From the holdings of the Department of Special Collections , Princeton University Library, Princeton, NJ

These documents are either in the public domain, or can only be used for educational and research purposes (Fair use) as per U.S. Copyright law (text below). By accessing this file, all users agree that their use falls in the public domain or within fair use as defined by the copyrightlaw. We do not charge any permission or usage fees for the publication of images of material inour collections, including those provided by Princeton University Library, via our website, catalog, or directly from staff.For more information regarding copyright, credit, and citation please review the information onthe Copyright, Credit and Citation Guidelines page of our website.
https://library.princeton.edu/special-collections/policies/copyright-credit-and-citation-guidelines. If you have any questions, including how to order higher-quality images from the collection, contact Special Collections Public Services via our Ask Us! Form. https://library.princeton.edu/special-collections/ask-us

Doctoral Dissertations, 2011-2026 Neuroscience

2025

Improved Rodent Behavioral Shaping: Modeling and Mitigating Timing Errors in Perceptual Decision-Making Tasks

Breda, Jessica Ruth

Princeton, NJ : Princeton University

http://arks.princeton.edu/ark:/88435/dsp0170795c11r

U.S. Copyright law: The copyright law of the United States (Title 17, §108, United States Code) governs the making of photocopies or other reproductions of copyrighted material. Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction is not to be used for any purpose other than private study, scholarship, or research. If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of fair use that user may be liable for copyright infringement. This institution reserves the right to refuse to accept a copying order if, in its judgment, fulfillment of the order would involve a violation of copyright law. https://www.copyright.gov/title17/title17.pdf

Improved Rodent Behavioral Shaping: Modeling and Mitigating Timing Errors in Perceptual Decision-Making Tasks

Jessica Ruth Breda

A Dissertation

Presented to the Faculty

of Princeton University

in Candidacy for the Degree

of Doctor of Philosophy

Recommended for Acceptance by

Princeton Neuroscience Institute

Adviser: Carlos D. Brody

Septemeber 2025

© Copyright by Jessica Ruth Breda, 2025.

All rights reserved.

Abstract

Training animals to perform cognitive tasks is a critical component of neuroscience research as it enables detailed behavioral and neurophysiological study under controlled conditions. Many tasks require animals to withhold movement during stimulus presentation to clearly signal attention and minimize motion-related artifacts in neural recordings. In rodent studies, this fixation is typically accomplished by maintaining a nose-poke in a central port throughout the stimulus presentation. If an animal prematurely terminates fixation before the response cue, the trial is considered a "violation". Typically, these events are excluded from all analyses, yet they account for a significant fraction of trials and are correlated with worse performance. This dissertation presents a detailed analysis of these violation events and proposes novel training strategies to mitigate them. In the first chapter, I extend a traditional binary decision-making model to explicitly incorporate violations as a distinct behavioral outcome. I find that past violation history predicts future violations. Moreover, these events negatively interact with stimulus sensitivity, thereby degrading task performance. Building upon these insights, in the second chapter I develop a novel training curriculum, FixGrower, that is designed to efficiently and effectively shape fixation behavior. Compared to legacy training methods, FixGrower significantly reduces training duration while achieving lower violation rates. This curriculum proves robust across various tasks and rodent species, demonstrating its potential as a field-wide standard. Collectively, this work represents a comprehensive characterization of violations and provides optimized methodologies with theoretical support to minimize them. These advancements improve experimental efficiency while enhancing the interpretability and reliability of behavioral and neurophysiological data. Lastly, this work provides a foundation to bridge basic neuroscience with clinical insight to elucidate the mechanisms of impaired motor control that arise in disorders like Parkinson's and Huntington's diseases.

Dedicated to Wib & Papa

for the gift of wonder

Acknowledgements

Finishing a PhD is likely one of the most difficult things I will do in my life, yet to do so is simple–just keep showing up each day and have an iota of belief that you can do it. While it is striking to witness over five years of work pile up into a hundred pages of writing, what is even more striking is thinking of all the people who supported this work and made it possible. First, thank you to my advisor Carlos Brody for giving me the freedom to follow the projects I was interested in. Moving away from Princeton during the 5th year of my PhD certainly wasn't part of the plan, yet you consistently supported my endeavors as I adapted to different life circumstances and I am forever grateful for your patience. Additionally, I am so grateful I got to witness you teach– I will never forget the time you showed up to MAT314 and decided to re-do an *entire lecture* because you watched the video and didn't think you were presenting it well. Your humility combined with your excitement to share knowledge is inspiring, and I hope to foster these same skills as I grow in my career.

Second, thank you to my informal second advisor Jonathan Pillow for taking the time to support the work I present in my first chapter. I am appreciative of the many meetings I had with you on the whiteboard walking through the fundamentals of GLMs and modeling more broadly. I approached you with a semblance of an idea and you helped me develop it far beyond what I thought was possible. While I certainly felt like an imposter, you never doubted my abilities and have greatly shaped the trajectory of my career because of this. I'd also like to thank my third thesis committee member Ilana Witten. Despite the fact we only meet yearly, you always quickly understood how the project was evolving and provided clear, concise and applicable feedback. I feel extremely fortunate to have gotten your insights in these meetings.

Next I'd like to thank the members of the Brody lab (and again to Carlos for bringing together such a fantastic group of people). First, to Chuck and Jovana for keeping everything

pieced together and moving along. I truly don't know what we would do without you and I feel extremely fortunate to have gotten a chance to work with and learn from you both. To the animal technicians for making these experiments possible especially– Emily R, Andres, Jess M, Jess T, Jack and Klaus.

To the senior graduate students, especially Sarah Jo, Diksha and Tyler for providing insights on projects, career, and the hidden curriculum of graduate school. Sarah Jo- thank you for geeking out with me on code and always being willing to answer my modeling questions. Diksha- thank you for reminding me I could learn anything in 3 years and I should just do what I want even if I don't know how yet. And Tyler, thank you for helping me focus and articulate my interests while helping me believe that they were, in fact, *interesting*. My second chapter wouldn't exist if it wasn't for you helping me realize I had a viable project sitting in my hands that checked all the boxes. I hope our paths continue to cross and we wind our ways through our next career steps.

To the postdoctoral researchers, especially Emily D and Julie. Emily– thank you for mentoring me during the first two years in the lab. You always, always believed in me and at some point, that got through and I started to believe in myself. Your dedication to making sure I had the resources I needed to succeed was (and still is) unyielding and your scientific critique was consistent, clear and crafted. I am so thrilled we get to bring things full circle and you are on my committee. Julie– thank you for mentoring me during the last year in the lab. I absolutely love working together and I am continually so impressed with your ability to decipher the massive amount of information I present you with and provide calm, clear and tactful insights. Your wise advice completely shaped my second chapter and I can't wait to publish it together. Thank you so much for your time and dedicated attention– I am excited to return the favor and see what we can create with HigherInf this year.

Finally, to Wynne and Jorge- we've really been through it all. From cleaning caked-on mouse feces to walking the beaches in Portugal, I couldn't have asked for two better people

to work with. Jorge- I have learned so much from you from the basics of convolution to code architecture. Thank you for always being so patient and willing to teach. Wynne- thank you for making me think deeper, and then deeper still. This thesis, and much of my work, has taken a completely different form because of your thoughtful comments. You both have pushed me to be a better scientist, developer and thinker. I will truly miss working with you and hope that our reunion in Salt Lake last month was only the first <3.

I would also like to thank people in the Princeton community outside of the Brody Lab. To Stefan Oline- thank you for teaching me the magic of the reverse calendar and to never optimize something you won't use. To Alex Riordan- thank you for your support and belief in starting a modeling project with Jonathan and the many stimulating conversations that followed. To Javier Masís- thank you for thinking deeply about animal training with me and helping me realize the potential of a curriculum-based project. To Iris Stone for taking the time to walk me through your work figure by figure and always having fantastic advice on navigating personal and professional next steps. To Caroline Jahn for teaching me to pace myself and strive to work at 80% on average. Thank you Olga Russakovsky and Ryan Adams in the Computer Science department for teaching some of the most interesting and informative classes I have ever taken. Finally, to my cohort, especially Addie, Eleni– that first year was wild and I am *so* glad we did it together.

Additionally, thank you to the community at Westminster University that instilled a love for neuroscience, computation and research especially, Brian Avery, Lesa Ellis, Sean Raleigh, and Will Deutchman. And most importantly to my undergraduate advisor Krista Todd– the opportunities you provided me with were unparalleled. I'm so grateful you didn't take that job at Home Depot (yet). You shaped the way I think and write for the better, despite the fact that we will almost, always, forever disagree in some way hehe.

Thank you to the Advanced Analytics group at Cedars Sinai for providing me with an opportunity to test myself in industry (and the privilege to finish my PhD with job security!).

To Matt Wells, Tod Davis, Ben Pang and Dr. Melissa Wong and the rest of the team– I am thrilled to start my next chapter working with you all and pushing the boundaries in health technology while making sure we enjoy the process.

Outside of research, thank you to all my amazing friends and family. To all the friends who came over to 217 Nassau deck hangs for dinners, birthday parties and celebrations. To the misfits– Eleni, Mayank, and Iris– for a covid Christmas to remember. To friends far and wide from all phases of life and who ventured to Brazil for our wedding. To Margot & Noah and Rosanise, for finding the time to plan adventures together amidst our busy grad school schedules. To Bob, thank you for giving us a place to land in Salt Lake. To Rosalyn, for being the best roommate and one of my dearest, wisest friends. To Rach and Nat- you are my rocks and I love you more than words. Thank you to my therapists– Beth Doherty, for giving me the space to grow and learn in all dimensions, and Lisa Gilbert, for helping me evolve as a partner, friend and person. To my mom and dad- Kathy and Dean. Thank you for always trusting my ability to figure it out and take the next best step even if you didn't exactly know what I was doing. To my brother Alex, I love you man. To my aunt Katie and uncle Eddie for deeply understanding the struggles of graduate work and love for the outdoors. To my grandparents, Wibby and Papa, for instilling a love of learning. From planting tulips to soldering circuit chips- my childhood was forever enhanced because of your love and this PhD is a culmination of many of the lessons you taught me.

Finally, and most importantly, to the love of my life Andre Beukers who I was fortunate enough to meet one month into my PhD. Life with you is a chaotic, beautiful journey. I feel so lucky to have shared time in Princeton together and now in this new phase of life in Salt Lake. You've influenced these projects innumerable ways from technical concepts to strategy to support. I wish I was less stubborn and followed your advice more often– just know that *I know* that. You are one of the deepest thinkers and most loyal companions. With so much gratitude– I love you, *te amo*.

Contents

Abstract 3

Acknowledgements 5

1 Introduction 12

1.1 Background 12

1.2 Behavioral Shaping & Violations 13

1.3 Chapter Overviews 14

2 Modeling of Violation Events 16

2.1 Introduction 17

2.2 Results 18

2.2.1 Parametric Working Memory Task and Violations 18

2.2.2 Violation History Predicts Future Violations 20

2.2.3 Violations Negatively Interact with Stimuli and Decrease Performance 24

2.2.4 Stimulus–Violation Interaction Features Improve Multinomial Model Over Binary 26

2.3 Discussion 27

2.3.1 Findings 27

2.3.2 Limitations & Future Directions 28

2.4 Methods 30

2.4.1 Dataset 30

2.4.2 Training Curriculum 31

2.4.3 Criteria for Determining Trained Behavior 34

2.4.4 Autocorrelation Analysis 36
2.4.5 Models 36
2.4.6 Feature Design & Choice Labels 38
2.4.7 Model Fitting 39
2.4.8 Model Comparison 40
2.4.9 Psychometric Analysis 41
2.4.10 Data & Code Availability 42

3 Mitigation Violations & Increasing Training Efficiency 43
3.1 Introduction 44
3.2 Results 46
3.2.1 Fixation Curricula Comparison 46
3.2.2 FixGrower Decreases Training Time 49
3.2.3 FixGrower Yields Low and Predictable Violation Rates 50
3.2.4 FixGrower Generalizes Over Time, Tasks and Species 53
3.2.5 Curriculum Features Modulate Vigor, Engagement Stability, and Motor Control During Learning 57
3.3 Supplemental Figures 66
3.4 Discussion 73
3.4.1 Findings 73
3.4.2 Rationale 75
3.4.3 Limitations & Future Directions 78
3.5 Methods 80
3.5.1 Animal Subjects 80
3.5.2 Behavior Training 81
3.5.3 Curricula 83

3.5.4 Excluded Animals 98
3.5.5 Simulated Fixation Growth 99
3.5.6 Trial Initiation Time 101
3.5.7 Statistical Analyses 101
3.5.8 Data & Code Availability 102

4 Conclusions 103
4.1 Summary 103
4.2 Future Directions 104
4.2.1 Violations as a Disengaged State: Toward Latent & Dynamic Models 104
4.2.2 A Standard Protocol for Shaping Fixation Behavior 104
4.2.3 Elucidating the Neural Basis of Violations 105

Bibliography 107

1

Introduction

The work in this dissertation focuses on characterizing, understanding, and preventing premature motor errors in rodent perceptual decision making tasks. I will begin this introduction by giving a high-level overview of the motivation for this thesis. I will then provide specific context to key ideas in the field that relate to this work the as well as an overview of each project.

1.1 Background

Consider a concert pianist preparing a challenging piece for a recital. Successful performance requires more than simply knowing which notes to play; the pianist must also master timing—knowing exactly when to press, sustain, or crucially, withhold each keystroke. Even the most knowledgeable pianist, equipped with perfect recall of musical scales, will fail to produce the recital piece if they cannot appropriately manage pauses or rest periods. Thus, the interplay of note identity and temporal motor control is critical to a successful performance.

In neuroscience and psychology, a parallel scenario unfolds in the laboratory, where animals are trained to perform perceptual decision-making tasks. Animals must learn both the identity of stimuli, analogous to recognizing musical notes, and the precise timing of their motor responses, determining when and how to act to obtain rewards. Just as a pianist translates musical notation into accurately timed keystrokes, trained animals must map sensory stimuli onto controlled and timely actions. If they do not master the temporal structure of the task, successful performance becomes impossible.

Although the discrimination of sensory stimuli has been extensively studied in these

tasks, comparatively less attention has been devoted to understanding and optimizing the timing-related aspects of motor control. For example, rodents trained on perceptual tasks are often required to remain motionless during stimulus presentation, initiating movement only after a designated response cue. Premature movements in these tasks, much as a pianist accidentally pressing keys during a rest, reflect timing errors that significantly disrupt task performance. In this thesis, I specifically address these motor-timing errors, investigating their underlying causes, examining how they interact with stimulus-driven responses, and developing strategies to mitigate them.

1.2 Behavioral Shaping & Violations

Formalized training of animals to perform laboratory tasks was popularized by B.F. Skinner, who pioneered the field of operant conditioning (Skinner, 1951). His foundational work established principles that continue to underpin contemporary behavioral shaping practices in neuroscience and psychology, wherein systematic reinforcement and progressive sub-tasks are used to guide animals toward learning complex cognitive behaviors.

A common subset of these tasks involves unrestrained, freely moving animals, who must maintain fixation at a designated location throughout stimulus presentation until a specific response cue is given. Such fixation requirements verify task engagement, ensure consistent stimulus delivery, and minimize movement-related artifacts that can interfere with neurophysiological data quality. Trials in which animals prematurely move and terminate fixation-termed "violations"—are routinely discarded from behavioral and neural analyses. Notably, in studies which report their frequency, violation trials are numerous. They typically account for one third of attempted trials in fully trained animals, yet, precise, field-wide estimates are scarce due to infrequent reporting in published literature (Ashwood et al., 2022; Constantinople et al., 2019; Cruz et al., 2022; Jang et al., 2024; Jaramillo & Zador, 2014; Licata et al., 2017; Pisupati et al., 2021). Additionally, anecdotal evidence from our laboratory

suggests that the best performing animals have the lowest violation rates. Thus, reducing violation rates is crucial for improving both data yield and task performance.

Critically, despite the clear benefits of reducing violations, current methods for training fixation behavior remain heuristic, inconsistently documented, and lacking quantitative characterization. As a result, we have only a limited understanding of why these errors arise, and precisely how they influence learning and performance. Addressing these knowledge gaps is essential—not only to accurately interpret animal behavior but also to develop standardized and efficient training protocols. Therefore, in this thesis, I focus explicitly on quantifying, understanding, and systematically reducing these premature motor errors called violations.

1.3 Chapter Overviews

In the first chapter, I investigate the behavioral impact of violations in trained animals performing a perceptual decision-making task. I use a dataset from animals trained on a Parametric Working Memory (PWM) task from Akrami et al. Akrami et al. (2018). In this task, rats compare the loudness of two tones separated by a delay while fixating for durations of 2-6 seconds. I systematically analyze the temporal structure and predictors of violation trials. I extend traditional binary choice models into a multinomial logistic regression framework that explicitly incorporates violations. My analyses demonstrate that violations impair subsequent task performance by reducing animals' sensitivity to stimuli, highlighting their significant influence on perceptual decision-making.

Building on these insights, in the second chapter I focus on reducing violations through improved behavioral shaping protocols. I design a novel training curriculum supported by theories in operant conditioning, reinforcement learning, and curriculum learning with the goal of efficiently establishing stable fixation behavior. When compared to the previous method used in our lab, this curriculum significantly decreases training time by 61%, reduces violation rates below those typically reported in the literature, and maintains robust

performance across both rats and mice. Importantly, these results hold even when complex stimulus rules are introduced.

Together, these projects highlight the importance of using explicit computational models of violations to guide behavioral shaping protocols in perceptual decision-making tasks. By carefully characterizing these premature movement errors and integrating them into formal behavioral models, this work deepens our understanding of their origins and impacts. Furthermore, targeted behavioral shaping informed by this enhanced understanding significantly improves both training efficiency and reliability.

2

Incorporating violation events into decision models reveals structured behavioral dynamics and improves prediction

2.1 Introduction

Understanding the factors that shape decision-making is a central aim in neuroscience, as it offers a window into which features are encoded, integrated, and acted upon by the brain. Formal behavioral models offer a principled way to quantify how stimuli and internal variables shape a subject's choice. Moreover, these models can guide neural analyses by clarifying which features are behaviorally relevant.

In the context of perceptual decision-making tasks, a common approach has been to model binary choices—e.g., left versus right responses—using logistic regression or its dynamic extensions. In a Parametric Working Memory (PWM) task, for instance, rats compare the loudness of two temporally separated stimuli and report their decision at lateral ports (Akrami et al., 2018). Results from binary models have revealed significant effects of both stimulus features and trial history, including prior stimuli, choices, and outcomes that can evolve over training (Akrami et al., 2018; Roy et al., 2021).

However, such models commonly ignore a substantial class of trials: *violations*, in which the animal initiates a trial but prematurely aborts it before the allowed response period. Violations account for approximately 15–30% of trials depending on the task and species (Akrami et al., 2018; Ashwood et al., 2022; Constantinople et al., 2019; Cruz et al., 2022; Jang et al., 2024; Jaramillo & Zador, 2014; Licata et al., 2017; Pisupati et al., 2021). And, despite their prevalence, violations are rarely modeled or systematically analyzed. Yet, recent findings suggest they are behaviorally and neuro-physiologically meaningful. For example, Jang et al. (2024) demonstrated that exciting acetylcholine neurons in the Dorsal Medial Striatum led to a significant increase in violations and a decrease in sensitivity to the task structure. These findings raise important and unaddressed questions about the origins of violations and their influence on subsequent behavior and task performance.

To fill this gap, we extend a binary logistic regression model to a multinomial frame-

work that models all three possible outcomes in the PWM task: left choice, right choice, and violation. Using data from trained rats, we first show that violations are frequent and negatively associated with task performance. We then identify features that predict violations and find that recent violation history—especially when modeled with an exponentially weighted filter—is a strong predictor. We also find that violations degrade performance on subsequent trials by reducing stimulus sensitivity, which we capture using interaction terms between stimuli and violation history. Finally, we show that including these features improves predictive accuracy even when evaluating the model on left and right choices only.

Together, these findings demonstrate that violations are structured, history-dependent events that shape ongoing decision behavior. Modeling them explicitly yields a more complete view of task engagement. Our work provides a principled extension of current behavioral models and opens the door for future analyses that incorporate multi-class choice behaviors into theories of learning, motivation, and decision making.

2.2 Results

2.2.1 Parametric Working Memory Task and Violations

Our analyses focus on data collected by Akrami et al. (2018) on rats performing a Parametric Working Memory (PWM) task. The structure of the task is as follows: a rat center-pokes and maintains fixation; stimulus A (s_a) sound plays, followed by a variable discrete delay of 2–6 seconds, then stimulus B (s_b) sound plays, followed by an additional brief delay, and finally a go-cue sound that indicates the animal can end fixation and report its choice(**Fig. 2.1a**). Rats must perform a delayed comparison of sound loudness between s_a and s_b reporting their decision at one of two lateral choice ports (e.g., if s_a is louder, go right). If an animal breaks center-poke fixation before the go-cue, the trial is aborted and considered a *violation*. We were interested in studying the effects of violations, as these trials are typically dropped from all behavioral and neural analyses and their impact is unknown.

To do this, we first recreated the dataset to address inaccuracies in how these trials were originally tracked in the published version (See *Methods*, **Fig. S2.1, S2.2**). We further filtered this dataset to include only sessions after animals were considered trained on the task (See *Methods*, **Fig. S2.3**; 1,866,003 trials, 5,210 sessions, N = 16 rats). Achieving this training criterion required substantial time, averaging 242 ± 116 sessions (8.0 ± 3.8 months).

Animals performed 358 ± 140 trials per session on average (**Fig. 2.1b**) with an average session performance of 81.9 ± 6.2% on valid trials (**Fig. 2.1c**). Notably, violations accounted for an average of 17.4 ± 9.2% of session trials (**Fig. 2.1d,e**). To investigate the relationship between average session violation and hit rates across animals, we conducted a simple linear regression, which revealed a negative association, explaining 19.0% of the variance (**Fig. 2.1f**). To account for within-animal variability, we next fit a linear mixed-effects model across all sessions with random intercepts for each animal. This model revealed a statistically significant negative relationship between session violation and hit rate (β = -0.342 ± 0.020; z = -17.3; $p < 0.0001$; **Fig. 2.1g**). Together, these results demonstrate that violations constituted a substantial portion of trials and were negatively associated with task performance.

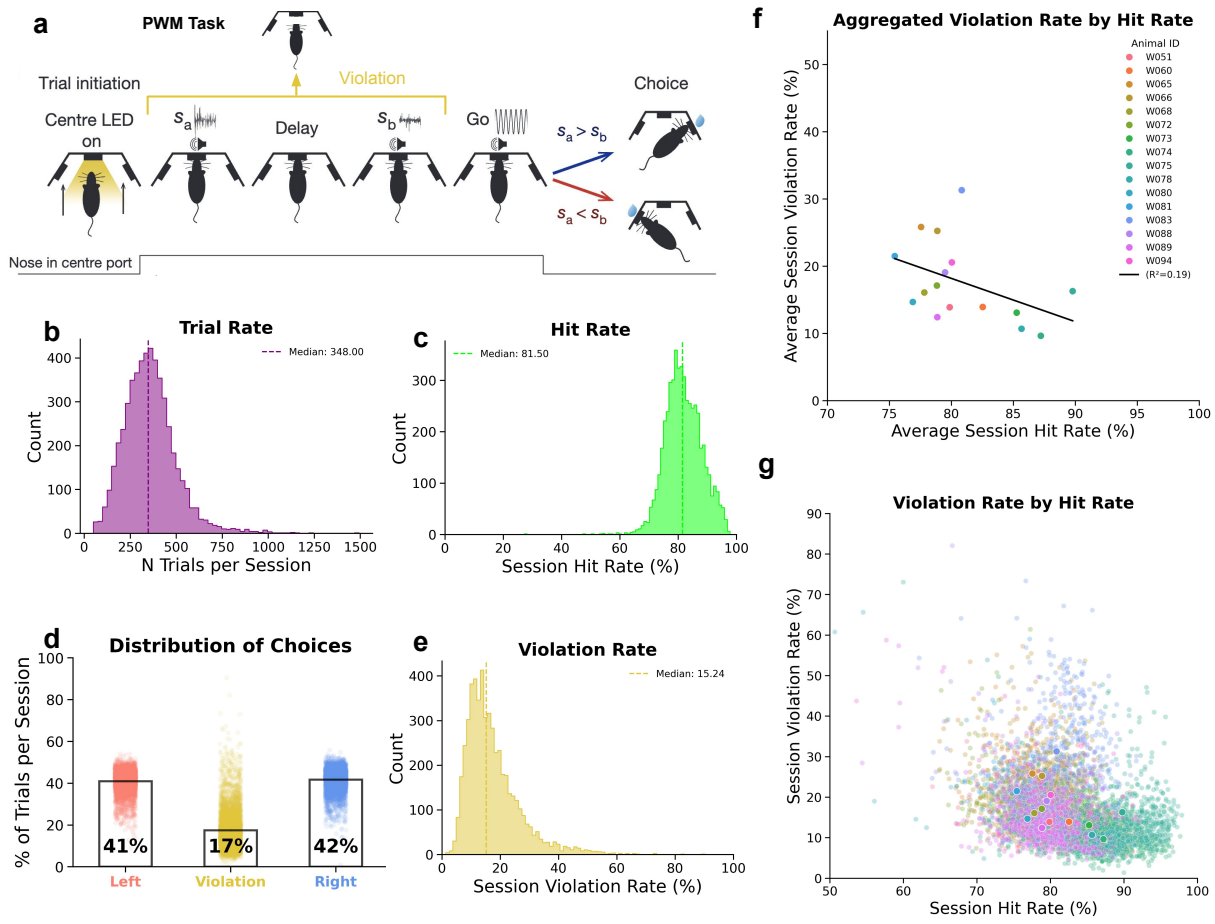


Figure 2.1: *Task and trained dataset overview.* **(A)** Overview of Parametric Working Memory (PWM) task from Akrami et al. 2018 with violation event annotation **(B)** Histogram of number of trials per session across all animals in trained dataset. Dashed line indicates median (N = 16 rats, n = 5,251 sessions). **(C)** As in (B), but for session hit rate on valid (non-violated) trials. **(D)** Percentages of trials in each session that were left, right or violation choices for the trained dataset. **(E)** As in (B) but for session violation rate. **(F)** Average session rate vs. session violation rate for each animal with the fit linear regression model in black. **(G)** Hit vs. violation rate for each animal and session that was used for the mixed-effects linear model. Shares legend with (F).

2.2.2 Violation History Predicts Future Violations

To model trial outcomes, we extended a binary logistic regression framework used in prior work (Akrami et al., 2018; Roy et al., 2021) (**Fig. 2.2a**) into a multinomial logistic regression, allowing us to predict left, right, or violation choices (**Fig. 2.2b**). This approach estimates a separate weight for each feature-class combination which was essential for iden-

tifying predictors of violations.

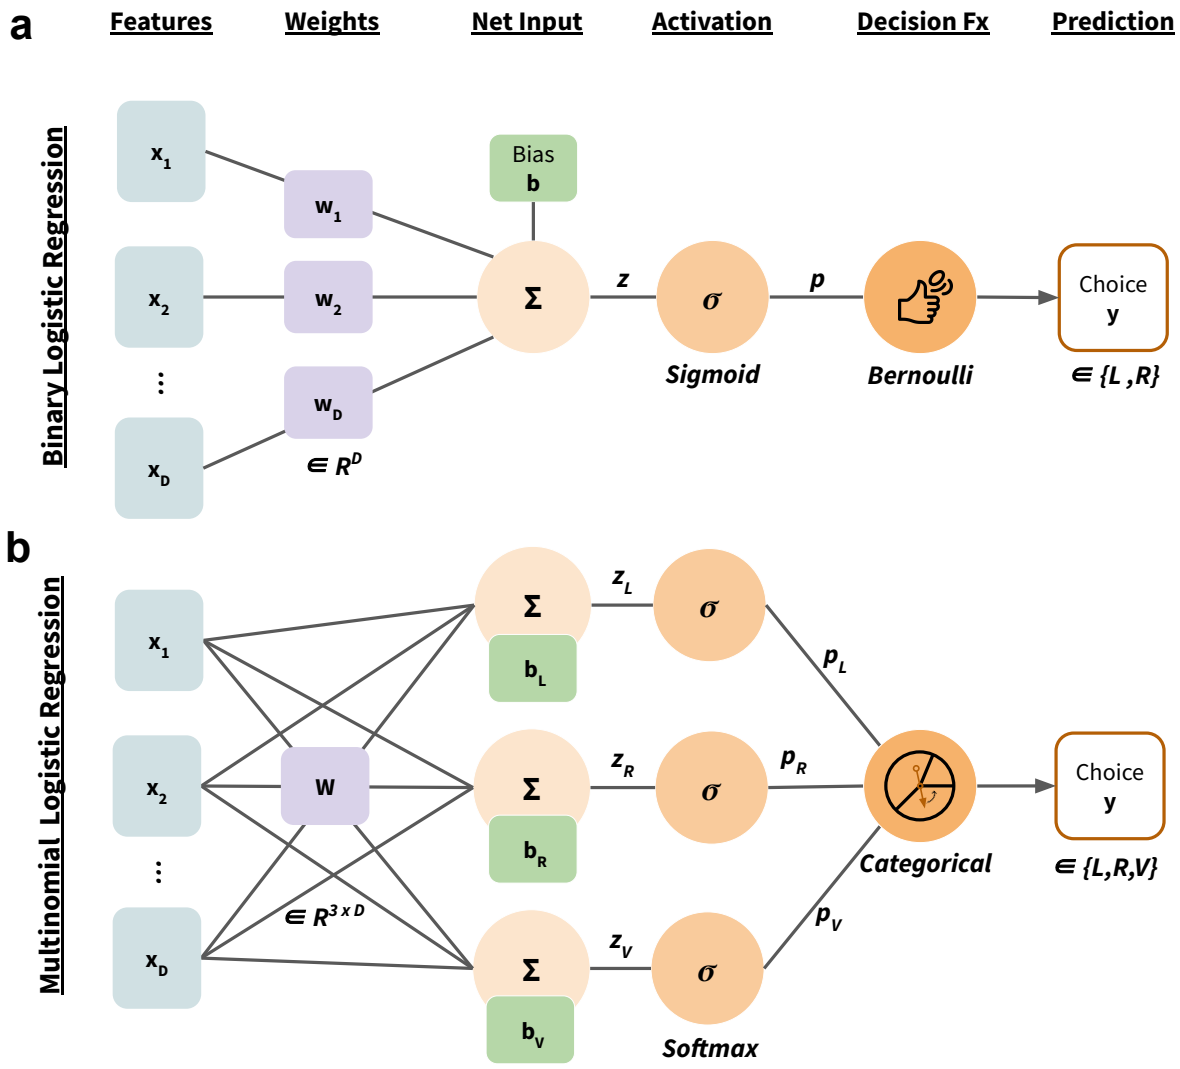


Figure 2.2: *Binary and multinomial logistic regression models.* **(A)** Diagrammatic overview of binary logistic regression model. A feature vector $\mathbf{x} \in \mathbb{R}^D$ is linearly combined with a weight vector $\mathbf{w} \in \mathbb{R}^D$ and bias term b to produce a scalar logit $z = \mathbf{w}^\top \mathbf{x} + b$. This is passed through a sigmoid activation function $\sigma(z)$, to produce a probability $p \in (0, 1)$, which is sampled from Bernoulli(p) to yield a binary choice $y \in L, R$. **(B)** Multinomial logistic regression with 3 classes- Left, Right and Violation. the same input vector $\mathbf{x} \in \mathbb{R}^D$ is mapped via a weight matrix $W \in \mathbb{R}^{3 \times D}$ and class-specific bias terms to produce a vector of logits $\mathbf{z} = W\mathbf{x} + \mathbf{b}$. These logits are passed through the softmax function $\sigma(\mathbf{z})$ to yield a probability distribution $\mathbf{p} = (p_L, p_R, p_V)$. The predicted label is then sampled from Categorical(p_L, p_R, p_V) to yield a choice $y \in L, R, V$.

Preliminary inspection of behavior suggested that violations often occurred in groups. To quantify this, we computed an autocorrelogram of violations and compared it to a shuffle-based null distribution. We found significant differences at nearly all lags, confirming that violations are temporally clustered (**Fig. 2.3a**).

To capture this effect in our model, we introduced a binary violation history regressor: a single-trial lag indicating whether the previous trial was a violation (**Fig. 2.3b(i)**). This feature had a large positive weight for the violation class, suggesting that rats were more likely to violate if they had just done so on the preceding trial (**Fig. 2.3b(ii)**).

We then extended this feature to capture longer-range effects by constructing an exponentially filtered violation history regressor (**Fig. 2.3c(i)**). This feature was created by convolving a binary history of previous violations with a decaying exponential kernel (see *Methods*, **Fig. S2.4**), allowing us to model the influence of multiple past trials. The decay time constant (τ) was optimized separately for each animal and ranged from 1 to 40 trials, with a median of 2 (**Fig. 2.3d,e**). This regressor also carried a large positive weight for the violation class (**Fig. 2.3c(ii)**), reinforcing the notion that recent violation history strongly predicts future violations.

Finally, we compared three model variants. First, was a baseline model that used a feature set previously designed for binary models of the task (**Table 2.1**; Akrami et al., 2018; Roy et al., 2021). We compared this to a baseline model with binary violation history added, and one with filtered violation history added. Adding binary or filtered violation history improved model performance by 0.012 and 0.019 bits per trial, respectively (**Fig. 2.3f**). These improvements corresponded to the data being, on average, 10^{81} and 10^{144} times more likely under the respective models for each animal. Together, these findings demonstrate that violation history, especially when distributed over several trials, strongly influences the likelihood of a violation on the current trial.

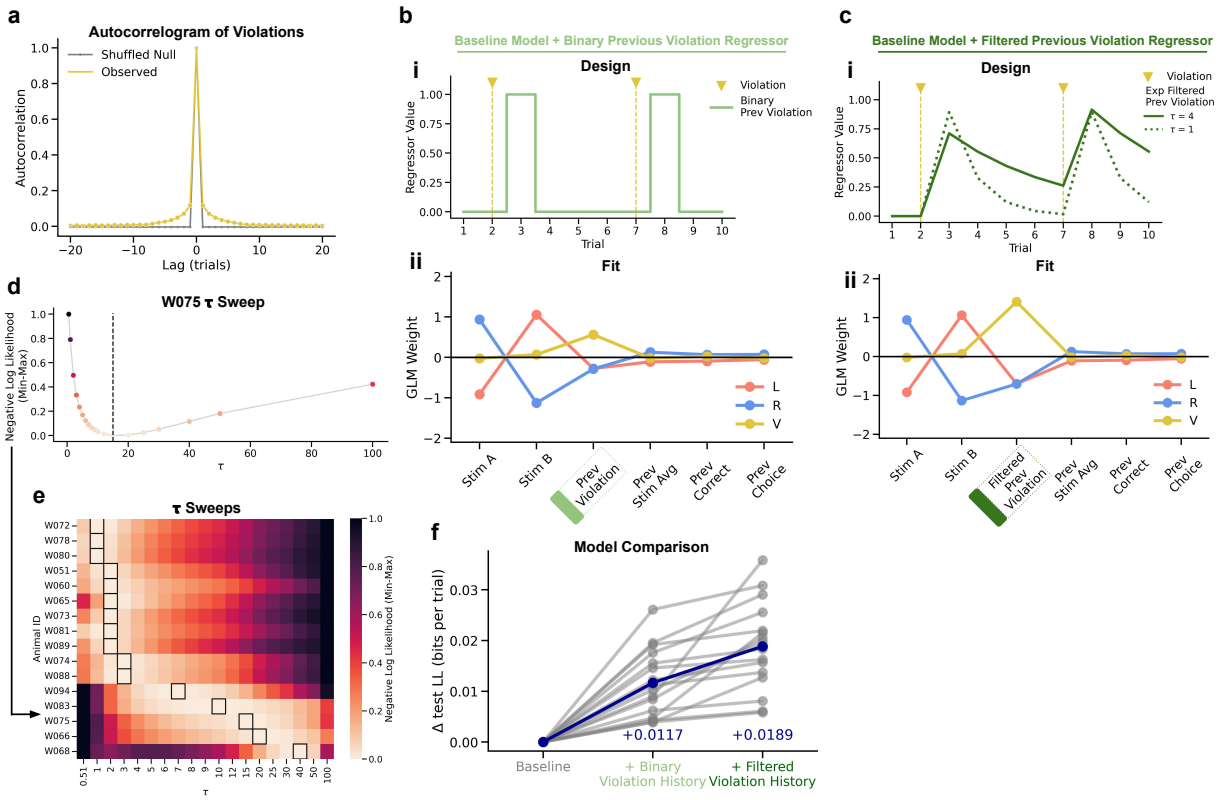


Figure 2.3: *Violation history predicts future violations.* **(A)** Autocorrelogram of violations (yellow) compared to a shuffled null distribution (gray), with 95% confidence intervals, computed over a ± 20 trial lag. Markers indicate significance from a two-tailed t-test with Bonferroni correction ($\bullet = p < 0.001$, $\times = p > 0.05$). **(B)(i)** Schematic of the binary violation history regressor (light green), indicating whether the immediately preceding trial was a violation (yellow). **(B)(ii)** Logistic regression weights from the multinomial model when including the binary history regressor (light green). Lines represent weights for each choice output averaged across all animals. **(C)(i)** Schematic of the exponentially filtered violation history regressor (dark green) for two different time constant values (τ). Solid line indicates $\tau = 4$. Dashed line indicates $\tau = 1$. The filter length spans $5 \times \tau$ trials to ensure full decay. **(C)(ii)** As in (B)(ii), but using the exponentially filtered violation history regressor (dark green). **(D)** Example sweep of τ values for a given animal showing test negative log-likelihood (*nll*) over τ. Gray dashed line marks the minimum *nll* and selected τ value for model in (C). Note, *nll* is scaled between 0 and 1 to allow for easy visual comparison across animals. **(E)** Heatmap of scaled test *nll* over τ for each animal. Gray box indicates minimum *nll* and selected τ for model in (C). **(F)** Change in model performance (test set log-likelihood, in bits per trial) relative to the baseline multinomial model. Gray lines represent individual animals (N = 16). Dark blue line and text indicate the mean.

2.2.3 Violations Negatively Interact with Stimuli and Decrease Performance

Having identified the variables that predict violations, we next examined how violations relate to overall task performance. Specifically, we asked whether violations reflect isolated motor errors or if they also impair stimulus-driven performance on subsequent trials. We observed that performance, measured by percent correct, significantly decreased following a violation trial (Not Previous Violation: mean ± SD 81.3 ± 4.1% vs. Previous Violation: 77.7 ± 3.5%, N = 16, $n_{sessions}$ = 5,210; Paired t-test, t = -6.121, p = 1.96 x 10-05; **Fig. 2.4a**).

To test whether our model could account for this effect, we introduced a stimulus–violation history interaction regressor. We constructed this by multiplying each stimulus regressor by a violation history regressor (binary or filtered; **Fig. 2.4b(i)**). The resulting models revealed interaction weights that were opposed in sign to the corresponding stimulus weights, suggesting that violations suppressed stimulus sensitivity and thus impaired performance (**Fig. 2.4b(ii)–(iii)**).

To assess whether this mechanism could reproduce the behavioral effect, we used each animal's fit weights to simulate psychometric curves for trials following a violation versus those that did not (see *Methods*). The simulated post-violation curves showed a significant decrease in slope (Not Previous Violation: mean ± SD 0.512 ± 0.104 vs. Previous Violation: 0.445 ± 0.090; Paired t-test, $t = -5.557$, $p = 5.49 \times 10^{-5}$) and a significant increase in lapse rate (Not Previous Violation: mean ± SD 0.041 ± 0.023 vs. Previous Violation: 0.067 ± 0.039; Paired t-test, $t = 4.603$, $p = 3.45 \times 10^{-4}$), mirroring the performance drop seen in the empirical data (**Fig. 2.4c-e**).

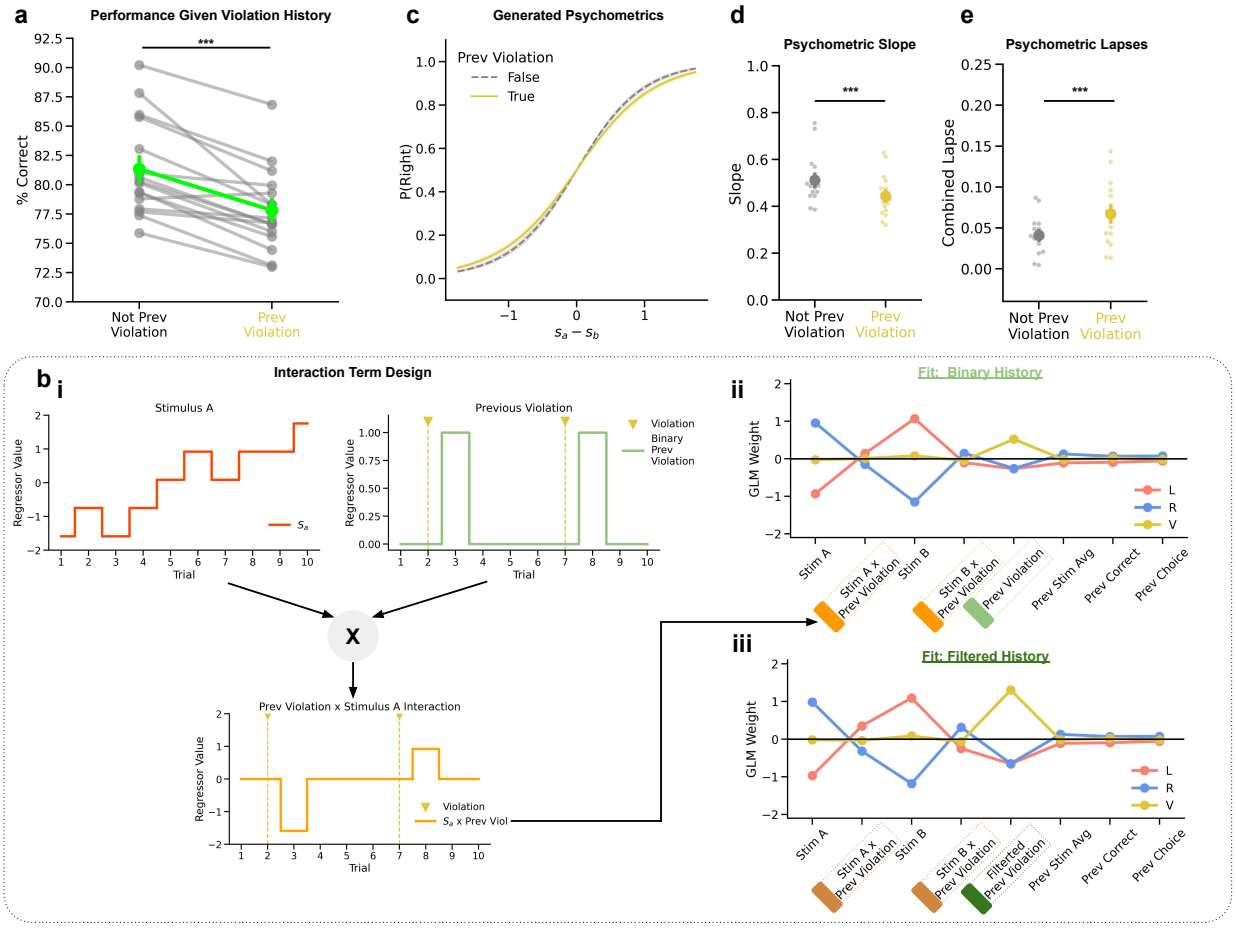


Figure 2.4: *Violation history impairs stimulus-driven performance.* **(A)** Performance measured as percent correct following non-violation versus violation trials. Gray lines indicate individual animals (N = 16). Lime green line indicates the mean. Error bars indicate standard error of the mean. **(B)(i)** Example schematic of the stimulus–by–violation-history interaction regressor (orange) for Stimulus A, formed by multiplying the standardized Stimulus A regressor (red) with the previous-violation history regressor (light green). **(B)(iii)** Similar to (ii): logistic-regression weights from the multinomial model including the filtered history term (dark green) and stimulus–history interaction terms (dark orange). **(C)** Simulated psychometric functions for trials following a violation (yellow) or not (gray), generated using each animal's fitted stimulus and interaction weights. **(D)** Psychometric slope for each condition in (C). **(E)** Combined lapse rate (upper + lower) for each condition in (C). *Statistics:* (A, D, E) Statistical comparisons were conducted within subjects using paired t-tests. *** indicates $p < 0.001$

Finally, when we compared models with and without interaction terms to the violation-history models presented in Figure 2.3 (**Fig. 2.5a**), we found that a model with filtered violation history *and* stimulus–violation history interaction terms fit best (**Fig. 2.5b**). Compared to the baseline multinomial model, this model produced a 0.020 bit-per-trial increase

in log-likelihood, which equates to the data being, on average, 10^{155} times more likely under this model for each animal. These results show that including violation-history interactions improves model fit by capturing additional variability in behavior, specifically by offsetting stimulus-driven choice tendencies on post-violation trials.

2.2.4 Stimulus–Violation Interaction Features Improve Multinomial Model Over Binary

Our final analysis asked whether explicitly modeling the violation class confers an advantage when evaluating only on left and right choices. In other words, is there any benefit over the simpler, binary model by using a multinomial model? For each rat we fit a binary logistic-regression model to left and right trials using the baseline feature set and compared it to the four multinomial variants previously described (**Fig. 2.5c**). The multinomial models were trained on left, right, and violation trials and tested on the same left and right trials as the binary model using a modified two-class softmax function to allow for direct comparison of the log-likelihoods (see Methods).

We found that when only violation history terms were included, the multinomial model performed worse than the binary model (binary violation history: –0.0006 bits/trial; filtered violation history: –0.0007 bits/trial; **Fig. 2.5d**). However, when stimulus–violation history interaction terms were included, the multinomial models performed better than the binary model. The filtered violation history with interaction terms performed best with an average improvement of 0.0007 bits per trial. This corresponds to the data being 10^6 times more likely under this model. Thus, a multinomial model exceeds a binary model only when it can capture the interaction between stimuli and violation history, underscoring the disruptive impact of recent violations on subsequent decision making.

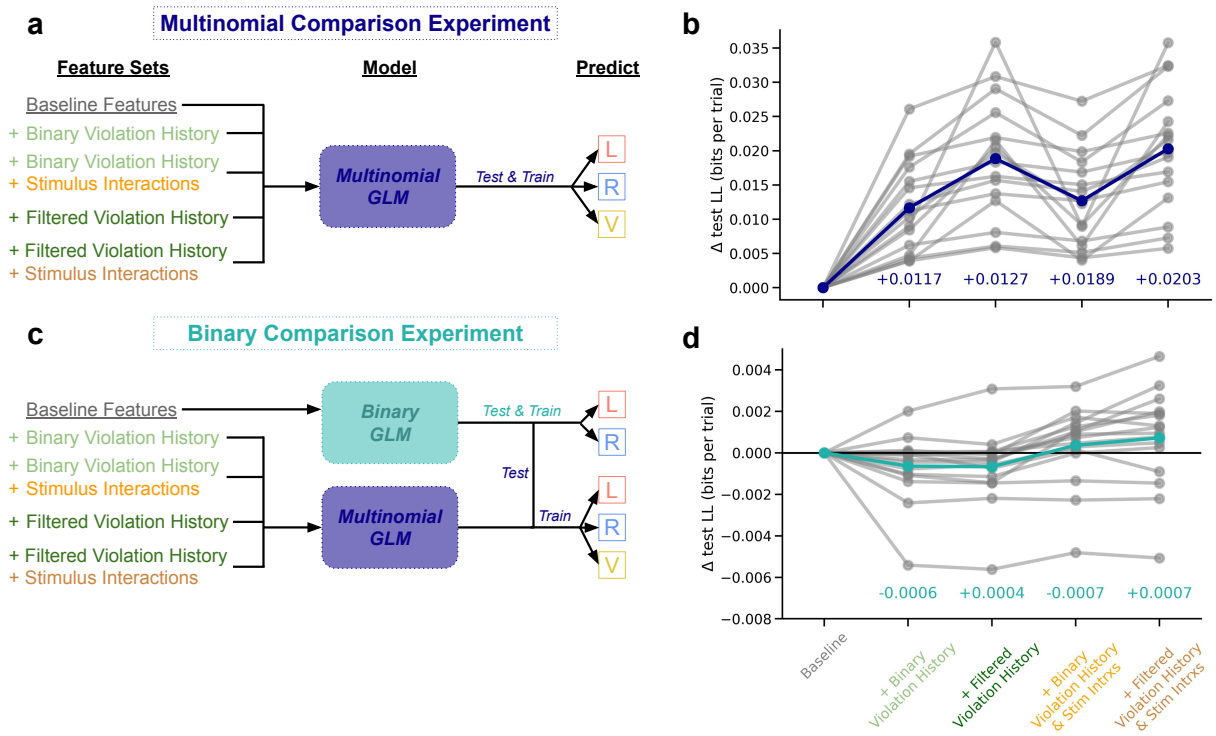


Figure 2.5: *Comparison of model performance across feature sets and evaluation paradigms.* **(A)** Schematic of the Multinomial Comparison Experiment. Multinomial GLMs (dark blue) are trained and tested on left, right, and violation trials using one of five feature sets: a baseline set (gray), violation history regressors (binary: light green; filtered: dark green), and their interactions with stimulus regressors (orange and dark orange). **(B)** Improvement in test log-likelihood (ΔLL, bits per trial) relative to the baseline features multinomial model, for each feature set in (A). Gray lines indicate individual animals (N = 16); dark blue line and text indicate the mean. **(C)** Schematic of the Binary Comparison Experiment. Binary GLMs (teal) are trained and tested on left/right trials. Corresponding multinomial GLMs are trained on left/right/violation trials but tested only on left/right trials using a reduced softmax. **(D)** Improvement in test log-likelihood (ΔLL) relative to the baseline binary model, for each feature set in (A, C). Gray lines show individual animals; teal line and text indicate the mean. The null model is binary.

2.3 Discussion

2.3.1 Findings

In this study, we extended a behavioral modeling approach to explicitly incorporate and analyze violation trials—instances in which rats prematurely break fixation during a Parametric Working Memory task (Akrami et al., 2018, **Fig. 2.1**). While traditional models of decision-making discard violation trials and focus on binary left and right choices, even

though violations constitute a substantial fraction of trials. Motivated by recent evidence that violations may reflect a tension between movement suppression and task performance (Jang et al., 2024), we developed a multinomial logistic regression framework that models left, right, and violation choices (**Fig. 2.2**).

Our results show that violations are not random lapses but structured, history-dependent events that significantly shape ongoing behavior. We found that recent violation history, particularly when modeled with an exponentially weighted filter, predicts future violations (**Fig. 2.3**). Furthermore, violations negatively interact with stimulus weights and reduce stimulus sensitivity, thereby impairing performance on subsequent trials (**Fig. 2.4**). Incorporating interaction terms between violation history and stimulus features improved model fit and captured behavioral effects observed in the data (**Fig. 2.5**). Critically, this extended model provided predictive advantages even when evaluating only left and right choices, underscoring the broader value of modeling violations explicitly in theories of decision-making.

2.3.2 Limitations & Future Directions

While the present study provides a principled extension of standard behavioral models to incorporate violations, several limitations remain that suggest directions for future work. First, our approach fits a single, static multinomial GLM to all trained sessions per animal, thereby assuming a stationary relationship between task variables and behavioral outcomes. However, prior work has shown that these feature weights can evolve both within and across sessions during learning (Cuturela et al., 2024; Roy et al., 2021). Moreover, animals can occupy multiple, discrete latent states with varying behavioral strategies and neural circuits within a session (Ashwood et al., 2022; Bolkan et al., 2022; Hulsey et al., 2024; Venditto et al., 2024). It would be of interest to extend our model to these frameworks to assess whether violations influence learning and whether they represent represent a discrete disengaged state.

A second limitation is that our current analysis treats all violation trials as a single

category, without distinguishing between those that occur early in the fixation period—before stimulus B—and those that occur later, after the second stimulus has begun. This distinction is meaningful: once stimulus B is presented, the animal has all the information needed to make a correct decision. It is plausible that these subtypes reflect distinct cognitive processes. For instance, early violations may reflect disengagement or low arousal, while late violations might arise from impulsivity near the decision point. Future work using higher-resolution behavioral data and latent state models could help disambiguate these possibilities.

Moreover, given the demonstrated impact of violations on subsequent task performance, an important next step is to examine how these behavioral fluctuations are represented in neural activity. Typically, violation trials are excluded from neural analyses, potentially overlooking meaningful sources of variability. Rather than treating them as irrelevant artifacts, violation history should be incorporated as a covariate when analyzing neural data to assess how these trials influence the neural dynamics that drive task behavior.

Finally, while violation events are informative, our findings demonstrate that they impair performance, even in fully trained animals. As such, minimizing their occurrence through curriculum design is a valuable objective in task shaping. For instance, adaptive training procedures that scaffold fixation demands could reduce the frequency of violations throughout training and promote more stable task performance.

In summary, this work highlights the importance of explicitly modeling violations in freely-moving perceptual decision making tasks. By extending a traditional binary model to account for these events, we reveal how recent trial history shapes performance. Moving forward, incorporating violations into both behavioral and neural analyses—while simultaneously designing training paradigms that limit their occurrence—offers a path toward more comprehensive and interpretable models of decision-making.

2.4 Methods

2.4.1 Dataset

Behavioral data analyzed in this study was originally collected as part of experiments described in Akrami et al. 2018, where rats performed a Parametric Working Memory task involving auditory discrimination of sound loudness. Briefly, sixteen male Long–Evans rats (*Rattus norvegicus*), aged between 6 and 24 months, were trained on the task, which required them to judge whether the first (s_a) or second stimulus (s_b) (both broadband noise stimuli varied between 60–92 dB) had greater loudness when presented with a multi-second delay between them. Rats indicated their choice by selecting either the left or right port, with correct responses rewarded with water. The original dataset is publicly available (https://tinyurl.com/PWMdata).

For the current analyses, the dataset was reconstructed from the original raw behavioral recordings to address inaccuracies identified in the previously published data, particularly concerning violation trials (originally referred to as *mistrials*; **Fig. S2.1**). Two specific issues were addressed. First, the original dataset stopped recording violation trials after session 200, which was prior to many animals reaching full training proficiency, thereby significantly limiting the data available for modeling trained behavior (**Fig. S2.1**a). Second, and more importantly, the original dataset erroneously classified certain events as violations; specifically, trials were incorrectly marked as violations if a rat did not initiate a trial within 200 seconds (**Fig. S2.1**b). Using raw files indexed using our MySQL database, the data was reprocessed and aligned to the old dataset (**Fig. S2.1**c,d), resulting in approximately a 9% reduction in labeled violation trials per session with a total of 3,034,200 trials over 9,968 session for all the animals (**Fig. S2.1**e-f).

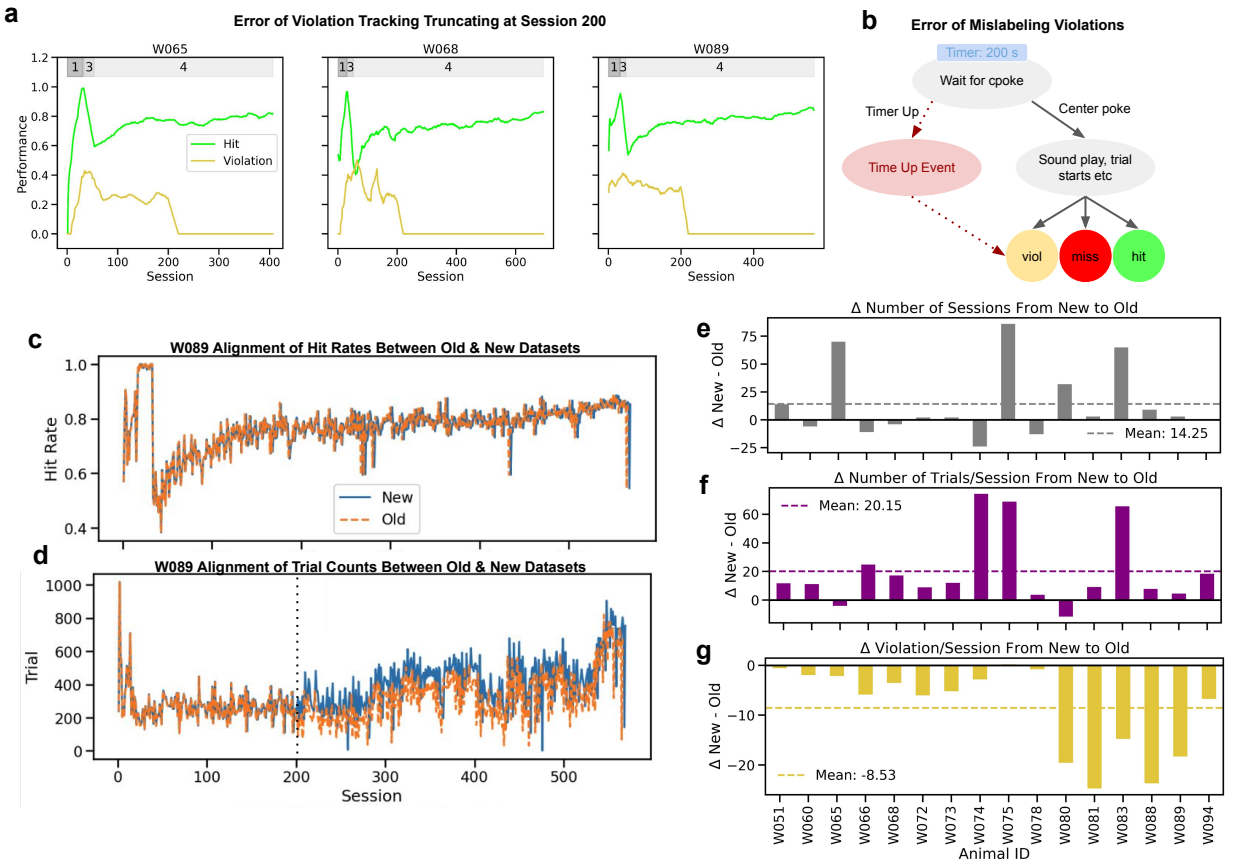


Figure S2.1: *Explanation, documentation and correction of errors in previously published data.* **(A)** Hit (lime green) and violation (yellow) data over sessions from the published dataset for three example animals showing the truncation of violation tracking at session 200. The gray bars at the top indicate the curriculum stage for the session. **(B)** State-machine logic that produced errors in mislabeling violations due to an animal not initiating a new trial within 200 seconds of the trial being ready. In the new dataset, these events are dropped. **(C)** Alignment between old (orange, dashed) and new (blue, solid) datasets with respect to hit rate for an example animal. **(D)** As in (C) but for trial counts. Note that trial counts are higher after session 200 (black dashed line) due to violation trials now being tracked. **(E)** Difference in the total number of sessions in the new dataset for all animals. **(F)** As in (E) but for trials per session. **(G)** As in (E) but for session violation rate.

2.4.2 Training Curriculum

Stage 2 “Always Reward" introduced auditory stimuli (s_a and s_b) by growing them into the fixation duration. This stage was 2 days long and animals grew from 10 ms to 8 seconds on average both days while the sounds played. Animals received rewards regardless of initial choice accuracy as long as they eventually poked in the correct port.

Stage 3 “Delayed Reward” required discrimination of auditory stimuli across two discrete delays (typically 2 and 3 seconds between stimuli). If animals answered incorrectly, they could retrieve the reward at the correct port after a penalty delay (1-5 seconds depending on the animal).

Stage 4 "Never Reward", represented the full task, requiring correct discrimination of stimuli across multiple discrete delays (commonly 2, 3, 4, and 6 seconds) to receive rewards. Lastly, in stage 4, a subset of animals (N=12) had psychometric trials designed to assess discrimination precision by varying the intensity of s_a while holding s_b constant (**Fig. S2.2**).

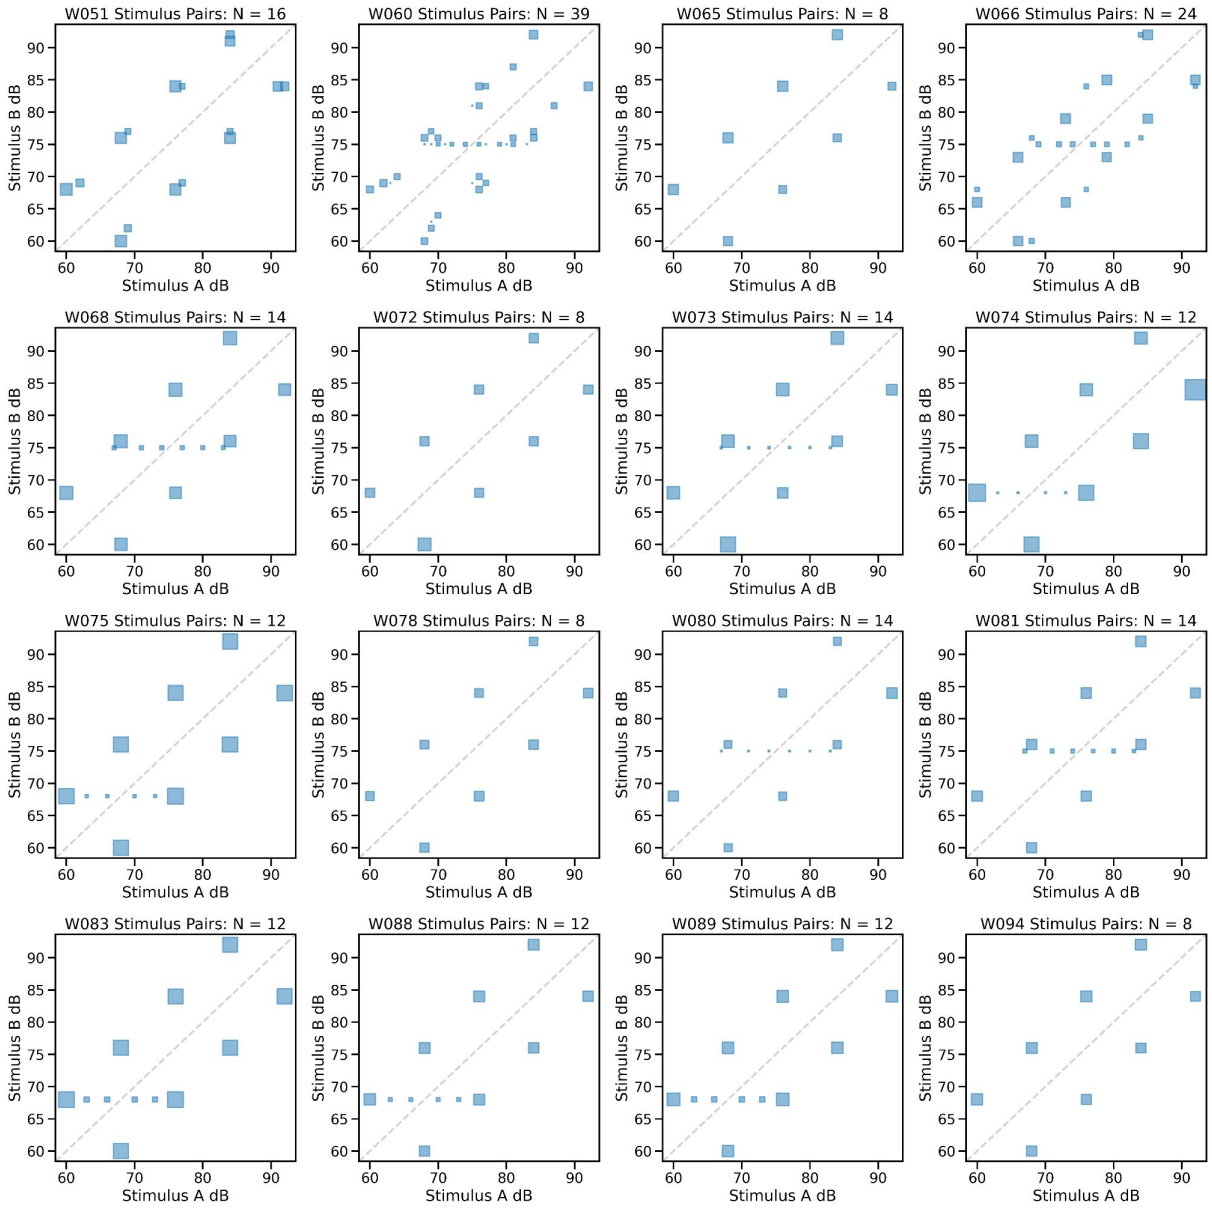


Figure S2.2: *Stimulus pairs sets for each animal.* Each square represents a stimulus pair (s_a- s_b) used in training stage 4 for a given animal. 12 of 16 animals had psychometric stimulus sets. The size of the square is represents to the number of trials the stimulus pair was presented relative to the entire dataset. Gray dashed line indicates decision boundary.

2.4.3 Criteria for Determining Trained Behavior

To determine when animals could be considered "trained" and thus exhibiting stable, stationary performance, quantitative criteria based on rolling averages of behavioral metrics were applied (**Fig. S2.3**). A rolling window analysis was conducted separately for hit rates (correct responses on *valid* non-violation trials) and violation rates (trials aborted prematurely). Rolling mean window sizes of 7, 14, 21, and 28 sessions were tested for hit rate analysis, and a window size of 21 sessions was selected to balance smoothness and responsiveness. Hit rate performance thresholds of 70%, 72.5%, and 75% accuracy were evaluated, with a threshold of 72.5% ultimately chosen as it best represented stable trained behavior. Animals were considered trained starting from the last session at which this threshold was crossed.

Violation rates were similarly analyzed using rolling mean window sizes of 7, 14, 21, and 28 sessions. A shorter rolling window of 7 sessions was selected to detect abrupt increases in violations indicative of potential behavioral instability. A threshold of 50% was applied, and animals exhibiting late-training violation rate spikes (i.e. when the animal is getting older and unstable) were truncated accordingly to exclude non-stationary behavioral data. This resulted in truncated sessions for 4 animals. In addition, we excluded any session containing fewer than 50 trials or a violation rate greater than 99% (n = 41 sessions). Only sessions and trials that met these criteria were retained for all subsequent analyses, thereby ensuring that the final dataset accurately reflected stable, trained behavior. This produced a dataset with 5,210 sessions across 16 animals totaling 1,866,003 trials.

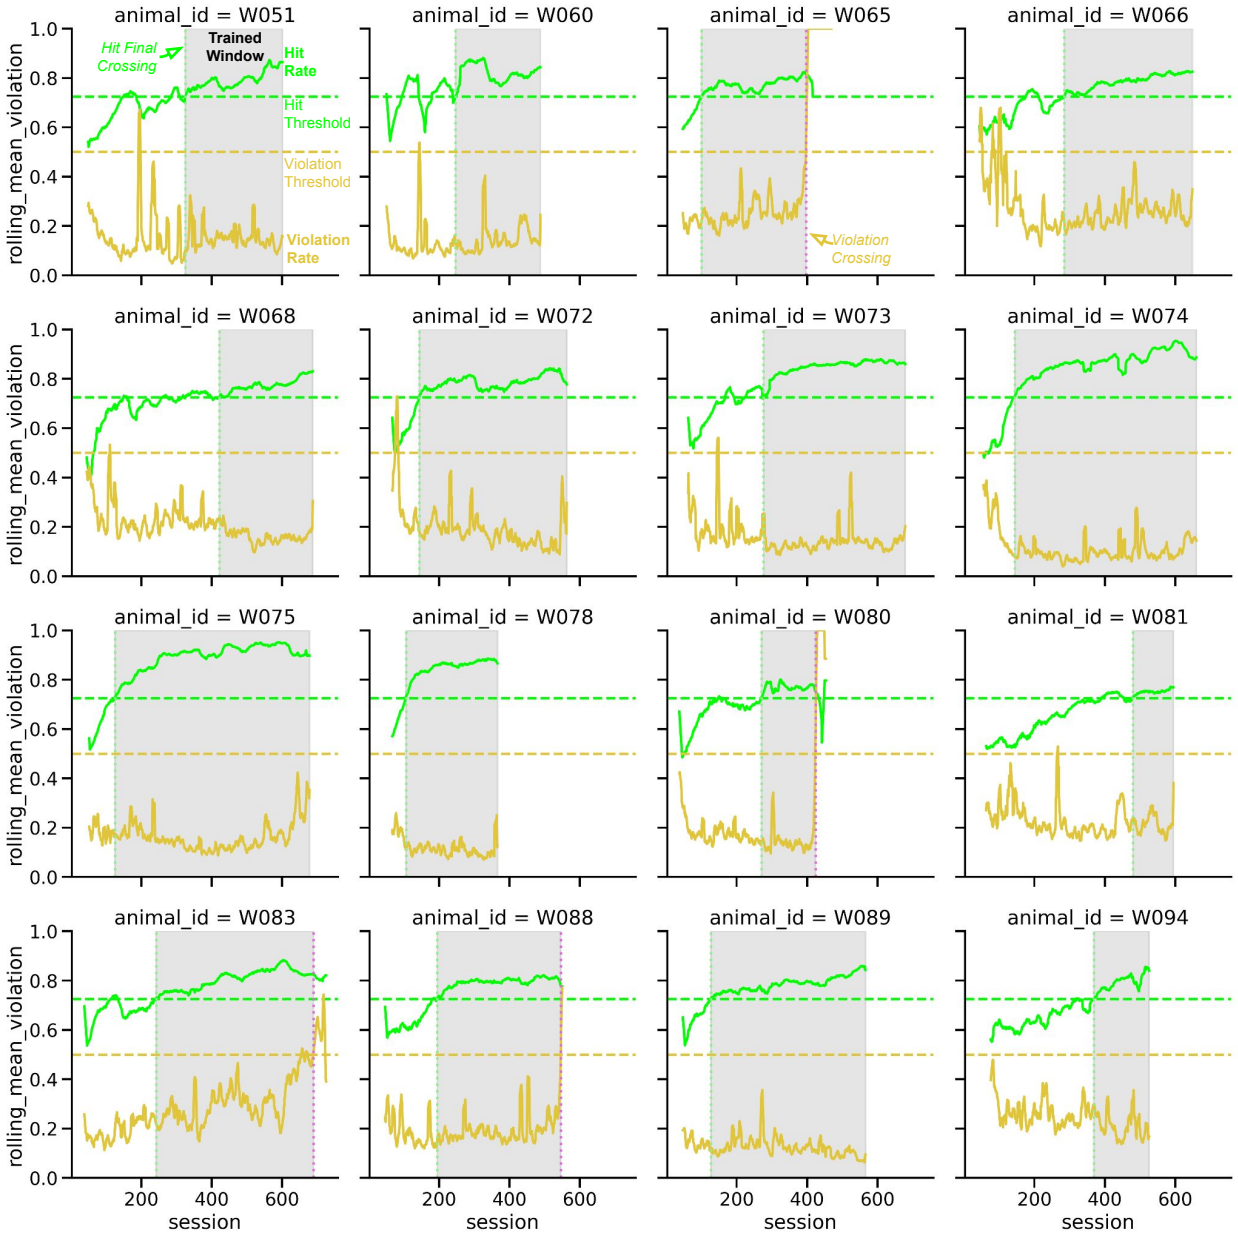


Figure S2.3: *Trained window for each animal over sessions.* Lime green solid line represents 21 day rolling hit rate. Green horizontal dashed line indicates a minimum hit threshold of 72.5%. Light green vertical line indicates the last crossing from low to high of the hit threshold when an animal is "trained". Yellow solid line represents 7 day rolling violation rate. Yellow horizontal dashed line represents a maximum violation threshold of 50%. Red vertical dashed line represents crossing of violation threshold at end of training (not applicable to all animals) Gray shading indicates sessions that are included in trained dataset given hit and violation thresholding.

2.4.4 Autocorrelation Analysis

We computed autocorrelograms of violation trials within each animal-session for ± 20 trial lags. We excluded sessions with no violations. Autocorrelation values were averaged across sessions and compared to a null distribution generated by shuffling violation labels within each session for 100 permutations. We computed 95% bootstrap confidence intervals for the observed mean autocorrelogram and percentile intervals for the null distribution. Statistical significance at each lag was determined using two-tailed p-values based on the empirical null distribution and corrected for multiple comparisons using the Bonferroni method.

2.4.5 Models

Binary Logistic Regression

To model the rats' left or right choices, we implemented a binary logistic regression model after dropping violation trials. The model quantified the probability of an animal making a right choice ($y = 1$) using a standard logistic (sigmoid) function:

$$P(y = 1) = \frac{1}{1 + e^{-Xw}}$$

where X represents the trial-by-trial design matrix composed of regressors capturing task related regressors and a bias term (intercept), and w is a vector of weights corresponding to each regressor. Model weights were fitted by minimizing the negative log-likelihood cost function, with a L2 regularization to prevent overfitting, which was computed as:

$$nll = -\sum_{i=1}^{N}[y_i(X_i w) - log(1 + e^{X_i w})] + \frac{1}{2\sigma^2}\sum_{j=1}^{D} w_j^2$$

where y_i denotes the observed binary choices (coded as 0 for left and 1 for right) for trial i, N is the number of trials, D is the number of regressors (excluding the bias term), and σ^2 represents the variance hyperparameter for Gaussian prior regularization that was tuned for

each animal. Optimization of model parameters was performed using the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm implemented in the `scipy.optimize.minimize` function in Python (3.10).

Multinomial Logistic Regression

To extend our analysis from binary left and right choices to include violation trials we implemented a multinomial logistic regression model. This approach generalizes the binary logistic regression model by employing a multinomial (softmax) logistic function, thereby enabling simultaneous prediction of the three discrete behavioral outcomes. The probability of each outcome class c for is expressed as:

$$P(y = c|X) = \frac{e^{XW_c}}{\sum_{k=1}^{C} e^{XW_k}}$$

where X represents the trial-by-trial design matrix composed of regressors capturing task related regressors and a bias term (intercept), and W is the matrix of regressor weights for each outcome class c where $C = 3$ (left, right, violation). Model weights were fit by minimizing the negative log-likelihood cost function with L2 regularization, which was computed as:

$$nll = -\sum_{i=1}^{N}\sum_{c=1}^{C} Y_{ic}(X_i W_c) + \sum_{i=1}^{N} log\left(\sum_{c=1}^{C} e^{X_i W_c}\right) + \frac{1}{2\sigma^2}\sum_{j=1}^{D}\sum_{c=1}^{C} W_{jc}^2$$

where Y_{ic} represents the observed outcome for trial i encoded as a one-hot vector, N denotes the total number of trials, D is the number of regressors (excluding the bias term), and σ^2 represents the variance hyperparameter for Gaussian prior regularization that was tuned for each animal. Consistent with the binary model, bias weights were not penalized by regularization and parameter optimization was performed with the BFGS algorithm in SciPy.

2.4.6 Feature Design & Choice Labels

The design matrix X is N trials by $D + 1$ features and represents the trial-by-trial stimulus and choice variables selected for the experiment with a bias column. The features used were similar to those used by Roy et al., 2021 (which we refer to as the "baseline" model) with the addition of features representing violation information. They were constructed as described in **Table 2.1**.

Feature	Description
Stimulus A (S_a)*	Standardized (mean 0, std 1) stimulus A loudness on current trial
Stimulus B (S_b)*	Standardized stimulus B loudness on current trial
*Previous Choice**	Choice animal made on previous trial where Left = -1, Right = 1 and Violation = 0; approximates side bias behavior
*Previous Correct**	Correct side on previous trial where Left = -1, Right = 1 and Violation = 0; approximates win-stay-lose-switch behavior
*Previous Stimuli**	Standardized average of stimulus A and B on previous trial. If the previous trial was a violation, it will be masked to 0
Previous Violation	Boolean representing if animal violated on previous trial
Filtered Previous Violation	Exponential filter of previous violation history given a specified time constant (τ)
Previous Violation & Stimulus Interaction	A multiplicative combination of violation history and standardized stimulus (e.g. previous violation x stimulus A). Can be either a binary history or filtered history

Table 2.1: Label and description of the features used to create design matrices in logistic regression modeling experiments. * indicates baseline model features.

To create an exponentially filtered violation history regressor, we convolved a decaying exponential kernel with the trial-wise previous violation boolean using a kernel length of five times the time constant (τ) to ensure sufficient decay where τ is tuned to each animal (**Fig. S2.4**). The convolution was performed independently within each session, and the resulting values were scaled by the maximum across all sessions for an animal to normalize the output between 0 and 1.

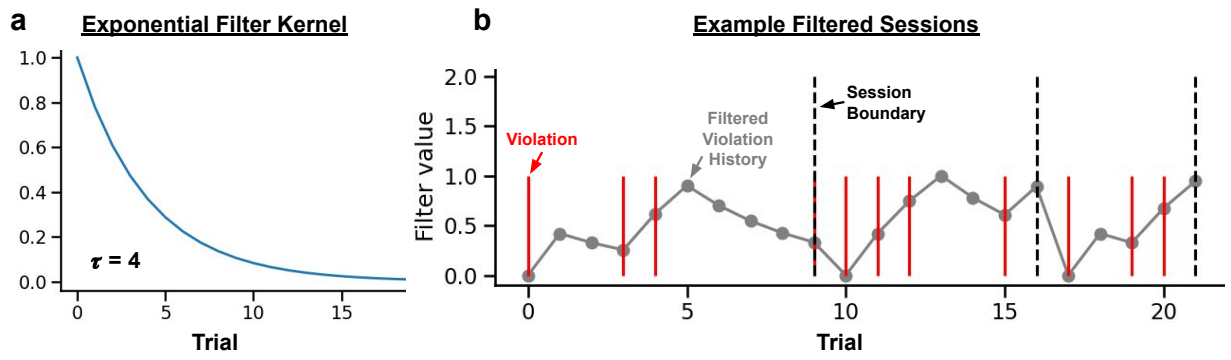


Figure S2.4: *Simulated example of creating exponentially filtered violation history.* **(A)** Exponential filter with time constant ($\tau = 4$). **(B)** Filter kernel in (A) applied to previous violation history over trials. Violation events are marked with red vertical lines, session boundaries are marked with black dashed lines and filter output is in gray. Note the filter resets at the beginning of a session and is scaled between 0 and 1 across all sessions

Finally, choice labels (y) were encoded as Left = 0 and Right = 1 for the binary model and one hot encoded as Left = [1 0 0], Right = [0 1 0] and Violation = [0 0 1] for the multinomial model.

2.4.7 Model Fitting

Once a design matrix was constructed for an animal, it was split by session into an 80/20 train and test sets. A standard experiment involved sweeping L2 penalty values in the set

$$\sigma = \{\, 0.07,\ 0.13,\ 0.25,\ 0.5,\ 1,\ 2,\ 4,\ 8,\ 16 \,\}$$

and selecting the value with the lowest test negative log-likelihood. To determine the best time constant τ value to use for a filtered history regressor, a grid search over τ and σ was performed for each animal. We again selected based on minimum test-set negative log-likelihood, with

$$\tau \ = \ \big\{\, 0.5, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 12, 15, 20, 25, 30, 40, 50, 100 \big\}. \tag{2.1}$$

2.4.8 Model Comparison

All models were fit separately for each individual animal on a single-fold validation. For each fit, we calculated the log-likelihood of the test set given the best fit parameters on the training set (LL_{model}). We additionally computed the log-likelihood of a null model (LL_{null}), which was the mean log-likelihood of the test choices under the empirical distribution of choice frequencies across the test set. The null model provides a baseline that reflects the predictive performance of a model with no sensitivity to features or history, serving as a lower bound for model comparison.

To compare models on a normalized scale, we computed bits per trial for each animal and model. Bits per trial quantify the improvement in predictive accuracy over the null model in units of information. Specifically, for each model, we computed the difference in log-likelihood between the model and the null model on the test set, normalized by the number of test trials:

$$\text{bits per trial} = \frac{LL_{\text{model}} - LL_{\text{null}}}{N_{\text{test}} \cdot \log(2)}$$

where LL_{model} and LL_{null} are the test log-likelihoods of the fitted and null models respectively, and N_{test} is the number of test trials. The null model's own bits per trial is computed as:

$$\frac{LL_{\text{null}}}{N_{\text{test}} \cdot \log(2)}$$

Next, to directly compare predictive performance on left/right decisions between binary and multinomial models, we evaluated multinomial models using a restricted softmax likelihood that excluded the violation class from the denominator. Specifically, we computed the negative log-likelihood on the test set using only the logits associated with left and right choices:

$$P(y = c \in \{\text{Left, Right}\} \mid X) = \frac{e^{XW_c}}{e^{XW_{\text{Left}}} + e^{XW_{\text{Right}}}}$$

This allowed us to assess whether training a multinomial model with violation regressors improves prediction of left/right choices, while holding evaluation metrics consistent with the binary framework.

2.4.9 Psychometric Analysis

To assess how violation history modulates stimulus-based decision-making, we simulated psychometric curves using the fitted weights from the multinomial logistic regression model. Specifically, we used the stimulus regressors (s_a, s_b), binary previous violation regressor (v_x), and their interaction terms to compute the decision variable z as:

$$z = [w_{s_a} + (w_{s_a \times v_x} \cdot v_x)] \cdot s_a + [w_{s_b} + (w_{s_b \times v_x} \cdot v_x)] \cdot s_b + w_{\text{bias}}$$

where w_{s_a}, w_{s_b} are the main effect weights for stimuli, $w_{s_a \times v_x}, w_{s_b \times v_x}$ are the weights for the interaction between violation history and stimulus, and w_{bias} is the bias term. We computed the probability of a rightward choice using the sigmoid function:

$$P(\text{right}) = \frac{1}{1 + e^{-z}}$$

Simulations were performed over a grid of $\delta = s_a - s_b$ values spanning from –2 to 2 to approximate the standardized stimulus on post-violation ($v_x = 1$) and post non-violation ($v_x = 0$) trials. To quantify the effect of violation history on stimulus sensitivity, we computed the slope and lapse rates from simulated psychometric curves generated using fitted model weights. The slope was calculated as the derivative of the psychometric curve when $\delta = 0$. Upper and lower lapse rates were defined as deviations from perfect performance at the extremes: $1 - p_{\text{right}}$ at $\delta = -2$ and p_{right} at $\delta = 2$. Upper and lower lapse rates were added

together to create a combined lapse rate. Each metric was compared across conditions using paired t-tests after verifying normality of each distribution with the Shapiro–Wilk test.

2.4.10 Data & Code Availability

All code for analyses can be found at https://tinyurl.com/jbredachapter2code. Data can be found on cup at: `/brody/jbreda/brody/jbreda/behavioral_analysis/violations_multinomial/data`

3

FixGrower: An Efficient And Robust Curriculum for Shaping Fixation Behavior in Rodents

3.1 Introduction

Training animals to perform perceptual decision making tasks in controlled laboratory environments has led to significant advancements in the fields of neuroscience and psychology (Oesch et al., 2024). Teaching animals these complex tasks takes a considerable amount of time, usually on the order of weeks to months, and this has led to a body of theoretical and behavioral research focused on optimizing task learning.

These types of tasks typically have two components to learn: a perceptual stimulus rule and a corresponding motor sequence. The motor sequence is composed of actions related to both trial engagement and timing (e.g. withholding movement during the stimulus) as well as choice reporting (e.g. moving to the left or right given the stimulus). Most efforts to improve training efficiency have focused on accelerating stimulus rule acquisition by using a variety of techniques such as passive stimulus exposure (Schmid et al., 2024), or difficulty scheduling (Reuschenbach et al., 2023; Soma et al., 2014). In contrast, inefficiencies in motor sequence learning have been comparatively understudied, despite their considerable impact on overall training duration and performance quality.

In perceptual decision-making tasks, these structured, sequential motor behaviors are essential for ensuring consistent engagement and reliable measurement of stimulus-driven choices. In freely-moving, unconstrained versions of these tasks in rodents, a common motor requirement is center-port fixation. This is when the animal must initiate and sustain a nose poke at a central port for a specified duration prior to and/or during stimulus presentation. Acquiring stable fixation behavior represents a surprisingly difficult challenge. For example, Schmid et al. (2024) reported that mice required 9–12 days of training to achieve a 250 ms pre-stimulus fixation hold. To learn a sound categorization rule following the hold took an additional 7–15 days. These findings underscore that the motor control demands associated with fixation can be substantial, sometimes requiring as much time as mastering the stimulus-

response association itself.

Despite the widespread use of fixation in rodent decision-making tasks, training procedures for this motor requirement are often idiosyncratic across laboratories and minimally described in the literature (to highlight the variety in reporting and training across similar labs see: Akrami et al., 2018; Constantinople et al., 2019; Jaramillo & Zador, 2014; Mah et al., 2023; Odoemene et al., 2018; Schmid et al., 2024; Scott et al., 2015). Nevertheless, several common elements are shared across approaches. First, the initial fixation requirement typically starts at an imperceptible threshold on the order of tens of milliseconds. Second, the required fixation duration increases gradually on a trial-by-trial basis. Third, failed fixation attempts are typically followed by either an inter-trial interval (ITI) delay or an additional penalty timeout that prevents an immediate retry after failure. Only after an animal consistently achieves the desired fixation duration are stimuli and the associated discrimination rules introduced.

However, even after animals consistently perform the stimulus discrimination rule, fixation errors— termed *violations*— represent a persistent challenge. These events are common, accounting for approximately 20–35% of initiated trials (Ashwood et al., 2022; Constantinople et al., 2019; Cruz et al., 2022; Jang et al., 2024; Jaramillo & Zador, 2014; Licata et al., 2017; Pisupati et al., 2021). As shown in Chapter 2, violations are not random lapses; they exhibit temporal structure, propagate across trials, and impair stimulus sensitivity on subsequent trials. In addition to their impact on behavior, violations also reduce data yield, as aborted trials are typically excluded from behavioral and neural analyses, thereby complicating dynamic models of decision making (Roy et al., 2021). Therefore, in the absence of standardized fixation training procedures, strategies for minimizing violations remain poorly defined.

In light of these challenges, we sought to develop an effective and efficient curriculum for shaping fixation behavior in freely moving rodents. Our goal was to design a robust,

practical procedure that could generalize across tasks and laboratories. We prioritized two outcomes: 1) minimize the time required to achieve a target fixation duration and 2) maintain low violation rates once at the target. There were three primary changes we implemented in this novel curriculum, which we refer to as "FixGrower". First, the starting fixation requirement used on the initial day of training was considerably long at 400 ms. Second, we kept the fixation requirement stable within a session and grew only between sessions. And, the magnitude of the growth of the fixation requirement was adaptively set by the performance of each animal in the preceding session. Finally, we implemented a violation penalty followed by an inter-trial interval only after animals reached their target fixation requirement of 2 seconds, thereby allowing them to immediately retry following a failure during fixation growth.

All together, when compared to a "Legacy" curriculum used in our laboratory and similar to those in the field, we demonstrate this new approach is 61% faster, while maintaining low violation rates of 13.8% on average. Critically, we demonstrate our results are robust over time, species and even when a stimulus rule is introduced during fixation. We find the principles we implemented are well supported by theories in operant conditioning, reinforcement learning and curriculum learning. We believe these findings will expedite the necessary, foundational aspects of early task shaping and will be of use to a variety of tasks in the field.

3.2 Results

3.2.1 Fixation Curricula Comparison

We trained freely moving rats to perform a center-port fixation task (**Fig. 3.1a**), which is a common motor requirement of perceptual decision making tasks. In our task, animals poked and held their nose in an illuminated center-port until an auditory white noise go-cue indicated the end of the fixation period. If animals successfully fixated until the go-cue, they could collect water reward at a randomly illuminated lateral choice port. If animals

violated the rule by failing to fixate until the go-cue, the trial was aborted. Depending on the curriculum and training stage, and a 2-second timeout penalty and inter-trial interval delay may follow a violation. Required fixation durations ranged from 10 milliseconds to 2 seconds.

We compared two different curricula for growing the fixation duration to a target hold of 2 seconds (**Fig. 3.1b**; **Table 3.1**). We evaluated the curricula using two metrics: (1) the number of days it took to reach the target and (2) violation rates in a pair of 5-day probe stages once the target was reached. The first probe stage tested performance at stable, long fixation holds of 2 second each trial. The second probe stage tested performance at variable fixation holds, sampled from an exponential with $\tau = 1.2\,\mathrm{s}$ each trial. Prior to the center-port fixation task, we started all animals on a simpler side-poking task to habituate them to the behavior box environment and stimuli, and ensure they were motivated to collect water reward (**Table 3.2**; see *Methods: Curricula: Side Poking*). We then randomly assigned animals into one of two curricula: Legacy ($N = 8$), which is based on a typical fixation curriculum that is used in our laboratory and similar to those in the field (**Fig. 3.1c**), and FixGrower ($N = 9$), which is the novel fixation curriculum we developed (**Fig. 3.1d**).

There were three differences between the two curricula (**Fig. 3.1e**). First was the *starting fixation requirement*. This is the duration of time animals must fixate in the center-port prior to the go-cue on their first day of the curriculum. The Legacy curriculum had a "short" fixation requirement of 10 ms, while FixGrower had a "long" fixation requirement of 400 ms.

The second difference was the *fixation growth algorithm*. For Legacy, the required fixation duration grew following each non-violation ("valid") trial by max(1 ms, 0.01%). Additionally, the first 20 trials of each session warmed up to the previous session's fixation maximum starting at 10 ms. In contrast, FixGrower had a stable fixation requirement within a session and only grew between sessions. The requirement grew to the average duration of an animal's valid center-pokes in the previous session.

The final difference between the curricula was the *violation penalty introduction*. Legacy had an "early" penalty such that there was a 2-second timeout and inter-trial interval delay following a violation trial from the first day in the curriculum. Whereas FixGrower had a "late" penalty such that there was no penalty for violations until the 2-second fixation target was reached. Therefore, prior to reaching the target, trials could be immediately restarted after a failed attempt. Once the target was reached, FixGrower had a single day of the violation timeout penalty being turned on (Stage 8) before progressing animals into the probe stages. In contrast, the Legacy curriculum move animals directly into the probe stages once they reached the target since the penalty was already on.

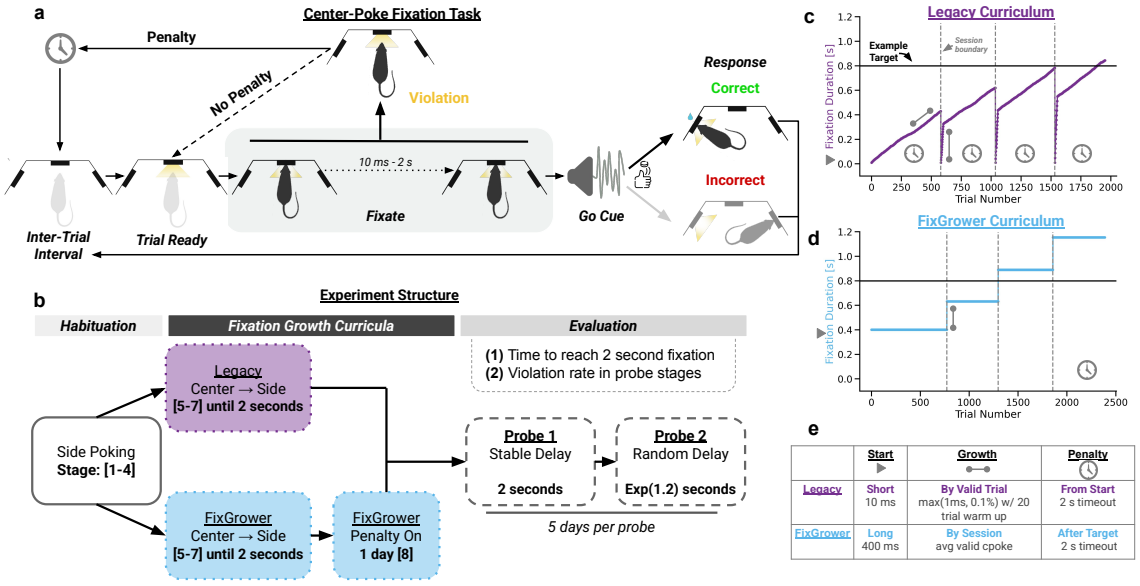


Figure 3.1: *Overview of task structure and curricula.* **(A)** Trial structure for center-port fixation task. Center light indicates trial ready, fixation occurs with nose in center hold and movement is allowed after a go-cue (200 ms white noise). The correct lateral choice is indicated by an illuminated side-port after the go-cue. If the animal prematurely removes their nose from center prior to the go-cue, the trial is considered a violation. A 2-second timeout penalty may follow a violation trial depending on the curriculum. An inter-trial interval period always follows a valid trial when animals fixate until the go-cue as well as trials when a violation penalty occurs. **(B)** Overview of curricula stages for our experiment. All animals have the same side poking habituation stages (gray bar) followed by differing fixation growth curricula (black bar) for Legacy (dark purple) and FixGrower (light blue) animals. Once the fixation target of 2 seconds is reached and violation penalty is turned on, all animals in the experiment go through 5 days of each probe stage to evaluate performance (gray bar). **(C)** Demonstration of how required fixation duration grows given the Legacy curriculum growth algorithm for four initial sessions (indicated by dashed vertical lines) with an example target duration of 0.8 seconds. **(D)** As in (C) but for the FixGrower curriculum growth algorithm. **(E)** Table summarizing differing features in the Legacy and FixGrower curricula.

3.2.2 FixGrower Decreases Training Time

Since the center-poke fixation task is a motor prerequisite for many perceptual decision-making tasks, reducing the training duration for this phase is desirable to expedite progression to the main task of interest. Therefore, we assessed whether training speed differed between the two curricula. We found that FixGrower-trained animals grew significantly faster to the target fixation requirement with an average of 11.4 ± 5.4 days compared to Legacy-trained animals which averaged 29 ± 18.3 days (*Mann–Whitney U* test, $U = 65.5$, $p = 0.005$; Post hoc power analysis; Cohen's $d = 1.33$; power $= 0.73$; **Fig. 3.2a–b**). To accomplish this, FixGrower animals grew significantly more session-to-session than Legacy animals (mean $\pm$ SD: **L:** 0.07 ± 0.17 seconds/session, median: 0.07, $n_{\text{Sessions}} = 245$ vs. **FG:** 0.16 ± 0.17 seconds/session, median: 0.20, $n_{\text{Sessions}} = 104$; *Mann–Whitney U* test, $U = 5654.0$, $p < 1.74 \times 10^{-12}$; **Fig. 3.2c**).

This increased efficiency was not due to differences in engagement or performance between the cohorts prior to starting the fixation growth. (**Fig. S3.1**). Nor was it due to an intrinsic growth ceiling in the Legacy growth algorithm given the number of trials performed (**Fig. S3.2**). In fact, simulations of normative performance demonstrated that Legacy animals could reach the fixation target significantly faster than FixGrower animals given the number of trials they were performing (5.8 ± 1.4 days, *Mann–Whitney U* test, $U = 5.000$, $p = 0.003$; **Fig. S3.2c**). The only notable difference between the two cohorts was that Legacy-trained animals had significantly more behavior box ("rig") switches than FixGrower-trained animals (mean $\pm$ SD: **L:** 1.62 ± 0.52 switches, **FG:** 1.0 ± 0.0), which can destabilize performance as the animal adjusts to a new environment (**Fig. S3.1d**). However, this is anticipated given the significantly longer training times for the Legacy curriculum cohort. Further analysis showed no correlation between switches and performance (**Fig. S3.3**). Altogether, these data demonstrate that the structure of the FixGrower curriculum facilitated

faster fixation growth and significantly reduced training times.

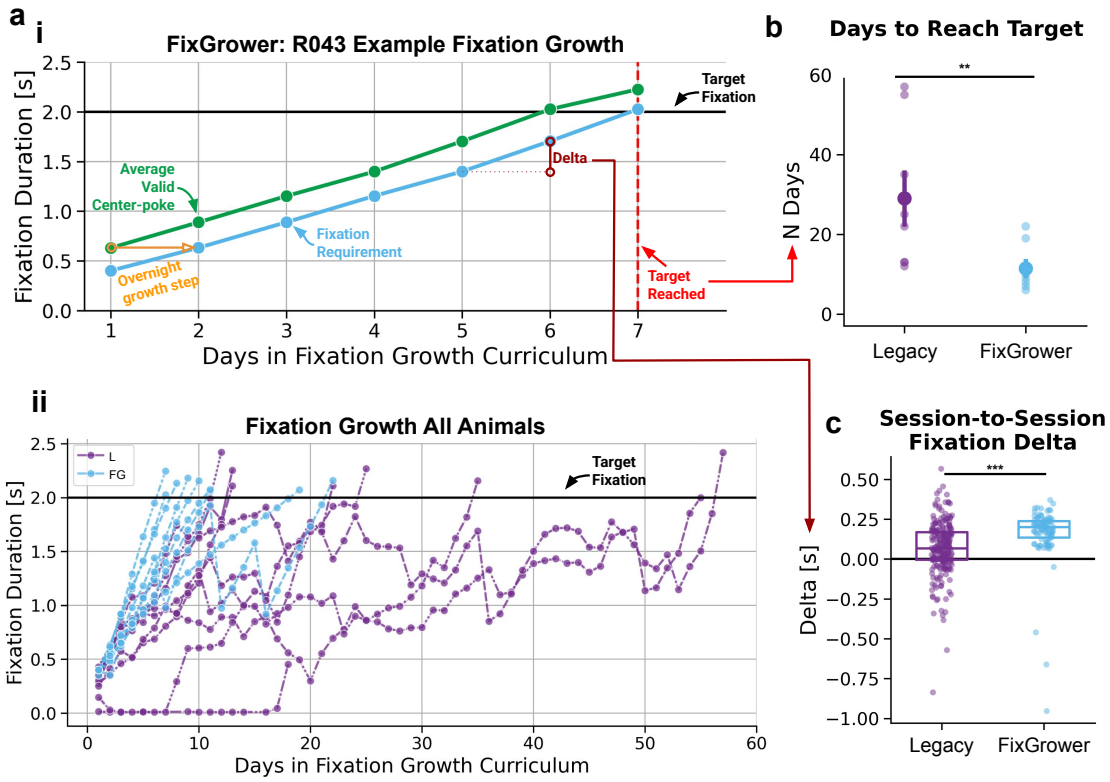


Figure 3.2: *FixGrower-trained animals grow faster to fixation target.* **(A)(i)** Example growth trajectory for a single FixGrower animal. Black line indicates the 2-second fixation target. Light blue line indicates the required fixation for the given day. Green line indicates the average valid center-poke duration for a day, which is then set as the fixation requirement the following day (orange). **(A)(ii)** Maximum fixation duration reached in a session for each animal colored by experimental condition over days in the fixation growth curriculum. Black line indicates target fixation duration of 2 seconds. **(B)** Summary of days to reach 2-second target by fixation experiment for Legacy-trained animals (dark purple) and FixGrower-trained animals (light blue). Large circular points represent cohort mean and error bars represent standard error of the mean. Individual animals indicated by smaller points. Example days to reach the target indicated by the red dashed line in (A)(i). **(C)** Summary of fixation delta for each session by experimental condition. Boxes extend between lower and upper quartiles with a line at the median. Example delta indicated by a maroon line with open circles in (A)(i). *Statistics:* In **(B)** and **(C)** $*p < 0.05$, $**p < 0.01$, $***p < 0.001$ and ns is non-significant tested using *Mann–Whitney U* tests. A mixed linear model with a random intercept for animal ID yielded the same results for **(C)**.

3.2.3 FixGrower Yields Low and Predictable Violation Rates

We next sought to evaluate fixation performance by assessing violation rates during two probe stages administered after animal were trained on their respective curricula to reach the 2-second fixation target. Each probe stage lasted five days. The first probe stage tested

performance under stable fixation demands: each trial had a 2-second hold requirement. This served two purposes: first, to assess performance at a sustained, effortful fixation duration; and second, to evaluate how Legacy-trained animals performed under a consistent, trial-invariant rule. Animals trained by both curricula performed well, with average violation rates of $15.5 \pm 5.3\%$ (median: 16.2%) for Legacy-trained and $15.6 \pm 8.2\%$ (median: 12.4%) for FixGrower-trained, and no significant difference in performance (*Mann–Whitney U* test, $U = 947.5$, $p_{\text{holm}} = 0.466$; **Fig. 3.3a**).

The second probe stage tested performance under variable fixation demands: on each trial, the required fixation duration was sampled from an exponential distribution with $\tau = 1.2$ seconds, truncated between 1 and 2 seconds. This stage was designed to challenge FixGrower-trained animals, who had previously been exposed only to sessions with stable requirements. Nevertheless, they adapted well, showing a mean violation rate of $12.2 \pm 7.7\%$ (median: 9.29%) and no significant difference from Legacy-trained animals (mean ± SD: $10.1 \pm 5.3\%$, median: 10.1%; *Mann–Whitney U* test, $U = 764.0$, $p_{\text{holm}} = 0.466$; **Fig. 3.3a**). Further, both cohorts showed a downward trend in violation rates over the ten total probe sessions (**Fig. 3.3b**), and both exhibited significantly lower violation rates in the second, random probe stage compared to the first, stable one (*Wilcoxon signed-rank* tests, **L:** $W = 1.0$, $p_{\text{holm}} = 0.046$; **FG:** $W = 2.0$, $p_{\text{holm}} = 0.046$; **Fig. 3.3a**). These data demonstrate that the FixGrower-trained animals not only generalized well to variable fixation demands but also improved with continued exposure, indicating ongoing refinement of motor control even after reaching the training criterion.

Lastly, we were interested in the relationship between training times and performance, as it is of significant cost for experimenters to wait with high uncertainty about an animal's training potential. To test this, we examined whether the number of sessions it took to grow to the target fixation could predict an animal's combined probe-stage performance. While no relationship was observed in Legacy-trained animals, we found that FixGrower-trained ani-

mals who reached the target faster had lower violation rates in the probe stages (**Fig. 3.3c**). A mixed-effects linear model comparing probe-stage violation rate and the number of days to reach the target fixation requirement revealed a significant difference in slope between the cohorts (interaction term $\beta = 0.016$, $p < 0.001$, see *Methods: Excluded Animals*). These findings suggest that, in the FixGrower curriculum, longer training durations predicted poorer probe-stage performance, highlighting an exploitable trade-off between training speed and behavioral stability in this curriculum.

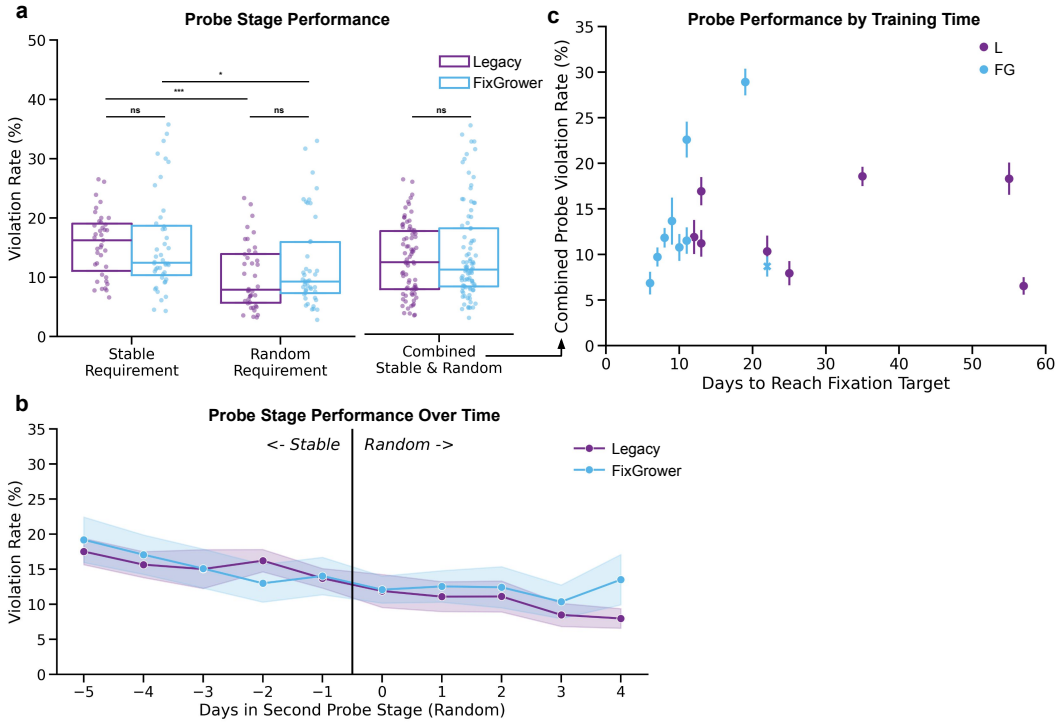


Figure 3.3: *FixGrower-trained animals have low violation rates in probe stages that are predicted by training time.* **(A)** Violation rate for the first, stable and second, random probe stages as well as the combined data from both stages for Legacy-trained (dark purple) and FixGrower-trained (light blue) animals. Dots represent individual animals and sessions. Boxes extend between lower and upper quartiles with a line at the median. Across-cohort statistics are computed using *Mann–Whitney U* test. Within-cohort statistics are computed using *Wilcoxon signed-rank* tests. Corrected for multiple comparisons with Holm method where * indicates $p < 0.05$, ** indicates $p < 0.01$, *** indicates $p < 0.001$ and ns is non-significant. **(B)** Violation rates over days in the probe stages for each cohort. Markers and lines represent the mean and shading indicates the standard error of the mean. For common alignment, day 0 is the first day of the random probe stage. **(C)** Relationship between days to reach target fixation of 2 seconds and average violation rate in the combined probe stages for each animal. Error bars are standard errors of the mean. FixGrower outlier animal has 'X' marker and is excluded from mixed-effects linear model as discussed in *Methods: Excluded Animals*.

3.2.4 FixGrower Generalizes Over Time, Tasks and Species

Given the observed potential of the FixGrower curriculum, we were interested in testing how robust and generalizable it was to determine its suitability for broader adoption in the field. We first started by assessing long-term task performance. Specifically, we examined whether animals trained on the FixGrower curriculum could sustain low violation rates over an extended period without further experimenter intervention. We selected a subset of animals trained on FixGrower from the preceding fixation experiment ($N = 8$) and monitored their performance for 100 days following completion of the probe stages. Each day, the required fixation duration varied randomly for each animal and was sampled from an exponential distribution ($\tau = 1.2$ seconds), truncated between 1 and 2 seconds (**Fig. 3.4a**–**b**). Across this prolonged evaluation, animals consistently maintained low and stable violation rates, averaging $9.89 \pm 8.26\%$ (median 7.27%; $n_{\text{sessions}} = 800$; **Fig. 3.4c**–**d**). These findings strongly indicate that animals trained with FixGrower reliably mastered and maintained fixation task performance.

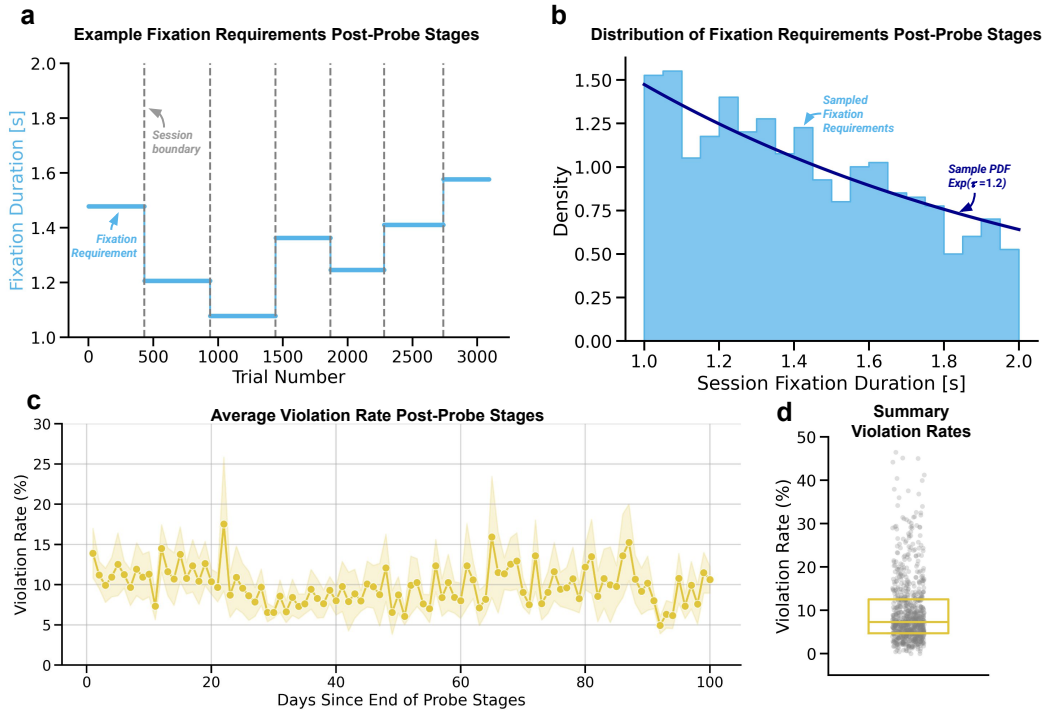


Figure 3.4: *FixGrower-trained animals maintain robust task performance over extended periods.* **(A)** Example sequence of fixation duration requirements sampled across trials from one representative animal (R053) after probe stage completion. Fixation durations were sampled from an exponential distribution ($\tau = 1.2$ seconds) truncated between 1 and 2 seconds, updated each session. Dashed vertical gray lines indicate session boundaries. **(B)** Distribution of fixation duration requirements aggregated across all sessions ($n = 800$). Histogram shows empirically sampled fixation durations (light blue bars), overlaid with the theoretical exponential PDF ($\tau = 1.2$ seconds) that fixation durations were drawn from. **(C)** Session average violation rates (%) plotted as a function of days post-probe stages completion. Each point reflects daily mean violation rates averaged across animals ($N = 8$), with shaded regions indicating standard error of the mean (SEM). **(D)** Summary plot displaying individual session violation rates aggregated over the entire 100-day assessment period. Box extends between lower and upper quartiles with a line at the median and each gray dot represents a single session for a single animal.

Next, to test the generalizability of FixGrower, we evaluated whether its low violation rates would persist when animals had to fixate while processing task-relevant stimuli. A new cohort of rats ($N = 13$) was trained on a sound discrimination task following FixGrower (see *Methods: Curricula: Sound Discrimination Task*). In this task, two 400 ms tones were separated by a 150 ms delay and flanked by additional delays (250 ms before and 450 ms after) for a total of 1.65 seconds of fixation (**Fig. 3.5a**). The first tone (3 kHz) served as a distractor, while the second tone indicated the rewarded side (12 kHz = right; 3 kHz = left;

Fig. 3.5b) and a 200 ms white noise go-cue indicated the end of the fixation period. Using FixGrower, rats averaged 6.6±2.2 days to grow to a common 1.45-second fixation target (min: 5, median: 6, max: 12; **Fig. 3.5c**). Once the sound rule was introduced, rats performed well in this task, achieving an average of $87.6 \pm 5.9\%$ correct ($n_{\text{sessions}} = 170$; $n_{\text{trials}} = 39{,}028$; **Fig. 3.5d**). Crucially, their violation rates remained low at $11.2 \pm 9.5\%$ (median: 8.33%; **Fig. 3.5e**), indicating that the FixGrower curriculum supports stable fixation even when the stimulus appears during the fixation period.

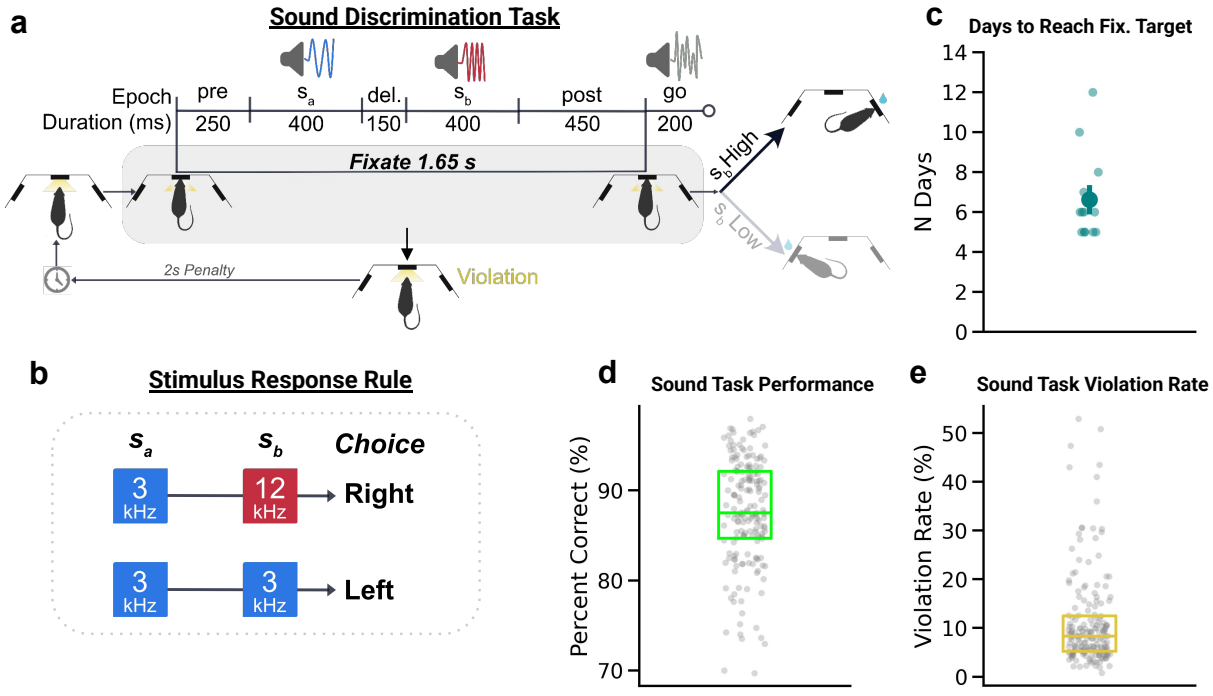


Figure 3.5: *FixGrower supports efficient and stable fixation in a sound discrimination task.* **(A)** Schematic of the sound discrimination task performed by a cohort of rats ($N = 13$) after completing FixGrower. Animals fixated for 1.65 seconds while auditory stimuli were presented: a non-informative distractor tone (s_a) followed by a second tone (s_b) indicating the correct choice (right for high frequency, 12 kHz; left for low frequency, 3 kHz). A 200 ms white noise go-cue signaled the end of fixation. **(B)** Stimulus set and response mapping used in the task. The frequency of s_b determined the rewarded port; s_a was always 3 kHz and non-informative. **(C)** Days required to reach the shared fixation target of 1.45 seconds (see *Methods: Curricula: Sound Discrimination Task*). Dots represent individual animals; the large dot shows the mean and error bars denote the standard error of the mean. **(D)** Task performance measured as percent correct on non-violation trials. Gray dots are individual animals and sessions; the green box shows the inter-quartile range with the median. **(E)** Violation rate for the sessions shown in (D). Gray dots are individual sessions; the yellow box shows the inter-quartile range with the median.

Lastly, we evaluated whether the FixGrower generalizes effectively from rats to mice, as previous literature has demonstrated that mice typically require longer durations to learn fixation tasks and exhibit higher average violation rates (Jaramillo & Zador, 2014; Odoemene et al., 2018; Schmid et al., 2024). We trained a cohort of mice ($N = 5$) to fixate to a target duration of 1.1 seconds (**Fig. 3.6a**). We chose this target because it is a substantial hold length that has been previously used in the field with reported training times (Odoemene et al., 2018). On average, these mice reached the 1.1-second fixation target in 4.4 ± 0.89 days (range: min = 3, median = 5, max = 5 days; **Fig. 3.6b**). Notably, a subset of these animals continued to grow beyond the initial target, achieving fixation durations exceeding 2 seconds within 9 days (**Fig. S3.4a**). Additionally, FixGrower-trained mice exhibited session-to-session fixation growth rates (0.25 ± 0.08 seconds/session, median: 0.21, $n_{\text{sessions}} = 22$; **Fig. 3.6c**) similar to those observed in FixGrower-trained rats (0.16 ± 0.17 seconds/session, median: 0.20, $n_{\text{Sessions}} = 104$). These results reflect substantial improvement over previously reported training times of approximately 10–12 days using a curriculum similar to Legacy (Odoemene et al., 2018).

We next assessed violation rates under modified conditions, using a short violation penalty (100 ms) and fixation durations ranging from approximately 1.5 to 2.7 seconds (see *Methods: Curricula: Mouse Fixation Experiment*, **Fig. S3.4b**). We determined how many sessions were required for animals to achieve violation rates below 20% while completing at least 150 trials per session (**Fig. 3.6d**). This 20% threshold represents the lower bound of violation rates commonly reported in prior mouse studies (Ashwood et al., 2022; Cruz et al., 2022; Jaramillo & Zador, 2014; Odoemene et al., 2018). Remarkably, all mice were sensitive to the short penalty and met this performance criterion within just 2–3 days (**Fig. 3.6e–f**), underscoring the effectiveness of the FixGrower curriculum in rapidly achieving low violation rates that match or exceed established benchmarks. Collectively, these findings indicate that FixGrower efficiently promotes rapid and robust fixation behavior in a variety of settings.

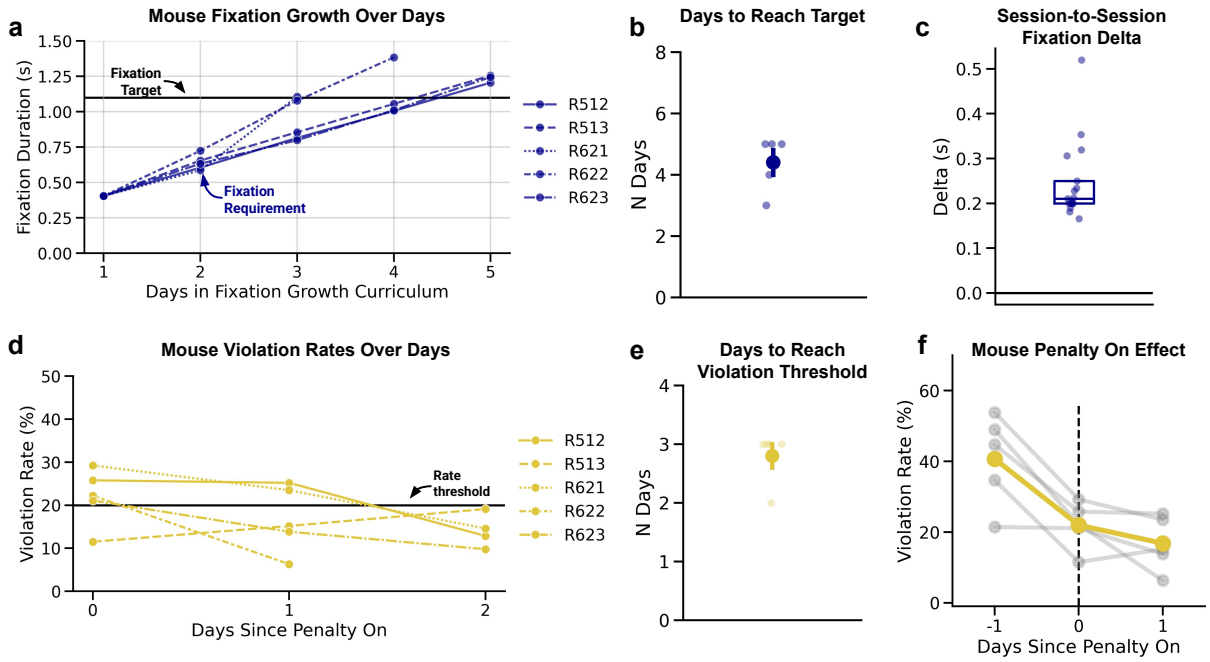


Figure 3.6: *FixGrower Curriculum generalizes to mice.* **(A)** Growth of session-maximum fixation duration across training days for individual mice ($N = 5$) using FixGrower. Each line represents one mouse. All animals were trained to a target of 1.1 seconds (black horizontal line). **(B)** Number of days each mouse required to reach the 1.1-second fixation target. Individual points represent mice; the large point indicates the cohort mean ± standard error of the mean (SEM). **(C)** Distribution of session-to-session fixation duration increases (delta), showing median (line), inter-quartile range (box), and individual animal and session values (dots). See **Fig. S3.4a** for delta intuition. **(D)** Violation rates for individual mice across days after introducing a 100 ms penalty. Each line represents one mouse, with a performance threshold of 20% indicated (black line). Mice complete this stage once they perform 150 trials below this threshold. **(E)** Days required for each animal to achieve violation rates below 20% while maintaining ≥ 150 trials per session. **(F)** Behavioral response to penalty onset. Violation rates are shown on the day before (–1), the day of (0), and the day after (+1) penalty activation. Gray lines represent individual mice, and the bold yellow line shows the group mean.

3.2.5 Curriculum Features Modulate Vigor, Engagement Stability, and Motor Control During Learning

Having shown that the FixGrower curriculum was more efficient than the Legacy curriculum and generalized across diverse settings, we next sought to understand the factors underlying this success. In our final set of analyses, we examined how specific curriculum features—*starting fixation duration, fixation growth algorithm*, and *violation penalty introduction*—may have influenced behavior to either accelerate training in FixGrower-trained

animals or impede it in Legacy-trained animals.

First, we focused on the delayed introduction of the violation penalty in FixGrower and its influence on motivation and trial initiation vigor. During fixation growth, the cost of failure is lower in FixGrower compared to in Legacy; in FixGrower, animals can immediately retry post-violation without additional delays, whereas in Legacy, animals incur both a 2-second violation penalty and an inter-trial interval delay. Notably, because fixation growth occurred on a per-trial basis in the Legacy curriculum, growth for these animals is constrained by the number of trials completed per session. We therefore hypothesized that the high cost of failure in the Legacy curriculum would suppress trial initiation, thereby slowing the rate of fixation growth.

Consistent with this hypothesis, Legacy-trained animals attempted significantly fewer trials per session compared to FixGrower-trained animals (**L:** mean ± SD 332 ± 109 vs. **FG:** 470 ± 184; $U = 5689$, $p = 2.96 \times 10^{-17}$; **Fig. 3.7a**), despite achieving a similar number of rewarded trials (**L:** 247 ± 80 vs. **FG:** 252 ± 110; $U = 12258.5$, $p = 0.372$; **Fig. 3.7b**). The higher trial initiation frequency in FixGrower animals was driven by them restarting trials post-violations 3.3 times faster on average than Legacy animals, even when accounting for experimenter-imposed delays (**L:** median 3.82 s, mean ± SD 16.6 ± 57.3 s, $n_{\text{trials}} = 17{,}908$ vs. **FG:** median 0.15 s, mean ± SD 5.00 ± 53.0 s, $n_{\text{trials}} = 22{,}622$; *Mann–Whitney* $U = 3.21 \times 10^8$, $p < 0.001$; **Fig. 3.7c**).

However, this increased vigor in FixGrower-trained animals was accompanied by higher violation rates during fixation growth, indicative of poor task learning (**L:** $21.8 \pm 11.6\%$, $n_{\text{sessions}} = 237$ vs. **FG:** $46.1 \pm 16.4\%$, $n_{\text{sessions}} = 110$; *Mann–Whitney* $U = 1977.5$, $p = 4.78 \times 10^{-37}$; **Fig. 3.7d**). Crucially, when we introduced the 2-second violation penalty in a single session prior to the probe stages, their violation rate immediately decreased from $47.4 \pm 16.5\%$ to $16.2 \pm 5.3\%$ (*Friedman* $\chi^2(2) = 13.5$, $p = 0.0011$; Day −1 vs. Day 0, *Wilcoxon signed-rank* $W = 0.039$, $p_{\text{Holm}} = 0.012$) and remained consistently low during the subsequent

probe session (18.4±8.9%; Day 0 vs. Day 1, *Wilcoxon signed-rank* $W = 0.039$, $p_{\text{Holm}} = 0.012$; **Fig. 3.7e**). These findings suggest that the elevated violation rates during fixation growth in animals trained with FixGrower were not due to impaired task comprehension, but rather reflected strategic behavior shaped by the minimal consequences of failure.

Correspondingly, the introduction of the penalty significantly increased the cost of failure, markedly slowing post-violation trial initiation times for FixGrower animals, bringing their behavior closer to that of Legacy animals (**L:** median 5.70 s, mean ± SD 23.0 ± 58.7 s, $n_{\text{trials}} = 2{,}960$ vs. **FG:** median 4.47 s, mean ± SD 23.5±66.9 s, $n_{\text{trials}} = 3{,}095$; $U = 5.04 \times 10^6$, $p = 7.82 \times 10^{-12}$; **Fig. 3.7f**). Moreover, we observed no significant difference between cohorts in the number of trials attempted during the probe stages (**L:** 345 ± 77 vs. **FG:** 359 ± 97; Welch's $t = 0.544$, $p = 0.595$; **Fig. S3.1a**). Collectively, these results demonstrate that large negative reinforcement early in learning is associated with reduced behavioral vigor, which—when coupled with the trial-dependent structure of the Legacy growth algorithm—likely contributed to slower fixation growth in the Legacy cohort.

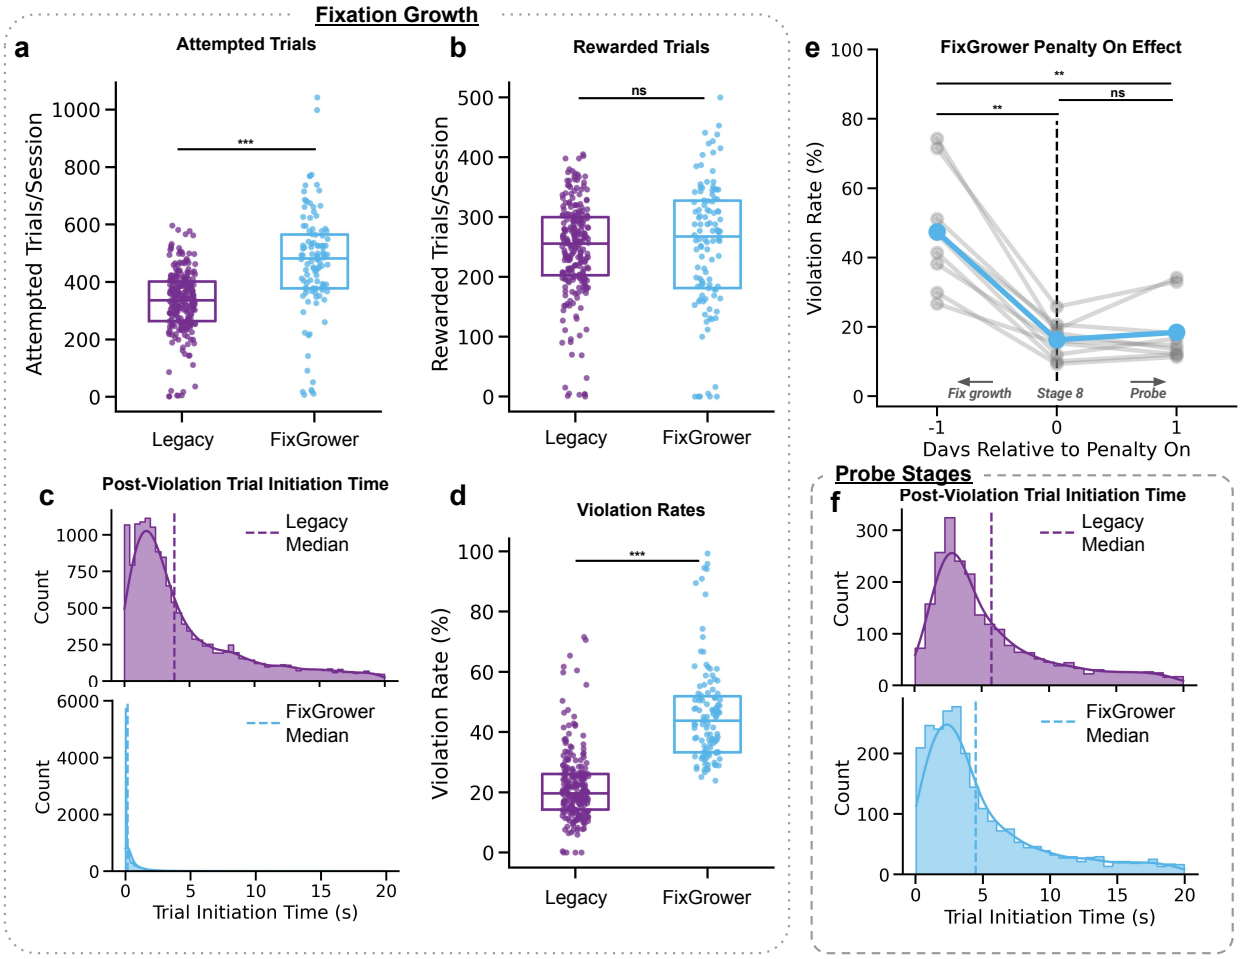


Figure 3.7: *FixGrower-trained animals have more vigor than Legacy-trained animals.* **(A)** Number of attempted trials per session for each cohort (Legacy: dark purple, FixGrower: light blue) during fixation growth where attempted trials are either violation or valid trials. **(B)** Same as (A) but number of rewarded trials per session where rewarded trials are valid hit trials. See *Fig. S3.1b* for hit rates). **(C)** Histogram and kernel density estimate of trial initiation times post-violation trials during fixation growth for each experimental cohort (Legacy: upper, FixGrower: lower) truncated to show initiations less than 20 seconds long. Note trial initiation times do not include incurred violation timeout penalty or inter-trial delay for Legacy-trained animals. Dashed lines indicate the cohort median. **(D)** Session violation rates during fixation growth for each experimental condition. **(E)** Violation rates for FixGrower-trained animals on the day before (-1), of (0) and after ($+1$) the 2-second violation penalty being turned on. Gray lines indicate individual animals. The light blue line indicates the group mean. Error bars indicate standard error of the mean. **(F)** As in (C) but with trials from the probe stages. *Statistics:* In **(A,B,D)** boxes extend between lower and upper quartiles with a line at the median and each dot represents a single session for a single animal. $^*p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$ and ns is non-significant using *Mann–Whitney U* tests (Note: mixed linear model with a random intercept for animal ID yielded the same results). In **(E)** the same p-value indicators are used with *Wilcoxon signed-rank* tests.

We next investigated how metrics of engagement and motivation evolved within individual sessions, given that task difficulty progressed differently for each cohort. In the Legacy curriculum, task difficulty incrementally increases throughout each session, as successful trials progressively raised the fixation requirement. In contrast, the FixGrower curriculum has peak task difficulty at session onset, where animals need to immediately adapt to the new fixation requirement set by the prior session. We thus predicted that engagement would decrease over sessions in Legacy-trained animals, while remaining stable or possibly improving in FixGrower-trained animals.

For Legacy-trained animals to optimally grow fixation duration, they ideally maximize their trial rate and minimize violation rates. We therefore analyzed these two metrics specifically during the first 50 and last 50 trials within sessions containing at least 200 trials as proxies of engagement. Consistent with our predictions, during fixation growth stages Legacy animals exhibited modest, yet significant increases in violation rates from early to late session (Early: $22.5 \pm 11.7\%$ vs. Late: $25.6 \pm 14.6\%$, $n_{\text{sessions}} = 216$; *Wilcoxon signed-rank* $W = 7803$, $p_{\text{holm}} = 3.22 \times 10^{-3}$; **Fig. 3.8a(i)**). In contrast, FixGrower animals showed significantly decreased violation rates from early to late session (Early: $66.1 \pm 14.1\%$ vs. Late: $35.6 \pm 13.7\%$, $n_{\text{sessions}} = 102$; $W = 8$, $p_{\text{holm}} = 6.67 \times 10^{-18}$; **Fig. 3.8a(i)**). Similarly, Legacy-trained animals displayed significantly increased trial initiation times from early to late session (Early: $11.4 \pm 4.08\,\text{s}$ vs. Late: $23.2 \pm 12.2\,\text{s}$, $n_{\text{sessions}} = 216$; $W = 920$, $p_{\text{holm}} = 3.96 \times 10^{-31}$; **Fig. 3.8b(i)**). Whereas, initiation times in FixGrower-trained animals only increased by half a second on average (Early: $14.5 \pm 13.9\,\text{s}$ vs. Late: $15.2 \pm 7.39\,\text{s}$, $n_{\text{sessions}} = 102$; $W = 1988$, $p_{\text{holm}} = 0.033$; **Fig. 3.8b(i)**).

Critically, these divergent within-session patterns were not observed during the subsequent probe stages, where task difficulty remained stable and matched between cohorts throughout each session. Under these matched conditions, both cohorts demonstrated significantly decreasing violation rates from early to late session (**L Early:** $25.2 \pm 12.7\%$ vs. **L**

Late: $12.6 \pm 7.84\%$, $n_{\text{sessions}} = 67$; $W = 104.5$, $p_{\text{holm}} = 4.91 \times 10^{-10}$; **FG Early:** $18.2 \pm 9.53\%$ vs. **FG Late:** $12.8 \pm 8.74\%$, $n_{\text{sessions}} = 71$; $W = 508.0$, $p_{\text{holm}} = 2.82 \times 10^{-5}$; **Fig. 3.8a(ii)**), consistent with improving proficiency, and significantly increasing trial initiation times (**L Early:** 15.5 ± 8.79 s vs. **L Late:** 26.5 ± 12.5 s, $n_{\text{sessions}} = 67$; $W = 190$, $p_{\text{holm}} = 6.13 \times 10^{-9}$; **FG Early:** 13.5 ± 6.42 s vs. **FG Late:** 21.1 ± 8.42 s, $n_{\text{sessions}} = 71$; $W = 338$, $p_{\text{holm}} = 7.20 \times 10^{-8}$; **Fig. 3.8b(ii)**), consistent with accumulating fatigue or satiety. Taken together, these data indicate that incremental within-session increases in task difficulty are associated with reductions in performance and engagement, factors which likely contributed to the significantly slower fixation growth observed in Legacy-trained animals.

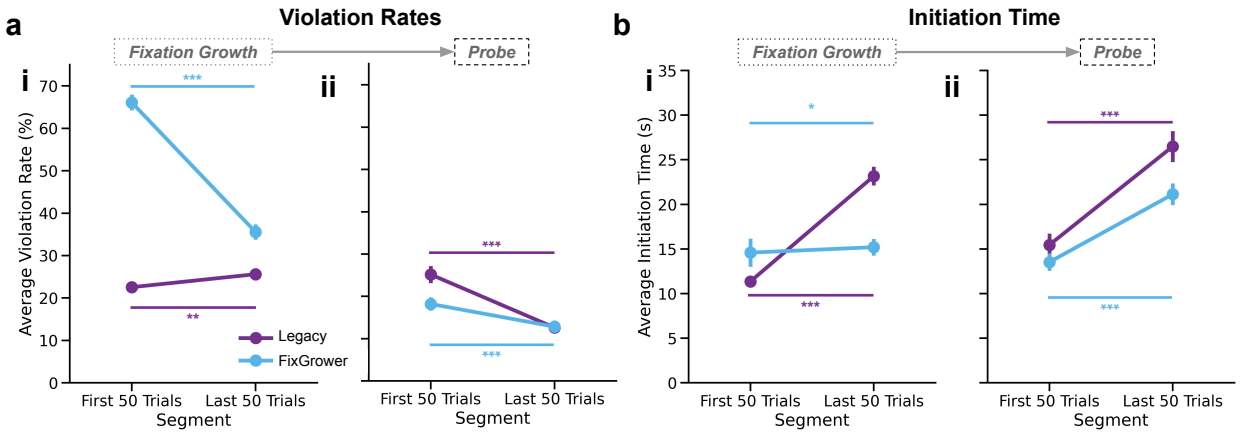


Figure 3.8: *FixGrower-trained animals have more stable within-session engagement than Legacy-trained animals.* **(A)(i)** Violation rates comparing the first 50 trials to the last 50 trials within each session during the fixation growth curriculum for each cohort (Legacy: dark purple, FixGrower: light blue). **(A)(ii)** As in (A)(i) but during the probe stages when task difficulty was stable and matched across cohorts. **(B)(i)** Trial initiation time following all trial types comparing the first 50 trials to the last 50 trials within each session during the fixation growth curriculum for each cohort. **(B)(ii)** Trial initiation times as in (B)(i) but during the probe stages when task difficulty was stable and matched across cohorts. *Statistics:* Circles represent group means and error bars indicate standard error of the mean. Statistical comparisons were conducted within groups (e.g., Legacy early vs. Legacy late) using *Wilcoxon signed-rank* tests. $*p < 0.05$, $**p < 0.01$, $***p < 0.001$ and ns is non-significant using Holm correction for multiple comparisons.

Lastly, we examined how differing starting fixation requirements and growth algorithms influenced center-poke timing, with a focus on impulsivity and motor control. The Legacy curriculum began each session with a minimal fixation requirement (10 ms) that incrementally increased after successful trials. While this structure eases animals into fixation, it

rewards short, rapid center pokes early on. Moreover, because fixation demands grew only after successful trials, previously rewarded movement patterns could become violations on subsequent trials, potentially disrupting learning through inconsistent reinforcement. In contrast, the FixGrower curriculum employed a stable, longer initial fixation requirement (400 ms) within each session, encouraging sustained fixation from the outset.

Given these structural differences, we hypothesized that Legacy-trained animals would display more impulsive center-poking behavior. We predicted this would manifest as temporal clustering around the go-cue, which would introduce behavioral noise and thereby hinder learning. To test this, we compared the timing of valid and violation pokes relative to the go-cue during the first week of fixation growth. This period was selected because it matched the time FixGrower-trained animals took to complete fixation growth, allowing for a temporally aligned behavioral comparison. Consistent with our hypothesis, poke-time distributions differed markedly between groups (**Fig. 3.9a**). Legacy-trained animals violated significantly closer to the go-cue compared to FixGrower-trained animals (**L:** mean $\pm$ SD $-0.211 \pm 0.082\,\mathrm{s}$, median: $-0.225\,\mathrm{s}$, $N = 8$; **FG:** $-0.622 \pm 0.193\,\mathrm{s}$, median: $-0.569\,\mathrm{s}$, $N = 9$; Welch's $t = 5.81$, $p = 1.12 \times 10^{-4}$; **Fig. 3.9b**). Additionally, a significantly higher fraction of violations in Legacy-trained animals occurred within the immediate 200 ms preceding the go-cue (**L:** $63.7 \pm 17.7\%$, $N = 8$; **FG:** $21.4 \pm 9.58\%$, $N = 9$; Welch's $t = 6.03$, $p = 1.04 \times 10^{-4}$; **Fig. 3.9c**), suggesting these animals were anticipating rather than reacting to the go-cue—likely a result of the curriculum differences highlighted above that fostered premature responses.

For valid pokes, FixGrower-trained animals exhibited significantly longer durations post-go-cue compared to Legacy-trained animals (**L:** $0.140 \pm 0.038\,\mathrm{s}$, median: $0.133\,\mathrm{s}$, $N = 8$; **FG:** $0.219 \pm 0.062\,\mathrm{s}$, median: $0.207\,\mathrm{s}$, $N = 9$; Welch's $t = -3.23$, $p = 6.27 \times 10^{-3}$; **Fig. 3.9d**), and these group differences were evident from the first day of fixation growth (Welch's t, $p_{\mathrm{bonferroni}} < 0.001$ for all days, valid and violation between cohorts; **Fig. 3.9e**). To assess

whether these differences could be attributed to the initial 20-trial "warm-up" in the Legacy curriculum, during which fixation requirements are closest to the go-cue, we reanalyzed the data excluding these trials. Results remained consistent, confirming that observed differences were not significantly influenced by early-session structure (**Fig. S3.5a–d**). Repeating analyses across the entire fixation growth phase (not just the first week) showed that Legacy-trained animals' valid poke durations gradually increased, eventually approaching those of FixGrower-trained animals (**L:** 0.198 ± 0.047 s, median: 0.192 s, $N = 8$; **FG:** 0.216 ± 0.068 s, median: 0.216 s, $N = 9$; Welch's $t = -0.650$, $p = 0.526$), whereas the violation poke metrics remained unchanged (**Fig. S3.5e–i**).

During the probe stages, all group differences in poke-timing disappeared (**Fig. 3.9f–i**). Legacy-trained animals exhibited poke distributions, average violation timing (**L:** -1.031 ± 0.070 s, median: -1.010 s, $N = 8$; **FG:** -1.036 ± 0.116 s, median: -1.008 s, $N = 9$; Welch's $t = 0.114$, $p = 9.10$), near-go-cue violations (**L:** $8.68 \pm 3.04\%$, $N = 8$; **FG:** $8.48 \pm 3.49\%$, $N = 9$; Welch's $t = 0.131$, $p = 0.897$), and valid poke durations (**L:** 2.15 ± 0.041 s, median: 0.195 s, $N = 8$; **FG:** 0.196 ± 0.071 s, median: 0.196 s, $N = 9$; Welch's $t = 0.705$, $p = 0.493$) similar to those of FixGrower-trained animals. Notably, FixGrower-trained animals showed consistent timing behavior across both the fixation growth and probe stages, indicating a stable motor strategy throughout training.

These findings suggest the Legacy curriculum placed animals at an initial disadvantage: by rewarding brief pokes and introducing shifting contingencies, it likely fostered impulsive habits and delayed the formation of stable fixation behavior. In contrast, the FixGrower curriculum provided consistent reinforcement for sustained fixation throughout training. This led to more deliberate poking, which directly contributed to increased training speed, as longer valid pokes led to larger growth increments for these animals.

In summary, the three curriculum features—*penalty timing, growth algorithm,* and *starting fixation requirement*—exerted multiple, interacting effects on behavior. Together, these

factors compounded to hinder learning in the Legacy curriculum while synergistically enhancing training efficiency and behavioral stability in FixGrower curriculum by modulating vigor, engagement and motor control.

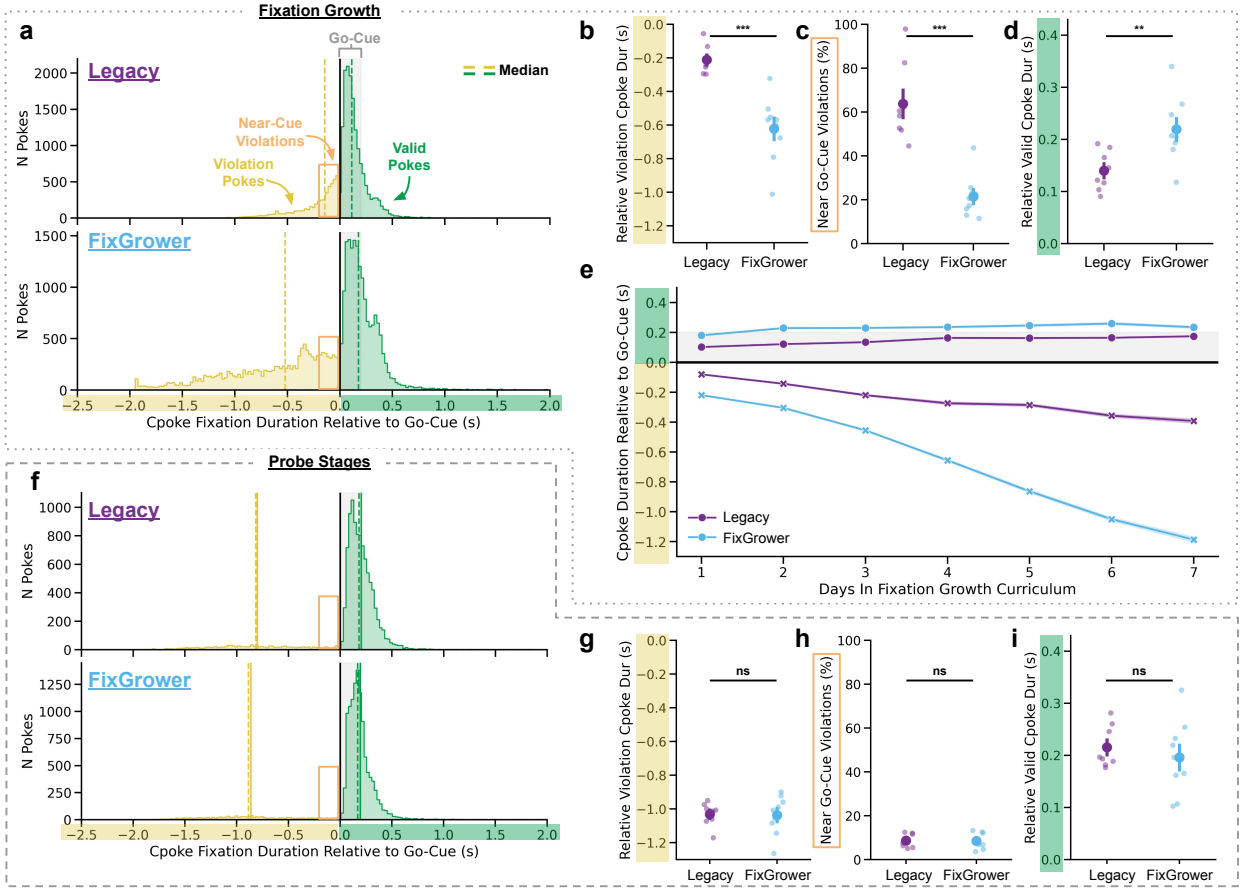


Figure 3.9: *Legacy-trained animals exhibit more impulsive center-poking behavior than FixGrower-trained animals during fixation training.* **(A)** Distributions of center-poke durations relative to the go-cue during the first 7 days of the fixation growth curriculum for Legacy (top, dark purple) and FixGrower (bottom, light blue) animals. Histograms show counts of all pokes with valid pokes colored green and violation pokes colored yellow. Dashed lines indicate respective medians. The orange box marks the "near-cue" period, defined as the 200 ms window immediately preceding the go-cue. **(B)** Mean violation poke durations relative to the go-cue for each cohort. **(C)** Percentage of violations occurring within the near-cue period defined in (A). **(D)** Mean valid poke durations (i.e., post-go-cue) relative to the go-cue for each cohort. **(E)** Daily averages of valid (green) and violation (yellow) poke times relative to the go-cue across the 7 days of fixation growth. Lines represent cohort means ± SEM. **(F)** Same as (A), but for the probe stages when curriculum differences were eliminated. **(G–I)** Same as (B–D), but calculated from probe session data. *Statistics:* All scatter/point plots display mean ± SEM, with individual animal data points overlaid. $*p < 0.05$, $**p < 0.01$, $***p < 0.001$, and ns is non-significant using Welch's t-tests.

3.3 Supplemental Figures

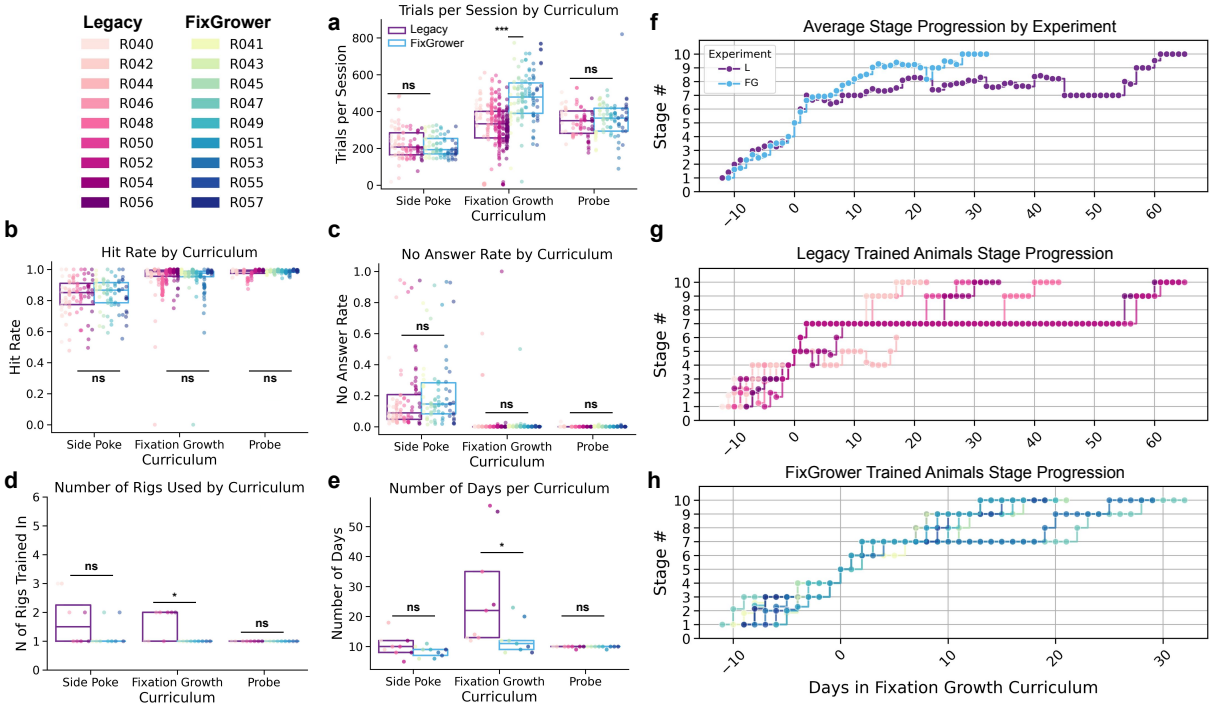


Figure S3.1: *Comparison of task metrics across Curricula Phases: Side Poke, Fixation Growth, Probe.* **(A)** Trials per session by curricula type compared by mixed linear model with a fixed effect for experiment and a random intercept for animal ID. Corrected for multiple tests with the Holm method. Experiment condition indicated by boxplot hue, individual animals indicated by dot color. Each dot indicates an animal's performance in a session within a given stage type. **(B)** Same as **(A)** but for hit rate where a hit is defined as a valid trial with a correct side poke to the light guide. **(C)** Same as **(A)** but for no answer rate where no answer is defined as a valid trial with no side poke during the answer period (8 seconds post go-cue). **(D)** Number of behavior boxes ("rigs") an animal trained in by stage type—see *Methods* for rig change rationale. Experiment condition indicated by box plot hue, individual animal indicated by dot color. Compared across experiment types by either *Mann–Whitney U* or *Welch's t* depending on Shapiro normality and corrected for multiple tests by the Holm method. **(E)** Same as (D) but for the number of days spent in each stage type. **(F)** Average stage progression by experiment group aligned to the start of the fixation growth curriculum (Stage 5). **(G)** Individual stage progressions for all Legacy-trained animals aligned to the start of the fixation growth curriculum. **(H)** Same as **(G)** but for FixGrower-trained animals. *Statistics:* non-significant indicated by ns, $p < 0.05$ indicated by *, $p < 0.01$ indicated by **, $p < 0.001$ indicated by ***.

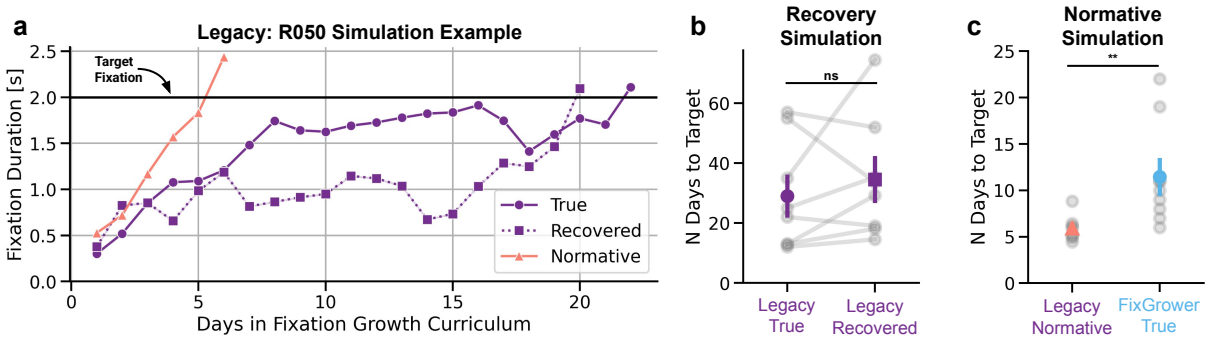


Figure S3.2: *Simulations of Legacy-trained animals' growth ceiling.* **(A)** Example simulation for a single animal (R050), showing fixation duration over days for the true data (dark purple circles), recovery simulation (dashed purple squares), and normative simulation (pink triangles). The black line marks the 2-second fixation target. See *Methods* for full simulation procedure. **(B)** Recovery simulation results for all Legacy-trained animals. Gray lines show individual animals (N = 9, $n_{\text{seeds}} = 30$) comparing true (circles) and recovered (squares) training durations. Purple markers indicate group means ± SEM. There was no significant difference in days to target between true and recovery conditions (*Wilcoxon signed-rank test*, $W = 12.5$, $p = 0.547$), supporting the validity of the simulation method. **(C)** Normative simulation results for Legacy-trained animals (pink triangles) compared to true FixGrower-trained animals (light blue circles). Gray dots indicate individual animals; colored markers represent group means ± SEM. Under normative conditions, Legacy-trained animals can reach the fixation target significantly faster than FixGrower-trained animals did, indicating they were not limited by an intrinsic growth-ceiling (*Mann–Whitney U*, $U = 5.000$, $p = 0.003$).

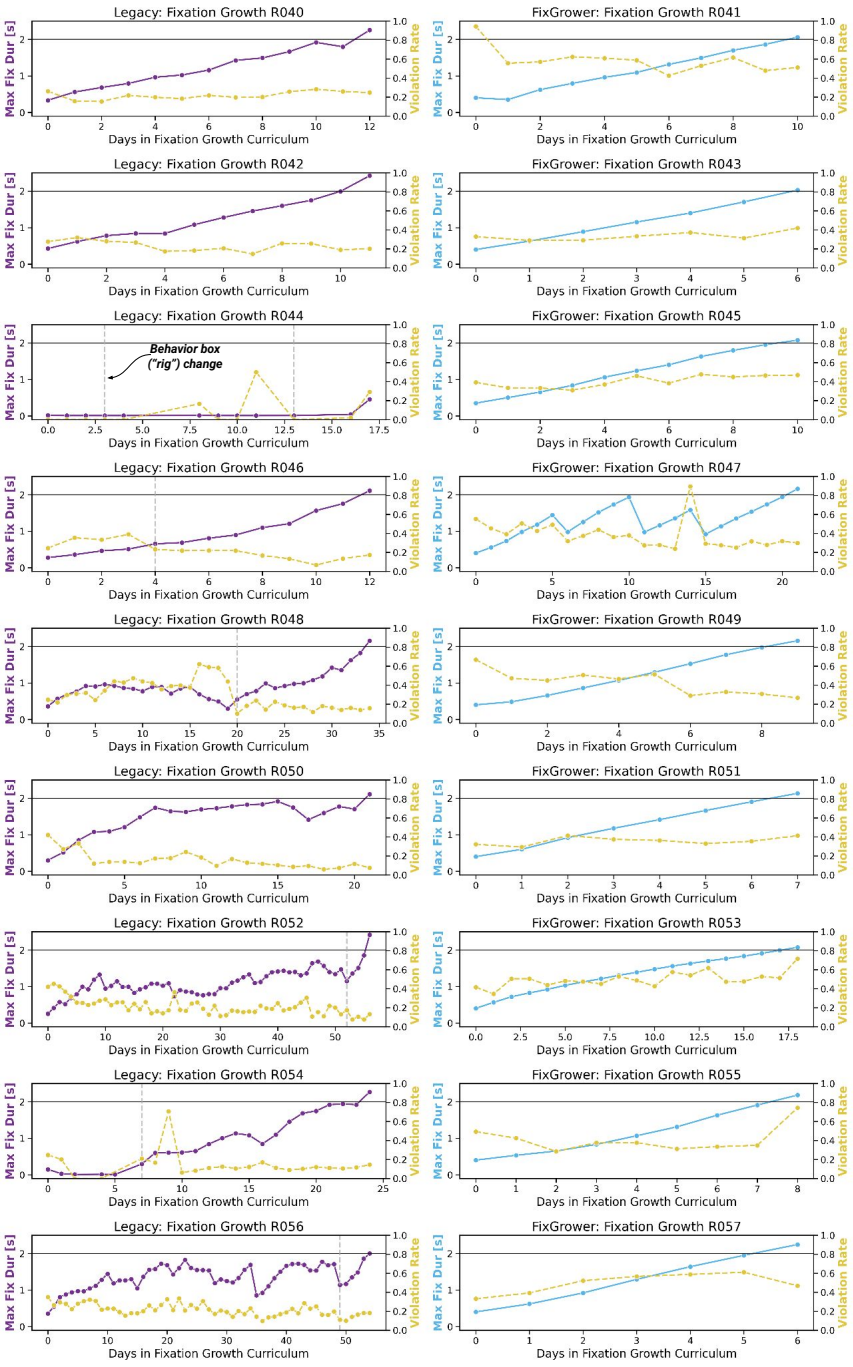


Figure S3.3: *Raw fixation, violation and behavioral apparatus data.* Data from sessions in fixation growth experiment for all Legacy and FixGrower-trained animals. The x-axis represents each day spent in the curriculum prior to reaching the 2-second fixation target. Left y-axis represents the maximum fixation duration reached in a given day, with the color (dark purple or light blue) indicating the experimental condition. Right y-axis indicates the violation rate for the given day. Gray dashed vertical lines indicate days in which a behavioral apparatus ("rig") was switched. Animal Legacy:R044 was dropped from analysis of growth timing and fixation performance due to equipment error and discontinued training. Animal FixGrower:R047 was an outlier in Figure 3.3 analysis.

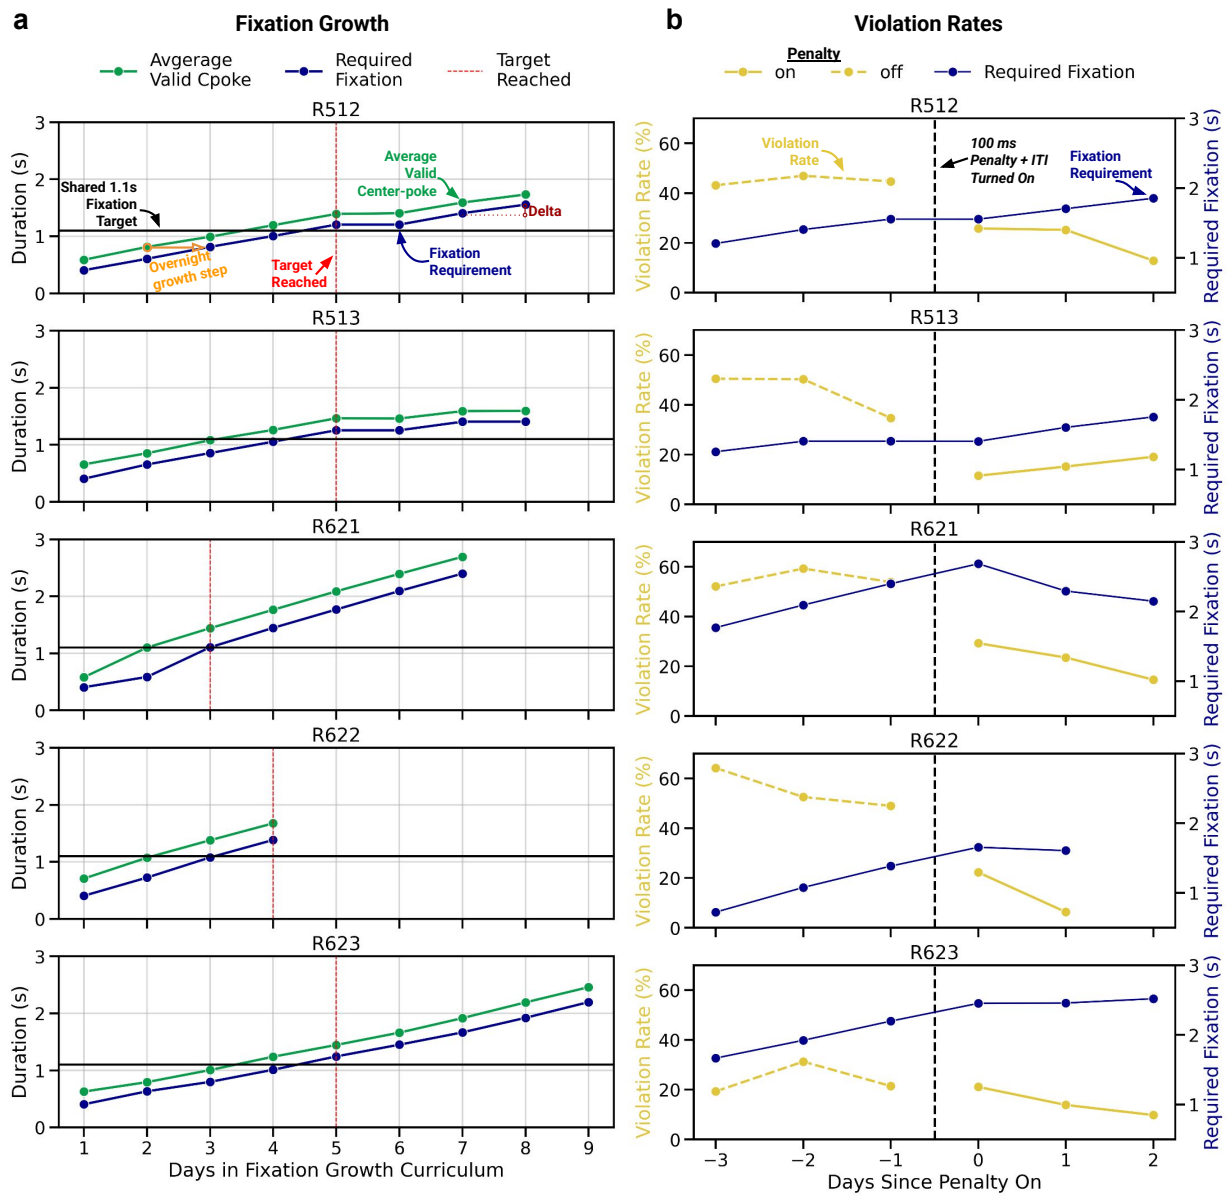


Figure S3.4: *Extended data from mouse fixation experiment.* **(A)** Fixation duration growth curves for individual mice across days in FixGrower. Green lines represent the average duration of valid center-pokes per day that was used to determine the next day's requirement (blue). Red dotted vertical lines mark the session in which each mouse first reached the shared 1.1-second fixation target (horizontal black line). The top panel (R512) includes annotations illustrating the curriculum's overnight stepwise growth (orange), and the session-to-session delta (maroon) used to assess growth rate. **(B)** Violation rate trajectories for each mouse before and after the onset of a 100 ms penalty and inter-trial interval (ITI). Yellow lines denote daily violation rates, with dashed lines representing pre-penalty sessions and solid lines post-penalty. The black dashed vertical line indicates the session when the penalty was introduced. Required fixation duration is overlaid in blue. Number of post-penalty sessions represents the number of days it took animals to reach a violation rate below 20% and perform at least 150 trials.

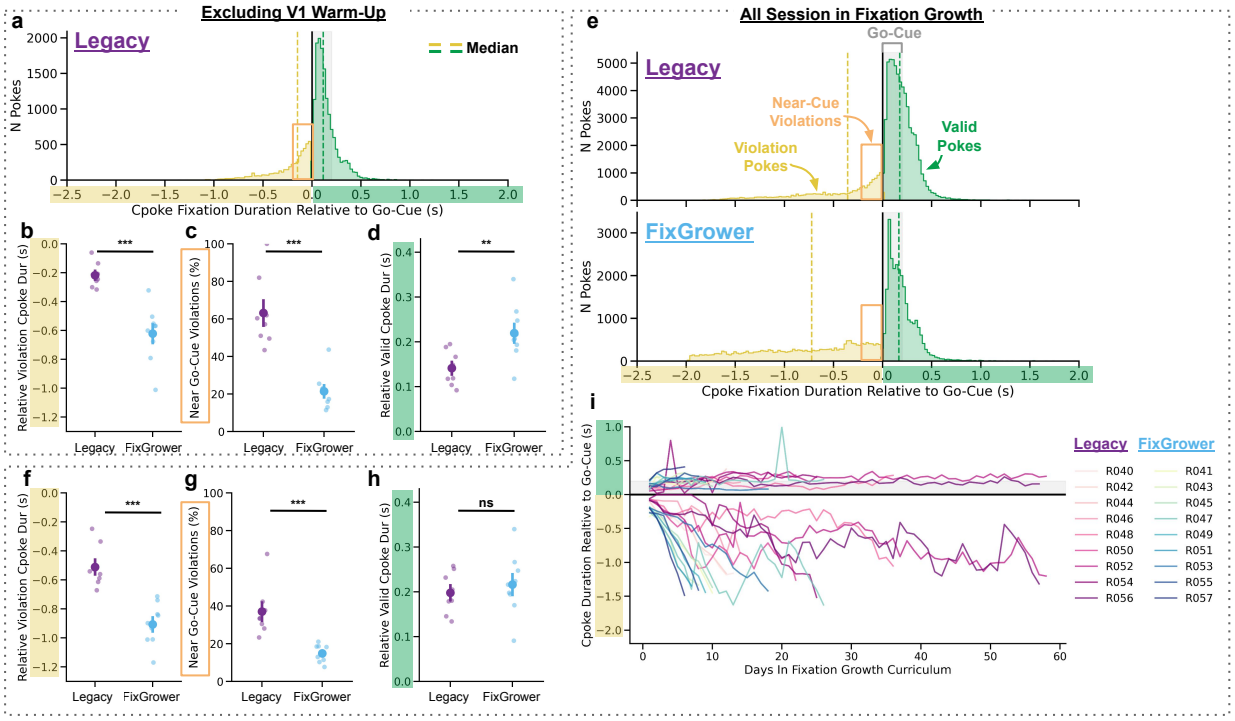


Figure S3.5: *Additional analysis accounting for warm-up exclusion and full fixation growth duration.* **(A)** Distribution of center-poke fixation durations relative to the go-cue for Legacy-trained animals after excluding the initial 20-trial "warm-up" period. Valid pokes are shown in green, violations in yellow. Dashed lines indicate medians. The orange box indicates the near-cue period (200 ms preceding the go-cue), used for quantification in (B, G). **(B–D)** Same metrics as in Figure 3.9B–D, computed after excluding the 20-trial warm-up for Legacy-trained animals. FixGrower animal data did not change from Figure 3.9. Data is only from the first 7 days of fixation growth. **(B)** Mean violation poke times relative to the go-cue; **(C)** percent of violations occurring within the near-cue window; and **(D)** valid poke durations relative to the go-cue. Statistical tests do not differ when warm-up trials are excluded. **(E)** Poke duration distributions for valid and violation trials across the full fixation growth period for both Legacy (top) and FixGrower (bottom) animals, which is an extension from the first 7 days shown in Figure 3.9. **(F–H)** Same as (B–D), but calculated from all sessions and trials during fixation growth. **(I)** Center-poke durations for average valid and violation pokes relative to the go-cue across fixation growth days for individual animals in both cohorts. Each line represents a single animal. *Statistics:* All scatter/point plots show mean ± SEM with individual animal data points overlaid. $p < 0.05$, $p < 0.01$, $p < 0.001$, and ns indicates non-significance using *Welch's* t-tests.

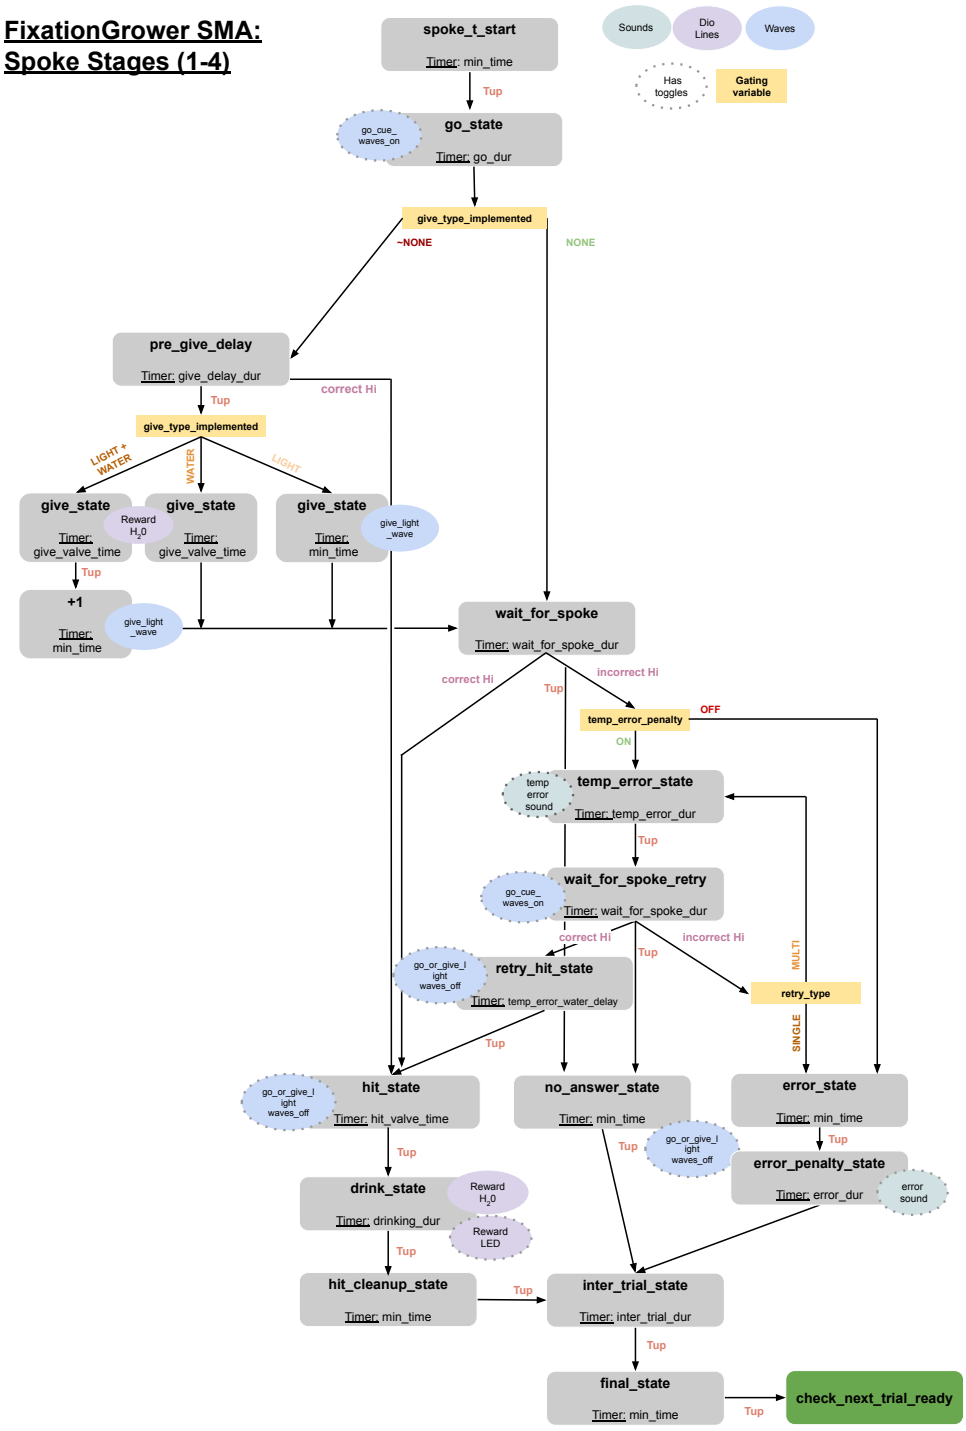


Figure S3.6: *Visualization of the state machine logic written in Bcontrol for the side-poking curriculum stages.*

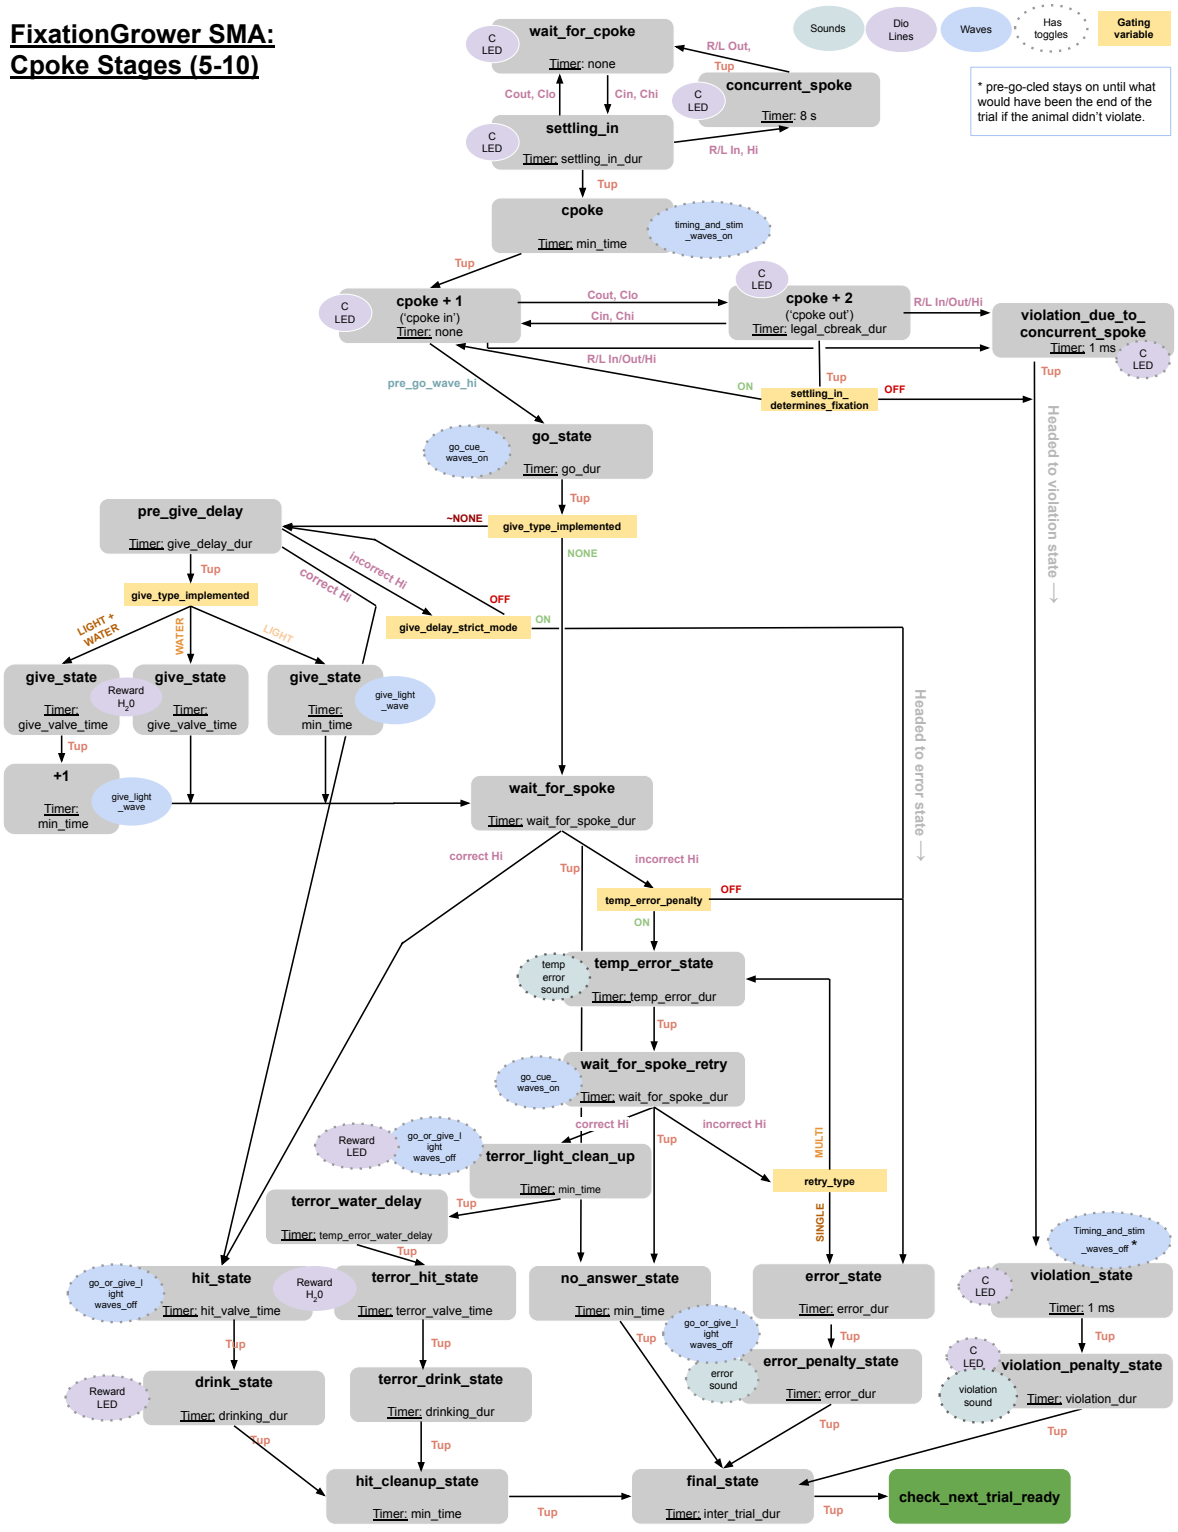


Figure S3.7: *Visualization of the state machine logic written in Bcontrol for the center-poking stages (Fixation Experiment stages 5–11).*

3.4 Discussion

3.4.1 Findings

Perceptual decision-making tasks require animals to learn both stimulus–response rules and structured motor behaviors. One such behavior is center-port fixation, where a freely moving animal must initiate and maintain a nose poke at a central port for a specified duration. Although fixation is a common task requirement, the procedures used to train this behavior vary widely and often require substantial time, presenting an opportunity for optimization and standardization. In response to this, we developed a novel curriculum for shaping fixation behavior in rodents, designed to minimize both training duration and premature-movement fixation errors termed "violations". Our approach, called "FixGrower", differed from commonly used methods in three key ways: it began with a long initial fixation duration of 400 ms, it increased the requirement adaptively across days rather than trials, and it delayed the introduction of violation penalties until the animal reached the target fixation duration.

To evaluate the effectiveness of this curriculum, we trained two cohorts of rats to reach a 2-second fixation target. One cohort ($N = 8$) followed a curriculum commonly used in our lab and similar to those in the field which we called ("Legacy", while the other ($N = 9$) was trained using the new FixGrower approach (**Fig. 3.1**). FixGrower-trained animals reached the target duration significantly faster than Legacy-trained, representing a 61% reduction in training time on average (**Fig. 3.2c**).

We then assessed violation rates under both stable and variable fixation conditions and found that FixGrower-trained animals maintained low violation rates, averaging 13.2% (**Fig. 3.3a**). Moreover, violation rates and training times were positively correlated in the FixGrower curriculum, providing experimenters with a straightforward metric for predicting cohort performance and setting training cutoffs (**Fig. 3.3c**). These improvements proved to

be robust: in a subset of animals trained for over 100 days following FixGrower completion, violation rates remained consistently low, averaging 9.9% (**Fig. 3.4**).

To confirm that these low violation rates were not the result of a simplified task structure, we trained a new cohort of rats on a sound discrimination task that following FixGrower training (**Fig. 3.5a–b**). Even with the added requirement of processing task-relevant stimuli during fixation, violation rates remained low at an average of 11.2% (**Fig. 3.5e**), which is notable given the findings presented in **Chapter 2** and additional research showing that having lower violations is associated with improved stimulus rule performance (Reuschenbach et al., 2023).

To our knowledge, only one study has reported faster fixation acquisition than our approach. Odoemene et al. (2018) trained mice to reach a 1.1-second fixation duration within 1–2 sessions by delivering a small water reward (0.5 μl) at the center port during shaping. However, without this mid-trial reward, their procedure, which was structurally similar to our Legacy curriculum, required 10–12 sessions to reach the same target. While their strategy is highly effective, it introduces both logistical and interpretive challenges: introducing center-port water solely for fixation growth adds complexity to behavioral setups, and rewarding animals prior to stimulus presentation may alter motivational state, affect neural activity, and/or complicate downstream learning.

To evaluate whether our approach could provide a practical alternative, we applied the FixGrower curriculum to a cohort of mice and trained them to a 1.1-second fixation duration. These animals reached the target in just 3–5 days (**Fig. 3.6b**), and a subset of these animals progressed to fixation durations beyond 2 seconds within 9 days (**Fig. S3.4a**). Importantly, their violation rates were equivalent to or lower than those published in the field (Ashwood et al., 2022; Cruz et al., 2022; Jaramillo & Zador, 2014; Odoemene et al., 2018)(**Fig. 3.6d**). These results underscore the utility of the FixGrower curriculum as a robust, scalable solution for shaping fixation behavior in both mice and rats—achieving high performance while

preserving standard equipment and task structure.

3.4.2 Rationale

Understanding why the FixGrower curriculum accelerated fixation training is essential for refining future protocols and contextualizing these results within broader frameworks of learning. To this end, we compared behavioral dynamics during the differing fixation growth stages with those in the probe stages, where task parameters were held constant across experimental cohorts (**Figs. 3.7–3.9**). All three curriculum elements—violation penalty introduction, fixation growth algorithm, and initial fixation requirement—appeared to shape learning by modulating motivation, engagement, and motor control. Moreover, we draw on theories from operant conditioning, reinforcement learning, and curriculum learning to interpret these findings.

During fixation growth, Legacy-trained animals completed 30% fewer trials per session than FixGrower-trained animals, despite achieving similar numbers of rewarded trials (**Fig. 3.7a–b**). This discrepancy was driven by slower trial re-initiation after violations in Legacy-trained animals (**Fig. 3.7c**); a pattern we hypothesized was influenced by the higher cost of failure imposed by the timeout penalty present only in the Legacy curriculum. When the penalty was introduced in the FixGrower curriculum after reaching target fixation, violation rates dropped markedly and trial initiation times slowed to Legacy-like levels (**Fig. 3.7e–f**). This behavioral inflection supports theories from operant conditioning suggesting that excessive early punishment can reduce motivation (Skinner, 1953), and is consistent with findings that animals re-engage more quickly post-failure when expected reward is high (Mah et al., 2023). Taken together, these results show that early shaping should prioritize rapid re-attempts post-failure and preserve sizable penalties until later in training so as to keep motivational drive high and facilitate initial learning.

Notably, decreasing vigor and motivation among Legacy-trained animals emerged not

only across sessions but within them. Over the course of individual sessions, these animals exhibited increasing violation rates and slower trial initiation, consistent with reduced engagement (**Fig. 3.8a(i), b(i)**). We hypothesize this arose from the Legacy curriculum's incremental fixation growth algorithm, which steadily increased task difficulty each trial, likely contributing to cumulative fatigue. Since both trial count and violation rate determine the growth magnitude in the Legacy curriculum, this within-session disengagement leads to slower fixation growth. In contrast, FixGrower-trained animals maintained stable engagement over time (**Fig. 3.8a(i), b(i)**), and both groups showed comparable behavior during probe sessions when difficulty was static (**Fig. 3.8a(ii), b(ii)**), reinforcing the idea that increasing task demands negatively impacted the Legacy cohort.

We further hypothesized that the variable demands of the Legacy curriculum impaired motor control. Because fixation duration increased only after successful trials, a fixation hold that was previously rewarded could be penalized on the very next attempt. This inconsistency mirrors the challenge of learning under non-stationary targets in reinforcement learning, where shifting objectives can destabilize learning by introducing ambiguity into the action–outcome mapping (Sutton & Barto, 2018). From this perspective, each increase in fixation requirement after a valid trial effectively changed the rule the animal was attempting to learn. Behaviorally, this was reflected in Legacy-trained animals exhibiting a smaller separation in timing between failed and successful pokes (**Fig. 3.9a–e**), and a greater concentration of violations immediately preceding the go-cue—suggesting anticipatory, impulsive behavior rather than deliberate, rule-based control.

One well-established solution to the problem of non-stationary targets in artificial learning systems is the use of a separate, slowly-updated target network (Mnih et al., 2015). This strategy stabilizes learning by keeping the goal ("target") fixed for a period of time, allowing the agent to learn from consistent feedback instead of constantly adjusting to shifting objectives. Analogously, the FixGrower curriculum achieved behavioral stability by holding

the fixation requirement constant within each session—effectively implementing a fixed "target" rule over a session. This provided animals with a consistent behavioral contingency to learn from, thereby promoting clearer reinforcement signals. Just as deep Q-learning benefits from steady targets, keeping task difficulty constant within sessions appears important for sustaining engagement and supporting efficient learning.

Finally, we propose that beginning with an extremely short (10 ms) fixation duration in the Legacy curriculum may have reinforced exploitation, non–goal-directed behaviors that slowed learning. The underlying assumption of the Legacy design—that fixation is a continuum from short to long holds—may not align with how animals conceptualize these actions. Instead, we hypothesize that short pokes and sustained holds constitute qualitatively distinct behaviors. By initially reinforcing rapid pokes, the Legacy curriculum may have inadvertently encouraged habits that later conflicted with the task's demands. This is reflected in the Legacy-trained animals' impulsive poke profiles and shorter valid hold durations following the go-cue (**Fig. 3.9a–e**). As these animals progressed to longer fixation requirements—and later into the probe stages—their post–go-cue hold durations increased (**Fig. S3.5e–i**), suggesting that they gradually abandoned the initial strategy of rapid, impulsive poking in favor of more deliberate fixation holds.

By contrast, FixGrower's initial fixation requirement was critically set to exceed the duration of a quick nose flick ($\sim$250 ms), yet remain attainable. This structure increased the salience of the goal behavior by only rewarding meaningful approximations of the desired behavior. This approach aligns with the curriculum learning principle of using carefully chosen examples that are informative yet achievable (Bengio et al., 2009). By shaping behavior through reinforcement of stable, goal-relevant responses, the FixGrower design facilitated more deliberate motor control and efficient learning.

Together, these findings highlight how subtle structural choices in early training can shape the trajectory of learning by influencing the clarity, consistency, and magnitude of

feedback. The Legacy curriculum imposed high penalties early and offered less stable reinforcement contingencies, leading to reduced engagement and noisier learning. In contrast, the FixGrower curriculum supported behavioral shaping through milder early consequences and clearer, more consistent reinforcement.

3.4.3 Limitations & Future Directions

A key limitation of this study is the inability to isolate the specific contributions and interactions of the three curriculum features that contributed to improved learning efficiency using the FixGrower curriculum. Future work could address this by conducting a systematic "knockout" analysis, in which animals or theoretical models are trained with subsets of the features to examine how behavioral metrics are affected.

A second limitation concerns task generalization across different delay durations. Some paradigms, such as those used in Akrami et al. (2018), incorporate delays exceeding 2 seconds, which may not be efficiently supported by overnight growth alone. To accommodate longer durations, we recommend supplementing overnight growth with intra-session increases—e.g., incrementing fixation duration every 100–150 trials once 1–2 seconds are reliably achieved. Experimenters should monitor motivational levels closely and consider implementing 5–10 warm-up trials to support adaptation to extended delays.

Third, not all behavioral tasks require or benefit from an auditory go-cue (e.g., Brunton et al., 2013). While we view the auditory cue as an important feature for supporting this fixation task, further investigation is needed to identify alternative salient cues or to develop strategies for phasing out the cue once fixation behavior is stable.

In summary, we present FixGrower: a novel curriculum-based approach for training center-port fixation in rodents that significantly reduces training duration while maintaining low rates of premature movement. Our results are robust across tasks and species, and grounded in principles from operant conditioning, curriculum learning, and reinforcement

learning. We anticipate that this work will be of broad utility for the efficient design and implementation of behavioral training protocols in psychology and neuroscience research.

3.5 Methods

3.5.1 Animal Subjects

All procedures were approved by the Institutional Animal Care and Use Committee at Princeton University and were carried out in accordance with National Institutes of Health standards.

Rats

We used male Long-Evans rats (*Rattus norvegicus*) between the ages of 4 and 7 months for this study. All rats were pair housed and had *ad libitum* access to standard rodent chow. Rats did not begin water restriction and training until their weight exceeded 250 g. For at least one week prior to training, animals were brought to the training room daily for weighing. Once on restriction, water was available during the behavioral session (28–60 μl for each correct trial depending on animal and stage) and for 1 hour at the end of the session. Rats were removed from water restriction if they showed signs of illness, failed to consume at least 3% of their body weight in water over three consecutive days, lost $\geq$10% body weight in 24 hours, lost $\geq$20% body weight over 2 weeks or less, or exhibited a statistically significant trend of $\geq$1.5% weight loss per week over 30 days.

Mice

We used male and female *B6.CAST-Cdh23Ahl+/Kjn* mice (The Jackson Laboratory, stock #002756). Animals were group-housed by sex in standard ventilated cages (2–5 animals per cage) under a 12-hour light/dark cycle with *ad libitum* access to standard rodent chow.

Upon arrival, animals were acclimated to the housing environment and handled using a non-aversive tunnel-based approach to reduce stress, as described by Gouveia and Hurst (2013). Briefly, a tunnel (Kent Scientific) was placed in each home cage to allow voluntary entry, and over the course of a week, mice were gently habituated to being picked up first

via the tunnel and then by hand.

Once mice were at least 8 weeks old, seven days prior to the start of behavioral training, they were weighed to establish a baseline body mass. At this time, they were transitioned to 2% citric acid (Thermo Scientific Chemicals, CAS: 77-92-9) to begin water restriction (Urai et al., 2021). From this point forward, mice were weighed daily, and a daily Body Condition Score (BCS) assessment was conducted to monitor animal welfare (Burkholder et al., 2012). Following the one-week acclimation to water restriction, animals began behavioral training. During behavioral sessions, water was available as a reward for correct trials (3–10 μl per trial, depending on the animal and training stage). If a mouse failed to consume at least max (1 mL, 4% of baseline body weight), it was supplemented with the remaining water to reach this threshold post training session. On non-training days, animals were provided *ad libitum* access to 2% citric acid water. If an animal did not train for seven consecutive days, it received 1 hour of *ad libitum* access to water. If an animal's weight dropped more than 10% below baseline, *ad libitum* water access was reinstated until the weight stabilized. Additionally, if an animal lost more than 1 g of body mass within 24 hours, dehydration was assessed via skin turgor. If dehydration was indicated, the animal was supplemented with free water.

3.5.2 Behavior Training

Apparatus

Behavioral training took place in custom-designed operant conditioning boxes housed within sound-attenuating chambers equipped with active ventilation. Each training box featured three conical nose ports arranged in a horizontal row (Left – Center – Right; **Fig. 3.1a**). Planar tweeters (Bohlender Graebener Neo3W) were mounted above the left and right nose ports to deliver auditory stimuli. Each port was equipped with a controllable white LED for visual cues and an infrared beam to detect nose pokes.

The left and right ports were fitted with solenoid valves (Lee Valve Company, LHDA1231115H) connected to sipper tubes that dispensed distilled water as a reward, sourced from a 10-gallon gravity-fed tank. The center port had an enlarged opening to facilitate nose fixation, allowing rats to maintain position comfortably.

All hardware was controlled by a Bpod 2.3 system (Sanworks) interfaced with a Windows computer running Bcontrol, a custom MATLAB-based behavioral training software (https://tinyurl.com/BrodyBcontrol, https://tinyurl.com/BrodyExperPort). Auditory stimuli were generated via an Analog Output Module (Sanworks), amplified by a mini class D stereo amplifier (Leapai LP-2020AD), and delivered through the planar tweeters.

Mouse training boxes were similarly configured but downsized in proportion to the animals' size. In the mouse setup, water reward was supplied from a 100 mL syringe located within the chamber and refilled after each session.

Training Overview

Eighteen rats participated in the primary fixation training study, divided into two cohorts that experienced differing fixation growth curricula: Legacy ($N = 9$) and FixGrower ($N = 9$). Rats were pair-housed such that each pair included one Legacy and one FixGrower animal. Water restriction and training commenced once both animals in a pair exceeded 250 g body weight.

Training sessions were conducted daily (7 days per week) and lasted between 1.5 and 3 hours per session. With few exceptions (e.g., equipment failure), each animal was trained in a consistent behavioral box throughout the experiment. An effort was made to distribute training such that both cohorts were represented across the available boxes over the course of the experiment. For example, if in one session a Legacy rat trained in Box #10 and a FixGrower rat in Box #20, a subsequent session might place a Legacy rat in Box #20 and a FixGrower rat in Box #10. This strategy was implemented to help mitigate any

potential box-specific effects. All animal handling—including weighing, training, and health checks—was performed by technicians blinded to the experimental conditions. The only experimental variable differing between cohorts was the specific structure of the fixation growth training curricula.

An additional cohort of 13 rats trained on a sound discrimination task followed the same training structure and procedures. The mouse cohort ($N = 5$) also adhered to this protocol with two exceptions: sessions lasted 1 hour, and training began once mice were at least 8 weeks old and had been on 2% citric acid for at least seven days.

3.5.3 Curricula

Overview of Stages for Fixation Experiment

Automated behavioral training was conducted using Bcontrol, a custom MATLAB-based software package (MathWorks, R2019b). A unified protocol file (`FixationGrower`) controlled the overall task structure and parameters for all animals. Stage-specific training progressions were governed by customized Training Section files, which defined the curricula and were the only components to differ between the experimental cohorts. Animals progressed through these three main experimental phases (**Fig. 3.1b**, **Table 3.1**):

(i) A **side poke curriculum** to facilitate reward learning and habituation.

(ii) A **fixation growth curriculum**, either "Legacy" or "FixGrower".

(iii) A set of **"probe" evaluation** stages following completion of fixation growth.

Each training stage is described in detail in the subsequent methods sections. Importantly, the only difference between curricula lies within the fixation growth stages (Stages 5–8 in **Table 3.1**, highlighted in red); all other training phases and parameters were identical across cohorts. For a detailed account of the underlying state-machine logic implemented in the protocol, refer to **Figs. S3.6** and **S3.7**.

Stage	Description	Shared	Legacy Only	FixGrower Only
1	Learn L Spoke	All Logic		
2	Learn R Spoke	All Logic		
3	Alternate L/R Spoke Blocks	All Logic		
4	L/R Spoke Blocks → Random	All Logic		
5	Center Fixation to Side Blocks with Growth	Trial structure, blocks, stage completion logic	Fix. start @ 10ms, grow by valid trial, violation penalty on	Fix. Start @ 400 ms, Grow overnight
6	Center Fixation to Side Blocks → Random with Growth	Trial structure, blocks, stage completion logic	Fixation warm up, grow by valid trial, violation penalty on	Grow overnight
7	Center Fixation to Side → Random with Growth	Trial structure, stage completion logic	Fixation warm up, grow by valid trial, violation penalty on	Grow overnight
8	FG Violation Penalty On	None- FG only		Violation penalty on
9	Probe: Stable Fix. Delay	All Logic		
10	Probe: Random Fix. Delay	All Logic		

Table 3.1: *Overall description of training stages. Similarities and differences are noted for each stage.* Red indicates the stages of the fixation growth curriculum with stages that differ between conditions.

Shared: Side Poking

The first 4 training stages are shared between cohorts and serve the purpose of habituating the animals to the behavior box and water reward while ensuring they are motivated to perform trials (**Table 3.2**).

Stage	Side(s)	Water	Retry	ITI	Completion
1	L	Seed 18 (3) µL, Reward 42 (10) µL	Y	N(40 s, 5 s)	> 40 t, NA < 15% within session
2	R	Seed 18 (3) µL Reward 42 (10)µL	Y	N(40 s, 5 s)	> 40 t, NA < 15% within session
3	L/R 15 t blocks	Reward 45 (6) µL	N	N(25 s, 3 s)	@EOD > 90 t, Hit > 87%, NA < 20%,
4	L/R 15 t blocks → rand w/ AB @ 70 t	Reward 30 (3) µL	N	N(15 s, 3 s)	@EOD > 175 t, Hit > 80%, NA < 20%

Table 3.2: *Summary of variable parameters across side poke learning stages.* The go-cue and side reward LED are present in all stages. Under "Sides," AB denotes anti-bias. Under "Water," values in parentheses indicate mouse reward volumes. Under ITI "N(X,Y)" indicates sampling from a normal distribution with mean X and standard deviation Y. Under "Completion," 't' refers to trials, and EOD (End of Day) indicates that stage transitions occur only at the end of the session and are implemented in the following session.

Stage 1– Learn Left (L) Poking

The structure for Stage 1 trials is the following: go-cue (200ms white noise at 70 dB, bilateral) → L side-port LED on → L side-port water seeded (30%, 18 μL (mouse: 3 μL)) → wait for animal to poke in L port (max 8 s) → animal pokes in L port → L side-port remaining water delivered (70%, 42 μL (mouse: 10 μL)) & LED turned off → trial inter-trial interval (ITI) sampled from a normal distribution with mean 40 s, standard deviation 5 s. If the animal does not poke in the left port during the wait-for-poke period it is considered a "no answer" trial. In this case, the remaining reward is not delivered, the L side light turns off, and the trial moves into the ITI period. On the trial following a no answer trial, water

is not seeded in the L port (we assume it still remains from the previous trial). If the animal pokes in the incorrect port (right) during the wait-for-poke period, there is no penalty and they can immediately retry within the same wait period. To complete this stage, animals must complete 40 trials and have a no answer rate $< 15\%$. They can move to the next stage within the same session. If the animal does not complete this stage on the current day's session, and they have yet to complete Stage 2, they will move into Stage 2 for the following session such that they alternate between Stage 1 and Stage 2 over days. If they have completed Stage 2, they will stay in Stage 1 until completed before moving to Stage 3.

Stage 2– Learn Right (R) Poking

The structure for Stage 2 is identical to that for Stage 1 except that the right (R) port is used. Additionally, the same stage completion logic follows such that if they fail to complete Stage 2 within the session, they will return to Stage 1 on the following day if Stage 1 has yet to be completed. If Stage 1 is completed, they will remain in Stage 2 until completed. Once Stage 1 and Stage 2 are *both* completed, the animal will move into Stage 3.

Stage 3– Learn Left & Right Poking in Blocks

In Stage 3, the animals will alternate between 15-trial blocks of left and right poking. The structure for Stage 3 trials is the following (assume in L block): go-cue sound $\rightarrow$ L side-port LED on $\rightarrow$ wait for animal to poke in L port (max 8 s) $\rightarrow$ animal pokes in L port $\rightarrow$ L side-port water reward delivered (100%, 45 μL (mouse: 6 μL)) $\rightarrow$ trial ITI sampled from a normal distribution with mean of 25 seconds, standard deviation of 3 seconds. In this stage, the ports are no longer seeded with water. If the animal pokes incorrectly during the wait-for-poke period, this is considered an "error" trial. Similar to a no answer trial, the side LED turns off immediately and the trial moves into the ITI period. To complete this stage, animals must do at least 90 trials (6 blocks) with $> 87\%$ accuracy and a no answer rate $< 20\%$. Note that accuracy is only computed on trials that were answered and stage advancement only occurs at the end of the session (i.e., overnight) for this stage and all

future stages.

Stage 4– Learn Left & Right Poking in Blocks to Random

In Stage 4, the animals will alternate between 15-trial blocks of left and right poking until they reach 70 trials. Then, each trial will be randomly drawn with an anti-biasing procedure which dynamically adjusts trial probabilities to counteract subject side-biases by considering past trial outcomes. It uses an exponential kernel with a decay parameter, τ (default = 30 trials), to estimate recent performance. It then modifies the probability of selecting left or right trials based on this history. The adjustment strength is controlled by β (default = 3), which is initialized with a linear warm-up over τ trials. The value of β is adjusted on a per-animal basis and ranged from 3 to 4.5 in our experiments.

The only adjustments to the trial structure from Stage 3 are the following: the water reward is 30 μL (mouse: 3 μL), and the ITI is sampled from a normal distribution with mean of 15 seconds, standard deviation of 3 seconds (**Table 3.2**). To complete this stage, animals must do at least 175 trials in a session with $> 80\%$ accuracy and $< 20\%$ no answer rate. Animals will then move into Legacy or FixGrower curriculum depending on the experiment.

Varied: Fixation Growth

The core of this experiment is a comparison between two training curricula designed to shape motor fixation behavior. Animals were trained to reach a target fixation duration of 2 seconds using one of two curricula: Legacy or FixGrower. The curricula differ in three key features: the initial fixation duration, the algorithm governing fixation growth, and the penalty for fixation violations (**Table 3.3**; **Fig. 3.1c-e**).

	Legacy	FixGrower
Initial Fixation Value	Every Day: 0.01 s	Day 1: 0.4 s Day > 1: previous session's average valid cpoke duration
Growth Algorithm	By valid trial: max(1 ms, 0.1%) following 20 trial warm up towards yesterday's fix. maximum	By session: To the average valid cpoke duration of the most recent session
Violation Penalty	2 second timeout followed by ITI	No penalty or ITI until stage 8, then 2 second timeout followed by ITI

Table 3.3: *Overview of training parameters that differ between the two fixation learning stages.*

Aside from these differences, both curricula follow the same overall stage progression, as outlined in **Table 3.4**. In the additional rat and mouse experiments presented in Figures 5 and 6, the FixGrower curriculum was used, with slight modifications to stage 8 that are described in the corresponding sections.

Stage	Side(s)	Completion
5	L/R 15 t blocks	@EOD 200 t, Hit > 85%, NA < 10%
6	L/R 15 t blocks → rand w/ AB @ 70 t	@EOD 250 t, Hit > 85%, NA < 10%
7	rand w/ AB	@EOD 200 t, Hit > 85%, NA < 10%, final fixation dur > target

Table 3.4: *Detailed overview of stage-wise variable changes shared by both Legacy and FixGrower curricula.* Stage 8 is excluded, as it applies only to FixGrower-trained animals. "AB" denotes anti-bias; "t" indicates trials; "EOD" (End of Day) signifies that stage transitions occur only overnight.

Legacy Fixation Growth (Stages 5–7)

Animals in the Legacy experimental cohort progressed through three structured stages to learn the center poke fixation-to-side poke reward motor sequence. These stages were adapted from standard shaping procedures used both in our lab and more broadly in the field (Akrami et al., 2018; Constantinople et al., 2019; Mah et al., 2023; Schmid et al., 2024; Scott et al., 2015). Specifically, they closely follow the "Grow Nose Poke" curriculum described in Akrami et al. (2018), with the primary modification being that the daily warm-up stage was performance-based rather than fixed. This adjustment was made to reduce early-stage dropout, which was frequently observed anecdotally due to animals disengaging from the task.

Legacy Stage 5– Center to Side, Blocks with Fixation Growth

In Legacy Stage 5, the animals are required to first poke in the center-port for a required fixation period before receiving water reward localized by LED light in one of the side-ports, which alternates in blocks of 15 trials. The structure for these trials is the following (assume in L block): Center-port LED indicates trial availability $\rightarrow$ animal pokes & holds nose in center $\rightarrow$ go-cue sound plays & Center LED off $\rightarrow$ L LED on & wait-for-poke period (max 8 s) $\rightarrow$ animal pokes in L port $\rightarrow$ L side-port water reward delivered (30 μL, mouse: 3 μL) $\rightarrow$ trial ITI sampled from a normal distribution with mean of 5 seconds and standard deviation of 2 seconds. If the animal leaves the center-port before the go-cue plays, or if the animal concurrently pokes in a side-port during fixation, the trial is considered a "violation." In this case, the center-port LED stays on only for the intended fixation period, the go-cue does not play, and there is a 2-second penalty timeout. Once the timeout is complete, the trial moves into the ITI period. If the animal stays in the center-port until the go-cue (i.e., a valid trial) but then pokes in the unlit side-port, it is an error trial. If the animal does not poke during the wait-for-poke period, it is a no answer trial. For both error and no answer trials, the

side LED turns off immediately and the trial moves into the ITI period.

For the first session in this stage, the starting fixation duration is 0.01 s. After a valid trial, the required fixation duration grows by min(1 ms, 0.1%). In other words, the fixation grows linearly by 1 millisecond until the fixation duration exceeds 1 second, and then grows exponentially by 0.1% of the current duration. Assuming the animal performed valid trials in the first session, for the second session—and all remaining sessions in this curriculum—there is a 20-trial fixation warm-up starting at 10 milliseconds and growing toward the previous session's final fixation value. The growth size in this warm-up block is computed as:

$$\frac{\big(\text{prev session final fixation value} - 0.01\,\text{s}\big)}{20}.$$

The same min-growth logic applies in the warm-up block: growth only occurs after valid trials. To complete this stage, animals must perform at least 200 trials in a session with $>$ 85% accuracy and $<$ 10% no answer rate. Note that accuracy is computed only on valid trials.

Legacy Stage 6– Center to Side, Blocks to Random with Fixation Growth

Legacy Stage 6 is identical to Legacy Stage 5 except that the blocked L/R trials end at trial 70 and then become randomly sampled with the anti-biasing procedure (as described in Stage 3). All other trial structure (fixation, violation rules, reward size, ITI distribution) remains the same. To complete this stage, animals must perform at least 250 trials in a session with $>$ 85% accuracy and $<$ 10% no answer rate.

Legacy Stage 7– Center to Side, Random with Fixation Growth

Legacy Stage 7 is identical to Stages 5 and 6 except that all trials are randomly sampled with anti-biasing from the start. To complete this stage, animals must have grown to the final fixation target duration (2 seconds) and performed at least 200 trials in that session, with $>$ 85% accuracy and $<$ 10% no answer rate. Note that only the final trial in the session

needs to reach $\geq$ 2-seconds fixation duration for the stage to be considered completed. From this stage, Legacy animals move directly into probe Stage 9.

FixGrower Fixation Growth (Stages 5–8)

Animals in the FixGrower experimental cohort progressed through 4 structured stages to learn the center poke fixation-to-side poke reward motor sequence as described below.

FixGrower Stage 5– Center to Side, Blocks with Fixation Growth

Identical to Legacy Stage 5, FixGrower Stage 5 requires animals to initiate a trial by poking into the center-port and maintaining fixation for a specified duration, after which a water reward is delivered at one of the side-ports, indicated by an illuminated LED. The rewarded side alternates in blocks of 15 trials. The structure for these trials is the following (assume in L block): Center-port LED indicates trial availability $\rightarrow$ animal pokes & holds nose in center $\rightarrow$ go-cue sound plays & Center LED off $\rightarrow$ L LED on & wait-for-poke period (max 8 seconds) $\rightarrow$ animal pokes in L port $\rightarrow$ L side-port water reward delivered (30 μL, mouse: 3 μL) $\rightarrow$ trial ITI sampled from a normal distribution with mean of 5 seconds and standard deviation 2 of seconds. If the animal leaves the center-port before the go-cue plays, or the animal concurrently pokes in a side-port during fixation, the trial is considered a violation.

This stage marks the beginning of divergence between the two curricula. In FixGrower Stage 5, if a fixation violation occurs, the center LED remains illuminated indefinitely, indicating that the trial is still active and allowing the animal to immediately retry. The center LED is only extinguished when the animal successfully maintains fixation until the go-cue is played. If the animal holds fixation until the go-cue (i.e., a valid trial) but then pokes into the unlit side-port, the trial is classified as an error. If the animal fails to initiate a trial during the wait-for-poke window, it is labeled a no answer trial. In both error and no answer trials, the side LED turns off immediately, and the trial transitions into the ITI period.

In the first session of this stage, the required fixation duration is set to 0.4 seconds. For all subsequent sessions, the fixation requirement is updated based on the animal's performance in a given session. If the animal had a bad session—defined as fewer than 20 valid trials—and has not yet begun growing (i.e., the required fixation duration is less than or equal to 0.4 seconds), the fixation requirement for the following session is reduced by 50 milliseconds to lower the entry barrier for initial trials. If the animal has a good session—defined as more than 50 valid trials—and the current fixation requirement is still below the 2-second target, the next day's requirement increases to match the average center-poke duration on valid trials from the current session, with an upper limit of 2.5 seconds. However, if the animal has a bad session after it has already begun growing (i.e., the required fixation duration exceeds 0.4 seconds), a reset rule is applied: the next day's fixation requirement is reduced to the midpoint between today's requirement and the average duration of failed fixation attempts. If the animal's performance does not qualify as either a "good" or "bad" session—defined as completing between 21 and 50 valid trials—the fixation duration remains unchanged. To complete this stage, animals must complete at least 200 trials in a single session with an accuracy greater than 85% and a no answer rate below 10%.

FixGrower Stage 6– Center to Side, Blocks to Random with Fixation Growth

FixGrower Stage 6 is identical to FixGrower Stage 5 except that the blocked L/R trials end at trial 70 and become randomly sampled with anti-biasing. To complete this stage, animals must perform at least 250 trials in a session with greater than 85% accuracy and less than 10% no answer rate.

FixGrower Stage 7– Center to Side, Random with Fixation Growth

FixGrower Stage 7 is identical to FixGrower Stages 5 and 6, with the exception that all trials are randomly sampled using anti-bias logic. To complete this stage, animals must perform at least 200 trials in a single session with a fixation duration exceeding the target, while maintaining greater than 85% accuracy and a no answer rate below 10%. For the primary

fixation curricula experiment, the target was 2-seconds. For the rat sound discrimination experiment, the target was 1.45 seconds. For the mouse fixation experiment, the target was 1.1 seconds. Unlike Legacy animals, FixGrower animals proceed to one additional stage before entering the probe phase.

FixGrower Stage 8– Violation Penalty On

FixGrower Stage 8 maintains the same trial structure as Stage 7, with the addition of a post-violation penalty timeout and a defined inter-trial interval. During the fixation period, a 150-millisecond free-retry window is implemented: if the animal enters and exits the center-port within this window, the trial immediately restarts without being counted as a violation. After this window, if the animal breaks fixation—either by exiting the center-port before the go-cue plays or by poking a side-port during fixation—the trial is classified as a violation. In such cases, the center-port LED remains on only for the originally intended fixation duration, the go-cue is not played, and a 2-second penalty timeout occurs followed by the ITI sampled from a normal distribution with mean of 5 seconds and standard deviation of 2 seconds.

Because animals have already reached the target fixation duration by this stage, the requirement is no longer increased. Instead, the required fixation duration is sampled independently for the session from an exponential distribution with a minimum of 1 second, a maximum of 2 seconds, and a time constant (τ) of 1.2 seconds. In the fixation curriculum experiment, after a single session in Stage 8, animals proceed to the fixation evaluation probe stages (Stage 9), regardless of performance.

Shared: Fixation Evaluation Probe Stages

There are two probe stages that all fixation experiment animals pass through once growth stages are complete to assess fixation performance, with each stage being 5 sessions in length.

Fixation Evaluation Stage 9– Probe Stable Fixation Delay

In Stage 9, trials are randomly sampled L/R with anti-biasing and have a stable fixation

duration of 2 seconds. The trial structure is as follows (assuming L trial): Center-port LED indicates trial availability → animal pokes & holds nose in center (2 seconds) → go-cue sound plays & Center LED off → L LED on & wait-for-poke period (max 8 s) → animal pokes in L port → L side-port water reward delivered (30 μL, mouse: 3 μL) → trial ITI sampled from a normal distribution with a mean of 4 seconds and standard deviation 1 second.

During the fixation period, a 150 milliseconds free-retry window is implemented: if the animal enters and exits the center-port within this window, the trial immediately restarts without being counted as a violation. After this window, if the animal breaks fixation—either by exiting the center-port before the go-cue plays or by poking a side-port during fixation—the trial is classified as a violation. In such cases, the center-port LED remains on only for the originally intended fixation duration, the go-cue is omitted, and a 2 seconds penalty timeout follows. This approach allows us to isolate violation trials that reflect active disengagement from the port, rather than brief exploratory movements as the animal settles into position.

If the animal stays in the center-port until the go-cue (i.e., a valid trial) but pokes in the unlit side-port, it is an error trial. If the animal does not poke during the wait-for-poke period, it is a no answer trial. For both of these trial types, the side LED turns off immediately and the trial moves into the ITI period. After 5 days in this stage, the animal moves into Stage 10.

Fixation Evaluation Stage 10– Probe Random Fixation Delay

Stage 10 is identical to Stage 9 except that the fixation duration is randomly sampled each trial from an exponential distribution with a minimum of 1 second, a maximum of 2 seconds, and a time constant (τ) of 1.2 seconds. After 5 days in this stage, the animal has completed the fixation experiment.

FixGrower Animals Post-Probe

A subset ($N = 8$) of FixGrower animals continued training after the end of the experiment in Stage 11 (**Fig. 3.4**). This stage was identical to Probe Stage 10 except that the fixation duration was set at the beginning of each session by sampling from an exponential distribution with a minimum of 1 second, a maximum of 2 seconds, and a time constant (τ) of 1.2 seconds (**Fig. 3.4a**). Unless a behavior box malfunctioned and animals needed to be moved to a different one, experimenters made no interventions for these animals for over 100 days.

Sound Discrimination Task

We trained a cohort of rats ($N = 13$) to perform a sound discrimination task following fixation growth using the FixGrower curriculum (**Fig. 3.5**). The only differences in the curriculum occurred in stage 8 and are described below. Moreover, we provide detailed methods of how we added sounds into the task.

Sound Discrimination Stage 8– Penalty On

After reaching a target fixation duration of 1.45 seconds through the FixGrower fixation curriculum Stages 5–7, animals entered a modified version of Stage 8. Unlike in the original curriculum, this stage was not limited to a single session. Instead, animals remained in this stage until they met the following performance criteria in a single session: at least 250 consecutive trials with a violation rate below 20% and a hit rate above 85%. Once this criterion was met, they advanced to the sound discrimination phase.

Although the target fixation duration during training was 1.45 seconds, rats commonly maintained fixation beyond the go-cue, by 200–250 ms. This natural tendency to overpoke meant that animals were already capable of sustaining fixations close to the full task requirement of 1.65 seconds. To streamline training and eliminate the need for an additional session solely to reach this duration, we introduced the auditory stimuli concurrently with an increase in the fixation requirement to 1.7 seconds in the next stage. This strategy allowed us

to reduce training time without compromising fixation stability during stimulus presentation.

Sound Discrimination Stage 9– Add Sounds

The goal of this stage was to introduce the two auditory tone stimuli (s_a and s_b) during the fixation period while maintaining low violation rates. Animals continued to receive rewards via light-guidance following the go-cue, which corresponded to the correct choice based on the sound rule. Stimuli were introduced progressively—first by increasing duration, then by increasing volume.

Initially, tones were presented at 35% of maximum volume (20 dB), beginning at a duration of 30 ms and increasing within-session toward a target of 400 ms. Task timing was as follows: 250 ms pre-s_a delay $\rightarrow$ 30 ms s_a $\rightarrow$ 940 ms pre-s_b delay $\rightarrow$ 30 ms s_b $\rightarrow$ 450 ms post-s_b delay, for a total fixation duration of 1.7 seconds. The inter-trial interval was sampled from a normal distribution with a mean of 1 second and a standard deviation of 0.5 seconds. Error choices triggered an additional 2-second timeout. Stimulus pairs were randomly drawn from two choice options and all tones were amplitude-matched to play at equal perceived volume (**Fig. 3.5b**):

1. **Left choice**: $s_a = 3$ kHz, $s_b = 3$ kHz
2. **Right choice**: $s_a = 3$ kHz, $s_b = 12$ kHz

On each trial, sound duration increased by 3 ms, while the delay between s_a and s_b decreased by 6 ms to maintain constant fixation time of 1.7 seconds. If the violation rate for a session exceeded 50%, sound growth was paused until the rate dropped below this threshold. If the violation rate remained below 40% for an entire session, the starting sound duration for the next session increased by 100 ms. This process repeated until the starting sound duration reached 400 ms, typically within four sessions. Final task timing was: 250 ms pre-s_a delay $\rightarrow$ 400 ms s_a $\rightarrow$ 200 ms pre-s_b delay $\rightarrow$ 400 ms s_b $\rightarrow$ 450 ms post-s_b delay (total fixation: 1.7 seconds).

After sounds played for their full duration, we gradually increased sound volume from

35% (20 dB) to 100% (80 dB). For the first 20 trials of each session, volume matched the previous session's value. After 20 trials, volume increased by 13% (6.15 dB) and was held constant for the remainder of the session. Animals advanced past this stage once they achieved 100% volume with a hit rate above 85% and a violation rate below 30%. On average, animals completed this stage in 8.9 ± 1.75 days (min: 7, max: 13 days).

As a precaution, we introduced the sounds conservatively to avoid startling the animals. However, our data suggest this approach may have been overly cautious: animals tolerated the stimuli well, even during early exposures. Based on this observation, we recommend combining sound duration and volume increases in future training protocols, as animals appear capable of adapting to both manipulations concurrently.

Sound Discrimination Stage 10– Wean Off Light Guide

The next stage of training involved phasing out the light-guidance cue so that animals relied solely on the sound rule to make correct choices. Task timing was the same as at the end of Stage 9, with two modifications: the pre-s_b delay was reduced to 150 ms (total fixation: 1.65 seconds), and the penalty for error responses increased to a 4 second timeout in addition to the ITI.

In the first session of the stage we reduced the proportion of light-guided trials from 100% to 60%. At the end of each session, we computed accuracy on non-guided trials. If performance exceeded 70%, the light-guided trial fraction was reduced by 15% in the subsequent session. This process continued until, within a single session, the proportion of light-guided trials was less than 25%, accuracy on non-guided trials was at least 75%, and the violation rate remained below 30%. Once all criteria were met, animals advanced to the final stage. On average, this stage took 12.3 ± 4.1 days (min: 6, max: 21 days) to complete.

Sound Discrimination Stage 11– Sound Rule

In the final stage, animals were required to perform the task using only the auditory rule, with no light-guidance. This is the stage we analyze in Figure 3.5. All task parameters

remained consistent with Stage 10 with the light guides turned off. Specifically, timing was: 250 ms pre-s_a delay $\rightarrow$ 400 ms s_a $\rightarrow$ 150 ms pre-s_b delay $\rightarrow$ 400 ms s_b $\rightarrow$ 450 ms post-s_b delay (total fixation: 1.65 seconds; **Fig. 3.5a**). Stimuli were randomly sampled from the set described earlier, requiring discrimination based on the frequency of s_b (**Fig. 3.5b**). The ITI remained normally distributed (mean = 1 second, SD = 0.5 second). Violations triggered a 2 second timeout in addition to the ITI, and error responses resulted in a 4 second timeout.

Mouse Fixation Task

We trained a cohort of mice ($N = 5$) to fixate using the FixGrower curriculum (**Fig. 3.6**). The only curriculum differences occurred in stage 8 and are described below.

Mouse Fixation Stage 8– Penalty On & Probe Violation Rates

After reaching a target fixation duration of 1.1 seconds through the FixGrower fixation curriculum Stages 5–7, mice entered a modified version of Stage 8. Unlike the original curriculum, the timeout duration following a violation was shortened to 100 ms and the ITI was sampled from a normal distribution with a mean of 2 seconds and a standard deviation of 0.5 seconds. Additionally, this stage was not limited to a single session and served as a benchmark for assessing violation rates (**Fig. 3.6c-d**). Mice remained in this stage until they met the following performance criteria in a single session: at least 150 trials with a violation rate below 20% and a hit rate above 85%. Once this threshold was met, animals completed the experiment.

3.5.4 Excluded Animals

Two animals from the fixation experiment were excluded from the final analyses, and their training progress is shown in (**Fig. S3.3**). Animal R044 (Legacy cohort) was excluded from all analyses related to the fixation growth and probe stages. This animal had substantial difficulty acquiring the center-poke behavior, requiring nine days and two behavior box changes before successfully executing a center-to-side motor sequence with a 10 ms hold. On

the third day of successful trials, the behavior box malfunctioned, and the animal's behavior failed to recover. Therefore, this animal is included only in analyses related to the side-poke training curriculum.

Animal R047 (FixGrower cohort) was excluded from analyses relating probe stage performance to training duration. This was the only animal in the dataset to trigger resetting logic in the fixation growth. R047 repeatedly engaged this mechanism during fixation growth, resulting in an unusually long training time (Days to Target = 22) despite exhibiting low probe-stage violation rates ($9.0 \pm 3\%$). As a result, this animal appeared as an outlier that broke the otherwise robust trend between training duration and performance observed in the FixGrower cohort. When included in the mixed-effects model, no significant interaction was detected between cohort and training time ($p = 0.221$). When excluded, a significant interaction emerged. This indicates that for animals that trigger the curriculum resetting logic, their training times may not be predictive of their final violation rate and alternative lines of evidence should be used to determine training potential.

3.5.5 Simulated Fixation Growth

To assess how quickly animals in the Legacy cohort could have reached the 2 second fixation target under different behavioral assumptions, we simulated fixation growth using the Legacy curriculum algorithm under two conditions: (1) a **normative** condition, in which no violations occurred, representing an upper bound on learning speed given each animal's typical session length, and (2) a **recovery** condition, which incorporated each animal's empirical violation patterns and session-level trial counts to validate the simulation against observed behavior.

Parameter Estimation

For each Legacy-trained animal, we estimated empirical session-level statistics using data from the fixation growth stages (Stages 5–7). These included the mean and variance of the

number of trials per session, as well as warm-up and post–warm-up violation rates. Warm-up trials were defined as the first 20 trials of a session if the previous session's final fixation duration exceeded 10 ms.

Violation rates were modeled using Beta distributions to capture their characteristic right-skew. Parameters for each animal's warm-up and post–warm-up distributions were estimated via the method of moments. Specifically, for a given empirical mean μ and variance σ^2, Beta distribution shape parameters α and β were derived using:

$$\alpha \;=\; \mu\Big(\frac{\mu\,(1-\mu)}{\sigma^2} \;-\; 1\Big), \quad \beta \;=\; (1-\mu)\Big(\frac{\mu\,(1-\mu)}{\sigma^2} \;-\; 1\Big).$$

Session trial counts were sampled from a truncated normal distribution, with mean and variance set to empirical values, rounded to the nearest integer, and constrained to a minimum of one trial.

Simulation Procedure

For each simulated session, we first sampled the number of trials based on the estimated normal distribution of session lengths for that animal. In the normative condition, both warm-up and post–warm-up violation rates were set to zero. In the recovery condition, violation rates were drawn from the animal's corresponding Beta distributions for that session.

Each trial was simulated sequentially following the Legacy fixation growth logic. For each trial, we drew a Bernoulli sample using the appropriate violation probability (warm-up or post–warm-up) for the selected condition to determine whether the trial was a violation. For the first 20 trials of a session (warm-up phase), fixation duration increased linearly after each valid trial, with the step size calculated as a fixed increment designed to evenly span the gap between 10 ms and the ending fixation duration of the previous session. In the post–warm-up phase, the fixation duration increased by max(0.1%, 1 ms) following each valid trial. Violation trials produced no change in the fixation duration.

The fixation duration at the end of each session was used as the warm-up target for the next session. Simulations continued until either the fixation duration exceeded 2 seconds or the simulation reached a maximum of 90 sessions. For each animal and condition, the simulation was repeated 30 times using a fixed random seed. Results were averaged across replicates to estimate the expected number of sessions required to reach the 2 second target in each condition.

3.5.6 Trial Initiation Time

The trial initiation time is defined as the time from the final poke out from the center-port on the previous trial to the time of the first poke into the center-port on the current trial. To control for experimenter-imposed delays, we subtract off any incurred inter-trial-interval or penalty timeout that occurred on the previous trial.

3.5.7 Statistical Analyses

Statistical analyses were performed using the `SciPy` and `Statsmodels` libraries in Python (3.10). Detailed statistical information is provided in the corresponding figure legends and Results sections. Unless otherwise specified, data are reported in-text as mean ± standard deviation (SD). Box plots show the first and third quartiles with the median indicated by a horizontal bar. Error bars and shaded regions in plots represent the standard error of the mean (SEM).

All statistical comparisons were conducted using non-parametric tests (e.g., *Mann–Whitney U* or *Wilcoxon signed-rank*) unless data were found to be normally distributed based on the *Shapiro–Wilk* test. When animals contributed multiple data points (e.g., session-level violation rates), mixed-effects linear models were used with animal ID included as a random intercept to account for within-subject variability. A significance threshold of $\alpha = 0.05$ was used for all tests. Where applicable, p-values were adjusted for multiple comparisons using the *Holm-Bonferroni* method.

3.5.8 Data & Code Availability

All data and code for analyses in this chapter can be found at https://tinyurl.com/jbredachapter3code.

4

Conclusions

4.1 Summary

Violations are instances where animals break fixation prematurely during perceptual decision-making tasks. These events are a common, yet often overlooked source of behavioral variability as they are typically discarded during analysis. However, this practice reduces the volume of usable data, complicates interpretation of internal decision processes, and overlooks potentially informative signals about these lapses that may shape subsequent neural activity. This thesis presented two detailed studies aimed at characterizing these violations behaviorally and computationally, and explored strategies to mitigate their occurrence.

In **Chapter 2**, I demonstrated that violations are structured, history-dependent events rather than random behavioral noise. Moreover, these events negatively interact with stimulus sensitivity, and diminish subsequent task performance. Integrating violations into a multinomial behavioral model improved predictive accuracy beyond a traditional binary choice model, marking a novel and detailed characterization of these behavioral phenomena.

In **Chapter 3**, I aimed to reduce violations through improved behavioral training protocols. Existing curricula lack detailed documentation and quantification, thereby providing only a limited baseline for efficiency and effectiveness. I introduced and validated a novel training curriculum ("FixGrower") against a widely used "Legacy" protocol. I achieved a substantial 61% reduction in training time while keeping violation rates below 15% on average. Crucially, these data demonstrated a clear relationship between reduced violation rates and decreased training duration, providing a reliable and data-driven approach to behav-

ioral training. This curriculum showed robustness across different experimental conditions, task modifications, and species. Grounded in theories of operant conditioning, reinforcement learning, and curriculum learning, these findings bolster the theoretical underpinnings of effective behavioral shaping protocols.

4.2 Future Directions

The studies presented here open exciting avenues for future research. Below, I will focus on three areas of interest

4.2.1 Violations as a Disengaged State: Toward Latent & Dynamic Models

In Chapter 2, preliminary modeling efforts initially aimed to characterize left/right choices using a GLM-HMM. However, the model fit was consistently poor and the states were non-interpretable. We hypothesized this may be due to the dropping of violation trials from the time series, which led to difficulties in fitting evolving states over trials. This led us to explore violation trials in more detail. Our findings suggest that violations reflect structured lapses rather than isolated random events. Therefore, a logical extension involves implementing a trinomial GLM-HMM approach, akin to that of Hulsey et al. (2024), to explicitly capture these events as a disengaged state. Further, employing dynamic modeling tools such as PsyTrack (Roy et al., 2021) could clarify how the impact of violations evolves throughout training. Such modeling would inform improved curricula and further reduce violations, enhancing behavioral stability.

4.2.2 A Standard Protocol for Shaping Fixation Behavior

While the FixGrower curriculum showed robustness across a variety of conditions within our laboratory, further validation across diverse settings is necessary. Recent initiatives like those of the The International Brain Laboratory et al. (2021) have demonstrated that standardized training pipelines bolster reproducibility and efficiency. Given the prevalence

of fixation requirements in freely moving perceptual decision-making tasks, adopting and refining a standardized curriculum would significantly advance behavioral training in the field. We believe FixGrower is a step toward this goal, offering a flexible framework adaptable to a range of experimental demands. In particular, it may serve as a valuable foundation not only for freely moving tasks, but for any paradigms that require animals to withhold movement—such as suppressing licks during stimulus presentation in head-fixed tasks. As such, FixGrower represents a promising starting point for building more unified, scalable approaches to training across behavioral neuroscience.

4.2.3 Elucidating the Neural Basis of Violations

Finally, understanding the causal factors and neural mechanisms underlying violations remains an open question. Given the maladaptive nature of these errors, what purpose do they serve, and what neural dynamics produce them? Motor control research, particularly focusing on basal ganglia pathways, provides key insights.

The basal ganglia are a collection of interconnected subcortical nuclei that regulate the initiation and suppression of voluntary actions, primarily via two major pathways: the direct pathway, which facilitates selected motor actions, and the indirect pathway, which suppresses competing or unwanted movements (Rocha et al., 2023). These pathways work in concert to enable precise motor control.

Across a spectrum of movement disorders, basal ganglia dysfunction illustrates how imbalances in direct and indirect pathway signaling can shift motor output. For example, in Parkinson's disease, degeneration of dopaminergic input leads to decreased excitatory drive in the direct pathway and increased inhibitory tone in the indirect pathway (DeLong & Wichmann, 2007; McGregor & Nelson, 2019; Redgrave et al., 2010). This then results in *hypo*kinetic symptoms such as bradykinesia and rigidity. Conversely, *hyper*kinetic symptoms observed in Huntington's disease patients are caused by the degradation of inhibitory neurons

in the indirect pathway (Albin et al., 1989; Plotkin & Surmeier, 2015; Reiner et al., 1988). This leads to decreased inhibitory tone in the indirect pathway which produces symptoms like chorea and tremors. Together, these clinical extremes underscore a core principle: increased inhibitory tone in the indirect pathway or reduced excitatory tone in the direct pathway suppresses movement, while the reverse facilitates it.

Rodent studies now provide causal evidence linking these circuit-level principles directly to violations. Cruz et al. (2022) demonstrated that optogenetic inhibition of indirect pathway inhibitory neurons in the dorsolateral striatum (DLS) led to a significant increase in violations in a time-interval categorization task. Complementing this, Jang et al. (2024) highlighted the role of cholinergic tone in shaping excitatory dopamine signaling dynamics in the dorsomedial striatum (DMS). They found that elevated acetylcholine levels altered the functional impact of dopamine—from reinforcing learning signals to initiating immediate movement. By optogenetically activating acetylcholine neurons during fixation in a temporal-wagering task, the authors induced a significant increase in violations and decreased overall task performance.

Together, these studies suggest that violations result from transient imbalances within basal ganglia circuits—whether through reduced inhibitory output, altered neuromodulatory dynamics or a combination of other mechanisms. These findings begin to bridge the gap between clinical and experimental models of motor disinhibition, illustrating how the same circuits that underlie movement disorders also shape trial-by-trial control in healthy systems. Understanding these mechanisms offers not only insight into the neural basis of behavioral errors but also a framework for exploring impulsivity and motor dysregulation.

Bibliography

Akrami, A., Kopec, C. D., Diamond, M. E., & Brody, C. D. (2018). Posterior parietal cortex represents sensory history and mediates its effects on behaviour. *Nature*, *554*(7692), 368–372. https://doi.org/10.1038/nature25510

Albin, R. L., Young, A. B., & Penney, J. B. (1989). The functional anatomy of basal ganglia disorders. *Trends in Neurosciences*, *12*(10), 366–375. https://doi.org/10.1016/0166-2236(89)90074-x

Ashwood, Z. C., Roy, N. A., Stone, I. R., Urai, A. E., Churchland, A. K., Pouget, A., & Pillow, J. W. (2022). Mice alternate between discrete strategies during perceptual decision-making. *Nature Neuroscience*, *25*(2), 201–212. https://doi.org/10.1038/s41593-021-01007-z

Bengio, Y., Louradour, J., Collobert, R., & Weston, J. (2009). Curriculum learning. *Proceedings of the 26th Annual International Conference on Machine Learning*, 41–48. https://doi.org/10.1145/1553374.1553380

Bolkan, S. S., Stone, I. R., Pinto, L., Ashwood, Z. C., Garcia, J. M. I., Herman, A. L., Singh, P., Bandi, A., Cox, J., Zimmerman, C. A., Cho, J. R., Engelhard, B., Pillow, J. W., & Witten, I. B. (2022). Opponent control of behavior by dorsomedial striatal pathways depends on task demands and internal state. *Nature neuroscience*, *25*(3), 345–357. https://doi.org/10.1038/s41593-022-01021-9

Brunton, B. W., Botvinick, M. M., & Brody, C. D. (2013). Rats and Humans Can Optimally Accumulate Evidence for Decision-Making. *Science*, *340*(6128), 95–98. https://doi.org/10.1126/science.1233912

Burkholder, T., Foltz, C., Karlsson, E., Linton, C. G., & Smith, J. M. (2012). Health Evaluation of Experimental Laboratory Mice. *Current protocols in mouse biology*, *2*, 145–165. https://doi.org/10.1002/9780470942390.mo110217

Constantinople, C. M., Piet, A. T., & Brody, C. D. (2019). An Analysis of Decision under Risk in Rats. *Current Biology*, *29*(12), 2066–2074.e5. https://doi.org/10.1016/j.cub.2019.05.013

Cruz, B. F., Guiomar, G., Soares, S., Motiwala, A., Machens, C. K., & Paton, J. J. (2022). Action suppression reveals opponent parallel control via striatal circuits. *Nature*, *607*(7919), 521–526. https://doi.org/10.1038/s41586-022-04894-9

Cuturela, L. I., Laboratory, T. I. B., & Pillow, J. W. (2024, December). Internal states emerge early during learning of a perceptual decision-making task. https://doi.org/10.1101/2024.11.30.626182

DeLong, M. R., & Wichmann, T. (2007). Circuits and circuit disorders of the basal ganglia. *Archives of Neurology*, *64*(1), 20–24. https://doi.org/10.1001/archneur.64.1.20

Gouveia, K., & Hurst, J. L. (2013). Reducing mouse anxiety during handling: Effect of experience with handling tunnels. *PloS One*, *8*(6), e66401. https://doi.org/10.1371/journal.pone.0066401

Hulsey, D., Zumwalt, K., Mazzucato, L., McCormick, D. A., & Jaramillo, S. (2024). Decision-making dynamics are predicted by arousal and uninstructed movements. *Cell Reports*, *43*(2), 113709. https://doi.org/10.1016/j.celrep.2024.113709

Jang, H. J., Ward, R. M., Golden, C. E. M., & Constantinople, C. M. (2024, May). Acetylcholine demixes heterogeneous dopamine signals for learning and moving. https://doi.org/10.1101/2024.05.03.592444

Jaramillo, S., & Zador, A. M. (2014). Mice and rats achieve similar levels of performance in an adaptive decision-making task. *Frontiers in Systems Neuroscience*, *8*. https://doi.org/10.3389/fnsys.2014.00173

Licata, A. M., Kaufman, M. T., Raposo, D., Ryan, M. B., Sheppard, J. P., & Churchland, A. K. (2017). Posterior Parietal Cortex Guides Visual Decisions in Rats. *Journal of Neuroscience*, *37*(19), 4954–4966. https://doi.org/10.1523/JNEUROSCI.0105-17.2017

Mah, A., Schiereck, S. S., Bossio, V., & Constantinople, C. M. (2023). Distinct value computations support rapid sequential decisions. *Nature Communications*, *14*(1), 7573. https://doi.org/10.1038/s41467-023-43250-x

McGregor, M. M., & Nelson, A. B. (2019). Circuit Mechanisms of Parkinson's Disease. *Neuron*, *101*(6), 1042–1056. https://doi.org/10.1016/j.neuron.2019.03.004

Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., & Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, *518*(7540), 529–533. https://doi.org/10.1038/nature14236

Odoemene, O., Pisupati, S., Nguyen, H., & Churchland, A. K. (2018). Visual Evidence Accumulation Guides Decision-Making in Unrestrained Mice. *The Journal of Neuroscience*, *38*(47), 10143–10155. https://doi.org/10.1523/JNEUROSCI.3478-17.2018

Oesch, L. T., Ryan, M. B., & Churchland, A. K. (2024). From innate to instructed: A new look at perceptual decision-making. *Current Opinion in Neurobiology*, *86*, 102871. https://doi.org/10.1016/j.conb.2024.102871

Pisupati, S., Chartarifsky-Lynn, L., Khanal, A., & Churchland, A. K. (2021). Lapses in perceptual decisions reflect exploration. *eLife*, *10*, e55490. https://doi.org/10.7554/eLife.55490

Plotkin, J. L., & Surmeier, D. J. (2015). Corticostriatal synaptic adaptations in Huntington's disease. *Current Opinion in Neurobiology*, *33*, 53–62. https://doi.org/10.1016/j.conb.2015.01.020

Redgrave, P., Rodriguez, M., Smith, Y., Rodriguez-Oroz, M. C., Lehericy, S., Bergman, H., Agid, Y., DeLong, M. R., & Obeso, J. A. (2010). Goal-directed and habitual control in the basal ganglia: Implications for Parkinson's disease. *Nature Reviews. Neuroscience*, *11*(11), 760–772. https://doi.org/10.1038/nrn2915

Reiner, A., Albin, R. L., Anderson, K. D., D'Amato, C. J., Penney, J. B., & Young, A. B. (1988). Differential loss of striatal projection neurons in Huntington disease. *Proceed-*

ings of the National Academy of Sciences of the United States of America, *85*(15), 5733–5737. https://doi.org/10.1073/pnas.85.15.5733

Reuschenbach, J., Reinert, J. K., Fu, X., & Fukunaga, I. (2023). Effects of Stimulus Timing on the Acquisition of an Olfactory Working Memory Task in Head-Fixed Mice. *Journal of Neuroscience*, *43*(17), 3120–3130. https://doi.org/10.1523/JNEUROSCI.1636-22.2023

Rocha, G. S., Freire, M. A. M., Britto, A. M., Paiva, K. M., Oliveira, R. F., Fonseca, I. A. T., Araújo, D. P., Oliveira, L. C., Guzen, F. P., Morais, P. L. A. G., & Cavalcanti, J. R. L. P. (2023). Basal ganglia for beginners: The basic concepts you need to know and their role in movement control [Publisher: Frontiers]. *Frontiers in Systems Neuroscience*, *17*. https://doi.org/10.3389/fnsys.2023.1242929

Roy, N. A., Bak, J. H., Akrami, A., Brody, C. D., & Pillow, J. W. (2021). Extracting the dynamics of behavior in sensory decision-making experiments. *Neuron*, *109*(4), 597–610.e6. https://doi.org/10.1016/j.neuron.2020.12.004

Schmid, C., Haziq, M., Baese-Berk, M. M., Murray, J. M., & Jaramillo, S. (2024). Passive exposure to task-relevant stimuli enhances categorization learning (S. Ostojic & M. J. Frank, Eds.). *eLife*, *12*, RP88406. https://doi.org/10.7554/eLife.88406

Scott, B. B., Constantinople, C. M., Erlich, J. C., Tank, D. W., & Brody, C. D. (2015). Sources of noise during accumulation of evidence in unrestrained and voluntarily head-restrained rats (N. Uchida, Ed.). *eLife*, *4*, e11308. https://doi.org/10.7554/eLife.11308

Skinner, B. F. (1951). How to teach animals. *Scientific American*, *185*(6), 26–29.

Skinner, B. F. (1953). *Science and human behavior*. Macmillan.

Soma, S., Suematsu, N., & Shimegi, S. (2014). Efficient training protocol for rapid learning of the two-alternative forced-choice visual stimulus detection task. *Physiological Reports*, *2*(7), e12060. https://doi.org/10.14814/phy2.12060

Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction, 2nd ed.* The MIT Press.

The International Brain Laboratory, Aguillon-Rodriguez, V., Angelaki, D., Bayer, H., Bonacchi, N., Carandini, M., Cazettes, F., Chapuis, G., Churchland, A. K., Dan, Y., Dewitt, E., Faulkner, M., Forrest, H., Haetzel, L., Häusser, M., Hofer, S. B., Hu, F., Khanal, A., Krasniak, C., … Zador, A. M. (2021). Standardized and reproducible measurement of decision-making in mice (N. Uchida & M. J. Frank, Eds.) [Publisher: eLife Sciences Publications, Ltd]. *eLife*, *10*, e63711. https://doi.org/10.7554/eLife.63711

Urai, A. E., Aguillon-Rodriguez, V., Laranjeira, I. C., Cazettes, F., Laboratory, T. I. B., Mainen, Z. F., & Churchland, A. K. (2021). Citric Acid Water as an Alternative to Water Restriction for High-Yield Mouse Behavior. *eNeuro*, *8*(1). https://doi.org/10.1523/ENEURO.0230-20.2020

Venditto, S. J. C., Miller, K. J., Brody, C. D., & Daw, N. D. (2024, October). Dynamic reinforcement learning reveals time-dependent shifts in strategy during reward learning. https://doi.org/10.1101/2024.02.28.582617