

2021

Gain signal manifest in prestimulus neural population dynamics underlies decision-making

<https://hdl.handle.net/2144/43704>

Boston University

BOSTON UNIVERSITY
COLLEGE OF ENGINEERING

Thesis

**GAIN SIGNAL MANIFEST IN PRESTIMULUS NEURAL
POPULATION DYNAMICS UNDERLIES
DECISION-MAKING**

by

PIERRE-OLIVIER BOUCHER

Honours B.S., University of Toronto, 2011
M.A., University of Waterloo, 2014

Submitted in partial fulfillment of the
requirements for the degree of
Master of Science

2022

© 2022 by
PIERRE-OLIVIER BOUCHER
All rights reserved

Approved by

First Reader

Chandramouli Chandrasekaran, Ph.D.
Assistant Professor of Psychological and Brain Sciences
Boston University, College of Arts and Sciences

Assistant Professor of Anatomy and Neurobiology
Boston University, School of Medicine

Assistant Professor of Biomedical Engineering

Second Reader

Kamal Sen, Ph.D.
Associate Professor of Biomedical Engineering

Third Reader

Michael Economo, Ph.D.
Assistant Professor of Biomedical Engineering

Every mental representation of a movement awakens to some degree the actual movement which is its object.

-William James, 1890

Acknowledgments

I would like to thank first and foremost my mentor, Chand. When you first told me what urgency was I think I understood maybe 10% of what you said. Yet here we are a year later, a completed thesis and a publication in the works. Now I maybe understand about 30% of what you say!

Thank you for taking the time to explain difficult concepts and for challenging me to reach your level. Thank you for your honest and blunt feedback such as when you told me that I lost steam during my proposal. That was important feedback that I incorporated for my defense. Thank you for telling me my defense was procedural, I immediately incorporated a story to tie it all together and the defense improved significantly. Thank you for not letting me get away easily with lackluster explanations such as when you held my feet to the fire to explain KiNeT analysis. Overall thank you for being fair. You have been one of my best mentors and I have acquired several new skills under your tutelage.

To my funding and emotional support, the Boucher-Roberge foundation. You have always cheered my work and provided emotional support at times of acute academic distress.

Thank you next to my committee members Dr. Kamal Sen and Dr. Michael Economo for agreeing to be on my committee and providing valuable feedback. Kamal you asked me if I was going to continue working on this research as you found it very interesting. Thank you for asking that, that is perhaps the best thing you could have said. Thank you Mike for recommending that I reach out to Chand. Without your recommendation I don't think I would have discovered Chand's work.

Thank you to Chand lab members Kenji Lee, Tian Wang, Nicole Carr, and Karen Marmon for your generous and insightful feedback. You guys listened patiently to several of my presentations and our conversations helped me to think deeper about

the concepts. Thank you Tian for the work you have contributed in this thesis and for explaining how some of the analyses work such as LFADS and the linear and logistic regressions.

This undertaking would have been much more arduous without the love and support and encouragement from all you wonderful and amazing people.

To the reader, my hope is that this thesis and the planned publication will make impactful contributions to systems neuroscience. We do not fully understand urgency or its implications but I believe that it could have broad appeal throughout neuroscience outside of just decision-making alone.

Pierre-Olivier Boucher

Biomedical Engineering Department

**GAIN SIGNAL MANIFEST IN PRESTIMULUS NEURAL
POPULATION DYNAMICS UNDERLIES
DECISION-MAKING
PIERRE-OLIVIER BOUCHER**

ABSTRACT

Decision-making is thought to be shaped by factors such as “urgency” and past outcomes. However, the effect of such factors on neural population dynamics and its link to decision-making behavior is largely unclear. Here, we addressed this gap by investigating the neuronal population dynamics in a heterogeneous population of dorsal premotor cortex (PMd) neurons recorded from monkeys performing a red-green checkerboard discrimination task.

We investigated effects of urgency by analyzing firing rates of neurons organized by reaction time (RT) and choice. Dimensionality reduction, regression, and decoding analyses suggested that prestimulus neural population state covaried with RT but not choice. Critically, effects were observed within a stimulus difficulty. Subsequent analysis suggested that faster RTs involved faster pre- and post-stimulus dynamics whereas slower RTs involved slower dynamics. This relationship between prestimulus state and RT but not choice suggests a gain signal that amplifies the sensory evidence and modulates prestimulus activity rather than a bias to choose a particular side. Furthermore, errors on the previous trial led to shifts in prestimulus state and slower RTs, suggesting that this gain signal is modulated by trial history and linked to an internal speed-accuracy tradeoff.

We used drift diffusion and recurrent neural network (RNN) modeling to test if variability in behavior and neural activity represents fluctuations in a gain signal.

Drift diffusion models with a trial-by-trial varying multiplicative gain signal on the sensory evidence provided the best description of the RT and choice behavior. Similarly, in optimized recurrent neural networks, trial-by-trial variation in multiplicative gains on the rectified linear unit (ReLU) nonlinearity were necessary to recapitulate network dynamics consistent with our PMd data.

Collectively, these results suggest that a gain signal dependent upon previous trial outcome alters the prestimulus state and is an important component of decision-related neural population dynamics and behavior.

Contents

1	INTRODUCTION	1
2	RESULTS	8
2.1	Behavioral results	8
2.2	PMd neurons demonstrate complex time-varying patterns of firing rates	11
2.3	PCA analyses of firing rates suggest an organized covariation with coherence and choice	13
2.4	PCA analyses organized by choice and RT reveal a lawful covariation of prestimulus firing rates with RT	16
2.5	Baseline spiking activity explains RT variability but does not predict eventual choice	19
2.6	The outcome of the previous trial influences baseline spiking activity	22
2.7	Behavioral models that include a gain signal outperform standard DDMs	25
2.8	Recurrent neural network models that include either a multiplicative or additive gain signal recapitulate PMd neural dynamics	28
3	Discussion	31
3.1	Summary	31
3.2	Prestimulus neural variation consistent with a variable time-dependent and stimulus-independent gain signal	33

3.3	Previous findings and our results suggest urgency impinges upon neural variability of motor planning	36
3.4	Theory and Origins	38
3.5	Future Directions	39
3.6	Conclusions	41
4	Methods	42
4.1	Subjects	42
4.2	Apparatus	43
4.3	Task	43
4.4	Training	45
4.5	Electrophysiological Recordings	46
4.6	Unit Selection and Classification	47
4.7	Effects of coherence on accuracy and reaction time (RT)	48
4.8	Fits generated by linear regression and logistic regression (decoder) .	49
4.9	General principal component analysis (PCA) procedure of PMd firing rates	52
4.10	Behavioral Modeling	52
4.11	Calculation of AIC and BIC	55
4.12	Recurrent neural network (RNN) modeling description	57
4.13	Kinematic analysis of neural trajectories (KiNeT)	58
4.14	Description of Latent Factors Analysis of Dynamical Systems (LFADS)	59
	Curriculum Vitae	67

List of Figures

2.1	Recording locations, techniques, task, and discrimination behavior . . .	9
2.2	Heterogeneity of neuronal firing patterns	12
2.3	Principal components of neural activity organized by coherence separates on choice and level of coherence	14
2.4	Principal components of neural activity organized by RT reveals that baseline neural activity covaries with RT before stimulus onset	18
2.5	LFADS reveals neural states that separate prior to stimulus onset as a function of RT but not as a function of choice	21
2.6	Principal components of neural activity organized by outcome shifts in neural state associated with outcome	24
2.7	Behavioral models that include an urgency term predict decision behavior better than models without it	27
2.8	Recurrent neural network models with additive or multiplicative gain replicate PMd prestimulus neural dynamics	30

List of Abbreviations

AIC	Akaike Information Criterion
BIC	Bayesian Information Criterion
DDM	Drift Diffusion Model
FR	Firing Rate
fMRI	Functional Magnetic Resonance Imaging
ISI	Interspike Interval
KiNeT	Kinematic analysis of Neural Trajectories
LFADS	Latent Factor Analysis of Dynamical Systems
LL	Log-Likelihood
PC	Principal Component
PCA	Principal Component Analysis
PES	Post Error Slowing
PMd	Dorsal Premotor Cortex
R^2	Variance Explained
RNN	Recurrent Neural Network
RT	Reaction Time
SAT	Speed-Accuracy Tradeoff
SEM	Standard Error of the Mean
SMA	Supplementary Motor Area

Chapter 1

INTRODUCTION

There are 10 minutes to make it to the airport but the GPS says you're still 12 minutes away. You see a yellow in the distance and you quickly floor it. You get to the intersection only to realize you've run a red light. The sight of the yellow and red lights result in patterns of neural activity that respectively lead you to respond quickly to your environment and process feedback (e.g. slow down after running the red). This process of choosing and performing actions in response to sensory cues is termed perceptual decision-making (Brody & Hanks, 2016; Brunton et al., 2013; Cisek, 2012; Gold & Shadlen, 2007; Kiani et al., 2013). It is an integral part of everyday life and is impacted in many neurological and psychiatric disorders. The high-level objective of this thesis is to develop a better understanding of how internal state factors into decision-related dynamics and behavior.

A large body of research in multiple species has attempted to delineate the neural events that lead to a perceptual decision (Briggman et al., 2005; Carandini & Churchland, 2013; Cisek et al., 2009; Hanks & Summerfield, 2017; Roitman & Shadlen, 2002). Classically, much of the research into decision-making has focused on the activity of single neurons in animals performing decision-making tasks. In these studies, the most consistent observation is that there is a build up of firing rates of single neurons before a choice is made by the animal (Hanks et al., 2014; Hanks et al., 2015; Latimer et al., 2015; Pereira et al., 2021; Roitman & Shadlen, 2002; Shadlen & Newsome, 1996; Shadlen & Newsome, 2001). This buildup of neural activity has been interpreted as

a neural signature of the normative diffusion decision model (DDM) posited by Ratcliff and colleagues, which posits that external sensory evidence is integrated over time (see recent reviews by Roger Ratcliff and colleagues, e.g. (Ratcliff et al., 2016)). However, this narrative is challenged by the fact that neurons in decision-related structures are heterogeneous, with firing rates demonstrating complex time-varying dynamics over trials (Jun et al., 2010; Mante et al., 2013) and that neural firing rates can ‘ramp-up’ comparatively instantaneously in tasks where evidence for a decision is presented piece-meal rather than all at once (Thura & Cisek, 2014). Further work from Cisek and others have suggested that internal factors such as “urgency”, the unspecific urge to act, can also lead to such buildups of neural activity (Cisek et al., 2009; Thura & Cisek, 2014; Thura et al., 2014). How external sensory evidence and urgency converge to modulate decision-related activity of populations of neurons and ultimately decision-making behavior is an open question. The specific goal of this thesis is to resolve this open question and develop a deeper understanding of the neural population dynamics underlying a decision.

The primary focus of this thesis is how urgency affects prestimulus population neural state. Our overall hypothesis is that urgency directly acts as a gain signal that modulates the rate of rise of neural population activity associated with accumulating sensory evidence, therefore modulating neural states related to motor preparation and choice (Murphy et al., 2016). Specifically, these fluctuations in gain will manifest as changes in the prestimulus baseline firing rate (Murphy et al., 2016). This hypothesis is derived from three pieces of evidence: 1) studies that show that trial-averaged activity of single neurons correlated with a choice increases faster when evidence is strongly in favor of the choice (Roitman & Shadlen, 2002; Shadlen & Newsome, 2001), 2) studies that employed a speed-accuracy tradeoff and showed that baseline firing rates were different for fast compared to slow blocks (Hanks et al., 2014; Heitz &

Schall, 2012), and 3) studies that demonstrate that urgency acts as a multiplicative gain signal on the firing rates of neurons associated with an increasing urge to act (Thura et al., 2014). Urgency and its theoretical accounts are backed by a recently growing body of modeling (Cisek et al., 2009; Ditterich, 2006) and neurophysiological evidence derived across species (Carland et al., 2019; Murphy et al., 2016; Thura & Cisek, 2014).

Theoretical accounts such as the affordance competition hypothesis (Cisek, 2007) purports that competing evidence accumulating neuronal populations ramp up their activity towards fixed bounds with the ‘winning’ population state being the first to reach a bound (Cisek, 2012; Cisek et al., 2009; Roitman & Shadlen, 2002). Thus such a node becomes the effector state from which broader network activation derives, bringing forth the ‘chosen’ motor action plan (Cisek, 2012; Cisek et al., 2009). Urgency is a theorized evidence-independent yet time dependent ubiquitous gain signal on competing, winner-take-all, evidence accumulator neuronal population states (Murphy et al., 2016). Theory and modeling support the idea that such a signal speeds up accumulator population states and drives decision-making behavior to maximize risk/reward rate (Ditterich, 2006).

The DDM is currently a 60 year gold standard model for describing behaviors such as choice and reaction time (RT) from a wide range of decision-making experiments (Hawkins et al., 2015; Ratcliff, 1978; Stone, 1960). DDMs posit that in a binary decision there are two bounds towards which evidence accumulators grow over time as they collect sensory evidence. Once a bound is crossed a decision is made. Recent modifications to the DDM which include a term for elapsed time, urgency models, have challenged the standard DDM (Carland et al., 2019; Cisek et al., 2009; Ditterich, 2006; Hawkins et al., 2015). Importantly, experimental paradigms that support DDMs tend to feature a time-stationary stimulus for which DDMs were orig-

inally envisioned (Carland et al., 2019; Stone, 1960). However, when tasks include non-stationary stimuli that change over the course of a single trial, urgency models tend to outperform standard DDMs (Carland et al., 2019). In fact even for stationary stimuli, urgency models and DDMs generally perform equally well (Carland et al., 2019). While contentious (Carland et al., 2019; Hawkins et al., 2015) growing modeling work supports a time-sensitive, evidence-independent gain signal that modulates decision-making.

Several urgency models demonstrate parity in how successfully they model behavioral results but the time-dependent mechanism differs between models (Chandrasekaran & Hawkins, 2019). Collapsing bound models are similar to a standard DDM however the thresholds for making a binary decision approach one another as a function of elapsed time. Additive and multiplicative urgency models propose an additional sensory-independent, yet-time dependent gain signal to a fixed bound DDM that is either added to the sensory evidence accumulation process or acts as a gain directly on the evidence accumulation process, respectively (Chandrasekaran & Hawkins, 2019; Cisek et al., 2009; Ditterich, 2006). Either way as time passes the gain signal grows in size and less evidence is needed to make a decision. So the mechanism by which the brain implements speeded decision-making, whether by lowering the ‘threshold’ that must be reached for a decision to be made or if a gain signal is added to or directly modifies the evidence accumulation process, remains unclear.

Results from three functional magnetic resonance imaging (fMRI) studies on speed-accuracy tradeoff (SAT) in humans provide evidence for baseline modification rather than threshold adjustment in order to implement speeded decision-making (Bogacz et al., 2010). Consistently across these three studies they found increased prestimulus activity, in the SAT block as compared to the normal task block, in pre-supplementary motor area (SMA) and striatum. No baseline increases were de-

tected in primary sensory or motor areas (Bogacz et al., 2010). Interestingly, SMA and striatum are associated with motor planning and preparation, possibly suggesting increased priming of motor signals prior to stimulus onset (Bogacz et al., 2010). Additionally faster RTs are associated with faster rates of firing rate buildup, but these different trajectories reach similar thresholds before movement onset (Roitman & Shadlen, 2002; Shadlen & Kiani, 2013; Thura & Cisek, 2014). Consistent with this, preparatory neural dynamics are found to covary with RTs and explain the speed of motor response more so than ongoing motor activity (Afshar et al., 2011). Given evidence of similar neural activity thresholds before movement onset and increased baseline activity when speed is emphasized suggests that the brain employs fixed evidence thresholds with a variable urgency signal in order to resolve speeded decision-making tasks.

To summarize thus far, theory suggests a signal that uniformly boosts neural activity to resolve neural competition. Such a signal when included in decision making models outperforms the gold standard model. Additionally, human fMRI research on SAT supports an urgency model that includes a time-variant signal that boosts baseline neural activity when speed is emphasized and that originates from motor preparation areas. Again, our hypothesis states that a variable (Thura et al., 2017) and time-dependent gain signal directly modulates the population activity associated with accumulating sensory evidence; directly influencing motor preparation yet indirectly affecting choice (Thura & Cisek, 2017). In order to test this hypothesis we had 2 monkeys indicate the dominant color in a red-green checkerboard discrimination task with stimuli of varying difficulty. As monkeys performed the task, we recorded neural activity from the dorsal premotor cortex (PMd), a region generally associated with the guiding of reach behavior (Churchland et al., 2006; Cisek & Kalaska, 2005).

Several pieces of evidence from this experiment would further support such a

hypothesis. We predict that: 1. prestimulus neural state will covary with and be explanative of RTs but will not covary with or be predictive of choice, 2. neural states associated with faster reaction times will have relatively faster neural dynamics before and after stimulus onset, 3. outcomes of previous trials will modulate the prestimulus neural state 4. behavioral results and neural dynamics will be best explained by behavioral and neurophysiological modeling that includes a multiplicative urgency signal.

A summary of the results which are generally supportive of our hypothesis and predictions, follows. Psychometric curves and RT distributions reveal that decision accuracy increases and RTs decrease with decreasing stimulus difficulty. Decision-modulated neurons in PMd reveal heterogeneous and complex modulations of firing rates; not easily understood at the single neuron level. A principal component analysis (PCA) of mean population firing rate organized by choice and sensory evidence reveals lawful low-dimensional neural activity that separates as a function of choice and stimulus difficulty. An additional PCA now conditioned on choice and RT demonstrates prestimulus neural state covariation with RT. Importantly, this covariation is observed between and within levels of stimulus difficulty. Further, neural dynamics associated with faster RTs evolve faster than dynamics associated with slower RTs. Furthermore, a single-trial analysis method affirms a separation of neural activity, before stimulus onset, for slow versus fast reaction times but not choice. Regressions and decoders further confirmed that spiking activity prior to stimulus onset is explanative of RTs but not predictive of choice. A final PCA conditioned on choice and previous outcome revealed distinct prestimulus neural states after an error as compared to after a correct response. Additionally, logistic regression revealed that prestimulus spiking activity is predictive of the previous trials outcome. Finally, behavioral and recurrent neural network (RNN) modeling affirm that models that include either a variable

additive or multiplicative urgency term most faithfully recapitulate behavioral data and neural dynamics.

Deliverable analyses discussed as per the proposal for this thesis and included here are: 1. whether within stimulus difficulty RT variability covaries with prestimulus neural activity, 2. whether prestimulus spiking activity is explanative of RT or predicts choice, 3. whether changes in neural state are driven by previous outcomes. Findings from the RNN modeling and single trial analysis reported here were completed by my colleague Tian Wang. Analyses on the speed of neural trajectories and behavioral modeling were performed by my mentor Dr. Chandrasekaran. This work was recently submitted in abstract form to a computational neuroscience conference (Cosyne).

Chapter 2

RESULTS

2.1 Behavioral results

The data analyzed in this proposal was collected at Stanford University by my thesis mentor, Dr. Chandrasekaran. The dataset was obtained from two macaque monkeys (O. and T.) trained to discriminate the dominant color of a central, static checkerboard composed of red and green squares and to report their decisions with arm movements (Fig 1A). Figure 1B depicts the trial timeline. The trial began when the monkey held the center target and fixated on the fixation cross. After a jittered holding time, a red and a green target appeared on either side of the central hold (target configurations were randomized). After an additional randomized target viewing time, the checkerboard appeared. The monkey’s task was to select the appropriate red or green target matching the dominant color of the checkerboard.

There were seven levels of stimulus difficulty parametrized by unsigned coherence (C , Figure 1C). The measure of color coherence, independent of the actual dominant color, is the absolute difference in the number of red (R) and green (G) squares normalized by the total number of squares in the checkerboard ($C = 100 \times |R - G| / (R + G)$). Signed color coherence (SC), dependent on the actual dominant color of the checkerboard, is defined as $SC = 100 \times (R - G) / (R + G)$. While animals performed the task, arm and eye movements were recorded. RTs were identified as the first time when hand velocity exceeded 10% of maximum velocity during a reach.

We analyzed the behavioral performance of the monkeys as a function of the dif-

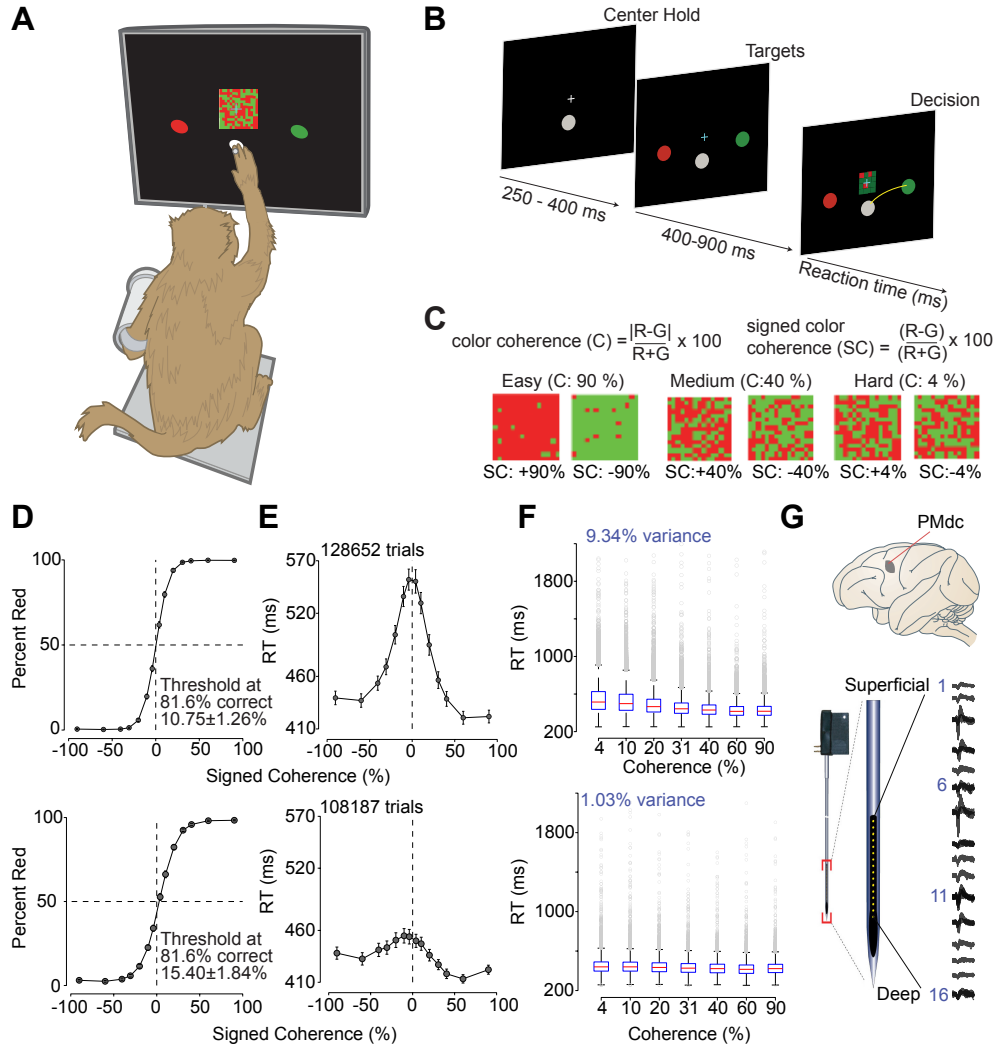


Figure 2-1: **Recording locations, techniques, task, and discrimination behavior:** (A) An illustration of the setup for discrimination. We loosely restrained the arm the monkey was not using with a plastic tube and cloth sling and taped a reflective infrared bead on the middle digit of the other hand and tracked it in 3D. We used the measured hand position to mimic a touch screen and to provide an estimate of instantaneous arm position; eye position was tracked using an infrared reflective mirror placed in front of the monkey's nose. (B) Timeline of the discrimination task. (C) Examples of different stimulus ambiguities used in the experiment parameterized by the color coherence of the checkerboard cue defined as $C = 100 \times |R - G| / (R + G)$. The corresponding SC is defined as $SC = 100 \times (R - G) / (R + G)$. Positive values of SC denote more red than green squares and vice versa. (D, E) Average discrimination performance (D) and RT (E) over sessions of the two monkeys as a function of the SC of the checkerboard cue. RT plotted here includes both correct and incorrect trials for each session and then averaged across sessions. Gray markers show measured data points along with $2 \times SEM$ estimated over sessions. The black line segments are drawn in between these measured data points to guide the eye and for many data points in d, the error bars lie within the marker. X-axes in both d, e depict the SC in %. Y-axes depict the percent responded red in d and RT in e. Also shown in d are discrimination thresholds ($M \pm SD$ over sessions) estimated from a Weibull fit to the overall percent correct as a function of coherence. The discrimination threshold is the color coherence level at which the monkey made 81.6% correct choices. 75 sessions for monkey T and 66 sessions for monkey O went into the averages.

Figure 2-1: **(F)** Box-and-whisker plot of RT as a function of unsigned checkerboard coherence. For each coherence, the central mark of the box indicates the median, the bottom and top edges of the box reflect the 25th and 75th percentiles, respectively, the maximum whisker length is specified as three times the interquartile range, and outliers are plotted as plus symbols. There is large variation of RTs both across and within coherences. **(G)** Location of caudal PMd (PMdc, location of electrode implantation) along with an example recording from a 16 electrode, 150- μm spacing U-probe. The brain in this figure is adapted from (Ghazanfar & Santos, 2004).

ficuity of the task (i.e. color coherence). In general across all sessions, the closer a checkerboard was to a 50/50 split the more errors were made (Fig 1D). We fit the proportion correct as a function of unsigned coherence using a Weibull function to estimate slopes and psychometric thresholds (average R^2 ; T=0.99 (over 75 sessions); O: 0.99 (over 66 sessions); slope (β , $M \pm \text{SD}$ over sessions, T: 1.30 ± 0.16 , O: 1.22 ± 0.16). Monkey T was better than monkey O in discriminating the color of the checkerboards (thresholds are computed on a per session basis and averaged over sessions at 81.6% correct, ($M \pm \text{SD}$): T, $10.77 \pm 1.26\%$, O: $15.42 \pm 1.87\%$, Wilcoxon rank sum comparing median thresholds, $p = 2.91 \times 10^{-23}$). While there is more RT variability and monkeys were generally slower for more ambiguous checkerboards, as the box plots in Fig 1F show, the spread between RTs is still quite large even for the easiest coherences. A linear regression between unsigned coherence and RT bolstered this qualitative observation. In Monkey T, coherence, which is a measure of the external sensory input, explained only $\sim 10\%$ of the reaction time variability, and in monkey O only $\sim 1\%$ of the reaction time variability. It should additionally be noted that there is no speed versus accuracy manipulation in this simple behavioral paradigm. Therefore, the remaining variability in RT is likely due to internal factors inherent in the animal and not driven by the task. Our hypothesis is that one of these internal factors is urgency.

2.2 PMd neurons demonstrate complex time-varying patterns of firing rates

Our database consisted of 996 units (546 units in T and 450 units in O, including both single neurons and multi units) recorded from PMd of the two monkeys over 141 sessions. Spike sorting was used to extract spiking activity from single units. Chosen units were included as they were well isolated from other units/separated from noise and modulated activity in at least one task epoch.

To analyze neural firing rates, we convolved spike times with a 20 ms wide Gaussian kernel for each trial in specific time windows. Checkerboard cue (time marker for stimulus presentation) aligned firing rates were plotted 100 ms before the onset of the stimulus and until the median RT for the color coherence, up to 600 ms after (Fig 2A). For movement aligned firing rates, spike times were plotted until the negative of the median RT for the color coherence up to 600 ms preceding movement onset, and 100 ms post movement onset (Fig 2B). Generally the lower the level of coherence, the longer the firing rate plot. Trials were averaged per unit and within coherence to visualize firing rate variability across neurons, coherence levels, and reach directions (Fig 2).

Figure 2 shows several examples of single neurons and the heterogeneity inherent in our dataset. While many neurons showed reassuring and classical patterns of firing rates documented by many researchers (Hanks et al., 2014; Latimer et al., 2015; Roitman & Shadlen, 2002; Shadlen & Newsome, 1996; Shadlen & Newsome, 2001), many other neurons showed heterogeneous and complex patterns of activity that included both increases and decreases in firing rate during the task. Note, these neurons still demonstrated covariation with coherence, choice, and RT suggesting that they were likely still a component of the decision formation process.

Canonically the firing rate activity, of neurons tuned to a preferred reach direc-

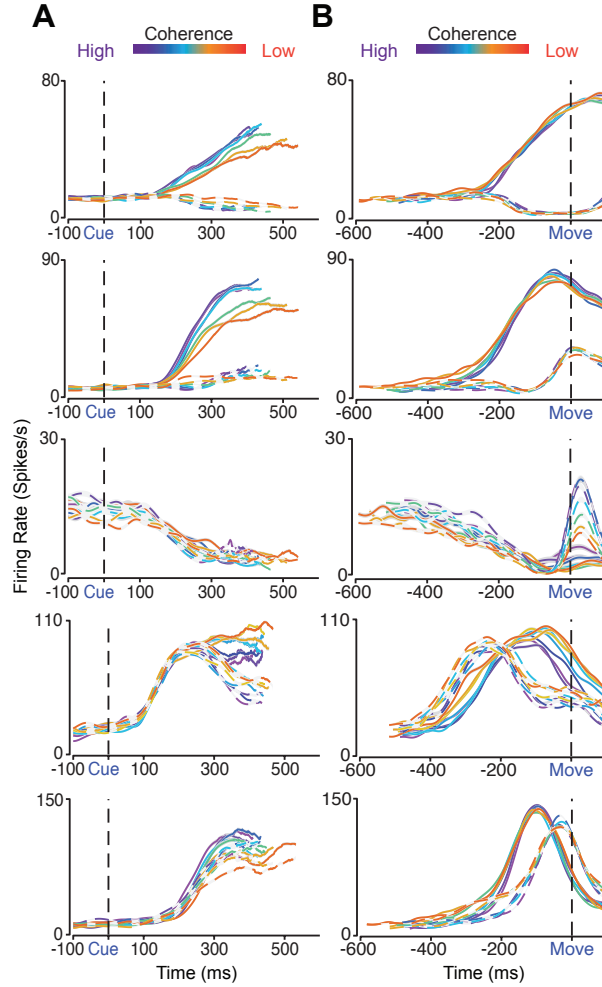


Figure 2-2: **Heterogeneity of neuronal firing patterns** (A) Firing rate (FR) activity across 7 levels of color coherence and both reach directions of 5 example single units in PMd from monkeys T and O aligned to stimulus onset (Cue). FR is plotted until the median RT for the color coherence (notice slightly different lengths of firing rate lines). (B) FR activity across 7 levels of color coherence and both reach directions of the same 5 single units from monkeys T and O aligned to movement onset (Move). FR is plotted until the negative of median RT for the color coherence. Spiking activity was convolved with a 20 ms wide Gaussian kernel for each trial in order to generate FRs. The gradation of color indicates the level of difficulty for the coherence from purple being the easiest coherence (mostly one color) to red being the most difficult coherence (a nearly 50/50 split in color). Different color dashed lines represent neuronal firing rates associated with a left reach whereas different color solid lines represent neuronal firing rates associated with a right reach. The black dashed line in each plot indicates either the onset of the stimulus or movement initiation. Error bars (gray outline visible around some lines) denote $2 \times SEM$.

tion, ramps up over time post stimulus presentation until some time around movement initiation. These neurons typically have little to no change in activity for the non-preferred reach direction (Roitman & Shadlen, 2002). While such neural activity is evident (Fig 2A, B, first two rows), single neuron activity presents in a multitude of ways post stimulus presentation (Fig 2A,B, rows 3-5). Therefore the firing rate activity of neurons involved in decision making generally comprise a complex, high dimensional array of responses. For example, the neuron shown in Fig. 2A (3rd row) exhibits a near constant firing rate post stimulus presentation and a rapid decrease in firing rate 200 ms prior to movement onset for both reach directions (Fig 2A, B, row 3). Additionally, different profiles of variable activity related to each reach direction can exist within one neuron. One such neuron ramps up over time post stimulus presentation for both reach directions, with these patterns being shifted in time from one another (Fig 2A, B, 5th row). In summary, this analysis of single neuron activity suggests that this PMd population is heterogeneous and that choice, coherence, and RT are encoded across this population. In the next section, we used PCA to examine how coherence, choice, and RT were represented in the shared activity of these neurons.

2.3 PCA analyses of firing rates suggest an organized covariation with coherence and choice

To make sense of time-varying firing rates and extract signals associated with decision-making in this heterogeneous population, I turned to PCA. I performed a PCA on trial-averaged firing rates conditioned on coherence and choice, and another conditioned on RT and choice. For each PCA, I created a single matrix of all 996 units and their average firing rate activity windowed about checkerboard onset and organized by level of condition within a reach direction. More formally, the windowed, trial

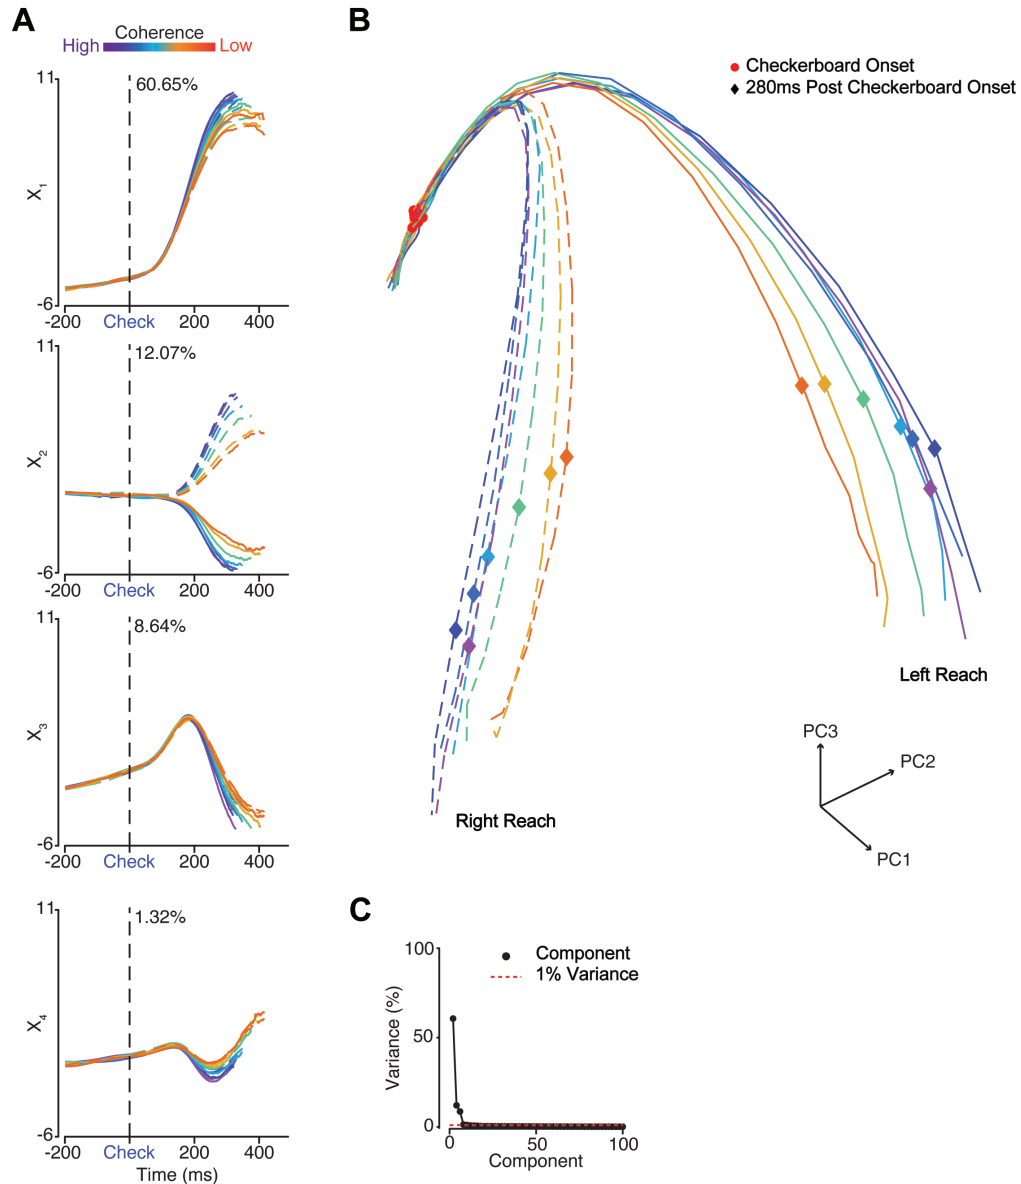


Figure 2-3: **Principal components of neural activity organized by coherence separates on choice and level of coherence (A)** The first 4 PCs ($X_{1,2,3,4}$) of trial averaged firing rates of all 996 neurons from monkeys T & O and all sessions organized across 7 levels of color coherence, both reach directions, and aligned to checkerboard onset. Spiking activity was convolved with a 20 ms wide Gaussian kernel for each trial in order to generate FRs. The gradation of color indicates the level of difficulty for the coherence from purple being the easiest coherence (mostly one color) to red being the most difficult coherence (a nearly 50/50 split in color). Dashed lines of different colors represent the PC associated with a right reach whereas solid lines of varying colors represent the PC associated with a left reach. The black dashed lines indicate the onset of the stimulus and the number to the right of the line indicates the amount of variance explained by that PC. **(B)** 3D structure of the first 3 PCs ($X_{1,2,3}$) aligned to checkerboard onset. PC 1/2/3 - principal component 1/2/3. Observe how neural activity separates as a function of choice and coherence < 200 ms post stimulus onset. Red circles indicate checkerboard onset. Notice that the plotting of PCs extends 200 ms before checkerboard onset. Different colored diamonds indicate 280 ms post checkerboard onset. Again the gradation of color indicates the level of difficulty for the coherence from purple being the easiest coherence (mostly one color) to red being the most difficult coherence (a nearly 50/50 split in color). Dashed lines of different colors represent the PC associated with a right reach whereas solid lines of varying colors represent the PC associated with a left reach.

Figure 2-3: (C) Scree plot of the percentage of variance explained by the first 50 components. Black dots represent each component plotted in order of the amount of variance explained. The red dotted line denotes 1% of variance and black dots that fall below this explain less than 1% of variance. The first 3 PCs already capture > 80% of the variance in the firing rate activity, with the 4th component only explaining just over 1% of the variance.

averaged, firing rate activity within a neuron for each condition level were vertically concatenated into a $\text{ctr} \times \text{N}$ matrix, combining all condition levels (c), windowed time per level of condition (t), reach directions (r), and neurons (N) into one matrix. The raw data was preprocessed by subtracting the mean of each column and then dividing by the 95th percentile of that column (i.e. softmax). We entered the data into the `pca` function in MATLAB (The MathWorks, Inc., Natick, MA, United States). Essentially the eigenvalues and eigenvectors of the covariance matrix for the $\text{ctr} \times \text{N}$ matrix are calculated such that the raw data can be projected along a single or multiple principal component/s.

To understand how neural activity was organized by coherence and choice, PCA was conducted first with 7 levels of coherence data across all 996 neurons (Fig 3). The first principal component (PC) explains $\sim 50\%$ of variance and appears to represent the change in firing rates over time post stimulus onset (Fig 3A, panel 1). The second PC explains $\sim 15\%$ of variance and separates clearly based on choice and then on stimulus difficulty within a choice at ~ 200 ms post stimulus onset (Fig 3A, panel 2). The first 4 PCs explain $> 80\%$ of the variance (Fig 3C). I plotted the first three PCs in a state space, and found that activity separates as a function of choice as well as coherence (Fig 3B). In this three dimensional space, activity separates faster for easier compared to harder choices. This result is reassuring as it recapitulates the findings from classical studies which demonstrated such features in selected single neurons (Ding & Gold, 2010; Roitman & Shadlen, 2002; Shadlen & Newsome, 2001). Our results here replicate this finding at the population level.

2.4 PCA analyses organized by choice and RT reveal a lawful covariation of prestimulus firing rates with RT

Organization by coherence and choice is reassuring, but does not provide insight into mechanism. I conducted a second PCA with the same neural data but grouped by RT bins, 11 levels representing a spectrum from faster to slower RTs (300-400 ms, to 600-1000 ms), and choice. The first 4 PCs explain $> 85\%$ of the variance (Fig 4E). The first PC again appears to represent the change in firing rates over time, post stimulus onset, accounting for $> 50\%$ of the variance (Fig 4A, panel 1). The second PC explaining $\sim 17\%$ of the variance, separates by choice and RT with choice signals emerging ~ 200 ms post stimulus, (Fig 4A, panel 2). These first two PCs appear remarkably similar to the first two PCs for the PCA conditioned on coherence (Fig 3A & 4A, panels 1 & 2). The third PC, explaining $\sim 11\%$ of the variance, separates on RT binning pre/post stimulus onset and separates by choice strongly ~ 200 ms post stimulus onset (Fig 4A, panel 3). The 4th PC accounting for $\sim 2.5\%$ of the variance separates, quite noticeably, by RT prior to stimulus onset and the separation persists until ~ 300 ms post stimulus onset, (Fig 4A, panel 4). Thus, multiple PCs (1,3, and 4), and perhaps suggestively not PC 2 (also see Fig 4D), separate prestimulus on RT grouping in the PCA done for RT bins, yet none of the PCs in the coherence PCA separate by coherence prior to onset of the stimulus. This is indicative of a prestimulus state that covaries with the monkeys' RT and is likely independent of the presented stimulus and the eventual choice.

A kinematic analysis of neural trajectories (KiNeT) (Remington et al., 2018) reveals that faster RTs involved faster pre- and post-stimulus dynamics whereas slower RTs involved slower dynamics as compared to a reference trajectory (trajectory associated with middle RT bin, cyan) (Fig 4C, top). Trajectories were also found to be spatially organized by RT (Fig 4C, bottom). The reference trajectory can be thought

of as a unit line, in that all other trajectories are interpreted in relation to it. Therefore trajectories that are below the reference trajectory in the speed analysis are faster than the reference in that they reached a point in the Euclidean space at an earlier time point than the reference trajectory. Given the design of this decision-making task on static stimuli one could reasonably predict that RT variability should be mostly driven by the amount of stimulus evidence. Yet there's a considerable amount of RT variability even within a stimulus difficulty and stimulus evidence accounts for maximally 10% of the variance. So, what could account for RT variability within a stimulus difficulty? Given we observe baseline neural state covariation with RT across coherences, it stands to reason that we may observe the same type of covariation within a stimulus difficulty. We propose that baseline neural activity driven by urgency can partially account for the RT variability that we observe within a stimulus coherence.

To address this question I performed seven PCAs, one for each coherence level, on trials selected within a stimulus difficulty to examine if covariation with baseline state and RT is also observed within a stimulus difficulty. The matrices were organized $RT_{tr} \times N$. RT_{tr} - windowed trials within a single coherence and separated on 11 reaction time bins, r - reach direction and N - neurons. We found that baseline neural activity separates by RT bin in a similar fashion to the previous PCA that was done with RT and across all 7 levels of coherence (Fig 4B, inset). We found similar levels of covariance between prestimulus neural state and RT variability for stimuli of all coherence levels (compare Fig 4B & Fig 4B, inset). This suggests that baseline neural state can partially account for variability in RT within a stimulus coherence and in the absence of sensory evidence. This implies that neural populations prime the decision-making process prior to any sensory evidence.

My preliminary analyses described above and in Figures 1-4 have provided three

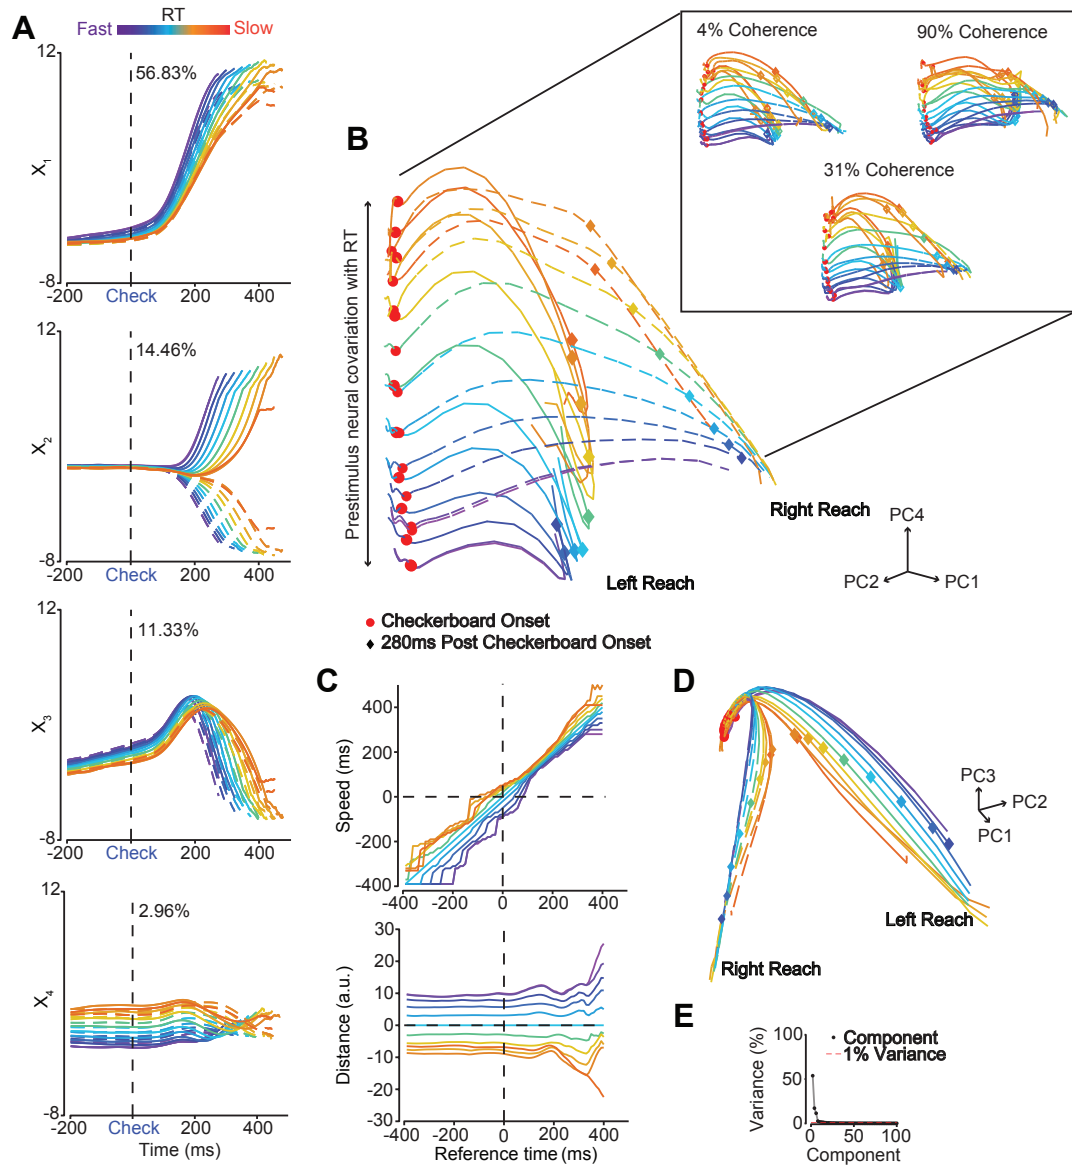


Figure 2-4: Principal components of neural activity organized by RT reveals that baseline neural activity covaries with RT before stimulus onset (A) The first 4 PCs ($X_{1,2,3,4}$) of trial averaged firing rates of all 996 neurons from monkeys T & O and all sessions organized across 11 RT bins, both reach directions, and aligned to checkerboard onset. Spiking activity was convolved with a 20 ms wide Gaussian kernel for each trial in order to generate FRs. The gradation of color indicates the length of RT with purple corresponding to faster RTs to red being the longest RTs. Dashed lines of different colors represent the PC associated with a right reach whereas solid lines of varying colors represent the PC associated with a left reach. The black dashed lines indicate the onset of the stimulus and the number to the right of the line indicates the amount of variance explained by that PC. **(B)** 3D structure of the 1st, 2nd and 4th PCs ($X_{1,2,4}$) aligned to checkerboard onset. PC 1/2/4- principal component 1/2/4. Red circles indicate checkerboard onset. Notice that the plotting of PCs extends 200 ms before checkerboard onset. Observe how neural activity separates as a function of RT bin, but not by choice, up to 200 ms before stimulus onset. Different colored diamonds indicate 280 ms post checkerboard onset. Inset: 3D structure of the 1st, 2nd and 4th PCs ($X_{1,2,4}$) for a subset of seven PCAs conditioned on RT bins and reach but performed within a single stimulus coherence. **(C)** "KiNeT" speed (top) and distance (bottom) analyses respectively showing that pre- and post-stimulus speed is faster for faster RTs as compared to a reference trajectory (cyan, middle trajectory) and that trajectories are spatially organized by reference time. *Abbreviation:* a.u. - arbitrary units.

Figure 2-4: (D) 3D structure of the first 3 PCs aligned to checkerboard onset. PC 1/2/3 - principal component 1/2/3. Remark how neural activity separates as a function of choice and coherence < 200 ms post stimulus onset (similar structure to Fig 3B). (E) Scree plot of the percentage of variance explained by the first 50 components. Black dots represent each component plotted in order of the amount of variance explained. The red dotted line denotes 1% of variance and black dots that fall below this explain less than 1% of variance. The first 4 PCs already capture > 85% of the variance in the firing rate activity, with the 5th component only explaining just over 1% of the variance.

tantalizing findings. First, RTs in monkeys performing decision-making tasks are highly variable perhaps due to internal factors such as “urgency”. Second, a heterogeneous population of neurons in PMd covary with choice, coherence and RT. Third, even before the stimulus onset, firing rates covary with RT suggesting that they could potentially be linked to the RT variability observed in behavioral results and could be an index of the internal factors that drive decision-making behavior. Below, I build on these findings to further understand how internal factors such as urgency, and external sensory input interact to drive decision-making neural responses and ultimately decision-making behavior.

2.5 Baseline spiking activity explains RT variability but does not predict eventual choice

An urgency model predicts and two PCAs suggest that prestimulus FR covaries with RT but not choice. We expect that more RT variance will be explained by the baseline spiking activity of a small subset of neurons (e.g. 20-30), from a single brain area, in a single session than by sensory evidence. However, predicting eventual choice using baseline spiking activity will be at chance level either in a single session or across all sessions. These results would be consistent with urgency and the hypothesis that pre-stimulus activity covaries with RT but not choice.

To determine whether spiking activity explains variance in RTs a linear regression, using binned spiking activity as the predictor and RT as the outcome, was performed. 1800 ms of spiking activity from each trial (600 ms prestimulus and 1200 ms post-stimulus) were divided into 20 ms bins (90 bins total) across all neurons within a

session. 90 matrices composed of a single time interval of spike counts (e.g. 0-20 ms) for all trials, across all neurons were linearly regressed with trial matched RTs. These fits were used to predict RTs per bin of spiking activity. These predicted RTs were used to calculate R^2 (variance explained) values per bin. The per bin R^2 values were then averaged across sessions. R^2 values were compared to the 1st and 99th percentile of R^2 values calculated from regressions of trial shuffled spiking activity and RTs (shuffled 500 times). To determine whether spiking activity can predict eventual choice a logistic regression, using binned spiking activity as the predictor and choice as the outcome was performed. The spiking activity was binned as before and 90 logistic regressions (1/bin) were computed in order to predict choice based on spiking activity. Accuracy per bin was calculated from the predictions of the logistic regression model. The per bin accuracy was then averaged across sessions and compared to the 99th percentile of accuracy of logistic regression models built on 500 shuffles of binned spiking activity (i.e. choices are randomly paired with spiking activity).

Spiking activity for each trial was regressed with corresponding RTs and the R^2 was calculated. Baseline spiking activity explains $\sim 25\%$ of the variance in RT in a single session (Fig 5B) and $\sim 15\%$ of the variance on average across all sessions (Fig 5C). This is considerably more variance explained than the 99th percentile of trial shuffled spiking activity regressed with random RTs (Fig 5B, C). Recall that sensory evidence was able to account for $\sim 10\%$ of the variance in RTs in monkey T and only $\sim 1\%$ of variance in RTs in monkey O. Here with at most 30 neurons, from a single brain area, in a single session prestimulus spiking activity explains more of the RT variability on average ($\sim 15\%$) than the actual sensory evidence. However, predicting eventual choice using baseline spiking activity is at chance level either in a single session (Fig 5F) or across all sessions (Fig 5G). These results are consistent with

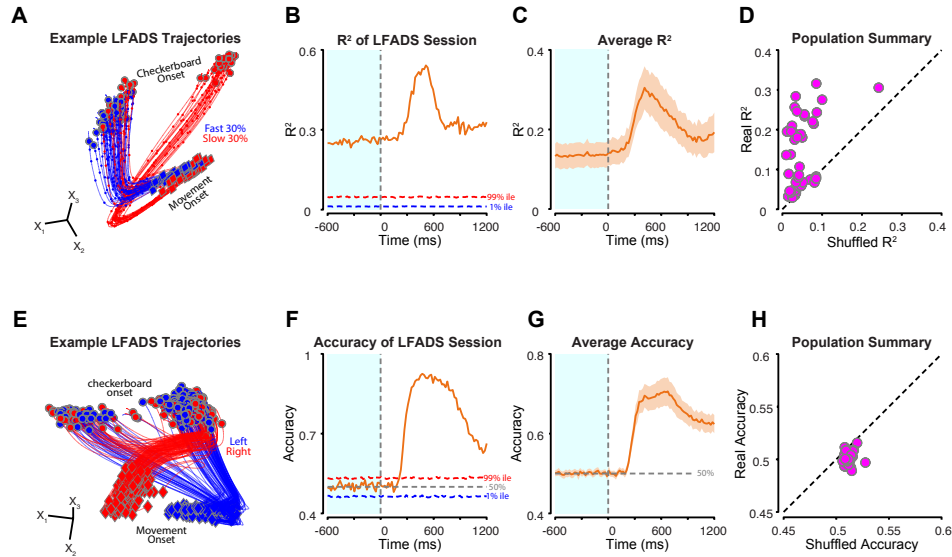


Figure 2-5: LFADS reveals neural states that separate prior to stimulus onset as a function of RT but not as a function of choice (A) LFADS trajectory plotted for the fastest 30% of trials and the slowest 30% of trials within the easiest coherence from a single session. Each trajectory is plotted from 200 ms before the checkerboard onset (large circles with gray border) to movement onset (diamonds with gray border). Small dots in the trajectories are the position of the first 3 latent factors every 50 ms. (B) Plot of R^2 values from 90, 20 ms bins calculated from regressions of each 20 ms bin of spiking activity, from all neurons in the LFADS session and easiest coherence, and RTs. The blue and red dotted lines represent the 1st (1% ile) and 99 (99% ile) percentiles of R^2 values calculated from averaged regressions of trial shuffled spiking activity and RTs (shuffled 500 times). (C) Plot of R^2 values, calculated as in (B), averaged across 51 sessions. Orange shaded area is the 99% confidence interval of the mean ($2.58 \times SEM$). (D) Scatterplot of the 99th percentile of the average shuffled R^2 values by the real average R^2 values from before stimulus onset (0-600 ms). Dotted line is where scatter points would fall if shuffled R^2 and real R^2 values were equivalent. Many of the points lie above this line suggesting that real prestimulus neural activity explains more of the RT variance than shuffled neural data. (E) LFADS trajectory plotted for left and right reaches within the easiest coherence from a single session. Each trajectory is plotted from 200 ms before the checkerboard onset (large circles with gray border) to movement onset (diamonds with gray border). (F) Spiking activity, from the same session that was used for LFADS, binned in 20 ms intervals was used to build a logistic regression model to predict eventual choice. Accuracy was calculated and plotted per 20 ms bin from the predictions of the logistic regression model. Also, all binned spiking activity was shuffled between trials such that choices were randomly paired with spiking activity. This was performed 500 times and logistic models and accuracy were calculated from each shuffle. The 1st and 99th percentile of the shuffled data were plotted as the blue and red dotted line respectively. (G) Logistic regression of spiking activity and choice was computed for each bin for all trials of the easiest coherence in a session. The mean accuracy of the models was calculated by bin across 51 sessions and plotted. Orange shaded area is the 99% confidence interval of the mean ($2.58 \times SEM$). (H) Scatterplot of the first 30 RT bins (20 ms each, 0-600 ms before checkerboard onset) by 99th percentile of shuffled accuracy on x-axis and real accuracy on y-axis. Dotted line is where scatter points would fall if shuffled and real accuracy values were equivalent. Many of the points lie on or below this line suggesting that real prestimulus neural activity is not predictive of choice.

urgency and the idea that prestimulus activity covaries with RT but not choice. Such a finding further bolsters our hypothesis that urgency modulates decision outcomes even before the stimulus is presented.

PCA results were performed on trial averaged firing rates. However, trial-averaging can obscure different causes of RT variability. In particular, RT variability could be caused by differences in speed of neural dynamics, further distance of starting point of the neural dynamics, or both. To better understand how decision-related dynamics evolve on single trials, we performed a single-trial analysis on simultaneously collected neural data from PMd. We used Latent Factor Analysis of Dynamical Systems (LFADS) (Pandarinath et al., 2018) to visualize single trial dynamics. Preliminary evidence from such an analysis (Fig 5A) demonstrates that fast and slow RTs differentiate in neural state space 200 ms prior to stimulus onset at the single trial level. Note, some of the initial states for slow RTs are also mixed in with fast RTs. Consistent, with the decoding analyses, initial neural states related to left and right reach are mixed prior to stimulus onset (Fig 5E). These results further bolster the notion that baseline spiking activity covaries with RTs but not choice. These results further bolster the notion on a trial-by-trial basis that faster RTs are associated with separate prestimulus neural states whereas different eventual choices are represented in similar states prestimulus.

2.6 The outcome of the previous trial influences baseline spiking activity

Evidence from this study and previous work (Murphy et al., 2016) demonstrate that baseline spiking activity is explanative of RT variability. So we became curious as to what could contribute to neural populations varying from trial-to-trial prior to any sensory evidence. One source of prestimulus neural variation could be post-error

slowing (PES) which is a behavioral phenomenon where RT for a trial following an error trial is slower than the RT following a correct response (Dutilh et al., 2012). We hypothesize that given that PES causes RT variability, that this variability would also be accompanied by shifts in trial-by-trial neural state. Hypothetically, such neural fluctuations could lead to a neural state in which the participant is slower to respond after an error. So we set out to demonstrate that post-error neural states were distinct from the neural state prior to an error or after a correct response.

First we performed a behavioral analysis to confirm that PES occurred. All trial sequences of correct, error, correct and their associated RTs were aggregated across both monkeys and all sessions (Figure 6D). All sequences with a RT < 200 ms were removed from the analysis (0.3% of trials). Correct trials following an error were significantly slower than correct trials immediately preceding the error trial ($M \pm SD$: 468 ± 118 ms, 449 ± 105 ms; Wilcoxon rank sum comparing median RTs, $p=4.99e-97$).

I performed a PCA on trial-averaged FRs from all 996 units conditioned on choice and trial outcome to assess whether errors led to a shift in the prestimulus neural state in the following trial. The 4D matrix was organized $t \times N \times r \times O$: t - windowed time per trial, N -neurons, r - reach direction, and O - outcome. The PCA revealed, reassuringly, the first three principal components we have come to expect. In other words the first PC is a change in FRs before a decision is made, the third PC is associated with choice and the second PC matches the third PC of the first two PCAs (Fig 6A, top three panels). Notably for this particular analysis along the fourth PC we observe a considerable separation between the post error trial and all other trial types. This separation in neural state space (Fig 6B) occurs before stimulus onset and endures at least 400 ms past checkerboard onset (Fig 6A, 4th panel). Such a result suggests that after an error there is a considerable change in neural state space

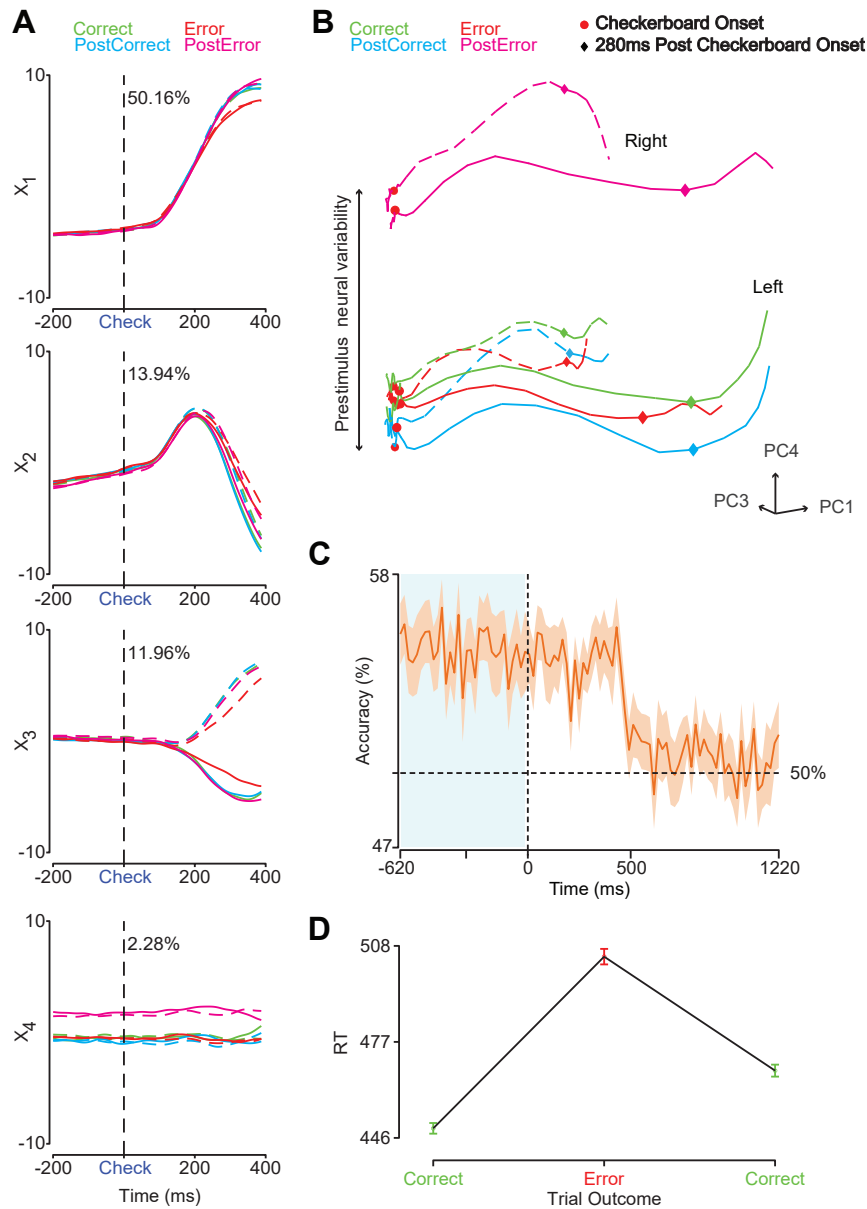


Figure 2-6: Principal components of neural activity organized by outcome shifts in neural state associated with outcome (A) The first 4 PCs ($X_{1,2,3,4}$) of trial averaged firing rates aligned to checkerboard onset of all 996 neurons from monkeys T & O and all sessions organized by choice and trial outcome. Spiking activity was convolved with a 20 ms wide Gaussian kernel for each trial in order to generate FRs. The four colors indicate different trial types with black indicating a correct trial, green a correct trial following a correct trial, red - error trial, and magenta - trial following an error trial. Solid and dashed lines of different colors indicate different reach directions. The black dashed lines indicate the onset of the stimulus and the number to the right of the line indicates the amount of variance explained by that PC. **(B)** 3D structure of the 1st, 2nd and 4th PCs ($X_{1,2,4}$) aligned to checkerboard onset. PC 1/2/4 - principal component 1/2/4. Red circles indicate checkerboard onset. Notice that the plotting of PCs extends 200 ms before checkerboard onset. Observe how neural activity separates as a function of outcome, but not by choice, up to 200 ms before stimulus onset. Different colored diamonds indicate 280 ms post checkerboard onset. **(C)** Accuracy of logistic regression of spiking activity from the current trial used to predict the outcome of the previous trial. Orange outline is the 99% confidence interval. **(D)** Average RTs from all correct, error, correct sequences found across both monkeys and all sessions. Error bars are $2 \times SEM$.

which could potentially explain why RTs after an error tend to be slower than RTs after a correct response. Such a result is also significant as it indicates that baseline activity is affected by the results of previous trials and it indicates that there are baseline fluctuations that covary with RT at all. It would suggest that these neural populations are shifting their activity proactively in order to account for prolonged deliberation and/or as a strategy to maximize reward rate.

A decoder was built to determine whether spiking activity from the current trial could predict previous trial outcome. We found that current trial spiking activity predicts the previous trial's outcome from before stimulus onset to ~ 500 ms after checkerboard onset at higher than chance levels (Fig 6C). ~ 500 ms post-checkerboard onset, similar to the overall mean response time, spiking activity predicts the outcome of the previous trial at chance levels. Overall these findings suggest that the outcome of the previous trial can influence the spiking activity of the current trial. Such a result further bolsters the claim of a trial-by-trial variation in neural activity and demonstrates a source of variance to prestimulus neural dynamics.

2.7 Behavioral models that include a gain signal outperform standard DDMs

Consistent with a stimulus-independent gain signal, prestimulus neural activity covaries with and is explanative of RT but does not predict choice. We predicted and found evidence consistent with the fact that variability in behavior and neural activity represents fluctuations in a gain signal. Therefore models of decision-making that include an urgency signal should outperform models without urgency and further support our findings thus far. To this end we employed a behavioral modeling toolbox developed by Dr. Chandrasekaran (Chandrasekaran & Hawkins, 2019), which allows for the selection and rigorous fitting of various behavioral models built off of

a standard DDM.

There were four general types of models that can be understood as modifications upon the standard DDM (Fig 7A:D). Conceptual descriptions of the models are given here; model equations and variable definitions are given in the methods. Figure 7A visualizes a modified DDM ($DDMS_vS_zS_t$) with a variable drift rate (S_v), variable start point (S_z), variable non-decision time (S_t) and 2 bounds (A & 0). The bounds of the model represent evidence thresholds to different choices. Starting point is the baseline condition from which evidence is accumulated. Drift rate is the rate at which evidence is accumulated during decision time. Evidence for one choice will move the evidence accumulator towards the respective bound and vice versa. A decision is made when the drift rate crosses one of the bounds. The drift rate (v) has a starting point (z) that varies from trial-to-trial, starting closer to one bound or the other. Non-decision time (T_{er}) is time required for stimulus processing and encoding, and motor preparation/execution (“non-decision” related processes). The predicted RT is the sum of the non-decision time and decision time.

The collapsing bounds (c) DDM ($cDDMS_vS_z$) (Fig 7B) is similar to the last model but rather than fixed bounds it has collapsing ($A(t)$ & $a - A(t)$) bounds that shrink as a function of decision time. As time increases, less evidence is needed to cross a boundary and therefore make a choice. The urgency (u) DDM ($uDDMS_bS_u$) (Fig 7C) is most similar to a standard DDM in that it has a non-variable drift rate (v), starting point (z), and non-decision time (T_{er}). However it differs in that the drift rate is scaled by a linear gain term ($\gamma(t)$) with variable intercept (S_b) and slope (S_u). Here as time elapses the gain signal grows and pushes the evidence accumulator closer to a bound with relatively less evidence. Finally, the nonlinear urgency (nlu) DMM ($nluDDMS_bS_u$) is identical to the urgency DDM but the variable slope of the gain term scales a sigmoidal function (Fig 7D). The nonlinear gain has the similar effect

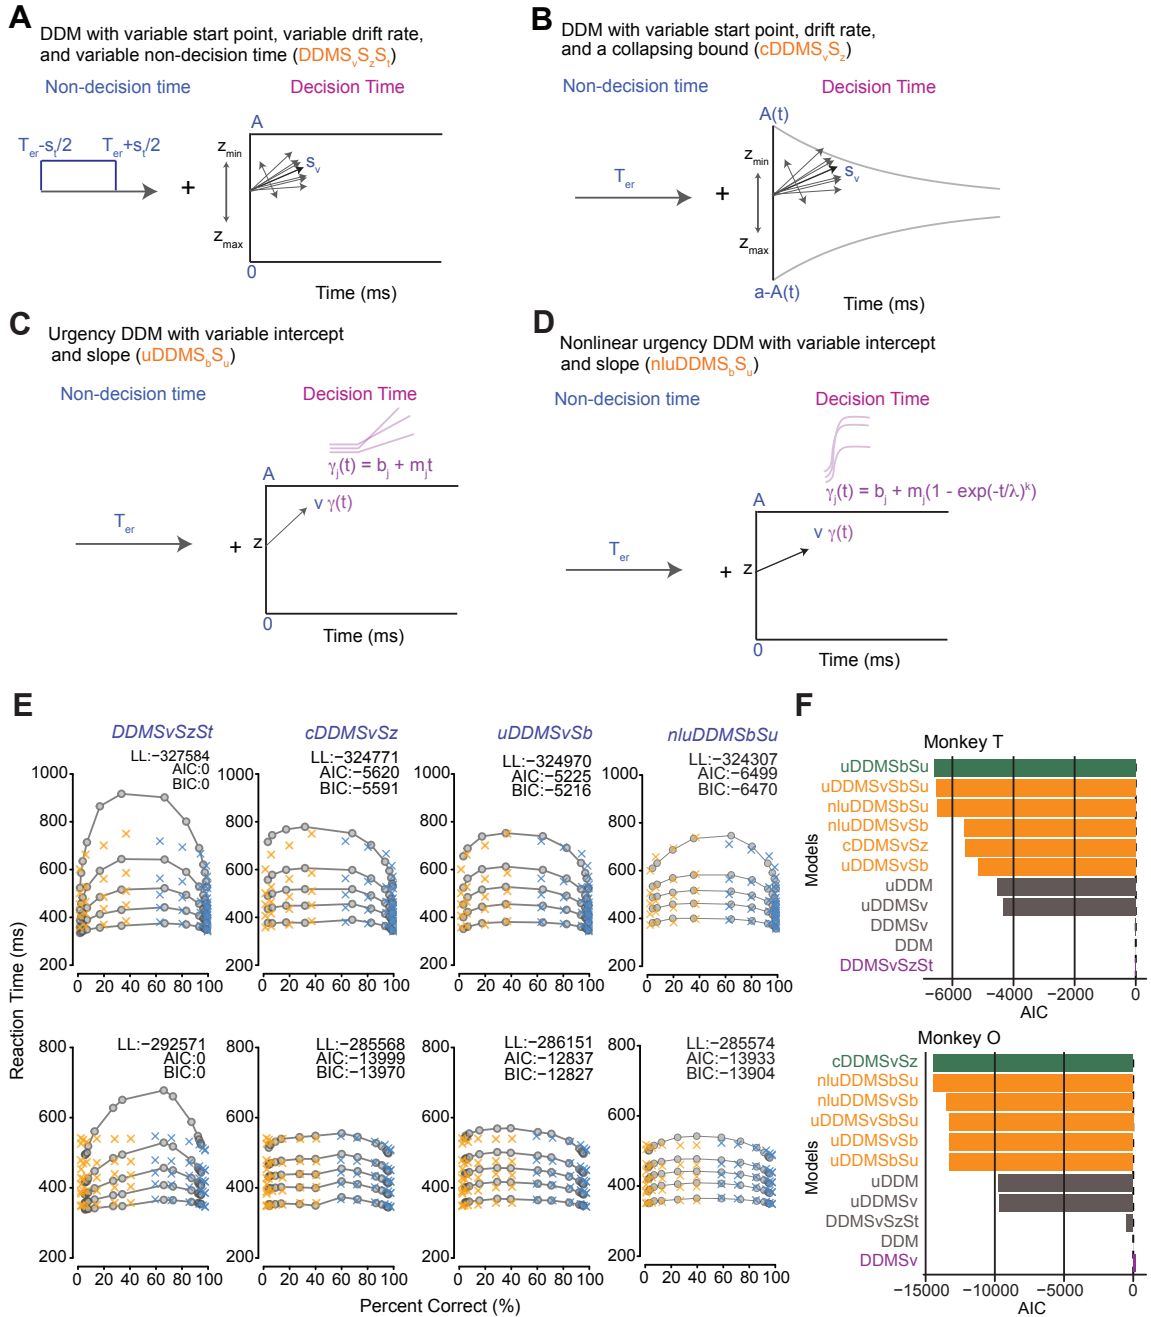


Figure 2-7: **Behavioral models that include an urgency term predict decision behavior better than models without it** (A) Diffusion Decision Model (DDM), models decision making with variable non-decision time (S_t), a variable start point (S_z), variable drift rate (S_v) and fixed bounds (A & 0). (B) DDM with fixed non-decision time (T_{er}), variable starting point (S_z), variable drift rate (S_v), and collapsing (c) bounds ($A(t)$ & $a - A(t)$). (C) Urgency (u) DDM with fixed non-decision time (T_{er}), a fixed drift rate (v), multiplicative urgency term ($\gamma(t)$), and fixed bounds (A & 0). The urgency term has a variable intercept (S_b) and slope (S_u), and acts as a gain on the drift rate. (D) Nonlinear urgency (nl) DDM with fixed non-decision time (T_{er}), fixed drift rate (v), multiplicative urgency term ($\gamma(t)$), and fixed bounds (A & 0). Here the urgency model has a variable intercept (S_b) and variable slope (S_u) which scales a nonlinear sigmoidal function.

Figure 2-7: **(E)** Quantile probability plots of the decision making models for monkey T (top row) and monkey O (bottom row). A quantile probability plot shows the percent correct on the x-axis and response time in milliseconds on the y-axis. Each column of x's represents a coherence, of which there are 7 levels. Blue x's are the proportion correct for that coherence, whereas yellow x's are the proportion of errors for that coherence. The x's are at the 10th, 30th, 50th, 70th, and 90th percentile of the RT distribution for the particular level of coherence and accuracy. The gray circles are the model predictions for the accuracy and RT distributions. The closer the gray circles are to the x's the better the model has predicted the real data. LL - Log Likelihood, AIC - Akaike Information Criterion, BIC - Bayesian Information Criterion **(F)** AIC scores for all tested models for both monkeys. The lower the AIC score the better the model is at predicting behavior. Models are listed in increasing order of AIC score (i.e. they become less negative).

of pushing the sensory accumulator towards a bound with relatively less evidence.

Overall models with urgency were substantially better at describing the behavior of the monkeys than models without urgency (e.g. compare $uDDM$ to $DDMS_vS_z$) (Fig 7E,F). Moreover, models with variable gain signals were even better than models without variability (e.g. compare $uDDMS_bS_u$ to $uDDM$) (Fig 7E,F). Together, the results of behavioral modeling corroborate the observation from the analysis of the neural population dynamics that a strong component of decision-making behavior is a stimulus-independent urgency signal which modulates the decision formation process.

2.8 Recurrent neural network models that include either a multiplicative or additive gain signal recapitulate PMd neural dynamics

Three RNN models (no gain, additive gain, and multiplicative gain) were built to test whether a gain signal on firing rates is necessary to recapitulate PMd population dynamics observed in our data. Figure 8A lays out the basic schematic for the RNN modeling (further details of modeling and full description of equations (Fig 8A) in methods). Noisy red and green input signals (u_t - input signal, Fig 8A) enter an RNN optimized (x_t - recurrent state, Fig 8A) to perform the red-green checkerboard discrimination task. The RNN outputs whether it 'reaches left or right' (y_t - sigmoid output function, Fig 8A) and an associated RT for each trial. The gain modulated or unmodulated 'firing rate' (r_t - rectified linear unit function (ReLU) firing rate

function, 8B) activity can be read out for each trial and used to analyze population dynamics. The no gain model directly passes the state of the RNN through a ReLU. Whereas the additive and multiplicative gain models add and multiply, respectively, a variable gain with the recurrent state.

The average firing rate of each unit in each RNN model was put into a PCA conditioned on choice and RT mirroring the PCA performed on real PMd data (Fig 4B). As shown in figure 8B, all 3 models recapitulated stereotypical post-stimulus population dynamics separating by choice and RT bin. Importantly, both the additive and multiplicative gain models demonstrated prestimulus neural covariation with RT, whereas the no gain model showed no prestimulus neural covariation with RT (Fig 8B). So only RNN models that included either an additive or multiplicative gain signal were able to recreate the full range of dynamics observed in real PMd population data.

Firing rate for each trial from each model was regressed with corresponding RTs and R^2 was calculated for each model in 20 ms bins of firing rate activity. Prestimulus firing rate explains $\sim 20\%$ of the variance in RT in both the additive and multiplicative gain model (Fig 8C). This is considerably more variance explained than the 99th percentile of trial shuffled spiking activity regressed with random RTs (Fig 8C). However the prestimulus firing rate activity from the no gain model doesn't explain any RT variance and is no better than the 99% or 1% shuffled firing rate activity with RTs (Fig 8C). Consistent with the modeled dynamics, linear regression analyses show that prestimulus activity in the additive and multiplicative gain RNNs, but not the 'no gain' RNN, explains RT .

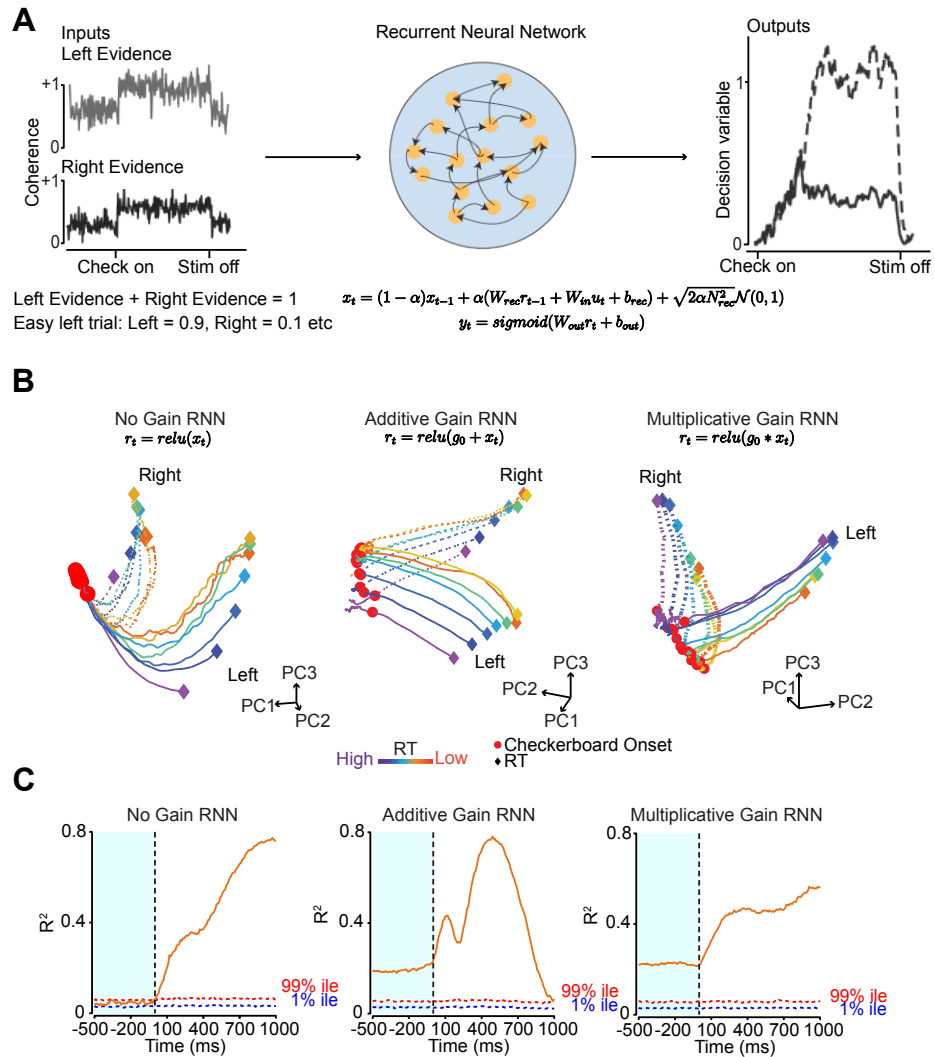


Figure 2-8: **Recurrent neural network models with additive or multiplicative gain replicate PMd pres-stimulus neural dynamics** (A) Schematic of RNN showing inputs, outputs, and state equations. (B) Trajectories of average firing rates of RNN units organized by choice and RT. Red dots are checkerboard onset; diamonds are 280 ms post-checkerboard onset. (C) RT variance explained (R^2) by RNN firing rate (20 ms time step, linear regression). Red and blue dotted lines represent the 99th and 1st percentile, respectively of shuffle controls.

Chapter 3

Discussion

3.1 Summary

Previous evidence has demonstrated that neural firing rates increase before a decision is made (Hanks et al., 2014; Latimer et al., 2015; Roitman & Shadlen, 2002; Shadlen & Newsome, 1996; Shadlen & Newsome, 2001). Typically this is equated to evidence accumulation (Ratcliff et al., 2016; Roitman & Shadlen, 2002). However further evidence challenges this as neurons in decision loci demonstrate heterogeneous and complex spiking activity (Jun et al., 2010; Mante et al., 2013; Okazawa et al., 2021) and build up of firing rate activity can also be modified by internal factors such as urgency (Cisek et al., 2009; Thura & Cisek, 2014; Thura et al., 2014). This thesis addresses how the spiking activity of a heterogeneous population of neurons is modulated by sensory evidence and urgency. Our hypothesis is that sensory evidence will modulate spiking activity associated with choice, whereas urgency, acts as a gain signal on the evidence accumulation process (Murphy et al., 2016), modulating prestimulus firing rates.

Several findings, from monkeys trained in a red-green checkerboard discrimination task, help give credence to the hypothesis of a stimulus-independent, time-dependent variable gain signal manifest in prestimulus neural activity that drives decision making behavior. First, prestimulus neural firing rates covaried with RT, between and within stimulus difficulties. Additionally $\sim 15\%$ of the variance in RTs was explained by prestimulus spiking activity from one brain area alone, whereas levels of stimulus

coherence maximally accounted for $\sim 10\%$ of RT variability. Furthermore faster RT trials were associated with faster neural dynamics. Second, prestimulus firing rates did not covary with choice and prestimulus spiking activity did not perform better than chance in predicting choice. The covariation of prestimulus neural activity but not with choice was even found in an analysis of single trials. Third, these findings were buttressed by behavioral and neurophysiological models that included variable gain terms outperforming models that lacked such terms. Finally, we found evidence that neural state changes dramatically as a function of PES and that the current trial's baseline spiking activity is influenced by the outcome of the previous trial. This suggests that trial history contributes to prestimulus neural state and therefore contributes to a stimulus independent gain signal. Fundamentally, the notable contribution of this thesis is that prestimulus neural dynamics contribute to the timing of a decision but not directly to the choice itself.

The dataset is rich with explorable avenues but in the interest of the audience's attention span (and the writer's sanity) this discussion will focus on four themes. First, we predicted that models that included a multiplicative gain signal would best recapitulate behavior and neural dynamics. While this prediction was true to an extent, modeling that included either an additive or multiplicative gain signal performed equally well and were indistinguishable. We will explore our own findings and previous findings to discuss why this could be the case. Second, previous evidence and our results suggest that urgency impinges upon prestimulus neural variability related to motor preparation. We will discuss sources of this prestimulus neural variation and its relation to urgency. Third, we add to the evidence that PES contributes to changes in prestimulus neural dynamics. This initial finding is suggestive of a multiplexing of neural activity involved in urgency signal variation. Finally, we will discuss decision-making theory, why urgency exists and where such a signal might originate.

This will naturally lead into ideas for future research.

3.2 Prestimulus neural variation consistent with a variable time-dependent and stimulus-independent gain signal

Behavioral results and neural dynamics were best explained by models that included a variable gain signal as compared to models that didn't have such a gain signal. We predicted that behavioral results and neurophysiological dynamics would best be replicated by modeling that included a multiplicative gain signal (Cisek et al., 2009). However our results from our models are unable to make reasonable or conclusive distinctions between additive and multiplicative gain models for behavioral results or neural dynamics.

Additive gain models, mathematically equivalent to collapsing bound models (Chandrasekaran & Hawkins, 2019), and multiplicative gain models have similar AIC scores and neural dynamics across both monkeys. While a cluster of models that include a multiplicative gain signal outperform the collapsing bound model for monkey T, the collapsing bound model slightly outperforms a non-linear multiplicative gain model for monkey O. Additionally, RNNs with an additive or multiplicative gain signal recapitulate prestimulus neural dynamics as observed with our PMd data. As well both models capture about the same amount of prestimulus variance in RTs ($\sim 20\%$). After stimulus presentation neither model captures variance in RT from spiking activity in quite the same way as the real spiking data. Yet each RNN seems to model some aspect of post stimulus variance. For example, the multiplicative gain model captures the growth of captured variance while the additive model peaks in variance explained at about the same time as in the real data. Perhaps such mixed results indicate that a model that combines a multiplicative and additive gain signal may be the best model.

These results and previous evidence (Heitz & Schall, 2012) suggest that there can be several ways in which the brain solves speeded decision-making tasks. Collapsing bound/additive gain models and multiplicative gain models are not mutually exclusive and one study found evidence that 2 non-human primates solved speed-accuracy tradeoff (SAT) problems consistent with models that included collapsing bounds and a growing urgency signal (Heitz & Schall, 2012). However, other studies do not support such a mixed model, but rather support a growing urgency signal model. These studies support baseline modulation as a means of adjusting decisions under SAT constraints (Bogacz et al., 2010; Hanks et al., 2014). These studies find baseline adjustments between different SAT conditions and that firing rates rise to a threshold that is constant across these different contexts. That being said it could be that subjects solve speeded decision-making tasks in an individualized manner (Carland et al., 2019) or alternatively that rise-to-threshold models are incomplete (Heitz & Schall, 2012).

Generally support for rise-to-threshold models derives from the qualitative assessment that firing rates tend to approach the same threshold in conditions that manipulate speed versus those that do not (e.g. (Roitman & Shadlen, 2002)). Such results offer support for fixed bounds by analogy and are by no means conclusive. As previous findings have demonstrated neural spiking activity is complex and diverse (Jun et al., 2010; Mante et al., 2013; Thura & Cisek, 2014). Yet readily understood in the language of dynamical systems (Mante et al., 2013). This study provides further evidence that the heterogeneous activity of individual units can be readily understood from their population dynamics where stereotyped neural behavior that is consistent across stimulus difficulties and choice is apparent. Where proportions of increasing, decreasing, and peri-movement neurons become difficult to interpret, population level activity is lawful and hints a low-dimensional neural code that belies decision-making

behavior.

In fact, in a recent study, neural populations increased in firing rate before a decision was made for a dot motion task, recapitulating past results. However in the other group of monkeys performing a human-monkey face morph discrimination task, the opposite effects were observed where firing rates decreased before a decision was made (Okazawa et al., 2021). Additionally these results were lawful in that with varying difficulty the firing rates either increased or decreased proportional to the difficulty of the respective tasks. These results are difficult to interpret with classic rise-to-threshold models. However when analyzed under the lens of dynamical systems, low-dimensional population dynamics emerged and it was found that the dynamics for both tasks existed on lawful rotated curved manifolds (Okazawa et al., 2021). Such results support the contention that spiking activity, and therefore decision-making, can be readily understood from the perspective of several latent factors (Afshar et al., 2011). Furthermore they suggest that a rise-to-threshold model for neural activity may be overly simplistic and not a suitable gauge of decision-making behavior.

Behavioral results, vigour of movement, and neurophysiology, suggest that individuals vary in their level of impulsiveness (Carland et al., 2019). Notably, there were no SAT manipulations in this study yet we found evidence for internal SAT mechanisms. The two monkeys replicated lawful behavioral results whereby as stimuli increase in difficulty both monkeys are less accurate and are slower to respond. Yet each monkey, as evidenced by their distinct RT distributions, respond to stimuli according to their own embodied risk/reward strategies. Relatedly, most individuals slow down after committing errors as observed in this study. However some individuals actually speed up after making errors (Purcell & Kiani, 2016). Again individuals demonstrate their different types of strategies towards maximizing reward. Specula-

tively this could support an individually weighted multiplicative gain signal (Carland et al., 2019) or multiple individually tuned mechanisms (e.g. collapsing bounds & multiplicative gain signal) (Heitz & Schall, 2012). It is the subject of future research as to whether there is only one way to solve speeded decision-making or it relies upon a confluence of individualized mechanisms.

3.3 Previous findings and our results suggest urgency impinges upon neural variability of motor planning

Prestimulus neural covariation with RT but not choice is surprising in a decision making task, especially from the lens of standard DDMs. Yet such findings are consistent with and supportive of theories of embodied cognition and results from delayed reach tasks. Embodied cognition theories essentially argue that the brain evolved as a motor output system that phylogenetically added more and more complex behaviors (Cisek, 2022). Essentially all behaviors and especially brain activity must be viewed from the context of how it supports or is derived from motor activity. Now delayed reach to target tasks demonstrate that during a delay period neural activity related to motor preparation can explain a large percent of variance in RTs (Afshar et al., 2011). While our results are consistent with this finding, our results are novel due to important differences between the two tasks. In the delayed reach to target task the monkey has forward knowledge of which target to press and can prepare the appropriate motor response during the delay period (30 -1000 ms) (Afshar et al., 2011). In our task while the targets are presented before the stimulus with a theoretically long enough delay period as to mitigate better or worse preparation between randomized delay periods (400-1000 ms), there is no indication of where the monkey should reach until the stimulus is shown. Presumably given the monkeys expertise in the task their brains are ‘aware’ that one of two sides will have to be

reached to. Within this context it makes sense that decisions are made based upon competing motor plans. All this to say that the results from a decision making task are highly consistent with a motor preparation study. This within the lens of embodied cognition suggests that decision making is essentially a competition between motor plans.

So prestimulus neural activity that we observe could be related to the preparation of two motor plans. This is fundamental as it suggests the preparation of multiple motor plans prior to any stimulus onset and that the state of progress of the motor plan relates to the ultimate RT. We can be confident in this assessment from a couple pieces of evidence. One PMd has been extensively studied and consistently shows involvement in motor preparation (Afshar et al., 2011; Cisek & Kalaska, 2005). Second there is evidence for the formation of multiple motor plans but one is chosen once sufficient evidence is presented to favor one choice (Cisek & Kalaska, 2005). So speculatively multiple motor plans are prepared prior to stimulus presentation and the state of those dynamics contribute to RT variability (Afshar et al., 2011).

Findings of trial-history dependent prestimulus neural covariation suggest potential sources of neural variation which may impinge upon prepared motor plans. Sources of variance could include the amount of time there is to prepare appropriate motor responses (Afshar et al., 2011), how many motor plans are being prepared simultaneously (Cisek & Kalaska, 2005), as well as the history of successes and failures of such motor programs. We have shown evidence for one of these potential sources of neural variability. Specifically, trial history contributes to variability in prestimulus neural states and the previous trial's outcome affects the current trials spiking activity. We demonstrate in a PES PCA that neural state is considerably different after an error as compared to after a correct answer. Thura and others similarly found that a percentage of recorded neurons decreased in firing rate after an error and persisted

after stimulus onset whereas a similar percentage of neurons increased in firing rate baseline activity and persisted post-stimulus (Thura et al., 2017) (Thura et al, 2016). Tantalizingly (Purcell & Kiani, 2016) also report prestimulus neural variation after an error in one monkey. However this monkey exhibited post error speeding rather than post error slowing. Yet there were differences in firing rate that corresponded with the relevant behavioral response; increased baseline firing rate after an error associated with faster responding. One can only assume that although the behavioral responses differ, there is a similar type of change in neural state as observed for post error slowing.

3.4 Theory and Origins

Affordance competition hypothesis states that when a decision maker is put into a dangerous situation, for example, the actor generally has multiple ways (‘affordances’) of engaging with the environment (e.g. fight or flight) (Cisek, 2007). The theory proposes that the brain prepares multiple motor action plans (Cisek & Kalaska, 2005) and calculates how soon or vigorously to engage a plan as well as the metabolic costs and probability of success for each plan (Cisek, 2007). The preparation and analysis of these motor plans constitute the affordance competition.

Urgency is a purported signal that is hypothesized to resolve the neural competition between separate plans as an evidence-independent time-dependent signal that promotes the activity of multiple motor plans (Pastor-Bernier et al., 2012). It boosts the signal for multiple plans to race them to the ‘critical threshold’ needed to bring about a decision one way or another (Thura & Cisek, 2014). However, what this arbitrary threshold is, or could be, as discussed previously, is unknown.

Speculatively, it could be that urgency is a drive on heterogeneous populations of neurons. So urgency drives the rate of neuronal firing uniformly across populations so

as to speed the building of low-dimensional manifold structure (Okazawa et al., 2021). Once a low-dimensional manifold is achieved then a decision is made and manifold activity related to other motor action plans die down.

Results from human fMRI SAT research, behavioral modeling, and computational modeling all support that there is a time-dependent gain signal that modifies the accumulation of sensory evidence. However where this signal originates from and in what form remains unclear. fMRI research suggests premotor areas such as pre-SMA and the striatum, the input nucleus of the basal ganglia, show increased activation in prestimulus blocks where speed is emphasized versus when it is not (Bogacz et al., 2010). Relatedly another study found evidence of an urgency type signal coming from the basal ganglia (Thura & Cisek, 2017). This study found that neural activity in basal ganglia did not relate to the deliberation process for the eventual choice but rather influenced the decision making process indirectly by acting as a gain signal for sensory evidence in PMd and primary motor cortex (Thura & Cisek, 2017). Activity in these regions prior to stimulus onset and their implication in non-deliberative processes of decision-making is consistent with research implicating the SMA, striatum and basal ganglia in motor planning (Nachev et al., 2008; Thura & Cisek, 2017; Thura et al., 2017) and with urgency theory (Cisek et al., 2009). The implication of these regions as potential sources of urgency make them attractive targets for further research on how their activity influences population dynamics.

3.5 Future Directions

Hypothetically a post correct state evolves faster than a post error state as post correct states are associated with a faster RT. As we have observed from the KiNeT analysis, faster RT states are associated with faster RTs. I propose another KiNeT analysis where the error state is used as the reference trajectory and post error and correct

states are compared to it. Such an analysis would be a key link in demonstrating that previous outcomes lead to faster or slower dynamics (i.e. variable urgency) and therefore corresponding RTs.

It is unclear whether there are specific subpopulations of neurons that are modulated by urgency or if urgency modulates neuronal activity uniformly. I propose to regress the spiking activity of different classes (i.e. as classified in (Chandrasekaran et al., 2017)) of neurons and randomized sets of neurons with RTs to determine if RT variability is better explained by certain classes of neurons or a ‘zoo’ of neurons. Such an analysis would reveal how urgency affects neuronal populations and drives decision-making behavior.

Can the current trial’s outcome (e.g. error) be predicted from prestimulus activity? Such an analysis, using logistic regression, could reveal that there are neural states that coincide with more impulsive/urgent behavior that lead to hasty incorrect responding.

Finally, poststimulus neural dynamics reveal structured and stereotyped multidimensional activity that covaries with RT and is predictive of choice. Similar but perhaps less well defined structure is found in prestimulus neural activity as it was found to covary with RT. However we believe that these dynamics belie further prestimulus activity structure to uncover. If this is the case then what kind of dynamics can be extracted from prestimulus activity? I propose an exploratory analysis where prestimulus neural activity would be conditioned on 2 of the following 3 variables: RT, delay between targets and stimulus, and previous trial outcome. Preparatory activity as suggested by Afshar and others (2011) has some level of structure related to latent dimensions however the preparatory activity exists in a larger subspace than neural activity associated with movement onset which is more stereotyped (Afshar et al., 2011). Importantly, prestimulus dynamics lead into post-stimulus dynamics so

how might prestimulus dynamics further influence behavior once a stimulus is shown?

3.6 Conclusions

Prestimulus neural activity was found to covary with RTs. This activity was explanatory of RTs but not predictive of choice. Neural states associated with faster or slower RTs had neural dynamics that evolved faster or slower respectively. These results are supported by behavioral models and RNN modeling revealing that when a time-dependent gain signal is included, behavioral results are better predicted and PMd neural dynamics are more accurately recapitulated. Finally, prestimulus neural states after an error differed considerably after an error as compared to after a correct response and there was evidence that previous trial outcome affected the current trial's prestimulus spiking activity. All together this data supports a variable trial-history and time-dependent yet stimulus-independent gain signal which controls the speed of prestimulus dynamics and acts as a gain on sensory evidence accumulation.

Chapter 4

Methods

Several method sections are adapted from Chandrasekaran and others 2017 as this thesis and that study share the same data set.

4.1 Subjects

Experiments were performed using two adult male macaque monkeys (*Macaca Mulatta*; monkey T, 7 years, 14 kg & monkey O, 11 years, 15.5 kg) trained to touch visual targets for a juice reward. Monkeys were housed in a social vivarium with a normal day/night cycle. Protocols for the experiment were approved by the Stanford University Institutional Animal Care and Use Committee as the experiments were conducted at Stanford University. Animals were initially trained to come out of their housing and to sit comfortably in a chair. After initial training, monkeys underwent sterile surgery where cylindrical head restraint holders (Crist Instrument Co., Inc., Hagerstown, MD, United States) and standard recording cylinders (Crist Instrument Co., Inc.) were implanted. Cylinders were placed surface normal to the cortex and were centered over caudal dorsal premotor cortex (PMdc) (+16, 15 stereotaxic coordinates). The skull within the cylinder was covered with a thin layer of dental acrylic.

4.2 Apparatus

Monkeys sat in a customized chair (Synder Chair System, Crist Instrument Co., Inc.) with their head restrained. The arm that was not used to respond in the task was gently restrained with a tube and cloth sling. Experiments were controlled and data collected using a custom computer control system (Mathworks' xPC target and Psychophysics Toolbox, The Mathworks, Inc., Natick, MA, United States). Stimuli were displayed on an Acer HN2741 monitor approximately 30 cm from the monkey. A photodetector (Thorlabs PD360A, Thorlabs, Inc., Newton, NJ, United States) was used to record the onset of the visual stimulus at a 1 ms resolution. A small reflective spherical bead (11.5 mm, NDI passive spheres, Northern Digital, Inc., Waterloo, ON, Canada) was taped to the middle finger, 1 cm from the tip, of the active arm of each monkey; right for T and left for O. The bead was tracked optically in the infrared range (60 Hz, 0.35 mm root mean square accuracy; Polaris system, NDI). Eye position was tracked using an overhead infrared camera with an estimated accuracy of 1° (ISCAN ETL-200 Primate Eye Tracking Laboratory, ISCAN, Inc., Woburn, MA, United States). To get a stable image for the eye tracking camera, an infrared mirror (Thorlabs, Inc.) transparent to visible light was positioned at a 45° angle (facing upward) immediately in front of the nose. This reflected the image of the eye in the infrared range while allowing visible light to pass through. A visor placed around the chair prevented the monkey from touching the juice reward tube, infrared mirror, or bringing the bead to its mouth.

4.3 Task

Experiments were made up of a sequence of trials that each lasted a few seconds. Successful trials resulted in a juice reward whereas failed trials led to a time-out of 2-4 s. A trial started when a monkey held its free hand on a central circular cue

(radius = 12 mm) and fixated on a small white cross (diameter = 6 mm) for ~350-400 ms. Then two isoluminant colored (red and green) targets appeared 100 mm to the left and right of the central hold cue. Targets were randomly placed such that the red target was either on right or the left trial-to-trial, with the green target opposite the red one. In this way color was not tied to reach direction. Following an additional center hold period (400-900 ms) a static checkerboard stimulus (15 x 15 grid of squares; 225 in total, each square: 2.5 mm x 2.5 mm) composed of isoluminant red and green squares appeared superimposed upon the fixation cross. The monkey's task was to move their hand from the center hold and touch the the target that matched the dominant color of the checkerboard stimulus for a minimum of 200 ms (full trial sequence; Fig 1B). For example, if the checkerboard stimulus was composed of more red squares than green squares the monkey had to touch the red target in order to have a successful trial. Monkeys were free to respond to the stimulus as quickly or slowly as they 'chose' There was no delayed feedback therefore a juice reward was provided immediately following a successful trial (Roitman & Shadlen, 2002).

The checkerboard stimulus was parameterized at 14 levels of red (R) and complementing green (G) squares ranging from nearly all red (214 R , 11 G) to all green squares (11 R , 214 G) (example stimuli; Fig 1C). These 14 levels are referred to as signed coherence (SC), defined as $SC = 100 \times (R - G)/(R + G)$ (R: 4%:90%, G: -4%:-90%). Correspondingly there are seven levels of color coherence, agnostic to the dominant color, defined as $C = 100 \times |R - G|/(R + G)$ (4-90%).

The hold duration between the onset of the color targets and onset of the checkerboard stimulus was randomly chosen from a uniform distribution from 400-900 ms for Monkey T and from an exponential distribution for Monkey O from 400-800 ms. Monkey O attempted to anticipate the checkerboard stimulus therefore an exponen-

tial distribution was chosen to minimize predictability.

4.4 Training

Operant conditioning was used to train the monkeys. Monkeys were rewarded for arm movements toward the screen and learned to take pieces of fruit from the screen. Once the monkey acquired the reach-reward association, the animal was then trained to hold a center cue for a juice reward. The position and color of the target was randomized so that the animal learned to reach to different positions on the screen. Then they were trained on a simplified task design where they briefly held the center cue and then the checkerboard stimulus, nearly all red or green, appeared for 400-600 ms before the two targets appeared. Monkeys were rewarded for touching the target that matched the dominant color of the checkerboard. The association of the checkerboard color and the correct target was reinforced by interleaving two-target ‘decision’ blocks with single-target blocks. Monkeys reliably reached to the target that matched the checkerboard stimulus after several weeks of training with the interleaved block design. The paradigm was then switched such that the targets appeared before the checkerboard onset. Monkeys were trained to maintain the center hold for 300-1800 ms post target onset and prior to checkerboard onset by rewarding them with juice. When the animal reliably held the center hold between target and checkerboard onset, the monkey was no longer rewarded for holding. Over subsequent weeks more difficult stimuli were presented to the monkey and the maximum hold time was tapered to 900 ms. The monkeys were encouraged to perform the task as quickly and as accurately as possible while impulsivity was discouraged by including timeouts after unsuccessful trials.

4.5 Electrophysiological Recordings

Electrophysiological recordings were guided by stereotaxic coordinates, known response properties of PMd, and neural responses to muscle palpation. Recordings were made anterior to the central sulcus, lateral to the precentral dimple and lateral to the spur of the arcuate sulcus. Electrodes were placed in the PMd contralateral to the dominant hand of the monkey (T: right arm, O: left arm). Recording chambers were placed surface normal to the cortex to align with the skull of the monkey and recordings were performed orthogonal to the surface of the brain. Estimates of upper and lower arm representation was confirmed with repeated palpation at a large number of sites to identify muscle groups associated with the sites.

Single electrode recording techniques were used for a subset of the electrophysiological recordings. Small burr holes in the skull were made using handheld drills. A Narishige drive (Narishige International USA, Inc., Amityville, NY, United States) with a blunt guide tube was placed in contact with the dura. Sharp FHC electrodes ($> 6 \text{ M}\Omega$) (FHC, Inc., Bowdoin, ME, United States) penetrated the dura and every effort was made to isolate, track, and stably record from single neurons.

180 μm thick 16-electrode linear multi-contact electrode (U-probe; interelectrode spacing: 150 μm , contact impedance: $\sim 100 \text{ k}\Omega$) recordings were performed similarly to single electrode recordings with some modifications. Scraping away any overlying tissue on the dura, under anesthesia, and a slightly sharpened guide tube aided in slow U-probe penetration ($\sim 2\text{-}5 \mu\text{m/s}$). U-probe penetration was stopped once a reasonable sample of neurons was acquired, potentially spanning multiple cortical layers. Neural responses were allowed to stabilize for 45-60 minutes before normal experimentation began. Monkey T had better recording yields on average (~ 16 units/session) than monkey O (~ 9 units/session). Additionally, lowering the electrode necessitated careful observation to ensure the electrode did not bend, break at the tip or exces-

sively dimple the dura. Therefore it was not possible to precisely localize the U-probes with a grid system between sessions.

4.6 Unit Selection and Classification

The electrophysiological recordings consist of 996 units (546 units in T and 450 units in O, including both single neurons and multi-units) recorded from PMd of the two monkeys as they performed the task over 141 sessions. Chosen units were included as they were well isolated from other units/separated from noise and modulated activity in at least one task epoch.

U-probes were useful for recording from isolated single neurons as U-probes are low impedance ($\sim 100\text{ k}\Omega$) with a small contact area. A conservative threshold was used to maximize the number of well defined waveforms and to minimize contamination from spurious non-neural events. Single neurons were delineated online by the ‘hoops’ tool of the Cerebus system software client (Blackrock Microsystems, Salt Lake City, UT, United States) after the electrodes had been in place for 30 - 45 minutes. When a spike was detected via thresholding, a 1.6 ms snippet was stored and used for subsequent evaluation of the clusters as well as modifications needed for spike sorting.

Some electrodes in U-probe recordings captured mixtures of 2 or more neurons, well separated from each other and noise. In the majority of cases the waveforms were separable and labeled as single units. These separations were verified by viewing the waveforms in principal component space using custom code in MATLAB (The MathWorks, Inc., Natick, MA, United States). MatClust the MATLAB based clustering toolbox or Plexon Offline Sorter (Plexon, Inc., Dallas, TX, United States) were used to adjust the clusters that were isolated online.

Recording activity labeled as ‘multi-units’ were mixtures of 2 or more neurons not separable using a principal components method or consisted of recordings with

waveforms only weakly separable from noise.

The number of interspike interval (ISI) violations after clustering and sorting was used to mitigate subjectivity in the classification of units. A unit was labeled as a single neuron if the percentage of ISI violations (refractory period of ≤ 2 ms) was $\leq 1.5\%$ otherwise it was labeled as a multi-unit. 801/996 PMd units were labeled as single neurons (T: 417, O: 384, mean ISI violation= 0.43%, ~ 0.13 additional spikes/trial). Therefore 195/996 units were labeled as multi-unit (T: 129, O: 66, mean ISI violation= 3.36%, ~ 1.4 additional spikes/trial).

Units from both monkeys were pooled together as the electrophysiological characteristics were similar. Change of mind trials ($\sim 2-3\%$) were excluded from averaging as the change in reach direction mid-movement execution made the assignment of choice ambiguous. Incorrect and correct trials arranged by choice were averaged together.

4.7 Effects of coherence on accuracy and reaction time (RT)

Behavior was analyzed by fitting psychometric and RT curves on a per-session basis and averaging the results across sessions. Behavioral data was analyzed in the same sessions as the electrophysiological data. In total there were 75 sessions for monkey T (128,989 trials) and 66 sessions for monkey O (108,344 trials). On average there were $\sim 1,500$ trials/session. Both incorrect and correct trials for each *SC* were included for estimating RT/session.

Data were fit to a psychometric curve to characterize how discrimination accuracy changed as a function of stimulus coherence. For each session a monkey's sensitivity to the checkerboard stimulus was estimated by estimating the probability (p) of a correct choice as a function of the color coherence of the checkerboard stimulus (c). The accuracy function was fit using a Weibull cumulative distribution function.

Weibull cumulative distribution function:

$$p(c) = 1 - 0.5e^{-\left(\frac{c}{\alpha}\right)^\gamma} \quad (4.1)$$

The discrimination threshold α is the color coherence level at which the monkey would make 81.6% correct choices. The threshold is the mean α across sessions. The parameter γ describes the slope of the psychometric function. Threshold and slope parameters were fit per session and averaged across sessions. Mean and standard deviation of threshold estimates are reported in Fig 1D. R^2 values from the fit are provided in the text.

Mean RT was calculated per *SC* on a session-by-session basis and averaged across sessions. Results are displayed in Fig 1E with error bars denoting $2 \times SEM$ and lines between the averages to guide the eyes. RT was also regressed with log coherence (c) per session. The fit coherence-RT model was used to predict RTs and calculate R^2 on a per session basis. R^2 values were averaged across sessions per monkey and are reported in Fig 1F as percentage of variance explained. The general framework and equations for linear regression and R^2 calculation are provided in the following section.

4.8 Fits generated by linear regression and logistic regression (decoder)

Throughout the thesis linear and logistic regressions (decoders) were used to determine the variance explained by spiking activity and whether spiking activity was predictive of certain outcomes, respectively. To determine whether spiking activity explains variance in RTs, 1800 ms of spiking activity from each trial (600 ms prestimulus and 1200 ms post-stimulus) were divided into 20 ms bins (90 bins total) across all neurons within a session. 90 matrices composed of a single time interval of spike

counts (e.g. 0-20 ms) for all trials, across all neurons were linearly regressed with trial matched RTs. Fits were used to predict RTs based on spiking activity, which were used to calculate R^2 values per bin. The per bin R^2 values were then averaged across sessions. R^2 values were compared to the 1st and 99th percentile of R^2 values calculated from regressions of trial shuffled spiking activity and RTs (shuffled 500 times).

Linear Regression:

$$RT(i) = b + mi \tag{4.2}$$

i - independent variable (i.e. coherence or binned spikes)

b - intercept

m - slope of regression

R^2 calculation:

$$R^2 = 1 - \frac{(RT_i - \widehat{RT}_i)^2}{(RT_i - \overline{RT}_i)^2} \tag{4.3}$$

\widehat{RT}_i - predicted RT

\overline{RT}_i - mean RT

To determine whether spiking activity can predict eventual choice or previous trial outcome a logistic regression, using binned spiking activity as the predictor and choice as the outcome was performed. The spiking activity was binned as before and 90 logistic regressions (1/bin) were computed in order to predict choice based on spiking activity. Accuracy per bin was calculated from the predictions of the logistic regression model. The per bin accuracy was then averaged across sessions and compared to the 1st and 99th percentile of accuracy of logistic regression models

built on 500 shuffles of binned spiking activity (i.e. choices are randomly paired with spiking activity).

Logistic Regression:

$$f(s_b) = \frac{1}{e^{-(as_b+b)}} \quad (4.4)$$

s_b - binned spiking activity

a - slope of linear equation

b - intercept of linear equation

Broyden-Fletcher-Goldfarb-Shanno quasi-Newton algorithm is used to find the optimal fit for the slope (a) and intercept (b) (Shanno, 1970). L2 regularization is used to simplify the model and decrease the collinearity of coefficients.

Cost function for logistic regression:

$$J = \frac{\lambda}{2} \sum \beta^2 \quad (4.5)$$

J - cost associated with coefficients

λ - penalty term

β - coefficients of the model

Peri-stimulus spiking activity was used to predict the choice in the current trial and outcome of the previous trial. 5 models were created via k-fold cross validation and loss was calculated for each model and accuracy is reported as $accuracy = 1 - mean(loss)$. The following equation is used to produce the outputs of the system. Either, -1 if $f(s_b) < 0.5$ or 1 if $f(s_b) > 0.5$.

4.9 General principal component analysis (PCA) procedure of PMd firing rates

PCA was used to examine firing rate variance in the recorded PMd neural population. PCA reveals dimensions that explain a large percentage of the data while making few assumptions about the underlying structure of the data. The dimensions extracted by PCA may not always be meaningful however they often align well with behavioral variables.

The general procedure for performing a PCA involved averaging firing rates (FR) within units (U) across trials organized by reach (Re) and then by coherence (C), reaction time (RT), or past outcome (PO). Unit firing rates were calculated per trial by convolving spike times aligned to checkerboard onset (\sim -600 ms: \sim 1200 ms) with a 20 ms wide Gaussian kernel. Typical matrix organization was FR x U x Re x C/RT/PO (\sim 1800 x 996 x 2 x 7/11/2). The 95th percentile of column (i.e. unit) was calculated on this data matrix. The per column (i.e. unit) mean was subtracted from the data matrix to center the data. Then the mean centered matrix was ‘softmax’ normalized by dividing by the 95th percentile of the ‘raw’ data matrix. Softmax normalization reduces the bias of units with high firing rates and ensures that each unit has roughly the same overall variability across stimulus coherences. This data was entered into the `pca` function in MATLAB. The score output from this function is used to project the dimensions that explain the most neural variance into state space.

4.10 Behavioral Modeling

Several behavioral models were built to test whether behavioral models that include urgency terms would outperform their non-urgency counterparts. Below are the equations, with explanations for the variables, from which the behavioral models

are built. The description of the standard drift diffusion model (DDM) (figure not shown) is presented first and presentation of the models continues in the order that they are presented in figure 7. Features such as variable drift rate, starting point, non-decision time, gain term can be mixed and matched with the any of the models. Model equations and explanations of variables are adapted from (Chandrasekaran & Hawkins, 2019). For further context on these models please refer to Chandrasekaran and Hawkins, 2019.

Standard DDM:

$$x(t + \Delta t) = x(t) + v\Delta t + s\sqrt{\Delta t} + N(0, 1) \quad (4.6)$$

$x(t)$ is the state of the decision-formation process at time t (i.e. decision variable), v is the drift rate (i.e. rate of sensory evidence accumulation), Δt is the time step of the process, s is the standard deviation of the moment-to-moment noise, Brownian motion, of the decision-formation process, and $N(0, 1)$ is a random sample from the standard normal distribution with a mean of 0 and standard deviation of 1. A response is made when $x(t + \Delta t) \geq a_{upper}$ or $x(t + \Delta t) \leq a_{lower}$. Whether a response is correct or not is determined from the boundary that was crossed and the sign of the drift rate. $v > 0$ implies upper boundary corresponds to the correct response and $v < 0$ implies lower boundary corresponds to the incorrect response. For simplicity $a_{lower} = 0$ & $a_{upper} = A$.

DDM with variable drift rate (v_{ij}), starting point (s_z), and non-decision

time ($T_{er,j}$):

$$x_j(t + \Delta t) = x_j(t) + v_{ij}\Delta t + s\sqrt{\Delta t} + N(0, 1) \quad (4.7)$$

$$x_j(0) \sim U(z - \frac{s_z}{2}, z + \frac{s_z}{2}) \quad (4.8)$$

$$v_{ij} \sim N(v_i, s_v) \quad (4.9)$$

$$T_{er,j} \sim U(T_{er} - \frac{s_t}{2}, T_{er} + \frac{s_t}{2}) \quad (4.10)$$

Similar to the standard DDM model, however here the decision variable ($x_j(t + \Delta t)$) has a variable drift rate (v_{ij}) which varies as a function of condition (i) and particular trial (j). The initial value of the decision variable ($x_j(0)$) has a variable starting point (z - midpoint between bounds, s_z - variable start point constant) determined from a uniform distribution (U). The variable condition and trial variable drift rate is pulled from a normal distribution (N) with a mean of the condition average drift rate (v_i) and standard deviation is the variable drift rate constant (s_v). The trial variable non-decision time ($T_{er,j}$) is pulled from a uniform distribution with a range determined by the constant non-decision time (T_{er}) and variable non-decision time constant (s_t).

DDM with collapsing bounds:

$$a_{lower}(t) = a(1 - e^{-(\frac{t}{\lambda})^k})(0.5 - a') \quad (4.11)$$

$$a_{upper}(t) = a - a_{lower}(t) \quad (4.12)$$

$a_{lower/upper}(t)$ - position of the lower/upper boundary at time t

a - initial position of upper boundary at $t=0$

a' - asymptotic boundary setting (where lower and upper boundary meet; $a' = 0.5$)

λ - scale parameter of Weibull distribution

k - shape parameter of Weibull distribution; the Weibull function is used here to

determine the shape of the collapsing bounds as the function can approximate different distributions (e.g. exponential) (Chandrasekaran & Hawkins, 2019).

DDM with variable urgency signal:

$$E(t) = v\Delta t + s\sqrt{\Delta t}N(0, 1) \quad (4.13)$$

$$x(t + \Delta t) = x(t) + E(t)\gamma(t) \quad (4.14)$$

$$\gamma_j(t) = b_j + m_j t \quad (4.15)$$

$$\gamma_j(t) = b_j + m_j(1 - e^{(\frac{-t}{\lambda})^k}) \quad (4.16)$$

$$b_j \sim U(b - \frac{s_b}{2}, b + \frac{s_b}{2}) \quad (4.17)$$

$$m_j \sim U(m - \frac{s_m}{2}, m + \frac{s_m}{2}) \quad (4.18)$$

$E(t)$ - momentary sensory evidence at time t

$\gamma(t)$ - magnitude of urgency signal at time t

b_j - trial variable urgency intercept (chosen from a uniform distribution)

s_b - variable intercept constant

m_j - trial variable urgency slope (chosen from a uniform distribution)

s_m - variable slope constant

λ - scale parameter of sigmoid

k - shape parameter of sigmoid

4.11 Calculation of AIC and BIC

Three different measures for model selection were calculated to test the goodness of fit of behavioral models for RT and accuracy data. First an approximation to the maximum likelihood estimation, Quantile Maximum Products statistic (Heathcote et al., 2002), is calculated and used to calculate Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC). AIC and BIC are similar in calculation,

differing in the size of the penalty term for model complexity, and are therefore used to distinct ends (Aho et al., 2014; Chandrasekaran & Hawkins, 2019). AIC (Akaike, 1974) is useful for model exploration when the expectation is that models will grow in complexity with increased sample size. AIC treats false negatives as worse than false positives with a lesser penalty for model complexity than BIC and is used to judge between models that are assumed to be incorrect (Aho et al., 2014; Chandrasekaran & Hawkins, 2019). BIC (Schwarz, 1978) on the other hand treats false positives more harshly with a larger penalty term for extra parameters. The expectation here is that the correct model has been chosen and that the model will stabilize with increased sample size (Aho et al., 2014; Chandrasekaran & Hawkins, 2019). Lower AIC & BIC and higher log-likelihood, respectively, indicate better model approximation of the actual data. These scores are not objective measures of a model's goodness of fit but rather are meant to be used as a comparison between different models.

Akaike Information Criterion (AIC):

$$AIC = 2k - 2\ln(\hat{L}) \tag{4.19}$$

k - number of model parameters

\hat{L} - maximized value of the likelihood function

Bayesian Information Criterion (BIC):

$$BIC = k\ln(n) - 2\ln(\hat{L}) \tag{4.20}$$

k - number of model parameters

n - sample size

\hat{L} - maximized value of the likelihood function

4.12 Recurrent neural network (RNN) modeling description

We used PsychRNN (Ehrlich et al., 2021), a Python package, to build RNNs to model dynamics observed from our PMd data. Three training-optimized RNNs that differed in gain signal performed a computerized red-green checkerboard discrimination task. The PsychRNN package is comprehensive in that it allows the experimenter to build the task as well as to build and train the RNN.

The three RNNs contain 100 hidden layer, recurrently- and fully-connected units with the activity of the network discretized in 10 ms time steps. The hidden layer weights were randomly initialized from a Gaussian distribution. The models were trained with Adam optimization (similar to stochastic gradient descent) using mean square error as the cost function. The gradient of the loss function was found using backpropagation. Networks were trained for 50,000 trials and results from the networks were derived from testing of 10,000 trials. The equations for the recurrent state (x_t), output (y_t), and firing rate (r_t) of the models are shown below:

Recurrent State (x_t):

$$x_t = (1 - \alpha)x_{t-1} + \alpha(W_{rec}r_{t-1} + W_{in}u_t + b_{rec} + \sigma(0, 1)\sqrt{2\alpha N_{rec}^2}) \quad (4.21)$$

Output Vector (y_t):

$$y_t = \text{sigmoid}(W_{out}r_t + b_{out}) \quad (4.22)$$

Rectified linear unit (ReLU) Firing Rate (r_t):

$$r_t = \text{relu}(x_t) \quad (4.23)$$

$$r_t = \text{relu}(g_0 + x_t) \quad (4.24)$$

$$r_t = \text{relu}(g_0 * x_t) \quad (4.25)$$

u_t - input

α - scaled identity matrix, recurrent synaptic weight initialization

W_{rec} - recurrent synaptic weight matrix

W_{in} - input synaptic weight matrix

b_{rec} - constant bias into recurrent units

b_{out} - constant bias into output units

W_{out} - output synaptic weight matrix

g_0 - variable gain term

$\sigma(0, 1)$ - constant to scale recurrent unit noise; 0 mean & standard deviation of 1

N_{rec} - recurrent neural noise

4.13 Kinematic analysis of neural trajectories (KiNeT)

KiNeT (Remington et al., 2018) was used in order to characterize how state space trajectories evolve over time in terms of relative speed and position as compared to a reference (ref.) trajectory. The analysis was performed on the dynamics from the PCA on unit averaged firing rates conditioned on choice and RT. This allowed the determination of whether certain RT states were associated with faster or slower dynamics. What follows is a brief description of the KiNeT algorithm and equations used to calculate speed and position. For a more detailed description of the KiNeT methodology please refer to Remington and colleagues, 2018.

First we chose the first 10 PCs ($\sim 90\%$ of variance) as our Euclidean coordinate system. Next we chose the trajectory associated with the middle RT bin (cyan, Fig 4C) as our ‘reference’ trajectory (Ω_{ref}). To find speed (t_i - time in relation to a reference, $t_i < 0$ faster, $t_i > 0$ slower) of trajectories we found the minimum Euclidean distance (e.g. $argmin || \cdot ||$) for all points on all non-reference trajectories from the corresponding points on the reference trajectory. Time (t_i) was used to find

the corresponding positions (s_i) on the non-reference trajectories and the normalized difference between the reference trajectory position and non-reference trajectory was used to find the distance (D_i). Equations and definitions of variables are detailed below.

Non-ref. time in relation to ref. time (i.e. speed) (t_i):

$$t_i[j] = \operatorname{argmin}_{\tau} ||\Omega_i(\tau) - s_{ref}[j]|| \quad (4.26)$$

Position of non-ref. trajectory in relation to ref. trajectory (s_i):

$$s_i[j] = \Omega_i(t_i[j]) \quad (4.27)$$

Distance of non-ref. trajectory from ref. trajectory (D_i):

$$D_i[j] = ||s_{ref}[j] - s_i[j]|| \quad (4.28)$$

i - non-ref. trajectories

j - location on trajectory

Ω_i - non-ref. trajectory

τ - time associated with position j on ref. trajectory

$\Omega_i(\tau)$ - position of non-ref. trajectory at corresponding time on ref. trajectory

s_{ref} - position of ref. trajectory at point j

argmin - where function achieves its minimum at point j

4.14 Description of Latent Factors Analysis of Dynamical Systems (LFADS)

Briefly, LFADS is a generative model which assumes that neuronal spiking activity is generated from an underlying dynamical system (Pandarinath et al., 2018).

This dynamical system is assumed to be relatively low-dimensional (i.e. considerably smaller than the number of neurons involved) and latent factors can be extracted and exploited to recreate spiking activity on single trials. This method uses a trained autoencoder to generate ‘initial conditions’ based on a trial’s neurons’ spike counts. This ‘latent code’ serves as the initial condition to the generator RNN. From the latent code the generator infers the latent factors of all the neurons in that trial. Please refer to Pandarinath and colleagues, 2018 for a full description of the LFADS method.

Here LFADS was used for a single session which recorded from 23 neurons. Our model consisted of 8 factors to recreate spiking activity of single trials. The first 3 factors were visualized in Fig 5 A & E.

Bibliography

- Afshar, A., Santhanam, G., Yu, B., Ryu, S., Sahani, M., & Shenoy, K. (2011). Single-trial neural correlates of arm movement preparation. *Neuron*, 71(3): 555–564. <https://doi.org/10.1016/j.neuron.2011.05.047>
- Aho, K., Derryberry, D., & Peterson, T. (2014). Model selection for ecologists: The worldviews of AIC and BIC. *Ecology*, 95(3): 631–636. <https://doi.org/10.1890/13-1452.1>
- Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6): 716–723. <https://doi.org/10.1109/TAC.1974.1100705>
- Bogacz, R., Wagenmakers, E.-J., Forstmann, B. U., & Nieuwenhuis, S. (2010). The neural basis of the speed–accuracy tradeoff. *Trends in Neurosciences*, 33(1): 10–16. <https://doi.org/10.1016/j.tins.2009.09.002>
- Briggman, K. L., Abarbanel, H. D. I., & Kristan, W. B. (2005). Optical imaging of neuronal populations during decision-making. *Science*, 307(5711): 896–901. <https://doi.org/10.1126/science.1103736>
- Brody, C. D., & Hanks, T. D. (2016). Neural underpinnings of the evidence accumulator. *Current Opinion in Neurobiology*, 37: 149–157. <https://doi.org/10.1016/j.conb.2016.01.003>
- Brunton, B. W., Botvinick, M. M., & Brody, C. D. (2013). Rats and humans can optimally accumulate evidence for decision-making. *Science*, 340(6128): 95–98. <https://doi.org/10.1126/science.1233912>
- Carandini, M., & Churchland, A. K. (2013). Probing perceptual decisions in rodents. *Nature Neuroscience*, 16(7): 824–831. <https://doi.org/10.1038/nn.3410>
- Carland, M. A., Thura, D., & Cisek, P. (2019). The urge to decide and act: Implications for brain function and dysfunction. *The Neuroscientist*, 25(5): 491–511. <https://doi.org/10.1177/1073858419841553>

- Chandrasekaran, C., & Hawkins, G. E. (2019). Chartr: An R toolbox for modeling choices and response times in decision-making tasks. *Journal of Neuroscience Methods*, 328: 108432. <https://doi.org/10.1016/j.jneumeth.2019.108432>
- Chandrasekaran, C., Peixoto, D., Newsome, W. T., & Shenoy, K. V. (2017). Laminar differences in decision-related neural activity in dorsal premotor cortex. *Nature Communications*, 8(1): 614. <https://doi.org/10.1038/s41467-017-00715-0>
- Churchland, M. M., Afshar, A., & Shenoy, K. V. (2006). A central source of movement variability. *Neuron*, 52(6): 1085–1096. <https://doi.org/10.1016/j.neuron.2006.10.034>
- Cisek, P. (2007). Cortical mechanisms of action selection: The affordance competition hypothesis. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1485): 1585–1599. <https://doi.org/10.1098/rstb.2007.2054>
- Cisek, P. (2012). Making decisions through a distributed consensus. *Current Opinion in Neurobiology*, 22(6): 927–936. <https://doi.org/10.1016/j.conb.2012.05.007>
- Cisek, P. (2022). Evolution of behavioural control from chordates to primates. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 377(1844). <https://doi.org/10.1098/rstb.2020.0522>
- Cisek, P., & Kalaska, J. F. (2005). Neural correlates of reaching decisions in dorsal premotor cortex: Specification of multiple direction choices and final selection of action. *Neuron*, 45(5): 801–814. <https://doi.org/10.1016/j.neuron.2005.01.027>
- Cisek, P., Puskas, G. A., & El-Murr, S. (2009). Decisions in changing conditions: The urgency-gating model. *Journal of Neuroscience*, 29(37): 11560–11571. <https://doi.org/10.1523/JNEUROSCI.1844-09.2009>
- Ding, L., & Gold, J. I. (2010). Caudate encodes multiple computations for perceptual decisions. *Journal of Neuroscience*, 30(47): 15747–15759. <https://doi.org/10.1523/JNEUROSCI.2894-10.2010>
- Ditterich, J. (2006). Evidence for time-variant decision making. *European Journal of Neuroscience*, 24(12): 3628–3641. <https://doi.org/10.1111/j.1460-9568.2006.05221.x>
- Dutilh, G., Ravenzwaaij, D. V., Nieuwenhuis, S., der Maas, H. L. V., Forstmann, B. U., & Wagenmakers, E. J. (2012). How to measure post-error slowing: A confound and a simple solution. *Journal of Mathematical Psychology*, 56(3): 208–216. <https://doi.org/10.1016/j.jmp.2012.04.001>

- Ehrlich, D. B., Stone, J. T., Brandfonbrener, D., Atanasov, A., & Murray, J. D. (2021). Psychrnn: An accessible and flexible Python package for training recurrent neural network models on cognitive tasks. *eNeuro*, 8(1): ENEURO.0427–20.2020. <https://doi.org/10.1523/ENEURO.0427-20.2020>
- Ghazanfar, A. A., & Santos, L. R. (2004). Primate brains in the wild: The sensory bases for social interactions. *Nature Reviews Neuroscience*, 5(8): 603–616. <https://doi.org/10.1038/nrn1473>
- Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, 30(1): 535–574. <https://doi.org/10.1146/annurev.neuro.29.051605.113038>
- Hanks, T., Kiani, R., & Shadlen, M. N. (2014). A neural mechanism of speed-accuracy tradeoff in macaque area lip. *eLife*, 3(3). <https://doi.org/10.7554/eLife.02260>
- Hanks, T. D., Kopec, C. D., Brunton, B. W., Duan, C. A., Erlich, J. C., & Brody, C. D. (2015). Distinct relationships of parietal and prefrontal cortices to evidence accumulation. *Nature*, 520(7546): 220–223. <https://doi.org/10.1038/nature14066>
- Hanks, T. D., & Summerfield, C. (2017). Perceptual decision making in rodents, monkeys, and humans. *Neuron*, 93(1): 15–31. <https://doi.org/10.1016/j.neuron.2016.12.003>
- Hawkins, G. E., Forstmann, B. U., Wagenmakers, E. J., Ratcliff, R., & Brown, S. D. (2015). Revisiting the evidence for collapsing boundaries and urgency signals in perceptual decision-making. *Journal of Neuroscience*, 35(6): 2476–2484. <https://doi.org/10.1523/JNEUROSCI.2410-14.2015>
- Heathcote, A., Brown, S., & Mewhort, D. J. K. (2002). Quantile maximum likelihood estimation of response time distributions. *Psychonomic Bulletin & Review*, 9(2): 394–401. <https://doi.org/10.3758/BF03196299>
- Heitz, R. P., & Schall, J. D. (2012). Neural mechanisms of speed-accuracy tradeoff. *Neuron*, 76(3): 616–628. <https://doi.org/10.1016/j.neuron.2012.08.030>
- Jun, J. K., Miller, P., Hernández, A., Zainos, A., Lemus, L., Brody, C. D., & Romo, R. (2010). Heterogenous population coding of a short-term memory and decision task. *Journal of Neuroscience*, 30(3): 916–929. <https://doi.org/10.1523/JNEUROSCI.2062-09.2010>
- Kiani, R., Churchland, A. K., & Shadlen, M. N. (2013). Integration of direction cues is invariant to the temporal gap between them. *Journal of Neuroscience*, 33(42): 16483–16489. <https://doi.org/10.1523/JNEUROSCI.2094-13.2013>

- Latimer, K. W., Yates, J. L., Meister, M. L. R., Huk, A. C., & Pillow, J. W. (2015). Single-trial spike trains in parietal cortex reveal discrete steps during decision-making. *Science*, 349(6244): 184–187. <https://doi.org/10.1126/science.aaa4056>
- 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. <https://doi.org/10.1038/nature12742>
- Murphy, P. R., Boonstra, E., & Nieuwenhuis, S. (2016). Global gain modulation generates time-dependent urgency during perceptual choice in humans. *Nature Communications*, 7(1). <https://doi.org/10.1038/ncomms13526>
- Nachev, P., Kennard, C., & Husain, M. (2008). Functional role of the supplementary and pre-supplementary motor areas. *Nature Reviews Neuroscience*, 9(11): 856–869. <https://doi.org/10.1038/nrn2478>
- Okazawa, G., Hatch, C. E., Mancoo, A., Machens, C. K., & Kiani, R. (2021). Representational geometry of perceptual decisions in the monkey parietal cortex. *Cell*, 184(14): 3748–3761.e18. <https://doi.org/10.1016/j.cell.2021.05.022>
- Pandarínath, C., O’Shea, D. J., Collins, J., Jozefowicz, R., Stavisky, S. D., Kao, J. C., Trautmann, E. M., Kaufman, M. T., Ryu, S. I., Hochberg, L. R., Henderson, J. M., Shenoy, K. V., Abbott, L. F., & Sussillo, D. (2018). Inferring single-trial neural population dynamics using sequential auto-encoders. *Nature Methods*, 15(10): 805–815. <https://doi.org/10.1038/s41592-018-0109-9>
- Pastor-Bernier, A., Tremblay, E., & Cisek, P. (2012). Dorsal premotor cortex is involved in switching motor plans. *Frontiers in Neuroengineering*, 5: 5. <https://doi.org/10.3389/fneng.2012.00005>
- Pereira, M., Megevand, P., Tan, M. X., Chang, W., Wang, S., Rezai, A., Seeck, M., Corniola, M., Momjian, S., Bernasconi, F., Blanke, O., & Faivre, N. (2021). Evidence accumulation relates to perceptual consciousness and monitoring. *Nature Communications*, 12(1): 3261. <https://doi.org/10.1038/s41467-021-23540-y>
- Purcell, B. A., & Kiani, R. (2016). Neural mechanisms of post-error adjustments of decision policy in parietal cortex. *Neuron*, 89(3): 658–671. <https://doi.org/10.1016/j.neuron.2015.12.027>
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85(2): 59–108. <https://doi.org/10.1037/0033-295X.85.2.59>

- Ratcliff, R., Smith, P. L., Brown, S. D., & McKoon, G. (2016). Diffusion decision model: Current issues and history. *Trends in Cognitive Sciences*, 20(4): 260–281. <https://doi.org/10.1016/j.tics.2016.01.007>
- Remington, E. D., Narain, D., Hosseini, E. A., & Jazayeri, M. (2018). Flexible sensorimotor computations through rapid reconfiguration of cortical dynamics. *Neuron*, 98(5): 1005–1019.e5. <https://doi.org/10.1016/j.neuron.2018.05.020>
- Roitman, J. D., & Shadlen, M. N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *The Journal of Neuroscience*, 22(21): 9475–9489. <https://doi.org/10.1523/JNEUROSCI.22-21-09475.2002>
- Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, 6(2): 461–464. <https://doi.org/10.1214/aos/1176344136>
- Shadlen, M. N., & Newsome, W. T. (1996). Motion perception: Seeing and deciding. *Proceedings of the National Academy of Sciences*, 93(2): 628–633. <https://doi.org/10.1073/pnas.93.2.628>
- Shadlen, M. N., & Newsome, W. T. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *Journal of Neurophysiology*, 86(4): 1916–1936. <https://doi.org/10.1152/jn.2001.86.4.1916>
- Shadlen, M., & Kiani, R. (2013). Decision making as a window on cognition. *Neuron*, 80(3): 791–806. <https://doi.org/10.1016/j.neuron.2013.10.047>
- Shanno, D. F. (1970). Conditioning of quasi-newton methods for function minimization. *Mathematics of Computation*, 24(111): 647. <https://doi.org/10.2307/2004840>
- Stone, M. (1960). Models for choice-reaction time. *Psychometrika*, 25(3): 251–260. <https://doi.org/10.1007/BF02289729>
- Thura, D., & Cisek, P. (2014). Deliberation and commitment in the premotor and primary motor cortex during dynamic decision making. *Neuron*, 81(6): 1401–1416. <https://doi.org/10.1016/j.neuron.2014.01.031>
- Thura, D., & Cisek, P. (2017). The basal ganglia do not select reach targets but control the urgency of commitment. *Neuron*, 95(5): 1160–1170.e5. <https://doi.org/10.1016/j.neuron.2017.07.039>
- Thura, D., Cos, I., Trung, J., & Cisek, P. (2014). Context-dependent urgency influences speed-accuracy trade-offs in decision-making and movement execution. *Jour-*

nal of Neuroscience, 34(49): 16442–16454. <https://doi.org/10.1523/JNEUROSCI.0162-14.2014>

Thura, D., Guberman, G., & Cisek, P. (2017). Trial-to-trial adjustments of speed-accuracy trade-offs in premotor and primary motor cortex. *Journal of Neurophysiology*, 117(2): 665–683. <https://doi.org/10.1152/jn.00726.2016>

CURRICULUM VITAE

