

MAPPING PERCEPTUAL DECISIONS TO CORTICAL REGIONS

PETER N. ZATKA-HAAS

A dissertation submitted in partial fulfilment
of requirements for the degree of

Doctor of Philosophy

UCL
Institute of Neurology

Supervisors:
Prof. Kenneth D. Harris
Prof. Matteo Carandini

December 2018

I, Peter Zátka-Haas, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

ABSTRACT

Perceptual decisions involve a complex interaction of several brain areas. The neocortex is thought to play a major role in this process, but it is unclear which cortical areas are causally involved, and what their individual roles are.

To explore this problem, we trained head-fixed mice to perform a two-alternative unforced-choice visual discrimination task. Mice were rewarded with water for turning a wheel to indicate which of two stimuli had higher contrast, or for holding the wheel still if no stimuli were present.

We developed a hierarchical Bayesian model of the choice behaviour and used this to quantify mouse behaviour in terms of perceptual states such as choice biases and stimulus sensitivities. We also used this model framework to quantify how these perceptual states vary across individual mice and across sessions.

Using widefield calcium imaging, we found robust sequential activation in primary visual, secondary visual, secondary motor, primary motor and somatosensory cortices in response to stimulus presentation. Optogenetic inactivation revealed that only the first two regions: visual (VIS) and secondary motor (MOs) areas, were causally relevant. VIS inactivation was effective earlier than MOs inactivation, which suggests a sequential causal role for these regions.

We observed a surprising effect of VIS inactivation which could only be explained by a downstream subtractive process which integrates information between the two hemispheres. We tested this idea by developing a mechanistic model which was fit to widefield fluorescence data, using the same Bayesian hierarchical framework used earlier. In this model, VIS activity enhances the decision variable associated with contraversive movements and suppresses the decision variable associated with ipsiversive movements. By contrast, activity in MOs enhances both. This model could predict average psychometric behaviour, trial-by-trial variation in choices within a stimulus condition, as well as simulate the effect of optogenetic inactivation.

This thesis therefore shines light on the cortical contributions towards visual discrimination behaviour. This work has implications for the neural processes underlying perceptual decision making more broadly.

IMPACT STATEMENT

The work in this thesis has contributed towards two publications:

Burgess, C.P.* , Lak, A.* , Steinmetz, N.A.* , **Zatka-Haas, P.*** , Bai Reddy, C., Jacobs, E.A.K., Linden, J.F., Paton, J.J., Ranson, A., Schröder, S., Soares, S., Wells, M.J., Wool, L.E., Harris, K.D., Carandini, M., (2017) “High-Yield Methods for Accurate Two-Alternative Visual Psychophysics in Head-Fixed Mice”. *Cell Reports* 20, 2513-2524.

Zatka-Haas P.*, Steinmetz, N. A.* , Carandini M., Harris K. D. (in preparation) “A mechanistic model of cortical contributions to visual discrimination in mice”

** Joint first-authors*

Additionally, this work has culminated in several presentations at conferences including SfN, COSYNE, FENS and Applied Vision Association.

The behavioural task and phenomenological model we have developed will be of methodological use to others in the field of perceptual decision making. A related form of the task is already in use across several labs within the International Brain Laboratory collaboration. Therefore, the results presented in this thesis are highly relevant for this collaboration.

This work will also benefit several disciplines outside of this field. For example, psychiatric illness can lead to disordered decision making. Our results shine light on the role of the neocortex in decision formation, and this may inform further clinical studies as to contribution of neocortex in disordered decisions.

Another possible benefit is in artificial intelligence research. Neuroscience and AI have benefited greatly from cross-collaboration. Our results may help in developing artificial agents which perform similar decision tasks.

ACKNOWLEDGEMENTS

I chose to work with Kenneth Harris because someone told me Kenneth was the smartest person they knew. I joined his lab and discovered several more people who are contenders for this title. I have learnt so much from so many people in this lab, and for this I am grateful.

I would first like to thank my supervisors Kenneth Harris and Matteo Carandini for their direction and guidance during this project. You both manage a large lab but always made time to see me. You have shown me the value of simple and elegant experiments combined with clear data analysis and presentation. I have learnt so much from our conversations, and I will undoubtedly carry these lessons with me in my future work.

Thank you especially to Nick Steinmetz, who has been a de facto supervisor during this project. When I joined the lab, I was struck by your quick and clear thinking, and ability to get a million things done. You have been patient with me when explaining difficult concepts and have always been available to talk about new ideas and help me with technical difficulties. You have also been greatly helpful as a mentor, helping me navigate the more challenging parts of the project, as well as help me think about my future career. I feel I owe a huge amount to you in this project. You have been an inspiration, and without you I don't think the project would have gotten as far as it did.

Thank you to Elina Jacobs, Mika Diamanti, and Andrew Peters for their support, particularly during the first year when I felt a bit lost in the project.

Thank you to Michael Krumin, Armin Lak, Pip Coen, Miles Wells, Laura Funnell and Hamish Forrest for helping with experiments, helpful discussions, and animal training. This project wouldn't have been possible with you! Thank you especially to Hamish for your friendship.

Thank you as well to Charu Reddy for being an amazing lab manager. You made me feel welcome, and you have created a friendly and engaging atmosphere. You've been a glue that has held the lab together.

Lastly, thank you to Robin Snowdon for supporting me through the ups and downs. You are my rock.

CONTENTS

Abstract.....	5
Impact statement	6
Acknowledgements.....	7
Contents	8
Table of Figures	11
Chapter 1 Introduction and Background.....	13
1.1 Philosophical preamble	13
1.2 Decision making as a model of perception.....	15
1.2.1 Theoretical models of decision making	17
1.3 Methods for studying the neural bases of decision making.....	21
1.3.1 Correspondence between neural and perceptual variability	21
1.3.2 Relating neural signals to theoretical models of decision.....	23
1.3.3 Animal models	24
1.4 The neural basis of perceptual decision making	25
1.4.1 A correspondence between early sensory coding and perception ...	25
1.4.2 Parietal/association cortex shows a link with decisions.....	26
1.4.3 Frontal motor areas: decision formation vs action generation	30
1.4.4 A fully distributed cortical process?	32
1.4.5 Subcortical structures	33
1.5 Conclusion and research questions.....	34
Chapter 2 A phenomenological model of mouse visual discrimination.....	37
2.1 Introduction	37
2.2 Methods	39
2.2.1 Ethics.....	39
2.2.2 Surgery.....	39
2.2.3 Water control	40
2.3 A behavioural task in mice	40
2.4 Logistic regression model of choice behaviour	45

2.5	A multi-level framework estimates behavioural variation	54
2.6	Discussion.....	65
2.6.1	Caveats	66
Chapter 3	Sequential activation of cortical areas	69
3.1	Introduction	69
3.2	Methods	70
3.2.1	Behaviour.....	70
3.2.2	Widefield calcium imaging.....	71
3.3	Stimulus and movement-specific activation across cortex	73
3.4	Activation is sequential.....	77
3.5	Discussion.....	79
3.5.1	Caveats	79
Chapter 4	A causal role for visual and secondary motor cortex	81
4.1	Introduction	81
4.2	Methods	82
4.2.1	Behaviour.....	82
4.2.2	Optogenetic inactivation.....	82
4.3	Visual and secondary motor areas causally contribute in this task	84
4.4	The causal role is sequential in time	89
4.5	A phenomenological model of the inactivation effect	94
4.6	Discussion.....	98
4.6.1	Caveats	100
Chapter 5	A mechanistic model of the cortical contribution to visual discrimination.....	103
5.1	Introduction	103
5.2	Method	104
5.2.1	Electrophysiological recording in VISp and MOs	104
5.3	A mechanistic model	104
5.4	Discussion.....	111

5.4.1	Mechanism.....	111
5.4.2	Function	112
5.4.3	Caveats	113
Chapter 6	Conclusions and Outlook	115
6.1	Working model	117
6.2	Future directions	119
6.2.1	Phenomenological model.....	120
6.2.2	Mechanistic model	121
6.2.3	Sub-optimal decision making	122
6.3	Closing remarks	123
Bibliography	125

TABLE OF FIGURES

Figure 1-1 The psychometric function	17
Figure 2-1 The behavioural setup.....	41
Figure 2-2 Task design	42
Figure 2-3 Psychometric data for 5 mice	44
Figure 2-4 Alternative pedestal representation of the psychometric data	45
Figure 2-5 Schematic of the phenomenological model	48
Figure 2-6 Example fits of the three phenomenological models	50
Figure 2-7 Model comparison of different models of contrast non-linearity ...	52
Figure 2-8 Z-space plot of the average data and model fit.....	53
Figure 2-9 Probabilistic Graphical representation of the hierarchical model ..	55
Figure 2-10 Posterior fit of the grand average data.....	60
Figure 2-11 Parameter variation across sessions and subjects.....	62
Figure 2-12 Model can capture variation in NoGo rate across sessions	64
Figure 3-1 Schematic of Allen Common Coordinate Framework.....	72
Figure 3-2 Stimulus-aligned cortical calcium fluorescence	74
Figure 3-3 Movement-aligned cortical calcium fluorescence	76
Figure 3-4 Stimulus-specific robust sequential activation of cortical areas	78
Figure 4-1 Unilateral and bilateral optogenetic inactivation of cortex on choice behaviour	85
Figure 4-2 Higher power optogenetic inactivation	88
Figure 4-3 Optogenetic inactivation at different onset times	90
Figure 4-4 Pulsed high-power optogenetic inactivation experiment.....	92
Figure 4-5 Effect of pulse laser illumination on cortical firing rate	93
Figure 4-6 A probabilistic graphical model of the optogenetic inactivation effect	96
Figure 4-7 Modelling the effect of optogenetic inactivation.....	97
Figure 5-1 Schematic of the mechanistic model	105
Figure 5-2 Electrophysiological recording of broadly-spiking VISp and MOs neurons in response to visual Gabor stimuli	107
Figure 5-3 Mechanistic model predicts behaviour	109
Figure 6-1 Working model of neural basis of two-alternative unforced-choice visual discrimination.....	118

CHAPTER 1 INTRODUCTION AND BACKGROUND

Decision making is a foundational component of animal cognition. Broadly speaking, decision making is defined as the process of evaluating among different alternatives based on some criteria and then committing to one alternative by making a choice. Decisions come in several different forms, depending on the nature of the alternatives being evaluated. This thesis is concerned with perceptual decision making, a decision process where sensory evidence in favour of each alternative is perceptually ambiguous.

A hallmark of perceptual decision making is the finding that for the same sensory stimulus, the choices a subject makes vary considerably. This is a striking observation, because it reveals fundamental limits inherent in a subject's perceptual experience. Psychophysics has an illustrious history in using perceptual decision tasks to quantify perceptual experience from purely behavioural data. However, we are interested in moving on from this, to understand the neural processes underpinning this experience. If we can understand the neural basis of perceptual decision making, we will therefore gain insight into the neural basis of perception itself.

The neural basis for perceptual decision making has attracted significant scientific interest over the past three decades. This has been due to technological innovation in neural recording and perturbation coupled with the development of sophisticated psychophysical tasks in animals. Furthermore, theoretical models of optimal decision making have made it possible to ascribe computational principles to patterns of activity observed in neural tissue. Despite great progress, the field is still in its infancy and many questions remain unanswered.

In this Chapter, we will summarise relevant background literature and discuss some gaps in the knowledge which motivate this research project. Finally, we will define the questions driving this thesis.

1.1 Philosophical preamble

Studies investigating the neural basis of perceptual decision making adopt a few important philosophical stances which are shared in systems neuroscience more broadly. Here we will briefly mention two important stances: the idea of reductive mechanistic explanations, and the idea of

computationalism/representationalism. We will outline their broad claim and mention some caveats.

Many systems/computational neuroscientists attribute ‘understanding the brain’ to establishing mechanisms of function (Kaplan, 2011). Mechanistic explanations focus on the idea that the model system has a ‘purpose’ or ‘function’ (a teleology), and explanation amounts to identifying the components and their causal relations which come together to achieve the function. These types of explanations in turn rely on reductionism: the idea of explaining a complex process in terms of constituent parts and relations between parts (Bechtel et al., 2001). Countless studies in neuroscience adopt this view, for example accounting for behaviour in terms of neural circuitry (e.g. Carandini, (2012)). Within this philosophy, “understanding” a system amounts to being able to manipulate the parts to achieve some goal. This framework also delineates how one can link statements which operate at different levels of description. For example, cognitive features of perceptual decision making may be ‘reduced’ to elementary cognitive processes. These cognitive processes may in turn be ‘reduced’ to patterns of neural network activation. This process continues, all the way down to fundamental physical laws, relying on bridge laws to link the levels together. However, this framework may break down if this reductive assumption doesn’t hold true or the phenomena relies on circular causality. This is exemplified in the concept of emergence. Emergent phenomena are phenomena which do not appear to be reducible to constituent parts. For example, groups of birds or fish coordinate together to move group even though no individual bird or fish makes the decision. It is an open question whether complex phenomena observed in the brain and elsewhere are truly emergent, or instead may be reducible with great effort (Chalmers, 2006).

Another major stance is computationalism/representationalism (Kriegeskorte and Douglas, 2018; Marr, 1982). This school of thought proposes that the brain can be understood as processing information as in a computational device. Computation requires two primary concepts: representation of information, and transformations which can be applied to representations to yield new representations which serves a function. Many studies ascribe neural activity to representations by showing that that neural activity correlates with an external variable, and carries this information forward through a causal chain (deCharms and Zador, 2000). One critique levied against the computational stance is that this

method would identify ‘representations’ in many non-cognitive physical structures (Grush, 2001). For example, the system of ion channels embedded within a cell membrane could be considered to ‘represent’ the electrochemical gradient across it because features of the membrane (e.g. channel density) correlates with the gradient, and this gradient causally impacts electrical functions of the cell. As another example, vibration in the air molecules between two people speaking shows both features of representation. Air molecule vibration correlates with the words spoken by the speaker. This vibration moves through the air and is a causal factor for someone to hear the speech. However, we would argue that this is not an interesting kind of representation. This therefore raises the question of whether representations truly exist, or instead exist as a useful metaphor in human study of these systems.

1.2 Decision making as a model of perception

Perceptual experience is unobservable as it is subjective. However, in the 19th century, Fechner (1860) introduced psychophysics as a discipline aimed at quantifying perceptual experience from observed behaviour. This approach relies on training subjects to distinguish perceptually-ambiguous stimuli and report their percept with a behavioural choice. The perceptual features of the subject could therefore be inferred by monitoring the decision behaviour. Decision tasks therefore form a bedrock upon which perception can be studied (Shadlen and Kiani, 2013). Furthermore, these tasks make it possible to identify neural correlates of perceptual states, because psychophysics simplifies this to identifying neural correlates of behaviour.

Numerous tasks have been developed, all sharing a common framework. An animal is presented with a stimulus, and then in response it must perform a specific action. This stimulus-response association is trained using reinforcement, in the form of rewards for correct actions and/or punishments for incorrect ones. The experimenter then presents stimuli of different ambiguity/difficulty and observes how the response varies in the animal - termed the ‘psychometric curve’. Psychophysics assumes that the observed psychometric curve shape reflects the perceptual limits of the animal, and therefore measures of the psychometric curve shape are measures of perception (Fechner, 1860).

Tasks are categorised by the set of actions used by the animal in signalling a choice. In Go/NoGo tasks, animals perform one action (e.g. a button press) when

a specific stimulus (e.g. a visual grating) is presented and do nothing otherwise (NoGo). While this type of task is simple, using them to make inferences about perception is problematic because NoGo behaviour could either reflect perception of stimulus absence or disengagement with the task itself. In two-alternative forced-choice tasks (2AFC), animals instead make an action on every trial, between two choices. 2AFC tasks therefore do not suffer from the same problem as Go-NoGo tasks when inferring perceptual states as disengagement with the task is reflected in a distinctive absence of action (Stüttgen et al., 2011). Hybrids of this include two-alternative unforced choice tasks (2AUC) which train animals on a 2AFC task but permit NoGo responses (Sridharan et al., 2014).

In 2AFC tasks, several perceptual features can be measured from the psychometric curve. For example: bias, sensitivity and lapse rate (Figure 1-1). Bias is quantified here as a shift in the psychometric curve and it reflects a tendency of the subject to choose one action over another independent of the stimulus. Sensitivity is quantified as the slope of the curve, and it reflects how sensitive the subject is to small changes in the stimulus. Lapse rate is quantified as the plateau performance on the easiest conditions, and it reflects an apparent irreducible error in task performance caused by confusion or disengagement in the task.

The time that a choice is made is also informative of the underlying perceptual process. In tasks where the subject is trained to make a rapid response soon after stimulus onset (a 'reaction time task'), longer reaction times are associated with more difficult stimulus conditions. Similarly, reaction time for error trials are typically longer than for correct trials (Ratcliff et al., 2016). The difference in the amount of time taken to reach a decision between two stimulus conditions is therefore informative about the means by which the two stimulus conditions impact the decision circuitry differently. Another approach to exploring the role of time in decision tasks is to limit the presentation time of the stimulus to pulses of varying duration (Gold and Shadlen, 2000). In these kinds of tasks, choices for trials with briefer stimulus viewing durations are associated with higher rates of error.

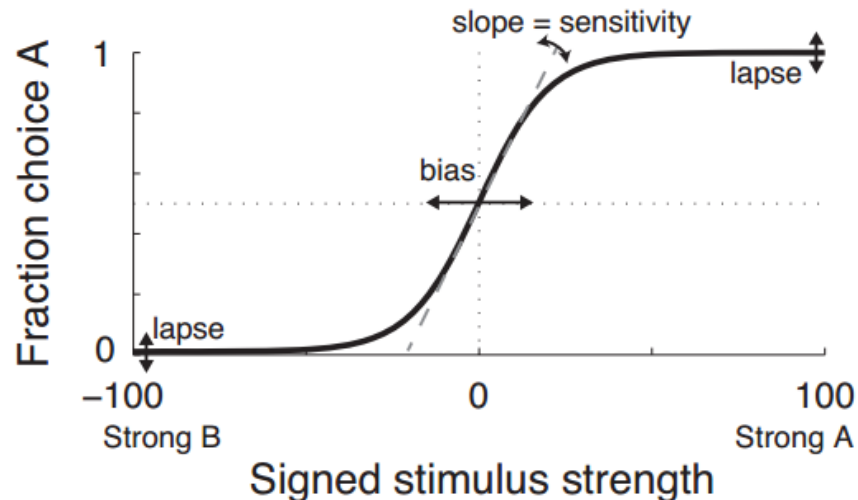


Figure 1-1 The psychometric function

This function quantifies the proportion of trials where the animal makes a choice A as a function of different stimulus strengths. Assuming that the proportion of choices varies smoothly with stimulus strength, the shape of this curve implies several perceptual features. The horizontal shift of the curve corresponds to a decision bias, the slope at the mid-point of the curve corresponds to perceptual sensitivity, and the asymptote of the curve reflects a lapse rate typically assumed to arise from task disengagement or confusion. From Gold and Ding, (2013).

1.2.1 Theoretical models of decision making

The concepts of bias, sensitivity, and lapse rate are idealised metrics of the perceptual process which rely on an assumption that the fraction of choices varies smoothly with stimulus strength. Several theories of decision making have been proposed which define a theoretical framework within which to understand decisions. Two approaches have been particularly insightful in the neural study of decision making: normative signal theory, and normative economic theory. Both seek to define how an optimal ‘ideal observer’ should solve the task but differ in how they account for uncertainty in the decision-making process.

Normative signal theories of decisions place the source of uncertainty in perceptual input (Lynn et al., 2015). Signal detection theory (SDT) proposes the existence of a Gaussian-distributed random decision variable (DV) represented within a deciding agent. The DV can be interpreted as the amount of evidence in favour of making a specific choice. The expected value of the DV is determined by sensory evidence in favour of that choice, while the variance reflects perceptual noise. A decision event is modelled as a sample from the DV distribution. If the DV value exceeds a pre-specified criterion, the choice is made (Tanner and Swets,

1954). By knowing the prior probability of a choice being correct, the reward amount for each choice, the decision criterion can be placed optimally to maximise expected reward (Lynn and Barrett, 2014).

Signal detection theory models decisions when stimuli are static, but what if stimuli are presented continuously in time? The drift-diffusion model (DDM) defines a decision variable (DV) whose value reflects the integrated evidence in favour of one choice over another (Laming, 1968; Ratcliff et al., 2016). This framework has been utilised extensively in models of decision making as it accounts for decisions and deliberation time. Perceptual uncertainty arises from the DV random walk, as well as variability in the DV increment magnitudes from instantaneous stimulus events. A choice is made when the DV value reaches a threshold. These models account for the time that a decision is made as the time when the DV breaches threshold. The value of this threshold sets a trade-off between response time and accuracy. A high threshold value will ensure that choices are only made when evidence is strong relative to noise, but this may take a long time to accumulate and therefore response time will be long. A low threshold will produce quick decisions but will lead to high error rates because brief sampling of the evidence will be noisy. Consequently, DDM defines a mathematical framework for relating choice identity and choice timing (measured by reaction time), and this framework is highly applicable for tasks requiring accumulation of evidence. In this case, both choice and choice timing data are crucial to constraining parameters of these models. This approach can also be utilised in tasks which do not have explicit accumulation of evidence. Decision field models are DDMs which model decisions arising from a static stimulus but accumulating internal 'urgency' signal, which predicts response times (Busemeyer and Townsend, 1993). DDMs could also be applied in tasks where subjects make repeated 'observations' of a fixed and perceptually-ambiguous stimulus, and where the number of repeated observations trades off with the accuracy in choice. In this sense, the subject is integrating static evidence over time. However, since the 'observation' events are neither measured nor experimentally predetermined, it may be difficult to sufficiently constrain a DDM which requires knowledge of the timing for each instantaneous observation event.

Perceptual features such as bias and sensitivity emerge explicitly from components of these theoretical decision models. In classic SDT, bias relates to the position of the criterion, and sensitivity arises from separation of the DV

distributions (Sridharan et al., 2014). In DDM, bias arises from offset in the initial DV value, and sensitivity arises from the gain of the stimulus increments. Lapse rate is not specified in either DDM or SDT model forms because lapses are thought to reflect disengagement with the task itself, not something internal to evaluating evidence. This relationship between psychometric bias and sensitivity, and components of theoretical models, makes it possible to understand observed behaviour in terms of states of an underlying theoretical process.

Normative economic theories of decisions place the source of uncertainty in choice payoff (Lynn et al., 2015). An agent makes a choice among unambiguous alternatives but is uncertain about the amount of reward that will be obtained for making that choice. However, with knowledge of the reward probabilities, optimal behaviour can be defined. The earliest normative economic decision theory of this kind is maximised expected value (MEV). This theory states that the optimal decisions are made by selecting the choice with the largest expected reward (an average of all possible reward states, weighted by the probability of that reward state occurring). However, two problems became clear: 1) this theory cannot account for individual preferences under identical expected values (e.g. risk aversion varies among individuals), and 2) under some probability distributions the expected value is infinite (St. Petersburg paradox). These problems led to its disuse. A new theory took its place which passes reward value through an individual-specific subjective utility function, and optimal behaviour (under a utility function) is achieved by maximising expected utility (MEU) (Hansson, 1994). Neumann et al. (1944) showed that, under four axioms of rational decision making, the deciding agent's behaviour is equivalent to MEU. These axioms are: completeness (preference is spread over a predefined set of possible choices), transitivity (if you prefer A to B, and B to C, then you prefer A to C), continuity (choice preference is a smoothly varying value), and independence of irrelevant alternatives (IIA - the relative preference between two choices is unchanged by the existence of other choices; Luce, (1959)). Under this theory, if animal behaviour obeys these axioms then they are necessarily performing MEU.

These normative models have been employed extensively in behavioural studies ranging from simple perceptual tasks in monkeys to financial risk decision making in humans. However, the fact that each approach assumes a single source of randomness could be considered a weakness. It is likely that both perceptual

uncertainty and economic/decision uncertainty play a role in behaviour. For example, if a 'signals' theory of decision is true, then behavioural variability should arise only from ambiguity in the stimulus. However in situations with identical stimuli, choices are biased by increasing the reward on one side (Summerfield and Tsetsos, 2012). These problems can be reduced by studying simple behaviours where uncertainty is specified to either perception or economic valuation. For example, if we want to study perceptual uncertainty, animals must be over-trained on fixed reward contingencies in the hope that they are not uncertain about potential payoff. Ultimately we want a unified understanding of decisions arising through uncertainty both in perception and in economic valuation, and so future modelling approaches will likely need to find a clean integration of the two sources of uncertainty (Lynn et al., 2015). An integrated framework is especially important as it is unclear whether sensory and economic uncertainty have distinct neural bases - these theoretical concepts might be two views of the same stochastic neural process.

The integration of normative signal and economic decision theories has generated some insight which the individual approaches lacked. For example, classic SDT decomposes the perceptual decisions into two distinct components: stimulus sensitivity, and choice bias. These two quantities are thought to be independent (Sridharan et al., 2014; Tanner and Swets, 1954). Consequently, studies have sought to identify separate neural signatures associated with each of these components. However, Lynn and Barrett (2014) showed that by framing SDT in terms of maximising expected utilities, bias and sensitivity are no longer orthogonal. Under situations with unequal base rates or payoffs, agents with a low sensitivity should be more biased than agents with high sensitivity. Similar work found a dependence between bias and sensitivity when redefining the problem in a Bayesian efficient coding framework (Wei and Stocker, 2016).

The normative signal models discussed above are models of decisions with two alternatives. This therefore makes the models applicable to 2AFC task regimes, but not applicable to tasks with multiple alternatives. Many choice situations an animal may face involve more than two choice options (Churchland et al., 2008), however there is no equivalent optimal decision rule as in the SPRT. Despite this, several attempts have been made to generalise SDT to multi-alternative decisions (Sridharan et al., 2014), and to generalise DDMs to multi-alternative choices (Roe et al., 2001). However, work is ongoing in decision theoretical work to define

models of multi-alternative decisions which are grounded in a normative framework.

1.3 Methods for studying the neural bases of decision making

Establishing the neural basis of perceptual decision making is challenging due to the difficulty in separating cause from effect. However, we can be more confident in our understanding by attacking this problem from multiple directions: 1) Establish a precise correspondence between neural variability and perceptual variability and 2) Fit this correspondence into a broader theoretical framework for decision process. In this section we will introduce these approaches. Afterwards, we will review literature which has used different combinations of these approaches to investigate the neural basis of decision making in monkeys and rodents.

1.3.1 Correspondence between neural and perceptual variability

As sensory information arrives in the brain, this information is transformed into a decision. Tracing this transformation through the brain requires identifying putative signals which correlate with the stimulus and/or decision. These approaches have arisen by the application of signal detection theory to neurophysiological data, and several breakthroughs have emerged from this union.

Any putative signal related to the stimulus (e.g. single neuron firing rates, or a pattern of activity across a population) should show two features. First, the signal should be tuned to features of the stimulus (e.g. visual contrast, or sound intensity). On its own this criterion is insufficient to show that this sensory signal is used in the decision process. Therefore, a second feature required is that the dependence of the signal on the stimulus closely mirrors the psychometric dependence of the decision on the stimulus. This is achieved using Neurometric-Psychometric (NP) comparison (Parker and Newsome, 1998; Stüttgen et al., 2011). This approach defines a putative stimulus-encoding rule for a neural signal (e.g. a neuron encodes a specific stimulus by its firing rate). Neurometric analysis then quantifies how well an ‘ideal observer’ could perform on the decision task by observing this neural signal alone (rather than the external stimulus). The decisions produced by this ideal observer can be quantified across stimulus conditions in a ‘neurometric’ curve. If the behavioural performance of the ideal observer (quantified in the neurometric curve) closely matches the behavioural

performance of the host animal (quantified in the psychometric curve), this suggests that the neural signal corresponds to a signal the animal may use. For example, the firing rates of neurons in visual area MT are tuned to visual motion stimuli. In tasks where subjects must discriminate visual motion, average choices vary as a function of strength of the visual motion. If an ideal observer of the MT signal could make decisions which match those of the host animal, this suggests a role for the neural activity in encoding the stimulus relevant for forming the decisions.

Any putative signal related instead to the upcoming decision should show slightly different features. The signal should correlate with the resulting decision the subject makes, and this correlation should also exist prior to any action being produced, otherwise ongoing motor-related signals trivially correlate with the decision. Signals related purely to encoding the stimulus (and not decisions) will however spuriously correlate with the decision across stimulus conditions because the decisions correlate with the stimulus by design. Therefore, a further criterion is that any decision signal should correlate with the upcoming decision even within trials of identical stimulus conditions. “Choice Probability” quantifies how variability in neural activity mirrors variability in choices within a fixed stimulus condition (Britten et al., 1996; Crapse and Basso, 2015). In the case of visual motion discrimination, if the firing rate of an area MT neuron correlates with the upcoming choice within trials of identical stimuli, this further suggests that the MT neuron is involved in decision formation.

We cannot assert a causal role of neurons in perception from NP-comparison or Choice Probability measures alone because they are correlation statements. To hint at a causal role, we must employ neural manipulation. If a neural structure and neural code forms the causal mechanism of decision formation, then artificially stimulating the structure in a manner consistent with the neural code should act as a complete replacement of the sensory stimulus. Similarly, disrupting activity in these structures should disrupt task performance. Lesions, microstimulation, pharmacological inactivation, and optogenetics are all examples of these approaches. However, two caveats should be made clear with this approach: due to the massively parallel circuitry of the brain, neural manipulation will likely have broad off-target effects in many brain areas (Otchy et al., 2015) or may be compensated for by other neural structures (Li et al., 2016), therefore unequivocal statements of causality restricted to a single neural structure cannot

be made. Also, this approach rests on a general assumption that neural structures perform distinct roles and therefore manipulation affects individual steps within a serial causal chain. This suffers from the mereological fallacy: ascribing to a part what only applies to the whole (Smit and Hacker, 2014). The neural bases, if they are distinct, almost certainly rely on each other's' existence to operate (deCharms and Zador, 2000). This critique seems to apply more to situations where neural manipulation affects behaviour. If neural perturbation has no effect on behaviour, this arguably suggests that the structure is not necessary for producing the behaviour.

1.3.2 Relating neural signals to theoretical models of decision

The complexity of the system makes correlational and perturbational studies hard to interpret. Therefore, a further requirement is that the signals observed and perturbed within the system should relate to a theoretical account of the decision process. A theory ideally constrains the interpretation of the data and proposes new experiments which would help develop the theoretical account of the decision process. Theories can operate at many different levels of abstraction (Kriegeskorte and Douglas, 2018). For example, phenomenological models are mathematical accounts of the relationship between measured quantities, and their success is defined by how well they can summarise the data. By contrast, mechanistic models account for the underlying causal process producing a function. Their success is determined whether components of the model match components of a real physical system, and whether their causal interactions reflect causal processes also evident in the physical system.

The normative signal and economic models described earlier are theoretical models defining relationships between theoretical quantities involved in a decision process. However, they can also be viewed as hypothetical semi-mechanistic models of neural events underlying decision behaviour. In this view, the theoretical quantities are assumed to have a physical realisation in the animal. For example, signal detection theory and drift-diffusion models propose the existence of a decision variable, and therefore a mechanistic account of the decision process would identify the decision variable with a specific signal in the brain (Gold and Ding, 2013; Gold and Shadlen, 2001). Likewise economic theories propose an internal estimate of expected utility which drive decisions (Glimcher and Fehr, 2013). These theories do not propose a full mechanistic account of the stimulus-to-action process, and therefore retain a degree of abstraction. Other

work has focused on developing models which are more mechanistic, and therefore contain more biologically plausible components. For example, recurrent neural network models can be tuned to exhibit attractor dynamics which categorise continuous sensory input into discrete choices (Wang, 2008). These mechanistic models show striking similarity to the decision variable of the DDM described earlier.

1.3.3 Animal models

Much of the foundational work in the neurophysiology of decision making has been done in non-human primates such as monkeys (Parker and Newsome, 1998). Monkeys can be trained to perform sophisticated decision tasks. One example task is the random dot task. In this task, monkeys observe a set of dots on a screen moving in one direction among a set of distractor dots moving in random directions. The coherence of dot motion can be controlled experimentally, thereby adjusting the uncertainty about the dot movement direction. The task for the monkey is to saccade to indicate the direction of the dot movement. While monkeys perform this task neural activity from several brain areas can be monitored during task performance. One problem with monkeys however is that training takes many months, and this limits the amount of data that can be acquired. There are also fewer tools available for genetically targeting specific cell types for neural recording and manipulation.

Rodents are now an attractive model to study for decision making (Carandini and Churchland, 2013). This is due to recent developments in high-throughput neural recording, cell-type-specific fluorescence imaging, and optogenetic manipulation. There was initially some doubt as to whether rodents could perform sophisticated decision tasks akin to those used in monkeys. The Carandini lab developed a visual discrimination task in mice where head-fixed mice turn a wheel to indicate whether the left or right side showed a Gabor stimulus (Burgess et al., 2017). Mice could perform this task very well, achieving performance which was only thought possible in monkeys. Additionally, training only takes about one month, which has increased experimental throughput. Similarly, the Brody lab developed a click-accumulation task in rats inspired by the random-dot task. Rats are trained to listen to a series of auditory clicks on the left and right side, and move to a nose-poke hole on the left or right corresponding to which side had the greater number of clicks (Brunton et al., 2013). The average number of clicks is experimentally controlled, but each trial is a random train of clicks drawn from a Poisson

distribution. The rat must therefore retain an accumulating measure of the difference in clicks from each side. These click-accumulation tasks are also amenable to theoretical analysis because, assuming that the behaviour follows from a theoretical evidence-accumulation process (e.g. DDM), the animal's time-varying 'decision variable' can be inferred from the relationship between the rat's reaction time and the experimentally-determined stimulus event times (Brunton et al., 2013).

1.4 The neural basis of perceptual decision making

The decision process can be conceived of as the following chain of events: sensation of stimulus, formation of decision, and generation of action. Several landmark studies over the past few decades have explored the neural basis underlying each stage. In this section we will first discuss some work demonstrating a neural basis for sensory perception in early sensory cortical areas. We will follow on with studies identifying decision signals with activity in parietal areas. Finally, we will discuss some investigation into frontal-motor regions thought to be involved in action generation, emphasising a role for classically 'motor' regions in early decision formation.

1.4.1 A correspondence between early sensory coding and perception

Early monkey work in the 1980-1990s demonstrated that visual area MT mediated performance of the random dot decision task. While monkeys performed the task, Newsome and Pare (1988) lesioned area MT in monkeys and found that their psychophysical thresholds were impaired. However, thresholds for a contrast discrimination task were unaffected. This suggested that MT is necessary for the perception of motion but not visual stimuli in general. This left open the question of whether the uncertainty in the visual stimulus reported perceptually (via highly variable choices) is mirrored neurally as variability in MT firing. Britten et al. (1992) recorded from MT neurons in monkeys performing a similar random dot task. The authors found that the neurometric dependence of the neural activity on the dot motion coherence closely mirrored the psychometric dependence of choices on coherence. Additionally, on trials with identical stimulus coherence, trial-by-trial variability in neural activity could weakly predict variability in choices (Choice Probability; Britten et al. (1996)). These findings were important because they were one of the first to demonstrate a precise correspondence between perceptual experience and neural activity. However, the weak Choice

Probability found in this area, coupled with the fact that neurometric sensitivity in this area was greater than psychometric sensitivity suggested that area MT was not itself the seat of decision formation. Recent work in monkeys and rodents has shown that (albeit weak) Choice Probability in early sensory areas probably arises from top-down feedback from other structures which form the decision (Nienborg and Cumming, 2009; Yang et al., 2016). Consistent with this, Liu and Pack (2017) trained monkeys to discriminate drifting gratings. Muscimol inactivation of MT had no effect in these monkeys. The authors then trained monkeys on a random dot discrimination task and found that muscimol inactivation of MT now impaired performance of interleaved drifting grating trials. This suggests that MT confers information related to the stimulus, which is passed downstream to structures which generate the actual decision.

In other sensory modalities, Romo and Salinas (1999) trained monkeys on a vibrotactile discrimination task. In this task, two vibrating mechanical stimuli are applied to the monkey's finger, one after another, with a time delay in-between. The monkey indicates with a button press or lever pull whether the second stimulus was at a higher vibrational frequency than the first. Previous work showed that neurons in primary somatosensory area (S1) fire in response to this vibration stimulation. However, it was unclear whether the activity here drives perception. The authors stimulated neural activity in S1, and found that this could act as a complete replacement for the vibrational stimulus for monkey perceptual decisions. Hernández et al. (2000) found that the sensitivity of S1 neural firing to the mechanical stimulus closely matched the psychophysical sensitivity of the monkey to the stimulus. Therefore, just like in visual area MT, these findings demonstrate a close correspondence between neural activity and perception in early sensory areas.

Early sensory areas therefore appear to relay information relating to the stimulus, but do not form the seat of decision formation. Subsequent studies have explored whether the downstream targets of the early sensory areas may instead relate more to the decision.

1.4.2 Parietal/association cortex shows a link with decisions

Several landmark studies identified decision-related signals in the Lateral Intraparietal area (LIP) of the Posterior Parietal Cortex (PPC) in monkeys performing the random dot task (Roitman and Shadlen, 2002; Shadlen and Newsome, 2001, 1996). The authors reported an important set of findings. Firstly,

LIP firing is topographically mapped to the target location of a saccade therefore seems related to action planning. Secondly, after onset of the visual stimulus, LIP firing ramped up gradually if the random dot stimulus moved in the direction associated with that LIP's saccade receptive field. Thirdly, the slope of the ramping was proportional to the coherence of the visual stimulus. Fourthly, ramping activity appeared to reach an upper bound when the saccade was made, thereby predicting the reaction time of the saccade. Finally, the activity of LIP persisted through a delay period where the monkey waits before making a saccade. In this delay period, there is no visual stimulus on-screen, nor any saccadic movement in the monkey. Therefore, LIP activity here does not reflect pure sensory or motor information. The authors suggested that LIP activity reflects the decision process itself. Huk and Shadlen (2005) trained monkeys on a similar task but instead of presenting a random dot stimulus of one coherence value, they momentarily perturbed this value within a trial. The authors found that the ramping of LIP activity was consistent with an integration of variable 'evidence' in favour of one choice. If LIP is the site of evidence integration, artificially increasing the activity should bias choices. Consistent with this, Hanks et al. (2006) found that choice towards one side were increased when microstimulating a subset of LIP neurons whose receptive fields covered one target. Additionally, occasional choices to the other side were slower. This is consistent with the evidence accumulation model of decision making. LIP activity also shows correlation with decision variables associated with decisions made under economic uncertainty (Platt and Glimcher, 1999), which suggests that LIP may flexibly represent different kinds of decisions.

However, subsequent work has cast doubt on the role of LIP in forming the decisions in the random dot task. In the original LIP-recording studies mentioned above, the measured neural activity shows Choice Probability measures for individual cells which are far lower than what would be expected from a decision signal (Crapse and Basso, 2015). Additionally, pharmacological inactivation of monkey LIP using muscimol has no effect on this task, but does bias saccades in a free-choice task (Katz et al., 2016). Other findings have shown that LIP activity may not correspond exactly to the theoretical decision variable of the accumulation models. In these models, the threshold of the decision variable sets the speed-accuracy trade-off. If accumulation-to-bound is a good account of LIP activity, then when a monkey emphasises speed/accuracy differently, then we would predict a corresponding shift in the boundary for LIP activity prior to

choices. Hanks et al. (2014) explored this in a motion discrimination task where monkeys were cued to emphasise speed or accuracy on different trials. Behaviourally, the monkeys succeeded in doing this, and a DDM could capture the behavioural changes as a shift in the decision boundary as predicted. However, when the authors recorded from LIP they found that the neural firing boundary was unaffected between these conditions. Taken together, these studies suggest that LIP does not drive saccadic decisions requiring integration of motion evidence but may instead play roles related to goal-directed saccades in other contexts.

Several rodent studies have helped illuminate the possible role for PPC in decision tasks. Many of these studies have employed the Poisson-click task outlined earlier, which is analogous to the random dot task in monkeys. Hanks et al. (2015) recorded from PPC and found neurons which showed a gradually-ramping firing prior to choice, akin to the ramping activity of monkey LIP in the random dot task. Just like LIP, high-dose muscimol inactivation of rat PPC had no effect on this decision task however did affect a free-choice task (Erlich et al., 2015). However other work points to PPC still causally relevant to some decision tasks. Licata et al. (2017) and Raposo et al. (2014) found that optogenetic inactivation of rat PPC biased visual, but not auditory, discrimination. This paints a picture where PPC acts as a higher visual area, not a seat of general decision making. Consistent with this, Goard et al. (2016) found that bilateral inactivation of mouse PPC impaired performance in a visual detection task, but only if inactivation was during the stimulus epoch.

Other work points to PPC playing a role in memory-guided navigation. Harvey et al. (2012) utilised 2-photon calcium imaging of PPC in head-fixed mice. Mice walk down a virtual corridor with a stimulus section (grating stimuli on the walls), a plain grey 'delay' section, and then a junction. Mice should turn left or right to indicate which wall had the higher contrast. During this task the authors found that PPC activity flowed in a sequence along a chain of anatomically intermingled cells, over the entire period of a trial. The sequence depended on the eventual choice made by the mouse and was maintained during the 'memory' phase of the task. The authors found that inactivating PPC impaired mice's ability to perform this task, but not on a modified task where the visual stimulus was presented through the entire corridor. This study offers a novel view of the PPC as mediating memory-guided decision making using serial information flow in long chains of neurons.

Caution must be made when interpreting the result of these inactivation studies, particularly in PPC. As PPC shares a border with visual areas, micro-stimulation or pharmacological/optogenetic inactivation may spread into these areas. If a study finds a positive result for PPC manipulation affecting behaviour, this may be entirely explained through unintended effects in visual processing. A further consideration is the role that inter-hemispheric communication can play in compensating for any unilateral perturbations of neural activity. Li et al. (2016) performed unilateral inactivation in mouse ALM during a tactile whisker-pole detection task. They show that performance in the task fully recovers after terminating unilateral ALM inactivation. However bilateral inactivation causes disruption for the rest of the trial even after inactivation is removed. When examining the spiking of ALM neurons, they find that neural firing recovers back to its pre-inactivated state for unilateral but not bilateral inactivation. Importantly they show that this process is dependent on signals from the other hemisphere via the corpus collosum, as a colossal bisection abolishes this compensation effect. This study demonstrates a broader point that unilateral inactivation of any cortical site may produce compensatory effects which may obscure interpretation.

A few studies suggest that PPC may represent multiple task features simultaneously, suggesting a more complex role for PPC than the aforementioned evidence-accumulation models. Raposo et al. (2014) recorded from rat PPC during a multi-sensory discrimination task and found a mixture of neural tuning to features of the task: stimulus modality and choice. Neurons which typically would be labelled as selective for a task feature were really the handful of cells in the tail end of a population distribution centred at zero selectivity. Furthermore, the selectivity to a mixture of both task features among the cell population was consistent with random. One interpretation of this is that PPC doesn't represent any specific factor in the task. However, the authors showed with modelling that different task features could be read-out from the population by taking different linear weightings of the population. Therefore, PPC seems to be capable of representing multiple task variables simultaneously, and this therefore makes cell-averaging misleading (as was done in many of the tasks mentioned earlier). Similar work in monkeys show individual LIP cells response to multiplexed features of the trial, not single features (Park et al., 2014). Both studies emphasise the need to shift the focus away from identifying cell representation of specific task

variables, instead studying how downstream targets might de-multiplex these quantities to support flexible behaviour.

What about other sensory modalities? In secondary somatosensory cortex (S2), Romo et al. (2002) recorded from neurons in monkeys performing the vibro-tactile discrimination task. The authors found that pre-movement neural activity reflected the comparison of the mechanical stimulus frequencies and was in line with the subsequent choice. This pointed to a possible neural substrate for the decision process, akin to LIP for the random dot task. However, since the stimulus is not continuous in time, S2 activity here is not thought to reflect a DDM process. Machens et al. (2005) proposed a mechanistic model which accounts for these findings in S2 as emerging from modifications to fixed points of recurrent circuit dynamics.

Parietal and other association cortices therefore show neural correlates expected of signals related to the decision. The multiple functions associated with PPC has led to proposals that PPC forms the seat of general decision making, independent of the sensory modalities present in the task. However, it remains unclear whether the activity causally drives the decision process.

1.4.3 Frontal motor areas: decision formation vs action generation

Sensation, decision and action are often viewed as distinct functions, however a number of studies point to a regime where these functions emerge in parallel. Gold and Shadlen (2000) trained monkeys on the random dot task, using saccades to indicate the direction of motion upwards or downwards. The frontal eye field (FEF) is known to mediate saccadic motor output (Schall and Thompson, 1999). Once the monkeys were trained on the task, the authors electrically stimulated the FEF to induce a saccade towards the right side. When the authors performed the microstimulation early into the random dot stimulus, the artificially-evoked saccade was deviated slightly towards the up/down target as instructed by the visual stimulus. The deviation increased with visual stimulus strength and viewing time prior to the microstimulation. This indicates that very early into the visual stimulus presentation, functional properties of the neural circuitry in motor areas are already adjusting. The stimulus-to-action sequence might therefore not contain an intermediate distinct 'decision' phase. However, this may be more specific to cases where the animal is over-trained, where non-trained decision behaviour may indeed rely on distinct decision and action formation steps.

Further work has shown that “decision” and “motor preparation” may actually be referring to the same thing. Kaufman et al. (2014) trained monkeys to perform complex arm movements instructed by a visual cue. After observing the cue, the monkeys had to wait for a delay period before starting the movement - during which the monkey is presumably deciding and preparing the appropriate action to make. The authors recorded from many neurons in dorsal pre-motor (PMd) and primary motor (M1) cortex during this task. They found that PMd population activity moved through different low-dimensional subspaces for the delay and movement periods. The subspace explored by the population during the delay period corresponded to the null space of the matrix of projections from PMd to M1. In other words, ongoing PMd recurrent dynamics were such that they only affected downstream M1 dynamics during the movement phase. This helps explain the finding that preparatory activity in PMd is as strong as activity during movement, and without any apparent inhibitory gating mechanism it was unclear why PMd preparatory activity doesn't generate movement. Mante et al. (2013) performed a similar analysis exploring monkey population dynamics in PFC during a coloured random-dot discrimination task. The authors found that population activity occupied different subspaces depending on whether the task was to discriminate colour or motion. The shape of these subspaces was such that irrelevant information (e.g. colour information during motion discrimination) did not affect population dynamics concerned with achieving the task. Using simulations of recurrent circuitry, the authors show that the process of integrating evidence, and selecting an appropriate action, can all occur within the same neural population moving through different subspaces. Both studies demonstrate sophisticated computational properties of neural populations. One caveat with this kind of dimensionality-reduction analysis is they often rely on the assumption that the firing rates of individual neurons are linearly combined in the population. This assumption is very unlikely to hold true exactly, especially given non-linearities observed within single cell firing (e.g. Branco and Häusser, (2011)), but useful approximate conclusions could be drawn from this approach nevertheless. These studies show that our views of decision as arising from separate sensation, decision and action stages may be incorrect. Instead these processes seem to emerge holistically, from the repeated interaction of “sensory” and “motor” cortical areas.

Recent rodent work points to frontal-motor areas as mediating the emergence of decision (Svoboda and Li, 2018). The frontal orienting field (FOF) is thought to be

the rat homolog for monkey FEF, based on common thalamic, collicular and oculomotor projections (Barthas and Kwan, 2016). Recent studies reveal a possible role for this structure in decision making. Erlich et al. (2015, 2011) found that muscimol inactivation in FOF induced a strong contralateral neglect in an auditory discrimination task requiring orienting movements. The FOF inactivation effects could be accounted for as the biasing of the output of an evidence accumulator process upstream. This points towards a view of FOF as a structure concerned with taking a continuous DV value and discretising it into a binary value for initiation an appropriate action. This interpretation is corroborated by Hanks et al. (2015) who estimated the ongoing decision variable for each trial from behavioural data. PPC showed smooth tuning to DV, whereas FOF showed abrupt step-like tuning to DV, consistent with the idea that FOF is the output of an accumulator. This however doesn't imply that FOF only plays a role 'after' evidence accumulation. Goard et al. (2016) found that optogenetic FOF inactivation impaired choices even when inactivating early in the stimulus epoch. Similarly, Kopec et al. (2015) showed that inactivating mouse FOF early into the stimulus epoch impaired performance in a memory-guided orienting task. Further emphasising an early role for FOF in decisions, Sul et al. (2011) found they could decode upcoming choice in rat FOF earlier than in any other brain area studied under a free-choice foraging task. In nearby motor areas Guo et al. (2014) found that inactivating anterior lateral motor cortex (ALM) biased choices in a mouse but only when inactivating during a delay/decision phase of a tactile decision task.

These studies demonstrate an early causal role for frontal-motor areas in decision making. This is reminiscent of the finding mentioned earlier, that frontal structures show changes to their functional properties as soon as the visual stimulus begins (Gold and Shadlen, 2000). The early role discussed here is inconsistent with the idea that frontal-motor areas are concerned just with triggering motor actions. Instead, decisions appear to emerge from a constant reverberation/interaction of activity between sensory, parietal and frontal-motor circuitry as soon as the stimulus begins (Cisek and Kalaska, 2010).

1.4.4 A fully distributed cortical process?

This holistic view predicts that decision-related signals should emerge simultaneously in all cortical areas involved. Hernández et al. (2010) recorded from several cortical areas spanning the sensory-parietal-motor axis while

monkeys compared the frequency of two mechanical vibration phrases. The authors found that stimulus-selective signals emerged earliest in sensory areas and later in motor areas. By contrast, decision-selective signals emerged in all areas (except S1) simultaneously. This is consistent with the idea that sensory information passes feedforward up a cortical hierarchy, and then decisions arise from reverberation along this hierarchy. A similar observation was made in monkeys performing visual pattern discrimination (Ledberg et al., 2007).

In partial agreement with this, recent work has shown decision signals to arise simultaneously in some regions, but at later times in other regions. Siegel et al. (2015) recorded from multiple cortical areas in monkeys performing a variant of random dot task where monkeys were cued to discriminate dot motion or dot colour. The authors could decode multiple aspects of the task in several cortical areas: cue identity, whether the monkey should discriminate motion or colour, the direction and colour of the dots, as well as the upcoming choice. Decoding of stimulus colour and motion emerged first in MT/V4, then IT/LIP, followed by FEF then PFC. By contrast, choice decoding followed a reverse order - first emerging simultaneously in LIP and PFC, and then following a serial order to FEF and MT/V4. This finding suggests that choice signals arise first in a fronto-parietal network, which are then relayed down to early sensory areas. This accounts for some studies claiming that decoding choice in early sensory areas is due to decision formation in these areas, instead this can be explained as arising from top-down projections from other areas. This point is made in recent papers exploring the origin of choice decoding in early sensory areas (Nienborg and Cumming, 2009; Yang et al., 2016).

Any observation of serial information flow could be viewed as a challenge to the 'holistic emergence' view discussed earlier. Nevertheless, the process of forming a decision and committing to an action involves a host of cortical areas engaging in feedforward and feedback signals.

1.4.5 Subcortical structures

The emphasis on cortex has largely been due to the ease of access. However, it's possible that cortical manipulation effects on choice occur because of downstream effects in subcortical targets which form the real basis for decisions. For example, the striatum has been implicated in valuation of choices in decision making (Schultz et al., 1997), and this structure has vast bi-directional connections with cortex. Znamenskiy and Zador (2013) investigated whether they

could induce behavioural effects on auditory discrimination by just targeting those cortical cells which project to the striatum. The authors injected rat striatum with a retrograde transport virus which conditionally expressed opsins ChR2 or Arch in auditory cortical cells with projections to striatum. They found that optogenetic excitation and inhibition of these cells respectively modulated choice bias towards and away from choices associated with the cortical tonotopy. It's therefore possible that the effects of parietal and frontal inactivation on choice biasing might be explained by cortico-striatal projections (Hintiryan et al., 2016).

The organisation of striatal circuitry seems well-suited to decision making (Wickens et al., 2007). Cortex converges to striatum medium spiny neurons (MSNs), which themselves receive dopaminergic input driving plasticity. MSNs also laterally inhibit each other, facilitating winner-take-all dynamics. In this view, MSNs encode value associated with each action, lateral inhibition induces action selection among alternatives, and dopaminergic encoding of reward prediction error adjusts value estimates depending on the reward history.

Superior colliculus (SC) also appears to play a role in perceptual decisions. Zénon and Krauzlis (2012) trained monkeys to detect changes to moving dot stimuli in specific regions of the visual field. They showed that inactivating SC impaired performance in this task but only to stimuli within the SC's visual field. The authors conclude that SC plays a specific role in this task, as visual cortex remains unaffected by the perturbation. In rodents, the story is similar. Kopec et al. (2015) found that inactivating SC in rats impaired performance on a memory-guided orienting task.

1.5 Conclusion and research questions

To conclude: we have several methods at our disposal for investigating the neural basis of decision making. With well-controlled behavioural tasks, we can measure perceptual features of an animal and link this to neural activity using recording and manipulation. Several papers have asserted a role for PPC in mediating decisions, however neural manipulation rules out PPC as being causally involved in some tasks. Choice decoding has been found in many cortical and subcortical areas, but it's still unclear whether decisions arise from single cortical sites and propagate elsewhere or arise from the interaction of distributed parts of the brain. Theoretical models of decision making have shown great promise in

explaining behaviour, however work still needs to be done to develop mechanistic models of neural circuitry underlying the process.

Therefore, the research questions defining my project are centred on investigating outstanding problems in the field:

- 1. Does decision making arise from activation of distinct cortical regions in a temporal sequence, or instead arise in a distributed fashion?**
- 2. Which cortical areas are causally necessary for the decision process?**
- 3. What are the functional contributions of cortical areas towards decision making?**

CHAPTER 2 A PHENOMENOLOGICAL MODEL OF MOUSE VISUAL DISCRIMINATION

This work contributed towards a publication,

Burgess, C.P.* , Lak, A.* , Steinmetz, N.A.* , **Zatka-Haas, P.*** , Bai Reddy, C., Jacobs, E.A.K., Linden, J.F., Paton, J.J., Ranson, A., Schröder, S., Soares, S., Wells, M.J., Wool, L.E., Harris, K.D., Carandini, M., 2017. High-Yield Methods for Accurate Two-Alternative Visual Psychophysics in Head-Fixed Mice. *Cell Reports* 20, 2513-2524. <https://doi.org/10.1016/j.celrep.2017.08.047>

*These authors contributed equally

2.1 Introduction

Perception is a fundamental feature of conscious experience. An open question in neuroscience is what the neural processes are which underlie perceptual experience. Answering this question requires identifying relations between perceptual states and neural activity, either through correlations or by establishing a causal relationship. Perceptual states are themselves unobservable, however by training perception-action associations in a behaving subject, we can infer perceptual states by observing the subject's behaviour.

Perceptual decision making a convenient behaviour to study because it is a relatively simple cognitive task with a long history in psychophysics (Fechner, 1860). Animals can be trained with rewards to perform specific actions in response to experimentally-controlled stimuli. By adjusting the difficulty of the task, we can observe how behaviour varies over this range of difficulty, thereby inferring properties of the underlying perceptual process. These tasks provide a foundation to further studying the neural processes involved.

Several features are desirable for a task design. Firstly, for practical purposes, tasks should be relatively easy to train, provide many trials per session, and be easily paired with neural recording and manipulation. These properties make mice an attractive model to study (Carandini and Churchland, 2013). Mice can be trained to perform a wide range of tasks involving olfaction, whisker deflection, audition and vision (reviewed in Hanks and Summerfield, (2017)).

Secondly, the task should be difficult enough that choices vary substantially during identical stimulus conditions. This is important to disentangle the neural signatures associated with the stimulus and associated with choice (Parker and Newsome, 1998). This requirement is fulfilled in tasks where the stimulus varies

by a continuously graded quantity (e.g. choose the stimulus of higher visual contrast, or sound intensity) and not by a discrete quality (e.g. choose between an image of cheese or a snake). Sensory detection and discrimination tasks and therefore well-suited for this purpose.

Thirdly, the type of choice action performed by the animal should sufficiently constrain inferred states of the perceptual system (Macmillan and Creelman, 2004; Stüttgen et al., 2011; Tanner and Swets, 1954). Example states are stimulus sensitivity, choice bias, or lapse rate (Busse et al., 2011; Gold and Ding, 2013). Go-NoGo tasks are unsuitable for this because non-perceptual choice biases can arise from the unbalanced motor effort required for each choice, as well as variable engagement in the task. Forced-choice paradigms improve on this by requiring active movement on every trial, therefore every choice has similar motor cost and requires constant engagement. However, forced-choice paradigms have two problems. Firstly, they confound decision bias with uncertainty bias. Decision bias arises when subjects choose one choice over another because their perception of the stimulus may be biased. Uncertainty bias, by contrast, arises when subjects opt for one “default” choice whenever they are uncertain of the correct option (García-Pérez and Alcalá-Quintana, 2013). Forced-choice paradigms also confound two types of bias induced by a neural manipulation. For example, if neural manipulation induces a choice bias in a 2AFC regime, this could arise from either biasing the detection of one stimulus alone, or from biasing the process of comparing between both stimuli. One task paradigm which addresses these problems is the two-alternative unforced-choice design. The design is similar to the forced-choice paradigm except subjects can abstain from selecting an option (NoGo) if they believe no option is correct. Previous work has shown that it is possible to distinguish choice bias and stimulus sensitivities in this task design if the task involves sensory discrimination and not sensory detection (Sridharan et al., 2014).

A further concern is in the quantification of perceptual states like bias and sensitivity. Estimating these quantities requires fitting a model to behavioural data, and the estimates may differ session to session and subject to subject. It is unclear whether this variation arises because of sampling noise, or because of true variation in the states across sessions and subjects. Hierarchical Bayesian models make it possible to estimate the variation associated with sampling noise,

and the variation due to real differences between sessions and subjects (Gelman et al., 2013; Lee, 2011; McElreath, 2018).

In this Chapter we will outline a two-alternative unforced visual discrimination task in mice. This task develops upon a previous forced-choice version of the task used in the Carandini/Harris lab, and forms part of a publication in Burgess et al (2017). We will then describe work done towards developing a phenomenological choice model which could describe the behaviour in terms of perceptual features such as stimulus sensitivity and choice bias. We will first outline a simple form of the model based on multinomial logistic regression, which extends the binary-choice logistic model of Busse et al. (2011) to tasks which permit multiple decision options. We will then expand this model to provide estimates for behavioural variation between subjects and sessions. This behavioural task and model form a foundation for the rest of this dissertation, as the task and models are used repeatedly in subsequent Chapters.

2.2 Methods

2.2.1 Ethics

The experimental procedures contained within this Chapter and the rest of the dissertation were conducted at UCL according to the UK Animals Scientific Procedures Act (1986) and under personal and project licenses granted by the Home Office following appropriate ethics review.

2.2.2 Surgery

Mice were implanted with a head-fixation plate and clear skull cap similar to that of Guo et al. (2014) and described previously in Burgess et al (2017). This permitted optical access to the neocortex through the skull useful for later imaging and inactivation experiments. The implantation surgery proceeded as follows. The dorsal surface of the skull was cleared of skin and periosteum, and the junction between cut skin and skull was sealed with cyanoacrylate. The exposed skull was prepared with a brief application of green activator to ensure strong connection between cement and bone (Super-Bond C&B, Sun Medical Co, Ltd, Japan). The junction between skin and skull was again covered, using dental cement (Super-Bond C&B). In most cases, a 3D printed 'cone' was attached to the head with cyanoacrylate and dental cement at this stage, surrounding the exposed skull and providing light isolation useful for later imaging/inactivation experiments. A thin layer of cyanoacrylate was applied to the skull and allowed to

dry. Two to four thin layers of UV-curing optical glue (Norland Optical Adhesives #81, Norland Products Inc., Cranbury, NJ; from ThorLabs) were applied to the skull and cured (~10 s per layer) until the exposed skull was covered (thin layers were used to prevent excessive heat production). A head-plate was attached to the skull over the interparietal bone with SuperBond polymer.

2.2.3 Water control

Mice were water deprived and their weight was kept no less than 80% of their pre-water-deprivation weight. After each day of performance on the behavioural task, mice were given surplus water (hydrogel) to meet the weight-dependent minimum requirements each day of 40 μ l/mg/day. If mice showed any signs of dehydration or ill-health, they were given water *ad libitum* and re-entered into water deprivation several days later.

2.3 A behavioural task in mice

Water-deprived mice were trained to perform a visual discrimination task. The task requires mice to discriminate between two visual grating stimuli of varying contrast levels and indicate which side has the higher contrast, or to abstain from moving if no stimulus is present.

Mice were head-fixed with their forepaws resting on top of a wheel which could be rotated leftwards or rightwards (Figure 2-1; Figure 2-2A). Surrounding the mouse were three screens positioned to the left, front and right side, and plastic tubing was positioned near the mouth for delivering water rewards. Trials began after 0.2-0.6sec of no wheel movement. A visual stimulus was then presented, comprising two visual gratings spanning 30 degrees (0.1 cycles per degree, oriented at 45 degrees) on the left and right screens. The left and right stimuli were located 180 degrees (azimuth) apart and therefore were positioned within the lateral field of each eye. Each grating could take on one of four contrast levels, and therefore both gratings together formed 16 stimulus conditions (Figure 2-2B). On each trial, one of the 16 conditions were presented randomly but with non-uniform probability. Trials with zero contrast on both sides were over-represented, making up ~30% of trials. Trials with equal contrast on both sides made up ~10% of trials. All other stimulus conditions were presented uniformly among the remaining 60% of trials. Onset of the visual stimulus also coincided with the onset of an auditory 'go cue' (12 kHz tone, 100 msec duration), marking the time at which the mouse could start responding.




Figure 2-1 The behavioural setup

(Left) Head-fixation mount. Mice are head-fixed with the forepaws on a wheel. A plastic half-cylinder is placed on top to further restrain mice and prevent the tail from moving near the head during neural recording/manipulation. In red is an articulated arm which holds thin plastic tubing to position near the mouse's mouth for water delivery.

(Right) The head-fixation mount is placed in the centre. Surrounding this mount are three iPad screens displaying the visual stimuli. On each screen is a Fresnel lens to ensure uniform light intensity to the mouse's eyes from all pixels on the screen. A photodiode is mounted facing one of the screens, to measure the timing of the screen frame updates. Image created by Lauren Wool 2018.

Wheel rotation was linked to the visual stimulus, such that wheel turns to the right would move the contents of both monitors the right, and vice versa (Figure 2-2C). Mice were rewarded with 0.7-2.5 μ l water for turning wheel ('Left' or 'Right') to bring the grating stimulus of higher contrast into the central screen, or for abstaining from turning the wheel ('NoGo') on zero contrast trials (Figure 2-2D). Left or Right responses were registered if the position of one grating stimulus entered the central screen within 1.5seconds, otherwise a NoGo was registered (Figure 2-2E). On trials with equal contrast, mice were rewarded randomly for turning Left or Right. If the response was incorrect, a 1 second white-noise sound timeout was played. After the trial ends, there was an inter-trial interval of 1 second.

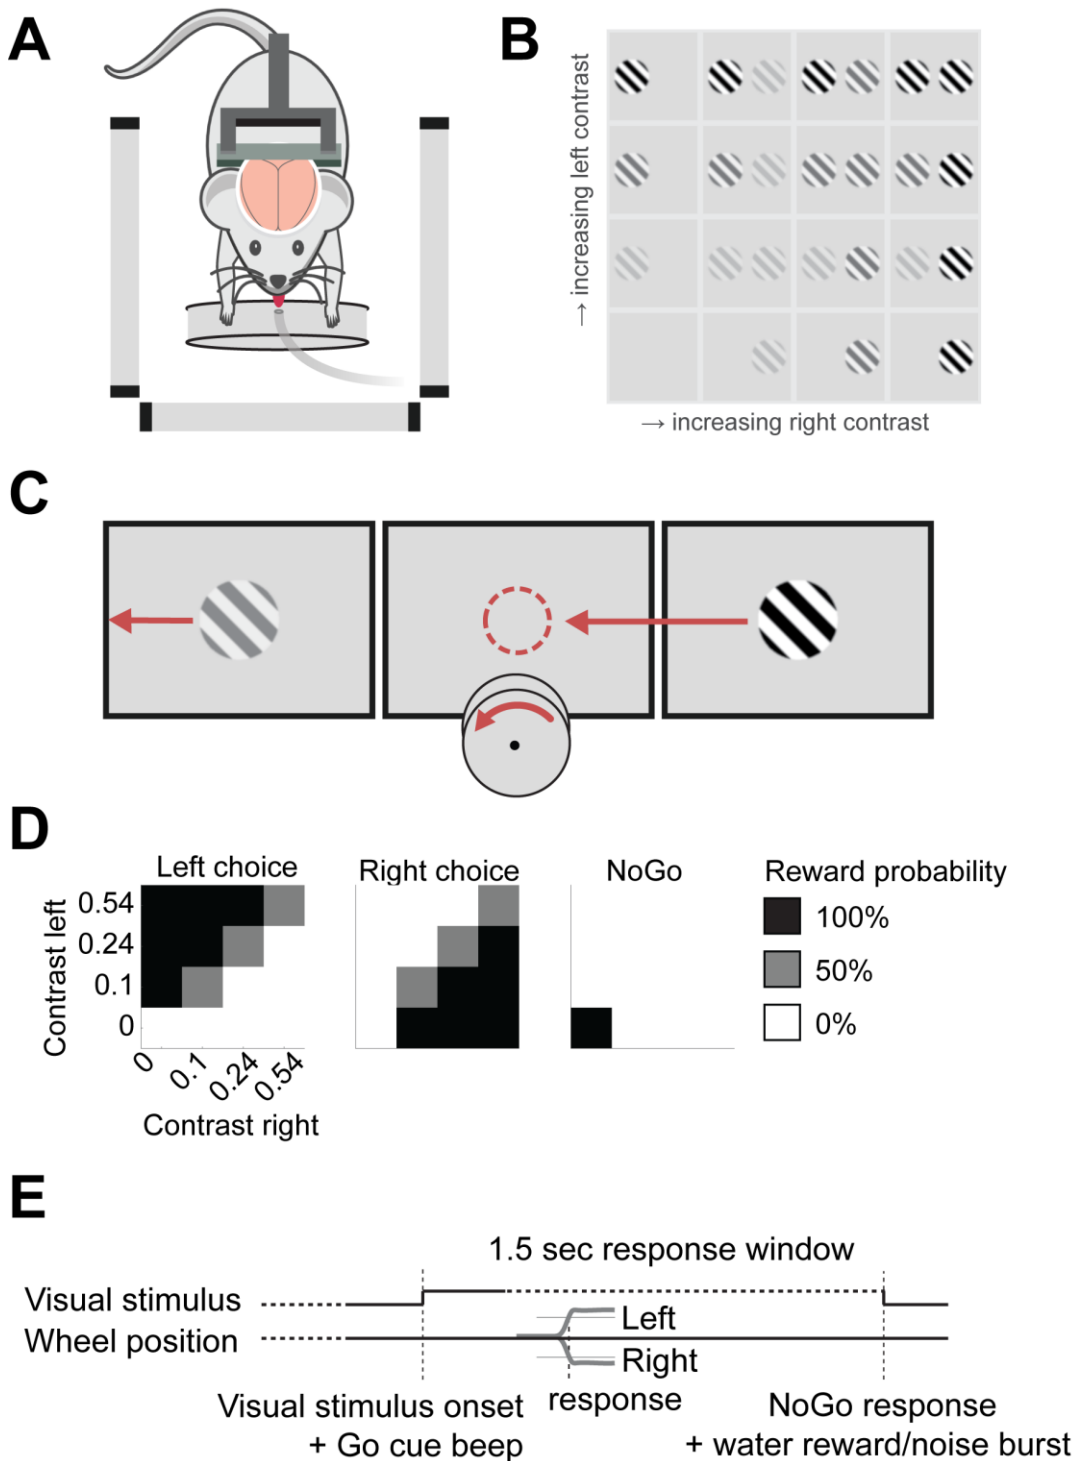


Figure 2-2 Task design

(A) Schematic of the head-fixed mouse. The forepaws rest on a wheel which can turn left and right. Screens display Gabor stimuli on the left and right sides. A water spout delivers 0.7-2.5 μ l water rewards for correct choices.

(B) The full set of contrast conditions displayed to the mouse. Each grey square represents the stimulus configuration spanning both left and right screens.

(C) Movement of the wheel is coupled to movement of the stimuli.

(D) Reward contingency for this task. Mice were rewarded with water for bringing the stimulus of higher contrast into the central screen ("Left" or "Right"). If no stimulus was present, then mice were rewarded for holding the

wheel still (“NoGo”). If the two contrasts were equal, then mice were rewarded randomly for making a Left or Right wheel turn.

(E) Task timeline. Trials began with a period of 0.2-0.6 seconds of no wheel movement. After this, the visual stimulus was presented alongside an auditory go cue. At this point the mice could respond immediately. Left or Right responses were registered if the wheel angle breached a threshold. A 1.5 second response window was used, and if no Left or Right responses was registered in this window then NoGo was registered. At the response time, a water reward or 1sec white noise timeout was given.

If mice made errors on easy Go trials (high contrast on one side only), this was assumed to arise because of confusion about the task rule or disengagement from the task itself. To ensure that mice maintained the task rule as well as stay engaged, easy trials were repeated until performance was correct. If mice mistakenly performed NoGo for more than 10 consecutive repeat trials, the session was stopped, and mice were returned to their home cage. Zero contrast trials were also repeated until mice correctly performed NoGo. This was an additional test to check that mice knew the task rule. In subsequent analyses, trials were excluded if the trial was a repetition trial. Mice were able to learn this task within 4 weeks (See Burgess et al, (2017) for training information).

After training, mice showed robust performance in the task (Figure 2-3, Figure 2-4). On easy trials (high left or right contrast), mice were 96.9% correct on average across 91 sessions. Trial counts for each session varied, ranging from 245 to 405 trials (25th and 75th percentile). Average reaction time for correct Left or Right choices was 258 ± 39 msec (median \pm median absolute deviation across trials pooled over sessions). Reaction time was calculated as the time when wheel movement was first detected, which was prior to the time when a response was registered.

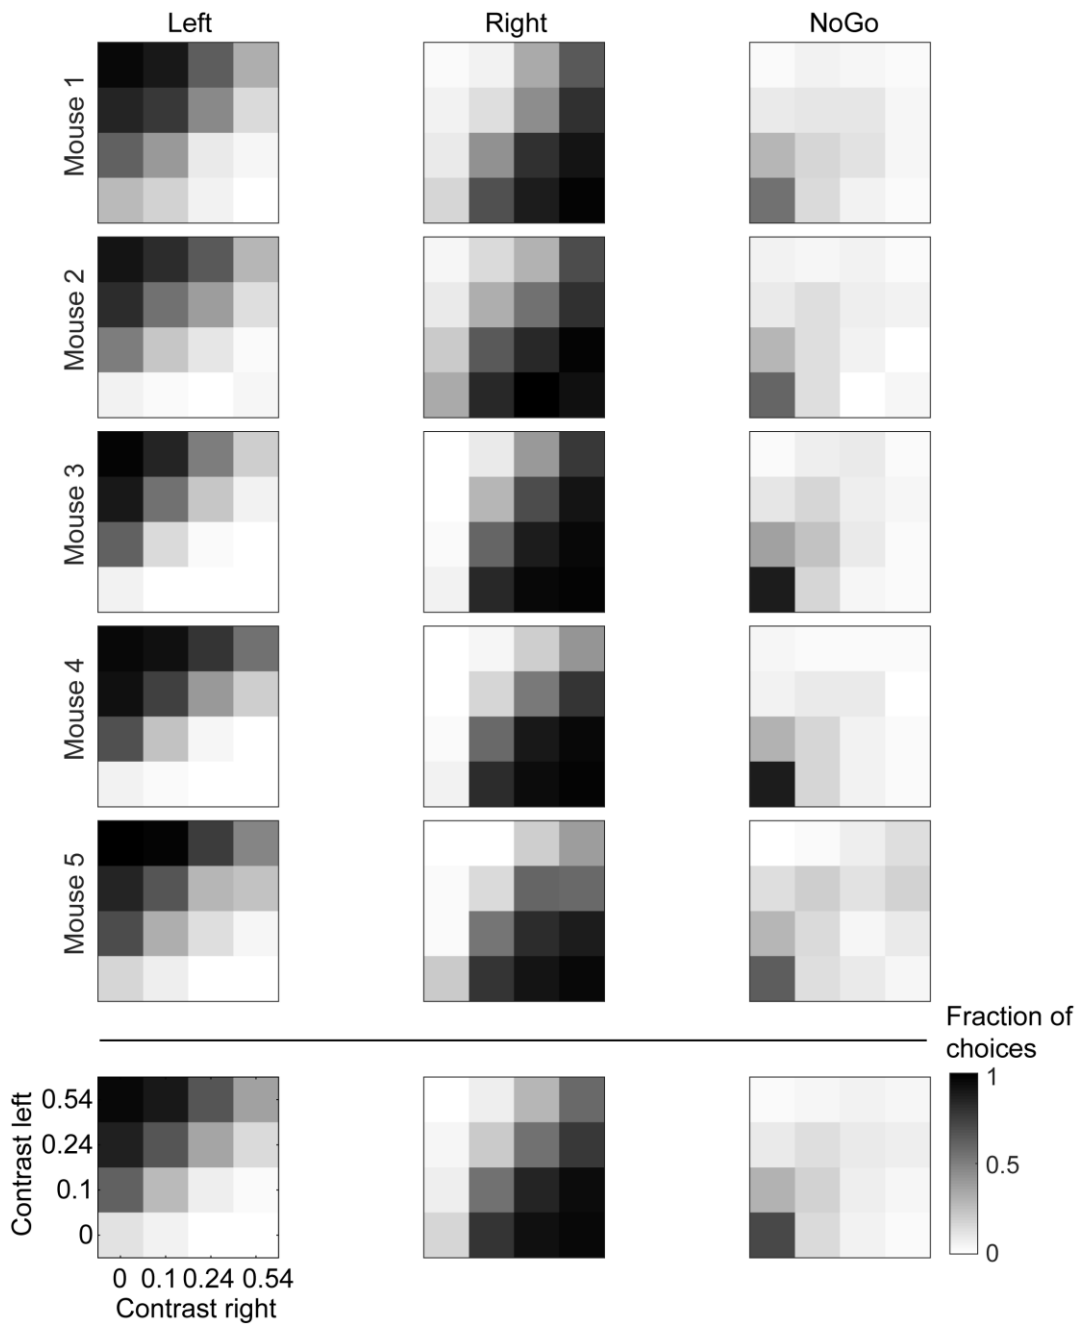


Figure 2-3 Psychometric data for 5 mice

The grey-scale map shows the fraction of Left, Right and NoGo choices as a function of visual contrast on the left (vertical axis) and right (horizontal axis). Each row corresponds to session-averaged data from one mouse, and the last row is the average across mice.

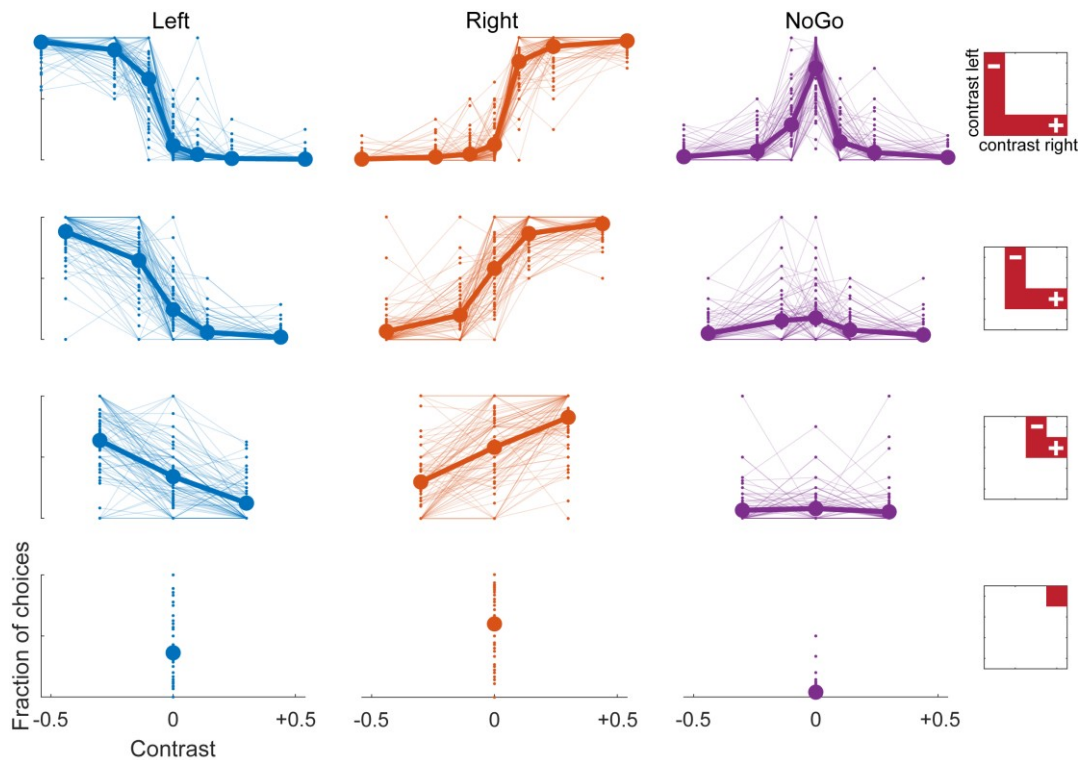


Figure 2-4 *Alternative pedestal representation of the psychometric data*

On each row is plotted the proportion of Left, Right and NoGo choices (vertical axis, spanning 0 to 1) as a function of contrast deviation from a minimum contrast value (horizontal axis). Negative contrast values correspond to stimuli on the left hemifield, and positive contrast values are stimuli on the right hemifield. Top row are the choices on trials where the minimum contrast on either side is 0. Second row is where the minimum contrast is 0.1. Third row is 0.24, and fourth row is 0.54. Large dots and thick connecting lines correspond to the average psychometric data plotted in Figure 2-3, small dots and thin lines are the data for each of 91 sessions in 5 mice.

2.4 Logistic regression model of choice behaviour

We developed a model of choice behaviour in order to summarise behaviour in terms of perceptual features such as choice bias and stimulus sensitivity. Note that the model only makes predictions for the choice the mouse makes, and not the decision time associated with the choice. While decision timing is an important and informative feature of the perceptual process, we chose to ignore reaction time within the model. This was because models accounting for reaction time (e.g. drift-diffusion models) are more applicable to tasks with sensory evidence accumulated over time, rather than static sensory stimuli. In addition, drift-diffusion models for multi-alternative decisions have typically employed

independent accumulators for each choice, but it is unclear how this would apply to tasks with NoGo as these choices have no associated decision time.

The model can be derived from normative economic theory and is equivalent to multinomial logistic regression. On each trial i , each choice (Left, Right, NoGo) has an associated unobserved random utility, u_L, u_R, u_{NG} (Greene, 2011). The expected utility for each choice is set by a fixed variable z_L, z_R, z_{NG} which captures the amount of subjective value of each choice to the deciding agent. For example, in our task, the expected utility of each choice is determined by the visual contrast, as that determines the payoff each choice would provide. There is an additional random component $\varepsilon_L, \varepsilon_R, \varepsilon_{NG}$ which captures all unmodeled features which determine the subjective utility on each trial.

$$\begin{aligned} u_L &= z_L + \varepsilon_L \\ u_R &= z_R + \varepsilon_R \\ u_{NG} &= z_{NG} + \varepsilon_{NG} \end{aligned}$$

This therefore makes the implicit assumption that variation in choices arises from the unmodeled features and not variation in expected utility. This contrasts with signal detection models which model variation in behaviour as noise in the decision variable coupled with a noise-less decision rule. However these models end up being mathematically equivalent (Lynn et al., 2015). Variation in behaviour can be attributed either to variation in perception of the stimulus, or variation in the subjective utility associated with the actions. In this framework, the unmodeled effects $\varepsilon_L, \varepsilon_R, \varepsilon_{NG}$ are given a Type-1 Extreme Value (Gumbel) distribution.

Given the utility associated with each choice, the agent makes the choice with the largest utility. For example, the probability of Left choices π_L can be derived,

$$\begin{aligned} \pi_L &= p(u_L > u_R \text{ and } u_L > u_{NG}) \\ &= p(u_L - u_R > 0 \text{ and } u_L - u_{NG} > 0) \\ &= p(z_L + \varepsilon_L - (z_R + \varepsilon_R) > 0 \text{ and } z_L + \varepsilon_L - (z_{NG} + \varepsilon_{NG}) > 0) \\ &= p(z_L - z_R > \varepsilon_L - \varepsilon_R \text{ and } z_L - z_{NG} > \varepsilon_L - \varepsilon_{NG}) \end{aligned}$$

Therefore, the probability of choosing Left is entirely determined by the difference in expected utility $z_L - z_R$ and $z_L - z_{NG}$. Therefore, adding a constant to each utility will have no effect. Subtracting z_{NG} from each expected utility gives a new utility,

$$Z_L = z_L - z_{NG}$$

$$Z_R = z_R - z_{NG}$$

These two variables Z_L and Z_R are interpreted as the expected utility of Left relative NoGo and Right relative to NoGo. Additionally, since the unmodeled effects ε are distributed with a Type-1 Extreme Value distribution, their differences $\varepsilon_L - \varepsilon_R$ and $\varepsilon_L - \varepsilon_{NG}$ follow a logistic distribution, making this process equivalent to logistic regression.

Therefore, it can be shown that the probability associated with each choice is derived from the softmax function of the expected utilities Z_L and Z_R ,

$$\pi_L = \frac{e^{Z_L}}{1 + e^{Z_L} + e^{Z_R}}$$

$$\pi_R = \frac{e^{Z_R}}{1 + e^{Z_L} + e^{Z_R}}$$

$$\pi_{NG} = \frac{1}{1 + e^{Z_L} + e^{Z_R}}$$

Rearranging the expressions shows that Z_L is equivalent to $\ln\left(\frac{\pi_L}{\pi_{NG}}\right)$ and Z_R is equivalent to $\ln\left(\frac{\pi_R}{\pi_{NG}}\right)$. This is equivalent to two independent binary logistic regression models for Left vs. NoGo and Right vs. NoGo. This model therefore assumes that the two log odds ratios are independent. This assumption is also known as the ‘independence of irrelevant alternatives’ assumption, and it assumes that the presence or absence of an extra choice does not change the relative pairwise preferences between all other choices (Luce, 1959). For example, the ratio of preference between Left and NoGo is unaltered by adding Right choices as an alternative.

This model framework can be used to fit arbitrary regressors to predict the behavioural choice on each trial (Figure 2-5). A simple model form is,

$$Z_L = b_L + s_L \times f(c_L)$$

$$Z_R = b_R + s_R \times f(c_R)$$

Where c_L and c_R are the visual contrast on the left and right hemifields. The contrast is passed through a chosen function $f(\cdot)$ and then multiplied by a free parameter s_L and s_R which scales the contrast into Z-space (log odds space). Another free parameter b_L and b_R adds to this value. With this simple model form, the s_L and s_R parameters capture how much the visual stimulus impacts the log

odds ratios. These parameters therefore are referred to as ‘sensitivity’ parameters because they capture the effect of the perceptual stimulus sensitivity on behaviour. The parameters b_L and b_R capture the choice preference independent of the stimulus, and therefore are referred to as ‘bias’ parameters.

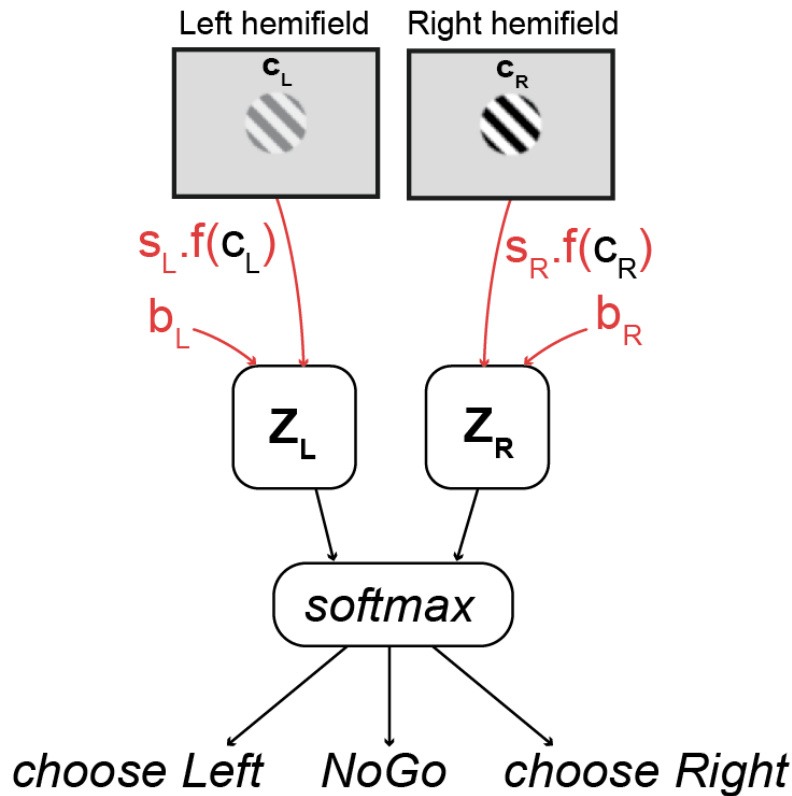


Figure 2-5 Schematic of the phenomenological model

The choices Left, Right, NoGo are treated as draws from a probability distribution over these choices. The probability mass for each choice is computed from the softmax computation of two decision variables Z_L and Z_R . Z_L represents the log odds of Left vs. NoGo, Z_R represents the log odds of Right vs. NoGo. Each decision variable is parameterised with a bias parameter b_L and b_R , and a sensitivity parameter s_L and s_R which multiplies the visual contrast on the appropriate hemifield. The effect of the visual contrast on the decision variables can be modelled with different functions with different shape parameters.

Our first objective was to determine which function $f(\cdot)$ may be suitable. We assessed different kinds of functional representation. We tested three different models (Table 2-1) which differed in the contrast transform f : linear $f(c) = c$, saturating non-linear $f_n(c) = c^n$, and Naka-Rushton (NR) saturating non-linear $f_{n,c_{50}}(c) = \frac{c^n}{c^n + c_{50}^n}$. The non-linear transformations contain shape parameters, in addition to the standard bias and sensitivity parameters. The Naka-Rushton

transformation has been used before to model the dependence of retinal (Evans et al., 1993) and V1 (Boynton et al., 1999) firing rate on visual contrast. The NR transformation is parameterised by two shape parameters: n and c_{50} , which control the slope and position of the input contrast c required to achieve $f(c) = 0.5$ transformed contrast.

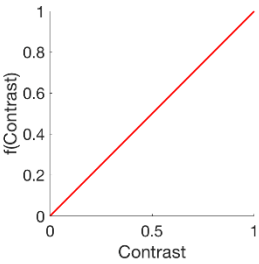
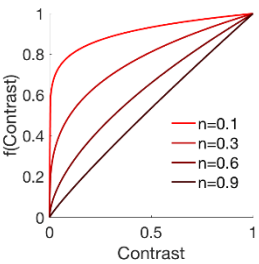
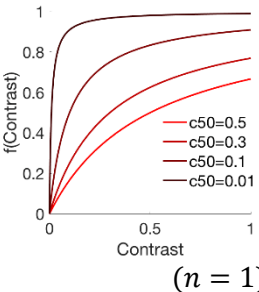
Model name	Model form	Contrast transformation	Parameter constraints
B	$Z_L = b_L$ $Z_R = b_R$	N/A	b_L $[-\infty, \infty]$ b_R $[-\infty, \infty]$
B & C	$Z_L = b_L + s_L \cdot c_L$ $Z_R = b_R + s_R \cdot c_R$		b_L $[-\infty, \infty]$ b_R $[-\infty, \infty]$ s_L $[0, \infty]$ s_R $[0, \infty]$
B & Cⁿ	$Z_L = b_L + s_L \cdot c_L^n$ $Z_R = b_R + s_R \cdot c_R^n$		b_L $[-\infty, \infty]$ b_R $[-\infty, \infty]$ s_L $[0, \infty]$ s_R $[0, \infty]$ n $[0, 1]$
B & NR	$Z_L = b_L + s_L \cdot \frac{c_L^n}{c_L^n + c_{50}^n}$ $Z_R = b_R + s_R \cdot \frac{c_R^n}{c_R^n + c_{50}^n}$		b_L $[-\infty, \infty]$ b_R $[-\infty, \infty]$ s_L $[0, \infty]$ s_R $[0, \infty]$ n $[0.3, 20]$ c_{50} $[0.001, 3]$

Table 2-1 Behavioural models which differ in the function applied to the visual contrast. The non-linear functions (rows 3 and 4) are parameterised with shape parameters.

The three models captured the primary characteristics of the psychometric data (Figure 2-6). The inclusion of a saturating non-linear function of contrast improves the fit and this can be seen by eye. However, the two non-linear models make very similar predictions and are difficult to distinguish.

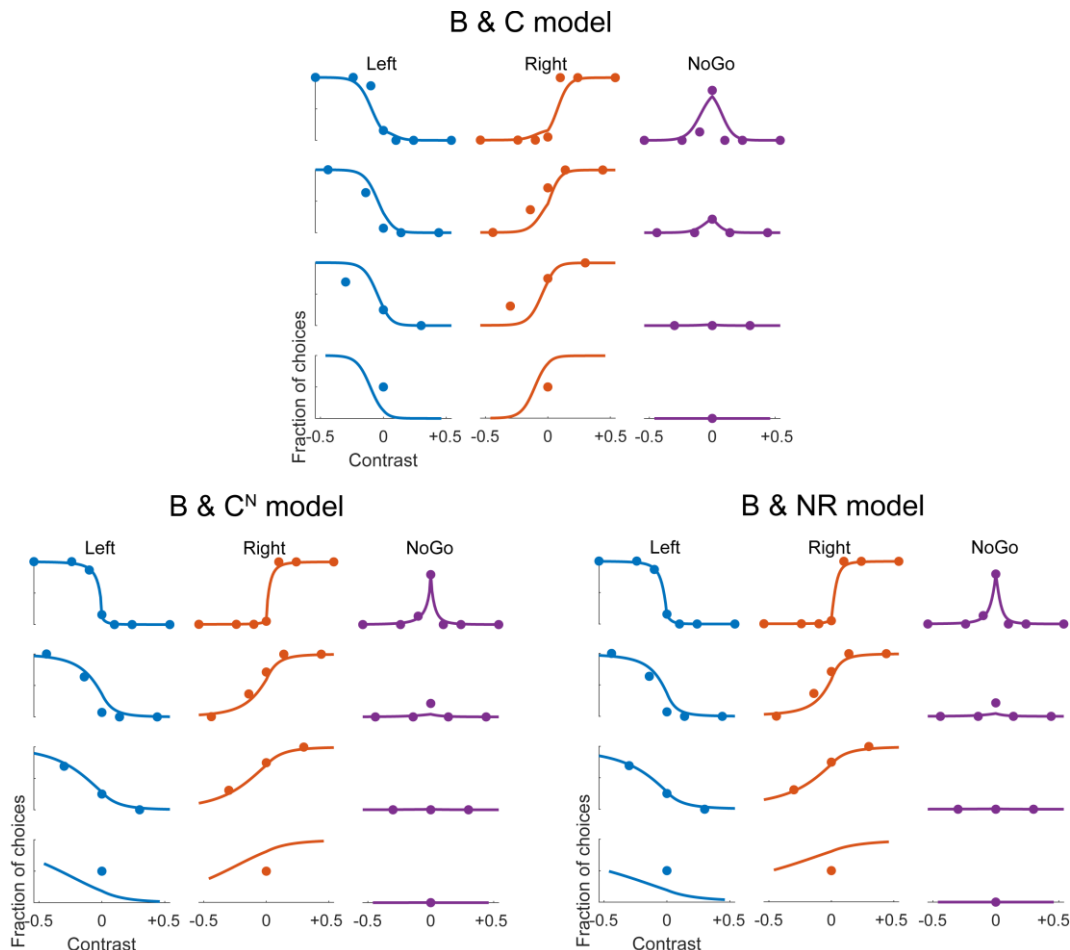


Figure 2-6 Example fits of the three phenomenological models

The contrast conditions are separated according to the minimum contrast on either side (pedestal) as in Figure 2-4. The dots are empirical fraction of Left, Right and NoGo choices for one session. The solid line is the prediction from the model fit.

To select the best model, we used two strategies. The first is out-of-sample prediction. Many models are able to overfit the data, which occurs when a model is capturing characteristic features of the data-generating process but also the noise. These types of models produce excellent fits to the data the model is fit to but will generalise poorly to other datasets which have the same ‘signal’ but different ‘noise’. To avoid this, the performance of a model can be tested on data which was not used to fit the model. Out-of-sample prediction achieves this

through several means: cross-validation and information criteria. Cross-validation involves fitting the candidate models to a subset of the data and testing them against held-out data. This process is repeated for different folds of the dataset, such that the model is eventually tested on every datapoint (but never simultaneously trained on it). This works very well for many models but becomes computationally intractable for models which take a long time to fit. For models such as these, another method is to utilise information criteria. These are quantities derived from information theory which estimate the out-of-sample prediction performance. One example information criteria which will be used later in this Chapter is the Widely Applicable Information Criterion (WAIC; Vehtari et al.(2017)). A general philosophical note about cross-validation and information criteria is that they assess the ability for models to account for the data in a phenomenological sense. They do not assess anything about the models' mechanistic explanatory power (Churchland and Kiani, 2016).

The second strategy is Occam's razor. This strategy can be applied to select among models which show an equal ability to fit the data (e.g. quantified by out-of-sample prediction) but which vary in the number of free parameters. Applying Occam's razor in this situation amounts to selecting the model with the fewest number of parameters.

We assessed the ability for these models to fit data on a held-out test set by using 20-fold cross-validation. We found that the two non-linear models (B & C^N and B & NR) offer improved fitting over the linear model (B & C) for most mice (Figure 2-7). However, we find that the two non-linear models were equally able to capture the dataset. Given that the B & C^N model contains fewer parameters, we selected this model for subsequent fitting in all mice.

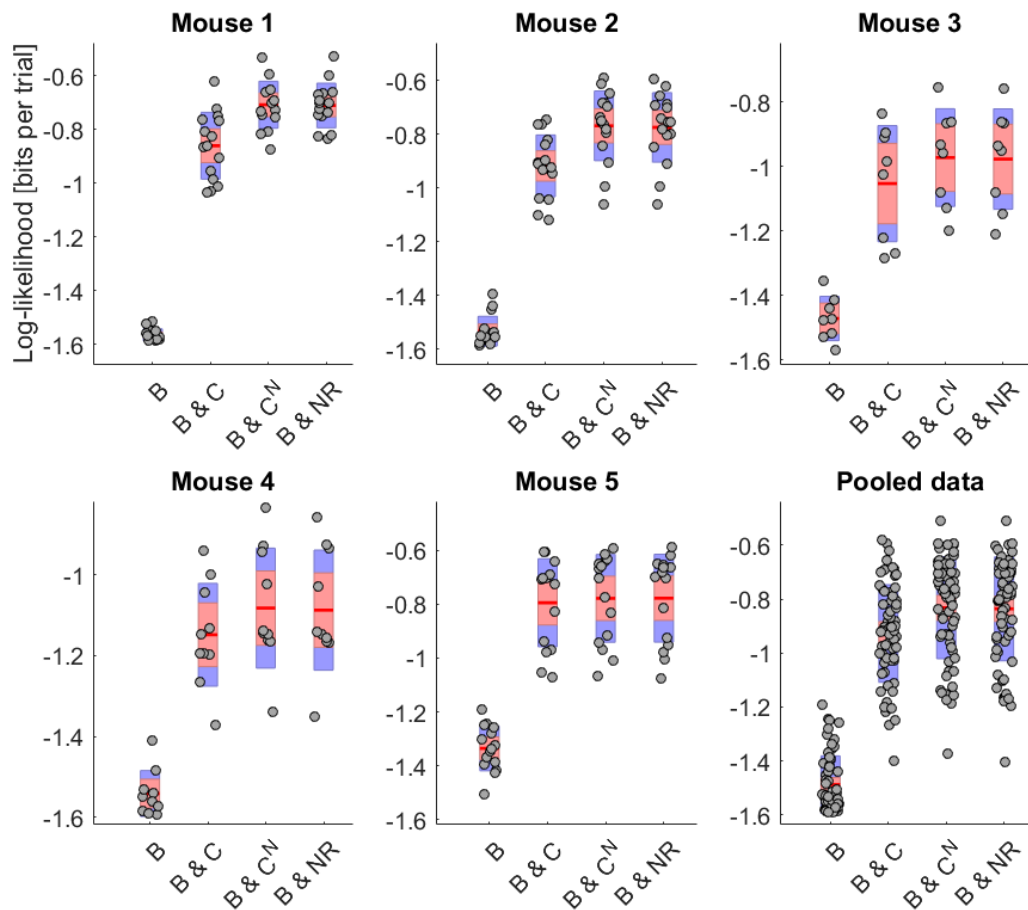


Figure 2-7 Model comparison of different models of contrast non-linearity

20-fold cross-validated model comparison of different contrast transformation for data from 5 mice. Log likelihood on test-set data is plotted for each mouse and each session individually. Each dot is one session. Red lines are the mean log likelihood across sessions. The red shaded region is the 95% confidence interval, and the blue shaded region is the standard deviation.

Another way to visualise the model fit is in the space of the proposed decision variables Z_L and Z_R (Figure 2-8). In this space the model's assumption of the independence of irrelevant alternatives is apparent. In the averaged data the IIA assumption is not strictly held, however this is more the case for stimulus conditions at extreme values of log odds ratios, which corresponds to probabilities around 0.99, and therefore uncertainty (via the limited number of samples) of probabilities in this area will translate to higher noise in the Z-space values.

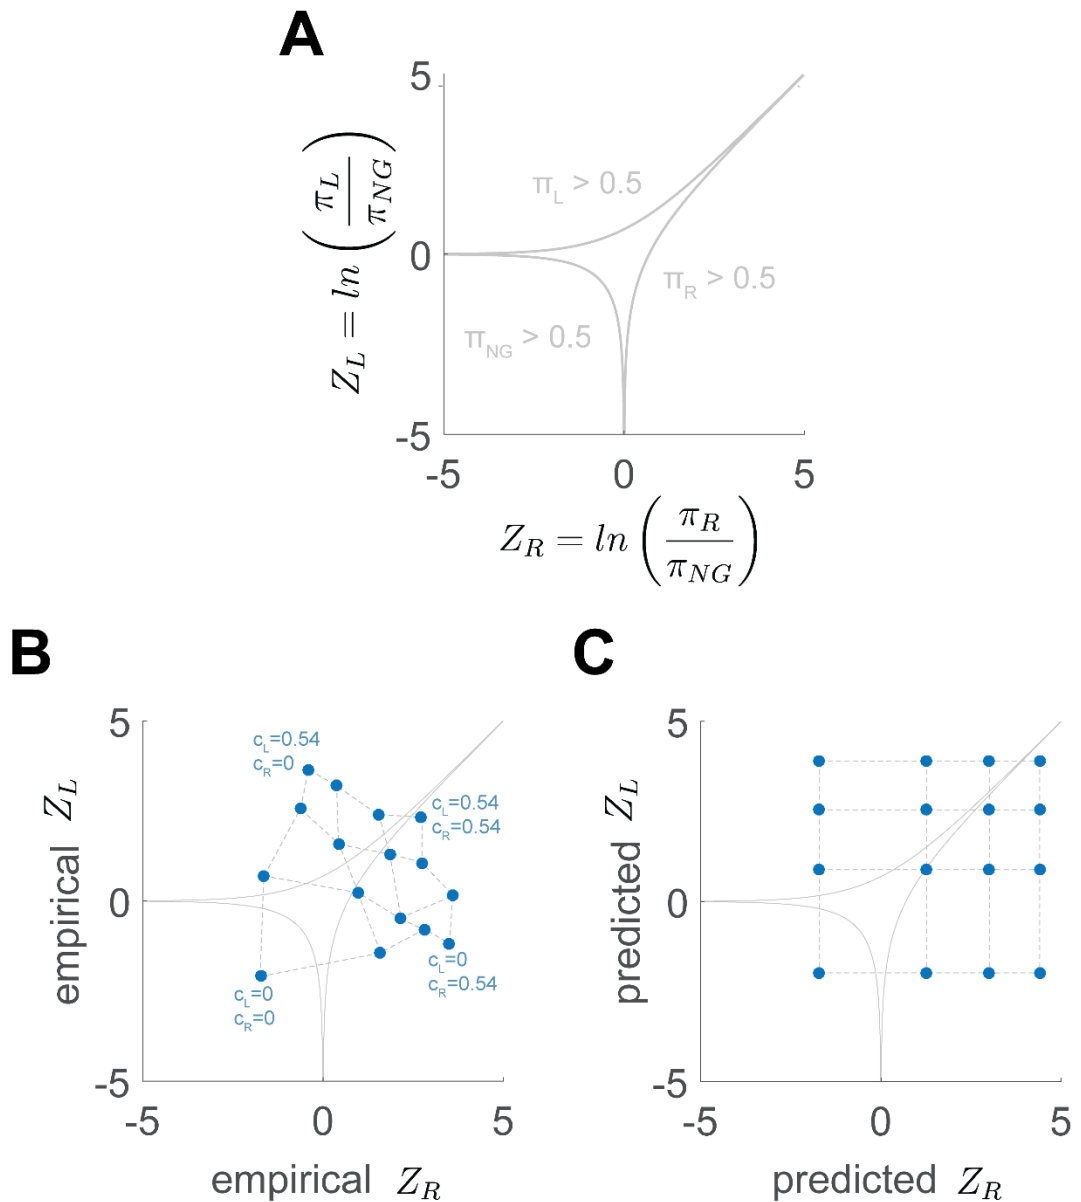


Figure 2-8 Z-space plot of the average data and model fit

(A) Schematic of the Z_L and Z_R space which is equivalent to the log odds ratio of Left vs. NoGo and Right vs. NoGo respectively. Grey lines mark the contours where the probability of Left, Right and NoGo is greater than 0.5.

(B) Empirical data (averaged across sessions) is plotted for each of the 16 stimulus conditions separately in blue. The contrast values span a 4x4 grid which covers this space. Contrast values are indicated inset.

(C) Prediction of the fit B & C^N logistic model for each of the 16 stimulus conditions. The square shape of the 4x4 grid in this space reflects the independence of irrelevant alternatives assumption present in the model.

2.5 A multi-level framework estimates behavioural variation

Behavioural data is acquired from several mice over many different sessions. Variation in the choice behaviour across sessions and mice may arise either from true variation in the bias, sensitivity and shape parameters, or due to sampling noise resulting from the fact that we have limited data for each session. In this section, we modified the logistic model outlined previously to provide an estimate of true behavioural variability, by estimating variability in the model parameters between sessions and mice.

To determine whether the parameters are truly varying between sessions, one strategy is to fit the logistic model to each session separately. Behavioural variation would be reflected as variation in the parameter fits. However, this strategy is problematic for two reasons. Firstly, the model may overfit the data of a single session. This will artificially increase the variation in the parameters between sessions, which could erroneously be interpreted as variation in perceptual features between sessions. Secondly, fitting the model to each session individually throws away information contained in the fact that the parameters for each session are likely to be similar for sessions from the same mouse.

We therefore employed a hierarchical Bayesian framework (Figure 2-9) which addresses these problems. These models define a prior probability distribution associated with each parameter. Bayes rule defines a procedure for obtaining the posterior probability distribution of the parameters, conditioned on the observed data. This approach therefore provides a measure of uncertainty associated with the parameter estimates. In addition, the prior distribution on the parameters associated with each session can be defined hierarchically. For example, the prior can be set such that the parameters for multiple sessions for each mouse are related to a common set of parameters specific for that mouse. Likewise, the parameters for multiple mice can be related to a common parameter for the average mouse.

Hierarchical Bayesian models can therefore provide a robust measure of behavioural variation. Similarly, since hierarchical models define a relationship between sessions, this modelling framework improves the stability of a fit to a single session. This arises because the parameter fit for a single session is also 'informed' by the fit for other sessions from the same mouse. This framework therefore strikes a balance between quantifying variation in parameters across

sessions and regularising the parameters such that they don't vary too much between sessions.




Figure 2-9 Probabilistic Graphical representation of the hierarchical model

White circles represent unobserved quantities whose distribution is inferred from the data. Large grey shaded regions reflect nesting of the model (i.e. trials are nested within sessions within subjects) which reflect the hierarchical dependence placed on the model parameters (reflected in the arrows). Grey squares are the known contrast on each trial, and the grey circle represents the choice on that trial.

For each trial i , choices $y^{(i)} \in \{Left, Right, NoGo\}$ are treated as draws from a probability distribution over choices Left, Right, NoGo, each with probabilities π ,

$$y^{(i)} \sim \text{Categorical}(\pi_L^{(i)}, \pi_R^{(i)}, \pi_{NG}^{(i)})$$

The probability mass associated with each choice is computed from the softmax function of two decision variables,

$$(\pi_L^{(i)}, \pi_R^{(i)}, \pi_{NG}^{(i)}) = \text{softmax}(Z_L^{(i)}, Z_R^{(i)})$$

Given this definition, the two decision variables are equivalent to the following log odds ratios as outlined in a previous section (Greene, 2011),

$$Z_L^{(i)} = \ln \left(\frac{\pi_L^{(i)}}{\pi_{NG}^{(i)}} \right)$$

$$Z_R^{(i)} = \ln \left(\frac{\pi_R^{(i)}}{\pi_{NG}^{(i)}} \right)$$

The value of the decision variables on each trial is set from a linear model which decomposes the decision variable into a stimulus-independent ‘bias’ value \mathcal{B}_L and \mathcal{B}_R (one for each decision variable), and a stimulus-dependent ‘sensitivity’ value \mathcal{S}_L and \mathcal{S}_R . Each expression contains the contrast of the appropriate hemifield (contrast left c_L and contrast right c_R) raised to an exponent \mathcal{N} which allows for a saturating non-linear contribution of contrast value to the decision variables,

$$Z_L^{(i)} = \mathcal{B}_L^{(i)} + \mathcal{S}_L^{(i)} \times \left(c_L^{(i)} \right)^{\mathcal{N}^{(i)}}$$

$$Z_R^{(i)} = \mathcal{B}_R^{(i)} + \mathcal{S}_R^{(i)} \times \left(c_R^{(i)} \right)^{\mathcal{N}^{(i)}}$$

The bias, sensitivity and shape parameters are indexed by trial i , because the parameters for each trial are constructed from a hierarchical prior which allows for per-session and per-subject deviations from a grand average set of parameters,

$$\mathcal{B}_L^{(i)} = b_L + b_L^{SESSION[i]} + b_L^{SUBJECT[i]}$$

$$\mathcal{B}_R^{(i)} = b_R + b_R^{SESSION[i]} + b_R^{SUBJECT[i]}$$

$$\mathcal{S}_L^{(i)} = s_L + s_L^{SESSION[i]} + s_L^{SUBJECT[i]}$$

$$\mathcal{S}_R^{(i)} = s_R + s_R^{SESSION[i]} + s_R^{SUBJECT[i]}$$

$$\mathcal{N}^{(i)} = n + n^{SESSION[i]} + n^{SUBJECT[i]}$$

Where (b_L, b_R, s_L, s_R, n) are 5 global parameters, estimated as the grand mean across all sessions and subjects. $(b_L^{SESSION[i]}, b_R^{SESSION[i]}, \dots)$ are the deviations in the 5 parameters, from the grand mean value, for the session corresponding to trial i . Likewise, $(b_L^{SUBJECT[i]}, b_R^{SUBJECT[i]}, \dots)$ are the deviations in the 5 parameters for the subject corresponding to trial i . For each level of the hierarchy (per-session and per-subject), the deviations of the 5 parameters are given a hyperprior, allowing the model to estimate the variance and covariance of the

parameter deviations within each level of the hierarchy. The model therefore explicitly estimates the variation in parameters which arise from true variation as opposed to sampling noise. For example, the model allows us to estimate how much $b_L^{SESSION}$ varies across sessions, as well as estimating how much covariance there is between $b_L^{SESSION}$ and $b_R^{SESSION}$ across sessions,

$$\begin{bmatrix} b_L^{SESSION} \\ b_R^{SESSION} \\ s_L^{SESSION} \\ s_R^{SESSION} \\ n^{SESSION} \end{bmatrix} \sim MVN \left(\begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \Sigma^{SESSION} \right)$$

$$\begin{bmatrix} b_L^{SUBJECT} \\ b_R^{SUBJECT} \\ s_L^{SUBJECT} \\ s_R^{SUBJECT} \\ n^{SUBJECT} \end{bmatrix} \sim MVN \left(\begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \Sigma^{SUBJECT} \right)$$

Where $\Sigma^{SESSION}$ and $\Sigma^{SUBJECT}$ are 5x5 covariance matrices reflecting the variance within each parameter across the sessions and subjects, as well as the covariance between parameters across sessions and subjects. These covariance matrices are each parameterised with standard deviation parameters $\sigma^{SESSION}$ and $\sigma^{SUBJECT}$, and 5x5 correlation matrices $\mathcal{R}^{SESSION}$ and $\mathcal{R}^{SUBJECT}$,

$$\Sigma^{SESSION} = \text{diag} \begin{bmatrix} \sigma_{b_L^{SESSION}} \\ \sigma_{b_R^{SESSION}} \\ \sigma_{s_L^{SESSION}} \\ \sigma_{s_R^{SESSION}} \\ \sigma_{n^{SESSION}} \end{bmatrix} \times \mathcal{R}^{SESSION} \times \text{diag} \begin{bmatrix} \sigma_{b_L^{SESSION}} \\ \sigma_{b_R^{SESSION}} \\ \sigma_{s_L^{SESSION}} \\ \sigma_{s_R^{SESSION}} \\ \sigma_{n^{SESSION}} \end{bmatrix}$$

$$\Sigma^{SUBJECT} = \text{diag} \begin{bmatrix} \sigma_{b_L^{SUBJECT}} \\ \sigma_{b_R^{SUBJECT}} \\ \sigma_{s_L^{SUBJECT}} \\ \sigma_{s_R^{SUBJECT}} \\ \sigma_{n^{SUBJECT}} \end{bmatrix} \times \mathcal{R}^{SUBJECT} \times \text{diag} \begin{bmatrix} \sigma_{b_L^{SUBJECT}} \\ \sigma_{b_R^{SUBJECT}} \\ \sigma_{s_L^{SUBJECT}} \\ \sigma_{s_R^{SUBJECT}} \\ \sigma_{n^{SUBJECT}} \end{bmatrix}$$

Finally, the grand-average parameters are given the following prior distributions, based on previous fits from the non-hierarchical logistic regression model outlined earlier.

$$(b_L, b_R) \sim \text{Normal}(0, 2^2)$$

$$(s_L, s_R) \sim \text{Normal}(5, 2^2)$$

$$n \sim \text{Normal}(0.5, 0.25^2)$$

The standard deviation parameters ($\sigma_{b_L^{SESSION}}$, $\sigma_{b_L^{SUBJECT}}$, etc) are given a weakly-informative Half Cauchy distribution which penalises extreme values. These priors

act to regularise variation, ensuring that the parameters don't vary too much across sessions and subjects.

$$\left(\sigma_{b_L^{SESSION}}, \sigma_{b_R^{SESSION}}, \dots \right) \sim \text{HalfCauchy}(0, 1)$$

$$\left(\sigma_{b_L^{SUBJECT}}, \sigma_{b_R^{SUBJECT}}, \dots \right) \sim \text{HalfCauchy}(0, 1)$$

The correlation matrices were given an LKJ prior (Lewandowski et al., 2009). The LKJ prior is a prior over correlation matrices, using a penalty factor 2 which penalises extreme positive and negative correlations between parameters. This also regularises variation but explicitly allows for variation in one parameter across sessions to ‘inform’ estimates of another parameter across sessions,

$$\mathcal{R}^{SESSION} \sim \text{LKJcorr}(2)$$

$$\mathcal{R}^{SUBJECT} \sim \text{LKJcorr}(2)$$

We fit this single model to a dataset comprising 91 sessions from 5 mice. Model fitting was achieved with numerical estimation of the joint posterior distribution over all parameters.

In many complex models, the joint posterior distribution over all parameters is not analytically solvable. Instead, one can numerically estimate the posterior distribution by drawing samples from it. If enough samples are drawn, the moments of the posterior distribution (true mean, true variance, etc) can be estimated from the samples (sample mean, sample variance, etc). This requires however that the posterior distribution can be evaluated at any given point in the parameter space, even if the full functional form cannot be produced. Markov Chain Monte Carlo (MCMC) methods are a class of algorithms which achieve numerical sampling of an arbitrarily complex posterior distribution. The general procedure of these methods is for the sampler to start at a random point in parameter space. The sampler moves around parameter space towards the region of parameter space with the higher posterior probability density. This is achieved by, at each iteration, making proposing a random direction as to where to go next, but only “accepting” the choice if that choice would increase the posterior probability of the location in parameter space. The sampler therefore operates in two phases: the first phase is to move from a random location towards the general vicinity of the posterior distribution. The second phase is to move around the region with high posterior density. The sampler will spend more time around the peak of the posterior distribution, but also occasionally explore the regions

towards the edges. The history of locations of the sampler can be treated as samples of the posterior distribution.

We perform the posterior distribution sampling using the probabilistic programming language Stan (Carpenter et al., 2017). This language uses Hamiltonian Monte Carlo (HMC) with No-U-Turn Sampling. HMC is a type of MCMC algorithm. No-U-Turn Sampling augments HMC by preventing the sampler from making any U-turns in parameter space.

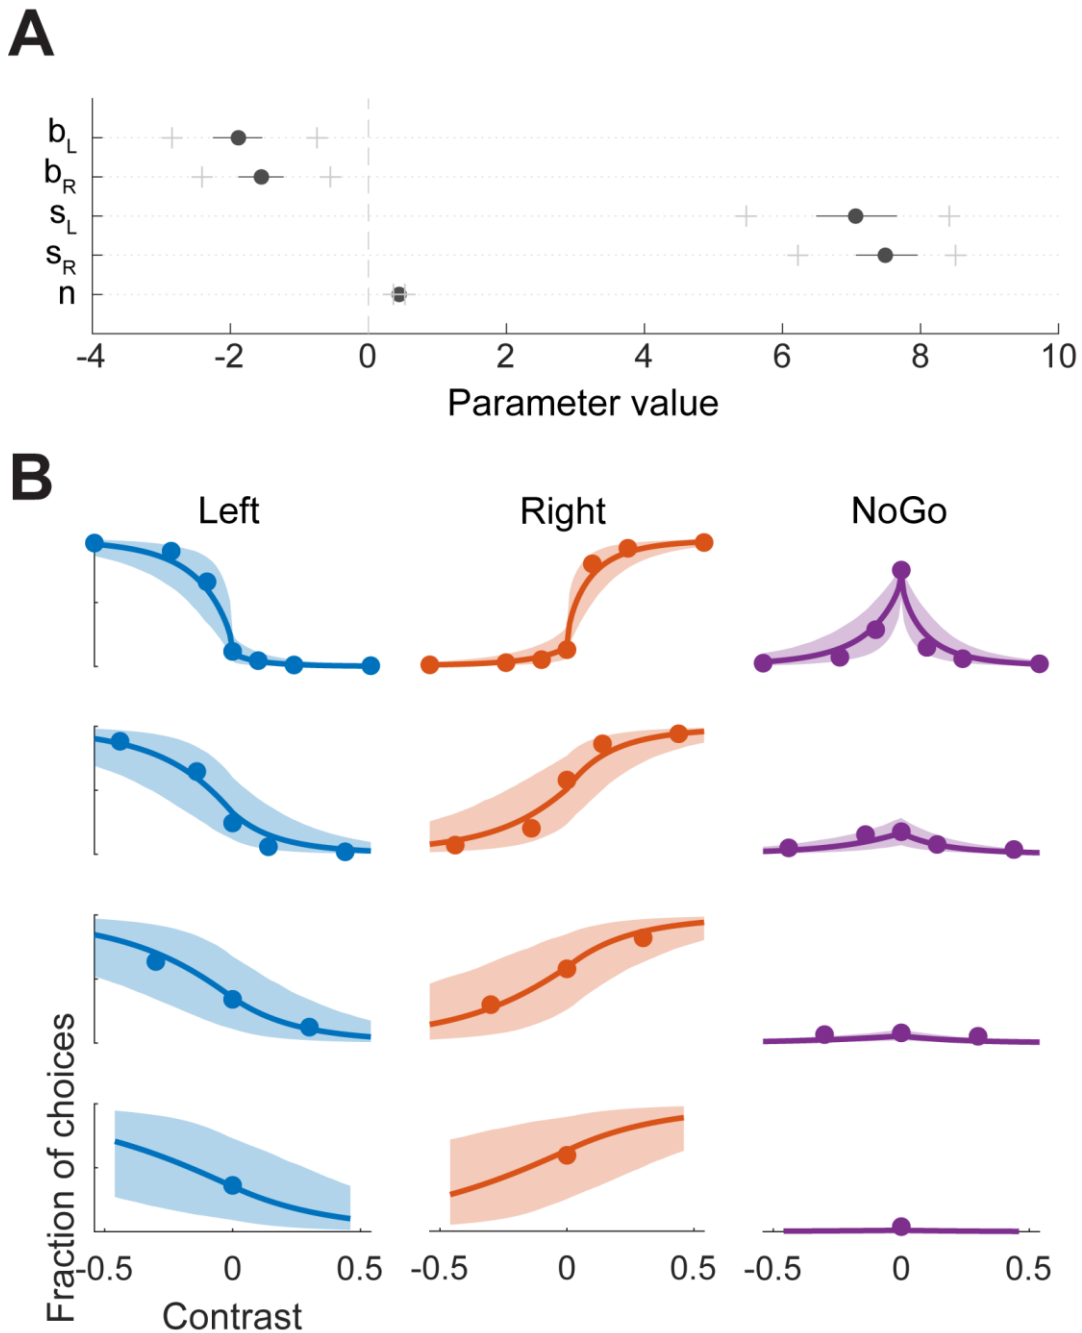


Figure 2-10 Posterior fit of the grand average data

(A) Posterior distribution of the 5 global parameters. The dots are the posterior mean. Solid dark lines span the 20th and 80th percentile, and crosses mark the 2.5th and 97.5th percentile.

(B) Fit of the global parameters to the grand average psychometric data. The plot uses the same convention as in Figure 2-4. Dots are the fraction of choices averaged over all sessions. The solid line and shaded region are the mean and 95% credible intervals of the posterior prediction from the global parameters of the hierarchical model.

The model fit could capture several features of the data. Firstly, the fit provided estimates of the grand average parameters, as well as uncertainty associated with each (Figure 2-10A). These parameters were able to capture the dependence of behaviour on the stimulus contrast (Figure 2-10B). The uncertainty in the parameter estimates also provided estimates of uncertainty in fraction of choices present in the psychometric curves.

Secondly, the model could capture variation in behaviour across sessions and subjects. The hierarchical structure of the model significantly improved fit of the dataset compared to a similar model without per-session and per-subject hierarchy (hierarchical: $33,533 \pm 312$, non-hierarchical: $37,025 \pm 21$, WAIC score \pm estimated standard error). For all parameters, the variation across subjects ($\sigma^{SUBJECT}$) was slightly larger than the variation across sessions within a subject ($\sigma^{SESSION}$) (Figure 2-11). The bias parameters showed the largest variation across sessions and subjects.

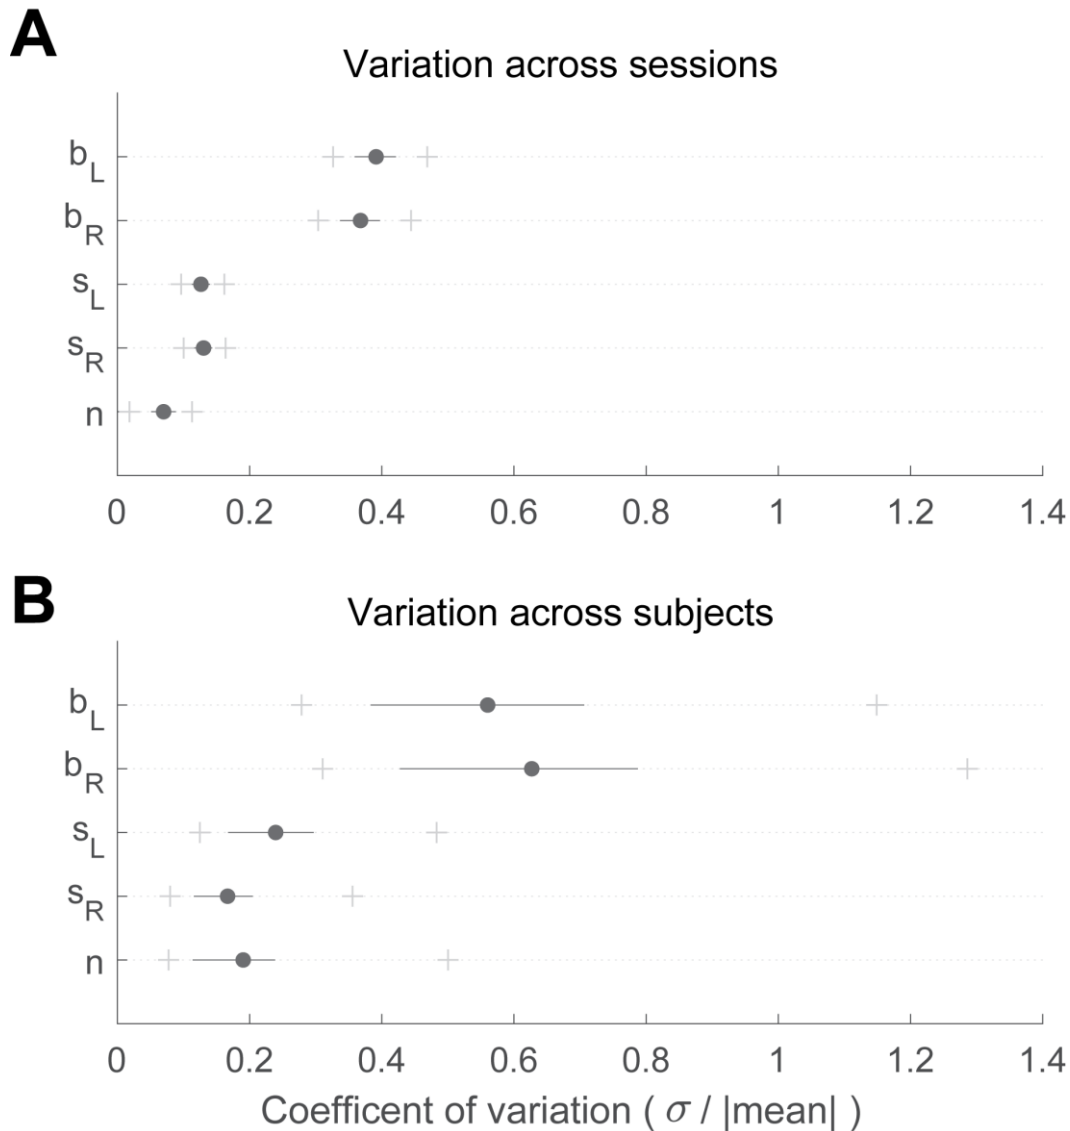


Figure 2-11 Parameter variation across sessions and subjects

(A) Posterior distribution of the per-session variation parameters $\sigma^{SESSION}$ normalised by the absolute value of the posterior mean of the corresponding global parameter. Single dots are the posterior mean. Solid grey lines span the 20th and 80th percentile, and crosses mark the 2.5th and 97.5th percentiles.

(B) Same plot but for the per-subject variation parameters $\sigma^{SUBJECT}$.

What does the high variability in the bias parameters correspond to in the behaviour? To explore this, we leveraged the fact that the behaviour on zero contrast trials is only determined by the bias parameters, because the sensitivity and shape parameters contribute on trials with non-zero contrast. Therefore, variation in the bias parameters must reflect variation in the choice behaviour on zero contrast trials. We found that different sessions varied substantially in the tendency to NoGo on zero contrast trials. The hierarchical model could capture

this NoGo variation successfully (Figure 2-12A), while regularising extremely high or low NoGo rates. Since the bias parameters are equivalent to the value of $\ln\left(\frac{\pi_L}{\pi_{NG}}\right)$ and $\ln\left(\frac{\pi_R}{\pi_{NG}}\right)$ on zero contrast trials, NoGo variation can only arise from positive correlation in the two bias parameters. Accordingly, the posterior fit of the model showed significant positive correlation in the two bias parameters across all sessions (0.49 posterior mean correlation, [0.27, 0.68] 95% credible intervals; Figure 2-12B).

The model was also able to capture variation in behaviour across the full set of contrast conditions. The variation present in the bias parameters, together with the smaller variation in the sensitivity and shape parameters, could capture behavioural variability across the contrast conditions and choices (Figure 2-12C-D).

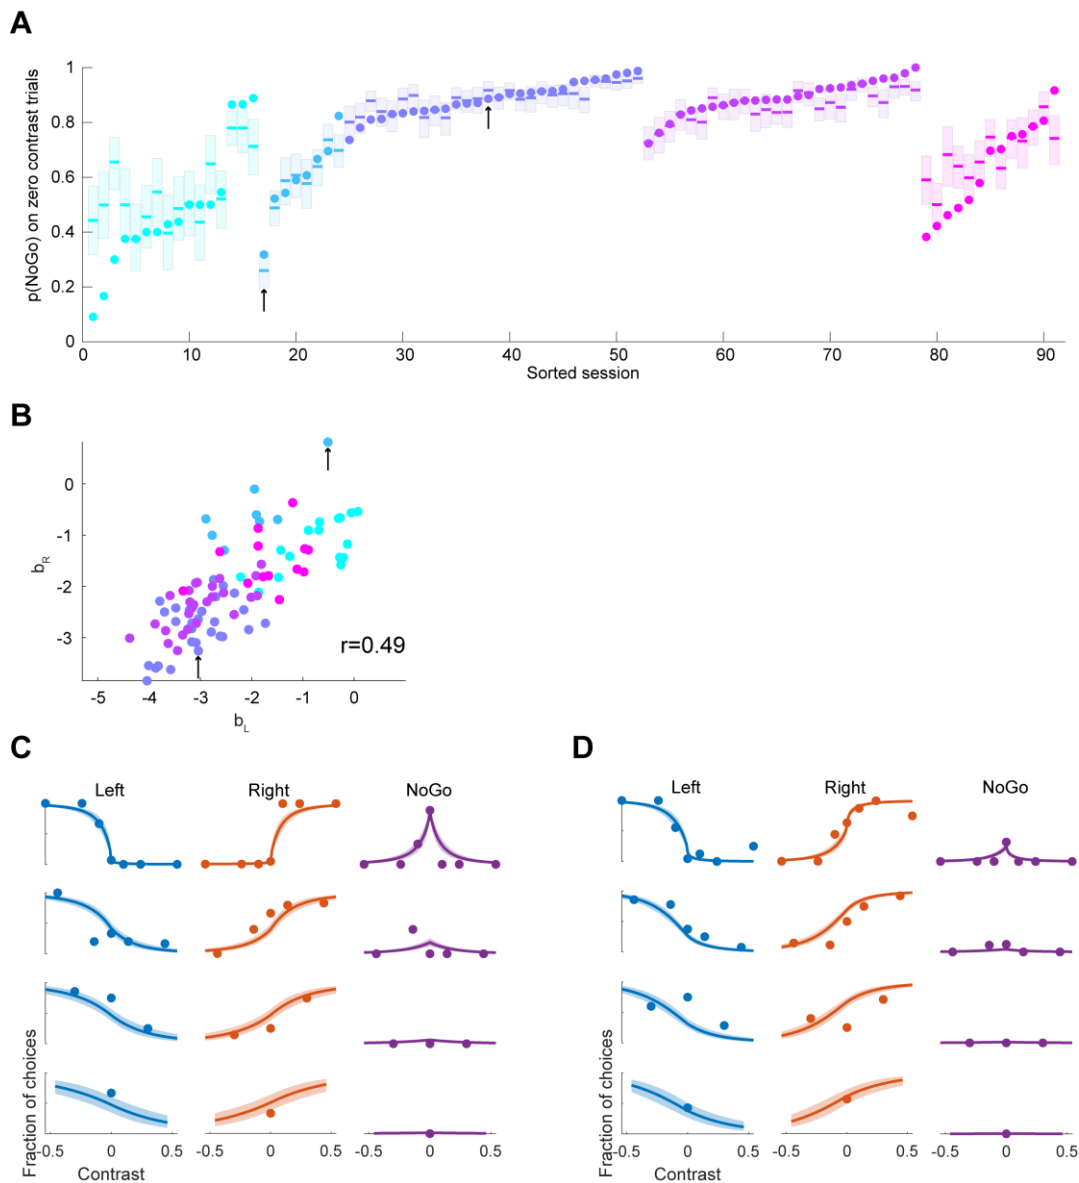


Figure 2-12 Model can capture variation in NoGo rate across sessions

(A) Each dot is the probability of NoGo on zero contrast trials for one session. Each shaded region is the mean and 95% credible interval of the model fit for that session. Each colour represents data from one mouse, and sessions are sorted within each mouse in ascending order of NoGo probability. The two sessions highlighted with black arrows are plotted below in (C) and (D).

(B) The posterior mean of the two bias parameters b_L and b_R is plotted for each session. The posterior mean correlation between the parameters is shown inset.

(C) Example model fit for one session with high NoGo bias. Same plotting convention as in Figure 2-10.

(D) Example model fit for one session with low NoGo bias.

2.6 Discussion

We have developed a two-alternative unforced-choice visual discrimination task for mice. The task requires mice to explicitly integrate visual information from both hemifields and indicate which side has the higher contrast using a wheel turn or abstain from turning if no stimulus is present. We have used behavioural data from this task to constrain a phenomenological model, which decomposes choices into perceptual states such as choice bias and stimulus sensitivity. We have also quantified how these states vary between sessions and mice. Using this method, we show that variation across sessions and subjects was primarily in the bias parameters, and that this could account for variation in the NoGo rates observed empirically.

Many rodent decision tasks employ either Go-NoGo (Goard et al., 2016; Guo et al., 2014; Li et al., 2016) or forced-choice (Akrami et al., 2018; Erlich et al., 2015; Hanks et al., 2015; Katz et al., 2016; Licata et al., 2017; Mante et al., 2013; Raposo et al., 2014; Znamenskiy and Zador, 2013) regimes. However, Go-NoGo and forced-choice tasks have several problems. Go-NoGo confounds perceptual biases with non-perceptual biases from unbalanced motor costs and disengagement. Forced-choice confounds decision bias with uncertainty bias, and also confounds bias associated with detecting one of the stimuli with bias associated with comparing between the stimuli. These problems limit how much one can infer perceptual states from the behavioural data and limit the possible interpretations of neural manipulation and recording.

Our task design solves a number of these issues. The combination of visual discrimination with a two-alternative unforced-choice design limits the role of uncertainty-related bias. Left-Right choice bias cannot be confounded with uncertainty bias because the perceptual state of uncertainty (i.e. being unsure that a stimulus is present at either location) is now reported with NoGo. Another form of uncertainty, where the subject knows that there is a stimulus present but is unsure which side (Left or Right) is correct, may still be affected by uncertainty-related biases.

Our task structure can be modelled as two independent binary decisions each using visual information from one hemifield. The Left vs. NoGo decision axis is informed only the left contrast. The Right vs. NoGo axis is informed only by the right contrast. Each of these decisions have their own perceptual states (bias and sensitivity). This decomposition makes it possible to interpret different kinds of

behavioural manipulation. For example, an inflated preference towards Left choices can arise either from a bias in the Left vs. NoGo decision (increased b_L) or a bias in a Left vs. Right comparison (increased b_L and decreased b_R). This model is related to models from Signal Detection Theory applied to multi-alternative decisions (Sridharan et al., 2014), and econometric logistic models of binary-choice decisions (Busse et al., 2011; Greene, 2011; Roe et al., 2001). However, we expand on these models by incorporating an estimate of non-linear functional dependence of the behaviour on the contrast value. A further advantage of the model framework is the ability to incorporate other regressors of interest, such as choice history (Akrami et al., 2018; Frund et al., 2014; Lueckmann et al., 2018) and arousal (Jacobs et al., 2018).

Our model also explicitly estimates variation in perceptual parameters over sessions and subjects. This type of multi-level modelling is fairly routine in human psychophysics and social science (Gelman et al., 2013; Lee, 2011; Rouder and Lu, 2005). However, to our knowledge this approach has not been used in rodent studies identifying perceptual states. Many rodent studies instead estimate bias and sensitivity for sessions separately (Busse et al., 2011; Licata et al., 2017). However, this has a high risk of model overfitting which makes parameter interpretation misleading. Our hierarchical model has regularisation built into the structure by relating sessions to each other by common subject-level parameters and relating subjects to each other by common global parameters.

2.6.1 Caveats

There are a few caveats to this work. Firstly, the task does not strictly control the precise perceptual strategy used to solve the task, nor the motor plan which the mice use to move the wheel. For example, mice may solve the discrimination task either by attending to both stimuli at once, or by attending to the one and then the other side. Likewise, mice may move the wheel using one or two paws, or turn the wheel with a brief strong movement, or multiple smaller movements. Our behavioural model successfully controls for inter-subject and inter-session variability, but only at the level of which choices tend to be made, and not at the level of how the motor plan is performed nor the unobserved perceptual strategy. Additionally, if perceptual strategies vary fundamentally between subjects, then it may be invalid to decompose behavioural performance into the same perceptual parameters (bias, sensitivity, shape) in the model.

Secondly, our task may suffer from the same motivational bias present in Go-NoGo tasks. We observe variation in NoGo rates across sessions as outlined earlier. This variation in NoGo rates across sessions may be accounted for either because of variation in long-term engagement in the task, or because of variation in perceptual biases. It is unlikely that this variation across sessions can be accounted for by variation in long-term engagement, because our repeated-trial design ensures that periods of disengagement terminate the session early. Short-term disengagement is harder to identify; however, we feel this kind of disengagement is rare because shrinking the proportion of zero-contrast trials leads to almost no NoGo choices (data not shown), hence we over-represent zero-contrast trials at ~30%. If short-term disengagement was regularly present, we would observe more NoGo choices than we currently do.

Thirdly, the model assumes that the decision on each trial can be described by a statistical model parameterised with a fixed set of numbers for each session. This ignores the possibility for perceptual biases and sensitivity to gradually change over the course of a single session. States can in principle vary across a session, for example motivational drift, water satiety, and muscular fatigue. These within-session changes in state would be reflected as changes in the model parameters yet our model ignores this variation and makes inferences only at the level of the session and not trial.

Fourthly, the hierarchical design of the model assumes that parameter variation across sessions within a subject is normally distributed. However, it's possible that a single mouse's perceptual states may be, for example, bimodally distributed across sessions. Similarly, variation in subject-level parameters may not be normally distributed around a global mean value, and instead may be clustered into multimodal groupings. Our model therefore imposes a unimodal variation between sessions and subjects, for the purposes of regularizing the fit. However, this model may be inappropriate for variation which is not Gaussian-like. Nevertheless, the success of the model fit in accounting for behavioural variation across sessions and subjects implies that the Gaussian approximation is justified for our dataset.

Finally, we assume that our task does not involve integration of evidence over time, and hence the behaviour can be appropriately described with a model which does not involve an accumulated decision variable. There are two ways that a task can be treated as involving integration of evidence. Firstly, tasks where the

identity of the correct choice can only be known from multiple stimulus events. For example, a task where the subject has to accumulate a number of Poisson clicks on the left and right side to determine which side has the highest average number of click events (e.g. Brunton et al., (2013)). Secondly, tasks where the identity of the correct choice could, in principle, be known from a single 'frame' of sensory evidence, but due to sensory ambiguity, continual observation of the same sensory stimuli improves choice accuracy. Our task is not of the first but could in principle involve an integration of evidence of the second type. This would be more applicable for low-contrast trials than high contrast trials. Consistent with this, reaction times for low contrast trials tend to be longer than for high contrast trials (data not shown). Such a situation would necessitate modifying the model to incorporate choice timing, as in drift-diffusion models. However, fitting such models requires knowledge of each instantaneous observation event of the static stimulus, which is currently not measurable for our task. Instead the task could be modified to incorporate brief pulsed stimulus events instead of a static Gabor stimulus, which would be more amenable to drift-diffusion modelling.

CHAPTER 3 SEQUENTIAL ACTIVATION OF CORTICAL AREAS

The work described in this chapter, and in subsequent chapters, contributed towards a publication,

Zatka-Haas P.*, Steinmetz N.*, Carandini M., Harris K.D., (in preparation) “A mechanistic model of cortical contributions to visual discrimination in mice”

* These authors contributed equally

Note: Widefield calcium imaging data was acquired by Nick Steinmetz. My contribution in this Chapter is in data analysis.

3.1 Introduction

Perceptual decision making recruits many brain areas involved in sensation, motor planning and action execution (Gold and Shadlen, 2007; Svoboda and Li, 2018). However, it is not clear exactly how this decision process unfolds neurally.

Evidence from observational and perturbational studies have supported the idea that cortical areas play key roles in perceptual decisions. In tasks where monkeys make eye movements to indicate the direction of a moving stimulus, choice-predictive signals have been observed in individual neurons from areas spanning the sensory-motor axis: visual area MT (Britten et al., 1992), the lateral intraparietal area (LIP; Roitman and Shadlen, 2002; Shadlen and Newsome, 2001), and visuomotor frontal eye field (FEF; Ding and Gold, 2012; Gold and Shadlen, 2000). In rodents, several studies have found choice-predictive signals for different tasks in early sensory areas (Nienborg and Cumming, 2009; Sachidhanandam et al., 2013; Yang et al., 2016), posterior parietal cortex (Harvey et al., 2012; Licata et al., 2017; Park et al., 2014; Raposo et al., 2014), and frontal-motor areas (Barthas and Kwan, 2016; Chen et al., 2017; Erlich et al., 2015; Goard et al., 2016; Guo et al., 2014; Hanks et al., 2015).

To what extent does the decision process arise from distinct brain areas operating sequentially, or from multiple areas active simultaneously (Cisek and Kalaska, 2010; Hunt and Hayden, 2017)? In tactile and visual discrimination, stimulus-specific activity is sequentially ordered from early sensory to premotor and motor areas (Le Merre et al., 2018), whereas choice-predictive activity arises simultaneously in many of the areas recorded (Hernández et al., 2010; Ledberg et al., 2007; Siegel et al., 2015). However, these studies have relied on tasks with an

explicit delay phase, where subjects have to hold their memory of a decision in mind for a prolonged period before performing the choice action. Therefore, it is unclear whether the sequential task-related activation arises because of neural processes inherent to decision making, or instead arises from structure imposed by a sequential task design.

This Chapter will outline work done towards investigating the activity of dorsal cortex while mice perform a decision task. Mice were trained on the two-alternative unforced-choice visual discrimination task as detailed in Chapter 2 and Burgess et al (2017). In this task, mice were permitted to respond immediately after stimulus onset, and therefore the task contains no delay period. While mice performed this task, we employed widefield calcium imaging to measure cortical activity. Imaging of the dorsal cortex revealed a robust sequential activation of visual, secondary motor, primary motor and somatosensory areas in response to presentation of the visual stimulus. We found that activation of visual and secondary motor areas was stimulus-dependent, whereas activation of primary motor and primary somatosensory areas was movement-dependent. This work sets up subsequent Chapters which will explore the causal relevance of this sequential activation in forming the upcoming choice.

3.2 Methods

3.2.1 Behaviour

Mice performed the behavioural task described in Chapter 2, but with a few modifications. Trials began after a period of no wheel movement (0.3-0.7seconds) which was 100msec longer than non-widefield imaging sessions. This helped reduce contamination of widefield signals associated with a visual stimulus on any given trial with residual GCaMP6s fluorescence on the previous trial. We also introduced a delay period (0.5-1.2 seconds) after stimulus onset before the auditory go cue, when the stimulus was present on the screen but could not be moved by the wheel. Mice tended to move the wheel before the go cue anyway, and so for the initial wheel movement the stimulus was motionless on the screen. After the delay period was complete, the ongoing wheel movement led to stimulus movement which would permit the mice to make a response and obtain rewards. This small task modification was important to ensure that stimulus-related cortical activity was not inter-mixed with activity related to an auditory go cue beep, and that movement-related activity was not inter-mixed with signals

related to movement of the stimulus on the screen. We exclude any trials where mice start moving after this go cue. For analysis in this study, ‘reaction time’ refers to the time from stimulus onset to the time of first wheel movement, even if this movement is prior to the delayed go cue.

3.2.2 Widefield calcium imaging

Widefield calcium imaging was performed in 5 transgenic mice expressing GCaMP6s in excitatory cells (Camk2a-TTA x TetO-GCaMP6s; mouse ID 3-5) or all neurons (Snap25-GCaMP6s; mouse ID 1-2) (Chen et al., 2013; Wekselblatt et al., 2016). These mouse lines do not show aberrant epileptiform activity shown in related mouse-lines (Steinmetz et al., 2017). We imaged using a macroscope (Scimedia THT-FLSP) with sCMOS camera (PCO Edge 5.5) and dual-wavelength illumination (Cairn OptoLED). The macroscope used 1.0x condenser lens (Leica 10450028) and 0.63x objective lens (Leica 10450027). Images were acquired from the PCO Edge with ~10-ms exposures and 2 x 2 binning in rolling shutter mode. Images were acquired at 70Hz, alternating between blue and violet illumination (35Hz each). The light sources were 470-nm and 405-nm LEDs (Cairn OptoLED, P1110/002/000; P1105/405/LED, P1105/470/LED). Excitation light passed through excitation filters (blue: Semrock FF01-466/40-25; violet: Cairn DC/ET405/20x), and through a dichroic (425 nm; Chroma T425lpxr). Excitation light then went through 3-mm-core optical fibre (Cairn P135/015/003) and reflected off another dichroic (495 nm; Semrock FF495- Di03-50x70) to the brain. Emitted light passed through the second dichroic and an emission filter (Edmunds 525/50-55 (86-963)) to the camera. Alternation was controlled with custom code on an Arduino Uno, and illumination was restricted to the ‘global’ phase of the rolling shutter exposures, i.e. only the times when all pixels of a frame were being exposed together. We de-noised the signal with singular value decomposition and normalised the signal to the mean fluorescence (over the entire recording session) at each pixel. The signal from the 405nm illumination frames was used to correct for parts of the 470nm signal that were due to changes in blood flow that obstruct the fluorescence signal (Ma et al., 2016) and the correction was performed with custom MATLAB code (<https://github.com/cortex-lab/widefield>). We then low-pass filtered the signal at 8.5Hz and applied a derivative filter to the fluorescence trace to approximate deconvolution of the calcium sensor’s time course from the underlying neural activity. When computing event-triggered averages of the fluorescence, pre-event baseline activity was removed, removing impact of long-term trends. To ensure that differences in fluorescence timing between

contralateral and ipsilateral stimuli were not due to differences in reaction time between contraversive and ipsiversive choices, we matched the reaction time distribution between Left and Right choices. This was achieved by pooling trials across sessions for each mouse and excluding Left and Right trials until a coarsely-binned reaction time histogram matched between the two choice types.

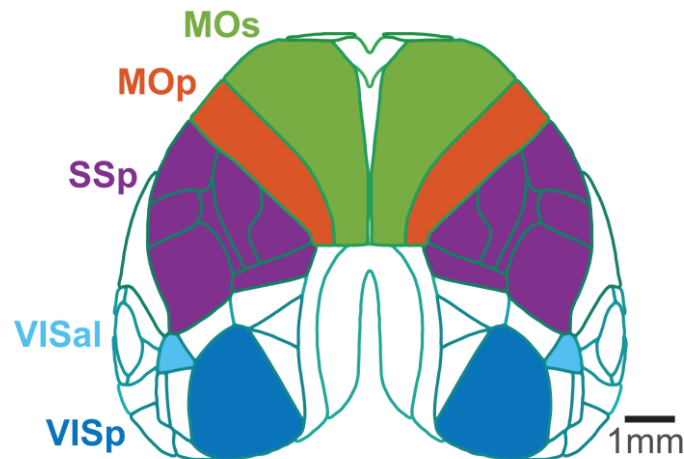


Figure 3-1 Schematic of Allen Common Coordinate Framework

Schematic of dorsal cortex from the Allen Common Coordinate Framework (CCF; Lein et al. (2007)). Primary visual (VISp), secondary visual (e.g. VISal), secondary motor (MOs), primary motor (MOp) and primary somatosensory (SSp) areas are highlighted.

ROIs were selected based on the Allen Common Coordinate Framework atlas (Figure 3-1; Lein et al. (2007)) aligned manually to each mouse, guided by skull features and retinotopic maps. ROI positions were manually adjusted to account for inter-mouse differences. Primary visual area (VISp) was selected as the peak of the most posterior-medial activated site in visual cortex in response to a contralateral stimulus. Secondary visual area (VISal) was selected as the centre of VISal according to the Allen CCF. VISal was selected as an exemplary secondary visual cortical area because it was furthest from the part of VISp activated by our visual stimuli, ensuring minimal contamination of fluorescence between these two ROIs. The secondary motor area (MOs) ROI was selected as centre of the most anterior site activated by the contralateral stimulus. Primary motor area (MOp) and primary somatosensory area (SSp) ROIs were selected as the most anterior-medial and posterior-lateral regions of the activated region spread over the MOp-SSp border.

3.3 Stimulus and movement-specific activation across cortex

Our first question was to ask which cortical areas are active during this task. We performed widefield calcium imaging in mice expressing GCaMP6s in excitatory neurons or in all neurons, while mice performed the task. We found robust activation in primary visual (VISp), secondary visual (e.g. VISal) and secondary motor (MOs) areas in response to stimuli presented on the contralateral side (Figure 3-2). This activation was present for trials where mice responded with movement, as well as NoGo trials. This indicates that the activation is specific to the stimulus presence and not movement. By contrast, we found broad bilateral activation of primary motor (MOp) and primary somatosensory (SSp) regions on trials where the mice moved the wheel, independent of whether there was a stimulus present. This bilateral activity began at movement onset (Figure 3-3), which suggests that the bilateral activation of MOp and SSp is related to ongoing motor signals or sensory feedback of ongoing movement. We observed weak and sustained activation of MOp from the beginning of the trial even on NoGo trials. On NoGo trials primary motor activity increased around the same time when Left or Right choices are made on other trials. This is consistent with mice performing NoGo using active muscular contraction to hold the wheel still. When comparing between correct Left and correct Right trials (right column of Figure 3-2 and Figure 3-3), widefield activity appeared to be selective for the choice only in the regions with sensory-related activity (VIS regions and MOs). The choice-selectivity in these areas arises because these regions are selective to the stimulus, and choice correlates with stimulus. By contrast, MOp and SSp regions did not show clear asymmetric sensitivity to the ongoing choice. Therefore, bilateral activation in these regions appears to be non-selective to Go movements in general.

These results highlight several cortical areas as being engaged in stimulus and movement-specific aspects of the task: unilateral visual and secondary motor areas relate to the stimulus, and bilateral primary motor and primary somatosensory areas relate to ongoing movement. In Chapter 5, the choice-predictive aspects of these signals will be further explored.

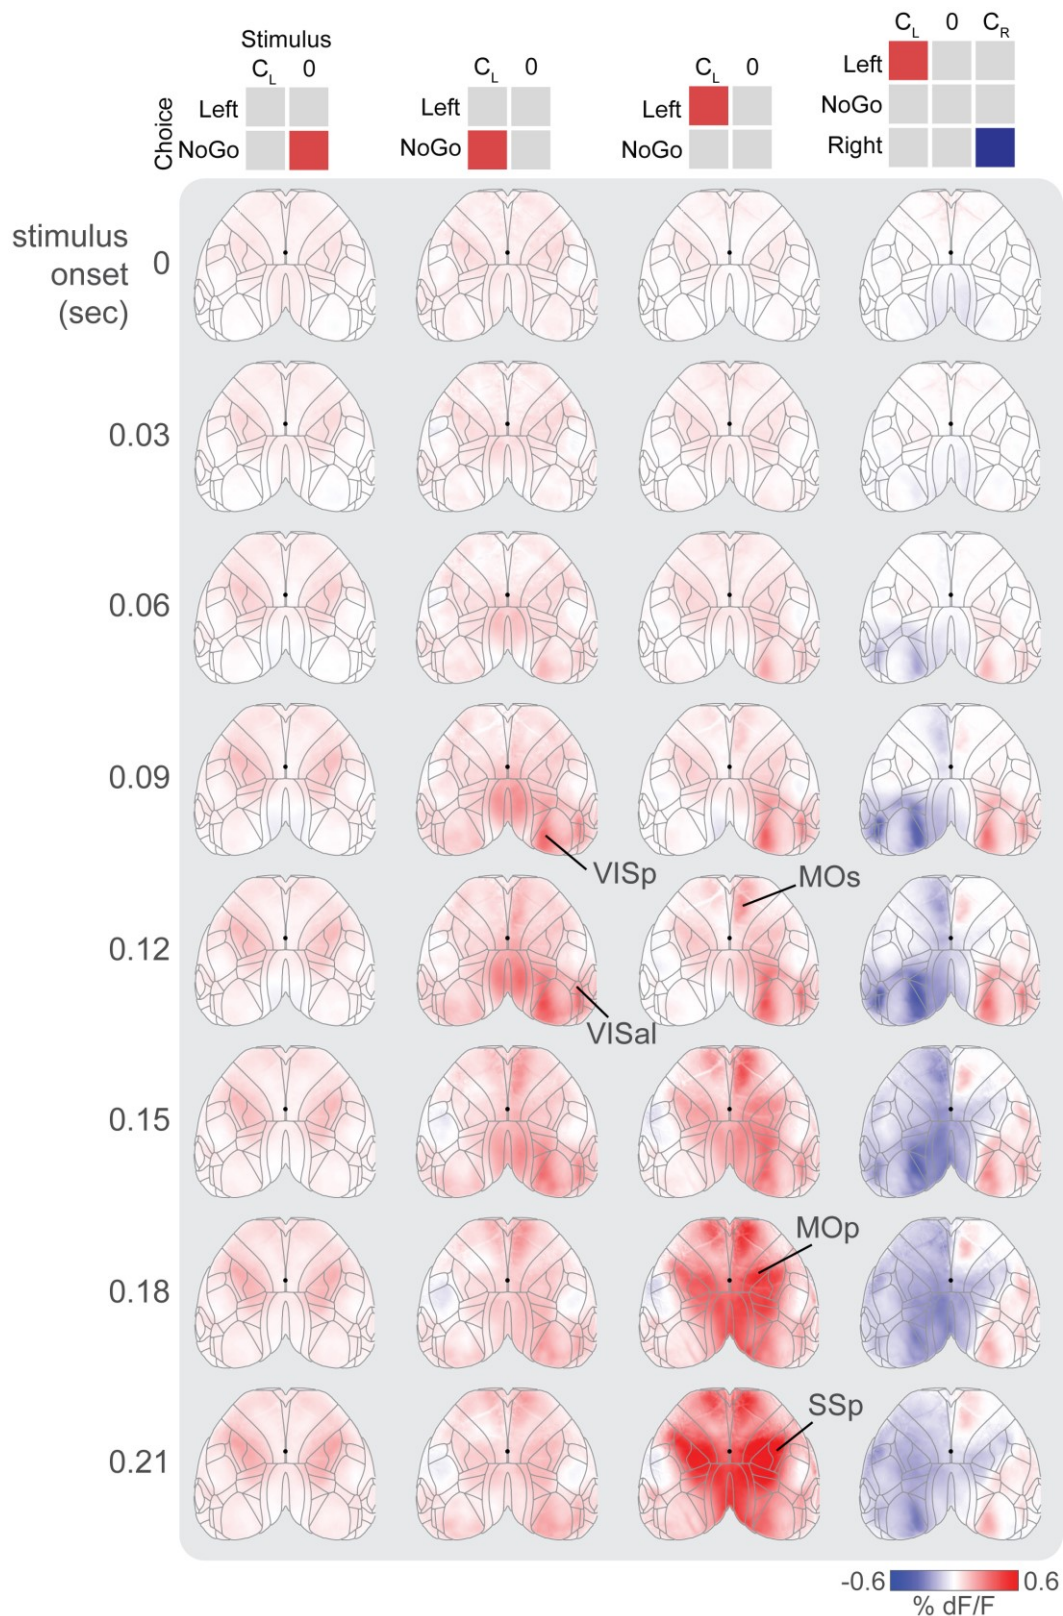


Figure 3-2 Stimulus-aligned cortical calcium fluorescence

Each colourmap shows cortical fluorescence (dF/F) averaged across 7 sessions from mouse 4 (*Camk2a-TTA x TetO-GCaMP6s*), aligned to stimulus onset. Each map is overlaid with an outline of cortical regions defined by the CCF, and the widefield fluorescence map is cropped to the outer edges of the CCF. Rows

indicate different times relative to stimulus onset. Columns correspond to a specific stimulus-response condition as indicated by the 2x2 grid on the top. Each column reflects the following conditions respectively: NoGo choice on zero contrast trials, NoGo choice on left contrast trials, Left choice on left contrast trials. The final column is the difference in fluorescence between two trial types: Left choice on left contrast (red) and Right choice on right contrast (blue).

The difference between the first two columns shows visual and secondary motor areas are active specifically in response to the stimulus. The difference between the second and third columns shows that large regions of the dorsal cortex (e.g. MOp and SSp) become bilaterally active specifically when mice move the wheel. The final column shows that visual and secondary motor areas are action-selective by virtue of being stimulus-selective, however primary motor and somatosensory regions are not clearly decision-specific.

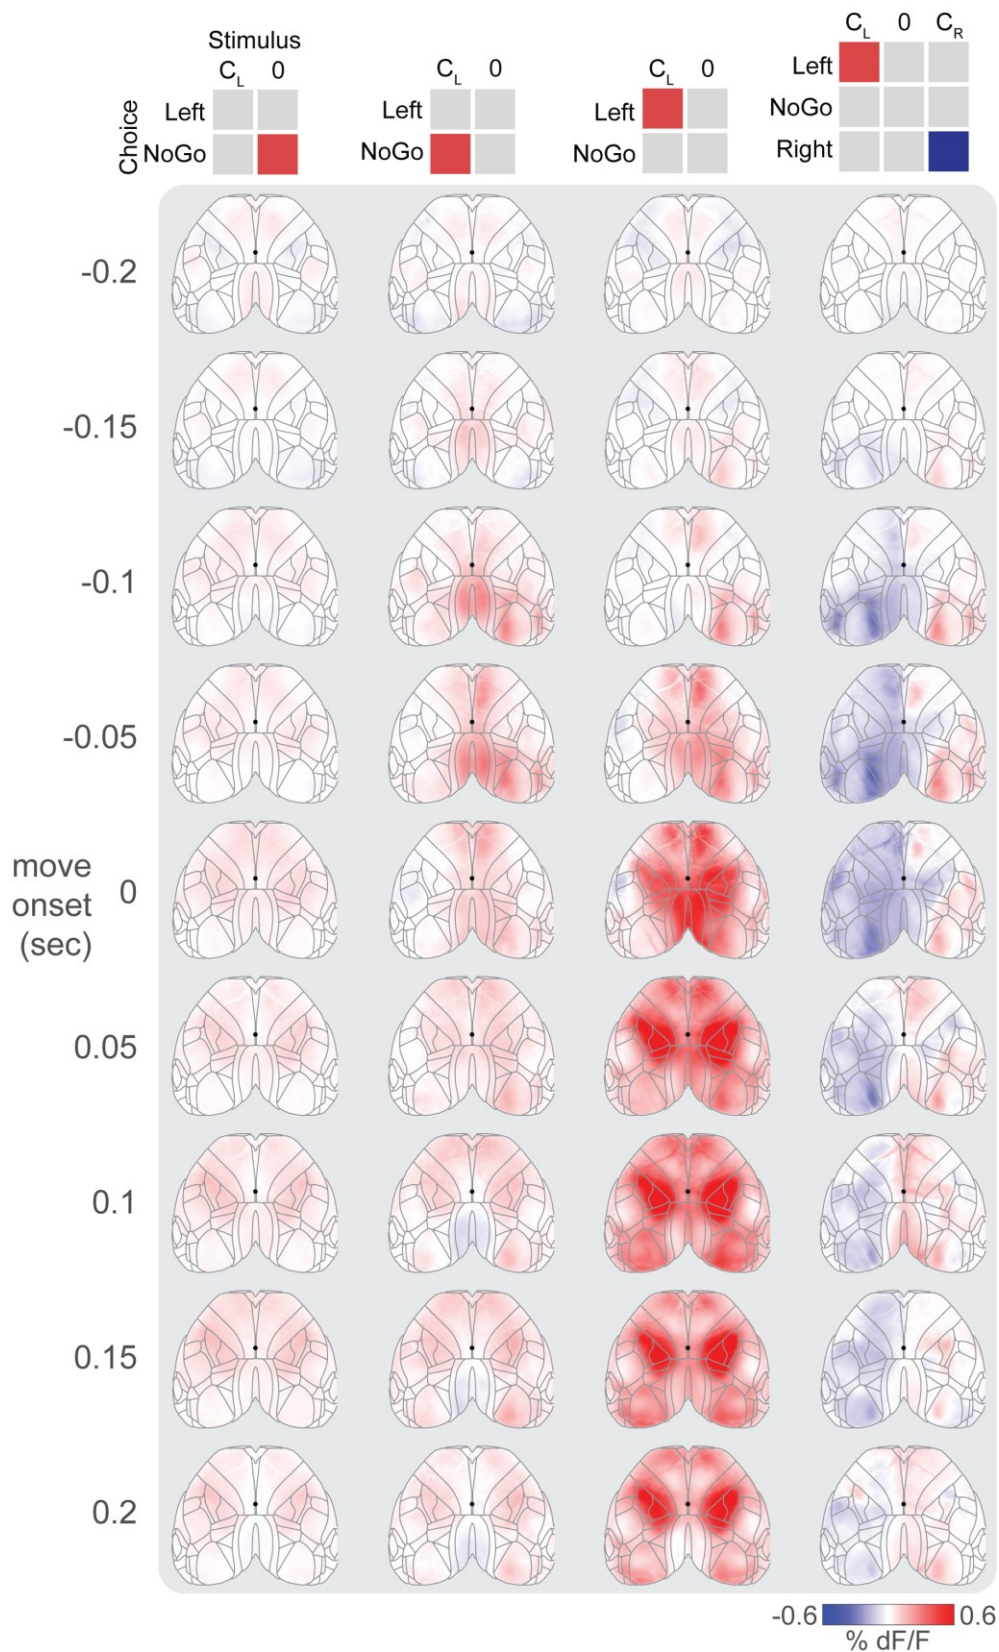


Figure 3-3 Movement-aligned cortical calcium fluorescence

Same plotting convention as in Figure 3-2 but aligned to movement onset. Movement “onset” for NoGo response trials is defined as a fixed delay after stimulus onset, corresponding to the median reaction time on Go trials.

3.4 Activation is sequential

We next asked whether these cortical regions show sequential activation or become active simultaneously. We analysed activity at single-pixel ROIs in primary visual (VISp), secondary visual (here defined as VISal), secondary motor (MOs), primary motor (MOp) and somatosensory (SSp) cortices.

We observed a reliable temporal sequence of activation, with onset (defined as time to 25% of peak activity) in VISp at 44 ± 2 ms (median \pm median absolute deviation across mice), VISal at 52 ± 4 ms and MOs at 90 ± 6 ms, before reaching areas such as MOp (142 ± 12 ms) and SSp (140 ± 10 ms) (Figure 3-4A-C). The early VISp, VISal, and MOs responses were absent for stimuli present on the ipsilateral side (Figure 3-4D), but the MOp/SSp responses were present for both contralateral and ipsilateral stimuli. We therefore conclude that presentation of a stimulus first elicits activity in contralateral visual cortex, which spreads to contralateral MOs then ipsilateral MOs, before reaching bilateral primary motor and somatosensory cortex, with the last areas to respond being ipsilateral visual cortex.

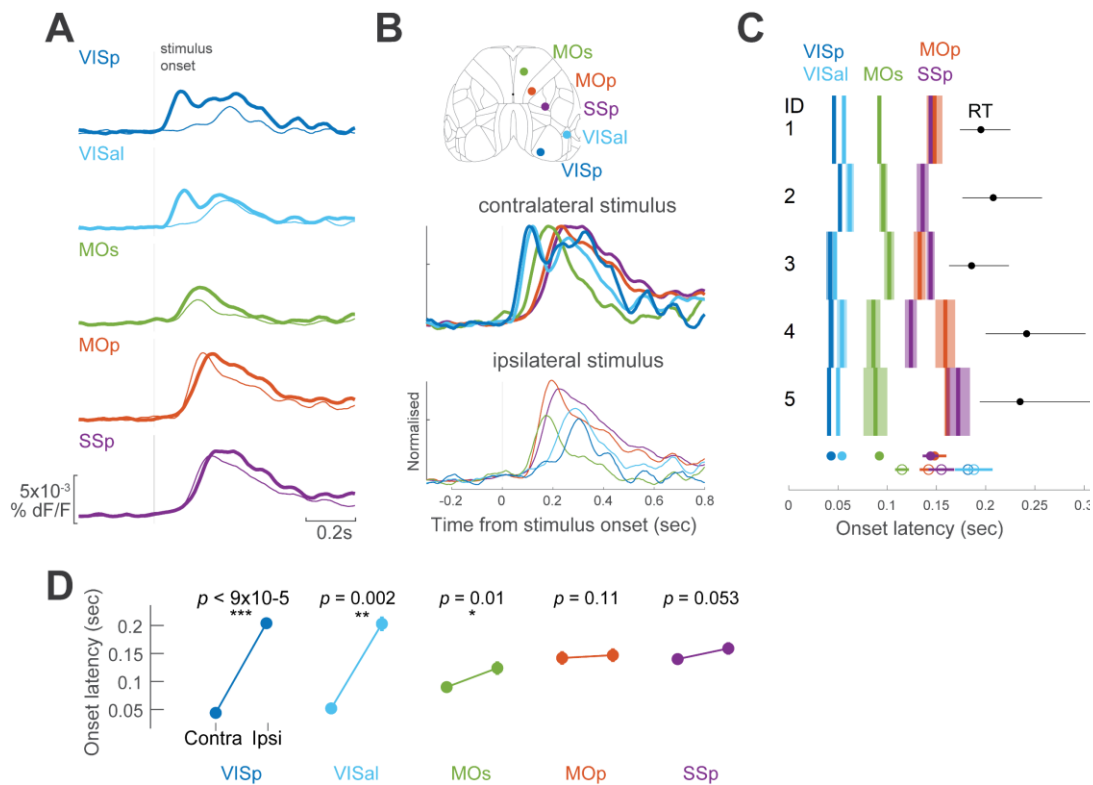


Figure 3-4 Stimulus-specific robust sequential activation of cortical areas

(A) Stimulus-locked calcium fluorescence at 5 ROIs for one example session. Vertical grey line marks the time of stimulus onset. Thick lines are the average response to contralateral stimuli over all correct trials. Thin lines are the same but for trials with ipsilateral stimuli.

(B) (Top) Position of 5 ROIs overlaid onto the Allen Common Coordinate Framework atlas (Lein et al., 2007). The normalised average fluorescence is shown for one session, for contralateral stimuli (middle) and ipsilateral stimuli (bottom).

(C) Summary of response onset latencies for 26 sessions in 5 mice. Each row contains data from a single mouse. Mice 1-2 correspond to Snap25-GCaMP6s, and mice 3-5 correspond to Camk2a-TTA x TetO-GCaMP6s. Onset latency is defined as the time from stimulus onset to fluorescence breaching 25% of the maximum. Solid vertical bars are the average latency across sessions, shaded regions are the median \pm median absolute deviation across sessions. For each mouse, the black line spans the 25th and 75th percentile of the reaction times pooled across sessions, and the black dot is the median reaction time. Filled coloured circles and lines are the average onset latencies for contralateral stimuli across mice (median \pm median absolute deviation). Open circles are the average onset latencies for ipsilateral stimuli across mice.

(D) Average latencies for trials with contralateral and ipsilateral stimuli across mice. Significance was determined by a paired t-test. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

3.5 Discussion

In this study we investigated the activation of multiple cortical areas in a two-alternative unforced choice visual discrimination task in the mouse. Widefield calcium imaging during task performance revealed a sequential activation of several cortical areas in this task, starting in visual, then secondary motor, and finally somatomotor areas.

Sensory information first arrives in VISp and higher visual areas on one hemisphere, then propagates to the same side MOs. This then passes cross-hemispherically producing bilateral MOs activation. After this, bilateral MOp and SSp regions become active and finally the VISp on the other hemisphere is activated. This sequence of events is consistent with a feed-forward signal from VISp to MOs, followed by a top-down feedback originating in MOs and moving down a cortical hierarchy. Sequential activation of several cortical areas has been demonstrated previously in monkey (Hernández et al., 2010; Ledberg et al., 2007; Siegel et al., 2015) and rodent (Le Merre et al., 2018) studies.

3.5.1 Caveats

There are a few caveats to consider in this work. Firstly, the fluorescence is a temporally-distorted measure of spatially-aggregate electrical activity in the neuron. Membrane depolarization in the dendritic tree, cell soma, or axon terminals, produces calcium influx through voltage-sensitive calcium channels (Grienberger and Konnerth, 2012). Calcium can also pass into the cell cytoplasm through calcium-permeable AMPA and NMDA receptors in the postsynaptic membrane or from intracellular stores. Intracellular calcium ions bind to GCaMP6s distributed throughout the cell membrane, which triggers conformational changes leading to seconds-long fluorescence when illuminated with light of a specific wavelength. The GCaMP6s fluorescence therefore peaks well after the trigger of calcium influx, and this delay corrupts any measure of absolute onset latency of activity. Nevertheless, it is still possible to compare the relative onsets between cortical areas, as all areas will presumably have the same delay.

Secondly, fluorescence at a single pixel reflects population activity of many thousands of cells. This therefore ignores differences between cells within a cortical region. However, as will be shown later using electrophysiological recording, putative excitatory cells (in the appropriate retinotopic position in VISp, and in MOs) show firing which monotonically increases with contrast. Therefore, among the VISp and MOs cells which respond to the stimulus, they all

seem to respond in the same way. In these cases, the population rate is therefore a suitable estimate of the average activity across the cells in these regions.

However, we have not recorded in VISal, MOp or SSp and therefore this may not hold in those regions.

Thirdly, the group-level analysis (Figure 3-4D) involves pooling fluorescence data across two mouse lines expressing GCaMP6s either in excitatory neurons alone, or in all neurons. This therefore neglects possible differences between excitatory and inhibitory neurons. Despite this problem, we did not observe qualitatively different fluorescence responses between the two mouse lines. One reason for this could be that excitatory and inhibitory neuron activity is correlated (e.g. Wehr and Zador, (2005)), which would give rise to similar fluorescence responses between the two mouse lines. A further possible reason for the similarity is that GABAergic neurons form ~20% of neocortical neurons in the mouse (Sahara et al., 2012). This smaller proportion of inhibitory neurons means the fluorescence signal is dominated by the excitatory neuron population. We therefore feel that this pooling is acceptable.

CHAPTER 4 A CAUSAL ROLE FOR VISUAL AND SECONDARY MOTOR CORTEX

Note: We discovered a problem with a dataset displayed in Figure 4-1, which resulted in having to repeat some experiments. The original data will be shown, and this problem will be pointed out. None of the conclusions drawn from this work are affected by this.

4.1 Introduction

We have seen in the previous Chapter that several cortical areas are robustly activated in a visual discrimination task. We have also shown that these regions show clear sequential activation. It is however unclear what the causal relevance of these signals are.

Despite several studies identifying clear task-related signals across many cortical areas, perturbation studies have produced some contradictory findings. For example, activity in monkey LIP predicts the choice in a motion discrimination task (Roitman and Shadlen, 2002), but muscimol inactivation of LIP has no effect on performance (Katz et al., 2016). Rat parietal cortex contains choice-predictive signals during an auditory accumulation task (Hanks et al., 2015), but muscimol inactivation does not impair auditory discrimination (Erlich et al., 2015). In a similar task, rat frontal orienting field (FOF) shows increasingly strong correlation to the upcoming choice during a memory phase (Kopec et al., 2015). However, optogenetic inactivation of FOF impairs performance most strongly during the earlier phase of the trial, when FOF activity has very weak predictive information for the upcoming choice. Therefore, while both observational and perturbation studies have provided insight into the role of different cortical areas, it remains ambiguous exactly which areas drive the decision process.

A further question concerns whether the causal role of different cortical areas is restricted to specific phases of the task. If cortical areas are causally relevant at distinct task phases, this would suggest a specific functional role for these areas in mediating the stimulus-decision-action sequence. By contrast, if cortical areas are causally necessary at the same time, this would suggest that decisions arise from a continuous reverberation of activity between the cortical regions. This has typically been explored in tasks with a delay period where the subject must hold their decision in memory before responding. In a memory-guided tactile detection

task, scanning optogenetic inactivation across the dorsal cortex reveals a restricted set of regions which are sequentially causal for the behaviour: barrel cortex during the sensation period, and frontal-motor area ALM during a memory period (Guo et al., 2014). In a visual detection task, the causal role of visual and parietal areas is restricted to the stimulus period, whereas the role of frontal-motor areas extends throughout the sensation, memory, and response periods (Goard et al., 2016). It is unclear whether the sequentially causal role for these areas reflects an intrinsic property of decision making, or instead arises because these task imposes a sequential structure (i.e. sensation then memory phases).

In this Chapter, we clarify the causal role cortical areas in this task using optogenetic inactivation. Inactivation of 52 sites across cortex highlighted a causal role for visual and secondary motor cortex but not primary motor and somatosensory areas, despite their large task-related activity observed in widefield calcium imaging. Using precisely-timed optogenetic inactivation of visual and secondary motor areas, we found distinct times in the trial when inactivation was effective, consistent with the timing of their responses as observed in widefield imaging. We also find distinct effects in choice behaviour when inactivating visual and secondary motor areas, which further suggests that these two regions contribute distinct roles towards the decision process. To further explore this, we modify the phenomenological model of Chapter 2 to model the effect of inactivation as perturbations to the bias and sensitivity parameters. We find that the distinct roles of visual and secondary motor areas can be attributed to different perturbations in biases and not stimulus sensitivity.

4.2 Methods

4.2.1 Behaviour

As before, the behavioural task used in this study is outlined in Chapter 2. However, there are a few differences. For the 52-coordinate inactivation experiment, trials began after a period of 0.2-0.6sec of no wheel movement. For all other inactivation experiments in this Chapter, there is no pre-trial quiescence period however trials are excluded post-hoc if wheel movement is detected -150 to +50 msec relative to stimulus onset.

4.2.2 Optogenetic inactivation

Ai32xPV-cre transgenic mice expressed Channelrhodopsin-2 (ChR2) in Pvalb-expressing cortical inhibitory interneurons. ChR2 is a light-activated ion channel

which non-selectively passes positively charged ions. Photo-illumination therefore depolarises the interneurons, which in turn inhibits local activity in excitatory neurons via release of GABA (Guo et al., 2014). This line was acquired from B6;129P2-Pvalbt(cre)Arbr/J crossed with Ai32.

While mice performed the task, we optogenetically inactivated several cortical areas through the skull using a blue laser. For the 52-coordinate inactivation experiment, we used 5 mice. For the remaining inactivation experiments, we used 6 mice: in the delayed-onset inactivation experiment we used 2 of the mice, in the pulse inactivation experiment we used all 6 mice, and for the higher-power inactivation experiment we used 5 of the 6.

In the 52-coordinate experiment, unilateral inactivation was achieved by mounting a fibre-optic cable on a moving manipulator. On every trial, the manipulator set the position of the fibre-optic cable to one of 52 different coordinates distributed across the cortex. Inactivation coordinates were defined stereotaxically from bregma. On ~75% of trials, the laser was switched on (473nm, 1.5mW, 40Hz sine wave) to inactivate the cortical site. Laser and non-laser trials, and the location of the cortical inactivation, was randomised. The duration of the laser was from visual stimulus onset, until a behavioural choice was made. For any given session, a single cortical site on the inactivation grid may only be inactivated a handful of times. This discouraged any adaptation effects that may occur on more frequent inactivation paradigms, however this approach does require combining data across sessions. The laser positioning was independent of laser power, so noise from the manipulator did not predict inactivation.

In subsequent inactivation experiments, a pair of mirrors mounted on galvo motors were used to orient the laser (462nm) to different points on the skull. We also introduced improved light isolation to ensure no light could reflect from the skull surface and be seen by the mouse. In the pulse inactivation experiment, we inactivated only visual (-4mm AP, \pm 2mm ML, coordinates relative to bregma) and secondary motor areas (+2mm AP, \pm 0.5mm ML), using a brief 25msec pulse at 15mW power. The onset time of the laser pulse was set randomly trial by trial, ranging -300ms to +300msec relative to stimulus onset. In the higher-power inactivation experiment, we inactivated visual and secondary motor areas, as well as primary motor areas (-0.5mm AP, \pm 0.5mm ML) for a fixed duration 1.5seconds, 40Hz sine wave using different laser powers (1.5, 2.9, 4mW).

4.3 Visual and secondary motor areas causally contribute in this task

We first asked whether the cortical areas activated in the widefield calcium imaging were also causally necessary for the task. To address this question, we performed scanning unilateral and bilateral optogenetic inactivation across the dorsal surface of the cortex while mice performed the task (Figure 4-1A). A laser illuminated the cortex through a clear skull preparation. Due to the sparse nature of the inactivation, each cortical site was inactivated for only a few trials on each session. It was therefore necessary to average across sessions to compare the effect of inactivation across many areas on behaviour. Unilateral and bilateral inactivation induced robust changes to choices at specific cortical sites.

Specifically, we observed a reduction in contraversive choices when inactivating visual and secondary motor areas unilaterally (Figure 4-1B). For example, for trials with high contrast on the left screen, unilateral inactivation of right VIS and right MOs impaired the mouse's ability to respond Left. This pattern was also observed when inactivating the left hemisphere VIS and MOs on trials with stimuli on the right side. In summary, the optogenetic inactivation acted to impair the mouse's ability to choose the action associated with the contralateral stimulus.

Interestingly, VIS inactivation also increased ipsiversive choices, even on trials where there is no stimulus on the ipsilateral side. This VIS effect is inconsistent with a simple blinding of the animal to the stimulus, as blinding should only act to increase NoGo choices. This effect was less apparent for MOs inactivation.

On equal contrast stimuli, where mice would normally respond Left and Right about equally, we observed robust changes to the choices away from the contralateral side when inactivating VIS and MOs (Figure 4-1B and Figure 4-1D; VIS: -0.35 $p < 0.0002$, MOs: -0.22 $p < 0.0002$; permutation test, shuffling laser trial identities within 91 sessions in 5 mice). We also observed an increase in choices towards the ipsilateral side of inactivation (Figure 4-1E; VIS: $+0.35$ $p < 0.0002$, MOs: $+0.18$ $p < 0.0002$).

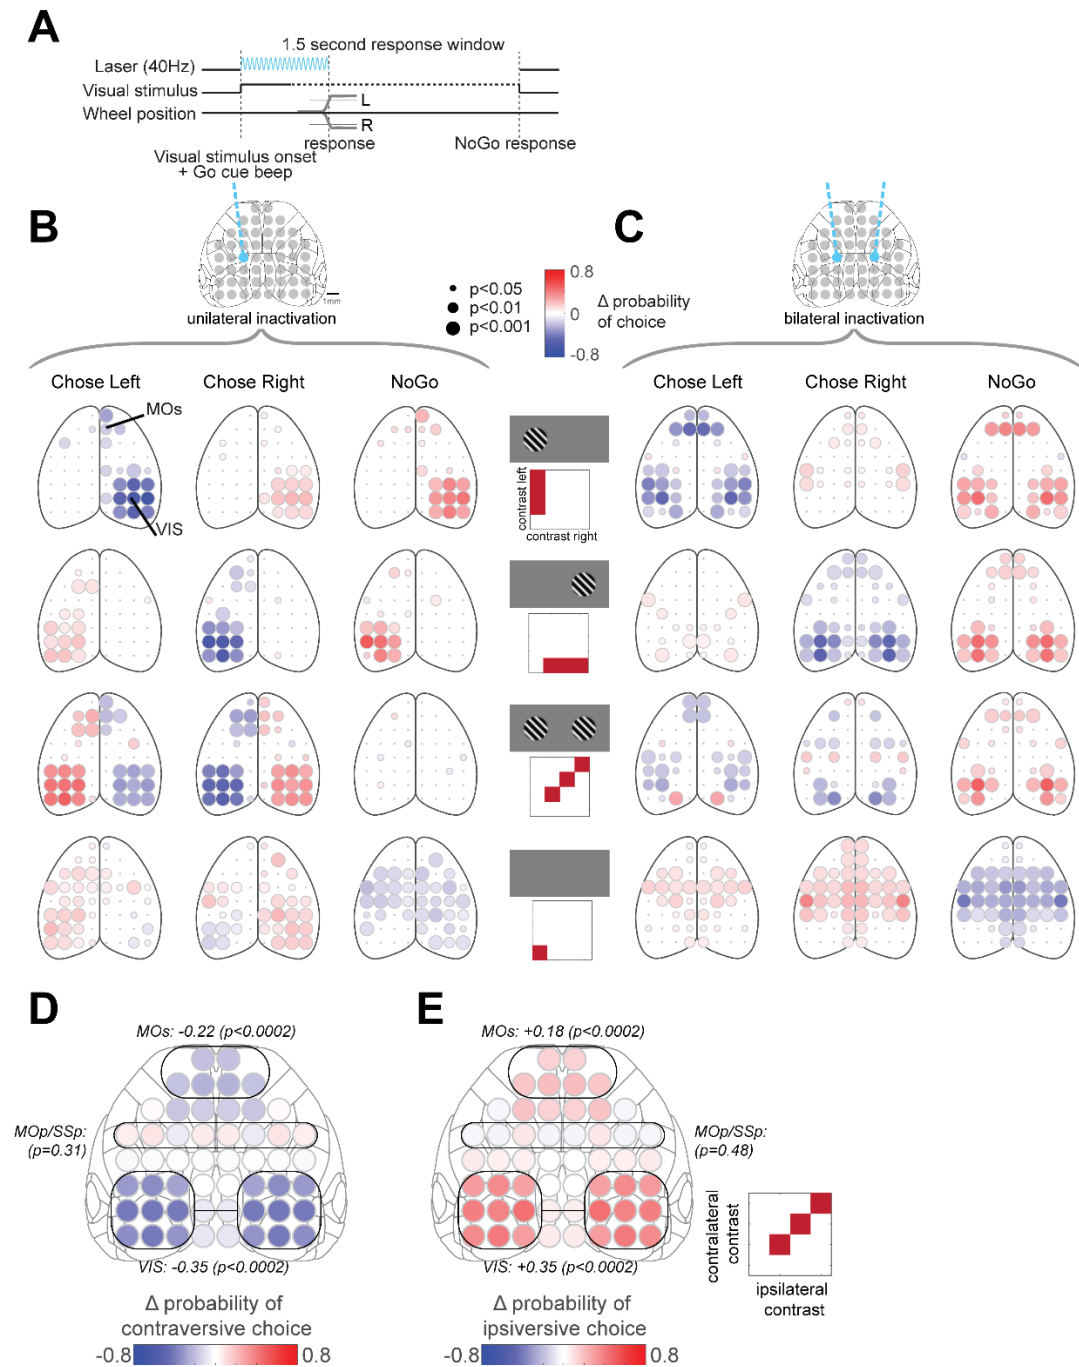


Figure 4-1 Unilateral and bilateral optogenetic inactivation of cortex on choice behaviour

(A) Timeline of laser inactivation during the trial. On ~75% of trials, a 473nm 40Hz 1.5mW laser was switched on during stimulus presentation and ceased when a response (or NoGo) was registered. The position of the laser varied randomly trial-to-trial over 52 different cortical sites.

(B) Each colourmap shows the change in the fraction of Left, Right and NoGo choices (columns) induced by unilateral laser inactivation averaged across 91 sessions in 5 mice. Each row shows the map for different stimulus conditions. The stimulus condition is represented by a square matrix schematic in the central column. The size of each dot reflects statistical significance based on a shuffle test where the identity of laser and no-laser trials is shuffled.

(C) Same plotting scheme as in (B) but for bilateral inactivation, showing

average data over 76 sessions in 4 mice. Note: the effect in primary motor/somatosensory cortex on zero contrast trials is artefactual (see text). (D) Summary map of the effect of unilateral inactivation on contraversive choices on trials with equal contrast on each side. The colourmap reflects the change in the probability of making a contraversive choice (Left choices for right hemisphere inactivation, and Right choices for left hemisphere inactivation). The change in proportion of contraversive choices is plotted symmetrically across the hemispheres. Black lines mark the grouping of coordinates used to test statistical significance. Significance is tested by shuffling the identities of laser and non-laser trials within each session. (E) Same plotting scheme as in (D) but plotting the change in ipsiversive choices (Left choices for left hemisphere, and Right choices for right hemisphere).

Bilateral inactivation of VIS and MOs also impaired the mouse's ability to respond to the contralateral stimulus (Figure 4-1C). However, the behavioural effect was not lateralised: bilateral inactivation of VIS and MOs reduced choices towards the contralateral stimulus and increased NoGo choices. These findings are consistent with an inter-hemispheric competition process driving a Left-Right decision. Unilateral inactivation imbalances this comparison, whereas bilateral inactivation does not.

We also observed an effect on reaction time. Inactivation of VIS and MOs increased reaction times for choices made towards the contralateral side of inactivation (VIS: +44.2msec $p < 0.0002$, MOs: +47.6msec; permutation test, $p = 0.0004$), and decreased reaction time for choices made towards the ipsilateral side (VIS: -33.5msec $p < 0.0002$, MOs: -3.9msec $p = 0.03$). We observed a weakly significant decrease in reaction time when inactivating primary motor and primary somatosensory areas (-12msec for contra choices $p = 0.003$ and -12msec for ipsi choices $p = 0.0014$), however they were in the same direction which suggested that the reaction time effect here was creating something like a startle response.

We observed a puzzling effect where inactivation of primary motor and somatosensory areas decreased NoGo choices. Unilateral inactivation increased ipsiversive choices, while bilateral inactivation increased both Left and Right choices. We hypothesised that inactivating somatosensory cortex was unpleasant to the mouse, and mice subsequently move the wheel to terminate the laser. We therefore performed a new experiment where we inactivated this region of cortex for a fixed 1.5second duration so mice couldn't terminate the laser early by moving the wheel. However, we still observed a decrease in NoGo choices (data not shown), which invalidated our original hypothesis. We discovered later that

improper light isolation was the cause, because improving light isolation removed the somatosensory effect (while the VIS and MOs effects remained). We discovered later that shining the laser light onto the somatosensory cortical region produced reflections of the laser light onto the ceiling of the experimental rig, possibly reflecting off the clear-skull implant. We propose therefore that mice were able to see this reflection and were influenced by it particularly on zero contrast trials because that reflection was the only visual “stimulus” present to the animal on that trial.

Due to the light-isolation problem, we repeated the inactivation experiment but with added light isolation to prevent laser light reflecting off the surface of the clear-skull implant. Since the 52-coordinate experiment required a lot of data, we opted to not inactivate all 52 coordinates again. Instead we focused on inactivating VIS, MOs and primary motor region MOp. We also modified the experiment to address some extra questions: 1) The results above suggest that primary motor areas play no role in this task. Is this because primary motor inactivation was not sufficiently strong, or because primary motor really has no effect? 2) Is the reason that VIS and MOs inactivation does not impair behaviour completely because the inactivation is incomplete or because there is redundancy with a non-inactivated structure elsewhere? To address these extra questions, we inactivated these regions with higher laser power (Figure 4-2A).

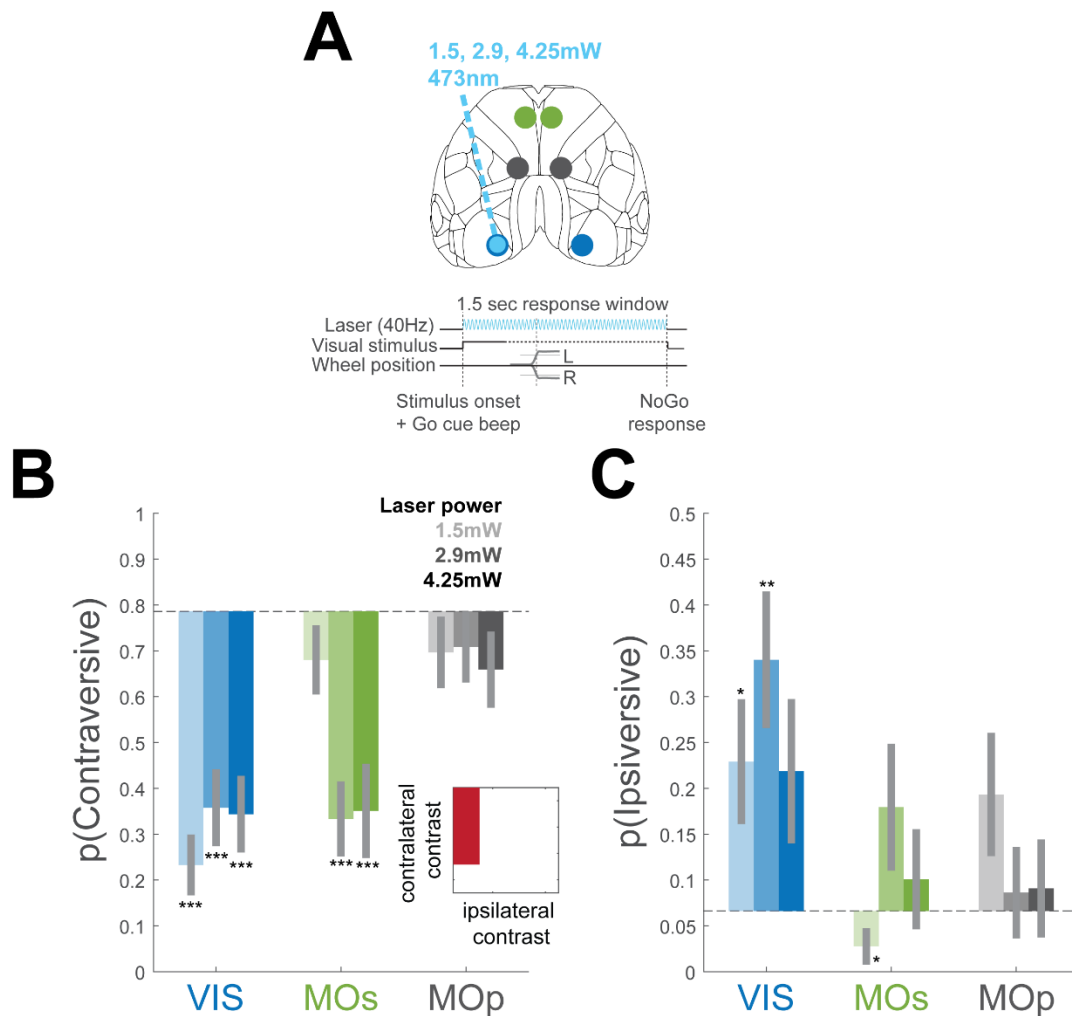


Figure 4-2 Higher power optogenetic inactivation

(A) Schematic of higher-power fixed-duration inactivation, focused on VIS (blue), MOs (green) and MOp regions (grey). Inactivation was performed at higher laser powers: 1.5, 2.9 and 4.25mW.

(B) The probability of moving contraversive to the inactivated hemisphere, on trials with visual stimuli only present on the side contralateral to the inactivated hemisphere. The dashed grey line represents the session-averaged (34 sessions in 6 mice) non-laser probability of moving to the correct side indicated by the stimulus. Bar values represent the session-averaged probability of moving towards the side indicated by the stimulus, on trials when inactivating the contralateral VIS (blue), MOs (green) and MOp (dark grey) at different laser powers (denoted inset). Error bars are the standard error in probabilities across sessions. Statistical significance for the change in the proportion of contraversive choices from the laser off condition was determined with a paired t-test. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

(C) Same plotting scheme as in (B) but plotting the probability of moving ipsiversive to the inactivated hemisphere, on trials with visual stimuli only present on the contralateral side. The dashed grey line represents the session-averaged non-laser probability of moving to the opposite side than what is indicated by the stimulus.

Unilateral inactivation of VIS and MOs, but not MOp, impaired the mouse's ability to respond to the contralateral stimulus (Figure 4-2B). Inactivating MOp at high

power had no effect, which suggests that it does not play a role in this task. For VIS, the effect was similar across laser powers, which suggests that VIS inactivation has saturated for these laser powers. For MOs inactivation, the reduction in contraversive choices was much stronger when inactivating at 2.9mW but this reduction did not increase at 4.25mW. Inactivating both VIS and MOs at high power (4.25mW) therefore impaired performance but this performance reached a floor at about 30% correct. This suggests that VIS and MOs do not individually completely determine the behavioural performance on the task.

As before, inactivation of VIS also significantly increased ipsiversive choices even on trials where there was no ipsilateral stimulus to respond to (Figure 4-2C). This was not present for MOs inactivation; however, it was weakly present in the previous experiment (Figure 4-1B), which may be explained by the improper light isolation before. These findings suggest that VIS and MOs are contributing different roles in the decision behaviour. We explore this in later sections.

4.4 The causal role is sequential in time

Visual and secondary motor areas therefore play a causal role in this task. We next asked whether this causal role was restricted to specific phases within the trial timing. To address this, we performed optogenetic inactivation of VIS and MOs at different times in the trial. This was explored with two experiments.

In the first experiment, we inactivated VIS and MOs with the same laser power and waveform as in the 52-coordinate experiment (1.5mW 40Hz sine wave for 1.5seconds), but the onset time of the laser was shifted to different time-points after stimulus onset. We found that as the onset time of inactivation was shifted later into the trial, the effect on performance was reduced. However, the time-course over which the performance impairment was reduced differed between VIS and MOs (Figure 4-3). The behavioural impairment was approximately half of its maximum when inactivating VISp at 50msec after stimulus onset. By contrast, the impairment was half of its maximum when inactivating MOs at 100msec after stimulus onset.

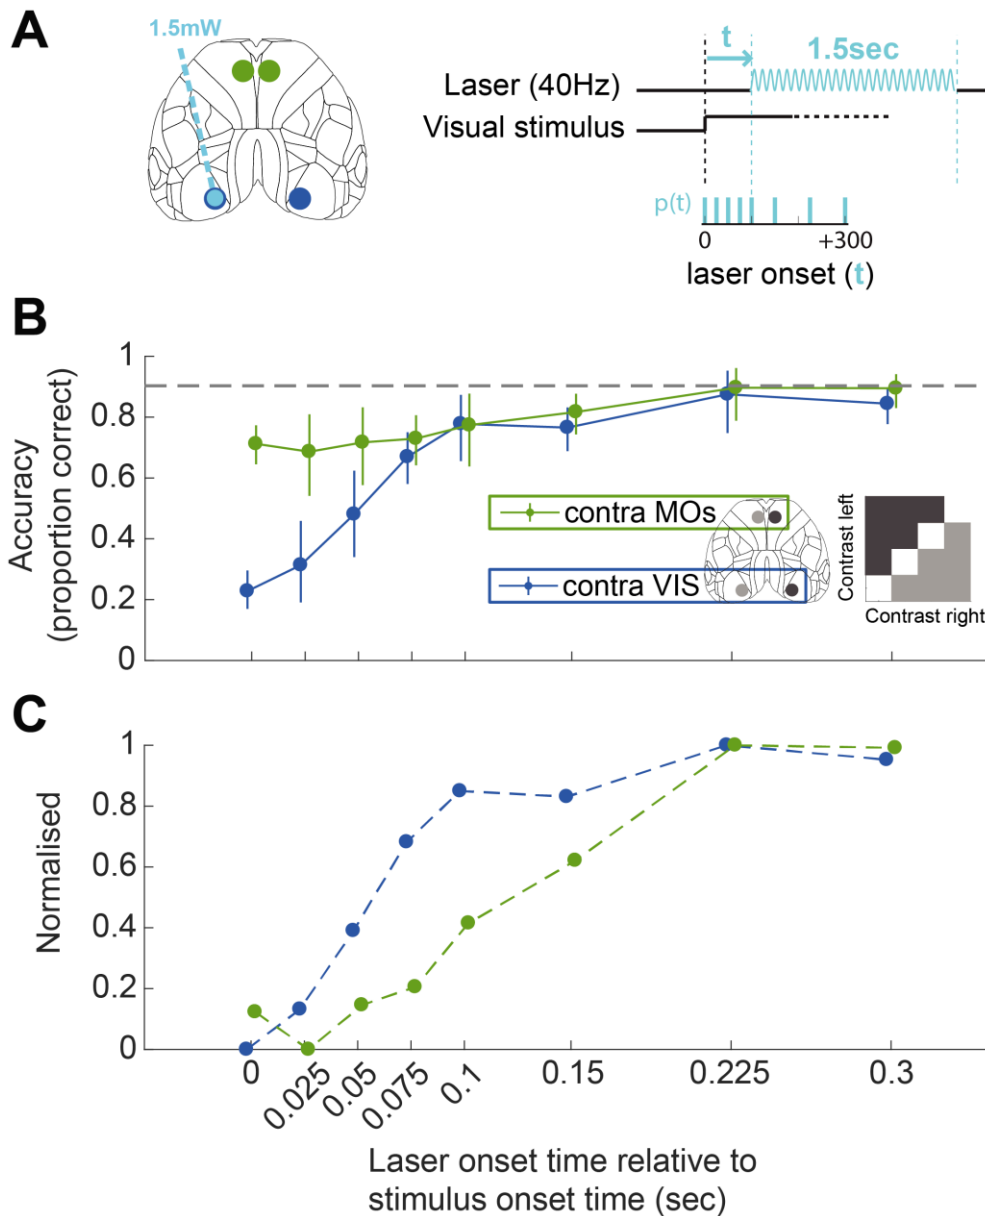


Figure 4-3 Optogenetic inactivation at different onset times

(A) Schematic of the inactivation experiment. VIS and MOs are inactivated with a 1.5second laser waveform (1.5mW 40Hz), but the onset time of the laser is chosen randomly among a set of delays after stimulus onset

(B) The blue dots and connecting line reflect the data from VIS inactivation, and green dots for MOs inactivation. The dots are the empirical performance and the error bars are the 95% binomial confidence intervals for the proportion of correct choices, pooling data over 23 sessions in 2 mice. Performance is calculated on trials where visual contrast is higher on one side than the other, and the contralateral hemisphere is inactivated. For example, 'contra VIS' comprises left VIS inactivation on $c_R > c_L$ trials, with right VIS inactivation on $c_L > c_R$ trials. The grey dashed line represents the non-laser performance.

(C) The same trace as in (B) but normalised by the maximum performance impairment.

In the second experiment, we inactivated VIS and MOs with a brief 25msec pulse at a high power of 15mW, and the onset time of the pulse was set randomly trial by trial between -300 to +300msec relative to stimulus onset. Since the laser onset time was drawn randomly trial by trial, a single time-point was only inactivated once. Therefore, it was necessary to smooth data across time-points and pool data across sessions and mice. We observed a characteristic impairment in performance and reaction time when briefly inactivating VIS and MOs. Inactivation of both VIS and MOs impaired accuracy and reaction time, but at different time-windows during the trial (Figure 4-4). Inactivation of VIS had a significant effect on the percentage of correct choices around the time of stimulus onset (-110 to +130 ms), while inactivation of MOs was significant only after stimulus onset (+52 to +174ms). Inactivation of both VIS and MOs also delayed the production of responses by up to 100ms. Again, the time windows for which MOs inactivation could have this effect were later than for VIS (-34 to +79ms for VIS, and +32 to +130ms for MOs).

The effect of pulsed laser inactivation on neural activity was not confined to the 25msec during laser presentation. In a separate experiment performed by Nick Steinmetz, VISp neuron activity was recorded extracellularly (Methods in Chapter 5) while the cortical region was illuminated with brief 10msec 4mW laser pulses. The brief laser pulse silenced activity in VISp for a sustained period up to ~100msec after pulse onset (Figure 4-5). We also observed rebound activity from 130 to 200msec after pulse onset which increased neural firing to about 2.5 spikes/second (2.5x baseline firing). Note that this rebound activity was much weaker than activity observed in VISp in response to visual stimuli which ranged from 8 spikes/second for weak stimuli to 30 spikes/second for strong stimuli (see later Figure 5-2). Given this long-lasting effect of pulsed illumination on the neural activity, the causal time-windows identified above are broader than what is probably the true case. Instead the true critical window corresponds to the late phase of the causal windows measured using the pulse inactivation. This corresponds approximately to +75 to +125msec for VIS and +125 to +175msec for MOs. The effect of inactivation at earlier times can be explained by a long-tailed effect of the pulse illumination which lasts sufficiently long to be active during the true (later) critical windows.

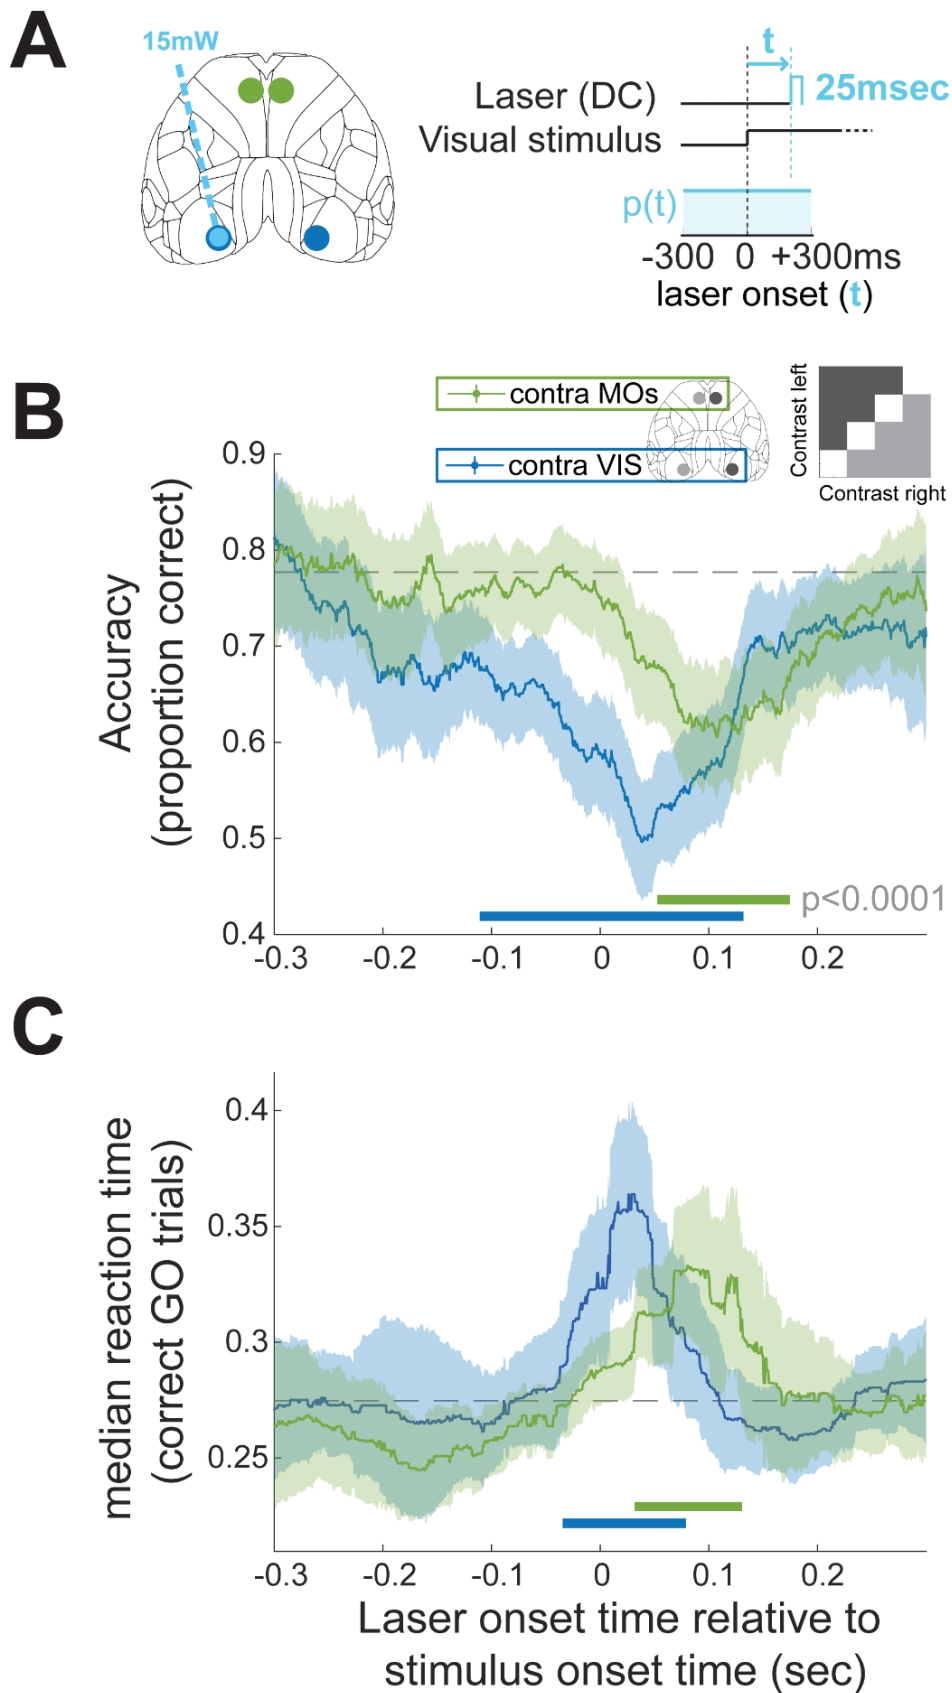


Figure 4-4 Pulsed high-power optogenetic inactivation experiment

(A) Schematic of the inactivation experiment. VIS and MOs regions were inactivated with a brief 25msec pulse at different times within the trial. The

onset time of the laser pulse was set randomly between -300 to +300msec relative to stimulus onset.

(B) The effect of pulsed inactivation on behavioural accuracy, on trials pooled over 65 sessions in 6 mice. Dashed-grey line is the accuracy on trials without inactivation. Solid blue and green lines reflect the accuracy on trials with inactivation in VIS and MOs regions respectively. Shaded regions are the 95% binomial confidence intervals for the proportion estimate. Left-hemisphere inactivation trials on $c_R > c_L$ contrast conditions are pooled with right-hemisphere inactivation trials on $c_L > c_R$ contrast condition. Data is smoothed with a 100msec sliding window. Solid blue and green bars indicate the timepoints when the accuracy is significantly different from the non-laser condition, based on a chi2 test of independence with false alarm rate of 0.0001.

(C) The effect of pulsed inactivation on reaction times for correct Go (Left/Right) trials. Shaded regions are the 95% confidence interval. Solid bars reflect the timepoints when the reaction time is significantly longer than the non-laser reaction time, based on a one-tailed Wilcoxon rank sum test with false alarm rate of 0.0001.

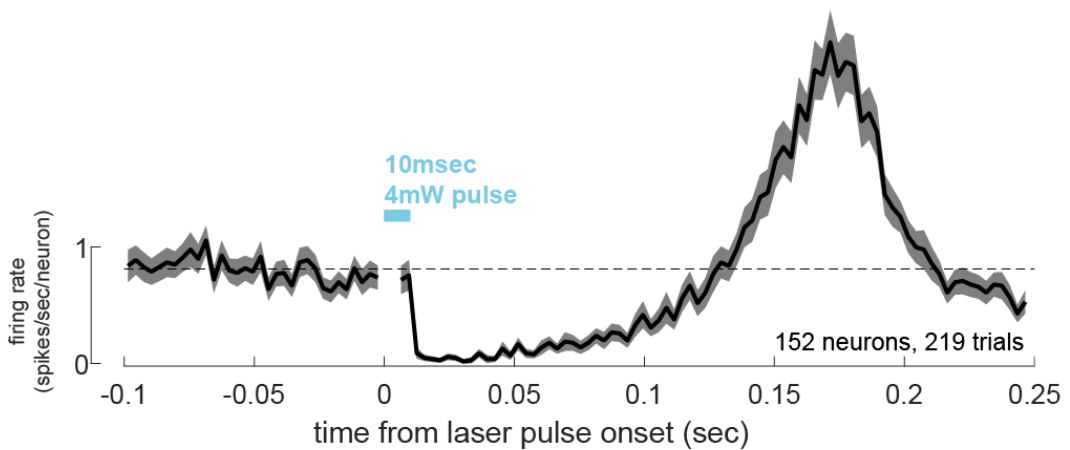


Figure 4-5 Effect of pulse laser illumination on cortical firing rate

The effect of pulsed inactivation on firing rates in a population of 152 broadly-spiking VISp neurons recorded over 219 repeated trials of inactivating with a 10msec 4mW pulse. Spiking data was acquired using Neuropixels probes (see Methods in Chapter 5). The firing rate is shown aligned to the time of laser pulse onset, averaged over 219 trials and 152 neurons. A laser pulse was used to briefly silence activity in these neurons. The firing rate at the beginning of the pulse onset is masked due to the presence of light-induced artefacts in the electrodes. The presentation of a brief 10msec inactivation pulse produced prolonged suppression of firing lasting approximately 100msec. This was met with rebound activity which lasted from about 130 to 200msec after pulse onset. This rebound activation peaked at about 2.5 spikes/sec. This was much smaller than the firing rate evoked by weak-contrast stimuli (8 spikes per second, see Figure 5-2 in Chapter 5).

Taken together, these results suggest that VIS and MOs contribute towards the decision process at different times. VIS has a critical window earlier than MOs. The MOs window is still earlier than movement onset by ~125-175msec. Interestingly, the shift in the critical window between these two regions is approximately 50msec, which matches the delay in activation observed between VISp and MOs from the widefield imaging data in the previous Chapter. Therefore, the sequential activation of these areas, beginning in VISp and then progressing to MOs, reflects a causal process which drives the decision task. By contrast, the late bilateral activation of MOp and SSp doesn't appear to contribute towards task performance.

4.5 A phenomenological model of the inactivation effect

So far, we have seen that VIS and MOs are sequentially causal in this task. A next question we asked was whether these two cortical areas contribute a distinct functional role for this behaviour. One way to explore this question is to ask whether the effect of inactivation on behaviour could be expressed as the perturbation to specific parameters of the behavioural model outlined in Chapter 2. The current model form comprises bias and sensitivity parameters. If optogenetic inactivation of VIS and MOs was accounted for as perturbations to different parameters in the behavioural model, that would suggest a functionally distinct role for these areas within the framework of the model.

As in Chapter 2, the probability of choosing Left (π_L), Right (π_R) and NoGo (π_{NG}) on each trial i is modelled with two expressions. The log odds of Left vs NoGo and the log odds of Right vs NoGo. Each expression is parameterised with a bias, sensitivity and shape parameter for the corresponding hemifield,

$$\ln\left(\frac{\pi_L^{(i)}}{\pi_{NG}^{(i)}}\right) = \mathcal{B}_L^{(i)} + \mathcal{S}_L^{(i)} \times (c_L^{(i)})^{\mathcal{N}^{(i)}}$$

$$\ln\left(\frac{\pi_R^{(i)}}{\pi_{NG}^{(i)}}\right) = \mathcal{B}_R^{(i)} + \mathcal{S}_R^{(i)} \times (c_R^{(i)})^{\mathcal{N}^{(i)}}$$

Where the prior distribution on the bias ($\mathcal{B}_L, \mathcal{B}_R$), sensitivity ($\mathcal{S}_L, \mathcal{S}_R$), and shape (\mathcal{N}) parameters are defined within a hierarchical framework containing per-subject and per-session deviations from the grand average values.

The effect of optogenetic inactivation was captured with an augmented model (Figure 4-6),

$$\ln\left(\frac{\pi_L^{(i)}}{\pi_{NG}^{(i)}}\right) = (\mathcal{B}_L^{(i)} + \Delta b_L^{(i)}) + (\mathcal{S}_L^{(i)} + \Delta s_L^{(i)}) \times (c_L^{(i)})^{\mathcal{N}^{(i)}}$$

$$\ln\left(\frac{\pi_R^{(i)}}{\pi_{NG}^{(i)}}\right) = (\mathcal{B}_R^{(i)} + \Delta b_R^{(i)}) + (\mathcal{S}_R^{(i)} + \Delta s_R^{(i)}) \times (c_R^{(i)})^{\mathcal{N}^{(i)}}$$

Where Δb_L , Δb_R , Δs_L , and Δs_R are new parameters which capture changes to the bias and sensitivity parameters when inactivating one cortical region. Each inactivated region has its own set of 4 delta parameters. The delta parameters are not contained within the per-session and per-subject hierarchical priors, and therefore act as fixed effects on the grand average values for the bias and sensitivity parameters. The model estimates of these delta parameters are interpreted as changes to the bias and sensitivity perceptual states as a result of optogenetic inactivation.

The delta parameters for each cortical region were given a regularizing prior around zero,

$$(\Delta b_L, \Delta b_R, \Delta s_L, \Delta s_R) \sim \text{Normal}(0, 2^2)$$

We fit this model to the behavioural data acquired on sessions with 1.5seconds of inactivation at higher laser powers. Trials with different laser power were pooled together.

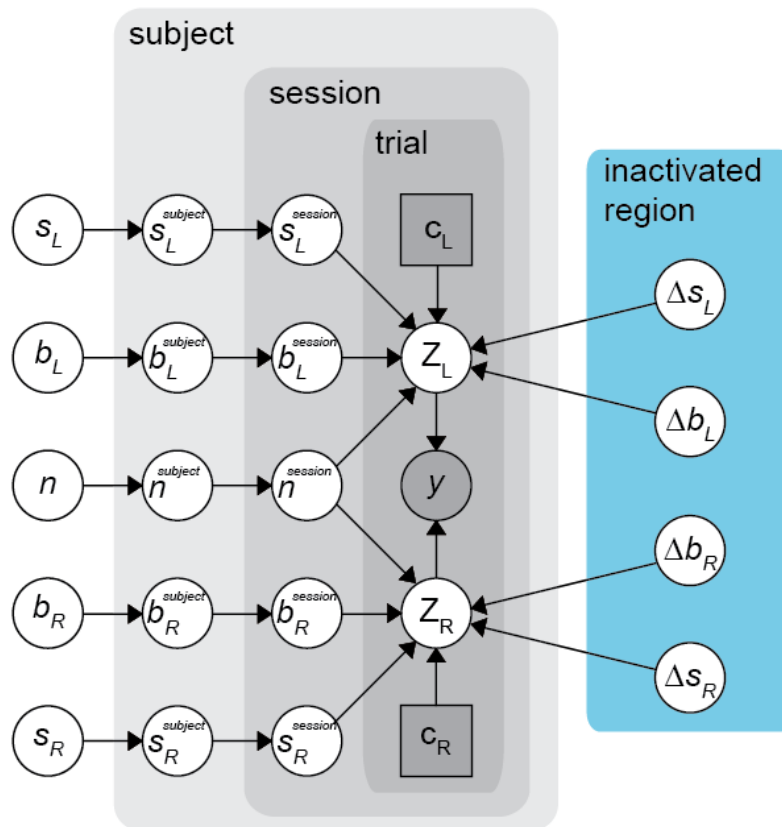


Figure 4-6 A probabilistic graphical model of the optogenetic inactivation effect

A graphical depiction of an augmented hierarchical model which accounts for the effect of optogenetic inactivation as perturbation to the bias and sensitivity parameters. In grey is the standard hierarchical model as outlined in Chapter 2. In blue are four new parameters for each inactivation region, reflecting the perturbation to bias and sensitivity parameters from their non-laser value.

The model was able to recapitulate the behavioural data during the trials with inactivation (Figure 4-7A-B). The behavioural effect of VIS and MOs inactivation could be captured with similar reduction in the sensitivity term to the contralateral stimulus (Figure 4-7C). By contrast, there were distinct effects on the bias parameters. Unilateral VIS inactivation decreased the bias associated with the contraversive choices (e.g. Left VIS decreased the log odds of Right vs NoGo) and increased the bias associated with the ipsiversive choices (e.g. Left VIS increased the log odds of Left vs NoGo). However, MOs inactivation decreased bias for both the contraversive and ipsiversive choices. This pattern was also observed when fitting an alternative model where the parameter perturbations were

contained within a hierarchy which modelled the per-subject variation in how the laser affected the parameter perturbations.

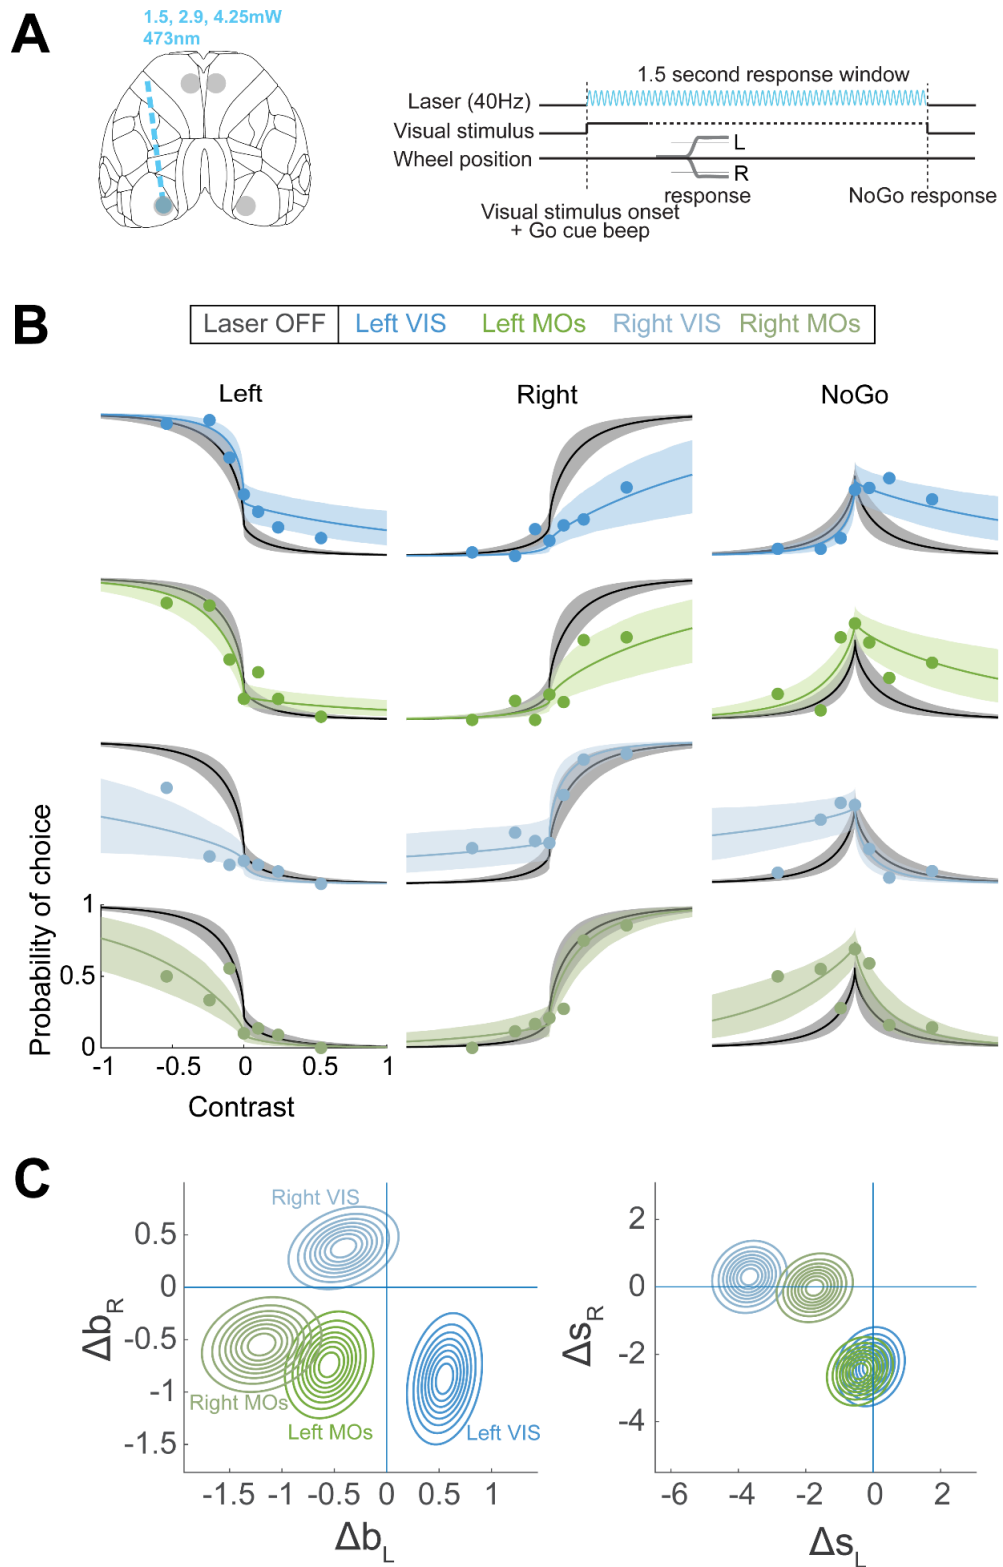


Figure 4-7 Modelling the effect of optogenetic inactivation

(A) Schematic of the inactivation experiment.

(B) Effect of inactivating VIS and MOs on the psychometric curve. Showing only

the subset of contrast conditions where the minimum contrast on either side is zero (zero pedestal). Dots are the empirical probabilities (averaged across sessions) for the choices on the detection stimulus subset. Black line and shaded region are the grand average model prediction with 95% credible intervals for non-inactivated trials. Coloured line and shaded region are the model fits for different inactivated regions: left VIS (dark blue), left MOs (dark green), right VIS (light blue), right MOs (light green).

(C) Posterior distribution of the bias (left) and sensitivity (right) parameter perturbations. The contours represent the fit of a 2D gaussian distribution to the samples of the posterior distribution acquired by MCMC.

These results suggest that VIS and MOs contribute distinct roles towards shaping the decision. VIS (or a structure downstream) contributes towards a competitive selection between Left and Right choices. This is because VIS inactivation primarily modifies behaviour by inducing a Left-Right bias. By contrast, MOs inactivation is more concerned with engaging the action informed by the stimulus, because inactivation of this region reduces Go choices but still in a manner which is sensitive to the contralateral stimulus.

4.6 Discussion

In this Chapter we have investigated the causal relevance of different regions in the mouse cortex. While mice performed the visual discrimination task, we optogenetically inactivated several cortical regions. We identified a causal role for visual and secondary motor areas in this task. We also demonstrated a sequential causal role for VIS and MOs using precisely-timed optogenetic inactivation. We observed two important features from the inactivation experiments. Firstly, inactivation did not totally abolish performance, which suggests that decision formation is not localised to these areas alone. Secondly, inactivation of VIS induced an increase in ipsiversive choices even when no ipsilateral stimulus was present. This suggests a distinct role for VIS compared to MOs, which we clarified by relating the inactivation effect to perturbations of parameters in the behavioural model. Inactivation of visual areas was associated with a Left-Right choice biasing, whereas inactivation of secondary motor areas was associated with basing a Go-NoGo decision process.

Note that these inferences were only made possible by our task design. Previous studies exploring the role of cortical areas during decision making have largely relied on Go/NoGo (Goard et al., 2016; Le Merre et al., 2018; Sachidhanandam et al., 2013) or 2AFC (Erlich et al., 2015; Hanks et al., 2006; Kopec et al., 2015;

Licata et al., 2017; Tsunada et al., 2016) task designs. Had we employed a 2AFC design, we could not have observed the effect where VIS and MOs inactivation suppresses contraversive choices, while only VIS inactivation increases ipsiversive choices. This is because in 2AFC designs, suppression of contraversive choices is equivalent to enhancement of ipsiversive choices by definition. This is apparent in previous studies. For instance, it is unclear whether the reduction in contralateral responses following FOF inactivation in Erlich et al. (2015) resulted from a decrease in propensity to make contralateral responses or an increase in propensity to make ipsilateral ones, or both.

The causal relevance of sensory and frontal-motor areas to different times in the decision process has been studied primarily in tasks with a memory-component (Goard et al., 2016; Guo et al., 2014; Kopec et al., 2015; Li et al., 2016). These memory tasks impose a sequential structure to the behaviour, which makes it difficult to ascertain whether the role of these cortical areas in decision making is necessarily sequential. Some evidence suggests that frontal-motor areas play an early role in the decision process (Gold and Shadlen, 2000; Sul et al., 2011). Our study shines light on this early role of frontal-motor areas during decision tasks, by showing that the role is indeed early (~100ms after stimulus onset), but sequentially after sensory cortex.

Is MOs involved in motor production? Microstimulation of frontal-motor areas in rats leads to a combined set of eye, eyelid, vibrissae, and head movements consistent with orienting (Donoghue and Wise, 1982; Hall and Lindholm, 1974). This rodent cortical area therefore interfaces with circuitry generating complex behavioural patterns. MOs neurons project directly to the spinal cord (Wang et al., 2018) which would be consistent with a role in motor production. However, in our experiment, inactivation of MOs does not paralyse wheel movements. Instead the wheel is turned in the wrong direction, and reaction times are increased. Therefore, perhaps frontal-motor areas are more related to decision formation than motor production. In Chapter 3 we found an early stimulus-specific activation of MOs. In this Chapter we found that inactivation of MOs early into the trial affected performance. This is consistent with previous findings where inactivation of frontal-motor areas during the stimulus phase of a visual detection task impairs performance (Goard et al., 2016). These findings suggest that MOs may have more of a sensory role and/or a role in motor planning related to an action informed by the stimulus.

MOs may have multiple roles depending on context. For example in posterior parietal cortex, neurons show multiplexed tuning to sensory-related and action-related task variables (Raposo et al., 2014). Perhaps MOs neurons similarly contain multiplexed representations of visual contrast and choice action in our task. Alternatively, different MOs neurons within the same region may confer different functions. In a self-initiated action task in the rat, different spatially-intermixed MOs neurons showed firing properties characteristic of different components of an accumulation-to-bound model. Some neurons were active with sustained firing which gradually changed over the trial, consistent with an accumulation signal. Other neurons were transiently active and correlated with the action, which is consistent with the input evidence pulses provided to the accumulator (Murakami et al., 2014).

MOp also appears to have no role in this task. MOp lesions have been shown to abolish fine motor skills of distal muscles while leaving proximal trunk muscles unaffected (Gharbawie et al., 2005). MOp inactivation in this task may have no effect because mice can still move the wheel using proximal trunk and forelimb muscles.

4.6.1 Caveats

There are several important caveats to this work. During the optogenetic inactivation experiments, we were limited by the spatial precision of the inactivation effect. In our experiments using 1.5mW laser power, this effect (defined as the region where cortical activity is at most half of baseline) has been estimated to be 1.2mm radius from epicentre (Guo et al., 2014). This large spread means we could not target different visual areas, nor target posterior parietal cortex independently of visual cortex or somatosensory cortex. Accordingly, we limit our interpretations of inactivation to all “visual areas” and do not distinguish early from higher visual areas. For the pulsed inactivation experiment, since the laser duration was reduced to 25msec, it was necessary to increase the laser power to 15mW to achieve any behavioural effect. This larger power may have larger spatial spread. Given this uncertainty, the different critical time-windows we observe between VIS and MOs locations may be more conservatively interpreted as different time-windows for the causal role of the frontal and occipital lobes.

We also did not perform any control experiments using mice which were ChR2-negative (e.g. a YFP control). This was primarily due to the labour-intensiveness

of the training and running experiments. However, the fact that inactivation of some cortical regions showed no significant effects on choice behaviour could be considered a negative control. It remains possible that laser illumination of secondary motor areas may penetrate the brain tissue and illuminate the retina directly, which could induce behavioural changes. However, we feel this is unlikely for two reasons. Firstly, pulsed-inactivation of secondary motor areas was effective during a window much later than would be expected if mice were seeing the laser light directly. Secondly, other studies which have photo-illuminated frontal-motor areas found that this was effective only for ChR2-positive mice (Guo et al., 2014).

A more general problem with studies invoking perturbations relates to how much we can interpret causality from these effects. One hope of perturbation studies is to identify components of a mechanistic causal chain which drives behaviour. However, behaviour is a process emerging from a complex dynamic whole (brain, body, and environment). It may be fallacious to attempt to ascribe the true causes of behaviour to only a part of this whole (a part of the brain), as behaviour only exists on the level of that whole (Gomez-Marin, 2017; Schaal, 2005). Related to this, the massively recurrent nature of brain connectivity means optogenetic perturbation of one part of this network may have widespread indirect effects across the brain (e.g. Otchy et al. (2015)). Likewise, other brain areas may compensate even for brief perturbations (e.g. Li et al. (2016)). Furthermore, subjects may engage slightly different motor plans or behavioural strategies in solving the task, which may translate to differences in the causal roles of different brain areas (Churchland and Kiani, 2016). These features make it difficult to derive meaningful conclusions from perturbation studies alone. The interpretation can be aided however by combining insight from observational and theoretical work. This will be explored in the next Chapter.

Our task has a fixed stimulus-action contingency. To disambiguate whether the optogenetic inactivation was impairing something about stimulus perception, or something about action selection, the stimulus-choice contingencies would need to be flipped. For example, if we observed the same inactivation result on trials where mice were trained with a flipped contingency (i.e. choose left when the contrast is higher on the right side), this would indicate that the inactivation effect operated primarily by affecting stimulus perception. If instead the

inactivation was affecting action selection, then we would expect the inactivation effect to change.

CHAPTER 5 A MECHANISTIC MODEL OF THE CORTICAL CONTRIBUTION TO VISUAL DISCRIMINATION

Note: Electrophysiological recordings were performed by Nick Steinmetz. My contribution in this Chapter is in data analysis and the model.

5.1 Introduction

So far, we have seen that visual and secondary motor areas are causally relevant for mouse visual discrimination. We have also shown data which hints at distinct roles for these two areas based on optogenetic inactivation. However, it is still unclear what the precise contribution is for activity in these areas towards forming the decision.

To address this problem, it is useful to map neural structures to components of a mechanistic theory of decisions. This approach has generated insight in previous studies using evidence accumulation tasks. Activity in monkey and rodent frontal and parietal areas has been linked to components of the drift-diffusion model (Erlich et al., 2015; Hanks et al., 2015; O'Connell et al., 2018; Shadlen and Newsome, 2001). For example, activity in parietal areas behaves like an accumulation signal, and activity in frontal areas behaves like a discretisation of the accumulator. However, as previously discussed, these models are of a 2AFC task design, which confounds several types of biases. Therefore, it is unclear what the mechanistic contributions are for cortical areas in a task without these confounds. Furthermore, it is unclear what the contribution may be for a task which does not require explicit integration of stimuli over time.

To ascertain the functional role of these sequentially-causal areas, we developed a mechanistic model of the cortical contribution of VIS and MOs towards forming a decision. In this Chapter we outline the structure of the mechanistic model, and we show that this model can predict average behaviour, trial-by-trial choice variability, and the effect of optogenetic manipulation. This model proposes distinct roles for visual and secondary motor areas in their contributions towards forming the decision.

5.2 Methods

5.2.1 Electrophysiological recording in VISp and MOs

Extracellular spikes were recorded using high-density Neuropixels probes (Phase 3A option 3; Jun et al., 2017) while mice performed a similar behavioural task. Mice were implanted with a clear-skull preparation as outlined in the Methods of Chapter 2. On the day of the recording or several days before, mice were anaesthetised, and a small craniotomy was performed through the clear skull cap at the site to be recorded, on the left hemisphere. The craniotomies for VISp were targeted in some cases using measured retinotopic maps in the same mice, and in other cases to the same position stereotaxically (-4AP, -1.7ML). The craniotomies for MOs were targeted stereotaxically (+2AP, -0.5ML). During the recording, mice were head-fixed, and the skull was covered with saline-based agar and silicon oil. Probes were inserted through the agar and into the brain, advanced at $\sim 10\mu\text{m}/\text{sec}$, and allowed to settle in the final position for 10-15 minutes before recording. The probes were externally referenced to an Ag/AgCl wire in the agar above the skull. Recording was performed with SpikeGLX software, and probes were configured for gain=500 in the AP band and 250 in the LFP band. Multi-channel voltage data was spike-sorted using Kilosort (Pachitariu et al., 2016) and manually curated with phy (<https://github.com/kwikteam/phy>). Cluster quality was assessed by visualising the waveform variation across spikes, auto- and cross-correlation of spiking over time, and spike amplitude. Putative excitatory cells were identified by broad spike waveform (Niell and Stryker, 2008).

5.3 A mechanistic model

We hypothesised that the distinct effects of VIS and MOs inactivation (Chapter 4) could be accounted for by optogenetic suppression of baseline firing rate. If downstream brain structures are subtractively comparing activity in right and left visual cortex, then even without any visual stimulation, suppressing baseline activity on one side could lead to increased ipsilateral responses. To test this hypothesis, we attempted to build a mechanistic model in which choices are made based on a weighted sum of cortical activity. The model was a modification of the phenomenological model introduced in Chapter 2, except the logistic component of the model is now taken literally as a hypothesis for the mechanistic causal processes underlying decision formation.

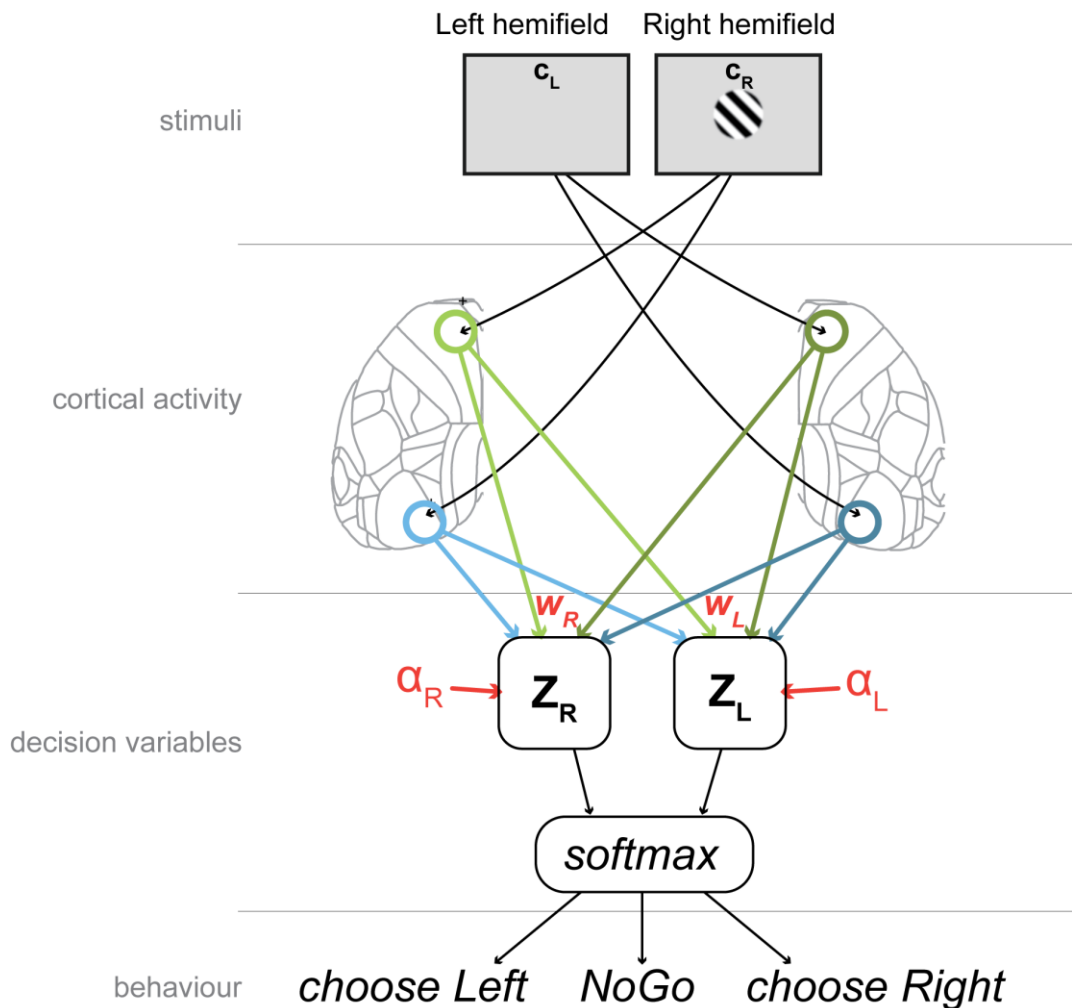


Figure 5-1 Schematic of the mechanistic model

Sensory stimuli on the left and right hemifields induces activity in contralateral VISp and MOs. Activity from both hemispheres is linearly summed (with connection weights w_R and an offset α_R) to form the Z_R decision variable, representing the log odds of Right vs. NoGo. The Z_L decision variable is similarly determined by activity in the four regions but weighted differently by w_L and with a different offset α_L . The offset parameters α_L and α_R capture all un-modelled effects which may otherwise shape the decision. The decision variables are pooled together via a competitive process (e.g. softmax) to determine the probability mass associated with each choice. A probabilistic decision process then selects one action. Note that the model is mechanistic with respect to the cortical activity, but leaves the instantiation of the decision variables, softmax computation, and probabilistic decision process to an abstract part of the model.

In this model, two decision variables (log odds of Left vs. NoGo and log odds of Right vs. NoGo) are computed from a weighted sum of neural activity in VIS and MOs (Figure 5-1),

$$\ln\left(\frac{\pi_L^{(i)}}{\pi_{NG}^{(i)}}\right) = \alpha_L^{(i)} + (\vec{f}^{(i)})^T (\vec{w}_L^{(i)})$$

$$\ln\left(\frac{\pi_R^{(i)}}{\pi_{NG}^{(i)}}\right) = \alpha_R^{(i)} + (\vec{f}^{(i)})^T (\vec{w}_R^{(i)})$$

Where \vec{f} is a 4-element vector containing the firing rate of each of the cortical areas: left VISp, right VISp, left MOs and right MOs. \vec{w}_L and \vec{w}_R are 4-element vectors of parameters, representing the weighting of the firing rates onto the two log odds ratio decision variables Z_L and Z_R respectively. α_L and α_R are offset parameters which estimates the effect of all other unmodelled brain regions onto the decision process. The two decision variables then feed into a softmax computation downstream, which computes the probability of choosing each action. This then goes through a decision process to select the action based on the computed probabilities. Our model is focused primarily on the cortical contribution towards the decision variables. The weightings of the neural activity onto the decision variables are therefore informative for the cortical contribution of that area to the decision. To improve fit stability, we enforced parameter symmetry between the left and right hemispheres (e.g. the weight of left VISp onto the log odds of Left vs. NoGo was the same as right VISp onto Right vs. NoGo). As in the phenomenological model, all parameters are modelled in a hierarchical Bayesian framework, capturing per-subject and per-session variation in the weights and offset parameters. The hyperpriors are the same as in the hierarchical model of Chapter 2. The global parameter α_L , α_R , \vec{w}_L and \vec{w}_R are given a weakly-regularising $Normal(0,4^2)$ prior.

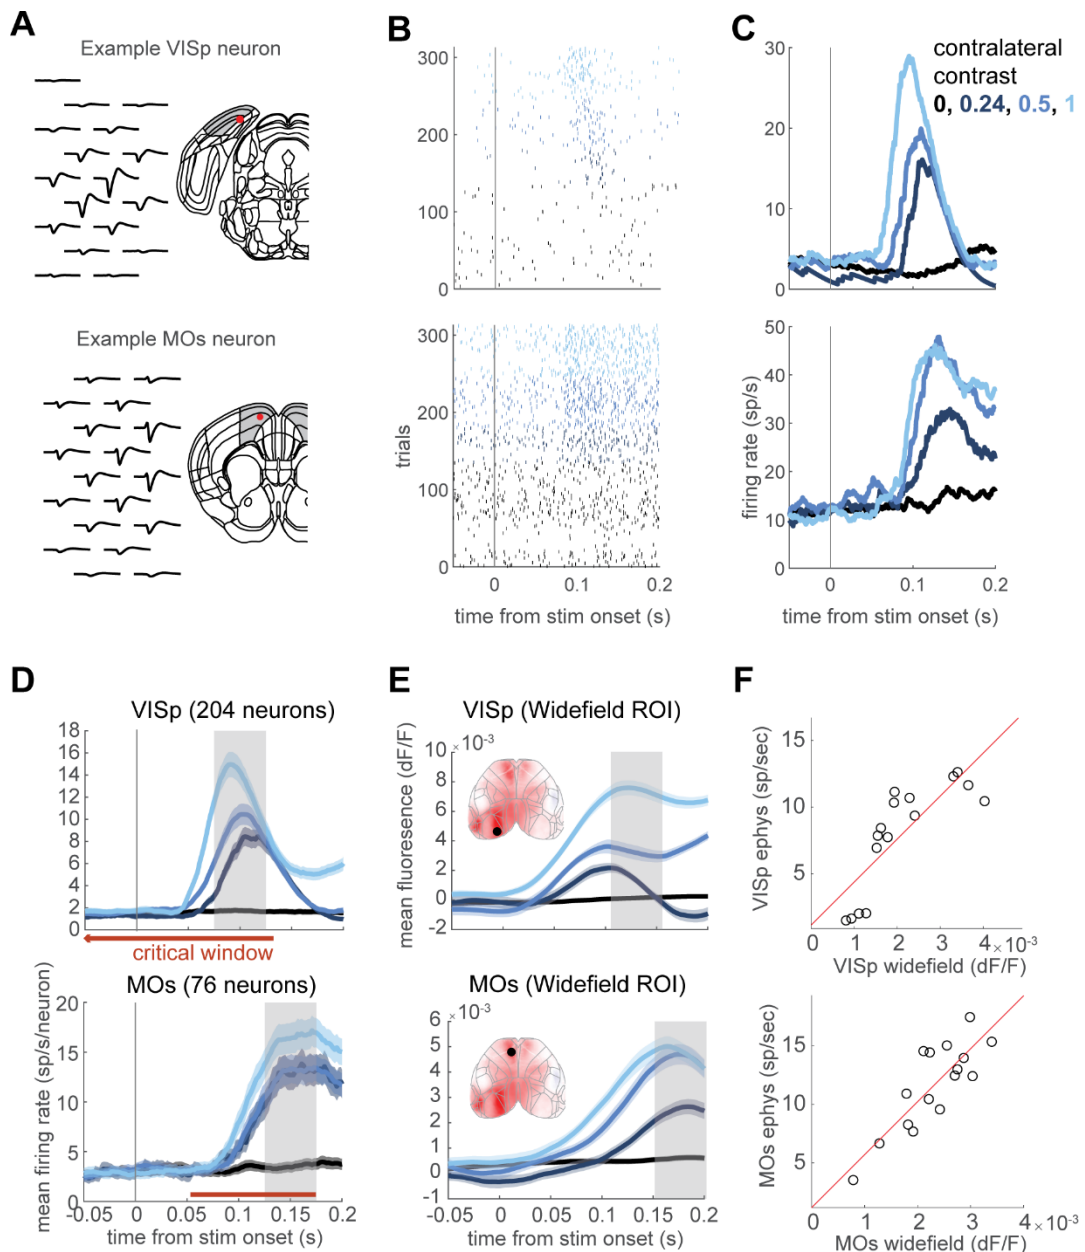


Figure 5-2 Electrophysiological recording of broadly-spiking VISp and MOs neurons in response to visual Gabor stimuli

(A) Example neuron in left VISp and left MOs. The waveforms are shown in black and the red dot marks the location of the neuron within an aligned Allen CCF atlas.

(B) Raster plots, showing the spiking activity aligned to stimulus onset. The colour reflects the contrast level presented to the contralateral hemifield.

(C) PSTHs showing an estimate for the firing rate for the example neurons, averaged across trials.

(D) PSTHs averaged over 204 neurons in VISp and 76 neurons in MOs. The shaded regions mark the time window when the firing rate is averaged for subsequent analyses: 75-125msec for VISp, and 125-175msec for MOs. Horizontal red line marks the time of the critical time window identified in the pulse inactivation experiment (Figure 4-4).

(E) Trial-averaged fluorescence of left VISp and left MOs ROIs in response to stimuli present on the contralateral side. Shaded regions mark the time

windows used for averaging in subsequent analyses. This window is 30msec after the window associated with electrophysiological data, to compensate for the slow onset of GCaMP6s.

*(F) Correlation of window-averaged fluorescence and firing rates for left VISp and left MOs. The 16 open circles correspond to the 16 possible contrast conditions (averaged over trials). Red line corresponds to the fit of a simple linear model $f(x) = b_0 + b_1*x$. The linear model is used to transform widefield fluorescence data into estimates of population firing rate.*

The neural activity contained within $\vec{f}^{(i)}$ is the estimated population firing rate of the four cortical regions on each trial. We estimate the population firing rate from widefield calcium fluorescence at the VISp and MOs ROIs described in Chapter 3. Since the baseline activity in widefield calcium fluorescence is poorly defined, the calcium fluorescence was calibrated to approximate population firing rates. To achieve this, we recorded extracellular spiking activity in VISp and MOs using Neuropixels probes in separate sessions, and computed trial-averaged firing rates for each of the contrast conditions over the end of the critical time-window identified from the pulse inactivation experiment (Figure 5-2A-E; VIS: 75-125msec, MOs: 125-175msec). Calcium fluorescence was also averaged over the same windows but 30msec later to allow for slower GCaMP6s kinetics. The transformation of widefield fluorescence to firing rate was computed by simple linear regression over the 16 contrast conditions between these two datasets (Figure 5-2F). This linear transformation was then applied to the fluorescence value for each individual trial, thereby providing an estimate of the population firing rate of the four cortical regions on every trial.

We fit the mechanistic model weights w_L , w_R and offsets α_L , α_R to the behavioural data using the trial-by-trial widefield activity in the four regions calibrated to population firing rate. The fitted model was able to capture the average probability of choices across different stimulus conditions (Figure 5-3A). The parameter weights indicated a distinct role for VISp and MOs in determining the value of the contraversive and ipsiversive decision variables (Figure 5-3B). Each hemisphere VISp had a positive weight onto the contraversive decision variable (e.g. left VISp was positively weighted onto Right vs. NoGo) and a negative weight onto the ipsiversive decision variable (e.g. left VISp was negatively weighted onto Left vs. NoGo). By contrast, each hemisphere MOs was positively weighted towards both the contraversive and ipsiversive decision variables.

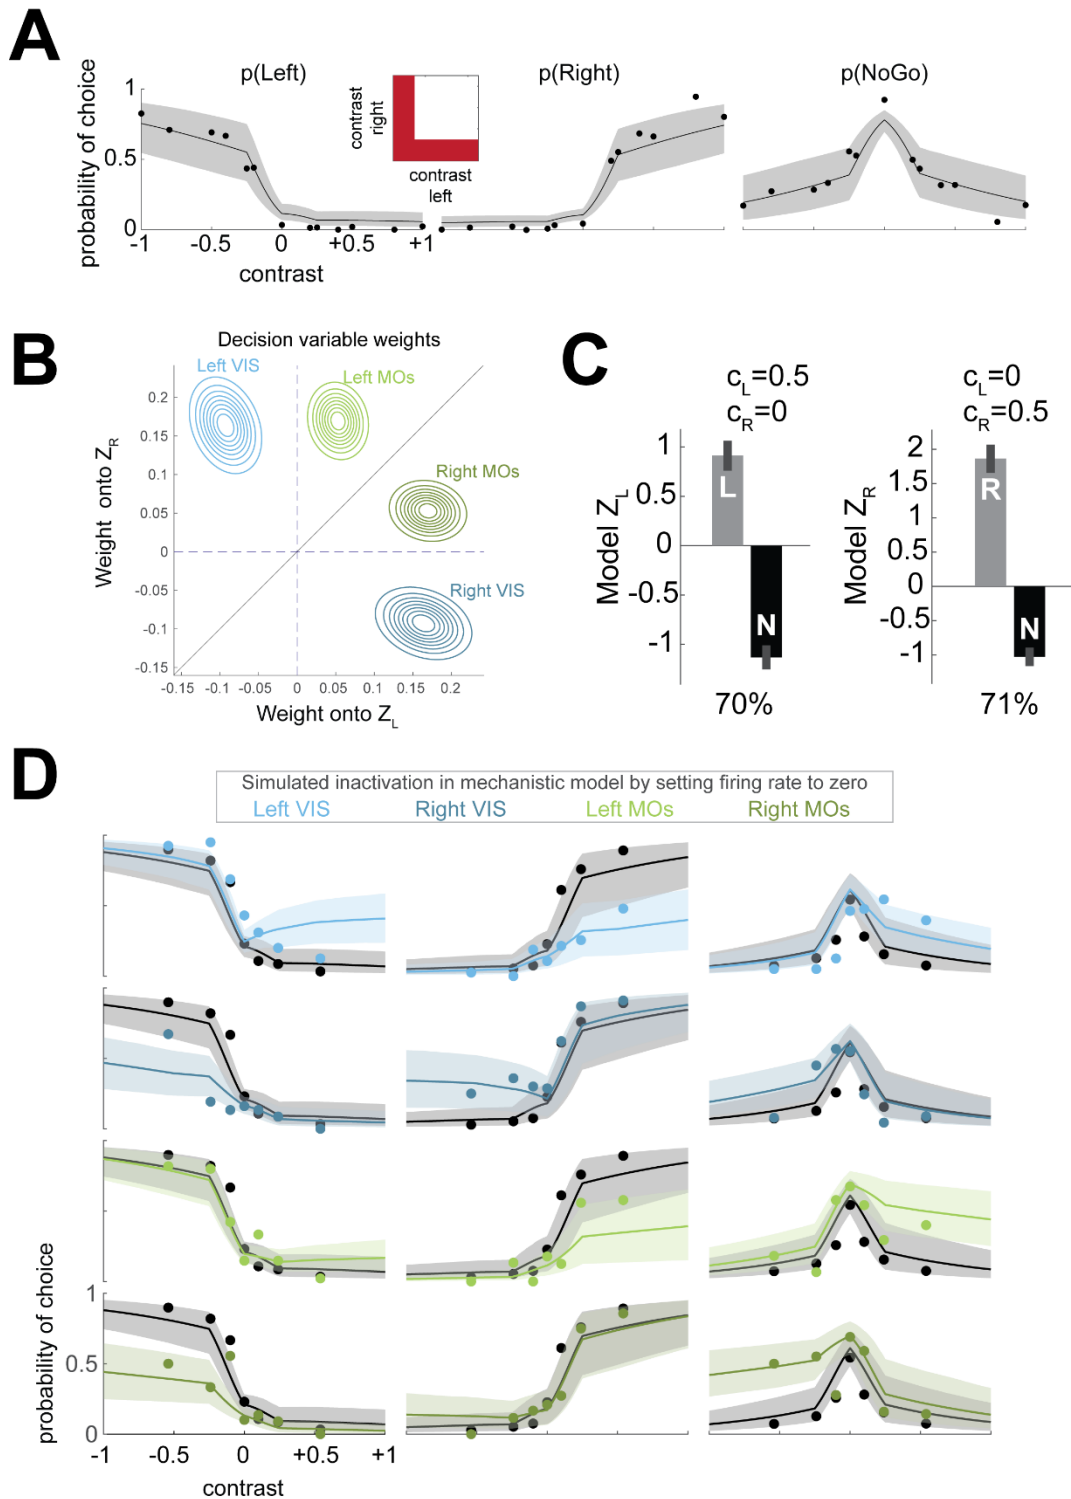


Figure 5-3 Mechanistic model predicts behaviour

(A) Model fit to the psychometric data. The probability of moving Left, Right and NoGo is shown as a function of contrast. A subset of contrast conditions is plotted, where a stimulus is only present on one side (shown inset). Positive contrast values reflect stimuli on the right side, and negative values reflect stimuli on the left side. Black dots are the session-averaged empirical fraction of choices made for each contrast conditions. The black lines and shaded regions are the posterior means and 95% credible intervals for the prediction of the model from the widefield data. For each contrast condition, the model

prediction is computed using the fitted parameters and a linearly-interpolated measure of the cortical firing rate, interpolating from the measured firing rates for the contrast conditions that were empirically tested.

(B) Posterior distribution of the mechanistic model weights for each of the four regions, left VIS (dark blue), left MOs (dark green), right VIS (light blue) and right MOs (light green) onto the two decision variables Z_L and Z_R . The parameters are constrained to be symmetrical about the diagonal (solid grey line). Lines reflect the contours of a 2D Gaussian distribution fit to MCMC samples of the posterior distribution.

(C) Trial-by-trial choice stochasticity within a single contrast condition is reflected in trial-by-trial variation in the decision variables. (Left) On trials with medium contrast on the left only, the value of the decision variable is shown for trials when mice chose Left (L) or NoGo (N). Bar height is the average of the posterior means for each of the trials, and error bar is the standard error across trials. Data is pooled across sessions and subjects. Below the plot is the cross-validated % correct of a logistic classifier which predicts the choice from the decision variable. The mechanistic weights are fit on half of the data, and classification accuracy is measured on the other half. (Right) Similar plot but for medium contrast on the right only, showing the value of for trials when mice chose Right (R) or NoGo (N).

(D) Simulating inactivation in the model by setting the neural activity to zero. The model fit is shown using the same plotting convention as in (A). Black dots are the session-averaged empirical fraction of choices made for the subset of stimulus conditions shown. Black lines and shaded regions are the model fit to the non-laser behavioural data. Coloured lines show the behavioural prediction when setting the neural activity of left VIS (dark blue), left MOs (dark green), right VIS (light blue), or right MOs (light green) to zero within the model. Coloured dots are the session-averaged empirical fraction of choices made when optogenetically inactivating these regions.

The model was also able to partly account for choice stochasticity for trials with identical stimulus conditions. For example, on trials with medium contrast on one side only, spontaneous variation in the decision variable associated with that side correlated with the choice eventually made by the mouse (Figure 5-3C). Trial-by-trial choice variability under identical stimulus conditions could therefore be partly attributed to variability in cortical activity which generates variability in the decision variables.

We further verified that the estimated weights reflected more than a correlation between neural activity and choice, by showing that simulating data from the model could predict the behavioural effect of optogenetic inactivation. To demonstrate that the mechanistic model could predict behaviour during optogenetic inactivation, it was first necessary to re-fit the offset parameters α_L and α_R to the non-laser trials contained within optogenetic inactivation sessions. For this purpose, we used the sessions which involved fixed-duration 1.5sec inactivation in VIS and MOs (Figure 4-2). Re-fitting the offset parameters was

necessary because the overall tendencies towards Go/NoGo were different between behavioural sessions during widefield imaging, and the non-laser trials during optogenetic inactivation sessions.

Using the re-fit offset parameters, and the weights estimated from the widefield dataset, we simulated behavioural predictions from the model for different contrast conditions. To compute predictions of the mechanistic model for different contrast conditions, the firing rate vector was set as the linearly-interpolated trial-averaged activity as measured previously (Figure 5-2). To simulate the effect of optogenetic inactivation in one of the four regions, the corresponding element of the firing rate vector was set to zero. We found that the behavioural prediction of this model could capture the primary features of the effect of optogenetic inactivation on real behaviour, despite these inactivation trials playing no role in the original model fit (Figure 5-3D). In this sense, the mechanistic model could cross-predict behaviour in a new dataset which it wasn't fit to. This demonstrates that the weights estimated from the mechanistic model reflect more than a correlation between neural activity and behaviour, as these weights also predict the effect of neural manipulation. These results therefore propose a mechanistic causal process by which cortical activity in VISp and MOs is integrated towards forming a decision.

5.4 Discussion

To bring together the findings in previous Chapters, we developed a mechanistic model of the cortical contributions towards decision formation. The model identified different roles for VISp and MOs in shaping the decision process: VISp activity enhances contraversive choices and suppresses ipsiversive choices, whereas MOs activity enhances both contraversive and ipsiversive choices. This model was able to predict average choice behaviour, as well as trial-by-trial variation in choice for identical stimulus conditions. Finally, we showed that the model could emulate the effects of optogenetic inactivation on choice behaviour.

5.4.1 Mechanism

The model makes testable predictions for the behavioural impact of perturbation at different stages of the mechanistic process. In our experiment, the effect of VIS and MOs inactivation is accounted for as a reduction in weighted input to the contraversive and ipsiversive decision variables. Therefore, stimulating activity in these regions should produce opposite effects: VIS stimulation should enhance

contraversive choices and suppress ipsiversive choices, whereas MOs stimulation should enhance both contraversive and ipsiversive choices. Similarly, the model makes predictions for the effects of multi-site perturbation. This model also provides a framework for understanding the effect of learning and serial dependence in choice. For example, learning may correspond to changes in cortical firing to stimuli (Hua et al., 2010) which propagates to the decision variables leading to changed behaviour. Similarly, serial dependence in choices could arise from serial dependence in the cortical inputs to the model. Indeed, PPC activity has been shown to represent and be causally involved in serial dependence of behaviour (Akrami et al., 2018).

The model also proposes the existence of two decision variables which arise from downstream convergence of activity from visual and secondary motor areas. These variables may be localised to particular subcortical structures or may be encoded in a distribution population. One possibility is the striatum, a structure implicated in decision tasks and orienting-type movements (Balleine et al., 2015, 2007; Jin and Costa, 2015; Tecuapetla et al., 2014; Wickens et al., 2007; Znamenskiy and Zador, 2013). The different roles of VIS and MOs in contraversive and ipsiversive actions could be mediated by specific projections in the striatum. For example, VIS could enhance contraversive choices by exciting the ipsilateral direct pathway, and could suppress ipsiversive choices by exciting the contralateral indirect pathway, or lateral inhibition within striatum. Likewise, MOs could enhance both types of movements by projecting to contralateral and ipsilateral direct pathways. Thalamus may also be a candidate structure, as thalamo-cortical loops have been implicated in decision tasks (W. Guo et al., 2017; Z. V. Guo et al., 2017; Halassa and Kastner, 2017). Another possibility is the midbrain superior colliculus, which has been implicated in visually-guided behaviour (Basso and May, 2017; Crapse et al., 2018; Kopec et al., 2015), and exhibits contralateral suppression (Lo and Wang, 2006) consistent with VIS activity suppressing the ipsiversive movement. There are many more possibilities. It therefore remains to be seen whether the decision variables can be identified with activity in particular subcortical structures, and what the neural instantiations of the softmax computation and probabilistic action selection are.

5.4.2 Function

The mechanistic model reveals distinct functional roles for visual and secondary motor areas in the decision process. Visual areas appear to not only favour the

action associated with the stimulus on the contralateral hemifield, but also suppress the action associated with the ipsilateral hemifield. Interestingly, this pattern does not match the finding from widefield calcium imaging that visual cortical responses were exclusively contralateral until late in the trial. This suggests that the suppression of ipsiversive choices is through downstream projections. By contrast, secondary motor areas appear to favour both types of action, albeit slightly favouring the action associated with that hemisphere. This cooperative role of MOs is consistent with the responses we observed in widefield imaging, where MOs activates during the decision period for both contralateral and ipsilateral trials. This is also consistent with a bilateral and mutually-reinforcing representation of movement plans in ALM (an area that is the same or nearby to our MOs) observed previously during a working memory task (Li et al., 2016).

5.4.3 Caveats

The mechanistic model is not complete mechanistic account of the circuit process underlying the decision. It is instead a hybrid mechanistic-statistical model which estimates the functionally-independent contributions of cortical activity towards forming the decision variables. Action selection is accounted for by a more abstract softmax computation and probabilistic action selection rule. However, the model can also be considered a hypothesis for the functionally-relevant circuit processes underlying the decision. Considering this as a hypothesis for the circuitry, the model appears to neglect inter-cortical interactions in forming the decision. However it's known for example that visual areas project to secondary motor areas (reviewed in Barthas and Kwan (2016)). It is therefore possible that the role of VIS in forming the decision arises in its effects through MOs activity. This suggestion implies that inactivation of MOs is equivalent to inactivation of VIS and MOs together. Nevertheless, we find that simulated inactivation of MOs alone in the model seems sufficient to account for the empirical behavioural effect of inactivating this area. This is true even though the model was itself fit to widefield data which would contain correlated activity reflective of any inter-cortical interactions. Therefore, it does not seem crucial to incorporate the inter-cortical interactions when accounting for the behavioural effects of inactivation. We therefore feel that the model is not only successful at estimating functionally-independent roles for the cortical areas, but it also doubles as a hypothesis for the circuit processes underlying decision formation.

In its current form, the mechanistic model does not take account of time due to the static nature of the stimulus and therefore cannot make predictions of reaction time. In tasks where sensory evidence must be accumulated over time, accumulation-to-bound models can leverage this gradual decision process to model the dynamics of internal decision variables and therefore predict reaction times (Brunton et al., 2013; Gold and Shadlen, 2001). In tasks without any accumulation in the stimulus, accumulation models can still be applied as a model of internal ‘urgency’ to perform an action (Busemeyer and Townsend, 1993). One way of incorporating time in our model is to define the decision variable as the time-integral of the weighted input cortical activity and define a threshold for committing to an action. This model would therefore be similar to the accumulation-to-bound models used previously, except this model would allow for more than two choices. Nevertheless, we have shown from the widefield imaging and optogenetic inactivation experiments that VIS and MOs contribute at distinct times within the trial. Previous work has also demonstrated that early visual areas contribute towards visual detection over a very brief time-window (Resulaj et al., 2018). Therefore, it may be sufficient for the model to consider only a single time window when modelling functionally-relevant activity.

CHAPTER 6 CONCLUSIONS AND OUTLOOK

The key questions driving my thesis were

- 1. Does decision making arise from activation of distinct cortical regions in a temporal sequence, or instead arise in a distributed fashion?**
- 2. Which cortical areas are causally necessary for the decision process?**
- 3. What are the functional contributions of cortical areas towards decision making?**

We have made the following progress towards addressing these questions.

Firstly, to investigate the nature of the decision process, we designed a task which would produce behavioural data amenable to modelling. Our task design (published in Burgess et al. (2017)) contains two crucial features which have not been used together in mice before: two-alternative unforced choices, and visual discrimination. The unforced choice structure made it possible to distinguish different types of biases and sensitivities (bias/sensitivity associated with one side, or associated with the difference between sides), and therefore better interpret the effect of neural manipulation. The visual discrimination aspect of the task made it possible to sufficiently constrain a multinomial logistic model, thereby obtaining stable estimates of the bias, sensitivity and contrast shape parameters. We also applied hierarchical modelling to quantify inter-subject and inter-session variability in behaviour. Using this we identified large correlated variation in the two bias parameters across sessions. This reflected the fact that mice varied session-to-session in their overall tendency to NoGo.

Secondly, we asked whether the decision process arises from sequential activation of distinct cortical areas, or instead arises from a more distributed activation spanning many cortical regions simultaneously. To explore this, we performed widefield calcium imaging. We found hotspots of activity in several cortical areas which obeyed a strict temporal order. Stimuli present on one side led to sequential activation of contralateral primary visual cortex (VISp), higher visual areas (e.g. VISal), secondary motor areas (MOs), primary motor (MOp) and primary somatosensory (SSp) regions. We also observed activation on the ipsilateral hemisphere which was earliest in MOs and latest in VISp. These findings suggest that the stimulus induces feedforward activation through a cortical hierarchy,

terminating in MOs. MOp and SSp then receive top-down feedback from MOs. Sequential cortical activation has been previously observed in monkey performing visual (Ledberg et al., 2007; Siegel et al., 2015) and somatosensory (Hernández et al., 2010) discrimination tasks. In the mouse, sequential activation has only been observed in a whisker detection task (Le Merre et al., 2018). Interestingly in this study, the authors found activity started in somatosensory cortex and then progressed to primary motor cortex, and frontal areas (defined in this study as medial pre-frontal cortex) were activated last. In our task by contrast, we found an early activation of frontal area MOs which was present even in trials where the mice didn't respond. Our study therefore demonstrates an early sensory-related role for MOs in mouse visual discrimination.

Thirdly, we asked whether the activation of these cortical regions reflected a causal process which was also sequential. We optogenetically inactivated several cortical areas while mice performed the task, an approach directly inspired by the scanning inactivation approach in Guo et al. (2014). We identified two areas: VIS and MOs, where inactivation affected behaviour. We also found that the time when inactivation was strongest was different between these two regions (first in VIS, then in MOs) but still early into the trial. Interestingly, MOp and SSp inactivation had no effect. Therefore, while we observe widespread activation of activity from early sensory to secondary motor, to primary motor and somatosensory regions, only the first two stages appear to be causally related to performance in the task. Activity in other regions may reflect more feedback processes and sensory feedback/motor efference from wheel movement. The causal role of early sensory and frontal-motor areas has been demonstrated before in Go/NoGo (Goard et al., 2016) and 2AFC (Erlich et al., 2015; Guo et al., 2014) tasks. However, many of these previous studies impose a sequential structure, via a memory period, which may misleadingly imply that cortical areas play a temporally sequential role. Our study demonstrates a sequential causal role for these areas even in a task without an explicit memory period.

Finally, our task design highlighted different functional roles for VIS and MOs in this task which we quantified using a phenomenological modelling approach and a mechanistic modelling approach. By quantifying the effect of inactivation as a perturbation to the bias and sensitivity parameters of the phenomenological model, we showed that inactivation of VIS and MOs both resulted in reduced sensitivity towards the contralateral stimulus. However, inactivation in these

regions had distinct effects of bias: VIS inactivation induced a bias towards the side of inactivation and away from the contralateral side, whereas MOs inactivation biased behaviour towards NoGo. To understand this effect further, we constructed a mechanistic model of the decision process. We proposed that the algebra of the logistic model could be taken literally as a hypothesis for the cortical contribution towards the behaviour. In this model, cortical activity is linearly summed across the VIS and MOs on both hemispheres to form two decision variables: the log odds ratio of Left vs. NoGo choices, and the log odds of Right vs. NoGo. By estimating these weights from widefield data, we found that VIS and MOs had qualitatively different weightings onto the decision variables. VIS activity enhanced the decision variable associated with the contraversive movement and suppressed the decision variable associated with the ipsiversive movement, whereas MOs activity enhanced both. We validated this model by showing that it could capture psychometric data of the animal, predict choice variability on trials with identical stimulus conditions, and could be used to simulate the effect of optogenetic inactivation.

The mechanistic model therefore constitutes a novel hypothesis for the contribution of cortical activity towards decision making. The model makes several testable predictions. For example, 1) unilateral inactivation of visual and secondary motor areas together should totally abolish behavioural sensitivity to the contralateral visual stimulus. 2) Optogenetic activation or microstimulation of visual areas should increase contraversive movements and suppress ipsiversive movements. 3) There should exist, subcortically, a neural signal which correlates with the log odds ratio of Left vs. NoGo choices and Right vs. NoGo choices across contrast conditions. 4) Neural activity which correlates with the log odds ratios should receive convergent input from visual and secondary motor cortex.

6.1 Working model

The results in this thesis point towards the following account of how decisions arise from cortical activity (Figure 6-1). Below are a set of statements about the functional processes underlying the decision, and underneath each statement are bullet points summarizing evidence consistent with each point. These statements are partly speculative, and the real neural process is likely not this simple.

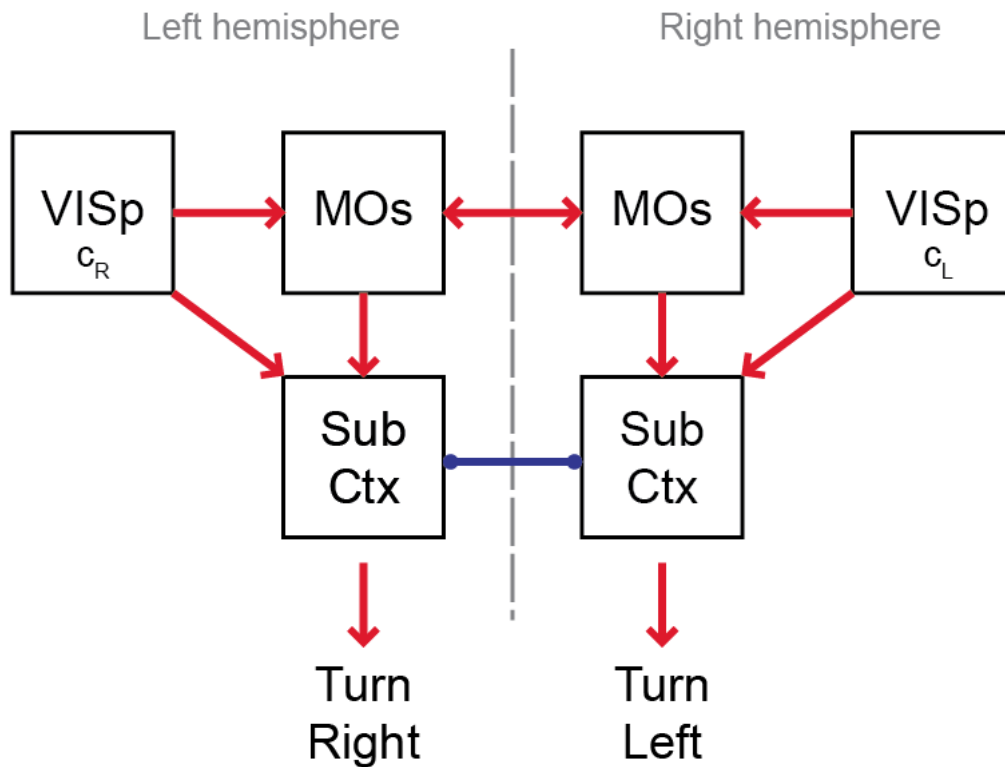


Figure 6-1 Working model of neural basis of two-alternative unforced-choice visual discrimination

Visual stimuli (c_L and c_R) arrive in the contralateral hemisphere primary visual cortex (VISp). Red arrows represent stimulation (by direct excitatory input or disinhibition). Blue line with filled circle represents suppression. See text for explanation.

Visual information arrives first in early visual cortex, which is relayed feed-forward to two targets: frontal-motor areas and subcortical structures.

- There are anatomical projections from VIS to MOs (Barthas and Kwan, 2016), and from VIS to several subcortical structures including thalamus, striatum and midbrain (Hintiryan et al., 2016).
- Widefield imaging and electrophysiology shows early feedforward activation of VIS followed by MOs on the same hemisphere.

The subcortical structures implement a competitive process, integrating visual information from both sides. This may function to resolve on which side the contrast was higher, or instead which action is appropriate for the stimulus.

- Optogenetic inactivation of VIS decreases contraversive choices and increases ipsiversive choices. This implies a subtractive downstream (subcortical) process which is affected by VIS inactivation.

- Subcortical structures like the superior colliculus have circuits implementing mutual inhibition (Knudsen, 2018; Lo and Wang, 2006).
- Mechanistic model weights suggest that the decision variables are enhanced by the contraversive visual information and suppressed by ipsiversive visual information.

The frontal-motor areas, by contrast, integrate visually-driven activity from the ipsilateral hemisphere's visual cortex with activity in the contralateral frontal-motor area. This signal may be akin to a 'go' signal, or 'urgency to move', but is still lateralised to Left/Right stimuli and/or choices.

- MOs receives excitatory projections from ipsilateral VIS and contralateral MOs (Barthas and Kwan, 2016).
- Widefield imaging shows MOs is active soon after VIS on the hemisphere contralateral to the stimulus. By contrast, on the ipsilateral hemisphere, the first region active is MOs.

This 'go' signal is relayed to subcortical structures, and if the signal is sufficiently strong, the subcortical structures trigger the motor plan.

- There are anatomical projections from MOs to subcortical structures such as spinal cord, superior colliculus, striatum and thalamus (Barthas and Kwan, 2016; Z. V. Guo et al., 2017; Hintiryan et al., 2016).
- Mechanistic model weights suggest that MOs activity enhances contraversive and ipsiversive ('go') movements via influence on downstream decision variables.

If a motor plan is triggered, an efference copy or sensory feedback signal is relayed to primary motor and somatosensory regions, possibly for ongoing feedback control of forelimb movements.

- Widefield imaging shows late bilateral activation of MOp and SSp.
- Inactivation in these regions has no effect, therefore these signals may be related to feedback/efference copy.

6.2 Future directions

The results obtained from this work have contributed towards our understanding of the neural basis of decisions. However, there are several aspects of the work which are worth expanding on in the future.

6.2.1 Phenomenological model

The current phenomenological model decomposes choices into choice biases and stimulus sensitivity. However, there are several additional states which have been shown to affect choice behaviour which future work could include.

In many decision tasks such as ours, optimal behaviour (defined as behaviour which maximises the rate of reward) requires that subjects make choices based only on the stimulus of the current trial. However, subjects exhibit history effects, where the behaviour on a current trial is influenced by a stimulus or action state on the previous trial. In a visual detection task, mice exhibit a win-stay lose-switch strategy where mice will tend to choose the action that was previously rewarded, or switch away from the action which was previously punished, independent of what the stimulus is (Busse et al., 2011). History effects have also been observed in monkeys, where perception of an oriented grating will be biased towards the angle of the previous stimulus (Fischer and Whitney, 2014). These perceptual biases have been attributed to midbrain dopaminergic circuits and basal ganglia (Lak et al., 2018), early visual cortex (John-Saaltink et al., 2016) and parietal cortex (Akrami et al., 2018). The phenomenological model which we have developed could therefore be expanded to include measures of history-dependent effects. This is particularly important in cases where estimates of the behavioural biases and stimulus sensitivities would change if we had an improved estimate of history effects. The hierarchical framework used in our model may be useful as well, because history effects are often estimated either by fitting on each session individually (Busse et al., 2011), or by concatenating data across many sessions (Akrami et al., 2018). The former method is hindered by overfitting, and the latter ignores session-by-session and subject-by-subject variation.

Another perceptual state is general arousal, which can be measured by proxy using the pupil diameter. In our task design, de-motivated mice will tend to perform NoGo more than Left or Right choices. Our current model can only attribute this to changes in bias, sensitivity and the non-linear contrast shape parameter. Motivation also interacts with history effects: human subjects with higher arousal states show stronger history effects (Urai et al., 2017). Therefore, the phenomenological model could be improved by incorporating both arousal and serial dependence. The model could also be improved by modelling parameter variation within a session. For example, as mice obtain water through the session,

satiety will increase gradually. This may lead to subsequent changes in behaviour over the course of a single session. The phenomenological model could therefore be improved by estimating these slow changes in behavioural state.

6.2.2 Mechanistic model

The mechanistic model can be expanded in several ways. For example, the mechanistic role of the downstream targets could be better defined. The model is currently missing a mechanistic account for how the value of the decision variables is transformed into the appropriate action, and what anatomical projections facilitate this process. Our model implies a softmax computation between the decision variables, as well as a probabilistic random draw to determine the action. Future work could therefore identify the neural mechanisms underlying these two stages. For example, as has been alluded to in the previous discussions, there are a few candidate regions (striatum, superior colliculus, etc) which could mediate this. Divisive normalization could implement something akin to a softmax computation (Carandini and Heeger, 2012), and lateral inhibition could implement winner-take-all dynamics to commit to one action (Maass, 2000). Investigating this will likely require experiments which record from several cortical and subcortical structures simultaneously, combined with careful manipulation of activity in specific regions and axonal projections. The highly complex recurrent connectivity may hinder interpretation, but this could be aided by theoretical models of circuit dynamics.

Both the phenomenological and mechanistic models could be further expanded to incorporate dynamics. The current models account for the decisions based on free parameters and inputs which are fixed in time. However, in reality, the presentation of the stimulus is continuous in time, leading to repeated feedforward activation of the cortical areas discussed in the previous Chapters. Similarly, ongoing recurrent dynamics in cortical and subcortical structures will likely alter the impact of early vs late sensory information on decision circuitry, and these dynamics eventually resolve to generate a decision at a specific time. To explore this, these models can be modified to incorporate time explicitly. This can be achieved in the mechanistic model by setting the decision variable values to be the time-integral of weighted activity in visual and secondary motor areas. A further development would be to translate the model into a drift-diffusion model, which would make specific predictions for the reaction time associated with each choice. Several drift-diffusion models have been developed in previous

tasks involving two choices (Ratcliff et al., 2016), however it is currently unclear what the equivalent model would be for multi-alternative unforced-choice tasks.

The mechanistic model could also be expanded to incorporate the state of individual cells within each cortical region. Currently the model is fit to widefield fluorescence, which is an estimate of the overall population rate. However individual cells within the population may contribute different roles, which may be averaged out in the population rate. For example, MOs neurons show diverse tuning towards task features in other tasks (Chen et al., 2017; Goard et al., 2016; Murakami et al., 2014), and therefore averaging over the neurons loses this information. This can be achieved by leveraging tools which can record electrophysiological activity in many neurons simultaneously (e.g. Jun et al., 2017). Nevertheless, the success of the mechanistic model could be attributed to the fact that the population rate was a suitable measure of the response of the population to the visual stimuli, reflecting the fact that individual neurons showed similarly monotonic tuning to visual contrast.

A further direction is to examine how these weights may vary with learning. Learning could take the form of learning the basic task rules, or instead learning to be more sensitive to stimuli after repeated exposure (i.e. perceptual learning). If the model is a good account of the mechanistic process underlying decisions, then it should be possible to understand learning in terms of changes to model weights or quantities. This question could be addressed by performing widefield calcium imaging in mice over the course of several weeks of training. The mechanistic model weights could be fit separately for each session, or instead a more sophisticated modelling approach could be used to model the drift in the parameters across sessions. Changes to the weights across training would suggest that learning occurs via the impact of cortex onto subcortical structures. Alternatively, if the weights do not change significantly with training, this would suggest the learning occurs either by changes to the cortical activity itself (e.g. Hua et al. (2010)) or through effects subcortically, after computation of the decision variables.

6.2.3 Sub-optimal decision making

The phenomenological and mechanistic models outlined in this thesis are based on a multinomial logistic framework, which is itself derived from normative economic theory. Normative models define behaviour which is optimal in some criterion, for example by maximizing reward rate. However, several studies have demonstrated

that observed decision behaviour violates optimality. For example, the independence of irrelevant alternatives (IIA) axiom states that the relative preference between any two alternatives is unaffected by changes in preference towards other alternatives (Luce, 1959). This assumption is consistently violated in tasks involving human and primate decisions (McNamara et al., 2014; Tversky and Kahneman, 1986). Taking inspiration from divisive normalization principles, Louie et al. (2013) showed that a model where the subjective value associated with each alternative is normalised by the total subjective utility across the alternatives was better able to account for IIA-violating human and monkey choice behaviour. Other examples of non-rational behaviour are loss aversion, probability distortion and intransitivity (Kahneman and Tversky, 1979). In loss aversion, when making decisions among alternatives that result in wins and losses, subjects will be disproportionately averse to losses. Probability distortion occurs when subjects over-estimate the frequency of low-probability events, and under-estimate the frequency of high-probability events. Intransitivity occurs when behaviour violates the principle that preferring A over B and B over C leads to preferring A over C. Tsetsos et al. (2016) developed a model of decisions where evidence for each alternative is noisy and must be accumulated over time. The authors found that behaviour which violates IIA and intransitivity could still be accounted for within a normative framework, by incorporating selective-gating of maximally-informative evidence samples (e.g. by attention). This shows that a normative account of decisions may still be applicable within a specific context.

Alternatives to normative theories are descriptive theories, which attempt to parsimoniously describe decisions as they are, rather than prescribing optimal decision rules. One successful example is Prospect theory, which models behaviour descriptively, incorporating parameters for loss aversion and probability distortion (Kahneman and Tversky, 1979). Models of this type are harder to fit to behavioural data because the larger number of parameters introduce redundancy and make the model underdetermined. However, with the combination of behavioural modelling and neural manipulations demonstrated in this thesis, it may be possible to determine the neural basis of non-rational behaviour and other cognitive properties observed in mice and humans.

6.3 Closing remarks

Perceptual decision making requires the coordinated interplay of sensory information, internal goals, and motor output. Accordingly, decision behaviour

recruits many brain areas mediating these three types of function. Our results shine light on the neocortical contribution towards this complex process. However, further work will be required to elucidate how the cortical and subcortical regions work together with the body and environment to achieve the behavioural goals.

BIBLIOGRAPHY

- Akrami, A., Kopec, C.D., Diamond, M.E., Brody, C.D., 2018. Posterior parietal cortex represents sensory history and mediates its effects on behaviour. *Nature* 554, 368-372. <https://doi.org/10.1038/nature25510>
- Balleine, B.W., Delgado, M.R., Hikosaka, O., 2007. The Role of the Dorsal Striatum in Reward and Decision-Making. *J. Neurosci.* 27, 8161-8165. <https://doi.org/10.1523/JNEUROSCI.1554-07.2007>
- Balleine, B.W., Dezfouli, A., Ito, M., Doya, K., 2015. Hierarchical control of goal-directed action in the cortical-basal ganglia network. *Curr. Opin. Behav. Sci., Neuroeconomics* 5, 1-7. <https://doi.org/10.1016/j.cobeha.2015.06.001>
- Barthas, F., Kwan, A.C., 2016. Secondary Motor Cortex: Where ‘Sensory’ Meets ‘Motor’ in the Rodent Frontal Cortex. *Trends Neurosci.* <https://doi.org/10.1016/j.tins.2016.11.006>
- Basso, M.A., May, P.J., 2017. Circuits for Action and Cognition: A View from the Superior Colliculus. *Annu. Rev. Vis. Sci.* 3, 197-226. <https://doi.org/10.1146/annurev-vision-102016-061234>
- Bechtel, W., Mandik, P., Mundale, J., Stufflebeam, R., 2001. Philosophy meets the neurosciences, in: *Philosophy and the Neurosciences: A Reader*. Wiley, Chapter 1.
- Boynton, G.M., Demb, J.B., Glover, G.H., Heeger, D.J., 1999. Neuronal basis of contrast discrimination. *Vision Res.* 39, 257-269. [https://doi.org/10.1016/S0042-6989\(98\)00113-8](https://doi.org/10.1016/S0042-6989(98)00113-8)
- Branco, T., Häusser, M., 2011. Synaptic Integration Gradients in Single Cortical Pyramidal Cell Dendrites. *Neuron* 69, 885-892. <https://doi.org/10.1016/j.neuron.2011.02.006>
- Britten, K.H., Newsome, W.T., Shadlen, M.N., Celebrini, S., Movshon, J.A., others, 1996. A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Vis. Neurosci.* 13, 87-100.
- Britten, K.H., Shadlen, M.N., Newsome, W.T., Movshon, J.A., 1992. The analysis of visual motion: a comparison of neuronal and psychophysical performance. *J. Neurosci.* 12, 4745-4765.
- Brunton, B.W., Botvinick, M.M., Brody, C.D., 2013. Rats and Humans Can Optimally Accumulate Evidence for Decision-Making. *Science* 340, 95-98. <https://doi.org/10.1126/science.1233912>

- Burgess, C.P., Lak, A., Steinmetz, N.A., Zatzka-Haas, P., Bai Reddy, C., Jacobs, E.A.K., Linden, J.F., Paton, J.J., Ranson, A., Schröder, S., Soares, S., Wells, M.J., Wool, L.E., Harris, K.D., Carandini, M., 2017. High-Yield Methods for Accurate Two-Alternative Visual Psychophysics in Head-Fixed Mice. *Cell Rep.* 20, 2513-2524.
<https://doi.org/10.1016/j.celrep.2017.08.047>
- Busemeyer, J.R., Townsend, J.T., 1993. Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. *Psychol. Rev.* 100, 432.
- Busse, L., Ayaz, A., Dhruv, N.T., Katzner, S., Saleem, A.B., Schölvinck, M.L., Zaharia, A.D., Carandini, M., 2011. The Detection of Visual Contrast in the Behaving Mouse. *J. Neurosci.* 31, 11351-11361.
<https://doi.org/10.1523/JNEUROSCI.6689-10.2011>
- Carandini, M., 2012. From circuits to behavior: a bridge too far? *Nat. Neurosci.* 15, 507-509. <https://doi.org/10.1038/nn.3043>
- Carandini, M., Churchland, A.K., 2013. Probing perceptual decisions in rodents. *Nat. Neurosci.* 16, 824-831. <https://doi.org/10.1038/nn.3410>
- Carandini, M., Heeger, D.J., 2012. Normalization as a canonical neural computation. *Nat. Rev. Neurosci.* 13, 51-62.
- Carpenter, B., Gelman, A., Hoffman, M.D., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M., Guo, J., Li, P., Riddell, A., 2017. Stan: A probabilistic programming language. *J. Stat. Softw.* 76.
- Chalmers, D.J., 2006. Strong and weak emergence. *Reemergence Emergence* 244-256.
- Chen, T.-W., Li, N., Daie, K., Svoboda, K., 2017. A Map of Anticipatory Activity in Mouse Motor Cortex. *Neuron* 94, 866-879.e4.
<https://doi.org/10.1016/j.neuron.2017.05.005>
- Chen, T.-W., Wardill, T.J., Sun, Y., Pulver, S.R., Renninger, S.L., Baohan, A., Schreiter, E.R., Kerr, R.A., Orger, M.B., Jayaraman, V., 2013. Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature* 499, 295.
- Churchland, A.K., Kiani, R., 2016. Three challenges for connecting model to mechanism in decision-making. *Curr. Opin. Behav. Sci., Computational modelling* 11, 74-80. <https://doi.org/10.1016/j.cobeha.2016.06.008>
- Churchland, A.K., Kiani, R., Shadlen, M.N., 2008. Decision-making with multiple alternatives. *Nat. Neurosci.* 11, 693-702. <https://doi.org/10.1038/nn.2123>

- Cisek, P., Kalaska, J.F., 2010. Neural Mechanisms for Interacting with a World Full of Action Choices. *Annu. Rev. Neurosci.* 33, 269-298.
<https://doi.org/10.1146/annurev.neuro.051508.135409>
- Crapse, T.B., Basso, M.A., 2015. Insights into decision making using choice probability. *J. Neurophysiol.* 114, 3039-3049.
<https://doi.org/10.1152/jn.00335.2015>
- Crapse, T.B., Lau, H., Basso, M.A., 2018. A Role for the Superior Colliculus in Decision Criteria. *Neuron* 97, 181-194.e6.
<https://doi.org/10.1016/j.neuron.2017.12.006>
- deCharms, R.C., Zador, A., 2000. Neural representation and the cortical code. *Annu. Rev. Neurosci.* 23, 613-647.
<https://doi.org/10.1146/annurev.neuro.23.1.613>
- Ding, L., Gold, J.I., 2012. Neural Correlates of Perceptual Decision Making before, during, and after Decision Commitment in Monkey Frontal Eye Field. *Cereb. Cortex* 22, 1052-1067. <https://doi.org/10.1093/cercor/bhr178>
- Donoghue, J.P., Wise, S.P., 1982. The motor cortex of the rat: Cytoarchitecture and microstimulation mapping. *J. Comp. Neurol.* 212, 76-88.
<https://doi.org/10.1002/cne.902120106>
- Erlich, J.C., Bialek, M., Brody, C.D., 2011. A Cortical Substrate for Memory-Guided Orienting in the Rat. *Neuron* 72, 330-343.
<https://doi.org/10.1016/j.neuron.2011.07.010>
- Erlich, J.C., Brunton, B.W., Duan, C.A., Hanks, T.D., Brody, C.D., 2015. Distinct effects of prefrontal and parietal cortex inactivations on an accumulation of evidence task in the rat. *eLife* 4, e05457.
<https://doi.org/10.7554/eLife.05457>
- Evans, L.S., Peachey, N.S., Marchese, A.L., 1993. Comparison of three methods of estimating the parameters of the Naka-Rushton equation. *Doc. Ophthalmol. Adv. Ophthalmol.* 84, 19-30.
- Fechner, G., 1860. *Elements of psychophysics*. Holt, Rinehart & Winston.
- Fischer, J., Whitney, D., 2014. Serial dependence in visual perception. *Nat. Neurosci.* 17, 738-743. <https://doi.org/10.1038/nn.3689>
- Frund, I., Wichmann, F.A., Macke, J.H., 2014. Quantifying the effect of intertrial dependence on perceptual decisions. *J. Vis.* 14, 9-9.
<https://doi.org/10.1167/14.7.9>

- García-Pérez, M.A., Alcalá-Quintana, R., 2013. Shifts of the psychometric function: Distinguishing bias from perceptual effects. *Q. J. Exp. Psychol.* 66, 319-337. <https://doi.org/10.1080/17470218.2012.708761>
- Gelman, A., Stern, H.S., Carlin, J.B., Dunson, D.B., Vehtari, A., Rubin, D.B., 2013. *Bayesian data analysis*. Chapman and Hall/CRC.
- Gharbawie, O.A., Gonzalez, C.L.R., Whishaw, I.Q., 2005. Skilled reaching impairments from the lateral frontal cortex component of middle cerebral artery stroke: a qualitative and quantitative comparison to focal motor cortex lesions in rats. *Behav. Brain Res.* 156, 125-137. <https://doi.org/10.1016/j.bbr.2004.05.015>
- Glimcher, P.W., Fehr, E., 2013. *Neuroeconomics: Decision Making and the Brain*. Academic Press.
- Goard, M.J., Pho, G.N., Woodson, J., Sur, M., 2016. Distinct roles of visual, parietal, and frontal motor cortices in memory-guided sensorimotor decisions. *eLife* 5, e13764. <https://doi.org/10.7554/eLife.13764>
- Gold, J.I., Ding, L., 2013. How mechanisms of perceptual decision-making affect the psychometric function. *Prog. Neurobiol., Conversion of Sensory Signals into Perceptions, Memories and Decisions* 103, 98-114. <https://doi.org/10.1016/j.pneurobio.2012.05.008>
- Gold, J.I., Shadlen, M.N., 2007. The Neural Basis of Decision Making. *Annu. Rev. Neurosci.* 30, 535-574. <https://doi.org/10.1146/annurev.neuro.29.051605.113038>
- Gold, J.I., Shadlen, M.N., 2001. Neural computations that underlie decisions about sensory stimuli. *Trends Cogn. Sci.* 5, 10-16. [https://doi.org/10.1016/S1364-6613\(00\)01567-9](https://doi.org/10.1016/S1364-6613(00)01567-9)
- Gold, J.I., Shadlen, M.N., 2000. Representation of a perceptual decision in developing oculomotor commands. *Nature* 404, 390-394. <https://doi.org/10.1038/35006062>
- Gomez-Marin, A., 2017. Causal circuit explanations of behavior: Are necessity and sufficiency necessary and sufficient?, in: *Decoding Neural Circuit Structure and Function*. Springer, pp. 283-306.
- Greene, W.H., 2011. *Econometric Analysis*, 7 edition. ed. Pearson Education, Boston; London.
- Grienberger, C., Konnerth, A., 2012. Imaging Calcium in Neurons. *Neuron* 73, 862-885. <https://doi.org/10.1016/j.neuron.2012.02.011>

- Grush, R., 2001. The semantic challenge to computational neuroscience. *Theory Method Neurosci.* 155-172.
- Guo, W., Clause, A.R., Barth-Maron, A., Polley, D.B., 2017. A Corticothalamic Circuit for Dynamic Switching between Feature Detection and Discrimination. *Neuron* 0. <https://doi.org/10.1016/j.neuron.2017.05.019>
- Guo, Z.V., Inagaki, H.K., Daie, K., Druckmann, S., Gerfen, C.R., Svoboda, K., 2017. Maintenance of persistent activity in a frontal thalamocortical loop. *Nature* 545, 181-186. <https://doi.org/10.1038/nature22324>
- Guo, Z.V., Li, N., Huber, D., Ophir, E., Gutnisky, D., Ting, J.T., Feng, G., Svoboda, K., 2014. Flow of Cortical Activity Underlying a Tactile Decision in Mice. *Neuron* 81, 179-194. <https://doi.org/10.1016/j.neuron.2013.10.020>
- Halassa, M.M., Kastner, S., 2017. Thalamic functions in distributed cognitive control. *Nat. Neurosci.* 20, 1669. <https://doi.org/10.1038/s41593-017-0020-1>
- Hall, R.D., Lindholm, E.P., 1974. Organization of motor and somatosensory neocortex in the albino rat. *Brain Res.* 66, 23-38. [https://doi.org/10.1016/0006-8993\(74\)90076-6](https://doi.org/10.1016/0006-8993(74)90076-6)
- Hanks, T., Kiani, R., Shadlen, M.N., 2014. A neural mechanism of speed-accuracy tradeoff in macaque area LIP. *eLife* 3, e02260. <https://doi.org/10.7554/eLife.02260>
- Hanks, T.D., Ditterich, J., Shadlen, M.N., 2006. Microstimulation of macaque area LIP affects decision-making in a motion discrimination task. *Nat. Neurosci.* 9, 682-689. <https://doi.org/10.1038/nn1683>
- Hanks, T.D., Kopec, C.D., Brunton, B.W., Duan, C.A., Erlich, J.C., Brody, C.D., 2015. Distinct relationships of parietal and prefrontal cortices to evidence accumulation. *Nature* 520, 220-223. <https://doi.org/10.1038/nature14066>
- Hanks, T.D., Summerfield, C., 2017. Perceptual Decision Making in Rodents, Monkeys, and Humans. *Neuron* 93, 15-31. <https://doi.org/10.1016/j.neuron.2016.12.003>
- Hansson, S., 1994. *Decision Theory: A Brief Introduction*. ResearchGate.
- Harvey, C.D., Coen, P., Tank, D.W., 2012. Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature* 484, 62-68. <https://doi.org/10.1038/nature10918>
- Hernández, A., Nácher, V., Luna, R., Zainos, A., Lemus, L., Alvarez, M., Vázquez, Y., Camarillo, L., Romo, R., 2010. Decoding a Perceptual Decision Process

- across Cortex. *Neuron* 66, 300-314.
<https://doi.org/10.1016/j.neuron.2010.03.031>
- Hernández, A., Zainos, A., Romo, R., 2000. Neuronal correlates of sensory discrimination in the somatosensory cortex. *Proc. Natl. Acad. Sci.* 97, 6191-6196. <https://doi.org/10.1073/pnas.120018597>
- Hintiryan, H., Foster, N.N., Bowman, I., Bay, M., Song, M.Y., Gou, L., Yamashita, S., Bienkowski, M.S., Zingg, B., Zhu, M., Yang, X.W., Shih, J.C., Toga, A.W., Dong, H.-W., 2016. The mouse cortico-striatal projectome. *Nat. Neurosci.* 19, 1100-1114. <https://doi.org/10.1038/nn.4332>
- Hua, T., Bao, P., Huang, C.-B., Wang, Z., Xu, J., Zhou, Y., Lu, Z.-L., 2010. Perceptual learning improves contrast sensitivity of V1 neurons in cats. *Curr. Biol. CB* 20, 887-894. <https://doi.org/10.1016/j.cub.2010.03.066>
- Huk, A.C., Shadlen, M.N., 2005. Neural Activity in Macaque Parietal Cortex Reflects Temporal Integration of Visual Motion Signals during Perceptual Decision Making. *J. Neurosci.* 25, 10420-10436.
<https://doi.org/10.1523/JNEUROSCI.4684-04.2005>
- Hunt, L.T., Hayden, B.Y., 2017. A distributed, hierarchical and recurrent framework for reward-based choice. *Nat. Rev. Neurosci.* 18, 172-182.
<https://doi.org/10.1038/nrn.2017.7>
- Jacobs, E.A.K., Steinmetz, N.A., Carandini, M., Harris, K.D., 2018. Cortical state fluctuations during sensory decision making. *bioRxiv* 348193.
<https://doi.org/10.1101/348193>
- Jin, X., Costa, R.M., 2015. Shaping action sequences in basal ganglia circuits. *Curr. Opin. Neurobiol., Motor circuits and action* 33, 188-196.
<https://doi.org/10.1016/j.conb.2015.06.011>
- John-Saaltink, E.S., Kok, P., Lau, H.C., Lange, F.P. de, 2016. Serial Dependence in Perceptual Decisions Is Reflected in Activity Patterns in Primary Visual Cortex. *J. Neurosci.* 36, 6186-6192.
<https://doi.org/10.1523/JNEUROSCI.4390-15.2016>
- Jun, J.J., Steinmetz, N.A., Siegle, J.H., Denman, D.J., Bauza, M., Barbarits, B., Lee, A.K., Anastassiou, C.A., Andrei, A., Aydın, Ç., Barbic, M., Blanche, T.J., Bonin, V., Couto, J., Dutta, B., Gratiy, S.L., Gutnisky, D.A., Häusser, M., Karsh, B., Ledochowitsch, P., Lopez, C.M., Mitelut, C., Musa, S., Okun, M., Pachitariu, M., Putzeys, J., Rich, P.D., Rossant, C., Sun, W., Svoboda, K., Carandini, M., Harris, K.D., Koch, C., O'Keefe, J., Harris, T.D., 2017.

- Fully integrated silicon probes for high-density recording of neural activity. *Nature* 551, nature24636. <https://doi.org/10.1038/nature24636>
- Kahneman, D., Tversky, A., 1979. Prospect Theory: An Analysis of Decision under Risk. *Econometrica* 47, 263-291. <https://doi.org/10.2307/1914185>
- Kaplan, D.M., 2011. Explanation and description in computational neuroscience. *Synthese* 183, 339.
- Katz, L.N., Yates, J.L., Pillow, J.W., Huk, A.C., 2016. Dissociated functional significance of decision-related activity in the primate dorsal stream. *Nature advance online publication*. <https://doi.org/10.1038/nature18617>
- Kaufman, M.T., Churchland, M.M., Ryu, S.I., Shenoy, K.V., 2014. Cortical activity in the null space: permitting preparation without movement. *Nat. Neurosci.* 17, 440-448. <https://doi.org/10.1038/nn.3643>
- Knudsen, E.I., 2018. Neural Circuits That Mediate Selective Attention: A Comparative Perspective. *Trends Neurosci.* <https://doi.org/10.1016/j.tins.2018.06.006>
- Kopec, C.D., Erlich, J.C., Brunton, B.W., Deisseroth, K., Brody, C.D., 2015. Cortical and Subcortical Contributions to Short-Term Memory for Orienting Movements. *Neuron* 88, 367-377. <https://doi.org/10.1016/j.neuron.2015.08.033>
- Kriegeskorte, N., Douglas, P.K., 2018. Cognitive computational neuroscience. *Nat. Neurosci.* 1. <https://doi.org/10.1038/s41593-018-0210-5>
- Lak, A., Okun, M., Moss, M., Gurnani, H., Wells, M.J., Reddy, C.B., Harris, K.D., Carandini, M., 2018. Dopaminergic and frontal signals for decisions guided by sensory evidence and reward value. *bioRxiv* 411413. <https://doi.org/10.1101/411413>
- Laming, D.R.J., 1968. Information theory of choice-reaction times.
- Le Merre, P., Esmaili, V., Charrière, E., Galan, K., Salin, P.-A., Petersen, C.C.H., Crochet, S., 2018. Reward-Based Learning Drives Rapid Sensory Signals in Medial Prefrontal Cortex and Dorsal Hippocampus Necessary for Goal-Directed Behavior. *Neuron* 97, 83-91.e5. <https://doi.org/10.1016/j.neuron.2017.11.031>
- Ledberg, A., Bressler, S.L., Ding, M., Coppola, R., Nakamura, R., 2007. Large-Scale Visuomotor Integration in the Cerebral Cortex. *Cereb. Cortex* 17. <https://doi.org/10.1093/cercor/bhj123>
- Lee, M.D., 2011. How cognitive modeling can benefit from hierarchical Bayesian models. *J. Math. Psychol.* 55, 1-7.

- Lein, E.S., Hawrylycz, M.J., Ao, N., Ayres, M., Bensinger, A., Bernard, A., Boe, A.F., Boguski, M.S., Brockway, K.S., Byrnes, E.J., Chen, Lin, Chen, Li, Chen, T.-M., Chin, M.C., Chong, J., Crook, B.E., Czaplinska, A., Dang, C.N., Datta, S., Dee, N.R., Desaki, A.L., Desta, T., Diep, E., Dolbeare, T.A., Donelan, M.J., Dong, H.-W., Dougherty, J.G., Duncan, B.J., Ebbert, A.J., Eichele, G., Estin, L.K., Faber, C., Facer, B.A., Fields, R., Fischer, S.R., Fliss, T.P., Frensley, C., Gates, S.N., Glattfelder, K.J., Halverson, K.R., Hart, M.R., Hohmann, J.G., Howell, M.P., Jeung, D.P., Johnson, R.A., Karr, P.T., Kawal, R., Kidney, J.M., Knapik, R.H., Kuan, C.L., Lake, J.H., Laramie, A.R., Larsen, K.D., Lau, C., Lemon, T.A., Liang, A.J., Liu, Y., Luong, L.T., Michaels, J., Morgan, J.J., Morgan, R.J., Mortrud, M.T., Mosqueda, N.F., Ng, L.L., Ng, R., Orta, G.J., Overly, C.C., Pak, T.H., Parry, S.E., Pathak, S.D., Pearson, O.C., Puchalski, R.B., Riley, Z.L., Rockett, H.R., Rowland, S.A., Royall, J.J., Ruiz, M.J., Sarno, N.R., Schaffnit, K., Shapovalova, N.V., Sivisay, T., Slaughterbeck, C.R., Smith, S.C., Smith, K.A., Smith, B.I., Sodt, A.J., Stewart, N.N., Stumpf, K.-R., Sunkin, S.M., Sutram, M., Tam, A., Teemer, C.D., Thaller, C., Thompson, C.L., Varnam, L.R., Visel, A., Whitlock, R.M., Wohnoutka, P.E., Wolkey, C.K., Wong, V.Y., Wood, M., Yaylaoglu, M.B., Young, R.C., Youngstrom, B.L., Yuan, X.F., Zhang, B., Zwingman, T.A., Jones, A.R., 2007. Genome-wide atlas of gene expression in the adult mouse brain. *Nature* 445, 168-176. <https://doi.org/10.1038/nature05453>
- Lewandowski, D., Kurowicka, D., Joe, H., 2009. Generating random correlation matrices based on vines and extended onion method. *J. Multivar. Anal.* 100, 1989-2001. <https://doi.org/10.1016/j.jmva.2009.04.008>
- Li, N., Daie, K., Svoboda, K., Druckmann, S., 2016. Robust neuronal dynamics in premotor cortex during motor planning. *Nature* 532, 459-464. <https://doi.org/10.1038/nature17643>
- Licata, A.M., Kaufman, M.T., Raposo, D., Ryan, M.B., Sheppard, J.P., Churchland, A.K., 2017. Posterior parietal cortex guides visual decisions in rats. *J. Neurosci.* 0105-17. <https://doi.org/10.1523/JNEUROSCI.0105-17.2017>
- Liu, L.D., Pack, C.C., 2017. The Contribution of Area MT to Visual Motion Perception Depends on Training. *Neuron* 0. <https://doi.org/10.1016/j.neuron.2017.06.024>

- Lo, C.-C., Wang, X.-J., 2006. Cortico-basal ganglia circuit mechanism for a decision threshold in reaction time tasks. *Nat. Neurosci.* 9, 956-963. <https://doi.org/10.1038/nn1722>
- Louie, K., Khaw, M.W., Glimcher, P.W., 2013. Normalization is a general neural mechanism for context-dependent decision making. *Proc. Natl. Acad. Sci.* 110, 6139-6144. <https://doi.org/10.1073/pnas.1217854110>
- Luce, R.D., 1959. *Individual Choice Behavior a Theoretical Analysis*. John Wiley and sons.
- Lueckmann, J.-M., Macke, J.H., Nienborg, H., 2018. Can Serial Dependencies in Choices and Neural Activity Explain Choice Probabilities? *J. Neurosci.* 38, 3495-3506. <https://doi.org/10.1523/JNEUROSCI.2225-17.2018>
- Lynn, S.K., Barrett, L.F., 2014. "Utilizing" Signal Detection Theory. *Psychol. Sci.* 25, 1663-1673. <https://doi.org/10.1177/0956797614541991>
- Lynn, S.K., Wormwood, J.B., Barrett, L.F., Quigley, K.S., 2015. Decision making from economic and signal detection perspectives: development of an integrated framework. *Front. Psychol.* 6. <https://doi.org/10.3389/fpsyg.2015.00952>
- Ma, Y., Shaik, M.A., Kim, S.H., Kozberg, M.G., Thibodeaux, D.N., Zhao, H.T., Yu, H., Hillman, E.M., 2016. Wide-field optical mapping of neural activity and brain haemodynamics: considerations and novel approaches. *Phil Trans R Soc B* 371, 20150360.
- Maass, W., 2000. On the Computational Power of Winner-Take-All. *Neural Comput* 12, 2519-2535. <https://doi.org/10.1162/089976600300014827>
- Machens, C.K., Romo, R., Brody, C.D., 2005. Flexible Control of Mutual Inhibition: A Neural Model of Two-Interval Discrimination. *Science* 307, 1121-1124. <https://doi.org/10.1126/science.1104171>
- Macmillan, N.A., Creelman, C.D., 2004. *Detection theory: A user's guide*. Psychology press.
- Mante, V., Sussillo, D., Shenoy, K.V., Newsome, W.T., 2013. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* 503, 78-84. <https://doi.org/10.1038/nature12742>
- Marr, D., 1982. *Vision: A computational investigation into the human representation and processing of visual information*, Henry Holt and Co. Inc N. Y. NY 2, 4-2.
- McElreath, R., 2018. *Statistical Rethinking: A Bayesian Course with Examples in R and Stan*. CRC Press.

- McNamara, J.M., Trimmer, P.C., Houston, A.I., 2014. Natural selection can favour 'irrational' behaviour. *Biol. Lett.* 10.
<https://doi.org/10.1098/rsbl.2013.0935>
- Murakami, M., Vicente, M.I., Costa, G.M., Mainen, Z.F., 2014. Neural antecedents of self-initiated actions in secondary motor cortex. *Nat. Neurosci.* 17, 1574-1582. <https://doi.org/10.1038/nn.3826>
- Neumann, J. von, Morgenstern, O., others, 1944. *Theory of games and economic behavior*. Princeton university press Princeton.
- Newsome, W.T., Pare, E.B., 1988. A selective impairment of motion perception following lesions of the middle temporal visual area (MT). *J. Neurosci.* 8, 2201-2211.
- Niell, C.M., Stryker, M.P., 2008. Highly selective receptive fields in mouse visual cortex. *J. Neurosci. Off. J. Soc. Neurosci.* 28, 7520-7536.
<https://doi.org/10.1523/JNEUROSCI.0623-08.2008>
- Nienborg, H., Cumming, B.G., 2009. Decision-related activity in sensory neurons reflects more than a neuron's causal effect. *Nature* 459, 89-92.
<https://doi.org/10.1038/nature07821>
- O'Connell, R.G., Shadlen, M.N., Wong-Lin, K., Kelly, S.P., 2018. Bridging Neural and Computational Viewpoints on Perceptual Decision-Making. *Trends Neurosci.* <https://doi.org/10.1016/j.tins.2018.06.005>
- Otchy, T.M., Wolff, S.B.E., Rhee, J.Y., Pehlevan, C., Kawai, R., Kempf, A., Gobes, S.M.H., Ölveczky, B.P., 2015. Acute off-target effects of neural circuit manipulations. *Nature* 528, 358-363.
<https://doi.org/10.1038/nature16442>
- Pachitariu, M., Steinmetz, N., Kadir, S., Carandini, M., Harris, K.D., 2016. Kilosort: realtime spike-sorting for extracellular electrophysiology with hundreds of channels. *bioRxiv* 061481. <https://doi.org/10.1101/061481>
- Park, I.M., Meister, M.L.R., Huk, A.C., Pillow, J.W., 2014. Encoding and decoding in parietal cortex during sensorimotor decision-making. *Nat. Neurosci.* 17, 1395-1403. <https://doi.org/10.1038/nn.3800>
- Parker, A.J., Newsome, W.T., 1998. Sense and the Single Neuron: Probing the Physiology of Perception. *Annu. Rev. Neurosci.* 21, 227-277.
<https://doi.org/10.1146/annurev.neuro.21.1.227>
- Piray, P., Dezfouli, A., Heskes, T., Frank, M.J., Daw, N.D., 2018. Hierarchical Bayesian inference for concurrent model fitting and comparison for group studies. *bioRxiv* 393561. <https://doi.org/10.1101/393561>

- Platt, M.L., Glimcher, P.W., 1999. Neural correlates of decision variables in parietal cortex. *Nature* 400, 233-238. <https://doi.org/10.1038/22268>
- Raposo, D., Kaufman, M.T., Churchland, A.K., 2014. A category-free neural population supports evolving demands during decision-making. *Nat. Neurosci.* 17, 1784-1792. <https://doi.org/10.1038/nn.3865>
- Ratcliff, R., Smith, P.L., Brown, S.D., McKoon, G., 2016. Diffusion Decision Model: Current Issues and History. *Trends Cogn. Sci.* 20, 260-281. <https://doi.org/10.1016/j.tics.2016.01.007>
- Resulaj, A., Ruediger, S., Olsen, S.R., Scanziani, M., 2018. First spikes in visual cortex enable perceptual discrimination. *eLife* 7, e34044. <https://doi.org/10.7554/eLife.34044>
- Roe, R.M., Busemeyer, J.R., Townsend, J.T., 2001. Multialternative decision field theory: A dynamic connectionist model of decision making. *Psychol. Rev.* 108, 370.
- Roitman, J.D., Shadlen, M.N., 2002. Response of Neurons in the Lateral Intraparietal Area during a Combined Visual Discrimination Reaction Time Task. *J. Neurosci.* 22, 9475-9489.
- Romo, R., Hernández, A., Zainos, A., Lemus, L., Brody, C.D., 2002. Neuronal correlates of decision-making in secondary somatosensory cortex. *Nat. Neurosci.* 5, 1217.
- Romo, R., Salinas, E., 1999. Sensing and deciding in the somatosensory system. *Curr. Opin. Neurobiol.* 9, 487-493. [https://doi.org/10.1016/S0959-4388\(99\)80073-7](https://doi.org/10.1016/S0959-4388(99)80073-7)
- Rouder, J.N., Lu, J., 2005. An introduction to Bayesian hierarchical models with an application in the theory of signal detection. *Psychon. Bull. Rev.* 12, 573-604. <https://doi.org/10.3758/BF03196750>
- Sachidhanandam, S., Sreenivasan, V., Kyriakatos, A., Kremer, Y., Petersen, C.C.H., 2013. Membrane potential correlates of sensory perception in mouse barrel cortex. *Nat. Neurosci.* 16, 1671-1677. <https://doi.org/10.1038/nn.3532>
- Sahara, S., Yanagawa, Y., O'Leary, D.D.M., Stevens, C.F., 2012. The fraction of cortical GABAergic neurons is constant from near the start of cortical neurogenesis to adulthood. *J. Neurosci.* 32, 4755-4761. <https://doi.org/10.1523/JNEUROSCI.6412-11.2012>

- Schaal, D.W., 2005. Naming Our Concerns about Neuroscience: A Review of Bennett and Hacker's Philosophical Foundations of Neuroscience. *J. Exp. Anal. Behav.* 84, 683-692. <https://doi.org/10.1901/jeab.2005.83-05>
- Schall, J., Thompson, K., 1999. Neural Selection and Control of Visually Guided Eye Movements. *Annu. Rev. Neurosci.* 22, 241-259. <https://doi.org/10.1146/annurev.neuro.22.1.241>
- Schultz, W., Dayan, P., Montague, P.R., 1997. A Neural Substrate of Prediction and Reward. *Science* 275, 1593-1599. <https://doi.org/10.1126/science.275.5306.1593>
- Shadlen, M.N., Kiani, R., 2013. Decision Making as a Window on Cognition. *Neuron* 80, 791-806. <https://doi.org/10.1016/j.neuron.2013.10.047>
- Shadlen, M.N., Newsome, W.T., 2001. Neural Basis of a Perceptual Decision in the Parietal Cortex (Area LIP) of the Rhesus Monkey. *J. Neurophysiol.* 86, 1916-1936.
- Shadlen, M.N., Newsome, W.T., 1996. Motion perception: seeing and deciding. *Proc. Natl. Acad. Sci.* 93, 628-633.
- Siegel, M., Buschman, T.J., Miller, E.K., 2015. Cortical information flow during flexible sensorimotor decisions. *Science* 348, 1352-1355. <https://doi.org/10.1126/science.aab0551>
- Smit, H., Hacker, P.M.S., 2014. Seven Misconceptions About the Mereological Fallacy: A Compilation for the Perplexed. *Erkenntnis* 79, 1077-1097. <https://doi.org/10.1007/s10670-013-9594-5>
- Sridharan, D., Steinmetz, N.A., Moore, T., Knudsen, E.I., 2014. Distinguishing bias from sensitivity effects in multialternative detection tasks. *J. Vis.* 14, 16. <https://doi.org/10.1167/14.9.16>
- Steinmetz, N.A., Buetfering, C., Lecoq, J., Lee, C.R., Peters, A.J., Jacobs, E.A.K., Coen, P., Ollerenshaw, D.R., Valley, M.T., Vries, S.E.J. de, Garrett, M., Zhuang, J., Groblewski, P.A., Manavi, S., Miles, J., White, C., Lee, E., Griffin, F., Larkin, J.D., Roll, K., Cross, S., Nguyen, T.V., Larsen, R., Pendergraft, J., Daigle, T., Tasic, B., Thompson, C.L., Waters, J., Olsen, S., Margolis, D.J., Zeng, H., Hausser, M., Carandini, M., Harris, K.D., 2017. Aberrant Cortical Activity in Multiple GCaMP6-Expressing Transgenic Mouse Lines. *eNeuro* 4, ENEURO.0207-17.2017. <https://doi.org/10.1523/ENEURO.0207-17.2017>
- Stüttgen, M., Schwarz, C., Jakel, F., 2011. Mapping Spikes to Sensations. *Front. Neurosci.* 5. <https://doi.org/10.3389/fnins.2011.00125>

- Sul, J.H., Jo, S., Lee, D., Jung, M.W., 2011. Role of rodent secondary motor cortex in value-based action selection. *Nat. Neurosci.* 14, 1202-1208. <https://doi.org/10.1038/nn.2881>
- Summerfield, C., Tsetsos, K., 2012. Building Bridges between Perceptual and Economic Decision-Making: Neural and Computational Mechanisms. *Front. Neurosci.* 6. <https://doi.org/10.3389/fnins.2012.00070>
- Svoboda, K., Li, N., 2018. Neural mechanisms of movement planning: motor cortex and beyond. *Curr. Opin. Neurobiol., Neurobiology of Behavior* 49, 33-41. <https://doi.org/10.1016/j.conb.2017.10.023>
- Tanner, W.P., Swets, J.A., 1954. A decision-making theory of visual detection. *Psychol. Rev.* 61, 401-409. <https://doi.org/10.1037/h0058700>
- Tecuapetla, F., Matias, S., Dugue, G.P., Mainen, Z.F., Costa, R.M., 2014. Balanced activity in basal ganglia projection pathways is critical for contraversive movements. *Nat. Commun.* 5, 4315. <https://doi.org/10.1038/ncomms5315>
- Tsetsos, K., Moran, R., Moreland, J., Chater, N., Usher, M., Summerfield, C., 2016. Economic irrationality is optimal during noisy decision making. *Proc. Natl. Acad. Sci.* 113, 3102-3107. <https://doi.org/10.1073/pnas.1519157113>
- Tsunada, J., Liu, A.S.K., Gold, J.I., Cohen, Y.E., 2016. Causal contribution of primate auditory cortex to auditory perceptual decision-making. *Nat. Neurosci.* 19, 135-142. <https://doi.org/10.1038/nn.4195>
- Tversky, A., Kahneman, D., 1986. Rational choice and the framing of decisions. *J. Bus.* S251-S278.
- Urai, A.E., Braun, A., Donner, T.H., 2017. Pupil-linked arousal is driven by decision uncertainty and alters serial choice bias. *Nat. Commun.* 8, 14637. <https://doi.org/10.1038/ncomms14637>
- Vehtari, A., Gelman, A., Gabry, J., 2017. Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Stat. Comput.* 27, 1413-1432.
- Wang, X.-J., 2008. Decision Making in Recurrent Neuronal Circuits. *Neuron* 60, 215-234. <https://doi.org/10.1016/j.neuron.2008.09.034>
- Wang, Z., Maunze, B., Wang, Y., Tsoulfas, P., Blackmore, M.G., 2018. Global connectivity and function of descending spinal input revealed by 3D microscopy and retrograde transduction. *bioRxiv* 314237. <https://doi.org/10.1101/314237>

- Wehr, M., Zador, A.M., 2005. Synaptic Mechanisms of Forward Suppression in Rat Auditory Cortex. *Neuron* 47, 437-445.
<https://doi.org/10.1016/j.neuron.2005.06.009>
- Wei, X.-X., Stocker, A.A., 2016. A new law of human perception. *bioRxiv* 091918.
<https://doi.org/10.1101/091918>
- Wekselblatt, J.B., Flister, E.D., Piscopo, D.M., Niell, C.M., 2016. Large-scale imaging of cortical dynamics during sensory perception and behavior. *J. Neurophysiol.* 115, 2852-2866. <https://doi.org/10.1152/jn.01056.2015>
- Wickens, J.R., Budd, C.S., Hyland, B.I., Arbuthnott, G.W., 2007. Striatal contributions to reward and decision making. *Ann. N. Y. Acad. Sci.* 1104, 192-212.
- Yang, H., Kwon, S.E., Severson, K.S., O'Connor, D.H., 2016. Origins of choice-related activity in mouse somatosensory cortex. *Nat. Neurosci.* 19.
<https://doi.org/10.1038/nn.4183>
- Zénon, A., Krauzlis, R.J., 2012. Attention deficits without cortical neuronal deficits. *Nature* 489, 434-437. <https://doi.org/10.1038/nature11497>
- Znamenskiy, P., Zador, A.M., 2013. Corticostriatal neurons in auditory cortex drive decisions during auditory discrimination. *Nature* 497, 482-485.
<https://doi.org/10.1038/nature12077>