

Distribution Agreement

In presenting this thesis or dissertation as a partial fulfillment of the requirements for an advanced degree from Emory University, I hereby grant to Emory University and its agents the non-exclusive license to archive, make accessible, and display my thesis or dissertation in whole or in part in all forms of media, now or hereafter known, including display on the world wide web. I understand that I may select some access restrictions as part of the online submission of this thesis or dissertation. I retain all ownership rights to the copyright of the thesis or dissertation. I also retain the right to use in future works (such as articles or books) all or part of this thesis or dissertation.

Signature:

Varun Saravanan

Date

Elucidating the role of dopamine in vocal learning in the Bengalese finch

By

Varun Saravanan
Doctor of Philosophy
Graduate Division of Biological and Biomedical Sciences
Neuroscience

Samuel J. Sober
Advisor

Gordon J. Berman
Committee Member

Dieter Jaeger
Committee Member

Robert Liu
Committee Member

Ilya Nemenman
Committee Member

Accepted:

Lisa A. Tedesco, Ph.D.
Dean of the James T. Laney School of Graduate Studies

Date

Elucidating the role of dopamine in vocal learning in the Bengalese finch

By

Varun Saravanan
B.Tech, Indian Institute of Technology, Madras, 2013

Advisor: Samuel J. Sober, Ph.D.

An abstract of
A dissertation submitted to the Faculty of the
James T. Laney School of Graduate Studies of Emory University
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy
in Neuroscience
2019

Abstract

Elucidating the role of dopamine in vocal learning in the Bengalese finch

By Varun Saravanan

All living organisms must learn to produce certain motor behaviors appropriate to the sensory stimuli encountered in their environment in order to survive. In complex organisms, such motor behavior produced is evaluated through further sensory feedback and is refined for future performance. This learning is referred to as sensorimotor learning. Sensorimotor learning can be distinguished into two broad categories – reinforcement learning that relies on external appetitive or aversive cues and sensorimotor adaptation in which the only feedback available to the organism is its own evaluation of the motor behavior performed. We have previously shown a role for dopamine in vocal reinforcement learning in Bengalese finches (*Lonchura striata* var. *domestica*). This dissertation again uses Bengalese finches, a highly accessible sensorimotor learner, to study the role of dopamine in sensorimotor adaptation.

We discovered, as we were analyzing data for a previous experiment, that some of the assumptions underlying the statistical tests we had used in the past were being violated due to the hierarchical nature of our data. When variability exists at multiple levels, as is the case with hierarchical data, the error has to be propagated appropriately in order to account for the total uncertainty in the measurement. We had previously been treating each data point independently in spite of the hierarchical structure which underestimates the total error. We show that using hierarchical bootstrapping, we can accurately quantify the uncertainty in the measurement. In addition, we show real world applications of the hierarchical bootstrap and demonstrate how it provides more accurate results than traditional statistical tests on hierarchical datasets. We then used hierarchical bootstrapping to show that birds displayed severe sensorimotor adaptation deficits following a dopamine lesion of Area X, a song-specific basal ganglia nucleus. Specifically, birds showed both an inability to shift their pitch adaptively to induced auditory feedback errors and a tendency to reduce pitch regardless of auditory error post-lesion. By building on the involvement of dopamine in reinforcement learning and showing its necessity for sensorimotor adaptation, this dissertation lays the foundation for uncovering the role of dopamine in sensorimotor learning.

Elucidating the role of dopamine in vocal learning in the Bengalese finch

By

Varun Saravanan
B.Tech, Indian Institute of Technology, Madras, 2013

Advisor: Samuel J. Sober, Ph.D.

A dissertation submitted to the Faculty of the
James T. Laney School of Graduate Studies of Emory University
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy
in Neuroscience
2019

ACKNOWLEDGEMENTS

True to its nature of being a “terminal degree”, I discovered that obtaining a PhD was a grueling process. I only made it this far by standing on the shoulders of giants and leaning heavily on the love and support of various people along the way. I would like to take this opportunity to thank all those who helped me along my journey from the bottom of my heart.

First and foremost I would like to thank my advisor, Dr. Samuel J. Sober. As the story goes of the rabbit who wrote a thesis titled “The superiority of rabbits over foxes and wolves” with the help of a very well-fed lion for an advisor, Sam has been an outstanding advisor who gave me the courage to frame my own research question and follow the science where it may lead me. Sam has been extraordinarily supportive. Every time I approached him with a problem, his first response was to ask me to not freak out and immediately start brainstorming ways in which to make the situation work. Much of my critical thinking and scientific temperament has been defined by my advisor. Finally, he made it a pleasure to work in the lab and I aspire to his people management skills in the future.

In the same vein, I would like to thank the members of my thesis committee: Dr. Gordon Berman – thank you for always making time for me to come to you with my confusions about particular concepts and helping me refine the computational section of my thesis; Dr. Dieter Jaeger – you were my first contact at Emory University, thank you for pushing me to perform at my best and for your constructive feedback during my time at Emory; Dr. Robert Liu – you were the first person I worked under at Emory, thank you for all your help and guidance as I was a new international student, your continuing support and big picture focus while on my thesis committee; Dr. Ilya Nemenman – I was first drawn to Emory by your unique perspective and would enjoy all scientific discussions with you, thank you for asking the hard questions and transforming me into a better scientist.

I would also like to thank all my lab mates in the Sober lab – both past and present: Kyle Srivastava, David Nicholson, Lukas Hoffmann, Alynda Wood, James McGregor, Andrea Pack, Amanda Jacob, Ben Kuebrich, Rachel Conn, Sena Agezo, MacKenzie Wyatt, Emily Berthiaume, Conor Kelly, David Hercules,

Laura Waters-Goggins, Connor Gallimore, Sean O'Connell, Abi Grassler and Sherry Ye. You guys have been amazing sources of support and fun through both good times and dark. I have had fun spending time with you all and will miss you as we all go our separate ways. I would particularly like to give a shout out to Conor Kelly and Lukas Hoffmann who trained me in the various experimental procedures I performed for my dissertation. I must also thank Elizabeth Amadei who was my first graduate mentor and remains a friend and role-model. I would also like to mention Baohua Zhou, David Hofmann and Ahmed Roman from the Nemenman group for helping me understand the technical side of my project and engaging in many a laborious blackboard session trying to fit math to confusing data. Along the same lines, I would also like to thank my friends and colleagues in the Biology department, Neuroscience graduate program and at Emory University in general. The community has been extremely welcoming and has helped make the journey to the Ph.D. that much easier.

I would also like to thank my parents, Saravanan Chandran and Sathyabama Vijayarangan. They have been two solid pillars of support that have never once wavered. They made me into the man I am today and I recognize that all of what I have achieved would not have been possible without their love and support. I would also like to thank my extended family for their support and words of encouragement through six long years.

My acknowledgements section would be incomplete without thanking my girlfriend turned fiancé turned wife, Olivia Moody. She has been with me throughout my time at Emory and gave my life purpose and direction even during times in my Ph.D. journey when experiments were not working and I was finding it difficult to motivate myself to keep going. She has been my rock and I would not be standing here today if she had not been at my side. Last but definitely not the least, I would like to thank a new addition to my life, my daughter Ileana Saravanan, for her smile that lights up the entire room and her laugh that never fails to brighten my day.

TABLE OF CONTENTS

TABLE OF FIGURES	3
LIST OF ABBREVIATIONS	5
1 CHAPTER I: INTRODUCTION AND LITERATURE REVIEW	7
1.1 Sensorimotor Learning	8
1.1.1 Reinforcement Learning	8
1.1.2 Sensorimotor Adaptation	14
1.1.3 Other classifications of Sensorimotor Learning	17
1.2 A role for Dopamine in Sensorimotor Learning	21
1.2.1 Dopamine before Reward Prediction Error (RPE)	21
1.2.2 Dopamine in Reinforcement Learning.....	22
1.2.3 Dopamine in Sensorimotor Adaptation	28
1.3 Songbirds as a model system	30
1.3.1 Overview.....	30
1.3.2 Reinforcement Learning in Songbirds	31
1.3.3 Sensorimotor Adaptation in Songbirds.....	34
1.4 Error Propagation in Hierarchical Data	35
1.4.1 Problem Identification and Early Solutions	36
1.4.2 Linear Mixed Models (LMMs) as a Solution	38
1.4.3 Bootstrapping and other Resampling methods as a solution	40
1.5 Dissertation Overview	41
2 CHAPTER II: DOPAMINE DEPLETION AFFECTS VOCAL ACOUSTICS AND DISRUPTS SENSORIMOTOR ADAPTATION IN SONGBIRDS	43
2.1 Abstract	43
2.2 Introduction	44
2.3 Materials and Methods	46
2.3.1 Experimental design and Statistical Analysis	46
2.3.2 6-OHDA Lesions:	51
2.3.3 Headphones attachment and assembly:	51
2.3.4 Histology:.....	52
2.3.5 Image and Lesion Analysis:.....	53
2.3.6 Pitch Quantification:	54
2.3.7 Error quantification:	55

2.3.8	Hypothesis testing with Bootstrap:	57
2.3.9	Validating our Results with Linear Mixed Models:	58
2.4	Results	59
2.4.1	6-OHDA lesions reduce dopaminergic innervation of Area X:.....	59
2.4.2	6-OHDA lesioned birds reduce pitch even in the absence of auditory error:	61
2.4.3	6-OHDA lesioned birds do not respond adaptively to pitch-shifted auditory error: 63	
2.4.4	No correlations between lesion extent and changes in pitch:	68
2.4.5	Washout is impaired by dopamine depletion:.....	69
2.5	Discussion	74
3	CHAPTER III: APPLICATION OF THE HIERARCHICAL BOOTSTRAP TO MULTI-LEVEL DATA IN NEUROSCIENCE	79
3.1	Abstract	79
3.2	Introduction	80
3.3	Materials and Methods	82
3.3.1	Traditional vs Summarized vs Bootstrap.....	82
3.3.2	Hypothesis testing using Bootstrap:.....	85
3.3.3	Design Effect (DEFF):.....	87
3.4	Results	88
3.4.1	Simulations	88
3.4.2	Examples.....	97
3.5	Discussion	104
4	CHAPTER IV: CONCLUSIONS AND FUTURE DIRECTIONS	110
4.1	Dopamine as a multiplexed signal	110
4.2	New Techniques to understand brain function	114
4.3	The problem of Information in an Interconnected Network	116
4.4	Future Experimental Proposals	119
4.5	Statistics and the Bootstrap	123
5	CHAPTER V: REFERENCES	126

TABLE OF FIGURES

Figure 1.1: Classifications of sensorimotor learning.	18
Figure 1.2: Response of dopamine neurons to rewards and reward predictions.	23
Figure 1.3: Neural pathways involved in and example spectrogram of birdsong.	32
Figure 1.4: An extreme example of pseudoreplication.	35
Figure 2.1: Songbird neuroanatomy and experimental design.....	50
Figure 2.2: Metric for quantifying the extent of our lesions in our population of birds. ..	60
Figure 2.3: Quantifying the effect of headphones without any pitch shifts on the average change in pitch of the bird with or without lesions.	62
Figure 2.4: Individual traces for unshifted birds.	62
Figure 2.5: Change in pitch in response to pitch shift errors through the headphones in unlesioned and 6-OHDA lesioned birds.	65
Figure 2.6: Individual traces for shifted birds.....	66
Figure 2.7: Analysis of change in pitch during washout for lesioned and unlesioned birds.....	71
Figure 2.8: Washout traces for individual birds.....	72
Figure 2.9: Dynamics of the change of pitch or $\Delta(\text{Pitch})$ during washout.....	73
Figure 3.1: Schematic of hierarchical datasets and pseudocodes for statistical methods considered.....	83
Figure 3.2: Results from the simulation in which there was no difference between the groups of two neurons.....	90

Figure 3.3: False positive rate and size of error bars quantified for the simulation in which there was again no difference between the groups but now the points were not independent.	93
Figure 3.4: Change of power with number of neurons and effect size.....	95
Figure 3.5: Reanalysis of generalization in the headphones learning paradigm.....	99
Figure 3.6: Hierarchical bootstrapping most accurately captures the signal in an experiment studying optogenetic control of behavior in flies.	103

LIST OF ABBREVIATIONS

6-OHDA	6-hydroxydopamine
AFP	Anterior Forebrain Pathway
AP	Anterior-Posterior axis
DLM	Dorsolateral nucleus of the anterior thalamus
DREADDs	Designer Receptors Exclusively Activated by Designer Drugs
DV	Dorsal-Ventral axis
fMRI	Functional Magnetic Resonance Imaging
HVC	Used as a proper name; songbird brain nucleus that projects to RA and Area X
L-DOPA	Levodopa or L-3,4-dihydroxyphenylalanine
LMAN	Lateral magnocellular nucleus of the anterior nidopallium
LMM	Linear Mixed Model
MDP	Markov Decision Process
ML	Medial-Lateral axis
nXIIts	Tracheosyringeal portion of the hypoglossal nucleus
OD ratio	Optical Density ratio
RA	Robust nucleus of the acropallium
RPE	Reward Prediction Error
SEM	Standard Error of the Mean
SNc	Substantia Nigra pars compacta
TD(λ)	Temporal Difference learning with discount factor λ

TH-DAB	A histology stain treated with antibodies to tyrosine hydroxylase and visualized using diaminobenzidine
VTA	Ventral Tegmental Area
WN	White noise

1 CHAPTER I: INTRODUCTION AND LITERATURE REVIEW

Sensorimotor learning refers to all types of learning in which organisms refine the motor behaviors they produce by using the sensory feedback received during motor production (Krakauer and Mazzoni 2011). It is the process by which Roger Federer learns to serve either down the T or out-wide through almost imperceptible variations in racquet angle as well as how human babies learn to use utensils over months of practice by first getting food all over their face and clothes. As such, understanding sensorimotor learning is an extraordinarily broad field of study and is usually addressed by subdividing sensorimotor learning into distinct categories. For the purpose of this dissertation, I will be classifying sensorimotor learning into two overarching types, namely *reinforcement learning* and *sensorimotor adaptation* (Wolpert, Diedrichsen, and Flanagan 2011). Any sensorimotor learning in which the feedback signal is an external and explicit appetitive or aversive cue is referred to as reinforcement learning. On the other hand, if the only feedback available to the organism is evaluation of its own sensory feedback then I will refer to such learning as sensorimotor adaptation. This dissertation builds on previous work that has shown that dopamine is involved in reinforcement learning in both primates (Ljungberg, Apicella, and Schultz 1992) and songbirds (Hoffmann et al. 2016, Ljungberg, Apicella, and Schultz 1992). The scientific question at the heart of this dissertation is therefore the following: “*Is dopamine involved in sensorimotor adaptation? If yes, how?*”

This introduction is structured into the following sections. I will first give an overview of reinforcement learning and sensorimotor adaptation as well as a brief survey of some of the other ways sensorimotor learning has been classified. I will then discuss explicitly the role of dopamine in each prefacing it with a short history of dopamine research prior to work implicating it in learning. Third, I will talk about the songbird as a model system in which to study the role of

dopamine in sensorimotor learning. Finally, I will talk about the hierarchical structure of the data we obtain and analyze from songbirds, the history of how this structure was discovered to be a problem for statistical tests and improved statistical approaches currently implemented to address this problem.

1.1 Sensorimotor Learning

1.1.1 Reinforcement Learning

In this sub-section, I will provide a brief overview of reinforcement learning as it relates to the work in this dissertation. First, I will outline the mathematical theory underlying reinforcement learning from a historical perspective culminating with the most recent advances in deep reinforcement learning technology in section 1.1.1.1. Second, I will provide examples of how reinforcement learning theory has driven experiments and discoveries in biology in section 1.1.1.2 and how those results motivated the experiments presented in this dissertation.

1.1.1.1 Mathematical Framework underlying Reinforcement Learning

Reinforcement learning has been studied independently (with some convergence) in biological contexts as well as mathematical frameworks. While biologists think of reinforcement learning in terms of psychological tasks and Pavlovian conditioning, largely initialized by the seminal work by Rescorla and Wagner (Rescorla and Wagner 1972, Rescorla 1969, Rescorla 1968, Wagner 1969, Wagner, Logan, and Haberlandt 1968), reinforcement learning has primarily been understood in the context of mathematical theory and computer science algorithms (Sutton and Barto 1998, Kaelbling, Littman, and Moore 1996). Broadly speaking, reinforcement learning involves the process of converging upon an optimal policy to follow for an agent in an environment using the cues (positive or negative) provided by the environment so as to maximize the positive feedback received over time. When the agent is in the environment, the agent is in a particular

state from which the agent may choose among a set of actions available to transition to a new state. Based on the action chosen, the state of the environment changes and as per a probability distribution unknown to the agent, a reward (or punishment) is delivered. The agent must use the knowledge so obtained over time to decide on a policy that controls which action to select for each state the agent visits. The optimal policy should maximize rewards and minimize punishments.

A historically significant example of reinforcement learning that has been analyzed in detail is the multi-arm bandit task (Berry and Fristedt 1985). Briefly, the one-arm bandit task is a single step environment where a single action results in a reward with a fixed probability unknown to the agent. The multi-arm bandit extends this by having multiple possible actions each with different probabilities of reward. The agent gets a fixed number of actions it may take in total. This task therefore delivers instant reward, or lack thereof, following the action and the environment is reset. However, typical tasks do not reset the environment following every action nor does every action lead to reward. In fact, most cases require multiple actions in a particular sequence to obtain a reward such as navigating a maze. In this case, it becomes necessary to assign credit temporally to the various actions taken along the way that resulted in the eventual delivery of the reward (Sutton 1984, Bradtke and Barto 1996, Tesauro 1992, Tsitsiklis and Van Roy 1997) and the framework of the multi-arm bandit is insufficient to solve such problems.

Markov Decision Processes (MDPs) have been widely used to model and study reinforcement learning with temporal credit assignment as described above (Bellman 1966, Howard 1960, Puterman 2014). An MDP is a process in which the transitions underlying the model are Markovian in nature (or a Markov chain). A Markov chain, in turn, is a process in which the probability of transitions from a given state are independent of the history of transitions made to reach that state. MDPs are extensions of Markov chains in that they allow the agent in the

environment to choose an action to perform in order to transition to the next state. Therefore, the outcome is partially random and partially under the control of the agent. MDPs have been studied extensively in a variety of contexts. Some of the properties that made MDPs particularly useful to study and model reinforcement learning are the following:

1. MDPs can be shown to have an optimal deterministic stationary policy (Bellman 1966).
This means that there exists an optimal policy solution for the agent to follow if found.
2. The optimal policy can be determined by finding the optimal value function. The value function in turn refers to the maximal temporally discounted reward the agent can obtain from the environment if the agent follows the optimal policy. The optimal value function can in turn be found by initializing randomly and updating it iteratively based on inputs from the environment until convergence (Bellman 1966, Bertsekas 1987).
3. The greedy policy, i.e., the policy that takes the best known value at a given point in time is shown to be optimal in a finite number of steps (Bertsekas 1987). Additionally, the greedy policy is shown to have an easy to evaluate stopping criterion based on the maximal difference between two successive value functions obtained through greedy iteration (Williams and Baird 1993).

However, as powerful as MDPs were, they were found to be prohibitively expensive to compute. In practice, solving MDPs requires a linear order of time on the number of actions and a polynomial order of time on the number of states per iteration or time step. As the number of states grows, MDPs become practically impossible to solve very quickly.

The computational cost of MDPs as described above stems from the fact that one may not know when a reward will be delivered. As a result, one has to hold in memory all the actions taken to reach the “final” state when reward is delivered. However, a way to reduce the computational

cost of MDPs is to simply retain the immediate reward associated with a particular action and the expected value of the new state being transitioned into (or more concretely, the difference between them which is called the reward prediction error). The agent can then use the prediction error to update the value function for all previous states with an active eligibility trace defined by a parameter λ . When $\lambda = 0$, the value of only the most recent state is updated. When $\lambda = 1$, the values of all previous states are updated. The algorithms that implement this type of learning are referred to as “model-free” learning methods and primarily take on forms such as TD(λ) learning where TD refers to Temporal Difference (Sutton 1984, Sutton 1988, Dayan and Sejnowski 1994) and Q-learning (Watkins and Dayan 1992, Peng and Williams 1994). Since I introduced the term “model-free”, I note that I will briefly discuss model-free versus model-based methods in Section 1.1.3.1.

The popularity of the temporal difference type of reinforcement learning received a further boost from studies in neuroscience that discovered a role for dopamine as the potential prediction error signal (Ljungberg, Apicella, and Schultz 1992, Schultz, Dayan, and Montague 1997, Schultz 1998). I will discuss these studies in detail in section 1.2.2. However reinforcement learning itself, following the efficiency of the above mentioned model-free methods and the advent of improved memory and computing capacities of modern day computers, underwent an explosion in capability and the range of problems such algorithms could solve. Reinforcement algorithms were used to solve common games such as Backgammon (Tesauro 1992, Tesauro 1995) as well as to explain behaviors in economic games as in those related to game theory (Erev and Roth 1998). Improvements were also made in the algorithms’ speed and efficiency (Sutton et al. 2000, Sutton, Precup, and Singh 1999, Crites and Barto 1996). Other notable improvements involved developing inverse reinforcement learning, i.e., the ability to predict the reward schedule by observing the assumed optimal behavioral policy (Abbeel and Ng 2004, Ng and Russell 2000, Ziebart et al. 2008)

and the extension of reinforcement algorithms to effectively handle multiple agents interacting with the same environment (Hu and Wellman 1998, Claus and Boutilier 1998, Bu, Babu, and De Schutter 2008). This was particularly relevant since animal subjects have been shown to not just learn from their own experiences but also by observing the experiences of others (Dawson and Foss 1965, Tomasello et al. 1987, Fiorito and Scotto 1992).

The most recent major advance in reinforcement learning algorithms came with the advent of “deep learning” in which a multi-layered neural network with several hundreds of thousands of artificial “neurons” and several tens of millions of parameters was trained to essentially halve the error rate as compared to then state-of-the-art algorithms for image classification (Krizhevsky, Sutskever, and Hinton 2012, LeCun, Bengio, and Hinton 2015). From its introduction in computer vision, deep learning has since found success in fields such as speech recognition (Hinton et al. 2012), sentence recognition for language translation (Sutskever, Vinyals, and Le 2014) and reinforcement learning (Mnih et al. 2015). Deep reinforcement learning has been used to train a system to successfully complete Atari 2600 games which have a much higher dimensional sensory input than tasks traditionally solvable by reinforcement learning algorithms (Mnih et al. 2015). One of the most visible successes of deep reinforcement learning was the successful ability of an algorithm to learn to beat both the European champion and the World champion in a game of Go, a game thought to be too complex for an algorithm to successfully master since the game has such a vast search space (Silver et al. 2016). The researchers then advanced this algorithm further by showing that they could train a new algorithm to teach itself to play Go simply using reinforcement learning without any external human input (Silver et al. 2017). As impressive as the recent exploits of deep reinforcement learning are, it is worth keeping in mind that these are very specific advances and are still vulnerable to adversarial attacks requiring incremental improvements such as the

algorithm that could complete the Atari 2600 games being unable to do so if the games screenplay was shifted laterally (Mnih et al. 2016, Van Hasselt, Guez, and Silver 2016, Marcus 2018). Additionally, while TD(λ) reinforcement learning may reveal insights into how the brain performs similar computations such as the need for an eligibility trace or the role of a temporal difference signal in learning (see section 1.2.2 for details), deep networks are more akin to a black box. As mentioned previously, deep networks have hundreds of thousands of computational units and several millions of parameters making interpretability of the algorithm/model inaccessible. Hence deep networks are currently of limited interest in neuroscience though there is some work attempting to change this (Shwartz-Ziv and Tishby 2017, Koh and Liang 2017).

1.1.1.2 Reinforcement Learning in Biology

Sensorimotor learning in the form of reinforcement learning has been studied in animals since the days of classical conditioning (Pavlov and Gantt 1928). However, as mentioned briefly above, it was the work of Rescorla, Wagner, and colleagues (Rescorla 1969, Rescorla 1968, Wagner 1969, Wagner, Logan, and Haberlandt 1968, Rescorla and Wagner 1972) that put a mathematical theory behind Pavlovian conditioning and contributed to refined experimentation to understand how reinforcement learning is implemented in organisms. Furthermore, operant conditioning was found to be a powerful tool to study the mechanisms underlying reinforcement learning (Skinner 1963). As mentioned briefly above, reinforcement learning in biology got its next big advance following the discovery of dopamine as a potential reward prediction error signal matching its use in mathematical algorithms (Ljungberg, Apicella, and Schultz 1992, Schultz 1998, Schultz, Dayan, and Montague 1997). I will elaborate on these studies in section 1.2.2.

1.1.2 Sensorimotor Adaptation

I refer to adaptive behavioral changes that result from evaluation of sensory feedback arising from one's own motor performance as "sensorimotor adaptation." Here, as opposed to an external reward signal in the case of reinforcement learning, the organism has to compute an error signal by comparing the produced motor output to the desired or ideal motor output. Due to the fact that the error signal is produced internally, studying sensorimotor adaptation in a laboratory setting can be challenging. However, a few paradigms have been developed in which the error signal can be controlled and I will expand on the results from those studies in this section.

One of the first and most intensively studied sensorimotor adaptation paradigms involves saccadic eye movements (McLaughlin 1967, Pélisson et al. 2010). Saccadic eye movements refer to the ability of the eye to make quick movements to stabilize at a point of interest in a visual field and bring it to focus. Sensorimotor adaptation in eye saccade movements was introduced by using the double-step target paradigm (McLaughlin 1967). Briefly, a target is presented to which the subject must make an eye saccade and while the subject is moving, the target is shifted either closer or further away. If trained repeatedly, the subject learns to saccade to the shifted position given the original position of the target. Saccades are fast movements happening on the time scale of hundreds of milliseconds (Fischer and Weber 1993, Fischer et al. 1993) and as such are too fast for error correction mid-flight either from visual information though some have reported small changes in trajectory (Gaveau et al. 2003) or input from proprioceptive feedback (Guthrie, Porter, and Sparks 1983, Lewis et al. 2001). Hence, almost all error correction is thought to happen upon completion of the saccade from the resulting sensory feedback that the movement made did not reach the intended target. A classic manifestation of adaptation is the after-effect, i.e., after the end of the task where the target had been shifted, subjects still compensate for an anticipated shift and

result in overshooting the target in the opposite direction (Gibson 1933). Saccade eye movements have been used to fit computational models of sensorimotor learning (Ethier, Zee, and Shadmehr 2008b, a, Findlay and Walker 1999, Robinson 1973) as well as to study the neural mechanisms underlying sensorimotor adaptation (Optican and Robinson 1980, Sparks 1986, Munoz 2002, Munoz et al. 2000).

Another experimental paradigm that has been used to study sensorimotor adaptation extensively is the prism goggles adaptation paradigm (Stratton 1897, 1896). Briefly, prism goggles are used over a subject's eyes to cause some manipulation of the visual input the subject receives. Typically these manipulations can take the form of image inversion (Stratton 1897, 1896, Foley Jr 1940), rotation or manipulation of curvature (Gibson 1933, 1937, Gibson and Radner 1937). However, it was following the study that showed that adult human subjects required active movement to completely adapt to a rotation in visual field, an adaptation absent without active movement that prism goggles began to be used extensively to study sensorimotor adaptation (Held and Bossom 1961, Held and Gottlieb 1958). Through studies on inter-manual transfer, i.e., learning on one arm generalizing to behavior for the other arm (Hamilton 1964) and effects of telling subjects about the manipulation beforehand (Welch 1972), a hypothesis formed that sensorimotor adaptation occurred due to comparison with an internal model of relationships between visual feedback and proprioception. As research on prism goggles adaptation continued (Rossetti et al. 1998, Welch et al. 1993), it became clear that the cerebellum was involved in this task since human patients with damage to cerebellar nuclei and non-human primates with cerebellar lesions showed deficits in adapting to the prism goggles rotational shift (Weiner, Hallett, and Funkenstein 1983, Martin et al. 1996, Baizer, Kralj-Hans, and Glickstein 1999, Morton and Bastian 2004). These studies led to the hypothesis that the cerebellum is responsible for generating internal models of

movement with which to compare to realize sensorimotor adaptation (Wolpert, Miall, and Kawato 1998, Blakemore, Frith, and Wolpert 2001, Ito 2008).

A task that has some of the same structure as the prism goggles adaptation task but comes with greater degree of experimental control is the center out reaching task in the presence of a rotational force-field. Georgopoulos popularized the concept of having subjects reach out to different directions from a central point (Georgopoulos et al. 1982) and further showed that neurons in the motor cortex were tuned for direction of the movement (Georgopoulos et al. 1983, Georgopoulos, Schwartz, and Kettner 1986). Following the development of the manipulandum (Shadmehr, Mussa-Ivaldi, and Bizzi 1993), it became possible to study how subjects responded to rotational challenges in a reaching task and that paradigm became a popular and reliable one to study sensorimotor adaptation (Brashers-Krug, Shadmehr, and Bizzi 1996, Caithness et al. 2004). Neural correlates underlying such learning strengthened the evidence for internal models built by the cerebellum (Diedrichsen et al. 2005, Tseng et al. 2007). Finally, while such tasks have formerly been limited to human and non-human primate subjects, it has recently been successfully demonstrated in rodents allowing for more specific investigations into the neural circuitry underlying learning of the task (Mathis, Mathis, and Uchida 2017) leveraging genetic tools available for rodent models. In all these tasks though, it seems as if an internal model is necessary to observe learning. I will discuss this in some detail in section 1.1.3.1.

The final paradigm I must mention that has been used to study sensorimotor adaptation relates to using shifted auditory feedback to affect changes in vocal output among organisms. Following the success of observing sensorimotor adaptation in response to altered visual feedback, there was an open question as to whether this was a uniquely visual phenomenon or other senses could also be similarly manipulated to drive adaptation. By using headphones to deliver altered

auditory feedback to humans, it was shown that the vowel sounds subjects made could be altered systemically (Houde and Jordan 1998, 2002). Since non-human primates do not have an extensive vocal repertoire, studies of altered auditory feedback was restricted to human subjects (Guenther 2006, Villacorta, Perkell, and Guenther 2007, Shum et al. 2011, Tremblay, Shiller, and Ostry 2003) until it was shown that songbirds are capable of adapting to pitch shifted auditory feedback of their songs in a similar manner (Hoffmann et al. 2012, Kelly and Sober 2014, Sober and Brainard 2012, 2009). As I will detail in Chapter 2, I used this paradigm to study the central question of my thesis: the role of dopamine in sensorimotor adaptation.

1.1.3 Other classifications of Sensorimotor Learning

While I have focused on reinforcement learning and sensorimotor adaptation above, there are many other ways by which one can classify different forms of learning. In the next section, I will discuss two other classifications that are relevant to my experiments. Note however that these classification schemes are not mutually exclusive and do not encompass all the ways in which one may break sensorimotor learning into smaller categories. This is represented diagrammatically in Figure 1.1.

1.1.3.1 Model-free versus Model-based learning

A theme that may have emerged from my discussion of the previous two sections is that reinforcement learning tends to primarily be discussed in terms of model-free algorithms while sensorimotor adaptation is primarily thought to be model-based. However, in order to fully classify algorithms as model-free or model-based, we must first understand the difference between the two categories.

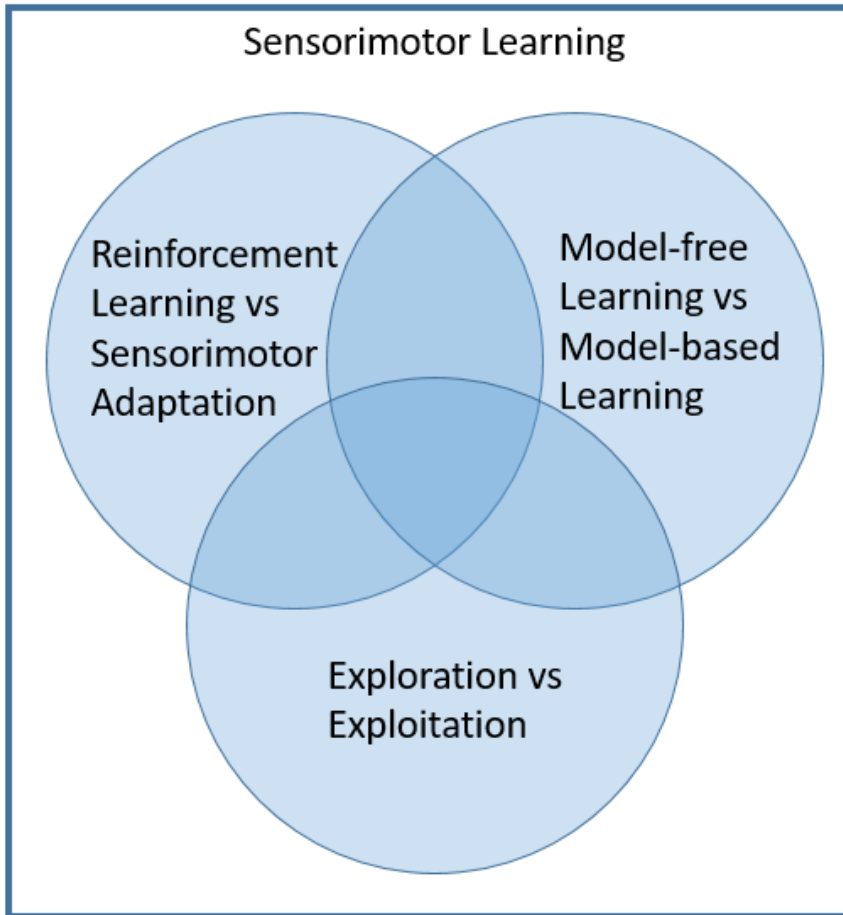


Figure 1.1: Classifications of sensorimotor learning.

Various classifications of sensorimotor learning discussed in Section 1.1. I primarily focus on the distinction between reinforcement learning and sensorimotor adaptation for the purpose of this dissertation (elaborated in section 1.1.1 and 1.1.2 respectively). Alternative classifications discussed briefly are model-free and model-based learning as well as exploration versus exploitation (see section 1.1.3). As the figure shows, all these classifications are of sensorimotor learning itself and therefore, all overlap with each other.

The primary difference between a model-free and a model-based algorithm lies in whether or not the underlying structure of the environment (the model), the transition probabilities between states for example, is learned or not. If one considers a reinforcement learning algorithm, a model-free algorithm will track only the received rewards and the estimated value function (i.e., the time-discounted sum of all future rewards from that state) at each state, updating said estimate each time it visits a particular state. A model-based algorithm on the other hand will also learn the structure of the environment such as the probabilities of reward at certain states as well as the new states

that each action among those available will lead to from the present state. Hence on the surface, model-free algorithms are very quick to compute but are slow to react to small changes in the environment such as cutting off of the most direct path in a maze. Model-based algorithms on the other hand are computationally expensive and take a long time to train but once trained, can adapt to such small changes very flexibly. In practice however, efficient and commonly implemented reinforcement learning algorithms are neither entirely model-free nor completely model-based but lie on a spectrum between the two extremes so as to minimize the weaknesses of each of the extremes (Watkins and Dayan 1992, Sutton 1991, Peng and Williams 1994, 1993).

While I described the difference between model-based and model-free algorithms using reinforcement learning algorithms, the definition can be applied to reinforcement learning as well as sensorimotor adaptation. As much as the cerebellum has been implicated in producing internal models necessary for sensorimotor adaptation, there have been experimental results observed that cannot be explained by model-based learning alone (Haith and Krakauer 2013). It has also been observed in animal reinforcement learning tasks that when attempting to model animal behavior, a purely model-free approach or a purely model-based approach work and fail according to context. There is an ongoing debate as to the usefulness of classifying behaviors as model-free or model-based but there is accumulating evidence that animals vary their strategies depending on the context (Dayan and Berridge 2014, Lee, Shimojo, and O'Doherty 2014, Russek et al. 2017, McDannald et al. 2012).

1.1.3.2 Exploration versus Exploitation

I wanted to conclude this section with a brief overview of another distinction often talked about in sensorimotor learning and reinforcement learning in particular – exploration versus exploitation. It makes the most sense to understand this distinction in terms of the classic multi-armed bandit

problem as discussed in section 1.1.1.1 above. Consider an agent in an environment where there are k gambling machines each with a different probability of payoff unknown to the agent. The agent is allowed to gamble N times. At the end of each gamble, the agent collects the payoff, if any, and chooses another machine to play till they reach their limit. What is the optimal strategy to pursue to maximize payoff? Since the probabilities are unknown, the agent must explore to get some sense of the payoff probabilities of some of the machines. But as the agent nears the limit N , the better strategy would be to continually play the machine with the current best estimate of payoff. So where does one decide to make the switch?

There have been in-depth mathematical formulations made to address this problem (Berry and Fristedt 1985, Kaelbling, Littman, and Moore 1996). However, in practice, ad-hoc methods such as a greedy algorithm (Edmonds 1971), random exploration or Boltzmann exploration (a weighted random exploration where the “temperature” can be reduced later to result in exploitation) are easier to implement in algorithms and seem to work as well as other methods. Exploration and exploitation have been studied in animal behavior as well (Watkins 1989, Harley 1981) though in recent times, the distinction has been characterized more as habitual (like exploitation) versus goal-seeking behavior which is closer to exploration (Balleine and O’doherly 2010, Keramati, Dezfouli, and Piray 2011, Schwabe and Wolf 2011, de Wit et al. 2012). I will also note that while I have discussed this exploration-exploitation trade-off in terms of reinforcement learning, one can imagine it occurring constantly in sensorimotor adaptation as one explores a lot initially when learning a new skill but as one gets better, starts exploiting learned tricks to perform at their best. Hence, the classifications expressed in both sections 1.1.3.1 and 1.1.3.2 are applicable to both reinforcement learning and sensorimotor adaptation. This point is made explicit in Figure 1.1.

1.2 A role for Dopamine in Sensorimotor Learning

Learning is an extraordinarily broad topic and there have been numerous brain structures, neuromodulators and neurotrophic factors that have been and are still actively studied in the context of various types of learning. For the purpose of this dissertation however, I will largely restrict myself to the role dopamine plays in sensorimotor learning. This section is divided into the following subsections. I will first provide a brief overview of the roles of dopamine prior to its discovery as a potential reward prediction error signal. I will then discuss its role in reinforcement learning and some other posited theories since. I will end the section with a brief overview of the potential role for dopamine in sensorimotor adaptation which forms the core of my dissertation.

1.2.1 Dopamine before Reward Prediction Error (RPE)

Parkinson's disease was first characterized by James Parkinson in 1817, over 200 years ago (Parkinson 1817, 2002). It was also one of the first major hypothesized roles for dopamine in the brain as researchers discovered that a common symptomology underlying Parkinson's disease was the loss of dopaminergic cells from the Substantia Nigra pars compacta (SNc) (Barbeau 1962, Bernheimer et al. 1973). Additionally, it was reported that manipulations of the nigro-striatal dopamine pathway caused movement deficits in both non-human primate and rodent models (Andén et al. 1966, Ungerstedt 1971, Ungerstedt and Arbuthnott 1970, Poirier and Sourkes 1965, Poirier et al. 1966). This first led to a hypothesis that dopamine was involved in the control of movement. However, studies were also reported that dopamine depleted rodents could still swim in a water maze task (Marshall, Levitan, and Stricker 1976) which was likened to the paradoxical kinesia seen in Parkinson's disease (Jarkowski 1925).

Another major role of dopamine stemmed from studies that linked it to reward and hedonia, (the feeling of liking or enjoying something). Rodents were found to be willing to press a lever for

intracranial direct current stimulation or self-administration of amphetamines. This process in turn could be blocked by dopamine blockers or antagonists administered to the subject (Lippa et al. 1973, Yokel and Wise 1975). Further experimentation also suggested a role for dopamine that acknowledged the contributions to both reward and movement (Stellar, Kelley, and Corbett 1983). However, as popular as the dopamine reward hypothesis became (Hernandez and Hoebel 1988, Wise and Rompre 1989), problems emerged in that dopamine depleted rats still seemed to enjoy eating their food (Berridge, Venier, and Robinson 1989). So, the terminology shifted to saying dopamine signaled the value of the reward which differed from hedonia in that what was now being tracked was not the fact that the subject enjoyed the reward but the degree to which the reward or other stimulus carried a positive or negative association. It was in this context that a role for dopamine as an RPE signal was discovered as I will discuss in section 1.2.2.

1.2.2 Dopamine in Reinforcement Learning


1.2.2.1 The Dopamine RPE hypothesis

Using electrophysiological recordings of neurons in the SNc of non-human primates, Schultz and colleagues investigated the role of dopamine and found that it was correlated with attention and arousal (Romo and Schultz 1990, Schultz 1986, Schultz and Romo 1990). However, two key studies in 1992 were most important for codifying the role of dopamine and the basal ganglia, with the ventral striatum in particular, in reward processing and reinforcement learning. Specifically, one study found that dopaminergic neurons in the VTA showed responses that would later be recognized as reward prediction error (RPE) like (Ljungberg, Apicella, and Schultz 1992) while the other found that neurons in the ventral striatum seemed to track expected value of future rewards in predictable contexts (Schultz et al. 1992). Separately, work in bees also implicated the dopamine system in reward processing (Montague, Dayan, and Sejnowski 1996). The researchers


who worked on bees collaborated with the Schultz group and designed an experiment to test the role of dopamine in reward processing. Dopamine was found to respond in a very particular fashion

Do dopamine neurons report an error in the prediction of reward?

No prediction
Reward occurs



Reward predicted
Reward occurs



Reward predicted
No reward occurs




Figure 1.2: Response of dopamine neurons to rewards and reward predictions.

CS refers to a conditioned stimulus that is predictive of the reward and R refers to reward. The panels show the raster plots of dopaminergic neuron activity across several trials and the histograms of binned firing rates are at the top of each panel. (Top panel) Dopamine neurons fire in response to an unexpected reward without a preceding CS. (Middle panel) Dopamine neurons respond to the CS when it occurs but do not respond noticeably to the actual reward delivery. (Bottom panel) Dopamine neurons respond to the CS as before but there is no reward delivered. The dopamine neurons show a dip in firing rate around the time the reward would have otherwise been delivered. From Schultz, Wolfram, Peter Dayan, and P. Read Montague. "A neural substrate of prediction and reward." *Science* 275, no. 5306 (1997): 1593-1599. Reprinted with permission from AAAS.

to reward presentation (see Fig. 1.2). Specifically, dopamine neurons were found to respond to an unexpected reward (Fig. 1.2 top panel). If the reward was paired reliably with a predictive stimulus, over time dopamine neurons began responding to the stimulus and responding less to the actual reward (Fig. 1.2 middle panel). Finally, if the reward was not delivered following the predictive stimulus, a drop in dopamine neuron firing was observed around the time the reward should have been delivered (Fig. 1.2 bottom panel). This led to a hypothesis that the phasic activity of dopaminergic neurons in the VTA closely tracked the reward prediction error signal in reinforcement learning (Schultz 1998, Schultz, Dayan, and Montague 1997, Glimcher 2011). As was being detailed in reinforcement learning algorithms, model-free reinforcement learning required the tracking of a reward prediction error (RPE) signal (Sutton 1984, Sutton 1988, Sutton and Barto 1998). Due to the nature of dopamine neuron firing described above, changes in dopamine concentration were hypothesized to convey that RPE signal and the basal ganglia was proposed to be the neural circuit that performed the required computations with that signal to realize reinforcement learning in biological systems.

Following the above breakthrough, extensive research was carried out to better characterize the dopamine signaling in reinforcement learning. Prior research had looked at predictive stimuli that always preceded a reward. A naturally open question to follow that finding was if the reward associated with a predictive stimulus was probabilistic, how does dopamine neuron firing encode uncertainty of future reward? It was found that dopamine neurons do indeed track reward probabilities and that they shifted their firing rates with the expected value of the reward and their gains with the uncertainty in reward delivery (Fiorillo, Tobler, and Schultz 2003, Tobler, Fiorillo, and Schultz 2005). Neural correlates for reinforcement learning were discovered using rodent models and human subjects (Pessiglione et al. 2006, Valentin, Dickinson, and O'Doherty 2007,

Samejima and Doya 2007) implicating several areas of cortex such as the orbitofrontal cortex, anterior cingulate cortex and medial prefrontal cortex in addition to the basal ganglia. These results further spurred interest in the neural basis of decision making (Körding and Wolpert 2006, Gold and Shadlen 2007).

With the advent of optogenetics (Boyden et al. 2005), i.e., genetic manipulations of targeted neuronal populations to insert a light-sensitive ion channel into neurons so as to manipulate their electrical activity with light pulses, it became possible to target dopaminergic cells and their targets in the striatum precisely both in time and in genetically targeted cell subpopulations in space. Early work showed results consistent with the RPE hypothesis for dopamine along with roles in movement control associated with Parkinsonism (Kravitz and Kreitzer 2012, Kravitz, Tye, and Kreitzer 2012). It was also shown that direct optical stimulation of dopaminergic neurons in an RPE-like fashion was sufficient to cause a learned association between an otherwise neutral cue and both a rewarding outcome (Steinberg et al. 2013) and an aversive outcome (Chang et al. 2016). Optogenetics was also capable of solving the long-standing criticism that most prior studies of dopamine in RPE were based on waveform identification of dopamine neurons which were not always reliable (Margolis et al. 2006, Lammel et al. 2008). By expressing the light-sensitive channels only in the dopaminergic neurons of interest, Uchida and colleagues were able to ensure they were characterizing dopaminergic cell firing and reported that the cells did indeed fire RPE-like signals (Cohen et al. 2012) and that the firing across the population of VTA dopaminergic neurons was fairly homogenous, i.e., both individual and population activity of the dopamine neurons could be explained using just two parameters (Eshel et al. 2016).

1.2.2.2 Problems with the Dopamine RPE hypothesis

As compelling as the RPE hypothesis of dopamine was, it was not without its flaws and complications (Dayan and Niv 2008, Berridge 2007, Watabe-Uchida, Eshel, and Uchida 2017). One of the primary sources of concern was the fact that dopamine neurons have a baseline firing rate of 3 to 5 Hz. They are capable of signaling positive RPE by increasing their firing rate up to 30 Hz or so but for negative RPE, they can only reduce it to zero. Hence, there seemed to be an inherent asymmetry in the level of signal between positive and negative RPE (Fiorillo 2013). However, whether this asymmetry is a problem is still an ongoing debate since there has been evidence that even if dopamine neurons cannot modulate their firing symmetrically, the amount of dopamine present in the ventral striatum still varies symmetrically between rewarding and aversive stimuli (Hart et al. 2014). It has also been proposed that a separate nucleus, the lateral habenula, is the one responsible for coding the negative RPE signal (Matsumoto and Hikosaka 2007, Tian and Uchida 2015).

Another common criticism of the dopamine RPE hypothesis is that not all dopamine neurons exhibit RPE-like responses. Several studies have reported that subsets of dopaminergic neurons show activation responses to both rewarding and aversive stimuli and the cues that predict them (Joshua et al. 2008, Fiorillo, Song, and Yun 2013, Matsumoto and Hikosaka 2009, Lerner et al. 2015). Since such neurons are found primarily in the SNc, it has been argued that such neurons are coding for behaviorally important stimuli and are not actually involved in updates to the estimate of the value function (Matsumoto and Hikosaka 2009, Watabe-Uchida, Eshel, and Uchida 2017, Berridge 2007). It has also been argued that the responses of the dopaminergic neurons actually have multiple phases one of which encodes for novelty or salience of the stimulus as

described above while another encodes for the RPE signal (Fiorillo, Song, and Yun 2013, Schultz 2016).

A related criticism arises from studies of mice that have a genetic mutation such that they do not produce dopamine in their brains from birth (Zhou and Palmiter 1995). Cannon and Palmiter showed that such mice were capable of developing a preference ratio for sucrose solution over plain water to the same degree as unmanipulated controls even though they had no dopamine in their brains at the time (Cannon and Palmiter 2003). It was also shown that such mice, if pre-treated with caffeine, were capable of learning a conditioned place preference for obtaining morphine (Hnasko, Sotak, and Palmiter 2005) and successfully navigating a T-maze (Robinson et al. 2005). It has been argued that these results call into question whether dopamine is necessary for reinforcement learning, thereby questioning the validity of the RPE hypothesis (Berridge 2007).

There have been several other criticisms of the dopamine RPE hypothesis. I have described previously how dopamine signals are correlated with movements and several studies have argued that the RPE signals observed are derived correlations from movement artifacts (Jin and Costa 2010) or that separate populations encode for RPE and movement (Howe and Dombeck 2016). Others have argued that since reinforcement learning in organisms seems to be neither purely model-free nor model-based (see section 1.1.3.1), RPE may be but one component of a larger error signal referred to as sensory prediction errors and that it is the sensory prediction errors that dopamine truly encodes (Gardner, Schoenbaum, and Gershman 2018, Momennejad et al. 2017). A final major criticism is that most of the studies in which the role of dopamine as an RPE signal is studied are highly artificial and not reflective of natural behavior (I will discuss in section 1.3.2 how songbird research has attempted to circumvent this problem). In addition, the subjects taking

part in the study are usually overtrained on the task and so the dopamine responses may not reflect learning. Coddington and Dudman reported recently that when dopamine responses were recorded as a rodent learned a novel association between sensory cues and appetitive rewards, the neurons showed both sensory-cue related responses and movement-initiation responses that resulted in apparent RPE correlates (Coddington and Dudman 2018). Further experiments will be required to contextualize these results in the larger RPE framework.

In spite of all these criticisms of the dopamine RPE hypothesis, there is a preponderance of evidence of dopamine conveying RPE-like signals in the ventral striatum of the basal ganglia in simple tasks (Watabe-Uchida, Eshel, and Uchida 2017). It is also being recognized that dopamine performs multiple other functions that will not be discussed in detail here including control of movement vigor (Beierholm et al. 2013, Niv et al. 2007, Panigrahi et al. 2015, Turner and Desmurget 2010), modulation of attention and manifestation of Attention Deficit Hyperactive Disorder (Cook Jr et al. 1995, Huang et al. 2015, Nieoullon 2002) as well as sleep (Monti and Monti 2007) and regulation of circadian rhythms (Korshunov, Blakemore, and Trombley 2017). Smart experimental designs will be required to study the interactions between these various proposed functions for dopamine to uncover a unified theory for the role of dopamine in the brain, if one exists.

1.2.3 Dopamine in Sensorimotor Adaptation

While there is an extensive literature on the role of dopamine in reinforcement learning as discussed in section 1.2.2, the role of dopamine in sensorimotor adaptation has been characterized much less. Dopaminergic neurons in the brain have been found to decrease their activity with age in humans and such a decline has been correlated with declines in motor function (Volkow et al. 1998). Furthermore, in prism goggles adaptation tasks (see Section 1.1.2), elderly subjects show a

distinct deficit in adaptation as compared to younger subjects (Bock 2005, Seidler 2006) though they did not show a deficit in recovering from the after-effect of the adaptation (Bock 2005). Furthermore, the deficit in sensorimotor adaptation was specific to adaptation tasks and did not extend to sequence learning tasks resulting in the hypothesis that there were age-related deficits in cerebellar mediated motor skills (Seidler 2006).

A study conducted in patients with Parkinson's disease (PD) revealed that patients not on L-DOPA treatment, therefore with less dopamine, showed stronger adaptation deficits than when on their L-DOPA treatment or against age-matched controls for a task in which participants were required to point directly ahead with rotational manipulation of their visual field using prism goggles (Paquet et al. 2008). A similar reduction in adaptation among PD patients was observed when adapting vocal motor behavior to an induced change in pitch of auditory feedback (Mollaei, Shiller, and Gracco 2013, Abur et al. 2018). A more finely controlled experiment involving a reaching task in a force-field through virtual reality reported that PD patients had trouble adapting to large perceivable sensory feedback errors but could adapt as well as controls to small gradually introduced errors even if the final error magnitude was quite large (Mongeon, Blanchet, and Messier 2013).

Hence, while there have been some studies linking dopamine to sensorimotor adaptation, they have largely been restricted to studies in PD patient populations. PD patients have a host of co-morbid conditions such as cognitive and executive deficits in addition to large motor deficits (Jankovic 2008, Dubois and Pillon 1996, Lees and Smith 1983, Cooper et al. 1991). In addition, PD patients also suffer other neuropathologies in addition to dopaminergic loss such as accumulation of Lewy bodies composed of aggregated α -synuclein and degeneration of several other nuclei such as the dorsal raphe, the locus ceruleus and motor vagal nucleus among others

(Jellinger 1991, Dickson 2012). As a result, the role of dopamine specifically in sensorimotor adaptation has been difficult to isolate. This question therefore forms the core of this dissertation and I detail my results in Chapter 2.

1.3 Songbirds as a model system

The experiments reported in this dissertation all used male Bengalese finches (*Lonchura striata* var. *domestica*), a type of songbird, as a model system in which to study sensorimotor learning. The qualities of the songbird as a model system in which to study sensorimotor learning are discussed below. This section is divided into the following subsections. In section 1.3.1, I will give a brief overview about the advantages of the songbird as a model system. In section 1.3.2, I will discuss the advances made in understanding reinforcement learning using songbirds and in section 1.3.3, I will discuss the same for sensorimotor adaptation in songbirds.

1.3.1 Overview

Human infants learn speech from their parents or other adults around them by the process of vocal imitation (Kuhl and Meltzoff 1996, Meltzoff and Moore 1983). Songbirds are among a select number of other organisms that similarly learn their vocalizations via vocal imitation (Price 1979, Doupe and Kuhl 1999, Lipkind et al. 2013). Furthermore, the song is actively maintained by auditory feedback such that deafening the birds causes song to deteriorate over time (Nordeen and Nordeen 1992, Sohrabji, Nordeen, and Nordeen 1990). Since song is a male-specific courtship behavior in Bengalese finches, only the males sing and so all experiments for this dissertation have been limited to male Bengalese finches. Songbirds have been found to have an extensive network of brain nuclei that are exclusively involved in song production and learning (Nottebohm, Stokes, and Leonard 1976, Scharff and Nottebohm 1991). As detailed in Figure 1.3 a), the song system has two major pathways for song control. One pathway which involves the HVC (used as a proper

name; see (Reiner et al. 2004) for current nomenclature in avian neuroscience) and the RA (Robust nucleus of the Acropallium) projects to the brainstem nucleus nXIIts which in turn projects to the muscles controlling the syrinx that produces song. Lesions of HVC or RA either severely impair or abolish song production (Nottebohm, Stokes, and Leonard 1976). Hence, this pathway is referred to as the motor production pathway with RA being considered the motor cortex analog and HVC being equated to the premotor cortex. The other major pathway is the Anterior Forebrain Pathway (AFP) which is thought to be involved in song learning and maintenance in adulthood but not production. The AFP includes the basal ganglia nucleus Area X, a part of the songbird thalamus DLM and a cortical nucleus LMAN, considered the output nucleus of the AFP. Lesions of LMAN were found to impair song learning in juveniles but not affect adult song production (Bottjer, Miesner, and Arnold 1984). Similarly, lesions of Area X were found to selectively impair song learning and prevent degradation of song following deafening (Sohrabji, Nordeen, and Nordeen 1990, Brainard and Doupe 2000). This dissociation of functional pathways in addition to the fact that invasive manipulations of these brain nuclei do not cause any other gross motor deficits in songbirds (Feenders et al. 2008) has made songbirds an excellent model system with which to study the neural mechanisms underlying sensorimotor learning.

1.3.2 Reinforcement Learning in Songbirds

As I mentioned in section 1.2.2.2, one of the major criticisms of the dopamine RPE hypothesis was the fact that the tasks that organisms performed in order to study RPE were highly simplistic. Specifically, such tasks were either not indicative of the organism's behavior in a natural environment or the organisms had been overtrained in the task prior to recording dopamine activity. Research in songbirds can sidestep that issue to a certain extent. For one, the data we analyze in songbirds consists of bird song which is a spontaneously produced natural behavior.

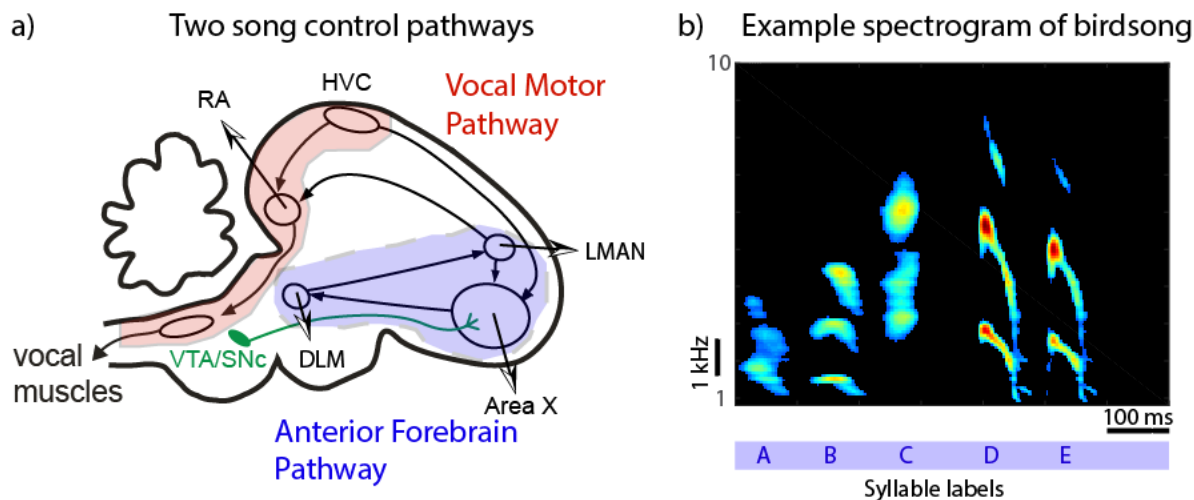


Figure 1.3: Neural pathways involved in and example spectrogram of birdsong.

a) A schematic of the bird brain showing the major song nuclei and their connectivity. Also shown are the classifications into the vocal motor pathway involved in song production and the anterior forebrain pathway (AFP) involved in song learning. b) An example spectrogram of a bird's song. The plot shows the frequency content in the bird's song over time. We can distinguish 5 syllables from the spectrogram shown and assign labels A through E to them. In this example, the pitch of syllables B, D and E can be quantified cleanly while those of A and C cannot. We restrict our analysis to data from syllables whose pitch can be quantified cleanly.

Birds vocalize spontaneously hundreds of times per day. We capture the complexity of their vocal repertoire by recording their song and analyzing the frequency components in the song over time as shown in a sample spectrogram in Figure 1.3 b). Furthermore, adult songbirds maintain their song through auditory feedback and change their songs in responses to auditory perturbations (Sakata and Brainard 2006, 2008). However, it was the development of the white noise (WN) learning paradigm that saw an explosion of research into the mechanisms underlying reinforcement learning in songbirds (Tumer and Brainard 2007).

Even though birds have a large degree of stereotypy in their songs and repeat it several hundreds of times per day, if one were to compute the pitch of each syllable every time they sang, one would find a fair degree of variability in the exact pitch sung. Tumer and Brainard showed that birds could learn actively on this variability. Specifically, they showed that if a portion of this variability, say the lower 50% of pitches produced, were to be targeted with an aversive stimulus in the form of a blast of WN, the birds would shift their pitch upwards over time so as to reduce

the proportion of pitches being targeted by WN (Tumer and Brainard 2007). This shift was bidirectional in nature and one could also make the birds shift further from their baseline pitch by systemically changing the threshold of pitches targeted with WN as the birds shifted their pitch (Andalman and Fee 2009, Warren et al. 2011). Furthermore it was showed that the AFP played a role in this learning since inactivation of the AFP resulted in a regression of the shifted pitch (Andalman and Fee 2009). When WN was introduced, it was originally meant for targeting changes in pitch. However, it was found to be effective in targeting durations of intersyllable gaps as well. Manipulations of the basal ganglia (Area X) impacted the changes in pitch but not the changes in temporal gaps (Ali et al. 2013). There was a lot of interest in uncovering the sensory nuclei involved in such reinforcement learning as well (Canopoli, Herbst, and Hahnloser 2014).

The connection to dopamine in reinforcement learning in songbirds was made when we showed that birds that had dopamine depleted in Area X, the basal ganglia nucleus involved in song learning, showed a deficit in successfully shifting their pitch away from those targeted by WN as compared to unlesioned or saline injected controls (Hoffmann et al. 2016). It was also shown that dopaminergic neurons in the VTA showed RPE like responses to avoiding the WN stimulus much like dopaminergic neurons in rodents or non-human primates in other situations (Gadagkar et al. 2016). Finally, it was also verified using optogenetics that direct optical stimulation or inhibition of VTA/SNc neurons projecting to Area X was sufficient to induce pitch shifts either towards or away from the range of pitches being targeted respectively (Xiao et al. 2018, Hisey, Kearney, and Mooney 2018). Active research in this area involves verifying whether the actor-critic framework of the basal ganglia as observed in other species is also applicable to songbirds (Chen et al. 2018).

1.3.3 Sensorimotor Adaptation in Songbirds

As I had mentioned briefly in section 1.1.2, Sober and Brainard developed a technique to implant customized miniature headphones onto songbirds through which one could provide almost real-time feedback (delay of 10 to 15 ms) of auditory inputs (Sober and Brainard 2009, Hoffmann et al. 2012). These headphones were intended to replace the auditory feedback of the bird and relay a pitch shifted feedback to the bird. Birds learned to correct for this introduced error in pitch by lowering their pitch in response to an upwards pitch shift and by raising their pitch in response to a downwards pitch shift. The birds also displayed the classical after-effect of the headphones or washout of the learned pitch after the pitch shift through the headphones was set back to zero. Hence, this headphones learning paradigm became a useful paradigm to use to study the neural mechanisms underlying sensorimotor adaptation in songbirds.

A curious effect of the headphones learning paradigm was that compensation scaled negatively with the size of the error (Sober and Brainard 2012). Specifically, in terms of percentage compensation, birds compensated most for a 0.5 semitone shift and the least for a 3 semitones shift. Speed of learning was also maximal for the smaller shifts. It was hypothesized that the degree of overlap between the history of produced pitches and the currently heard pitches through the shifted auditory feedback dictated the percent compensation and speed of learning (Sober and Brainard 2012). This was verified through a study of juvenile song birds who had learned most of their song but had a greater degree of variability than fully grown adults (Kelly and Sober 2014). Additionally, computational models using Bayesian inference could successfully capture the dynamics of the learning in response to various error sizes (Zhou et al. 2018, Hahnloser and Narula 2017).

For this dissertation, I am primarily interested in the role dopamine plays, if any, in sensorimotor adaptation. In order to address this question, I replicated the dopamine depletion paradigm of Area X described for reinforcement learning (Hoffmann et al. 2016) and measured the adaptation the birds showed in response to pitch shifted auditory feedback through headphones (Hoffmann et al. 2012, Sober and Brainard 2009). The results of my experiments are discussed in detail in Chapter 2.

1.4 Error Propagation in Hierarchical Data

Before I conclude this introduction, I must give some importance to the problem of error

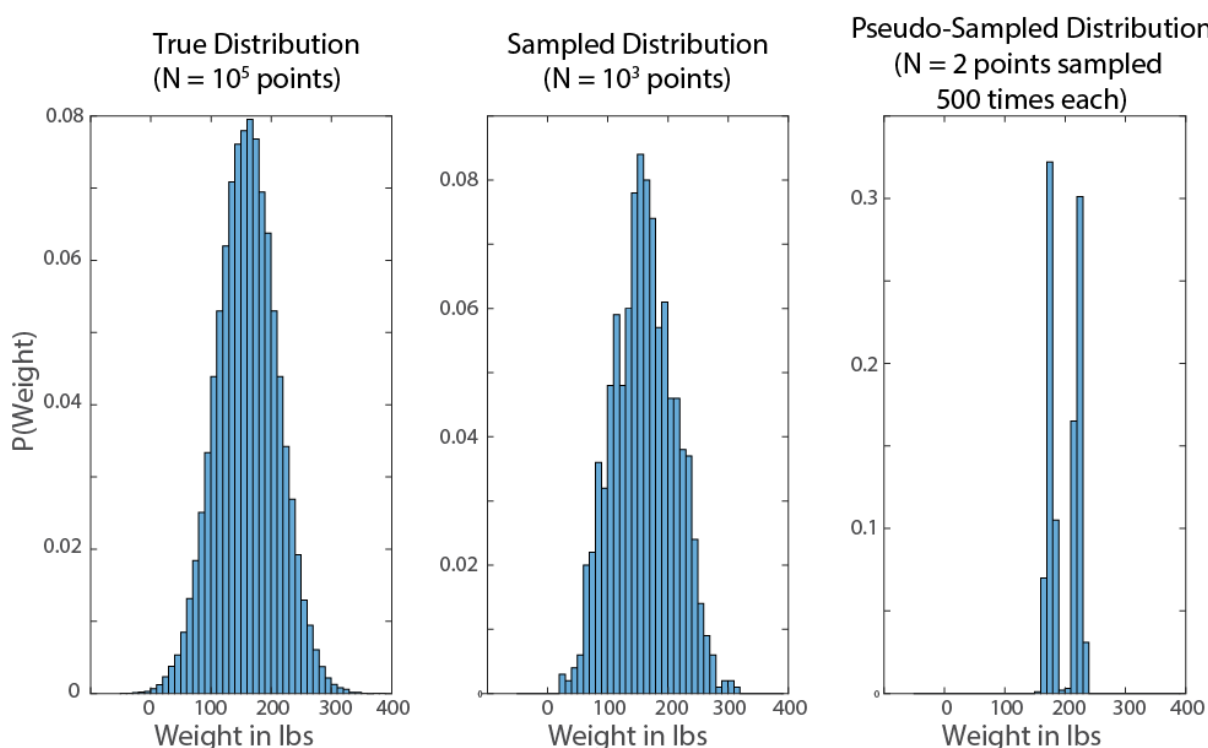


Figure 1.4: An extreme example of pseudoreplication.

Left panel: The true distribution of weight in a population. This sample was generated by assuming a normal distribution with mean 160 and standard deviation of 50 (the negative weights are because of the symmetric simulation). *Middle panel:* An independent samples draw from the distribution of 1000 points. As is evident from comparing to the true distribution, the sample is a close approximation to the true distribution. *Right panel:* An example of extreme pseudoreplication. Two samples were drawn from the original population at random and those samples were then resampled 500 times (simulated as drawing from a normal distribution with standard deviation of 5 around the true value as the mean) and combined to be presented as 1000 individual draws. As is evident from comparing to the left panel, it is not a good approximation of the true distribution and cannot be fixed simply by resampling the two points again.

propagation in hierarchical datasets. Statistical tests and analyses of hierarchical data have been a long-standing problem in several scientific domains, since one of the fundamental assumptions of such tests is violated, namely the assumption that all data points are independent. By the very nature of being hierarchical, data within a particular level are often not independent from each other. If one observes a typical songbird study, the study contains order of 10 birds, each bird sings order of 10 types of syllables and each syllable is repeated tens to thousands of times per day. The covariance one could expect between the pitch of one particular syllable of one particular bird at two different time points is very different from that expected if the syllables were different or if the birds themselves were different. Hence, propagating errors through the various levels and computing statistical tests correctly is of utmost importance. In this section, I will discuss how this problem was first identified as well as the early approaches made to address it. I will then discuss Linear Mixed Models (LMMs) and their applications in addressing these issues. I will conclude with a discussion of bootstrapping and how I used it to address error propagation in hierarchical data.

1.4.1 Problem Identification and Early Solutions

One of the earliest identifications of such a problem and the statistical inferences being made about the data from such experiments came from an article detailing pseudoreplication in ecological field experiments (Hurlbert 1984). Replication or repeated measures means that one draws independent samples from an unknown underlying probability distribution several times. If the samples are independent, then as the number of samples increases, the distribution of samples approaches that of the true underlying distribution. However, if the samples are not independent, then the previous statement about approaching the underlying distribution with increasing number of samples need not be true. An extreme example where the assumption is demonstrably false is shown in Figure

1.4 where the weight of two individuals is sampled repeatedly and therefore cannot approximate the true distribution of weight of individuals in the population regardless of the number of samples.

While the problem of pseudoreplication in ecological studies was not as extreme as described above, it could take on several forms and identifying these accurately and accounting for them accordingly was challenging. Hurlbert in his 1984 review classified pseudoreplication in ecological studies under 4 major categories that are briefly discussed below (Hurlbert 1984, Heffner, Butler, and Reilly 1996).

1. *Simple pseudoreplication*: This was the most common form of pseudoreplication observed at the time. This involved samples from a single experimental “unit” where a unit could be a single sampling location in a study comparing samples of fish across several water bodies, for example, being analyzed as if they had been replicated across multiple experimental units.
2. *Temporal pseudoreplication*: This is when a single experimental unit is sampled repeatedly through time and each individual sample is treated as an independent unit. The forms of pseudoreplication discussed above and some aspects of the data I analyze for my results fall under this category. However, as I will detail later, due to the nature of my analysis, I mostly do not address this category.
3. *Sacrificial pseudoreplication*: This happens when multiple samples under multiple experimental units are pooled together prior to analysis and treated as if they were all collected as part of a single experiment. This ignores the fact that there are two levels of variance in this dataset, namely, variance within a single experimental unit and variance between experimental units. This deals with the hierarchical nature of data which is what I will primarily address in my results in Chapter 3.

4. *Implicit pseudoreplication*: This is where studies present subsampled treatments within a single experimental unit which were not intended to be tested for differences originally as different groups and present graphs depicting non-overlap between the standard errors or confidence intervals. Additionally, statistical tests are typically not applied but “significance” is nevertheless discussed.

Since the identification of pseudoreplication, several methods have been proposed for addressing it (Stewart-Oaten, Murdoch, and Parker 1986, Millar and Anderson 2004) including resampling methods (Crowley 1992) which I will detail in section 1.4.3. Additionally, while pseudoreplication was first identified in ecological studies, it has been recognized in many other fields including neuroscience (Lazic 2010) and is part of the reason journals have been pushing for reporting exact p-values in addition to degrees of freedom and test statistic values. Furthermore, while there has been pushback against pseudoreplication (Oksanen 2001, Schank and Koehnle 2009), the fields at large have recognized and embraced it and are reducing the incidence of pseudoreplication in their studies (Heffner, Butler, and Reilly 1996, Kroodsma et al. 2001). It is still the case however that in neuroscience many studies that should be correcting for this are not doing so.

1.4.2 Linear Mixed Models (LMMs) as a Solution

Multilevel analysis (Snijders 2011, Hox, Moerbeek, and Van de Schoot 2017) emerged as a prime solution to the problem of pseudoreplication, particularly to problems in the categories of temporal and sacrificial pseudoreplication. Multilevel analysis acknowledges the fact that data, and therefore variance, exist at various levels that each must be treated separately. In order to do so, multilevel analysis typically takes the form of linear regressions where the contributions of different levels are regressed as “random” effects while the factor of interest is accounted for as the “fixed” effect (see (Roux 2002) for a full glossary of terms). The most popular form used in

biological sciences and in neuroscience in particular is the LMM (Aarts et al. 2015, Aarts et al. 2014).

While LMMs had been actively researched and used in various settings for decades before, it was only in the 1990s that computational hardware and software made large scale implementation of LMMs feasible for the vast majority of biological experiments. Even then, its popularity did not catch on much till the early 2000s when it was proposed as a potential fix for the pseudoreplication problem in studies involving fisheries (Millar and Anderson 2004). Since then, LMMs have been used regularly in biological sciences including neuroscience when dealing with time series data which would otherwise fall under temporal pseudoreplication (Wykes et al. 2012, Howe et al. 2013) and when dealing with hierarchical data collecting a large number of trials from a small number of subjects which would otherwise be a form of sacrificial pseudoreplication (Arlet et al. 2015, Pleil et al. 2016, Liang et al. 2015).

In spite of the numerous examples of studies that do use LMMs, they are still not widely used in a large number of areas where pseudoreplication may still be occurring, including the analysis of songbird vocal behavior. A potential reason for this is that in spite of the availability of software that can build LMMs relatively easily (*fitlme* in MATLAB; *mixedlm* from the *Statsmodels* library in Python), LMMs can be quite tricky to build accurately for testing a complex hypothesis from a hierarchical dataset and often times require the help of a statistician to accurately build and compute the appropriate LMM. However, more importantly, LMMs also assume that all hierarchical structure present in the dataset is linear in nature. In practice, this is often not the case and therefore the LMMs may not fit the data as well as one would like. A model that makes less assumptions about the structure of the data is therefore more desirable.

1.4.3 Bootstrapping and other Resampling methods as a solution

The bootstrap is a resampling method that has been popular since the 1980s (Efron 1981, 1992, Efron and Tibshirani 1994) and was in fact proposed as a viable fix to the problem of pseudoreplication as early as 1992 by Crowley's seminal review (Crowley 1992). The bootstrap as a statistical method has fallen in and out of popularity over the years since. It has been used in fields as wide ranging as meta-analysis of ecological studies (Adams, Gurevitch, and Rosenberg 1997), phylogenetic analyses (Garland Jr et al. 1993, Garland Jr, Midford, and Ives 1999) and analysis of genetic diversity among crops (Mohammadi and Prasanna 2003) for example. However, its relative simplicity and lack of an explicit hypothesis testing step have made it less preferable to alternatives such as the multi-factor ANOVA which claim to also perform hypothesis testing in addition to accurately dealing with hierarchical data (Anderson and Braak 2003, Anderson 2001). Additionally, the bootstrap was shown to be more conservative, i.e., produce larger error bars, than standard methods and is therefore less likely to result in a positive result (Adams, Gurevitch, and Rosenberg 1997). These factors may have contributed to the bootstrap not gaining as widespread use as LMMs in tackling pseudoreplication, particularly in the biological sciences.

However, the bootstrap did find widespread use in statistics. Following work by Efron and Tibshirani in popularizing the bootstrap, several algorithms were developed for the bootstrap to apply to various contexts. Some of them, such as residual bootstrap models resembled LMMs in many aspects (Davison and Hinkley 1997, Carpenter, Goldstein, and Rasbash 2003). Non-parametric and multi-layered sampling techniques were developed for dealing with hierarchical data and were shown to be more accurate in estimating uncertainty and confidence intervals than other commonly used methods (Field and Welsh 2007, Harden 2011, Thai et al. 2013, Huang

2018). Given its successful application to a wide variety of other fields, I argue in Chapter 3 that its use in neuroscience is most warranted.

Bootstrapping is a very simple and easy to understand, yet powerful resampling procedure. In essence, bootstrapping involves resampling with replacement among the given population and computing a metric of interest (as long as it obeys the central limit theorem) from that resampled data. This process is repeated N times so that at the end, one has N values for the metric of interest. If N is sufficiently large, the mean of these N values will be very close to the true value for the metric of interest in the original population and the 67% confidence interval gives an accurate estimate for the error in measurement of said metric (Efron and Tibshirani 1994). Building on previous work on hierarchical bootstrapping (Efron, Halloran, and Holmes 1996, Shimodaira 2004, 2002), I detail in Chapter 3 how I used bootstrapping to account for the error in songbird data accurately and show how a result published previously from our lab (Hoffmann and Sober 2014) did not have the statistical power required to state one of its conclusions if pseudoreplication is appropriately accounted for in addition to its utility for a second example in behavioral experiments with flies (Cande et al. 2018). I also use the same technique to quantify my own experimental results presented in Chapter 2.

1.5 Dissertation Overview

As I hope to have conveyed through this introduction, it is still an open question as to whether dopamine plays a role in sensorimotor adaptation and if it does, how it does so. Furthermore, neuroscience is an area of study that is highly susceptible to pseudoreplication due to the hierarchical nature of the data analyzed. Therefore, the studies in this dissertation are arranged into the following chapters.

1. Chapter 1 provides an introduction and review of the literature identifying specific gaps in the scientific knowledge that this dissertation hopes to address.
2. Chapter 2 details results from an experiment designed to test the role of dopamine in sensorimotor adaptation in songbirds. I found that dopamine depletion produces two not mutually exclusive effects in songbirds. First, there is an effect on song production in that the average pitch of the song reduces over time following dopamine depletion. Second, there is a clear deficit in sensorimotor adaptation since songbirds are unable to respond adaptively to an induced shift in pitch through the headphones following dopamine depletion. A full discussion of the interpretation of these results is detailed in section 2.5.
3. Chapter 3 details the use of hierarchical bootstrapping to solve the problem of pseudoreplication in songbird neuroscience due to the hierarchical nature of the data. I also detail false positive rates of various previous approaches in a simulation example and provide a couple of examples of the utility of the hierarchical bootstrap in analyzing hierarchical datasets in neuroscience.
4. Chapter 4 provides an overarching discussion for the entire dissertation and concludes with some future lines of research that have been opened as a result of this dissertation.
5. A comprehensive list of references cited in this dissertation is provided at the end.

2 CHAPTER II: DOPAMINE DEPLETION AFFECTS VOCAL ACOUSTICS AND DISRUPTS SENSORIMOTOR ADAPTATION IN SONGBIRDS¹

2.1 Abstract

Dopamine is hypothesized to convey error information in reinforcement learning tasks with explicit appetitive or aversive cues. However, during motor skill learning feedback signals arise from an animal's evaluation of sensory feedback resulting from its own behavior, rather than any external reward or punishment. It has previously been shown that intact dopaminergic signaling from the ventral tegmental area – substantia nigra compacta complex (VTA/SNc) is necessary for vocal learning when songbirds modify their vocalizations to avoid hearing distorted auditory feedback (playbacks of white noise). However, it remains unclear whether dopaminergic signaling underlies vocal learning in response to more naturalistic errors (pitch-shifted feedback delivered via headphones). We used male Bengalese finches (*Lonchura striata* var. *domestica*) to test the hypothesis that the necessity of dopamine signaling is shared between the two types of learning. We combined 6-hydroxydopamine (6-OHDA) lesions of dopaminergic terminals within Area X, a basal ganglia nucleus critical for song learning, with a headphones learning paradigm that shifted the pitch of auditory feedback and compared their learning to that of unlesioned controls. We found that 6-OHDA lesions affected song behavior in two ways. First, over a period of days lesioned birds systematically lowered their pitch regardless of the presence or absence of auditory errors. Second, 6-OHDA lesioned birds also displayed severe deficits in sensorimotor learning in response to pitch-shifted feedback. Our results suggest roles for dopamine in both motor production and

¹ A version of this Chapter has been published in eNeuro (Saravanan et al. 2019b).

auditory error processing, and a shared mechanism underlying vocal learning in response to both distorted and pitch-shifted auditory feedback.

2.2 Introduction

Complex organisms perform sensorimotor learning to modulate behavior in response to sensory feedback. This process uses feedback from past performances arising from either explicit reward/punishment cues (e.g. food reward, electric shocks) or from self-evaluation of the performance (e.g. hearing one's own voice during speech or song). While prior work has taken a number of approaches to taxonomizing different forms of sensorimotor learning, including distinguishing model-based and model-free learning (Mohan, Morasso, and Metta 2011, Wolpert, Ghahramani, and Jordan 1995, Haith and Krakauer 2013) and habitual versus goal-directed behavior (Balleine and O'doherty 2010, Redgrave et al. 2010), here we focus on an orthogonal distinction into two broad components: error-based learning that relies on self-evaluation and reinforcement learning that relies on cues from the environment (Wolpert, Diedrichsen, and Flanagan 2011). Classic studies have linked dopamine to reinforcement learning as a reward prediction error signal that conveys information about explicit rewards and punishments (Schultz, Dayan, and Montague 1997, Glimcher 2011). However, the question of whether dopamine is also involved in error-based learning in the absence of external rewarding or aversive cues has been harder to address. Some studies have reported deficits in error-based learning in patients with Parkinson's disease (Paquet et al. 2008, Mollaei, Shiller, and Gracco 2013), but since Parkinson's disease is associated with cognitive and executive deficits in addition to larger motor deficits (Jankovic 2008, Dubois and Pillon 1996, Lees and Smith 1983, Cooper et al. 1991), the specific role of dopamine has been difficult to isolate.

Songbirds have emerged as an effective model system in which to study the role of dopamine in sensorimotor learning. Songbirds spontaneously produce songs hundreds of times per day. Like human speech, song is learned during development (Lipkind et al. 2013, Wilbrecht and Nottebohm 2003) and actively maintained by auditory feedback through adulthood (Sakata and Brainard 2006, 2008, Sober and Brainard 2009, Kuebrich and Sober 2015). Additionally, songbirds have a well-defined neural circuitry dedicated to song production and song learning (Sohrabji, Nordeen, and Nordeen 1990, Brainard and Doupe 2000, Scharff and Nottebohm 1991). Dopaminergic neurons from the ventral tegmental area/substantia nigra pars compacta (VTA/SNc) complex innervate Area X, a basal ganglia nucleus essential for song learning, and have been hypothesized as a way for auditory error information to enter the song system (Mandelblat-Cerf et al. 2014, Peh, Roberts, and Mooney 2015, Bottjer 1993, Soha, Shimizu, and Doupe 1996) (see Fig. 2.1). Researchers examining vocal control employ two primary methods to induce song learning in adult songbirds: through distorted auditory feedback (Tumer and Brainard 2007) and through pitch shifts played through custom-made headphones (Sober and Brainard 2009). It remains unclear to what extent the two paradigms share underlying neural mechanisms. Dopamine has been shown to be involved in changing the pitch of the song in response to distorted auditory feedback. Specifically, birds display deficits in learning to avoid distorted feedback under dopamine depleted conditions (Hoffmann et al. 2016, Hisey, Kearney, and Mooney 2018). Neural recordings of dopaminergic neurons revealed prediction error type responses when birds were required to avoid such auditory distortions while singing (Gadagkar et al. 2016), and pitch-contingent optical stimulation of dopaminergic terminals in Area X evoked changes in the pitch of the birds' song (Xiao et al. 2018, Hisey, Kearney, and Mooney 2018). Here, we tested the hypothesis that there are common neural mechanisms underlying both learning paradigms by

studying the role of dopamine in birds when they respond to a pitch shifted version of their own auditory feedback (Sober and Brainard 2009).

We tested the role of dopamine in error-based learning by selectively lesioning dopaminergic terminals in Area X using 6-hydroxydopamine (6-OHDA). Since the cell bodies of dopaminergic neurons in VTA/SNc that innervate Area X are intermingled with those projecting to the rest of the songbird basal ganglia (Person et al. 2008), we injected 6-OHDA directly into Area X to avoid introducing general motor or song production deficits. We fitted the birds with custom-built headphones through which we introduced a shift in pitch (either upwards or downwards) of the bird's auditory feedback (Hoffmann et al. 2012, Sober and Brainard 2009) to measure how birds changed their pitch over time in response to this induced sensory error and how self-guided error correction was affected by dopamine manipulations.

2.3 Materials and Methods

All 16 animals used for this study were adult (range of ages: 105 to 217 days post hatch; median age: 141 days post hatch) male Bengalese finches (*Lonchura striata* var. *domestica*). Throughout the study, the animals were housed in isolated sound attenuating chambers (referred to as sound boxes) on a 14-10 hour light-dark cycle. All singing analyzed for this paper was undirected song, i.e., songs sung in the absence of a female. All procedures were approved by Emory University's Institutional Animal Care and Use Committee.

2.3.1 Experimental design and Statistical Analysis

Songbirds display significant bird-by-bird variability in the amount of learning displayed, and so most experimental designs include a within-bird control to measure the amount of learning within a bird before and after a manipulation of interest (Hoffmann et al. 2016, Hisey, Kearney, and Mooney 2018). However, in the case of headphones as we use here (described in Headphones

Attachment and Assembly below), the only way to secure the headphones to the birds for the duration of the experiment is to cement them to the skull. Although this method ensures that the headphones fit comfortably around the ear canals and remain in place for the duration of the experiment, cementing the headphones to the skull prevents access to the brain, thereby preventing us from examining learning in the same animals pre- and post-lesion. As a result, we designed a group comparison study to test the role of dopamine in sensorimotor adaptation. We performed pitch shift experiments on 6 unlesioned birds (3 each for upward shifts and downward shifts) and 8 lesioned birds (4 for upward pitch shift and 4 for downward pitch shift). As detailed below, virtual auditory feedback through the headphones was delivered almost in real time and was meant to replace the natural auditory feedback that birds would otherwise receive. All pitch shifts were 1 semitone in magnitude (equally split between +1 and -1 semitone shifts). Each experiment consisted of 3 days of baseline (unshifted auditory feedback through headphones) followed by 14 days of pitch shifted auditory feedback. At the end of the shift period, we turned off the shift in pitch (i.e. set the pitch shift to zero semitones as in the baseline epoch) and recorded the birds' activity for 6 to 7 days. During this period, unlesioned birds typically reverse the effects of the pitch shift (Sober and Brainard 2009). We refer to this period as "washout." Washout data were collected for all 6 unlesioned birds. Due to technical difficulties associated with keeping the headphones attached for extended periods of time, washout data was collected for only 4 out of the 8 lesioned birds (2 for upward pitch shifts and 2 for downward). In addition, we performed control experiments with 2 unlesioned birds fitted with headphones and no pitch shift and 8 lesioned birds without any pitch shifts (5 with headphones and zero pitch shift throughout; 3 with no headphones). To minimize the number of animals we used, our unlesioned bird group consisted of data reanalyzed from Sober and Brainard, 2009. All data that have not been labeled explicitly

as “Data reanalyzed from a previous study” are new data collected for the purpose of this study. Furthermore, since we showed previously (Hoffmann et al. 2016) that animals injected with saline instead of 6-OHDA were statistically indistinguishable from unlesioned birds, we did not include a saline injected control group in this study. Note that of the 8 birds whose data were reanalyzed from Sober and Brainard (2009), the raw data for 2 animals – the unlesioned birds with no headphones shift – were unavailable. However, we were able to extract the daily mean pitch values from each animal’s data from an eps version of the original figure summarizing the data. The resulting figure that shows the mean change in pitch and error bars for the group was produced from the 2 data points for each day.

For our lesioned group, we reduced the dopaminergic innervation of Area X (Fig. 2.1), a song specific nucleus of the basal ganglia, using 6-hydroxydopamine (6-OHDA) microinjections as described in detail previously (Hoffmann et al. 2016). Briefly, we used stereotactic surgeries to target Area X with a 4 x 3 grid of microinjections of 6-OHDA (see 6-OHDA lesions below). Following 6-OHDA surgery, the birds were allowed to recover in their sound boxes for 4 to 5 days which also served as a period to allow the 6-OHDA to cause degeneration of striatal innervation (Jeon, Jackson-Lewis, and Burke 1995). Subsequently, the headphones (Hoffmann et al. 2012) were fitted to the birds and set to initially provide unshifted auditory feedback (zero pitch shift). Following headphones attachment, the birds typically did not sing for 2 to 4 days (see Fig. 2.1c for a timeline schematic). Once they started singing again (defined as at least 30 song bouts produced over the entire day), we began recording a 3 day baseline period. Following the 3 days of baseline, the birds were recorded for 14 days during a period of shift. As described previously (Sober and Brainard 2009, Kelly and Sober 2014), the pitch shift was a 1 semitone shift (either upwards or downwards) played back to the bird through the headphones. The auditory feedback

through the headphones was almost real-time (delay of around 10 ms) and was intended to replace the bird's natural auditory feedback. In order to do so, the volume is set to be at least 2 log units greater in sound intensity than the bird's own feedback. For the birds that had no pitch shift through the headphones, they continued with zero shift as they were in baseline for the equivalent 14 days. Following this 14 day period, we recorded the birds' activity for 6 to 7 days of washout. Owing to the difficulties of keeping the headphones attached and functional for long periods of time, we were not able to collect washout data for every animal. Analysis of washout data was therefore necessarily limited to data from birds that did have data collected for the washout period.

Note that one of our 6-OHDA lesioned birds in the -1 semitone shift group was subjected to an extended baseline period of 6 days rather than the 3-day period used for all other animals. Excluding data from this bird did not change any of our results significantly. Therefore all results reported include this bird, treating the last three days of baseline equivalent to days 1 through 3 of baseline for every other bird.

Birds with lesions that were not fitted with headphones were returned to their sound boxes post-surgery and were recorded for the duration of the experiment. In this case, since they did not have a break in singing due to placement of fully assembled headphones, the baseline was defined as days 6 through 8 post lesion and the "shift" period was defined as day 9 through 22 post lesion to keep the timelines comparable between groups.

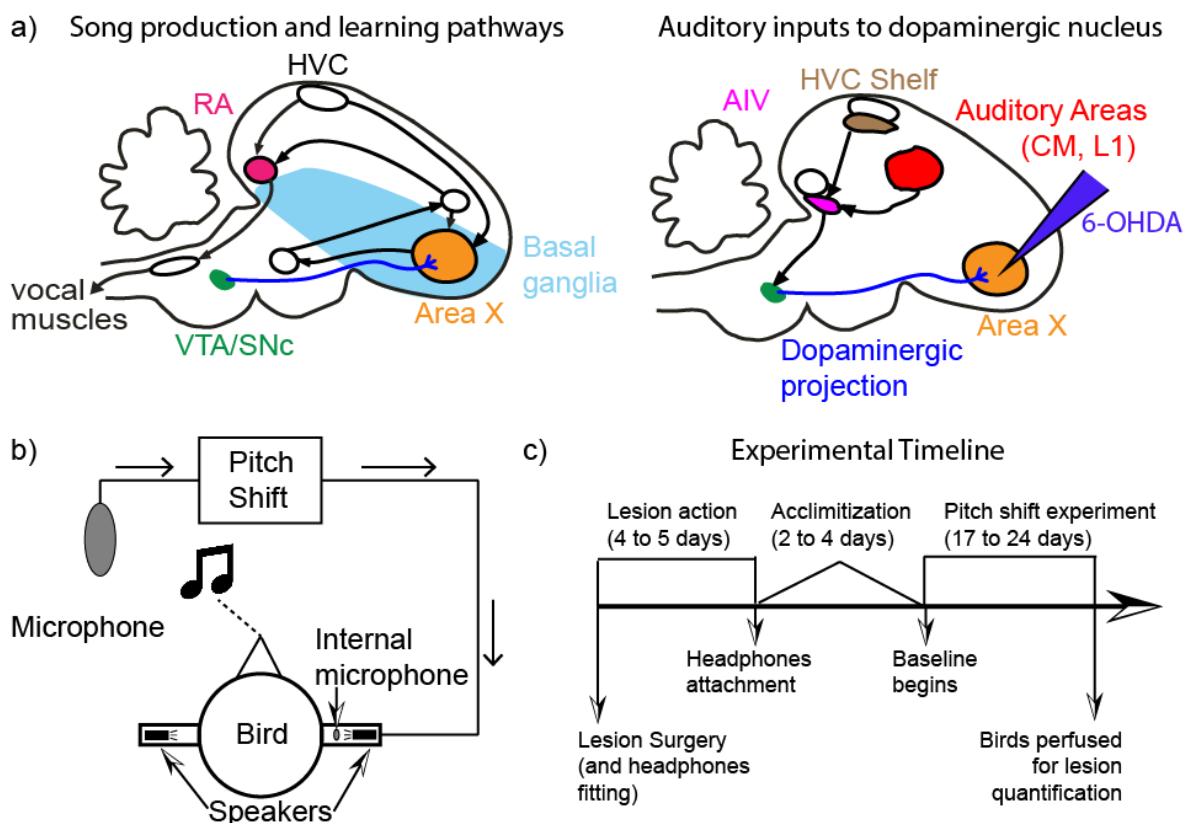


Figure 2.1: Songbird neuroanatomy and experimental design.

a) A theory for the role of dopamine in sensorimotor learning in songbirds. The left panel shows the brain nuclei in the songbird primarily involved in song production and learning. Area X, a songbird basal ganglia nucleus critical for song learning, receives dense dopaminergic projections from the VTA/SNc complex. The right panel shows the nuclei involved in auditory processing in the songbird. There are other inputs (not shown) to the VTA/SNc complex from auditory areas and the ventral basal ganglia (vBG). One of the known pathways for auditory information to influence song learning is through the dopaminergic projections to Area X. We target these projections when we perform 6 hydroxydopamine (6-OHDA) lesions into Area X as depicted. b) A schematic for how the custom-built headphones introduce a pitch shifted auditory error to the birds. Briefly, a cage microphone records all sounds made within the cage and sends it through a pitch shifting program which is subsequently played back to the bird through miniature speakers attached to the headphones. The headphones also have an internal microphone to record output from the headphones speakers and to calibrate sound intensity. c) A detailed timeline for each of our experiments (see Materials and Methods).

2.3.2 6-OHDA Lesions:

We performed the lesions using stereotactic surgeries as described in detail previously (Hoffmann et al. 2016). Briefly, birds were anesthetized using ketamine and midazolam and positioned at a beak angle of 20 degrees below horizontal. Isoflurane was used to sustain anesthesia following the first hour of surgery. All stereotactic coordinates were relative to the landmark Y_0 , the posterior border to the divergence of the central sinus in songbirds. Small craniotomies were performed above the coordinates AP 4.75 to 6.4; ML 0.75 to 2.3 on both sides (all coordinates are in mm). 6-OHDA (Tocris; conjugated with HBr) was injected bilaterally in a 4 x 3 grid at AP coordinates 5.1, 5.5, 5.9 and 6.3 and ML coordinates 0.9, 1.55 and 2.2 with a DV coordinate between 3.08 and 3.18 from the surface of the brain. For each injection, the glass pipette was lowered into the brain slowly allowing for time for rebounding of tissue, and following the injection, the electrode was left in place for at least 30 seconds before withdrawal at a similarly slow pace. Additionally, we initially performed one final injection at AP 4.8, ML 0.8 and DV 2.6 from the surface of the brain targeting the tail portion of Area X but dropped this injection in later birds as the targeting was not reliable and the injection required a larger craniotomy to perform. 13.8 nL of 6-OHDA was injected in the slow setting (23 nL/s) at each injection site using a Drummond Scientific (Broomall, PA) Nanoject II auto-nanoliter injector.

2.3.3 Headphones attachment and assembly:

The methodology is described in detail in (Hoffmann et al. 2012). Briefly, each set of headphones was custom-fit to an individual bird under anesthesia. If attached on a bird that also had a 6-OHDA lesion, both lesion and headphones fit adjustment were performed back-to-back in the same surgery. Once the headphones had been successfully fitted for the bird, the electronics (a speaker on each side and a miniature microphone on one side to record headphones output and calibrate

volume) were assembled offline. The fully assembled headphones were then refitted to the bird 4-5 days post-surgery. We used a flexible tether with a commutator to power the headphones and read the electronic signals.

2.3.4 Histology:

Following the end of the experiment, headphones were removed and the birds were deeply anesthetized with ketamine and midazolam before performing perfusions using 10% formalin. The brains were postfixed overnight in formalin and then cryoprotected in 30% sucrose for 1 to 4 days prior to slicing into 40 μm sections on a freezing sliding microtome. Alternating sections were either immunoreacted with tyrosine hydroxylase antibody and visualized with diaminobenzidine (TH-DAB) or Nissl stained. TH-DAB was used to quantify the extent of lesions in the 6-OHDA birds, while Nissl was used to verify that there had been no necrosis and to assist in identifying boundaries of Area X in adjacent TH-DAB sections. For the TH-DAB reaction, all incubations were carried out on a shaker at room temperature and all chemicals were dissolved in 0.1M phosphate buffer (PB) unless otherwise noted. Fixed sections were treated sequentially with 0.3% hydrogen peroxide to suppress endogenous peroxidases and 1% sodium borohydride to reduce exposed aldehydes and improve background staining before incubating overnight in a tyrosine hydroxylase antibody solution (Millipore Cat# MAB318, RRID:AB_2201528, 1:4000; 0.3% Triton X-100; and 5% normal horse serum). Tissue was then incubated in biotinylated anti-mouse secondary antibody (Vector Laboratories Cat# BA-2000, RRID:AB_2313581, 1:200 and 0.3% Triton X-100) followed by avidin-biotin-complex (ABC) solution (Vector Laboratories Cat# PK-4000, RRID:AB_2336818). Tissue was exposed to DAB solution (Amresco E733; 5 mg DAB per tablet; 2 tablets in 20 ml of purified water) for approximately 5 min. Sections were mounted, air-dried, delipidized with ethanol and citrisolv, and coverslipped with Permount (Fisher scientific,

SP15-500). For the Nissl stained sections, Nissl stain was applied on mounted, air-dried tissue, which was delipidized with ethanol and citrisolv, and coverslipped with Permount. Stained sections were imaged using a slide scanner (Meyer Instruments PathScan Enabler IV; 24 bit color, 7200 dpi, “sharpen more” filter, brightness, and contrast level 50) and the resulting images were analyzed using ImageJ (RRID:SCR_003070).

2.3.5 Image and Lesion Analysis:

TH-DAB stained sections were used for lesion quantification by analysis through a custom written macro in ImageJ. The analysis was based on a metric of optical density described in detail in (Hoffmann et al. 2016). Briefly, the macro allowed us to demarcate the boundary of Area X in every section that it is present. We also used a circle of diameter 0.5 mm to mark a section of representative striatum outside of Area X in the same section. We then defined the optical density ratio (OD ratio) as the ratio between the optical density of Area X in the section to that of striatum in the section as follows:

$$OD\ ratio = \frac{OD_{Area\ X}}{OD_{striatum}}$$

One of the established ways of identifying Area X in songbirds has been that Area X is darker than the surrounding striatum when stained with TH-DAB (Soha, Shimizu, and Doupe 1996, Hoffmann et al. 2016, Bottjer 1993). Due to this property, we used the cumulative distribution of the optical density ratio in saline injected birds to define our threshold for lesions. Any section in our group of 6-OHDA lesioned birds with an OD ratio less than the 5th percentile of the saline injected birds sections counted towards the overall proportion of lesioned sections. Additionally, we used a two-sample Kolmogorov-Smirnov test to test whether the lesioned and saline populations were indeed drawn from separate distributions. We also used the threshold procedure described above to quantify lesion extent for individual animals. We then asked whether lesion extent was

significantly correlated with vocal behavior metrics such as baseline variance, change in variance from baseline to end of shift and change in pitch at the end of shift.

2.3.6 Pitch Quantification:

All our analysis was performed using an extracted value of pitch for every instance in which a bird sings a particular syllable. Briefly, birds have multiple syllables within their song and they typically repeat their song hundreds of times per day during the course of the experiment. We call each time they sing a particular syllable an iteration of that syllable. We restricted our analysis to roughly 30 song files per day between 10 am to 12 pm and have shown earlier that the choice of time window does not qualitatively affect our results (Hoffmann and Sober 2014, Kelly and Sober 2014, Sober and Brainard 2009). To quantify pitch, for each syllable we specify a time during the syllable (relative to syllable onset) during which the syllable is relatively flat and clear in the frequency vs time space and can be reliably quantified across iterations across days. The pitch we extract represents a weighted average of the frequencies with the highest power in the lowest harmonic of the syllable. In order to make comparisons between different syllables whose base frequency can vary widely, we convert the pitches into semitones as shown below:

$$s = 12 * \log_2 \left(\frac{pitch}{baseline} \right)$$

where s is the change in pitch in semitones, $pitch$ is the observed pitch and $baseline$ is the average pitch across the 3 days of baseline for that particular syllable. For all group analysis, the means reported are the means over all birds and over all syllables weighted by the proportion of times they sang each syllable. This was chosen to account for the fact that syllables that are sung more often are exposed a greater number of times to the shifted auditory feedback. Pitch quantification was performed using custom-written scripts in MATLAB (RRID:SCR_001622).

2.3.7 Error quantification:

For each of our groups, we had between 4 to 8 birds, each bird performed between 4 to 12 different syllables whose pitch could be quantified, and each syllable was repeated between 40 to 600 times per day. As a result, while we have several thousands of data points towards establishing the position of the mean pitch change per group for each day, the structure of the data is hierarchical and error accumulates at different levels (birds, syllables and iterations). Grouping all the data together and estimating the standard error of the mean underestimates the error by ignoring the non-independence between data points due to the hierarchical structure. On the other extreme, aggregating points and simply using individual birds or syllables does not allow us to use all of our data effectively. This is a complex problem that different studies, including our own prior efforts have used varying methods to address (Tian and Brainard 2017, Aarts et al. 2014, Galbraith, Daniel, and Vissel 2010, Sober and Brainard 2012). To more accurately quantify the error in our groups and better account for the variance arising from finite data samples, we use a hierarchical bootstrapping approach (Efron and Tibshirani 1994, Crowley 1992). In its simplest form, bootstrapping involves generating N ($N = 10^4$ throughout this paper) random subsamples of the dataset by sampling with replacement from the original data and computing a metric of interest for each subsample. This results in having a distribution of the metric of interest, the 67% confidence interval of which provides an accurate estimate of the uncertainty in measurement of that metric in the original dataset (Efron 1981, 1992, Efron and Tibshirani 1994). For example, if one wanted to obtain the uncertainty in measuring the kurtosis of the data, one would generate bootstrap subsamples and calculate the kurtosis for each subsample. The standard deviation of the population of kurtosis values so obtained gives an accurate estimate of the uncertainty of the kurtosis in the original data. In the special case of estimating a population of means (which is the metric of interest

in all instances in this paper), the uncertainty in measurement referred to above corresponds to the standard error of the mean of the dataset. However, bootstrapping by itself does not solve the problem of non-independence in hierarchical data. Crucially, to address this issue the resampling described above has to be done separately over each level of the hierarchy. This means that to generate a single subsample, we first resampled among the birds, then for each selected bird, we resampled among its syllables and finally for each syllable, we resampled among its iterations. Finally, we acknowledged that Bengalese finches can vary greatly in their syllable repertoires from one bird to the next. While all birds typically have an order of 10 syllables, some birds repeat one or two syllables with a much higher frequency than any other syllable while others represent each syllable equally. Since the bootstrapping procedure was used to calculate uncertainty of measurement due to sampling from a limited number of birds, we posited that each syllable would be equally likely in hypothetical new birds. Therefore, we set the number of iterations of a particular syllable that could occur in a bootstrapped subsample to be independent of the frequency of occurrence of that syllable in the actual data. All the data for the subsample were then combined and their mean was calculated for the subsample. Note that this procedure only applies to our estimate of measurement uncertainty (not the mean pitch values), since the means reported in the results are calculated from the actual data collected. This process was then repeated N times. In order to also account for the error in estimation of the mean of each syllable during baseline, the resampling was performed on pitch measurements recorded in hertz (Hz) and the measurements were converted to semitones just prior to calculating the mean pitch for each subsample. A similar procedure was followed for quantifying error during washout. To account for the error in estimation of pitch on the last day of pitch shift, the subtraction of the mean pitch on the final day of shift through the washout period was performed following the resampling. Our error

quantification was performed using custom written scripts in MATLAB (available at https://github.com/soberlab/Dopamine_Headphones_Paper_code).

2.3.8 Hypothesis testing with Bootstrap:

In addition to using bootstrapping to compute error estimates as described above, we also used a bootstrapping approach to test whether vocal pitches were significantly different across time or experimental conditions by computing direct posterior probabilities for individual hypotheses. Hence, we report our results in terms of direct probabilities of a sample being greater than or equal to another sample or fixed value in lieu of p-values. Specifically, we resample the distribution for each group and calculate the mean 10^4 times to produce a distribution of resampled means to calculate the variance associated with having a finite number of samples.

These resampled distributions were used to compute whether the two distributions of vocal pitches were significantly different. For all instances in this paper, we use two-way tests with $\alpha = 0.05$. This means that a probability is significant if the probability supporting the hypothesis, $p < \alpha/2$ or if $p > (1 - \alpha/2)$, i.e., if $p < 0.025$ or if $p > 0.975$. In the case of computing the probability of the mean of a group being different from a constant, one can calculate the proportion of the population of bootstrapped means (as defined in Error quantification above) being greater than or equal to said constant. For example to compute the probability that the mean shift in pitch of a particular group is significantly different from zero, one would compute the proportion of the population of bootstrapped means that are greater than or equal to zero. If this proportion is less than 0.025 then the pitch of the group of interest is significantly below zero while if the proportion is greater than 0.975 then the pitch of the group is significantly above zero.

We used a similar approach to compute significant differences between two groups of interest. In this case, we compute a population of bootstrapped means for each group. From these

two bootstrapped populations, we compute a joint probability distribution between the bootstrapped means of the two groups. The null hypothesis representing no difference between the two groups would correspond to a circle centered about the unity line. Therefore, to test the difference between the two groups, we compute the volume of the joint probability distribution on one side of the unity line (including the unity line itself) to quantify the probability of one group being greater than or equal to the other group. If the probability computed is greater than 0.975, then the first group is statistically greater than the second group. Alternatively, if the probability computed is less than 0.025, then the first group is statistically less than the second group. We computed multiple comparisons between groups by computing differences between 2 groups at a time and applied a Bonferroni correction to the threshold for significance. Our statistical tests were performed using custom scripts written in MATLAB which have been made available at https://github.com/soberlab/Dopamine_Headphones_Paper_code.

2.3.9 Validating our Results with Linear Mixed Models:

To ensure that our results were robust to our choice of error quantification and design, we also separately reported frequentist statistical tests on our results. Since our data are hierarchical (see Error quantification above), the recommended way to perform frequentist statistics on our data is through linear mixed models (Aarts et al. 2014, Aarts et al. 2015). Specifically, we built linear mixed models by using bird identity and syllable identity within a bird as variable effects and tested for significance of fixed effect factors. Concretely, our linear mixed models were of the form:

$$Pitch_{ijk} = \beta_{0jk} + \beta_1 * x_{ij} + \varepsilon_{ij}$$

$$\beta_{0jk} = \beta_{00k} + b_{0jk}$$

$$\beta_{00k} = \beta_{000} + c_{00k}$$

where x_{ij} refers to the condition of the shift (± 1 semitone or 0 semitone) and is the fixed effect while b_{0jk} accounts for the bird identity and c_{00k} accounts for syllable identities within a bird which are both variable effects. The code for hypothesis testing using LMMs was also done in MATLAB and are available at https://github.com/soberlab/Dopamine_Headphones_Paper_code.

2.4 Results

We performed pitch shift experiments on 6 unlesioned birds (3 each for upward shifts and downward shifts) and 8 lesioned birds (4 for upward pitch shift and 4 for downward pitch shift). Following the end of the pitch shift, we also collected data during the “washout” period, i.e., when the pitch shift is set back to zero and the bird typically reverts its pitch back to baseline. All 6 unlesioned birds had washout data collected for 6 days following the end of shift. Of the 8 6-OHDA lesioned birds, 4 had data for washout for 7 days each (we were unable to record washout data for the other 4 lesioned animals due to technical problems associated with long-term use of the headphones). In addition, we performed control experiments with 2 unlesioned birds fitted with headphones who heard unshifted (zero pitch shift) auditory feedback and 8 birds who received 6-OHDA lesions but did not undergo any pitch shifts (see Materials and Methods for complete details).

2.4.1 6-OHDA lesions reduce dopaminergic innervation of Area X:

We quantified the lesion extent using a metric developed as part of our prior work (Hoffmann et al, 2016). Specifically, we used sections of Area X stained with diaminobenzidine (DAB), a chromogen that conjugates to antibodies specific for tyrosine hydroxylase (TH), the rate limiting enzyme involved in catecholamine synthesis and a reliable marker for dopaminergic and noradrenergic innervation (Figure 2.2). TH-DAB does not follow the Beer-Lambert law and varies in stain intensity even within the same animal (Van Eycke et al. 2017). As a result, quantification

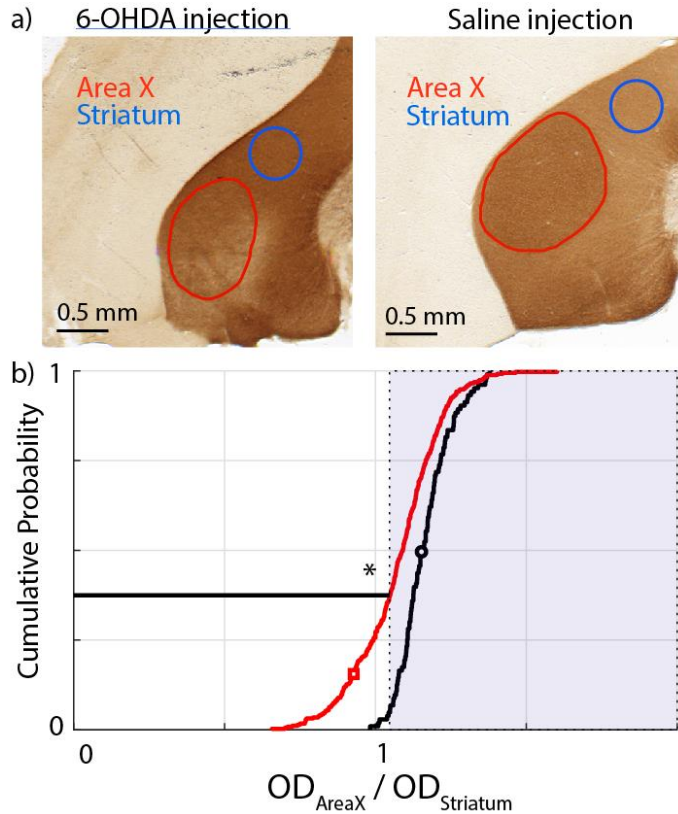


Figure 2.2: Metric for quantifying the extent of our lesions in our population of birds.

We used an optical density ratio (OD ratio) between Area X and the surrounding basal ganglia (see Materials and Methods) and compared the cumulative ratios between a saline injected population (N = 4 birds) and our 6-OHDA lesioned population (N = 16 birds). a) Examples of 6-OHDA lesioned (left) and saline injected (right) sections. The red trace demarcates the Area X boundary. The blue circle is chosen to represent a uniformly stained section of the rest of the striatum. The ratio for each section is calculated as the OD ratio between these two regions. b) Cumulative distribution plots for the saline injected birds (black trace) and the 6-OHDA lesioned birds (red trace). The shaded portion represents ratios that are greater than the 5th percentile for the saline injected birds. By this metric, 37.5% of all 6-OHDA lesioned sections have a smaller OD ratio. The black and red symbols correspond to the examples shown in a). The * represents a statistically significant difference between the red trace and the black trace (KS test; $p < 0.05$; see Results for full description).

is typically performed between hemispheres within one section comparing a lesioned to an unlesioned hemisphere. However, we had to perform bilateral lesions for our experiments since song learning is not known to be lateralized in Bengalese finches. To quantify lesion extent, we used the fact that Area X has denser dopaminergic innervation and thus stains darker by TH-DAB than the surrounding striatum (Soha, Shimizu, and Doupe 1996, Bottjer 1993). Specifically, we quantified an optical density ratio (OD ratio) for a batch of birds that had been injected with saline into Area X (N = 4 birds; data reanalyzed from Hoffmann et al, 2016) and produced a cumulative distribution plot of the ratio across all sections for these birds. We then defined the 5th percentile of that distribution as the threshold for defining lesioned sections (see Materials and Methods). When we produced a similar cumulative distribution plot of the OD ratio for all 16 of our 6-OHDA lesioned birds, around 37.5% of all sections were below the threshold defined above (Fig. 2.2b). This was somewhat smaller than the lesion extent for the cohort of birds in (Hoffmann et al. 2016) in which 50% of lesioned sections were below the threshold. However, the lesions were qualitatively similar between the two groups. In addition, the population of OD ratios for the 6-OHDA lesioned birds was consistently below that for the saline injected birds as verified by a two-sample Kolmogorov-Smirnov test ($K = 0.3467$; $p = 5.75 \times 10^{-9}$). We have also previously shown through High Performance Liquid Chromatography (HPLC) analysis that such 6-OHDA lesions have no discernible effect on the existing low levels of noradrenergic innervation of Area X (Hoffmann et al. 2016).

2.4.2 6-OHDA lesioned birds reduce pitch even in the absence of auditory error:

We showed earlier that in unlesioned animals, the headphones do not cause changes in vocal pitch in the absence of any shifts in feedback pitch (Sober and Brainard 2009). As shown in Figure 2.3a, the mean pitch across days 12 through 14 of the experiment for these birds was found to be $0.02 \pm$

0.07 semitones (all measures of mean pitch reported are mean \pm SEM). Since this particular dataset

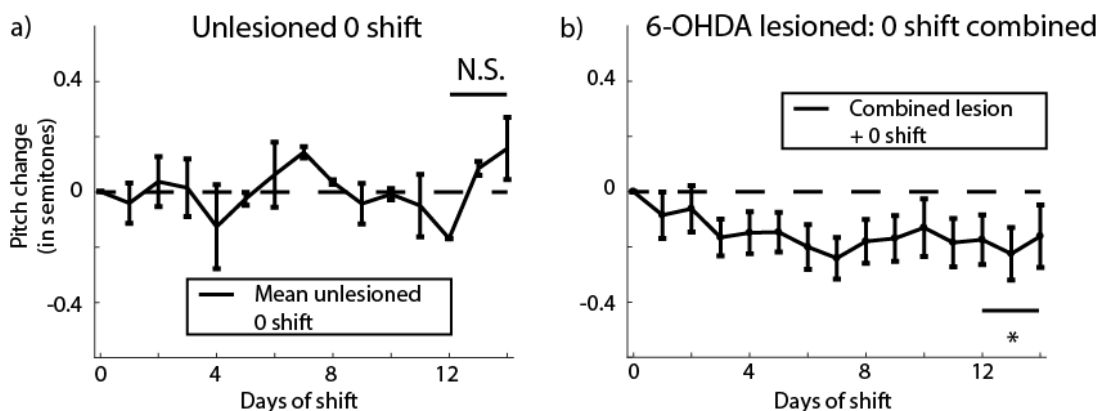


Figure 2.3: Quantifying the effect of headphones without any pitch shifts on the average change in pitch of the bird with or without lesions.

a) Mean change in pitch of song for 2 unlesioned birds with headphones but no shifts through the headphones (analyzed from data extracted from Supp. Fig. 6 from Sober and Brainard, 2009). b) Mean change in pitch for 6-OHDA lesioned birds combining both birds with headphones but no shift in pitch ($N = 5$ birds) or without headphones ($N = 3$ birds) for a total of 8 birds. The group averages for the two groups and the individual traces for all 8 birds is shown in Figure 2.4b. N.S. represents “not significantly different from zero” while the * represents a significant difference when comparing the last 3 days of shift combined from zero ($p < 0.05$).

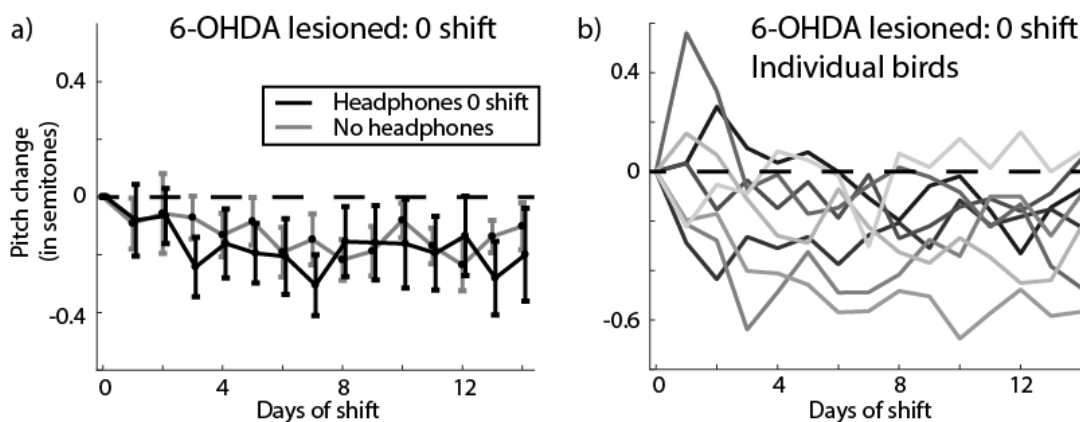


Figure 2.4: Individual traces for unshifted birds.

a) Mean change in pitch for 6-OHDA lesioned birds either with headphones but no shift in pitch (black trace; $N = 5$ birds) or without headphones (gray trace; $N = 3$ birds). b) Mean change in pitch for individual lesioned birds subjected to zero pitch shift either with or without headphones.

only consists of the 6 data points shown in Figure 2.3a, it did not make sense to perform a bootstrap analysis (here SEM is measured across 6 data points; see Materials and Methods). Instead we used a one sample t-test and found that this distribution was not significantly different from zero ($t = 0.35$; $df = 5$; $p = 0.74$).

Data from 8 birds with 6-OHDA lesions but without any pitch shift revealed an unexpected systematic lowering of vocal pitch after dopamine depletion. Of those, 5 birds had headphones that conveyed unshifted auditory feedback (i.e. no pitch shift) and 3 birds had no headphones attached. When we analyzed the mean pitch change for each day for these two groups, we found them within error bars of each other for all 14 days of the experiment, and their pitch change across days 12 through 14 (-0.20 ± 0.14 with headphones; -0.16 ± 0.06 without headphones) were statistically indistinguishable (probability of resampled mean pitch with headphones greater than that without headphones was $p = 0.098$; see *Hypothesis testing with Bootstrap in Materials and Methods*). As a result, we combined the data from the 2 groups for the remainder of our analyses (the means for individual groups and traces for individual birds are shown in Fig. 2.4). The resulting mean shift in pitch during the course of the experiment is shown in Figure 2.3b. The overall shift in pitch over days 12 through 14 for this combined group was -0.19 ± 0.08 semitones. This decrease in pitch was statistically significant (probability of resampled mean pitch greater than or equal to zero was $p = 0.0029$), demonstrating, unexpectedly, that 6-OHDA lesions of Area X impacted song production by reducing the average pitch over time even in the absence of pitch-shifted auditory feedback.

2.4.3 6-OHDA lesioned birds do not respond adaptively to pitch-shifted auditory error:

In unlesioned animals, birds respond to a pitch shift through the headphones in an adaptive manner. Specifically, when subjected to a +1 semitone pitch shift through the headphones, the unlesioned

birds compensated adaptively by lowering their pitch (mean pitch change over days 12 to 14 for $N = 3$ birds was -0.40 ± 0.07 semitones; blue trace, Fig. 2.5a; probability of resampled mean pitch greater than or equal to zero was $p < 10^{-4}$; limit due to resampling 10^4 times) and when subjected to a -1 semitone shift in pitch, the unlesioned birds increase their pitch (mean pitch change over days 12 to 14 for $N = 3$ birds was 0.36 ± 0.11 semitones; red trace, Fig. 2.5a; probability of resampled mean pitch greater than or equal to zero was $p = 0.9996$, recall that in our bootstrapping analysis we conclude that distributions are significantly different if the probability that one is greater than or equal to the other is less than 0.025 or greater than 0.975; see Methods; traces for individual birds are shown in Fig. 2.6a). The result of plotting adaptive change in pitch (inverting y-axis for +1 semitone shift birds) for unlesioned birds is shown in Figure 2.5c (black trace). A direct comparison between the populations of -1 semitone shift and +1 semitone shift birds revealed a complete non-overlap among posterior distributions of sampled means (probability of resampled mean pitch for +1 semitone shift greater than or equal to that for -1 semitone shift was $p < 10^{-4}$; limit due to resampling 10^4 times). This resampling-based analysis reaffirms our initial finding (Sober and Brainard 2009) that unlesioned birds respond adaptively to pitch-shifted auditory errors and compensate accordingly for them, despite the fact that this earlier paper did not take into account the hierarchical nature of the data and the resulting propagation of uncertainty when computing statistical significance.

For 6-OHDA lesioned birds however, all birds decreased their pitch over time regardless of the direction of pitch shift through the headphones (Fig. 2.5b), similar to what we observed in lesioned birds with no pitch shifts (Fig. 2.3b). The +1 semitone shift group had a final pitch change of -0.38 ± 0.16 semitones (probability of resampled mean pitch greater than or equal to zero was $p = 0.0040$) while the -1 semitone shift group changed to a final pitch of -0.46 ± 0.19 semitones

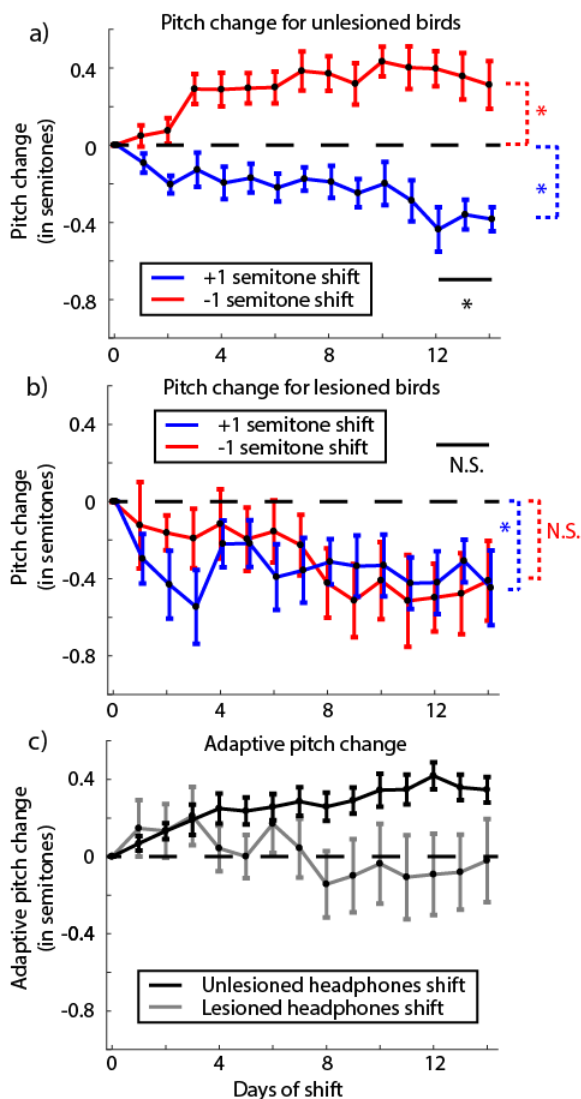


Figure 2.5: Change in pitch in response to pitch shift errors through the headphones in unlesioned and 6-OHDA lesioned birds.

a) Change in pitch from baseline over the period of pitch shift for unlesioned birds broken up by the direction of introduced shift in pitch (data reanalyzed from Sober and Brainard, 2009). The graph shows that birds increase their pitch over time in response to a downward pitch shift (red trace; $N = 3$ birds) and decrease their pitch to an upwards pitch shift (blue trace; $N = 3$ birds). Traces for individual birds are shown in Figure 4-1a. b) Same graph as in a) quantified for 6-OHDA lesioned birds ($N = 4$ birds for each trace). Individual birds are shown in Figure 4-1b. c) Adaptive change in pitch (see Results) for unlesioned birds (black trace; $N = 6$ birds) and 6-OHDA lesioned birds (gray trace; $N = 8$ birds). For a) and b), the * and N.S. in black represent significant and not significant differences respectively between the two shift conditions while the color coded differences check difference of each group from zero (see Results and Table 1).

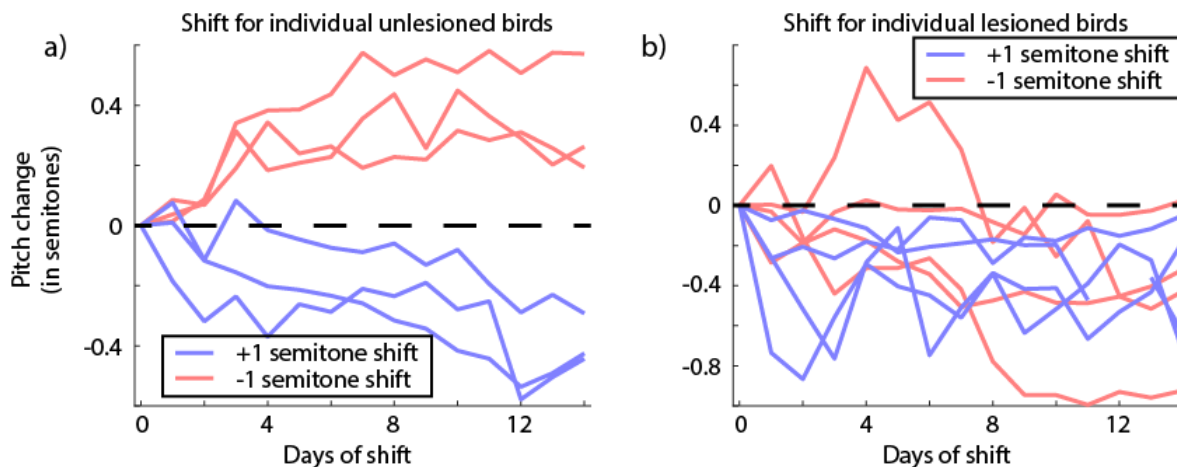


Figure 2.6: Individual traces for shifted birds.

a) Mean change in pitch for individual unlesioned birds subjected to a ± 1 semitone pitch shift. b) Mean change in pitch for individual lesioned birds subjected to a ± 1 semitone pitch shift. Note that one bird subjected to a +1 semitone shift has a discontinuity at shift day 12 since the bird did not sing at all that day. Also note how one bird in the -1 semitone shift group is at or slightly above zero by the end of the shift. This bird is the reason for the group not being statistically significantly below zero (this bird also had an extended baseline of 6 days; see Materials and Methods).

(probability of resampled mean pitch greater than or equal to zero was $p = 0.0747$) relative to the baseline (traces for individual birds are shown in Fig. 2.6b). The two groups were not statistically different from each other (probability of resampled mean pitch of +1 semitone shift group being greater than or equal to that of -1 semitone shift group was $p = 0.26$). We also compared each group to the no shift group and did not find statistically significant results (probability of resampled mean pitch of no shift group being greater than or equal to that of -1 semitone shift group was $p = 0.62$; probability of resampled mean pitch of no shift group being greater than or equal to that of +1 semitone shift group was $p = 0.91$). All statistical comparisons have been summarized in Table 1. Furthermore, when we quantified the adaptive change in pitch for this group, the final change in pitch was close to zero (gray trace, Fig. 2.5c). This suggests that following 6-OHDA lesions, birds do not respond adaptively to the auditory error. Instead, the birds seem to reduce their pitch

over time regardless of the direction or presence of pitch-shifted auditory error. Note that as was mentioned above and shown in Table 1, there was not a statistically significant difference between the Lesioned -1 semitone shift group and zero. This was due to the fact that while birds subjected to the -1 semitone shift did reduce their pitch on average, a few syllables for each bird increased their pitch, resulting in a group effect that fell short of significance. Since our error quantification treats the contribution from each syllable equally, the effects of individual syllables add up resulting in a not statistically significant difference (see *Error Quantification* under *Materials and Methods*).

Table 1: Statistical tests summary

Hypothesis tested - Bayesian Probability of group on left being \geq column heading (see <i>Hypothesis testing with Bootstrap</i> in <i>Materials and Methods</i>)			
Groups Compared	Zero	Lesioned +1 semitone shift	Lesioned -1 semitone shift
Lesioned 0 shift	0.0029	0.91	0.62
Lesioned +1 semitone shift	0.0040		0.26
Lesioned -1 semitone shift	0.0747		

Since the hierarchical bootstrapping as we have performed here to calculate statistical tests and standard errors has not been widely applied to such datasets in neuroscience previously, we also analyzed our data using hierarchical linear mixed models (LMMs) (Aarts et al. 2015, Aarts et al. 2014). LMMs have been widely applied to datasets involving large numbers of samples from a small number of subjects such as non-human primate studies (Pleil et al. 2016, Arlet et al. 2015) and rodent studies (Liang et al. 2015) or to analyze repeated measures or time series data (Wykes

et al. 2012, Howe et al. 2013). Specifically, we built LMMs to test the effects of the shift condition while controlling bird identity and specific syllables within each bird as variable effects (see section 2.3.9). For the unlesioned birds the linear mixed model revealed a strong effect of the shift condition ($t = 7.17$; $p = 7.92 * 10^{-13}$) on final pitch at the end of the shift period. For the 6-OHDA lesioned birds, the effect of the shift condition (+1 semitone shift vs -1 semitone shift vs no shift) was not significant ($t = 1.91$; $p = 0.056$). Also, when we combined the shift groups and compared them to the no shift groups, the effect was not statistically significant ($t = 1.47$; $p = 0.14$). That these models give us the same statistically significant results as our bootstrapping procedure gives us an independent verification of our error calculation and statistics.

2.4.4 No correlations between lesion extent and changes in pitch:

We measured the extent of 6-OHDA lesions by quantifying the proportion of histological sections that fell below the 5th percentile of section OD ratio for saline injected birds (see Methods). We can use this same threshold to obtain a rough metric of the lesion extent for each bird. Using this lesion extent, we computed correlations between the lesion extent and a variety of metrics of changes in pitch during the experiment (Table 1). However, we saw no significant correlations.

Table 2: Correlations between lesion extent and changes in song metrics

Lesion extent versus:	Pearson's correlation, r	Correlation significance, p
Final pitch change	0.4261	0.1466
Baseline variance	0.296	0.3261
Final variance	-0.0498	0.8716
Percent increase in variance	-0.4272	0.1454

2.4.5 Washout is impaired by dopamine depletion:

Following the end of the shift period, we turned the pitch shift through the headphones back to zero and recorded the birds' songs for an additional 6-7 days. During this period, birds without lesions typically revert their pitch back towards baseline levels (Sober and Brainard 2009). Hence, we refer to this period as washout. We first collected washout data from the birds that had 6-OHDA lesions and headphones but no shifts. As stated earlier, by days 12 through 14 of the shift period, these birds had a mean pitch of -0.20 ± 0.13 semitones. By days 6 and 7 of the washout period, their pitch had changed to -0.34 ± 0.15 semitones (Fig. 2.7a; traces for individual birds are shown in Fig. 2.8a). The probability of the resampled mean pitch during the end of the shift period being greater than or equal to that during the end of the washout period was $p = 0.67$. Therefore, although the change was not statistically significant, the mean pitch did drop further during washout. In order to quantify how much the pitch changes in response to the end of the sensory perturbation (pitch shift), we subtracted the mean pitch for each syllable on the last day of pitch shift throughout the entire washout period and quantified the resulting deviation in pitch (Fig. 2.9a). This emphasizes the dynamics of how the pitch changes or $\Delta(\text{Pitch})$ over time during washout in response to the end of the shift. The resulting change in pitch was found to be -0.12 ± 0.11 semitones (probability of resampled mean pitch greater than or equal to zero was $p = 0.22$).

Unlesioned birds displayed a robust return to baseline following the end of the pitch shift period as shown in Figure 2.7b (see traces for individual birds in Fig. 2.8b). For birds subjected to a -1 semitone shift, they reduced their pitch from 0.36 ± 0.11 semitones at the end of shift to 0.17 ± 0.08 semitones during the last 2 days of washout (probability of mean resampled pitch during washout being greater than or equal to that at the end of shift was $p = 0.08$). Equivalently, birds subjected to a +1 semitone shift increased their pitch from -0.40 ± 0.07 semitones at the end of the

shift period to -0.20 ± 0.05 semitones by the end of the washout period (probability of mean resampled pitch during washout being greater than or equal to that at the end of shift was $p = 0.98$). We also computed the dynamics underlying the $\Delta(\text{Pitch})$ over time during the washout period by subtracting the pitch for each syllable on the last day of shift through the washout period (Fig. 2.9b). Birds subjected to a +1 semitone shift, having reduced their pitch during the shift increased their pitch during washout. The last 2 days of washout had a mean change relative to the last day of shift of 0.17 ± 0.07 semitones (probability of resampled mean pitch lesser than or equal to zero was $p = 0.0003$). Similarly, birds subjected to a -1 semitone shift reduced their pitch back towards baseline during washout by -0.22 ± 0.11 semitones relative to the last day of shift (probability of resampled mean pitch greater than or equal to zero was $p = 0.0064$).

For our 6-OHDA lesioned birds, only 4 out of 8 birds had data for 7 days of washout due to difficulties in keeping the headphones attached (2 each for upward and downward shifts). We repeated the analysis for washout for these birds as described above for lesioned no shift and unlesioned birds. First, the mean change in pitch from the last day of shift through the washout period is shown in Figure 2.7c. Birds subjected to a +1 semitone shift returned their pitch back towards baseline increasing their pitch from -0.31 ± 0.19 semitones at the end of the shift period to -0.20 ± 0.14 semitones by the end of the washout period (blue trace in Fig. 2.7c, probability of mean resampled pitch during washout being greater than or equal to that at the end of shift was $p = 0.75$). Contrary to expectations however, the birds subjected to a -1 semitone shift drifted further away from baseline reducing their pitch from -0.16 ± 0.22 semitones at the end of the shift to -0.38 ± 0.30 semitones by the end of the washout period (red trace in Fig. 2.7c, probability of mean resampled pitch during washout being greater than or equal to that at the end of shift was $p = 0.35$). The traces for individual birds are shown in Figure 2.8c.

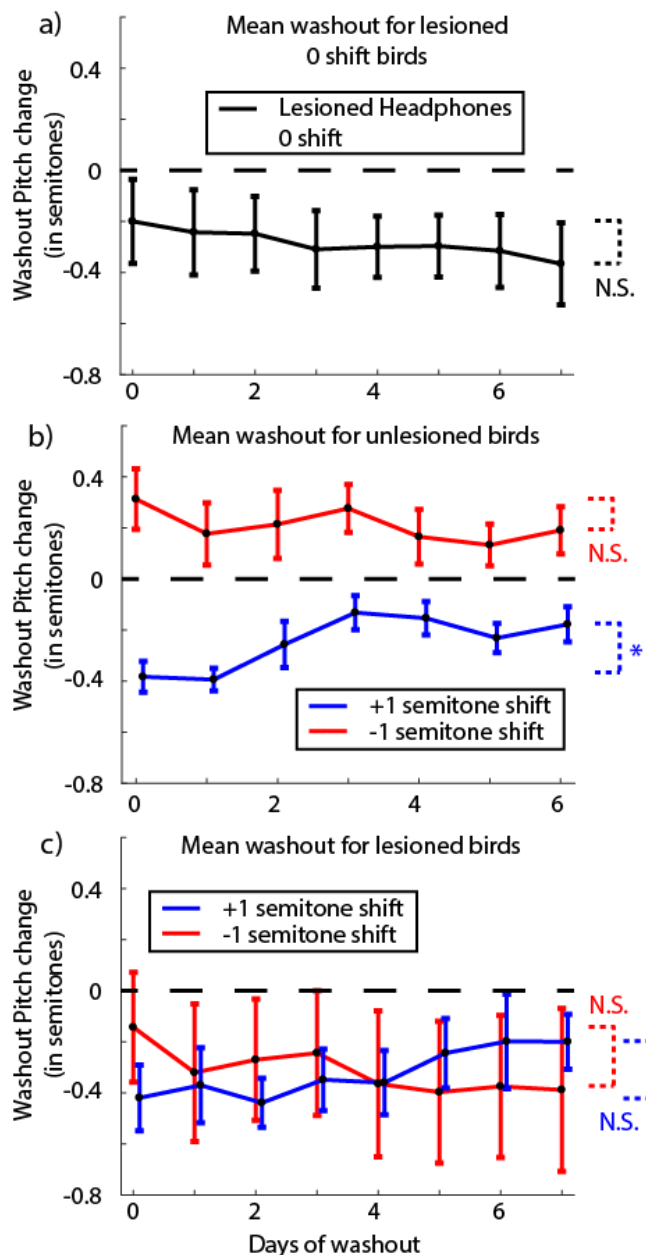


Figure 2.7: Analysis of change in pitch during washout for lesioned and unlesioned birds.

a) Mean change in pitch during “washout” for lesioned birds with headphones but no pitch shift (N = 5 birds). Day 0 refers to the last day of the shift period. Pitch shift is turned off at the end of this day. Individual bird traces are shown in Figure 5-1a. b) Mean change in pitch during washout for unlesioned birds (N = 3 birds for each trace). Individual bird traces are shown in Figure 5-1b. c) Mean change in pitch during washout for 6-OHDA lesioned birds (N = 2 birds for each trace). The extremely large error bars are due in part to the bimodal nature of the data (see individual birds in Fig. 5-1c). The statistical tests check the last three days of the shift period against the last two days of washout with * representing a significant difference ($p < 0.05$) and N.S. representing “not significant” (see Results for full tests).

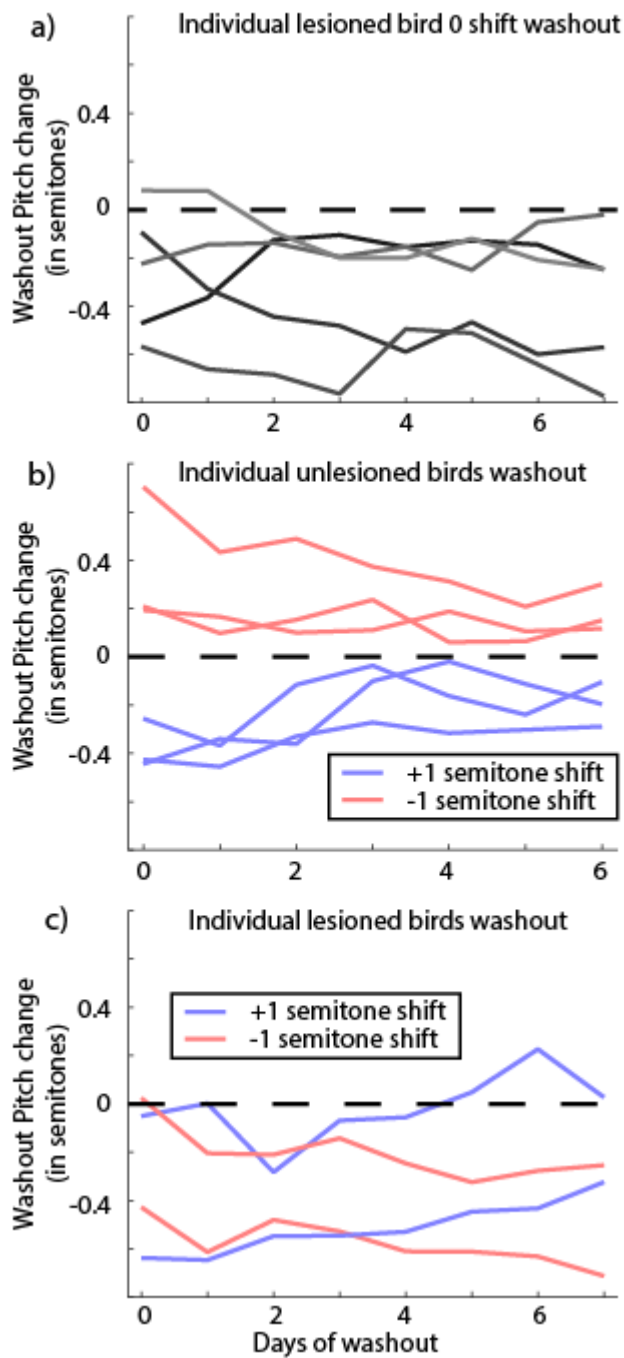


Figure 2.8: Washout traces for individual birds.

a) Individual birds that had a 6-OHDA lesion, with headphones but no pitch shift. Each color is a separate bird. b) Washout traces for individual birds that were unlesioned and subjected to a ± 1 semitone pitch shift. c) Washout traces for individual 6-OHDA lesioned birds subjected to a ± 1 semitone pitch shift.

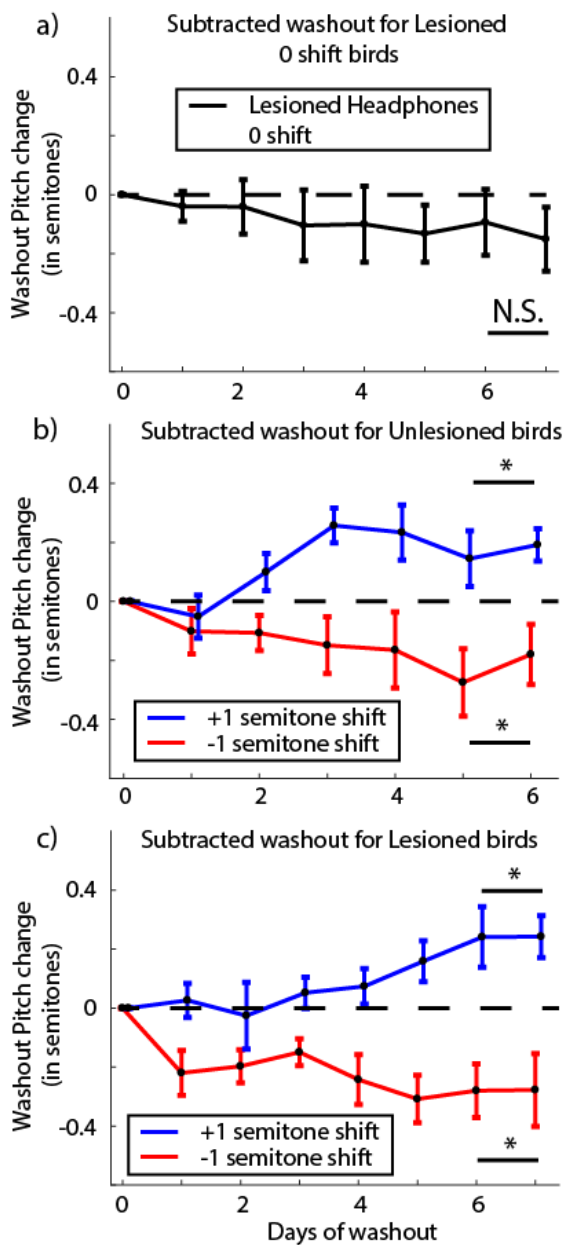


Figure 2.9: Dynamics of the change of pitch or $\Delta(\text{Pitch})$ during washout.

Note that this figure shows the same data as Figure 2.7, but with pitch data plotted relative to the pitch on the final shift day rather than to the experiment's baseline period as in Figure 2.7, a) $\Delta(\text{Pitch})$ during washout for lesioned no shift birds ($N = 5$ birds). b) The same analysis as in a) for unlesioned birds subjected to ± 1 semitone shift ($N = 3$ birds each). c) The same analysis as in a) for lesioned birds subjected to ± 1 semitone shift ($N = 2$ birds each). The * and N.S. refer to a significant difference versus not respectively for each group compared to zero over the last two days of washout.

Curiously, when we quantified the change in pitch in response to the end of the sensory perturbation subtracting the pitch change through the last day of shift through the washout period as before (i.e. measured the direction of pitch changes during washout, without considering the magnitude or direction of the pitch changes at the end of the shift period), the dynamics of the change in pitch was very similar to that seen in unlesioned birds (Fig. 2.9c). Lesioned birds subjected to a +1 semitone shift, averaging across the last 2 days of washout, shifted their pitch 0.24 ± 0.06 semitones with respect to the last day of shift (probability of resampled mean pitch lesser than or equal to zero was $p = 0.0003$). Lesioned birds subjected to a -1 semitone shift on the other hand, changed their pitch by -0.28 ± 0.11 semitones with respect to the last day of shift (probability of resampled mean pitch greater than or equal to zero was $p = 0.0182$). This result once again shows the dual effects we are observing following dopamine depletion. First, while not statistically significant, the pitch continued to drop for birds with unshifted auditory feedback. On the other hand, washout results between lesioned and unlesioned shift birds were very different in that washout was severely impaired in lesioned birds but confusingly followed the same dynamics for the $\Delta(\text{Pitch})$ over time following the end of the pitch shifted auditory feedback.

2.5 Discussion

Our results reveal two key effects of dopamine manipulation on the control of birdsong. First, all birds subjected to a 6-OHDA lesion of Area X displayed a drop in average vocal pitch which appeared between a week and two weeks post-lesion (Fig. 2.3b and Fig. 2.5b). Second, 6-OHDA lesioned birds displayed a severe deficit in sensorimotor learning as is evidenced by the lack of difference in response to a +1 or -1 semitone shift in pitch (Fig. 2.5b and gray trace in Fig. 2.5c).

While our primary finding seems to be one that implicates a role for dopamine in motor production, i.e., ability to produce higher pitched renditions of syllables in a bird's repertoire, there

is also a clear role for dopamine in learning the adaptive response to a sensory perturbation. It is true that when subjected to a +1 semitone pitch shift, there was no difference in mean change of pitch between lesioned (-0.38 ± 0.16 semitones) and unlesioned (-0.40 ± 0.07 semitones) birds (Fig. 2.5a and b, blue traces). However, when subjected to a -1 semitone pitch shift, while the adaptive response would be to raise their pitch, lesioned birds lowered their pitch (red trace in Fig. 2.5b). In addition, even for the lesioned birds subjected to a +1 semitone shift, their final change in pitch was not statistically different from the pitch drift seen in lesioned birds with no pitch shift (compare black trace in Fig. 2.3b with blue trace in Fig. 2.5b). This impairment in sensorimotor learning is reminiscent of deficits in learning in persons with Parkinson's disease (Paquet et al. 2008, Mollaei, Shiller, and Gracco 2013) and rodent models of dopamine depletion in striatum and motor cortex (Hosp and Luft 2013, Hosp et al. 2011, Shiotsuki et al. 2010). Hence our results suggest two factors at play, namely, motor production and sensorimotor learning. Disentangling these has been a hard problem in neuroscience (Beninger 1983, Wise 2004) since manipulations that affect motor learning also degrade motor production, complicating efforts to isolate learning mechanisms (Ungerstedt 1968, Cenci and Lundblad 2007, Iancu et al. 2005). Here, we isolated the lesions' effects on motor production by including the lesioned no shift group.

We have previously reported that 6-OHDA lesions of Area X do not produce any changes in number of songs produced or in any general motor behavior (Hoffmann et al. 2016). We similarly did not observe any qualitative difference in song quality or motor behavior between lesioned birds reported in this study and the birds reported in the 2016 study except the systematic drop in average pitch of songs sung post-lesion. Note however that the lesioned birds reported in this study were recorded from for 2 to 3 weeks longer post-lesion than those from the 2016 study

due to differences in time required to complete the behavioral experiments post-lesion. It therefore seems likely that this extended timeframe was necessary to observe the aforementioned pitch drop.

Vigor has been characterized as motivation (Salamone and Correa 2012, Salamone et al. 2007), speed of movements, or both (Mazzoni, Hristova, and Krakauer 2007, Turner and Desmurget 2010). A reduction in motor vigor following dopamine depletion could explain the systematic drop in pitch we observed. Dopamine has been shown to be associated with vigor in humans and other mammalian systems (Beierholm et al. 2013, Niv et al. 2007, Panigrahi et al. 2015, Berke 2018). In our experiments, we found that following 6-OHDA lesions of Area X the average pitch across all syllables for each bird dropped by roughly 11 to 13 days post-lesion. Higher pitched syllables require a combination of greater muscle activation and higher air sac pressure to be produced (Elemans et al. 2008, Goller and Suthers 1996, Riede, Fisher, and Goller 2010, Elemans et al. 2015) suggesting that higher pitched renditions of a particular syllable are more effortful to produce than lower pitched ones. We thus hypothesize that while unlesioned birds are capable of flexibly changing their pitch in a bidirectional fashion, dopamine lesioned birds will display a deficit in raising their pitch due to the increased effort required to do so. A related observation supporting our interpretation of our results is that birds sing at an elevated pitch when singing directed songs to females (Sakata, Hampton, and Brainard 2008, Leblois, Wendel, and Perkel 2010). Since it has also been reported that dopamine levels in Area X are elevated during directed song (Sasaki et al. 2006), this fits with the overall trend in our results.

Studies that have targeted individual syllables for pitch changes following dopamine depletions have not reported a systematic drop in pitch post-lesion (Hoffmann et al. 2016, Hisey, Kearney, and Mooney 2018). Our study does not necessarily contradict these results since those studies reported a deficit in learning post-lesion by either combining upwards and downwards

shifts (Hoffmann et al. 2016) or only driving pitch changes in one direction (Hisey, Kearney, and Mooney 2018). Additionally, for the birds reported in this study, while the average pitch across all syllables for each bird dropped, some individual syllables did increase their pitch. Furthermore, as noted above the birds in the present study were recorded for a longer period of time post-lesion than those reported previously.

The results from our washout data from the 6-OHDA lesioned birds are challenging to interpret. It is true that the lesioned birds subjected to a +1 semitone shift did return their pitch towards baseline and washout seemed to be unaffected for these birds (blue trace, Fig. 2.7c). Previous studies have reported that washout was not affected by dopamine depletion in tasks where birds shifted the pitch of a single syllable to avoid distorted auditory feedback (Hoffmann et al. 2016, Hisey, Kearney, and Mooney 2018). However, the birds subjected to a -1 semitone shift reduced their pitch resulting in their mean pitch moving further away from the baseline pitch (red trace, Fig. 2.7c). This suggests that washout is severely impaired in dopamine depleted birds. On the other hand, curiously, the change in pitch over time analyzed during washout in response to the end of the shift period was very similar between lesioned and unlesioned birds (compare Fig. 2.9b and c). We speculate that the lesion effects reported above could reflect either an inability to adaptively modulate motor output in response to error signals or from miscalculations in computing the error in the first place.

Adaptive sensorimotor learning in songbirds in response to induced auditory pitch shifts has been an effective paradigm to study the computational principles underlying sensorimotor learning (Kelly and Sober 2014, Sober and Brainard 2012, 2009). Bayesian inference works well to explain how unlesioned birds respond to auditory errors based on their prior experience of singing (Zhou et al. 2018, Hahnloser and Narula 2017). However, since 6-OHDA lesioned birds

exhibit drops in vocal pitch regardless of the direction of feedback pitch shift, any model that performs an adaptation to an error signal will fail to replicate the data without an additional mathematical mechanism to drive pitch downward in the presence of a reduced dopamine signal. One potential modification to the model would be to add a “relaxation state” into which the system relaxes in the absence of dopamine (Shadmehr and Arbib 1992, Shadmehr and Mussa-Ivaldi 1994). However, apart from the mean pitch, which did drop consistently across groups following 6-OHDA lesions, we did not find any other consistent relationships among other moments such as variance, skewness and kurtosis or overall probability distributions of produced pitch that could be used to constrain a revised Bayesian model to explain our results. Future work might therefore investigate the hypothesis that dopamine lesions disrupt sensorimotor learning by degrading the brain’s ability to perform Bayesian inference.

To conclude, our experiments show that dopamine plays a critical role in the brain’s ability to modulate vocal production in response to auditory errors. Future experiments will focus on disentangling specific roles for dopamine in sensorimotor learning by manipulating the dopamine signal at a faster temporal resolution. Results from such experiments could help fill gaps regarding the roles of tonic and phasic dopamine (Grace 1991) for example and the timeline of error correction. Eventually, results from various such experiments can be used to impose mathematical constraints on a computational model detailing the quantitative role of dopamine in such sensorimotor learning.

3 CHAPTER III: APPLICATION OF THE HIERARCHICAL BOOTSTRAP TO MULTI-LEVEL DATA IN NEUROSCIENCE

3.1 Abstract

A common feature in several types of neuroscience datasets is the presence of hierarchical data structures (such as recording the activity of multiple neurons in multiple animals across multiple trials). Due to such hierarchical structure, the measurements constituting the dataset are not independent, even though the traditional statistical analyses often applied in such cases (e.g. student's t-test) treat them as such. The hierarchical bootstrap has been shown to be an effective tool to accurately analyze such data and while it has been used extensively in the statistical literature, it is not as widespread in neuroscience despite the ubiquity of hierarchical datasets. We use simulated neural data to show that traditional statistical tests can result in a false positive rate of over 45% even if the Type-I error rate is set at 5%. While summarizing data across the non-independent points (or lower levels) can potentially fix this problem, this methodology greatly reduces the statistical power of the dataset. The hierarchical bootstrap, when applied sequentially over the levels of the hierarchical structure, keeps the Type-I error rate within the intended bound and retains more statistical power than summarizing methods. We conclude by demonstrating the effectiveness of the method in two real-world examples, first analyzing singing data in male Bengalese finches (*Lonchura striata* var. *domestica*) and second quantifying changes in behavior under optogenetic control in flies (*Drosophila melanogaster*). We present the hierarchical bootstrap as an intuitive and powerful tool to analyze such hierarchically nested datasets.

3.2 Introduction

It is commonplace for studies in neuroscience to collect multiple samples from within a category (e.g. multiple neurons from one animal) to boost the sample size. A recent survey found that of 314 papers published in prominent journals covering neuroscience research over an 18 month period in 2013-14, roughly 53% of those studies had nested datasets featuring such hierarchical data (Aarts et al. 2014). When data are collected this way, the resulting data points are not independent. However, commonly deployed statistical tests like the Student's t-test and ANOVA treat all data points as independent. This results in an underestimation of uncertainty in the dataset and a corresponding underestimation of the p-value (Musca et al. 2011, Hahs-Vaughn 2005, Arceneaux and Nickerson 2009). This arises due to the problem of intraclass correlation or ICC (Walsh 1947, Kish 1965) and pseudoreplication (Hurlbert 1984, Lazic 2010) in which variance within a cluster (or lower hierarchical level) and variance between clusters (or higher levels) are not propagated appropriately. To illustrate this problem, consider a hypothetical example in which one measures changes in dendritic spine size during learning. Since one can typically only measure from a few animals each in different treatment conditions, researchers usually increase sample sizes by measuring multiple spines from each neuron and by measuring multiple neurons within an animal. The hierarchical nature of such datasets can result in different samples not being statistically independent from each other: in the above example, spines measured from the same neuron may be more similar than spines measured across different neurons and even more so than spines measured from different animals within the same treatment condition. Such data points should not be treated as statistically independent but still very frequently are.

Linear Mixed Models (LMMs) can be used to account for the variance across different levels (Aarts et al. 2015, Aarts et al. 2014) and have recently been used to do so in several studies

(Arlet et al. 2015, Pleil et al. 2016, Liang et al. 2015, Machado et al. 2015). However, LMMs assume that all hierarchical structure present is linear which is often not true for typical datasets. Additionally, concerns have been raised about the bias and reliability of parameters returned by LMM fits when the number of clusters is small as is also often the case in neuroscience datasets (Maas and Hox 2005, Huang 2018, Gehlbach et al. 2016).

The hierarchical bootstrap (Efron 1981, 1992, Efron and Tibshirani 1994, Carpenter, Goldstein, and Rasbash 2003) is a statistical method that has been applied successfully to a wide variety of clustered datasets including census and polling data, education and psychology, and phylogenetic tree data (Efron, Halloran, and Holmes 1996, Harden 2011, Huang 2018). Unlike LMMs, the hierarchical bootstrap is relatively agnostic to the underlying structure present in the data and has consistently performed better at quantifying uncertainty and identifying signal than traditional statistics (Field and Welsh 2007, Harden 2011, Thai et al. 2013) though some concerns have been raised that the bootstrap may be excessively conservative in a limited subset of cases (Hillis and Bull 1993, Adams, Gurevitch, and Rosenberg 1997). However, the use of the hierarchical bootstrap in neuroscience is limited even though its application is increasingly warranted.

This paper is divided into two parts. In the first, we simulate a typical dataset studied in neuroscience and use it to illustrate how the Type-I error is inflated in hierarchical datasets when applying traditional statistical methods but can be averted using the hierarchical bootstrap. In the second, we demonstrate the use of the hierarchical bootstrap in two real-world examples using singing data from songbirds (Hoffmann and Sober 2014) and optogenetic control of behavior in flies (Cande et al. 2018). In both cases, the data have a strong hierarchical structure and our

analyses highlight the need to use appropriate statistical tests when analyzing hierarchical datasets in neuroscience.

3.3 Materials and Methods

The simulations for this paper were run in the Jupyter Notebooks environment using Python (version 3.7.2) and importing the following libraries: NumPy (version 1.15.4), SciPy (version 1.1.0), Matplotlib (version 3.0.2) and Pandas (version 0.23.4). Reanalysis of data from Hoffmann and Sober (2014) was performed using MATLAB (version 2017a). The codes for both simulation and the data analysis will be made available on Github post-publication.

3.3.1 Traditional vs Summarized vs Bootstrap

Throughout this paper, we compare 3 statistical methods that we refer to by shorthand as “Traditional”, “Summarized” and “Bootstrap” respectively. Throughout this paper, when we refer to the “Bootstrap” method, we mean a hierarchical bootstrap procedure. We will detail what each of those terms mean here (see Fig. 3.1 for schematics of each). For the sake of clarity, let us consider a fictitious example. Suppose our dataset involves recording the neural activity of neurons in the amygdala when an individual was exposed to an aversive auditory cue either in the presence or absence of a drug of interest believed to reduce anxiety. Each neuron was recorded for around one hundred trials of exposure to the auditory cue and the process was repeated for several hundreds of neurons in both the presence and absence of the drug (see Fig. 3.1a). We could add a layer of complexity by considering that the experiment was repeated across several individuals but for the sake of simplicity, let us assume that all the data were collected from a single individual. In the “Traditional” method every data point (i.e. the firing rate of every neuron to every instance of the auditory cue) is treated as independent, regardless of the hierarchical structure present in the

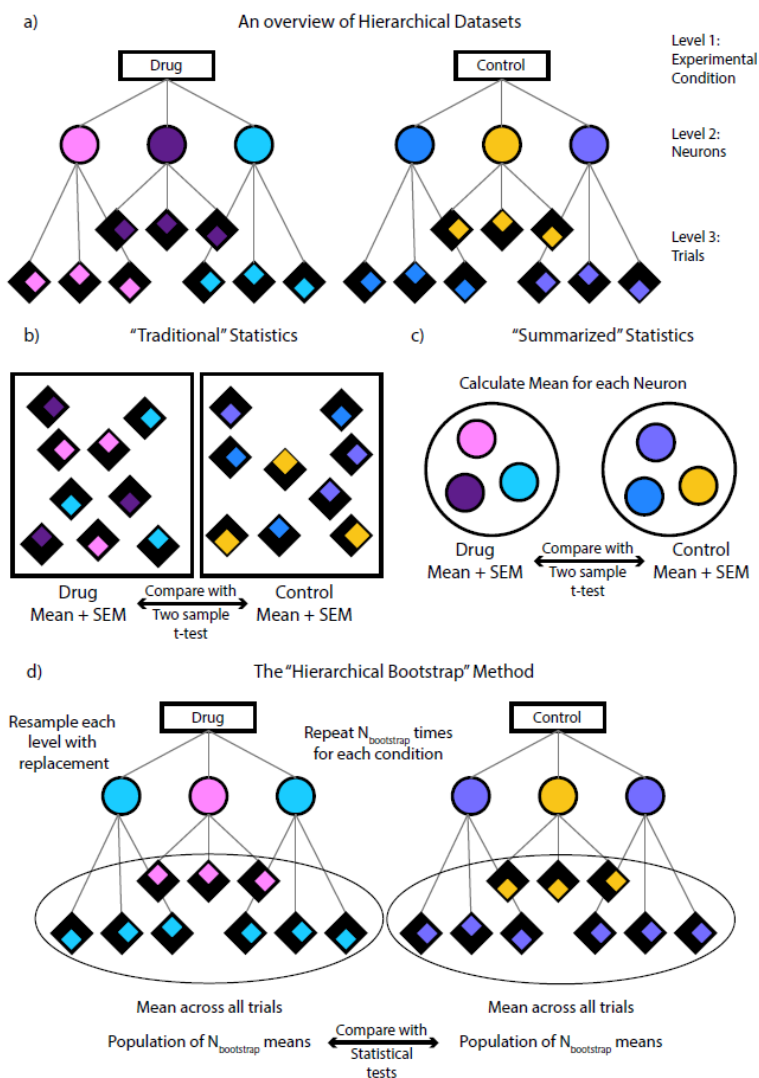


Figure 3.1: Schematic of hierarchical datasets and pseudocodes for statistical methods considered.

a) An example of a hierarchical dataset. Here the dataset is divided into 3 levels with the first level containing the experimental groups to be compared, the second containing the individual neurons and the third containing the neuronal firing rate during each trial. Each neuron is color coded and the trials per neuron are distinguished by the position of the colored diamond. b) In "Traditional" statistics, the means for each group is computed across all the trials and are then compared using a two sample t-test. c) In "Summarized" statistics, the mean for each neuron is computed first. These means are then used to compute an overall mean for each group and the groups are compared using a two-sample t-test. d) In the "Hierarchical Bootstrap" method, we create new datasets $N_{\text{bootstrap}}$ times by resampling with replacement first at the level of neurons followed by trials within a neuron. We then compute the mean across all trials every time we perform resampling. The final statistic is computed on this population of resampled means (see Methods for details).

dataset (see Fig. 3.1b). All the data points are used to calculate the mean and the uncertainty in the estimate of the mean, namely the standard error of the mean (SEM) and a Student's t-test is used to ascertain statistically significant differences between the mean firing rate of the neurons in the presence versus absence of the drug of interest. The "Summarized" method, on the other hand, acknowledges the possibility that repeated trials within the same neuron may be more similar to each other than trials across neurons. As a result, the mean firing rate for each neuron is calculated first and the mean of the group is calculated as the mean of the population of mean firing rates for each neuron in the group and the SEM is computed from this population of means (see Fig. 3.1c). Note that the mean for each group in this case is equal to that in the "Traditional" case if and only if the number of trials recorded for every neuron within a group is the same, i.e., if every neuron is represented equally. A Student's t-test is thus applied to the population of mean firing rates between the two groups. An additional complication that we circumvent in our toy example by considering all the data to be obtained from a single subject is the decision as to which level one must summarize the data. In the case of multiple subjects, one may summarize either at the level of individual neurons or individual subjects. While summarizing at the level of subjects is the most appropriate way to avoid non-independence between data points, it can seriously reduce sample size and therefore power. Finally in the "Bootstrap" method, we perform the hierarchical bootstrap on the two groups to compute posterior distributions of the range of means possible from each group (see Fig. 3.1d), as follows. First, we sample with replacement (i.e., we sample from the current distribution in such a way that replications of previously drawn samples are allowed) from the neurons in the group. Then, for the neurons selected, we then sample with replacement the individual trials for the number of times each neuron was recorded. We then compute the mean firing rate across the group for that resampled population and repeat the entire process $N_{bootstrap}$

times ($N_{bootstrap} = 10^4$ for all instances in this paper unless otherwise noted). The mean for each group in this case is identical to that computed in the “Traditional” method. The 67% confidence interval (or equivalently, the standard deviation) of the population of means so obtained gives an accurate estimate of the uncertainty in the mean. Note that the mean is a special case where this uncertainty is more commonly referred to as the Standard Error of the Mean or SEM. We can then compute the probability of one group being different from the other using the population of resampled means obtained above for each group (see Hypothesis testing with Bootstrap below for complete details).

3.3.2 Hypothesis testing using Bootstrap:

We described above how bootstrap samples can be used to compute the uncertainty in measuring the mean of a population. However, the bootstrap can be used more broadly to measure the uncertainty in any metric of interest as long as it obeys the law of large numbers and scales linearly in the probability space. In addition, the bootstrap is used to compute posterior distributions of the range of values possible for the metric of interest from the data of limited sample size. As a result, the distribution of bootstrap samples hence computed can be used to compute probabilities that the data supports particular hypotheses of interest directly. We will describe below how this can be done with an example. Note that while this is not the only way of hypothesis testing using bootstrapping, we found this to be a particularly simple and effective way of doing so.

As will be done several times in this paper, suppose we wished to evaluate the support for the hypothesis that the mean of particular sample was significantly different from a fixed constant, zero for our example. In order to do so, we would compute the proportion of means in our population of bootstrapped means of the sample that were greater than or equal to zero. If we set the acceptable false positive (Type-I error) rate to α ($\alpha = 0.05$ throughout this paper), then if the

computed proportion was greater than $1 - \alpha/2$ (or $p > 0.975$ for $\alpha = 0.05$) we would conclude that the sample of interest had a mean significantly greater than zero. Alternatively, if the computed proportion was less than $\alpha/2$ (or $p < 0.025$ for $\alpha = 0.05$) then we would conclude that the sample of interest had a mean significantly less than zero. Any proportion $\alpha/2 \leq p \leq (1 - \alpha/2)$ would indicate a relative lack of support for the hypothesis that the mean of the sample of interest is different from zero. In the case of multiple comparisons, we use the Bonferroni correction to adjust the threshold for significance accordingly.

We would also like to make a distinction between the probabilities we referred to above and the p-values typically associated with statistical tests. p-values refer to the probability of obtaining a result as extreme or more extreme than those obtained under the assumption that the null hypothesis is true. As such, they do not provide a direct measure of support for the hypothesis one truly wishes to test, i.e., the likelihood of the alternate hypothesis being true. The 'p' referred to in the bootstrapping procedure above however, provides such a direct probability of the tested hypothesis being true. For the rest of this paper, in order to distinguish the direct probabilities obtained using the bootstrapping procedure from p-values reported from traditional statistical tests, we will use 'p_{boot}' to refer to bootstrap probabilities and 'p' to refer to p-values from other tests.

While it is not performed in this paper, the procedure described above can also be used to compare the means of two different groups using their respective samples. In this case, we would compute a joint probability distribution of the two samples with each sample forming the two axes of a 2-D plot. In this case, the null hypothesis would be a circle centered on the line $y = x$. Therefore, to test if the two groups are different, one would compute the total density of the joint probability distribution on one side of the unity line. If the volume computed is greater than $1 - \alpha/2$ then the first group is significantly greater than or equal to the second while if the volume

computed is less than $\alpha/2$, the second group is significantly greater than or equal to the first with all other volumes indicating no significant differences between the groups. We can also extend this formulation to comparisons between multiple groups by performing pairwise-comparisons between the groups and adjusting the threshold for significance accordingly (by Bonferroni correction for example).

3.3.3 Design Effect (DEFF):

When one analyzes data from hierarchical datasets, the unique information provided by each additional data point at the lowest level of the hierarchy depends on the average number of samples in the cluster and the relative variance within and between clusters. This relationship was mathematically quantified using the Intra-cluster correlation (a.k.a. intra-class correlation) or ICC. ICC is a useful metric that provides a quantitative measure of how similar data points are to each other within an individual cluster in a hierarchical dataset (Walsh 1947, Kish 1965). While there are some differences in how it is calculated, in general it is defined as the following ratio:

$$ICC \text{ or } \rho = \frac{s_{between}^2}{s_{between}^2 + s_{within}^2}$$

Where $s_{between}^2$ represents the variance across cluster means while s_{within}^2 represents the variance within clusters. Hence, the ICC is a metric that varies from zero to one where a measure of zero represents no clustering of data and every data point being independent while a measure of one represents perfect reproduction of samples within clusters, i.e., all points within a cluster are exactly the same. Kish further formalized the relationship between ICC and the adjusted effect size that was termed the “Design Effect” or DEFF with a corresponding correction to be applied to the standard error of the mean computed from the dataset termed DEFT, defined as the square root of DEFF (Kish 1965, McCoach and Adelson 2010). Formally, DEFF was defined as:

$$DEFF = \frac{\text{var}(data)}{\text{var}(data \text{ if independent})} = 1 + \rho * (\bar{n}_j - 1)$$

Where \bar{n}_j represents the average sample size within each cluster and ρ is the ICC. Hence, as the number of samples within a cluster increases, the DEFF increases resulting in a need for a larger correction (increase) to the standard errors for accurate quantification. Conversely, as the number of samples within clusters increase, the standard error of the mean is underestimated potentially resulting in underestimation of the p-values and inflation of the Type-I error rate.

3.4 Results

Our results section has been organized into two sub-sections: Simulations and Examples. In the Simulations sub-section, we show results from simulations that illustrate the utility of the hierarchical bootstrap and in the Examples sub-section, we highlight the differences in results when analyzing data in two examples with and without the hierarchical bootstrap. Throughout the results section, we will compare statistical tests we refer to by shorthand as “Traditional”, “Summarized” and “Bootstrap” respectively. See *Traditional vs Summarized vs Bootstrap in Materials and Methods* and Figure 3.1 for a detailed description of the differences between the three conditions. Also note that whenever we refer to the “Bootstrap” in this section, we mean the hierarchical bootstrap unless otherwise specified.

3.4.1 Simulations

We used simulations of neuronal firing in order to highlight the key characteristics of the hierarchical bootstrap as applied to nested data in neuroscience and the differences between the bootstrap and other more commonly used statistical tests. Specifically, we were interested in whether the bootstrap displayed a conservative bias for independent and non-independent datasets as well as in quantifying the bootstrap’s statistical power compared to other techniques. While

these results may be derived from other mathematical results previously published (Davison and Hinkley, 1997; Carpenter et al., 2003), we found them instructive to depict explicitly.

3.4.1.1 The hierarchical bootstrap is more conservative than Traditional and Summarized methods for independent data points:

It has been reported earlier that the bootstrap has a conservative bias (Adams, Gurevitch, and Rosenberg 1997, Hillis and Bull 1993), resulting in larger error bars than strictly necessary for the chosen threshold of Type-I error α (here set to 0.05). It has also been argued that this is not a bug or bias in the algorithm, but rather a more generic property of hypothesis testing by resampling (Felsenstein and Kishino 1993, Efron, Halloran, and Holmes 1996) and newer algorithms have claimed to reduce bias further (Shimodaira 2004, 2002). Here we tested the conservative bias of the bootstrap by running a simulation checking for significant differences when none existed and quantified the proportion of cases that returned a significant difference. Given that we set α to 0.05, we would expect a 5% false positive rate if there was no bias in the algorithm.

We simulated a situation in which we recorded the activity of 1000 neurons over 100 trials each. The neurons were simulated using a Poisson random number generator with an average firing rate of 5Hz (each trial was considered to be 1 second of activity). Note that while we have set the problem up as a hierarchical problem, since the average firing rate for each neuron was kept constant both across neurons and within a neuron, each sample in this dataset is indeed independent. Therefore, we would not expect differences between the Traditional and Summarized methods. We then split these 1000 neurons into two groups of 500 each randomly and computed the mean firing rate for each group. We then tested whether the means were significantly different from each other using the Traditional, Summarized and Bootstrap methods. We repeated this analysis 10000 times and plotted the proportion of trials that resulted in significant differences for




Figure 3.2: Results from the simulation in which there was no difference between the groups of two neurons.

a) Proportion of significant results when comparing the 2 groups with each statistical method at α of 0.05. As expected, both the Traditional and Summarized methods give roughly 5% false positive results. However, the bootstrap gives a much smaller proportion of significant results suggesting a conservative bias. b) The size of SEMs computed using each of the methods. The bootstrap does give an error bar roughly 1.4 times that of the other two metrics.

each of the methods in Figure 3.2a. The error bars were computed by bootstrapping the results obtained from the simulation runs. As shown in the figure, both the Traditional and Summarized methods resulted in a proportion of significant results close to and not significantly different from 5% as expected (Traditional – 4.86 ± 0.21 %; probability of proportion of significant results being greater than or equal to 0.05 was $p = 0.26$; Summarized – 4.95 ± 0.22 %; probability of proportion of significant results being greater than or equal to 0.05 was $p = 0.42$). By contrast, when using the bootstrap method, the proportion of significant results was significantly lower than the expected 5% at 0.66 ± 0.08 %. Even when we increased the value of α to 0.1, the proportion of significant results was still only 2.19 ± 0.15 % (probability of proportion of significant results being greater than or equal to 0.05 was $p < 10^{-4}$ in both cases; limit due to resampling 10^4 times).

We also computed the standard error of the mean (SEM) in each case and reported the results in Figure 3.2b. As shown, the error bars for both the Traditional and Summarized methods are almost identical at $1.002 \pm 0.002 * 10^{-2}$ for Traditional and $0.994 \pm 0.023 * 10^{-2}$ for Summarized respectively. The error bars computed using the Bootstrap method are roughly 1.4 times larger at $1.407 \pm 0.035 * 10^{-2}$. Since the effect size is inversely proportional to the uncertainty in the dataset (Coe 2002), which is captured here by the error bars, we conclude that the larger error bars do partially account for the drop in proportion of significant results observed and that the bootstrap seems to have a conservative bias for independent datasets.

3.4.1.2 *The bootstrap does not have a conservative bias in a hierarchical dataset:*

If the bootstrap does have a strong conservative bias regardless of the nature of the data (hierarchical or independent), it may not be the right metric with which to address the problem of statistical analysis in hierarchical datasets. In the first experiment, we tested a situation in which the variance was the same between levels resulting in all data points being independent in spite of

a hierarchical structure. Here, we abolished that independence by adding Gaussian noise to the first level creating a truly hierarchical dataset. Specifically, we used the same situation as before where we had 1000 neurons each with a mean firing rate of 5Hz, that were split randomly into 2 groups. However, each neuron now had a mean firing rate of 5Hz plus Gaussian random noise of width 3Hz. Neurons were still drawn independently, so there should be no difference between the two groups. We also varied the number of trials per neuron to study its effect on the false positive rate. We simulated the experiment 1000 times for each value of number of trials and computed the false positive rate from each. As before, we used bootstrapping on the obtained results to estimate error bars and to test for significant differences away from 0.05. Given the relationship between the number of points within a cluster to the Design Effect (DEFF; see Intra-cluster correlation in *Materials and Methods*), we would expect the false positive rate to increase with the number of trials per neuron (Snijders 2011, Snijders and Bosker 1993, Aarts et al. 2014).

As shown in Figure 3.3a, the false positive rate for the traditional method does increase with the number of trials per neuron rising from around 46% for 10 trials to almost 96% in the case of 3000 trials per neuron (probability of resampled proportions being greater than or equal to 0.05 was $p > 0.9999$ in all cases; limit due to resampling 10^4 times). On the other hand, both the summarized and bootstrap methods stayed remarkably similar in value and were not significantly different from 0.05 in all cases (adjusting for threshold of significance with Bonferroni corrections for 3 comparisons). This was a marked departure from Figure 3.2a where we saw that the bootstrap had a significant conservative bias.

We also computed the estimate for the SEM using all 3 cases for each number of trials simulated, and the result is shown in Figure 3.3b. As shown, the SEM estimate remains fairly constant for both the summarized and bootstrap methods but decreases with an increase in the




Figure 3.3: False positive rate and size of error bars quantified for the simulation in which there was again no difference between the groups but now the points were not independent.

a) The proportion of significant results using each statistical method as a function of the number of trials per neuron. As expected, the false positive rate for the traditional method rises with increasing number of samples within each neuron. The Summarized and Bootstrap methods on the other hand have almost identical false positive rates close to the theoretical 5% in all cases. b) The size of SEMs computed for all 3 methods as a function of the number of trials per neuron. While those for both summarized and bootstrap stay roughly the same, those for the traditional method reduces with increasing number of trials. Note that for both traces, since the Bootstrap and Summarized almost perfectly overlap, the green trace has been thickened for visualization.

number of trials per neuron in the traditional case. Furthermore, the SEM estimate for the traditional case starts out much lower than either the summarized or bootstrap case suggesting that the increased false positive rate is at least partially due to the underestimation of error in the traditional case.

3.4.1.3 The Bootstrap balances intended Type-I error rate and power better than Traditional and Summarized methods at low sample sizes:

As we saw in the previous section, both the summarized and bootstrap methods bind the Type-I error rate at the intended 5% and the estimate of the SEM is roughly the same for both methods. What then is the advantage of the bootstrap method over simply using the summarized method? The answer lies in the fact that the summarized methods result in a loss of statistical power, i.e. the ability to detect a statistical difference when one truly exists, particularly for low sample sizes of the upper hierarchical levels and for small effect sizes. We used simulations to calculate the power for each of the three methods and the results are shown in Figure 3.4.

Since power depends on the effect size and the number of samples, we chose to study the change in power with respect to the number of neurons per group (N) and the effect size for these simulations ($\Delta mean$). In order to do so, we varied N between 1 and 16, keeping the number of trials per neuron constant at 100 each. We kept the mean firing rates of one group of neurons at 5Hz as before and varied the mean firing rate of the other group by $\Delta mean$, adding an additional 3Hz random Gaussian noise to each neuron in both groups. Since the previous simulations did not estimate the false positive rate when the number of neurons was as low, we first kept the mean firing rate for both groups of neurons equal at 5Hz, simulating 1000 times for each value for the number of neurons per group. The result is shown in Figure 3.4a. As shown, the false positive rate for the traditional method stays around or above 80% (blue trace in Fig. 3.4a), while that for the

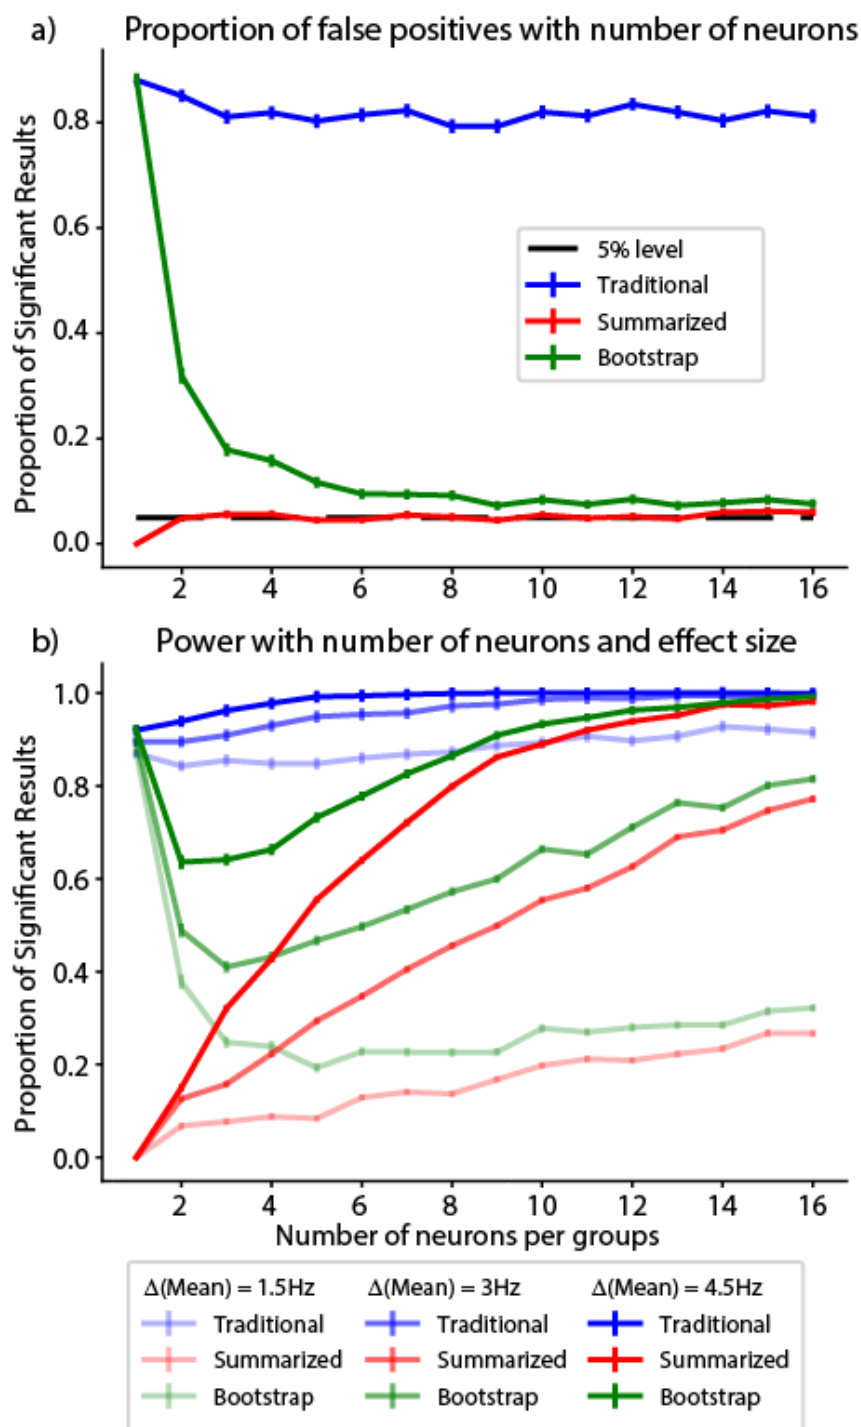


Figure 3.4: Change of power with number of neurons and effect size.

a) The false positive rate when there was no difference between the mean firing rates for the two groups of neurons. b) The proportion of significant results or power when the difference in mean firing rates (Δmean) between the two groups of neurons was 1.5Hz (light), 3Hz (medium) and 4.5 Hz (dark) respectively.

summarized method hugs the expected 5% line (red trace in Fig. 3.4a) except for the special case of 1 neuron per group where you can never achieve significance since you are comparing two points. The behavior of the bootstrap highlights the fundamental characteristic of the bootstrap and is therefore worth exploring in detail (green trace in Fig. 3.4a). The essence of the bootstrap is to provide a reliable range for your metric under the assumption that the limited dataset you have captures the essential dynamics of the underlying population. When there is only one neuron, the bootstrap assumes that trial level data is the true distribution and therefore has a false positive rate equal to that of the traditional method. As the number of neurons increase, one gets a better sampling of the true underlying distribution and correspondingly, the bootstrap tends towards a 5% error rate with increasing number of neurons as the weight of data points shifts from individual trials to trials across neurons with increasing number of neurons. Therefore, if the data collected does not accurately represent the dynamics of the underlying distribution, the bootstrap cannot provide accurate estimates of population metrics.

We then computed the power for the three methods as a function of the number of neurons per group and the difference in mean firing rate between the groups. Accordingly, we repeated the simulations described above changing the mean firing rate of one of the groups to 6.5Hz, 8Hz and 9.5Hz (light, medium and dark traces in Fig. 3.4b respectively). Since there is an actual difference between the groups in this case, the ideal plot will have a very high proportion of significant results barring adjustments for extremely low sample sizes. As shown, the traditional method has the most power, but as was seen in Figure 3.4a, also has a very high false positive rate for this type of data (blue traces in Fig. 3.4b). The summarized method has the lowest power among the three methods, but does catch up for large effect sizes and with increasing group sizes (red traces in Fig. 3.4b). The bootstrap is between the two extremes and has more power than the summarized metric

particularly for small effect sizes and small group sizes (green traces in Fig. 3.4b). As a result, we see that the bootstrap helps retain statistical power while also being sensitive to the Type-I error rate. However, as shown in Figure 3.4a, the bootstrap can weight trials within levels more heavily than one would expect if the number of samples in the upper levels is very low and one must therefore be mindful when dealing with very low sample sizes that their data collected may not represent the true distribution in the population.

3.4.2 Examples

We now present two real-world examples of the utility of the hierarchical bootstrap as applied to behavioral data collected from experiments in songbirds (Hoffmann and Sober 2014) and flies (Cande et al. 2018). These examples provide concrete instances of why one should use the appropriate statistical tests depending on the nature of their data and how the popular tests can result in more false positives or less statistical power than one desires.

3.4.2.1 The bootstrap highlights the prevalence of false positives using traditional statistical methods in strongly hierarchical datasets as in vocal generalization in songbirds:

As described above, although the bootstrap provides a better compromise between statistical power and false-positive rate than the Traditional or Summarized methods, its use is not widespread in the neuroscience literature, including in some of our own prior published work. To illustrate the practical importance of these issues, and to encourage other authors to critically re-evaluate their prior analyses, we here present a case in which we have used the bootstrap to reexamine one of our prior results – which used both Traditional and Summarized methods – and found the choices made can significantly affect the outcome of our analyses. As a reminder, when discussing Traditional or Summarized statistical tests, we will report a p-value denoted by ‘p’ which yields a significant result if $p < 0.05$. When talking about the Bootstrap tests however, we will report a p_{boot}

which in turn yields a significant result if $p_{\text{boot}} < 0.025$ or $p_{\text{boot}} > 0.975$. In addition, p_{boot} provides a direct probability of the hypothesis being true.

Our lab studies vocal behavior in songbirds, which yields datasets with an inherently hierarchical structure: each bird sings a variety of syllables and each syllable is repeated a different number of times. In many of our studies, we examine changes in the pitch of these syllables in response to manipulations. We performed an experiment studying generalization of vocal learning in songbirds in response to an induced auditory pitch shift on a particular (target) syllable (Hoffmann and Sober 2014). In these studies, the auditory feedback of one syllable was shifted in pitch and relayed to the bird through custom-built headphones with very short (~ 10 ms) latency, effectively replacing the bird's natural auditory feedback with the manipulated version (Hoffmann et al. 2012, Hoffmann and Sober 2014, Sober and Brainard 2009). Note that while the headphones provided auditory feedback throughout the song, only the feedback for the single syllable targeted for pitch shift was shifted in pitch. We reported that in addition to birds changing the pitch of the target syllable in response to the pitch shift, the birds also changed the pitch of other syllables that had not been shifted in pitch. Specifically, we reported that syllables of the same type (acoustic structure) as the target syllable changed pitch in the same direction as the target syllable (“generalization”) while syllables of a different type than the target syllable changed pitch in the direction opposite to that of the target syllable (“anti-adaptive generalization”; see Fig. 5a). Since in Hoffmann and Sober (2014) we employed traditional and summarized (at a syllable level) statistics when analyzing generalization, we decided to reanalyze the data from that study to ask if the generalization observed was still statistically significant when statistical tests were computed using the hierarchical bootstrapping procedure. In order to do so, we first recapitulated the results reported by computing statistics on the last 3 days of the shift period using the traditional

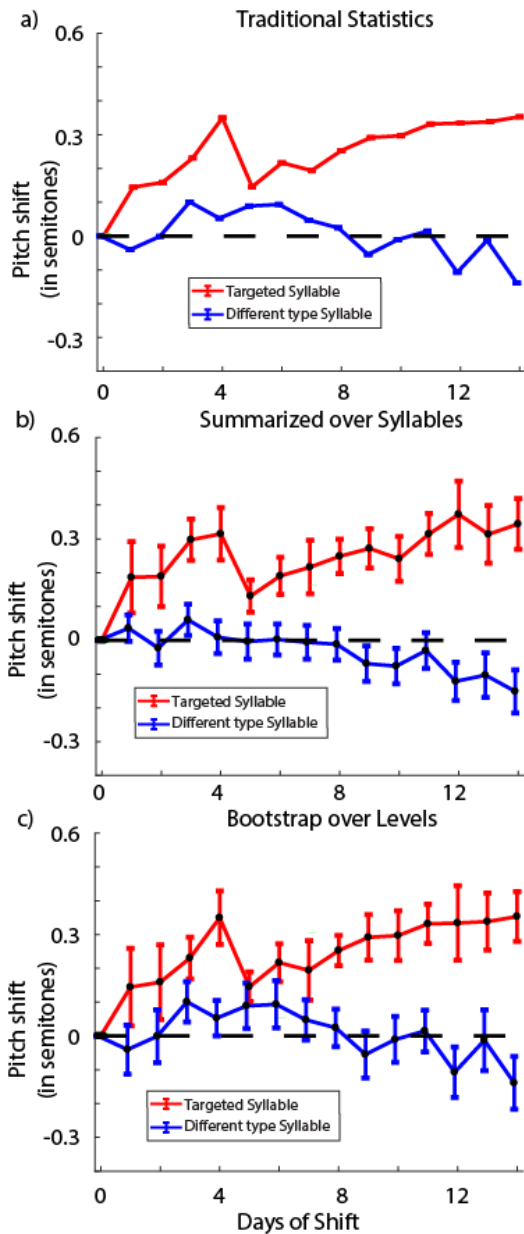


Figure 3.5: Reanalysis of generalization in the headphones learning paradigm.

a) The results when quantified using traditional statistics. As shown, both target and different type syllables differ significantly from zero over the last 3 days of the shift with target syllable moving adaptively and the different type syllables showing anti-adaptive generalization respectively. b) The results when quantified using summarized statistics when summarized over the syllables. In this case, the target syllable is significantly different from zero while the different type syllables are just over the threshold for significance. c) The results when bootstrap is applied over the hierarchical levels. The target syllable is significantly different from zero but the different type syllables are not.

and summarized methods as was reported earlier (Hoffmann and Sober 2014). We focus our reporting on changes in the target syllable and anti-adaptive generalization in different type syllables for the purpose of this example.

When we computed the change in pitch over the last 3 days of the shift period for the syllable targeted with auditory pitch shifts, we found that the birds compensated for the pitch shift of 1 semitone by 0.341 ± 0.007 (mean \pm SEM in all cases) semitones with traditional statistics (one sampled t-test comparing to zero; $t = 47.3$; $p < 2 \cdot 10^{-308}$; limit due to smallest number in MATLAB; red trace in Fig. 5a) and by 0.34 ± 0.08 semitones with summarized statistics (one sampled t-test comparing to zero; $t = 4.25$; $p = 0.004$; red trace in Fig. 5b). We did see anti-adaptive generalization in different type syllables of -0.087 ± 0.003 semitones with traditional statistics (one sampled t-test; $t = 23.9$; $p = 4 \cdot 10^{-125}$; blue trace in Fig. 5a). With summarized statistics, the different type syllables changed by -0.12 ± 0.06 semitones (one sampled t-test; $t = 2.00$; $p = 0.053$; blue trace in Fig. 5b) and was not statistically significant, though it was close to the threshold for significance. When we reanalyzed the data using bootstrapping over the hierarchical levels, we found that we did not have enough statistical power to claim that the anti-adaptive generalization was statistically significant. As expected, the targeted syllable shifted significantly away from zero to a final value of 0.34 ± 0.12 semitones (probability of resampled mean being greater than or equal to zero was $p_{\text{boot}} = 0.995$; red trace in Fig. 5c). As a reminder, p_{boot} gives the probability of the hypothesis tested being true. Therefore, a value of 0.5 indicates minimal support for either the hypothesis (or its opposite) while values close to 1 (or 0) represent strong support for (or for the opposite of) the hypothesis. Different type syllables however shifted to a final value of -0.09 ± 0.09 (probability of resampled mean being greater than or equal to zero was $p_{\text{boot}} = 0.25$; blue trace in Fig. 5c). Hence, this result shows that the anti-adaptive generalization was too small an effect to detect with the

sample size in the original study suggesting that the generalization effects observed were driven largely by a couple of individual birds rather than a population wide effect.

We will note that we also reanalyzed generalization in same type syllables and did not find significant generalization there either. Therefore, it is reasonable to say that we did not perform the generalization experiment on sufficient number of birds to adequately power the study. However it is worth noting that while the results did not meet the threshold for statistical significance, reporting probabilities in support of the hypotheses provides more information than simply determining whether or not a statistical threshold was met. This is particularly relevant since we found that if we only looked at the last two days of shift, the probability of same type syllables being greater than or equal to zero increased to $p_{\text{boot}} = 0.91$. While still not at the threshold for statistical significance, this coupled with the fact that independent studies have reported adaptive generalization for same type syllables in songbirds (Tian and Brainard 2017) suggests this may yet be a true effect. A final confound is that due to a discrepancy in our data archiving, our recapitulation of the old analysis for the paper yielded a slightly different p-value ($p = 0.053$) for summarized analysis of different type syllables than was originally reported in the original paper ($p < 0.05$; records indicate it was ~ 0.048). The point of this analysis is therefore not to replicate the exact findings but to highlight how choices made for statistical analyses can define the interpretation of one's results.

3.4.2.2 The hierarchical bootstrap captures the signal better than traditional or summarized methods in optogenetic control of behavior in flies:

We wanted to test the utility of the hierarchical bootstrap in an independent example, and so we chose to analyze the data from an experiment studying the role of descending neurons in control of behavior in flies (Cande et al. 2018). Studies involving optogenetics are another area where

hierarchical datasets are the norm. Each fly, since it can be tracked over extended periods of time, will exhibit each behavior multiple times within the period of observation. Additionally, multiple flies can be tracked simultaneously, and the behavior is typically averaged across flies across trials for each experimental group. In this study, the authors used optogenetics to activate descending neurons in flies and studied the corresponding changes in behavior displayed. In order to do so, the authors first created a two-dimensional representation of the behavior of the flies in the absence of any manipulations, as has been described in detail previously (Berman et al. 2014, Cande et al. 2018). They then mapped the behavior of experimental animals both in the presence and absence of light stimulation as well as control animals that were not fed retinol, a binding co-factor needed for functionality of the light-activated channels, both in the presence and absence of light stimulation onto the behavioral representation. The resulting map of behavior for one class of descending neurons is shown in Figure 3.6. In order to assess whether the light stimulation caused a statistically significant change in the frequency of behavior observed, the authors argued that the frequency of behavior had to be significantly greater during optical stimulation than during periods of no stimulation within experimental animals. In addition, the frequency of behavior during optical stimulation had to be greater in experimental animals than in control animals. The authors used Wilcoxon rank summed tests coupled with Sidak corrections (Šidák 1967) for multiple comparisons in order to test for statistically significant differences. This would fall under the category of traditional methods as we have described previously. We compared the regions obtained from the original analysis with regions we obtained when using summarized or bootstrap statistics on this dataset and the result is shown in Figure 3.6. As shown, the traditional method seems to overestimate the region of significant differences and includes a false positive area that is separate from the main region where signal is present. The summarized method, on the other

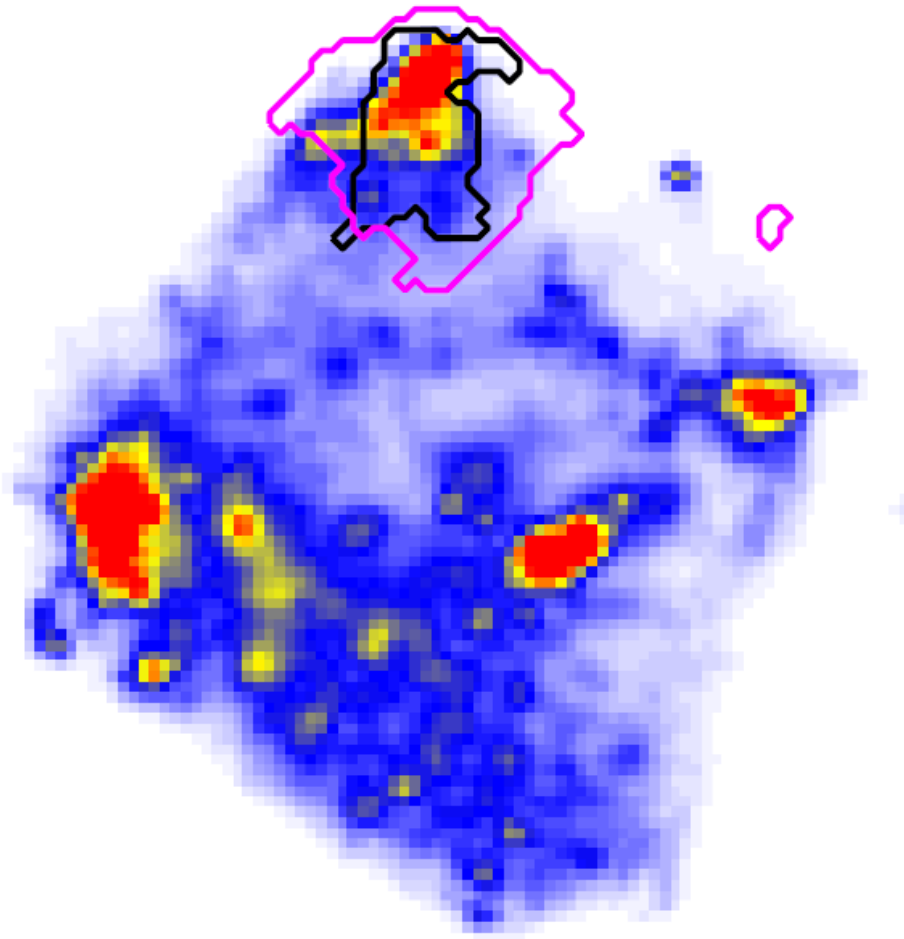


Figure 3.6: Hierarchical bootstrapping most accurately captures the signal in an experiment studying optogenetic control of behavior in flies.

The differences in frequency of behavior mapped onto a two-dimensional space when a particular group of descending neurons in the experimental flies were manipulated using optogenetic control. In this particular case, the descending neurons targeted were controlled head grooming behavior represented at the top of the map and so, the animals display elevated frequencies of head grooming during light stimulation when compared to control flies or when the light was turned off. The magenta trace shows the statistically significant differences after accounting for multiple comparisons when using the traditional method. As shown, the magenta trace seems to overestimate the signal present and captures some false regions as well as shown in the upper right region. The black trace represents the areas of significant difference as defined by using the hierarchical bootstrap and it matches the region expected by the authors as per video analysis very well. The summarized method did not return any regions that were statistically different between groups even though there was a clear signal present in the data (see videos and other data in Cande et al., 2018).

hand, does not identify any regions as being statistically significant despite video evidence suggesting the clear presence of some signal in the data (see Cande et al., 2018). The hierarchical bootstrap returns a concise region tightly mapped to the region the authors would have expected based on analysis of behavioral videos. Hence, this is another example showing that the hierarchical bootstrap can be a powerful tool to accurately quantify results in neuroscience where a majority of datasets analyzed are hierarchical in nature and therefore the data points are not independent.

3.5 Discussion

The hierarchical bootstrap is a powerful statistical tool that was developed to quantify the relationship between school class sizes and achievement (Carpenter, Goldstein, and Rasbash 2003) and has since been used to quantify effect sizes in a wide variety of fields. The use of the hierarchical bootstrap in neuroscience however is still limited in spite of the need for it being increasingly clear. Through our simulations, we have shown the utility of the hierarchical bootstrap in neuroscience by examining the shortcomings of more common statistical techniques in a typical example one might encounter. While the results of our simulations may be inferred from other mathematical work on the subject (Carpenter, Goldstein, and Rasbash 2003, Davison and Hinkley 1997, Field and Welsh 2007, Goldstein 2011), to our knowledge our results have not been shown explicitly in previous work. We first illustrated that the bootstrap does not have a conservative bias for hierarchical datasets in which the assumption of independence between data points is violated (Fig. 3.3). We then showed that the bootstrap performs better than summarized statistical measures by not sacrificing as much statistical power especially at low sample sizes and small effect sizes (Fig. 3.4). Finally, we showed real world applications of applying the hierarchical bootstrap to two

datasets from songbirds and flies to demonstrate the advantages of the hierarchical bootstrap over other more commonly used statistical analysis techniques in neuroscience (Figs. 3.5 and 3.6).

“Pseudoreplication” refers to studies in which samples were either not replicated or the replicated samples were not independent, yet statistical tests performed treated the data points as independent replicates (Hurlbert 1984). While pseudoreplication was first extensively reported on in ecological studies (Heffner, Butler, and Reilly 1996, Hurlbert 1984), it has since been identified as a common problem in other fields, including neuroscience (Lazic 2010). While resampling methods including bootstrapping were originally suggested as tools by which one could overcome the pseudoreplication problem (Crowley 1992), the bootstrap was argued to have a conservative bias resulting in larger error bars than necessary (Hillis and Bull 1993, Adams, Gurevitch, and Rosenberg 1997). Since then, however, several versions of the bootstrap algorithm have been developed to apply to hierarchical data and have been found to be unbiased and more robust for calculation of uncertainty in clustered data than other statistical methods (Field and Welsh 2007, Harden 2011, Huang 2018, Carpenter, Goldstein, and Rasbash 2003, Goldstein 2011, Davison and Hinkley 1997). In order to test the bootstrap for any potential bias in a typical example we might encounter in neuroscience, we produced simulations to quantify differences in mean firing rates between two groups of neurons when there was no difference between the groups. We illustrated that the bootstrap produced a false positive rate significantly below the expected 5% (Fig. 3.2a) and had larger error bars (Fig. 3.2b) than other statistical methods when the data were independent. However, when the independence between data points was abolished by introducing a hierarchical structure, the bootstrap was not statistically different from the expected 5% false positive rate (green bars in Fig. 3.3a) and the error bars computed were similar to those computed using

summarized statistics (red and green bars in Fig. 3.3b) suggesting that the hierarchical bootstrap is robust to bias for applications in neuroscience.

Among the reasons Linear Mixed Models (LMMs) gained in popularity for statistical testing was the fact that they could accommodate hierarchical datasets by controlling for various levels as “random” effects while still using all available data thereby minimizing loss in statistical power (Snijders 2011, Snijders and Bosker 1993, Hox, Moerbeek, and Van de Schoot 2017, Roux 2002, Aarts et al. 2014). Though we did not directly compare the loss in power between bootstrapping and LMMs, we showed that the bootstrap also does not lose power to the degree that using summarized statistics does (see green traces versus red traces in Fig. 3.4b) while also keeping the false positive rate within the intended bound (see green trace in Fig. 3.4a). Additionally, unlike LMMs which assume linearity, the hierarchical bootstrap as applied in this paper does not make assumptions about the relationships underlying the latent variables that define the hierarchical structure. However, as we saw in Figure 3.4a, the bootstrap assumes that the data collected captures essential characteristics of the population distribution. One may have to adjust the resampling procedure to ensure that the distribution of resampled data points most accurately matches the population distribution one wishes to study.

We then used the hierarchical bootstrap on two independent examples to showcase its utility in analyzing hierarchical datasets in neuroscience. First, we reanalyzed data from Hoffmann and Sober, 2014 in which we used both Traditional and Summarized statistical analysis to conclude that songbirds generalize changes in pitch targeted on a single syllable anti-adaptively to syllables of a different type. When reanalyzed with the bootstrap however, we found that the anti-adaptive generalization of different type syllables (blue trace in Fig. 3.6c) did not meet the threshold for statistical significance. This was a particularly striking result as the original study did report

statistically significant changes from zero even while using summarized statistics (Hoffmann and Sober 2014). A probable reason for the differences between the summarized and bootstrap methods for this dataset stems from a decision point regarding the level to which one must summarize the data when using summarized statistics (we avoided this decision in the simulations by assuming all data came from a single subject). The summary statistics reported were summarized at the level of syllables for this dataset. However, in order to truly make sure all points are independent, one must summarize at the highest level, i.e., at the level of individual birds in the dataset. The differences in results between the summarized and bootstrap methods here suggest that the generalization effects were driven largely by strong effects in a subset of birds as opposed to a population-wide effect and that, by failing to take the hierarchical nature of the dataset into account, the original authors overestimated their statistical power and chose too low an N. Further evidence for this interpretation comes reanalysis of data from a separate study looking at learning birds display in response to pitch shift of their entire song through custom-built headphones (Sober and Brainard 2009) using the hierarchical bootstrap. Since the changes in pitch were far more systemic across birds in this experiment, we did not see any changes in statistically significant results (Saravanan et al. 2019b).

Second, we used the hierarchical bootstrap on an independent experiment studying the role of descending neurons in controlling behavior in flies using optogenetics (Cande et al. 2018). As shown in Figure 3.6, the hierarchical bootstrap performs better than the traditional and summarized statistical methods in isolating the true signal in the experiment. The traditional method includes areas that are likely false positives and the summarized method does not return any statistically significant areas.

We would also like to reiterate another advantage of the direct probabilities returned by the bootstrap (p_{boot}) over traditionally reported p-values. p-values represent the probability of obtaining results at least as extreme as the ones obtained under the assumption that the null hypothesis is true. It is a cumbersome definition that has led to numerous misconceptions regarding what it actually means (Halsey et al. 2015, Wasserstein and Lazar 2016). The value returned by the bootstrap however, p_{boot} , provides a direct probability in support of a particular hypothesis. In the case of the generalization in songbirds example, we found that the probability of same-type syllables generalizing was 0.85. This means that if we measured data from more birds drawn from the same distribution, we will see adaptive generalization in 85% of cases which is much higher than chance (50%). Hence, the hierarchical bootstrap method can provide a measure of the relative support for the hypothesis which is both easier to understand and can be useful information for both positive and negative results in research.

To conclude, neuroscience research is at a crossroads wherein on the one hand exciting new technologies are being built promising bigger and more complex datasets to help understand brain function (Yizhar et al. 2011, Burns et al. 2013, Vogelstein et al. 2018), and on the other we have rising concerns over the incorrect use of statistical tests (Ioannidis 2005, Nieuwenhuis, Forstmann, and Wagenmakers 2011, Greenland et al. 2016) and the lack of reproducibility of a number of past findings (Baker 2016, Gerlai 2019, Miłkowski, Hensel, and Hohol 2018). We propose the hierarchical bootstrap as a powerful but easy-to-implement method that can be scaled to large and complicated datasets, that returns a direct probability in support of a tested hypothesis reducing the potential for misinterpretation of p-values and that can be checked for correct implementation through sharing of analysis code. As we have shown through this paper, we

believe that widespread use of the bootstrap will reduce the rate of false positive results and improve the use of appropriate statistical tests for a given type of dataset.

4 CHAPTER IV: CONCLUSIONS AND FUTURE DIRECTIONS

The central question at the heart of this dissertation was: “*Is dopamine involved in sensorimotor adaptation? If yes, how?*” Through my experiments detailed in Chapter 2, I have found evidence that dopamine is involved in sensorimotor adaptation in that dopamine depletion caused severe disruptions to sensorimotor adaptation. However, the interpretation was complicated by the fact that we also observed a concurrent effect on motor production. Hence the question as to how dopamine is involved in sensorimotor adaptation has not been answered in its entirety by my dissertation and future experiments will be required to elucidate this. This chapter details some of the future experiments that could be performed to understand the computations dopamine performs in sensorimotor adaptation.

4.1 Dopamine as a multiplexed signal

As I have shown in Chapter 2, dopamine seems to perform more than one single function. Specifically, dopamine seems to be involved in both processing sensory and reward prediction errors (Gardner, Schoenbaum, and Gershman 2018, Berke 2018, Coddington and Dudman 2018, Schultz 2016, Schultz, Dayan, and Montague 1997) as well as modulating motor output (Barbeau 1962, Mazzoni, Hristova, and Krakauer 2007, Panigrahi et al. 2015, Saravanan et al. 2019a). This is even without considering the hypothesized roles for dopamine in sleep and wakefulness regulation (Monti and Monti 2007, Korshunov, Blakemore, and Trombley 2017), mood disorders (Nestler and Carlezon Jr 2006, Diehl and Gershon 1992), addiction (Johnson and Kenny 2010, Volkow et al. 2011) and other neurological disorders (Cook Jr et al. 1995, Huang et al. 2015, Howes, McCutcheon, and Stone 2015, Grace 2016). All this is leading to an increasing perception that dopamine is a multiplexed signal with multiple roles and functions that also targets several different regions of the brain. If we assume the assertion that dopamine is a multiplexed signal to

be true, that naturally leads to two follow-up questions: 1) how does the signal get multiplexed? and 2) how is the multiplexed signal decoded at the relevant target area?

There has been quite a bit of work into understanding what sort of multiplexing happens both at the level of the dopamine signals from the VTA themselves (Kremer et al. 2018, Lee et al. 2019) as well as at the level of the neurons where the dopamine neurons project to like the striatum and the somatosensory and motor cortices (Ramakrishnan et al. 2017). However, even these studies observe the existence of multiplexed neuronal signals at the inputs and outputs of the neurons but not the mechanism of encoding and decoding a multiplexed signal in itself. One of the roadblocks to such studies has been methodological. Until recently, we did not have the tools required to even consider building an experiment capable of answering this question. However, by combining optogenetic stimulation of the VTA in a biologically plausible manner over a timescale typical of learning tasks with local field potential (LFP) recording of activity in the medial prefrontal cortex (mPFC), Lohani and colleagues recently showed how dopamine activity causes weak and heterogeneous responses at the level of individual neurons but sustained changes in ensemble activity, gamma-theta coupling and gamma oscillations over multiple timescales (Lohani et al. 2019). Further work in this direction could help understand how different target brain regions of the dopaminergic neurons interpret the multiplexed signal they receive specific to their functions.

One of the most commonly discussed mechanisms hypothesized to underlie dopamine signal decoding involves the different affinities of the D1 and D2 families of dopamine receptors. In general, D1 receptors are thought to have a low affinity for dopamine requiring concentrations in the micro-molar range to be activated while D2 receptors have a high affinity and are saturated at concentrations in the nano-molar range (Kebabian and Calne 1979, Gingrich and Caron 1993)

though this simplistic view has been challenged in recent years (Yapo et al. 2017). Yet fundamentally, D1 receptors belong to the family of G-protein coupled receptors that stimulate the production of adenylyl cyclase and are generally thought to lead to neuron excitation while D2 receptors inhibit the same and are generally thought to lead to neuron inhibition. Hence, one could plausibly envision a role for these two types of receptors, combined with their differing affinities for dopamine, in decoding a multiplexed dopamine signal. Given the predominant separation of dopamine receptors in the dorsal striatum and their corresponding direct and indirect pathways, this hypothesis seems plausible. However, it has recently been shown that such a distinction is not as clear in the ventral striatum (Kupchik et al. 2015). Additionally, the picture is complicated in the songbird as D1 and D2 receptors are known to co-localize on the same neurons in the striatum making the distinction between the two pathways less clear (Farries, Ding, and Perkel 2005, Carrillo and Doupe 2004, Kubikova, Wada, and Jarvis 2010). We know from experiments in this dissertation that songbirds display as strong a role of dopamine in learning and movement as in many other common model systems. How then is this signal still correctly interpreted by the post-synaptic neuron?

While we performed lesions for the experiments described in this dissertation, systemic infusion of D1 and/or D2 specific agonists and antagonists are another way to study the potential effects of dopamine on the two types of receptors. Recently, a study in songbirds found that while D1 receptor antagonists caused the birds to display deficits in reinforcement learning, D2 receptor antagonists caused the number of songs sung to drop precipitously so much so that most birds did not sing enough songs per day to be included for analysis (Hisey, Kearney, and Mooney 2018). Such results point to a potential for the different receptor types to be involved in disambiguating the multiple roles dopamine seems to perform. Future experiments in this direction would involve

using optogenetics (Boyden et al. 2005) or DREADDs (Armbruster et al. 2007) to specifically target individual types of receptors and measure their effects of activation or blocking during various behavioral tasks.

Additionally, there has been new evidence to suggest that the dopaminergic projections themselves may not be uniform in their function with respect to targets. The textbook version of the dopaminergic terminals from neurons originating in the VTA/SNc is that they spread diffusely to targets throughout the brain including but not limited to the ventral striatum or Nucleus Accumbens, the dorsal striatum and areas of the prefrontal cortex. Mohebi and colleagues have shown that while the underlying neural activity in the VTA/SNc seems to be the same, the release of dopamine at the axon terminals seems different among the various areas. In particular, they found that dopamine levels correlated with an animal's motivation to engage in particular tasks in the Nucleus Accumbens but not in other areas. Additionally, while they observed such an increase in *dopamine release* in the Nucleus Accumbens, they did not observe a corresponding change in the *neural activity of dopaminergic neurons* projecting to the Accumbens suggesting an external mechanism of release (Mohebi et al. 2019). Future experiments can focus on understanding the cause of this observed dopamine release, one of the candidates for which is the cholinergic interneurons present in the striatum (Threlfell et al. 2012, Ding et al. 2010, English et al. 2012).

I mentioned above that one of the limitations in understanding dopamine signaling thus far has been methodological and that we are beginning to develop tools to address this. I will expand on this slightly in the following section. Another major source of problems in interpreting results from various studies in neuroscience involves the interconnected nature of the brain and our limited tools to such complex circuits with extensive feedback loops. I will expand on this idea in section 4.3.

4.2 New Techniques to understand brain function

The majority of work into the role of dopamine so far has been performed using electrophysiological studies involving either recordings of single unit neurons with a single electrode or through single or multi-unit recordings using multi-electrode arrays. These methods are extraordinarily intensive and yet have an extremely small throughput in terms of the number of units one can simultaneously record from in relation to the population in the area being recorded from. Techniques such as voltammetry and microdialysis have also been used but suffer from less temporal resolution and in the case of microdialysis, limited spatial resolution as well.

Increasing the number of units from which one can simultaneously record neural activity has been a long standing goal in neuroscience. The hypothesis behind the goal is that increasing the number of units one can record from simultaneously will provide a more accurate snapshot of the state of the brain at various time points and therefore will lead to a better understanding of the computations being performed. This idea has led to several grand projects such as the Blue Brain project (Markram 2006), the Human connectome project (Van Essen et al. 2013) and others. In terms of methodological advances, one of the most prominent and quickly growing fields has been that of optical imaging. While electrophysiological recordings provide a temporally detailed account of the neural activity, the spatial extent that each electrode can cover is limited. This limitation can be mitigated with a little sacrifice on the temporal resolution through optical imaging using techniques such as voltage sensitive dye imaging and calcium based fluorescent imaging. While the use of voltage sensitive dyes has not been widespread due to the fact that the technique can only be used for a limited time on the surface of the brain, calcium based fluorescent imaging has been advancing rapidly. Calcium based fluorescent imaging is based on the principle that cells increase their intracellular calcium concentrations when they fire action potentials. By inserting

proteins that are sensitive to this increase in calcium concentration and go into a configuration that allows them to fluorescence tracking the calcium concentration, we can get a proxy for the electrical activity of target neurons using an optical signal. This led to the development of one-photon based fluorescent imaging better known as fiber photometry (Guo et al. 2015, Meng et al. 2018) as well as two-photon based fluorescent imaging (Tian et al. 2009). The disadvantage of fiber photometry is that it only provides as much coverage as the electrical methods mentioned previously since it only provides one pixel of signal. On the other hand, two photon methods currently require animal subjects to be head-fixed for successful recording though the technology is improving to allow for recording from freely behaving animals in some cases (Zong et al. 2017). I previously mentioned that in both cases, the temporal resolution is not quite as good as that of electrical signals but newer versions of the GCaMP protein promise to show better temporal resolution and higher signal to noise ratios (Inoue et al. 2019).

Techniques such as that described above could theoretically allow us to genetically target a subgroup of neurons, dopamine neurons from the VTA or striatal D2-receptor carrying medium spiny neurons for example, and record their collective neural activity over an extended period of time while the animal is performing a behavior of interest. Such experiments coupled with studies that uncover the structure of inputs and outputs to the brain regions of interest could help us contextualize the signals observed to their corresponding inputs and the behavioral outputs observed.

Another technological frontier has been the development of tools that can record neural activity for extended periods of time. Current protocols typically record neural activity on the timescale of days to at most months. This is due to glial scarring and other issues with long term recording that cause the signal to noise ratio to drop leading to a loss of viable signal (Polikov,

Tresco, and Reichert 2005). Carbon fiber electrodes (Huffman and Venton 2009) and other nanoelectronic probes (Luan et al. 2017) are being actively researched for use in chronic recording and stimulation paradigms. Long term recording of neural activity could allow us to quantify changes in neural activity in response to learning or repeated exposure to certain tasks or cues.

4.3 The problem of Information in an Interconnected Network

Deep neural networks are thought to be the category of machine learning algorithms that are closest to the structure of the brain in a nervous system. Like the brain, they receive inputs and must use the inputs to produce an output. Like the brain, they consist of multiple layers feeding into one another between the input and the output layer. Also like the brain, the inner workings of the network are not well understood especially as the number of layers between the input and the output increases.

One of the most frequent criticisms against deep neural networks is that they often feel like black boxes transforming the inputs into the required outputs with little to no understanding of how they do so. As the number of hidden layers increases and the number of nodes within a layer increases it becomes increasingly difficult to understand, even with post-hoc analysis, what the particular contribution of each layer (let alone each node within a layer) is in the input-output transformation. As complex as deep neural networks are though, there are no connections between nodes within a layer, no reciprocal connections or feedback loops in a simple feedforward network and there is a single source of input with a single category of expected output. The brain however has all of these in addition to the features already mentioned.

Everything we know about the brain has been gleaned from recording neural activity in a brain region of interest and correlating that activity either to sensory inputs experienced by or motor outputs produced by the subject. If we go back to the neural network analogy, this works

well to understand the activity at the input and output layers which in the case of the brain, correspond to the sensory and motor areas respectively. However, if we want to understand the activity of neurons in layers in between these two extremes, things become more complicated since we are probing a multi-dimensional space with a limited dimensional dataset (Ganguli and Sompolinsky 2012). In order to be successful in doing so, more data as described in section 4.2 might help but in addition, we need new theories for analyzing and interpreting high dimensional data (Gao and Ganguli 2015).

Perhaps the most useful and powerful tool we currently have to assess the relationship between neural activity and behavior observed is through mutual information (Shannon 1948). In lay terms, mutual information can be thought of as the measure of the relatedness between two variables similar to the correlation. While the correlation measures exclusively linear relationships between the variables, mutual information can pick up non-linear relationships as well. Measures of mutual information have been used to show that sensory neurons encode information not just in the rate at which they fire (referred to as rate coding) but also in the precise timing of the spikes they produce (Bialek et al. 1991, Strong et al. 1998, van Steveninck et al. 1997) which is referred to as temporal coding. Following these initial discoveries, evidence for temporal coding has been found in encoding stimulus location (Panzeri et al. 2001) and more recently, in the motor cortex like areas in songbirds (Tang et al. 2014) as well as in actual muscle control of respiration (Srivastava et al. 2017) and flight control in moths (Conn, Putney, and Sponberg 2019). Such advances have required the motor studies field to rethink long-held assumptions regarding coding of neural activity to produce movements (Sober et al. 2018). Additionally, these advances point to new lines of inquiry for the role of dopamine in sensorimotor control. If motor output is controlled by the precise timing of spikes from motor units, is the dopamine signal that modulates and corrects

the output also correspondingly precise in time? If it is precise in time, how do the receptors involved convey precisely timed signals when the dynamics of activation of the receptor and subsequent downstream processes (dopamine receptors are metabotropic receptors) are much slower than that of action potentials? Does co-release of glutamate help navigate this since glutamate can activate ionotropic receptors which in turn act much faster? There has been some evidence for co-release of glutamate and dopamine particularly in the ventral striatum (Stuber et al. 2010, Fortin et al. 2012).

As powerful as information theory is, it still lacks the theoretical foundations for one case that is abundantly common in neuroscience, namely, how does one compute the amount of information being transmitted between two points where both points are simultaneously receivers and transmitters? If one considers neurons as the points in the above example, neurons receive inputs from hundreds to thousands of neurons while themselves projecting to hundreds to thousands more. Hence these neurons simultaneously receive and transmit information. Statistical tools have not been developed to handle such a complex network with such extensive bidirectional information exchange. This is the other major methodological hurdle to understanding the workings of neural communication.

There is one more often overlooked aspect of neural communication. The action potential is the most prominent and recognizable form of output signal of a neuron and so, has received the most attention and research into understanding its role. However, in focusing on the relationship between action potentials and behavioral outputs, less emphasis has been placed on the elaborate dendritic trees and the numerous synapses that are provided as inputs to each neuron with their own action potentials. There is evidence that the location of synapses on dendrites is much more precise than previously acknowledged (Bloss et al. 2016) and can cause complex relationships

such as direction selectivity in starburst amacrine cells (Vlasits et al. 2016) or modulate the input-output gain of a neuron directly suggesting that studies examining action potentials arising from current injection into the soma may be missing crucial details regarding network function (Jarvis, Nikolic, and Schultz 2018). Given these factors, it becomes essential to either identify underlying constraints that are common among dendritic trees and synaptic placement such that inputs to all neurons can be appropriately adjusted or map out connections and take them into account when performing network level analysis relating brain activity to behavior.

4.4 Future Experimental Proposals

I have spent much of the previous sections in this Chapter mentioning major challenges to enhancing our understanding of the brain and central nervous system and proposed some experiments and lines of work that could help address these gaps in knowledge. In this subsection, I will focus on particular experiments that my experimental results will directly benefit from and how they may be performed.

As I briefly mentioned in section 2.5 as part of Chapter 2, a plausible hypothesis from our data is that dopamine is involved in actively elevating the pitch of syllables sung by birds and is required to produce higher pitched renditions of particular syllables. This could be tested both by dopamine antagonist infusion as well as direct dopamine or L-DOPA infusion to increase the concentration of dopamine in Area X. A lot of emphasis has been placed on effects of dopamine depletion on learning and such but less is known about dopamine infusions. This is particularly relevant for ties to interventions in patients affected by Parkinson's disease since the first order of treatment is that of L-DOPA infusions.

A further direction that ultimately proved beyond the scope of this dissertation was to build a computational model that details a mathematical role for dopamine in sensorimotor learning. A

collaborator, Zhou and colleagues were able to successfully model unmanipulated songbird behavior in response to pitch-shifted auditory feedback as was performed in Chapter 2 using Bayesian Inference (Zhou et al. 2018). While it did not shed light on the role of dopamine in the computation, it did provide insights into how songbirds reconciled conflicting sensory feedback from multiple sources and integrated them optimally to produce a successful strategy for change. I tried to adapt this model to fit the data from the dopamine depleted birds in Chapter 2 but was unsuccessful since a Bayesian filter model necessarily has to change in the adaptive direction. As was shown, all the birds reduced their pitch over time regardless of the direction of pitch shift through the headphones. While this could have been explicitly built-in into the model, the lack of any other consistent trends in the data precluded the possibility of testing any such model for accuracy. Therefore, more data for additional constraints will be required before a successful model incorporating a role for dopamine in the computation can be built.

Enhancing the timescale over which dopamine is manipulated or recording the neural activity from the circuits of interest are among the types of additional data that could help address this problem. For instance, future work in our lab involves the use of fiber photometry (see section 4.2 above) to record the activity of dopaminergic terminals in Area X during singing and other learning behavior. Such experiments could provide data on the types of signals provided during song and how they change with perceived sensory errors. By combining them with a dopamine depletion paradigm, we could also study how the signal changes under depleted conditions shedding light on how dopamine controls an upward shift in pitch of the produced syllables. Additionally, the use of DREADDs to control the uptake and activity of dopamine receptors could help understand the role dopamine plays in learning at a functionally relevant timescale.

As was eluded to briefly in Chapter 2, there is also some controversy surrounding the nature of the white noise blasts used to drive reinforcement-type learning in songbirds. While some labs drive the pitch-contingent learning by using white noise blasts which are referred to as aversive (Tumer and Brainard 2007, Hoffmann et al. 2016), other labs use a version of a different syllable played at the same time and refer to such feedback as distorted auditory feedback (Andalman and Fee 2009, Gadagkar et al. 2016). Furthermore, it was shown that when birds are not singing, if a perch is triggered to play back such distorted auditory feedback when the bird lands and another perch is not, the birds seem to preferentially land on the perch that triggers the feedback suggesting that they may not find the feedback itself aversive (Murdoch, Chen, and Goldberg 2018). Another line of research in our lab is focused on using a non-auditory aversive cue to see if non-auditory cues may also drive pitch-contingent vocal learning. This cue is an electric stimulation delivered to the back of the neck muscles and is hence more clearly aversive. The discovery of similar neural circuits underlying such learning to both auditory and non-auditory cues would suggest that the cues may have been aversive in a singing context and shed light on how birds perform reinforcement learning. On the other hand, separate neural mechanisms for such learning may reveal differences in the brain for how auditory versus non-auditory cues are processed and the parallel systems available for effecting the same changes in motor output, namely a change in pitch.

Future experiments will also be required to dissociate motor production effects from motor learning effects either in songbirds or other species. As has been made clear through decades of research in patients with Parkinson's disease, when patients exhibit severe motor deficits, deficits in motor learning can be extremely hard to identify let alone study, understand and eventually, treat. In the experiments conducted in Chapter 2, we once again observed concurrent effects on

motor production, namely lowering of the pitch sung on average, and on motor learning, namely the failure to change their pitch in an adaptive manner to the induced pitch shift. Future experiments could aim to dissociate these effects such as by waiting for the pitch to stabilize following dopamine depletion and then beginning the pitch shift experiments. In section 2.5, we argued that the motor production effect we observed could be attributed to a lack in motor vigor following dopamine depletion where vigor could mean motivation or speed of movement or both. As mentioned in section 4.1, a recent study found that while the amount of dopamine released in the nucleus accumbens correlated with a rodent's state of motivation, there were no appreciable differences in the firing rate of dopaminergic neurons projecting there to cause such an increase (Mohebi et al. 2019). This suggests that while dopamine is involved in vigor or motivation, an external circuit may be responsible for modulating the level of dopamine and therefore the vigor of the subject. Future experiments could be aimed at isolating the circuit responsible for the modulation of dopamine in the nucleus accumbens and once identified, testing its direct relationship to the control of motivation or vigor.

Finally, as I briefly elaborated on in section 4.3, a major push going forward has to be centered on both discovering better mathematical tools to analyze the vast amounts of data that will be gathered by the new techniques being developed and developing mathematical theories underlying the working of the brain that can both explain past results observed and make predictions for how the system might behave in different circumstances which can then be tested through experiments. As post-Einstein physicists are acknowledging, there may not be a "Theory of Everything" or TOE that unifies all 4 types of forces observed in the physical world, namely gravitational, electromagnetic, strong and weak forces. Even the Grand Unified Theory has only been theorized but not observed in practice. However, the complex field of electromagnetism could

be adequately summarized into 4 equations through the seminal work of James Maxwell (Maxwell 1865). Similarly, it is possible that we may not achieve a unified theory of the working of the brain. But that should not prevent us from attempting to summarize the vast amounts of experimental data obtained thus far.

4.5 Statistics and the Bootstrap

As I have shown in Chapter 3, a large fraction of datasets in neuroscience are hierarchical and with the advent of new techniques such as those mentioned in section 4.2, the datasets are about to become much larger and even more nested than they currently are. In this time, particularly with the reproducibility crisis and propagation of statistical misnomers (such as misnomers regarding the p-value), a comprehensive statistical training is required across the board for trainees and current researchers. While a robust statistical training program for new trainees may potentially create a new generation that knows how to perform statistical tests correctly, without concurrent training for current researchers this effort is bound to fail. This is because regardless of classes and other outside knowledge trainees bring to labs that they grow in, statistical habits of their advisors will be the ones they receive the most exposure to and are most likely to emulate. Additionally, due to the system of blinded peer-review, even if new researchers attempt to publish methods with updated statistical procedures, the reviewers must have the knowledge to assess the method's appropriateness for the data and either approve or recommend changes correspondingly. Hence one of the most immediate requirements to address this issue is an academy-wide retraining on the use of appropriate statistical tests for the corresponding types of data.

As was shown in Chapter 3, when traditional statistical methods that assume independence between data points are applied to hierarchical datasets where the data are not independent, the false positive rate is much higher than the acceptable Type-I error rate that was set. As a result, the

onus falls on future researchers to verify whether a given study in the past was performed correctly and whether their results have a large enough effect size to be detected when the appropriate statistical tests are applied. As we saw in Chapter 3 though, the larger the effect size, the more likely the result is to be statistically significant even with the adjusted tests.

In Chapter 3, we introduced the hierarchical bootstrap as a potential alternative to traditional statistical tests that could handle hierarchical datasets and their non-independence between data points appropriately. However, we only addressed the problems of bias and other issues that have been raised with the bootstrap briefly (Hillis and Bull 1993, Crowley 1992). Further analysis of bias and other potential issues should be explored and alternatives suggested if the bootstrap is found to have significant problems in handling hierarchical data. Additionally, we presented one method of calculating statistics using the bootstrap but other methods should be explored and their relative strengths and weaknesses assessed. Finally, what makes the bootstrap as powerful as it is popular is that it is a very simple algorithm to explain and understand. Regardless, the use of the bootstrap will increase greatly with the advent of scripts or statistical packages that can perform the hierarchical bootstrapping as needed for a given dataset automatically as opposed to having to manually write or adapt code by oneself for their particular dataset.

In conclusion, the experiments in this dissertation revealed a role for dopamine in sensorimotor adaptation. Further experiments will be required to understand how dopamine impacts sensorimotor adaptation and across what timescales. I also showed that when dealing with hierarchical datasets, the non-independence between data points makes traditional statistical tests that assume independence between data points return extremely large false positive rates. I showed how using hierarchical bootstrapping can help overcome this problem and showed two real world

examples of applying hierarchical bootstrapping to previously published datasets, one of which the results did not change while for the other the results did. I propose that for future experiments involving hierarchical datasets, the bootstrap or other appropriate statistical tests such as LMMs ought to be deployed from the outset.

5 CHAPTER V: REFERENCES

- Aarts, E., C. V. Dolan, M. Verhage, and S. van der Sluis. 2015. "Multilevel analysis quantifies variation in the experimental effect while optimizing power and preventing false positives." *BMC Neurosci* 16:94. doi: 10.1186/s12868-015-0228-5.
- Aarts, E., M. Verhage, J. V. Veenvliet, C. V. Dolan, and S. van der Sluis. 2014. "A solution to dependency: using multilevel analysis to accommodate nested data." *Nat Neurosci* 17 (4):491-6. doi: 10.1038/nn.3648.
- Abbeel, Pieter, and Andrew Y Ng. 2004. "Apprenticeship learning via inverse reinforcement learning." Proceedings of the twenty-first international conference on Machine learning.
- Abur, Defne, Rosemary A Lester-Smith, Ayoub Daliri, Ashling A Lupiani, Frank H Guenther, and Cara E Stepp. 2018. "Sensorimotor adaptation of voice fundamental frequency in Parkinson's disease." *PloS one* 13 (1):e0191839.
- Adams, Dean C, Jessica Gurevitch, and Michael S Rosenberg. 1997. "Resampling tests for meta-analysis of ecological data." *Ecology* 78 (4):1277-1283.
- Ali, F., T. M. Otchy, C. Pehlevan, A. L. Fantana, Y. Burak, and B. P. Olveczky. 2013. "The basal ganglia is necessary for learning spectral, but not temporal, features of birdsong." *Neuron* 80 (2):494-506. doi: 10.1016/j.neuron.2013.07.049.
- Andalman, A. S., and M. S. Fee. 2009. "A basal ganglia-forebrain circuit in the songbird biases motor output to avoid vocal errors." *Proc Natl Acad Sci U S A* 106 (30):12518-23. doi: 10.1073/pnas.0903214106.
- Andén, Nils-Erik, Annica Dahlström, Kjell Fuxe, and Knut Larsson. 1966. "Functional Role of the Nigro-Neostriatal Dopamine Neurons*." *Acta pharmacologica et toxicologica* 24 (2-3):263-274.
- Anderson, Marti, and Cajo Ter Braak. 2003. "Permutation tests for multi-factorial analysis of variance." *Journal of statistical computation and simulation* 73 (2):85-113.
- Anderson, Marti J. 2001. "A new method for non-parametric multivariate analysis of variance." 26 (1):32-46. doi: doi:10.1111/j.1442-9993.2001.01070.pp.x.
- Arceneaux, Kevin, and David W Nickerson. 2009. "Modeling certainty with clustered data: A comparison of methods." *Political Analysis* 17 (2):177-190.
- Arlet, Malgorzata, Ronan Jubin, Nobuo Masataka, and Alban %J Biology letters Lemasson. 2015. "Grooming-at-a-distance by exchanging calls in non-human primates." 11 (10):20150711.
- Armbruster, Blaine N, Xiang Li, Mark H Pausch, Stefan Herlitze, and Bryan L Roth. 2007. "Evolving the lock to fit the key to create a family of G protein-coupled receptors potently activated by an inert ligand." *Proceedings of the National Academy of Sciences* 104 (12):5163-5168.
- Baizer, Joan S, Ines Kralj-Hans, and Mitchell Glickstein. 1999. "Cerebellar lesions and prism adaptation in macaque monkeys." *Journal of neurophysiology* 81 (4):1960-1965.
- Baker, Monya. 2016. "1,500 scientists lift the lid on reproducibility." *Nature News* 533 (7604):452.
- Balleine, Bernard W, and John P O'doherty. 2010. "Human and rodent homologies in action control: corticostriatal determinants of goal-directed and habitual action." *Neuropsychopharmacology* 35 (1):48.
- Barbeau, Andre. 1962. "The pathogenesis of Parkinson's disease: a new hypothesis." *Canadian Medical Association Journal* 87 (15):802.

- Beierholm, U., M. Guitart-Masip, M. Economides, R. Chowdhury, E. Duzel, R. Dolan, and P. Dayan. 2013. "Dopamine modulates reward-related vigor." *Neuropsychopharmacology* 38 (8):1495-503. doi: 10.1038/npp.2013.48.
- Bellman, Richard. 1966. "Dynamic programming." *Science* 153 (3731):34-37.
- Beninger, R. J. 1983. "The role of dopamine in locomotor activity and learning." *Brain Res* 287 (2):173-96.
- Berke, J. D. 2018. "What does dopamine mean?" *Nat Neurosci* 21 (6):787-793. doi: 10.1038/s41593-018-0152-y.
- Berman, Gordon J, Daniel M Choi, William Bialek, and Joshua W Shaevitz. 2014. "Mapping the stereotyped behaviour of freely moving fruit flies." *Journal of The Royal Society Interface* 11 (99):20140672.
- Bernheimer, H, Wt Birkmayer, O Hornykiewicz, K Jellinger, and F Seitelberger. 1973. "Brain dopamine and the syndromes of Parkinson and Huntington Clinical, morphological and neurochemical correlations." *Journal of the neurological sciences* 20 (4):415-455.
- Berridge, Kent C. 2007. "The debate over dopamine's role in reward: the case for incentive salience." *Psychopharmacology* 191 (3):391-431.
- Berridge, Kent C, Isabel L Venier, and Terry E Robinson. 1989. "Taste reactivity analysis of 6-hydroxydopamine-induced aphagia: implications for arousal and anhedonia hypotheses of dopamine function." *Behavioral neuroscience* 103 (1):36.
- Berry, Donald A, and Bert Fristedt. 1985. "Bandit problems: sequential allocation of experiments (Monographs on statistics and applied probability)." *London: Chapman and Hall* 5:71-87.
- Bertsekas, Dimitri P. 1987. *Dynamic Programming: Determinist. and Stochast. Models*: Prentice-Hall.
- Bialek, William, Fred Rieke, RR De Ruyter Van Steveninck, and David Warland. 1991. "Reading a neural code." *Science* 252 (5014):1854-1857.
- Blakemore, Sarah-J, Chris D Frith, and Daniel M Wolpert. 2001. "The cerebellum is involved in predicting the sensory consequences of action." *Neuroreport* 12 (9):1879-1884.
- Bloss, Erik B, Mark S Cembrowski, Bill Karsh, Jennifer Colonell, Richard D Fetter, and Nelson Spruston. 2016. "Structured dendritic inhibition supports branch-selective integration in CA1 pyramidal cells." *Neuron* 89 (5):1016-1030.
- Bock, Otmar. 2005. "Components of sensorimotor adaptation in young and elderly subjects." *Experimental Brain Research* 160 (2):259-263.
- Bottjer, Sarah W %J Journal of neurobiology. 1993. "The distribution of tyrosine hydroxylase immunoreactivity in the brains of male and female zebra finches." 24 (1):51-69.
- Bottjer, Sarah W, Elizabeth A Miesner, and Arthur P Arnold. 1984. "Forebrain lesions disrupt development but not maintenance of song in passerine birds." *Science* 224 (4651):901-903.
- Boyden, Edward S, Feng Zhang, Ernst Bamberg, Georg Nagel, and Karl Deisseroth. 2005. "Millisecond-timescale, genetically targeted optical control of neural activity." *Nature Neuroscience* 8 (9):1263.
- Bradtke, Steven J, and Andrew G Barto. 1996. "Linear least-squares algorithms for temporal difference learning." *Machine learning* 22 (1-3):33-57.
- Brainard, M. S., and A. J. Doupe. 2000. "Interruption of a basal ganglia-forebrain circuit prevents plasticity of learned vocalizations." *Nature* 404 (6779):762-6. doi: 10.1038/35008083.
- Brashers-Krug, Thomas, Reza Shadmehr, and Emilio Bizzi. 1996. "Consolidation in human motor memory." *Nature* 382 (6588):252.

- Bu, Lucian, Robert Babu, and Bart De Schutter. 2008. "A comprehensive survey of multiagent reinforcement learning." *IEEE Transactions on Systems, Man, Cybernetics, Part C* 38 (2):156-172.
- Burns, Randal, Kunal Lillaney, Daniel R Berger, Logan Grosenick, Karl Deisseroth, R Clay Reid, William Gray Roncal, Priya Manavalan, Davi D Bock, and Narayanan Kasthuri. 2013. "The open connectome project data cluster: scalable analysis and vision for high-throughput neuroscience." Proceedings of the 25th International Conference on Scientific and Statistical Database Management.
- Caithness, Graham, Rieko Osu, Paul Bays, Henry Chase, Jessica Klassen, Mitsuo Kawato, Daniel M Wolpert, and J Randall Flanagan. 2004. "Failure to consolidate the consolidation theory of learning for sensorimotor adaptation tasks." *Journal of Neuroscience* 24 (40):8662-8671.
- Cande, Jessica, Shigehiro Namiki, Jirui Qiu, Wyatt Korff, Gwyneth M Card, Joshua W Shaevitz, David L Stern, and Gordon J Berman. 2018. "Optogenetic dissection of descending behavioral control in *Drosophila*." *Elife* 7:e34275.
- Cannon, Claire Matson, and Richard D Palmiter. 2003. "Reward without dopamine." *Journal of neuroscience* 23 (34):10827-10831.
- Canopoli, Alessandro, Joshua A Herbst, and Richard HR Hahnloser. 2014. "A higher sensory brain region is involved in reversing reinforcement-induced vocal changes in a songbird." *Journal of Neuroscience* 34 (20):7018-7026.
- Carpenter, James R, Harvey Goldstein, and Jon Rasbash. 2003. "A novel bootstrap procedure for assessing the relationship between class size and achievement." *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 52 (4):431-443.
- Carrillo, Gabriela D, and Allison J Doupe. 2004. "Is the songbird Area X striatal, pallidal, or both? An anatomical study." *Journal of Comparative Neurology* 473 (3):415-437.
- Cenci, M Angela, and Martin Lundblad. 2007. "Ratings of L-DOPA-induced dyskinesia in the unilateral 6-OHDA lesion model of Parkinson's disease in rats and mice." *Current protocols in Neuroscience* 41 (1):9.25. 1-9.25. 23.
- Chang, Chun Yun, Guillem R Esber, Yasmin Marrero-Garcia, Hau-Jie Yau, Antonello Bonci, and Geoffrey Schoenbaum. 2016. "Brief optogenetic inhibition of dopamine neurons mimics endogenous negative reward prediction errors." *Nature Neuroscience* 19 (1):111.
- Chen, Ruidong, Pavel A Puzerey, Andrea C Roeser, Tori E Riccelli, Archana Podury, Kamal Maher, Alexander Farhang, and Jesse Goldberg. 2018. "Songbird ventral basal ganglia sends performance error signals to dopaminergic midbrain." Available at SSRN 3231856.
- Claus, Caroline, and Craig Boutilier. 1998. "The dynamics of reinforcement learning in cooperative multiagent systems." *AAAI/IAAI* 1998:746-752.
- Coddington, Luke T, and Joshua T Dudman. 2018. "The timing of action determines reward prediction signals in identified midbrain dopamine neurons." *Nature neuroscience* 21 (11):1563.
- Coe, Robert. 2002. "It's the effect size, stupid." Paper presented at the British Educational Research Association annual conference.
- Cohen, Jeremiah Y, Sebastian Haesler, Linh Vong, Bradford B Lowell, and Naoshige Uchida. 2012. "Neuron-type-specific signals for reward and punishment in the ventral tegmental area." *Nature* 482 (7383):85.
- Conn, Rachel, Joy Putney, and Simon Sponberg. 2019. "Timing is (almost) everything in a comprehensive, spike-resolved flight motor program." *bioRxiv*:602961.

- Cook Jr, Edwin H, Mark A Stein, Matthew D Krasowski, Nancy J Cox, Deborah M Olkon, John E Kieffer, and Bennett L Leventhal. 1995. "Association of attention-deficit disorder and the dopamine transporter gene." *American journal of human genetics* 56 (4):993.
- Cooper, J. A., H. J. Sagar, N. Jordan, N. S. Harvey, and E. V. Sullivan. 1991. "Cognitive impairment in early, untreated Parkinson's disease and its relationship to motor disability." *Brain* 114 (Pt 5):2095-122.
- Crites, Robert H, and Andrew G Barto. 1996. "Improving elevator performance using reinforcement learning." *Advances in neural information processing systems*.
- Crowley, Philip H. 1992. "Resampling methods for computation-intensive data analysis in ecology and evolution." *Annual Review of Ecology and Systematics* 23 (1):405-447.
- Davison, Anthony Christopher, and David Victor Hinkley. 1997. *Bootstrap methods and their application*. Vol. 1: Cambridge university press.
- Dawson, Betty V, and BM Foss. 1965. "Observational learning in budgerigars." *Animal behaviour*.
- Dayan, Peter, and Kent C Berridge. 2014. "Model-based and model-free Pavlovian reward learning: revaluation, revision, and revelation." *Cognitive, Affective, & Behavioral Neuroscience* 14 (2):473-492.
- Dayan, Peter, and Yael Niv. 2008. "Reinforcement learning: the good, the bad and the ugly." *Current opinion in neurobiology* 18 (2):185-196.
- Dayan, Peter, and Terrence J Sejnowski. 1994. "TD (λ) converges with probability 1." *Machine Learning* 14 (3):295-301.
- de Wit, Sanne, Poppy Watson, Helga A Harsay, Michael X Cohen, Irene van de Vijver, and K Richard Ridderinkhof. 2012. "Corticostriatal connectivity underlies individual differences in the balance between habitual and goal-directed action control." *Journal of Neuroscience* 32 (35):12066-12075.
- Dickson, Dennis W. 2012. "Parkinson's disease and parkinsonism: neuropathology." *Cold Spring Harbor perspectives in medicine* 2 (8):a009258.
- Diedrichsen, Jörn, Yasmin Hashambhoy, Tushar Rane, and Reza Shadmehr. 2005. "Neural correlates of reach errors." *Journal of Neuroscience* 25 (43):9919-9931.
- Diehl, David J, and Samuel Gershon. 1992. "The role of dopamine in mood disorders." *Comprehensive psychiatry* 33 (2):115-120.
- Ding, Jun B, Jaime N Guzman, Jayms D Peterson, Joshua A Goldberg, and D James Surmeier. 2010. "Thalamic gating of corticostriatal signaling by cholinergic interneurons." *Neuron* 67 (2):294-307.
- Doupe, Allison J, and Patricia K Kuhl. 1999. "Birdsong and human speech: common themes and mechanisms." *Annual review of neuroscience* 22 (1):567-631.
- Dubois, Bruno, and Bernard %J Journal of neurology Pillon. 1996. "Cognitive deficits in Parkinson's disease." 244 (1):2-8.
- Edmonds, Jack. 1971. "Matroids and the greedy algorithm." *Mathematical programming* 1 (1):127-136.
- Efron, Bradley. 1981. "Nonparametric estimates of standard error: the jackknife, the bootstrap and other methods." *Biometrika* 68 (3):589-599.
- Efron, Bradley. 1992. "Bootstrap methods: another look at the jackknife." In *Breakthroughs in statistics*, 569-593. Springer.

- Efron, Bradley, Elizabeth Halloran, and Susan Holmes. 1996. "Bootstrap confidence levels for phylogenetic trees." *Proceedings of the National Academy of Sciences* 93 (23):13429-13429.
- Efron, Bradley, and Robert J Tibshirani. 1994. *An introduction to the bootstrap*: CRC press.
- Elemans, C. P., A. F. Mead, L. C. Rome, and F. Goller. 2008. "Superfast vocal muscles control song production in songbirds." *PLoS One* 3 (7):e2581. doi: 10.1371/journal.pone.0002581.
- Elemans, CPH, Jeppe Have Rasmussen, Christian T Herbst, Daniel Normen Düring, Sue Anne Zollinger, Henrik Brumm, K Srivastava, Niels Svane, Ming Ding, and Ole Næsbye Larsen. 2015. "Universal mechanisms of sound production and control in birds and mammals." *Nature communications* 6:8978.
- English, Daniel F, Osvaldo Ibanez-Sandoval, Eran Stark, Fatuel Tecuapetla, György Buzsáki, Karl Deisseroth, James M Tepper, and Tibor Koos. 2012. "GABAergic circuits mediate the reinforcement-related signals of striatal cholinergic interneurons." *Nature neuroscience* 15 (1):123.
- Erev, Ido, and Alvin E Roth. 1998. "Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria." *American economic review*:848-881.
- Eshel, Neir, Ju Tian, Michael Bukwich, and Naoshige Uchida. 2016. "Dopamine neurons share common response function for reward prediction error." *Nature neuroscience* 19 (3):479.
- Ethier, Vincent, David S Zee, and Reza Shadmehr. 2008a. "Changes in control of saccades during gain adaptation." *Journal of Neuroscience* 28 (51):13929-13937.
- Ethier, Vincent, David S Zee, and Reza Shadmehr. 2008b. "Spontaneous recovery of motor memory during saccade adaptation." *Journal of Neurophysiology*.
- Farries, Michael A, Long Ding, and David J Perkel. 2005. "Evidence for "direct" and "indirect" pathways through the song system basal ganglia." *Journal of Comparative Neurology* 484 (1):93-104.
- Feenders, Gesa, Miriam Liedvogel, Miriam Rivas, Manuela Zapka, Haruhito Horita, Erina Hara, Kazuhiro Wada, Henrik Mouritsen, and Erich D Jarvis. 2008. "Molecular mapping of movement-associated areas in the avian brain: a motor theory for vocal learning origin." *PLoS One* 3 (3):e1768.
- Felsenstein, Joseph, and Hirohisa Kishino. 1993. "Is there something wrong with the bootstrap on phylogenies? A reply to Hillis and Bull." *Systematic Biology* 42 (2):193-200.
- Field, Christopher A, and Alan H Welsh. 2007. "Bootstrapping clustered data." *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 69 (3):369-390.
- Findlay, John M, and Robin Walker. 1999. "A model of saccade generation based on parallel processing and competitive inhibition." *Behavioral & Brain Sciences* 22 (4):661-674.
- Fiorillo, Christopher D. 2013. "Two dimensions of value: dopamine neurons represent reward but not aversiveness." *Science* 341 (6145):546-549.
- Fiorillo, Christopher D, Minryung R Song, and Sora R Yun. 2013. "Multiphasic temporal dynamics in responses of midbrain dopamine neurons to appetitive and aversive stimuli." *Journal of Neuroscience* 33 (11):4710-4725.
- Fiorillo, Christopher D, Philippe N Tobler, and Wolfram Schultz. 2003. "Discrete coding of reward probability and uncertainty by dopamine neurons." *Science* 299 (5614):1898-1902.
- Fiorito, Graziano, and Pietro Scotto. 1992. "Observational learning in *Octopus vulgaris*." *Science* 256 (5056):545-547.

- Fischer, Burkhardt, and Heike Weber. 1993. "Express saccades and visual attention." *Behavioral & Brain Sciences* 16 (3):553-567.
- Fischer, Burkhardt, Heike Weber, Monica Biscaldi, Franz Aiple, Petra Otto, and Volker Stuhr. 1993. "Separate populations of visually guided saccades in humans: reaction times and amplitudes." *Experimental Brain Research* 92 (3):528-541.
- Foley Jr, John P. 1940. "An experimental investigation of the effect of prolonged inversion of the visual field in the rhesus monkey (*Macaca mulatta*)." *The Pedagogical Seminary & Journal of Genetic Psychology* 56 (1):21-51.
- Fortin, Guillaume M, Marie-Josée Bourque, Jose Alfredo Mendez, Damiana Leo, Karin Nordenankar, Carolina Birgner, Emma Arvidsson, Vladimir V Rymar, Noémie Bérubé-Carriere, and Anne-Marie Claveau. 2012. "Glutamate corelease promotes growth and survival of midbrain dopamine neurons." *Journal of Neuroscience* 32 (48):17477-17491.
- Gadagkar, V., P. A. Puzerey, R. Chen, E. Baird-Daniel, A. R. Farhang, and J. H. Goldberg. 2016. "Dopamine neurons encode performance error in singing birds." *Science* 354 (6317):1278-1282. doi: 10.1126/science.aah6837.
- Galbraith, S., J. A. Daniel, and B. Vissel. 2010. "A study of clustered data and approaches to its analysis." *J Neurosci* 30 (32):10601-8. doi: 10.1523/JNEUROSCI.0362-10.2010.
- Ganguli, Surya, and Haim Sompolinsky. 2012. "Compressed sensing, sparsity, and dimensionality in neuronal information processing and data analysis." *Annual review of neuroscience* 35:485-508.
- Gao, Peiran, and Surya Ganguli. 2015. "On simplicity and complexity in the brave new world of large-scale neuroscience." *Current opinion in neurobiology* 32:148-155.
- Gardner, Matthew PH, Geoffrey Schoenbaum, and Samuel J Gershman. 2018. "Rethinking dopamine as generalized prediction error." *Proceedings of the Royal Society B* 285 (1891):20181645.
- Garland Jr, Theodore, Allan W Dickerman, Christine M Janis, and Jason A Jones. 1993. "Phylogenetic analysis of covariance by computer simulation." *Systematic Biology* 42 (3):265-292.
- Garland Jr, Theodore, Peter E Midford, and Anthony R Ives. 1999. "An introduction to phylogenetically based statistical methods, with a new method for confidence intervals on ancestral values." *American Zoologist* 39 (2):374-388.
- Gaveau, Valérie, Olivier Martin, Claude Prablanc, Denis Péllisson, Christian Urquizar, and Michel Desmurget. 2003. "On-line modification of saccadic eye movements by retinal signals." *Neuroreport* 14 (6):875-878.
- Gehlbach, Hunter, Maureen E Brinkworth, Aaron M King, Laura M Hsu, Joseph McIntyre, and Todd Rogers. 2016. "Creating birds of similar feathers: Leveraging similarity to improve teacher–student relationships and academic achievement." *Journal of Educational Psychology* 108 (3):342.
- Georgopoulos, Apostolos P, Roberto Caminiti, John F Kalaska, and Joseph T Massey. 1983. "Spatial coding of movement: a hypothesis concerning the coding of movement direction by motor cortical populations." *Experimental Brain Research* 7 (32):336.
- Georgopoulos, Apostolos P, John F Kalaska, Roberto Caminiti, and Joe T Massey. 1982. "On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex." *Journal of Neuroscience* 2 (11):1527-1537.

- Georgopoulos, Apostolos P, Andrew B Schwartz, and Ronald E Kettner. 1986. "Neuronal population coding of movement direction." *Science* 233 (4771):1416-1419.
- Gerlai, Robert. 2019. "Reproducibility and replicability in zebrafish behavioral neuroscience research." *Pharmacology Biochemistry and Behavior* 178:30-38.
- Gibson, James J. 1933. "Adaptation, after-effect and contrast in the perception of curved lines." *Journal of experimental psychology* 16 (1):1.
- Gibson, James J. 1937. "Adaptation, after-effect, and contrast in the perception of tilted lines. II. Simultaneous contrast and the areal restriction of the after-effect." *Journal of Experimental Psychology* 20 (6):553.
- Gibson, James J, and Minnie Radner. 1937. "Adaptation, after-effect and contrast in the perception of tilted lines. I. Quantitative studies." *Journal of experimental psychology* 20 (5):453.
- Gingrich, Jay A, and Marc G Caron. 1993. "Recent advances in the molecular biology of dopamine receptors." *Annual review of neuroscience* 16 (1):299-321.
- Glimcher, P. W. 2011. "Understanding dopamine and reinforcement learning: the dopamine reward prediction error hypothesis." *Proc Natl Acad Sci U S A* 108 Suppl 3:15647-54. doi: 10.1073/pnas.1014269108.
- Gold, Joshua I, and Michael N Shadlen. 2007. "The neural basis of decision making." *Annu. Rev. Neurosci.* 30:535-574.
- Goldstein, Harvey. 2011. "Bootstrapping in multilevel models." *Handbook of advanced multilevel analysis*:163-171.
- Goller, F., and R. A. Suthers. 1996. "Role of syringeal muscles in controlling the phonology of bird song." *J Neurophysiol* 76 (1):287-300. doi: 10.1152/jn.1996.76.1.287.
- Grace, A. A. 1991. "Phasic versus tonic dopamine release and the modulation of dopamine system responsivity: a hypothesis for the etiology of schizophrenia." *Neuroscience* 41 (1):1-24.
- Grace, Anthony A. 2016. "Dysregulation of the dopamine system in the pathophysiology of schizophrenia and depression." *Nature Reviews Neuroscience* 17 (8):524.
- Greenland, Sander, Stephen J Senn, Kenneth J Rothman, John B Carlin, Charles Poole, Steven N Goodman, and Douglas G Altman. 2016. "Statistical tests, P values, confidence intervals, and power: a guide to misinterpretations." *European journal of epidemiology* 31 (4):337-350.
- Guenther, Frank H. 2006. "Cortical interactions underlying the production of speech sounds." *Journal of communication disorders* 39 (5):350-365.
- Guo, Qingchun, Jingfeng Zhou, Qiru Feng, Rui Lin, Hui Gong, Qingming Luo, Shaoqun Zeng, Minmin Luo, and Ling Fu. 2015. "Multi-channel fiber photometry for population neuronal activity recording." *Biomedical optics express* 6 (10):3919-3931.
- Guthrie, Barton L, John D Porter, and David L Sparks. 1983. "Corollary discharge provides accurate eye position information to the oculomotor system." *Science* 221 (4616):1193-1195.
- Hahnloser, R. H., and G. Narula. 2017. "A Bayesian Account of Vocal Adaptation to Pitch-Shifted Auditory Feedback." *PLoS One* 12 (1):e0169795. doi: 10.1371/journal.pone.0169795.
- Hahs-Vaughn, Debbie L. 2005. "A primer for using and understanding weights with national datasets." *The Journal of Experimental Education* 73 (3):221-248.
- Haith, A. M., and J. W. Krakauer. 2013. "Model-based and model-free mechanisms of human motor learning." *Adv Exp Med Biol* 782:1-21. doi: 10.1007/978-1-4614-5465-6_1.
- Halsey, Lewis G, Douglas Curran-Everett, Sarah L Vowler, and Gordon B Drummond. 2015. "The fickle P value generates irreproducible results." *Nature methods* 12 (3):179.

- Hamilton, Charles R. 1964. "Intermanual transfer of adaptation to prisms." *The American Journal of Psychology*.
- Harden, Jeffrey J. 2011. "A bootstrap method for conducting statistical inference with clustered data." *State Politics & Policy Quarterly* 11 (2):223-246.
- Harley, Calvin B. 1981. "Learning the evolutionarily stable strategy." *Journal of theoretical biology* 89 (4):611-633.
- Hart, Andrew S, Robb B Rutledge, Paul W Glimcher, and Paul EM Phillips. 2014. "Phasic dopamine release in the rat nucleus accumbens symmetrically encodes a reward prediction error term." *Journal of Neuroscience* 34 (3):698-704.
- Heffner, Robert A, Mark J Butler, and Colleen Keelan Reilly. 1996. "Pseudoreplication revisited." *Ecology* 77 (8):2558-2562.
- Held, Richard, and Joseph Bossom. 1961. "Neonatal deprivation and adult rearrangement: complementary techniques for analyzing plastic sensory-motor coordinations." *Journal of comparative and physiological Psychology* 54 (1):33.
- Held, Richard, and Norman Gottlieb. 1958. "Technique for studying adaptation to disarranged hand-eye coordination." *Perceptual and Motor Skills* 8 (3):83-86.
- Hernandez, Luis, and Bartley G Hoebel. 1988. "Food reward and cocaine increase extracellular dopamine in the nucleus accumbens as measured by microdialysis." *Life sciences* 42 (18):1705-1712.
- Hillis, David M, and James J Bull. 1993. "An empirical test of bootstrapping as a method for assessing confidence in phylogenetic analysis." *Systematic biology* 42 (2):182-192.
- Hinton, Geoffrey, Li Deng, Dong Yu, George Dahl, Abdel-rahman Mohamed, Navdeep Jaitly, Andrew Senior, Vincent Vanhoucke, Patrick Nguyen, and Brian Kingsbury. 2012. "Deep neural networks for acoustic modeling in speech recognition." *IEEE Signal processing magazine* 29.
- Hisey, E., M. G. Kearney, and R. Mooney. 2018. "A common neural circuit mechanism for internally guided and externally reinforced forms of motor learning." *Nat Neurosci* 21 (4):589-597. doi: 10.1038/s41593-018-0092-6.
- Hnasko, Thomas S, Bethany N Sotak, and Richard D Palmiter. 2005. "Morphine reward in dopamine-deficient mice." *Nature* 438 (7069):854.
- Hoffmann, L. A., C. W. Kelly, D. A. Nicholson, and S. J. Sober. 2012. "A lightweight, headphones-based system for manipulating auditory feedback in songbirds." *J Vis Exp* (69):e50027. doi: 10.3791/50027.
- Hoffmann, L. A., V. Saravanan, A. N. Wood, L. He, and S. J. Sober. 2016. "Dopaminergic Contributions to Vocal Learning." *J Neurosci* 36 (7):2176-89. doi: 10.1523/JNEUROSCI.3883-15.2016.
- Hoffmann, L. A., and S. J. Sober. 2014. "Vocal generalization depends on gesture identity and sequence." *J Neurosci* 34 (16):5564-74. doi: 10.1523/JNEUROSCI.5169-13.2014.
- Hosp, Jonas A, Ana Pekanovic, Mengia S Rioult-Pedotti, and Andreas R Luft. 2011. "Dopaminergic projections from midbrain to primary motor cortex mediate motor skill learning." *Journal of Neuroscience* 31 (7):2481-2487.
- Hosp, Jonas Aurel, and Andreas R Luft. 2013. "Dopaminergic meso-cortical projections to M1: role in motor learning and motor cortex plasticity." *Frontiers in neurology* 4:145.
- Houde, John F, and Michael I Jordan. 1998. "Sensorimotor adaptation in speech production." *Science* 279 (5354):1213-1216.

- Houde, John F, and Michael I Jordan. 2002. "Sensorimotor Adaptation of Speech I." *Journal of Speech, Language, and Hearing Research*.
- Howard, Ronald A. 1960. "Dynamic programming and markov processes."
- Howe, MW, and Daniel A Dombeck. 2016. "Rapid signalling in distinct dopaminergic axons during locomotion and reward." *Nature* 535 (7613):505.
- Howe, William M, Anne S Berry, Jennifer Francois, Gary Gilmour, Joshua M Carp, Mark Tricklebank, Cindy Lustig, and Martin %J Journal of Neuroscience Sarter. 2013. "Prefrontal cholinergic mechanisms instigating shifts from monitoring for cues to cue-guided performance: converging electrochemical and fMRI evidence from rats and humans." 33 (20):8742-8752.
- Howes, Oliver, Rob McCutcheon, and James Stone. 2015. "Glutamate and dopamine in schizophrenia: an update for the 21st century." *Journal of psychopharmacology* 29 (2):97-115.
- Hox, Joop J, Mirjam Moerbeek, and Rens Van de Schoot. 2017. *Multilevel analysis: Techniques and applications*: Routledge.
- Hu, Junling, and Michael P Wellman. 1998. "Multiagent reinforcement learning: theoretical framework and an algorithm." ICML.
- Huang, Francis L. 2018. "Using cluster bootstrapping to analyze nested data with a few clusters." *Educational and Psychological Measurement* 78 (2):297-318.
- Huang, Jian, Zhaomin Zhong, Mingyong Wang, Xifeng Chen, Yicheng Tan, Shuqing Zhang, Wei He, Xiong He, Guodong Huang, and Haiping Lu. 2015. "Circadian modulation of dopamine levels and dopaminergic neuron development contributes to attention deficiency and hyperactive behavior." *Journal of Neuroscience* 35 (6):2572-2587.
- Huffman, Megan L, and B Jill Venton. 2009. "Carbon-fiber microelectrodes for in vivo applications." *Analyst* 134 (1):18-24.
- Hurlbert, Stuart H. 1984. "Pseudoreplication and the design of ecological field experiments." *Ecological monographs* 54 (2):187-211.
- Iancu, Ruxandra, Paul Mohapel, Patrik Brundin, and Gesine Paul. 2005. "Behavioral characterization of a unilateral 6-OHDA-lesion model of Parkinson's disease in mice." *Behavioural brain research* 162 (1):1-10.
- Inoue, Masatoshi, Atsuya Takeuchi, Satoshi Manita, Shin-ichiro Horigane, Masayuki Sakamoto, Ryosuke Kawakami, Kazushi Yamaguchi, Kouhei Otomo, Hiroyuki Yokoyama, and Ryang Kim. 2019. "Rational Engineering of XCaMPs, a Multicolor GECI Suite for In Vivo Imaging of Complex Brain Circuit Dynamics." *Cell*.
- Ioannidis, John PA. 2005. "Why most published research findings are false." *PLoS medicine* 2 (8):e124.
- Ito, Masao. 2008. "Control of mental activities by internal models in the cerebellum." *Nature Reviews Neuroscience* 9 (4):304.
- Jankovic, Joseph. 2008. "Parkinson's disease: clinical features and diagnosis." *Journal of neurology, neurosurgery & psychiatry* 79 (4):368-376.
- Jarkowski, Jean. 1925. *'Kinésie paradoxale' des parkinsoniens*: Masson.
- Jarvis, Sarah, Konstantin Nikolic, and Simon R Schultz. 2018. "Neuronal gain modulability is determined by dendritic morphology: a computational optogenetic study." *PLoS computational biology* 14 (3):e1006027.
- Jellinger, Kurt A. 1991. "Pathology of Parkinson's disease." *Molecular and Chemical Neuropathology* 14 (3):153-197.

- Jeon, B. S., V. Jackson-Lewis, and R. E. Burke. 1995. "6-Hydroxydopamine lesion of the rat substantia nigra: time course and morphology of cell death." *Neurodegeneration* 4 (2):131-7.
- Jin, Xin, and Rui M Costa. 2010. "Start/stop signals emerge in nigrostriatal circuits during sequence learning." *Nature* 466 (7305):457.
- Johnson, Paul M, and Paul J Kenny. 2010. "Dopamine D2 receptors in addiction-like reward dysfunction and compulsive eating in obese rats." *Nature neuroscience* 13 (5):635.
- Joshua, Mati, Avital Adler, Rea Mitelman, Eilon Vaadia, and Hagai Bergman. 2008. "Midbrain dopaminergic neurons and striatal cholinergic interneurons encode the difference between reward and aversive events at different epochs of probabilistic classical conditioning trials." *Journal of Neuroscience* 28 (45):11673-11684.
- Kaelbling, Leslie Pack, Michael L Littman, and Andrew W Moore. 1996. "Reinforcement learning: A survey." *Journal of artificial intelligence research* 4:237-285.
- Kebabian, John W, and Donald B Calne. 1979. "Multiple receptors for dopamine." *Nature* 277 (5692):93.
- Kelly, C. W., and S. J. Sober. 2014. "A simple computational principle predicts vocal adaptation dynamics across age and error size." *Front Integr Neurosci* 8:75. doi: 10.3389/fnint.2014.00075.
- Keramati, Mehdi, Amir Dezfouli, and Payam Piray. 2011. "Speed/accuracy trade-off between the habitual and the goal-directed processes." *PLoS computational biology* 7 (5):e1002055.
- Kish, Leslie. 1965. "Survey sampling."
- Koh, Pang Wei, and Percy Liang. 2017. "Understanding black-box predictions via influence functions." Proceedings of the 34th International Conference on Machine Learning-Volume 70.
- Körding, Konrad P, and Daniel M Wolpert. 2006. "Bayesian decision theory in sensorimotor control." *Trends in cognitive sciences* 10 (7):319-326.
- Korshunov, Kirill S, Laura J Blakemore, and Paul Q Trombley. 2017. "Dopamine: a modulator of circadian rhythms in the central nervous system." *Frontiers in cellular neuroscience* 11:91.
- Krakauer, John W, and Pietro Mazzoni. 2011. "Human sensorimotor learning: adaptation, skill, and beyond." *Current opinion in neurobiology* 21 (4):636-644.
- Kravitz, Alexxai V, and Anatol C Kreitzer. 2012. "Striatal mechanisms underlying movement, reinforcement, and punishment." *Physiology* 27 (3):167-177.
- Kravitz, Alexxai V, Lynne D Tye, and Anatol C Kreitzer. 2012. "Distinct roles for direct and indirect pathway striatal neurons in reinforcement." *Nature Neuroscience* 15 (6):816.
- Kremer, Yves, Jérôme Flakowski, Clément Rohner, and Christian Lüscher. 2018. "VTA dopamine neurons multiplex external with internal representations of goal-directed action." *bioRxiv*:408062.
- Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E Hinton. 2012. "Imagenet classification with deep convolutional neural networks." Advances in neural information processing systems.
- Kroodsma, Donald E, Bruce E Byers, Eben Goodale, Steven Johnson, and Wan-Chun Liu. 2001. Pseudoreplication in playback experiments, revisited a decade later. Academic Press.
- Kubikova, Lubica, Kazuhiro Wada, and Erich D Jarvis. 2010. "Dopamine receptors in a songbird brain." *Journal of Comparative Neurology* 518 (6):741-769.

- Kuebrich, B. D., and S. J. Sober. 2015. "Variations on a theme: Songbirds, variability, and sensorimotor error correction." *Neuroscience* 296:48-54. doi: 10.1016/j.neuroscience.2014.09.068.
- Kuhl, Patricia K, and Andrew N Meltzoff. 1996. "Infant vocalizations in response to speech: Vocal imitation and developmental change." *The journal of the Acoustical Society of America* 100 (4):2425-2438.
- Kupchik, Yonatan M, Robyn M Brown, Jasper A Heinsbroek, Mary Kay Lobo, Danielle J Schwartz, and Peter W Kalivas. 2015. "Coding the direct/indirect pathways by D1 and D2 receptors is not valid for accumbens projections." *Nature neuroscience* 18 (9):1230.
- Lammel, Stephan, Andrea Hetzel, Olga Häckel, Ian Jones, Birgit Liss, and Jochen Roeper. 2008. "Unique properties of mesoprefrontal neurons within a dual mesocorticolimbic dopamine system." *Neuron* 57 (5):760-773.
- Lazic, Stanley E. 2010. "The problem of pseudoreplication in neuroscientific studies: is it affecting your analysis?" *BMC neuroscience* 11 (1):5.
- Leblois, A., B. J. Wendel, and D. J. Perkel. 2010. "Striatal dopamine modulates basal ganglia output and regulates social context-dependent behavioral variability through D1 receptors." *J Neurosci* 30 (16):5730-43. doi: 10.1523/JNEUROSCI.5974-09.2010.
- LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton. 2015. "Deep learning." *Nature* 521 (7553):436.
- Lee, Rachel S, Marcelo G Mattar, Nathan F Parker, Ilana B Witten, and Nathaniel D Daw. 2019. "Reward prediction error does not explain movement selectivity in DMS-projecting dopamine neurons." *eLife* 8:e42992.
- Lee, Sang Wan, Shinsuke Shimojo, and John P O'Doherty. 2014. "Neural computations underlying arbitration between model-based and model-free learning." *Neuron* 81 (3):687-699.
- Lees, A. J., and E. Smith. 1983. "Cognitive deficits in the early stages of Parkinson's disease." *Brain* 106 (Pt 2):257-70.
- Lerner, Talia N, Carrie Shilyansky, Thomas J Davidson, Kathryn E Evans, Kevin T Beier, Kelly A Zalocusky, Ailey K Crow, Robert C Malenka, Liqun Luo, and Raju Tomer. 2015. "Intact-brain analyses reveal distinct information carried by SNc dopamine subcircuits." *Cell* 162 (3):635-647.
- Lewis, R, D Zee, M Hayman, and R Tamargo. 2001. "Oculomotor function in the rhesus monkey after deafferentation of the extraocular muscles." *Experimental Brain Research* 141 (3):349-358.
- Liang, Zhifeng, Glenn DR Watson, Kevin D Alloway, Gangchea Lee, Thomas Neuberger, and Nanyin %J Neuroimage Zhang. 2015. "Mapping the functional network of medial prefrontal cortex by combining optogenetics and fMRI in awake rats." 117:114-123.
- Lipkind, D., G. F. Marcus, D. K. Bemis, K. Sasahara, N. Jacoby, M. Takahasi, K. Suzuki, O. Feher, P. Ravbar, K. Okanoya, and O. Tchernichovski. 2013. "Stepwise acquisition of vocal combinatorial capacity in songbirds and human infants." *Nature* 498 (7452):104-8. doi: 10.1038/nature12173.
- Lippa, Arnold S, Seymour M Antelman, Alan E Fisher, and Don R Canfield. 1973. "Neurochemical mediation of reward: a significant role for dopamine?" *Pharmacology Biochemistry and Behavior* 1 (1):23-28.

- Ljungberg, Tomas, Paul Apicella, and Wolfram Schultz. 1992. "Responses of monkey dopamine neurons during learning of behavioral reactions." *Journal of neurophysiology* 67 (1):145-163.
- Lohani, Sweyta, Adria K Martig, Karl Deisseroth, Ilana B Witten, and Bitu Moghaddam. 2019. "Dopamine modulation of prefrontal cortex activity is manifold and operates at multiple temporal and spatial scales." *Cell reports* 27 (1):99-114. e6.
- Luan, Lan, Xiaoling Wei, Zhengtuo Zhao, Jennifer J Siegel, Ojas Potnis, Catherine A Tuppen, Shengqing Lin, Shams Kazmi, Robert A Fowler, and Stewart Holloway. 2017. "Ultraflexible nanoelectronic probes form reliable, glial scar-free neural integration." *Science Advances* 3 (2):e1601966.
- Maas, Cora JM, and Joop J Hox. 2005. "Sufficient sample sizes for multilevel modeling." *Methodology* 1 (3):86-92.
- Machado, Ana S, Dana M Darmohray, Joao Fayad, Hugo G Marques, and Megan R Carey. 2015. "A quantitative framework for whole-body coordination reveals specific deficits in freely walking ataxic mice." *Elife* 4:e07892.
- Mandelblat-Cerf, Y., L. Las, N. Denisenko, and M. S. Fee. 2014. "A role for descending auditory cortical projections in songbird vocal learning." *Elife* 3. doi: 10.7554/eLife.02152.
- Marcus, Gary. 2018. "Innateness, alphazero, and artificial intelligence." *arXiv preprint arXiv:05667*.
- Margolis, Elyssa B, Hagar Lock, Gregory O Hjelmstad, and Howard L Fields. 2006. "The ventral tegmental area revisited: is there an electrophysiological marker for dopaminergic neurons?" *The Journal of physiology* 577 (3):907-924.
- Markram, Henry. 2006. "The blue brain project." *Nature Reviews Neuroscience* 7 (2):153.
- Marshall, John F, Deborah Levitan, and Edward M Stricker. 1976. "Activation-induced restoration of sensorimotor functions in rats with dopamine-depleting brain lesions." *Journal of Comparative and Physiological Psychology* 90 (6):536.
- Martin, TA, JG Keating, HP Goodkin, AJ Bastian, and WT Thach. 1996. "Throwing while looking through prisms: I. Focal olivocerebellar lesions impair adaptation." *Brain* 119 (4):1183-1198.
- Mathis, Mackenzie Weygandt, Alexander Mathis, and Naoshige Uchida. 2017. "Somatosensory cortex plays an essential role in forelimb motor adaptation in mice." *Neuron* 93 (6):1493-1503. e6.
- Matsumoto, Masayuki, and Okihide Hikosaka. 2007. "Lateral habenula as a source of negative reward signals in dopamine neurons." *Nature* 447 (7148):1111.
- Matsumoto, Masayuki, and Okihide Hikosaka. 2009. "Two types of dopamine neuron distinctly convey positive and negative motivational signals." *Nature* 459 (7248):837.
- Maxwell, James Clerk. 1865. "VIII. A dynamical theory of the electromagnetic field." *Philosophical transactions of the Royal Society of London* (155):459-512.
- Mazzoni, P., A. Hristova, and J. W. Krakauer. 2007. "Why don't we move faster? Parkinson's disease, movement vigor, and implicit motivation." *J Neurosci* 27 (27):7105-16. doi: 10.1523/JNEUROSCI.0264-07.2007.
- McCoach, D Betsy, and Jill L Adelson. 2010. "Dealing with dependence (Part I): Understanding the effects of clustered data." *Gifted Child Quarterly* 54 (2):152-155.
- McDannald, Michael A, Yuji K Takahashi, Nina Lopatina, Brad W Pietras, Josh L Jones, and Geoffrey Schoenbaum. 2012. "Model-based learning and the contribution of the

- orbitofrontal cortex to the model-free world." *European Journal of Neuroscience* 35 (7):991-996.
- McLaughlin, Samuel C. 1967. "Parametric adjustment in saccadic eye movements." *Perception & Psychophysics* 2 (8):359-362.
- Meltzoff, Andrew N, and M Keith Moore. 1983. "Newborn infants imitate adult facial gestures." *Child development*:702-709.
- Meng, Chengbo, Jingheng Zhou, Amy Papaneri, Teja Peddada, Karen Xu, and Guohong Cui. 2018. "Spectrally resolved fiber photometry for multi-component analysis of brain circuits." *Neuron* 98 (4):707-717. e4.
- Miłkowski, Marcin, Witold M Hensel, and Mateusz Hohol. 2018. "Replicability or reproducibility? On the replication crisis in computational neuroscience and sharing only relevant detail." *Journal of computational neuroscience* 45 (3):163-172.
- Millar, Russell B, and Marti J Anderson. 2004. "Remedies for pseudoreplication." *Fisheries Research* 70 (2-3):397-407.
- Mnih, Volodymyr, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. 2016. "Asynchronous methods for deep reinforcement learning." International conference on machine learning.
- Mnih, Volodymyr, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, and Georg Ostrovski. 2015. "Human-level control through deep reinforcement learning." *Nature* 518 (7540):529.
- Mohammadi, SA, and BM Prasanna. 2003. "Analysis of genetic diversity in crop plants—salient statistical tools and considerations." *Crop science* 43 (4):1235-1248.
- Mohan, Vishwanathan, Pietro Morasso, and Giorgio %J Neurocomputing Metta. 2011. "The distribution of rewards in sensorimotor maps acquired by cognitive robots through exploration." 74 (17):3440-3455.
- Mohebi, Ali, Jeffrey R. Pettibone, Arif A. Hamid, Jenny-Marie T. Wong, Leah T. Vinson, Tommaso Patriarchi, Lin Tian, Robert T. Kennedy, and Joshua D. Berke. 2019. "Dissociable dopamine dynamics for learning and motivation." *Nature*. doi: 10.1038/s41586-019-1235-y.
- Mollaei, F., D. M. Shiller, and V. L. Gracco. 2013. "Sensorimotor adaptation of speech in Parkinson's disease." *Mov Disord* 28 (12):1668-74. doi: 10.1002/mds.25588.
- Momennejad, Ida, Evan M Russek, Jin H Cheong, Matthew M Botvinick, ND Daw, and Samuel J Gershman. 2017. "The successor representation in human reinforcement learning." *Nature Human Behaviour* 1 (9):680.
- Mongeon, David, Pierre Blanchet, and Julie Messier. 2013. "Impact of Parkinson's disease and dopaminergic medication on adaptation to explicit and implicit visuomotor perturbations." *Brain and cognition* 81 (2):271-282.
- Montague, P Read, Peter Dayan, and Terrence J Sejnowski. 1996. "A framework for mesencephalic dopamine systems based on predictive Hebbian learning." *Journal of neuroscience* 16 (5):1936-1947.
- Monti, Jaime M, and Daniel Monti. 2007. "The involvement of dopamine in the modulation of sleep and waking." *Sleep medicine reviews* 11 (2):113-133.
- Morton, Susanne M, and Amy J Bastian. 2004. "Prism adaptation during walking generalizes to reaching and requires the cerebellum." *Journal of neurophysiology*.

- Munoz, Douglas P. 2002. "Commentary: saccadic eye movements: overview of neural circuitry." In *Progress in brain research*, 89-96. Elsevier.
- Munoz, Douglas P, Michael C Dorris, Martin Paré, and Stefan Everling. 2000. "On your mark, get set: brainstem circuitry underlying saccadic initiation." *Canadian journal of physiology and pharmacology* 78 (11):934-944.
- Murdoch, Don, Ruidong Chen, and Jesse H Goldberg. 2018. "Place preference and vocal learning rely on distinct reinforcers in songbirds." *Scientific reports* 8.
- Musca, Serban C, Rodolphe Kamiejski, Armelle Nugier, Alain Méot, Abdelatif Er-Rafiy, and Markus Brauer. 2011. "Data with hierarchical structure: impact of intraclass correlation and sample size on Type-I error." *Frontiers in Psychology* 2:74.
- Nestler, Eric J, and William A Carlezon Jr. 2006. "The mesolimbic dopamine reward circuit in depression." *Biological psychiatry* 59 (12):1151-1159.
- Ng, Andrew Y, and Stuart J Russell. 2000. "Algorithms for inverse reinforcement learning." *Icml*.
- Nieouillon, André. 2002. "Dopamine and the regulation of cognition and attention." *Progress in neurobiology* 67 (1):53-83.
- Nieuwenhuis, Sander, Birte U Forstmann, and Eric-Jan Wagenmakers. 2011. "Erroneous analyses of interactions in neuroscience: a problem of significance." *Nature neuroscience* 14 (9):1105.
- Niv, Y., N. D. Daw, D. Joel, and P. Dayan. 2007. "Tonic dopamine: opportunity costs and the control of response vigor." *Psychopharmacology (Berl)* 191 (3):507-20. doi: 10.1007/s00213-006-0502-4.
- Nordeen, Kathy W, and Ernest J Nordeen. 1992. "Auditory feedback is necessary for the maintenance of stereotyped song in adult zebra finches." *Behavioral and neural biology* 57 (1):58-66.
- Nottebohm, Fernando, Tegner M Stokes, and Christiana M Leonard. 1976. "Central control of song in the canary, *Serinus canarius*." *Journal of Comparative Neurology* 165 (4):457-486.
- Oksanen, Lauri. 2001. "Logic of experiments in ecology: is pseudoreplication a pseudoissue?" *Oikos* 94 (1):27-38.
- Optican, LANCE M, and DAVID A Robinson. 1980. "Cerebellar-dependent adaptive control of primate saccadic system." *Journal of Neurophysiology* 44 (6):1058-1076.
- Panigrahi, B., K. A. Martin, Y. Li, A. R. Graves, A. Vollmer, L. Olson, B. D. Mensh, A. Y. Karpova, and J. T. Dudman. 2015. "Dopamine Is Required for the Neural Representation and Control of Movement Vigor." *Cell* 162 (6):1418-30. doi: 10.1016/j.cell.2015.08.014.
- Panzeri, Stefano, Rasmus S Petersen, Simon R Schultz, Michael Lebedev, and Mathew E Diamond. 2001. "The role of spike timing in the coding of stimulus location in rat somatosensory cortex." *Neuron* 29 (3):769-777.
- Paquet, F., M. A. Bedard, M. Levesque, P. L. Tremblay, M. Lemay, P. J. Blanchet, P. Scherzer, S. Chouinard, and J. Filion. 2008. "Sensorimotor adaptation in Parkinson's disease: evidence for a dopamine dependent remapping disturbance." *Exp Brain Res* 185 (2):227-36. doi: 10.1007/s00221-007-1147-1.
- Parkinson, James. 1817. *An Essay on the Shaking Palsy*: Dawson.
- Parkinson, James. 2002. "An essay on the shaking palsy." *The Journal of neuropsychiatry and clinical neurosciences* 14 (2):223-236.

- Pavlov, Ivan Petrovitch, and W Gantt. 1928. "Lectures on conditioned reflexes: Twenty-five years of objective study of the higher nervous activity (behaviour) of animals."
- Peh, W. Y., T. F. Roberts, and R. Mooney. 2015. "Imaging auditory representations of song and syllables in populations of sensorimotor neurons essential to vocal communication." *J Neurosci* 35 (14):5589-605. doi: 10.1523/JNEUROSCI.2308-14.2015.
- Péllisson, D, N Alahyane, M Panouilleres, and C Tilikete. 2010. "Sensorimotor adaptation of saccadic eye movements." *Neuroscience & Biobehavioral Reviews* 34 (8):1103-1120.
- Peng, Jing, and Ronald J Williams. 1993. "Efficient learning and planning within the Dyna framework." *Adaptive Behavior* 1 (4):437-454.
- Peng, Jing, and Ronald J Williams. 1994. "Incremental multi-step Q-learning." In *Machine Learning Proceedings 1994*, 226-232. Elsevier.
- Person, A. L., S. D. Gale, M. A. Farries, and D. J. Perkel. 2008. "Organization of the songbird basal ganglia, including area X." *J Comp Neurol* 508 (5):840-66. doi: 10.1002/cne.21699.
- Pessiglione, Mathias, Ben Seymour, Guillaume Flandin, Raymond J Dolan, and Chris D Frith. 2006. "Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans." *Nature* 442 (7106):1042.
- Pleil, Kristen E, Christa M Helms, Jon R Sobus, James B Daunais, Kathleen A Grant, and Thomas L %J Addiction biology Kash. 2016. "Effects of chronic alcohol consumption on neuronal function in the non-human primate BNST." 21 (6):1151-1167.
- Poirier, LJ, and TL Sourkes. 1965. "Influence of the substantia nigra on the catecholamine content of the striatum." *Brain* 88 (1):181-192.
- Poirier, LJ, TL Sourkes, G Bouvier, R Boucher, and S Carabin. 1966. "Striatal amines, experimental tremor and the effect of harmaline in the monkey." *Brain* 89 (1):37-52.
- Polikov, Vadim S, Patrick A Tresco, and William M Reichert. 2005. "Response of brain tissue to chronically implanted neural electrodes." *Journal of neuroscience methods* 148 (1):1-18.
- Price, Philip H. 1979. "Developmental determinants of structure in zebra finch song." *Journal of comparative and physiological psychology* 93 (2):260.
- Puterman, Martin L. 2014. *Markov decision processes: discrete stochastic dynamic programming*: John Wiley & Sons.
- Ramakrishnan, Arjun, Yoon Woo Byun, Kyle Rand, Christian E Pedersen, Mikhail A Lebedev, and Miguel AL Nicolelis. 2017. "Cortical neurons multiplex reward-related signals along with sensory and motor information." *Proceedings of the National Academy of Sciences* 114 (24):E4841-E4850.
- Redgrave, Peter, Manuel Rodriguez, Yoland Smith, Maria C Rodriguez-Oroz, Stephane Lehericy, Hagai Bergman, Yves Agid, Mahlon R DeLong, and Jose A Obeso. 2010. "Goal-directed and habitual control in the basal ganglia: implications for Parkinson's disease." *Nature Reviews Neuroscience* 11 (11):760.
- Reiner, Anton, David J Perkel, Laura L Bruce, Ann B Butler, Andrés Csillag, Wayne Kuenzel, Loreta Medina, George Paxinos, Toru Shimizu, and Georg Striedter. 2004. "Revised nomenclature for avian telencephalon and some related brainstem nuclei." *Journal of Comparative Neurology* 473 (3):377-414.
- Rescorla, Robert A. 1969. "Conditioned inhibition of fear resulting from negative CS-US contingencies." *Journal of comparative physiological psychology* 67 (4):504.
- Rescorla, Robert A 1968. "Probability of shock in the presence and absence of CS in fear conditioning." *Journal of comparative physiological psychology* 66 (1):1.

- Rescorla, Robert A, and Allan R Wagner. 1972. "A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement." *Classical conditioning II: Current research theory* 2:64-99.
- Riede, T., J. H. Fisher, and F. Goller. 2010. "Sexual dimorphism of the zebra finch syrinx indicates adaptation for high fundamental frequencies in males." *PLoS One* 5 (6):e11368. doi: 10.1371/journal.pone.0011368.
- Robinson, David A. 1973. "Models of the saccadic eye movement control system." *Kybernetik* 14 (2):71-83.
- Robinson, Siobhan, Suzanne M Sandstrom, Victor H Denenberg, and Richard D Palmiter. 2005. "Distinguishing whether dopamine regulates liking, wanting, and/or learning about rewards." *Behavioral neuroscience* 119 (1):5.
- Romo, Ranulfo, and Wolfram Schultz. 1990. "Dopamine neurons of the monkey midbrain: contingencies of responses to active touch during self-initiated arm movements." *Journal of neurophysiology* 63 (3):592-606.
- Rossetti, Yves, Gilles Rode, Laure Pisella, Alessandro Farné, Ling Li, Dominique Boisson, and Marie-Thérèse Perenin. 1998. "Prism adaptation to a rightward optical deviation rehabilitates left hemispatial neglect." *Nature* 395 (6698):166.
- Roux, AV Diez. 2002. "A glossary for multilevel analysis." *Journal of Epidemiology & Community Health* 56 (8):588-594.
- Russek, Evan M, Ida Momennejad, Matthew M Botvinick, Samuel J Gershman, and Nathaniel D Daw. 2017. "Predictive representations can link model-based reinforcement learning to model-free mechanisms." *PLoS computational biology* 13 (9):e1005768.
- Sakata, J. T., and M. S. Brainard. 2006. "Real-time contributions of auditory feedback to avian vocal motor control." *J Neurosci* 26 (38):9619-28. doi: 10.1523/JNEUROSCI.2027-06.2006.
- Sakata, J. T., and M. S. Brainard. 2008. "Online contributions of auditory feedback to neural activity in avian song control circuitry." *J Neurosci* 28 (44):11378-90. doi: 10.1523/JNEUROSCI.3254-08.2008.
- Sakata, J. T., C. M. Hampton, and M. S. Brainard. 2008. "Social modulation of sequence and syllable variability in adult birdsong." *J Neurophysiol* 99 (4):1700-11. doi: 10.1152/jn.01296.2007.
- Salamone, J. D., and M. Correa. 2012. "The mysterious motivational functions of mesolimbic dopamine." *Neuron* 76 (3):470-85. doi: 10.1016/j.neuron.2012.10.021.
- Salamone, J. D., M. Correa, A. Farrar, and S. M. Mingote. 2007. "Effort-related functions of nucleus accumbens dopamine and associated forebrain circuits." *Psychopharmacology (Berl)* 191 (3):461-82. doi: 10.1007/s00213-006-0668-9.
- Samejima, Kazuyuki, and Kenji Doya. 2007. "Multiple representations of belief states and action values in corticobasal ganglia loops." *Annals of the New York Academy of Sciences* 1104 (1):213-228.
- Saravanan, Varun, Lukas A Hoffmann, Amanda L Jacob, Gordon J Berman, and Samuel J Sober. 2019a. "Dopamine depletion affects vocal acoustics and disrupts sensorimotor adaptation in songbirds." *bioRxiv*:600874.
- Saravanan, Varun, Lukas A Hoffmann, Amanda L Jacob, Gordon J Berman, and Samuel J. Sober. 2019b. "Dopamine depletion affects vocal acoustics and disrupts sensorimotor adaptation in songbirds." *eneuro*:ENEURO.0190-19.2019. doi: 10.1523/eneuro.0190-19.2019.

- Sasaki, A., T. D. Sotnikova, R. R. Gainetdinov, and E. D. Jarvis. 2006. "Social context-dependent singing-regulated dopamine." *J Neurosci* 26 (35):9010-4. doi: 10.1523/JNEUROSCI.1335-06.2006.
- Schank, Jeffrey C, and Thomas J Koehnle. 2009. "Pseudoreplication is a pseudoproblem." *Journal of Comparative Psychology* 123 (4):421.
- Scharff, C., and F. Nottebohm. 1991. "A comparative study of the behavioral deficits following lesions of various parts of the zebra finch song system: implications for vocal learning." *J Neurosci* 11 (9):2896-913.
- Schultz, W., P. Dayan, and P. R. Montague. 1997. "A neural substrate of prediction and reward." *Science* 275 (5306):1593-9.
- Schultz, Wolfram. 1986. "Responses of midbrain dopamine neurons to behavioral trigger stimuli in the monkey." *Journal of neurophysiology* 56 (5):1439-1461.
- Schultz, Wolfram. 1998. "Predictive reward signal of dopamine neurons." *Journal of neurophysiology* 80 (1):1-27.
- Schultz, Wolfram. 2016. "Dopamine reward prediction-error signalling: a two-component response." *Nature Reviews Neuroscience* 17 (3):183.
- Schultz, Wolfram, Paul Apicella, Eugenio Scarnati, and Tomas Ljungberg. 1992. "Neuronal activity in monkey ventral striatum related to the expectation of reward." *Journal of neuroscience* 12 (12):4595-4610.
- Schultz, Wolfram, and Ranulfo Romo. 1990. "Dopamine neurons of the monkey midbrain: contingencies of responses to stimuli eliciting immediate behavioral reactions." *Journal of neurophysiology* 63 (3):607-624.
- Schwabe, Lars, and Oliver T Wolf. 2011. "Stress-induced modulation of instrumental behavior: from goal-directed to habitual control of action." *Behavioural brain research* 219 (2):321-328.
- Seidler, Rachael D. 2006. "Differential effects of age on sequence learning and sensorimotor adaptation." *Brain research bulletin* 70 (4-6):337-346.
- Shadmehr, R., and M. A. Arbib. 1992. "A mathematical analysis of the force-stiffness characteristics of muscles in control of a single joint system." *Biol Cybern* 66 (6):463-77.
- Shadmehr, R., and F. A. Mussa-Ivaldi. 1994. "Adaptive representation of dynamics during learning of a motor task." *J Neurosci* 14 (5 Pt 2):3208-24.
- Shadmehr, Reza, Ferdinando A Mussa-Ivaldi, and Emilio Bizzi. 1993. "Postural force fields of the human arm and their role in generating multijoint movements." *Journal of Neuroscience* 13 (1):45-62.
- Shannon, Claude Elwood. 1948. "A mathematical theory of communication." *Bell system technical journal* 27 (3):379-423.
- Shimodaira, Hidetoshi. 2002. "An approximately unbiased test of phylogenetic tree selection." *Systematic biology* 51 (3):492-508.
- Shimodaira, Hidetoshi. 2004. "Approximately unbiased tests of regions using multistep-multiscale bootstrap resampling." *The Annals of Statistics* 32 (6):2616-2641.
- Shiotsuki, Hiromi, Kenji Yoshimi, Yasushi Shimo, Manabu Funayama, Yukio Takamatsu, Kazutaka Ikeda, Ryosuke Takahashi, Shigeru Kitazawa, and Nobutaka Hattori. 2010. "A rotarod test for evaluation of motor skill learning." *Journal of neuroscience methods* 189 (2):180-185.

- Shum, Mamie, Douglas M Shiller, Shari R Baum, and Vincent L Gracco. 2011. "Sensorimotor integration for speech motor learning involves the inferior parietal cortex." *European Journal of Neuroscience* 34 (11):1817-1822.
- Shwartz-Ziv, Ravid, and Naftali Tishby. 2017. "Opening the black box of deep neural networks via information." *arXiv preprint arXiv:00810*.
- Šidák, Zbyněk. 1967. "Rectangular confidence regions for the means of multivariate normal distributions." *Journal of the American Statistical Association* 62 (318):626-633.
- Silver, David, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, and Marc Lanctot. 2016. "Mastering the game of Go with deep neural networks and tree search." *Nature* 529 (7587):484.
- Silver, David, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, and Adrian Bolton. 2017. "Mastering the game of go without human knowledge." *Nature* 550 (7676):354.
- Skinner, Burrhus F. 1963. "Operant behavior." *American psychologist* 18 (8):503.
- Snijders, Tom AB. 2011. *Multilevel analysis*: Springer.
- Snijders, Tom AB, and Roel J Bosker. 1993. "Standard errors and sample sizes for two-level research." *Journal of educational statistics* 18 (3):237-259.
- Sober, S. J., and M. S. Brainard. 2009. "Adult birdsong is actively maintained by error correction." *Nat Neurosci* 12 (7):927-31. doi: 10.1038/nn.2336.
- Sober, S. J., and M. S. Brainard. 2012. "Vocal learning is constrained by the statistics of sensorimotor experience." *Proc Natl Acad Sci U S A* 109 (51):21099-103. doi: 10.1073/pnas.1213622109.
- Sober, Samuel J, Simon Sponberg, Ilya Nemenman, and Lena H Ting. 2018. "Millisecond spike timing codes for motor control." *Trends in neurosciences* 41 (10):644-648.
- Soha, J. A., T. Shimizu, and A. J. Doupe. 1996. "Development of the catecholaminergic innervation of the song system of the male zebra finch." *J Neurobiol* 29 (4):473-89. doi: 10.1002/(SICI)1097-4695(199604)29:4<473::AID-NEU5>3.0.CO;2-5.
- Sohrabji, F., E. J. Nordeen, and K. W. Nordeen. 1990. "Selective impairment of song learning following lesions of a forebrain nucleus in the juvenile zebra finch." *Behav Neural Biol* 53 (1):51-63.
- Sparks, David L. 1986. "Translation of sensory signals into commands for control of saccadic eye movements: role of primate superior colliculus." *Physiological reviews* 66 (1):118-171.
- Srivastava, Kyle H, Caroline M Holmes, Michiel Vellema, Andrea R Pack, Coen PH Elemans, Ilya Nemenman, and Samuel J Sober. 2017. "Motor control by precisely timed spike patterns." *Proceedings of the National Academy of Sciences* 114 (5):1171-1176.
- Steinberg, E. E., R. Keiflin, J. R. Boivin, I. B. Witten, K. Deisseroth, and P. H. Janak. 2013. "A causal link between prediction errors, dopamine neurons and learning." *Nat Neurosci* 16 (7):966-73. doi: 10.1038/nn.3413.
- Stellar, James R, Ann E Kelley, and Dale Corbett. 1983. "Effects of peripheral and central dopamine blockade on lateral hypothalamic self-stimulation: Evidence for both reward and motor deficits." *Pharmacology Biochemistry and Behavior* 18 (3):433-442.
- Stewart-Oaten, Allan, William W Murdoch, and Keith R Parker. 1986. "Environmental impact assessment: "Pseudoreplication" in time?" *Ecology* 67 (4):929-940.

- Stratton, George M. 1896. "Some preliminary experiments on vision without inversion of the retinal image." *Psychological review* 3 (6):611.
- Stratton, George M. 1897. "Vision without inversion of the retinal image." *Psychological review* 4 (4):341.
- Strong, Steven P, Roland Koberle, Rob R de Ruyter van Steveninck, and William Bialek. 1998. "Entropy and information in neural spike trains." *Physical review letters* 80 (1):197.
- Stuber, Garret D, Thomas S Hnasko, Jonathan P Britt, Robert H Edwards, and Antonello Bonci. 2010. "Dopaminergic terminals in the nucleus accumbens but not the dorsal striatum corelease glutamate." *Journal of Neuroscience* 30 (24):8229-8233.
- Sutskever, Ilya, Oriol Vinyals, and Quoc V Le. 2014. "Sequence to sequence learning with neural networks." *Advances in neural information processing systems*.
- Sutton, Richard S. 1988. "Learning to predict by the methods of temporal differences." *Machine learning* 3 (1):9-44.
- Sutton, Richard S. 1991. "Dyna, an integrated architecture for learning, planning, and reacting." *ACM SIGART Bulletin* 2 (4):160-163.
- Sutton, Richard S, and Andrew G Barto. 1998. *Introduction to reinforcement learning*. Vol. 135: MIT press Cambridge.
- Sutton, Richard S, David A McAllester, Satinder P Singh, and Yishay Mansour. 2000. "Policy gradient methods for reinforcement learning with function approximation." *Advances in neural information processing systems*.
- Sutton, Richard S, Doina Precup, and Satinder Singh. 1999. "Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning." *Artificial intelligence* 112 (1-2):181-211.
- Sutton, Richard Stuart. 1984. "Temporal credit assignment in reinforcement learning."
- Tang, Claire, Diala Chehayeb, Kyle Srivastava, Ilya Nemenman, and Samuel J Sober. 2014. "Millisecond-scale motor encoding in a cortical vocal area." *PLoS biology* 12 (12):e1002018.
- Tesauro, Gerald. 1992. "Practical issues in temporal difference learning." *Advances in neural information processing systems*.
- Tesauro, Gerald %J Communications of the ACM. 1995. "Temporal difference learning and TD-Gammon." 38 (3):58-69.
- Thai, Hoai-Thu, France Mentré, Nicholas HG Holford, Christine Veyrat-Follet, and Emmanuelle Comets. 2013. "A comparison of bootstrap approaches for estimating uncertainty of parameters in linear mixed-effects models." *Pharmaceutical statistics* 12 (3):129-140.
- Threlfell, Sarah, Tatjana Lalic, Nicola J Platt, Katie A Jennings, Karl Deisseroth, and Stephanie J Cragg. 2012. "Striatal dopamine release is triggered by synchronized activity in cholinergic interneurons." *Neuron* 75 (1):58-64.
- Tian, Ju, and Naoshige Uchida. 2015. "Habenula lesions reveal that multiple mechanisms underlie dopamine prediction errors." *Neuron* 87 (6):1304-1316.
- Tian, L. Y., and M. S. Brainard. 2017. "Discrete Circuits Support Generalized versus Context-Specific Vocal Learning in the Songbird." *Neuron* 96 (5):1168-1177 e5. doi: 10.1016/j.neuron.2017.10.019.
- Tian, Lin, S Andrew Hires, Tianyi Mao, Daniel Huber, M Eugenia Chiappe, Sreekanth H Chalasani, Leopoldo Petreanu, Jasper Akerboom, Sean A McKinney, and Eric R Schreiter. 2009. "Imaging neural activity in worms, flies and mice with improved GCaMP calcium indicators." *Nature methods* 6 (12):875.

- Tobler, Philippe N, Christopher D Fiorillo, and Wolfram Schultz. 2005. "Adaptive coding of reward value by dopamine neurons." *Science* 307 (5715):1642-1645.
- Tomasello, Michael, Maryann Davis-Dasilva, Lael CamaK, and Kim Bard. 1987. "Observational learning of tool-use by young chimpanzees." *Human evolution* 2 (2):175-183.
- Tremblay, Stéphanie, Douglas M Shiller, and David J Ostry. 2003. "Somatosensory basis of speech production." *Nature* 423 (6942):866.
- Tseng, Ya-weng, Jörn Diedrichsen, John W Krakauer, Reza Shadmehr, and Amy J Bastian. 2007. "Sensory prediction errors drive cerebellum-dependent adaptation of reaching." *Journal of neurophysiology*.
- Tsitsiklis, John N, and Benjamin Van Roy. 1997. "Analysis of temporal-difference learning with function approximation." *Advances in neural information processing systems*.
- Tumer, E. C., and M. S. Brainard. 2007. "Performance variability enables adaptive plasticity of 'crystallized' adult birdsong." *Nature* 450 (7173):1240-4. doi: 10.1038/nature06390.
- Turner, R. S., and M. Desmurget. 2010. "Basal ganglia contributions to motor control: a vigorous tutor." *Curr Opin Neurobiol* 20 (6):704-16. doi: 10.1016/j.conb.2010.08.022.
- Ungerstedt, Urban. 1968. "6-Hydroxy-dopamine induced degeneration of central monoamine neurons." *European journal of pharmacology* 5 (1):107-110.
- Ungerstedt, Urban. 1971. "Adipsia and aphagia after 6-hydroxydopamine induced degeneration of the nigro-striatal dopamine system." *Acta Physiologica Scandinavica* 82 (S367):95-122.
- Ungerstedt, Urban, and Gordon W Arbuthnott. 1970. "Quantitative recording of rotational behavior in rats after 6-hydroxy-dopamine lesions of the nigrostriatal dopamine system." *Brain research* 24 (3):485-493.
- Valentin, Vivian V, Anthony Dickinson, and John P O'Doherty. 2007. "Determining the neural substrates of goal-directed learning in the human brain." *Journal of Neuroscience* 27 (15):4019-4026.
- Van Essen, David C, Stephen M Smith, Deanna M Barch, Timothy EJ Behrens, Essa Yacoub, Kamil Ugurbil, and Wu-Minn HCP Consortium. 2013. "The WU-Minn human connectome project: an overview." *Neuroimage* 80:62-79.
- Van Eycke, Y. R., J. Allard, I. Salmon, O. Debeir, and C. Decaestecker. 2017. "Image processing in digital pathology: an opportunity to solve inter-batch variability of immunohistochemical staining." *Sci Rep* 7:42964. doi: 10.1038/srep42964.
- Van Hasselt, Hado, Arthur Guez, and David Silver. 2016. "Deep reinforcement learning with double q-learning." Thirtieth AAAI Conference on Artificial Intelligence.
- van Steveninck, Rob R de Ruyter, Geoffrey D Lewen, Steven P Strong, Roland Koberle, and William Bialek. 1997. "Reproducibility and variability in neural spike trains." *Science* 275 (5307):1805-1808.
- Villacorta, Virgilio M, Joseph S Perkell, and Frank H Guenther. 2007. "Sensorimotor adaptation to feedback perturbations of vowel acoustics and its relation to perception." *The Journal of the Acoustical Society of America* 122 (4):2306-2319.
- Vlasits, Anna L, Ryan D Morrie, Alexandra Tran-Van-Minh, Adam Bleckert, Christian F Gainer, David A DiGregorio, and Marla B Feller. 2016. "A role for synaptic input distribution in a dendritic computation of motion direction in the retina." *Neuron* 89 (6):1317-1330.
- Vogelstein, Joshua T, Eric Perlman, Benjamin Falk, Alex Baden, William Gray Roncal, Vikram Chandrashekar, Forrest Collman, Sharmishta Seshamani, Jesse L Patsolic, and Kunal Lillaney. 2018. "A community-developed open-source computational ecosystem for big neuro data." *Nature methods* 15 (11):846.

- Volkow, Nora D, Ruben C Gur, Gene-Jack Wang, Joanna S Fowler, Paul J Moberg, Yu-Shin Ding, Robert Hitzemann, Gwenn Smith, and Jean Logan. 1998. "Association between decline in brain dopamine activity with age and cognitive and motor impairment in healthy individuals." *American Journal of psychiatry* 155 (3):344-349.
- Volkow, Nora D, Gene-Jack Wang, Joanna S Fowler, Dardo Tomasi, and Frank Telang. 2011. "Addiction: beyond dopamine reward circuitry." *Proceedings of the National Academy of Sciences* 108 (37):15037-15042.
- Wagner, Allan R, Frank A Logan, and Karl Haberlandt. 1968. "Stimulus selection in animal discrimination learning." *Journal of experimental psychology* 76 (2p1):171.
- Wagner, AR. 1969. "Stimulus validity and stimulus selection in associative learning." *Fundamental issues in associative learning*:90-122.
- Walsh, John E. 1947. "Concerning the effect of intraclass correlation on certain significance tests." *The Annals of Mathematical Statistics* 18 (1):88-96.
- Warren, Timothy Little, Evren C Tumer, Jonathan D Charlesworth, and Michael S Brainard. 2011. "Mechanisms and time Course of Vocal Learning and Consolidation in the Adult Songbird." *American Journal of Physiology-Heart and Circulatory Physiology*.
- Wasserstein, Ronald L, and Nicole A Lazar. 2016. "The ASA's statement on p-values: context, process, and purpose." *The American Statistician* 70 (2):129-133.
- Watabe-Uchida, Mitsuko, Neir Eshel, and Naoshige Uchida. 2017. "Neural circuitry of reward prediction error." *Annual review of neuroscience* 40:373-394.
- Watkins, Christopher JCH, and Peter Dayan. 1992. "Q-learning." *Machine Learning* 8 (3-4):279-292.
- Watkins, Christopher John Cornish Hellaby. 1989. "Learning from delayed rewards." King's College, Cambridge.
- Weiner, Mark J, Mark Hallett, and H Harris Funkenstein. 1983. "Adaptation to lateral displacement of vision in patients with lesions of the central nervous system." *Neurology* 33 (6):766-766.
- Welch, Robert B. 1972. "The effect of experienced limb identity upon adaptation to simulated displacement of the visual field." *Perception & Psychophysics* 12 (6):453-456.
- Welch, Robert B, Bruce Bridgeman, Sulekha Anand, and Kaitlin E Browman. 1993. "Alternating prism exposure causes dual adaptation and generalization to a novel displacement." *Perception & Psychophysics* 54 (2):195-204.
- Wilbrecht, Linda, and Fernando Nottebohm. 2003. "Vocal learning in birds and humans." *Mental retardation and developmental disabilities research reviews* 9 (3):135-148.
- Williams, Ronald J, and Leemon C Baird. 1993. Tight performance bounds on greedy policies based on imperfect value functions. Citeseer.
- Wise, R. A. 2004. "Dopamine, learning and motivation." *Nat Rev Neurosci* 5 (6):483-94. doi: 10.1038/nrn1406.
- Wise, Roy A, and Pierre-Paul Rompre. 1989. "Brain dopamine and reward." *Annual review of psychology* 40 (1):191-225.
- Wolpert, D. M., J. Diedrichsen, and J. R. Flanagan. 2011. "Principles of sensorimotor learning." *Nat Rev Neurosci* 12 (12):739-51. doi: 10.1038/nrn3112.
- Wolpert, D. M., Z. Ghahramani, and M. I. Jordan. 1995. "An internal model for sensorimotor integration." *Science* 269 (5232):1880-2.
- Wolpert, Daniel M, R Chris Miall, and Mitsuo Kawato. 1998. "Internal models in the cerebellum." *Trends in cognitive sciences* 2 (9):338-347.

- Wykes, Robert C, Joost H Heeroma, Laura Mantoan, Kaiyu Zheng, Douglas C MacDonald, Karl Deisseroth, Kevan S Hashemi, Matthew C Walker, Stephanie Schorge, and Dimitri M %J Science translational medicine Kullmann. 2012. "Optogenetic and potassium channel gene therapy in a rodent model of focal neocortical epilepsy." 3004190.
- Xiao, L., G. Chattree, F. G. Oscos, M. Cao, M. J. Wanat, and T. F. Roberts. 2018. "A Basal Ganglia Circuit Sufficient to Guide Birdsong Learning." *Neuron* 98 (1):208-221 e5. doi: 10.1016/j.neuron.2018.02.020.
- Yapo, Cedric, Anu G Nair, Lorna Clement, Liliana R Castro, Jeanette Hellgren Kotaleski, and Pierre Vincent. 2017. "Detection of phasic dopamine by D1 and D2 striatal medium spiny neurons." *The Journal of physiology* 595 (24):7451-7475.
- Yizhar, Ofer, Lief E Fenno, Thomas J Davidson, Murtaza Mogri, and Karl Deisseroth. 2011. "Optogenetics in neural systems." *Neuron* 71 (1):9-34.
- Yokel, Robert A, and Roy A Wise. 1975. "Increased lever pressing for amphetamine after pimozide in rats: implications for a dopamine theory of reward." *Science* 187 (4176):547-549.
- Zhou, B., D. Hofmann, I. Pinkoviezky, S. J. Sober, and I. Nemenman. 2018. "Chance, long tails, and inference in a non-Gaussian, Bayesian theory of vocal learning in songbirds." *Proc Natl Acad Sci U S A* 115 (36):E8538-E8546. doi: 10.1073/pnas.1713020115.
- Zhou, Qun-Yong, and Richard D Palmiter. 1995. "Dopamine-deficient mice are severely hypoactive, adipsic, and aphagic." *Cell* 83 (7):1197-1209.
- Ziebart, Brian D, Andrew L Maas, J Andrew Bagnell, and Anind K Dey. 2008. "Maximum entropy inverse reinforcement learning." Aaai.
- Zong, Weijian, Runlong Wu, Mingli Li, Yanhui Hu, Yijun Li, Jinghang Li, Hao Rong, Haitao Wu, Yangyang Xu, and Yang Lu. 2017. "Fast high-resolution miniature two-photon microscopy for brain imaging in freely behaving mice." *Nature methods* 14 (7):713.