

DIGITAL ACCESS TO SCHOLARSHIP AT HARVARD
DASH.HARVARD.EDU

HARVARD LIBRARY


The Structure of Mouse Behavior

Citation

Wiltschko, Alexander Bame. 2016. The Structure of Mouse Behavior. Doctoral dissertation, Harvard University, Graduate School of Arts & Sciences.

Link

http://nrs.harvard.edu/urn-3:HUL.InstRepos:33493569

Terms of use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material (LAA), as set forth at

https://harvardwiki.atlassian.net/wiki/external/NGY5NDE4ZjgzNTc5NDQzMGIzZWZhMGFlOWI2M2EwYTg

Accessibility

https://accessibility.huit.harvard.edu/digital-accessibility-policy

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. Submit a story

The Structure of Mouse Behavior

A dissertation presented

by

Alexander Bame Wiltschko

to

The Division of Medical Sciences

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Neurobiology

Harvard University

Cambridge, Massachusetts

April 2016

©2016 Alexander Bame Wiltschko

All rights reserved.

Dissertation Advisor: Dr. Sandeep R. Datta Alexander Bame Wiltschko

The Structure of Mouse Behavior

Abstract

Complex animal behaviors are likely built from simpler modules, but their systematic identification in mammals remains a significant challenge. Here we use depth imaging to show that three-dimensional (3D) mouse pose dynamics are structured at the sub-second timescale by using a newly developed 3D imaging and machine learning-based automated phenotyping system, which we term Motion Sequencing (MoSeq). Computational modeling of these fast postural dynamics effectively describes mouse behavior as a series of reused and stereotyped modules with defined transition probabilities, which collectively encapsulate the underlying structure of mouse behavior within a given experiment.

By deploying MoSeq in a variety of experimental contexts, we show that it unmasks strategies employed by the brain to generate specific adaptations to changes in the environment, and captures both predicted and previously-hidden phenotypes induced by genetic, neural, and pharmacological manipulations. We directly compare the predictive power of behavioral representations built by MoSeq against traditional measurements of behavior, including speed, length, and allocentric position, and demonstrate MoSeq is able to discriminate between subtle pharmacological manipulations of behavior, while traditional methods are not. This work demonstrates that mouse body language is built from identifiable components and is organized in a predictable fashion; deciphering this language establishes a framework for characterizing the influence of environmental cues, genes, neural activity and pharmacology on behavior.

Table of contents

Abstract iii

List of figures viii

Dedication ix

Acknowledgments xi

1 General introduction and background 1

2 Mapping Sub-Second Structure in Mouse Behavior 5

2.1 Introduction 5

2.2 Results 6

2.2.1 3D Imaging Captures Inherent Structure in Mouse Pose Dynamics . . 6

2.2.2 Mouse Behavior Can Be Described and Predicted with Modules and Transitions 12

2.2.3 Model-Identified Behavioral Modules Are Stereotyped and Distinct . . 19

2.2.4 Using the AR-HMM to Characterize Baseline Patterns of Behavior . . 24

2.2.5 Using the AR-HMM to Characterize the Nature of Behavioral Change 27

2.2.6 Sub-Second Architecture of Behavior Reflects the Influence of Genes and Neural Activity 32

2.3 Discussion 38

2.3.1 The AR-HMM Automatically and Systematically Captures Known and New Phenotypes 39

2.3.2 The AR-HMM Identifies Possible Mechanisms for Behavioral Change 40

2.3.3 Potential Neural Underpinnings for Modules and Transitions 42

2.3.4 Unsupervised Behavioral Characterization via Modeling of 3D Data: Strengths, Caveats and Future Directions . . . 43
2.3.5 Mouse Body Language: Syllables and Grammar . . . 46
2.4 Acknowledgments . . . 48
2.5 Author contributions . . . 49

3 Computer Vision and Machine Learning Methods for the Segmentation of Freely-Moving Mouse Behavior . . . 50
3.1 Methods Summary . . . 50
3.2 Methods . . . 51
3.2.1 Mouse Strains, Housing and Habituation . . . 51
3.3 Behavioral Assays . . . 52
3.3.0.1 Innate Exploration . . . 52
3.3.0.2 Stimulus-Driven Innate Behaviors . . . 52
3.3.0.3 Optogenetics . . . 53
3.3.1 Camera Setup and Initialization . . . 54
3.3.2 Data Acquisition . . . 55
3.3.3 Data Preprocessing . . . 55
3.3.4 Model-Free Data Analysis Methods . . . 61
3.3.5 Data Modeling and Fitting Methods . . . 63
3.3.6 Model Training and Parameter Tests . . . 68
3.3.7 Evaluating the Log-likelihood of Held-out Test Data . . . 69
3.3.8 Predictive comparison metric to stratify model capacities . . . 70
3.3.9 Analysis of the Model Consistency . . . 72
3.3.10 Analysis on Model Output . . . 73
3.3.11 Statemap Visualizations . . . 74
3.3.12 Hinton Diagram . . . 74
3.3.13 Computing Entropy Rate and Mutual Information . . . 75
3.3.14 Cross-likelihood Analysis of AR-HMM States . . . 75
3.3.15 Visualizing Synthetic Mice . . . 77
3.3.16 Statistical Methods . . . 78

3.3.17 Technical definitions of the generative models . . . 81
3.3.18 Prior on the transition matrix . . . 82
3.3.19 Generation of the discrete state sequence . . . 82
3.3.20 Prior on the autoregressive parameters . . . 82
3.3.21 Generation of the 3D pose sequence principal components . . . 83
3.3.22 Technical description of the inference algorithms . . . 84
3.3.22.1 Sampling $x \mid \pi, \theta, y$ . . . 85
3.3.22.2 Sampling $\theta \mid x, y$ . . . 87
3.3.22.3 Sampling $\beta, \pi \mid x$ . . . 88
3.4 Author contributions . . . 90

4 Motion Sequencing for Quantitative Behavioral Pharmacology . 91

4.1 Introduction . . . 91
4.2 Methods . . . 99
4.2.1 Drug preparation and administration . . . 99
4.2.2 Behavioral protocol . . . 100
4.2.3 Behavioral Fingerprint Extraction . . . 101
4.2.4 Pharmacological classification . . . 101
4.2.5 Held-out classification . . . 102
4.3 Results . . . 102
4.3.1 Psychoactive Drugs as a Testbed for Behavioral Fingerprints . . . 102
4.3.2 MoSeq, and Other Model-Based Segmentations of Action . . . 105
4.3.3 Comparing the Utility of A Portfolio of Behavioral Fingerprints . . . 107
4.3.3.1 Discriminating Drug Class . . . 110
4.3.3.2 Discriminating Drug Identity . . . 114
4.3.3.3 Discriminating Drug Identity and Dose . . . 115
4.3.4 Discovering Relationships Between Drugs with MoSeq . . . 117
4.3.5 Individual Syllables Can Discriminate Drugs . . . 119
4.4 Discussion . . . 121
4.4.1 MoSeq produces discriminative and interpretable representations of behavior . . . 123

4.4.2 Drugs with similar pharmacological profiles are distinct, but related, in the MoSeq representation . . . 125
4.4.3 Individual syllables distinguish broad classes of drugs . . . 126
4.4.4 Developing better behavioral representations are a keystone task for neuroscience and pharmacology . . . 127
4.5 Author contributions . . . 128

5 Discussion and Future Directions . . . 129
5.1 Contribution of Present Work . . . 129
5.1.1 Building a Computational Model for Ethology . . . 129
5.1.2 Summarizing Behavior With Behavioral Syllables . . . 131
5.1.3 Behavioral Syllables Are Regulated by Neurological Manipulations At Widely Varying Timescales . . . 132
5.1.4 Directly Targeting Behavior with Pharmaceuticals . . . 134
5.2 Future Directions . . . 136
5.2.1 Neural Substrates of Behavioral Syllables . . . 136
5.2.2 Opportunities in Learned Behavior . . . 138
5.2.3 Completing Our View of Musculoskeletal Behavior . . . 140
5.2.4 Parsing Social Behavior . . . 143
5.2.5 Application to Other Species . . . 145

References . . . 156

Listing of figures

2.1 Depth imaging reveals block structure in 3D mouse pose dynamics data. . . . 7
2.1 (Continued) 8
2.2 Image extraction workflow. 9
2.3 Block structure, autocorrelation and compressibility in mouse depth imaging data. 13
2.3 (Continued) 14
2.4 Additional template-matched trajectories in pose space. 15
2.5 Modeling mouse behavior. 17
2.5 (Continued) 18
2.6 Reused behavioral modules within mouse pose dynamics data. 20
2.7 Characterizing behavioral modules identified by the AR-HMM. 22
2.8 The physical environment influences module usage and spatial pattern of expression. 25
2.9 Subjective and objective analysis of AR-HMM identified behavioral modules. 26
2.10 Odor-driven innate avoidance alters transition probabilities. 30
2.11 The AR-HMM disambiguates wild-type, heterozygous and homozygous Ror1β mice. 34
2.12 Optogenetic perturbation of motor cortex yields both neomorphic and physiological modules. 36

3.1 Sensitivity analysis of model-free parameters. 57

4.1 Three-Dimensional Imaging Can Be Used to Build A Variety of Behavioral Fingerprints. 103
4.2 Drug Class, Identity And Dose Are Qualitatively Discernible In Mouse Behavioral Fingerprints. 108

4.3 MoSeq Behavioral Fingerprints Reliably Distinguish Pharmaceutical Treatment Classes. 111
4.4 MoSeq-based Behavioral Fingerprints Can Reliably Distinguish Drug Identity. 113
4.5 MoSeq can Correctly Identify both Drug And Dose 116
4.6 Using a MoSeq-based Linear Classifier to Characterize ”Unknown” Drug Effects 118
4.7 Decision Trees Built Using Syllable Usage Rates Can Classifying Individual Drugs 120

To my Sister, who softens my sharp edges and is wiser than she knows.

To my Mother, who has encouraged and nurtured me my entire life.

To my Father, my first teacher and guiding moral light, rest in peace. I miss you every day.

To my partner in life and best friend, Katy, I am so grateful to be with you, and incredibly excited for the adventure still ahead of us.

Acknowledgments

The time I have spent in graduate school has been graced by the presence, advice and conversation of so many wonderful peers and mentors.

First, my advisor Professor Bob Datta, whose love of science is truly infectious – I hope that I have absorbed some of his incredible vigor, rigor and passion in the pursuit of knowledge. Also, before arriving at graduate school, I did not appreciate the art and necessity of scientific storytelling. Bob is a consummate communicator, and I am honored to have seen him, up close, turn raw fact into shareable knowledge. These two qualities, rigor and clarity, are rarely found in such concentration in the same individual, and I am glad to know and have learned from him.

Through the course of my research, I worked very closely with Matthew Johnson, first when he was a student at MIT, and then after he joined the Datta Lab as a postdoc. Matt has helped me immensely to grow intellectually, and he has constantly pushed me to be a more complete and more thorough thinker. The greatest gift he has given me is that of a broadened intellectual horizon – there is so much to be learned in the world of mathematics, probability and machine learning that I am aware of now because of him. He is also perhaps the nicest person I have known, and I am glad to be his friend.

Among my cohort, and among the years before and after me, I'm glad to call Willie Tobin,

Ari Morcos, Stan Pashkovski, and Daniel Bear my friends. Their conversation and advice, about science and about life, has made the tougher spots bearable and the high points that much better.

I have been pleased to have worked with two extremely talented and motivated co-op students, Ralph Peterson and Jesse Katon, who have done an enormous amount of the experimental work in this dissertation. Both of them have worked extremely hard, acquired new skills (both learned computer programming from scratch), and maintained a positive and industrious attitude during it all. I am excited to see where their futures take them.

At the Datta lab, I was fortunate enough to be seated between Tari Tan and Masha Bloom, two extremely talented and extremely hard-working scientists with excellent senses of humor and admirable resolve. One bay over, was Giuliano Iurili, Stan Pashkovski and Andrew Giessel. I always valued our conversations, and our shared excitement for fusing computational thinking and neurobiology. One bay over, the other direction, is Paul Greer, whose advice I sought whenever I could about science, and about life as a scientist. He's a phenomenal person with excellent scientific taste.

In the middle of my PhD, I hit the end of my abilities to analyze the behavioral data I was collecting, and so Bob sent me to a computer vision conference, CVPR, to learn more. I wandered between the sessions, and ended up in conversation with a professor named Neil Lawrence, who recommended I go talk to his buddy Ryan Adams, a professor at Harvard. Across the river, in the engineering school, Ryan introduced me to the world of machine

learning, which has become an essential component of this dissertation, and also, I hope, my professional life to come. In Ryan's lab, I spent time with some incredibly smart folks, and I'm very glad to have known Scott Linderman, David Duvenaud and Dougal Maclaurin.

I would also like to thank my dissertation advisory committee, Bernardo Sabatini, John Assad and Chris Harvey, for their scientific guidance in the course of this project, as well as my examiners, Tim Gardner, Aravi Samuel and Mark Andermann, for their thoughts and criticisms of this dissertation.

Karen Harmin has kept me on track and on time throughout my stay in graduate school, and I would be more than a little bit lost without her.

I thank my Mother and Sister for their love, realism and encouragement. I thank my Father for the love, wisdom and constancy he gave while he was with us.

Last, I thank my partner Katy for her love, patience, and humor. I couldn't have paired up with someone better.

Chapter 1

General introduction and background

Innate behaviors are sculpted by evolution into stereotyped forms that enable animals to accomplish particular goals (such as exploring or avoiding a predator). Ultimately understanding how neural circuits create these patterned behaviors requires a clear framework for characterizing how behavior is organized and evolves over time. One conceptual approach to addressing this challenge arises from ethology, which proposes that the brain builds coherent behaviors by expressing stereotyped modules of simpler action in specific sequences (Tinbergen, 1951).

Although behavioral modules have traditionally been identified, one at a time, through careful human observation, recent technical advances have enabled more comprehensive characterization of the components of behavior. For example, in invertebrates, behavioral modules and their associated transition probabilities can now be discovered systematically through automated machine vision, clustering and classification algorithms (Berman et al., 2014; Croll, 1975; Garrity et al., 2010; Stephens et al., 2008; Vogelstein et al., 2014). Furthermore, identifying behavioral modules and transition probabilities has uncovered context-specific strategies used by invertebrate brains to adapt behavior to changes in the

environment, which include both the emission of new behavioral modules (such as when the animal switches from "exploring" to "mating"), and the generation of new behaviors through re-sequencing existing modules. In *C. elegans*, for example, neural circuits that respond to appetitive cues alter transition probabilities between a core set of locomotor-related behavioral modules, thereby creating new behavioral sequences that enable taxis towards attractive odorants (Gray et al., 2005; Pierce-Shimomura et al., 1999). Similar observations have been made for sensory-driven behaviors in fly larvae (Garrity et al., 2010).

Comparable systematic approaches to discovering behavioral modules have not yet been implemented in mice. Instead, traditional behavioral classification approaches have been instantiated *in silico*, enabling machine vision algorithms to replace tedious and unreliable human scoring of videotaped behavior (de Chaumont et al., 2012; Jhuang et al., 2010; Kabra et al., 2013). These approaches are powerful, and have been successfully used to quantify components of innate exploratory, grooming, approach, aggressive and reproductive behaviors. However, because they depend upon the prior specification, by human observers, of what constitutes a meaningful behavioral module, the insight from these methods is bounded by human perception and intuition. Currently-available approaches therefore focus on identifying a small number of pre-specified modules within a given experiment, rather than on discovering new behavioral modules (which potentially encapsulate novel patterns of action), describing the global structure of behavior, or predicting future actions based upon those in the past.

Quantitative unsupervised approaches for phenotyping animal behavior have been recently proposed for invertebrates. For example, the work of Berman et al. (2014) uses a technique that shares aspects with the one proposed in this thesis, particularly in the machine vision-based preprocessing steps of raw behavioral data. With their method, aligned images of an animal (a fly, instead of a mouse) are dimensionally-reduced using principal components analysis (PCA), and then augmented with local temporal information using a wavelet transform, and then are projected into a 2D space using the t-SNE clustering algorithm. They discover that clusters in their behavioral representation correspond to stereotyped motifs of action, but that most behavior is not stereotyped. They highlight individual behavioral differences between flies, indicating that fly behavior is composed of both stereotyped, and also highly variable and individualized aspects, mixed together over time. This approach uses off-the-shelf software and techniques, making it easily implementable, and provides insights into the similarity or dissimilarity between individuals, populations, and treatments applied to flies. However, their approach has a few hindrances which limit its utility as a general-purpose behavioral phenotyping tool. First, the use of t-SNE limits the embedded dimension of behavior to 2 or 3 dimensions. Further, a t-SNE clustering must be done once for all data, and new data cannot be easily projected into an existing embedding without the use of extra machine learning techniques. And also, the t-SNE algorithm is sensitive to hyperparameters, which strongly influence the number of clusters observed and their content. Overall, the approach produces excellent visualizations and interpretable groupings of behavioral

motifs which are extremely useful in the context of using behavior as a tool to study genetics or environmental perturbations. However, the number of behaviors discovered by the technique and their explicit content are not defined in a hypothesis-driven or introspectible manner, thus the technique is of less use for the study of behavior itself.

Systematically describing the structure of behavior in mice — and understanding how the brain alters that structure to enable adaptation — requires overcoming three challenges. First, it is not clear which features are important to measure when identifying candidate behavioral modules. Mice interact with the world by expressing complex three-dimensional (3D) pose dynamics, but because these are difficult to capture, most current methods track two-dimensional parameters such as the position, velocity, or 2D contour of the mouse (de Chaumont et al., 2012; Jhuang et al., 2010; Kabra et al., 2013; Spink et al., 2001; Ou-Yang et al., 2011). Second, given that behavior evolves on many timescales in parallel, it is not clear how to objectively identify the relevant spatiotemporal scales at which to modularize behavior. Finally, even stereotyped modules of behavior exhibit moment-to-moment and animal-to-animal variability (Colgan, 1978). This variability raises significant challenges for identifying the number and content of behavioral modules, or with associating observed actions with specific behavioral modules.

In this dissertation I describe a novel method, based upon recent advances in machine vision and learning, that identifies behavioral modules and their associated transition probabilities without human supervision. This approach uses 3D imaging to capture the

pose dynamics of mice as they freely behave in a variety of experimental contexts; these data reveal a surprising regularity that appears to divide mouse behavior into recognizable behavioral motifs that are organized at the sub-second timescale. A computational model then takes advantage of the observed fast temporal structure to describe mouse behavior as a series of reused modules, each a brief and stereotyped 3D trajectory through pose space that is connected in time to other modules through predictable transitions. This combined 3D imaging/modeling approach is used to explore how the underlying structure of behavior is altered after distinct environmental, genetic or neural manipulations, and we show that this method can detect both predicted changes in action and new features of behavior that had not been previously described. This work reveals that defining behavioral modules based upon structure in the 3D behavioral data itself — rather than using *a priori* definitions for what should constitute a measurable unit of action — can yield key information about the elements of behavior, offer insight into adaptive behavioral strategies used by mice, and enable discovery of subtle alterations in patterned action.

Chapter 2

Mapping Sub-Second Structure in Mouse Behavior

2.1 Introduction

In this chapter we develop a computer vision and machine learning methodology for automatically segmenting mouse behavior. In several orthogonal analyses, we observed a rhythmicity in the 3D pose data recorded from freely-moving, untrained mice with a 300-900 millisecond periodicity. We explored temporal structure at this timescale, and directly instantiated the hypothesis that behavior is composed of reused motifs of motion using recently-developed advances in probabilistic graphical models and Bayesian time-series modeling. Using these computational techniques, we discovered that behavior may be reliably segmented into a finite number of reused motifs of motion, which we term ”behavioral syllables”, that each last less than a second and encapsulate a stereotyped motion, or pose trajectory. Using our behavioral syllable segmentation as a lens into behavior, we investigated the effect of context, sensory stimuli, genetic alteration and optogenetic neural circuit perturbation on behavior. We find that the deployment of and transition structure between these behavioral syllables is sensitive to both subtle and

dramatic experimental alterations.

2.2 Results

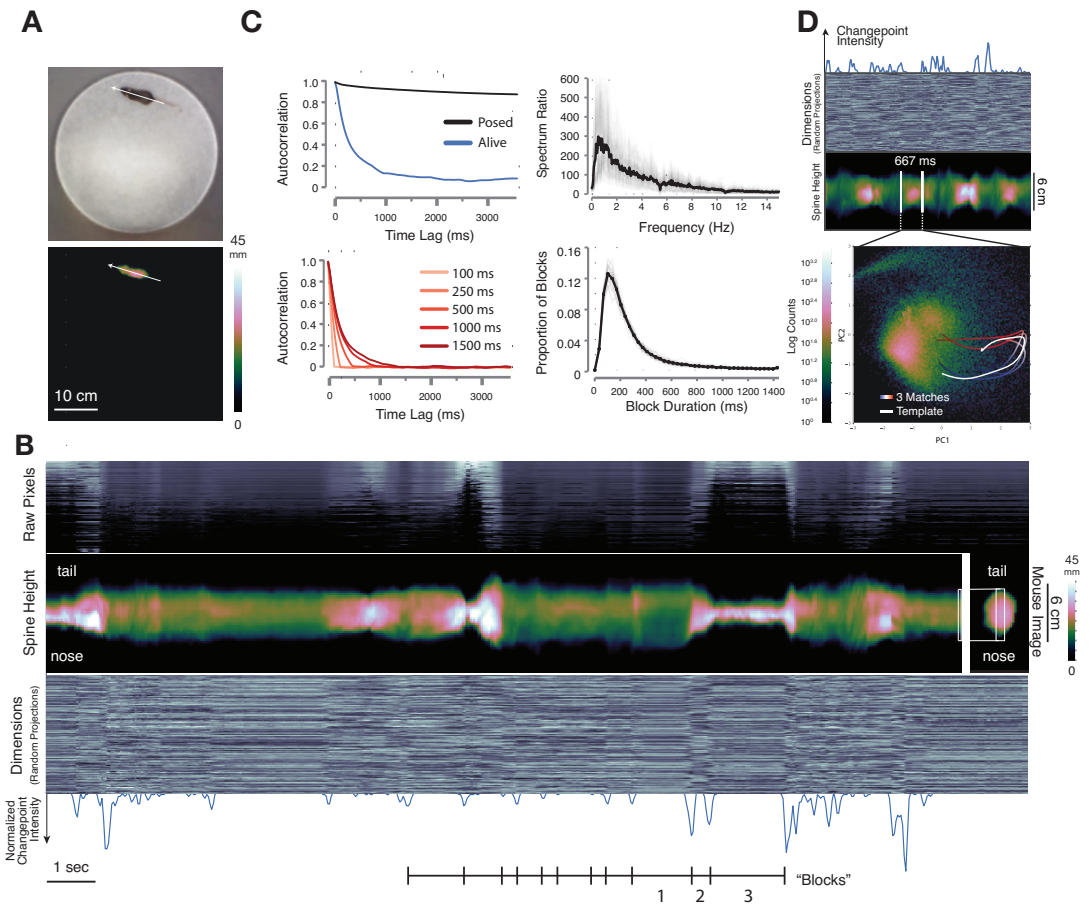
2.2.1 3D Imaging Captures Inherent Structure in Mouse Pose Dynamics

We wished to develop a method that would allow unsupervised phenotyping of mice based upon patterns of 3D movement. However, it is not clear whether spontaneous behaviors exhibited by mice have a definable underlying structure that can be used to characterize action as it evolves over time. To ask whether such a structure might exist, we measured how the shape of a mouse's body changes as it freely explores a circular open field (see Chapter 3 for more details).

We used a single depth camera placed above the arena to capture these 3D pose dynamics, and then extracted the image of the mouse from the arena, corrected imaging artifacts due to parallax effects, centered and aligned the mouse along the inferred axis of its spine, and then quantified how the mouse's pose changed over time (Figure 2.1A, Figure 2.2). Plotting these 3D data over time revealed that mouse behavior is characterized by periods during which pose dynamics evolve slowly, punctuated by fast transitions that separate these periods; this pattern appears to divide the behavioral imaging data into blocks typically lasting 200-900 ms (Figure 2.1B, 2.3A). This temporal structure is apparent in the raw imaging pixels (Figure 2.1B, top), the inferred shape of the mouse's spine (Figure 2.1B, middle), and dimensionally-reduced data that randomly samples from the depth images on the sensor (Figure 2.1B, bottom, see Chapter 3 for experimental procedures).

Figure 2.1 ***(following page)*** **|** A. Mouse imaged in the circular open field with a standard RGB (top) and 3D depth camera (bottom, mm = mm above floor). Arrow indicates the inferred axis of the animal's spine. B. Raw pixels of the extracted and aligned 3D mouse image (top panel, sorted by mean height), compressed data (bottom panel, 300 dimensions compressed via random projections arrayed on the Y axis, pixel brightness proportional to its value), and height at each inferred position of the mouse's spine (middle panel, "spine" data extracted from the mouse as indicated on the right, mm = mm above floor) each reveal sporadic, sharp transitions in the pose data over time. Note that the cross-sectional profile of the spine with respect to the camera varies depending upon the morphology of the mouse; when reared this profile becomes smaller, and when on all fours it becomes larger. Changepoints analysis (bottom panel, blue trace = normalized changepoint probability) identifies approximate boundaries between blocks. Blocks encode a variety of behaviors including locomotion (1), a pointing rear with the mouse's body elongated with respect to the sensor (2), and a true rear (3). C, upper left. Autocorrelation analysis performed on the top 10 dimensions of principal components data reveals that temporal correlation in the mouse's pose stabilizes after about 400 milliseconds (τ = 340 ± 58 ms). C, lower left. Shuffling behavioral data in blocks of 500 milliseconds or shorter destroys autocorrelation structure (shuffle block size indicated). C, upper right. Spectral power ratio between a behaving and dead mouse (mean plotted in black, individual mice plotted in grey) reveals most frequency content is represented between 1 and 6 Hz (mean = 3.75 ± .56 Hz). C, lower right. Changepoints-identified block duration distribution (mean = 358 ms, SD 495 ms, mean plotted in black, individual mice in gray, n=25 mice, 500 total minutes imaging,). D. Projecting mouse pose data (top panels, random projections and spine data depicted as in B) into Principal Component (PC) space (bottom) reveals that blocks of pose data encode reused trajectories (density of all recorded poses colormapped behind trajectories). Tracing out the path associated with a block highlighted by changepoint detection (top) identifies a trajectory through PC space (white). Similar trajectories identified through template matching (time indicated as progression from blue to red), are superimposed. Note that this procedure uses the first 10 PCs to identify matched trajectories, although only the first two PCs are depicted here.

Figure 2.1 | (Continued)



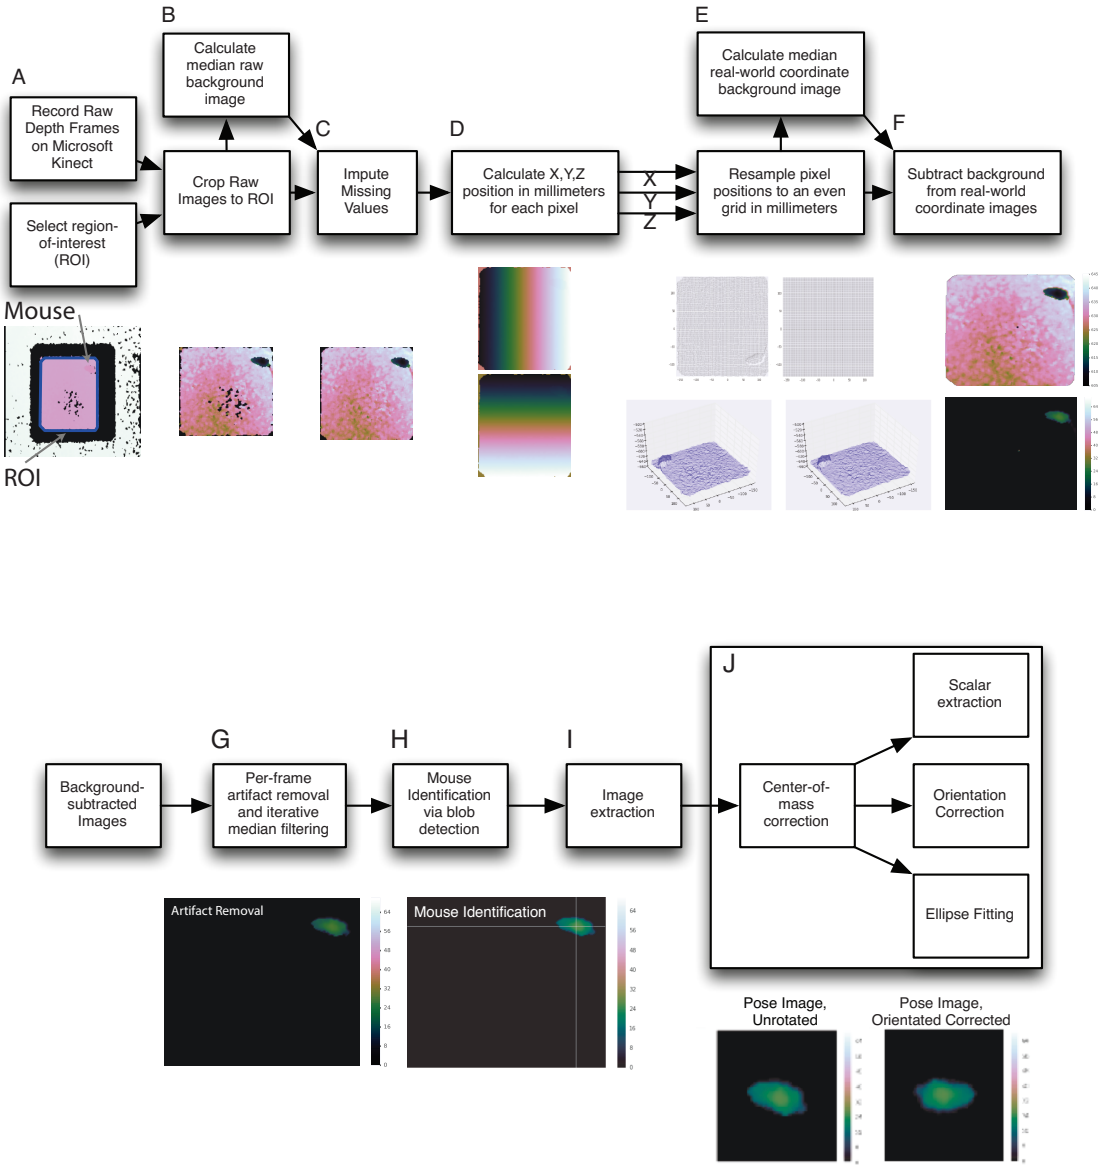


Figure 2.2 | Workflow (shown here with text descriptions of each step above and graphics corresponding to each step below) used to convert depth images of freely-behaving mice into mouse pose dynamics data in which each mouse is aligned along the axis of its spine. This process is described in detail in the Supplemental Experimental Procedure. In brief, after recording depth data within a behavioral apparatus (A), background values are computed (B) and missing pixel values are imputed (C). The mouse is then resampled into real-world coordinates to correct for parallax artifacts (D-E), and the background contributed by the apparatus is subtracted away (F). These images are then filtered, (G) and the image of the mouse is identified and extracted as a region of interest (H-I). The unrotated image is then aligned longitudinally based on fitting an ellipse to the boundary of the image (the major axis of which we refer to as the axis of its spine) (J), enabling quantitative analysis of egocentric pose dynamics while the animal freely translates in space. This pre-processed imaging data is then subject to wavelet decomposition and PCA for dimensionality reduction before submission to the modeling algorithms.

This structure is absent in data acquired from anesthetized or dead mice, but is present for the entire duration of each experiment in the pose dynamics data of mice exploring behavioral arenas of different shapes and after exposure to a wide variety of sensory cues, suggesting that it is a fundamental and ubiquitous feature of mouse behavior (data not shown, Figure 2.3).

To characterize this fast temporal structure we performed three separate quantitative analyses. First, approximate boundaries between blocks in the behavioral imaging data were identified by a changepoints algorithm, which revealed that the mean block duration was about 350 ms, roughly matching the timescale of the blocks apparent upon visual inspection (Figures 2.1B,C, Chapter 3). Second, we performed temporal autocorrelation analysis on the pose dynamics data, which demonstrated that autocorrelation in the mouse's pose largely dissipated after 400 ms ($\tau = 340 \pm 58$ ms, Figures 2.1C, 2.3B). This pattern of autocorrelation reflects specific behavioral dynamics organized at sub-second timescales, as it was destroyed by shuffling the behavioral data at timescales of 500 ms or less, and was not observed in synthetic mouse behavioral data designed to evolve with either random walk or Levy flight characteristics (Figures 2.1C, 2.3C). Third, we used a Wiener filter analysis to compare power spectral densities taken from live and dead mice; this approach identifies frequencies that must be changed in imaging data taken from a dead mouse to match the power spectrum of a live mouse. Nearly all of the frequency content differentiating behaving from dead mice was concentrated between 1 and 6 Hz (measured by spectrum ratio, mean 3.75 ± .56 Hz, Figures 2.1C, 2.3B). Taken together,

the qualitative appearance of block structure in the pose dynamics data, along with the convergent results obtained with these three quantitative analyses, demonstrate that mouse pose dynamics exhibit structure at the sub-second timescale. The observed temporal structure within the pose dynamics data suggests a timescale at which continuous behavior may be naturally segmented into meaningful components, as visual inspection of 3D movies revealed that each of the sub-second blocks of behavior appears to encode a recognizable action (e.g., a dart, a pause, the first half of a rear). To explore the possibility that these sub-second actions are stereotyped (and therefore reproducibly performed at different times during an experiment), the 3D mouse imaging data was subjected to wavelet decomposition followed by principal component analysis (PCA), which transformed each block in the pose dynamics data into a continuous trajectory through principal component (PC) space (Figures 2.1D).

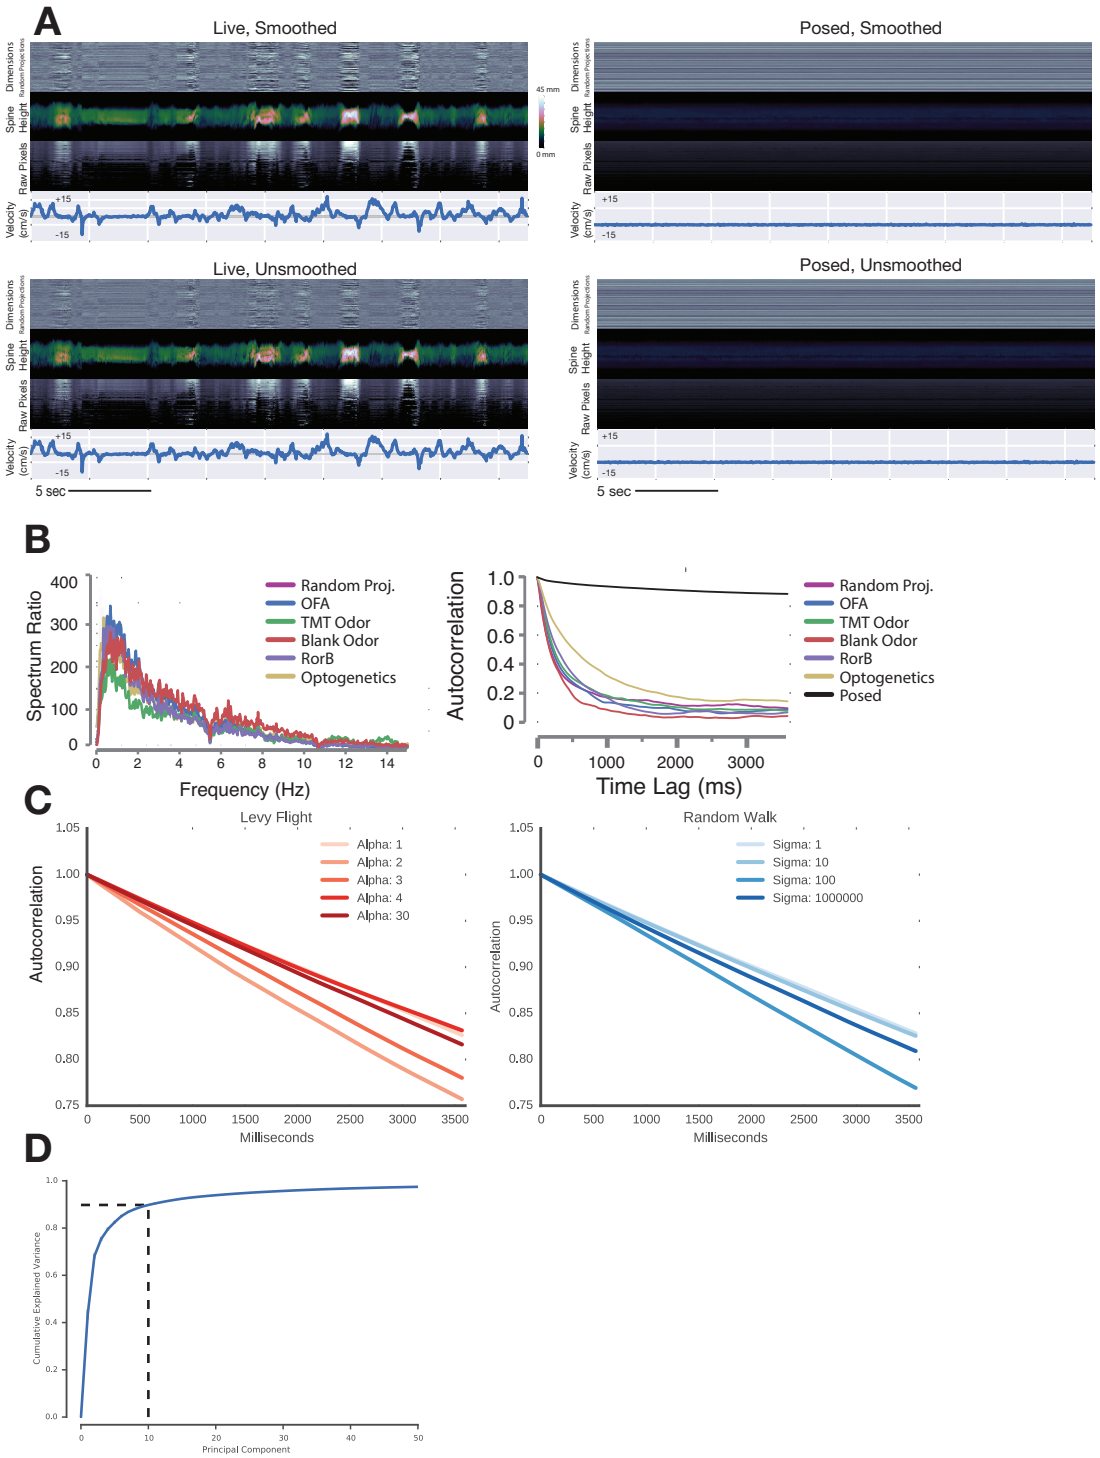
By scanning the behavioral data using a template matching method (using Euclidean distance among the first 10 PCs, which explain 88% of the data variance, Figure 2.3D), additional instances were identified in which each template action was reused (Figure 2.1D, Chapter 3, see Figure 2.4 for additional examples). These anecdotal observations suggest that mice create complex behaviors through the serial expression of stereotyped and reused behavioral modules (Tinbergen, 1951).

2.2.2 Mouse Behavior Can Be Described and Predicted with Modules and Transitions

Although our analysis suggests a timescale at which behavioral modules might exist, and qualitative inspection of 3D video and PCA trajectories is consistent with the possibility that sub-second blocks of behavior correspond to reused modules, current methods do not allow for the systematic identification of candidate behavioral modules in mice. Indeed, available approaches neither reveal whether dividing behavior into modules actually helps to explain the overall structure of behavior, nor identify the most likely number of modules expressed within any given dataset, or the content and durations of those modules.

Figure 2.3 ***(following page)*** **|** A. Block structure is present in random projections data, spine data and raw pixel data derived from aligned mouse pose dynamics data (random projections and spines are depicted as in Fig. 1A, raw data in bottom panels are heatmapped based upon imaging pixel value; values are sorted by mean height, and rarely-used pixels are omitted). Each dimension of random projections data uses all pixels from the original mouse image. Live mice exhibit significant block structure in imaging data (left panels), while dead mice do not (right panels). Smoothing refers to the modest median filtration applied during pre-processing, which does not visibly affect the structure of the data (Supplemental Experimental Methods). B. Dimensionality reduction does not significantly affect autocorrelation structure or spectral content of mouse pose dynamics data, nor does the particular experimental setup. The result of the Wiener filtering analysis (left) is similar for both a 10-dimensional PCA and a random projection representation of depth data, as well as across experimental conditions (OFA = open field assay, TMT Odor = square box with the fox odor TMT in one quadrant, Blank Odor = square box without introduced odorants, Ror1β = OFA data with mice mutant, heterozygous or wild-type at the Ror1β locus, Optogenetics = optogenetic experiment in the OFA with unilateral motor cortex implantation of an optical fiber). The autocorrelation of random projections and PCA data (right), both of which represent the same depth data, exhibit approximately the same falloff, demonstrating that data compression does not influence fine-timescale autocorrelation structure in the imaging data. The particular arena in which the mouse was recorded, or the experimental manipulation to which the mouse was subjected also does not affect the autocorrelation structure, indicating it is a core feature of our data. Note that the slightly different autocorrelation structure observed in the optogenetics data can be explained by the presence of an optical fiber implanted in the head of the animal, which alters the imaged pose dynamics of the animal. C. This correlation structure is not observed if mouse poses evolve as if taking a Levy flight (left) or random walk (right), suggesting that live mice express a specific sub-second autocorrelation structure. The alpha parameter and sigma parameter control the rate of dispersion for data generated from the Levy flight and Brownian motion processes, respectively. D. The first 10 principal components capture 88 percent of the variance in raw mouse imaging data. This number of dimensions was used for data analysis in the AR-HMM.

Figure 2.3 | (Continued)



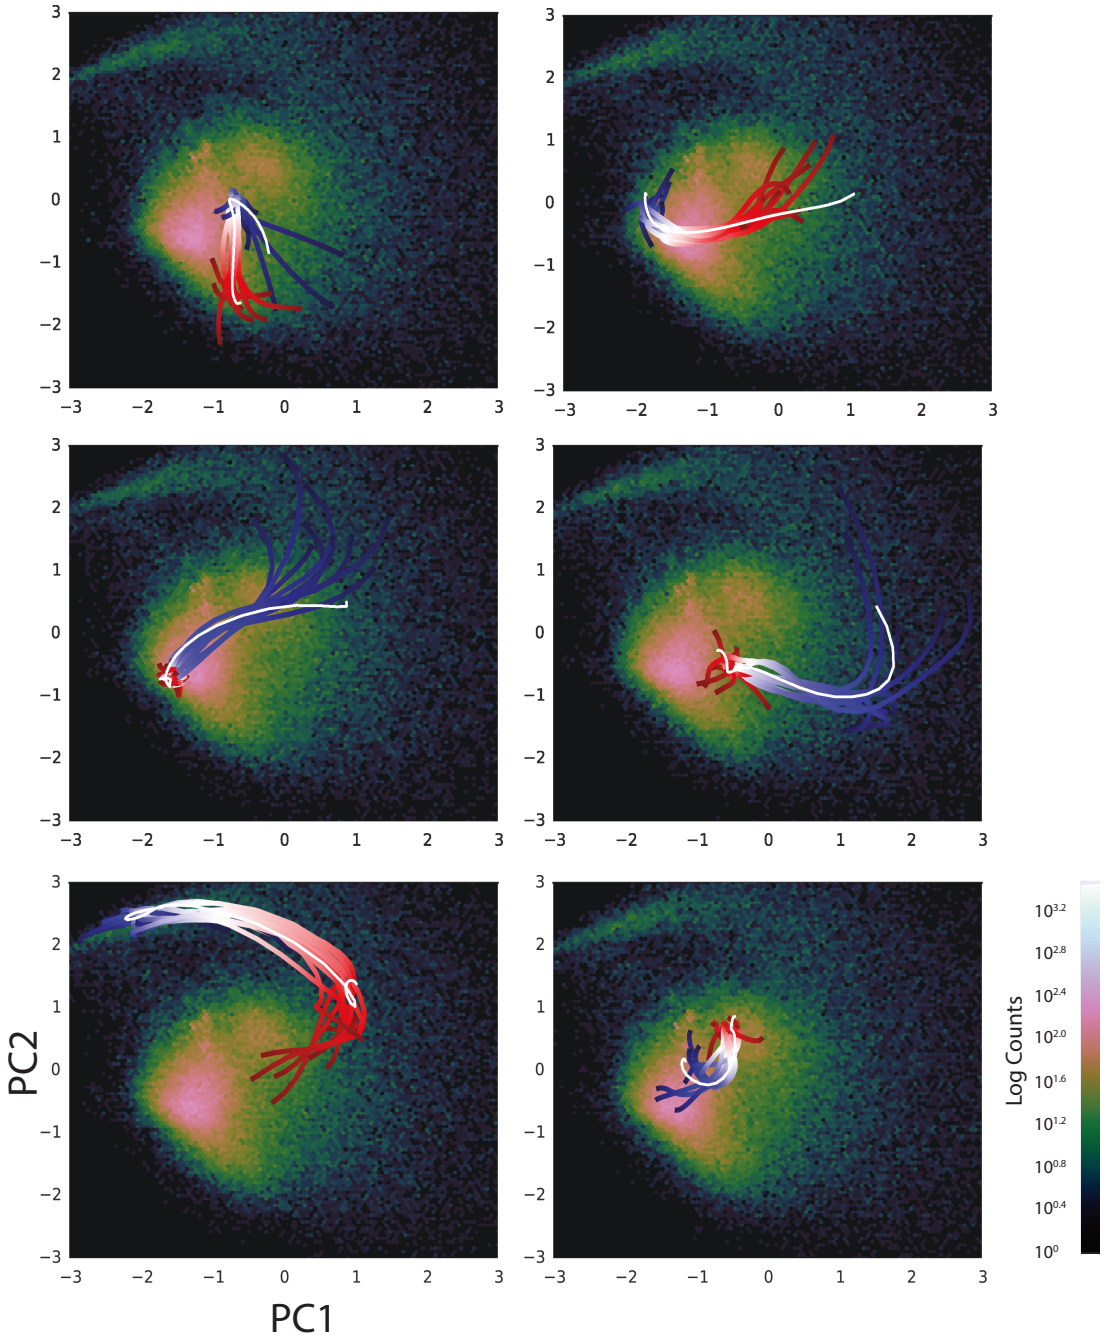
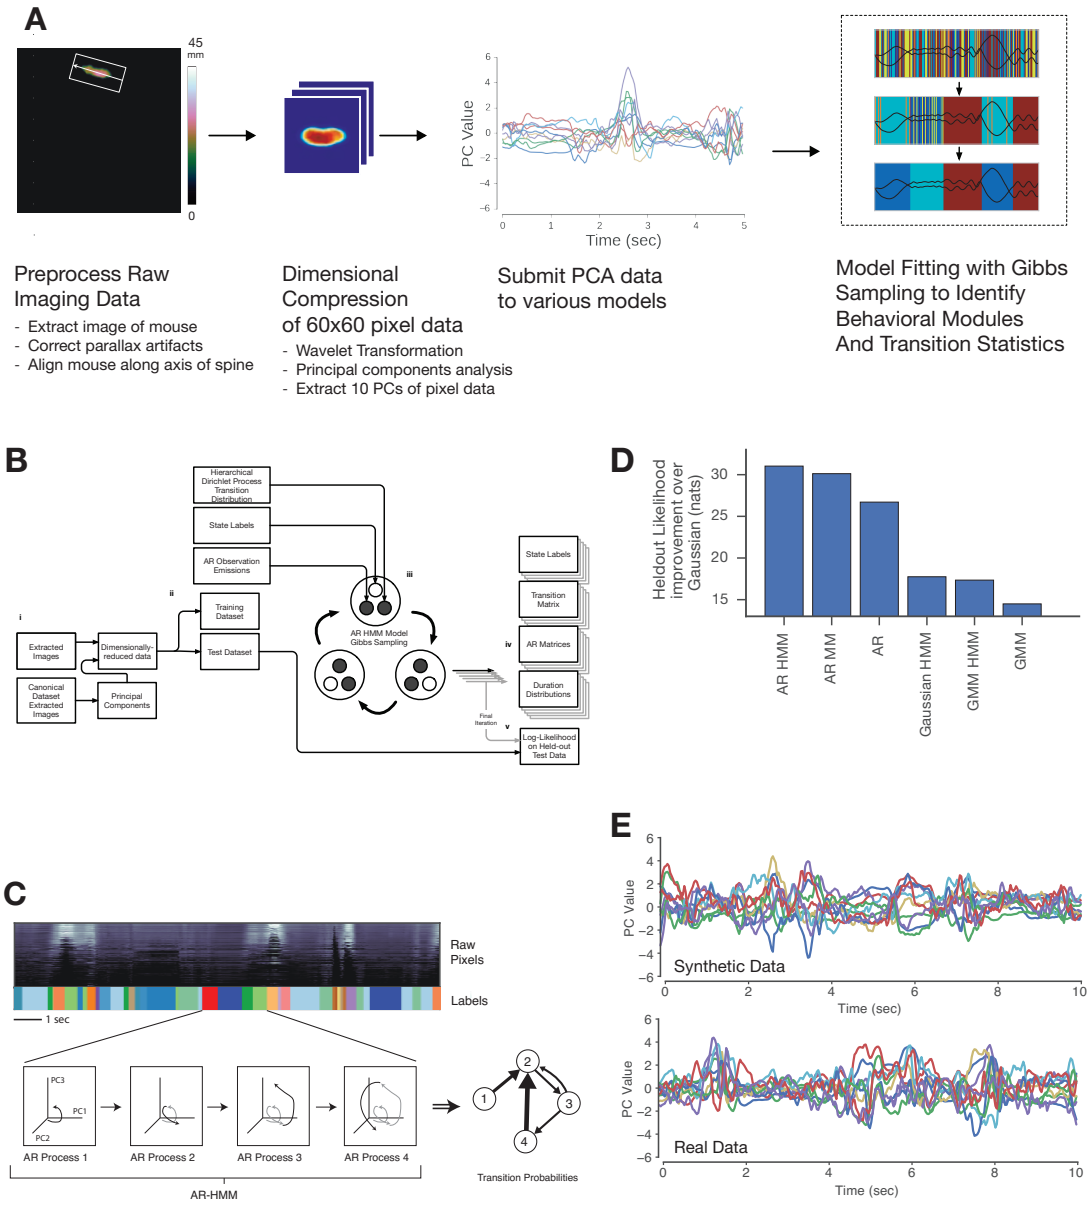


Figure 2.4 | Scanning the behavioral data using a template matching method identified additional instances in which given stereotyped motifs were reused (with time proceeding from blue to red). Six additional templates and their nine closest matched trajectories are plotted in the first two dimensions of principal components. The complete pose data in principal component space is represented as a background density plot. Note that 10 PCs were used to identify examples matching the seed templates, although only 2 PCs are depicted here for clarity.

To address these issues we built a family of computational models, each of which proposes a unique underlying structure for behavior, and asked which of these models best predicts the pose dynamics of freely behaving mice (Figures 2.6A, 2.5); we reasoned that the model that most closely fit behavioral data (to which the model had not been exposed) would reveal key features of the underlying organization of behavior, and could be used to characterize its components. After pre-processing the imaging data, the top 10 PCs of the data (Figure 2.5E) were used to fit each model; this use of PCs (which directly reflect the pixel data) as a basis for modeling minimized potential biases from feature engineering. Models were fit using Bayesian nonparametric and Markov Chain Monte Carlo techniques that can automatically identify structure within large datasets, including the optimal state number for a given dataset and model (see Chapter 3, Figure 2.5). Each model was trained on one set of pose dynamics data, and then tested for its ability to explain a separate set of held-out data, a metric canonically used to compare unsupervised learning models (Hastie et al., 2005). Our simplest model proposed that each behavioral module is a single fixed pose with no defined transition probabilities between the modules. We iteratively added structure to build increasingly complex models, which incorporated modules with more elaborate internal structures (ranging from mixtures of poses to smooth pose trajectories), allowed predictable transitions between specific modules (by embedding the modules within a Markov model), or both (Figure 2.5). Where relevant, the fitting procedures were explicitly focused to search for behavioral modules at the sub-second timescale matching the temporal structure identified using our model-free methods; this approach provided an

Figure 2.5 ***(following page)*** **|** A. Sketch of the part of the data processing pipeline that includes model fitting. The mouse is first identified, extracted and aligned from depth camera video frames along the major axis of the mouse, which approximates its spine. Parallax artifacts are removed. The tracked mouse images are then subjected to a wavelet decomposition, and then reduced to a principal component representation, yielding a multidimensional time series. This principal component time series is then modeled with one of a family of models, which is then fit using Gibbs sampling iterations; this procedure segments the principal component time series into reused behavioral modules. B. The workflow for fitting behavioral data with the AR-HMM, the best-performing of the family of models composed to represent different hypotheses about the underlying structure of behavior. The input to the workflow extracted and aligned mouse images, and the output is a label sequence equal in length to the original data timeseries that assigns a behavioral module identity for every recorded frame of behavioral data. This process is described in detail in the Supplemental Experimental Procedures. (i) First, after pre-processing image data are compressed from 3600 to 10 dimensions via wavelet decomposition and PCA. (ii) Data are randomly split into "train" and "test" segments (in a 70:30 ratio), in order to evaluate the performance of the model on unseen data, and to guard against overfitting. (iii) The parameters of the model are randomly initialized (parameters describing each module's characteristic trajectory through pose space, duration distribution, and transition statistics relative to other modules), and fit iteratively with Gibbs sampling. Gibbs sampling alternates between several updates: the algorithm first attempts to segment the imaging data into modules given a fixed set of transition statistics and a fixed description of the AR parameters that describe any given module, and then the algorithm switches to fixing the segmentation and updating the transition matrix and the AR parameters. Filled circles represent fixed parameters, while open circles represent parameters that are varied in a particular iteration. (iv) The product of this process is the parameters described above: "AR Matrices," the AR parameters that describe the specific trajectory through pose space that defines any given behavioral module, "Duration Distributions," the parameters describing the duration distribution for each behavioral module (in model classes in which this distribution is explicitly modeled), "Transition Matrix," the probability that any given module will precede or follow any other module, and "State Labels," the behavioral module inferred for each frame of 3D video data. (v) The quality of the fit in terms of predicting behavior is evaluated using the likelihood of data that was held out from training. C. A schematic representation of the aspects of behavior captured by the AR-HMM. The AR-HMM takes in a sequence of mouse behavioral data (fed into the model as 10 PCs, which are derived from pixel data), and generates a sequence of re-used labels over time ("Labels"), such that every moment of the mouse's behavior is associated with a single behavioral module. Each behavioral module is internally modeled as an autoregressive process that can be visualized as a trajectory through pose space ("AR Process 1," AR Process 2," etc.). The transition probabilities between modules are represented in a transition matrix, which may be represented as a state map, where each behavioral module is a node, and the probability of transition between two nodes is represented by the thickness of a directed arrow between those nodes ("Transition Probabilities"). D. The AR-HMM outperforms alternative models as measured by next-frame average held-out log likelihood. E. Data generated from a fit AR-HMM model are qualitatively similar to PCs derived directly from imaging data.

Figure 2.5 | (Continued)



important – and previously unavailable – constraint, given the multiple possible timescales upon which behavior evolves simultaneously.

The model that best fit previously-unseen behavioral data describes mouse behavior as a sequence of modules (each capturing a brief motif of 3D body motion) that switch from one to another at the sub-second timescale identified by our model-free analysis of pose dynamics (Figure 2.6A, 2.5D, Chapter 3). We refer to this model as an AR-HMM, as each behavioral module was modeled as a vector autoregressive (AR) process capturing a stereotyped trajectory through pose space, and the switching dynamics between different modules were modeled using a hidden Markov model (HMM). In other words, the model is a hierarchical description of behavior, with the ”internals” of each module reflecting the mouse’s pose dynamics over short timescales, and the longer-timescale relationships between behavioral modules (i.e., the possible module sequences) governed by the transition probabilities specified by an HMM (Figures 2.5B, 2.5C). The observation that the AR-HMM outperforms alternative models (Figure 2.6A) demonstrates that modularity and transition structure at fast timescales are critical for describing mouse behavior, a key prediction from ethology.

2.2.3 Model-Identified Behavioral Modules Are Stereotyped and Distinct

The AR-HMM systematically identifies modules and their transition probabilities from behavioral data without human supervision; this suggests that the AR-HMM can be used to identify behavioral modules and their associated transition probabilities – and thereby expose the underlying structure of behavior – in a wide variety of experimental contexts.

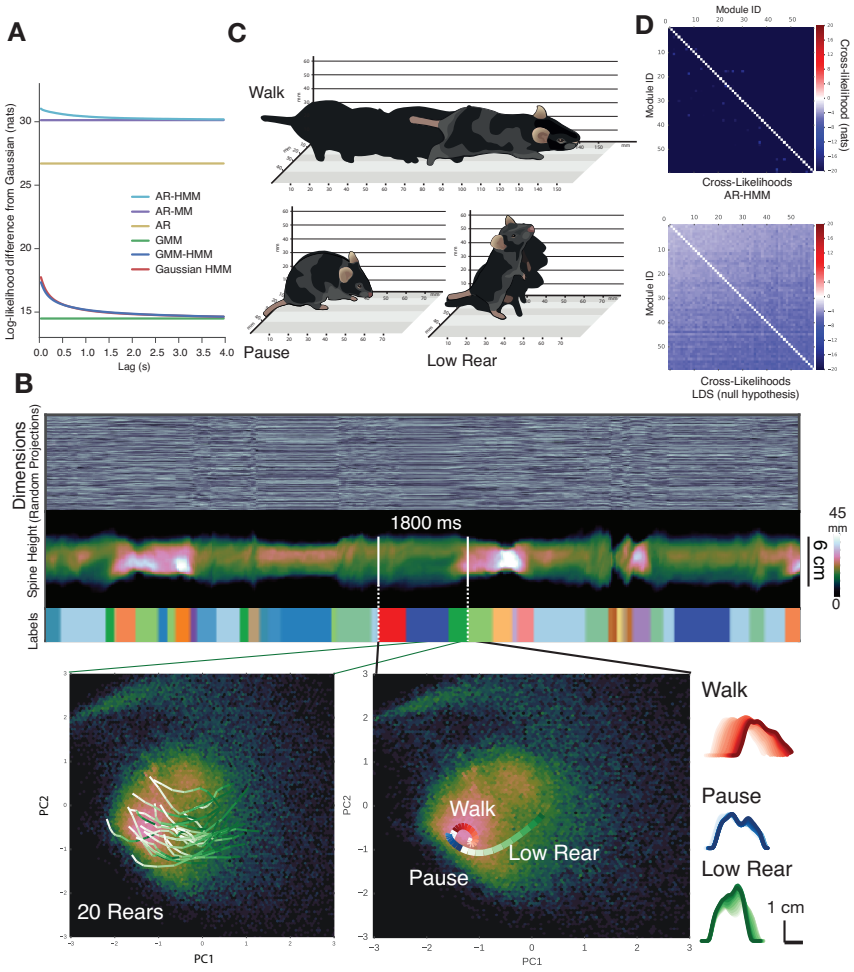


Figure 2.6 | A. Predictive performance comparison of computational models describing possible structures for behavior (details of each model and the comparison metric provided in Chapter 3). Models range from a Gaussian model (which proposes that mouse behavior is built from modules, each a single Gaussian in pose space) to an AR-HMM (which proposes that mouse behavior is built from modules, each of which encodes an autoregressive trajectory through pose space, and which transition from one to another with definable transition statistics; AR-MM = autoregressive mixture model, AR = autoregressive model, GMM = Gaussian mixture model, GMM-HMM = GMM hidden Markov model, Gaussian HMM = Gaussian hidden Markov model). Model performance plotted as the log likelihood (Y-axis) ascribed to held-out test data at some time lag (X-axis) into the future (after subtracting Gaussian model performance). B. The AR-HMM parses behavioral data into identifiable modules (top panels, marked ”labels”, each module is uniquely color coded). Multiple data instances associated with a single behavioral module (between green lines, encoding a rear) each take a stereotyped trajectory through PC space (bottom left, trajectories progress from white to green over time); multiple trajectories define behavioral sequences (bottom center). Each trajectory within a sequence encodes a different elemental action (side-on view of the mouse calculated from depth data, bottom right, time indicated as increasingly darker lines, from module start to end). C. Isometric-view illustrations of the 3D imaging data associated with walk, pause and low rear modules. D. Cross-likelihood analysis depicting the likelihood that a data instance assigned to a particular module is well-modeled by another module. Cross-likelihoods were computed for the open field dataset (above, see Chapter 3, units are nats, where enats is the likelihood ratio) and for module-free synthetic data whose autocorrelation structure matches actual mouse data (below).

We therefore performed a series of control analyses to establish whether the AR-HMM can indeed reliably identify behavioral modules encoding repeatedly-used and stereotyped motifs of distinguishable behavior that are organized at sub-second timescales.

Although the AR-HMM is tuned to identify modules at a particular timescale, it is possible that after training the model could fail to effectively capture temporal structure in behavior. However, the AR-HMM successfully identified modules at the fast behavioral timescale defined by the model-free methods, as the distribution of module durations was similar to the duration distribution for changepoints-identified blocks (Figure 2.7A). Importantly, the ability of the AR-HMM to identify behavioral modules depended upon the inherent sub-second organization of mouse pose data, as shuffling the behavioral data in small chunks (i.e., < 500 milliseconds) substantially degraded model performance, whereas shuffling the data in bigger chunks had little effect (Figure 2.7B).

We then asked whether model-identified modules encode repeatedly-used and stereotyped motifs of behavior. The pose trajectories associated with a specific model-identified behavioral module took similar paths through PC space, and visual inspection of the 3D movies associated with multiple instances of this module revealed they all encoded a stereotyped rearing behavior (Figure 2.6B, 2C). In contrast, data instances drawn from different behavioral modules traced distinct (and stereotyped) paths through PC space (Figure 2.6B, see Figure 2.7C for additional examples). Furthermore, visual inspection of the 3D movies assigned to different modules reveals that each encoded a coherent pattern of 3D motion that post hoc can be distinguished and labeled

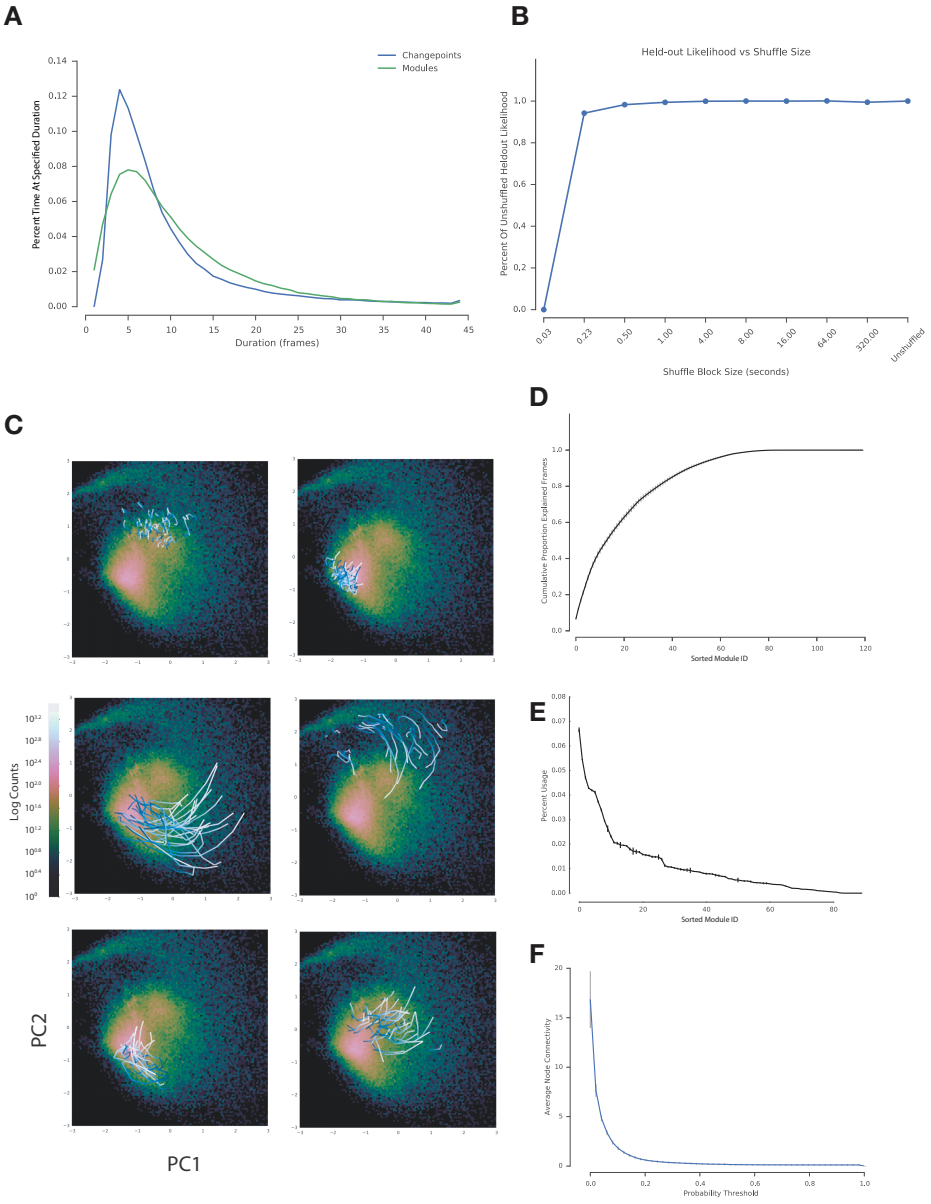


Figure 2.7 | A. Duration distributions for changepoint-identified blocks and AR-HMM-identified modules are similar. B. Shuffling behavioral data at fast timescales lowers AR-HMM performance, as measured by the ratio of held-out likelihood before and after shuffling. C. Data instances drawn from different behavioral modules trace distinct paths through PCA space. Six separate representative behavioral modules are plotted in the first two principal components, overlaid on a density indicating the distribution of all recorded poses. D. 95 percent of frames were explained by only 51 behavioral modules; 99 percent of frames were explained by 65 behavioral modules in the open field dataset (error bars indicate ± S.E. calculated with bootstrap samples). E. Modules (X axis) sorted by usage (Y axis) with Bayesian credible intervals indicated as error bars, which are smaller than SEs computed with the nonparametric bootstrap (see Fig. 2). F. Module interconnectivity is sparse. With no thresholding, each module is interconnected with on average 16.85±0.95 other modules.

with descriptors. The modules identified by the AR-HMM are distinct from each other, as a cross-likelihood analysis demonstrated that the imaging data associated with a given module are best assigned to that module, and not to any of the other behavioral modules (Figure 2.6D, Chapter 3). In contrast, the AR-HMM failed to identify any well-separated modules in a synthetic mouse behavioral dataset that lacks modularity but otherwise matches all multidimensional and intertemporal correlations of the real data, demonstrating that the AR-HMM does not discover modularity where none exists (Figure 2.6D). Furthermore, rerunning the AR-HMM training process from random starting points generated highly similar behavioral modules (r^2 = .94 ± 0.03, n = 15 restarts, see Chapter 3); in comparison, models with lower held-out likelihood scores have lower consistency (r^2 < .4), suggesting that these alternatives fail to reliably identify underlying structure in behavior. The AR-HMM output was robust to the specific training data used, as models created from subsets of a larger dataset representing a single experiment were highly similar to each other (r^2 > .9). These findings demonstrate that the AR-HMM converges on a consistent set of behavioral modules regardless of the specific training data (within a given experiment) or how the model is initialized. Finally, the modules of behavior and the associated transition probabilities identified by the AR-HMM appear to fully capture the richness of mouse behavior, as a 3D movie of a behaving mouse generated by a trained AR-HMM was qualitatively difficult to distinguish from a 3D movie of behavior exhibited by a real animal (Figure 2.5E); in contrast, movies generated by more poorly performing models appeared discontinuous and were easily distinguished from real

animals (data not shown). Taken together with our model-free and model-based analyses and controls, the observation that after training our model can synthesize a convincing replica of 3D mouse behavior from learned modules and transition probabilities is consistent with the hypothesis that mouse behavior is organized into distinct sub-second modules that are combined to create coherent patterns of action.

2.2.4 Using the AR-HMM to Characterize Baseline Patterns of Behavior

The AR-HMM identifies two key features of mouse behavior (from the perspective of 3D pose dynamics): which behavioral modules are expressed during behavior, and how those modules transition into each other over time to create action.

The AR-HMM identified 60 reliably-used behavioral modules from a circular open field dataset, which is representative of normal mouse exploratory behavior in the laboratory (51 modules explained 95 percent of imaging frames, and 65 modules explained 99 percent of imaging frames, Figures 2.8A, 2.7D, 2.7E; subjective categorization of the 51 most-used modules is shown in Figure 2.9A). These modules were connected to each other over time in a highly non-uniform manner, with each module preferentially linked to some modules and not others (Figure 2.8B, Figure 2.7F; average node degree after thresholding transitions that occur with < 5% probability, 4.08 ± .10). This specific transition pattern among modules constrained the module sequences observed in the dataset (17,000/ 75,000 possible sequences of three modules (”trigrams”) given the total data size) demonstrating that certain module sequences were favored over others, and that mouse behavior is therefore predictable (per frame entropy rate without

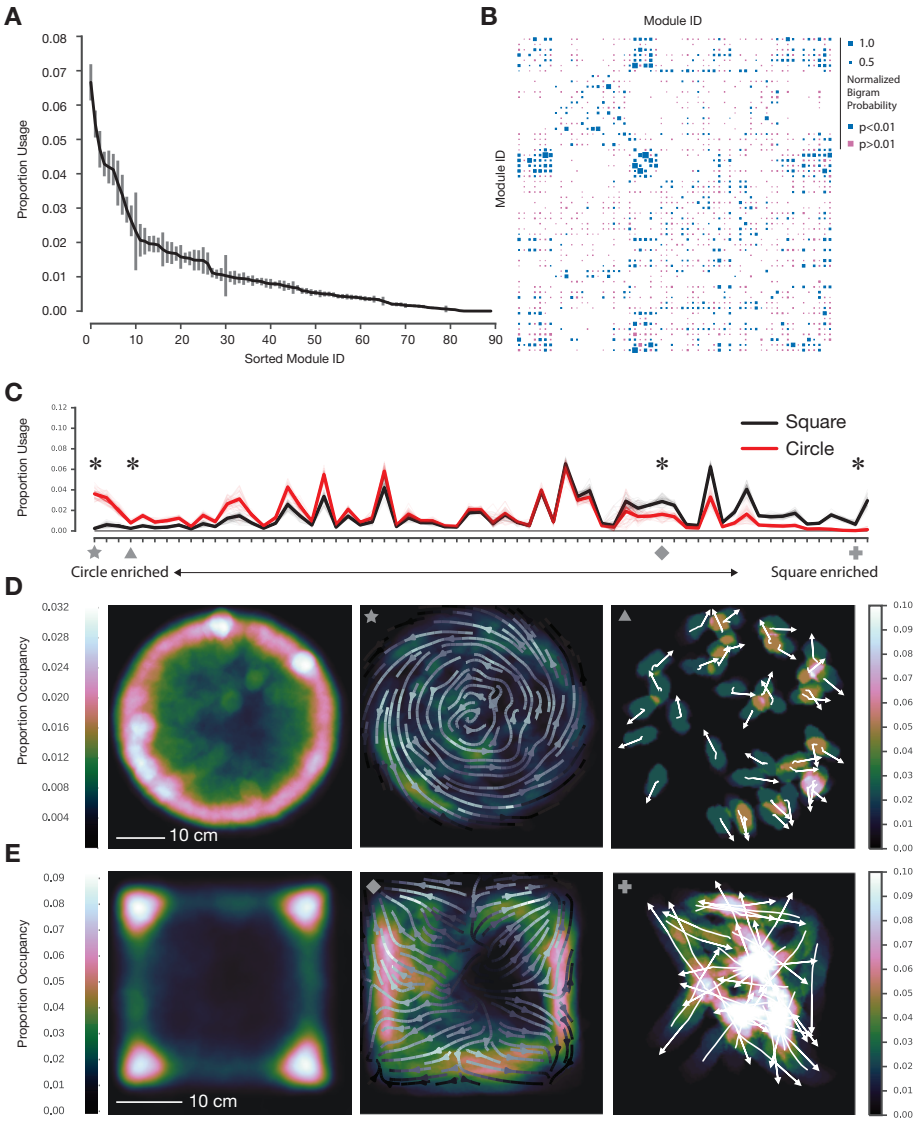


Figure 2.8 | A. Modules identified by the AR-HMM, sorted by usage (n = 25 mice, 20 minutes per mouse, data from circular open field, error bars are SEs calculated using bootstrap estimation, n=100 bootstrap estimates, see Figure 2.7E for Bayesian credible intervals). B. Hinton diagram depicting the probability that any pair of modules is observed as an ordered pair. C. Module usage, sorted by context (see Chapter 3 for sorting details). Mean usages across animals depicted with dark lines, with bootstrap estimates depicted in fainter lines (n=100). Marked modules discussed in main text and shown in panel D: star = circular wall-hugging locomotion ("thigmotaxis"), triangle = outward-facing rears, diamond = square thigmotaxis, cross = square dart. Usage for all asterisk-marked modules significantly modulated by context. D. Occupancy plot of mice in circular open field (left, n=25, 500 minutes total) indicating cumulative spatial positions across all experiments. Occupancy plot depicting deployment of circular thigmotaxis module (middle, average orientation across the experiment indicated as arrow field) and circle-enriched outward-facing rear module (right, orientation of individual animals indicated with arrows). E. Occupancy plot of mice in square box (left, n=15, 300 minutes total) indicating cumulative spatial positions across all experiments. Occupancy plot depicting a square-enriched thigmotaxis module (middle, average orientation across the experiment indicated as arrow field), and square-specific darting module (right, orientation of individual animals indicated with arrows).

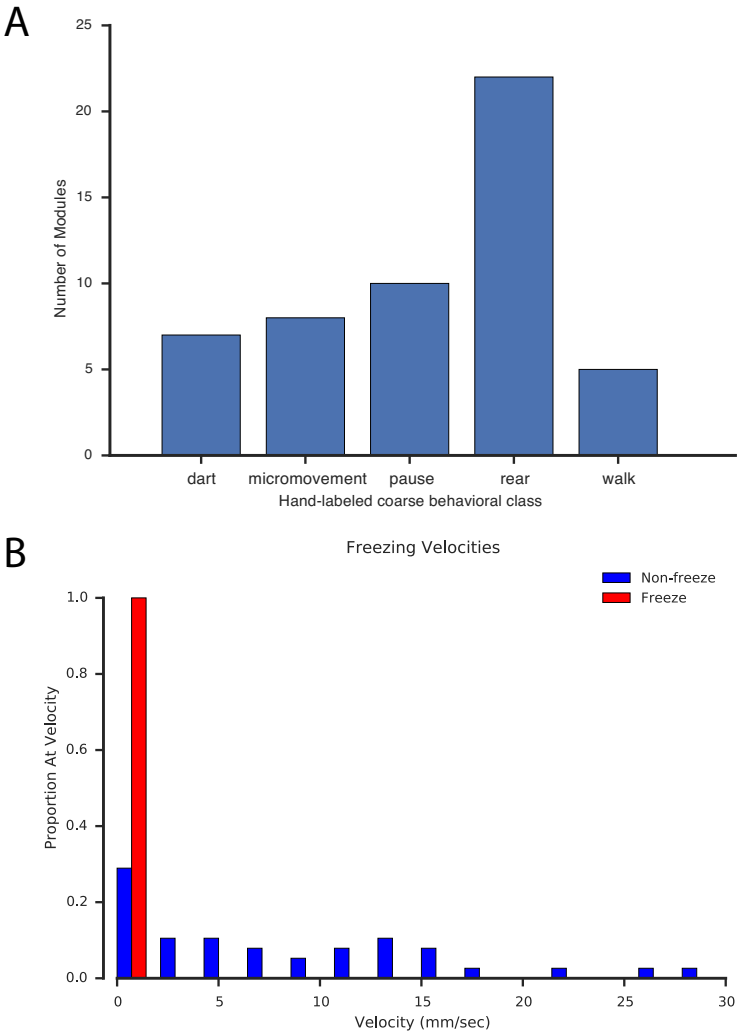


Figure 2.9 | A. Histogram depicting the number of modules that fall into one of five coarse labels of behavior, as scored by a human observer; this coarse subjective description is intended to offer a sense of the distribution of behaviors parsed by the AR-HMM. Note that these module descriptors are completely subjective, and applied post-hoc after modeling. Furthermore, inspection revealed that the modules within each broad category could be further distinguished from the others within that category; the “rear” category, for example, included a wide variety of behaviors in which the mouse’s head was transiently or perdurantly elevated, including true rears, elevated pointing behaviors, head-bobs, and sniffs. B. Histogram depicting the average velocity of the modules that were interconnected after TMT exposure (red), which individually and collectively encode “freezing,” compared to all other modules in the dataset.

self-transitions 3.78 ± .03 bits, with self-transitions .72 ± .01 bits, entropy rate in a uniform matrix 6.022 bits; average mutual information without self-transitions 1.92 ± .02 bits, with self-transitions 4.84 bits ± .03 bits. Note that while estimating coarse changes in specific trigram frequencies is possible, accurately estimating higher-order k-gram transition statistics is difficult as the amount of data required grows exponentially with k. We therefore focus our analysis on lower-order statistics such as module usage frequencies and temporal interactions between pairs of modules.

2.2.5 Using the AR-HMM to Characterize the Nature of Behavioral Change

We tested whether the AR-HMM could effectively capture changes in behavior (both predicted and unpredicted) elicited by a range of simple experimental interventions designed to probe the influence of the environment, genes, or neural circuit activity on behavior. We first asked how mice behavior adapts to changes in apparatus shape. We imaged mice within a small square box and then co-trained our model with both the circular open field and square box data, thereby enabling direct comparisons of modules and transitions under both conditions; the modules identified by this co-training procedure did not erroneously lump together data that would otherwise be distinguishable, as each behavioral module's mean pose trajectory was stereotyped across the different experimental arenas (data not shown). Although mice tended to explore the corners of the square box and the walls of the circular open field, the overall usage of most modules was similar between these apparatuses, consistent with exploratory behavior sharing many common features across arenas (Figure 2.8C).

However, the AR-HMM also identified a small number of behavioral modules that were deployed selectively in just one context, consistent with the idea that different physical environments drive expression of new behavioral modules (Figure 2.8C). For example, one circular arena-specific module encoded a behavior in which mice walk near the arena wall with a body posture that matches the curvature of the wall, while within the square box mice expressed a context-specific module that encodes a dart out of the center of the square (Figures 2.8D and E). Several behavioral modules were also differentially enriched (but not exclusively expressed) in one context or the other. In the circular arena, for example, mice preferentially executed a rear characterized by pointing outwards while pausing near the center of the open field, whereas in the smaller square box mice preferentially expressed a high rear in the corners of the box (Figure 2.8D, data not shown).

These results demonstrate that the AR-HMM can effectively capture predictable changes in behavior resulting from altering the physical environment (like walking along a curved wall or rearing in a corner). Importantly, these results also demonstrate that the AR-HMM can unmask arena-specific patterns of behavior that are expressed in the center of both arenas, away from the physically-constraining walls (like the darting across the center of the square box and the outward pointing behavior expressed in the circle); this surprisingly suggests that arena shape influences mouse behavior in a manner that extends significantly beyond the predictable changes in action at the walls themselves. Taken together, these experiments reveal that the AR-HMM can suggest strategies used by the mouse brain to adapt to new physical environments: in the case of a change in environmental geometry, this

strategy includes the recruitment of a limited set of context-specific behavioral modules into baseline patterns of action, and a broad rewriting of where in space modules are expressed with respect to the arena boundaries.

The small number of behavioral modules distinguishing the circular and square arenas suggests only modest differences in the global pattern of behavior in these two experiments. To ask how the AR-HMM captures changes in the underlying structure of behavior after an overt change in behavioral state, we exposed mice to an ethologically-relevant olfactory cue, the aversive fox odor trimethylthiazoline (TMT), which was delivered to one quadrant of the square box via an olfactometer. This odorant profoundly changes mouse behavior, inducing odor investigation, escape, and freezing behaviors that are accompanied by increases in corticosteroid and endogenous opioid levels (Fendt et al., 2005; Wallace & Rosen, 2000).

Consistent with these known effects, mice sniffed the odor-containing quadrant and then avoided the quadrant containing the predator cue, displaying prolonged periods of immobility traditionally described as freezing behavior (Figure 2.10A, Figure 2.9B). Given that TMT-induced behaviors are dramatically different than those observed at baseline, one might predict that TMT should induce new behavioral modules that underlie the generation of these new actions. Surprisingly, the AR-HMM revealed that the TMT-induced suite of new behaviors was best explained by the same set of behavioral modules that were expressed during normal exploration; several modules were up- or down-regulated after TMT exposure, but new modules were not introduced or eliminated relative to control (Figure 2.10B). We therefore asked whether the changes in observed

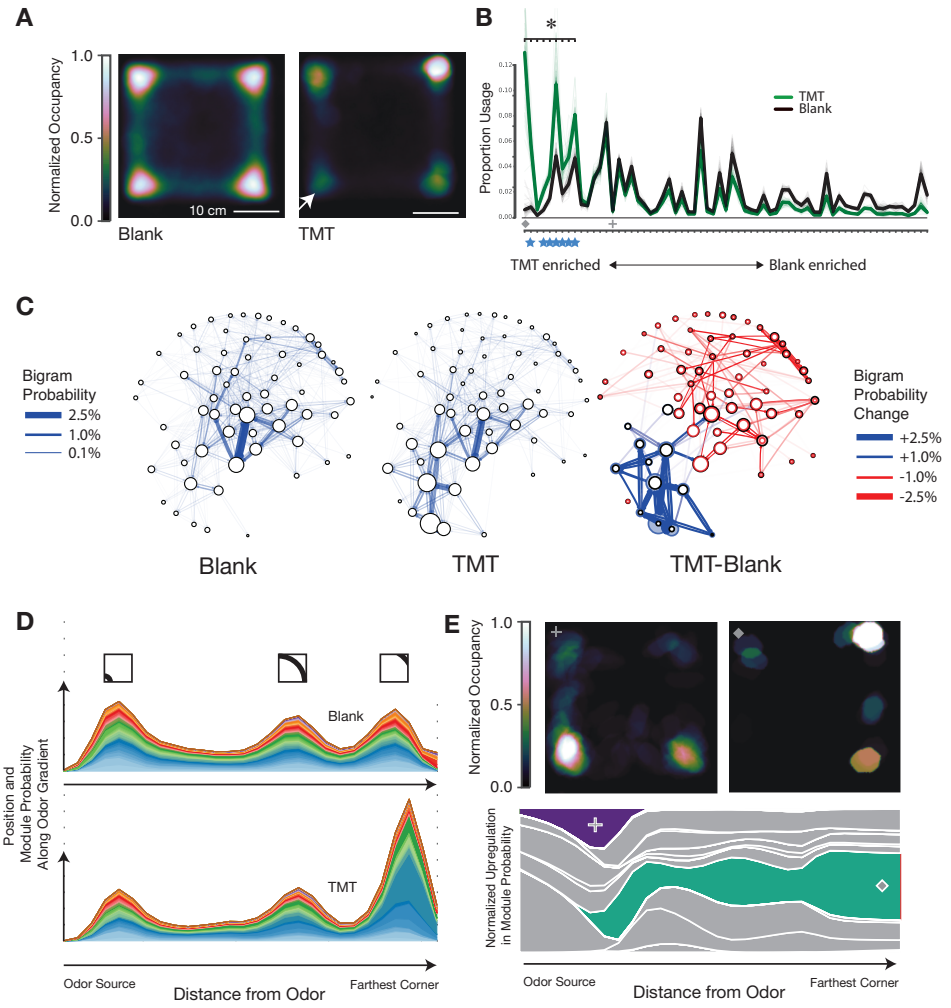


Figure 2.10 | A. Occupancy plot of mice under control conditions (n=24, 480 minutes total) and after exposure to the fox-derived odorant trimethylthiazoline (TMT, 5% dilution, n=15, 300 minutes total, model co-trained on both conditions) in the lower left quadrant (arrow). Plots normalized such that maximum occupancy = 1. B. Module usage plot sorted by "TMT-ness" (Dark lines depict mean usages, bootstrap estimates depicted in fainter lines, sorting as in Figure 3). Marked modules discussed in main text and panel E: cross = sniff in TMT quadrant, diamond = freeze away from TMT. Blue stars indicate freezing modules. Asterisk indicates statistically-significant regulation (Wald test, Holm-Bonferroni corrected, $p < 10\text{-}4$). C, left and middle. Behavioral state maps for mice exploring a square box before and after TMT exposure, with modules depicted as nodes (usage proportional to the diameter of each node), and bigram transition probabilities depicted as directional edges. Graph layout minimizes the length of edges and is seeded by spectral clustering to emphasize local structure. C, right. Statemap depiction of the difference between blank and TMT. Usage differences are indicated by the newly-sized colored circles (upregulation indicated in blue, downregulation indicated in red, previous usages in control conditions indicated in black). Altered bigram transition probabilities are indicated in the same color code; only those significant transition probabilities ($p<0.01$) are depicted. D. Mountain plot depicting the joint probability of module expression and spatial position, plotted with respect to the TMT corner (X-axis); note that the "bump" two-thirds of the way across the graph occurs due to the two corners equidistant from the odor source (see inset for approximate position in square box, modules are color coded). E. Occupancy plot (upper) indicating spatial position at which mice after TMT exposure emit an investigatory sniffing module (left) or a pausing module (right). Mountain plot (lower) indicating the differential deployment of these two modules (purple, green; other modules in grey) with respect to distance from the odor source.

behavior were the consequence of altered connections between behavioral modules. Plotting the module transitions altered after exposure to TMT defined two neighborhoods within the behavioral statemap; the first included an expansive set of transitions that was modestly downregulated by TMT, and the second included a focused set of transitions that was upregulated by TMT (Figure 2.10C). Many of these newly-interconnected modules encoded different forms of freezing behavior (average velocity .14 ± .54 mm/sec, for other modules 34.7 ± 53 mm/sec, Figure 2.9B). In addition, the TMT-initiated modulation of transition probabilities altered the expression of specific behavioral sequences; for example, the most commonly-observed sequence of three freezing modules was expressed 716 times after TMT exposure (in 300 minutes of imaging), as opposed to just 17 times under control conditions (in 480 minutes of imaging). The stimulus-evoked rewriting of transition probabilities was accompanied by an increase in the overall predictability of mouse behavior (per frame entropy rate fell from 3.92 ± .02 bits to 3.66 ± .08 bits without self-transitions, and from .82 ± .01 bits to .64 ± .02 bits with self-transitions) consistent with the mouse enacting an avoidance strategy that was more deterministic in nature than locomotor exploration. Proximity to the odor source also governed the usage of specific behavioral modules (Figures 2.10D, 4E). For example, a set of freezing-related modules tended to be expressed in the quadrant most distant from the odor source, while the expression of an investigatory rearing module (whose overall usage was not altered by TMT) was specifically enriched within the odor quadrant (Figures 2.10D, 4E). Although TMT is known to induce dramatic changes in behavior, it has not been possible to

systematically identify those specific behavioral features altered in response to this odorant, or to place those altered features in context with normal patterns of exploration. Analysis by the AR-HMM suggests that the strategy used by the mouse brain to adapt to the presence of TMT overlaps with – and yet is distinct from – that used to accommodate physical changes in the environment. As was true for a changed physical environment, exposure to TMT alters the spatial deployment of modules and sequences to support particular patterns of action; in contrast, the complete cohort of behaviors elicited by TMT, including seemingly ”new” behaviors such as freezing, are the consequence of altered transition structure between individual modules. Behavioral modules are not, therefore, simply reused over time, but instead act as flexibly interlinked components whose specific sequencing and deployment in space has profound consequences for the generation of adaptive behavior.

2.2.6 Sub-Second Architecture of Behavior Reflects the Influence of Genes and Neural Activity

As described above, the AR-HMM shows the fine-timescale structure of behavior to be sensitive to persistent changes in the physical or sensory environment. However, manipulation of individual genes or neural circuits influences behavior across a range of spatiotemporal scales and with variable penetrance and reliability; these changes in behavior may or may not be effectively captured by a classification method designed to characterize the sub-second structure of action. We therefore directly asked whether the AR-HMM could systematically reveal the behavioral consequences of manipulating the

mouse genome or activity within the nervous system. To explore this possibility, we used the AR-HHM to characterize the phenotype of mice mutant for the retinoid-related orphan receptor (Ror1β) gene, which is expressed in neurons in the brain and spinal cord. This mouse was selected for analysis because adult homozygous mutant animals permanently exhibit abnormal gait, which would be expected to be observed during a brief open-field experiment (Andre et al., 1998; Eppig et al., 2015; Liu et al., 2013; Masana et al., 2007).

Analysis with the AR-HMM revealed that littermate control mice are nearly indistinguishable from fully inbred C57/Bl6 mice, whereas homozygous mutant mice express a unique behavioral module encoding a waddling gait (Figure 2.11A,C, see Chapter 3 for statistical details, data were co-trained with both circular open field and TMT datasets to facilitate comparisons with C57 animals). Conversely, the expression of five behavioral modules encoding normal forward locomotion in wild-type and C57 mice was downregulated in Ror1β mutants (Figure 2.11A, average during-module velocity = 114.6 ± 76.3 mm/sec).

Previously unobserved phenotypes in the Ror1β mutant mice were also uncovered by the AR-HMM, as the expression of a set of four modules that encoded brief pauses and headbobs was upregulated in mutant mice (Figure 2.11A, average during-module velocity = 8.8 ± 25.3 mm/sec); this pausing phenotype had not been previously reported in the literature. Furthermore, heterozygous mice, which have no reported phenotype (Andre et al., 1998; Eppig et al., 2015; Liu et al., 2013; Masana et al., 2007), exhibit wild-type running wheel behavior (Masana et al., 2007), and appear normal by eye – were also

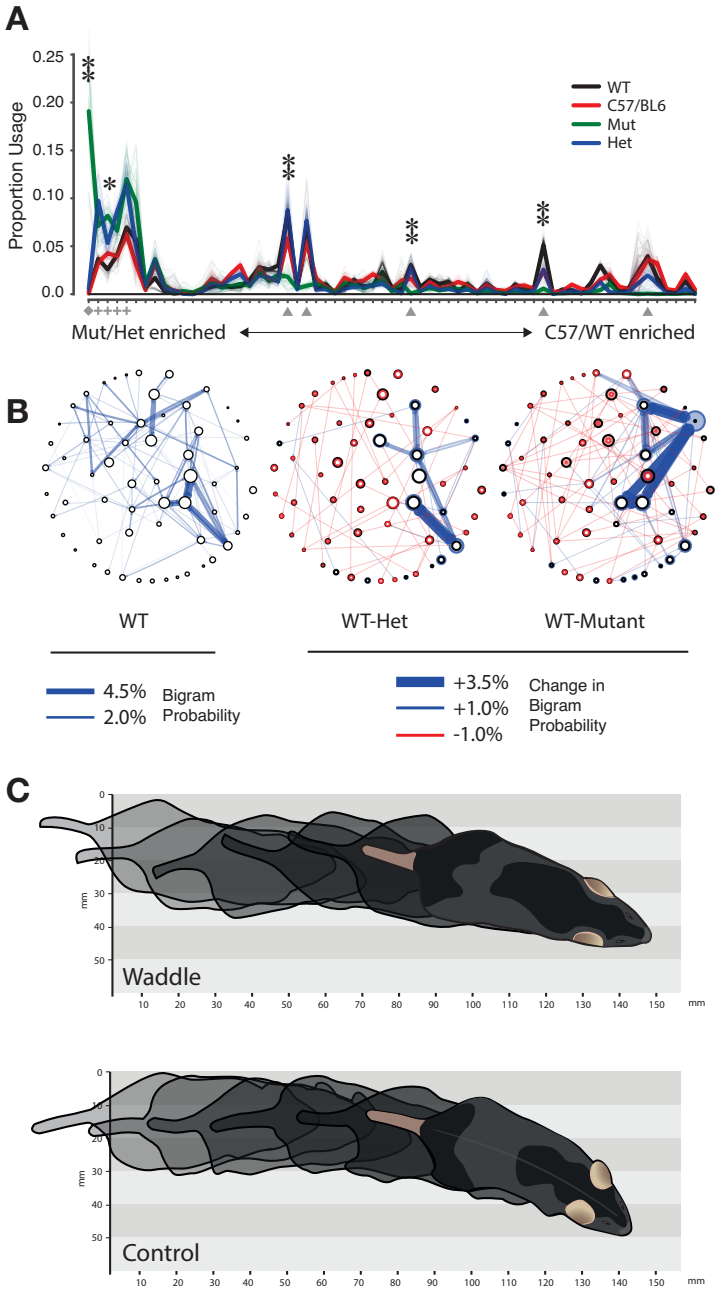


Figure 2.11 | A. Usage plot of modules exhibited by littermate Ror1β mice (n=25 C57/BL6, n = 3 +/+, n = 6 +/- , n = 4 -/-, open field assay, 20 minute trials), sorted by ”mutant-ness” (sorting and depiction as in Figure 3). Wild-type module usage is not statistically different from C57, but differs significantly from homozygote and heterozygote (Hotelling two-sample t-squared test, p < 10-15). Marked modules described in main text, diamond = waddle, triangle = normal locomotion, cross = pause. Single asterisk indicates significant usage difference between mutant and wildtype, p < 0.05, double asterisk indicates p < 0.01 under Wald test, Holm-Bonferroni corrected. B. State map depiction of baseline OFA behavior for +/+ animals as in Figure 4C (left); difference state maps as in Figure 4C between the +/+ and +/- genotype (middle), and +/+ and -/- genotype (right); all depicted transitions that distinguish genotypes are statistically significant, p<.01. C. Illustration of the ”waddle” module in which the hind limbs of the animal are elevated during walking.

found to express a fully-penetrant mutant phenotype: they overexpressed the same set of pausing modules that were upregulated in the full Ror1β mutants, while failing to express the more dramatic waddling phenotype (Figure 2.11A). Differences between wild-type animals and both heterozygotes and mutants were also observed in transition probabilities associated with these pausing modules (Figure 2.11B).

The AR-HMM therefore describes the pathological behavior of Ror1β mice as the combination of a single neomorphic waddling module, decreased expression of normal locomotion modules, and increased expression of a small group of physiological modules encoding pausing behaviors; heterozygous mice express a defined subset of these behavioral abnormalities, whose penetrance is not intermediate but equals that observed in the mutant. These results suggest that the sensitivity of the AR-HMM allows fractionation of severe and subtle behavioral abnormalities within the same litter of animals, facilitates comparisons amongst genotypes and enables discovery of new phenotypes. In the case of the Ror1β animals, a permanent alteration in DNA sequence is translated into an ongoing change in the overall sub-second statistical structure of behavior, one that is expressed continuously over the lifetime of the animal. We also wished to characterize how transient changes in activity in specific neural circuits influence the moment-to-moment structure of behavior; furthermore, given that the relationship between neural circuit activity and behavior can vary on a trial-to-trial basis, we wanted to probe the ability of the AR-HMM to afford insight into the probabilistic relationships between neural circuit activity and behavior.

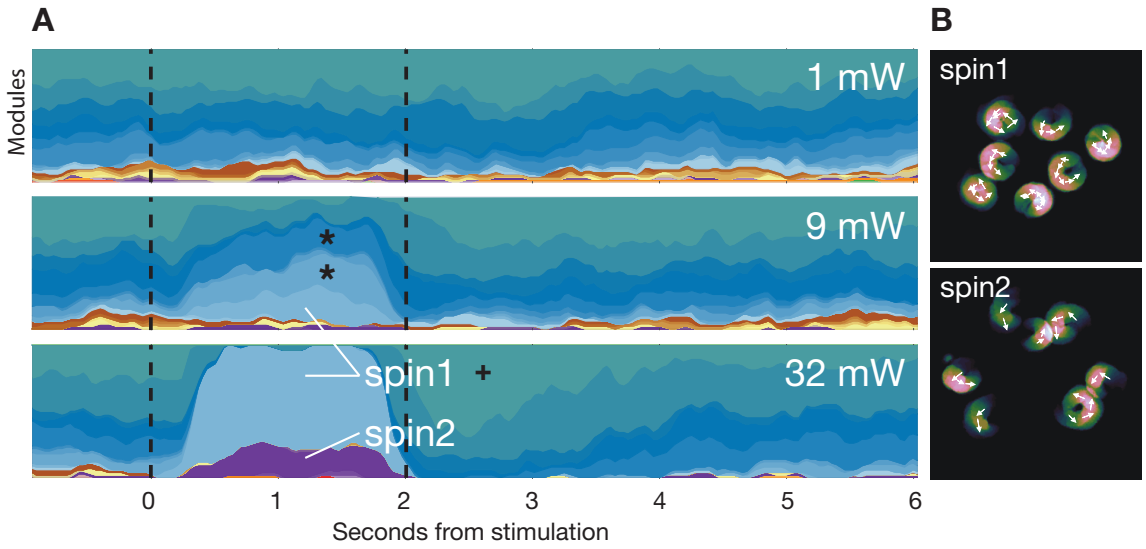


Figure 2.12 | A. Mountain plot depicting the probability of expression of each behavioral module (assigned a unique color on the Y axis) over time (X axis), with light stimulation indicated by dashed vertical lines (each plot is the average of 50 trials). Note that modest variations in the baseline pattern of behavior, due to trial structure, are captured before light onset. Stars indicate two modules expressed during baseline conditions that are also upregulated at intermediate powers (11 mW) but not high powers (32 mW, Wald test, Holm-Bonferroni adjusted $p < 10\text{-}5$); cross indicates pausing module upregulated at light offset (Wald test, Holm-Bonferroni adjusted, $p < 10\text{-}5$). B. Average position of example mice (with arrows indicating orientation over time) of the two modules induced under the highest stimulation conditions. Note that A and B are generated from one animal and that the observed modulations are representative of the complete dataset (n= 4 mice, model was trained separately from previous experiments).

To address these questions we unilaterally expressed the light-gated ion channel Channelrhodopsin-2 in a subset of layer 5 corticostriatal neurons in the right hemisphere and assessed behavioral responses before, during, and after two seconds of light-mediated activation of motor cortex (Glickfeld et al., 2013).

At negligible power levels no light-induced changes in behavior were observed, whereas at the highest power levels the AR-HMM identified two behavioral modules whose expression was reliably induced by the light, as on nearly every trial either one or the other module was expressed (Figure 2.12A). As would be expected, both of these modules encode forms of spinning-to-the-left behavior, and neither of these modules was

expressed during normal mouse locomotion (Figure 2.12B). In addition we noted that approximately 40 percent of the time the overall pattern of behavior did not return to baseline for several seconds after the end of optogenetic stimulation. This deviation from baseline was not due to continued expression of the same spinning modules that were triggered at light onset; instead, mice often paradoxically expressed a pausing module at light offset (average during-module velocity = .8 ± 7 mm/sec, Figure 2.12A, see cross).

The behavioral changes induced by high intensity optogenetic stimulation were reliable, as on nearly every trial the animal emitted one of the two spinning modules. To ask whether the AR-HMM could characterize the pattern of behavior observed when behavioral modules are generated more probabilistically, the levels of light stimulation were reduced; we identified conditions under which one of the two spinning modules was no longer detected, and the other was expressed in only 25 percent of trials. Under these conditions the AR-HMM detected the upregulation of a distinct set of behavioral modules, each of which was expressed in a fraction of trials (Figure 2.12A, see stars). Unlike the spinning modules triggered by high intensity stimulation, these modules were not neomorphic; rather, these modules were normally expressed during physiological exploration, and encoded distinct forms of forward locomotion behavior (data not shown). Interestingly, although each of these individual light-regulated modules was emitted probabilistically on any given trial, low intensity neural activation reliably influenced behavior across all trials when the behavioral modules were considered in aggregate (Figure 2.12A).

Thus both low and high intensity optogenetic stimulation dramatically alters behavior, but the identity of these behaviors and their relative probabilities of expression vary across light levels. This effect would not have been apparent without the ability of the AR-HMM to distinguish new from previously-expressed behavioral modules, and to identify the specific behavioral module induced at light onset on each trial; indeed under low light stimulation the behavioral phenotype apparent by eye on any given trial often appears to be an extension of normal mouse exploratory behavior. Taken together, these data demonstrate that the AR-HMM can identify and characterize both obvious and subtle optogenetically-induced phenotypes, distinguish ”new” optogenetically-induced behaviors from upregulated expression of ”old” behaviors, and reveal the trial-by-trial relationships between neural activity and action.

2.3 Discussion

It has been long hypothesized that innate behaviors are composed of stereotyped modules, and that specific sequences of these modules encode coherent and adaptive patterns of action (Bizzi et al., 2000; Brown, 1911; Drai et al., 2000; Lashley, 1951; Sherrington, 1916; Tinbergen, 1951). However, most efforts to explore the underlying structure of mouse behavior have relied on ad hoc definitions of what constitutes a behavioral module, and have focused on specific behaviors rather than systematically considering behavior as a whole. As a consequence, we lack insight into the global organization of mouse behavior, into the relationships between currently expressed actions and past or future

behaviors, and into the strategies used by the brain to generate behavioral change. Furthermore, we lack a comprehensive framework for characterizing the influence of individual genes or neural circuits on behavior. Here we use 3D imaging to identify a sub-second spatiotemporal scale at which mouse behavior may be organized. Using this finding as a constraint, we then built a family of computational models, each of which represents a different hypothesis for the potential structure of behavior, and compared the ability of these models to predict and explain mouse behavior. The best performing model (the AR-HMM) specifically searches for modularity at sub-second timescales similar to those observed in the pose dynamics data, and quantitatively describes behavior as a series of sub-second modules with defined transition probabilities. This combined 3D imaging/modeling approach can be used to automatically identify the behavioral modules expressed during a variety of experiments and to systematically discover how the architecture of behavior is altered as the mouse adapts to a changing world.

2.3.1 The AR-HMM Automatically and Systematically Captures Known and New Phenotypes

Our experiments reveal that the AR-HMM identifies both predicted changes in action and new features of behavior that had not been previously described. The high sensitivity of the AR-HMM – illustrated by the identification of a previously-undescribed phenotype in heterozygous Ror1β mice – raises the possibility that the AR-HMM could be useful for extracting subtle phenotypes from mouse models, including those in which single gene alleles are mutated; such patterns of mutation are common in human disease, but rarely

explored in mouse models. Furthermore, the ability of the AR-HMM to place changes in behavior into context – illustrated by the ability of the AR-HMM to identify the modules expressed after modest optogenetic stimulation on a trial-by-trial basis – suggests that the AR-HMM may be useful for relating unreliable or noisy patterns of neural activity to the probabilistic expression of specific actions. The AR-HMM (or similar approaches based upon unsupervised machine learning) may therefore be useful for discovering the behavioral consequences of a wide variety of experimental manipulations, particularly those enabled by the ever-growing toolbox of gene editing and optogenetic techniques.

2.3.2 The AR-HMM Identifies Possible Mechanisms for Behavioral Change

The AR-HMM suggests three regulatory strategies that may be used by the brain to implement behavioral adaptation. First, behavioral modules and their transitions appear to be selectively – and to some extent independently – vulnerable to alteration. None of the physiological and pathological deviations from baseline described here, from freezing to a waddling gait, caused global changes in the underlying structure of action; instead, new behaviors were well-described as changes in a small number of specific modules or transitions. This suggests that the brain can focally manipulate individual modules or transitions to generate new behaviors, and furthermore that mice can accommodate pathological actions (such as waddling) without catastrophic alterations in behavioral patterning.

Second, dramatically new behaviors can be created by altering the transitions statistics between modules alone – without invoking new behavioral modules – as was observed in

animals exposed to TMT. This strategy has been shown to underlie sensory-driven negative and positive taxis behaviors in bacteria, worms and flies (Berg et al., 1972; Garrity et al., 2010; Gray et al., 2005; Pierce-Shimomura et al., 1999); here we show that this strategy is conserved in mice, and used by vertebrate nervous systems to create complex patterns of action in response to an external cue. In the specific case of the TMT response, the induced behaviors extend beyond taxis to adaptations like freezing, suggesting that the restructuring of transition probabilities may be a general mechanism for creating new patterns of action.

Third, we find that modulation of where in space behavioral modules are expressed supports the generation of specific adaptive behaviors. For example, the rearing module used by mice to investigate TMT is not significantly upregulated or resequenced relative to control, and yet its spatial pattern of expression in the quadrant containing TMT facilitates detection (and therefore avoidance) of the aversive odorant. Characterizing where in space behavioral modules are expressed also reveals that changes in physical context – such as the difference between a square and circular arena – elicit "state"-like changes in mouse behavior that extend beyond predictable changes in action at apparatus boundaries. Because the training data for the AR-HMM does not include any explicit allocentric parameters (such as the mouse's spatial position within the apparatus), the ability of the model to uncover meaningful relationships between allocentric space and egocentric pose dynamics is an important validation that its segmentation of behavior is informative.

2.3.3 Potential Neural Underpinnings for Modules and Transitions

The observations generated by the AR-HMM lead to several predictions about the neural control of behavior. One such prediction is that specific behavioral modules and their associated transition probabilities will have explicit neural correlates, whose pattern of activity should reflect the 2-5 Hz timescale at which modules are expressed; it is tempting to speculate that neural correlates representing transition probabilities between behavioral modules will be encoded in higher-order neural circuits tasked with behavioral sequencing, while neural correlates for the behavioral modules themselves might be encoded in central pattern generators or related circuit motifs in the brainstem or spinal cord. The relevant neural circuits may include both evolutionarily-ancient regions of the brain involved in releasing innate behaviors, such as the amygdala, hypothalamus and brainstem, as well as other areas, such as the striatum, that regulate fine-timescale behavioral sequencing (Aldridge et al., 2004; Swanson, 2000). Testing these predictions will require simultaneous characterization of neural activity and assessment of behavior; future embellishments of the modeling approach described here may allow for inference of joint structure between dense neural and behavioral data, and therefore be useful for revealing mechanistic relationships between the dynamics of neural activity and action.

2.3.4 Unsupervised Behavioral Characterization via Modeling of 3D Data: Strengths, Caveats and Future Directions

Although there are multiple possible approaches to acquiring 3D pose dynamics data, we chose to implement our imaging system using a single standard depth camera because such cameras are widely available, adaptable to a variety of different experimental circumstances, and can characterize behavior under most lighting conditions and in animals with any coat color (due to the use of infrared light as an illumination source). The "unsupervised" modeling approach taken here is also transparent, insofar as the assumptions of the model are explicitly stated (with no constraints supplied by the researcher other than the structure of the model, and the specification of a single parameter that acts as a tunable "lens" to focus the model on behavior at a particular timescale). This makes clear the precise bounds of human influence on the output of the AR-HMM, and insulates key aspects of that output from the vagaries of human perception and intuition. Given that inter-observer reliability in scoring even single mouse behaviors can be low (from 50 to 70 percent, with reliability falling as the number of scored behaviors increases), developing methods free from observer bias is essential for informatively characterizing mouse behavior (Garcia et al., 2010).

In addition, the generative modeling and inferential fitting methods described herein offer several practical advantages over the approaches typically used to analyze mouse behavior (Crawley, 2003), including: the explicit time-series modeling of behavioral data (as opposed to simple clustering of dimensionally-reduced data); the ability to directly

inspect and explore each behavioral module; the flexible discovery of previously-unobserved behavioral modules (rather than characterization of behavior from the perspective of "known" phenotypes); and the ability to generate synthetic behavioral data, thereby allowing quantitation of how well a given model predicts the structure of behavior. This quantitative framework for comparing the performance of different methods for dividing up and measuring behavior is critical for advancing behavioral neuroscience, as it enables objective evaluation of alternative models for behavior, and assessment of future extensions that incorporate inevitable improvements in camera resolution, model structure, and fitting procedures.

On the other hand, our conclusions regarding the underlying structure of behavior (including its timescale) are limited by the simplicity of our experimental manipulations, which were designed to expose differences in motor outputs. Furthermore "mouse behavior" as described by the model is restricted to the imaged pose dynamics of the mouse at a particular spatiotemporal resolution and within a controlled laboratory experiment. Because the AR-HMM directly models the pixel data (after pre-processing), comparisons can only be made between mice of roughly similar size and shape. Furthermore, there are clearly important physical features of mouse behavior (operating at a variety of spatiotemporal scales) not captured in the pose data and therefore not modeled – these range from individual joint dynamics and paw position to sniffing, whisking and breathing. In the future, complementary datastreams that capture different facets of behavior could be integrated with 3D pose data to generate more comprehensive

behavioral models.

In addition, the modeling approach itself has several important limitations. The AR-HMM uses Bayesian nonparametric approaches to identify the most likely number of modules that describe behavior at a particular temporal scale. However, this insight also comes at a cost: as the amount of data fed to the algorithm increases, the number of discovered modules necessarily rises. This challenge parallels the well-described phenomenon in ethology in which the number of discovered behaviors increases in subjects that have been observed either more frequently or for longer (Colgan, 1978). This monotonic (although sublinear) relationship between data size and the number of discovered modules limits comparisons of behavior across experiments without co-training models, as was done here. One potential approach to address this challenge could be the incorporation of a "canonical" behavioral dataset against which additional data could be compared; such a framework (in which the "canonical" modules are either frozen or flexible) may enable analysis of new behavioral experiments within a fixed frame of reference.

Finally, the AR-HMM cannot explicitly disambiguate those features of behavior that are the consequence of the biomechanics of the mouse – for example, transitions between specific modules that are impossible due to physical constraints – from those that are the consequence of the action of the nervous system. Given that the nervous system and the body in which it is embedded co-evolved to facilitate action, ultimately disentangling the relative contributions of each to the organization of behavior may be difficult (for

discussion of this issue, see for example (Tresch & Jarc, 2009)). However, changes in the structure of behavior that are induced by experimental intervention arise principally from the action of the nervous system; the observed context-dependent flexibility of the transition statistics between modules, taken with the ability of the mouse to emit new behavioral modules in response to internal or external cues, together suggest that – at the spatiotemporal scale captured by our methods – the nervous system plays a key role in regulating the overall structure of behavior.

2.3.5 Mouse Body Language: Syllables and Grammar

Candidate behavioral modules have been recognized in a variety of different contexts and on a wide range of spatiotemporal scales, and accordingly researchers have given them a diverse set of names, including motor primitives, behavioral motifs, motor synergies, prototypes and movemes (Anderson & Perona, 2014; Flash & Hochner, 2005; Tresch et al., 1999). The behavioral modules identified by the AR-HMM here find their origin in switching dynamics that are expressed on timescales of 100s of milliseconds. Mouse behavior is clearly also organized at the varied and interdigitated timescales at which internal state (e.g. neural activity, endocrine function and development) and external state (e.g. daily, monthly, seasonal and annual variation in behavior) unfold. The behavioral modules we have characterized therefore likely exist at an intermediate hierarchical level within the overall structure of behavior, albeit one that captures many of the behavioral changes induced by experimental manipulations carried out at both short and long timescales.

The modules identified by the AR-HMM do not exist in isolation; instead they are given behavioral meaning through a transition structure that governs their sequencing. The observation of both modularity and transition structure within the pose dynamics of mice suggests strong analogies to birdsong, which is also hierarchically organized and composed of identifiable modules whose sequence is governed by definable transition statistics; importantly, birdsong has also been well described using generative modeling and low-order Markov processes (Berwick et al., 2011; Markowitz et al., 2013). By analogy to birdsong we therefore propose to refer to the modules we have identified as behavioral "syllables," and the statistical interconnections between these syllables as behavioral "grammar." Such a grammatical structure has been previously proposed for restricted subsets of mouse behavior (such as grooming) in which behavioral modules were defined on an ad hoc basis (Berridge et al., 1987; Fentress & Stilwell, 1973); through unsupervised identification of behavioral syllables here we show that the notion of a regulatory grammar is general, and can be used to explain a wide variety of behavioral phenotypes. As is true for birdsong, the grammar we describe is highly restricted in nature (as only low-order interactions are modeled) and lacks the richness and flexibility of context-dependent grammars that have been explored in human language (Berwick et al., 2011). Despite this limitation, the experiments described herein expose an underlying structure for mouse body language organized at the sub-second timescale; this structure encapsulates mouse behavior (as detected by a depth sensor) within a given experiment, and reveals a balance between stochasticity and determinism that is dynamically

modulated as the mouse varies its pattern of action to adapt to challenges in the environment.

2.4 Acknowledgments

We thank Michael Greenberg, Vanessa Ruta, David Ginty, Jesse Gray, Mike Springer, Aravi Samuel, Chris Harvey, Bernardo Sabatini, Rachel Wilson, Andreas Schaefer and members of the Datta Lab for helpful comments on the manuscript, Joseph Bell for useful conversations, and Alexandra Nowlan and Christine Ashton for laboratory assistance. We thank David Roberson and Clifford Woolf for assistance with behavioral experiments, and Ofer Mazor and Pavel Gorelik from the Research Instrumentation Core Facility for engineering support. We thank Sigrid Knemeyer for illustration assistance. Core facility support is provided by NIH grant P30 HD18655. AW is supported by an NSF Graduate Research Fellowship and is a Stuart H.Q. & Victoria Quan Fellow. GI is supported by the Human Frontiers Science Program. MJJ is supported by a fellowship from the Harvard/MIT Joint Grants program. SP is supported by an NSF Graduate Research Fellowship and is a Stuart H.Q. & Victoria Quan Fellow. RPA is supported by NSF IIS-1421780. SRD is supported by fellowships from the Burroughs Wellcome Fund, Searle Scholars Program, the Vallee Foundation, the McKnight Foundation, the Khodadad Program, by grants DP2OD007109 and RO11DC011558 from the National Institutes of Health, and by the SFARI program and the Global Brain Initiative from the Simons Foundation.

2.5 Author contributions

ABW and SRD conceived of the project, designed the experiments and analyses. Matthew Johnson (MJJ) created the computational model, ABW and MJJ built the computational modeling pipeline. ABW processed the data and ABW and MJJ ran the analyses. ABW, Ralph Peterson (REP) and Jesse Katon (JMK) collected the data. Stan Pashkovski (SLP) and Victoria Abraira (VEA) contributed reagents and experimental methodologies. Giuliano Iurili (GI) and REP ran the optogenetics experiments. Ryan Adams (RPA) supervised the creation of the computational model.

Chapter 3

Computer Vision and Machine Learning Methods for the Segmentation of Freely-Moving Mouse Behavior

3.1 Methods Summary

Experiments, unless otherwise noted, were performed on C57/BL6 male mice, aged 6-8 weeks, on the dark cycle of a reverse 12 hour light / 12 hour dark cycle. All mice were videotaped with a Microsoft Kinect depth camera and custom recording software. Recorded depth video was subjected to custom algorithms to extract, parallax correct and align the 3D image of the mouse, and then the time-series data were dimensionally compressed using a wavelet transformation and PCA before being subjected to both model-free and model-based analysis. Autocorrelation analysis was performed using correlation distance between time-lagged data points to correct for sensor-specific noise. Spectral analysis was carried out using the Welch periodogram method, and approximate changepoints were identified using a filtered derivative algorithm. Template matching was achieved using a minimal Euclidean distance criterion.

Established Bayesian time-series modeling methods were modified to identify

behavioral modules. The core AR-HMM was built from three components: 1) an observation distribution, describing the fast-timescale pose dynamics of each behavioral module; 2) a duration distribution, to allow variability in the length of emission of each behavioral module; and 3) a transition distribution (when relevant), which summarizes the interconnectivity over time of identified behavioral modules. Each behavioral module has its own distinct pose dynamics, its own characteristic duration distribution, and its own transition structure in relation to all other behavioral modules. Models were fit using Markov Chain Monte Carlo methods such as Gibbs sampling. The quality of fit was evaluated by testing the model on previously held-out data, which allows direct comparison between alternative model constructions.

Statistical significance of differences in behavioral module usage and transition usage were evaluated by computing both bootstrap samples over the data (on a per-mouse basis) and Bayesian credible intervals. Group comparisons were performed using the Hotelling t-squared statistic, single hypothesis tests using the Wald test. Both tests were adjusted with the Holm-Bonferroni step-down procedure to correct for multiple comparisons.

3.2 Methods

3.2.1 Mouse Strains, Housing and Habituation

Unless otherwise noted, all experiments were performed on 6-8 week old C57/BL6 males (Jackson Laboratories). Mice from the Ror1β and Rbp4-Cre strains (Jackson Laboratories) were habituated and tested identically to the reference C57/BL6 mice. Mice

were brought into our colony at 5 weeks of age, where they were group-housed for one week in a reverse 12 hours light/12 hours dark cycle. On the day of testing, mice were brought into the laboratory in a light-tight container, where they were habituated in darkness for 20 minutes before testing.

3.3 Behavioral Assays

3.3.0.1 Innate Exploration

For the open field assay (OFA), mice were habituated as noted above and then placed in the middle of a circular 17 inch diameter enclosure with 15 inch-high walls (US Plastics), immediately after which 3D video recording was begun. Mice were allowed to freely explore the enclosure for the 20 minute experimental period. Mice whose behavior was assessed in a square box were handled identically to the OFA.

3.3.0.2 Stimulus-Driven Innate Behaviors

To assess innate behavioral responses to volatile odorants, an odor delivery system was developed that spatially isolates odors in specific quadrants of a square box. Each 12 inch x 12 inch box was constructed of 0.25 inch black matte acrylic (Altech Plastics), with 3/4 inch holes patterning the bottom of the box in a cross formation, and a 1/16 in thick glass cover (Tru Vue). These holes were tapped and connected via PTFE tubing to a vacuum manifold (Sigma Aldrich) that provides negative pressure to isolate odors within quadrants. Odor was injected into the box through 1/2 inch NPT-3/8 inch pipe-fittings (Cole-Parmer). Filtered

air (1.0 L/min) was blown over odorant-soaked blotting paper (VWR) placed at the bottom of Vacutainer syringe vials (Covidien). The odorized airstream was then passed through corrugated PTFE tubing (Zeus) into one of the four pipe-fittings in a corner of the odor box. The ability of the odor box to isolate odors within specified quadrants was verified by visualizing vaporized odor or smoke through sheet illumination of the box with a low-power handheld HeNe laser. This approach allowed us to tune the vacuum flow and odor flow rates to achieve odor isolation, which was verified using a photoionization device (Aurora Scientific). To eliminate the possibility of cross contamination between experiments, the odor boxes were soaked in a 1% Alconox solution overnight, then thoroughly cleaned with a 70% ethanol solution. Mice were habituated to the experimental room for 30 minutes before the initiation of the experiment. Under control conditions, dipropylene glycol with air (1.0 L/min) was delivered to each of the four corners of the apparatus before a single mouse was placed in the center of the box and allowed to freely explore while 3D video records were acquired for 20 minutes. The same cohort of animals was tested for odor responses by subsequently repeating the experiment with odorized air delivered to one of the four quadrants. All 3D video recordings were performed in total darkness. TMT (Pherotech) was used at 5% concentration.

3.3.0.3 Optogenetics

Four adult male Rbp4-Cre (The Jackson Laboratory) mice were anesthetized with 1.5% isoflurane and placed in a stereotaxic frame (Leica). Microinjection pipettes (O.D. 10-15 μm) were inserted into the left motor cortex (coordinates from Bregma: 0.5 AP, -1 ML,

0.60 DV). 0.5 µl of AAV5.EF1a.DIO.hChR2(H134R)-eYFP.WPRE.hGH (1012 infectious units/mL, Penn Vector Core) was injected in each mouse over 10 minutes followed by an additional 10 minutes to allow diffusion of viral particles away from the injection site. After the injection, a bare optical fiber with a zirconia ferrule (O.D. 200 µm, 0.37 numerical aperture) was inserted 100 µm above the injection site and secured to the skull with acrylic cement (Lang). Twenty-eight days following the viral injection, mice were placed in a circular arena and the optical implant was coupled to a laser pump (488 nm, CrystaLaser) via a patch-chord and a rotary joint (Doric Lenses). The laser was directly controlled from a PC. After allowing the mouse 20 minutes of familiarization to the arena, optogenetic stimuli were delivered. The laser power, the pulse width, the inter-pulse interval and the inter-train interval were controlled by custom-made software (NI Labview). Each train of laser pulses consisted of 30 pulses (pulse width: 50 ms) at 15 Hz. The interval between successive trains was set to 18 seconds. 50 trains were delivered for each laser intensity. The animal was progressively exposed to higher laser intensities over the course of the experiment.

3.3.1 Camera Setup and Initialization

Mice were tracked in 3D using a Kinect for Windows (Microsoft). A boom tripod (Manfrotto) was used to suspend the camera above the recording arena, affording a stable top-down view of the mouse. The Kinect has a minimum working distance (in Near Mode) of 0.5 meters; by quantitating the number of missing depth pixels within an imaged field, the optimal sensor position was found to be between 0.6 and 0.75 meters depending on ambient light conditions and assay material. Thermal effects can cause degradation in

the signal coming off the sensor; after each recording session, the camera was therefore unplugged from its power source to allow the sensor to deactivate and cool. The Kinect sensor has a small indicator LED that was masked using electrical tape so as to not emit spurious illumination.

3.3.2 Data Acquisition

Data from the Kinect was sent to an acquisition computer (hand-assembled, 16GB RAM, Intel i7, 512GB SSD) via USB. Custom Matlab software was used to interface the Kinect via the official Microsoft .NET API that retrieves depth frames at a rate of 30 frames per second and saves them in raw binary format (16-bit signed integers), along with the timestamp of each frame, in milliseconds, to an external hard-drive. After acquisition, we indicate a region-of-interest (ROI) delimiting the boundaries of the experimental arena using Matlab GUI tools, and save the polygon specifying the ROI containing the arena.

3.3.3 Data Preprocessing

All data analysis performed with Python scientific computing tools and OpenCV for image analysis (Pedregosa et al., 2011; Pérez & Granger, 2007; Van Der Walt et al., 2011) on r3.8xlarge EC2 virtual machines running an AMI customized from a base Ubuntu 14.04 Linux AMI. Tracking the evolution of an imaged mouse's pose over time requires identifying the mouse within a given video sequence, segmenting the mouse from the background (in this case the apparatus the mouse is exploring), orienting the isolated image of the mouse along the axis of its spine, correcting the image for perspective

distortions, and then compressing the image for processing by the model.

To isolate our analysis to the experimental arena in which the mouse is behaving, a region-of-interest (ROI) was identified for further analysis that was manually traced along the outside edge of any imaged arena; pixels outside this ROI were by default set to zero to prevent spurious object detection. The raw imaging data were filtered with an iterative median filter (Arias-Castro & Donoho, 2009), which is well suited to removing correlated noise from the sensor within the Kinect. Filter settings were identified that maximized between-pose discriminability while minimizing within-pose differences by using dead mice in rigor mortis that were placed into distinct poses (Figure 3.1A); based upon this analysis, a three pixel median filter over space, and then a three-then-five frame iterative median filter over time was applied to the imaging data.

To subtract the background image of the arena, the median value of the first 30 seconds of data from any imaging stream were taken, and subtracted from all video frames; any spurious values less than zero were by default reset to zero. The image was then binarized and any objects that did not survive three iterations of morphological opening were eliminated. The mouse was defined as the largest object within the arena that survived the subtraction and masking procedures. The centroid of the mouse was identified as the center-of-mass of the preprocessed image; an ellipse was fit to its contour to detect its overall orientation. In order to orient the mouse along the axis of the spine, a random forest classifier was trained on a set of manually oriented mouse images to facilitate automated identification of the head and tail. Although this algorithm was nearly always

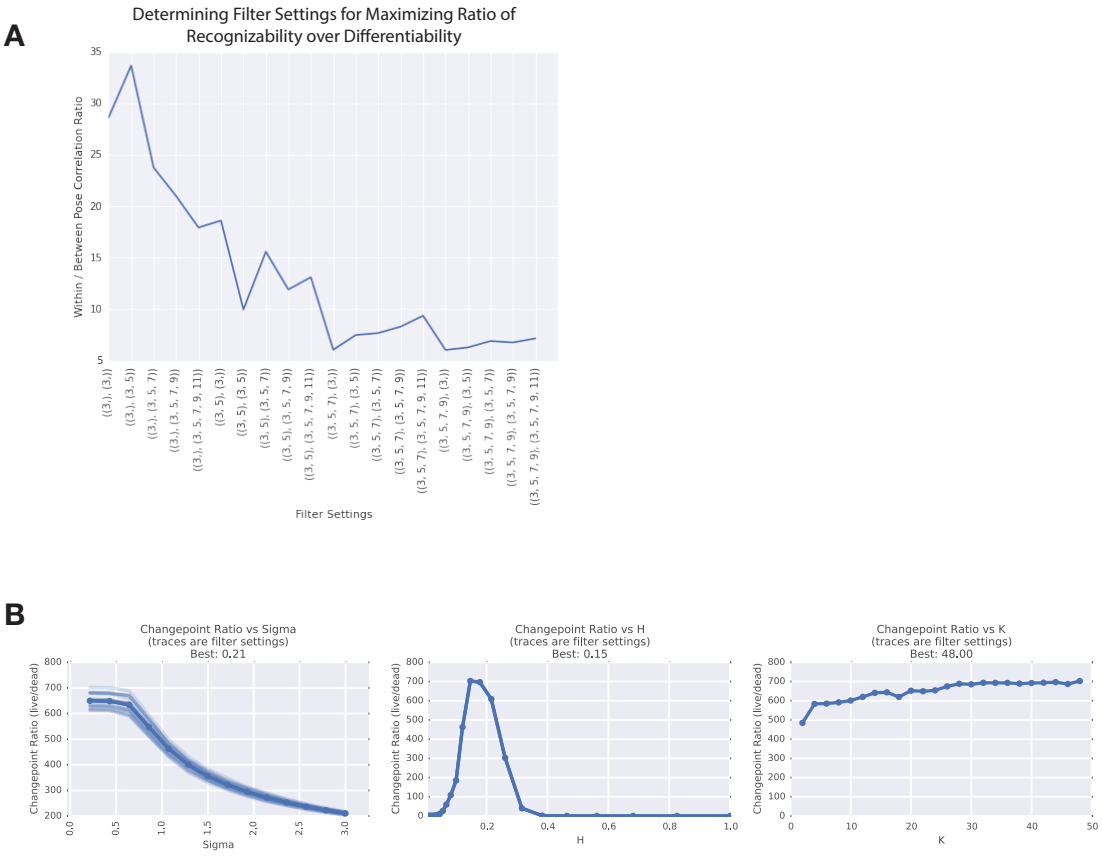


Figure 3.1 | A. To denoise data from the Kinect a filtering approach was used in which a median filter was applied iteratively both in space and in time; this approach has been shown to effectively maintain data structure while smoothing away noise (see Supplemental Experimental Procedures). To determine filter settings, we imaged dead mice that were differentially posed in rigor mortis; ideal filter settings would distinguish mice that were posed differently, but be unable to distinguish data from the same mouse. Filter settings are indicated as ((pixels), (frames)) with the numbers within each parenthesis referring to the iterative settings for each round of filtering. To assess filter performance, a within/between pose correlation ratio (Y axis) was computed, in which the mean spatial correlation for all frames of the same pose was divided by the mean spatial correlation for all frames of different poses. This revealed that light filtering (with settings ((3), (3,5))) optimized discriminability in the data. B. Optimal parameters for changepoints analysis were identified by grid scanning. By maximizing the number of changepoints identified in a live mouse, while fixing the number in a dead mouse to be zero, clear optimal values were identified via grid scanning for sigma and H (left two panels). This changepoint ratio was not highly sensitive to K; a setting of 48 (at the observed maximum) was therefore chosen.

correct, its output was manually supervised to ensure accuracy. Note that this procedure does not directly identify the spine, but that the ellipse-based alignment procedure has the effect of aligning the animal along this anatomic axis. During episodes when the mouse was reared directly upwards (and therefore appears as a circle against the sensor, which is orthogonal to the axis of the mouse's spine under these conditions) neither a human nor our algorithm can identify the spinal axis; because we model pose dynamics based upon dimensionally-reduced representations of the pixels (rather than extracted scalar features, like orientation or length), the presence of this symmetry does not adversely affect model performance, and synthetic data generated from the model appropriately captures these rearing episodes without "flipping inch the mouse along its spinal axis.

The position of the mouse within the arena enabled extraction of allocentric data about the mouse, including the centroid, head and tail positions of the mouse, orientation, length, width, height, and each of their first derivatives with respect to time. Critically, none of these scalar parameters are fed to the model; we obtain these data for the sake of convenience and to understand how the egocentric behaviors captured by the model might relate to allocentric parameters like position or velocity.

To correct perspective distortion in the X and Y axes, a tuple of (x,y,z) coordinates was generated for each pixel in real-world coordinates, and then those coordinates were resampled to fall on an even grid in the (x,y) plane using Delaunay triangulation. This procedure corrected for the slight differences in the appearance of the mouse when at the left or right edges of the arena; these differences derive from parallax effects caused by

imaging the animal from a single fixed vantage point from above, but are normalized by the interpolation procedure. All extracted images and scalars were saved in HDF5 format. The extraction pre-processing which transforms raw data into the extracted HDF5 format runs at approximately real-time (e.g. 20 minutes of preprocessing is required for a 20 minute recording session) on a 32-core r3.8xlarge EC2 instance.

The full extracted image size of a mouse covered an area of 120 mm2, where each pixel is 2 mm on a side; each frame to be submitted for analysis contained 3,600 16-bit integers (which were cast to floating point before any image analysis) in a 60x60 array. The information captured in this image was often either highly correlated (neighboring pixels) or uninformative (pixels on the border of the image that never contain the mouse's body). To both reduce redundant information and make modeling computationally tractable, each image was dimensionally reduced. A five-level wavelet decomposition was first applied, thereby transforming the image into a representation in which each dimension captured and pooled information at a single spatial scale; in this transformation, some dimensions coded explicitly for fine edges on the scale of a few millimeters, while others encoded broad changes over spatial scales of centimeters (Mallat, 1989). This wavelet decomposition expanded the dimensionality of the image; to reduce this dimensionality principal components analysis was applied to these vectors, and projected the wavelet coefficients into ten dimensions, which captures 88% of total variance (Figure 2.3) and which were submitted to our modeling algorithms. Principal components were built using a canonical dataset of 25 C57 BL/6 mice, aged 6 weeks, recorded for 20 minutes each in

the OFA, and all datasets were projected into this common pose space.

For some model-free depictions of behavioral data, the random projections technique was used to reduce the data dimensionality. This approach also has the benefit of representationally normalizing the data to changes in the size of the mouse on the sensor. For example, when a mouse rears up the number of pixels that describe the mouse changes drastically, leading to obvious striations in the raw pixel data over time that may visually overstate the overall change in the behavior of the mouse. Because, as described below, each random projection dimension takes into account all of the pixels within the 60x60 array, representations of the random projections data are not dominated by these effects, enabling more holistic inspection of changes in the data over time. Random projections is an approach that produces new dimensions derived from an original signal, with dimensionality D_{orig}, by randomly weighting each original dimension, and then summing each dimension according to that weighting, producing a single number per data point. This procedure is repeated several times, with new random weightings, to produce a set of "randomly projected" dimensions. The Johnson–Lindenstrauss lemma shows that distances between points in the original dataset with dimensionality D_{orig} are preserved in the randomly projected dimensions, D_{proj}, where $D_{proj} < D_{orig}$. As random projections is primarily employed as a visualization aid, we set $D_{proj} << D_{orig}$, where $D_{proj} = 300$, $D_{orig} = 3600$. For visualization purposes we do not sort the random projections dimensions in any way.

3.3.4 Model-Free Data Analysis Methods

Standard methods for computing autocorrelation were modified, as sensor-specific noise results in a monotonically decreasing autocorrelogram, even for a mouse that is posed in rigor mortis. Average correlation of all 10 dimensions of the mouse's pose data was therefore used as the comparator between time-lagged versions of the time-series signal in question, resulting in a nearly flat autocorrelation function of value 1.0 for a dead mouse, and a declining autocorrelation function for a behaving mouse. The recording of the dead mouse was performed after the animal had entered rigor mortis; the posed dead mouse was placed in the open field arena and recorded for 5 minutes. The rate at which this autocorrelogram declines in a behaving mouse may be characterized as a time-constant, tau, of an exponentially decaying curve. Tau was fit using the Levenberg-Marquardt algorithm (non-linear least squares) using the SciPy optimization package.

Power-spectral density (PSD) analysis was performed on behavioral data to further analyze its time domain structure (Welch, 1967) using a Wiener filter (Wiener, 1949), and implemented by taking the ratio of the PSD of a behaving mouse over the PSD of a dead mouse. The PSD was computed using the Welch periodogram method, which takes the average PSD over a sliding window across the entire signal; a 33-second window was used, sliding in 0.8 second increments.

Plotting the mouse depth data over time (as either raw data or as random projections) yields obvious striations, each a potential boundary between blocks or modules. To automate the identification of these approximate boundaries between blocks, a simple

changepoint identification technique was used, called the filtered derivative algorithm (Basseville et al., 1993). Briefly, the algorithm calculated the derivative of the per-frame unit-normalized random projections with a lag of k=4 frames; this signal was binarized using a threshold h=0.15 mm, summed across each of D=300 random projection dimensions, and then the resulting one-dimensional (1D) signal was smoothed with a Gaussian filter with a standard deviation of sigma=0.43 frames. Changepoints were identified as the local maxima of this smoothed 1D time-series. As specified above, this procedure depends in part upon the values of the parameters k, h and sigma; these were identified as those values that maximize the number of changepoints in the behaving mouse while yielding no change points in a dead mouse. Optimizing the difference between live and dead mice was largely insensitive to k, but sensitive to h and sigma; however, grid searching revealed clear optimal values for h and sigma, which are used in the analysis above (Figure 3.1B). Note that changepoints methods provide access to the overall timescale of behavior, but the specific boundaries between blocks identified by these approaches are highly approximate, as they only consider local data structure and ignore higher-order interactions between modules (including transition structure).

To ask whether any reasonably long snippet of behavior (greater than just a few frames) was ever ”repeated” (without reliance on a underlying model for behavior), a template matching procedure was devised to identify similar trajectories through PCA space; here similarity was defined both quantitatively (minimal Euclidean distance in pose PC space) and qualitatively (via inspection of spines and PC trajectories). To identify similar

trajectories, the Euclidean distance between some target snippet, the ”template”, and every possible snippet of equal length was computed. The most similar snippets were selected, ignoring snippets discovered that were shifted less than 1 second from each other (to ensure selection of behavioral snippets that occur distanced in time from each other).

3.3.5 Data Modeling and Fitting Methods

The modeling approach used herein takes preprocessed depth-camera video data (e.g., the 10 top PCs) as input and yields a fit model as output (Figure 2.5A), where the fit model includes estimates of several key parameters, including the number of behavioral modules observed within a given set of experimental data, the autoregressive parameters that describe the motion expressed by the mouse within a given module, and the transition parameters that describe how often any particular module precedes or follows any other module; furthermore, for each video frame an assignment is made to the most likely associated behavioral module. To model pre-processed data, we modified established and well-documented modeling and fitting approaches that have provided insight into complex and dense datasets with time-series structure (Fox et al., 2008, 2009). These approaches include Bayesian nonparametric methods, which allow identification of the most likely number of behavioral modules and lags in the AR process without human supervision. For convenience and clarity, we provide a summary of these methods and their specific application to the 3D behavioral data below, together with formal mathematical descriptions of our modeling and fitting procedures.

In brief, three-dimensional pose data were modeled using generative probabilistic

modeling, which is often used to model complex dynamical processes (Bishop, 2006a; Koller & Friedman, 2009; Murphy, 2012). By choosing probability models that reflect stylized versions of the physical dynamics that might give rise to the 3D pose dynamics data within a given experiment, this framework allowed the instantiation and testing of hypotheses about different organizations of behavior. These hypotheses ranged from simple progressions of poses described as Gaussians in pose space to the hierarchical AR-HMM that includes both autoregressive states and switching dynamics. Furthermore, using a generative framework for modeling behavior addresses a key limitation in alternative methods used for behavioral classification: distinct models for behavior can now be quantitatively and objectively compared based upon each model's ability to describe the statistics of held-out test data.

To test the hypothesis that behavior is composed of a series of behavioral modules that are stereotyped in form and which are interconnected in time with defined transition statistics, a family of discrete-time hidden Markov models (HMMs) were instantiated (Bishop, 2006a; Koller & Friedman, 2009; Murphy, 2012). Here, at each point in time (e.g., for every frame of imaging data), an HMM posits that the mouse is within a discrete state that can be given a label. Each hidden Markov state corresponds to a prototypical brief three-dimensional motif of motion the animal undertakes while within that state. The Markovian state of the dynamical process is then composed of both the latent discrete component, which identifies the behavioral mode of the animal, and a number of lagged values of the observation sequence, which are used by the autoregressive model to predict

the short-timescale behavior of the animal based on the behavioral mode. This model structure is often called a switching vector-autoregressive (SVAR) model or autoregressive HMM (AR-HMM) (Fox et al., 2009; Murphy, 2012). Different experimental conditions were allowed to share the same library of state-specific VAR dynamics but learned their own transition patterns as well as any unique VAR dynamical modes, allowing the model to reveal changes in the parameters due to changes in the experiment.

Compositionally altering model structure allowed objective comparisons between different underlying models for how mouse behavior is organized over time. Removing the discrete switching dynamics captured in the transition matrix and replacing them with a mixture model generated an alternative model in which the distribution over each discrete state does not depend on its previous state. This would be the case if animals had a set of behavioral modules from which to choose, and the likelihoods of expressing any given one of them did not depend on the order in which they appear. This simplification resulted in the autoregressive mixture model (AR-MM). Alternatively, replacing the conditionally autoregressive dynamics with simple state-specific Gaussian emissions resulted in a Gaussian-emission HMM (G-HMM); this model explored the hypothesis that each behavioral module is best described by a simple pose, rather than being a dynamical trajectory. Applying both simplifications yielded a Gaussian mixture model (G-MM), in which behavior is simply a sequence of poses over time in which the probability of expressing any given pose does not depend on the prior pose. Removing the switching dynamics yielded a pure autoregressive (AR) or linear dynamical system (LDS) model, in

which behavior was described as a trajectory through pose space without any reused discrete behavioral modules at all. By fitting each of these alternative models and comparing model performance, the model structure that best explains the behavioral dynamics captured by the 3D imaging was identified (in this case the parent AR-HMM); this model (of the alternatives) thus best reflected the underlying organization of action (Figure 2.5).

To fit models we utilized algorithms that have been well validated in a variety of contexts, including specifically for fitting models of the class used herein (Bishop, 2006; Koller and Friedman, 2009; Murphy, 2012 Robert, 2013, r14413). In brief, to estimate the parameters of any given model, approximate Bayesian inference was performed using Gibbs sampling, a Markov Chain Monte Carlo (MCMC) inference algorithm (Robert and Casella, 2013). The Gibbs sampling algorithm employed here has a natural alternating structure (Figure 2.5C), directly analogous to the alternating structure of the popular expectation-maximization (EM) algorithm. Applied to the AR-HMM, the algorithm first segments the imaging data into modules given a fixed set of transition statistics and a fixed description of the AR parameters that describe any given module, and then the algorithm switches to fixing the segmentation and updating the transition matrix and the AR parameters. These two steps are alternated as an iterative process.

To employ Bayesian inference methods to fit the model, unknown quantities, including the transition matrix and the autoregressive parameters that describe each state, were treated with a uniform representation as latent random variables. In particular, weak prior

distributions (except for the sticky parameter) were placed on these quantities and their posterior distributions after conditioning on observed 3D imaging data were investigated. For the autoregressive parameters, we used Automatic Relevance Detection (ARD) algorithms, which incorporate a prior with a Lasso-like penalty to encourage uninformative lag indices to have their corresponding regression matrix coefficients tend to zero (Gelman et al., 2014). For the transition matrix, a Hierarchical Dirichlet Process (HDP) prior was used, which served to regularize the number of discrete latent states. In addition, the transition matrix prior also included a sticky bias, which is a single nonnegative number that controlled the tendency of the discrete states to self-transition (Fox et al., 2008, 2009). Because this parameter controlled the timescale of the inferred switching dynamics, this parameter was set such that the output of the model inference algorithms matches (as closely as possible) the model-free duration distribution determined by our changepoint analysis. In this sense the sticky bias acts as a kind of "lens inch which focuses the model on the organization of behavior at a specific temporal scale. It is important to note that despite the fact that this parameter was tuned — indeed this parameter defined our prior over the timescale of behavior — it was not clear *a priori* that the model could identify a duration distribution that was similar to the changepoints distribution (or that a duration distribution mode would match the specific sub-second timescale found in the model-free analysis) for any sticky parameter value.

3.3.6 Model Training and Parameter Tests

Datasets from the open-field, odor, and genetic manipulation experiments were modeled jointly in a single model to increase statistical power and facilitate comparisons between results. Because the neural implants associated with the optogenetics experiment modestly altered the physical profile of the animal, these data were modeled separately. In all experiments, the first 10 principal components for the wavelet decomposition coefficients of each frame of each imaged mouse were gathered. Data were then randomly subdivided and assigned either a ”train” or a ”test” label, with 70% of the data assigned to the training subset, and 30% assigned to the test subset. The mice labeled ”test” were held-out from the training process, and used to test generalization performance via measurement held-out likelihood. This approach allowed us to directly compare algorithms whose composition reflected different underlying structures for behavior.

Models were trained on data using the procedures described above; modeling was robust to both initialization settings and to parameter and hyperparameter settings (with the exception of the sticky bias parameter kappa, see below). Specifically, the number of lags used in our AR observation distributions and the number of used states in our transition matrix with an HDP prior was found to be robust to the particular hyperparameter settings on both. The hyperparameters of our sparsifying ARD prior were varied by several orders of magnitude, and held-out likelihood, the number of used lags, and the number of used states varied negligibly. The hyperparameters of our HDP prior were also varied by several orders of magnitude and again no change to the number of

used states or held-out likelihood was observed. All jointly-trained data shared observation distributions, but each treatment class was allowed its own transition matrix. Each model was updated through 1000 iterations of Gibbs sampling; upon the last iteration of Gibbs sampling the model output was saved; all further analysis was performed on this final update.

The ”stickiness” of the duration distribution of our behavioral modules — defined by the kappa setting of the model — influenced the average duration of behavioral modules discovered by the AR-HMM; this allowed us to control the temporal scale at which behavior was modeled. As discussed in the main text, autocorrelation, power spectral density, and the changepoint algorithm identified switching dynamics at a specific sub-second temporal scale (as encapsulated by the changepoints duration distribution and reflected by the spectrogram and autocorrelogram). The kappa stickiness parameter of the time-series model was therefore empirically set to best match the duration distribution discovered by changepoint detection. To find the kappa setting at which these distributions were best matched, the Kolmogorov-Smirnov distance was minimized between the inter-changepoint interval distribution and the posterior behavioral module duration distribution through a dense grid search. Using alternative distance metrics did not affect the results.

3.3.7 Evaluating the Log-likelihood of Held-out Test Data

A crucial feature of a useful model is its ability to ”generalize”, meaning its ability to usefully describe phenomena in similar data that it has not previously seen. This is

intimately related to the problem of over-fitting; if a model performs well on training data, but poorly on some new data, then the model has been over-fit. The higher the generalization performance, which was measured as the log-likelihood of held-out test data, the less the model has been over-fit. Model comparisons were performed by evaluating the log-likelihood of held-out test data under the AR-HMM and each member of the family of alternative models. For each Markov Chain Monte Carlo (MCMC) sample of the AR-HMM transition matrix and autoregressive parameters, the log likelihood of a held-out data sequence was computed by summing over every possible hidden state sequence the joint probability of the data sequence and that hidden state sequence (Bishop, 2006).

3.3.8 Predictive comparison metric to stratify model capacities

We use two metrics to compare model performance: the aggregate held-out likelihood score used to compare models in 2.5D, and a temporal horizon-based predictive comparison metric in Figure 2.6A. While the held-out likelihood score used in Figure 2.5D compares the models' overall prediction performance, the metric shown in Figure 2.6A shows separately the benefit of each component of the model's structure. There are two possible versions of the horizon-based prediction metric, both based on predictive likelihood but differing in what data is used to make the prediction: Given the data up to time t, score how well the model predicts what happens at time t+k (for various k>0 and averaging over t). Symbolically, this metric is

$$\log p(y_{t+k} \mid y_1, \ldots, y_t) \tag{3.1}$$

Given the data up to time t and an instant of data before time t+k for the short timescale model, score how well the model predicts what happens at time t+k. More precisely, an ”instant of data” is taken to be the amount of data used in the short-timescale AR model. Symbolically, this metric for the AR-HMM, AR-MM, and AR models is

$$\log p(y_{t+k} \mid y_1, \ldots, y_t, y_{t+k-3}, y_{t+k-2}, y_{t+k-1}) \tag{3.2}$$

where the short-timescale data is not included for the G-HMM, G-MM, or Gaussian models (i.e. the ”pose” models).

Metric #1 shows overall prediction capability, much like the comparison in 2.5D, and hence lumps together the benefit from the AR model component and that from the discrete hidden state model component. However, Metric #2 actually separates these effects out: the amount by which the AR-HMM beats the AR-MM is precisely due to the discrete Markov state, since both models use the same short-timescale information. We therefore use Metric #2 for Fig. 2. Metric #2 demonstrates that the AR-MM persistently beats the AR model, corresponding to the fact that it is a more expressive dynamics model for short timescales. Metric #2 also shows the G-HMM and G-MM (i.e. the ”pose models”) as persistently lower because they do not model the short timescale information. Finally, we note that Metric #2 can be computed exactly, while #1 requires approximation with Monte Carlo or importance

sampling methods.

3.3.9 Analysis of the Model Consistency

To evaluate model consistency the following procedure was adopted, which enabled comparisons between model fits. For any two fits, each containing a set of behavioral modules, the modules between the two fits were matched by the Frobenius norm between the AR matrices for each module, and then the Hungarian algorithm (also called Munkres matching) was used to generate a mapping that associated a module in the first fit with a single module in the second fit. The number of frames each fit assigned to the module was then plotted as a single point, where the x-axis represents frames assigned to that module in the first fit, and the y-axis represents frames assigned to that module in the second fit. We evaluate how well two fits matched as the goodness-of-fit r^2 value between these points and the $y = x$ line. Both algorithmic consistency, in which the data do not change, but the random seed that initializes the model is different, and data consistency, in which different bootstrapped datasets are used to train the model, were evaluated with this procedure. Results for algorithmic consistency are described in the main text (r^2 = 0.93 ± 0.03). To test data consistency, we randomly sampled a subset of n=24, 20, 10 and 5 mice from a dataset with 25 mice (20 minutes each in the OFA), and refit the AR-HMM model 15 times. For n=24: r^2 = 0.95 ± 0.02; n=20: r^2 = 0.93 ± 0.04; n=10: r^2 = 0.92 ± 0.04; n=5: r^2 = 0.88 ± 0.09.

3.3.10 Analysis on Model Output

Fit models included a number of states that were infrequently used (see Figure 2.7E). Because we wished to focus on those changes in behavior that are most significant, we concentrated our analysis on the set of modules that explain 95 percent of the total frames in our dataset (e.g., the top 51 modules in the dataset) and in the main text only describe modules within this group. Visual inspection reveals that many of the frames assigned to states above this threshold are corrupted by sensor-specific noise that is not eliminated by our pre-processing. Because the model attempts to cluster together data with similar structure, this thresholding procedure has a practical de-noising effect and mitigates ambiguity injected into the imaging by the depth camera itself. Future use of our approach with different depth cameras may require distinct pre-processing and thresholding methods, as alternative cameras will have unique noise characteristics.

To represent the grammatical relationship between behavioral syllables, two analyses were performed. First, the probabilities of transition from any given module to any other module (in a pairwise fashion) were computed; second, the bigram probability was computed, which is the probability that two syllables were found occurring one after the other (a "bigram" of modules), as a fraction of all observed bigrams. To calculate this value for each pair (i, j) of modules, first a square nxn matrix, A, was created where n is the number of total modules in the label sequence. The label sequences that were saved at the last iteration of Gibbs sampling were then scanned, incrementing the entry $A_{i,j}$ for every time syllable i directly preceded syllable j. At the end of the label sequence, each

entry was divided by the number of total bigrams observed.

In order to visually organize those modules that were specifically up-regulated or selectively expressed as a result of a manipulation, a selectivity index was assigned to each module. Specifically, where $p(condition)$ indicates the percent usage of a module in a condition, modules in the circular open field versus square box comparison were sorted by $(p(circle) - p(square)/(p(circle) + p(square))$. In the comparison between blank odor and fox odor (TMT), modules were sorted by $(p(tmt) - p(blank))/(p(tmt) + p(blank))$. In the comparison of Ror1β mutant, heterozygous mutant, wildtype and C57/BL6, modules were sorted by $(p(het) + p(mut))/(p(wt) + p(c57))$.

3.3.11 Statemap Visualizations

To lay out each graph in a reproducible manner (up to global rotation of the figure), the position of the nodes was initialized using the spectral layout algorithm and fine-tuned node positions using the Fructherman-Reingold iterative force-directed layout algorithm (Koren, 2005); both algorithms were used as implemented in the NetworkX and GraphViz software package. To simplify the visual representation, only edges for which the bigram value was more than some multiple of bootstrap-estimated standard error units (indicated in figure legends) above zero were drawn.

3.3.12 Hinton Diagram

To succinctly display all bigram probabilities between all syllables, the full bigram probability matrix was displayed using a Hinton diagram. In this representation, the value

at each index of a rectangular matrix is shown as a 2D geometric object (such as a square or circle) at the spatial index of the matrix's element, where the area of the geometric element is proportional to the value of the matrix at that index. This has the property of representing larger values as more visually salient than smaller values. Additionally, we encoded additional properties of the matrix, such as the statistical significance of a particular bigram, using color of each geometric element. Because the indices of syllables in our AR-HMM model are chosen randomly by our model, we resorted the rows and columns of the bigram probability matrix using spectral clustering to better represent syllables which tend to transition amongst themselves. This re-sorting did not affect the data itself, only its presentation.

3.3.13 Computing Entropy Rate and Mutual Information

Let A denote the estimated transition matrix, so that $A_{ij} = \mathbb{P}[x_{t+1} = j \,|\, x_t = i]$ for $i, j = 1, 2, \ldots, n$ and all time indices t, and let π denote the steady-state distribution of the Markov chain, so that $\pi_i = \mathbb{P}[x_t = i]$. The entropy rate of the transition matrix was calculated as $-\sum_{i,j} \pi_i A_{ij} \ln A_{ij}$. The mutual information between states was calculated as $-\sum_i \pi_i \ln \pi_i + \sum_{i,j} \pi_i A_{ij} \ln A_{ij}$.

3.3.14 Cross-likelihood Analysis of AR-HMM States

To evaluate whether discovered AR-HMM states represent meaningful clusters of data in the underlying dataset, we measured how well, on average, each state's autoregressive parameters explain its assigned data segments compared to how well, on average, they

explain data segments assigned to other states. This cross-likelihood measure can also be viewed in terms of likelihood ratio testing: given a data segment to assign to one of two possible states, if that data segment was assigned according to a likelihood ratio test, the corresponding cross-likelihood values measure how far on average such data segments are from the likelihood ratio test decision boundary. Cross-likelihoods near or above 1 (i.e., when their logarithms are near or above 0) indicate that the corresponding states' autoregressive models can explain each others' data segments; the states' assigned data segments cannot be readily or reliably discriminated, at least using the states' autoregressive parameters and a likelihood ratio test. Conversely, likelihoods well below 1 indicate that the states' assigned data segments are easily discriminated.

For any two states $i, j \in \{1, 2, \ldots, n\}$ with assigned data segments

$$\{y_k^{(i)}\}_{k=1}^{N_i} \triangleq \{y_{t_a:t_b} \,:\, x_{t_a:t_b} = i, x_{t_a-1} \neq i \neq x_{t_b+1}\} \tag{3.3}$$

$$\{y_k^{(j)}\}_{k=1}^{N_j} \triangleq \{y_{t_a:t_b} \,:\, x_{t_a:t_b} = j, x_{t_a-1} \neq j \neq x_{t_b+1}\} \tag{3.4}$$

and autoregressive parameters $\theta^{(i)}$ and $\theta^{(j)}$, respectively, we computed the (i, j) entry of the log cross-likelihood matrix as

$$\mathrm{CL}_{ij} \triangleq \frac{1}{N_j} \sum_{k=1}^{N_j} \log \frac{p\left(y_k^{(j)} \mid \theta^{(i)}\right)}{p\left(y_k^{(j)} \mid \theta^{(j)}\right)}. \tag{3.5}$$

While these cross-likelihood ratios can be interpreted on an absolute scale in terms of classification of data segments using likelihood ratio testing, it is also useful to compare

this measure to a null model (i.e., a reasonable baseline hypothesis that should not have meaningful discrete state structure). This baseline was constructed by first fitting with Gibbs sampling a linear dynamical system (LDS) with 60-dimensional latent state to the same open field dataset on which AR-HMM was fit; while an LDS can capture the correlation structure in the data, both in time and across components, its only state is continuous and thus it by definition lacks modularity. Synthetic data were generated from the fit LDS; an AR-HMM was then fit to the synthetic data, thereby producing the corresponding cross-likelihood plot from the inferred states. In this specific case the AR-HMM hyperparameters were adjusted to enable more than one state to be inferred; otherwise, the Bayesian inference procedure only instantiates one autoregressive state to explain the LDS synthetic data, a result that is consistent with our non-parametric methods appropriately recognizing state number in the underlying data.

3.3.15 Visualizing Synthetic Mice

To generate synthetic mice from the AR-HMM and each model in the family of alternative models, the sticky kappa biasing term in each model was replaced with an explicit duration model so as to more accurately represent the timescales in each model's inferred label sequence. For each model fit, a corresponding autoregressive hidden semi-Markov model (AR-HSMM) was instantiated with the same autoregressive and transition parameters but with a negative binomial duration model for each state, which was also fit using Gibbs sampling. Each AR-HSMM generates a 10D principal coordinate sequence; using the principal components analysis fit to the original data, these are mapped into

images to produce the “dream mouse” videos.

3.3.16 Statistical Methods

To evaluate the model’s performance from a frequentist perspective, the inference algorithm is treated as a (stochastic) point estimation procedure, and how that point estimate might change if the estimation procedure were applied to a new random dataset was explored. In particular, a bootstrap procedure was used to estimate the distribution of these point estimates under new random datasets (see below). In addition, tools from hypothesis testing were used to summarize the degree to which the measured module usage differences could be explained by random fluctuation or finite-sample effects. To evaluate the model’s performance from a Bayesian perspective, the posterior uncertainty estimated by our MCMC inference algorithm was reported in terms of both raw samples and coordinate-wise posterior standard deviation error bars.

For the purposes of evaluating the procedure as an estimator that might be applied to new random datasets, the last MCMC sample was treated as a (stochastic) point estimate. To estimate the distribution of that point estimate under new random datasets, a bootstrap was performed: the underlying population distributions over sequences for each group (e.g., blank, TMT) were approximated by their empirical distributions, and sampling with replacement from those distributions generated new random datasets. Under the assumption that the data are generated from a stationary random process, this bootstrap procedure provides a consistent estimate of the distribution of the quantities estimated by the model inference, one that has been well validated in a wide range of contexts

(Wasserman, 2013).

To solve the label matching problem in the bootstrap samples, the new fits were initialized from the first fit on the full (un-bootstrapped) dataset; this way the variations in the bootstrap fits reflected the variations due to new random datasets and not any label matching effects (see the model consistency section above for methods that enable independent evaluation of label matching in fits with distinct random initializations). Figures in Chapter 2 show the bootstrap samples and resulting standard error estimates for the module usages in each experiment. In the main text, all additive statistical errors and error bars are bootstrap-estimated standard errors (±SE) unless otherwise indicated. We also computed Bayesian uncertainty estimates from MCMC samples. In Figure 2.7E we show the last 1000 sampled module usage vectors in a 1250 iteration run of our MCMC algorithm. The small variation indicates that the posterior mode being explored by the MCMC algorithm is highly peaked. To summarize the degree to which the differences in module usage vector estimates might be explained by random fluctuation and finite sample sizes, null hypotheses were formulated in terms of asymptotic normality assumptions on the distributions of the estimates, which was validated using Probability-Probability plots (data not shown). To define a single overall null hypothesis that the usage vectors of two populations are the same up to statistical variation, let u and v denote vectors of estimated usage proportions for two populations (e.g., blank and TMT). We test the null hypothesis that $u - v$ is distributed according to a multivariate normal distribution

$$H_0 \ : \ u - v \sim \text{Normal}(0, \Sigma) \tag{3.6}$$

where σ is estimated from the bootstrap samples. We used a standard test statistic, Hotelling's T-squared statistics, for H_0.

In addition to testing overall hypotheses about whether module usage vectors are the same in two populations, we also formulated p-values for hypotheses on the individual module usages to bound the (estimated) Type I family-wise error rate. To define these families of module-wise null distributions and p-values, asymptotic normality assumptions were again employed: letting u and v denote the estimated module usage vectors for two populations and letting ui and vi denote their ith coordinates, respectively, for each module i we define the hypothesis

$$H_0^{(i)} \ : \ u_i - v_i \sim \text{Normal}(0, \sigma_i^2) \tag{3.7}$$

The p-values in the optogenetic stimulation experiment were computed differently from the statistics described above. Let $u_i^(t))$ denote the frequency (across the 50 trials for a single fit) with which module i is used at frame index t relative to stimulus onset for t[-1,6] seconds (with 30 frame indices per second). We modeled each $u_i^(t))$ as a random variable with a distribution that may depend on the stimulus; that is, we partitioned the time relative to stimulus onset into three subsets corresponding to times when the laser is on $(0t < 2)$, times just after the laser turns off $(2t < 3)$, and all other times. The family of hypotheses

that

$$H_0^{(i,\text{ on vs off})} : u_i^{(t_1)} - u_i^{(t_2)} \sim \text{Normal}(0, \sigma^2_{i,\text{ on vs off}}) \qquad 0 \leq t_1 < 2, t_2 \in [-1, 0) \cup (3, 6], \tag{3.8}$$

$$H_0^{(i,\text{ offset vs off})} : u_i^{(t_1)} - u_i^{(t_2)} \sim \text{Normal}(0, \sigma^2_{i,\text{ offset vs off}}) \quad 2 \leq t_1 < 3, t_2 \in [-1, 0) \cup (3, 6]. \tag{3.9}$$

was then tested, where each null distribution's variance was estimated from the sample variance across time indices. P-values were then computed using a t-test and adjusted with the Holm-Bonferroni step-down procedure.

3.3.17 Technical definitions of the generative models

In this section, the priors and generative models implemented in this paper are formally defined following standard notation. These generative models are adaptations of models successfully used previously to characterize structure in complex datasets (Fox et al., 2009; Murphy, 2012)

3.3.18 Prior on the transition matrix

A sticky HDP prior was placed on the transition matrix π with concentration parameters $\alpha, \gamma > 0$ and sticky parameter $\kappa > 0$,

$$\rho_j \overset{\text{iid}}{\sim} \text{Beta}(1, \gamma) \qquad \beta_i = (1 - \rho_i) \prod_{j<i} \rho_j \tag{3.10}$$

$$\pi_i \overset{\text{iid}}{\sim} \text{DP}(\alpha\beta + \kappa\delta_i) \qquad i = 1, 2, \ldots, \tag{3.11}$$

where δ_{ij} is 1 when $i = j$ and is 0 otherwise and π_i denotes the ith row of π. Gamma priors are placed on α and γ, setting $\alpha \sim \text{Gamma}(1, 1/100)$ and $\gamma \sim \text{Gamma}(1, 1/100)$.

3.3.19 Generation of the discrete state sequence

Given the transition matrix, the prior on a discrete state sequence x was

$$x_t \,|\, x_{t-1}, \pi \sim \pi_{x_{t-1}}, \quad t = 2, 3, \ldots, T \tag{3.12}$$

where x_1 is generated by the stationary distribution under π.

3.3.20 Prior on the autoregressive parameters

The autoregressive parameters $\theta = \{\theta^{(i)}\}_{i=1}^{\infty} = \{(A^{(i)}, b^{(i)}, \Sigma^{(i)})\}$ for each state $i = 1, 2, \ldots$ were sampled from a Matrix Normal Inverse-Wishart prior:

$$(A, b),\ \Sigma \sim \text{MNIW}(\nu_0, S_0, M_0, K_0) \tag{3.13}$$

or equivalently

$$\Sigma \sim \text{InvWishart}(\nu_0,\, S_0) \tag{3.14}$$

$$\text{vec}((A, b)) \sim \text{Normal}(\text{vec}(M_0),\, \Sigma \otimes K_0) \tag{3.15}$$

where $\otimes$ denotes a Kronecker product and (A, b) denotes the matrix formed by appending b to A as a column.

$$K_0 = \text{diag}(k_1, \ldots, k_{KD}) \qquad k_i \overset{\text{iid}}{\sim} \text{InvGamma}(1/25,\, 1/25). \tag{3.16}$$

3.3.21 Generation of the 3D pose sequence principal components

Given the autoregressive parameters and the discrete state sequence, the data sequence y was generated according to an affine autoregression:

$$y_t \sim \text{Normal}\left(A^{(x_t)}\widetilde{y}_{t-1} + b^{(x_t)},\, \Sigma^{(x_t)}\right) \qquad t = K+1, K+2, \ldots, T \tag{3.17}$$

where the covariance matrices Σ are dense and $\widetilde{y}$ denotes a vector of K lags:

$$\widetilde{y}_t \triangleq \left[y_{t-K}^{\mathsf{T}} y_{t-K+1}^{\mathsf{T}} \cdots y_{t-1}^{\mathsf{T}}\right]^{\mathsf{T}}. \tag{3.18}$$

The alternative models are special cases of the AR-HMM and were constructed by adding constraints. In particular, the Gaussiaion-emission HMM (G-HMM) corresponds

to constraining $A^{(i)} = 0$ for each state index i. Similarly, the autoregressive mixture (AR-MM) and Gaussian mixture (GMM) correspond to constraining the transition matrix to be constant across rows, $\pi_{ij} = \pi_{i'j} = \pi_j$ for each i and $i\prime$, in the AR-HMM and G-HMM, respectively.

3.3.22 Technical description of the inference algorithms

As discussed above, the Gibbs sampling inference algorithm alternates between two principal stages: updating the segmentation of the data into modules given a fixed transition matrix and autoregressive parameters, and updating the transition matrix and autoregressive parameters given a fixed segmentation. Mathematically, updating the segmentation is sampling the label sequence x conditioned on the values of the data y, the autoregressive parameters θ, and the transition matrix π; that is, sampling the conditional random variable $x \mid \theta, \pi, y$. Similarly, updating the transition matrix and autoregressive parameters given the segmentation is sampling $\pi \mid x$ and $\theta \mid x, y$, respectively.

For inference in the AR-HMM the weak-limit approximation to the Dirichlet process (Fox, 2009; Ishwaran & Zarepour, 2002) is used, in which the infinite model is approximated by a large finite one that shares symmetries with the infinite case. That is, choosing some finite approximation parameter L, β and π are modeled using finite

Dirichlet distributions of size L

$$\beta \sim \text{Dir}(\gamma/L, \cdots, \gamma/L) \tag{3.19}$$

$$\pi_k \sim \text{Dir}(\alpha\beta_1, \cdots, \alpha\beta_j + \kappa\delta_{kj}, \cdots, \alpha\beta_L) \tag{3.20}$$

where π_k denotes the ith row of the transition matrix. This finite representation of the transition matrix allows the state sequence x to be resampled as a block and for large L provides an arbitrarily good approximation to the infinite Dirichlet process (Ishwaran & Zarepour, 2002).

Using a weak limit approximation, the Gibbs sampler for the AR-HMM iterates resampling the conditional random variables

$$x \,|\, \pi, \theta, y, \qquad \theta \,|\, x, y, \qquad \text{and} \qquad \beta, \pi \,|\, x. \tag{3.21}$$

For simplicity, throughout this section notation for conditioning on hyperparameters and the superscript notation for multiple observation sequences is suppressed.

3.3.22.1 Sampling $x \,|\, \pi, \theta, y$

Sampling the state labels x given the dynamical parameters, π and θ, and the data y corresponds to segmenting the 3D video sequence and assigning each segment to a behavioral mode that describes its statistics.

Given the observation parameters θ and the transition parameters π, the hidden state

sequence x is Markov with respect to a chain graph. The standard HMM backward message passing recursions are

$$B_t(k) = p(y_{t+1:T} \,|\, \theta, \pi, x_t = k) \tag{3.22}$$

$$= \sum_{j=1}^{K} p(x_{t+1} = j \,|\, x_t = k, \pi) p(y_{t+1} \,|\, x_{t+1} = j, \theta) B_{t+1}(j) \tag{3.23}$$

for $t = 1, \ldots, T-1$ and $k = 1, \ldots, K$, where $B_T(k) = 1$ and $y_{t+1:T} = (y_{t+1}, y_{t+2}, \ldots, y_T)$. Using these messages, the conditional distribution of the first state x_1, marginalizing over all the future states $x_{2:T}$ is

$$p(x_1 = k \,|\, \pi, \theta, y) \propto p(x_1 = k \,|\, \pi) p(y_1 \,|\, x_1 = k, \theta) B_1(k), \tag{3.24}$$

which can be sampled efficiently. Given a sampled value $\bar{z}_1$, the conditional distribution of the second state x_2 is

$$p(x_2 = k \,|\, \pi, \theta, y, x_1 = \bar{z}_1) \propto p(x_2 = k \,|\, x_1 = \bar{z}_1, \pi) p(y_2 \,|\, x_2 = k, \theta) B_2(k). \tag{3.25}$$

Therefore after passing HMM messages backward the state sequence can be recursively sampled forwards.

3.3.22.2 Sampling $\theta \mid x, y$

Sampling the autoregressive parameters θ given the state sequence x and the data sequence y corresponds to updating each mode's dynamical parameters to describe the 3D video data segments assigned to it.

To resample the observation parameters θ conditioned on a fixed sample of the state sequence x and the observations y one can exploit conjugacy (Bernardo & Smith, 2009; Bishop, 2006b) between the autoregressive likelihood and the MNIW prior. That is, the conditional also follows the MNIW distribution:

$$p(A_k, \Sigma_k \mid x, y, S_0, \nu_0, M_0, K_0) = p(A_k, \Sigma_k \mid S_n, \nu_n, M_n, K_n) \tag{3.26}$$

where (S_n, ν_n, M_n, K_n) are posterior hyperparameters that are functions of the elements of y assigned to state k as well as the preceding lagged observations:

$$S_n = S_0 + S_{yy^\intercal} + (M_0 K_0^{-1} M_0^\intercal - M_n K_n^{-1} M_n^\intercal) \tag{3.27}$$

$$M_n = (M_0 K_0^{-1} + S_{y\tilde{y}^\intercal}) K_n \tag{3.28}$$

$$K_n = (K_0^{-1} + S_{\tilde{y}\tilde{y}^\intercal})^{-1} \tag{3.29}$$

$$\nu_n = \nu_0 + n \tag{3.30}$$

where

$$S_{yy^\mathsf{T}} = \sum_{t:x_t=k} y_t y_t^\mathsf{T} \qquad S_{\tilde{y}\tilde{y}^\mathsf{T}} = \sum_{t:x_t=k} \tilde{y}_t \tilde{y}_t^\mathsf{T} \tag{3.31}$$

$$S_{y\tilde{y}^\mathsf{T}} = \sum_{t:x_t=k} y_t \tilde{y}_t^\mathsf{T} \qquad n = \#\{t : x_t = k\}. \tag{3.32}$$

Therefore resampling $\theta \,|\, x, y$ includes three steps: 1) collecting statistics from the data assigned to each state, 2) forming each state's posterior hyperparameters, and 3) updating each state's observation parameter by simulating a draw from the appropriate MNIW. Simulating $(A, \Sigma) \sim \mathrm{MNIW}(S_n, \nu_n, M_n, K_n)$ proceeds as

$$\Sigma \sim \mathrm{InvWishart}(S_n, \nu_n) \tag{3.33}$$

$$A = M_n + \Sigma^{\frac{1}{2}} G K_n^{-\frac{1}{2}} \quad \text{where} \quad G_{ij} \overset{\text{iid}}{\sim} \mathcal{N}(0,1). \tag{3.34}$$

3.3.22.3 Sampling $\beta, \pi \,|\, x$

Sampling the transition parameters π and β given the state sequence x corresponds to updating the probabilities of transitions among behavioral modules to reflect the transition patterns observed in the state sequence. Updating β encourages redundant behavioral modes to be pruned from the model, while updating each π_{ij} fits the transitions observed from state i to state j.

Resampling the transition parameters β and π, which are draws from the weak limit approximation to the (sticky) HDP, is performed using an auxiliary variable sampling

scheme (Teh et al., 2006; Fox, 2009). That is, $\beta, \pi \,|\, x$ can be generated by first sampling auxiliary variables $m \,|\, \beta, x$. Then $\beta, \pi \,|\, x, m$ can be generated by first sampling from the marginal $\beta \,|\, m$ and then the conditional $\pi \,|\, \beta, x$.

The matrix of transition counts in the sampled state sequence x is

$$n_{kj} = \#\{t : x_t = k,\ x_{t+1} = j,\ t = 1, 2, \ldots, T-1\}. \tag{3.35}$$

Suppressing conditioning notation for simplicity, the auxiliary variables $m = \{m_{kj} : k, j = 1, 2, \ldots, K\}$ are sampled via

$$m_{kj} = \sum_{l=1}^{n_{kj}} b_{kjl} \quad \text{where} \quad b_{kjl} \overset{\text{iid}}{\sim} \text{Bernoulli}\left(\frac{\alpha\beta_j}{\alpha\beta_j + \kappa} \frac{\alpha\beta_j + \kappa\delta_{kj}}{\alpha\beta_j + l + \kappa\delta_{kj}}\right) \tag{3.36}$$

where $\text{Bernoulli}(p)$ denotes a Bernoulli random variable that takes value 1 with probability p and takes value 0 otherwise. Note that the update for the HDP-HMM without a sticky bias corresponds to setting $\kappa = 0$ in these updates.

Given the auxiliary variables, the update to β is a Dirichlet-multinomial conjugate one, where

$$\beta \,|\, m \sim \text{Dir}(\gamma/K + m_{\cdot 1}, \gamma/K + m_{\cdot 2}, \cdots, \gamma/K + m_{\cdot K}) \tag{3.37}$$

where $m_{\cdot j} = \sum_{k=1}^{K} m_{kj}$ for $j = 1, 2, \ldots, K$. The update to $\pi \,|\, \beta, x$ is similar, with

$$\pi_k \,|\, \beta, x \sim \text{Dir}(\alpha\beta_1 + n_{k1}, \cdots, \alpha\beta_j + n_{kj} + \kappa\delta_{kj}, \cdots, \alpha\beta_K + n_{kK}). \tag{3.38}$$

3.4 Author contributions

The attribution of work included in this chapter is identical to that specified in Chapter 2, with special note that ABW is responsible for the development of computer vision algorithms, and MJJ is responsible for the development of modeling techniques.

Chapter 4

Motion Sequencing for Quantitative Behavioral Pharmacology

4.1 Introduction

The ultimate output of the brain is behavior. In nearly all organisms, sensory, affective and cognitive systems trigger or shape patterns of behavior, enabling the generation of adaptive responses to the outside world. Thus understanding how genes and circuits interact to create action – and how diseases and drugs can modify those relationships – requires a deep understanding of how behavior is organized across spatiotemporal scales and across time.

However, the main research tool deployed in behavioral neuroscience is the behavioral assay, a narrowly focused lens in which one specific feature of behavior is characterized at a time. This simplification involves two forms of human engineering that introduce different types of bias. First, behavioral assays are typically based upon anthropocentric inferences about the ”meaning” of various actions expressed by model organisms in the laboratory. Thus when trying to develop therapeutics for depression, researchers

commonly use the forced swim test (Porsolt et al., 1978), in which the failure of a rodent to swim is thought to be analogous to the amotivational states that characterize clinical depression; this assay has been validated because these periods of immobility can be reversed after pretreatment with known antidepressants (Castagné et al., 2006). The forced swim test has been useful for developing antidepressants similar to those used to validate the assay, but has not led to therapeutics that act through novel mechanisms. This finding – and observations suggesting alternative explanations for the periods of immobility – raises the possibility that the forced swim assay may not actually capture the rodent equivalent (if such a thing exists) of learned helplessness. As a consequence, the value of the learned swim test for characterizing genes and circuits related to affective state, a necessary antecedent for building new classes of therapeutics, is unclear. Second, most behavioral assays distill complex time-series behavior – the evolution of allocentric and egocentric actions over time — into a single average metric (Crawley, 2003, 2007); in the case of the forced swim test, this metric is the amount of time the animal remains immobile when facing the threat of drowning. Although this metric captures the feature of behavior that researchers *a priori* believe is meaningful (from their anthropocentric perspective), it is also a human-engineered choice, one that eliminates insight into the impact of plunging mice into water on global patterns of behavior.

This dependence on traditional behavioral assays has rendered many problems in behavioral neuroscience refractory to analysis — there often simply is not enough information recovered from narrowly tailored assays and single metrics to allow

researchers to fully understand the behavioral function of a gene or a neural circuit, or to develop new therapeutics aimed at altering pathological behaviors. In recognition of these limits, researchers are increasingly turning to characterizing animal behavior more comprehensively through the generation of behavioral ”fingerprints” (Rihel & Schier, 2012; Rihel et al., 2010; Vogelstein et al., 2014; Berman et al., 2014). In general, behavioral fingerprints are collections of individual metrics that create a snapshot of the overall pattern of behavior generated by a given animal, one that is useful for summarizing how an animal's behavior changes overall under distinct experimental conditions, and for distinguishing the behavioral responses of individual animals.

Although behavioral fingerprints are often collections of metrics that emerge from commonly-used behavioral assays (such as test batteries that aim to probe as many individual behavioral domains as possible), recent advances in machine vision have facilitated the creation of behavioral fingerprints that include large collections of parameters related to animal movement (Spink et al., 2001; Kabra et al., 2013; Rihel & Schier, 2012; Roberds et al., 2011; Rihel et al., 2010). Although these fingerprints include objective data (such as animal velocity), these collections of statistics are not free from bias: they often propagate the biases inherent in their component assays, and the very act of building a fingerprint from spatiotemporally diverse datastreams requires researchers to make many arbitrary choices, including engineering the features that make up any given fingerprint, and weighing how individual components are balanced within the fingerprint itself. Because of the bespoke nature of each behavioral fingerprint (which is usually

constructed on a case-by-case basis by an individual researcher, behavioral core or contract research organization), it is difficult to know whether any given fingerprint is more or less informative that any other, or whether the tradeoffs that are inevitably made between data quantity and cost or time are worthwhile.

An alternative to this commonly-used approach to behavioral analysis arises from ethology, whose explicit goal is to identify an underlying structure for behavior that reflects how it evolves over time, and which can be used to infer underlying the functional properties of the neural circuits (and genes) that create action. Underlying ethology is the notion of the behavioral module – a repeatedly used and stereotyped motifs of behavior that represent a unit of action; behavioral modules can be used to hierarchically organize complex behaviors from simpler parts, and to describe and predict how behavior evolves over multiple spatiotemporal scales (Tinbergen, 1951). For example, the continuous call of songbirds, such as zebra finches, has been segmented into non-overlapping and re-used syllables that can be usefully used to describe the time-series structure of vocal behavior (Thorpe, 1958). The identification of these behavioral modules led directly to the characterization of neural dynamics that underpin the production of bird song depended critically on segmenting song, and to the identification of neurons that code for the duration and identity of syllables and their transitions (Albert & Margoliash, 1996; Hahnloser et al., 2002; Olveczky et al., 2005). Likewise, a collection of simple movements has been defined in the nematode *C. elegans*, whose expression is controlled by defined neurons and whose sequencing creates adaptive behavioral outputs (Gray et al.,

2005). This ethological approach has revealed similar mechanistic insight into a wide variety of patterned behaviors across domains and organisms, ranging from wake/sleep cycles in humans to courtship in the fruit fly *D. melanogaster* (Konopka & Benzer, 1971; Vitaterna et al., 1994; Takahashi et al., 1994).

In all these examples, the segmentation of continuous behavior into modules was achieved by researchers who subjectively identified likely structure in the behavioral data; high-throughput systems based upon machine learning techniques have recently been developed that can automatically translate these ”supervised” segmentations of behavior into behavioral fingerprints (Jhuang et al., 2010; Ohayon et al., 2013; Dankert et al., 2009). These automated systems instantiate human intuition about the structure of behavior, although they also propagate the anthropocentric biases held by humans about what constitutes a meaningful behavior. A related set of approaches, which is less subject to bias uses mathematical and computational techniques to identify patterns within behavioral data, thereby revealing its underlying structure. This type of unsupervised analysis has been successfully performed in *D. melanogaster* to identify behavioral differences between fly strains during walking, and in *C. elegans* to create a look-up table that relates individual genes to specific patterns of action (Stephens et al., 2008, 2011; Brown et al., 2013; Vogelstein et al., 2014). These unsupervised methods for identifying the spatiotemporal structure in behavioral data are powerful, as the units of behavior identified comprehensively account for the global pattern of action expressed by a given organism in an experiment in a manner largely free of human influence, and the inherent

structure in behavior that is organized by these algorithms likely relates in a meaningful manner to the genes and neural circuits that created that pattern of behavior.

We have recently developed a method, which we refer to as Motion sequencing (MoSeq), that segments three-dimensional (3D) mouse behavior into its constituent parts using unsupervised machine learning (Chapters 2,3 and Wiltschko et al. (2015)). This approach is based upon the observation that the 3D pose dynamics of freely-behaving mice have inherent sub-second structure that appears to naturally segment action into potential behavioral modules. The MoSeq algorithm leverages this structure to identify ”behavioral syllables” – a set of reused and stereotyped 3D behavioral motifs – and their associated transitions, or ”behavioral grammar.” Because it takes advantage of inherent structure in every dataset it explores, MoSeq can identify and characterized new patterns of behavior that have not been observed previously, unlike supervised methods which depend upon prior observation of quantified behaviors. One feature unique to MoSeq (compared to prior supervised and unsupervised methods) is that it explicitly addresses the problem of the cardinality of behavior – rather than pre-specifying the number of behavioral syllables, MoSeq uses Bayesian non-parametric approaches to statistically determine the optimal number of syllables within a given experiment (Colgan, 1978). Importantly, behavioral syllables and grammar are semantically meaningful; unlike methods that collect a large number of arbitrarily-specified behavioral features, syllables and grammar collectively encapsulate the total observed pattern of action a mouse expressed within a given experiment.

MoSeq has successfully been used to characterize the structure of a variety of adaptive behaviors, to reveal the influence of specific genes and neural circuits on behavior, and to identify new phenotypes that had not been previously identified using traditional behavioral classification approaches. Although it was initially designed as a method to provide a comprehensive description of ongoing behavior – and to enable predictions about future behavior –- these finding suggest that MoSeq may have significant utility as a behavioral fingerprint. However, whether MoSeq (or related modeling-based segmentations of behavior) is suitable for use as a behavioral fingerprint is unknown. Furthermore, it is not clear whether MoSeq has any advantage over traditionally constructed behavioral fingerprints that agglomerate a variety of distinct parameters about behavior; given the computational expense of MoSeq, it is possible that simpler models or representations of behavior might be equally or even more useful as fingerprints. Indeed, to our knowledge, no principled framework has ever been used to broadly compare the utility of different representations of behavior in the task for which fingerprints were initially designed, to identify differences between experimental conditions.

We therefore wished to build an experimental framework that would enable us to compare the relative abilities of each fingerprint to distinguish and identify patterns of behavior. To generate a relevant behavioral dataset to make comparisons we imaged, in the circular open field, a large cohort of mice that had been treated with one of a variety of different neuro- and psychoactive drugs, at one of two separate doses. We chose to explore the ability of fingerprints to identify behavioral differences elicited by

pharmacological agents in part because phenotyping is a major roadblock in the drug development process: in this context it is not clear how to best define a behavioral ”signature” that is characteristic for a drug (and which may be revealing about the molecules and neural circuits targeted by that agent), nor is it obvious how to characterize patterns of behavior elicited by potential new therapeutics if those behaviors are substantially different from those observed previously. We also chose this approach because it would allow us to observe both dramatic and subtle changes (elicited by e.g., dopamine agonists and antidepressants, respectively) in a relatively neutral experimental context that would not be expected to itself elicit dramatic changes in behavior. After obtaining these data (using a 3D camera), we built a standard linear classifier, enabling us to compare the ability of different behavioral fingerprints to identify differences between treatment conditions, using the ability of the classifier to correctly predict the drug and dose a given mouse had been given as our unit of measure.

Using this approach, we compared the ability of MoSeq and other computational models for mouse behavior (similar to but simpler than MoSeq) to predict the identity and dose of the drug that was administered to individual animals. We also asked whether MoSeq was able to build more discriminative fingerprints than several traditional behavioral fingerprints (which are typically include metrics that capture the mouse’s velocity, height, length, etc.). We find that MoSeq outperforms alternative models for mouse behavior as well as traditional behavioral fingerprints, consistent with behavioral syllables representing a comprehensive, semantically-meaningful, and useful description

of the totality of the mouse's behavior in a given experiment. We then used MoSeq to characterize similarities between behavioral patterns elicited by various drugs, and find that MoSeq correctly identifies both drugs that share underlying mechanisms of action, and drugs that tend to be co-prescribed (which therefore may alter behavior in a similar manner albeit through a distinct mechanism). Finally, we use MoSeq to reveal the specific sub-second behavioral motifs that are most effective at distinguishing individual drugs from each other. Taken together, this work demonstrates that MoSeq is a high-throughput and dense phenotyping approach that can tease apart subtle behavioral changes induced by a wide range of pharmacological agents, suggesting its broad applicability to the automated discovery of relationships between treatment and behavior. These results further suggest that behavioral components identified by MoSeq – such as the behavioral syllables – may themselves represent a new form of druggable target that could be used to identify new classes of therapeutics for neurodevelopmental and neuropsychiatric disease.

4.2 Methods

4.2.1 Drug preparation and administration

IP, standard volume, solution changed based on drug solubility. All drugs were delivered as a 0.2ml IP injection, dissolved in saline, except when compound solubility required DMSO. Twelve animals were tested for each dose (Table 4.1).

Table 4.1 | Drug identities, classes and doses

Drug	Class	Low Dose (mg/kg)	High Dose (mg/kg)
Alprazolam	Benzodiazepine	0.5	2
Atomoxetine	Benzodiazepine	0.25	2
Bupropion	Antidepressant	10	40
Chlorpromazine	Antipsychotic	0.25	1
Citalopram	SSRI	0.25	10
Clozapine	Antipsychotic	0.25	10
Diazepam	Benzodiazepine	0.5	2
Fluoxetine	SSRI	10	30
Haloperidol	Antipsychotic	0.1	6
Methamphetamine	Stimulant	0.5	2
Methylphenidate	Stimulant	5	20
Modafinil	Stimulant	32	128
Phenelzine	Antidepressant	2.5	12.5
Risperidone	Antipsychotic	0.25	1
Venlafaxine	SNRI	2.5	10

4.2.2 Behavioral protocol

Male mice, aged 6-8 weeks, weighing between 28-30 mg, were used. Mice were habituated to the reverse light/dark cycle for a period of 5 days prior to the behavioral testing. On the day of the experiment, mice were habituated to the animal room in the lab for 10 minutes and then to the behavior room for an additional 10 minutes. Behavioral testing was carried out in a sound isolated room, under dim red lighting. Mice were injected with the drug or control solution and recording was begun immediately, lasting one hour per mouse. Mice were not reused.

4.2.3 Behavioral Fingerprint Extraction

Identification and extraction of mouse from raw Kinect depth video performed as previously described. Histogram-based fingerprints generated on a per-mouse basis with dimensionality equal to the maximum number of syllables in motion sequencing output (160). Upper and lower bounds set empirically such that the tail ends of the histogram were 0, thus not removing any signal from the data. All data preprocessing performed using custom code written in the Python language, using the open-source NumPy, SciPy and Pandas libraries (Van Der Walt et al., 2011; Pedregosa et al., 2011; McKinney et al., 2010). Visualizations built using the Matplotlib and Seaborn libraries (Hunter et al., 2007). All analysis performed in IPython notebooks and versioned using git (Pérez & Granger, 2007). The AR-HMM model used to generate the MoSeq fingerprint was fit using identical methods to (Wiltschko et al., 2015). We performed a grid search over the hyperparameter κ, the sticky bias that determines the average duration of syllables, and chose the value ($\kappa \approx 1.13e7$) which maximized discriminative performance on a validation dataset.

4.2.4 Pharmacological classification

The logistic regression linear classifier was trained as a one-versus-rest binary classification problem using binary entropy loss. The algorithm was implemented in the C/C++ library liblinear, with bindings implemented in the scikit-learn Python library (Pedregosa et al., 2011). The inverse regularization strength parameter, C, was optimized

on the AUC of control classification performance on a validation set via hold-2-out cross-validation. The first held-out sample was used as the validation set, and the second held-out sample was reserved as the test set, and used only to evaluate final performance. AUC, ROC and confusion matrices were calculated with scikit-learn.

4.2.5 Held-out classification

For each of the 16 drugs in the dataset, a training dataset was constructed by removing all mice that were injected with a particular drug. A classifier was trained using the optimal C previously derived from cross-validation for the MoSeq data representation, and then the held-out mice were presented to the classifier. A single row of the held-out confusion matrix represents the proportion of a given drug that was assigned to another drug. This procedure is repeated n=16 times, holding out each drug in sequence, and training a new classifier on the remaining data. In this procedure, it is impossible for the classifier to correctly identify the held-out drug, because it has no representation of the held-out drug class.

4.3 Results

4.3.1 Psychoactive Drugs as a Testbed for Behavioral Fingerprints

We wished to compare the ability of different 2D and 3D behavioral fingerprints to discriminate between patterns of behavior. Addressing questions of this type requires the assembly of a large reference behavioral dataset, one that includes both dramatic and subtle behavioral differences amongst subjects.

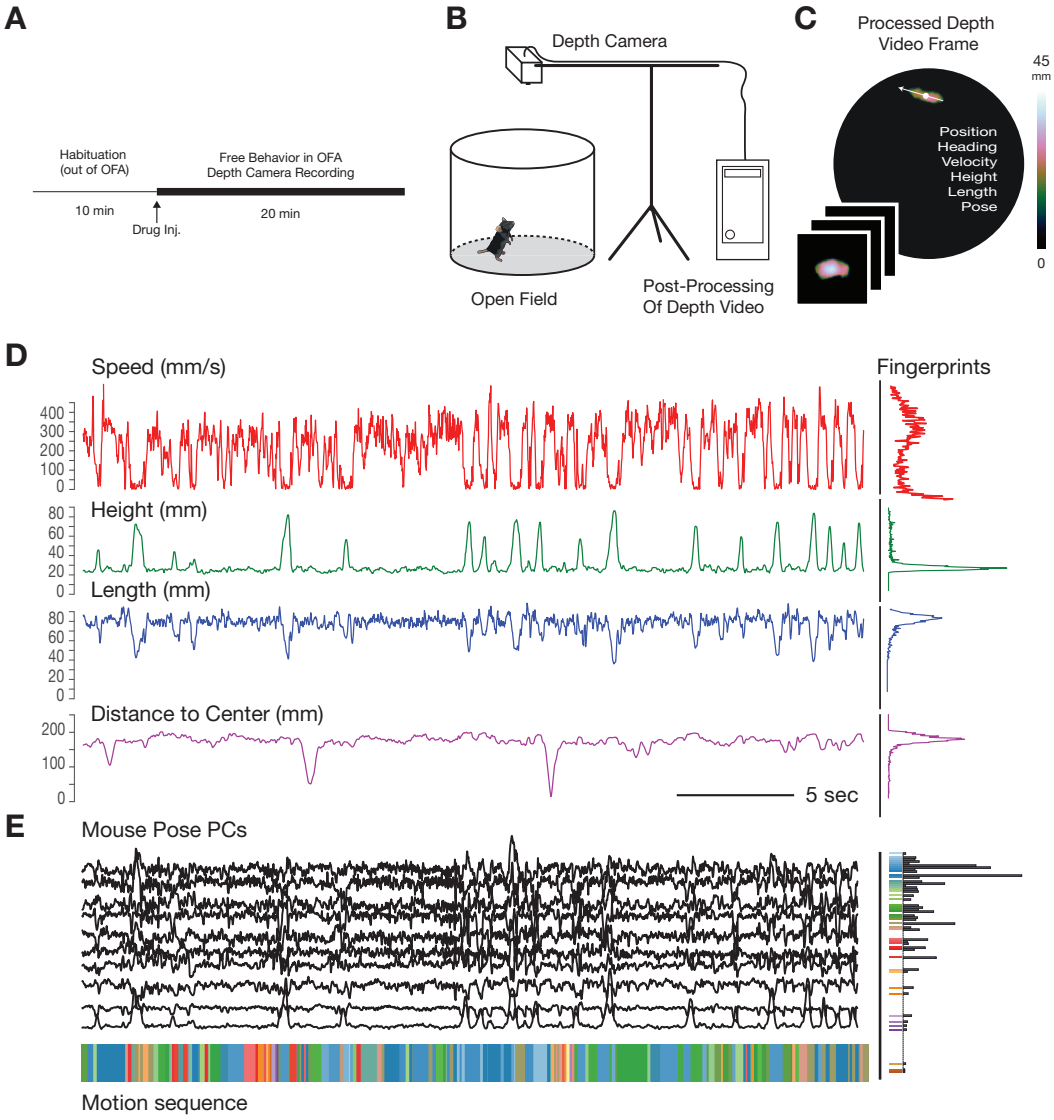


Figure 4.1 | A. After a brief habituation period in the room in which behavioral imaging is performed, mice are intraperitoneally injected with control or experimental pharmacological agents and then moved to the circular open field, after which recording is immediately begun. B. Schematic of recording setup. A depth camera is suspended above the OFA, recording the 3D contour of the mouse. Post-processing of frames using machine vision identifies the mouse's position, heading and velocity, and an aligned image of the mouse's 3D pose is extracted from the depth frame. C. Representative speed, height, length and distance-to-center time-series data (left). To distill these data into a characteristic fingerprint, data from each modality are used to build a histogram of the time-series data over a recording session (right). D. Creating a Motion Sequencing (MoSeq) fingerprint. Principal components (PCs) of the extracted and aligned mouse pose image are generated (top left), and segmentation into behavioral syllables is performed using the PC data (bottom left), where each color corresponds to a unique syllable. Counting the frequency of each syllable over the entire recording session yields a MoSeq fingerprint (right). Note that this fingerprint can also be engineered to include data about the transition statistics between individual behavioral syllables, although in the analysis performed here these additional data did not improve classification performance substantially.

We reasoned that acutely exposing mice to a panel of psycho- or neuroactive pharmacological agents would allow us to actuate a variety of neural circuits and neuromodulator systems, thereby eliciting many different patterns of action (Table 4.1, Figure 4.1A). For instance, mice treated with a high dose of a stimulant exhibit a higher mean speed, and anxiolytics increase a mouse's exploration away from walls and enclosed spaces. We therefore used intraperitoneal injection to deliver one of 16 drugs drawn from 6 broad pharmacological classes, at either a high or low dose (Figure 4.2A), to a large cohort of mice (n=360). Because we wished to generate behavioral differences in a manner that would be largely independent of the specific experimental context, we then placed mice in the circular open field, a neutral behavioral arena that lacks conspecifics or stimuli (other than the geometry of the field itself); we do not pre-habituate mice to this apparatus before imaging, and as a consequence the main behavior exhibited by untreated wild-type mice is simple exploratory locomotion.

To simultaneously enable the construction of traditional behavioral fingerprints as well as 3D-based fingerprints, after drug exposure mice were imaged using a depth camera, which allows simultaneously measurement of both conventional 2D parameters and moment-to-moment 3D mouse pose data (Figures 4.1B,4.2C). This system captured each mouse's speed, height, length and distance from the arena center (Figures 4.1C, 4.1D, left). In addition, the 3D image of the mouse was aligned and then subjected to MoSeq and related forms of behavioral analysis (see methods). MoSeq takes advantage of methods in computational inference to explicitly model three key aspects of behavior: the

number of behavioral syllables that most parsimoniously describes the observed patterns of mouse behavior, the fast postural dynamics ”within” each syllable (for example, the stereotyped pose dynamics that underlie the first half of a rear or a brief turn to the left), and the grammar that governs the sequencing of individual syllables. Rather than describing behavior as a collection of individual statistics (such as velocity, ellipticity and position), MoSeq coherently and holistically describes action by simultaneously identifying its component parts and the statistical structure that governs how those parts are connected in time.

4.3.2 MoSeq, and Other Model-Based Segmentations of Action

The computational model underlying MoSeq makes two assumptions about mouse behavior: first, that each ”syllable” can be described as a trajectory through pose space (which is a space defined by all of the poses a mouse exhibits during an experiment) and second, that each syllable switches from one to another via dynamics that can be described with a Markov Model. A number of simpler models that describe behavior could also be proposed (such as that each syllable is a simple Gaussian distribution of poses in pose space). These alternative models are less effective than MoSeq at predicting the moment-to-moment behavior of mice (a standard method for comparing model performance), which suggest that these alternative models (which are faster and simpler to fit) provide less rich descriptions of behavior. However, it is not clear that the most accurate description of moment-to-moment behavior is the most discriminative – it is possible that simpler models are as capable as MoSeq at telling different behaviors apart,

without all of the computational bells and whistles.

We therefore used our database of drug-treated mice to identify the modeling method that was best able to discriminate mice that had been exposed to different drugs. MoSeq and simpler behavioral models were trained on the drug dataset, and then the observed frequency of each behavioral syllable (whose number varied from model to model) was used to compose a series of behavioral fingerprints. These fingerprints were then submitted to a linear classifier, which was initially ”trained” to identify the specific drug a given mouse was exposed to (by co-submitting both the behavioral fingerprint and the identity of the drug treatment to the classifier) and then ”tested” for whether the classifier could correctly identify the drug a given mouse was exposed to based upon the behavioral fingerprint alone. The accuracy of a given classifier is often expressed as an Area-Under-the-Curve (AUC) measurement, which summarizes how accurately a given classifier predicted the treatment class of a given mouse (with 1 being perfect performance, and .5 being performance at chance). AUC analysis revealed that, of the various behavioral models compared, MoSeq was best able to identify the drug treatment class to which an individual mouse belonged (data not shown). This suggests that, of the models tested, the naturalistic description of behavior articulated by MoSeq (with syllables described as smooth changes between poses, and switching between syllables described as a first order Markov process) is both the most predictive (as shown via held-out likelihood analysis, Wiltschko et al. (2015)) and has the most utility in terms of capturing differences in behavior that distinguish individual mice.

4.3.3 Comparing the Utility of A Portfolio of Behavioral Fingerprints

These results suggested that we could extend the drug classification framework to ask whether the MoSeq fingerprint was more or less able to identify the drug treatment class than behavioral fingerprints containing more traditional components. Each of the 2D measurements of behavior that emerge from our imaging (whether it is velocity or distance from center) captures different features and dynamics of mouse behavior; when these metrics are summarized across the duration of the experiment by binning and counting all values into a histogram, these metrics can be used to compose a ”traditional” behavioral fingerprint. We measured four parameters (speed, position, length and height, each commonly used in traditional forms of behavioral analysis in the open field) for each mouse from our dataset, and then created a histogram that described each parameter’s distribution during each 20 minute experiment.

We then asked whether each of these four metrics alone revealed any drug-specific features of behavior, which would be expected given that distinct drug classes differentially alter aspects of mouse behavior (and trigger a spectrum of behavioral effects, including the stimulation or suppression of arousal, the induction of stereotyped repetitive behaviors and increases in exploration). This is visually apparent for drugs in the stimulant class, which as a group increase the likelihood of high-speed movements and proximity to the arena edge (indicated by brighter color on the right side of both the position and speed histograms, Figures 4.2B,4.2C). Conversely, the high dose of the antipsychotic and sedative haloperidol creates a ”speckled” pattern in the position

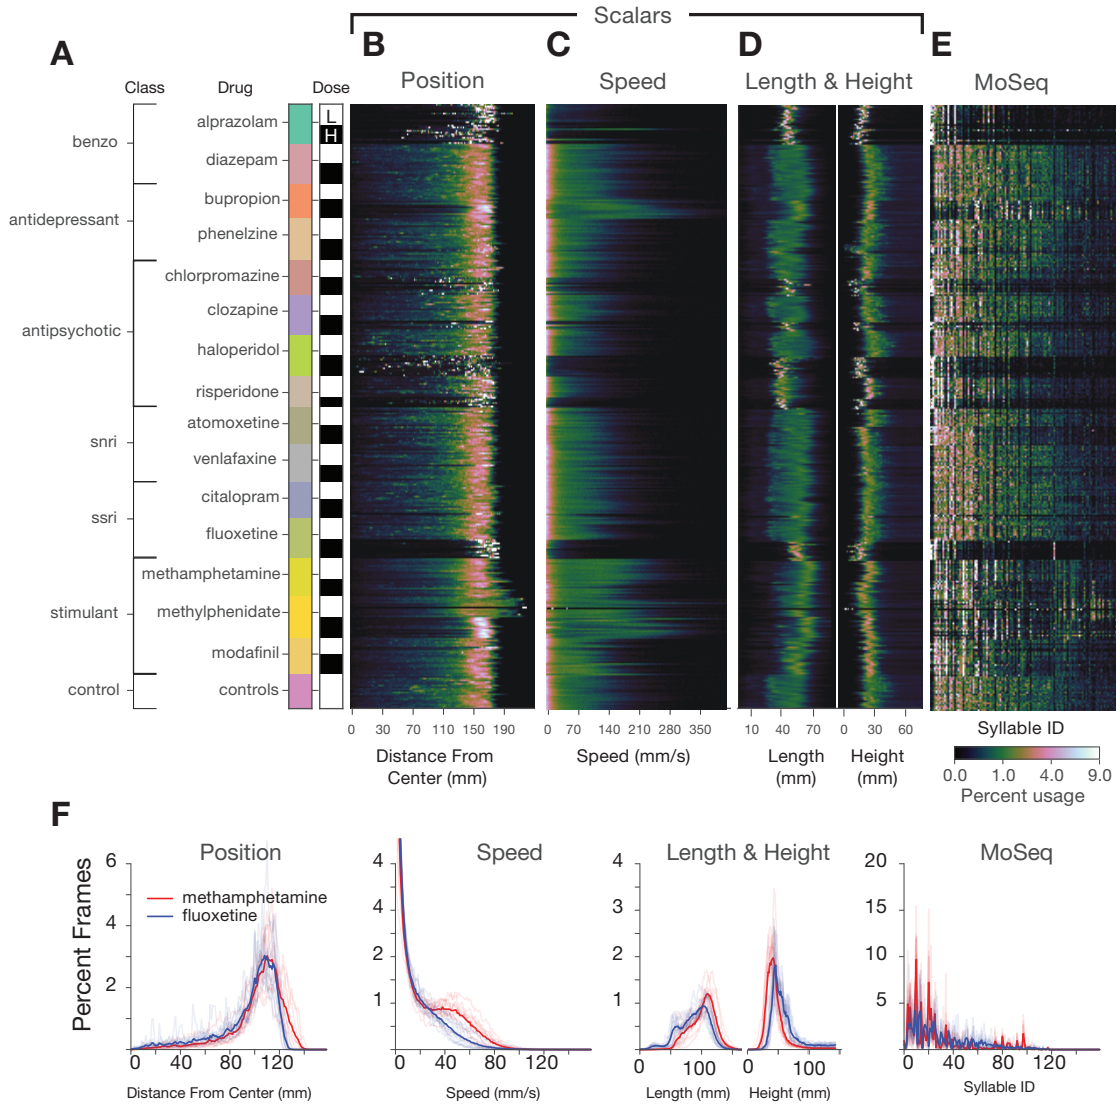


Figure 4.2 | A. All fingerprint types for all mice are shown, organized hierarchically by class, drug and dose (top, rows represent individual mice). In Tata representation many different drugs and drug/dose combinations, appear as qualitatively distinct blocks, which are apparent for each class of behavioral fingerprint. However, not all drug/dose pairs are visually distinguishable. B. Comparing fingerprints for two drugs, methamphetamine and fluoxetine (both at high dose), reveals differences fingerprints based on position in the arena, overall velocity, simple pose, and MoSeq. Individual mice are shown at low opacity, and the population mean for each drug shown as a thick line.

fingerprint and a high probability of zero speed. This is a result of treated mice lying prone and immobile in a stereotyped pose for long stretches during the experiment, also visible as a shifted and narrowed length and height fingerprint (Figure 4.2D). The drugs alprazolam, risperidone and fluoxetine, all drugs with sedative or anxiolytic effects, elicit similar behavioral responses. Note, however, that from inspection of these histograms it is not obvious that many of the drugs (e.g., fluoxetine, venlafaxine) induce behavioral changes that differ from control. Comparing all mice from just two drugs, methamphetamine, a stimulant, and fluoxetine, an anxiolytic SSRI (both at high dose), and plotting the mean data for all types, shows statistically-significant differences at the population level (Figure 4.2F). Consistent with a large volume of prior data, these results suggest that traditional behavioral metrics can successfully identify mice treated with drugs that have dramatically different behavioral effects; these results also suggest that our experimental approach generated the desired diversity of behaviors in our reference dataset.

When plotted as histograms, the four scalars that were measured appear to be smoothly varying and seem to fall naturally into categories; plotting the MoSeq fingerprint (composed as the usage of individual behavioral syllables identified by the algorithm, ordered by the most frequently-used syllables across all mice to least) is less interpretable in this data representation (Figure 4.2E). Thus it is unclear whether MoSeq fingerprints provide as much information as conventional metrics, or would be better suited to identifying treatment classes. To quantitatively test whether MoSeq can outperform

conventional behavioral metrics in a classification task, we considered each scalar histogram as a behavioral fingerprint (referred to as a position fingerprint, speed fingerprint, or length and height fingerprint), and in addition constructed a master fingerprint that assembled all of the data from all the scalar histograms into a single fingerprint (referred to as a scalar fingerprint). Note that, unlike traditionally obtained metrics which typically are composed of simple averages, the scalar fingerprint is composed of multiple histograms describing data distributions (including the histogram for height, which results from our 3D imaging and is typically not available from beam break measurements); we chose this approach to maximize the information density in the traditional fingerprints. We then, as before, built a classifier that learned to label a fingerprint from a single mouse as belonging to a single drug class, drug identity or drug identity and dose.

4.3.3.1 Discriminating Drug Class

We then investigated the discriminability of behavior induced by different pharmacological classes of drug when summarized into various fingerprints by position, by speed, by length and height, by the concatenated scalars and by MoSeq; this first, crude classification was performed as one would expect that any fingerprinting method that was reasonably accurate should be able to distinguish one category of drug from another (i.e., tell a stimulant from an anti-depressant, Figure 4.3). The ability of each fingerprint (after submission to a linear classifier) to correctly identify the class of drug to which individual mice were exposed was plotted as a confusion matrix. In this representation the actual

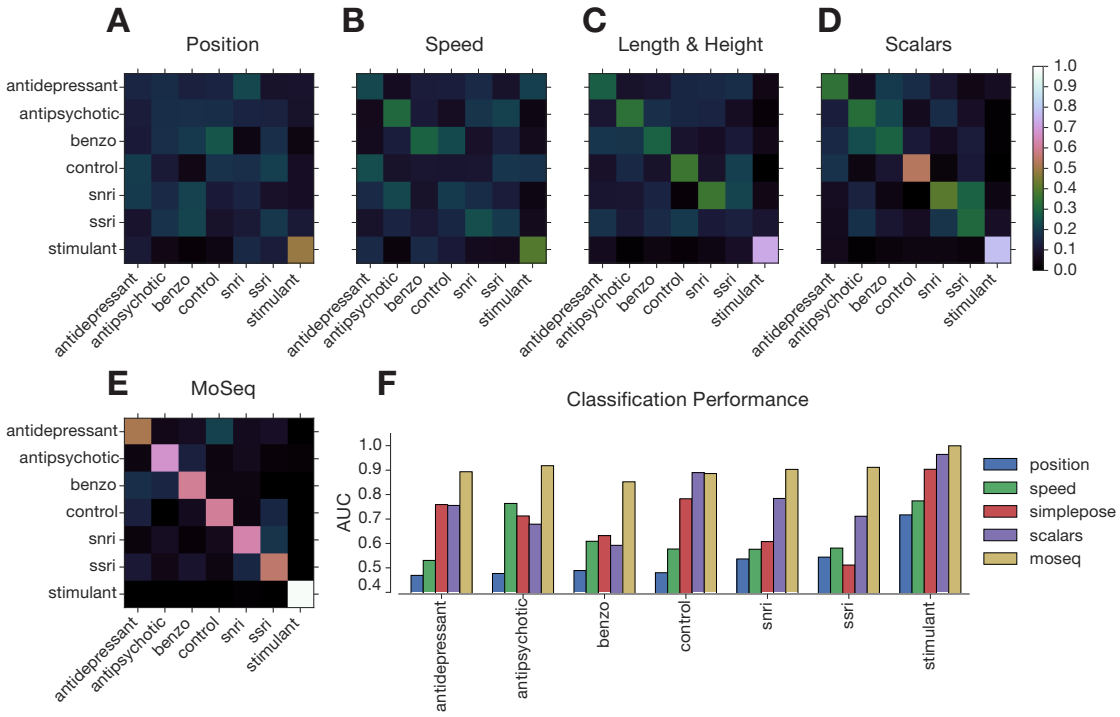


Figure 4.3 | A.Normalized confusion matrices of drug class for classifiers trained on position, speed, simple pose, scalar and MoSeq fingerprints (see Methods), with the actual treatment class on the Y axis, and the classifier predicted class on the X axis (with class assignments plotted with the indicated color code); perfect classifier performance in this data representation would be indicated as a white line across the diagonal. Stimulants are discriminable using any of the tested behavioral fingerprints, but only the concatenation of all scalar behavioral data, and MoSeq, reliably distinguish control from treatment classes. Importantly, MoSeq affords the best overall ability to distinguish drug treatment classes of the possible behavioral fingerprints. B. Area-Under-the-Curve analysis of classifiers trained with each behavioral fingerprint, organized based upon drug treatment class reveals that MoSeq equals or outperforms all alternative representations of behavior.

treatment class is indicated on the Y axis, and treatment class predicted by the classifier is indicated on the X axis; thus perfect performance would generate a perfectly white diagonal with no off-diagonal shading, while poor classifier performance will have significant off-diagonal signal. Consistent with visual inspection of the scalar data, stimulants were distinguishable with every fingerprint type, indicating that as a class, they influence body morphology, overall arousal, and exploratory activity in a manner easily captured by speed and length measurements. However all of the conventional fingerprints are significantly less able to identify non-stimulant classes of drugs, suggesting that conventional metrics are not broadly useful for identifying the class of drug mice were treated with. MoSeq, in contrast, outperforms all alternative fingerprints for all drug classes, and overall enables accurate prediction of the drug class in 85 percent of mice. These results are consistent with AUC analysis, in which MoSeq outperformed the scalar fingerprint for all drug classes except for stimulants (which induce dramatic and obvious behavioral changes, see Figure 4.2), and for untreated mice. Note that we built additional scalar fingerprints (which include higher-order metrics such as acceleration, angular velocity and ellipticity) and MoSeq fingerprints (which include information about the transition statistics that interconnect syllables in time), but these did not significantly improve classifier performance and so were not included in this or subsequent analysis (data not shown).

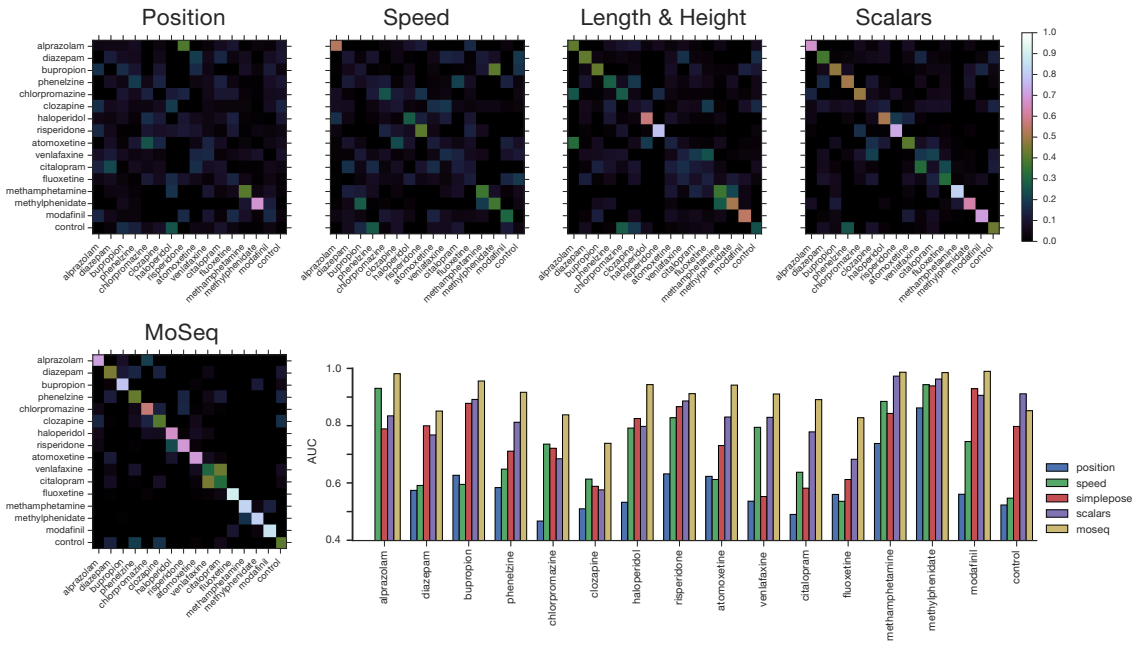


Figure 4.4 | A. Normalized confusion matrices (as in Figure 4.3) of drug identity for linear classifiers trained on each indicated fingerprint type. Traditional scalar-based behavioral fingerprints can correctly identify mice exposed to pharmaceuticals, like stimulants and sedatives, that primarily affect locomotion and overall arousal, but only MoSeq successfully classifies all individual drugs (with the exception of confusing citalopram and venlafaxine). B. AUC analysis summarizing linear classifier performance across drug treatments using each indicated behavioral fingerprint. MoSeq-based classifier performance has an AUC above 0.8 for every drug except clozapine.

4.3.3.2 Discriminating Drug Identity

The notion of pharmacological class is not rigorous, as each ”class” includes drugs that are placed into that class either because of their known molecular mechanisms or because of their psychological or neurological effects; in our dataset there are clearly drugs that can be considered to span classes (e.g. SSRIs and non-reuptake anti-depressants). The poor performance of scalar-based fingerprints may be due to the ambiguous nature of these classes, raising the possibility that scalar fingerprints may be better able to discriminate between treatments with more specific categorization. To directly test this possibility, we asked whether a fingerprint-based classifier could identify the specific drug each mouse was given, among fifteen drugs (irrespective of dose) and one control treatment (Figure 4.4). As was observed in the drug class-based classifier, MoSeq exhibited high performance as a behavioral fingerprint. For example, classifier fed the MoSeq fingerprint was able to identify the drug each mouse was treated with more than 85 percent of the time; AUC analysis revealed nearly perfect MoSeq performance for many drugs, with the worst performance for clozapine (AUC .70). in contrast, on average none of the scalar-based fingerprints improved their classification performance when asked to identify the specific drug to which a mouse was exposed. While the performance overall and the drug-by-drug classification efficiency of the full scalar fingerprint is significantly below that of MoSeq, the use of the full scalar fingerprint enabled near-perfect identification of several stimulants and risperidone.

The drug identity-based classification framework also presents a higher-resolution view

into the errors and successes of each fingerprint type. For instance, the MoSeq-based classifier often confused citalopram (an SSRI used in the clinic as an anti-depressant) and venlafaxine (an SNRI also used as an anti-depressant); these confusions between individual drugs in part explain the inability of the MoSeq classifier to always distinguish between SNRIs and SSRIs (Figure 4.3). Similarly mice exposed to phenelzine, one of two anti-depressants, were confused at low rates with control mice, which contributed to the confusion in the class-based classifier between anti-depressants and control (Figures 4.3,4.4). Conversely, this analysis reveals circumstances in which MoSeq was uniquely successful. For example, the stimulants (methamphetamine, methyphenidate, modafinil) elicited similar behavioral patterns when assessed via scalars (Figure 4.2), and accordingly these drugs were imperfectly classified based upon individual or concatenated scalar fingerprints; in contrast, a MoSeq-based classifier very rarely misclassified these drugs amongst themselves, and never confused them with a drug of another class (Figures 4.3,4.4). Taken together, these results demonstrate that traditional scalar based behavioral fingerprints (even those generously composed with height histograms) are outperformed by MoSeq fingerprints, suggesting that this ethologically-based representation of behavior has more utility in teasing apart subtle behavioral variation.

4.3.3.3 Discriminating Drug Identity and Dose

Given the ability of MoSeq fingerprints to disentangle drug identity from a simple open field experiment, we further increased the specificity of the task by asking whether MoSeq-based classifiers could successfully identify both the drug and dose a given mouse

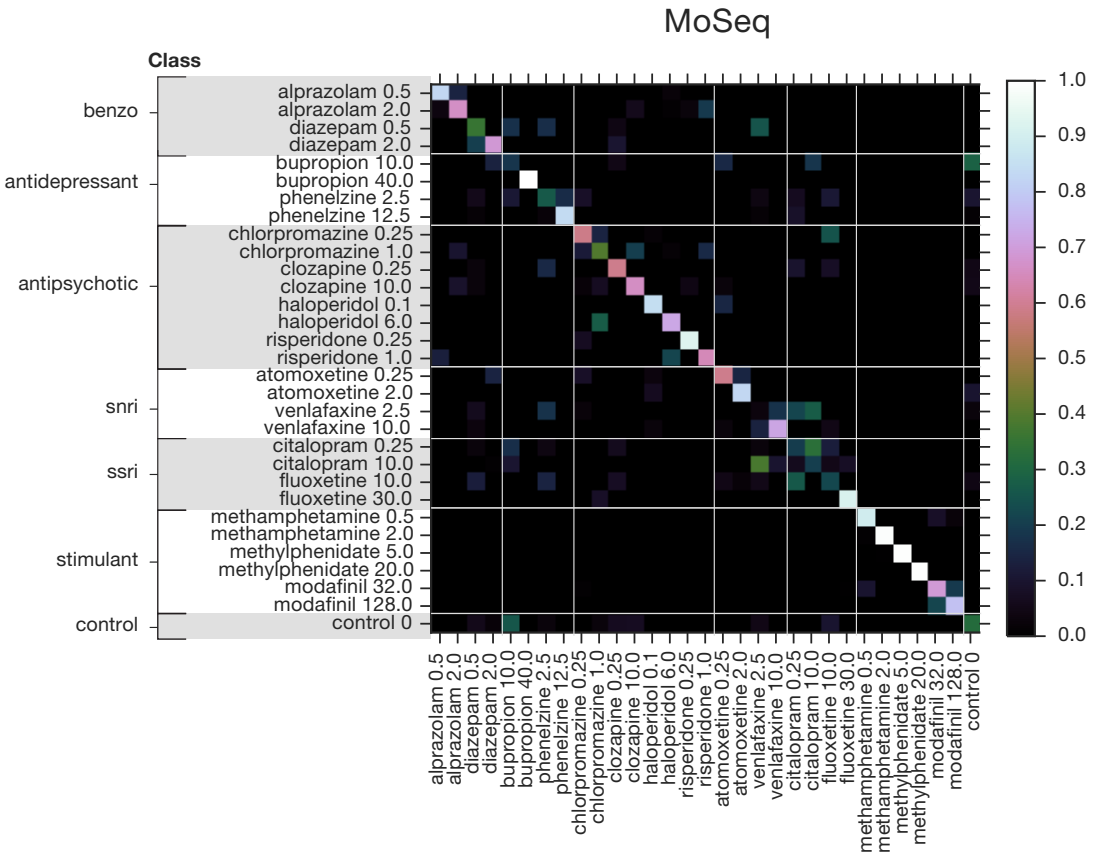


Figure 4.5 | Normalized confusion matrix for a MoSeq-based classifier tasked with predicting the drug and dose to which a given mouse was exposed. While many specific drug/dose combinations can be correctly identified, in those cases in which the classifier is confused, often the confusion is between different doses of the same drug. Across different classes of drug, confusions occur most frequently among antidepressants (see block structure in matrix). Antipsychotics are most frequently confused amongst themselves, while stimulants are rarely misclassified.

was exposed to (when simultaneously confronted with behavioral patterns representing all drug-dose pairs, Figure 4.5). Surprisingly, MoSeq was able to disambiguate most drugs and doses. Nearly all of the failures occurred when the classifier mis-identified a lower dose of a specific drug for a higher dose of that same drug, which occurred for two drug pairs (e.g. diazepam, modafinil). The only cross-drug confusions were amongst SNRIs and SSRIs (previously observed in the drug alone-based classifier), in which individual drugs (venlafaxine, citalopram and low dose fluoxetine) were often confused for each other.

These observations suggest that the failures in drug classification using MoSeq-based fingerprints may be useful for identifying drugs that modify either the same molecular target, similar molecular targets, or (more interestingly) modify patterns of behavior that are shared by drugs that do not share a mechanism of action. For example, high- and low-dose citalopram are indistinguishable, as are venlafaxine and citalopram; these drugs are all prescribed for depression and based upon our classifier analysis generate similar behavioral fingerprints. This suggests the possibility that new, unknown compounds could be characterized by their similarities to existing compounds when viewed through the lens of a MoSeq drug classifier.

4.3.4 Discovering Relationships Between Drugs with MoSeq

To simulate the scenario of a novel compound, we iteratively removed each drug from the training dataset, thereby building a classifier that had never been exposed to the behavioral data for a given drug; this ablation was performed for each drug in the dataset,

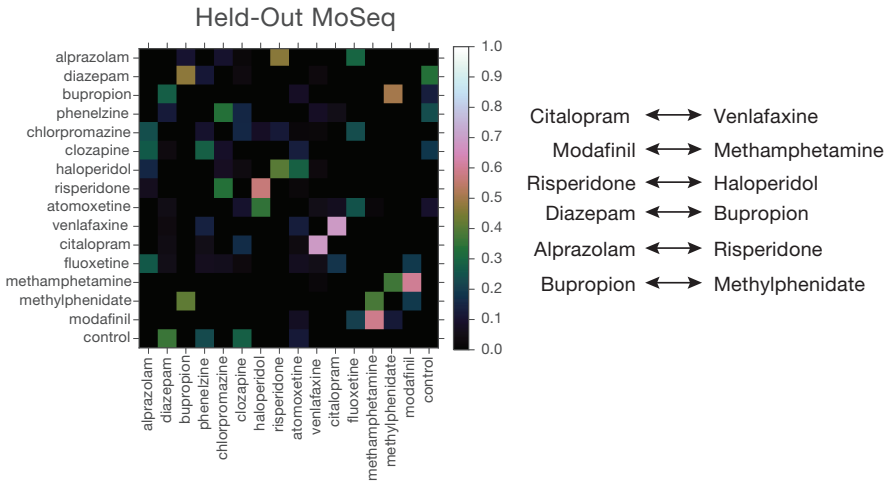


Figure 4.6 | Held-out confusion matrix for drug identities of a MoSeq-based classifier; in this configuration, the diagonal has no weight because the classifier was trained on every drug except the one being withheld, and then tested only on the withheld drug, in order to reveal similarities between compounds. The Y axis represents the drug each mouse was exposed to, and the X axis the drug the classifier believes the animal has been exposed to (after training on data that exclude mice exposed to the test drug). High-similarity pairs of drugs are highlighted on the right. Note that this analysis highlights the misclassifications noted from Figure 4.4 , such as the strong similarity between citalopram and venlafaxine.

resulting in a family of classifiers that were each missing a single drug. We then assessed which drugs were most often confused with each of the probe drugs, reasoning that these confusions may shed light on common molecular or circuit-level targets that influence behavior. The confusion matrix resulting from repeating this procedure for every drug in our dataset has no weight on the diagonal, indicating that each drug was unknown to the classifier at testing time (Figure 4.6). Confusions that were prominent in the previous task are now amplified, highlighting the relationships between different drugs in our common behavioral dataset. The most prominent confusion in the held-out classifier is the previously-confused pair, citalopram and venlafaxine, both anti-depressants (although with distinct mechanisms of action). Importantly, drugs with the similar behavioral effects

in the clinic were also confused – methamphetamine is grouped with modafinil, a non-addictive stimulant with an unknown mechanism of action, and risperidone and haloperidol, both antipsychotics with sedative-like effects, were also grouped. Two pairs, primarily used as anxiolytics in the clinic, diazepam/bupropion and alprazolam/risperidone, were also found to be similar. Finally, bupropion, an anti-depressant occasionally prescribed for ADHD, and methylphenidate, commonly prescribed for ADHD were linked via the held-out fingerprint classifier. Thus MoSeq-based fingerprints allows objective identification of similarities between pharmacological agents; these similarities include the identification of drugs that share mechanisms of action, but more importantly also includes drugs that appear generate common patterns of behavior (and may target similar neural circuits, regardless of their molecular mechanism of action).

4.3.5 Individual Syllables Can Discriminate Drugs

The high discrimination performance using MoSeq behavioral fingerprints was achieved by using the entire high-dimensional behavioral representation. It is possible that different drug treatments affect the usage frequency of all syllables diffusely. However, in prior work (Wiltschko et al., 2015), we found that odor-, genetic- and optogenetic-based modifications of behavior induced changes in only a small set of the animal's emitted syllables. To investigate whether drugs also focally affect behavior syllables, and whether drug-induced behavior can be discriminated on the basis of subsets of the animal's total behavioral syllabic repertoire, a decision tree-based classifier was fit on all drugs and

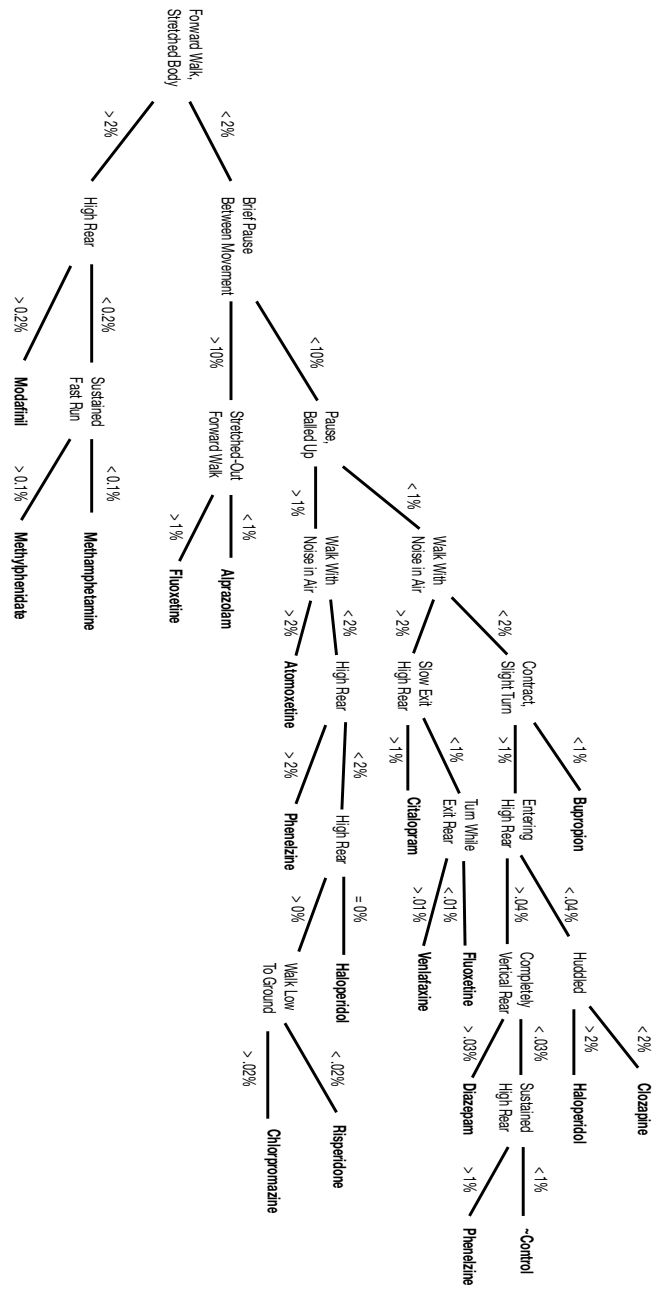


Figure 4.7 | By hierarchically sorting mice based upon the frequency of use of one syllable at a time, the behavioral effects of most drugs can be reliably distinguished. Low-reliability decision points have been pruned for visual clarity. Each node represents a ”decision” of whether a single syllable was used more or less than a given threshold, and ”leaves” of the decision tree (nodes with no branches to the right) are drug classification decisions, with drug names in bold. Mice that end in leaves towards the top of the tree have had more syllables downregulated (relative to control) than mice that end in leaves at the bottom of the tree. For example, all stimulants are nodes at the bottom of the first decision branch. The first syllable node is a forward locomotion syllable, which is up-regulated in all stimulants.

doses, combined together. A decision tree classifier discriminates between drugs by iteratively checking whether a given syllable is used above or below some learned threshold, until it reaches a leaf in the tree associated with a single drug (Figure 4.7). We found that individual drugs may be segmented from each other on the basis of the usage of single syllables, as shown schematically as the learned ”flow chart” derived from the trained decision tree, in Figure 4.7, and have also given each syllable a human label. All three stimulants are segmented from the rest of the pharmaceuticals immediately, based on the usage of a single locomotive syllable, and are themselves distinguished on the basis of an exploratory high rear, and a very high velocity dart. Conversely, the sedatives haloperidol and risperidone are defined by the absence of exploratory rearing behavior and the presence of pausing or stalling behaviors. Overall, most drugs in this analysis are defined primarily by the absence or downregulation of certain syllables, not the presence or upregulation of them.

4.4 Discussion

In this study we directly compared the predictive power of behavioral representations. To do so, we recorded the behavior of mice injected with a high or low dose of one of 16 drugs across 6 classes, and summarized the complicated time-series dynamics of their behavior using each of five different behavioral representations, including traditional speed- or rough-morphology-based metrics, or MoSeq, a technique we have recently developed. Although drugs that primarily affect arousal have been well-characterized by

speed-based behavioral representations, it was unclear to us if we could successfully discriminate non-sedative, non-stimulant drugs. Using each of these behavioral representations, we attempted to classify what treatment a mouse had received using only information in each behavioral representation, allowing an even comparison of the predictive power of each representation type. Although representations besides MoSeq were generally able to classify stimulants, only MoSeq was able to classify all drug classes, drug identities, as well as even drug identities and dose. Given MoSeq's strong classification performance in the regime where it was tasked to identify which treatment a mouse was given from of a known set of drugs, we attempted to emulate the scenario of a completely novel drug. We sequentially held out each drug in our dataset, and retrained our classification system to predict the class of a completely held-out drug (where identifying it successfully had now become impossible), in order to see which drugs with which the system would confuse it. The most frequent confusions made in this experiment were between drugs that either are prescribed for the same indication in the clinic, or drugs that share a similar mechanism of action. Last, we wished to inspect whether individual behavioral syllables, or repeated motifs of action, were indicative of the treatment a mouse has received. We *a priori* did not know whether the deployment frequency of syllables were diffusely changed by drug action, or whether a small set of syllables were focally altered, e.g., whether or not there exists individual "stimulant" or "antipsychotic" syllables. We constructed an alternative classification model, a decision forest, which, although less performant than the logistic regression model used to compare

representations, is more interpretable. The decision tree builds a ”flow chart” for classifying drugs, examining the usage frequency of one syllable at a time. When inspecting the result of a trained decision forest, we found that a single syllable, when used more than 2% of total experiment time, positively identified the mouse as having received a stimulant with high accuracy. When examined, the syllable was found to be a fast forward locomotion syllable, which is an intuitive indicator of stimulant-induced mouse behavior. Additional single syllables differentiated between sub-types of stimulants were less obvious, but still distinct, for instance the frequency of rearing up on the side of the arena was a robust differentiator between modafinil and methylphenidate. Among a variety of alternative representations of behavior, MoSeq is the most discriminative behavioral fingerprint for classifying drug class, identity and dose, suggesting broad applicability to the automated discovery of relationships between treatment and behavior.

4.4.1 MoSeq produces discriminative and interpretable representations of behavior

Histogram-based representations of locomotory behavior (e.g. speed, position) are only discriminative for arousal-altering drugs, such as stimulants, and sedatives. MoSeq, however, is able to distinguish a broader set of pharmaceuticals, again using only information drawn from gross locomotor behavior. We hypothesize that the distinction between MoSeq’s performance and alternative representations is MoSeq builds its representation by directly modeling the temporal dynamics of behavior. Simpler representations only represent how often a particular behavioral metric occurs, such as how often the mouse runs at 120 ms/s, and not how the mouse’s speed is modulated over

time. MoSeq's attempts to model not just the short-timescale dynamics of the mouse's pose, but also the cardinality and dynamics of the stereotyped movements the mouse makes. Thus, comparatively, individual dimension of histogram-based representations are information-poor – a single dimension of e.g. the position representation, perhaps the time spent at 10-11cm from the center of the arena, is both highly correlated with nearby dimensions (time spent 9-10cm, and 11-12cm from the center), and also difficult to interpret in isolation. MoSeq, in contrast, offers highly interpretable individual dimensions – each dimension of a MoSeq representation is a repeated motif of action, that can be examined in the original video of the experiment, overlaid to produce "crowd movies", or otherwise aligned for inspection. Our result that a single syllable differentiates stimulants from all other classes of drugs indicates that syllables are not only interpretable in isolation, but can individually contribute predictive validity to the MoSeq approach. We note that the predictive power of the concatenation of all scalar metrics, for some drugs, is nearly as performant as MoSeq for discriminating class, and drug identity, and is a better representation than any isolated histogram-based representation. It is possible that the addition of even higher-resolution behavioral features (such as e.g. the full skeletal pose of the animal, or its foot-paw positions over time) would be nearly as predictive as MoSeq, but it is possible that MoSeq, when given such additional features, would receive a similar improvement in predictive power.

4.4.2 Drugs with similar pharmacological profiles are distinct, but related, in the MoSeq representation

Psychoactive pharmaceuticals can be roughly related by their effect in ameliorating certain cognitive or motor symptoms, like anxiety or tremor. A surprising feature of the current study is that each drug tested can produce a nearly unique signature, distinguishing it from all other drugs in the dataset, even the same drug at a different dose. The most common type of misclassification, however, is confusion between the same drug at a different dose, or between drugs from within the same broad pharmacological class, or those that share a common clinical purpose. We observed this phenomena more strongly in the classification task that explicitly amplified confusions between drugs. Given that nearly all drugs and doses in our dataset are distinguishable, and rare, or forced, mistakes occur between related compounds, we reason that MoSeq reveals that precise representations of behavior can succinctly organize descriptions of the neurological effect of drugs. Further, given that drugs with similar clinical indications are similar in their MoSeq representation, we expect that our syllable-based representation of behavior could be useful in screening for novel drugs that produce a known and desired neurological effect. This capability would be useful in conjunction with cell- and receptor-based assays for forward-screens in drug and target discovery.

4.4.3 Individual syllables distinguish broad classes of drugs

Innately-driven changes in behavior, such as predator-driven fear behaviors, mate-driven reproductive behaviors, are constructed by highly stereotyped reorganizations of small subsets of existing behaviors. Such ethologically-relevant cues produce rapid and concentrated (both in time, and in terms of total behavioral repertoire affected) changes to overt action. We hypothesized that perhaps changes caused by pharmaceuticals would share this property, and not result in overly diffuse changes to behavior. Concretely, we expected related drugs to modulate the deployment of small and shared subsets of behavioral syllables. We trained a decision forest to distinguish individual drug identities from individual mouse behavioral fingerprints. The classification of the decision forest was lower compared to logistic regression, but its classification decisions were more readily interpretable. The decision forest learns a ”flow chart” for classifying which drug an animal was given, by looking at one dimension of the input at a time, and applying a threshold. Depending on whether the dimension is above or below the learned threshold, either of two paths down the flow chart are followed. We found after learning such a decision tree on our behavioral fingerprint data that a single syllable distinguished all drugs in the stimulant class from all other drugs. Inspecting the raw video data of this syllable, we found it to be a fast forward locomotion syllable, that when used more than 2% of total behavior, indicated strongly that a stimulant had been administered to the animal. All three stimulants were distinguished from the rest of the dataset with a simple usage threshold applied to a single syllable. We note that modafinil, a stimulant, was

distinguished from the rest of the dataset, together with methyphenidate and methamphetamine. While the receptor pharmacology of methylphenidate and methamphetamine are comparatively well-understood and narrow, the molecular mechanism of modafinil on the brain are poorly understood, and affect the neuromodulatory system broadly. The separation of stimulants by a single syllable, and that modafinil is also grouped with stimulants in this task, leads us to two conclusions regarding behavioral syllables. First, that syllables may have some concrete neural substrate, and may be directly instantiated and controlled in the brain, and are real neurological entities. Second, that syllables themselves, may be more desirable as a druggable target than neurotransmitter receptors. That the receptor pharmacology of modafinil is so complicated, while the syllabic behavioral pharmacology revealed here is relatively simple, indicates that characterizing psychoactive pharmaceuticals by their precise behavioral effects, instead of their receptor-binding activity, might lead to new insights and more efficient drug discovery. Behavior is the sole output of the brain, and so a comprehensive and rich summary of locomotor behavior may be reasonably viewed as a more powerful readout than individual receptor affinities for psychoactive compound development.

4.4.4 Developing better behavioral representations are a keystone task for neuroscience and pharmacology

David Marr, in his posthumously published work *Vision* (Marr, 1982), described his ideal approach to studying the neuroscientific mechanisms of vision starting with a primary

prerequisite task: discover and describe good representations of the visual world. With succinct, information-rich descriptions of visual phenomena, then proceed into the brain to study the processes and mechanisms by which neural circuits digest information from the visual world and transmute it into adaptive action. The neuroscientific study of movement is no different – understanding behavior requires representations which enable researchers to either control or predict behavior. In our prior study, we constructed MoSeq to enable the moment-to-moment forecasting of behavior, and in the current study, we show that individual syllables are directly modulated by drugs. Given its utility in both elucidating the sub-second structure of freely-moving behavior, and in summarizing and discriminating the effects of a broad range of behavioral perturbations, MoSeq may serve as a general-purpose representation of an animal's behavior.

4.5 Author contributions

ABW and SRD conceived of the project, designed the experiments and analyses. ABW built the computational pipeline, and MJJ and ABW built the machine learning methodology. ABW, Ralph Peterson (REP) and Jesse Katon (JMK) collected the data. ABW performed all analyses.

Chapter 5

Discussion and Future Directions

5.1 Contribution of Present Work

By fusing novel computer vision, machine learning and behavioral neuroscientific methodologies, this dissertation contains work identifying a novel rhythmicity and structure in mouse behavior with periodicity of *circa* 300 milliseconds. Perturbations in this fine-timescale structure directly reflects a wide variety of experimental manipulations, from near-instantaneous optogenetic manipulation of neural circuits, to pharmaceutical manipulation of the animal's neuromodulatory state, to permanent alteration of an animal's genes.

5.1.1 Building a Computational Model for Ethology

We constructed a model that directly instantiates the core hypothesis of ethology: behavior is composed of reused motifs of action, which we term "behavioral syllables". This model took the form of a Hidden Markov Model (HMM) with special structure – we took care in the definition of behavioral syllables' behavioral content, the duration of each

syllable, their number, and the grammar that governs their transitions. We modeled the observation distributions of individual syllables of the HMM with an autoregressive (AR) process. This is in contrast to the typical use of Gaussian observations for states. The use of Gaussians to describe mouse behavior is equivalent to saying that behavior progresses as the abrupt teleportation between static poses (and those poses have some noise that is completely independent over time). Instead, with the use of AR observation distributions, we posited that individual behavioral syllables are stereotyped trajectories through pose space. We also directly modeled the duration distribution of each syllable, reasoning that some reused behavioral components are more ”sticky” than others and last longer, but also that reused behavioral components would not last for a strictly defined amount of time, but instead have some variability in their duration. This indirectly allowed for the mouse to perform the same action with different amounts of ”vigor” (rearing up more quickly, contracting into a ball slowly), or simply dwell in poses for different durations (grooming behavior can last variable amounts of time, forward locomotion may be short, or last the time it takes to traverse an experimental arena). We also learned the most parsimonious number of behavioral syllables, given the duration of the experiment, and the number of mice observed. It has been posited that the number of uniquely observable behavioral states grows sub-linearly with the time spent observing the animal (Colgan, 1978), and we instantiated this hypothesis directly by positing there are in fact an infinite number of states, but we are only able to observe a finite number of them in a finitely long experiment. The mathematical technique we used for reasoning about infinitely numbered

states with finite data is called a ”Hierarchical Dirichlet Process” (HDP). Last, we expected that syllables transition into each other over time with some regularity – a mouse might pause, then contract, then enter into a rear, then exit the rear, and then extend, and perhaps each of these small actions transition into each other with some reliability across different mice. We modeled the transition structure using an HMM. Practically, this meant that given that a mouse is emitting some single syllable, we can infer which syllable will come next without using any information further in the past. Putting all these modeling decisions together, we created a Sticky Hierarchical Dirichlet Process Autoregressive Hidden Markov Model (Sticky HDP-AR-HMM). This model directly embodied the specific hypotheses regarding the modularity and dynamics of behavior as described above, and will learn all of the free parameters (the specific durations of each state, the number of syllables, their pose dynamics, their transition structure) directly from recorded data using a statistical fitting technique called Gibbs sampling. These techniques individually were not new, but their specific combination, and especially their application to the study of behavior, was, to our knowledge, novel.

5.1.2 Summarizing Behavior With Behavioral Syllables

We fit our model on mouse pose data, resulting in a ”parse”, or ”segmentation” of continuous behavior. Concretely, the only output of the model is, for each frame of behavioral data, a label of what syllable the mouse was emitting. This is akin to the output of computerized speech transcription – each moment of an audio recording of a single person speaking is associated with a single word, and moments adjacent in time are likely

ascribed to the same word. All biological analyses in this dissertation started from this raw material, and summarized it or related it to exogenous variables in some way.

We used this behavioral segmentation to study the effects of biological perturbations on mouse behavior. To do so, we converted our behavioral segmentation, which is of variable length (as long as the experiment lasts), into a ”transition matrix”, or ”behavioral fingerprint”, of fixed length. We calculated the transition matrix, which indicates the probability of transitioning from one syllable into another, for all possible pairs, by counting how often each syllable is followed by another. We constructed an animal's syllabic ”behavioral fingerprint” by simply counting how often each syllable was used by an animal across the length of an experiment.

5.1.3 Behavioral Syllables Are Regulated by Neurological Manipulations At Widely Varying Timescales

Using our ”behavioral fingerprint”, we revealed that context creates statistically-discernible new behaviors. This result is unsurprising, because the animal may interact in novel ways with physically distinct environments. We considered this experiment to be an initial validation of the method. We then performed an experiment where we delivered a smell derived from a fox, the molecule TMT. This ethologically relevant stimulus did not evoke new behaviors from mice. What is classically called ”fear behavior” seemed to be the reorganization of pre-existing, but rare, behavioral syllables. The persistent state of fear was created by the stimulus-driven mutation of behavioral grammar. After receiving the stimulus of the fox odor TMT, once a mouse enters a

”huddled” or ”scrunched up” state (a pose classically associated with mouse fear behavior), it was extremely likely to transition to similar states, and extremely unlikely to transition into locomotive or exploratory states. The brain may thus create adaptive behaviors by shifting transition probabilities between groups of behavioral syllables.

Lifelong experimental perturbations, the manipulation of an animal’s genes, were also evident in the alteration of the usage of sub-second duration behavioral syllables. The Ror1β mouse, made known to us through a collaboration with the Ginty lab and post-doc Victoria Abraira, has a striking ”duck walk” motor phenotype. The exact mechanism of this phenotype is unknown, but the gene is expressed in the brain and spinal cord (Liu et al., 2013; Andre et al., 1998; Masana et al., 2007). We found that a single syllable accounts for the very apparent duck walk phenotype in the homozygous mutant, indicating that our method can effectively compress a complicated, but strong, phenotype, automatically into a parsimonious description. Interestingly, the heterozygous mutant, for which there is no published phenotype, shared other upregulated syllables with the mutant. These upregulated syllables were ”head bobs” and pauses that occur interstitially between the waddling bouts. In the heterozygous mutant, which does not waddle, these upregulated syllables occurred between normal locomotion. Other than these syllables, the heterozygous mutant was completely indistinguishable from wildtype. This is an important and novel finding for translational research, because much of human psychiatric disease manifests as heterozygous mutations, and our methodology is sensitive enough to detect these subtle effects. Instantaneous perturbations, such as optogenetic stimulation of

motor cortex, also manifest in behavioral syllables. We find our method encapsulates the optogenetically-induced (and previously described) ”strong turn” into a single syllable. Thus, our model fully recovers the experimental structure.

Using our behavioral syllable segmentation as a lens into behavior, in Chapter 2 (and the methods supporting it in Chapter 3), we investigated the effect of context, sensory stimuli, genetic alteration and optogenetic neural circuit perturbation on behavior. We found that the deployment of and transition structure between, or ”grammar”, of behavioral syllables is sensitive to both subtle and dramatic experimental alterations.

5.1.4 Directly Targeting Behavior with Pharmaceuticals

In these modeling efforts, experiments and analyses, we developed a model built to accurately forecast mouse behavior. But, we applied it biologically as a way to discriminate between treatments by first destroying much of the information in the segmentation, through the creation of behavioral fingerprints. The modeling approach we took was computationally expensive, and it is possible that, if the goal is simply to discriminate between behavioral perturbations, a simpler or computationally cheaper method would perform just as well. To address this concern directly, we computed alternative summaries of behavior (those classically used in behavioral pharmacology and genetics), such as the distribution of the mouse’s velocity or position in the arena. These classical fingerprints, along with a syllable-based fingerprint, were used to unblind a pharmaceutical treatment of mice. In this task, we designed a machine learning classifier to predict the drug class, drug identity, or drug and dose, given to an individual mouse,

using only a behavioral fingerprint of that mouse's behavior.

We found that syllable-based fingerprints far outperformed any other behavioral fingerprint for unblinding drug class. No other fingerprint could successfully unblind all drug identities, but surprisingly, syllable-based fingerprints were also able to successfully predict the drug identity and its dose, on a mouse-per-mouse basis. Although built for forecasting moment-to-moment behavior, our behavioral segmentation method yields the most information-rich summaries of behavior, with applicability to instantaneous neurological, brief ethological and pharmaceutical, and lifetime genetic manipulations of behavior.

Further, we observed that using behavioral syllables, we could relate drugs with similar behavioral effects in humans. For instance, the drugs modafinil (or Provigil, which is used to keep fighter pilots awake on long flights) and methamphetamine were closely related in our analysis. Modafinil's receptor pharmacology is currently a mystery, in contrast with methamphetamine, which has been well-characterized and studied for many decades. Through receptor binding-affinity analysis, these two drugs appear quite dissimilar, yet they cause similar behavioral effects. Drugs are typically described and developed in the context of their properties on the scale of proteins, but they succeed and have societal impact on the scale of animal behavior. Given that our analysis is capable of relating pharmacologically-dissimilar, yet behaviorally-similar, drugs, we reason that perhaps behavioral syllables may themselves be directly druggable targets.

5.2 Future Directions

5.2.1 Neural Substrates of Behavioral Syllables

The hyper-ultradian, sub-second timescale described in Chapter 2 and exploited in our modeling efforts is certainly one of many behavioral rhythms, among ultradian, circadian, monthly, seasonal, yearly and lifetime periodicities. There are also certainly faster behavioral rhythms – whisking can occur at up to 20 Hz (Voigts et al., 2008), and the mouse's average heart rate is 10 Hz (Davies & Morris, 1993). The 300 millisecond duration corresponds roughly to the length of a mouse step (Clarke & Still, 1999), the duration of a single average mouse breath (Davies & Morris, 1993), the duration of spike transmission through the cortico-striatal loop (Oldenburg & Sabatini, 2015), and delta-band LFP oscillations.

Although obviously all locomotor behavior has a neural substrate (at the very least, the muscular activations that directly modify skeletal pose), we speculate that perhaps behavioral syllables may be not just useful fractionations of observed action, but also may have direct neural representations. Prior work has found that in flies and worms, individual neurons can regulate the deployment, duration and transition between specific elements of finite behavioral "dictionaries" (Berg et al., 1972; Garrity et al., 2010; Gray et al., 2005; Pierce-Shimomura et al., 1999; Vogelstein et al., 2014). In the songbird, individual song syllables (we directly borrow this term "syllable" from the songbird literature for our mouse behavioral syllables) have concrete neural representations of their

identity, duration, and patterning over time (Albert & Margoliash, 1996; Hahnloser et al., 2002; Olveczky et al., 2005; Berwick et al., 2011; Markowitz et al., 2013). Further, there has been moderate success in identifying reorderable elements of behavior in the rat, through disruption of grooming behavior via striatal lesioning (Berridge et al., 1987). Thus, there is precedent for the brain representing behavior at the high level of a reused ”module”.

The work of (Bizzi et al., 1995; d’Avella & Bizzi, 1998; Bizzi et al., 2000) has shown that, in the spinal cord, the flexion and tension of ”synergies” of muscles are coincident, a neurological strategy that reduces the dimensionality of behavior. In this way, the spinal cord presents to the brain a greatly simplified interface for the body, that helps mitigate the complexity of puppeteering each muscle in the body through the environment to reach goals. Thus, when the nervous system enacts a particular adaptive behavior, the brain does not need to build and manipulate representations of individual muscles.

Given the utility of behavioral syllables in discriminating direct changes to the animal’s neural activity or neuromodulatory state, their suggestive neural timescale, and the existence of neural representations of behavioral modules in other organisms, we anticipate the brain may be able to construct and emit complex behavior through the use of representations of behavioral syllables, an even higher level of abstraction than muscle synergies. Such a representation is advantageous in its parsimony, and thus its metabolic cost.

To record evidence of neural firing coincident with and perhaps coding for behavioral syllables requires the solution to a few technical hurdles. For instance, one would need to

achieve synchronization of both the electrophysiological behavioral recording modalities, and develop the proper computer vision tracking of a mouse with an implanted headstage. However, we do not believe any of the challenges are insurmountable given current technology. In line with prior literature in the rat and songbird, we anticipate that neural representations of individual syllables (or any of their properties, such as their duration, vigor, or transition structure) will be most likely found in motor cortex (Graziano, 2016) or dorsal striatum (Aldridge et al., 2004; Kravitz & Kreitzer, 2012).

If correlations are found between single-neuron firing, or local field potential activity, and unlearned behavioral syllables, we will be curious to see whether such activity leads or lags actual emitted behavior in time. Are syllabic representations of behavior useful for behavioral planning, and thus causal of behavior, or for behavioral modification, and thus more likely to be useful as an efference copy in the nervous system?

5.2.2 Opportunities in Learned Behavior

In this dissertation, we have only used behavioral syllables to analyze unlearned behavior. However, they certainly could have use in the context of learning and memory. The neuroscientific study of the relation of neural activity to learned behavior proceeds first by imposing restraints on the animal. These restraints are used in aid of obtaining high-quality measurements of both the behavior and neural activity, and can be of a physical or psychological nature. Physical restraints, such as head-fixing an animal, can aid in e.g. the two-photon microscopy of brain tissue (Helmchen & Denk, 2005), or in precise measurement of whisker location (Voigts et al., 2008). Psychological restraint,

such as over-training on an operant task, provides precise moments in time to which valence may be attached (if the task is such that the animal is working towards achieving a reward or avoiding a punishment), with which experimenters may align the firing of neurons. In this way, analysis of neural activity may yield whether or not some area of the brain is more or less responsive to, or ”codes” for, some precisely-measured events in the world. Without these restraints, statements such as ”dopamine neurons appear to emit an alerting message about the surprising presence or absence of rewards” (Schultz, 1998) are impossible, because dopamine neuron activity must be recorded precisely and concomitantly with reward events. Thus, freely-moving, or untrained behavior is used sparingly in the behavioral and systems neurosciences when temporal precision is required.

However, the use of behavioral syllables not only provides events upon which to align neural data, but these alignment opportunities are nearly continuous in time. Rigid trial structure is not required to detect whether a neuron fires coincidentally with the onset of a particular behavior – one need simply to look up all the onset times of each behavioral syllable, and align each neuron’s firing activity to those time events. Statistical care is of course required, given the number of alignment opportunities in a typical behavioral dataset (for a 20 minute dataset, recorded at 30 frames/second, yielding 36,000 frames, the mouse on average transitions through 3,600 behavioral syllables, drawn from a set of $\sim$60 unique syllables, see Chapter 2).

Currently, a behavioral dataset is parsed into its constituent syllables after the

experiment has ended. However, with some engineering effort, one may imagine ”closing the loop”, and inferring which syllables an animal is emitting live, during the experiment. One may then design experiments directly testing whether a syllable may be reinforced, for instance through optogenetic stimulation of dopaminergic nuclei in the brain at the moment of syllabic emission (Lammel et al., 2012; Danjo et al., 2014). One may then investigate whether behaviors are uniformly reinforceable, or if some, such as mating or grooming behaviors (Aldridge & Berridge, 1998), occur in difficult-to-modify syntactic chains, and are incapable of being reinforced (Breland & Breland, 1961).

5.2.3 Completing Our View of Musculoskeletal Behavior

The animal’s 3D body pose is the raw material used to create our segmentation of free movement into behavioral syllables. We note, however, that here we have imaged only the top of the animal at 30 frames/second. The advent of affordable depth cameras and cheap computation were an enabling technology for obtaining even a partial view of the animal’s 3D pose. However, the technologies we exploited (off-the-shelf Microsoft Kinect depth cameras) were not initially designed with the task of mouse phenotyping in mind. In the future, we anticipate that bespoke imaging and biological technologies could increase the accuracy of the technique, and enable further discoveries in quantitative ethology.

The imaging modality of the initial version of the Microsoft Kinect, which was the only sensor used in this dissertation, used structured light, or a ”constellation” of infrared-band dots presented at an angle to the imaging viewpoint, to triangulate the distance of every visible location in the scene. The closer an object moved to the camera, the more one

section of the dot pattern shifted to the right. This depth sensing modality, although returned as a moderately-sized image (640 x 480 pixels), had an effective resolution of half that. Much of the innovation of our computer vision algorithms revolved around accounting for the non-uniform noise characteristics of the sensor. For instance, if one pixel was labeled as ”missing”, very often nearby pixels would be missing as well, rendering the noise extremely non-uniform and non-Gaussian. Traditional image denoising strategies require this assumption, which necessitated the development of novel techniques.

In ongoing experiments, we have deployed the second version of the Microsoft Kinect, which uses a separate depth sensing modality, termed ”time-of-flight” (ToF). ToF sensors image depth truly on a pixel-per-pixel basis. An extremely brief infrared light pulse is sent out to the scene, and each pixel acts as ”clock”, counting the nanoseconds until a photon returns to the pixel. The ToF of the returning photon can easily be converted to distance within some working range (the clocks in each pixel are oscillators, and thus can only measure the ToF of a photon modulo their periodicity. This oscillation period directly translates into the depth of each pixel being placed within a fixed range of distances). The camera sends out 30 such pulses per second, yielding a 30 frame/second depth video. The accuracy and reliability of this camera is much higher, with the caveat that it is more sensitive to reflective surfaces (the sensor cannot distinguish between a photon that returned directly from the mouse’s body, or one that bounced off the body, then off a shiny surface, and then reached the sensor). This problem is easily alleviated with the addition

of non-reflective coatings on behavioral apparatuses, such as those used on stealth vehicles, or on the insides of astronomy telescopes.

Both versions of the Kinect record effectively the top half of the mouse's body, and return no information about the mouse's feet. This makes reasoning about the animal's full skeletal pose nearly impossible. We have experimented with touch-screens using frustrated total internal reflectance (FTIR) technology, which yields paw-, and even toe-level, resolution. Combining the touch screen with the Kinect imaging might yield enough information to computationally infer the mouse's body pose. Early work, which we hope to extend, has indicated that the use of particle filters and a custom 3D rasterizer would be successful in allowing us to computationally infer the full pose of the animal.

One may imagine a direct route to obtaining the full 3D skeletal pose of the animal, by using video X-ray cameras, in a stereo configuration, to directly observe the skeleton of the animal in such a way that the 3D position of each joint may be calculated. The amount of radiation required to achieve these images may limit the duration of the experiment substantially, but such data streams, if collected together with other imaging modalities, may provide concrete "ground truth" datasets for inferring the full skeletal configuration of the animal.

A more distant goal is to directly obtain the muscular flexion and tension state of the animal, which is a more direct measurement of the behavioral actuations the brain may achieve (the brain has no direct control over the position of bones, but may tense and relax muscles to indirectly achieve a desired pose or pose trajectory). With some ideal and

hypothetical technology, an experimenter could record the state of every muscle in the mouse's body, at an electrophysiological sampling rate at or near the refractory period of motor neurons. In practice, muscles tense and flex in stereotyped combinations (Bizzi et al., 1995), so the entire system may perhaps be inferred by the recording of strategically chosen muscles or muscle groups. It is perhaps possible to induce and record activity-gated fluorescence of such carefully chosen muscles using genetically-encoded indicators. If the muscles and the fluorescence spectra are chosen carefully, one might imagine a low-resolution color recording of a mouse, where the color intensities of such a movie are deconvolved to then infer the full pose and muscular state of the animal.

If any meaningful steps are taken towards recording the full 3D pose of the animal, a challenge to our current definition of behavioral syllables will rise quickly – under our model, the animal is presently said to only emit one syllable at a time. However, a mouse may run and sniff simultaneously, just as we can walk and brush our teeth at the same time. If we were to assign a unique syllable to each combination of limb-based behaviors, the possible configurations would combinatorially explode. Thus, a factorial representation of behavioral syllables will be required as sensor technology develops. The mathematics for such a model does exist (Ghahramani & Jordan, 1997), but would have to be integrated with our current system.

5.2.4 Parsing Social Behavior

Mice are innately social creatures, and have evolved complicated behavioral and chemical strategies for managing the collaborative aspects of propagating their genes, finding food

and avoiding predators. In future work, we hope to extend the ideas presented here to social behavior. However, there are difficult technical barriers do doing so. First, the computer vision algorithms required to consistently and reliably track the identities of multiple mice as they interact are presently under-developed, and it is possible that the development of new sensors, instead of algorithms, is a necessary prerequisite. For instance, the fighting or reproductive pursuit behaviors of mice occur extremely rapidly, and are likely too fast for the resolution of the current system. A stop-gap solution might be physically segregating mice, but allowing them to interact chemically via nose-touches through a ”fence”. In this way, each mouse may be tracked independently, but still be physically and chemically aware of other mice, and thus engage in social behaviors.

Should the computer vision and sensor problems be solved, there are still open machine learning problems in the context of understanding social behavior. Instead of trying to parse a single stream of action into behavioral syllables, and describing the grammar of behavior, the challenge is now to parse a behavioral ”conversation”. The most active thread of research in developing models of social behavior is work underway to automatically discover strategies in NBA basketball games, from the raw data of the position of all ten players and the ball, recorded at 30 frames/second. Although the application might seem unrelated, the models under development are extensions and embellishments of those used in this dissertation.

5.2.5 Application to Other Species

The broadest horizon for this work is the extension to species other than the mouse. We note that there is no fundamental barrier preventing us from applying this technique in humans. The underlying statistical engine – the souped-up Hidden Markov Model that parses behavioral time-series – is agnostic to the source of the data it segments. This dissertation developed computer vision techniques for extracting mouse pose data from a depth data stream, and indeed, these efforts are specific to mice. However, tracking the pose of humans has historically received a great deal more research and investment attention, and thus a researcher interested in investigating the behavioral syllables of humans would be able to take advantage of this prior work.

We have, in a very preliminary fashion, applied this work both to rats (using the 1st-generation Kinect), and also to humans (using wrist-attached accelerometers, not pose data), but the research is still quite early. Currently, work by Scott Linderman in the lab of Ryan Adams is extending our approach to larval zebrafish. Being a simpler organism, it is perhaps unsurprising that the animal's behavior can be effectively captured by many fewer states than in the mouse. We anticipate that, in the human, there will be many more distinct behavioral syllables than we find in the mouse. However, it is likely that extensions to our modeling approach will be required for in-depth understanding of human behavioral syllables. As mentioned above, it is apparent that humans can easily perform more than one distinct action at a time, in a large number of possible combinations, like texting and walking, or sipping coffee and skipping to the end of a thesis. Thus, extensions of our

model to allow for overlapping or composable behavioral syllables might be a prerequisite for meaningful behavioral predictions.

References

Albert, C. Y. & Margoliash, D. (1996). Temporal hierarchical control of singing in birds. *Science*, 273(5283), 1871–1875.

Aldridge, J. W. & Berridge, K. C. (1998). Coding of serial order by neostriatal neurons: a "natural action" approach to movement sequence. *The Journal of Neuroscience*, 18(7), 2777–2787.

Aldridge, J. W., Berridge, K. C., & Rosen, A. R. (2004). Basal ganglia neural mechanisms of natural movement sequences. *Canadian journal of physiology and pharmacology*, 82(8-9), 732–739.

Anderson, D. J. & Perona, P. (2014). Toward a science of computational ethology. *Neuron*, 84(1), 18–31.

Andre, E., Conquet, F., Steinmayr, M., Stratton, S. C., Porciatti, V., & Becker-André, M. (1998). Disruption of retinoid-related orphan receptor β changes circadian behavior, causes retinal degeneration and leads to vacillans phenotype in mice. *The EMBO journal*, 17(14), 3867–3877.

Arias-Castro, E. & Donoho, D. L. (2009). Does median filtering truly preserve edges better than linear filtering? *The Annals of Statistics*, (pp. 1172–1206).

Basseville, M., Nikiforov, I. V., et al. (1993). *Detection of abrupt changes: theory and application*, volume 104. Prentice Hall Englewood Cliffs.

Berg, H. C., Brown, D. A., et al. (1972). Chemotaxis in escherichia coli analysed by three-dimensional tracking. *Nature*, 239(5374), 500–504.

Berman, G. J., Choi, D. M., Bialek, W., & Shaevitz, J. W. (2014). Mapping the stereotyped behaviour of freely moving fruit flies. *Journal of The Royal Society Interface*, 11(99), 20140672.

Bernardo, J. M. & Smith, A. F. (2009). *Bayesian Theory*, volume 405.

Berridge, K. C., Fentress, J. C., & Parr, H. (1987). Natural syntax rules control action sequence of rats. *Behavioural Brain Research*, 23(1), 59–68.

Berwick, R. C., Okanoya, K., Beckers, G. J., & Bolhuis, J. J. (2011). Songs to syntax: the linguistics of birdsong. *Trends in cognitive sciences*, 15(3), 113–121.

Bishop, C. M. (2006a). *Pattern recognition and machine learning*. springer.

Bishop, C. M. (2006b). *Pattern Recognition and Machine Learning*. Springer.

Bizzi, E., Giszter, S. F., Loeb, E., Mussa-Ivaldi, F. A., & Saltiel, P. (1995). Modular organization of motor behavior in the frog's spinal cord. *Trends in neurosciences*, 18(10), 442–446.

Bizzi, E., Tresch, M. C., Saltiel, P., & d'Avella, A. (2000). New perspectives on spinal motor systems. *Nature Reviews Neuroscience*, 1(2), 101–108.

Breland, K. & Breland, M. (1961). The misbehavior of organisms. *American psychologist*, 16(11), 681.

Brown, A. E., Yemini, E. I., Grundy, L. J., Jucikas, T., & Schafer, W. R. (2013). A dictionary of behavioral motifs reveals clusters of genes affecting caenorhabditis elegans locomotion. *Proceedings of the National Academy of Sciences*, 110(2), 791–796.

Brown, T. G. (1911). The intrinsic factors in the act of progression in the mammal. *Proceedings of the Royal Society of London. Series B, containing papers of a biological character*, 84(572), 308–319.

Castagné, V., Porsolt, R. D., & Moser, P. (2006). Early behavioral screening for antidepressants and anxiolytics. *Drug development research*, 67(9), 729–742.

Clarke, K. & Still, J. (1999). Gait analysis in the mouse. *Physiology & behavior*, 66(5), 723–729.

Colgan, P. W. (1978). *Quantitative ethology*. Wiley.

Crawley, J. N. (2003). Behavioral phenotyping of rodents. *Comparative medicine*, 53(2), 140–146.

Crawley, J. N. (2007). *What's wrong with my mouse?: behavioral phenotyping of transgenic and knockout mice*. John Wiley & Sons.

Croll, N. A. (1975). Components and patterns in the behaviour of the nematode caenorhabditis elegans. *Journal of zoology*, 176(2), 159–176.

Danjo, T., Yoshimi, K., Funabiki, K., Yawata, S., & Nakanishi, S. (2014). Aversive behavior induced by optogenetic inactivation of ventral tegmental area dopamine neurons is mediated by dopamine d2 receptors in the nucleus accumbens. *Proceedings of the National Academy of Sciences*, 111(17), 6455–6460.

Dankert, H., Wang, L., Hoopfer, E. D., Anderson, D. J., & Perona, P. (2009). Automated monitoring and analysis of social behavior in drosophila. *Nature methods*, 6(4), 297–303.

d'Avella, A. & Bizzi, E. (1998). Low dimensionality of supraspinally induced force fields. *Proceedings of the National Academy of Sciences*, 95(13), 7711–7714.

Davies, B. & Morris, T. (1993). Physiological parameters in laboratory animals and humans. *Pharmaceutical research*, 10(7), 1093–1095.

de Chaumont, F., Coura, R. D.-S., Serreau, P., Cressant, A., Chabout, J., Granon, S., & Olivo-Marin, J.-C. (2012). Computerized video analysis of social interactions in mice. *Nature methods*, 9(4), 410–417.

Drai, D., Benjamini, Y., & Golani, I. (2000). Statistical discrimination of natural modes of motion in rat exploratory behavior. *Journal of neuroscience methods*, 96(2), 119–131.

Eppig, J. T., Blake, J. A., Bult, C. J., Kadin, J. A., Richardson, J. E., Group, M. G. D., et al. (2015). The mouse genome database (mgd): facilitating mouse as a model for human biology and disease. *Nucleic acids research*, 43(D1), D726–D736.

Fendt, M., Endres, T., Lowry, C. A., Apfelbach, R., & McGregor, I. S. (2005). Tmt-induced autonomic and behavioral changes and the neural basis of its processing. *Neuroscience & Biobehavioral Reviews*, 29(8), 1145–1156.

Fentress, J. C. & Stilwell, F. P. (1973). Grammar of a movement sequence in inbred mice. *Nature*, 244, 52–53.

Flash, T. & Hochner, B. (2005). Motor primitives in vertebrates and invertebrates. *Current opinion in neurobiology*, 15(6), 660–666.

Fox, E., Jordan, M. I., Sudderth, E. B., & Willsky, A. S. (2009). Sharing features among dynamical systems with beta processes. In *Advances in Neural Information Processing Systems* (pp. 549–557).

Fox, E. B. (2009). *Bayesian Nonparametric Learning of Complex Dynamical Phenomena*. Ph.D. thesis, MIT, Cambridge, MA.

Fox, E. B., Sudderth, E. B., Jordan, M. I., & Willsky, A. S. (2008). An hdp-hmm for systems with state persistence. In *Proceedings of the 25th international conference on Machine learning* (pp. 312–319).: ACM.

Garcia, V. A., Junior, C. F. C., & Marino-Neto, J. (2010). Assessment of observers' stability and reliability—a tool for evaluation of intra-and inter-concordance in animal behavioral recordings. In *Engineering in Medicine and Biology Society (EMBC), 2010 Annual International Conference of the IEEE* (pp. 6603–6606).: IEEE.

Garrity, P. A., Goodman, M. B., Samuel, A. D., & Sengupta, P. (2010). Running hot and cold: behavioral strategies, neural circuits, and the molecular machinery for thermotaxis in c. elegans and drosophila. *Genes & development*, 24(21), 2365–2382.

Gelman, A., Carlin, J. B., Stern, H. S., & Rubin, D. B. (2014). *Bayesian data analysis*, volume 2. Taylor & Francis.

Ghahramani, Z. & Jordan, M. I. (1997). Factorial hidden markov models. *Machine learning*, 29(2-3), 245–273.

Glickfeld, L. L., Andermann, M. L., Bonin, V., & Reid, R. C. (2013). Cortico-cortical projections in mouse visual cortex are functionally target specific. *Nature neuroscience*, 16(2), 219–226.

Gray, J. M., Hill, J. J., & Bargmann, C. I. (2005). A circuit for navigation in caenorhabditis elegans. *Proceedings of the National Academy of Sciences of the United States of America*, 102(9), 3184–3191.

Graziano, M. S. (2016). Ethological action maps: A paradigm shift for the motor cortex. *Trends in cognitive sciences*, 20(2), 121–132.

Hahnloser, R. H., Kozhevnikov, A. A., & Fee, M. S. (2002). An ultra-sparse code underliesthe generation of neural sequences in a songbird. *Nature*, 419(6902), 65–70.

Hastie, T., Tibshirani, R., Friedman, J., & Franklin, J. (2005). The elements of statistical learning: data mining, inference and prediction. *The Mathematical Intelligencer*, 27(2), 83–85.

Helmchen, F. & Denk, W. (2005). Deep tissue two-photon microscopy. *Nature methods*, 2(12), 932–940.

Hunter, J. D. et al. (2007). Matplotlib: A 2d graphics environment. *Computing in science and engineering*, 9(3), 90–95.

Ishwaran, H. & Zarepour, M. (2002). Exact and approximate sum representations for the dirichlet process. *Canadian Journal of Statistics*, 30(2), 269–283.

Jhuang, H., Garrote, E., Yu, X., Khilnani, V., Poggio, T., Steele, A. D., & Serre, T. (2010). Automated home-cage behavioural phenotyping of mice. *Nature communications*, 1, 68.

Kabra, M., Robie, A. A., Rivera-Alba, M., Branson, S., & Branson, K. (2013). Jaaba: interactive machine learning for automatic annotation of animal behavior. *nature methods*, 10(1), 64–67.

Koller, D. & Friedman, N. (2009). *Probabilistic graphical models: principles and*

techniques. MIT press.

Konopka, R. J. & Benzer, S. (1971). Clock mutants of drosophila melanogaster. *Proceedings of the National Academy of Sciences*, 68(9), 2112–2116.

Koren, Y. (2005). Drawing graphs by eigenvectors: theory and practice. *Computers & Mathematics with Applications*, 49(11), 1867–1888.

Kravitz, A. V. & Kreitzer, A. C. (2012). Striatal mechanisms underlying movement, reinforcement, and punishment. *Physiology*, 27(3), 167–177.

Lammel, S., Lim, B. K., Ran, C., Huang, K. W., Betley, M. J., Tye, K. M., Deisseroth, K., & Malenka, R. C. (2012). Input-specific control of reward and aversion in the ventral tegmental area. *Nature*, 491(7423), 212–217.

Lashley, K. S. (1951). *The problem of serial order in behavior*. Bobbs-Merrill.

Liu, H., Kim, S.-Y., Fu, Y., Wu, X., Ng, L., Swaroop, A., & Forrest, D. (2013). An isoform of retinoid-related orphan receptor β directs differentiation of retinal amacrine and horizontal interneurons. *Nature communications*, 4, 1813.

Mallat, S. G. (1989). A theory for multiresolution signal decomposition: the wavelet representation. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 11(7), 674–693.

Markowitz, J. E., Ivie, E., Kligler, L., & Gardner, T. J. (2013). Long-range order in canary song. *PLoS Comput Biol*, 9(5), e1003052.

Marr, D. (1982). Vision: A computational investigation into the human representation and processing of visual information, henry holt and co. *Inc., New York, NY*, 2.

Masana, M. I., Sumaya, I. C., Becker-Andre, M., & Dubocovich, M. L. (2007). Behavioral characterization and modulation of circadian rhythms by light and melatonin in c3h/hen mice homozygous for the rorβ knockout. *American Journal of Physiology-Regulatory, Integrative and Comparative Physiology*, 292(6), R2357–R2367.

McKinney, W. et al. (2010). Data structures for statistical computing in python. In *Proceedings of the 9th Python in Science Conference*, volume 445 (pp. 51–56).

Murphy, K. P. (2012). *Machine learning: a probabilistic perspective*. MIT press.

Ohayon, S., Avni, O., Taylor, A. L., Perona, P., & Egnor, S. R. (2013). Automated multi-day tracking of marked mice for the analysis of social behaviour. *Journal of neuroscience methods*, 219(1), 10–19.

Oldenburg, I. A. & Sabatini, B. L. (2015). Antagonistic but not symmetric regulation of primary motor cortex by basal ganglia direct and indirect pathways. *Neuron*, 86(5), 1174–1181.

Olveczky, B. P., Andalman, A. S., & Fee, M. S. (2005). Vocal experimentation in the juvenile songbird requires a basal ganglia circuit. *PLoS Biol*, 3(5), e153.

Ou-Yang, T.-H., Tsai, M.-L., Yen, C.-T., & Lin, T.-T. (2011). An infrared range camera-based approach for three-dimensional locomotion tracking and pose reconstruction in a rodent. *Journal of neuroscience methods*, 201(1), 116–123.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al. (2011). Scikit-learn: Machine learning in python. *The Journal of Machine Learning Research*, 12, 2825–2830.

Pérez, F. & Granger, B. E. (2007). Ipython: a system for interactive scientific computing. *Computing in Science & Engineering*, 9(3), 21–29.

Pierce-Shimomura, J. T., Morse, T. M., & Lockery, S. R. (1999). The fundamental role of pirouettes in caenorhabditis elegans chemotaxis. *The journal of neuroscience*, 19(21), 9557–9569.

Porsolt, R. D., Anton, G., Blavet, N., & Jalfre, M. (1978). Behavioural despair in rats: a new model sensitive to antidepressant treatments. *European journal of pharmacology*, 47(4), 379–391.

Rihel, J., Prober, D. A., Arvanites, A., Lam, K., Zimmerman, S., Jang, S., Haggarty, S. J., Kokel, D., Rubin, L. L., Peterson, R. T., et al. (2010). Zebrafish behavioral profiling links drugs to biological targets and rest/wake regulation. *Science*, 327(5963), 348–351.

Rihel, J. & Schier, A. F. (2012). Behavioral screening for neuroactive drugs in zebrafish. *Developmental neurobiology*, 72(3), 373–385.

Roberds, S. L., Filippov, I., Alexandrov, V., Hanania, T., & Brunner, D. (2011). Rapid, computer vision-enabled murine screening system identifies neuropharmacological potential of two new mechanisms. *Frontiers in neuroscience*, 5.

Schultz, W. (1998). Predictive reward signal of dopamine neurons. *Journal of neurophysiology*, 80(1), 1–27.

Sherrington, C. S. (1916). *The integrative action of the nervous system*. CUP Archive.

Spink, A., Tegelenbosch, R., Buma, M., & Noldus, L. (2001). The ethovision video tracking system–a tool for behavioral phenotyping of transgenic mice. *Physiology & behavior*, 73(5), 731–744.

Stephens, G. J., de Mesquita, M. B., Ryu, W. S., & Bialek, W. (2011). Emergence of long timescales and stereotyped behaviors in caenorhabditis elegans. *Proceedings of the National Academy of Sciences*, 108(18), 7286–7289.

Stephens, G. J., Johnson-Kerner, B., Bialek, W., & Ryu, W. S. (2008). Dimensionality and dynamics in the behavior of c. elegans. *PLoS Comput Biol*, 4(4), e1000028.

Swanson, L. W. (2000). Cerebral hemisphere regulation of motivated behavior. *Brain research*, 886(1), 113–164.

Takahashi, J. S., Pinto, L. H., & Vitaterna, M. H. (1994). Forward and reverse genetic approaches to behavior in the mouse. *Science (New York, NY)*, 264(5166), 1724.

Teh, Y. W., Jordan, M. I., Beal, M. J., & Blei, D. M. (2006). Hierarchical Dirichlet processes. *Journal of the American Statistical Association*, 101(476), 1566–1581.

Thorpe, W. H. (1958). The learning of song patterns by birds, with especial reference to the song of the chaffinch fringilla coelebs. *Ibis*, 100(4), 535–570.

Tinbergen, N. (1951). The study of instinct.

Tresch, M. C. & Jarc, A. (2009). The case for and against muscle synergies. *Current opinion in neurobiology*, 19(6), 601–607.

Tresch, M. C., Saltiel, P., & Bizzi, E. (1999). The construction of movement by the spinal cord. *Nature neuroscience*, 2(2), 162–167.

Van Der Walt, S., Colbert, S. C., & Varoquaux, G. (2011). The numpy array: a structure for efficient numerical computation. *Computing in Science & Engineering*, 13(2), 22–30.

Vitaterna, M. H., King, D. P., Chang, A.-M., Kornhauser, J. M., Lowrey, P. L., McDonald, J. D., Dove, W. F., Pinto, L. H., Turek, F. W., & Takahashi, J. S. (1994). Mutagenesis and mapping of a mouse gene, clock, essential for circadian behavior. *Science*, 264(5159), 719–725.

Vogelstein, J. T., Park, Y., Ohyama, T., Kerr, R. A., Truman, J. W., Priebe, C. E., & Zlatic, M. (2014). Discovery of brainwide neural-behavioral maps via multiscale unsupervised structure learning. *Science*, 344(6182), 386–392.

Voigts, J., Sakmann, B., & Celikel, T. (2008). Unsupervised whisker tracking in unrestrained behaving animals. *Journal of neurophysiology*, 100(1), 504–515.

Wallace, K. J. & Rosen, J. B. (2000). Predator odor as an unconditioned fear stimulus in rats: elicitation of freezing by trimethylthiazoline, a component of fox feces. *Behavioral neuroscience*, 114(5), 912.

Wasserman, L. (2013). *All of statistics: a concise course in statistical inference*. Springer Science & Business Media.

Welch, P. D. (1967). The use of fast fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms. *IEEE Transactions*

on audio and electroacoustics, 15(2), 70–73.

Wiener, N. (1949). *Extrapolation, interpolation, and smoothing of stationary time series*, volume 2. MIT press Cambridge, MA.

Wiltschko, A. B., Johnson, M. J., Iurilli, G., Peterson, R. E., Katon, J. M., Pashkovski, S. L., Abraira, V. E., Adams, R. P., & Datta, S. R. (2015). Mapping sub-second structure in mouse behavior. *Neuron*, 88(6), 1121–1135.