (Schoenbaum and Setlow, 2001; Schultz, 2015) These studies have detailed how cortical structures also play a role in the reward system. It has been shown that the primary motor cortex (M1) modulates for reward (Roesch and Olson, 2004; An et al., 2018(a); Tarigoppula et al., 2018). Since M1 is shown to also encode for movement and force, it makes for a compelling case to record from there when developing a reinforcement learning brain machine interface (Marsh et al., 2015; An et al., 2018(b)).

Some studies observe the rate at which a neuron fires and have used it to gauge how excitable it is to reward (Roesch and Olson, 2004; Bayer and Gilmcher 2005). This would work in the case if neurons always fired in the same behavior, but it has been shown that neurons can have variable spiking activity in response to the repeated presentation of an identical stimulus (Tolhurst et al., 1983; Shadlen and Newsome, 1998). Other studies have shown that this varied response from one trial to the next can be shared across neurons (Gawne and Richmond, 1993; Zohary et al, 1994; Lee et al, 1998). This trial-to-trial variability can diminish the signal-to-noise ratio of a neuron population and was thought to limit the amount of information transferred in it (Zohary et al, 1994; Abbot and Dayan, 1998, Moreno-Bote et al., 2014). However, correlations have been shown to be more indicative for sensory coding, and it has been shown to be modulated by learning, behavior, and stimulus (Gutnisky and Dragoi 2008; Cohen and Maunsell, 2009; Mitchell et al., 2009; Douglas and Cohen, 2016)

In the current study, we look at M1 correlated variability of nonhuman primates (NHPs) analyzed during a grip force task and compare the results between rewarding and nonrewarding trials. The goal of our work is to show that noise correlation in M1 is modulated by reward.

Chapter 2 Method

2.1 Behavioral task

One male bonnet macaque (NHP S, Macaca Radiata) and one female rhesus macaque (NHP P, Macaca Mulatta) were trained to manually perform or passively observe a grip force task (GFT) to obtain a fruit juice reward. In this, a projection of a virtual 3-D environment consisting of a table, a cylinder target, and a simulated robotic arm (Barret WAM) is presented to the subjects. This simulation was visualized using rviz from the Robot Operating System built on a PC running Ubuntu Linux. A force transducer is situated in front of the NHPs. During manual tasks, the NHPs had to reach out with their right arm and grasp the transducer and apply the desired force. During the observation task, the NHPs would sit and watch as the desired force is delivered automatically by the simulation. Each of these tasks had a uniformly random distribution of rewarding and nonrewarding trial structure.

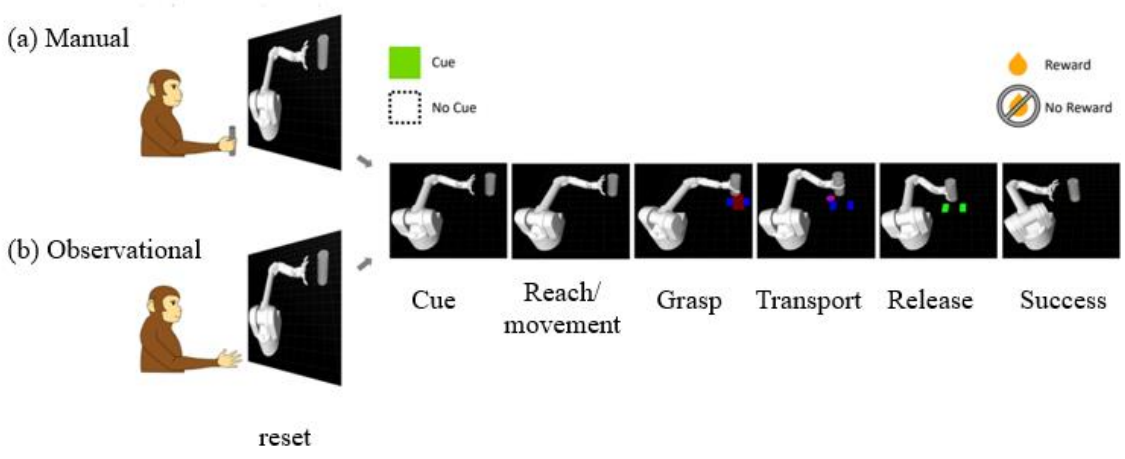

Fig. 1 Grip force task (GFT). Scene description below figures. (a) Manual task. (b) Observational task. (a) and (b) are reproduced with permission from An et al., 2018(a).

There are 7 scenes to the GFT; reset, cue presentation, reaching, grasping, transporting, releasing, and success/reward delivery. The simulated robot arm resets itself a predetermined distance away from the target in the first scene. The cue presentation scene is represented by green squares flying in from the top left side to the center of the simulated environment. The number of squares represents the amount of fruit juice reward that will be delivered after a successfully completed trial. In the case of a nonrewarding trial, no cues are displayed. Data used in this analysis only consisted of two levels of reward: either no reward or one. This cue is present throughout the trial. Reaching is automatically performed by the simulation with the end effector of the robot moving towards the cylinder target. When the robot has positioned itself just before the target, the task will move into the grasping scene. During this scene, two blue rectangles are displayed in the foreground in front of the target. They are separated a specified distance apart, and area of the squares represent the accepted range of force that the NHPs will have to apply to move to the next scene. The NHPs would then reach out and grasp the transducer in front of them. The amount of force they apply is represented as a red bar that grows between the blue rectangles. Once the bar has entered and stayed within the accepted range, the simulation moves onto the transportation scene. The simulated arm will automatically pick up and move the target to a random location on the table, so long as the desired force is still applied. After the transportation is complete, the NHPs would then have to let go of the force transducer to proceed to the releasing scene. The success scene is immediately after the release, and the reward is delivered if the trial was of the rewarding type. After the success scene, the GFT will move onto the next trial and start the scenes over.

If the NHPs were to fail a trial (either by applying too much or not enough force, thus falling out of the desired force range during grasp and transport scenes), they would have to repeat the trial at the same reward level. For example, if the NHPS were to fail a nonrewarding trial during manual trials, the next trial would be nonrewarding as well. This was implemented to help deter the NHPs from purposefully failing nonrewarding trials. Only successful trials were included in this analysis.

2.2 Surgery

After the NHPs were proficiently trained for the GFT, they were implanted with a 96-channel platinum microelectrode arrays (10 x 10 Utah array with 400 μm electrode separation, Blackrock Microsystems) in M1 in the left hemisphere, contralateral to the right hand. The implantation technique is detailed in previous work (Chhatbar et al. 2010; Marsh et al. 2015).

2.3 Neural Data Recording

After the NHPs recovered for two to three weeks, neural signals from M1 were recorded using a multichannel acquisition processor system (MAP, Plexon Inc). The signals were amplified, bandpass filtered from 170 Hz to 8 kHz, sampled at 40 kHz, and thresholded to determine possible spiking activity. The Sort Client on the MAP further sorted the spikes using principle component-based methods with waveform matching. The sorted spikes were further processed through Offline Sorter (Plexon Inc). Well isolated clusters were considered and any with less than realistic inter-spike interval histograms or with spike shapes that did not follow convention were eliminated.

2.4 Session analysis

After the final sort of the neural data offline, three manual and three observational sessions from NHP S and P each were used. NHP S manual sessions had 131, 79, and 71 units respectively while their observational sessions had 133, 128, and 81 units. NHP P manual sessions consisted of 79, 86, and 88 units while their observational sessions had 98, 125, and 70 units. Each session was analyzed individually, and the correlations are then pooled.

2.5 Neuron population response to reward

The average response of the population was calculated before movement of the simulated robot arm and around success scenes. 1000 ms before the reaching scene and 500 ms before and after the success scene were used for response analysis, and 100 ms time bins with a sliding window of 5 ms were used to calculate the response. This meant that there was a 95% overlap between a time bin and those adjacent to it. The firing rate, *R*, for bin number *b* of unit *n* during trial *k* is found in with

$$R_n^k(b) = \frac{S_n^k(b)}{T}, \tag{1}$$

where *S* is the spike count in the bin and *T* is the size of the bin in ms. The binned firing rates are then averaged across all *K* trials with

$$\overline{R_n(b)} = \frac{\sum_{k=1}^{K} R_n^k(b)}{K}, \tag{2}$$

to find the average response of a unit during the recording session. The average unit responses are then averaged to find the population response using

$$\overline{R(b)} = \frac{\sum_{n=1}^{N} \overline{R_n(b)}}{N}, \tag{3}$$

where *N* is the total number of units present in the recording session. This is depicted in Fig. 2 for an example unit during cue period for rewarding trials. Fig. 2(c) shows the average response across all trials.

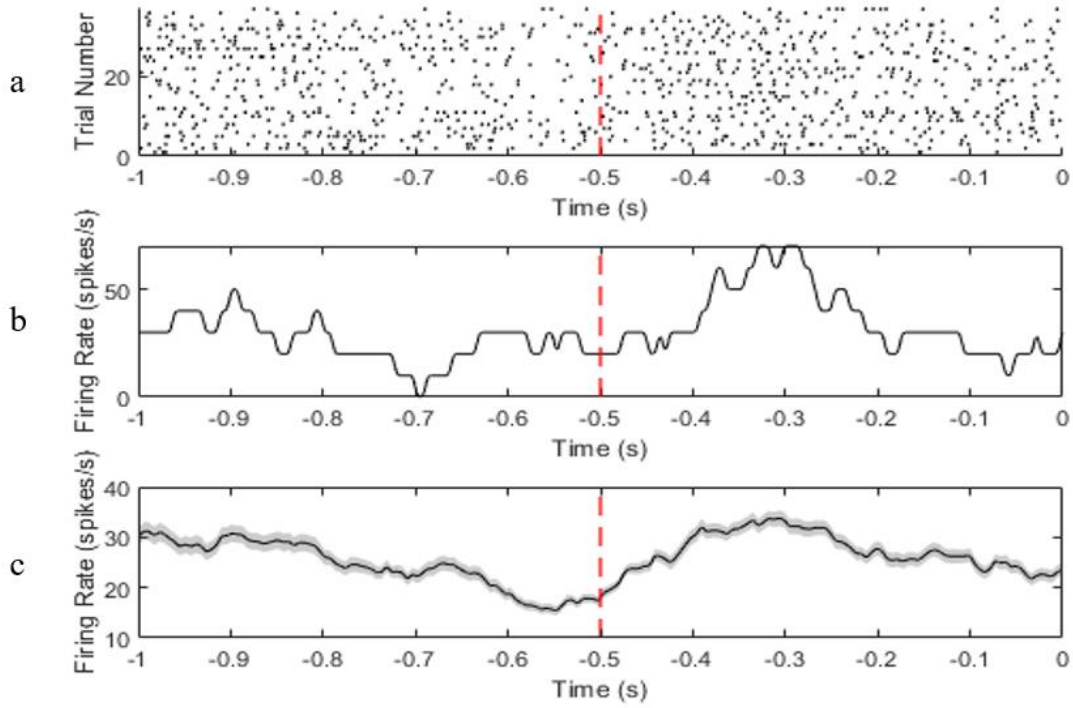

Fig. 2 Example of a unit's average response peri-cue for rewarding trials. Red dashed line represents the time when the cue is presented. (a) Raster of unit spike times. (b) Response from an example trial. (c) Average response across trials.

2.6 Noise correlation

The spike times of each unit were extracted and separated into spike trains. Since rewarding trials will have the cue moving in at the beginning of the cue scene, the analysis takes 500 ms before this as pre-cue period and 500 ms after the cue has fully arrived as post-cue period. For nonrewarding trials, 500 ms before and after the cue scene is used. For the success scene, 500 ms before and after are considered as pre-reward and post-reward periods. Even though nonrewarding trials do not have reward delivered, the periods are still called the same.

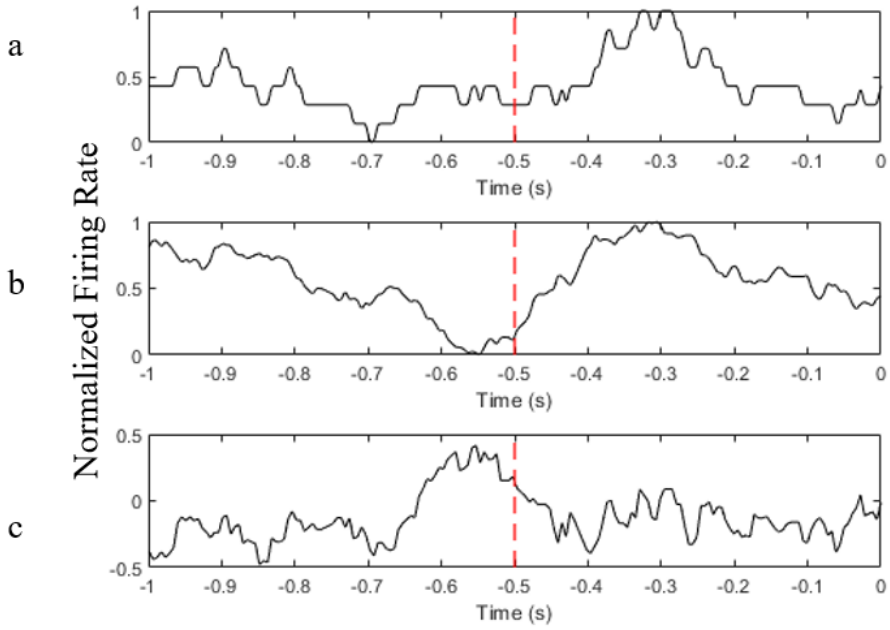

Fig. 3 Visual example of how the noise of a unit's response for a given trial is calculated. (a) Unit response for a trial. (b) Average of the unit response. (c) The average response is subtracted from the unit response to find the noise.

To find the noise, we subtracted the mean of the normalized response of a unit during either a rewarding or a nonrewarding trial of a session from that unit's normalized response during that trial. An example of this is illustrated in Fig. 3 where (c) is the noise associated with the example unit during the example trial. Mathematically, the spike trains from each session was first binned as previously described in section 2.5 using Eq. 1. Each bin was then Min-Max normalized by

$$R_{norm_n}^{\ k}(b) = \frac{R_n^k(b) - R_{min_n}}{R_{max_n} - R_{min_n}}, \tag{4}$$

where R_{min_n} is the minimum value for the bin across the recording session and $\mathrm{R_{max_n}}$ is the maximum. Then, the mean response for each unit across trial types was then calculated using

$$\overline{R_{norm_n}(b)} = \frac{\sum_{k=1}^{K} R_{norm_n}^{k}(b)}{K}, \quad (5)$$

where *K* is the total number of trials for the given trial type. The noise of a unit during any given trial was determined by subtracting the mean response of the unit for that trial type from the response during that trial. This is expressed as

$$R_{noise_n}^{k}(b) = R_{norm_n}^{k}(b) - \overline{R_{norm_n}(b)}\,. \quad (6)$$

where $R_{noise_n}^{\boldsymbol{k}}$ represents all the bins *b* in trial *k* for unit *n*. It is then split into the different trial periods. The total noise of a neuron during a trial period for a recording session was formed by concatenated back to back the noise across all the trials. For example, neuron n would be made with

$$R_{noise_n} = [R_{noise_n}^{1}; R_{noise_n}^{2}; R_{noise_n}^{3}; \ldots; R_{noise_n}^{m}]. \quad (7)$$

The noise correlation between a pair of neurons is the Pearson correlation coefficient between their trial-to-trial concatenated noises. This is accomplished with

$$r_{noise_{(n1,n2)}} = \frac{cov(R_{noise_{n1}}, R_{noise_{n2}})}{\sigma_{noise_{n1}} * \sigma_{noise_{n2}}}. \quad (8)$$

This was performed by using MATLAB's *corr( )* function. The *p-values* were assessed to determine if the correlation coefficients between the unit pairs were significantly different than zero, and coefficients found to be significant were recorded. The sessions were then checked for unit pairs that were present in both rewarding and nonrewarding trials. Those that were not were eliminated.

Taking the average correlation coefficients across the population would underestimate the actual correlation coefficient (Silver and Dunlap, 1987). To get a closer

estimation of the mean correlation coefficient, the r_{noise} values were first Fisher z-transformed using

$$z_{noise} = arctanh(r_{noise}) = \frac{1}{2} log\left(\frac{1+r_{noise(n1,n2)}}{1-r_{noise(n1,n2)}}\right). \quad (9)$$

The average of z_{noise} was found, and this new value was inverse transformed it back to an r_{noise} value using the hyperbolic tangent function

$$\overline{r_{noise}} = \tanh(\overline{z_{noise}}) = \frac{e^{2z_{noise}}-1}{e^{2z_{noise}}+1}. \quad (10)$$

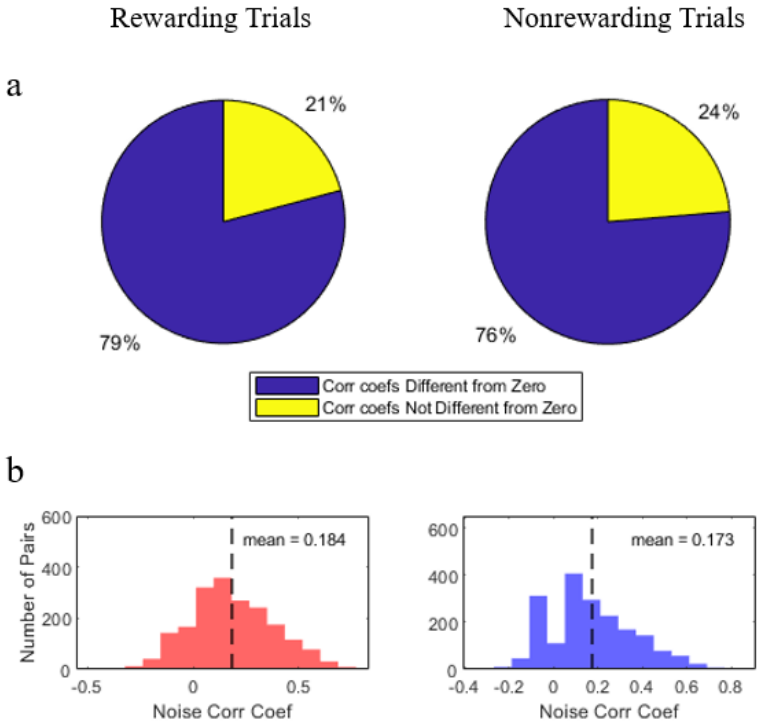

Fig. 4 Example of significant correlations from an NHP S manual session. (a) Population of noise correlations in a population that are statistically different from zero. (b) Finding the mean noise correlation of the blue percentage from (a)

2.7 Firing rate vs noise correlation

It has been reported that an increase in firing rates could cause an increase in correlations in unit pairs, and the presentation of a stimulus has been shown to increase

firing rates in NHPs. This is particularly important in noise correlation analysis as different stimuli can elicit responses of varying strength which can affect the correlational structure. One way to relate a unit pair's correlation coefficient to their average firing rates is to take the geometric mean rate (GMR). This was chosen as it was shown that the GMR estimated correlations better than the arithmetic mean (Schultz et al., 2015). Post-cue and post-reward periods are used for this section. To find the GMR, we take the results from Eq. 2, the average response of a unit across trials, and find the mean firing rate of the unit by using

$$\overline{R_n} = \frac{\sum_b^B \overline{R_n(b)}}{B}, \tag{11}$$

where *B* is the total number of bins. The GMR between two units is determined by

$$GMR_{n1,n2} = \sqrt{\overline{R_{n1}} * \overline{R_{n2}}}. \tag{12}$$

The GMR of a pair of units during a trial period is then plotted against their noise correlation. All unit pairs from the recording sessions of an NHP were combined and separated into 10 bins, each having an equal number of pairs. The average GMR and noise correlation are plotted for each of the bins.

2.8 Significantly modulating units

The analysis of section 2.6 was repeated for units who significantly modulated their average firing rate across trial types and then for those within trial periods. Unit activity has been shown be different between rewarding and nonrewarding trials (Marsh et al., 2015), so it is worth observing units that significantly change their rates across cue and reward period during rewarding trials. The goal of this section was to observe how the noise correlation changes if only units that showed a significantly difference in their firing rates between the lack of and presentation of a stimulus (cue or reward) were considered.

Post-cue and post-reward periods of rewarding and nonrewarding trials were used to determine significant difference in this section.

The unit activity was binned, and Min-Max normalized as in section 2.5 and 2.6. Then the mean response of the unit was determined. The mean responses were separated into trial periods and tested between post-cue periods and then between post-reward periods for a change in response across rewarding and nonrewarding trials. The Wilcoxon signed-rank test was used to determine the significant change. This test was used as firing rate distribution was not assumed to have any distribution, and because responses from the same units were being compared. This test was performed with MATLAB's *signrank( )* function. Units with a $p\text{-value} < 0.05$ was determined as significantly modulating across trial types. Those that did not meet this were eliminated from the analysis. Since the units are tested for cue period and then for reward period, they could be in both peri-cue and peri-reward noise correlation analysis, in one or the other, or in neither.

After units were found to be significantly modulating, the same analysis in section 2.6 was performed on them; that is subtracting mean response from individual trial responses, separate and concatenate trial periods of similar trial type, finding the noise correlation of the neuron population, Fisher z-transformation, averaging, and then inverse transforming back to correlation value.

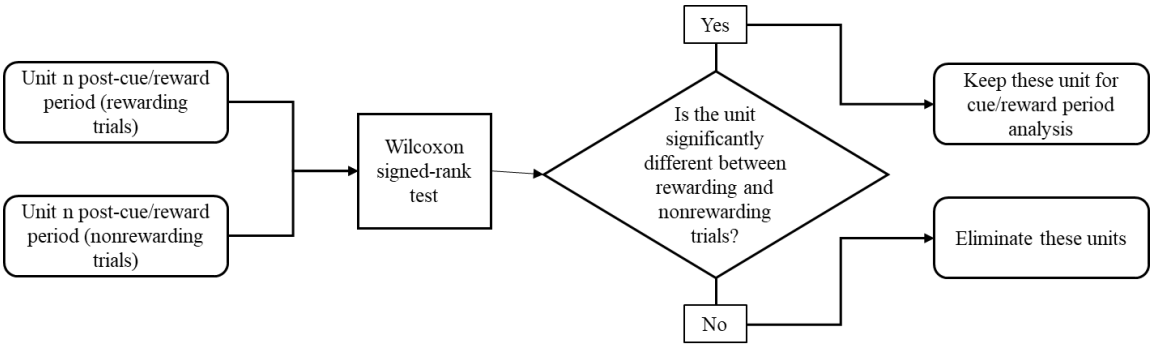

Fig. 5 Work flow for eliminating units that do not modulate their activity across trial types. Unit N would be checked separately for (a) cue and (b) reward periods.

2.9 Mean-matched noise correlation

The results of section 2.8 could explain the relationship between firing rates and noise correlation. Although, other studies have shown that the noise correlations can be stimulus dependent as well as being firing rate independent (Kohn and Smith, 2005, Banyai et al., 2018). The goal of this section is to show that even though noise correlation and firing rates are not entirely independent of one another, changes in noise correlation are not trivially explained by dissimilarities in firing rate.

One way to accomplish this would be to control the firing rates of the pairs of neurons, but this would require controlling evoked responses by first cataloging the responses to stimuli used. As this is excessive and difficult to do the GMR distributions of rewarding and nonrewarding trials were equated instead. First, the GMR of unit pairs having significant noise correlations found with Eq. 11 are separated into trial periods. The trial periods were compared between rewarding and nonrewarding and the one with the maximum value was used to determine the behavior of a histogram used to plot the distributions, 40 bins equally spaced up to the maximum values found. Rewarding and nonrewarding distributions were plotted together and the area of overlap between them was

taken down as a new, marginal distribution. The rewarding and nonrewarding distributions are subsampled to match this marginal distribution, and the correlation coefficient of the subsampled data is used to complete the noise correlation analysis.

Chapter 3 Results

3.1 Population response to the presence or lack of cue and reward stimuli

To understand how the population of units respond to stimuli, the average response is first observed. The units were binned and averaged across the population for all successful trials. They are plotted along with their standard error of the mean (SEM), which is shaded in the figures below. Each session is displayed separately for both the peri-cue and peri-reward time periods.

Fig. 6 – 9 show the average population response of example sessions from NHP S and P for rewarding, red line, and nonrewarding trials, blue line. The other sessions population responses are in the appendix, Fig. A.1-A.4. The shaded area represents the standard error of the mean (SEM). The plots in the left column are for peri-cue period and the right for peri-reward. The vertical dashed black line represents when a stimulus was, or would have been, delivered during that period.

NHP S follows the general trend of previously described subjects. This is that for a preferred cue there is an increase in the firing rate. This is seen in all peri-cue plots for all manual and observational sessions for the GFT of NHP S. Average responses show that there is a trend similar for the peri-reward period across all manual and observational sessions as well. Firing rates during this period tended to decrease for rewarding trials when compared to nonrewarding ones.

On the other hand, NHP P's average population responses are varied. The average response in the cue periods across the manual and observation sessions is distinguishable by the slight increase in firing rates before 500 ms. The rewarding periods are like that of NHP S for manual sessions, as firing rates are generally lower for rewarding trials.

The differences that are observed in these trial periods show that population modulation for reward occurs in M1 during the GFT during cueing and reward periods. This is especially true for the manual task. The interactions between the neurons is further analyzed in the sections below.

Average population response for NHP S Manual GFT

(a) Peri-cue period (pre-movement) *(b) Peri-reward period (success)*

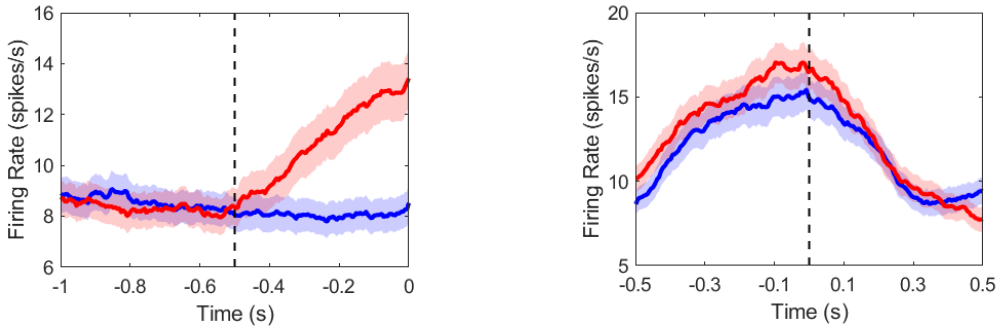

Fig. 6 Average population response for a manual session for NHP S. Red for rewarding and blue for nonrewarding. Shaded region represents the SEM. (a) Peri-cue period. (b) Peri-reward period

Average population response for NHP S Observational GFT

(a) Peri-cue period (pre-movement) *(b) Peri-reward period (success)*

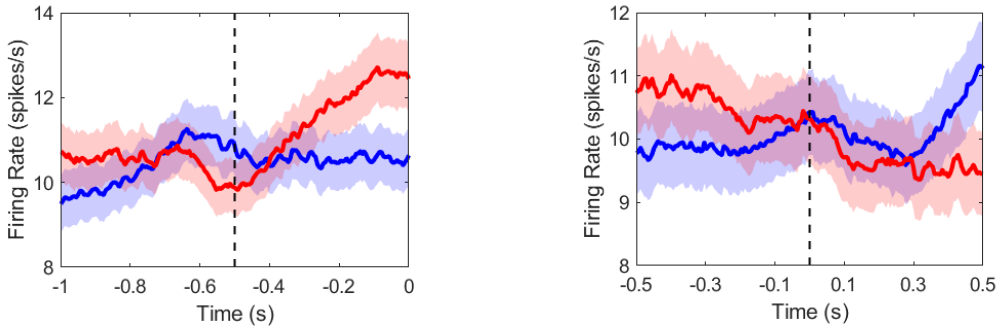

Fig. 7 Average population response for an observational session for NHP S. Red for rewarding and blue for nonrewarding. Shaded region represents the SEM. (a) Peri-cue period. (b) Peri-reward period

Average population response for NHP P Manual GFT

(a) Peri-cue period (pre-movement) *(b) Peri-reward period (success)*

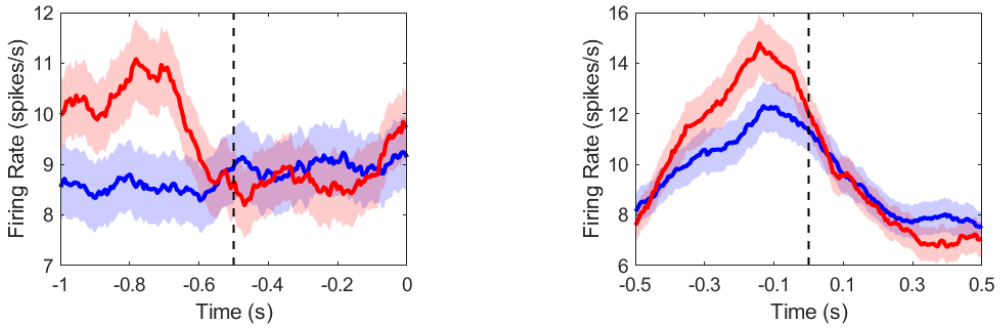

Fig. 8 Average population response for a manual session for NHP P. Red for rewarding and blue for nonrewarding. Shaded region represents the SEM. (a) Peri-cue period. (b) Peri-reward period

Average population response for NHP P Observational GFT

(a) Peri-cue period (pre-movement) *(b) Peri-reward period (success)*

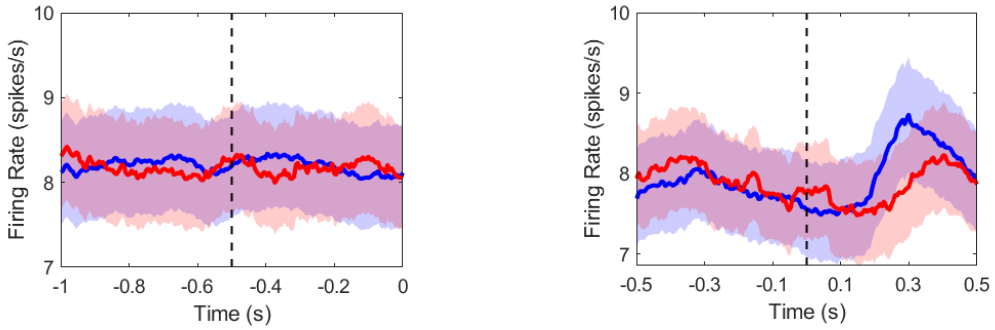

Fig. 9 Average population response for an observational session for NHP P. Red for rewarding and blue for nonrewarding. Shaded region represents the SEM. (a) Peri-cue period. (b) Peri-reward period

3.2 Average noise correlations are higher for rewarding trials

The binned data is Min-Max normalized, and the mean response of a unit is subtracted from their individual trial responses to extract the noise. These resulting trial noise responses are concatenated from trial-to-trial and repeated for all units. The concatenated noises are compared using Pearson's correlation coefficient and significant

unit-pairs are determined and recorded. After Fisher z-transforming the significant unit pair correlations the average value is found and then reverse Fisher z-transformed back to correlation. The average correlation coefficients for pooled from all sessions are plotted along with the SEM, Fig. 10 and 13. Red bars represent rewarding trials and blue for nonrewarding. The significant difference between average coefficients of rewarding and nonrewarding trial periods was tested with Wilcoxon signed rank test and the results of this test is displayed above the trial period. One asterisk represents a signed rank test with *p-value* < 0.05, two for *p-value* < 0.01, and three asterisks for *p-value* < 0.001. These represents the noise correlation distributions that are found in Fig. 11, 12 and 14, 15. In these figures, the color of the asterisk represents which distribution is higher.

The noise correlation for both NHPs during manual and observational tasks show that average correlation value varies in strength and significance across trial periods and trial types for all sessions. The average coefficients generally fell into an expected range, 0.01 – 0.3 as reported by other work (Cohen & Kohn, 2011). One trial period that is seen to have more consistency is post-cue. For most manual and observational sessions, rewarding trials had a higher or significantly higher average correlation coefficient compared to nonrewarding trials.

The NHPs do differ in the post-reward period. NHP S tended to have lower correlation for rewarding trials compared to nonrewarding ones, and the opposite is true for NHP P. This may be explained by behavioral differences between the subjects, and more on this is presented in the discussion section.

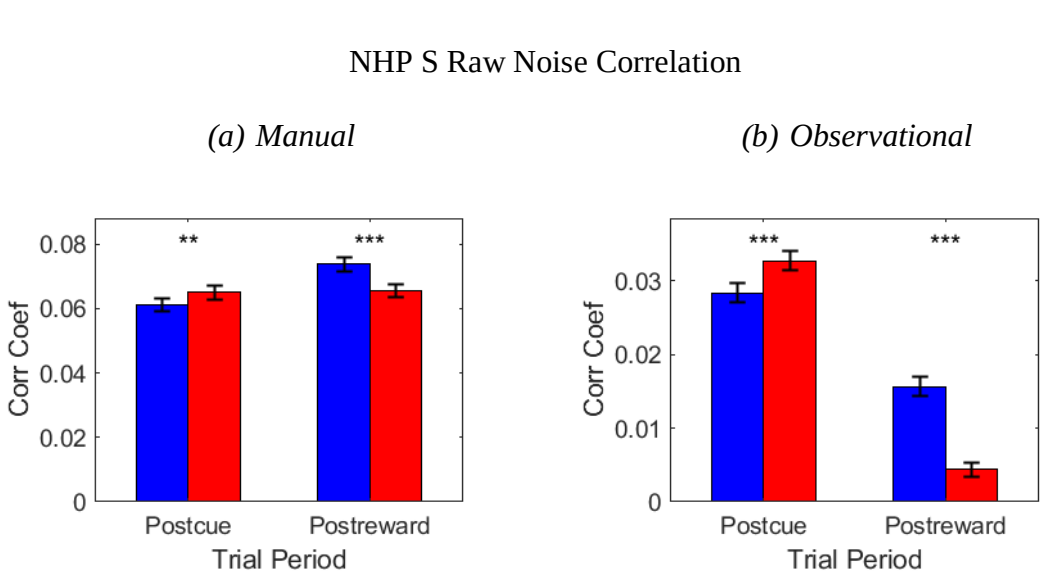

Fig. 10 Average noise correlation coefficient for NHP S manual and observational sessions separated by trial periods. Red for rewarding and blue for nonrewarding. (a) Pooled manual sessions. (b) Pooled observational sessions.

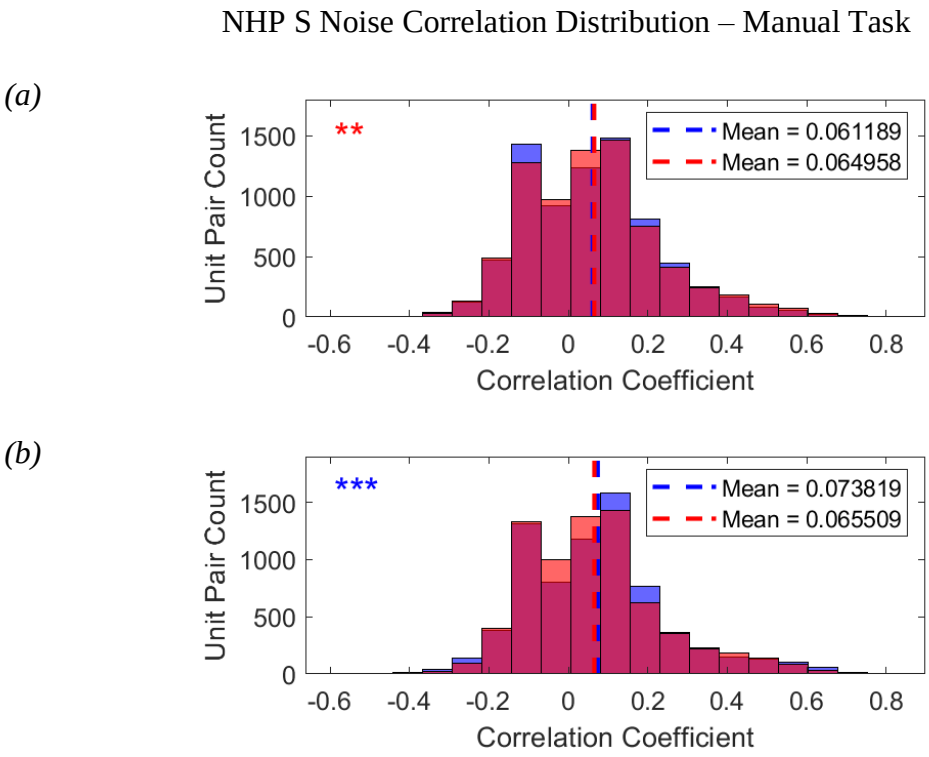

Fig. 11 Noise correlation distribution of NHP S during manual task. Red represents rewarding trial distribution and blue for nonrewarding. (a) depicts distribution during the post-cue period and (b) for post-reward period.

NHP S Noise Correlation Distribution – Observational Task

(a)

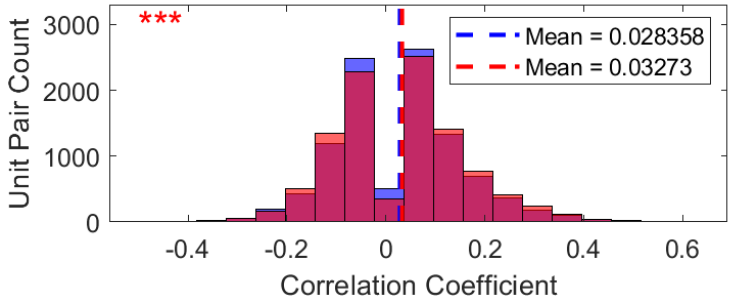

(b)

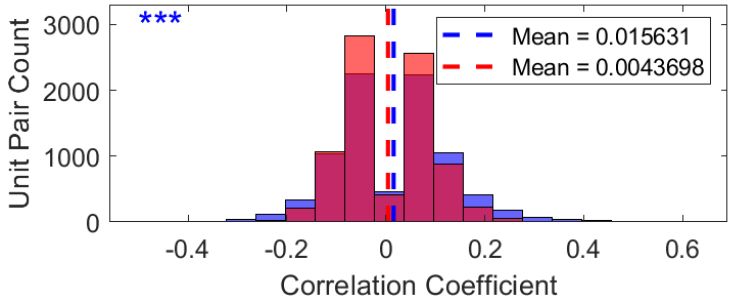

Fig. 12 Noise correlation distribution of NHP S during observational task. Red represents rewarding trial distribution and blue for nonrewarding. (a) depicts distribution during the post-cue period and (b) for post-reward period.

NHP P Raw Noise Correlation

(a) Manual *(b) Observational*

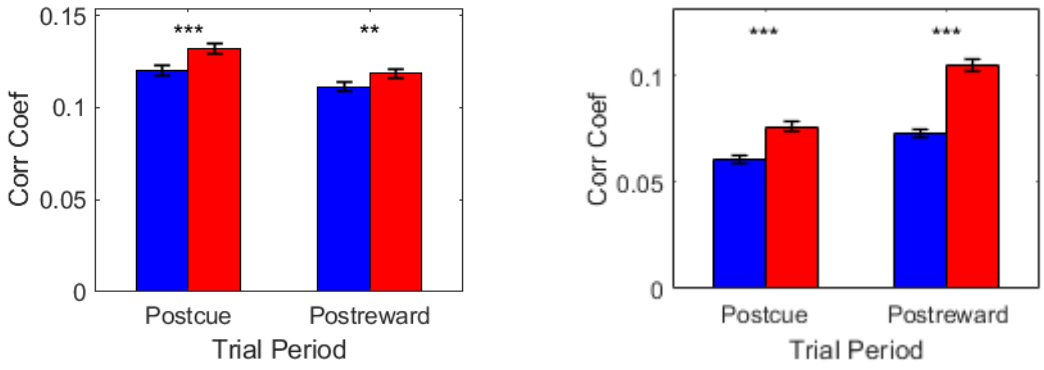

Fig. 13 Average noise correlation coefficient for NHP P manual and observational sessions separated by trial periods. Red for rewarding and blue for nonrewarding. (a) Pooled manual sessions. (b) Pooled observational sessions.

NHP P Noise Correlation Distribution – Manual Task

(a)

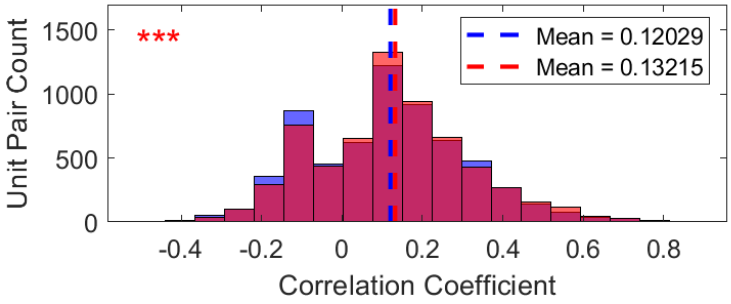

(b)

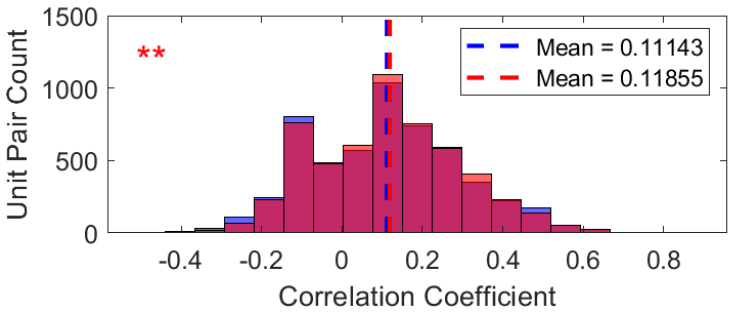

Fig. 14 Noise correlation distribution of NHP P during manual task. Red represents rewarding trial distribution and blue for nonrewarding. (a) depicts distribution during the post-cue period and (b) for post-reward period.

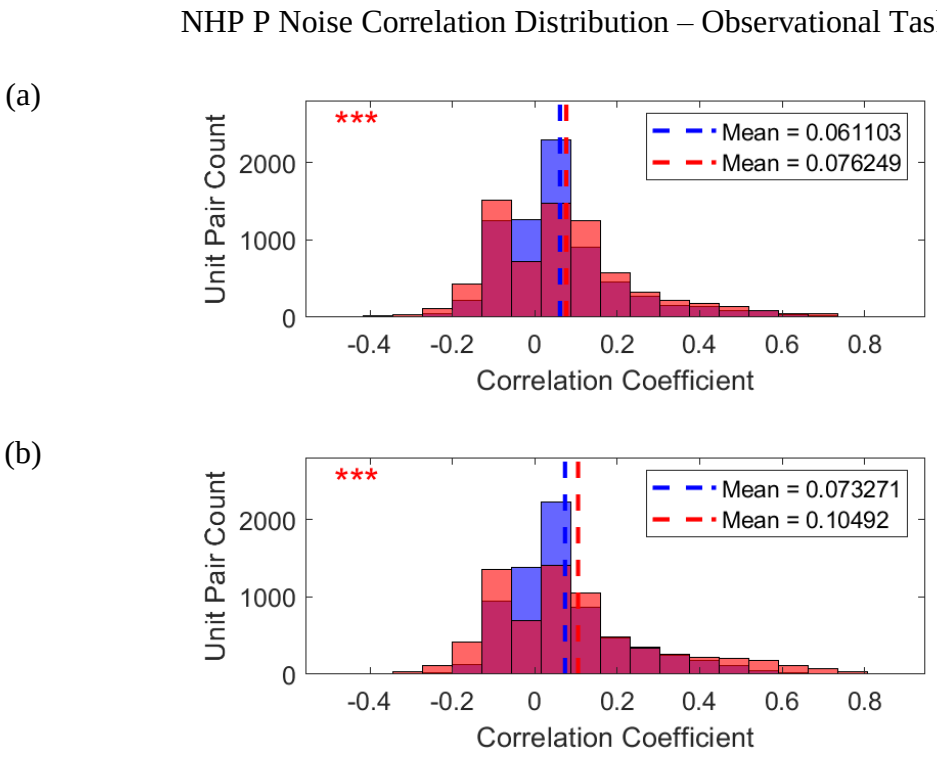

Fig. 15 Noise correlation distribution of NHP S during observational task. Red represents rewarding trial distribution and blue for nonrewarding. (a) depicts distribution during the post-cue period and (b) for post-reward period.

3.3 Firing rates and Correlations

To visually observe what relations firing rates and correlation have, they are plotted together. First, the significant correlations for post-cue and post-reward periods of rewarding and nonrewarding trials are plotted against their GMR. Then, these points are split into ten bins with an equal number of markers. The GMR and the output correlations are averaged and then the averages are plotted with their SEM. The difference between rewarding and nonrewarding plots were tested with Wilcoxon signed rank test to determine which distribution was significantly higher. Red lines are for the rewarding distribution and the blue lines are for the nonrewarding. Asterisks are assigned like that of section 3.2.

The relationship between unit pairs GMR and their correlations differ slightly between the NHPs. Post-cue periods for the NHPs were observed to have significantly higher distribution for rewarding trials compared to nonrewarding, except for NHP S's observation sessions. Post-reward period results are more varied with some sessions showing a higher distribution for nonrewarding and others showing no significant differences.

There seems to be a linear relationship displayed in some of the figures below. NHP S's observational post-cue period and all NHP P's trial period plots show a positive linear relationship between GMR and the resulting correlation coefficients. NHP S's manual sessions show a negative linear relationship.

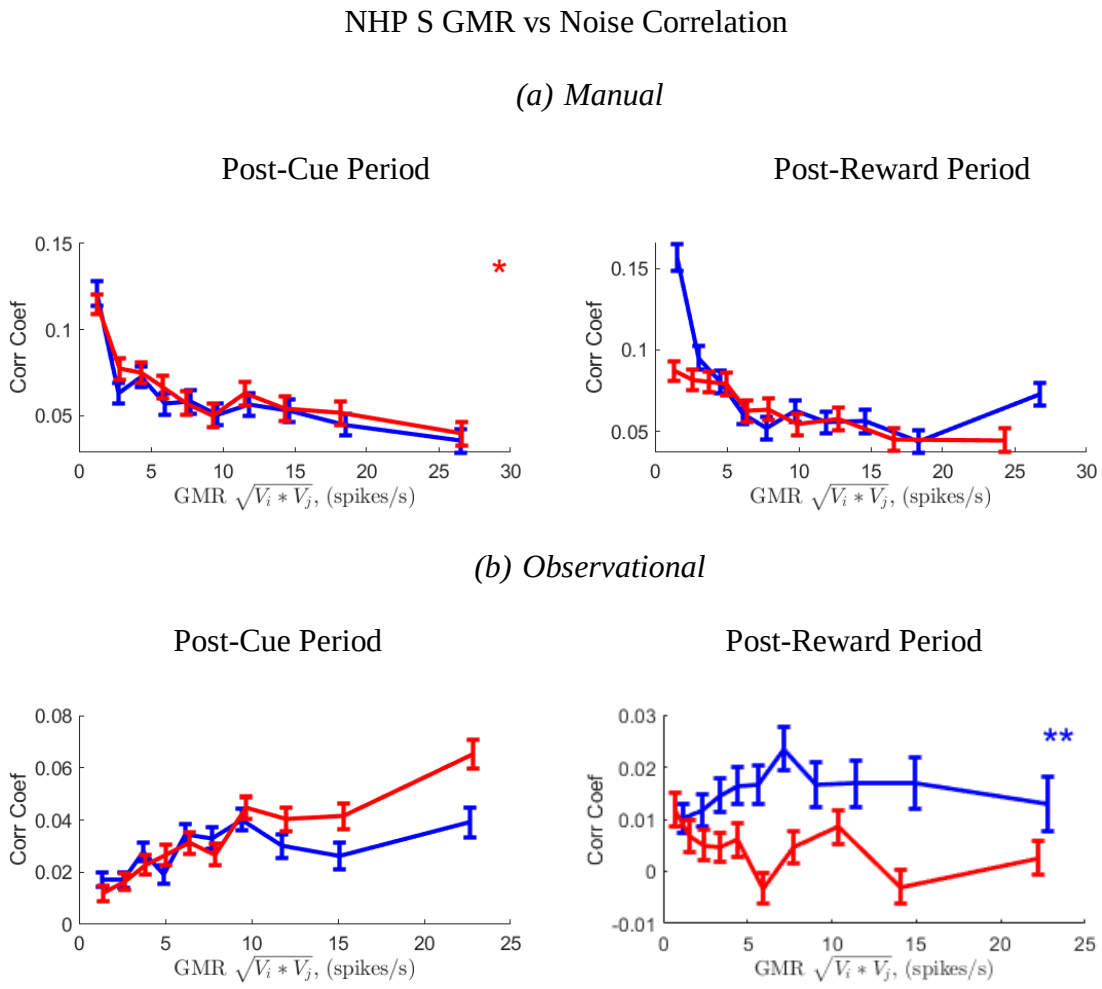

Fig. 16 GMR, $\sqrt{V_i * V_j}$, vs average correlation coefficient for NHP S. The subplots on the left represents post-cue period and the ones on the right for post-reward. (a) Manual sessions. (b) Observational sessions

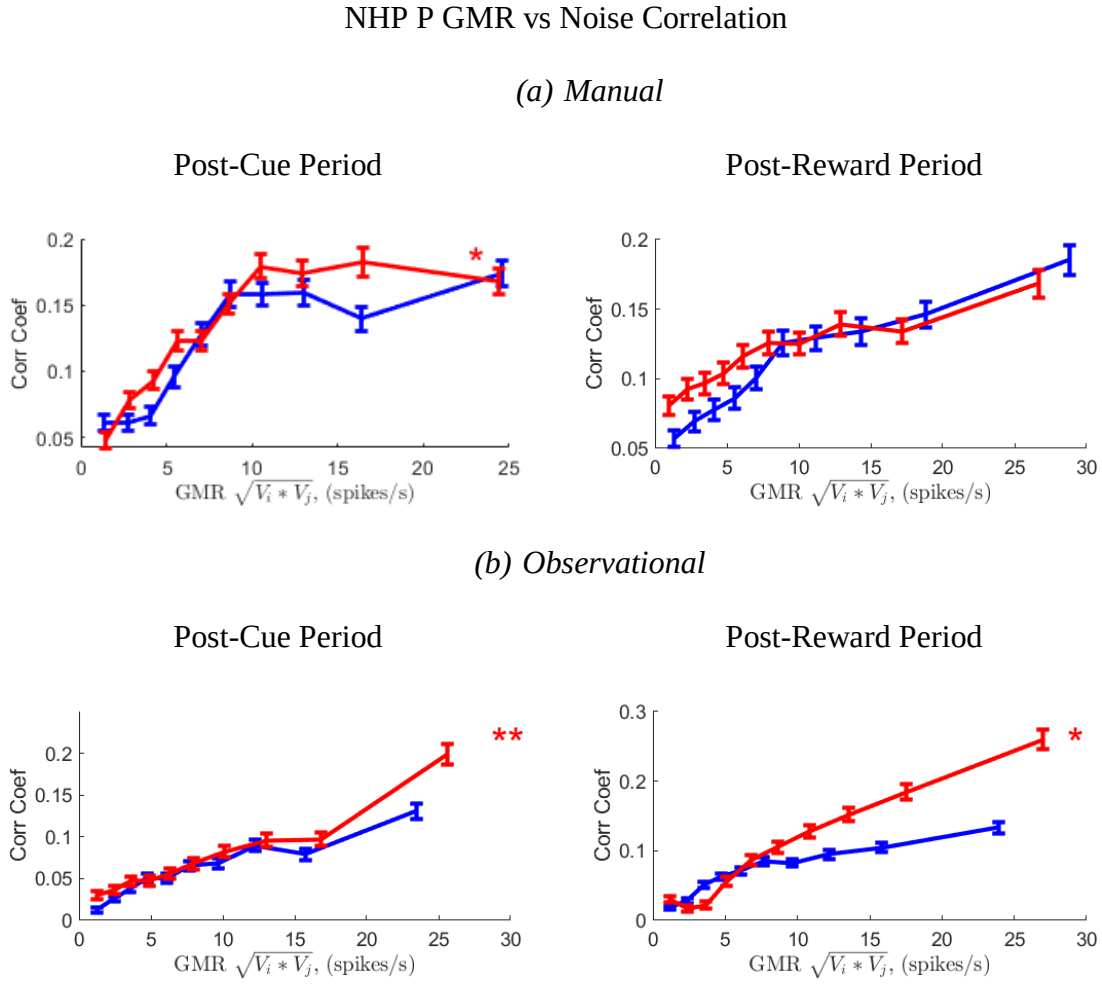

Fig. 17 GMR, $\sqrt{V_i * V_j}$, vs average correlation coefficient for NHP P. The subplots on the left represents post-cue period and the ones on the right for post-reward. (a) Manual sessions. (b) Observational sessions

3.4 Significantly modulating units

To observe how firing rate modulation across trial types effect the noise correlation, units are separated into different subpopulations, those that significantly changed their firing rates between rewarding and nonrewarding trials for post-cue/reward periods and those that do not significantly change. Significant differences were checked with Wilcoxon signed-rank test. The percentage of units for each subpopulation is displayed in Fig. 18 and

19 with the total number of units in the pooled sessions above the pie charts. The blue portion of the pie chart represents units who have modulated their firing rates, and the portion in yellow are those that did not change. Those that do not modulate for reward were eliminated from this analysis. The rest were analyzed in the same manner as previously described.

We see a difference for NHP S's manual sessions, where the noise correlation distribution is no longer significantly different between rewarding and nonrewarding trials. The average across the other NHP sessions show that for post-cue period rewarding trials had higher noise correlations than nonrewarding trials, marked by the asterisks above the bars for the trial period. NHP S and P still differed in post-reward periods as S had lower correlations in rewarding trials while P had higher when compared to nonrewarding trials. These results tell us that only considering firing rates as a metric for reward would result in

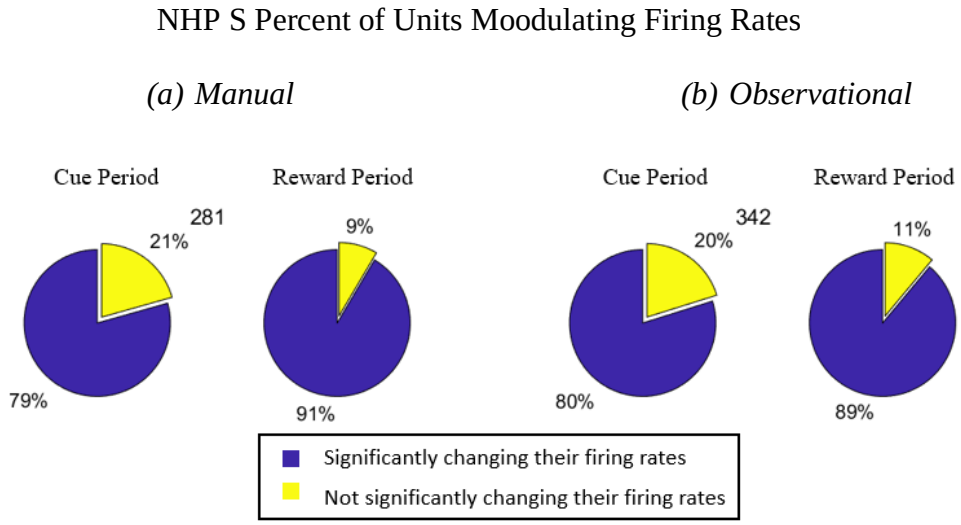

Fig. 18 Percent of units that have modulated their mean firing rates across rewarding and nonrewarding trials for NHP S pooled from sessions. Number in the top right is total number of units. (a) Manual sessions. (b) Observation sessions.

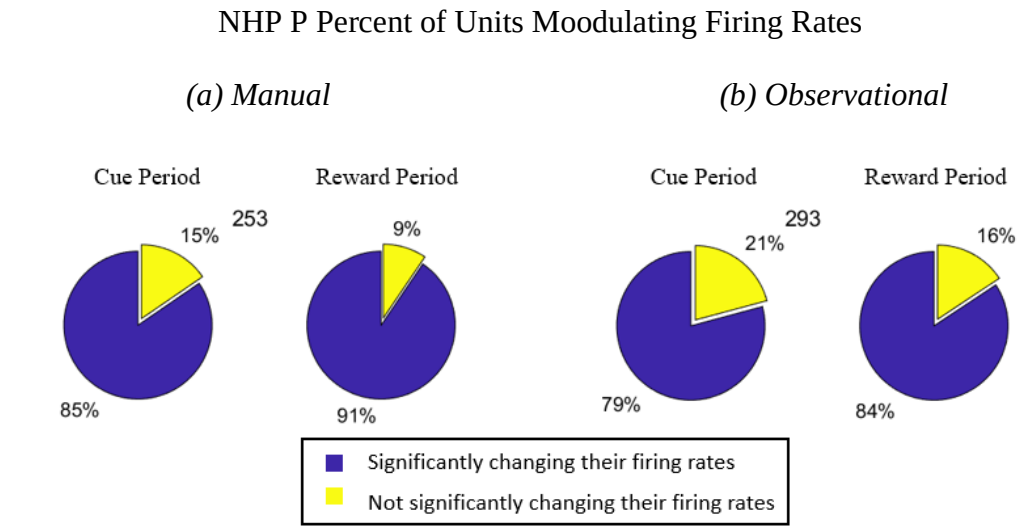

Fig. 19 Percent of units that have modulated their mean firing rates across rewarding and nonrewarding trials for NHP P pooled from sessions. Number in the top right is total number of units. (a) Manual sessions. (b) Observation sessions.

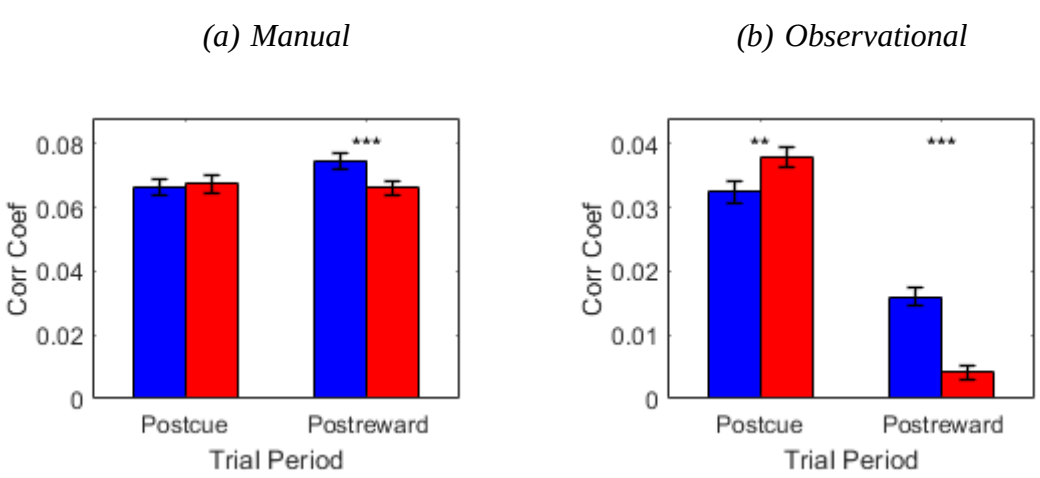

Fig. 20 Pooled average noise correlation for NHP S separated by trial periods. Red for rewarding and blue for nonrewarding. (a) Manual sessions. (b) Observation sessions.

NHP P Noise Correlation of Signiicant Units

(a) *Manual* ***(b)*** *Observational*

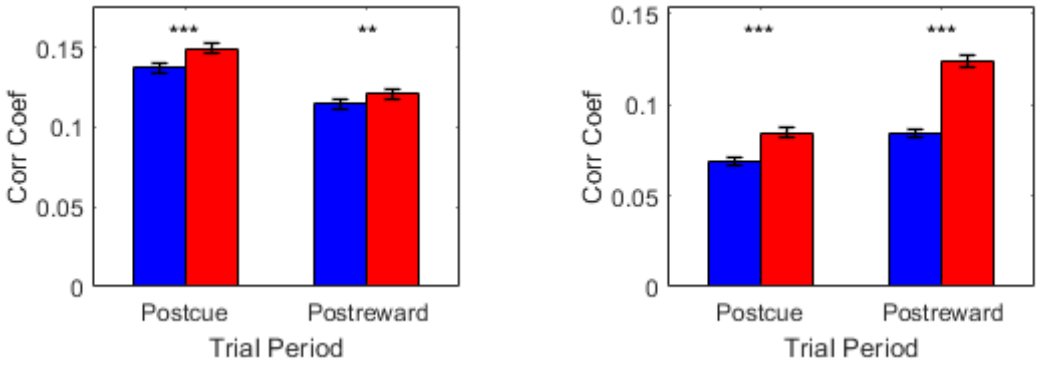

Fig. 21 Pooled average noise correlation for NHP P separated by trial periods. Red for rewarding and blue for nonrewarding. (a) Manual sessions. (b) Observation sessions.

Another way we can look at this is by analyzing those unit pairs that we have eliminated, the yellow portion of the pie charts in Fig. 18 and 19. Again, this gives us an understanding of how correlations are affected if firing rates are the only factor used to look for rewarding signals. To accomplish this, the units from the population who did not significantly change their firing rates were passed through the noise correlation analysis. The results of which are shown in Fig. 22 and 23.

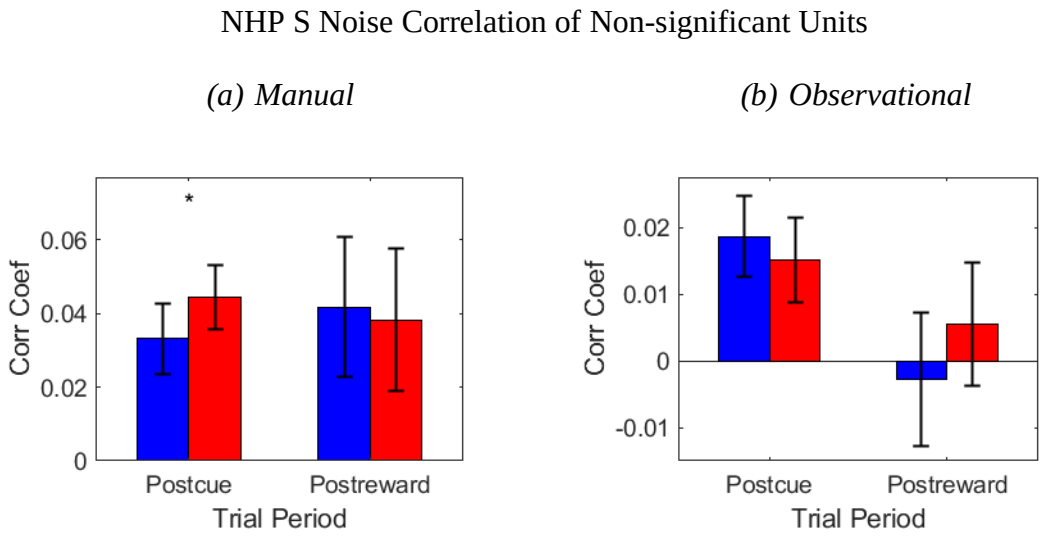

Fig. 22 Noise correlation distribution of non-significantly modulating units from NHP S. Red is for rewarding trials and blue for nonrewarding. (a) is the manual task and (b) is observational task.

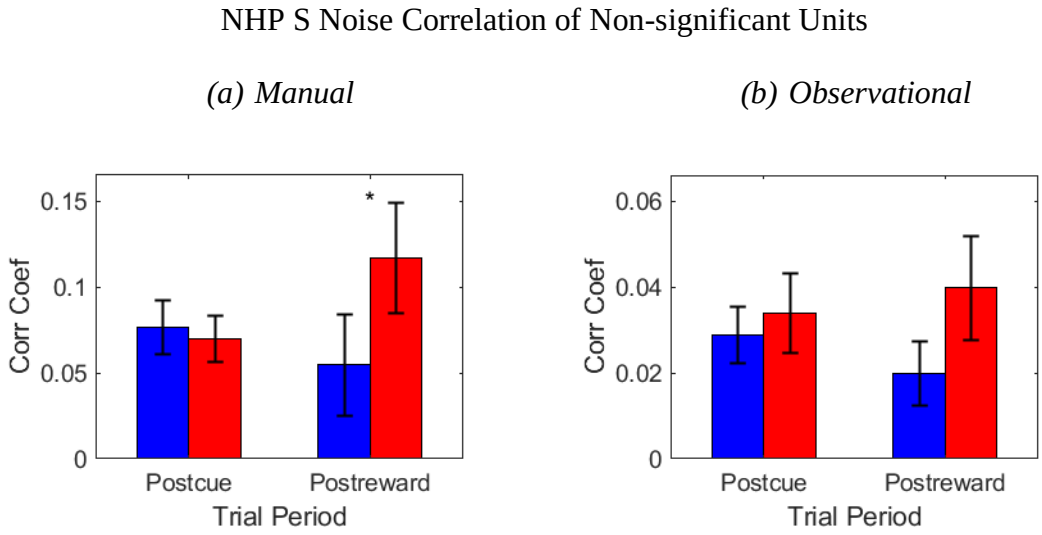

Fig. 23 Noise correlation distribution of Non-significantly modulating units from NHP P. Red is for rewarding trials and blue for nonrewarding. (a) is the manual task and (b) is observational task.

3.5 Noise correlations are not easily explained by changes in firing rates.

To show that the noise correlations are stimulus dependent while being independent of changes in firing rates, the distributions of GMR are matched. This is accomplished by plotting the two distributions on top of one another, showing how different they can be. Then the larger distribution is subsampled to match that of the smaller distribution, resulting in a new marginal distribution. This is illustrated in Fig. 24 through 27. Their (a) represents the actual firing rate distributions and (b) the mean matched ones. As in previous sections, the number of asterisks represents the level of difference among the two distributions. The significance was tested with the Wilcoxon rank sum test. This was chosen because the no assumptions were made of the GMR distributions, and the same unit pairs were not guaranteed to be present in the mean matched distributions. The unit pairs that are present in the new distributions were matched to their noise correlation and the noise correlation distribution was recorded in Fig. 28 and 29. The same trends seen previous are present here as well. That being

NHP S Manual – Mean Matching of Firing Rate Distribution

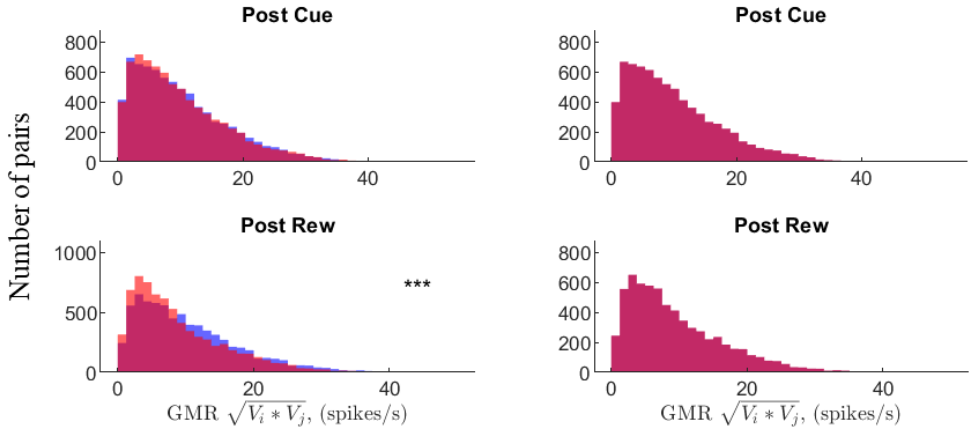

(a) Raw Firing Rate Distribution *(b) Mean Matched Firing Rate Distribution*

Fig. 24 Pooled raw and mean matched firing rate distributions for NHP S manual sessions. Red for rewarding trials distribution, blue for nonrewarding trials distribution.

NHP S Observational – Mean Matching of Firing Rate Distribution

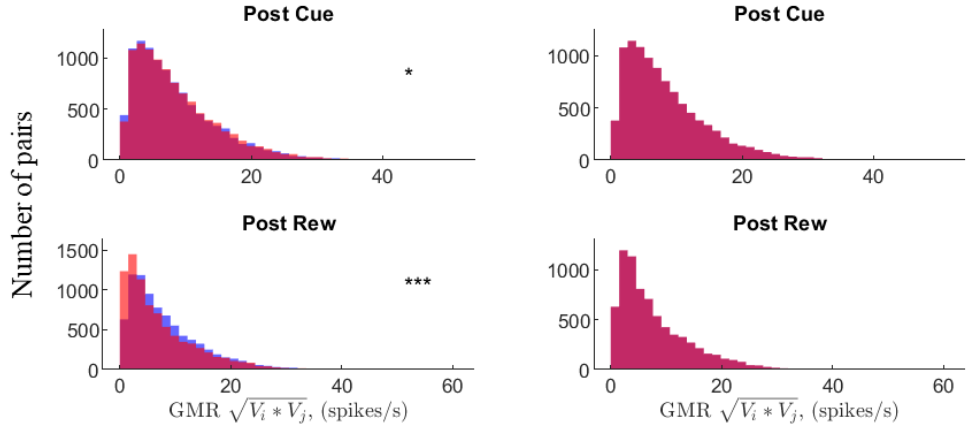

(a) Raw Firing Rate Distribution *(b) Mean Matched Firing Rate Distribution*

Fig. 25 Pooled raw and mean matched firing rate distributions for NHP S observational sessions. Red for rewarding trials distribution, blue for nonrewarding trials distribution.

NHP P Manual – Mean Matching of Firing Rate Distribution

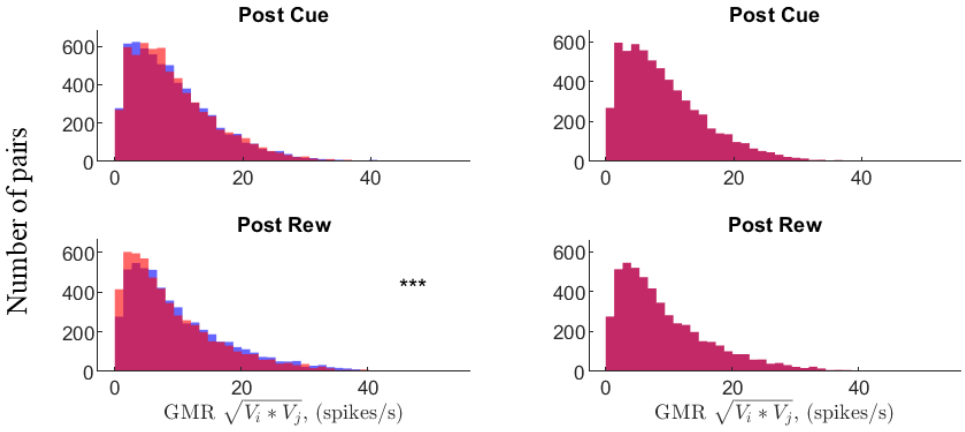

(a) Raw Firing Rate Distribution *(b) Mean Matched Firing Rate Distribution*

Fig. 26 Pooled raw and mean matched firing rate distributions for NHP P manual sessions. Red for rewarding trials distribution, blue for nonrewarding trials distribution.

NHP P Observational – Mean Matching of Firing Rate Distribution

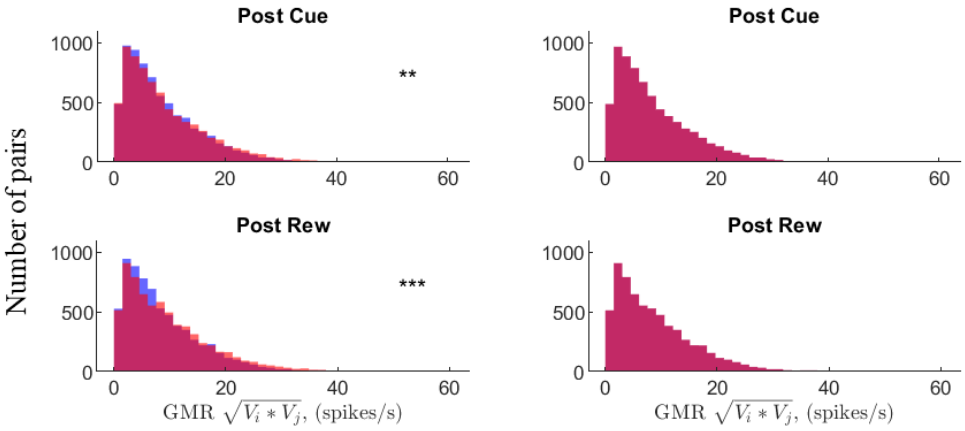

(a) Raw Firing Rate Distribution *(b) Mean Matched Firing Rate Distribution*

Fig. 27 Pooled raw and mean matched firing rate distributions for NHP P observational sessions. Red for rewarding trials distribution, blue for nonrewarding trials distribution.

NHP S – Mean Matched Noise Correlation

(a) Manual *(b) Observational*

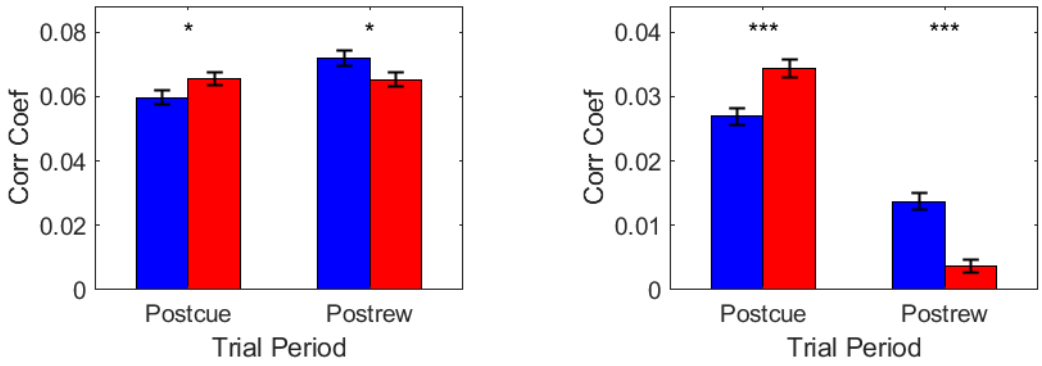

Fig. 28 Average noise correlation from the pooled mean matched unit pairs for NHP S. (a) Manual sessions. (b) Observational sessions.

NHP P – Mean Matched Noise Correlation

(a) Manual *(b) Observational*

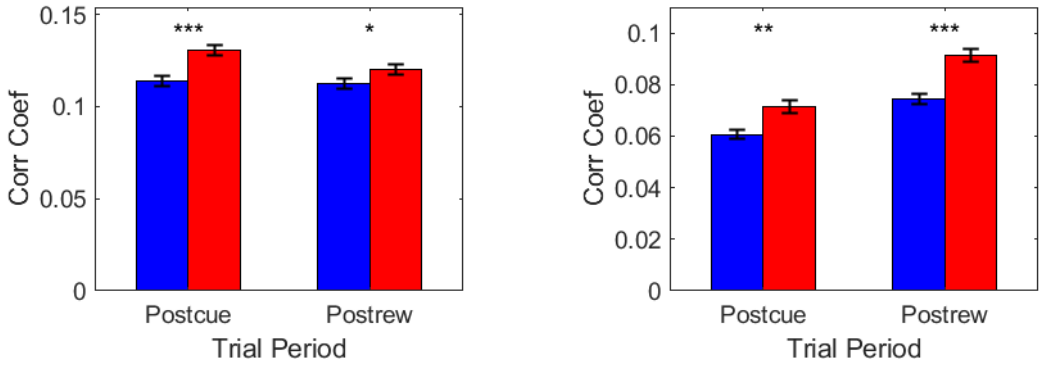

Fig. 29 Average noise correlation from the pooled mean matched unit pairs for NHP P. (a) Manual sessions. (b) Observational sessions.

Chapter 4 Discussion

In this work, we showed that raw noise correlations in NHPs performing a grip force task are generally higher in rewarding trials than that of nonrewarding when presented with a reward cue. This was consistent in both manual and observational tasks and was seen to be as a general case with subpopulation analysis. Our results are also not easily explained by the dissimilarities in firing rates. This suggest that cortical neurons in M1 can discriminate between rewarding and nonrewarding stimuli by changing the correlation structures.

The results are interesting as there is a large body of work that suggests that correlations decrease with the presentation of a stimulus or attention. Previous correlation analyses have been primarily focused on the visual cortex of the brain, where eliciting specific responses of a neuron is better understood. It was seen that a presented stimulus, generally in the preference of the neurons, was observed to decrease the correlation of the population for anesthetized and awake, behaving subjects (Zohary et al., 1994; Kohn and Smith., 2005; Huang and Lisberger, 2009). Some experiments vary in the window of measurement and state of the subject, which has been shown to change estimates of correlation (Bair et al., 2001; Reich et al., 2001; Mitchell et al., 2009;).

Previous work found that during center out reaching task (COT), rewarding trials showed a lower noise correlation for rewarding trials when compared to nonrewarding. The behavioral task of the study had NHP subjects either make a lateral movement towards one target to the right with visual endpoint feedback or passively observe the endpoint move towards the target. The trials started with a center hold of the endpoint along with a color cue describing what type the trial was, red for rewarding and blue for nonrewarding.

In the current study, movement was not restricted before trial starts, i.e. center hold. One possible explanation of the contrasting results is that the NHPs in this experiment were free to move their arm throughout the experiment. They were also allowed to grasp the force transducer at any given time during reset, cue, reach, release, and success scenes and would still be able to complete the trials successfully if they correctly grasped during grasp and transport scenes. This permission of freedom leads to behavioral differences in the NHPs under study and is seen in some the force profiles from the task data. NHP P periodically grabs the force transducer throughout the experiment, seeming to want to obtain the reward as soon as possible. NHP S on the other hand generally waits until the appropriate time. Their respective force outputs during rewarding trials are depicted in Fig. 30 and 31.

Force outputs of NHP S for Successful Trials

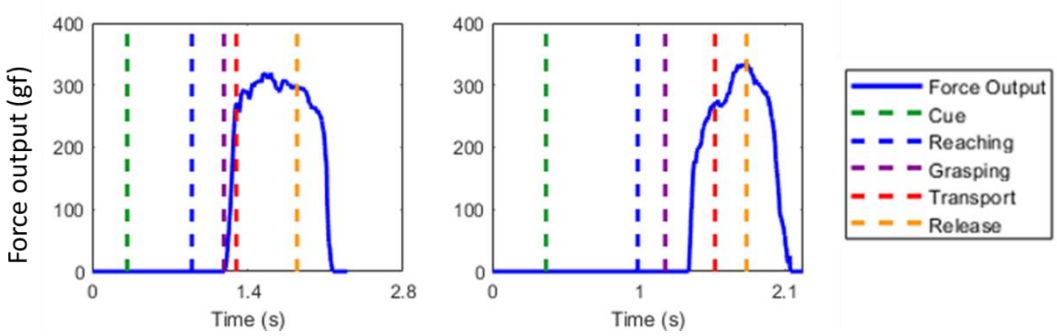

Fig. 30 Grip force output examples from NHP S. Time 0s is the reset scene and the start of the trial. The dashed lines are the difference scenes during the task.

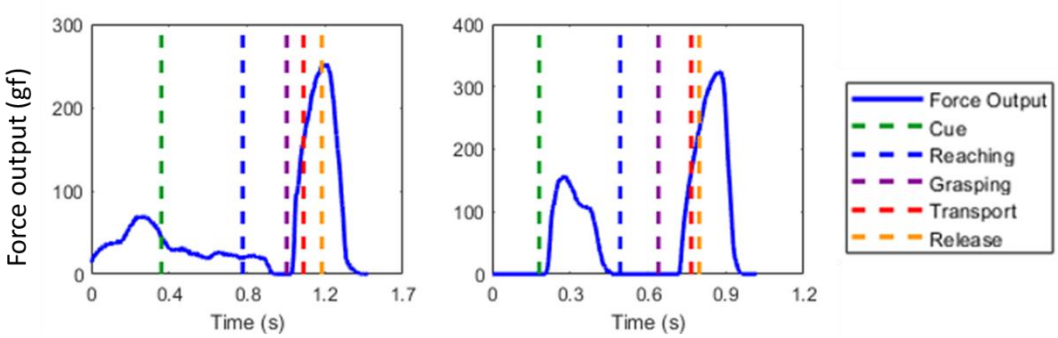

Fig. 31. Grip force output examples from NHP P. Time 0s is the reset scene and the start of the trial. The dashed lines are the difference scenes during the task.

Sometimes, NHP P does not let go before grasping scene and fails the trial, as evident in the percentage of failed trials (average 20% over all sessions). On the other hand, NHP S generally waits through the scenes until application of force will lead to the successful completion of the trial. This results in NHP S having a much lower failure rate (about 8% overall) This kind of behavior difference may explain the differing noise correlation results in the post-reward periods as well. NHP P seems fixated on getting reward, while NHP S seems to focus on completing the trials correctly to maximize reward delivered. Another difference between the COT and the GFT is that there is no cue present for nonrewarding trials in the GFT. Marsh et al. showed in 2015 that expectations of reward modulates the activity in M1. It is reasonable to imagine that reward expectations are different between a case where the subject is told explicitly that reward will not be delivered for a successful trial and the case where they must infer the same outcome implicitly through a lack of cue.

It was previously thought that if the stimulus was attended by a subject, then the shared variability would always decrease (Gutnisky and Dragoi, 2008; Cohen and Maunsell, 2009; Mitchell et al., 2009; Gu et al., 2011; Herrero et al., 2013). This theory has since been further developed with the observation that an increase or decrease in

correlation values depends on a neuron and what role they play in the task (Ruff and Cohen, 2014). Correlated variability was also previously thought to be detrimental to population statistics and that a decrease in neural correlations lead to more accurate decoding (Zohary et al., 1994). Though this assumption holds true for a population of homogeneous neurons, meaning they all share common inputs, it rarely is the case as neurons in cortical regions can have diverse responses (Monier et al., 2003; Truccolo et al., 2008; Ecker et al. 2011). Reduction of correlations do not necessarily mean better information transfer between the neurons. A population of diversely tuned neurons can lead to a more efficient encoding of the stimuli and more information can be carried through the correlations. (Shamir and Sompolinsky, 2006; Ecker et al., 2011; Goris et al., 2015). Some studies have shown that an increase in correlations can improve encoding accuracy for models as well, while others have shown that correlations can be stimulus dependent (Josic et al., 2009; Ecker et al., 2011).

It has been suggested that correlations increase with firing rates (de la Rocha et al, 2007). As this is a simplistic assumption for neuronal correlations, some studies have tried to control firing rates by mean matching the GMR of unit pairs (Kohn and Smith, 2005; Banyai et al 2018). Using the arithmetic mean could have been useful as well as it can help determine the differences in Fano factor using matched rates (Churchland et al., 2010), but the GMR has been shown to better predict noise correlations (Schulz et al., 2015). With the GMR matched between rewarding and nonrewarding trials, the results of section 3.5 showed that the same correlation trends were still observed as those of raw correlations. This indicates that the correlations are stimulus dependent and are less likely to be the cause of changes in rates.

Chapter 5 Conclusion and future work

The results from this experiment attempts to further the understanding of how stimuli are modulated in a population of M1 neurons. The higher noise correlation for rewarding trials during post-cue period are shown to be stimuli dependent and consistent with manual and observational tasks. The differences in post-reward period between subjects could result from a difference in behavior. Previous work showed that a different task resulted in an opposite effect, suggesting that noise correlations are task dependent as well. This nature of correlations could be helpful in developing a reinforcement learning BMI.

One possible way to extend this work is to observe noise correlations at: different periods of the trial, failed trials, multiple levels of reward, and higher levels of correlations such as triplets, quadruplets, etc. Lastly, design of an experiment that reduces the effects of external factors would solidify the results here. This could include restrictions of when the force can be outputted by the subjects.

References

An, J., Yadav, T., Hessburg, J. P., & Francis, J. T. (2018a). Reward Modulates Local Field Potentials, Spiking Activity and Spike-Field Coherence in the Primary Motor Cortex. *bioRxiv*, 471151.

An, J., Yadav, T., Ahmadi, M. B., Tarigoppula, V. S. A., & Francis, J. T. (2018b, July). Near perfect neural critic from motor cortical activity toward an autonomously updating brain machine interface. In *2018 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)* (pp. 73-76). IEEE.

Abbott, L. F., & Dayan, P. (1999). The effect of correlated variability on the accuracy of a population code. *Neural Computation, 11*(1), 91-101. doi:10.1162/089976699300016827

Bányai, M., Lazar, A., Klein, L., Klon-Lipok, J., Stippinger, M., Singer, W., & Orbán, G. (2019). Stimulus complexity shapes response correlations in primary visual cortex. *Proceedings of the National Academy of Sciences*, *116*(7), 2723-2732.

Bair, W., Zohary, E., & Newsome, W. T. (2001). Correlated firing in macaque visual area MT: Time scales and relationship to behavior. *Journal of Neuroscience,* 21(5), 1676-1697.

Bayer, H. M., & Glimcher, P. W. (2005). Midbrain dopamine neurons encode a quantitative reward prediction error signal. *Neuron*, 47(1), 129-141. doi:10.1016/j.neuron.2005.05.020

Buzsáki, G. "Large-scale recording of neuronal ensembles." *Nature neuroscience* 7.5 (2004): 446.

Churchland, M. M., Yu, B. M., Cunningham, J. P., Sugrue, L. P., Cohen, M. R., Corrado, G. S., Newsome, W. T., Clark A. M., Hosseini, P., Scot, B. B., Bradley, D. C., Smith, M. A., Kohn, A., Movshon, J. A., Armstrong, K. M., Moore, T., Chang, S. W. Snyder, L. H., Lisberger, S. G., Priebe, N. J., Finn, I. M., Fester, D., Ryu, S. I., Santhanam, G., Sahani, M., & Shenoy, K. V. (2010). Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nature Neuroscience*, 13(3), 369-U325. doi:10.1038/nn.2501

Chhatbar, P. Y., von Kraus, L. M., Semework, M., & Francis, J. T. (2010). A bio-friendly and economical technique for chronic implantation of multiple microelectrode arrays. *Journal of neuroscience methods*, 188(2), 187-194.

Cohen, M. R., & Kohn, A. (2011). Measuring and interpreting neuronal correlations. *Nature neuroscience*, 14(7), 811.

Cohen, M. R., & Maunsell, J. H. R. (2009). Attention improves performance primarily by reducing interneuronal correlations. *Nature Neuroscience*, 12(12), 1594-U1148. doi:10.1038/nn.2439

Cohen, M. R., & Maunsell, J. H. R. (2011). Using Neuronal Populations to Study the Mechanisms Underlying Spatial and Feature Attention. *Neuron*, 70(6), 1192-1204. doi:10.1016/j.neuron.2011.04.029

Cohen, M. X., & Frank, M. J. (2009). Neurocomputational models of basal ganglia function in learning, memory and choice. *Behavioural Brain Research*, 199(1), 141-156. doi:10.1016/j.bbr.2008.09.029

de la Rocha, J., Doiron, B., Shea-Brown, E., Josic, K., & Reyes, A. (2007). Correlation between neural spike trains increases with firing rate. *Nature, 448*(7155), 802-U806. doi:10.1038/nature06028

Ecker, A. S., Berens, P., Tolias, A. S., & Bethge, M. (2011). The Effect of Noise Correlations in Populations of Diversely Tuned Neurons. *Journal of Neuroscience, 31*(40), 14272-14283. doi:10.1523/jneurosci.2539-11.2011

Gawne, T. J., & Richmond, B. J. (1993). HOW INDEPENDENT ARE THE MESSAGES CARRIED BY ADJACENT INFERIOR TEMPORAL CORTICAL-NEURONS. *Journal of Neuroscience, 13*(7), 2758-2771.

Goris, R. L. T., Movshon, J. A., & Simoncelli, E. P. (2014). Partitioning neuronal variability. *Nature Neuroscience, 17*(6), 858-865. doi:10.1038/nn.3711

Gu, Y., Liu, S., Fetsch, C. R., Yang, Y., Fok, S., Sunkara, A., . . . Angelaki, D. E. (2011). Perceptual Learning Reduces Interneuronal Correlations in Macaque Visual Cortex. *Neuron, 71*(4), 750-761. doi:10.1016/j.neuron.2011.06.015

Gutnisky, D. A., & Dragoi, V. (2008). Adaptive coding of visual information in neural populations. *Nature*, 452(7184), 220-U263. doi:10.1038/nature06563

Herrero, J. L., Gieselmann, M. A., Sanayei, M., & Thiele, A. (2013). Attention-Induced Variance and Noise Correlation Reduction in Macaque V1 Is Mediated by NMDA Receptors. *Neuron*, 78(4), 729-739. doi:10.1016/j.neuron.2013.03.029

Hosp, J. A., Pekanovic, A., Rioult-Pedotti, M. S., & Luft, A. R. (2011). Dopaminergic Projections from Midbrain to Primary Motor Cortex Mediate Motor Skill Learning. *Journal of Neuroscience*, 31(7), 2481-2487. doi:10.1523/jneurosci.5411-10.2011

Huang, X., & Lisberger, S. G. (2009). Noise correlations in cortical area MT and their potential impact on trial-by-trial variation in the direction and speed of smooth-pursuit eye movements. *Journal of Neurophysiology*, 101(6), 3012-3030.

Josic, K., Shea-Brown, E., Doiron, B., & de la Rocha, J. (2009). Stimulus-Dependent Correlations and Population Codes. *Neural Computation*, 21(10), 2774-2804. doi:10.1162/neco.2009.10-08-879

Kohn, A., & Smith, M. A. (2005). Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. *Journal of Neuroscience*, 25(14), 3661-3673.

Kohn, A., Coen-Cagli, R., Kanitscheider, I., & Pouget, A. (2016). Correlations and neuronal population information. *Annual review of neuroscience*, *39*, 237-256.

Kapogiannis, D., Campion, P., Grafman, J., & Wassermann, E. M. (2008). Reward-related activity in the human motor cortex. *European Journal of Neuroscience, 27*(7), 1836-1842. doi:10.1111/j.1460-9568.2008.06147.x

Lee, D., Port, N. L., Kruse, W., & Georgopoulos, A. P. (1998). Variability and correlated noise in the discharge of neurons in motor and parietal areas of the primate cortex. *Journal of Neuroscience, 18*(3), 1161-1170.

Liu, J., & Newsome, W. T. (2006). Local field potential in cortical area MT: Stimulus tuning and behavioral correlations. *Journal of Neuroscience, 26*(30), 7779-7790. doi:10.1523/jneurosci.5052-05.2006

Macke, J. H., Berens, P., Ecker, A. S., Tolias, A. S., & Bethge, M. (2009). Generating spike trains with specified correlation coefficients. *Neural computation, 21*(2), 397-423.

Marsh, B. T., Tarigoppula, V. S. A., Chen, C., & Francis, J. T. (2015). Toward an Autonomous Brain Machine Interface: Integrating Sensorimotor Reward Modulation and Reinforcement Learning. *Journal of Neuroscience, 35*(19), 7374-7387. doi:10.1523/jneurosci.1802-14.2015

McNiel, D. B., Choi, J. S., Hessburg, J. P., & Francis, J. T. (2016, Aug 16-20). Reward Value is Encoded in Primary Somatosensory Cortex and can be Decoded from Neural Activity During Performance of a Psychophysical Task. Paper presented at the 38th Annual International Conference of the *IEEE-Engineering-in-Medicine-and-Biology-Society (EMBC),* Orlando, FL.

Mitchell, J. F., Sundberg, K. A., & Reynolds, J. H. (2009). Spatial Attention Decorrelates Intrinsic Activity Fluctuations in Macaque Area V4. *Neuron, 63*(6), 879-888. doi:10.1016/j.neuron.2009.09.013

Molina-Luna, K., Pekanovic, A., Rohrich, S., Hertler, B., Schubring-Giese, M., Rioult-Pedotti, M. S., & Luft, A. R. (2009). Dopamine in Motor Cortex Is Necessary for Skill Learning and Synaptic Plasticity. *Plos One, 4*(9), 10. doi:10.1371/journal.pone.0007082

Monier, C., Chavane, F., Baudot, P., Graham, L. J., & Fregnac, Y. (2003). Orientation and direction selectivity of synaptic inputs in visual cortical neurons: A diversity of combinations produces spike tuning. *Neuron, 37*(4), 663-680. doi:10.1016/s0896-6273(03)00064-3

Moreno-Bote, R., Beck, J., Kanitscheider, I., Pitkow, X., Latham, P., & Pouget, A. (2014). Information-limiting correlations. *Nature Neuroscience, 17*(10), 1410-1417. doi:10.1038/nn.3807

Ramakrishnan, A., Byun, Y. W., Rand, K., Pedersen, C. E., Lebedev, M. A., & Nicolelis, M. A. L. (2017). Cortical neurons multiplex reward-related signals along with sensory and motor information. *Proceedings of the National Academy of Sciences of the United States of America, 114*(24), E4841-E4850. doi:10.1073/pnas.1703668114

Reich, D. S., Mechler, F., & Victor, J. D. (2001). Independent and redundant information in nearby cortical neurons. *Science,* 294(5551), 2566-2568.

Roesch, M. R., & Olson, C. R. (2004). Neuronal activity related to reward value and motivation in primate frontal cortex. *Science, 304*(5668), 307-310. doi:10.1126/science.1093223

Ruff, D. A., & Cohen, M. R. (2016). Attention increases spike count correlations between visual cortical areas. *Journal of Neuroscience, 36*(28), 7523-7534.

Schoenbaum, G., & Setlow, B. (2001). Integrating orbitofrontal cortex into prefrontal theory: Common processing themes across species and subdivisions. *Learning & Memory, 8*(3), 134-147. doi:10.1101/lm.39901

Schultz, W. (1998). Predictive reward signal of dopamine neurons. *Journal of Neurophysiology, 80*(1), 1-27.

Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science, 275*(5306), 1593-1599. doi:10.1126/science.275.5306.1593

Schultz, W. (2015). NEURONAL REWARD AND DECISION SIGNALS: FROM THEORIES TO DATA. *Physiological Reviews, 95*(3), 853-951. doi:10.1152/physrev.00023.2014

Shadlen, M. N., & Newsome, W. T. (1998). The variable discharge of cortical neurons: Implications for connectivity, computation, and information coding. *Journal of Neuroscience, 18*(10), 3870-3896.

Shamir, M., & Sompolinsky, H. (2006). Implications of neuronal diversity on population coding. *Neural Computation*, 18(8), 1951-1986. doi:10.1162/neco.2006.18.8.1951

Silver, N. C., & Dunlap, W. P. (1987). Averaging correlation coefficients: should Fisher's z transformation be used?. *Journal of applied psychology*, *72*(1), 146.

Tarigoppula, V. S. A., Choi, J. S., Hessburg, J. H., McNiel, D. B., Marsh, B. T., & Francis, J. T. (2018). Primary Motor Cortex Encodes A Temporal Difference Reinforcement Learning Process. *bioRxiv*, 257337.

Tolhurst, D. J., Movshon, J. A., & Dean, A. F. (1983). THE STATISTICAL RELIABILITY OF SIGNALS IN SINGLE NEURONS IN CAT AND MONKEY VISUAL-CORTEX. *Vision Research*, 23(8), 775-785. doi:10.1016/0042-6989(83)90200-6

Truccolo, W., Friehs, G. M., Donoghue, J. P., & Hochberg, L. R. (2008). Primary motor cortex tuning to intended movement kinematics in humans with tetraplegia. *Journal of Neuroscience*, *28*(5), 1163-1178. doi:10.1523/jneurosci.4415-07.2008

Yager, L. M., Garcia, A. F., Wunsch, A. M., & Ferguson, S. M. (2015). THE INS AND OUTS OF THE STRIATUM: ROLE IN DRUG ADDICTION. *Neuroscience*, *301*, 529-541. doi:10.1016/j.neuroscience.2015.06.033

Zohary, E., Shadlen, M. N., & Newsome, W. T. (1994). CORRELATED NEURONAL DISCHARGE RATE AND ITS IMPLICATIONS FOR PSYCHOPHYSICAL PERFORMANCE. *Nature*, *370*(6485), 140-143. doi:10.1038/370140a0

Appendix

A. Other Figs from the NHPs

The zero time in Fig. A.1 to A.4 are centered at the previously described cue presentation and reward delivery. In Fig. A.3 top left, the session shows that for the peri-cue period, firing rates for nonrewarding trials were higher than that of rewarding ones. This goes against convention as desirable stimuli is thought to elicit a stronger response. This could be explained by the way the subject behaves during some trials as it seems to be more driven by getting reward than by completing the trials to get maximum reward.

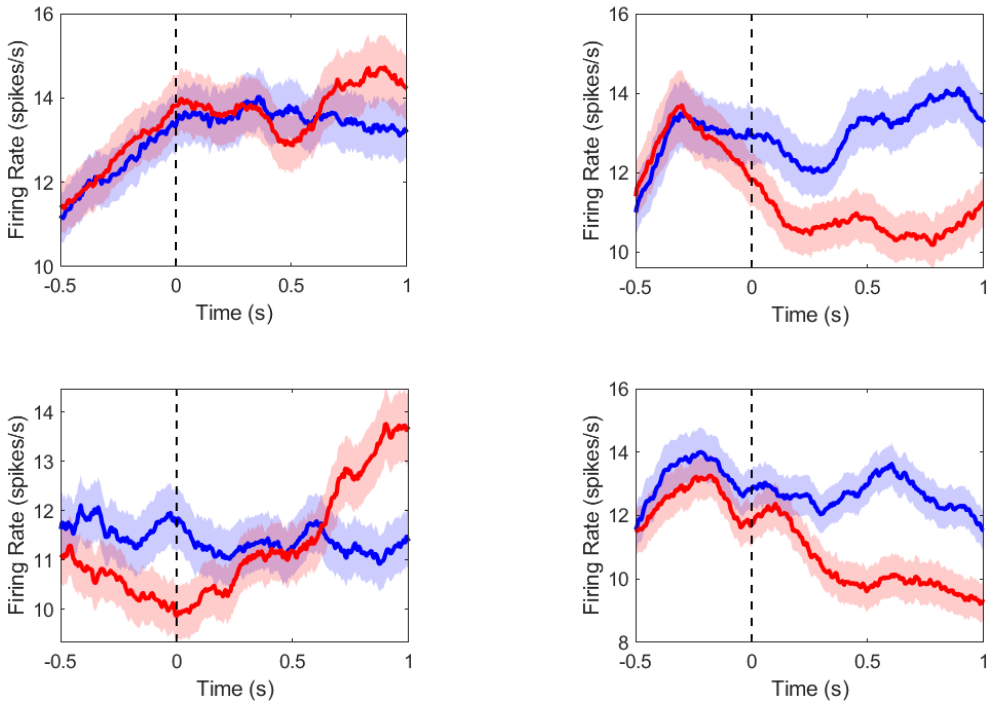

Fig. A.1 Average population firing rates for NHP S manual sessions not shown previously. (a) Manual sessions. (b) Observational sessions

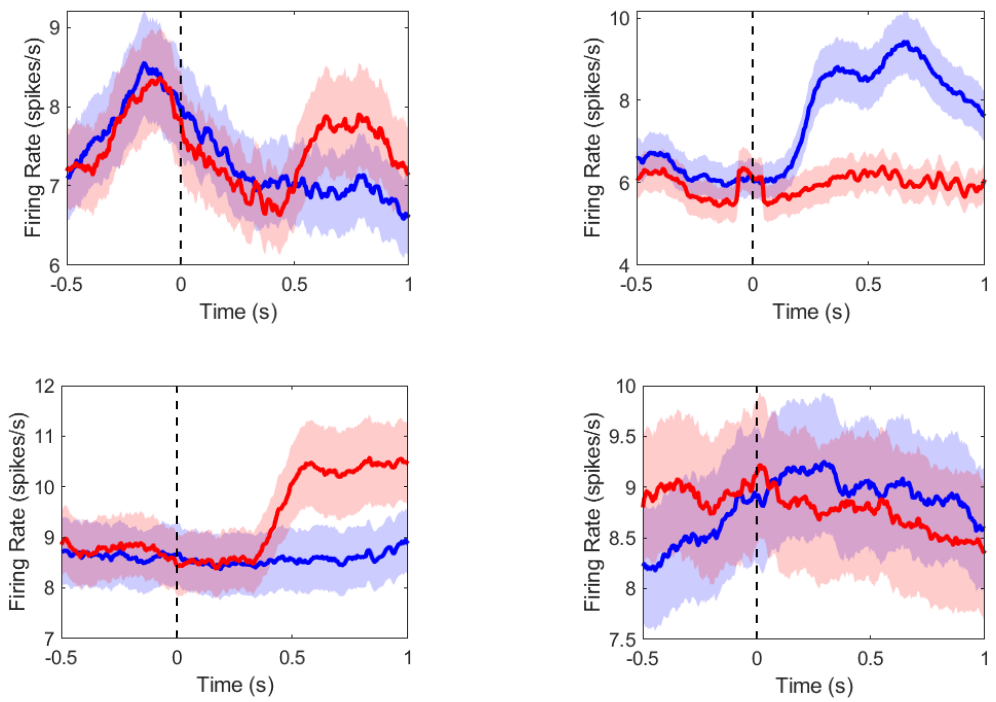

Fig. A.2 Average population firing rates for NHP S observational sessions not shown previously. (a) Manual sessions. (b) Observational sessions

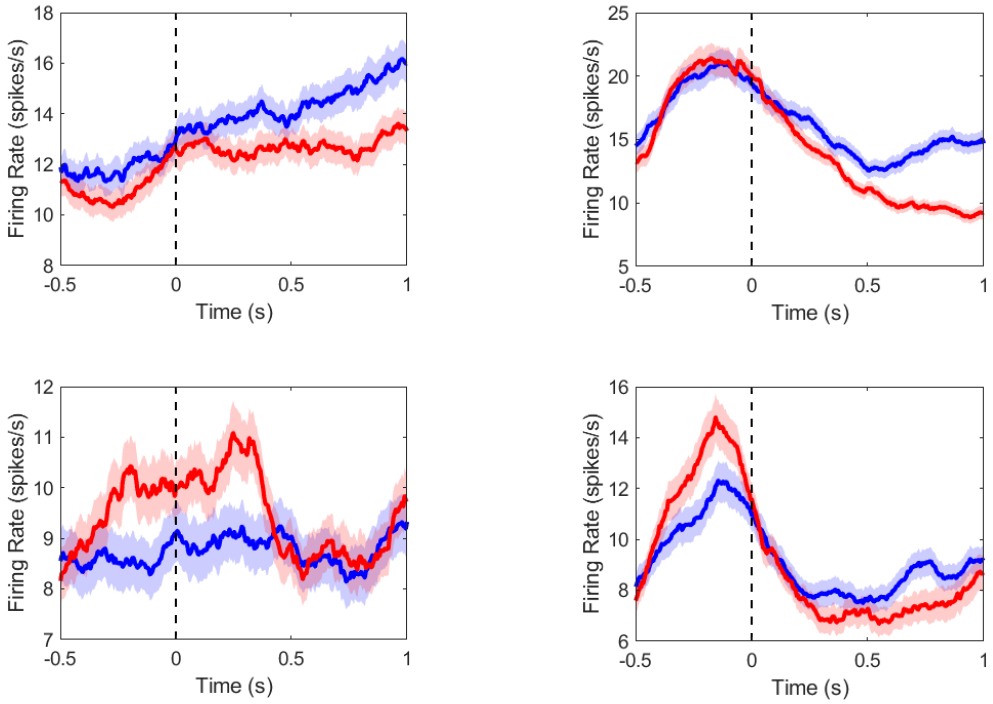

Fig. A.3 Average population firing rates for NHP P manual sessions not shown previously. (a) Manual sessions. (b) Observational sessions

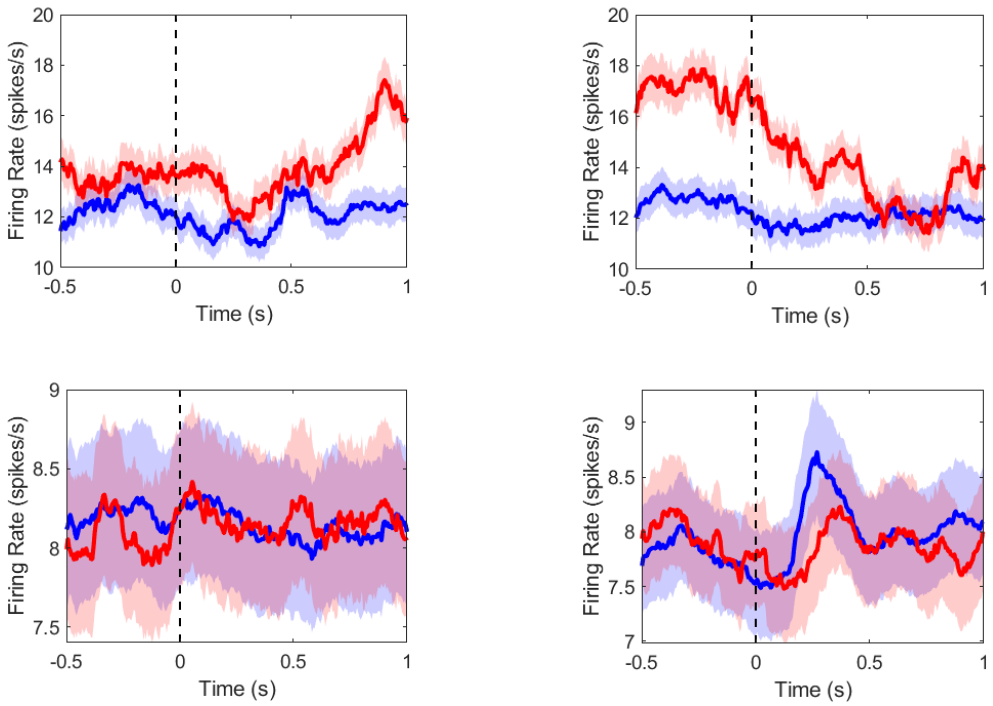

Fig. A.4 Average population firing rates for NHP P observational sessions not shown previously. (a) Manual sessions. (b) Observational sessions

B. Validation of code

To validate that the codes were not giving random results, spike trains of differing correlations were generated and tested. The method for generating these spike trains are detailed in Macke et al., 2009. In summary, the generation uses a dichotomized gaussian distribution model. The firing rates will be normally distributed, and the occurrence of a spike in time will result in possible simultaneous spike times in other simulated neurons.

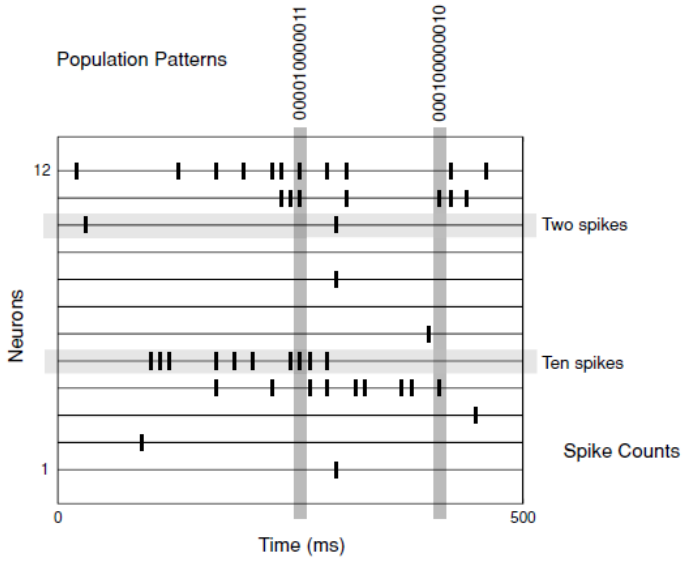

Fig. B.1 Macke et al., 2009. Illustration of how spike times can be correlated.

Correlations are modulated through generating a covariance matrix of size N x N, where N is the number of simulated neurons. This method does not allow for direct control of actual correlations, but for validation, the covariance was set to high and low values for simulated “rewarding” and “nonrewarding”.

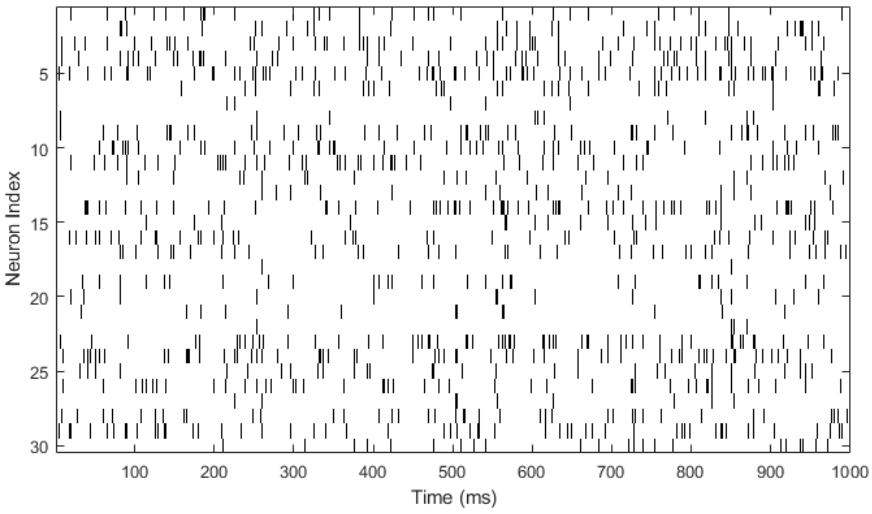

Fig. B.2 Example of simulated raster.

This model does not consider the absolute or relative refractory periods, so the model is not physiologically realistic. The resulting firing rate distribution and average noise correlation are displayed in Fig.

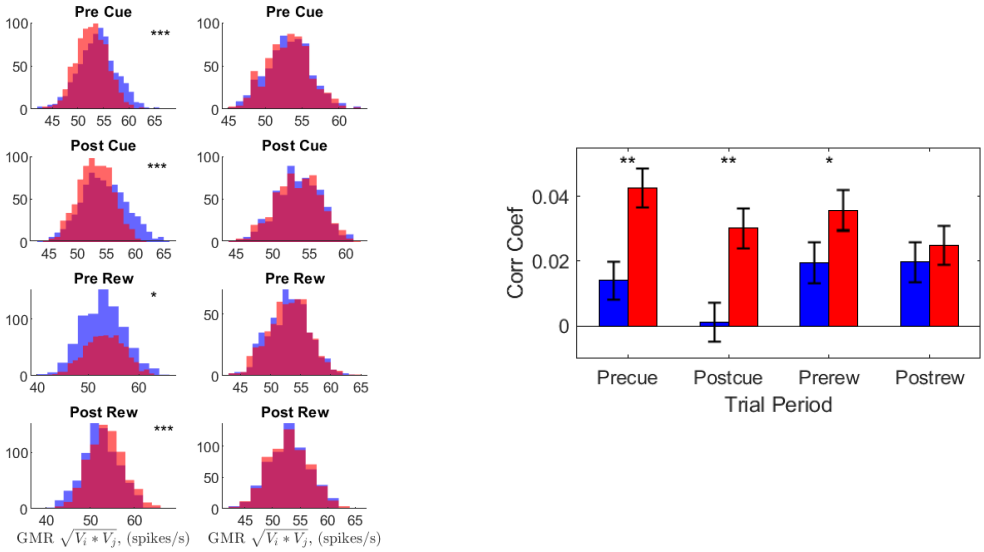

a b

Fig. B.3 Analysis of the simulated population where "rewarding trials" had higher correlations. (a) Actual and mean matched FR distribution of simulated population. (b) Mean noise correlation of simulated population.

The correlation trends are as expected. To further inspect if the code is not biased for rewarding trials, toy data was generated with it known that nonrewarding would have higher correlation values.

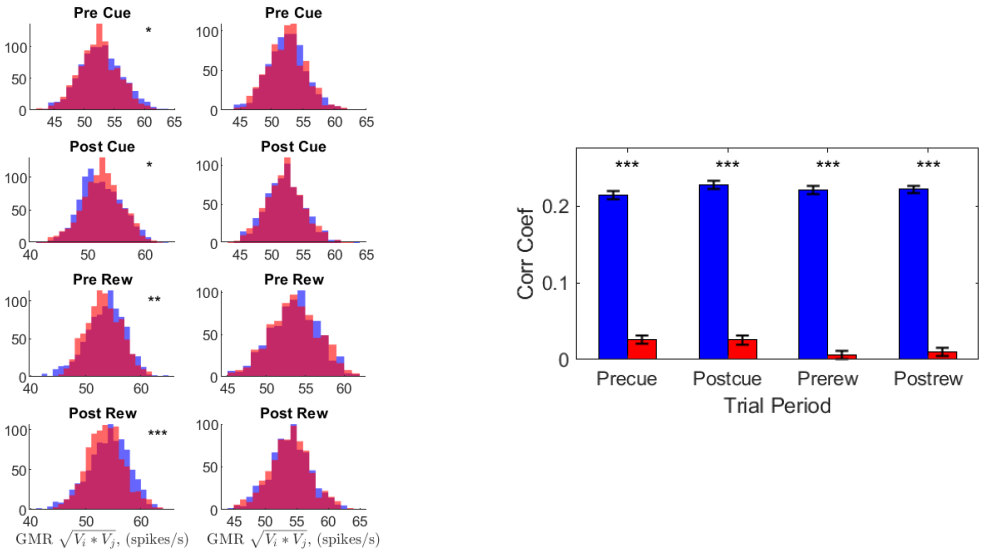

a b

Fig. B.4 Sanity check with simulated "nonrewarding trials." Results still showed which would be result in higher mean noise correlations. (a) Actual and mean matched FR distribution of simulated population. (b) Mean noise correlation of simulated population