




ADVERTIMENT. L'accés als continguts d'aquesta tesi queda condicionat a l'acceptació de les condicions d'ús establertes per la següent llicència Creative Commons:  <https://creativecommons.org/licenses/?lang=ca>

ADVERTENCIA. El acceso a los contenidos de esta tesis queda condicionado a la aceptación de las condiciones de uso establecidas por la siguiente licencia Creative Commons:  <https://creativecommons.org/licenses/?lang=es>

WARNING. The access to the contents of this doctoral thesis it is limited to the acceptance of the use conditions set by the following Creative Commons license:  <https://creativecommons.org/licenses/?lang=en>



Neural Network Dynamics Underlying Flexible Adjustment of Temporal Evidence Weighting in Perceptual Decision-Making

Citlalli Vivar Aburto

Supervised by Klaus Wimmer

Centre de Recerca Matemàtica

Doctorate Programme in Mathematics

UAB

Universitat Autònoma de Barcelona

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy in Mathematics

at the

Universitat Autònoma de Barcelona

January 2025

Acknowledgements

This thesis would not have been possible without the support of my friends, family, and colleagues. I am deeply grateful to all of you for your encouragement, cheers, and understanding throughout this journey. I extend my sincere thanks to my colleagues from the Barccsyn community, the Pinto Lab, Alex Huk, Mike Shadlen, and Carlos Gutierrez for their critical feedback and insightful comments, which greatly shaped the direction of this work. I would also like to express my gratitude to the CRM administration for their assistance with visa issues, conferences, fun chats, and coffee breaks. Finally, I am thankful to the FPI for their financial support, which made this research financially possible.

In particular, I would like to thank Klaus Wimmer for his generosity, guidance, and support during the course of this research. I am also grateful to Eszter, Emma, Ging Cee, Melanie, and Duygu for their uplifting take on life and their unique perspectives, which have made this experience much brighter. To Nico and Daria for being a home in what has now become a familiar land. To Butch and Botas for the necessary cuddles. To my brothers, Andrés and Bolívar, whose adventurous spirits never fail to inspire me. Thank you for always being there with a obvious solution to my problems and free medical advice. Lastly, but never least, to my parents, Bolívar and Susana, for their energetic personalities, endless love, generosity, and for being an ever-present source of inspiration through their dedication and perseverance.

Abbreviations

MT Medial Temporal Area

CP Choice Probability

RDM Random Dot Motion

LIP Lateral Intraparietal Area

FEF Frontal Eye Field

PK Psychophysical Kernel

RNN Recurrent Neural Networks

PCA Principal Component Analysis

PTA Pulse-Triggered Average

DDM Drift Diffusion Model

PSTH Peri-Stimulus Time Histogram

PFC Prefrontal Cortex

SBI Simulation Based Inference

Abstract

In perceptual decision-making, sensory information is accumulated over time to reach a categorical decision. This accumulation process can follow different strategies, where certain moments of the stimulus are weighted more heavily in the decision than others. Such temporal weighting strategies have been linked to various underlying mechanisms, including bounded accumulation, which prioritizes early stimulus information (primacy weighting); perfect integration, which treats all stimulus information uniformly (uniform weighting); and leaky integration, which favors only the later part of the stimulus (recency weighting). Moreover, experimental evidence suggests that both primates and humans can dynamically adjust their temporal weighting strategy in response to stimulus statistics, prioritizing later or early stimulus information as the task requires. Nevertheless, the proposed mechanisms fall short in explaining this flexibility, as none can generate the full diversity of experimentally observed temporal weighting patterns.

In this thesis, we combine data analysis, computational modeling, and recurrent neural networks to develop a comprehensive framework that proposes one single mechanism that can capture the full range of flexibility in the temporal weighting patterns, characterize the behavioral and neural correlates of this flexibility, and examine the optimality of these patterns.

First, we present experimental evidence of how temporal weighting flexibility is represented in the behavior of both humans and macaques. We found that this flexibility is acquired through exposure to the stimulus statistics with the state of vigilance previous to the presentation of the stimulus having a significant impact on effectively performing the task. Additionally, we analyzed middle temporal (MT) recordings from behaving macaques and found that, while the average firing rate of the neural population remains stable despite changes in stimulus statistics, stimulus-related and choice-related activity depend on the stimulus statistics in a non-trivial way.

Second, we developed a two-area firing rate model consisting of a sensory circuit and a decision circuit, connected by bottom up and top down, with a modulation signal that controls the attractor dynamics of the decision circuit. This model successfully replicates the experimentally observed flexibility in temporal weighting. Mechanistically, the modulation signal initiates the decision-making process by pushing the network into a competitive regime and modulates the decision dynamics by either accelerating or delaying the choice, similar to an urgency signal. By controlling the initiation of evidence integration and altering the signal's shape and/or duration, we were able to generate a broader range of temporal weighting patterns, extending beyond what was observed experimentally.

Third, we use Recurrent Neural Networks (RNNs) to explore whether the temporal weighting patterns observed in experimental evidence are optimal. Non-uniform weighting strategies emerge as optimal adaptations to internal noise when RNNs specialize in a single stimulus condition. These patterns are robust across networks, with task-specific low- or high-dimensional dynamics. When trained in a sequential setup similar to the macaques, the RNN adopted a rigid temporal weighting strategy biased towards the initial training condition, failing to adapt to new task demands. To overcome this limitation, we introduced contextual signals representing the task conditions which allowed simultaneous training across all conditions in a single RNN. This approach allowed the RNN to exhibit flexibility comparable to independently trained networks (aligning temporal weighting with stimulus statistics), faster learning, and generalization to novel tasks. These findings suggest that contextual signals may be the mechanism underlying the context switching observed in experimental data.

Contents

1	Introduction	1
1.1	Perceptual decision-making	2
1.1.1	Motion perception	2
1.1.2	Motion signals in MT	4
1.1.3	Decision signals in LIP	7
1.2	Temporal weighting	9
1.2.1	Flexible temporal weighting	10
1.3	Computational models of perceptual decision-making	13
1.3.1	Bounded accumulation models	13
1.3.2	Biophysical models	15
1.3.3	Firing rate models	18
1.3.4	Recurrent Neural Networks	23
2	Objectives	25
3	Materials and methods	27
3.1	Dataset	27
3.1.1	Experimental design	27
3.1.2	Electrophysiology	30
3.1.3	Data analysis	31
	Time to fixate	31
	Psychophysical Kernel	31
	Pulse triggered average and Peri-Stimulus Time His- togram	32
	Choice probability	33
	LIP correlation analysis	33
	Psychometric curve	34
3.2	Two area network model	34
3.2.1	Reduced two-variable model	34
3.2.2	Two area hierarchical spiking neural network	36
3.2.3	Two area hierarchical neural network	38
	Sensory circuit	38

Decision circuit	39
Modulation Signal	40
Simulation parameters	41
3.2.4 Simulation Based Inference	42
3.3 Recurrent Neural Networks	44
Sequentially trained RNNs	45
Principal component analysis	45
Time-lagged autocorrelation	46
Task Variance	47
3.4 Software	47
4 Experimental insights into flexible temporal weighting	49
4.1 Behavioral signatures of flexible temporal weighting	50
4.1.1 Session-to-session temporal weighting update	50
4.1.2 Time to fixate	51
4.2 Neural activity underlying flexible temporal weighting	52
4.2.1 LIP activity as a measure for engagement	52
4.2.2 Firing rate of neurons in MT area	54
4.2.3 Pulse triggered average in MT area	55
4.2.4 Choice probability in MT area	55
Conclusion	58
5 Neural network dynamics underlying the adjustment of temporal evidence weighting in perceptual decisions: a two-area neural network model	59
5.1 Two-area neural network model	60
5.2 Impact of the sensory circuit on the decision dynamics	62
5.3 Modulatory signals can influence the decision dynamics	65
5.4 Impact of circuit connectivity on temporal weighting	70
5.5 Effect on choice probability	73
5.6 Model fitting using simulation based-inference	76
Conclusion	78
6 Optimality of temporal evidence weighting: a RNN approach	81
6.1 RNN architecture and training	82
6.2 Independently trained RNNs	84
6.2.1 RNNs adapt their temporal weighting strategies to the task demands	84
6.2.2 Temporal weighting patterns are sensitivity to noise levels	85

6.2.3	Different temporal weighting strategies emerge robustly across multiple networks	86
6.2.4	Training length affects the temporal weighting strategies	87
6.2.5	Low-dimensional task-related dynamics	88
6.2.6	Neural activity underlying the decision build-up	92
6.3	Sequentially trained RNN	94
6.3.1	Sub-optimal temporal weighting	96
6.3.2	Training length does not influence temporal weighting	97
6.4	Context learning	99
6.4.1	Single-module RNNs emerge from flexible stimulus statistics tasks	101
6.4.2	Task variance and firing rates are driven by a small set of units	102
6.4.3	Modular network dynamics	103
6.4.4	Context signals facilitates generalization of temporal weighting strategies and dynamic motifs	106
	Conclusion	110
7	Discussion	111
	Mechanisms of flexible temporal weighting	111
	Origin of choice-related activity in visual brain areas .	113
	Optimality of flexible temporal weighting	115
8	Supplementary Figures	117
	Bibliography	123

Introduction

Behavior arises from the choices and decisions individuals make. As a concept, decision-making has been defined across multiple disciplines. For instance, in economics, it is often assumed that individuals are rational, making decisions that maximize their utility based on available information, preferences, and the perceived value. However, behavior is rarely rational; it is often, if not always, influenced by perception, which introduces uncertainty into the decision-making process long before the outcomes valued. In my own research journey, decision-making has played a crucial role. As an economist by training, utility theory was my "totopos and salsa", a fundamental framework for understanding behavior. While my perspective at the time was limited, it provided a strong foundation that allowed me to transition into computational neuroscience, where decision-making is explored through a wider lens of perceptual uncertainty.

Perceptual decision-making focuses on the variability of input signals that influence choices by examining how noisy sensory stimuli are processed and transformed into categorical choices. This field, has extensively studied, since the late 1980s, the neural and computational mechanisms underlying perceptual decisions across a variety of species, including rodents, monkeys and humans. In many of these studies the tasks are designed to maximize performance and the stimulus input remain relatively stable over time, promoting a strategy to sensory information being weighted equally across time. Integrating the full stimulus will maximize the reward leading to each moment of sensory input contributes equally to the final choice. However, in real-life situations, to make a decision, the brain may prioritize and adapt to efficiently interpret the most reliable information available leading to a wide range of temporal weighting patterns. This thesis centers around understanding how the brain dynamically adjusts its temporal weighting strategies in response to a changing environment, a significant aspect of decision-making that while studied experimentally, lacks a clear mechanistic explanation. We contribute to the field by presenting a comprehensive framework that proposes a mechanism for the flexible temporal weighting, characterizes this process beyond the primary experimental analysis, and allows studying the optimality of the observed patterns.

In this introduction, we begin by reviewing the key literature in motion perception, one of the corner stones of perceptual decision-making, and how this motion signals are processed by both sensory and decision related areas. We then explain what temporal weighting is in decision-making, why it matters and introduce the primary method for measuring temporal weighing flexibility. Finally, we discuss how decision-making has been modeled in the field, covering accumulation models, spiking network models, firing rate models and task-optimized recurrent neural networks, and how each contributes to our understanding of temporal weighting flexibility.

1.1 Perceptual decision-making

Perceptual decisions-making is the process by which animals transform noisy sensory signals into categorical choices (Hanks and Summerfield, 2017). These sensory signals, such as vision, smell, taste, pain, touch, hearing, balance, temperature, and proprioception, are detected by the sensory system receptors which provide the first neural representation of the external world (Kandel et al., 2000). For visual stimuli, processing begins in the retina, where the visual stimulus is converted into neural signals. This signals are then passed on to the visual cortex, including areas V1, V2, V3, V4, Medial Temporal Area (MT), and Medial Superior Temporal (MST). Neurons in the visual cortex extract low-level features of the stimulus such as edges, angles, color, orientation, direction of movement, speed, and depth, by responding selectively to specific aspects of the stimulus based on the properties of their receptive fields. In particular, MT area specializes in motion detection, with its neurons exhibiting directional selectivity. Afterwards, other parietal and frontal cortices, such as the Lateral Intraparietal Area (LIP) or Frontal Eye Field (FEF), gather and integrate the incoming visual inputs and send feedback signals to early stages of processing for action planning (Kandel et al., 2000; Roitman and Shadlen, 2002; Huk, Katz, et al., 2017; Gold and Shadlen, 2007).

1.1.1 Motion perception

One common categorical perceptual decision tasks is the classical Random Dot Motion (RDM) discrimination task. This task is widely used in perceptual decision-making because it allows for capturing the neural activity associated

with the accumulation of sensory evidence. At a neural level, it relies on the directional selectivity of MT neurons, which respond strongly when the stimulus moves in their preferred direction, linking neural responses to behavioral choices. In the initial design of the task (Newsome and Pare, 1988; Newsome, Britten, et al., 1989; Britten, Shadlen, et al., 1992; Britten, Newsome, et al., 1996), monkeys were trained to detect the direction of coherent (organized) motion within a field of moving dots while recordings were made from MT neurons. Some dots moved coherently in a specific direction, while others move randomly. The strength of the motion signal was defined by the coherence: at 0%, all dots moved randomly, and at 100% all dots moved in the same direction (Fig. 1.1). The task was structured as a two-alternative forced choice (2AFC) design, and monkeys had to discriminate the motion direction of the stimulus and respond with a saccade toward the perceived direction (left or right).




Fig. 1.1: Random Dot Motion Discrimination Task: Diagram and Performance.

The first row illustrates the stochastic motion stimulus with field of dots at varying coherence levels: 0% (random motion), 50% (mixed signal and noise), and 100% (uniform motion signal). The second row (left) illustrates the psychometric function for a single MT neuron, with the x-axis representing the strength of the motion signal (coherence) and the y-axis representing the monkey's performance, fitted with sigmoidal functions. The second row (right) illustrates the responses of a directionally selective MT neuron at three different motion correlations, in 60 trials. The hatched bars represent responses to motion in the neuron's preferred direction; the solid bars indicate responses to motion in the null direction (opposite to the preferred). *Adapted from Figures 1, 4, 5 in Britten, Shadlen, et al., 1992.*

The performance of the monkey in the RDM task is represented by the psychometric curve (Fig. 1.1), which shows that the accuracy increases with higher motion signal coherence. As coherence decreases, it becomes harder to accurately discriminate the direction of motion of the stimulus, with performance approaching chance level (50%). Lower coherence elicits smaller neuronal responses, especially as the coherence moves farther from the neuron's preferred direction (Fig. 1.1).

1.1.2 Motion signals in MT

The medial temporal MT area also known as V5, is located at the intersection of the temporal and occipital lobes in the cerebral cortex. It is reciprocally connected with primary and secondary visual areas, including V1, V2, and V4 (Kandel et al., 2000), making it a central hub for integrating visual information integration, particular motion-related signals (Born and Bradley, 2005). Neurons in MT are highly specialized for motion processing and have larger receptive fields than those in V1 allowing for detecting motion over a wider area of the visual field. Each neuron is tuned to a specific direction of motion, with its preferred direction eliciting the strongest response (Fig. 1.2), the shape of this tuning curve is roughly Gaussian. This direction selectivity allows MT to encode the trajectory of moving objects, and coordinate eye movements in response to motion (Born and Bradley, 2005).




Fig. 1.2: MT neurons direction selectivity and tuning properties. On the left, the neuronal responses to a bar of light swept across the receptive field in different directions. Each trace shows the spiking activity of the neuron as the bar was swept in the direction indicated by the arrow. The neuron's preferred direction was up and to the right Adapted from Figure 1 in Born and Bradley, 2005. On the right, tuning of an MT neuron to orientation of a stationary flash slit. Adapted from Figure 4 in Albright, 1984.

The role of MT in motion perception has been established through a combination of recordings, lesion studies, and stimulation experiments (Newsome and Pare, 1988; Britten, Shadlen, et al., 1992; Britten, Newsome, et al., 1996; Salzman et al., 1992). These studies have shown that MT neurons carry sensory information about motion direction, with direction selectivity emerging approximately 100 ms after the onset of the RDM stimulus (Fig. 1.3), reflecting the time required to process motion information. The firing rate of MT neurons increases proportionally with the strength of the stimulus (N. A. Steinemann et al., 2023) and varies in response to different levels of motion coherence (Britten, Shadlen, et al., 1993). Stronger motion coherence lead to an increase in MT's firing rate (Fig. 1.3, solid color lines) while weaker motion coherence elicit a weaker firing rate (Fig. 1.3, dotted color lines). Lesion studies highlight MT's critical role in motion perception, as damage to the area results in a decrease in performance as a result of the inability to perceive motion accurately (Katz et al., 2016; Newsome and Pare, 1988, Baker et al., 1991, Hess et al., 1989, Vaina et al., 2001).




Fig. 1.3: Response of MT neurons during RDM stimulus presentation. Average responses from direction selective neurons in area MT to motion in the preferred and antipreferred directions. After a transient, MT responds at a nearly constant rate. *Adapted from Figure 5 in Gold and Shadlen, 2007*

In classical studies, the RDM task was performed while recording from individual MT neurons. For each recording, each neuron was presented with the preferred direction (Newsome, Britten, et al., 1989; Britten, Newsome, et al., 1996). They found that neuronal activity correlates with the monkey's choice, with stronger firing rates in direction selective neurons increasing the likelihood of a decision in the neuron's preferred direction. These neurons showed a strong relationship to the behavior, suggesting that MT signals guide motion perception.

To measure the sensitivity of individual neurons to the monkey's behavior, they used Choice Probability (CP), a signal detection theory measurement that quantifies the correlation between a subject's choice and random fluctuations in a single neuron's firing rate or how well an ideal observer can predict

the choice (Britten, Newsome, et al., 1996). In the RDM task, Britten, Newsome, et al., 1996 found that the mean CP was 0.55, where 0.50 represents chance and 1.0 represents perfect correlation, this indicated that individual MT neurons weakly but significantly predict the monkey's direction decisions. In this classical interpretation, the values of CP suggests that neurons in MT are involved in the task, as their activity is correlated, trial by trial, with the monkey's decisions on the same sensory signals, potentially reflecting a contribution to behavioral choice (Britten, Newsome, et al., 1996). Furthermore, CP links these trial-to-trial fluctuations in neuronal activity to fluctuations in internal percepts as indicated by behavior (Parker and Newsome, 1998).

CP has been observed in many brain areas (Nienborg and Cumming, 2009; Celebrini and Newsome, 1994; M. R. Cohen and Newsome, 2009; Romo et al., 2002, Uka and DeAngelis, 2004, among others), regardless of whether they report sensory evidence or are involved in mediating decisions (Crapse and Basso, 2015). It has also been found that CP values tend to increase along the processing hierarchy, with sensory areas often showing values closer to chance while higher values are observed in decision-related areas (Crapse and Basso, 2015). Given the wide range of CP values and the variation across different brain areas, its emergence has been attributed to two interpretations (Wimmer et al., 2015; Crapse and Basso, 2015). The bottom-up interpretation suggest that CP reflects a causal relationship, with variability in the choice partly attributed to variability in the response of sensory neurons (Wimmer et al., 2015; Britten, Newsome, et al., 1996), and is formalized in Shadlen, Britten, et al., 1996 as a feedforward network model. Alternatively, the top-down interpretation proposes that variability in sensory neurons correlates with choice due to trial-to-trial fluctuations in top-down signals. These signals modulate the magnitude of CP and escape the control of the experimenter and cause trial-to-trial responses variability that is not necessarily noise (Wimmer et al., 2015).

Work by Nienborg and Cumming, 2009 suggests that CP cannot only be attributed to the causal effect of sensory activity on decision-making (bottom-up effect). As the impact of stimulus fluctuations on neural activity decreases over time and CP increases to a plateau, this supports the idea of a non-causal relationship mediated by top-down signals. Wimmer et al., 2015 explain the CP timecourse with a hierarchical network model composed of reciprocally connected sensory and integration circuits. The model integrates sensory evidence transiently, as seen in MT area firing rate patterns, leading to the

decaying bottom-up CP (Fig. 1.4, blue line), while top-down signals, as a result of the interaction with the decision circuit, produce a rising (Fig. 1.4, green line), and sustained CP (Fig. 1.4, black line).




Fig. 1.4: Bottom-up correlations together with top-down signals can lead to sustained CP. CP obtained in the network with top-down connections receiving non-replicate stimuli ($\sigma = 1$, $b_{FB} = 1$; black line) is sustained throughout the stimulus interval (shading). Two complementary contributions to CP are revealed by using replicate stimuli, thus removing bottom-up correlations (green line) or by removing the top-down connections ($b_{FB} = 0$; blue line). Adapted from Figure 3 in Wimmer et al., 2015.

1.1.3 Decision signals in LIP

During motion discrimination tasks, motion direction signals encoded in MT are used to make binary decisions, with subjects saccading in the direction of motion (Roitman and Shadlen, 2002). These signals are processed in high-level areas like LIP, and Frontal Eye Field (FEF) (Huk, Katz, et al., 2017; Roitman and Shadlen, 2002). Anatomically, LIP receives input from MT and project to the FEF and superior colliculus (SC) (Kandel et al., 2000). This anatomical position first motivated the recordings from LIP as it was thought the spiking activity of the neurons in the region could represent the accumulation of noisy evidence, acquired over time, leading to a decision (Huk, Katz, et al., 2017; Shadlen and Newsome, 2001). The role of LIP in decision-making has been extensively studied (Roitman and Shadlen, 2002; Kiani, Hanks, et al., 2008; Britten, Shadlen, et al., 1992; Shadlen and Newsome, 2001; Katz et al., 2016; Yates, Katz, et al., 2020; N. Steinemann et al., 2024). Its spiking activity has been directly correlated with both guiding decision-making and motor planning, as neurons respond to visual stimuli associated with planned saccadic eye movements (Roitman and Shadlen, 2002).

In the RDM task, during stimulus presentation, LIP activity gradually increases over time (Fig. 1.5, left), ramping up as if integrating evidence from sensory areas leading to the formation of the monkey's choice (Shadlen, Britten, et al., 1996; Shadlen and Newsome, 2001; N. Steinemann et al.,

2024; Roitman and Shadlen, 2002). This ramping up activity is dependent on the difficulty of the direction and strength of motion of the task (Roitman and Shadlen, 2002; N. Steinemann et al., 2024) (Fig. 1.5, left; strong coherence signals (blue) lead to a steep increase in firing rate, while weak coherence signals (orange) suppress activity) and is strongly correlated with decision outcomes (Shadlen and Newsome, 2001). LIP's ramp-like activity predicts the monkey's decision (Roitman and Shadlen, 2002); as the stimulus's direction of motion indicates the saccade target, the activity modulates accordingly, reflecting the upcoming response (Shadlen, Britten, et al., 1996; Shadlen and Newsome, 2001). During saccade preparation, motor buildup signals are observed in LIP. These signals ramp-up, with their slope proportional to the strength of the motion stimulus, reflecting LIP's role in accumulating motion evidence (Huk, Katz, et al., 2017) (Fig. 1.5, Right). The decision process is terminated when the neurons associated with the chosen target reach a critical firing rate, triggering a saccade in the chosen target's direction (N. Steinemann et al., 2024). Following this, LIPs activity decreases abruptly, suggesting that LIP accumulates evidence to a threshold level that marks the completion of the decision process (Roitman and Shadlen, 2002). The activity pattern of LIP supports the classical view of perceptual decision making, where decision areas integrate evidence provided by sensory stimuli up to a threshold and has been studied with accumulation models (see Section 1.3.1)




Fig. 1.5: Population responses from LIP. Average firing rates of the subset of neurons that represent the left target. Responses for each coherence and direction. Motion strength and direction are indicated by color (legend) and aligned to motion onset (left) or saccade initiation (right). The gray bars indicate the stimulus. LIP first registers a signal related to the strength and direction of motion in saccade-aligned response and stimulus-aligned response. *Adapted from Figure 2 in N. A. Steinemann et al., 2023.*

1.2 Temporal weighting

In the classic literature, the RDM task consists of a continuous stimulus with a fixed duration, during which evidence about the environment's state is continuously accumulated over time (Gold and Shadlen, 2007). Ideally, to maximize task accuracy, evidence should be weighed uniformly across time (Bogacz et al., 2006). In this optimal scenario, all sensory evidence contributes equally with the same weight in the decision-making process, resulting in uniform or flat temporal weighting which has been observed experimentally in some tasks (Brunton et al., 2013). To quantify this temporal weighting pattern, the Psychophysical Kernel (PK) is used. The method measures the time-course of the stimulus's impact on the decision using reverse correlation or logistic regression analysis (Neri, 2004; Okazawa et al., 2018; Nienborg and Cumming, 2009; Levi, Yates, et al., 2018). In this thesis, we will use the logistic regression analysis approach (section 3.1.3).

In many tasks, the weight assigned to sensory information over time is not constant. Rather, sensory evidence is integrated sub-optimally, with an unequal weighting over time (Drugowitsch et al., 2016) or accumulated discontinuously during decision making (Wyart, De Gardelle, et al., 2012). As a result, subjects assign more weight to some stimulus epochs than others. For instance, they place more weight to early evidence known as primacy weighting (Kiani, Hanks, et al., 2008; Wimmer et al., 2015; Yates, Park, et al., 2017), more weight to evidence later in time, referred as recency weighting (Cheadle et al., 2014; Drugowitsch et al., 2016; Tsetsos et al., 2012; Wyart, Myers, et al., 2015; Levi, Yates, et al., 2018) or more weight on intermediate portions of the stimulus known as non-monotonic weighting (Keung et al., 2019). These deviations from uniform weighting have been attributed to various factors. One factor is the specific characteristics of the stimulus, with shorter stimuli eliciting primacy and longer stimuli evoking recency (Bronfman, Brezis, and Usher, 2016). Another factor is the magnitude of noisy stimulus fluctuations, where increasing the noise cause a shift from primacy to recency weighting (Prat-Ortega et al., 2021). Task difficulty may also play a role, specifically on whether evidence is noisy at the level of sensory information or at the level of category information (uncertain category membership; Lange et al., 2021).

1.2.1 Flexible temporal weighting

An important factor could be that temporal weighting is flexible and adapts to the demands of the task (Levi, Yates, et al., 2018), making it highly context dependent (Talluri, Urai, Bronfman, et al., 2021). Evidence of this flexibility has been observed in various experimental settings (Levi, Yates, et al., 2018; Talluri, Urai, Tsetsos, et al., 2018; Luu and Stocker, 2018; Talluri, Urai, Bronfman, et al., 2021), where subjects adjust their weighting strategy or switching between weighting strategies depending on the task demands. In this thesis, we explore how temporal weighting flexibility emerges as a response to changes in stimulus statistics. For this, we use data from Levi, Yates, et al., 2018 and Levi, Zhao, et al., 2023. To study this flexibility, Levi, Yates, et al., 2018 designed a stimulus consisting of seven consecutive pulses, each lasting 150 ms, with five coherence levels: strong left, weak left, zero, weak right, and strong right (See Section 3.1 for details) and manipulated the stimulus statistics in three conditions: Flat condition, where motion strength remained constant on average across pulses (Fig. 1.6, Flat); Late condition, with reduced motion strength in the first three pulses and increased motion strength in the later pulses (Fig. 1.6, Late); and Early condition, with increased motion strength in the first three pulses followed by a decrease (Fig. 1.6, Early).




Fig. 1.6: Temporal stimulus statistics. Subjects viewed a sequence of 7 motion pulses of 150 ms duration. The average stimulus strength per pulse (5 levels with the sign encoding stimulus direction) is varied in three different conditions: flat condition, with constant average motion strength; late condition, with reduced motion strength in the first 3 pulses so that late pulses are most informative; and early condition, with inverse time course.

Levi, Yates, et al., 2018 found that both humans and monkeys could perform the task accurately (Fig. 1.7, C and D, respectively) with minor differences between species and could flexible adapt their temporal weighting patterns to the task demands, producing either early or late weighting depending on the stimulus statistics (Fig. 1.7, A and B). When early epochs of the stimulus had a higher mean motion strength (Flat condition and Early condition, Fig. 1.7, blue and red lines in A and B) compare to the rest of the stimulus, subjects

adopted an early weighting strategy (primacy), placing more weight on the initial information. In contrast, when the mean motion strength was higher in later epochs of the stimulus (Late condition, Fig. 1.7, yellow lines in A and B), subjects shifted to late weighting strategy (recency), giving more importance to the later information. In particular, humans were able to more flexibly adjust to the stimulus statistics in the Flat condition, exhibiting a kernel with a more uniform distribution compared to monkeys (Flat condition, Fig. 1.7, blue line in A vs B).




Fig. 1.7: Comparison of temporal weighting and psychometric functions within species across stimulus conditions. A, B) Temporal weighting profiles for macaques (A) and humans (B) averaged over all sessions in the early-, flat-, and late-stimulus conditions, fitted by a linear model (semitransparent lines) to capture the overall trend of the weights. Error bars represent ± 1 SEM. C, D) Psychometric behavior of macaques (C) and humans (D) averaged over all sessions in the early-, flat-, and late-stimulus conditions, fitted by a logistic function to capture the dependence of choice on stimulus strength. Error bars represent ± 1 SEM. Adapted from Figure 3 in Levi, Yates, et al., 2018.

Additionally, in a follow up paper, the authors (Levi, Zhao, et al., 2023) measure the timecourse of the CP with two methods: Logistic regression (logReg) and latent factors (vLGP) (Fig. 1.8). Overall, they found larger values, around 0.6, of CP, not so characteristic of sensory areas (see Section 1.1.2 and Section 7). CP was consistently highest late during the stimu-

lus period, regardless of the changes in temporal weighting strategy across conditions (Fig. 1.8, A, B, C). In particular, the authors found a positive correlation between the task-relevant direction sensitivity and the stimulus (Fig. 1.8, D). However, they also found no correlation between task-relevant direction sensitivity and the choice (Fig. 1.8, E). This suggests that the choice-related activity is distributed across multiple uncorrelated dimensions. It also suggests that neurons encoding sensory stimuli differ from those encoding choices, with MT neurons in the sensory dimension carrying minimal information about the choices(Levi, Zhao, et al., 2023).




Fig. 1.8: Choice-correlated activity in MT distinct from motion encoding or psychophysical readout of motion signals. timecourse of population-decoded choice probability during flat (a), late (b), and early (c) conditions. Solid versus dashed lines denote stimulus versus choice dimensions, respectively. Darker traces in the foreground denote CP calculated from latent factors, while semi-transparent traces denote logistic regression-based CP. (d) Individual neuron weights in the logReg model to predict direction (left) and choice (right) as a function of task-relevant direction sensitivity (d'). The d' value of a neuron was strongly correlated with its weight on the direction axis ($r = 0.651$) and uncorrelated with its weight on the choice axis ($r = 0.039$). (e) d' value of a neuron was strongly correlated with the vLGP loading weights for the stimulus dimension ($r = 0.484$), and was only weakly correlated with the loading weights for the choice dimensions ($r = 0.106$). Adapted from Figure 4 in Levi, Zhao, et al., 2023.

Multiple experimental findings in different species have demonstrate that perceptual decisions involve a temporally flexible weighting process (Levi, Yates, et al., 2018; Levi, Zhao, et al., 2023; Talluri, Urai, Bronfman, et al., 2021; Jazayeri and Movshon, 2007; Luu and Stocker, 2018; Talluri, Urai, Tsetsos, et al., 2018). However, the mechanisms underlying this flexibility remains

unclear, as current modeling framework can not explain the experimentally observed flexibility in temporal weighting patterns. For instance, primacy is explain as a consequence of bounded accumulation (Huk and Shadlen, 2005; Kiani, Hanks, et al., 2008) as information is accumulated until it reaches a bound and then ignored as irrelevant. Early stimulus epochs contribute more to decisions than late (Levi, Yates, et al., 2018). In contrast, recency weighting is consistent with leaky integration (Usher and McClelland, 2001), where information that is accumulated early on, has more time to leak away (Levi, Yates, et al., 2018). In this model, early sensory evidence contributes less to decisions compared to late. Each of this mechanisms are mutually exclusive; meaning that a single model is unable to account for both patterns. In the next section, we take a closer look at three main types of perceptual decision models and at their temporal weighting strategies.

1.3 Computational models of perceptual decision-making

In this section, we examine how temporal weighting is represented in the key computational models (see Section 1.3), including bounded accumulation models (see Section 1.3.1), biophysically inspired frameworks (see Section 1.3.2), and firing rate models (see Section 1.3.3) and why each of this models are not suited to explain the temporal weighting flexibility. Complementary, we talk about the role of recurrent neural networks in perceptual decision making and flexible temporal weighting.

1.3.1 Bounded accumulation models

Bounded accumulation models, are among the most commonly used frameworks in perceptual decision making due to their tractability and ability to explain behavior. These models are particularly well-suited to explain reaction times and accuracy in two-choice tasks and provide a mechanistic explanation of how neural activity builds up over time in decision related areas (such as LIP, see Section 1.1.3) during perceptual decision tasks. They focus on evidence accumulation, where noisy information is integrated over time until the accumulated evidence or decision variable, for one of the choices reaches a decision bound. At this point, decision is made and action

is initiated (Cisek and Thura, 2022; Lim et al., 2017; Kiani, Hanks, et al., 2008). Among these models, the Drift Diffusion Model (DDM) is one of the most widely used to explain the evidence accumulation in the RDM task. In this task, motion information is presented in two opposite directions, represented as h_1 (right) and h_2 (left) in Fig. 1.9. The DDM (Shadlen and Kiani, 2013; Gold and Shadlen, 2007; Kiani and Shadlen, 2009) interprets this sensory evidence for decision making as the difference in firing rates between pools of direction-selective neurons in a decision-related area, like LIP. Each is tuned to one of the two motion directions, forming the decision variable. At each moment, this difference in neural activity is modeled as a stochastic process, illustrated by the irregular trace in Fig. 1.9, with its mean proportional to the motion strength (or coherence) of the stimulus. The mean of this difference corresponds to the drift rate of the diffusion process (e) (Shadlen and Kiani, 2013). In tasks involving the RDM, lower motion coherence reduces the signal-to-noise ratio, increasing the probability of an error, leading to longer reaction times as the decision variable takes longer to reach the bound. Likewise, higher motion coherence improves the signal-to-noise ratio, allowing for faster evidence accumulation and shorter reaction times and larger accuracy (Cisek, Puskas, et al., 2009).




Fig. 1.9: Drift diffusion with symmetric bounds Noisy momentary evidence for/against a choice h_1 (left) and h_2 (right) is accumulated (irregular trace) until it reaches an upper or lower decision bound, leading to a choice in favor of h_1 or h_2 . Adapted from Figure 2 in Shadlen and Kiani, 2013.

While bounded accumulation models provide a solid framework for explaining behavior, they have limitations. These models are not able to explain the temporal weighting flexibility found in nature (see Section 1.2.1). In its most basic form, the model assumes pure evidence accumulation, where the accumulation process continues indefinitely due to the absence of decision bound or because the threshold is set to infinity (Ratcliff and McKoon, 2008).

As a result, sensory evidence is accumulated continuously and uniformly over time, leading to optimal weighting (Bogacz et al., 2006) or uniform temporal weighting (Zhang et al., 2009; Brunton et al., 2013). When bounds are introduced, as in the DDM, evidence accumulation favors initial information, resulting in primacy weighting (Kiani, Hanks, et al., 2008; Zhang et al., 2009) because parts of the stimuli displayed after reaching the bound do not contribute to the integration. Other modifications to the bounds introduce different dynamics in the decision process and temporal weighting patterns. With reflecting bounds, each boundary is hit, there is a partial loss of accumulated evidence. This is attributed to not incorporating momentary the evidence around the bound (Zhang, 2012). As a result, early evidence is partially lost and the decision becomes increasingly influenced by later evidence, leading to recency weighting (Zhang et al., 2009). In contrast, with collapsing bounds, decisions are triggered with less evidence as time passes, terminating the evidence accumulation process earlier than a model with fixed boundaries (Hawkins et al., 2015; Cisek, Puskas, et al., 2009). As a result, this results in more transient evidence integration, leading to a primacy weighting (Prat-Ortega et al., 2021). Additionally, Prat-Ortega et al., 2021 demonstrated that these dynamics of evidence accumulation are invariant to changes in stimulus parameters, highlighting the distinct properties of each bound type.

The DDM and variants successfully explain a wide range of experimental data (psychometric curves, speed-accuracy trade-off, etc. (Ratcliff and McKoon, 2008; Palmer et al., 2005; Zhang, 2012)). However, these models are phenomenological, meaning they account for behavioral outcomes with precision without providing insight into the underlying neural mechanisms. They also, by design, fail to capture the complex dynamics of the brain during perceptual decision making tasks. To address this gap, we explore biophysical models, which incorporate the dynamics of specific neural circuits and spiking properties of the neurons providing a more mechanistic perspective on the decision dynamics.

1.3.2 Biophysical models

In contrast to bounded accumulation models (see Section 1.3.1), biophysical models offer a mechanistic framework for studying cognitive processes (Brunel and X.-J. Wang, 2001; XJ, 2002; Compte et al., 2000) by incorpo-

rating realistic biophysical features of the neurons, such as spiking activity, synaptic inputs, ion channels, or membrane potentials. In particular, in perceptual decision making, these models emphasize the role of individual neurons' spiking activity in information representation. By modeling neural circuits using spiking neurons (e.g., integrate-and-fire models (Larry F Abbott, 1999), cortical network models (XJ, 2002; Wimmer et al., 2015), etc), these models capture the high degree of variability and noise present at the single neuron data (Gerstner and Kistler, 2002), linking this activity to behavioral outcomes in decision related areas (Wimmer et al., 2015) like LIP and Prefrontal Cortex (PFC).

Among the various biophysical models used in perceptual decision-making, two stand out in capturing the mechanisms underlying the decision process during to simulate a two-choice decision task. The recurrent cortical network model utilizes networks of spiking neurons with recurrent connections to simulate how sensory evidence is integrated over time. In this model (Fig. 1.10), two neural population, each selective to one of the two motions directions of the stimulus, have strong recurrent excitatory connections mediated by NMDA receptors that generate self-sustained persistent activity (slow dynamics that enhances the network's stability X.-J. Wang, 1999). These populations compete against each other through feedback inhibition provided by interneurons (XJ, 2002). The interaction between slow recurrent excitation and winner-take-all competition, mediated by feedback inhibition, generates attractor dynamics. In this process, the firing rates of both populations gradually ramp up, leading to the sustained activation of one population while the activity of the other is suppressed (XJ, 2002). This attractor dynamics explain the decision-correlated neural activity with slow temporal integration that resembles that in LIP (see Section 1.1.3), as well as behavioral outcomes such as choice and reaction times.

Hierarchical spiking networks (Wimmer et al., 2015), build on the principles of recurrent cortical network models (XJ, 2002), with a two-layered architecture that includes circuit specialization in sensory and decision-making areas. In this model (see Section 3.2.2 for details), the sensory circuit is responsible for encoding and processing incoming stimuli, and consists of two neural populations selective to opposite motion directions that are coupled to two choice associated populations in the integration circuit (Fig. 3.3, A, middle row). The decision circuit accumulates sensory evidence until the network reaches the attractor state associated with one of the choices through




Fig. 1.10: Recurrent cortical network models. Two pyramidal cell groups (A and B), each of which is selective to one of the two stimuli (mimicking motion to the right or left). Each pyramidal neural group has strong recurrent excitatory connections that can sustain persistent activity triggered by the preferred stimulus. The two neural groups compete through feedback inhibition from interneurons. *Adapted from Figure 1 in XJ, 2002*

winner-take-all competition between two decision-encoding populations (Fig. 3.3, A, top row). Both circuits are connected through feedforward and top-down feedback (Fig. 3.3, A, green lines) and have lateral excitatory and inhibitory (Fig. 3.3, A, population I) recurrent connections within each circuit (Wimmer et al., 2015). This model characterizes the dynamics of sensory processing and decision integration, including the timecourse of the stimulus and choice related activity and their contribution to the sustained time-course of the CP(see Section 1.1.2, Fig. 1.4). As the network accumulates sensory information, the early component increases. However, as the network approaches the attractor state, the impact of sensory activity fluctuations on the upcoming decision diminishes (Wimmer et al., 2015) and the choice-related component increases.

These models are flexible enough to capture multiple aspects of the mechanisms underlying the decision process, such as decision dynamics (XJ, 2002), choice probability (CP) (Wimmer et al., 2015), and behavioral outcomes. However, they have limitations in capturing the observed temporal flexibility of the experimental data. In XJ, 2002, the PK shows primacy (Prat-Ortega et al., 2021), as well as in the hierarchical spiking network model by Wimmer et al., 2015, despite the sustained CP timecourse observed in the network (see Fig. 1.11). The primacy has been attributed to multiple mechanisms. On one hand, this could be a result of decision-related mechanisms, such as the accumulation to a bound (Wimmer et al., 2015; Kiani, Hanks, et al., 2008), where information is integrated as it comes, or it could reflect the adaptation of the network to the stimulus (Yates, Park, et al., 2017).




Fig. 1.11: Psychophysical kernel in the hierarchical spiking network model. Temporal weighting pattern resulting from the hierarchical neural network showing primacy on every component interaction. *Adaptation from Figure 3 in Wimmer et al., 2015.*

Biophysical models are an excellent starting point for our work, as they provide mechanistic insights into decision dynamics and offer a biologically grounded perspective on how sensory evidence is processed within neural circuits. However, one downside is that these models are mathematically complex, making them difficult to tractably implement. They require numerous parameters and equations to account for the intricate dynamics (Kriegeskorte and Douglas, 2018), which makes them computationally expensive and challenging to scale to multiple brain areas.

1.3.3 Firing rate models

Firing rate models provide a simplified yet biologically plausible approach to understanding neural dynamics by focusing on the dynamics of neural populations rather than the individual spikes of neurons. These models capture the overall network activity while maintaining some key biological processes and avoiding the complexity of detailed biophysical models. In particular, the rate model from Wong and X.-J. Wang, 2006 behaves qualitatively equivalent to the spiking neural network (XJ, 2002). However, instead of describing a neuron's spike time it describes the population firing rate, this models advantage is its simplicity, faster computations and analytical tractability. The model reduces the spiking neural network from 2000 spiking neurons (XJ, 2002) to a two-variable model (Wong and X.-J. Wang, 2006) using the mean field approach (Fig. 1.12) .

The two-variable rate model consists of two excitatory populations (r_1 and r_2), each selective for the motion of a stimulus in a two-choice task. Each population receives external inputs (I_1 and I_2 , Fig. 1.12) and exhibits cross-coupling and positive recurrent connectivity (w_+ , Fig. 1.12). The populations compete via mutual inhibition, mediated by inhibitory interneurons. Detail




Fig. 1.12: Reduction of a biophysical neuronal decision-making model. Reduction of a biophysical neuronal decision-making model. The original model (top) has strong recurrent excitation between neurons with similar stimulus selectivity, and effective inhibition between them via shared inhibition. NS and I denote the nonselective excitatory (black) and inhibitory (green) pools of cells, respectively. Arrows, Excitatory connections; circles, inhibitory connections. I_1 and I_2 are inputs from external stimulus to selective neural populations 1 (blue) and 2 (red). Brown arrows, Background noisy inputs. w denotes enhanced excitatory connections within each selective neural pool. The numbers on the right displays the total number of dynamical equations involved in the model. The mean-field approach reduces 2000 spiking neurons into four neural units (with a total of 11 dynamical variables), Simplify the linear input-output relation (F-I curve) of the cells. The final reduced two-variable model (bottom) consists of two neural units, endowed with self-excitation and effective mutual inhibition. Adapted from Figure 1 in Wong and X.-J. Wang, 2006

equations of this model are available in Section 3.2.1. The model also displays attractor dynamics where the firing rates of the populations gradually ramp up due to competition until one population dominates and wins. The coupling between the populations increases with greater recurrent connectivity in the attractor network, which resembles the attractor dynamics found in LIP (Hart and Huk, 2020). Eventually, the system reaches the steady state, with the outcome of the competition, where the firing rate of the winning population stabilize, and the competing population's activity is suppressed. This model can exhibit different steady states, with high neural activity in the presence or absence of an external stimulus (Wong and X.-J. Wang, 2006).

Wong and X.-J. Wang, 2006 characterize the steady states of this system in a bifurcation diagram (Fig. 1.13), where the synaptic gating of both populations is equal ($S_1 = S_2$ (middle curve of Fig. 1.13),), as a function of the absolute stimulus strength μ_0 (Hz) (Wong and X.-J. Wang, 2006). Before stimulus presentation, the system is in an stable steady state (black square of Fig. 1.13, one symmetric steady state with only one attractor (Wong and X.-J. Wang, 2006)). As the stimulus strength increases, the steady state loses its stability (dotted line between black square and open square of Fig. 1.13). When the saddle steady state appears (open square of Fig. 1.13, the network goes from monostable to bistable, a coexistence between one symmetric and two asymmetric attractor states (Wong and X.-J. Wang, 2006)). In this transition is where competition occurs, bistability relies on strong recurrent connections and long-range inhibition (Wong and X.-J. Wang, 2006). The state can transition toward the upper or lower stable state (filled circles of Fig. 1.13, where the upper and lower circle correspond to asymmetric states in S ; S_1 is high and S_2 is low, or vice versa, a bifurcation of the system or an attractor basin). In this case, S_1 wins the competition (winner label of Fig. 1.13) and S_2 loses. The system regains its stability where the saddle point turns into an attractor (arrow with double asterisks of Fig. 1.13). After removal of stimulus, a sustained persistent activity of the decision is stored in working memory (black triangle of Fig. 1.13). In addition to being influenced by the stimulus strength, attractor network models have been shown to have their decision dynamics affected by several factors. A modulatory increase in the recurrent excitation and in lateral inhibition of the system may turn the decision state from an accumulator to attractor dynamics, in which the state is relatively insensitive to new evidence (Bronfman, Brezis, Moran, et al., 2015; Usher and McClelland, 2001). Neuromodulators, like norepinephrine, when released in the cortex, mediate attention (Usher and




Fig. 1.13: Bifurcation diagram. A) Bifurcation diagram of a selective population with stimulus strength 0 as a parameter ($c = 0\%$). Bold lines represent stable steady states; dashed lines represent saddle steady states. The spontaneous state before stimulus presentation is denoted by the filled square. With a $0 \rightarrow 30$ Hz stimulus, the spontaneous stable state loses stability, and a saddle steady state appears (open square). The state either transitions toward the upper or lower stable state (filled circles). The population wins the competition if the upper branch is chosen, and loses otherwise. When the stimulus is removed, hysteresis of the upper stable branch allows the activity to persist (memory storage of a decision choice). The arrow with an asterisk indicates the point where the spontaneous state loses stability. The arrow with double asterisks indicates the point where the saddle point turns into an attractor. *Adapted from Figure 10 in Wong and X.-J. Wang, 2006.*

Davelaar, 2002), affecting the trade-off between speed and accuracy and the accumulation of evidence, ultimately pushing the system into or out of an attractor state (Bogacz et al., 2006; Usher and Davelaar, 2002). Theoretical studies have suggested that top-down signals from higher level decision areas could provide a substrate for changing the time-course of evidence integration (Wimmer et al., 2015; Haefner et al., 2016). Lastly, changes in the input like an external ramping input may change the attractor dynamics (Finkelstein et al., 2021; Inagaki et al., 2022).

The model we propose in this thesis builds on the computational framework of the firing rate model by Wong and X.-J. Wang, 2006, combined with the hierarchical architecture described in Wimmer et al., 2015. The novelty of our model lies in incorporating a modulatory signal that influences the attractor dynamics within the decision circuit while also providing a mechanistic explanation for the flexible temporal weighting patterns observed

in experimental data (see Section 1.2). Details about the nature of the modulatory signal can be found in Section 7.

In particular, our modulatory signal has a similar effect on the system's dynamics as the external ramping input described in Inagaki et al., 2022 and Finkelstein et al., 2021. However, our model specifically focuses on providing a mechanistic explanation for the dynamics underlying flexible temporal weighting. The ramping input signal an increasing commitment to a selected action, similar to an urgency signal (Cisek and Kalaska, 2010; Thura and Cisek, 2014). It's role is to push the system towards the decision as time progresses. This is achieved by dynamically shifting the energy landscape over time (Fig. 1.14) , deforming it and increasing the separation between attractor basins (Finkelstein et al., 2021) The separation between the attractors is controlled by the amplitude of the ramping input: weaker inputs result in greater fluctuations, moving between attractor states, while stronger ramping inputs stabilize the system (Inagaki et al., 2022).




Fig. 1.14: Illustration of an external ramping model In the external ramping model, the energy landscape is reshaped over time by an external ramping input. The ball is always near the bottom of the valley. When attractors are still shallow early in the delay epoch (middle landscape), distractors/perturbations can kick the ball above the ridge to move it into another attractor (light blue ball and purple dotted line). When attractors are deep prior to action (right landscape), the same kick cannot move the ball out of the valley. *Adapted from Figure 3 in Inagaki et al., 2022*

In terms of modeling temporal weighting strategies, attractor models operating in the typical regime with strong attractor dynamics (XJ, 2002) make it increasingly difficult to reverse an upcoming choice as the network evolves towards an attractor, leading to early temporal weighting (Prat-Ortega et al., 2021). However, increasing stimulus fluctuations or duration can promote transitions between decision states, resulting in a crossover between weighting primacy and recency (Prat-Ortega et al., 2021). While such models are

flexible enough to capture these temporal weighting patterns, they cannot account for all observed behaviors on their own. Further testing has shown that when a Wong and X.-J. Wang, 2006-inspired network (see Section 5.1, Fig. 5.1) is presented with changing stimulus statistics (Levi, Yates, et al., 2018), primacy persists despite these alterations in stimulus statistics. Our model aims to provide a mechanistic explanation for the dynamics underlying flexible temporal weighting. By incorporating a modulatory signal into the decision circuit, we clarify how these dynamics support the flexible adaptation of temporal weighting.

1.3.4 Recurrent Neural Networks

The models discussed above focus on explaining mechanisms, dynamics, or behavior. However, Recurrent Neural Networks (RNNs) take a different approach. While also being computational models inspired by biological recurrent circuits, RNNs are design to perform tasks rather than explaining the underlying dynamics. Through training, RNNs continuously learn and adapt their internal states to meet any task requirements. They excel at capturing complex dynamics, uncovering hidden features, and identifying patterns that are not readily observable.

In computational neuroscience, RNNs have been widely use to study multiple cognitive tasks (Sussillo, 2014; Driscoll et al., 2024; Barak, 2017; Yang and X.-J. Wang, 2020; Rajan et al., 2016; Mante et al., 2013). For instance, Yang, Joglekar, et al., 2019 trained a single network to perform 20 cognitive tasks—including working memory, categorization, perceptual decision-making, context-dependent decision making, inhibitory control, and multi-sensory integration. The authors showed that a single RNN can effectively learn the unique rules of each task while identifying both the distinct and shared features. Tasks requiring similar computational elements reused dynamical motifs, such as attractors or decision boundaries (Driscoll et al., 2024) allowing the RNN to generalize knowledge across tasks and quickly learn new ones. This ability to generalize reflects a learning process in line with human cognition (Driscoll et al., 2024).

The flexibility of RNNs is particularly useful for comparing behaviors between humans and animals (Yang and X.-J. Wang, 2020) but this flexibility comes at the cost of interpretability (Yang and Molano-Mazón, 2021). RNNs function

as 'black-box' models, making it challenging to interpret the internal dynamics. Additionally, the architecture of RNNs does not align with the brain's organization (Yang and Molano-Mazón, 2021) which limits the ability to derive mechanistic insights from them. Furthermore, the RNNs rely heavily on gradient-based optimization techniques, such as backpropagation-through-time, these methods are not biologically plausible partly because it is not causal, as information flows backward in time contradicting the forward, time-dependent nature of biological neural signals (Yang and Molano-Mazón, 2021).

Despite their limitations, RNNs serve as an excellent tool for testing a wide range of decision-making conditions, like the changing stimulus statistics task in Levi, Yates, et al., 2018, the sequential training of the monkey (see Section 1.2), and context-dependent learning (Yang, Joglekar, et al., 2019; Mante et al., 2013; Song et al., 2016). Simulating these task allows us to explore how the temporal weighting patterns observed in experimental data (see Section 1.2) compare to the RNN's weighting patterns, evaluate their optimality, and complement other mechanistic derivations.

Objectives

We have established that primates can flexibly adapt their temporal weighting strategies based on stimulus statistics, exhibiting a wide range of temporal weighting patterns. This flexibility remains unexplained by current perceptual decision-making models, as fixed mechanisms such as bounded accumulation or leaky integration cannot replicate the patterns observed experimentally. The main objective of this thesis is to provide a comprehensive framework for understanding the mechanisms of flexible temporal weighting in perceptual decision-making, organized into three central aims:

1. Investigate behavioral and neural signatures of flexible evidence weighting in experimental data. In Section 4, we analyze the behavior of humans and monkeys as well as recordings from the MT and LIP area of behaving monkeys. Our goals are:

- To characterizing previously unexplored behavioral correlates of the temporal weighting patterns.
- To link neuronal activity of sensory and decision areas to the varying stimulus statistics and explore its adaptation to the different statistical contexts.

2. Build a computational model that explains the neural dynamics underlying flexible weighting. In Section 5, we propose a two area network computational model that captures the interactions between decision and sensory circuits while performing the varying stimulus statistics task. Our goals are:

- To illustrate the dynamics that each neural circuit plays in the decision-making process.
- To provide a unified mechanism that accounts for the flexibility in temporal weighting patterns and can be generalized to diverse contexts.

- To explore the role of circuit connectivity strength in weighting flexibility and its impact on choice probability.

3. **Explore whether flexible weighting is an optimal strategy that arises in trained Recurrent Neural Networks.** In Section 6, we explore the optimality of these weighting patterns by training and testing multiple recurrent neural networks on the task with changing stimulus statistics. Our goals are:

- To investigate if weather recurrent neural networks can exhibit temporal weighting flexibility, meet the task demands, and produce patterns that align with optimality.
- To explore which training scenarios—three individual networks, sequential training to resemble the original experimental environment, or a single network with combined knowledge—best support the emergence of temporal weighting flexibility.
- To assess the network’s ability to generalize its learning to novel environments.

Materials and methods

3.1 Dataset

3.1.1 Experimental design

Full experimental details of the data set can be found in Levi, Yates, et al., 2018 and Levi, Zhao, et al., 2023. Behavioral data was recorded from two rhesus macaque monkeys (Nancy (N), and Leo (L)) and three humans subjects (A,H,K). Subjects participated in an motion discrimination task where they viewed a motion stimulus that consisted of a sequences of seven motion pulses. The stimulus was designed this way to control its time course by independent modifying the motion strength of each of the motion pulses. They were required to discriminate the net direction of the motion stimulus with a saccade to one of two response targets, either left or right, corresponding to the perceived direction.

The trial began with a fixation point in the center of the screen, followed by two targets appearing on each side of the screen that remained visible for the duration of the trial. The motion stimulus was then presented consisting of seven pulses, each with a 150 milliseconds duration, making the total duration of the stimulus 1050 milliseconds. Following the motion stimulus, there was a brief random delay period and finally, a go cue indicating subjects to respond with a saccade (see Fig. 3.1 A).

The motion strength, or coherence, of each of the seven pulse (i) was determined by the proportion of the 19 Gabor patches (see Fig. 3.1 B) moving in the same direction. The value was taken from a Gaussian distribution, $X_i \sim N(\mu_k, \sigma)$ where k represented the trial type (strong left, weak left, zero-mean, weak right and strong right) and μ_k the corresponding mean values indicating the direction(-50% , -10% , 0% , 10% , 50%) and σ the variance (15%). The sum of the seven motion strengths determined the net direction, which was rewarded when the saccade matched the direction, regardless of the trial type (see Fig. 3.1 C).




Fig. 3.1: Motion discrimination task and Gabor patch stimulus. A) Subjects fixated on a central point through the appearance of targets and motion stimulus until the disappearance of the fixation point (“go”). Choices were made with saccades to the target corresponding to the perceived net direction of motion. Initial fixation time, target-on duration, and time until fixation point disappearance were randomly varied. B), An example frame of the Gabor motion pulse stimulus. The stimulus is composed of 19 Gabor patches, where motion strength is denoted by the proportion of coherently drifting Gabors out of the total number elements in the stimulus. C), Motion pulse values are generated from Gaussian distributions spanning a large range of possible motion strengths in either direction. A single trial consists of seven motion pulses, each randomly drawn from one of the Gaussians. Example trials with pulses drawn from each Gaussian (strong left/right, weak left/right, and zero-mean) are presented in cartoon form where the number of arrows represents the number of coherently drifting Gabor elements. Reproduced from Levi, Yates, et al., 2018 with permission of the publisher under a Creative Commons Attribution 4.0 International License.

Complementary to the classic motion discrimination task where the coherence remained constant over the stimulus duration (Gold and Shadlen, 2007; Britten, Newsome, et al., 1996; Britten, Shadlen, et al., 1992 and Newsome and Pare, 1988; see Introduction for a detailed review on the topic) this experimental paradigm focused on the temporal manipulations of the stimulus where μ_k had freedom to change from pulse to pulse. This manipulation could be classified on three conditions: flat-stimulus, where the mean motion strength per pulse was constant (just as the classic literature), late-stimulus, where the late pulses had a higher mean motion strength compared to the initial pulses and early-stimulus where early pulses had a higher mean motion strength compared to later pulses (Fig. 3.2 C). In all stimulus conditions,

the motion strength was shown for all the five trial categories previously mentioned with an addition of 5-10% of 'frozen noise trials' (repetitions of the exact same stimulus), where the net direction always summed zero.




Fig. 3.2: Temporal weighting profiles. A)-C), Top: schematic of the Gaussian distributions that generate the motion pulses. In the flat-stimulus A), Gaussians remain stationary over time. In the late-stimulus B) and early-stimulus C) conditions, the distribution means for signal trials are varied over time. Bottom: example sessions for each stimulus condition. Motion pulse values are drawn from their color-matched Gaussians on each pulse such that the mean of many trials (bold line) reflects the temporal structure of the mean of the Gaussians. Motion pulse values in individual trials (semitransparent traces) vary considerably, in accordance with the variance of color-matched Gaussians. Reproduced from Levi, Yates, et al., 2018 with permission of the publisher under a Creative Commons Attribution 4.0 International License.

The subjects performed the task in session batches and only one condition per session: first completing the flat-stimulus, then late-stimulus and lastly the early -stimulus sessions. Monkeys in particular, were trained extensively on the flat-stimulus before undergoing any temporal manipulation. In total monkeys performed 71 sessions (23 flat-stimulus, 22 late-stimulus and 26 early-stimulus, table 3.1 , in total 87,812 trials table 3.2). For simplification, throughout the rest of this work, we will refer to the flat-stimulus, late-stimulus, and early-stimulus as the flat condition, late condition, and early condition, and motion strength as coherence.

Condition	Nancy	Leo	Total
Flat	10	13	23
Late	11	11	22
Early	15	11	26

Tab. 3.1: Number of Sessions by Monkey

Condition	Value
50	14,654
20	13,776
0	22,692
-20	13,567
-50	14,569
Frozen	8,554
Total	87,812

Tab. 3.2: Total Number of trials by coherence

3.1.2 Electrophysiology

Neural data were recorded from the two monkeys with multielectrode arrays, full details can be found in Levi, Zhao, et al., 2023 and Yates, Park, et al., 2017. Middle Temporal Area (MT) recordings were taken while performing all the sessions, 594 units were identified: 161 in the flat stimulus; 223 in the late stimulus and 210 during the early-stimulus. Recordings from the Lateral Intraparietal Area (LIP) recordings were taken only during the flat condition (see Yates, Park, et al., 2017). In the MT recordings, we identified the neurons as stimulus selective neurons by using a independent two-sample t-test to compare the mean spike rates during stimulus presentation for strong coherence, positive and negative, on each recording session. Stimulus selective neurons were those that exhibited a significant difference in firing rates between strong left coherence and strong right coherence, in total we used 361 stimulus selective neurons (table 3.3, in total 87,812 trials 3.2).

Condition	Stimulus selective	Non stimulus selective	Total neurons
Flat	102	59	161
Late	116	107	223
Early	143	67	210
Total	361	233	594

Tab. 3.3: Stimulus selective neurons by condition

3.1.3 Data analysis

Time to fixate

To calculate the time to fixate, we used the behavioral data from both monkey and human subjects. The time to fixate was defined as the difference between the time interval between when the fixation target appeared and when fixation was acquired. We analyzed data for each subject (2 monkeys: N and L; 3 humans: A, H, and K) across two conditions: Early and Late, which correspond to contrasting demands (integrate Early vs. integrate Late). This analysis included all trials and sessions. Section 4.1.2 shows the mean session time to fixate by subject and stimulus condition along with its corresponding SEM. Statistical significance between conditions was tested using a two-sample Kolmogorov-Smirnov test (`kstest2` in matlab).

Psychophysical Kernel

We use the Psychophysical kernels (PK) to quantify the temporal weighting strategies or patterns. This method involves measuring the impact of each of the seven stimulus pulses on the net direction using a logistic regression model:

$$P(Y) = \frac{1}{1 + e^{-(\mathbf{x}\beta)}} \quad (3.1)$$

where $Y \in \{1, 0\}$ is the reported net direction or choice for every trial and X is a the matrix of the 7 pulses on each trials, augmented by a column of ones to capture the bias. The result is seven temporal weights in the β vector, one for each motion pulses.

To quantify the slope of the PK or the β vector, we use a linear regression analysis $y = \beta_0 + \beta_1 x + \epsilon$ where: y is the previously calculated temporal weights β , β_1 represents the slope coefficient. A positive slope, $\beta_1 > 0$, indicates that the kernel exhibits recency, signaling that the last stimulus epochs have a stronger influence on the integration. A zero slope, $\beta_1 = 0$, indicates uniform integration, signaling similar weight for every epoch. Lastly, a negative slope, $\beta_1 < 0$, indicates primacy, where earlier epochs have a larger weight in the

integration process. Following regression analysis, we normalized the slope values to facilitate comparison. Following Prat-Ortega et al., 2021, β_1 was normalized to quantify the shape of the PK by fitting linear function of time $PK = \beta_0 + k\beta_1 t$ where β_1 was the PK slope, and $k = \frac{1}{2\text{var}(t)}$ which is the factor that normalizes the PK slope to the interval between -1 (decaying PK, primacy) to +1 (increasing PK, recency).

Pulse triggered average and Peri-Stimulus Time Histogram

To measure the relationship between the time-varying pulse strength and the spike rate of neurons in MT of the monkeys we used the Pulse Triggered Average (PTA) methodology described in Yates, Park, et al., 2017. The pulse-triggered average (PTA) measures MT's activity that results from each additional pulse. For this, we created two binned vectors of 10 ms bins: one corresponding to the pulse stimulus and the other representing the spike rate, \vec{y} , (calculated as the binned spike count divided by the bin length). In the stimulus vector, \vec{x} , all trials were concatenated, so the total length T corresponds to the full duration of the session. The vector \vec{x} is zero everywhere except at the pulse onsets, where it has positive values when to motion of the pulses is moving in the preferred direction and negative in the anti-preferred direction. To capture the temporal dependencies of the stimulus vector, we use a lagged design matrix, D , in which $D(i, j)$ is the stimulus, \vec{x} at the i th bin at the j th lag. We compute a D for each of the seven pulses $X = [D1, D2, \dots, D7]$ and estimated the PTA using ordinary least squares regression.

$$PTA = (X^T X)^{-1} (X^T \vec{y}) \quad (3.2)$$

We calculated the PTA for each recording session provided each recording session had more than 4 neurons (see table 3.4)

Condition	Sessions (more than 4 neurons)	Total neuron count
Flat	16	144
Late	14	210
Early	18	190
Total	48	544

Tab. 3.4: Sessions by condition for PTA

To calculate the Peri-Stimulus Time Histogram (PSTH) for each stimulus condition, spike trains from all sessions with more than four neurons were first smoothed using a 100-ms boxcar filter and then binned into 10 ms intervals. For each trial and condition, the spike counts within each bin were summed, and normalized by dividing the spike count in each bin by the maximum spike count observed across all bins for that condition.

Choice probability

To calculate the choice probability (CP) in the neural activity of MT (Britten, Newsome, et al., 1996; Nienborg and Cumming, 2009; Wimmer et al., 2015), we analyzed spike recordings from single neurons, focusing on trials with zero coherence, frozen noise, and the combination of both. For each neuron, we calculated the moving average of spike data using a window of 350 ms, starting from stimulus onset to offset (1050 ms in total). This moving average time series (neurons x 701 ms) was then used to calculate the area under curve (AUC) of the receiver operating characteristic (ROC) of each neuron. The CP shown on our data analysis section (Section 4.2.4), is the average of the AUC for all stimulus selective neurons by condition (Section 3.1.2).

LIP correlation analysis

For measuring the correlation between the time to fixate (Section 3.1.3) and the neural activity of LIP, we analyzed data from 200 neurons¹. The recordings were retrieved from the data of Yates, Park, et al., 2017, as they use the same experimental set up as in Levi, Yates, et al., 2018 with the limitation that they only performed in the Flat condition. For each cell, we calculated the pearson correlation coefficient between the time to fixate (ms) and the firing rate during the time window between the monkey's eye enter the fixation point and 450 ms afterward. The correlation was taken across trials and statistical significance was determined with p-values < 0.05 using MATLAB function corr. The number of significant and non-significant correlation coefficient was normalized by the number of cells (n = 200) to provide the fraction of significant cells relative to entire population.

¹some cells were recoded in the same session

Psychometric curve

The psychometric function relates the subjects' performance to the net motion coherence of the stimulus. It was measured by calculating the proportion of rightward choices (P) as a function of the net motion (x) on each trial.

$$P(Y) = \gamma + (1 - 2\gamma) \frac{1}{1 + e^{-\beta(x-\alpha)}} \quad (3.3)$$

The net motion was categorized in 61 bins evenly spaced covering the full stimulus distribution, $-.60$ to $.60$, inclusive, for specific detail on the bins see the caption of the related figures.

3.2 Two area network model

For our model we develop a firing rate model that combines the two area hierarchical spiking neural network described in Wimmer et al., 2015 with the mean-field reduction described in Wong and X.-J. Wang, 2006. Our firing rate model is able to capture the same phenomena as just described in the spiking network but at a population level.

3.2.1 Reduced two-variable model

This model is a mean-field reduction of the spiking neural network proposed by XJ, 2002 (details of the mean-field reduction can be consulted in Wong and X.-J. Wang, 2006). The reduced two-variable model consists of two excitatory neural populations, selective for rightward and leftward motion directions, labeled by $i = 1, 2$. Each population receives external input and competes via mutual inhibition mediated by inhibitory interneurons. The dynamics can be completely described by the following equations:

$$\frac{\partial S_i}{\partial t} = -\frac{S_i}{\tau_S} + (1 - S_i)\gamma H_i \quad (3.4)$$

$$H_i = \frac{-ax_i - b}{1 - \exp[-d(ax_i - b)]} \quad (3.5)$$

$$\begin{aligned}x_1 &= J_{N,11}S_1 - J_{N,12}S_2 + I_0 + I_1 + I_{noise,1} \\x_2 &= J_{N,22}S_2 - J_{N,21}S_1 + I_0 + I_2 + I_{noise,2}\end{aligned}\tag{3.6}$$

$$\begin{aligned}r_1 &= H_1 \\r_2 &= H_2\end{aligned}\tag{3.7}$$

$$I_i = J_{A,ext}\mu_0 \left(1 \pm \frac{c'}{100\%}\right)\tag{3.8}$$

$$\tau_{AMPA} \frac{\partial I_{noise,i}(t)}{\partial t} = -I_{noise,i}(t) + \eta_t \sqrt{\tau_{AMPA} \sigma_{noise}^2}\tag{3.9}$$

Where equation 3.4 describes the synaptic gating variable of NMDA S_i for each population i , τ_S is the time constant for synaptic gating, and H_i nonlinear input-output function (rate) for population i . Equation 3.5 describes the function of the firing rate of each population where $a =$, $b =$, $d =$ and x_i is the total synaptic input for population i . This variable is defined by equation 3.6. The total synaptic input of each population x_i , depends on the synaptic gating variable of NMDA S_i and the coupling within and across populations, this is indicated as $J_{N,ij}$. The cross-coupling between populations has a negative sign in front of $J_{N,12}$, $J_{N,21}$, indicating that the overall effective connectivity between the two selective populations is inhibitory and the positive sign in the recurrent coupling within population, indicated by $J_{N,11}$, $J_{N,22}$ is excitatory. I_i represent the visual motion stimulus to the population i and depends on the 3.8. The rate of both populations is given by equation 3.7. I_{noise} is the noise term model in equation 3.9 as Ornstein–Uhlenbeck process, and I_0 is an unspecific background input common for both populations. This includes direct background input to a selective population but also indirect background inputs from these nonselective cells. Equation 3.8 represents the visual motion stimulus to the population i and depends on the coherence of the stimulus c , $J_{A,ext}$ average synaptic coupling and the absolute stimulus strength μ_0 . Lastly, equation 3.9 describes the noise process as an Ornstein–Uhlenbeck process.

3.2.2 Two area hierarchical spiking neural network

The model is a hierarchical network of spiking neurons proposed by Wimmer et al., 2015. The spiking network includes an integration circuit, such as LIP or FEF areas, and a sensory circuit, such MT area (see Fig. 3.3 A) linked by bottom-up feedforward connections and top-down feedback connections. The integration or decision circuit accumulates sensory evidence, from a classical perceptual discrimination task with temporally uniform sensory evidence of RDM task using fixed-duration, and produces a binary choice reflecting if the stimuli moving to the left or right. The dynamics of the sensory circuit (see Fig. 3.3 B, middle row) are determined by two neural populations $E1$ and $E2$ referred to in red and blue respectively. Each population is selective to one direction of the stimulus's motion and the stimulus strength is the input to the sensory circuit (see Fig. 3.3 B, bottom row). Meanwhile, the decision circuit operates through a winner-take-all competition between two populations (D1 vs D2, Fig. 3.3 B, top row) that encode the choice.




Fig. 3.3: Model architecture, single-trial response, population dynamics and modulation signals **A)** Network model composed of a sensory circuit (MT) with two opposite stimulus selective excitatory populations E1 and E2 that are coupled to two choice associated populations D1 and D2 in an integration circuit (for example, LIP, FEF). There are feedforward and top-down feedback connections (strength b_{FB}) between the two circuits as well as lateral excitatory and inhibitory (population I) recurrent connections within each circuit (connections are represented by lines with a width proportional to the synaptic efficacy and connection probability). The stimulus inputs currents to neurons in E1 and E2 (red and blue traces show two examples, s.d. s), each selective to a motion of an RDK. **B)** Response of the network to an example zero-coherence stimulus. Bottom traces show the population averaged stimulus currents into E1 and E2. Rastergrams show the spiking activity of neurons in E1 and E2 (middle) and in D1 and D2 (top), sorted by rate. Traces below the rastergrams show the corresponding instantaneous population rates. Top-down connections were set to zero ($b_{FB} = 0$). *Reproduced from Wimmer et al., 2015 with permission of the publisher under a Creative Commons Attribution 4.0 International License.*

3.2.3 Two area hierarchical neural network

For our full model, we use the mean-field reduction (Section 3.2.1) of the spiking neural network described in Section 3.2.2 to model the population dynamics of a two area network model that performs the changing stimulus statistics task described in Section 3.1. The dynamics of our model can be described by the following equations:

Sensory circuit

$$\frac{\partial S_i^E}{\partial t} = -\frac{S_i^E}{\tau_G} + (1 - S_i^E)\gamma H_i^E \quad (3.10)$$

$$H_i^E = \frac{-aX_i^E - b}{1 - \exp[-d(aX_i^E - b)]} \quad (3.11)$$

$$\begin{aligned} x_1^E &= J_{N,11}^E S_1^E - J_{N,12}^E S_2^E + J_{\text{TopDown},1} S_1^D + I_1 + I_{\text{noise}_{\text{stim},1}} + \\ &\quad + I_0^E + I_{\text{noise},1}^E \\ x_2^E &= J_{N,22}^E S_2^E - J_{N,21}^E S_1^E + J_{\text{TopDown},2} S_2^D + I_2 + I_{\text{noise}_{\text{stim},2}} + \\ &\quad + I_0^E + I_{\text{noise},2}^E \end{aligned} \quad (3.12)$$

$$\begin{aligned} I_0^E &= I_0^E - A_{\text{decay}} \\ A_{\text{decay}} &= c_A \cdot (1 - \exp(-\frac{I_0^E}{\tau_A})) \end{aligned} \quad (3.13)$$

In the Flat and Late condition:

$$I_i = J_{A,\text{ext}}\mu_0 (1 \pm c(t)) \quad (3.14)$$

In the Early condition:

$$\begin{aligned} I_i &= J_{A,\text{ext}}\mu_0 (1 \pm k_e c(t)) \\ k_e &= a_2 t_e^2 + b_2 t_e + c_2 \end{aligned} \quad (3.15)$$

$$\tau_{AMPA}^E \frac{\partial I_{noise,i}^E}{\partial t} = -I_{noise,i}^E + \eta_t \sqrt{\tau_{AMPA}^E \sigma_{noise}^E} \quad (3.16)$$

Decision circuit

$$\frac{\partial S_i^D}{\partial t} = -\frac{S_i^D}{\tau_S} + (1 - S_i^D) \gamma H_i^D \quad (3.17)$$

$$H_i^D = \frac{-aX_i^D - b}{1 - \exp[-d(aX_i^D - b)]} \quad (3.18)$$

$$\begin{aligned} x_1^D &= J_{N,11}^D S_1^D - J_{N,12}^D S_2^D + J_{BottomUp,1} S_1^E + I_0^D + I_{modulation} + I_{noise,1}^D \\ x_2^D &= J_{N,22}^D S_2^D - J_{N,21}^D S_1^D + J_{BottomUp,2} S_2^E + I_0^D + I_{modulation} + I_{noise,2}^D \end{aligned} \quad (3.19)$$

$$\tau_{AMPA}^D \frac{\partial I_{noise,i}^D}{\partial t} = -I_{noise,i}^D + \eta_t \sqrt{\tau_{AMPA}^D \sigma_{noise}^D} \quad (3.20)$$

Equations 3.10 and 3.17 describe the synaptic gating variables for population i in the sensory and decision circuit, respectively. Similarly, equations 3.18 and 3.18 define the firing rate function for these two circuits, and the noise in each circuit is modeled as an Ornstein–Uhlenbeck process, as described in equations 3.16, and 3.20.

In the sensory circuit, the total synaptic input x_i^E , is described in equation 3.12, and depends on several factors: the synaptic gating variable S_i^E , the cross coupling terms $J_{N,12}^E, J_{N,21}^E$, the recurrent coupling $J_{N,11}^E, J_{N,22}^E$, the circuit's noise $I_{noise,i}^E$, the stimulus noise $I_{noise_{stim},i}$, that is a Ornstein–Uhlenbeck process with a $\tau_{AMPA} = 0.013$ and a $\sigma_{noise} = \frac{0.023}{\sqrt{2}}$ and the unspecific background input of the circuit I_i^E . This background input I_0^E , defined in 3.13 incorporates the adaptation decay of the firing rate from the MT area from the experimental findings (section 4.2.2). The adaptation decay A_{decay} depends on the transient activity constant c_A that scales the adaptation, while the decay time constant τ_A controls how quickly the adaptation occurs. The exponential term captures the saturation of the firing rate, resembling the MT's transient dynamics. Lastly, depends on the visual motion stimulus I_i , described in equation 3.14 for the Flat and Late condition. It is influenced

by the coherence of the stimulus $c(t)$ which changes over time, the average synaptic coupling $J_{A,ext}$, and the absolute stimulus strength μ_0 . The stimulus duration is S_t while the trial duration is N_t . The stimulus coherence $c(t)$ was simulated with the same characteristics as the experimental data (3.1). For the Early condition, described in equation 3.15, the visual motion stimulus also depends on the coherence, average synaptic coupling, and absolute stimulus strength, but additionally includes a quadratic time series k_e optimized to reflect the unique shape of the PTA in the Early condition (section 4.2.3).

In the decision circuit, the total synaptic input x_i^D , is described in equation 3.12, and depends on several factors: the synaptic gating variable S_i^D , the cross coupling terms $J_{N,12}^D$, $J_{N,21}^D$, the recurrent coupling $J_{N,11}^D, J_{N,22}^D$, the circuit's noise $I_{noise,i}^D$, the unspecific background input of the circuit I_0^D , and the modulation signal $I_{modulation}$.

Modulation Signal

The modulation signal describes a time-dependent modulation process that affects only the decision circuit. In the initial stage of our modulation signal (see Section 5.3), the overall modulation term $I_{modulation}$ (Fig. 8.1) is equal to the unspecific background input I_0 from the Wong and X.-J. Wang, 2006 model. This allows the system to have attractor dynamics. The onset of the modulation signal is determined by the point at which greater variance in the stimulus statistics begins. In the Flat and Early conditions, this corresponds to the first pulse, whereas in the Late condition, it occurs at the third pulse. The transition between the lowest value of the modulation signal ($I_{u1} = 0$) and the largest value of the modulation signal ($I_u = I_0$) occurs over 150 ms (ac_t). The offset of the modulation signal is determined by a subsequent decrease in stimulus information (Fig. 8.1). After the stimulus offset, the signal maintains a symbolic value that corresponds to storing in memory the decision to facilitate attractor dynamics in the system.

$$I_{modulation} = \begin{cases} \frac{I_u - I_{u1} \cdot \frac{1}{s_c}}{1 + e^{-b(t - ac_t)}} + I_{u1} \cdot \frac{1}{s_c}, & t \in [\text{onset}, \text{onset} + ac_t] \\ I_u, & t \in [\text{onset} + ac_t, \text{offset}] \\ \frac{2}{3} \cdot I_u, & t \in [\text{offset} + S_t, N_t] \end{cases} \quad (3.21)$$

The steepness of the transition is controlled by the parameter b , the time at which the transition occurs is determined by a and c reflects the timing or values associated with the stimulus condition (f for Flat, l for Late, and e for Early).

Modulation Signal with preparatory activity

For this modulatory signal we include preparatory activity (Fig. 8.2, see Section 5.3). The onset of the preparatory activity starts 450 ms (pr_t) prior to the stimulus presentation and mainly affects I_{u1} as now has specific values for each condition (in the Flat and Early condition $I_{u1} = 0.02$; in the Late $I_{u1} = 0.005$).

$$I_{modulation} = \begin{cases} I_{u1}, & t \in [\text{onset} - pr_t, \text{onset}] \\ \frac{I_u - I_{u1} \cdot \frac{1}{s_c}}{1 + e^{-b(t - a_c)}} + I_{u1} \cdot \frac{1}{s_c}, & t \in [\text{onset}, \text{onset} + ac_t] \\ I_u, & t \in [\text{onset} + ac_t, \text{offset}] \\ \frac{2}{3} \cdot I_u, & t \in [\text{offset} + S_t, N_t] \end{cases} \quad (3.22)$$

Modulation Signal, smoothen

Lastly, we smoothen the modulation signal for visualization purposes (Fig. 8.3, see Section 5.3).

$$I_{modulation} = \begin{cases} \frac{I_{u1} \cdot \frac{1}{s_c}}{1 + e^{-b(t - a_{fix})}}, & \text{for } t \in [0, \text{peak_fix}] \\ \frac{I_u - I_{u1} \cdot \frac{1}{s_c}}{1 + e^{-b(t - a_c)}} + I_{u1} \cdot \frac{1}{s_c}, & \text{for } t \in [\text{peak_fix}, \text{peak_sc}] \\ \frac{I_u \cdot \frac{2}{3} - I_u}{1 + e^{-b(t - a_{2c})}} + I_u, & \text{for } t \in [\text{peak_sc}, N_t] \end{cases} \quad (3.23)$$

Simulation parameters

Parameter	Flat Condition	Late Condition	Early Condition
I_u	0.05	0.05	0.05
I_{u1}	0.01	0.01	0.01
s_f	1		
s_l		2	
s_e			0.5
b	0.1	0.1	0.1
a_{fix}	100	100	100
a_f	200	-	150
a_l		250	-
a_e	-		150
a_{2f}	350	-	250
a_{2l}		450	-
a_{2e}	-		250
peak_fix	800	800	800
peak_sf	1300	-	
peak_sl		1400	-
peak_se	-		1200

Tab. 3.5: Parameter values for the modulation signal in the Flat, Late, and Early conditions.

3.2.4 Simulation Based Inference

To quantitatively fit the behavioral data from the motion discrimination task in the flat, early and late conditions, we use statistical inference based on maximum likelihood estimation. Since the likelihood of our mechanistic neural network model cannot be computed directly, we use a simulation-based inference approach to estimate it from the model simulations (Ganmer et al., 2020; Boelts et al., 2022). Our parameter optimization procedure works as follows: (i) Start with initial values for the network parameters and simulate the network model with 15,000 trials for each stimulus condition. (ii) Fit a logistic regression model to the model responses for each stimulus condition. (iii) Compute the log-likelihood for the three logistic regression models (for the flat, early and late condition). (iv) The loss function is defined as the summed negative log-likelihood of the three stimulus conditions. This loss function is then minimized using Bayesian Adaptive Direct Search (BADs) using the Python library PyBADs (Acerbi and Ma, 2017; Singh and Acerbi, 2024). The network parameters that are optimized –in addition to the shape of the modulation signal – are the recurrent connectivity strengths in integration circuit, the noise level of sensory circuit and integration circuit, the mean background current in the integration circuit, and the bottom-up

Parameter	Value
N_t	3150 ms
S_t	1050 ms
dt	100 ms
τ_G	10 ms
τ_S	100 ms
γ	0.641
a	270 Hz/nA
b	108 Hz
d	0.154 s
$J_{N,12}^E, J_{N,21}^E$	0.1 nA/Hz
$J_{N,11}^E, J_{N,22}^E$	0.7 nA/Hz
$J_{N,12}^D, J_{N,21}^D$	0.05467 nA/Hz
$J_{N,11}^D, J_{N,22}^D$	0.28699 nA/Hz
$J_{\text{bottomUp},i}$	0.7 nA/Hz
$J_{\text{TopDown},i}$	0.01 nA/Hz
I_0^E	0.35154 nA
I_0^D	0.276 nA
$J_{A,\text{ext}}$	0.00052 nA/Hz
μ_0	60
c_A	0.03
τ_a	0.22
τ_{AMPA}^E	0.02
τ_{AMPA}^D	0.033
σ_{noise}^E	$\frac{0.01}{\sqrt{2}}$
σ_{noise}^D	$\frac{0.033}{\sqrt{2}}$
a_2	2
b_2	-0.8
c_2	1
t_e	N equally spaced values between [1.2, 2.15]

Tab. 3.6: Parameter values used for the simulation.

and top-down connectivity strength between the sensory and the integration circuit.

3.3 Recurrent Neural Networks

The Recurrent Neural Networks (RNN) described in this section shared the same network architecture as in Yang, Joglekar, et al., 2019. With $N_{rec} = 256$ recurrent units, a batch size of 64, and a sequence length of 20, the network processes 64 sequences simultaneously, each sequence being 20 time steps long. The continuous-time dynamics of the RNN is defined by:

$$\tau \frac{dr}{dt} = -r(t) + f(\mathbf{W}_r r(t) + \mathbf{W}_x x(t) + \mathbf{b}_r). \quad (3.24)$$

Equation 3.24 models how the neuronal activity in an RNN evolves over time. the variable $r(t)$, referred to as the hidden state, represents the time-dependent activity of the neurons within the recurrent network. This activity is influenced by four factors: (1) the decay of the neuronal activity over time and the tendency of the neuronal activity to diminish in the absence of input, represented by $-r(t)$; (2) the recurrent interactions within the network, which captures the network's recurrent connectivity matrix (W_r) effect on the hidden state ($r(t)$) represented by $W_r r(t)$; (3) an external input term ($x(t)$), which integrates sensory stimuli through the connectivity matrix (W_x) represented by $W_x x(t)$; and a bias term (b_r). These components are transformed through a Rectified Linear Unit (ReLU) function ($f(\cdot)$), a non linear activation function where $f(x) = \max(x, 0)$, this function is asymmetric and non-saturating at high values and ensures the neural responses are non-negative. The network dynamics are modulated by the neuronal time constant $\tau = 100 \text{ ms}$, this value mimics the slower synaptic dynamics on the basis of NMDA receptors (Yang, Joglekar, et al., 2019). This system 3.24 can be discretize using the Euler method with a time step of Δt :

$$\begin{aligned} \mathbf{r}(t + \Delta t) &= \mathbf{r}(t) + \Delta \mathbf{r} \\ &= \mathbf{r}(t) + \frac{\Delta t}{\tau} [-\mathbf{r}(t) + f(\mathbf{W}_r \mathbf{r}(t) + \mathbf{W}_x \mathbf{x}(t) + \mathbf{b}_r)] \\ &= \left(1 - \frac{\Delta t}{\tau}\right) \mathbf{r}(t) + \frac{\Delta t}{\tau} f(\mathbf{W}_r \mathbf{r}(t) + \mathbf{W}_x \mathbf{x}(t) + \mathbf{b}_r) \end{aligned} \quad (3.25)$$

Lastly, the final output of the network, z_r , is nonlinear a read out from the hidden state represented by $z_r = g(W_{out}r(t))$. Here W_{out} is the output weight matrix, connecting the hidden activity to the output layer and $g(x) = \frac{1}{1+e^{-x}}$ is the logistic function that ensures that the output values are bounded between 0 and 1.

The network, the training process and testing are implemented using TensorFlow. The input to the network was design using the NeuroGym package where we design a virtual environment that resemble the experimental task from Section 3.1. The loss function of the network was a cross-entropy loss, which calculates the difference between the predicted logits (raw scores the model gives for each class before turning them into probabilities) and the target labels (the correct answers for each input). The model assigns higher probabilities to the correct classes while penalizing incorrect predictions. The training was performed with Adam optimizer, a variant of stochastic gradient descent that adjusts its learning rates for each parameter during training. The learning rate was 0.001 and the training length was 6000 epochs, unless specified otherwise. The recurrent connection matrix is initialized at random. The testing was in the same environment for 15,000 trials.

Sequentially trained RNNs

Network	A	B	C	D
Input weight	No	No	No	6000
Connectivity matrix	No	Yes	Yes	6000
Output weight	No	No	Yes	6000
Training length	600	6000	600	600

Tab. 3.7: Training Protocol of Fig. 6.12

Principal component analysis

We analyzed the activity of the trained and tested networks using Principal Component Analysis (PCA), this method allowed us to reduce the dimensionality of the activity data and plot it in the principal component space. Initially, we reshaped the activity matrix from *trials x time x neurons* to *time x neurons* for trials of an specific coherence (high coherence and low coherence) and calculated the PCA with the first three components for the

independently, and sequentially trained networks and five components for the contextual training networks.

For the independently trained network, we analyze dynamics using two principal components (PCs) (Fig. 6.6 and Fig. 6.16) and present the mean distribution of PC activity across all high-right coherence trials, broken down by each trial epoch (Fig. 6.7). To further explore the dynamics across 10 RNNs, we examine the time course of PCs eigenvalues, using the trial-averaged eigenvalues at each time point (Fig. 6.8 A). For two network examples, we visualize the dynamics of individual PCs across conditions (Fig. 6.8 B and C).

For the context-trained RNN, we investigate how fixation dynamics adapt in novel environments by analyzing the fixation PCa distribution (Fig. 6.19) on each environment. Additionally, the heatmap (Fig. 6.20) provides a visual representation of the PCA-transformed state space for the context trained RNN under novel environments. Each row in the heatmap corresponds to a different novel environment (ranging from 100% Flat to 100% Late), while columns represent trial epochs. The heatmap's color intensity reflects the magnitude of the data's projection along the PC space, summarizing the changes in variance and structure within the reduced-dimensional space as the network encounters novel environments.

Time-lagged autocorrelation

For the time-lagged autocorrelation analysis, we use the Trial-averaged temporal activity of the neurons, sorted by the weight of the first principal component (calculated in Fig. 6.8 A). The neuronal activity is derived from our sample RNN (Fig. 6.2). Neurons selective for the left stimulus are placed at the bottom, while those selective for the right stimulus are positioned at the top. Next, we compute the difference in trial-averaged activity for each neuron, comparing the left versus right trials during high-coherence stimuli. We then calculate the time-lagged autocorrelation between the population states of each network, following the methodology in Murray et al., 2017. This method quantifies the similarity between activity patterns at two different time points during the trial. It does so by measuring the within-stimulus-condition Pearson correlation coefficient between the population activity patterns at time t_1 and time t_2 . The correlation is computed separately for each stimulus condition.

Task Variance

To determine whether the units in the context signal RNN were selectively activated by specific stimulus conditions (Flat, Late, Early), or whether the same neurons were shared across conditions, we analyzed the task variance. For this, we computed the variance of the neural activity across stimulus conditions (trials) at each time point for each of the $N_{rec} = 256$ units, focusing on activity during the stimulus presentation, we then averaged the variance across all time points to obtain the task variance for each unit. Next, we normalized the variance for each neuron by dividing it by the maximum variance observed for that condition. Finally, we sorted the neurons by similarity as well as the conditions by similarity, clustering together neurons and conditions with similar activity patterns.

3.4 Software

All analyses were performed using Python. In particular, it's main libraries (NumPy, Pandas, Matplotlib, SciPy, Scikit-learn, TensorFlow, Keras, PyTorch). Additionally: (1) for the two area hierarchical model we used costume made packages for parallel processing prepared by Jose Mari Esnaola, Post-doc at Wimmer lab at the time (2) For the RNNs we used Neurogym package (<https://github.com/neurogym/neurogym>) (3) For the pulse triggered average we used MATLAB and modified code from the Yates lab (<https://github.com/jcbyts/mtlipglm>).

Experimental insights into flexible temporal weighting

In this section, we use multiple analysis techniques to investigate how temporal weighting strategies manifest in both behavior and neural recordings when the stimulus statistic changes over time. Our analysis centers around new aspects of previously published data recorded in the laboratory of Alex Huk (Levi, Yates, et al., 2018 and Levi, Zhao, et al., 2023) and aims to characterize the relation between the stimulus statistics and the modulation signal. We begin by examining the behavioral correlates, focusing first on the session-to-session changes of the temporal weighting patterns to better understand how the evidence integration strategies are acquired over time. Next, we explore the monkeys' behavior prior to the stimulus presentation, specifically the time taken to fixate (from the onset of the fixation cross until the monkey initiates a trial by saccading to it). This time-to-fixate serves as a measure of engagement (inspired in Mochol et al., 2021). Following the behavioral analysis, we examine the neural correlates of the time to fixate in the decision areas, we analyze neural activity from LIP while a monkey was performing the Flat stimulus condition (Yates, Park, et al., 2017) to use it as an indicator of task engagement that can actually be linked to the decision making process. Next, we turn to analyze the neural activity of MT neurons, focusing on the population firing rates, the choice probability and MT's activity resulting from each stimulus pulse. We found, despite the overall population firing rate remaining stable among stimulus conditions, the effect of each additional pulse of information on MT activity shows a transient pattern. Together, these findings provide important insights and constraints for the development of the neural network model

4.1 Behavioral signatures of flexible temporal weighting

4.1.1 Session-to-session temporal weighting update

We first analyzed how the temporal weighting strategies of the monkey update from session to session. In our task, the monkey subjects performed the task in batches, with each recording session containing only one stimulus condition and only switching condition after at least 10 recording sessions—first Flat, then Late, and finally Early. This design allowed us to follow changes in temporal weighing strategies over time. For this, we calculated the PKs (see Section 3.1.3) and their slopes for every recording session (Fig. 4.1). We found that the PK slope, for one monkey, in the Flat condition was negative but close to zero, indicating primacy or uniform weighting in the PKs. Interestingly, when the subject switched to the Late condition, the slope remained negative for a few sessions, however, as the monkey had more exposure to the condition, the slope changed to positive values, reflecting a shift toward recency weighting. Similarly, when subsequently switching to the Early condition, the slope initially maintained positive values and in later sessions, progressively, shifted to negative values, even more pronounced than in the Flat condition. Comparing the PK shapes from the first to the last session within each condition, we observed a variety of weighting strategies. In the Flat condition, the PK shifted from a uniform weighting strategy to slight primacy (moderate Pearson correlation of 0.64). In the Late condition, the PK shifted from slight primacy to slight recency (Pearson correlation of 0.15). In the Early condition, the PK shifted from slight recency to primacy (Pearson correlation of -0.25). Additionally, we compared the PKs around the condition switch — last session PK to the first PK — and found a high Pearson correlation between the PKs: 0.97 from Flat to Late and 0.8 from Late to Early. This supports our finding that the weighting updates are not immediate but instead update gradually over multiple sessions. Our results indicate that the flexibility of the temporal weighting strategies is acquired and perfected through exposure to the stimulus conditions. Additionally, bias toward the prior condition influences the strategy update temporarily until the task demands are met.




Fig. 4.1: Changes in PKs across recording session. A) PK slope for each recording day. Not all sessions were recorded on consecutive days, except during transitions to a new stimulus condition. Flat condition in blue, Late condition in yellow and Early condition in red. B) PK corresponding to the first recording session of the condition. C) PK corresponding to the last recording session of the condition.

4.1.2 Time to fixate

We first examined the pre-stimulus behavior of the monkeys and humans as a proxy of the subject state of vigilance. We refer to this measure as the "time to fixate", defined as the interval between the onset of the fixation point and when the subjects acquired fixation. Shorter fixation times reflect high engagement, while longer time to fixate suggest less engaged. By analyzing the time to fixate across the multiple stimulus condition, we aimed to identify how the task engagement changes in relation to the stimulus conditions. At an individual level, humans showed short fixation times, with over 70% of trials acquiring fixation under 1 second, indicating high engagement (Fig 4.2 A). When examining the mean time to fixate by subject and condition, we found that the subject's engagement was higher by having smaller mean time to fixate in the Early condition than in the Late condition for both species (Fig 4.2 B). When examining the mean time to fixate by condition and species, we

found that subjects' engagement was higher, as indicated by a shorter mean time to fixate, during the Early condition compared to the Late condition, this difference was significant, for both species. However, monkeys generally showed a shorter mean time to fixate and smaller error bars than humans, likely due to their extensive training and the higher number of trials and sessions completed. Individual differences in the magnitude of the mean time to fixate did not affect the overall trend. Our results suggest the subjects' state of vigilance may be higher in the Early condition, as they need to engaged early on during the stimulus presentation to effectively perform the task.




Fig. 4.2: Engagement in human subjects and mean time to fixate by condition in all subjects. A) Histogram of times to fixate for one individual human subject, $n = 300$ trials. B) Mean time to fixate for Early (red) and Late (Yellow) conditions. Significance at p - value < 0.05 , using Kolmogorov - Smirnov test

4.2 Neural activity underlying flexible temporal weighting

4.2.1 LIP activity as a measure for engagement

We continued with the engagement measures by analyzing neural data from area LIP in macaques performing the task in the Flat condition. The recordings were retrieved from Yates, Park, et al., 2017, as the authors use the same experimental set up as Levi, Yates, et al., 2018 with the limitation that

they only presented the Flat stimulus and recorded LIP's activity during the presentation of the stimulus. We tested whether the task engagement of the macaques is reflected in neural activity by calculating the Pearson correlation of pre-stimulus neural activity for each of the recorded 200 LIP neurons and time to fixate (in ms). Specifically, we used the firing rate during the 450 ms period after the monkey's eye entered the fixation point and prior to the presentation of the stimuli (see Section 3.1.3 for methods). A positive correlation suggested that as fixation time increases, the firing rate of the LIP neurons also increases whereas a negative correlation suggests that as fixation time increases, the firing rate of the LIP neurons decreases.

We found that the pre-stimulus activity in LIP was correlated with task engagement, with 33% of recorded neurons (Fig. 4.3 black bars) showing a significant correlation. Moreover, we observed both positive and negative correlations suggesting pre-stimulus modulation of the firing rates in LIP, a decision-related area.




Fig. 4.3: Correlation between time to fixate and LIP firing rate. Pearson correlation coefficient between time to fixate and firing rate taken within a 450 ms window after macaque acquire eye fixation. The correlation is taken across trials for $n = 200$ LIP neurons. Correlation values range from -0.4 to 0.4, indicating that most cells show little to no strong relationship between fixation time and firing rate. The central peak around 0 suggests that the majority of cells exhibit minimal correlation between the two variables. 33% of cells showed significant correlation coefficients (black bars).

4.2.2 Firing rate of neurons in MT area

Next, we explore the impact of the changing stimulus statistics of the task on the neural activity of MT. For this, we use the neural recordings of MT (Section 3.1) obtained through multi-array electrodes that captured the spike activity of individual neurons during each trial. MT neurons show a burst of activity rapidly after motion onset (Fig. 4.4 A). The recordings captured activity from both stimulus-selective and non-selective neurons to the specific coherence of the stimulus ($\pm 50\%$, $\pm 20\%$, $\pm 0\%$). We classified each neuron based on their response to stimulus coherence (Section 3.1.2). Stimulus selective neurons exhibited a significant difference in firing rates between strong left and strong right coherence (Fig. 4.4 AB), first two plots, one left selective and other right selective) while non selective neurons did not show a significant difference in the firing rate (Fig. 4.4 B, last plot). Lastly, the average firing rate of MT population showed a transient and sustained response following stimulus presentation, that did not changed between conditions, indicating that the overall MT's activity was not, on average, influenced by the stimulus statistics (Fig. 4.4 C).




Fig. 4.4: Average firing rate by condition. **A)** Raster plot of a sample MT neuron, early condition for strong left coherence. **B)** Trial-averaged firing rate of three sample neurons, 2 stimulus selective and one non-selective. Indigo line corresponds to positive coherence and yellow to negative coherence. **C)** Normalized population firing rate by condition.

4.2.3 Pulse triggered average in MT area

To characterize the temporal representation of the stimulus in MT, we used the Pulse-Triggered Average (PTA) (see section 3.1.3 for methods), which shows the changes in MT's activity resulting from each stimulus pulse. At a single-session level, we observed differences between the stimulus conditions (Fig. 4.5, A). In the Flat condition, MT response was transient, with the first pulse driving the activity and each consecutive pulse contributing less to the overall activity in a decreasing manner. In the late condition, we observed similar dynamics to those in the Flat condition, although with a slower decrease. The change in the stimulus statistics did not affected MT's response, as there was no additional increase in in the activity in the 4-5th pulse. Lastly, in Early condition, we observed a u-shaped dynamic: the first pulse elicited a strong response, which then decreased until the fifth pulse and increases again after. These condition-dependent PTA patterns were also observed at a population level, with very similar profiles to the individual sessions (Fig. 4.5, B).




Fig. 4.5: Pulse-triggered average of MT activity by stimulus condition. A) Pulse-triggered average of three sample sessions with more than four neurons. B) Session-averaged pulse-triggered average across all sessions with more than four neurons.

4.2.4 Choice probability in MT area

Next, we examined the CP across stimulus-selective neurons for all stimulus conditions for trials with average coherence zero (Fig. 4.6, solid line), and

frozen noise (Fig. 4.6, dotted line) (see Section 3.1.3 for methods). This analysis aims to bridge the experimental observations with the predictions of our proposed model (see Section 5). Our model predicts the characteristic CP time course for the different stimulus conditions, as shown in Fig. 5.10 and explained in detail in Section 5.5. Here, we test these predictions. In the Flat condition, the CP was sustained throughout the trial, with higher values at the beginning of the trial. In the Late condition, the CP time course is remarkably different in coherence-zero trials, where it is practically absent during the first half of the trial before it starts to increase. In the Early condition, we observed sustained CP throughout the trial in coherence-zero trials. Note that in frozen noise trials, we did not observe the same pattern, but CP was overall much more variable, likely due to an insufficient number of trials in this condition (see Section 3.1).




Fig. 4.6: Choice probability in MT area by stimulus condition. Choice probability time course with a 350 ms moving window that starts at motion onset for trials with average coherence 0 (solid line) and frozen noise trials (dotted line). Flat condition in blue, Late condition in Yellow, Early condition in Red.

To test whether the CP time course is different across Flat, Late and Early conditions, we computed, for each neuron and condition, the CP for the first (0 - 500 ms) and second half of the trial (500 - 1000 ms). We then fitted a repeated measures model and analyzed the variance using repeated measures ANOVA (Matlab functions `fitrm` and `ranova`). For trials with zero coherence (4.7), post-hoc analysis showed a difference between conditions for CP in the first half of the trial (one-way ANOVA, $F(2, 358) = 10.13$, $p = 0.0000527$), with significantly lower CP values in the Late condition compared to Flat ($p = 0.0000251$) and Early ($p = 0.0133$). In contrast, the CP in the last half of

the trial did not differ between conditions (one-way ANOVA, $F(2, 358) = 0.36$, $p = 0.70$).




Fig. 4.7: Average choice probability in MT area, coherence 0 trials. Average choice probability (CP) time course for zero-coherence trials, with a 500 ms window dividing the CP time course into two segments. Error bars represent the standard error of the mean, computed using bootstrap.

Lastly, we repeated this analysis for the CPs obtained from frozen noise trials (4.8). In this case, the repeated measures ANOVA did not reveal a significant interaction between condition and time ($F(2, 358) = 2.42$, $p = 0.090$), nor a main effect of time ($F(1, 358) = 0.31$, $p = 0.58$). Thus, we did not find systematic differences of CP across conditions or time intervals. A likely explanation for this is the overall higher variability in the CP values for frozen noise stimuli (4.8) and fewer trials (see Section 3.1.2).




Fig. 4.8: Average choice probability in MT area by condition, frozen noise trials. Average choice probability (CP) time course for frozen noise trials, with a 500 ms window dividing the CP time course into two segments. Error bars represent the standard error of the mean, computed using bootstrap.

Conclusion

Motivated by the experimental work on Levi, Yates, et al., 2018, Levi, Zhao, et al., 2023 and Yates, Park, et al., 2017, we further investigated how temporal weighting strategies manifest in both behavior and neural recordings when the stimulus statistic change over time. A closer examination of the behavioral correlates revealed that temporal weighing strategies gradually adapt over time, updating slowly to the new stimulus condition and becoming more consistent as exposure to the task increases (Fig 4.1).

The level of engagement, measured by time to fixation, varies across stimulus conditions with conditions with greater information content at the end of the trial having longer fixation times (Fig 4.2). This suggests that monkeys may be better prepared to integrate information starting from the onset of the stimulus in the early condition. The analysis of the neural correlates showed also evidence of pre-stimulus task engagement in a decision-related area, the lateral intraparietal (LIP) area, where neurons exhibited lower firing rates when the time to fixate was shorter (Fig. 4.3).

Moreover, we found that the firing rate of MT area is not directly influenced by the stimulus statistics (Fig. 4.4). However, the marginal effect of each additional pulse of information on MT activity exhibits transient dynamics that do not align with the time course of the stimulus statistics across conditions (Fig. 4.5). This motivates us to test in detail the potential contribution of MT to the integration dynamics in the three conditions in a neural network model in the next chapter.

Lastly, we performed an analysis of the population choice probability (CP) and found that the time course depends on both the stimulus statistics and the type of trial. Specifically, the Flat and Late conditions differ from each other, as do the Late and Early conditions, but only during the first 500 ms of the CP time course and only for trials with zero coherence (Fig. 4.6). This findings are consistent with the model predictions (Section 5.5).

Neural network dynamics underlying the adjustment of temporal evidence weighting in perceptual decisions: a two-area neural network model

Experimental work by Levi, Yates, et al., 2018, has demonstrated that primates can flexibly adapt their temporal weighting strategy based on the stimulus statistics (Section 1.2). For example, sensory stimuli that provide more information at the beginning of a trial result in an early weighting strategy or primacy, while stronger sensory evidence in later stimulus epochs causes a late weighting strategy or recency. However, the neural mechanisms that allow for this flexible adjustment to stimulus statistics remain unknown. In this chapter, we develop a neural network model to shed light on the potential mechanisms underlying flexible temporal weighting at a neuronal level.

Our model takes into account the two stages of decision-making: a first stage in which visual neurons represent and transform incoming information and a second stage in which a decision circuit accumulates evidence over time (Gold and Shadlen, 2007). In particular, we use a two-area firing rate model, consisting of a sensory circuit and a decision circuit with both bottom-up and top-down connectivity, following Wimmer et al., 2015. The sensory circuit is modeled after the middle temporal area (area MT), which receives visual input and is directly constrained by the empirically observed time course of MT activity (section 4.2.1). This circuit is reciprocally connected to the decision circuit, modeled after areas such as the lateral intraparietal cortex (area LIP) and the frontal eye field (area FEF), which integrates information as described by Wong and X.-J. Wang, 2006.

Initially, we considered three mechanisms as possible sources for temporal weighting flexibility: alteration to the sensory circuit, gain modulation within the decision circuit, and, the variations in the strength of connectivity between the circuits. However, experimental evidence from the MT area, along with results from our model simulations, indicated that the sensory circuit was not responsible for the changes in the temporal weighting patterns. Therefore, we then focused on two key aspects: a modulation signal that adjusts the attractor dynamics of the decision circuit, effectively altering the PK shape, and changes in the connectivity between the two circuits, specifically the strength of the top-down connections. We found that flexibility in temporal weighting could be explained by gain modulation of the decision circuit through a time-varying modulation signal, similar to an urgency signal, which effectively altered the PKs by driving the network into a competition regime that accelerated or delayed the decision integration. Our model successfully reproduced the experimentally observed primacy for the Early and Flat stimulus conditions, as well as late weighting for the Late stimuli condition. By controlling the timing and duration of the evidence integration initiation, and by adjusting the onset and duration of the modulation signal, we observed a wide range of PK changes, including primacy, recency, and other combinations. In contrast, changes in the top-down connectivity between the circuits had minimal effects on the PK but significantly influenced the choice probability (CP, Section 1.1.2). Increasing the strength of top-down connectivity mainly affected the late component of the CP, which reflects the integration of the sensory information, regardless of the stimulus statistics. These findings suggest that gain modulation in the decision circuit is the main mechanism underlying the flexibility in temporal weighting strategies. Finally, we tested whether our two-area model, incorporating the gain modulation signal to the decision circuit and the constraints on the sensory circuit, could quantitatively account for the experimental data. Using simulation-based inference (SBI), we found that the model provided an excellent fit to the observed data.

5.1 Two-area neural network model

To study different mechanisms underlying the flexibility of temporal weighting strategies, we developed a two-area firing rate model composed of a sensory circuit (modeled after MT area) and a decision circuit (modeled after LIP or FEF areas) with reciprocal connectivity, following Wimmer et al., 2015.

This model is characterized by a two-area hierarchical network of spiking neurons. The sensory circuit encodes sensory information from a classical perceptual discrimination task with temporally uniform sensory evidence ¹. The decision circuit then accumulates this sensory evidence and, through winner-take-all competition —where competing neural populations inhibit each other and only the most active prevails — produces a binary choice reflecting whether the stimulus is moving to the left or right. This model captures the time course of stimulus encoding and the integration of the choice. Additionally, it can explain the time course of choice probability (CP), which reflects the likelihood that an ideal observer can predict the monkey's behavior based on the activity of individual neurons (Section 1.3.2 and Section 3.2.2 for details of the model). However, spiking neural networks are challenging to interpret due to their high dimensionality, sensitivity to changes in parameters, and computational demands (Dayan and Laurence F Abbott, 2005, chapter 7).

To address these challenges, we can reduce the model with a mean-field approach (Wong and X.-J. Wang, 2006 and Section 1.3) to a simpler and tractable firing rate model, with fewer parameters (Fig. 5.1 A)), which allows for exploring the essential dynamics. We then simulate the model with the input of the 3 stimuli with different statistics: Flat, Late, and Early, with the same coherence distribution as in the experimental setup. We found that the variations in the stimulus statistics do not alter the dynamics of the decision circuit (Fig. 5.1 C)), resulting in primacy weighting irrespective of the stimulus condition (Fig. 5.1 D). This inability to capture the observed flexibility in the temporal weighting patterns is expected because the model has a fixed decision dynamics (Fig. 5.1 A, bottom row). This is due to the winner-take-all dynamics of the decision circuit, which resulted in early integration despite the changing information of the stimuli.

In summary, the current model does not capture or explain the underlying mechanisms of flexibility in decision-making. To address this, we explore changing three key components of the model in a task-dependent manner: the sensory circuit, the decision circuit, and their connectivity. If flexibility originated in the sensory circuit, we should observe changes in how sensory information is encoded at a population and single-trial level. If it came from the decision circuit, we should observe shifts in decision dynamics and

¹two-alternative forced-choice motion discrimination task using fixed-duration random dot kinematograms




Fig. 5.1: Reduced model architecture, dynamics, and temporal weighting patterns. **A)** Firing rate model architecture, composed of a sensory circuit (MT) with two opposite stimulus selective populations E1 and E2 that are coupled to two choice- associated populations D1 and D2 in a decision circuit (for example, LIP, FEF). Both circuits are connected via feedforward (FF) and feedback (FB). The stimulus input is model after Levi, Yates, et al., 2018, with 3 stimulus condition consisting of 3 different stimulus statistics (Flat, Late and Early) and 5 coherence strength (section 3.1). **B)** Single-trial dynamics of the Late condition for a weak left coherence trial. Bottom row show the stimulus input of each population (E1 in yellow, and E2 in indigo) , middle row the sensory circuit dynamics and top row the dynamics of the decision circuit. **C)** Sample population dynamics for the same condition, both circuits show attractor dynamics from the start of the stimulus presentation. **D)** Psychophysical kernels for the Flat (blue) condition, Late (yellow), and Early (red) all showing primacy.

competition patterns. If top-down connectivity were the source, we should expect alterations in the time course of reverberation between circuits. In the following sections, we introduce an adaptation of the neural attractor network model designed to enable flexible decision-making, providing a robust and mechanistic explanation for the observed temporal weighting flexibility in evidence integration.

5.2 Impact of the sensory circuit on the decision dynamics

We start by investigating the role of the sensory circuit in temporal weighting flexibility. In the experiment, subjects' temporal weighting strategies varied from condition to condition; however, neural recordings from the

MT area—which our sensory circuit is modeled after—have shown that the population activity does not change accordingly. In all stimuli conditions, the average firing rate of MT area exhibits a fast increase, approximately 50 ms after stimulus onset, followed by a slow decrease into a sustained activity that lasted until the end of the stimuli (Section 4.2.2, Fig. 4.4). This transient neural response of MT is well documented in the literature and has been observed in recordings of individual neurons recording (Britten, Newsome, et al., 1996) and population responses (Yates, Park, et al., 2017) as well as, changes in the stimuli motion or speed (Ernst et al., 2021). However, when considering the impact of each stimulus pulse on MT activity (PTA), we observed distinct responses depending on the condition (Section 4.2.3 and Levi, Zhao, et al., 2023). In particular, the Flat and Late conditions showed a large response in the first pulse, followed by a steady decay in subsequent pulses. This decay did not correspond to the stimulus statistics for the Late condition, where the most information of the stimulus was placed after the fifth pulse. Strikingly, the Early condition showed two increases: one corresponding to the first pulse, as anticipated, and an unexpected one in the fifth pulse, which continued to rise until the end of the stimulus presentation. In sum, several aspects of the empirically observed MT activity appear incompatible with the observed changes in the corresponding temporal weighting patterns, presenting a significant challenge for understanding MT’s impact on the decision dynamics.

These results also present a challenge for our model: our model must accurately capture the transient dynamics of MT and the unexpected PTA activity, and, at the same time, it needs to explain the PKs in light of these dynamics. Rather than building a detailed model of the stimulus-driven responses in area MT, we included the time course of MT’s neural responses as a constraint in our model. We adjusted the sensory circuit’s average firing rate to align with MT’s neural response by modifying the unspecific background input, common to both populations, which can be expressed as:

$$I_0^E = I_0^E - A_{decay} \quad (5.1)$$

$$A_{decay} = c_A \cdot (1 - \exp(-\frac{I_0^E}{\tau_A})) \quad (5.2)$$

where I_o^E represents the unspecific background input of the sensory circuit, c_A is the transient activity constant, and τ_A is the decay time constant. By incorporating this decay, the average population firing rate of the sensory circuit closely resembles those observed in the MT neurons (Fig. 5.2 B).




Fig. 5.2: Sensory Circuit Firing Rate and PTA. A) Model schematic of the affected sensory circuit. B) The Average firing rate for each stimuli condition shows transient dynamics, characteristic of MT area activity. C) Psychophysical kernels for each condition, showing primacy despite the stimulus statistics.

Moreover, PTA patterns of the sensory circuit should resemble the patterns observed in the experimental data. To achieve this, we adjusted the impact of the Early condition, since it was the only condition with a unique pattern, in our model's stimuli encoding using a non-linear scaling factor k_e for each time point of the stimulus which the following equation can express:

$$k_e = a \cdot eX^2 + b \cdot eX + c \quad (5.3)$$

$$I_i = j_{A,ext} \cdot \mu_0 + (1 \pm k_e \frac{c'}{100\%}) \quad (5.4)$$

where $\pm j_{A,ext} \cdot \mu_0 \cdot k_e \frac{c'}{100\%}$ represents the modulation of the incoming stimulus and k_e the time series of that modulation. This modulation of the Early stimulus encoding allowed our circuit to resemble the activity observed in the experimental data (Fig. 5.2 C) without affecting the PK shape (Fig. 5.2 D).

In summary, by incorporating the experimental findings as constraints in our sensory circuit model, we have a more realistic sensory circuit that is able to simulate the behavior of real MT neurons. Our simulations show that

the changes to the sensory circuit do not have strong effect on the temporal weighting patterns, as the PKs remain in primacy. The lack of effect of these constraints on temporal weighting strongly suggests that the mechanisms underlying flexible weighting do not originate in the sensory circuit. In the following sections we will focus on the gain modulation of the decision circuit and the connectivity between circuits as potential mechanisms for the flexibility in the temporal weighting.

5.3 Modulatory signals can influence the decision dynamics

Unlike the sensory circuit, which is responsible for encoding sensory information, the decision circuit plays a crucial role in the evidence integration process. We propose a time-varying modulation signal, I_m , as a potential mechanism for flexible temporal weighting. This signal alters the dynamical regime of the decision circuit similarly to how a change in neural gain does (Niyogi and Wong-Lin, 2013). In particular, the signal affects the unspecific background input of the model and represents an external input which could originate from other brain region like the superior colliculus (Inagaki et al., 2022). A neural gain change could be realized through neuromodulators (Luo and Pinto, 2022; Aston-Jones and J. D. Cohen, 2005), changes in E/I balance (Niyogi and Wong-Lin, 2013), balanced synaptic inputs (Lo et al., 2015), bottom-up / top-down dynamics (Haefner et al., 2016; Wimmer et al., 2015) or common excitatory back-ground inputs (Roxin and Ledberg, 2008). This modulation signal influences the integration of the stimuli by accelerating or delaying the onset of the decision process, similar to an urgency signal. Increasing I_m shifts the network from a low-activity state to an increasingly competitive regime with bistable dynamics (i.e, gradually reshapes the landscape from shallow point attractors to two attractors, Fig. 1.14). We will show that this signal allows to flexibly adjust the evidence integration by controlling the onset and the dynamics of the decision process.

To model the modulation signal, we reduced the unspecific background input I_0^D of the decision circuit (Section 3.2.3 for model details) so that $I_0^D + I_m = I_0$. This adjustment places the circuit in a competitive regime, consistent with Wong and X.-J. Wang, 2006. First, we set I_m to reflect the stimulus statistics of each stimuli condition, meaning that the signal start

aligned with the increase in stimulus information and maintain activity after the offset of the stimulus (see section 3.2.3 for mathematical details). The modulation signal is design to ramp up immediately after its onset to account for temporal scaling (i.e control the dynamics of ongoing patterns of neural activity J. Wang et al., 2018) and self-initiated processes, in this case how fast it reaches the attractor (Inagaki et al., 2022). We reasoned that, given the experimental setup of our task, subjects are aware of the condition they are performing. This may lead them to adapt their weighting strategy to meet the task demands. A strong indication of this is the observed session-to-session variability in the PK slope (see section 4.1.1). In the flat condition, the modulation signal in our model does not differ from traditional input signals described in the literature (Wimmer et al., 2015; Wong and X.-J. Wang, 2006). In the Late condition, the modulation signal starts with the increase of the stimulus information, this pushes the dynamics of the decision circuit to later integration of the stimulus while the sensory circuit dynamics remain unaffected. In the Early condition, the integration of the decision starts at the beginning of the trial, however, it only integrates for the time when the information of the stimulus is higher and holds the decision in memory until the decision time (Fig. 5.3, first row).




Fig. 5.3: Modulatory signal and temporal weighting pattern. A) I_0 and I_m time course for each condition: Flat in blue, Late in yellow, and Early in red. B) cartoon of how I_m affect the decision circuit. C) Single-trial dynamics for strong left coherence, for the decision and sensory circuit. D) Single trial of the sensory input for each condition

When examining the population dynamics of each circuit for trials with strong coherence moving to the left, the differences in population activity become more pronounced (Fig. 5.4 A and B). We categorize the activity within the populations based on their selectivity to left (solid line) and right (dotted line) stimuli. The left selective population has higher activity in response to

the strong left coherence compared to right selective. Since the sensory circuit (Fig. 5.4 B) only encodes the sensory information, there is no competition between population E1 (yellow) and E2 (purple). This contrast with the dynamics in the decision circuit, where competition is observed between population D1 (yellow) and D2 (purple) based on their stimulus preferences and triggered by the stimulus statistics. In the Flat condition, the firing rate in the sensory circuit align with the onset of the stimulus, where activity of the left selective population E1 is constantly higher than that of the right selective population E1. In the decision circuit (Fig. 5.4 A), this results in the competition starting at the beginning of the trial, with both populations separating (D1 vs D2) with D1 ending with the higher firing rate. In the Late condition, activity in the sensory circuit increases towards the end of the trial, corresponding to an increase in information from the stimulus. In the decision circuit, competition and integration of the stimulus is delayed until the increase in the sensory circuits activity. In contrast, in the Early condition, activity in the sensory circuit starts at the beginning of the trial and then decreases, aligning with the stimulus information. In the decision circuit, competition starts with the initial pulses in the sensory circuit and decreases as the trial unfolds. However, this activity reaches an steady state that facilitates the correct choice by the end of the trial.

The resulting PK curves from our simulation did not exactly match those observed in the experimental data; however, they fell into the same weighting category (Fig. 5.4 C). The Flat and Early conditions exhibited primacy, while the Late condition showed the highest weight at the fifth pulse followed by a steep decrease. Despite the PK fit not aligning closely with the monkey data, the adaptation of the PKs to the task demands suggests that the modulation signal is a plausible mechanism underlying the flexibility of temporal weighting patterns. To improve the PK fit, we then use a modulation signal that includes preparatory activity (section 3.2.3 for methods), accounting for variations in the time to fixate found on section 4.1.2. This was modeled as an increase in preparatory modulation with a latency of up to 450 ms, reflecting the correlation between LIP activity and task engagement (Section 4.2). In the Early condition a shorter time to fixate indicated higher engagement, resulting in a higher preparatory modulation (Fig. 5.5, red solid line), this was extrapolated to the Flat condition due to similar stimulus statistics in the first pulses (Fig. 5.5, blue solid line), while in the Late condition, the longer time to fixate indicated less engagement, resulting in a lower preparatory modulation (Fig. 5.5, blue solid line). By including the preparatory activity




Fig. 5.4: Population dynamics of the decision and sensory circuits, along with temporal weighting patterns, for all conditions when presented with strong leftward coherence. **A)** Decision dynamics in the Flat, Late, and Early conditions: The left-selective population (yellow line) shows higher activity in response to strong left coherence compared to the right-selective population (indigo line). Both populations compete until their activity separates, reflecting decision integration as the activity ramps up. In the Flat condition, ramp-up starts at the 2nd pulse; in the Late condition, it begins around the 5th pulse when stimulus strength increases; and in the Early condition, activity starts at stimulus onset and decreases after the 4th pulse as stimulus strength decreases. **B)** Sensory circuit dynamics in the Flat, Late, and Early conditions: The left-selective population (yellow line) shows stronger activity than the right-selective population (indigo line) due to the strong leftward coherence and population selectivity, with transient dynamics. In the Flat condition, activity rises at stimulus onset, then decreases slightly in the second pulse. In the Late condition, left-selective activity increases until the 5th pulse, aligning with the stimulus strength change. In the Early condition, activity rises at stimulus onset and decreases after the 4th pulse. **C)** A modulation signal reflecting the stimulus statistics accounts for temporal weighting flexibility in the same categorical range to the monkeys (gray line): primacy in the Flat (blue) and Early (red) conditions, and recency in the Late (yellow) condition.

in the modulation signal we increased the primacy of the kernels in the Flat and Early condition (Fig. 5.5 B, solid blue and red line) with little effect on the PK of the Late condition (Fig. 5.5 B, solid yellow line).

Next, we smoothed the modulation signal (Fig. 5.6 A) to ensure the changes in the I_M transitioned gradually rather than abruptly (section 3.2.3 for methods). This slightly improved the fit between the simulated PKs and the empirical data in the Flat and Early condition. In particular, the weight of the PK in first two pulses increased and closer to the monkeys PK (Fig. 5.6 B). In the Late condition, the PK showed recency, as with previous modulation signals, but the fit was not improved by smoothing (Fig. 5.6 B).




Fig. 5.5: Effect of time to fixate on the modulation signal and temporal weighting strategies by condition. **A)** Comparison of the two version of the modulation signal for all stimulus conditions: the light color represents the modulation signal from Fig.5.3, while the solid color indicates the modulation signal with preparatory activity. **B)** Comparison of the resulting PKs across all conditions: Monkey PKs are shown in gray, solid colors represent PKs with the modulation signal including preparatory activity, and light colors represent PKs without the modulation signal. Blue for Flat, yellow for Late, and red for Early.

Lastly, we examined whether altering the onset of the modulation signal could produce multiple weighting patterns. We found that the varying the start of the modulation signal resulted in a wide range of temporal weighting strategies (Fig. 5.7). We attributed this to changes in the timing of the decision making process, in particular in the start of the attractor dynamics of the decision circuit. An early onset of the modulation signal prioritizes the initial sensory information, resulting in a primacy effect, while a later onset suggest that the sensory evidence arrives later in the trial, resulting in recency effect (Fig. 5.7 B).

In conclusion, the modulation signal in the decision circuit accounts for the flexibility in the temporal weighting strategies by controlling the initiation and the dynamics of evidence integration. By incorporating the preparatory modulation and smoothing the modulation signal we improved the fit in the Flat and Early condition. Lastly, changes in the signal's onset and/or duration produced a wide range of psychophysical kernel. Our results suggest attractor dynamics as a plausible substrate underlying flexible evidence weighting in perceptual decision making.




Fig. 5.6: Temporal weighting patterns for each condition under a smoothly modulated signal A) Smoothed modulation signal for each condition B) Psychophysical kernels by condition, gray line monkey, blue Flat, yellow late, red Early




Fig. 5.7: Range of temporal weighting strategies resulting from delaying the signal onset A) Modulation signal of the Late condition (yellow line) when the signal onset is delayed or advanced . B) Temporal weighting patterns corresponding to each of the signal onsets, color matched to the modulation signal.

5.4 Impact of circuit connectivity on temporal weighting

In this section, we examine how the connectivity between the two circuits—the sensory and the decision circuits— affects the temporal weighting flexibility. We consider the connectivity as bottom-up, when the direction is feed-forward from the sensory to the decision circuit and reflects the sensory encoding of the stimuli, and top-down, when feedback from the decision circuit flows into the sensory circuit and reflects the choice-related activity (Nienborg and Cumming, 2009). For illustrative purposes, we began by isolating the effects of each connectivity direction in the temporal weighting

patterns in our full model, which includes the modulation signal, and the sensory circuit constraints. First, we remove the bottom-up connectivity leaving only the top-down connectivity or the feedback effect. As expected, without the stimulus information flowing to the decision circuit, the PK values flattened, indicating the absence of decision dynamics (Fig. 5.9 A). Next, we remove the top-down connectivity while maintaining the bottom-up strength as in the full model. This allowed us to isolate the bottom-up (feed-forward) effect on the temporal weighting patterns. In the absence of feedback, the PK values remained stable, reflecting the influence of feed-forward processing alone (Fig. 5.9 B).




Fig. 5.8: Connectivity effect on the temporal weighting strategies by condition. A) Comparison of top-down effect in the temporal weighting patterns by condition, dotted line indicates the no bottom-up effect, strong line indicates the simulation of the full model and gray indicates the monkey PKs B) Comparison of bottom-up effect in the temporal weighting patterns by condition, dashed line indicates the no top-down effect, strong line indicates the simulation of the full model and gray indicates the monkey PKs.

Furthermore, we conducted a sensitivity analysis to systematically vary the connection strengths and examine their impact on temporal weighting patterns. We tested this by systematically increasing the strength of each connection keeping the others unchanged, running the simulation with the new parameters, and measuring the resulting PKs. In the bottom-up scenario (Fig. 5.9 B), increasing the connectivity strength resulted in larger primacy in the Flat and Early condition, with increased weights across all the pulses as well. In the Late condition, the effect was most noticeable in the 5th pulse, where only large connectivity strength values caused primacy. This resulted

in an initial increase in the first pulse, followed by a decrease, and then a second increase in the 5th pulse. These extreme changes also impacted the CP's time course (Fig. 5.11). In the top-down scenario (Fig. 5.9 B), increasing the top-down strength resulted in greater primacy in the Flat and Early PKs. In the Late condition, the weight of the fifth pulse increased with the strength but did not reach the weights observed in the monkeys. In the highest connectivity strength, these changes were localized to the pulses of the stimulus with the most information while the rest of the PK remained similar to that of the full simulation.




Fig. 5.9: Impact of bottom-up and top-down strength variations on PK flexibility. A) Effect of bottom-up connection strength on the temporal weighting profile, shown separately for each condition. B) Effect of top-down connection strength on the temporal weighting profile, shown separately for each condition. The color scale represents the connection strength, with lighter shades for weaker strengths and darker shades for stronger strengths. Monkey data PK in gray, and * indicates the parameters used in the full model.

Our results indicate that the bottom-up strength may have a more direct impact on the overall shape of the PKs, while the top-down strength appears

to have a more localized effect. In both cases, increasing the connectivity strength led to higher weight for the pulses carrying the most information, with faster increase and larger variations happening when the bottom-up connectivity was strong (Fig. 5.9, Late condition with connectivity strength of 0.7). In contrast, top-down increases produced more robust weighting patterns. The specific dynamics of the circuits during these simulation with increased connectivity strengths are not shown, however, the CP time course indicate that changes in the connectivity strength influence the underlying circuit dynamics (Fig. 5.11).

5.5 Effect on choice probability

One interesting feature of the two-area models is their ability to separate the choice probability time course into stimulus-related and decision-related activity (Wimmer et al., 2015). This distinction emerges from the bi-directional connectivity between the circuits. The early, bottom-up component of choice probability arises from stimulus related-activity, while the late, top-down component of choice probability is due to the decision-related activity feedback from the decision circuit to the sensory circuit. Each component is obtained by simulating the removal of the other connectivity to the circuit. For example, to isolate the bottom-up effect, we set the top-down connectivity to zero, removing the interaction from the decision circuit to the sensory circuit, allowing the early component to exclusively reflect the stimulus-related activity. In the classic version of the RDM task (Britten, Newsome, et al., 1996), where the stimulus statistics remain constant in time, and the early and late components of CP are often of comparable magnitude yielding flat CP throughout its entire time course. The early component is aligned with the stimulus onset and shows a characteristic transient effect that decreases over time while the late component increases over time (Section 1.1.2) as the decision builds up. This relationship becomes less straight forward when the statistical properties of the stimulus change over time, as is the case of our task. The complexity further increases when we account for the representation of sensory information in visual areas —by considering the firing rates and the effect each additional pulse of information has on the neural population— and the decision circuit’s dynamics affected by the time-varying modulation signal, as our full model does.

In this section, we investigated the bottom-up and top-down contributions to CP in our model in the different evidence integration regimes (flat, early, late). The model proposed in this thesis focuses on the population dynamics rather than individual neurons spikes, so traditional CP measurements based on single-neuron activity cannot be applied (Section 1.1.2). Therefore, we use the choice-dependent firing rates as a proxy for the CP. This measure calculates the difference in firing rates between left and right choices when the stimulus coherence is zero, reflecting the changes in the time course of the choice-related activity independently of the strength of the stimulus. In the reduced model (Section 5), that is when the sensory circuit is unconstrained by the experimental findings and the decision circuit lacks the modulatory signal, we observed that despite the significant differences in stimulus statistics across condition, the CP is similar in all cases (Fig. 5.10 A), showing an increase over time. There is a clear distinction in the time course of the early and the late CP component. The overall increase in CP over time is driven by the late component. When the sensory circuit of our model is constrained so that it resembles the experimentally observed time course of MT neurons (Section 5.2), the time course of the CP was sustained over time and almost flat across conditions (Fig. 5.10 B). The distinction between the early and late components is maintained, with the early component increasing rapidly at the beginning while the late component increases rapidly until it reaches a plateau. It is important to note that this model was not affected by the modulation signal and shows primacy PK in all stimulus conditions (Fig. 5.4 C).

Finally, in the full model (Section 5.3), with both the modulation signal and the sensory circuit restrictions (Fig. 5.10 C), the CP time course varied across conditions. In the Flat condition, the CP initially rose quickly before decreasing to a plateau, while maintaining a clear distinction between the early and late components of the CP time course. In the Late condition, the CP increased in the second half of the stimulus presentation; however, the early and late components remained indistinguishable until the last frame. The absence of CP in the first half of the trial is a consequence of the delayed decision onset with early stimulus information being ignored in the integration process (5.4 A). Finally, the CP in the Early condition behaved similarly to the CP in the Flat condition. In sum, these results show a non-trivial relationship between the dynamics of the sensory and the decision circuits and CP. The findings predict a characteristic CP time course that

can be tested against experimental observations of the Flat, Early and Late stimulus conditions (Section 4.2.4).




Fig. 5.10: Choice-dependent firing rates. A) Choice-dependent firing rates of the reduced model B) Choice-dependent firing rates of the experimentally restricted model C) Choice-dependent firing rates of the full model. Difference in choice-dependent firing rates used as a proxy for choice probability (CP) across conditions. The black line indicates the early component of the CP, and the grey line the late component. Blue line shows the CP during the Flat-stimulus condition, the yellow line represents the Late-stimulus condition, and the red line represents the Early-stimulus condition.

Furthermore, we tested the overall sensitivity of the CP's time course to variations in top-down and bottom-up connectivity to characterize the dynamics of the CP. To examine the impact of increasing top-down connectivity on the overall CP, we incremented its strength while keeping the bottom-up connectivity constant. The change in the parameters primarily influenced the late part of the CP by enhancing the feedback effect of our model. As the strength increased, the late component dominated the early component, regardless of the stimulus condition (Fig. 5.11 A). In contrast, increasing the bottom-up connectivity magnified the feed-forward effect, resulting in a rise in the initial part of the CP. This increase was similar in the Flat and Early condition where the early component of the CP increased with the strength in the connectivity. Interestingly, only one extreme case of bottom-up strength

led to a change in the dynamics of the Late conditions' CP, causing the early component to diverge from the late one at the beginning of the time course (Fig. 5.11 B). This effect can also be seen in the PK shape (Fig. 5.9 B), where the weight in the initial pulses show increased primacy. Overall, our results confirm that the changes in the connectivity strength are limited to affecting the corresponding component: early in the bottom-up and late in the top-down.




Fig. 5.11: Connectivity effect in Choice probability A) Effect of top-down connectivity strength on CP, shown separately for each condition. B) Effect of bottom-up connectivity strength on CP, shown separately for each condition. The color scale represents the connection strength, with lighter shades for weaker strengths and darker shades for stronger strengths.

5.6 Model fitting using simulation based-inference

Finally, we wondered whether our neural network model could also fit the monkey behavioral data quantitatively. Model fitting typically involves maximum likelihood estimation. However, mechanistic models as ours are poorly suited for statistical inference, precisely because the likelihood for

a given observation is intractable. We thus resorted to simulation-based inference (Cranmer et al., 2020). The core idea of simulation-based inference is to estimate the intractable likelihood, for example using neural networks to emulate the mechanistic model. This approach has been successfully used to fit the parameters of generalized drift-diffusion models (Boelts et al., 2022).




Fig. 5.12: Model fit with two-step modulation signal. **A)** Optimal two-stage modulation signal obtained using simulation-based inference. **B)** Psychophysical kernels for the three stimulus conditions obtained for the neural network model with optimized parameters. The monkey PK is added for comparison (gray line).

Here, we took a simplified approach and approximated the choices of the mechanistic model by a logistic regression model (for which we can then compute the likelihood) and by running gradient-free optimization (Methods, section 3.2.4). We fitted the monkey data by optimizing the parameters of the integration circuit, the onset, and height of the modulation signal, and the feed-forward and feedback connectivity between the sensory and the integration circuit. For our standard modulation signal (Fig. 5.6), we optimized the timing of the modulation signal and its intermediate level for each stimulus condition separately and the other network parameters jointly for all conditions (in total 12 parameters). We found a very good match of the experimental PKs and the model PKs for all three stimulus conditions (Fig. 5.12).

We repeated the fitting procedure with a more parsimonious modulation signal which is just a linear ramping input (Fig. 5.3). In this case we optimized the onset and the slope of the modulation signal separately for each condition and the remaining model parameters jointly for each condition (in total 13 parameters). Again, we found a very good match to the experimental




Fig. 5.13: Model fit with linear-ramping modulation signal. **A)** Optimal linear-ramping modulation signal obtained using simulation-based inference. **B)** Psychophysical kernels for the three stimulus conditions obtained for the neural network model with optimized parameters. The monkey PK is added for comparison (gray line).

data (Fig. 5.12). Remarkably, the best fit to the data is obtained when the linear ramping signal starts before the stimulus onset in the early and flat stimulus conditions (but not in the late condition), similar to what we found with the two-step modulation signal (Fig. 5.12).

In sum, these results show that the two-area network model can quantitatively fit the experimentally observed timecourse of evidence integration. The key mechanism to adjust the temporal weighting in the model is a modulation or urgency signal, either two-step or ramping. In both cases, the modulation signal already starts before the stimulus onset and modulates the firing rates of the decision circuit. This is consistent with the finding that LIP firing rates are modulated prior to stimulus onset depending on the level of task engagement (Fig. 4.3).

Conclusion

Our computational model initially explored three potential mechanisms for flexible temporal weighting: modulation of the sensory circuit, modulation of the decision circuit, and the connectivity strength between these circuits. Experimental evidence (particularly, the transient dynamics observed in MT area and PTA activity) suggested that the sensory circuit modulation alone could not account for the observed flexible patterns. Implementing this as

restrictions within our model's sensory circuit maintained a primacy effect in the PK patterns in all stimulus conditions, similar to traditional models.

Next, by introducing a modulation signal to exclusively alter the attractor dynamics within the decision circuit, we successfully modeled flexible temporal weighting strategies by adjusting the start of the evidence integration process. Lastly, by adjusting the signal's onset and duration of the modulation signal, we produced a broad range of PKs (primacy, recency and other combinations), supporting attractor dynamics as a plausible mechanism for flexible evidence weighting in perceptual decision-making.

Lastly, increasing the connectivity strength impacted the shape of the PKs in distinct ways. Stronger bottom-up connectivity led to widespread changes in the PKs weight, enhancing primacy and increasing the weight around the time with the most stimulus information. In contrast, stronger top-down connectivity had a more localized effect, concentrating mostly on the pulse carrying the most information.

One prediction of our model is that the decision-related activity in the sensory circuit (choice probability) should be different across the different conditions. The CP time course is reflecting the different circuit dynamics and can be decomposed in a bottom-up component affecting the earlier time course and a later component reflecting decision-build up.

Optimality of temporal evidence weighting: a RNN approach

We used Recurrent Neural Networks (RNNs) to investigate whether the observed flexible temporal weighting in the experimental data (Introduction 1.2) might reflect optimal behavior. At first glance, the observed primacy and recency kernels seem suboptimal because subjects should report the average motion direction across the 7 stimulus frames, giving equal weight to all frames. Here, we demonstrate that non-uniform evidence weighting arises as an optimal solution in the Early and Late conditions of our task when evidence accumulation is subject to internal noise.

We further study how RNNs learn to flexible adapt their evidence weighting to the different task conditions and what dynamical mechanism supports this flexibility. RNNs, as computational models, allow us to study neural dynamics associated with complex cognitive tasks (Yang and Molano-Mazón, 2021), by using supervised learning to find an optimal solution to a given task under certain architectural constraints. For the optimization, one specifies the network inputs and outputs and the learning algorithm optimizes the network parameters to approximate the desired input-output mapping. We begin by exploring whether RNNs can learn the three task conditions: Flat, Late, Early. Initially, we train one RNN on each task condition independently to analyze and characterize their unique behaviors and performances. Individually, the networks learned the task, and showed temporal weighting that reflected the stimulus statistics when evidence accumulation is subject to noise. Next, we modified the training environment to mirror the experimental conditions faced by the monkeys by training a single RNN in the three conditions in a sequential order. The sequentially trained network was able to solve the task, however, the network failed to adapt its temporal weighting strategy to the requirements of each task and was biased towards the first learned strategy, in our case flat weighting, that was not optimal in all conditions.

Finally, we trained an RNN with additional context inputs (Yang and X.-J. Wang, 2020; Driscoll et al., 2024), which helped the network to learn the task demands of each condition (with different stimulus statistics). The context consisted of the three task conditions, and the network was trained simultaneously in all of them. We found that the network learned the task faster than individually trained networks and showed flexible temporal weighting that closely followed the stimulus statistics. Additionally, we observed that the network is capable to generalize the temporal weighting strategies to novel environments. Finally, we studied the network dynamics underlying this remarkable generalization capabilities.

6.1 RNN architecture and training

All our networks shared the same architecture, consisting of an input layer, a hidden layer, and an output layer, with a batch size and hidden units defined according to the literature (Yang, Joglekar, et al., 2019) and a sequence length that captured the task-time dependencies over two trials ¹. The network receives two key pieces of information as input: a fixation signal, marking the start of each trial and continuing until the decision period begins, and a stimulus input (7 pulses as in the experimental data (see 3.1)) that defines the condition the network is training for²: Flat, Late or Early that the network was trained to discriminate. In this section, we use a minimal version of the task which consists only of 3 phases: a fixation period, stimulus presentation, and a decision period. The dt of the task is 150 ms making the total duration 1350 ms and 9 bins. This input structure is illustrated for the Flat condition in the first column of Fig. 6.2. The second column of Fig. 6.2 illustrates a fully connected network³ that is randomly initialized to provide an unbiased starting point for learning the task, with a ReLU activation function that introduces non-linearities to the model. The connectivity matrix, the weights of the fixation signal, and the output weights for a given trial, out of 15,000, are illustrated in the third column of Fig. 6.2.

¹batch size= 64 , hidden units= 256, sequence length = 20, trial sequence equal to 9 bins of 150 ms each. While the batch size and number of hidden units do not drastically affect the training performance, they are computationally expensive. However, the sequence length directly impacts the RNN's ability to capture time dependencies and retain previous trial information. In computational neuroscience, shorter sequences are often preferred to avoid introducing history or serial bias into the model

²In the case of contextual learning, the contextual signal was added on to the inputs of the network as a binary value, see in Context learning 6.4

³not graphically accurately




Fig. 6.1: Architecture of the RNN and task structure, analog to the monkey task.

Two inputs are fed into the RNN: (1) a binary fixation signal indicating fixation status, which is "on" before and during stimulus presentation and "off" during the decision period, and (2) a noisy signal from the 7-pulse stimulus, which starts 150 ms after fixation onset and its average motion direction indicates left or right motion direction. The network output weight is either negative or positive, corresponding to a left or right choice.

Each network was trained over 6,000 epochs using an Adam optimizer, with a 0.001 loss rate, which optimized the network's weights by minimizing its loss function—the difference between the predicted outcomes and the actual target values. This minimization is equivalent to maximizing the accuracy of the model on every training epoch, which aligns to ideal observer's definition of optimality for this task: giving equal weight to all frames, only maximizing the amount of correct trials to maximize the reward.

Once trained, the weights are fixed and the network is tested on 20,000 new stimuli from the same environment⁴ and produced the response; -1 if the net motion of the stimulus was moving to the left and 1 otherwise. For more details of the network's architecture and training see Section 3.3 and for the extended version on the motivation of the usage of RNNs for cognitive task see Section 1.3.4. The following chapters will delve into the conditions under which the networks successfully learn the task, the optimal strategies for learning when trained sequentially, and the dynamics of rule learning.

⁴with the exception of the context training generalization, see Section 3.3 for the details

6.2 Independently trained RNNs

6.2.1 RNNs adapt their temporal weighting strategies to the task demands

We initiated the training with three networks, each trained independently on one of the three tasks and the same internal noise level ($\sigma = 3$). In this initial approach, summarized in Fig. 6.2, we found that the networks learned their respective task effectively: the training accuracy saturated within the first 300 epochs without significant differences between conditions, and the training loss steadily decreased across all conditions with a slightly higher value in the early condition.




Fig. 6.2: Single-task RNN shows kernel flexibility. A) Results of training three independent RNNs, each on a specific stimulus statistic: Flat (blue), Late (yellow), and Early (red) conditions. Each RNN was trained over 6000 epochs. B) The training accuracy, the proportion of correct predictions on each epoch, saturates within the first 100 epochs. C) The training loss saturates within the first 1000 epochs. D) The psychophysical kernel shows that each trained and tested network used a different temporal weighting strategy

Each network adopted a distinct temporal weighting strategy (Fig. 6.2 D). The Flat-condition trained network exhibited a uniform pattern across all the stimuli presentation, indicating that the RNN weighted each stimulus pulse equally in the decision. Additionally, the Late-condition trained network displayed a recency effect, assigning greater weight to the last three pulses of the stimulus. In contrast, the Early-condition trained RNN showed a

primacy effect, with higher weights in the first 3 pulses. These PK patterns seem to follow the stimulus statistics of the experimental design and closely resemble the PKs found in the experimental data in the Late and Early conditions. Having learned the optimal temporal weighting strategies under each condition, our next step was to understand under which conditions we could expect our findings to hold true. We next look into a series of factors that could determine these patterns when the network is learning: noise, network-to-network variability, and training length.

6.2.2 Temporal weighting patterns are sensitivity to noise levels

First, we explore the effect of the noise on the temporal weighting flexibility. For this, we varied the internal noise in the RNN and measured the kernels' sensitivity. At low noise levels, ($\sigma = 0.5$, Fig. 6.3 lighter color lines) the performance of the RNN shows high accuracy resembling a nearly step-function on the psychometric curve (Fig. 6.3, B). The PKs are nearly uniform across stimulus conditions, with slight increases in the pulses containing the most stimulus information (Fig. 6.3 A). This is confirmed by the slope of the PK being close to zero (Fig. 6.3 C). At high noise levels, ($\sigma = 3$, Fig. 6.3 bold color lines) the performance of the RNNs goes down as show by flatter psychometric curves despite the stimulus condition (Fig. 6.3,B). In the PKs this is reflected as an overall decrease in weights and in their weighting pattern we observe a clear change that aligns more with the task demands, resembling more the distribution of the stimulus (see Fig. 1.6). In particular, the Flat condition exhibits a uniform temporal weighting pattern, while the Late and Early conditions shift to strong recency and primacy, respectively (Fig. 6.3 A).

The results of the sensitivity of the PKs to the noise level are summarized the slope of the PK in Fig. 6.3 C. In the Flat condition, the slope remains stable and close to zero, indicating consistent uniform PKs. In the Late condition, the slope becomes more positive as the noise increases, reflecting a shift towards recency of the PK. Lastly, in the Early condition, the slope becomes more negative as the noise increases, indicating a stronger the primacy effect. In sum, at low noise levels, the PKs across all conditions become similar and essentially uniform, aligning with the optimal solution regardless of the




Fig. 6.3: Noise sensitivity can significantly impact temporal weighting flexibility. **A)** The psychophysical kernel is sensitive to the RNN noise, showing a range of kernel patterns. **B)** the psychometric curve becomes less steep as a result of the higher noise, indicating a decrease in performance due to increased uncertainty in responses. Kernel slope sensitivity to noise is summarized in panel. For **A)** and **B)**, colors indicate the noise level, faintest color for $\sigma = 0.5$, darkest color $\sigma = 3.0$ **C)**, larger noise changes result in clearer temporal weighting patterns.

stimulus statistics. However, at higher noise levels, this pattern changes to reflect the stimulus statistics.

6.2.3 Different temporal weighting strategies emerge robustly across multiple networks

We continue our exploration of other parameters that could affect the optimal temporal weighting strategies by looking into network-to-network variability. To establish that the temporal weighting patterns observed were not unique to a single network, we trained 10 different RNNs on each stimulus condition, measured their temporal weighting strategy and performance to test the robustness of our findings. We found that the 10 networks exhibit similar weighting pattern within a specific condition (Fig. 6.4 A), aligning with our single network training findings. When trained in the Flat condition (Fig. 6.4 A, Flat), the networks commonly used a uniform pattern to weigh the stimuli. In contrast, when trained in the Late condition (Fig. 6.4 A, Late), the networks demonstrate a recency effect, with minimal variations. Lastly, when trained in the Early condition (Fig. 6.4 A, Early), the networks showed a primacy effect, with the highest weights placed on the first three pulses.

Behaviorally, the networks demonstrated robust performance across all 10 networks, consistently showing high accuracy (Fig. 6.4, B)). However, we observed suboptimal idiosyncratic biases, particularly when coherence is zero. In this case, performance should be at chance level, but the networks still show slight deviations, indicating a small bias despite otherwise robust performance with the Flat conditions showing the highest variance and the Early condition the lowest.




Fig. 6.4: Similar psychophysical kernels across multiple RNNs. A) Psychophysical kernel of each condition-trained network, each color indicates one of the 10 networks B) Psychometric curve with sigmoidal fit, each color corresponds to one of the 10 networks.

Overall, we observed robustness in the RNN's temporal weighting patterns across all 10 networks, regardless of the condition, demonstrating consistent temporal weighting patterns. This findings are consistent with previous work (Song et al., 2016; Sussillo et al., 2015), where multiple networks achieve the same behavioral performance however, they may differ in their connectivity and dynamics (Section 6.2.5).

6.2.4 Training length affects the temporal weighting strategies

Finally, we examine the training length of the networks as a potential parameter that influences the optimal temporal weighting strategies. The training length is closely related to the optimization process, with the RNN weights being updated on every training epoch. A shorter training period might fail

to capture accurately the temporal weighting patterns resulting in suboptimal weighting, whereas an extensive training could result in overfitting. To capture the sensitivity of the temporal weighting patterns to the training length, we analyze the slope of the PK on multiple training epochs. We expect the PKs to show a slope of zero in the Flat condition (uniform weighting), a positive slope in the Late condition (recency), and a negative slope in the Early condition (primacy) upon completing training.

We found that the emergence of the PK is sensitive to the training length (Fig. 6.5). Regardless of the stimulus condition, the PK changes rapidly during the initial epochs, reflecting the networks' weighting updates from a randomly connected state (Fig. 6.5 B, gray PKs) to follow the stimulus statistics of corresponding stimulus after 1000 epochs (Fig. 6.5 B). These changes were quantifiable by the PK slope (Fig. 6.5 A), with the patterns beginning to emerge after 400 epochs and stabilizing after 1000 epochs in the Late and Early conditions, and taking a bit longer in the Flat condition. Our results indicate that the networks learn the task conditions easily but require sufficient training, at least 1000 epochs, to reliably consolidate the temporal weighting patterns. This slow consolidation of temporal weighting patterns can also be found in our experimental data (Section 4.1.1).

In conclusion, independently trained networks showed significant sensitivity to noise, producing PKs that resemble the stimulus statistics of the experimental design under high noise levels. These temporal weighting patterns had low variability across different networks, indicating the robustness of the patterns. Lastly, these patterns only emerged when the training length was sufficiently long to produce reliable PKs.

6.2.5 Low-dimensional task-related dynamics

We next explored the dynamics of the independently trained networks for each task condition. For this, we use principal component analysis (Section 3.3) to identify the underlying dynamics of each network. First, we focused on the variance explained for each condition-trained network (Fig. 6.6 A), which quantifies the proportion of total variance in each network's activity explained by each component. Our analysis shows that the first two component captures over 60% of the explained variance in the Flat condition (Fig. 6.6 A, blue bars), over 50% in the Late condition (Fig. 6.6 A,




Fig. 6.5: Kernel slope evolution during training. A) Kernel slope changes for 20,000 epochs. B) Sample of PKs shape for multiple epochs, Flat condition in blue, Late condition in yellow, Early condition in red

yellow bars), and only over 30% in the Early condition (Fig. 6.6 A, red bars). The small amount of explained variance by the first two component in the Early condition suggests that the dynamics are complex in this condition. Additionally, we explore the trajectories in time of the first two principal components for low and high coherence trials (Fig. 6.6). Each trajectory illustrates the activity of the corresponding network when coherence is -20 % vs 20% or -50% vs 50% (Fig. 6.6 B and C respectively). For both coherence levels, on average, in the Flat condition, the two trajectories separate right after the fixation period, with the trajectories being smooth and consistent in each time bin (Fig. 6.6, B , first row). In the Late condition, the two trajectories separate but remain in the same latitude until the decision period (Fig. 6.6, B , middle row). These trajectories are consistent with the notion that the RNNs implement a decision variable in their low-dimensional activity. However, in the Early condition, the two components do not separate or converge to any given direction (Fig. 6.6, B , bottom row). The latter confirms that in the Early condition, the network dynamics are not low dimensional.




Fig. 6.6: PCA dynamics for low and high coherence. A) Explained variance for the same trained networks as in Fig. 6.2. B) and C) Trajectories for low and high coherence trials. Each point represents one of the time bins of the trial (150 ms each, 9 time bins). The average trajectory for left and right trial choice in indigo and orange respectively, start of trial, fixation, marked by a black square and decision by a rhombus. Sample left trials are shown in light gray and right trials in dark gray.

To shed light on the time-related dynamics of the networks, we focus on the distribution of the PCs in each time bin (Fig. 6.7). First, we analyzed the activity of positive high coherence trials of our three independently trained networks (Fig. 6.7 A). During fixation, each independent network exhibited different fixation dynamics, with minimal differences between the Flat and Late conditions. During stimulus presentation and decision time, the Flat condition showed activity buildup in a consistent direction (right), indicating attractor-like dynamics, as expected from the previous figure. Similarly, the Late condition exhibited stable activity with a clear directional increase at 750 ms of the stimulus presentation, corresponding to the fifth pulse of the stimulus, which persisted until the end of the trial, also suggesting attractor-like dynamics. In contrast, the Early condition demonstrated a pattern of switching from positive to negative, changing in each stimulus bin, aligning with our previous findings. Additionally, when examining another of the 10 networks (Fig. 6.7 B), we observed similar attractor-like dynamics for the Flat

and Late conditions. However, the Early condition did not show consistent build-up activity by the end of the trial.




Fig. 6.7: Average network dynamics. **A)** Mean distribution of PC activity for all high-right coherence trials separated by time bin **.B)** Same as in **A)** but for other sample network. Blue indicates the flat condition, yellow the late condition and red for the early condition.

Next, we separated the effect of each component's eigenvectors on the dynamics of the network (Fig. 6.8). For this we focused on trials with coherence equal to zero, this captured the effect of the decision since the choice is random and the stimuli was weak. When averaging the trial for all 10 networks (Fig. 6.8 A), we found that in the Flat condition, the networks increase their activity gradually as the trial unfolds, and it's similar within networks. For the Late condition, the activity remains stable with minor fluctuations till the end of the stimuli and the decision period. Lastly, the Early condition shows larger variability and an unpredictable activity pattern overall. Our results suggest that the networks may have different underlying dynamics despite using the same temporal weighting strategy. Additionally, we look at the dynamics of our initial network and found that the first component captured the dynamics since the start of the trial (Fig. 6.8 B), which could potentially reflect the response to the stimuli while the second component has more activity by the end.




Fig. 6.8: Average dynamics for 10 neural networks and time course of the eigenvalues. **A)** violin plot of the trial average dynamics for the 10 networks by condition **B)** eigenvector values for initial networks, **C)** eigenvector values for sample networks by condition.

6.2.6 Neural activity underlying the decision build-up

Lastly, we look at the neural activity of the network in time. First, we focused on the high coherence stimuli trials and sorted the neurons by the weights of the PCA first component and separated them in left (-50%) and right (50%) coherence (Section 3.3). We found that certain neurons have a preference for left and right coherence as their firing rate is larger (Fig. 6.9 A). This is clearer when we look at the averaged activity of the top 10% neurons, which correspond to the left selective, and the bottom 10% which correspond to the right selective (Fig. 6.9 B).

We use the difference of the trial-averaged temporal activity of the neurons, left minus right trials (Fig. 6.10 A) to show the multiple activation patterns depending on the condition. During the Flat condition, activity increases rapidly at the beginning of the trial and continues as the trial unfolds (Fig. 6.10 A, top row). Neurons exhibit persistent activity during stimuli presentation, but only a few remain active during the decision period. For the Late condition, activity starts to increase by the end of the stimuli presentation (Fig. 6.10 A, middle row). Similarly to the Flat condition, only a few neurons are active during the decision period. In the Early condition, neurons show an on-off




Fig. 6.9: Neuron activity in the networks. A) Trial-averaged temporal activity of neurons, sorted by the weight of the first principal component calculated for Fig. 6.8 A of our sample RNN (Fig. 6.2) of Left selective neurons on the bottom and right selective neurons on the top. B) Average activity for 10% top and bottom neurons

pattern regardless of the part of the trial they are in (Fig. 6.10 A, bottom row). This is consistent with the oscillatory-like dynamics of the condition.

Next, to characterize the monotonic decision build-up in the Flat and Late conditions and the non-monotonic switching in the Early condition, we calculated the cross-correlation between the population states of each network. (Section 3.3). At different time bins we found high correlation for the flat condition due to the broad distribution of overall firing rates across neurons, while the Late condition shows a high correlation only in the last time bins before decision time. The early condition shows a slight correlation at the beginning of the trial and an anticorrelation at the end.

In summary, our results demonstrate that the RNNs exhibit different underlying dynamics depending on the stimulus conditions: low-dimensional dynamics in the Late and Flat conditions, and high-dimensional dynamics in the Early condition. This dynamics are supported by dimensionality reduction analysis (Fig. 6.6 and Fig. 6.8) and network specific (Fig. 6.7). The decision build-up on each condition differed in time course (Fig. 6.10, Fig. 6.9) with the Flat condition aligning with the start of the trial and the Late condition aligning with the change in stimulus statistics. Surprisingly, the decision build-up in the Early condition was more complex to interpret, as the analysis showed an unstable correlation pattern.




Fig. 6.10: Trial average activity sorted by PCA1 weights and correlation. A) Difference in trial average activity, left minus right trials during high coherence stimuli **B)** Time lagged auto correlation between the population states of each network

6.3 Sequentially trained RNN

Up until now, our results come from networks that were trained independently on each of the three task conditions, starting with naive networks that random initial connectivity. This training approach contrast with the training procedure used in the experimental task, where subjects were trained sequentially in batches: first on the Flat condition, then on the Late condition, and finally on the Early condition. In this section, we explore if a RNN is able to show flexible temporal weighting and optimal weighting strategies when trained sequentially, similar to the human and monkey subjects of Levi, Yates, et al., 2018. For this, we first trained one single network in the Flat condition followed by the Late and Early condition for 6000 epochs each (Fig. 6.11 A, see Section 3.3) without constraints on the input-output mapping of the network. Our training resulted in the network effectively learning the first tasks condition: the training accuracy saturated within the first 100 epochs and the training loss steadily decreased and plateaued as training progressed

(Fig. 6.11 B, C). However, when switching to a new task condition, the network's accuracy and loss values remained constant (Fig. 6.11 B, C, middle and bottom row). This suggests that the network fails to adapt to the new stimulus statistics and may continue using the temporal weighting strategy learned in the first task in for all subsequent tasks, leading to suboptimal temporal weighting patterns.

When analyzing the temporal weighting patterns, we found that the Flat condition exhibited a uniform pattern, indicating that the RNN equally weighted every pulse of the stimulus, just as in the independently trained network (6.11, D, blue line vs gray line). In the Late condition, the kernel remained mostly flat until the last pulse, which was weighted more heavily and matched the last pulse of the PK when independently trained (6.11, D, yellow line vs gray line). Lastly, the Early condition displayed an M-shaped pattern with light primacy, and gradual change differing largely from the previous findings (6.11, D, red line vs gray line).




Fig. 6.11: Sequentially trained networks. A) Results of training one RNN on the three conditions sequentially. The network is firstly trained on the flat condition, then the late condition, and lastly, the early condition. Each task was trained over 6000 epochs. B) The training accuracy, saturates within the first 100 epochs, and it does not diminish when changing to the next condition. C) The training loss shows a steady decrease before saturation and do not not change with the condition. D) Psychophysical kernel shows multiple weighting strategies that do not match the flexibility found on the independently trained kernels (gray).

Further inspection of 10 sequentially trained RNNs (see Fig. 8.5 and Fig. 8.6 in Supplementary Materials, Section 8), compared to 10 independently trained RNNs (see Fig. 6.4), revealed that the mean slope of the psychometric

curve in the Flat condition showed little variation between the two training methods, indicating that the Flat PK in Fig. 6.11 D is likely optimal. However, this findings are not surprising since the sequentially training always starts with the Flat condition. A more interesting scenario is the slope of the psychometric curve is reduced in the Late condition when trained sequentially and diminishes further in the Early condition compared to when networks are trained independently. This may indicate that the PKs are sub-optimal in the Late and Early condition (Fig. 6.11 D) and this sub-optimality may increase as the network learns a new task. In summary, the temporal weighting strategies resulting from the sequential training of the network differ from those of the independently trained networks, suggesting that the network becomes increasingly stuck in a suboptimal solution with each task switch.

6.3.1 Sub-optimal temporal weighting

Our initial unrestricted approach to sequential training resulted in sub-optimal temporal weighting. To determine whether optimal behavior could emerge under specific conditions during sequential learning, we modified the input-output mapping of our network. Specifically, we adjusted the weight update process in the RNN—focusing on the connectivity matrix, or output weights⁵—to ensure that only specific parts of the network were trainable. Additionally, we limited the training length to explore whether of the network’s lack of flexibility arise from over-specialization.

When restricting training to only the connectivity matrix weights, we improved the training accuracy and loss (see Fig. 8.7 in Supplementary Materials, Section 8), indicating that the network effectively learned the new task. However, the temporal weighting strategies remained unchanged compared to the restricted approach (Fig. 6.12, connectivity matrix⁶). Additionally, when extending the restriction to include also the output weights, we observed a similar accuracy and loss (see Fig. 8.8 in Supplementary Materials, Section 8), accompanied by smoother PKs in the Late and Early conditions

⁵We also consider the scenario where input weights are updated, however, this modification didn’t allowed for accurate depiction of the stimulus into the network

⁶Minor differences in the shape of the PK are attributed to the variability between simulations

(Fig. 6.12, connectivity matrix - output weights⁷). Neither of the restrictive conditions improved the suboptimality of the PK.

To address our over-specialization concerns, we restricted the RNN's training length to 600 epoch on each condition before switching to the next task condition (Fig. 6.12, All short training). The shorter training duration changed the shape of the Flat PK by decreasing the second pulse and increasing the seventh pulse while maintaining the other pulses consistent with the results observed in the independently trained networks (Fig. 6.12, All short training, blue line). The Late condition exhibited a pronounced recency effect, with the first and last pulses showing similar patterns to those observed in the independently trained networks (Fig. 6.12, All short training, yellow line). In contrast, the Early condition exhibited an inconsistent trend: while the first pulse had a higher weight, it was followed by a sharp drop in the second pulse, followed by an increase in the third pulse and a decrease in subsequent pulses (Fig. 6.12, All short training, red line). It is relevant to note that the shorter training resulted in greater variability in the shape of the PK, affecting the consistency of the PK shapes. This is because the shorter training period directly influenced the network's learning, dynamics and the resulting PK. It is important to note that the shorter training resulted in greater variability in the shape of the PK. This can be attributed to the reduced training time, which influenced the network's learning dynamics and the resulting PK patterns.

6.3.2 Training length does not influence temporal weighting

Our previous results suggests that the sequentially trained RNN struggle to flexible adapt its temporal weighting strategy to the task demands. This limitation may arise from over-specialization or the networks being stuck in a local minima. To address this concern, we explore whether varying the training length affects the temporal weighting patterns. The improvement in the PKs observed during short training (see Fig. 6.12), and evidence from experimental data in monkeys, suggests that flexible temporal weighting can be achieved with relatively short task exposure to different stimulus statistics

⁷Minor differences in the shape of the PK are attributed to the variability between simulations




Fig. 6.12: Comparison of psychophysical kernels across conditions during sequential training with restricted network features. **A** PKs resulting from sequential training where all network parameters were trainable. **B** PKs resulting from sequential training where only the connectivity matrix weights were trainable. **C** PKs resulting from sequential training where both the connectivity matrix weights and the output weights trainable. **D** PKs resulting from sequential training with a fixed duration of 600 epochs per condition before switching to the next stimuli condition. The Flat condition is shown in blue, the Late condition in yellow, and the Early condition in red. PKs resulting from the independently trained network shown in gray for comparison.

(see Section 4.1.1). To test this, we trained our connectivity matrix-restricted sequential RNN (see Section 6.3) in various training epochs, up to 20,000 epochs. At each epoch, we measured the slope of the PKs to track changes in temporal weighting patterns over training time when trained sequentially. We found that when the network was initially trained on the Flat condition, switching to the Late condition was challenging, even with extended training. The slope of the PK showed minimal deviation from zero or from the slope observed in the Flat condition (Fig. 6.13 A, blue and yellow line). This effect amplifies to the last training condition⁸, the Early condition (Fig. 6.13 A, red line). Some individual examples of the PK can be seen in (Fig. 6.11 B), where the difference between PKs is minimal.

⁸This was confirmed by changing the order of the training condition




Fig. 6.13: Training length minimally impacts temporal weighting flexibility. **A** PK slope measured on each training epoch; slopes for the Flat, Late, and Early conditions should be interpreted as a group at each epoch, reflecting the sequential training and the progression of the temporal weighting patterns. **B** Sample PKs at three distinct training epochs, exemplifying the limited effect of the extended training duration on temporal weighting flexibility. The Flat condition is shown in blue, the Late condition in yellow, and the Early condition in red.

Based on our results and existing literature, we conclude that sequential training leads to sub-optimal behavior by interfering with the network's ability to adapt its temporal weighting strategy to new stimulus statistics. Evidence of this is the slope of the PKs consistently remaining close to zero, regardless of the stimulus statistics and the training length. This suggests that the lack of flexibility is not due to over-training, but rather the network being trapped in a local minimum solution, making it difficult to adapt to new task demands after the initial task is learned.

6.4 Context learning

Although RNNs can outperform subjects in certain tasks (Driscoll et al., 2024), they still lack the flexibility needed for rapid learning and efficient task switching (Driscoll et al., 2024) as demonstrated by our sequential training network (see Section 6.3). In this section, to address this limitation by using

context learning, or rule learning (Mante et al., 2013). Context learning allows for extracting context-dependent signals from the environment, similar to how in nature, animals select relevant sensory information for decision-making (Mante et al., 2013). In RNNs, context learning allows a single network to learn the multiple tasks, related and unrelated tasks (Driscoll et al., 2024; Yang, Cole, et al., 2019) by providing contextual signals in the form of the of task-specific inputs.

In our task, the subjects do not have a contextual clue indicating in which task or rule they should follow. However, the task is done in session blocks, all Flat trials for multiple sessions before moving on to Late and so on. This knowledge allows the monkeys to know which task they perform and update their integration strategy, so the context signal could be potentially generated internally. Some quantifiable evidence of this has been shown in our experimental analysis (see Section 4.1.2) where we found context signal in the shape of pre-stimulus activity unique to the Early and Late condition. We trained a single network on the three stimulus conditions simultaneously for 5000 epochs (see Section 3.3). Each stimulus condition was a contextual signal that was an additional input to the network and remaining active until the decision time. This context signal had a binary value: one when the corresponding task was active and zero otherwise (One-hot encoding, Fig. 6.14 A). The contextual signal allowed us to test (1) whether the RNN could learn all three tasks, (2) to characterize the underlying network dynamics, and (3) test whether the RNNs generalize.

Context signals facilitate adaptation to the task demands by encouraging the network to adjust its temporal weighting strategies flexibly across to each conditions. This allows the network to respond dynamically to changes in stimulus statistics. As a result, the PKs of the RNN trained with context signals are practically identical to the PKs of the independently trained networks showing primacy in Early condition, recency in Late condition and uniform weighting in the Flat condition (Fig. 6.14 G).




Fig. 6.14: Context training network. **A)** Additionally to the two standard inputs (fixation and stimulus signal), the network receives a context signal that indicates which condition it's receiving: flat, late, early. **B)** and **C)** The network is trained over 5000 epochs, the training loss shows saturation and a constant loss value (**B)** and the training accuracy, saturates within the first 200 epochs (**C)**. **D),E)** and **F)** Input weights to each neuron with the context signal weights for each condition (**D)**, connectivity matrix (**E)**, and output weights (**F)**, and of the trained network not affected directly by the input context signal **G)** Temporal weighting strategies of each condition (flat in blue, late in yellow and early in red), compared to the independently trained network (gray) on their respective task. **H)** Normalized the task variance by condition of each unit, x-axis: recurrent units; y-axis: different tasks conditions, maximum normalized value is 1, sorted by similarity.

6.4.1 Single-module RNNs emerge from flexible stimulus statistics tasks

Cognitive tasks often involve the reuse of elements across similar tasks, with the brain often recruiting different regions or modules, specialized (visual, motor, language, cognitive, sensory, etc) in a particular process to solve the task (Yang, Cole, et al., 2019). Similarly, in RNNs, this modularity suggests that features learned in one task can be reused to adapt to unfamiliar tasks (Yang, Joglekar, et al., 2019) and each of the modules's element are thought as dynamic motifs (Driscoll et al., 2024). Networks trained on multiple perceptual decision-making tasks (Yang, Cole, et al., 2019; Driscoll et al., 2024) have demonstrated both modularity and compositionality, where dynamic motifs play a key role in facilitating generalization to novel tasks. However,

these studies focus on scenarios with stable stimulus statistics, leaving open questions about how networks adapt to dynamic environments.

For our generalization analysis, we start by measuring the modularity with the similarity of the recurrent units in the task variance (see Section 3.3). We hypothesize that if an RNN use a different group of recurrent units to solve our flexible stimulus statistics tasks, we would have 3 neuronal modules, indicating that the units are specialized and therefore do not share dynamic motifs that facilitate generalization. On the contrary, if the RNN had only one neuronal module, then all tasks would share dynamic motifs, and the network would be able to generalize to novel tasks due to the context signal modifying its dynamics. We found that the larger variance is shared by the same group of neurons across the three tasks (Fig. 6.14 H), with the activity driving the variance clustered in a single module. This suggests that the three tasks share a dynamic motif, which could facilitate generalization to novel tasks.

6.4.2 Task variance and firing rates are driven by a small set of units

To assess variability in neural responses during the task and changes in firing rates within the network, we analyzed the correlation between task variance and the absolute difference in choice-dependent firing rates. Specifically, we compared firing rates when a coherence of 0.5 was presented with Left versus Right choices. Our analysis revealed that units with higher task variance also exhibited greater differences in firing rates between choices

To measure the variability in the neural responses during the task and changes in firing rates within the network, we analyzed the correlation between the task variance and the absolute difference in choice-dependent firing rates. Specifically, we compared firing rates during stimulus presentation when a coherence of 0.5 was presented with Left versus Right choices (Fig. 6.15). Our analysis revealed that units showing large task variance also tend to have a large difference in the firing rate. This activity was concentrated in few units and contrasts with a concentrated cluster with low activity for all conditions. Particularly, the Flat and Early conditions elicit a higher firing rate range compared to the Late condition. This result aligns with the timecourse

of stimulus statistics and the expected integration process during stimulus presentation. In the Late condition, decisions occur later, causing firing rate differences to emerge later on, resulting in a smaller average firing rate difference.




Fig. 6.15: Firing rate variability and task variance across different recurrent units. The distribution of the task variance (y-axis) and the absolute difference in firing rate (x-axis) of individual recurrent units in the three task conditions—flat (blue), late (yellow), and early (red).

6.4.3 Modular network dynamics

Evidence of the lack of modularity is also apparent in the state-space analysis, as there is only a minimal variation in the neural dynamics of each condition. State-space analysis has been shown to be a useful tool to study the neural dynamics of the population level (Carnevale et al., 2015). For this analysis, we examined the neural dynamics at the population level using PCA. We found that the first two principal components accounted for around 50% of the explained variance (Fig. 6.16 A). Analysis of the temporal dynamics of these components revealed little variation between low and high coherence trials and similar patterns between conditions during stimulus presentation, but different at the decision time (Fig. 6.16 B, C, D). The trajectory of the PCAs indicates attractor-like dynamics, this result contrasts with what we observed in the Early condition of the independently trained network (Fig. 6.16 D).

To explore how training conditions impact neural response patterns, we reexamined the neural activity of the network using the same method as employed during independent training. We found distinct neuronal preference for the direction of the stimulus. In the Flat condition, neural activity




Fig. 6.16: Low dimensional for a context trained network. **A)** Explained variance by the first 5 principal components. **B)** Trajectories for low and high coherence trials. Each point represents one of the time bins of the trial (150 ms each, 9 time bins). The average trajectory for left and right trial choice in indigo and orange respectively, start of trial, fixation, marked by a black square and decision by a rhombus. Sample left trials are shown in light gray and right trials in dark gray.

started right after the fixation period. In the Late condition, activity began from the fifth pulse and persisted until the end of the trial. However, with context training, the Early condition's neural activity was entirely different from independently trained, with organized patterns of activity that present though out the trial and closely resembling the findings in the flat condition (Fig. 6.14 A).

To characterize the temporal dynamics of neural activity within the population (Fig. 6.17 B, see Section 3.3). First we measure the difference in choice dependent average activity during strong coherence (as in Section 6.10 from the the independently trained networks) and based on this activity, we calculate the correlation between population states at all time points. Our results show significant differences in the dynamics between the contextually trained network and the independently trained network, in the Early condition. In this condition (Fig. 6.17 B, Early), the activity patterns of the context trained

network remain constant and stable as the trial unfolds, in contrast to the independently trained network, where we observed evidence of anticorrelated activity by the end of the trial. In the Flat condition (Fig. 6.17 B, Flat), we observed, as in the independently trained network, the activity patterns show decision build up after the 2 pulse of the stimulus that lasts up until the decision time. In the Late condition (Fig. 6.17 B, Late), activity patterns increase after the fifth pulse of the stimulus presentation, showing decision build up until the decision time. Overall, we found a certain degree of similarity in the temporal dynamics correlation among the tasks, despite the difference in the stimulus statistics (Fig. 6.17 B). These findings strongly suggest that the three tasks share dynamic motives that could facilitate generalization to novel tasks.




Fig. 6.17: Effect of context training in the trial average activity and cross correlation A) Difference in trial-averaged activity, left minus right trials during high coherence stimuli sorted by the weights of the first principal component. B) Correlation between the population firing rates at different time points (i.e., time-lagged autocorrelation).

6.4.4 Context signals facilitates generalization of temporal weighting strategies and dynamic motifs

To evaluate the generalizability of the dynamic motifs in our network, we tested our context signal trained network (see Section 6.4) on a novel environment. This novel environment was a combination of the context signals given by the Flat, Late and Early condition. For instance, during training (see Section 6.4), the network received a binary context signal for one of the three stimulus for instance, [1,0,0], where "1" indicated that the Flat context was active, and the other two conditions (Late and Early) were inactive (Fig. 6.14). During testing, the context signal was a linear combination of two of the three contexts. The testing contextual signal was varied in 10% increments (e.g., from 100% Flat and 0% Late, to 90% Flat and 10% Late, and so on, up to 0% Flat and 100% Late). This changing environment allowed us to quantify small variations in temporal weighting strategies. For instance, [0.7, 0.3, 0] where 0.7 corresponds to a context signal of Flat condition, 0.3 to the Late condition, and 0 to the Early condition. In this example, if the network indeed generalizes, we would expect the PK pattern to be between the uniform (Flat context) and the recency (Late context).

Our results revealed that the network successfully generalized to a novel environment in two key ways. First, by adapting the temporal weighting strategies to each of the each specific demands of the testing condition. Second, by extrapolating the dynamic motifs to the novel context. The network's ability to adjust its temporal weighting strategies was especially impressive. For instance, when transitioning from the Flat to the Late condition (Fig. 6.18 A), the network started to shift its temporal strategy gradually from uniform weighting (Fig. 6.18 A, blue kernel) to recency weighting (Fig. 6.18 A, yellow kernel). This adaptation of the temporal weighting strategy was quantified by calculating the PK slope for each transition combination (Fig. 6.18 B). In the Flat condition, the slope began at zero, indicating a uniform PK. As the network transitioned to the novel environment, in this case the Late condition (Fig. 6.18 B, top row from blue to yellow), the slope slowly increased to +50%, indicating recency. In contrast, when transitioning to Early condition it decrease to -50%, indicating primacy (Fig. 6.18 B, top row from blue to red). Similarly, when transitioning from the Late condition to Early (Fig. 6.18

B, middle row from yellow to red) and Flat condition (Fig. 6.18 B, middle row from yellow to blue) the slope started in positive values, characteristic of recency and gradually decayed to negative values (primacy in Early) and zero (uniform in Flat) respectively. In the Early condition (Fig. 6.18 B, bottom row, red dot), the initial negative slope indicated primacy, which then increased to reflect uniform and recency patterns when transitioning to the Flat (Fig. 6.18 B, bottom row, from red to blue) and Late conditions (Fig. 6.18 B, bottom row, from red to yellow), respectively.




Fig. 6.18: Context training allows the network to generalize temporal weighting strategies when tested in a novel context A) Kernel generalization resulting from testing the context signal network on a sample of the novel context that switching from Flat to Late condition, passing by every combination in 10% increments. Blue indicates the Flat condition and yellow the Late condition, transition colors in between indicate the percentage contribution from each condition. B) Changes in the PK slope for each novel context and condition for each 10% increment or decrease in the linear combination of the novel environment. Top row is the transition from the Flat condition (blue dot) to the Late (yellow dot) and Early (red dot), Middle row is the transition from the Late condition (yellow dot) to the Flat (blue dot) and Early (red dot), bottom row is the transition from the Early condition (red dot) to the Late (yellow dot) and Flat (blue dot)

Having established the network's flexibility in the temporal weighting strategies, we now focus on examining the dynamic motifs for the relevant contexts.

For example, how the network switch from Flat to Late and then to Early. Our findings are divided into two key areas: the fixation dynamics and the changes in the dynamics across different transition combinations. We first focus on the pre-stimulus activity of the network, also known as fixation dynamics. In networks trained independently, we observed that each network exhibited distinct fixation dynamics. This variation highlighted the differences in the internal state of the network prior to stimulus presentation. We found this variability present when testing in a novel context (Fig. 6.7). As the network transitions between novel environments, the fixation dynamics primarily updates in the dimension of the first principal component, with few variation in the second component. Specifically, the transition from the Flat condition to the Late condition is smooth, with the final states opposite sides of the PC space. In contrast, the transition from Late to Early condition exhibits more variation in the second component, indicating abrupt changes in the fixation dynamics in the last novel environment (100% Early, Fig. 6.19). In sum, the pre-stimulus adaptation of the network has demonstrated that the context signal of the novel environment is reflected in the network's preparatory activity.




Fig. 6.19: Fixation dynamics of the generalized kernels A) Fixation dynamics of each novel environment. Each color indicates the novel context signal the network was tested in and matches Fig. 6.18. The label indicates the color that correspond to the training condition. Blue indicates the Flat condition and yellow the Late condition, transition colors in between indicate the percentage contribution from each condition.

Next, we examine whether this effect extends to the network's dynamics during the trial (Fig. 6.20). For this, we analyze the trial-average state activity of the first two principal components as the network transitioned to a novel environment in all relevant combinations (Flat to Late, Late to Early). Our

findings showed that the most significant changes occurred in the time frame related to the integration of the stimulus. The transition from Flat to the Late (Fig. 6.20, left plot) condition begins with the integration of the stimulus starting at the second pulse (300 ms) and continuing until the end of the trial. As the Late-related contextual signal of the novel environment increases, the network delays the integration of the stimulus to the sixth pulse. During the transition from Late to Early condition (Fig. 6.20, right plot), the integration of the stimulus initiates towards the begin of the trial. As the network faces a stronger Early-related signal, this activity begins increasing closer to the stimulus onset. Once the Early-related signal is 100 %, the activity onsets from the second pulse until the decision time.




Fig. 6.20: Dynamic motifs of context signals in the RNN Both figures, illustrate the patterns of network dynamics associated with each novel context signals from condition to condition. **Flat to Late)** Trial-average state activity of the first two principal components when the network transitions from Flat to Late for each trial time bin. **Late to Early)** Trial-average state activity of the first two principal components when the network transitions from Late to Early for each trial time bin.

In summary, the RNN trained with contextual signals can learn the the task condition as effectively as independently trained RNNs while demonstrating advantages in the dynamics. In contrast to the independent trained RNNs, the contextual RNN shows low-dimensional dynamics across all conditions, regardless of the stimulus statistics, with only a few neurons driving the activity. In particular, the firing rate activity adapts to the specific demands of each task: in the Flat and Early conditions, activity spreads broadly from trial onset to the end, while in the Late condition, activity peaks around the fifth pulse, aligning with the stimulus statistics. Lastly, when training the network as a contextual RNN, a single module emerged, meaning that the same

neurons were active across all tasks. There was no modularity, as the task variance and firing rates were driven by a small set of neurons. This suggests that the three tasks share dynamic motifs that facilitate generalization to unfamiliar environments as the temporal weighting strategies adapt to meet the new environments with minimal training. Evidence of this generalization is reflected in changes in preparatory activity and the timing of the emergence of dynamic motifs.

Conclusion

Our findings reveal that the non-uniform temporal weighting patterns observed in the monkey can arise as an optimal response to internal noise in evidence accumulation. When trained on each task condition, RNNs successfully adapted their temporal weighting strategies to match stimulus statistics, demonstrating robust patterns across networks, despite differences in underlying dynamics. Notably, the Flat and Late conditions displayed low-dimensional dynamics, while the Early condition exhibited high-dimensional dynamics. Analysis of neural activity patterns indicates that only few neurons drive the neural activity at crucial times, with distinct activation and correlation patterns emerging for each condition. In the Flat condition, activity rises quickly and remains correlated throughout the trial. In the Late condition, activity ramps up near the end of the stimulus presentation, showing high correlation just before decision time. The Early condition displays an oscillatory on-off firing pattern, and anticorrelation toward the end of the trial.

To resemble the experimental training set up, the RNN was trained sequentially across condition, the network learned the task but did not adjust its weighting flexibly to the task demands, defaulting to the first learned strategy (flat weighting). To address this sub-optimal behavior, we introduced context learning, which provided the network with a context signal about the current task condition. This allowed the network to train in all conditions (Flat, Late, and Early) simultaneously. This context-driven approach enabled faster learning, flexible temporal weighting that aligned with stimulus statistics, one module of neurons that solving of the conditions, low dimensional dynamics, and effective generalization to new environments, facilitated by the context signal.

Discussion

Mechanisms of flexible temporal weighting

Flexible temporal weighting has been consistently observed across multiple experimental studies, perceptual decision-making tasks and species (Levi, Yates, et al., 2018, Levi, Zhao, et al., 2023, Talluri, Urai, Tsetsos, et al., 2018, Lange et al., 2021), suggesting the existence of a flexible and adaptable decision-making mechanism. This thesis provides new insights into this question by proposing a mechanisms and characterizing the neural dynamics of this flexibility. From a modeling perspective, in hierarchical firing rate network models that include reciprocally connected sensory and decision circuits, the origin and contribution of each circuit to the flexibility in temporal weighting patterns can be analyzed separately.

In the sensory circuit, one hypothesis is that different PKs are realized through progressive attenuation where the attenuation of neural responses at sensory processing stages could account, at least in part, for the decrease in evidence sensitivity over time influencing the temporal weighting strategy (Wilming et al., 2020, Yates, Park, et al., 2017). Another hypothesis suggests that the gain of sensory circuits is modulated over time, influencing the temporal weighting strategy (Levi, Yates, et al., 2018, Wilming et al., 2020). However, our model shows that including a more realistic sensory circuit that accounts for the transient dynamics of MT (Section 4.2.2) and gain modulation of MT (Section 4.2.3) does not affect the temporal weighting patterns, as it only produce primacy (as in Wimmer et al., 2015) with no flexibility between stimulus conditions despite the differences in the stimulus statistics (Section 5.2) .

While the model has demonstrated that the modulation of the sensory circuit does not influence the temporal weighting strategy, the decision circuit plays a crucial role in shaping this flexibility. We propose a modulation signal that alters the attractor dynamics of the decision circuit as the mechanisms responsible for the flexible weighting patterns (Section 5.3). This signal affects the decision circuit in two ways: it initiates the integration of the decision process by pushing the network into a competition regime and alters the decision

dynamics by accelerating or delaying the choice, similar to an urgency signal (Cisek, Puskas, et al., 2009). By adjusting the strength, shape, and timing of the signal, the model can flexibly shift between weighting strategies, with a single underlying mechanism accounting for the experimentally observed PK patterns (primacy in the Flat and Early conditions, recency in the Late condition), and other PK patterns. We designed the modulation signal to reflect the stimulus statistics of the task. However, we also show that simplifying the signal's shape still allows us to achieve categorically the experimentally observed PK patterns. To improve the PK fit, we fit the experimental data quantitatively using Simulation Based Inference (SBI) analysis (see Section 5.6). The best fit is obtained when incorporating a modulation signal that shows pre-stimulus activity. And this may be consistent with the time-to-fixate (see Section 4.1.2) and the related changes of firing rates that we found in LIP (see Section 4.2), supported in two ways. Behaviorally, the time to fixate in the Late condition was longer than in the Early condition, which we attribute to increased engagement and readiness to process the upcoming relevant stimulus. This, in turn, impacted the strength of the pre-stimulus modulation signal. Neurally, we observed a positive correlation between the activity of neurons in LIP and the time to fixate in the Flat condition, indicating that the modulation signal is present in a decision-related area. To confirm if the time course of the modulation signal matches the one in the decision related areas, further experimental work is needed. This would offer critical insights into the pre-stimulus modulation of LIP activity in the Late and Early conditions, helping to confirming validate the model's modulation signal across different task contexts.

The exact nature of the modulation signal remains unknown. Further work needs to identify possible sources of this modulation signal for instance it could reflect ramping input from other brain regions (Finkelstein et al., 2021) or alternatively, the modulation signal might involve a neuromodulator, such as acetylcholine (ACh) or noradrenalin (NA). Both of these neuromodulatory signals are known to influence cognitive and sensory processes. Acetylcholine, for instance, is released in the cortex by cholinergic neurons originating from the basal forebrain, which project long axons that spread across wide cortical areas, allowing for modulation of both localized and global cortical activity. Experimental evidence (Luo and Pinto, 2022) measuring choline concentration in the cortex revealed ACh operates on a dual timescales. Fast, localized release within the medial prefrontal cortex correlates with successful stimulus detection, while slower, widespread release across cortical

areas appears to be associated with overall task engagement. In contrast, noradrenaline (NA) is mainly produced by neurons in the locus coeruleus, and control signals originating from the basal ganglia (Joshi et al., 2016), is also thought to underlie 'urgency signals' that limit deliberation time in reaction time tasks to control the speed-accuracy trade-off (Cisek, Puskas, et al., 2009, Carland et al., 2019). Future modeling work could address how these modulatory signals might vary across different brain regions. For example, one could use larger-scale brain models (Mejias and X.-J. Wang, 2022, Zou et al., 2023) to study perceptual decision making and incorporate the modulation provided by neuromodulators (like ACh from Luo and Pinto, 2022) to study their effect on the temporal integration dynamics. This may lead to a more comprehensive understanding of the modulation in the cortex and the dynamics of the decision-making processes in the brain.

We also explored the connectivity between circuits as another potential source of the temporal weighting flexibility. Sensitivity analysis (see Section 5.4) revealed that each direction of the connectivity had a specific effect over the magnitude of the PK patterns. Stronger bottom-up connectivity altered the PK by enhancing primacy and the weight around the most informative stimulus pulse of each stimulus condition. In contrast, stronger top-down connectivity localized the effect to the pulse carrying the most information. Our analysis showed that PKs are sensitive to connectivity values in either direction. However, the connectivity strength does not account for temporal weighting flexibility, as it primarily influences weight magnitude rather than PK shape. That said, the connectivity analysis provided valuable insights into the dynamics of decision-making through the CP.

Origin of choice-related activity in visual brain areas

Our model extends and modifies the hierarchical spiking neural network proposed in Wimmer et al., 2015 (see Sections: 1.3.2 and 3.2.2 for details of the original model and Section 5 for our model). In the model of Wimmer et al., 2015, the classical perceptual decision making task, equal to the Flat condition of our task, has a CP timecourse with sustained activity throughout the stimulus duration, driven initially by the bottom-up component reflecting the sensory stimulus, which gradually decreases over time as the top-down component reflecting the upcoming choice, increases toward the end of the trial (Fig. 1.4). In our reduced model, equivalent to the Wimmer et al., 2015 but simplified to a firing rate model, we observed the same pattern in the

CP components, a clear separation between the bottom-up and top-down, yet the overall CP does not vary across stimulus condition (Fig. 5.10, A). This suggests that the previous model is limited in capturing the condition-dependent dynamics of variations in the stimulus statistics.

In contrast, in our full model (see Section 5.3, Fig. 5.10), predicts a change in the timecourse of the CP with changing stimulus statistics across condition. In the Early and Flat condition, CP shows transient activity but with clear separation of the bottom-up and top-down components. This is not the case in the Late condition, where the timecourse of the CP aligns with the Late stimulus statistics, increases only after halfway through the stimulus presentation, with no difference between the bottom up and top down components. When we were able to confirm this model prediction in the MT recordings (see Section 4.2.4, Fig. 4.6). In MT neurons, the CP timecourse varies depending on the trial type. In zero coherence trials, the CP in both the Flat and Late condition support our model's predictions, with transient activity and highest CP at the end of the trial, respectively.

One limitation of our model is the assumption that the CP's choice (top-down) and stimulus (bottom-up) components are perfectly aligned, implying that the encoding of motion direction in sensory circuit directly correlates with behavioral choices. However, recent work by Levi, Zhao, et al., 2023, using our same dataset, challenges this assumption. They found that stimulus encoding in MT occurs along a single dimension, capturing neural responses to motion features, while choice-related activity is distributed across multiple uncorrelated dimensions. This misalignment shows that the neurons involved in sensory encoding differ from those driving behavioral decisions, and that the neurons in the stimulus dimension carry minimal information about choices (Levi, Zhao, et al., 2023). These findings challenge the traditional measurement of the CP (see Section 1.1.2) as a method to capture the choice related activity. Classical CP analysis focuses on individual neurons, optimizing for their individual responses, and reflects the correlation between neural activity and behavioral choice. However, traditional CP mainly captures the stimulus dimension and overestimates the amount of choice-related activity because it fails to account for the distributed choice-related signals that are not aligned with the stimulus axis. Further modeling work needs to address how to incorporate the misalignment between stimulus and choice-related activity while also explaining the experimentally observed CP.

In the RNNs section (see Section 6), we examined whether the temporal weighting patterns observed in the experimental data (see Section 1.2.1) were optimal in the classical literature sense, where all stimulus frames are equally weighted in the decision-making process, resulting in an optimal uniform PK (Bogacz et al., 2006). Our results show that the PK patterns in independently trained RNNs are not uniform and reflect the task demands when the noise level is sufficiently high. Intuitively, the RNN is putting more emphasis on portions of the stimulus with a higher signal to noise ratio. While the PK patterns were consistent across many networks, the networks did not exhibit the same neural dynamics. The Flat and Late conditions exhibited consistently low-dimensional dynamics (see Section 6.2.5) with the first PCA capturing a large portion of the variance, consistent with the existence of a one-dimensional decision variable (N. A. Steinemann et al., 2023). Surprisingly, the RNNs trained on the Early condition showed high dimensionality dynamics. These high dimensional dynamics were accompanied by an on-off pattern in the trial average firing rate in the network neurons (see Section 6.2.6). These results were unexpected given the similarity of the PKs in the Flat and Early conditions and suggest that the Early condition may recruit a particular set of features unique to condition, regardless of the training order.

In contrast to the individually trained RNNs, the sequentially trained RNNs displayed limitations in the temporal weighting flexibility, bias toward the first-trained stimulus condition (Flat) and significant difficulty in switching between temporal weighting strategies when exposed to a new condition. In contrast to the RNNs, monkeys exhibited a gradual adaptation in their PK patterns when switching between conditions, maintaining their previous strategy for several sessions before fully adopting the new weighting strategy (section 4.1.1). The limitations of sequential training in neural networks are well-documented in the literature (Kirkpatrick et al., 2017) particularly in the context of catastrophic forgetting and bias toward initial conditions. One solution for this is to expose the model to a variety of conditions throughout training, preventing it from fixating on the first learned pattern (Hammam et al., 2024). This approach aligns with our contextual trained RNN, where we use multiple contextual signals during training to ensure that the network could learn the task without becoming biased towards any particular stimulus condition. When trained exclusively in the Flat, Late and Early conditions as

context signals, our network was able to generalize its knowledge to other temporal weighting patterns, resulting in a wide range of PK patterns often transitory between conditions (i.e showing the full spectrum from primacy to recency). This finding suggest that contextual signals may be the mechanism underlying the rapid context switching observed in the experimental data.




Fig. 8.1: Initial Modulation signal, Fig. 5.3, see Section 3.2.3




Fig. 8.2: Modulation signal with preparatory activity Fig. 5.5, see Section 3.2.3




Fig. 8.3: Modulation signal with preparatory activity smoothed for visualization purposes Fig. 5.6, see Section 3.2.3




Fig. 8.4: Environment in which independent RNNs are trained and tested Each RNN is exposed to an environment where stimulus values vary by condition. Trials begin with a fixation period without stimulus, followed by 7 stimulus pulses and a decision period, during which fixation is broken




Fig. 8.5: Psychophysical kernels across 10 sequentially trained RNNs. A Psychophysical kernel of each condition, sequentially trained. Each color indicates one of the 10 networks




Fig. 8.6: Psychometric curves across 10 sequentially trained RNNs. A Psychometric curve with sigmoidal fit, each color indicates one of the 10 networks




Fig. 8.7: Kernel comparison for when training certain features of the networks.




Fig. 8.8: Kernel comparison for when training certain features of the networks.




Fig. 8.9: Environment in which independent RNNs is trained The RNN is exposed to an environment where stimulus values vary according to the condition, while receiving a context signal that remains active throughout the trial, except during the decision period. This contextual signal indicates which condition is currently active. Trials begin with a fixation period without stimulus, followed by 7 stimulus pulses and a decision period, during which fixation and context are interrupted




Fig. 8.10: Condition and Neuron similarity Row similarity shows the clustering of the stimulus conditions (row similarity), while Column similarity shows the clustering of the units. (see Fig. 6.14 H)




Fig. 8.11: State space resulting from generalization to novel conditions

Bibliography

- Abbott, Larry F (1999). „Lapicque’s introduction of the integrate-and-fire model neuron (1907)“. In: *Brain research bulletin* 50.5-6, pp. 303–304 (cit. on p. 16).
- Acerbi, Luigi and Wei Ji Ma (2017). „Practical Bayesian Optimization for Model Fitting with Bayesian Adaptive Direct Search“. In: *Advances in Neural Information Processing Systems* 30, pp. 1834–1844 (cit. on p. 42).
- Albright, Thomas D (1984). „Direction and orientation selectivity of neurons in visual area MT of the macaque“. In: *Journal of neurophysiology* 52.6, pp. 1106–1130 (cit. on p. 4).
- Aston-Jones, Gary and Jonathan D Cohen (2005). „An integrative theory of locus coeruleus-norepinephrine function: adaptive gain and optimal performance“. In: *Annu. Rev. Neurosci.* 28.1, pp. 403–450 (cit. on p. 65).
- Baker, CL, Robert F Hess, and J Zihl (1991). „Residual motion perception in a "motion-blind" patient, assessed with limited-lifetime random dot stimuli“. In: *Journal of Neuroscience* 11.2, pp. 454–461 (cit. on p. 5).
- Barak, Omri (2017). „Recurrent neural networks as versatile tools of neuroscience research“. In: *Current opinion in neurobiology* 46, pp. 1–6 (cit. on p. 23).
- Boelts, Jan, Jan-Matthis Lueckmann, Richard Gao, and Jakob H Macke (2022). „Flexible and efficient simulation-based inference for models of decision-making.“ In: *eLife* 11 (cit. on pp. 42, 77).
- Bogacz, Rafal, Eric Brown, Jeff Moehlis, Philip Holmes, and Jonathan D Cohen (2006). „The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks.“ In: *Psychological review* 113.4, p. 700 (cit. on pp. 9, 15, 21, 115).
- Born, Richard T and David C Bradley (2005). „Structure and function of visual area MT“. In: *Annu. Rev. Neurosci.* 28.1, pp. 157–189 (cit. on p. 4).
- Britten, Kenneth H, William T Newsome, Michael N Shadlen, Simona Celebrini, and J Anthony Movshon (1996). „A relationship between behavioral choice and the visual responses of neurons in macaque MT“. In: *Visual neuroscience* 13.1, pp. 87–100 (cit. on pp. 3, 5, 6, 28, 33, 63, 73).

- Britten, Kenneth H, Michael N Shadlen, William T Newsome, and J Anthony Movshon (1992). „The analysis of visual motion: a comparison of neuronal and psychophysical performance“. In: *Journal of Neuroscience* 12.12, pp. 4745–4765 (cit. on pp. 3, 5, 7, 28).
- (1993). „Responses of neurons in macaque MT to stochastic motion signals“. In: *Visual neuroscience* 10.6, pp. 1157–1169 (cit. on p. 5).
- Bronfman, Zohar Z, Noam Brezis, Rani Moran, et al. (2015). „Decisions reduce sensitivity to subsequent information“. In: *Proceedings of the Royal Society B: Biological Sciences* 282.1810, p. 20150228 (cit. on p. 20).
- Bronfman, Zohar Z, Noam Brezis, and Marius Usher (2016). „Non-monotonic temporal-weighting indicates a dynamically modulated evidence-integration mechanism“. In: *PLoS computational biology* 12.2, e1004667 (cit. on p. 9).
- Brunel, Nicolas and Xiao-Jing Wang (2001). „Effects of neuromodulation in a cortical network model of object working memory dominated by recurrent inhibition“. In: *Journal of computational neuroscience* 11, pp. 63–85 (cit. on p. 15).
- Brunton, Bingni W, Matthew M Botvinick, and Carlos D Brody (2013). „Rats and humans can optimally accumulate evidence for decision-making“. In: *Science* 340.6128, pp. 95–98 (cit. on pp. 9, 15).
- Carland, Matthew A, David Thura, and Paul Cisek (2019). „The urge to decide and act: implications for brain function and dysfunction“. In: *The Neuroscientist* 25.5, pp. 491–511 (cit. on p. 113).
- Carnevale, Federico, Victor de Lafuente, Ranulfo Romo, Omri Barak, and Néstor Parga (2015). „Dynamic control of response criterion in premotor cortex during perceptual detection under temporal uncertainty“. In: *Neuron* 86.4, pp. 1067–1077 (cit. on p. 103).
- Celebrini, Simona and William T Newsome (1994). „Neuronal and psychophysical sensitivity to motion signals in extrastriate area MST of the macaque monkey“. In: *Journal of Neuroscience* 14.7, pp. 4109–4124 (cit. on p. 6).
- Cheadle, Samuel, Valentin Wyart, Konstantinos Tsetsos, et al. (2014). „Adaptive gain control during human perceptual choice“. In: *Neuron* 81.6, pp. 1429–1441 (cit. on p. 9).
- Cisek, Paul and John F Kalaska (2010). „Neural mechanisms for interacting with a world full of action choices“. In: *Annual review of neuroscience* 33.1, pp. 269–298 (cit. on p. 22).
- Cisek, Paul, Geneviève Aude Puskas, and Stephany El-Murr (2009). „Decisions in changing conditions: the urgency-gating model“. In: *Journal of Neuroscience* 29.37, pp. 11560–11571 (cit. on pp. 14, 15, 112, 113).
- Cisek, Paul and David Thura (2022). „Models of Decision-Making Over Time“. In: *Oxford Research Encyclopedia of Neuroscience* (cit. on p. 14).

- Cohen, Marlene R and William T Newsome (2009). „Estimates of the contribution of single neurons to perception depend on timescale and noise correlation“. In: *Journal of Neuroscience* 29.20, pp. 6635–6648 (cit. on p. 6).
- Compte, Albert, Nicolas Brunel, Patricia S Goldman-Rakic, and Xiao-Jing Wang (2000). „Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model“. In: *Cerebral cortex* 10.9, pp. 910–923 (cit. on p. 15).
- Cranmer, Kyle, Johann Brehmer, and Gilles Louppe (2020). „The frontier of simulation-based inference.“ In: *Proceedings of the National Academy of Sciences of the United States of America* 117.48, pp. 30055–30062 (cit. on pp. 42, 77).
- Crapse, Trinity B and Michele A Basso (2015). „Insights into decision making using choice probability“. In: *Journal of Neurophysiology* 114.6, pp. 3039–3049 (cit. on p. 6).
- Dayan, Peter and Laurence F Abbott (2005). *Theoretical neuroscience: computational and mathematical modeling of neural systems*. MIT press (cit. on p. 61).
- Driscoll, Laura N, Krishna Shenoy, and David Sussillo (2024). „Flexible multitask computation in recurrent networks utilizes shared dynamical motifs“. In: *Nature Neuroscience*, pp. 1–15 (cit. on pp. 23, 82, 99–101).
- Drugowitsch, Jan, Valentin Wyart, Anne-Dominique Devauchelle, and Etienne Koechlin (2016). „Computational precision of mental inference as critical source of human choice suboptimality“. In: *Neuron* 92.6, pp. 1398–1411 (cit. on p. 9).
- Ernst, Udo A, Xiao Chen, Lisa Bohnenkamp, Fingal Orlando Galashan, and Detlef Wegener (2021). „Dynamic divisive normalization circuits explain and predict change detection in monkey area MT“. In: *PLoS computational biology* 17.11, e1009595 (cit. on p. 63).
- Finkelstein, Arseny, Lorenzo Fontolan, Michael N Economo, et al. (2021). „Attractor dynamics gate cortical information flow during decision-making“. In: *Nature neuroscience* 24.6, pp. 843–850 (cit. on pp. 21, 22, 112).
- Gerstner, Wulfram and Werner M Kistler (2002). *Spiking neuron models: Single neurons, populations, plasticity*. Cambridge university press (cit. on p. 16).
- Gold, Joshua I and Michael N Shadlen (2007). „The neural basis of decision making“. In: *Annu. Rev. Neurosci.* 30.1, pp. 535–574 (cit. on pp. 2, 5, 9, 14, 28, 59).
- Haefner, Ralf M, Pietro Berkes, and József Fiser (2016). „Perceptual decision-making as probabilistic inference by neural sampling“. In: *Neuron* 90.3, pp. 649–660 (cit. on pp. 21, 65).
- Hammam, Ahmed, Bharathwaj Krishnaswami Sreedhar, Nura Kawa, Tim Patzelt, and Oliver De Candido (2024). „Structuring a Training Strategy to Robustify Perception Models with Realistic Image Augmentations“. In: *arXiv preprint arXiv:2408.17311* (cit. on p. 115).

- Hanks, Timothy D and Christopher Summerfield (2017). „Perceptual decision making in rodents, monkeys, and humans“. In: *Neuron* 93.1, pp. 15–31 (cit. on p. 2).
- Hart, Eric and Alexander C Huk (2020). „Recurrent circuit dynamics underlie persistent activity in the macaque frontoparietal network“. In: *Elife* 9, e52460 (cit. on p. 20).
- Hawkins, Guy E, Birte U Forstmann, Eric-Jan Wagenmakers, Roger Ratcliff, and Scott D Brown (2015). „Revisiting the evidence for collapsing boundaries and urgency signals in perceptual decision-making“. In: *Journal of Neuroscience* 35.6, pp. 2476–2484 (cit. on p. 15).
- Hess, RH, CL Baker, and J Zihl (1989). „The "motion-blind" patient: low-level spatial and temporal filters“. In: *Journal of Neuroscience* 9.5, pp. 1628–1640 (cit. on p. 5).
- Huk, Alexander C, Leor N Katz, and Jacob L Yates (2017). „The role of the lateral intraparietal area in (the study of) decision making“. In: *Annual review of neuroscience* 40.1, pp. 349–372 (cit. on pp. 2, 7, 8).
- Huk, Alexander C and Michael N Shadlen (2005). „Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making“. In: *Journal of Neuroscience* 25.45, pp. 10420–10436 (cit. on p. 13).
- Inagaki, Hidehiko K, Susu Chen, Kayvon Daie, et al. (2022). „Neural algorithms and circuits for motor planning“. In: *Annual review of neuroscience* 45.1, pp. 249–271 (cit. on pp. 21, 22, 65, 66).
- Jazayeri, Mehrdad and J Anthony Movshon (2007). „A new perceptual illusion reveals mechanisms of sensory decoding“. In: *Nature* 446.7138, pp. 912–915 (cit. on p. 12).
- Joshi, Siddhartha, Yin Li, Rishi M Kalwani, and Joshua I Gold (2016). „Relationships between pupil diameter and neuronal activity in the locus coeruleus, colliculi, and cingulate cortex“. In: *Neuron* 89.1, pp. 221–234 (cit. on p. 113).
- Kandel, Eric R, James H Schwartz, Thomas M Jessell, et al. (2000). *Principles of neural science*. Vol. 4. McGraw-hill New York (cit. on pp. 2, 4, 7).
- Katz, Leor N, Jacob L Yates, Jonathan W Pillow, and Alexander C Huk (2016). „Dissociated functional significance of decision-related activity in the primate dorsal stream“. In: *Nature* 535.7611, pp. 285–288 (cit. on pp. 5, 7).
- Keung, Waitsang, Todd A Hagen, and Robert C Wilson (2019). „Regulation of evidence accumulation by pupil-linked arousal processes“. In: *Nature Human Behaviour* 3.6, pp. 636–645 (cit. on p. 9).
- Kiani, Roozbeh, Timothy D Hanks, and Michael N Shadlen (2008). „Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment“. In: *Journal of Neuroscience* 28.12, pp. 3017–3029 (cit. on pp. 7, 9, 13–15, 17).

- Kiani, Roozbeh and Michael N Shadlen (2009). „Representation of confidence associated with a decision by neurons in the parietal cortex“. In: *science* 324.5928, pp. 759–764 (cit. on p. 14).
- Kirkpatrick, James, Razvan Pascanu, Neil Rabinowitz, et al. (2017). „Overcoming catastrophic forgetting in neural networks“. In: *Proceedings of the national academy of sciences* 114.13, pp. 3521–3526 (cit. on p. 115).
- Kriegeskorte, Nikolaus and Pamela K Douglas (2018). „Cognitive computational neuroscience“. In: *Nature neuroscience* 21.9, pp. 1148–1160 (cit. on p. 18).
- Lange, Richard D, Ankani Chattoraj, Jeffrey M Beck, Jacob L Yates, and Ralf M Haefner (2021). „A confirmation bias in perceptual decision-making due to hierarchical approximate inference“. In: *PLoS Computational Biology* 17.11, e1009517 (cit. on pp. 9, 111).
- Levi, Aaron J, Jacob L Yates, Alexander C Huk, and Leor N Katz (2018). „Strategic and dynamic temporal weighting for perceptual decisions in humans and macaques“. In: *eneuro* 5.5 (cit. on pp. 9–13, 23, 24, 27–29, 33, 49, 52, 58, 59, 62, 94, 111).
- Levi, Aaron J, Yuan Zhao, Il Memming Park, and Alexander C Huk (2023). „Sensory and choice responses in MT distinct from motion encoding“. In: *Journal of Neuroscience* 43.12, pp. 2090–2103 (cit. on pp. 10–12, 27, 30, 49, 58, 63, 111, 114).
- Lim, Koeun, Wei Wang, and Daniel M Merfeld (2017). „Unbounded evidence accumulation characterizes subjective visual vertical forced-choice perceptual choice and confidence“. In: *Journal of neurophysiology* 118.5, pp. 2636–2653 (cit. on p. 14).
- Lo, Chung-Chuan, Cheng-Te Wang, and Xiao-Jing Wang (2015). „Speed-accuracy tradeoff by a control signal with balanced excitation and inhibition“. In: *Journal of Neurophysiology* 114.1, pp. 650–661 (cit. on p. 65).
- Luo, Jiaqi Keith and Lucas Pinto (2022). „Genetically tagging cholinergic diversity“. In: *Neuron* 110.22, pp. 3650–3652 (cit. on pp. 65, 112, 113).
- Luu, Long and Alan A Stocker (2018). „Post-decision biases reveal a self-consistency principle in perceptual inference“. In: *Elife* 7, e33334 (cit. on pp. 10, 12).
- Mante, Valerio, David Sussillo, Krishna V Shenoy, and William T Newsome (2013). „Context-dependent computation by recurrent dynamics in prefrontal cortex“. In: *nature* 503.7474, pp. 78–84 (cit. on pp. 23, 24, 100).
- Mejias, Jorge F and Xiao-Jing Wang (2022). „Mechanisms of distributed working memory in a large-scale network of macaque neocortex“. In: *Elife* 11, e72136 (cit. on p. 113).
- Mochol, Gabriela, Roozbeh Kiani, and Rubén Moreno-Bote (2021). „Prefrontal cortex represents heuristics that shape choice bias and its integration into future behavior“. In: *Current Biology* 31.6, pp. 1234–1244 (cit. on p. 49).

- Murray, John D, Alberto Bernacchia, Nicholas A Roy, et al. (2017). „Stable population coding for working memory coexists with heterogeneous neural dynamics in prefrontal cortex“. In: *Proceedings of the National Academy of Sciences* 114.2, pp. 394–399 (cit. on p. 46).
- Neri, Peter (2004). „Estimation of nonlinear psychophysical kernels“. In: *Journal of Vision* 4.2, pp. 2–2 (cit. on p. 9).
- Newsome, William T, Kenneth H Britten, and J Anthony Movshon (1989). „Neuronal correlates of a perceptual decision“. In: *Nature* 341.6237, pp. 52–54 (cit. on pp. 3, 5).
- Newsome, William T and Edmond B Pare (1988). „A selective impairment of motion perception following lesions of the middle temporal visual area (MT)“. In: *Journal of Neuroscience* 8.6, pp. 2201–2211 (cit. on pp. 3, 5, 28).
- Nienborg, Hendrikje and Bruce G Cumming (2009). „Decision-related activity in sensory neurons reflects more than a neuron’s causal effect“. In: *Nature* 459.7243, pp. 89–92 (cit. on pp. 6, 9, 33, 70).
- Niyogi, Ritwik K and KongFatt Wong-Lin (2013). „Dynamic excitatory and inhibitory gain modulation can produce flexible, robust and optimal decision-making“. In: *PLoS computational biology* 9.6, e1003099 (cit. on p. 65).
- Okazawa, Gouki, Long Sha, Braden A Purcell, and Roozbeh Kiani (2018). „Psychophysical reverse correlation reflects both sensory and decision-making processes“. In: *Nature communications* 9.1, p. 3479 (cit. on p. 9).
- Palmer, John, Alexander C Huk, and Michael N Shadlen (2005). „The effect of stimulus strength on the speed and accuracy of a perceptual decision“. In: *Journal of vision* 5.5, pp. 1–1 (cit. on p. 15).
- Parker, Andrew J and William T Newsome (1998). „Sense and the single neuron: probing the physiology of perception“. In: *Annual review of neuroscience* 21.1, pp. 227–277 (cit. on p. 6).
- Prat-Ortega, Genís, Klaus Wimmer, Alex Roxin, and Jaime de la Rocha (2021). „Flexible categorization in perceptual decision making“. In: *Nature Communications* 12.1, p. 1283 (cit. on pp. 9, 15, 17, 22, 32).
- Rajan, Kanaka, Christopher D Harvey, and David W Tank (2016). „Recurrent network models of sequence generation and memory“. In: *Neuron* 90.1, pp. 128–142 (cit. on p. 23).
- Ratcliff, Roger and Gail McKoon (2008). „The diffusion decision model: theory and data for two-choice decision tasks“. In: *Neural computation* 20.4, pp. 873–922 (cit. on pp. 14, 15).
- Roitman, Jamie D and Michael N Shadlen (2002). „Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task“. In: *Journal of neuroscience* 22.21, pp. 9475–9489 (cit. on pp. 2, 7, 8).

- Romo, Ranulfo, Adrián Hernández, Antonio Zainos, Luis Lemus, and Carlos D Brody (2002). „Neuronal correlates of decision-making in secondary somatosensory cortex“. In: *Nature neuroscience* 5.11, pp. 1217–1225 (cit. on p. 6).
- Roxin, Alex and Anders Ledberg (2008). „Neurobiological models of two-choice decision making can be reduced to a one-dimensional nonlinear diffusion equation“. In: *PLoS Computational Biology* 4.3, e1000046 (cit. on p. 65).
- Salzman, C Daniel, Chieko M Murasugi, Kenneth H Britten, and William T Newsome (1992). „Microstimulation in visual area MT: effects on direction discrimination performance“. In: *Journal of Neuroscience* 12.6, pp. 2331–2355 (cit. on p. 5).
- Shadlen, Michael N, Kenneth H Britten, William T Newsome, and J Anthony Movshon (1996). „A computational analysis of the relationship between neuronal and behavioral responses to visual motion“. In: *Journal of neuroscience* 16.4, pp. 1486–1510 (cit. on pp. 6–8).
- Shadlen, Michael N and Roozbeh Kiani (2013). „Decision making as a window on cognition“. In: *Neuron* 80.3, pp. 791–806 (cit. on p. 14).
- Shadlen, Michael N and William T Newsome (2001). „Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey“. In: *Journal of neurophysiology* 86.4, pp. 1916–1936 (cit. on pp. 7, 8).
- Singh, Gurjeet Sangra and Luigi Acerbi (2024). „PyBADs: Fast and robust black-box optimization in Python“. In: *Journal of Open Source Software* 9.94, p. 5694 (cit. on p. 42).
- Song, H Francis, Guangyu R Yang, and Xiao-Jing Wang (2016). „Training excitatory-inhibitory recurrent neural networks for cognitive tasks: a simple and flexible framework“. In: *PLoS computational biology* 12.2, e1004792 (cit. on pp. 24, 87).
- Steinemann, Natalie, Gabriel M Stine, Eric Trautmann, et al. (2024). „Direct observation of the neural computations underlying a single decision“. In: *Elife* 12, RP90859 (cit. on pp. 7, 8).
- Steinemann, Natalie A, Gabriel M Stine, Eric M Trautmann, et al. (2023). „Direct observation of the neural computations underlying a single decision“. In: *eLife* 12 (cit. on pp. 5, 8, 115).
- Sussillo, David (2014). „Neural circuits as computational dynamical systems“. In: *Current opinion in neurobiology* 25, pp. 156–163 (cit. on p. 23).
- Sussillo, David, Mark M Churchland, Matthew T Kaufman, and Krishna V Shenoy (2015). „A neural network that finds a naturalistic solution for the production of muscle activity“. In: *Nature neuroscience* 18.7, pp. 1025–1033 (cit. on p. 87).
- Talluri, Bharath Chandra, Anne E Urai, Zohar Z Bronfman, et al. (2021). „Choices change the temporal weighting of decision evidence“. In: *Journal of Neurophysiology* 125.4, pp. 1468–1481 (cit. on pp. 10, 12).

- Talluri, Bharath Chandra, Anne E Urai, Konstantinos Tsetsos, Marius Usher, and Tobias H Donner (2018). „Confirmation bias through selective overweighting of choice-consistent evidence“. In: *Current Biology* 28.19, pp. 3128–3135 (cit. on pp. 10, 12, 111).
- Thura, David and Paul Cisek (2014). „Deliberation and commitment in the premotor and primary motor cortex during dynamic decision making“. In: *Neuron* 81.6, pp. 1401–1416 (cit. on p. 22).
- Tsetsos, Konstantinos, Nick Chater, and Marius Usher (2012). „Salience driven value integration explains decision biases and preference reversal“. In: *Proceedings of the national academy of sciences* 109.24, pp. 9659–9664 (cit. on p. 9).
- Uka, Takanori and Gregory C DeAngelis (2004). „Contribution of area MT to stereoscopic depth perception: choice-related response modulations reflect task strategy“. In: *Neuron* 42.2, pp. 297–310 (cit. on p. 6).
- Usher, Marius and Eddy J Davelaar (2002). „Neuromodulation of decision and response selection“. In: *Neural Networks* 15.4-6, pp. 635–645 (cit. on pp. 20, 21).
- Usher, Marius and James L McClelland (2001). „The time course of perceptual choice: the leaky, competing accumulator model.“ In: *Psychological review* 108.3, p. 550 (cit. on pp. 13, 20).
- Vaina, Lucia M, Jeffrey Solomon, Sanjida Chowdhury, Pawan Sinha, and John W Belliveau (2001). „Functional neuroanatomy of biological motion perception in humans“. In: *Proceedings of the National Academy of Sciences* 98.20, pp. 11656–11661 (cit. on p. 5).
- Wang, Jing, Devika Narain, Eghbal A Hosseini, and Mehrdad Jazayeri (2018). „Flexible timing by temporal scaling of cortical responses“. In: *Nature neuroscience* 21.1, pp. 102–110 (cit. on p. 66).
- Wang, Xiao-Jing (1999). „Synaptic basis of cortical persistent activity: the importance of NMDA receptors to working memory“. In: *Journal of Neuroscience* 19.21, pp. 9587–9603 (cit. on p. 16).
- Wilming, Niklas, Peter R Murphy, Florent Meyniel, and Tobias H Donner (2020). „Large-scale dynamics of perceptual decision information across human cortex“. In: *Nature communications* 11.1, p. 5109 (cit. on p. 111).
- Wimmer, Klaus, Albert Compte, Alex Roxin, et al. (2015). „Sensory integration dynamics in a hierarchical network explains choice probabilities in cortical area MT“. In: *Nature communications* 6.1, pp. 1–13 (cit. on pp. 6, 7, 9, 16–18, 21, 33, 34, 36, 37, 59, 60, 65, 66, 73, 111, 113).
- Wong, Kong-Fatt and Xiao-Jing Wang (2006). „A recurrent network mechanism of time integration in perceptual decisions“. In: *Journal of Neuroscience* 26.4, pp. 1314–1328 (cit. on pp. 18–21, 23, 34, 40, 59, 61, 65, 66).

- Wyart, Valentin, Vincent De Gardelle, Jacqueline Scholl, and Christopher Summerfield (2012). „Rhythmic fluctuations in evidence accumulation during decision making in the human brain“. In: *Neuron* 76.4, pp. 847–858 (cit. on p. 9).
- Wyart, Valentin, Nicholas E Myers, and Christopher Summerfield (2015). „Neural mechanisms of human perceptual choice under focused and divided attention“. In: *Journal of neuroscience* 35.8, pp. 3485–3498 (cit. on p. 9).
- XJ, WANG (2002). „Probabilistic decision making by slow reverberation in cortical circuits“. In: *Neuron* 36, pp. 1–20 (cit. on pp. 15–18, 22, 34).
- Yang, Guangyu Robert, Michael W Cole, and Kanaka Rajan (2019). „How to study the neural mechanisms of multiple tasks“. In: *Current opinion in behavioral sciences* 29, pp. 134–143 (cit. on pp. 100, 101).
- Yang, Guangyu Robert, Madhura R Joglekar, H Francis Song, William T Newsome, and Xiao-Jing Wang (2019). „Task representations in neural networks trained to perform many cognitive tasks“. In: *Nature neuroscience* 22.2, pp. 297–306 (cit. on pp. 23, 24, 44, 82, 101).
- Yang, Guangyu Robert and Manuel Molano-Mazón (2021). „Towards the next generation of recurrent network models for cognitive neuroscience“. In: *Current opinion in neurobiology* 70, pp. 182–192 (cit. on pp. 23, 24, 81).
- Yang, Guangyu Robert and Xiao-Jing Wang (2020). „Artificial neural networks for neuroscientists: a primer“. In: *Neuron* 107.6, pp. 1048–1070 (cit. on pp. 23, 82).
- Yates, Jacob L, Leor N Katz, Aaron J Levi, Jonathan W Pillow, and Alexander C Huk (2020). „A simple linear readout of MT supports motion direction-discrimination performance“. In: *Journal of neurophysiology* (cit. on p. 7).
- Yates, Jacob L, Il Memming Park, Leor N Katz, Jonathan W Pillow, and Alexander C Huk (July 2017). „Functional dissection of signal and noise in MT and LIP during decision-making“. In: *Nature Neuroscience* 20, pp. 1285–1292 (cit. on pp. 9, 17, 30, 32, 33, 49, 52, 58, 63, 111).
- Zhang, Jiaxiang (2012). „The effects of evidence bounds on decision-making: theoretical and empirical developments“. In: *Frontiers in psychology* 3, p. 263 (cit. on p. 15).
- Zhang, Jiaxiang, Rafal Bogacz, and Philip Holmes (2009). „A comparison of bounded diffusion models for choice in time controlled tasks“. In: *Journal of mathematical psychology* 53.4, pp. 231–241 (cit. on p. 15).
- Zou, Licheng, Nicola Palomero-Gallagher, Douglas Zhou, Songting Li, and Jorge F Mejias (2023). „Distributed evidence accumulation across macaque large-scale neocortical networks during perceptual decision making“. In: *bioRxiv*, pp. 2023–12 (cit. on p. 113).

List of Figures

1.1	Random Dot Motion Discrimination Task: Diagram and Performance	3
1.2	MT neurons direction selectivity and tuning properties	4
1.3	Response of MT neurons during RDM stimulus presentation. . .	5
1.4	Bottom-up correlations together with top-down signals can lead to sustained CP	7
1.5	Population responses from LIP	8
1.6	Temporal stimulus statistics	10
1.7	Comparison of temporal weighting and psychometric functions within species across stimulus conditions	11
1.8	Choice-correlated activity in MT distinct from motion encoding or psychophysical readout of motion signals	12
1.9	Drift diffusion with symmetric bounds	14
1.10	Recurrent cortical network models	17
1.11	Psychophysical kernel in the hierarchical spiking network model. .	18
1.12	Reduction of a biophysical neuronal decision-making model . .	19
1.13	Bifurcation diagram	21
1.14	Illustration of an external ramping model	22
3.1	Motion discrimination task and Gabor patch stimulus	28
3.2	Temporal weighting profiles	29
3.3	Model architecture, single-trial response, population dynamics and modulation signals	37
4.1	Changes in PKs across recording session	51
4.2	Engagement in human subjects and mean time to fixate by condition in all subjects	52
4.3	Correlation between time to fixate and LIP firing rate	53
4.4	Average firing rate by condition	54
4.5	Pulse-triggered average of MT activity by stimulus condition . .	55
4.6	Choice probability in MT area	56
4.7	Average choice probability in MT area, coherence 0 trials	57
4.8	Average choice probability in MT area, frozen noise trials	57

5.1	Firing rate model architecture, dynamics, and temporal weighting patterns	62
5.2	Sensory Circuit Firing Rate and PTA	64
5.3	Modulatory signal and temporal weighting pattern	66
5.4	Population dynamics of the decision and sensory circuits, along with temporal weighting patterns, for all conditions when presented with strong leftward coherence	68
5.5	Effect of time to fixate on the modulation signal and temporal weighting strategies by condition	69
5.6	Temporal weighting patterns for each condition under a smoothly modulated signal	70
5.7	Range of temporal weighting strategies resulting from delaying the signal onset	70
5.8	Connectivity effect on the temporal weighting strategies by condition	71
5.9	Impact of bottom-up and top-down strength variations on PK flexibility	72
5.10	Choice-dependent firing rates	75
5.11	Connectivity effect in Choice probability	76
5.12	Model fit with two-step modulation signal.	77
5.13	Model fit with linear-ramping modulation signal.	78
6.1	Architecture of the RNN and task structure, analog to the monkey task	83
6.2	Single-task RNN show kernel flexibility	84
6.3	Noise sensitivity can significantly impact temporal weighting flexibility	86
6.4	Similar psychophysical kernels across multiple RNNs	87
6.5	Kernel slope evolution during training	89
6.6	PCA dynamics for low and high coherence	90
6.7	Average network dynamics	91
6.8	Average dynamics for 10 neural networks and time course of the eigenvalues	92
6.9	Neuron activity in the networks	93
6.10	Trial average activity sorted by PCA1 weights and Cross correlation	94
6.11	Sequentially trained networks	95
6.12	Comparison of [psychophysical kernels across conditions during sequential training with restricted network features	98
6.13	Training length minimally impacts temporal weighting flexibility	99

6.14	Context training network	101
6.15	Firing rate variability and task variance across the recurrent units	103
6.16	Low dimensional dynamics for a context trained network	104
6.17	Effect of context training in the trial average activity and cross correlation	105
6.18	Context training allows the network to generalize temporal weighing strategies when tested in a novel context	107
6.19	Fixation dynamics of the generalized kernels	108
6.20	Dynamic motifs of context signals in the RNN	109
8.1	Initial Modulation signal, Fig. 5.3, see Section 3.2.3	117
8.2	Modulation signal with preparatory activity Fig. 5.5, see Section 3.2.3	117
8.3	Modulation signal with preparatory activity smoothed for visualization purposes Fig. 5.6, see Section 3.2.3	117
8.4	Environment in which independent RNNs are trained and tested	118
8.5	Psychophysical kernels across 10 sequentially trained RNNs . .	118
8.6	Psychometric curves across 10 sequentially trained RNNs	118
8.7	Kernel comparison for when training certain features of the networks	119
8.8	Kernel comparison for when training certain features of the networks	119
8.9	Environment in which context RNN is trained	120
8.10	Condition and Neuron similarity	120
8.11	State space resulting from generalization to novel conditions . .	121

