

A Neural Model of Context Dependent Decision Making in the Prefrontal Cortex

Sugandha Sharma (s72sharm@uwaterloo.ca)

Brent J. Komer (bjkomer@uwaterloo.ca)

Terrence C. Stewart (tcstewar@uwaterloo.ca)

Chris Eliasmith (celiasmith@uwaterloo.ca)

Center for Theoretical Neuroscience, 200 University Ave. W.
Waterloo, ON N2L3G1 Canada

Abstract

In this paper, we present a spiking neural model of context dependent decision making. Prefrontal cortex (PFC) plays a fundamental role in context dependent behaviour. We model the PFC at the level of single spiking neurons, to explore the underlying computations which determine its contextual responses. The model is built using the Neural Engineering Framework and performs input selection and integration as a nonlinear recurrent dynamical process. The results obtained from the model closely match behavioural and neural experimental data obtained from macaque monkeys that are trained to perform a context sensitive perceptual decision task. The close match suggests that the low-dimensional, nonlinear dynamical model we suggest captures central aspects of context dependent decision making in primates.

Keywords: context dependent decision making; decision making; neural engineering framework; neural dynamics; theoretical neuroscience

Introduction

Decision making is a complex process that involves the aggregation of many sensory signals in order to produce a response. What kinds of sensory information are relevant as well as what to do with that information can be highly dependent on the context of the situation. Prefrontal cortex (PFC) contributes to representing and maintaining relevant contextual information, ignoring irrelevant sensory information. Neurons in PFC have been found to be involved in selective attention and selection of an intended action (Tanji & Hoshi, 2008). Thus the relevant sensory evidence can be thought to be integrated in PFC before a decision is made in a contextual situation.

Early selection models have been used in the past to explain the underlying context dependent computations in PFC. These models suggest that the irrelevant inputs are filtered out before the integration stage in PFC thus leading to a larger effect of relevant as compared to the irrelevant sensory information on context dependent behaviour. However, in a recent experiment of context dependent behaviour in monkeys, no evidence of early selection was found (Mante et al., 2013). Instead it was found that the relevant sensory input is selected late by the same PFC circuitry which integrates sensory evidence to make a choice.

In this paper we present a spiking neural model of context dependent decision making based on the above findings. The model is built using the Neural Engineering Framework (NEF; Eliasmith & Anderson (2003)) and is meant to provide insight into the biological mechanisms behind context dependent computations in PFC.

Related Work

To explore context dependent decision making in prefrontal cortex, Mante et al. (2013) designed an experiment in which macaque monkeys were trained to perform two different perceptual discriminations on the same set of visual stimuli based on the given context. On the basis of the results of the experiment, they built a learned recurrent network model to characterize the mechanisms underlying context dependent behaviour.

Experiment

Two monkeys (monkey A and monkey F) were trained on a perceptual discrimination task. On each trial of the task, the monkey was shown a random dot display consisting of dots of two different colours. A fraction of the dots moved coherently in one of the two directions (left or right), while the remaining dots were transiently flashed at random locations. The monkey was provided a contextual cue to instruct it to either discriminate the direction of motion (i.e., the coherent direction of motion of the dots) or the colour (i.e., which of the two colours are on the largest number of dots). The monkey reported its choice with a saccade to one of the two visual targets provided. After each trial the monkey was rewarded for correct responses with a small quantity of juice.

To vary the difficulty of the colour discrimination, the relative number of dots of the two colours was parametrically changed, while keeping the total number of dots constant. The fraction of colour 1 to colour 2 dots i.e., colour coherence (100% coherence: only one colour, 0% coherence: equal number of dots of two colours) was fixed throughout a trial. To vary the difficulty of the motion discrimination, the fraction of dots moving coherently i.e., motion coherence (100% coherence: all dots moving in the same direction, 0% coherence: all dots moving randomly) was parametrically changed. A total of 163,187 trials were performed and the trials varied with respect to the colour coherence and motion coherence of the random dot display, and the contextual cue provided.

Extracellular responses were recorded from the neurons in the frontal eye field (an area of PFC involved in the selection and execution of saccadic eye movements) while the monkeys performed this task. The responses of the population of neurons as a whole were analyzed to understand the neural mechanism underlying context-dependent selective integration of sensory stimuli.

A learned recurrent network model

To identify the mechanism underlying context-dependent behaviour, Mante et al. (2013) trained a network of recurrently connected neurons to solve a task analogous to the one solved by monkeys. Their model reproduced the main features of the data, and they analyzed the trained network to determine the selection mechanism. The results of their analysis indicate that the rich dynamics of PFC responses during selection and integration of inputs can be understood with just two features - a line attractor and a selection vector, which are defined only at the level of the neural population. It's important to note that their results indicate that computations in PFC emerge from collective dynamics of large populations of neurons. Further details about this model can be found in Mante et al. (2013) paper.

However, their model results do not fit the reported data very closely, significantly deviating from the response dynamics reported for the monkey experiments. As well, because the model was learned from data, it is unclear precisely what the nature of the underlying computation is. The identification of two degrees of freedom (i.e., memory and selection) is helpful, but is evident from performing dimensionality reduction on the original data, without a model. Here we propose a specific, low-dimensional, nonlinear computation and demonstrate a superior fit to the original data. Additionally, we show how such computations can be implemented in biologically realistic spiking neurons.

A nonlinear dynamical model

The basic principle of the model we describe is that four different pieces of information must be represented by the same group of neurons. The empirical data indicates that the neurons in the frontal eye field change their activity based on the context (whether the response should be based on colour or motion), colour coherence (the proportion of dots of one colour over another), motion coherence (the proportion of dots moving in one direction over another), and the actual choice being made. We denote this combination of four values as the vector x .

Rather than this information being held separately by different groups of neurons, it is distributed across a single population. This means that each neuron in the population is sensitive to some combination of those four values. Some neurons may be sensitive to colour and motion, while others may respond primarily due to the final choice, and others may be equally sensitive to all four values. In general, we assume that each neuron has some particular combination of values for which it will fire most strongly, consistent with the standard idea of a preferred stimulus or a preferred direction vector (Schwartz et al., 1988). If this vector is e (for "encoder"), then the amount of current flowing into a particular neuron i will be $\alpha_i e_i \cdot x + J_i$, where α and J_i are a gain and bias respectively. This current is used as the input to the standard leaky integrate-and-fire (LIF) neuron model to produce spiking output (it is these spikes which flow through the con-

nections and not the continuous signals). The encoder, gain, and bias are randomly chosen, so different neurons have different background firing rates and different maximum firing rates (set here to be uniformly distributed between 20Hz and 120Hz) along their preferred direction.

The purpose of this model is to combine the three inputs (context, colour coherence, and motion coherence) to produce an appropriate fourth value (the appropriate choice). To provide the input, we directly stimulate the neurons based on the sensory inputs of the given trial (we are not simulating the visual system). From these inputs, the neural population must compute the resulting choice of action. To capture this computation, the neurons are recurrently connected to compute a nonlinear dynamical system that appropriately integrates the inputs, and stores the resulting choice.

We achieve this by writing a mathematical expression for how the choice value should be dynamically computed given the sensory input, and then solving for the neural connectivity that optimally approximates that mathematical expression. This mathematical expression is our hypothesis about how information is integrated in PFC to produce a choice. Specifically, writing the dynamics of the population as representing the 4D state variable, x (where x_4 is the final choice):

$$\dot{x}_4 = f(x) \quad (1)$$

we must determine the function f . In a simple PFC memory model, $f(x) = 0$. However, here there are clearly context interactions which must be accounted for. We propose that the nonlinearity in the recurrent connection is given by

$$f(x) = (1 + x_1)x_2 + (1 - x_1)x_3 \quad (2)$$

where x_1 is the task context (1 if the decision should be based on motion, and -1 if it should be based on colour), x_2 is the motion stimulus (between -1 for completely leftward motion, and +1 for completely rightward motion), and x_3 is the colour stimulus (between -1 for all one colour, and +1 for all the other colour). In short, this nonlinearity essentially gates either colour or motion information into a working memory integrator, depending on the context signal. Consequently, the model is a simple hypothesis of multiplicative gating in PFC. Importantly, this multiplicative effect is achieved at the network level using normal synapses between LIF neurons. Mathematically, it basically forms a line attractor, however, this needs to be explored further in future work.

To simulate the model, we used the software package Nengo, a Python-based simulation engine and environment for constructing neural networks that includes facilities for generating NEF-style networks (Bekolay et al., 2014).

A schematic of the model is shown in Figure 1. The central group of neurons consists of 1000 randomly generated LIF neurons with preferred direction vectors e such that each neuron responds differentially to the four x values being represented. These neurons are connected recurrently back to themselves in an all-to-all manner (for a total of 1 million synapses), with the synaptic connection strengths optimized

via the NEF to closely approximate equation 2. We also included a separate group of neurons that just represent the final choice (x_4). The synaptic connection weights from the “pfc” neurons to the “choice” neurons are also found using the NEF, where the target function was simple $choice = x_4$, providing a simple readout of the model’s choice.

For each trial, the inputs were set as follows:

- Contextual cue (x_1) - This indicates the context for each trial and takes a value of 1 and -1 for motion and colour contexts respectively. Recall, motion context implies discriminating the direction of motion of the dots (i.e., make a decision based on motion coherence) whereas the colour context implies discriminating the colour of the dots (i.e., make a decision based on colour coherence).
- Colour coherence (x_2) - For each trial one of the six signed colour coherence values: [-0.50, -0.18, -0.06, 0.06, 0.18, 0.50] is randomly provided to the model as an input. The sign of the values indicates different colours i.e., a negative sign corresponds to colour 1 (red) while a positive sign corresponds to colour 2 (green).
- Motion coherence (x_3) - For each trial one of the six signed motion coherence values: [-0.50, -0.15, -0.05, 0.05, 0.15, 0.50] was randomly provided to the model as an input. The sign of the values indicates the direction of motion i.e., a negative sign indicates motion towards left while a positive sign indicates motion towards right.

Different values of motion and colour coherence mentioned above were used to vary the strength of the motion and colour signals between trials. The output of the model is a choice (i.e., the decision made for a particular trial given the context). Each trial in the model is 750ms long (i.e., inputs are provided to the model for 750ms per trial) followed by a delay period of 750ms between consecutive trials. This delay period was added to avoid the influence of the previous trial on the current trial, and can be thought of as analogous to the period when the monkey was being rewarded between trials.

In the experiment, the monkeys were able to quickly switch between behavioural contexts. This implies the presence of rapid modulation or ‘gating’ mechanisms in the PFC that select relevant sensory information for decision making. In our model, multiplicative gating is used to control the dynamics of “pfc”, such that only the relevant evidence in a given context gets integrated. For example in motion context ($contextual_cue = 1$), only the motion coherence values get integrated over time and the colour coherence values are irrelevant to the final decision. This is analogous to the actual experiment because once the monkey knows the context, it pays selective attention to only the relevant evidence and the irrelevant evidence mostly gets ignored. When the integrated relevant evidence passes beyond a threshold, the model is considered to have made a choice.

Note that instead of early selection (i.e., selection of relevant sensory input before the integration stage in PFC), input

selection and integration happens within the same PFC circuitry as a part of a single dynamical process.




Figure 1: A schematic diagram of the model structure built using Nengo GUI (a graphical user interface to aid in the design and testing of Nengo networks). Here “pfc” contains 1000 neurons and “choice” contains 200 neurons.

Results

The model runs through a series of trials consecutively, and the choice the model makes for each trial is recorded, along with the particular condition variables for that trial (context, motion coherence, colour coherence, and correct choice). Figure 2 shows the spike trains and decoded values from one such trial. The spike trains are shown only for 10 random neurons picked from a population of 1000 neurons in “pfc”, and show that different neurons fire at different times according to their preferred stimulus. The spike trains from each individual neuron are recorded for each trial. This allows the model output to be compared to the monkey data on both the behavioural and neural level. Note that we compare our model to the performance of monkey A since monkey A data has primarily been discussed by Mante et al. (2013). Additionally we fix our model parameters and use the data from the same model for behavioural and neural comparisons.

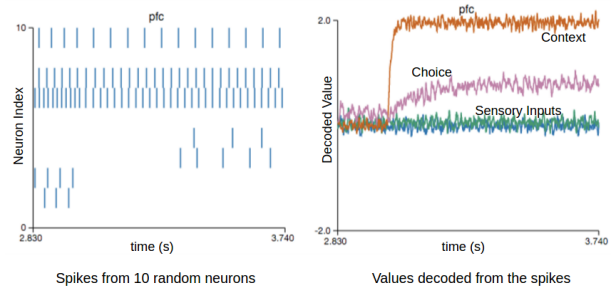


Figure 2: Spike trains and the decoded values from “pfc”. The decoded values show the context, choice and the sensory inputs (i.e., colour and motion coherence).

Behavioural comparison

A comparison of the model output to the behavioural data of the monkey is shown in Figure 3. The figure plots the performance of the model as a function of motion and colour coherence in a given behavioural context. The performance is measured in terms of % of choices made to the right, which

implies choosing ‘right’ direction in motion context and % of choices made to green which implies choosing the ‘green’ colour (i.e., colour 2) in colour context. The plot on the top left shows that in the motion context, the model makes very few choices towards right for negative values of the motion coherence (recall that negative values simulate dots moving to the left). However the % choices made to the right increases with an increase in the motion coherence, and is highest when the motion coherence takes its maximum positive value (recall that positive values simulate dots moving to the right). The plot on the top right shows that the colour coherence does not have much effect on the choices in motion context. Overall, the two plots show that in the motion context, the choices of the model depend strongly on motion coherence of the dots whereas they depend only weakly on colour coherence. The opposite pattern is observed in the colour context (bottom two plots) in which the choices of the model are largely affected by colour coherence and weakly by motion coherence.

These results show that the model successfully differentiates the relevant sensory evidence in each context and ignores the irrelevant evidence. The result also matches well with the monkey data obtained experimentally as shown in Figure 2. The overall % average error of our model with respect to the experimental data is 7.33%, while that for the learned recurrent network model of Mante et al. (2013) is 7.45%. Importantly, in their model the response is completely independent of the irrelevant input (e.g., the %choices to the right is constant with respect to the motion coherence in colour context). However, in our model, the irrelevant input has a weak effect on the response which is consistent with the experimental data. Our model captures this subtle feature since it is built using biologically realistic LIF neurons. Moreover, our model has only one parameter which is a constant transform (0.45) applied to scale the motion and colour coherence inputs (this parameter was not tuned in the results we present). The synaptic time constant (200ms) was fixed based on biological data (Flint et al., 1997) and the neural firing rates were set to be uniformly distributed between 20Hz to 120Hz. The connection strengths were not fit to the experimental data, but to the equation that we are trying to approximate i.e., Equation 2, using the NEF (Eliasmith & Anderson, 2003). On the other hand, the learned recurrent network model has many parameters which are learned through training.

Analysis of neural spike data

In order to facilitate a quantitative comparison to the monkey data, identical analysis methods as described in the supplementary information from Mante et al. (2013) are used on our spike data. All the analysis is performed on the spike data obtained from the central population of neurons i.e., ”pfc” in Figure 1.

First, we determine the condition average population response by grouping the trials based on the unique combination of condition variables. There are five different condition variables: context (motion or colour), motion coherence (range of six values), colour coherence (range of six

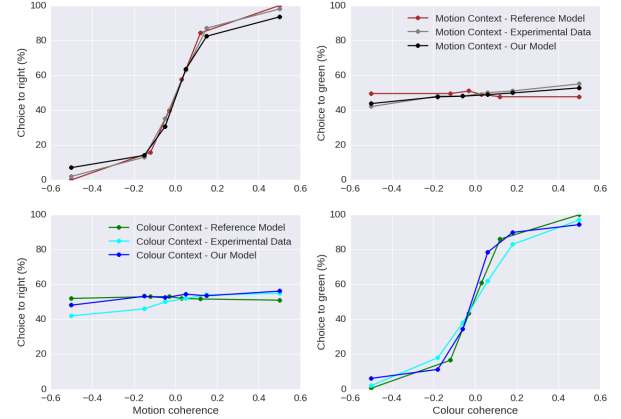


Figure 3: Behavioural output of the model. Results are averaged across 14,688 trials. Reference model and experimental data used for the comparison are taken from Mante et al. (2013) and were averaged across 14,400 and 163,187 trials respectively.

values), choice (two options), and correctness (whether or not the monkey made the correct choice for the given condition). This amounts to a total of 288 possible unique conditions ($2 \times 6 \times 6 \times 2 \times 2$).

The spike response of each neuron is averaged across all trials having the same unique conditions, and then smoothed by a gaussian filter with $\sigma = 40ms$. The result is then z-scored for each neuron to give the overall response across all times a zero mean and unit standard deviation. This results in a two dimensional data matrix X of size $N_{neuron} \times (T \cdot N_{condition})$ containing the condition averaged population response. Here N_{neuron} is the total number of neurons in the neural population, $N_{condition}$ is the total number of unique conditions and T is the number of time samples. Note that for our model $N_{neuron} = 1000$, $N_{condition} = 288$ and $T = 750$.

Next, we perform principal component analysis (PCA) on the condition averaged response (X) to identify the dimensions in state space that capture the most variance. The PCA is done along the number of neurons and the principle components (PCs) are vectors of length $N_{neurons}$. We use the first twelve PCs to define a de-noising matrix D of size $N_{neuron} \times N_{neuron}$ using equation 3.

$$D = \sum_{a=1}^{12} v_a v_a^T \quad (3)$$

Next, we perform linear regression to determine the effect of various task variables (i.e., choice, motion, colour and context) on the response of each neuron. In order to perform linear regression, the response of neuron i at time t is represented as a linear combination of task variables as shown in equation 4, where $r_{i,t}(k)$ is the z-scored spike response of neuron i at time t on trial k . The values of the four task variables also correspond to trial k . The regression coefficients $\beta_{i,t}(v)$ for $v = 1$ to 4 describe how the corresponding task variable affects the

firing rate of neuron i at a given time t during the trial.

$$r_{i,t}(k) = \beta_{i,t}(1)choice(k) + \beta_{i,t}(2)motion(k) + \beta_{i,t}(3)colour(k) + \beta_{i,t}(4)context(k) + \beta_{i,t}(5) \quad (4)$$

In order to estimate the regression coefficients, a set of condition matrices is defined, one for each neuron. Each matrix F_i is of the size $N_{coeff} \times N_{trial}$, meaning there is a row for each of the five regression coefficients and a column for each trial. The first four rows of F_i contain the task variables for each trial, while the last row contains only ones, and is used solely to estimate $\beta_{i,t}(5)$. The regression coefficients are then estimated using equation 5.

$$\beta_{i,t} = (F_i F_i^T)^{-1} F_i r_{i,t} \quad (5)$$

The matrix $\beta_{i,t}$ of length N_{coeff} is then rearranged to form $\beta_{v,t}$ of length N_{neuron} which is de-noised by projecting it to the subspace spanned by the first 12 PCs using equation 6.

$$\beta_{v,t}^{pca} = D \beta_{v,t} \quad (6)$$

$\beta_{v,t}^{pca}$ is of length N_{neuron} and contains the de-noised regression coefficients. Next step is to define the time-independent de-noised regression vectors. To do this we determine the time at which the corresponding set of vectors $\beta_{v,t}^{pca}$ has the maximum norm, and select the value of the regression coefficients at that time (equation 7).

$$\beta_v^{max} = \beta_{v,t}^{pca,max} \quad (7)$$

β_v^{max} is of length N_{neuron} and is orthogonalized using QR decomposition as shown in equation 8.

$$B^{max} = [\beta_1^{max} \beta_2^{max} \beta_3^{max} \beta_4^{max}] = QR \quad (8)$$

The first four columns of Q become a set of orthogonalized regression vectors β_v^\perp which define the task-related axes of choice, motion, colour, and context. The condition averaged population response (X) is then reorganized to form a set X_c of $N_{condition}$ matrices of size $N_{neuron} \times T$. Then we project X_c onto the task related axes (Equation 9) to allow for the visualization of the task related representations in ‘‘pfc’’.

$$p_{v,c} = (\beta_v^\perp)^T X_c \quad (9)$$

Neural comparison

Figure 4 shows the plots of the neural output of the model along the various task-related axes in two different contexts. The response trajectories start from a point in state space close to the center of the plots, which corresponds to the population response at the beginning. When the sensory input is provided, the response moves away from the center along the axis of choice. Movement of the population response along the choice axis is caused by integration of evidence while the sensory input is being provided to the model. Additionally, the population response moves away from the choice axis,

along the axis of motion/colour. The magnitude of the response along the motion/colour axes reflects the strength of the sensory evidence (motion/colour coherence), while its direction (up or down) reflects the sign of the sensory evidence.

Mante et al. (2013) obtained the neural data from two monkeys (monkey A and monkey F) trained to perform the same task. The trajectory for monkey A is shown in Figure 4 while the trajectory for monkey F can be found in the extended data of Mante et al. (2013) paper (Extended Data Figure 7). The response trajectories obtained for both monkeys varied significantly from each other. However the following three key features of population response were shared between the two monkeys:

- In both contexts, movement along the same choice axis corresponds to the integration of relevant evidence (Figure 4, plots a2 and f2).
- The motion evidence evokes response along the axis of motion which is different from the axis of choice in both contexts (Figure 4, plots a2 and d2).
- Motion evidence is strongly represented regardless of whether it is relevant (Figure 4, plot a2) or irrelevant (Figure 4, plot d2).

The response trajectories obtained from our model are very similar to the ones obtained from the experimental data. Importantly, they exhibit the three key features of the population response that were shared among the experimental data from the two monkeys. Plots a1 and f1 in Figure 4 show that the model trajectories show movement along the choice axis which corresponds to the integration of relevant evidence (i.e., motion in plot a1 and colour in plot f1). The movement along the choice axis in plots c1 and d1 is not as much because they plot the irrelevant sensory evidence in their respective contexts. Plots a1 (motion context) and d1 (colour context), both show that the motion evidence evokes response along the axis of motion which is different from the axis of choice. Moreover, these plots also show that motion evidence is strongly represented in both contexts regardless of whether it is relevant or not.

Figure 5 shows a comparison of our model to the learned recurrent network model of Mante et al. (2013). Note that the scaling on the choice axis in the figure is 1.5:1 for our model and the monkey data, and 6:1 for the reference model. The response trajectory obtained by our model is closer to the trajectory obtained from the experimental data. While our model is able to produce a curved trajectory as in the experimental data, the model of Mante et al. (2013) doesn’t produce a curved trajectory. This might be because the dynamics of their model doesn’t capture the fact that neural response in a given context is weakly dependent on irrelevant sensory input (Figure 3).

Conclusion and Future Direction

The main purpose of this work was to test a hypothesis (Equation 2) regarding the computations underlying context-

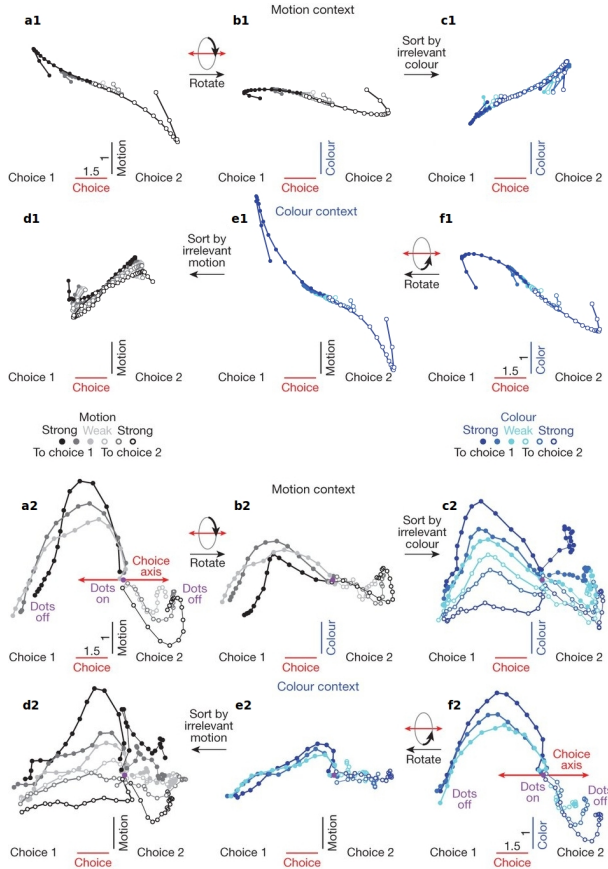


Figure 4: Neural Output of the Spiking Model (plots a1 - f1). Plots a2 - f2 show the neural response trajectories from the experimental data (monkey A). Only the responses from correct trials are plotted. Units are arbitrary. **a1, a2**: Effect of relevant motion input and choice in motion context where trials are sorted by the relevant sensory input (motion). **b1, b2**: Same data as in a1, a2 rotated by 90° about the choice axis to show the projection on to the colour axis. **c1, c2**: Same trials as in a1, a2 but resorted by irrelevant sensory input (colour). **d-f**: Responses plotted in the colour context, analogous to **a-c**. Display format taken from Mante et al. (2013).

dependent behaviour in the frontal eye fields of the prefrontal cortex (PFC). We have presented a spiking neural model to explain how neural and behavioural data can be captured by a simple nonlinear dynamical system. One limitation of the current work is that we directly stimulate the neurons based on the sensory inputs of the task without simulating the visual system required for the task. However, this should not have a major impact on the underlying computations in PFC.

In this model we simulated a task where the sensory information is constant within each trial and the decision is made only at the end of the trial. However, animals are able to change their behaviour in a constantly changing environment. The future work would involve investigation of neural mechanisms of such dynamic choices where the sensory information is constantly changing (even within a given trial) and a




Figure 5: Comparison of our model and the reference model to the experimentally recorded data for the motion context. Data is sorted by motion coherence and plotted on the choice and colour axes. Scaling on the choice axis is 1.5:1 for our model and the monkey data, and 6:1 for the reference model.

decision can be made at any time. One such experiment has been previously performed with monkeys where they were trained to solve the tokens task (Thura & Cisek, 2014). Building a scalable and robust model which can not only perform the task discussed in this paper, but also a more dynamic task like the tokens task can provide insights to more precisely explain the computational mechanisms underlying context dependent decision making in dynamic environments.

Acknowledgments

This work was supported by CFI and OIT infrastructure funding, the Canada Research Chairs program, NSERC Discovery grant 261453, ONR grant N000141310419, AFOSR grant FA8655-13-1-3084 and OGS graduate funding.

References

- Bekolay, T., Bergstra, J., Hunsberger, E., DeWolf, T., Stewart, T. C., Rasmussen, D., ... Eliasmith, C. (2014). Nengo: A python tool for building large-scale functional brain models. *Frontiers in Neuroinformatics*, 7(48). doi: 10.3389/fninf.2013.00048
- Eliasmith, C., & Anderson, C. H. (2003). *Neural engineering: Computation, representation, and dynamics in neurobiological systems*. MIT press.
- Flint, A. C., Maisch, U. S., Weishaupt, J. H., Kriegstein, A. R., & Monyer, H. (1997). Nr2a subunit expression shortens nmda receptor synaptic currents in developing neocortex. *The Journal of Neuroscience*, 17(7), 2469–2476.
- Mante, V., Sussillo, D., Shenoy, K. V., & Newsome, W. T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature*, 503(7474), 78–84.
- Schwartz, A. B., Kettner, R. E., & Georgopoulos, A. P. (1988). Primate motor cortex and free arm movements to visual targets in three-dimensional space. i. relations between single cell discharge and direction of movement. *The Journal of Neuroscience*, 8(8), 2913–2927.
- Tanji, J., & Hoshi, E. (2008). Role of the lateral prefrontal cortex in executive behavioral control. *Physiological reviews*, 88(1), 37–57.
- Thura, D., & Cisek, P. (2014). Deliberation and commitment in the premotor and primary motor cortex during dynamic decision making. *Neuron*, 81(6), 1401–1416.