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

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

---

Doctoral Dissertations, 2011-2026 Neuroscience

2024

Behavioral Dynamics and Neural Computation for Sensory Navigation in C. elegans

Chen, Kevin S

Princeton, NJ : Princeton University

---

http://arks.princeton.edu/ark:/88435/dsp01bv73c382q

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

Behavioral Dynamics and Neural Computation for Sensory Navigation in *C. elegans*

Kevin S. Chen

A DISSERTATION
PRESENTED TO THE FACULTY
OF PRINCETON UNIVERSITY
IN CANDIDACY FOR THE DEGREE
OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE
BY THE
PRINCETON NEUROSCIENCE INSTITUTE

ADVISERS: PROFESSOR ANDREW M. LEIFER & PROFESSOR JONATHAN W. PILLOW

NOVEMBER 2024

© Copyright by Kevin S. Chen, 2024.
All rights reserved.

Abstract

Sensory navigation is crucial for survival and is a common trait across species of all scales. Studying navigation strategies in animals strengthens our understanding of behavioral algorithms and the underlying neural computation. In this thesis, we investigated sensory navigation in the nematode worm *C. elegans* by combining a novel experimental setup, statistical modeling, neural perturbation and measurements. The integrative approach provides scientific insight into the worm's navigation strategies, learning-dependent behavior, and the underlying neural mechanisms. In the following three chapters, we first built an odor flow chamber and sensory array to control and monitor airborne odor environments, enabling the study of navigation in small animals, such as *C. elegans* and *Drosophila* larvae. Secondly, we measured learning-dependent odor-guided navigation in worms and developed a novel statistical model to analyze detailed navigation trajectories. We characterized how learning alters the sensorimotor transformation in worms, demonstrate that the statistical model outperforms classic chemotaxis metric, and revealed distributed neural mechanisms through optogenetic perturbation and genetic ablation. Thirdly, we examined behavioral dynamics along the worm's chemotaxis trajectory by fitting an extended model with latent states, forming an input-output hidden Markov model for navigation. We found that this state-dependent model better describes the data, reveals complex state-dependent navigation strategies, and pinpoints neural circuits involved in the behavioral dynamics. Additionally, in the last two chapters, we introduced algorithms for stochastic network dynamics that generate state-dependent behavior and analyzed the underlying dynamics. Theoretical projects include a latent variable model for signal progressing in neural imaging, a dynamical model that captures adaptive neural response, and a theoretical framework to analyze nonequilibrium networks. Together, this thesis demonstrates the integration of precise measurements, statistical models, and theory-driven experiments, providing insights into animal navigation algorithms and neural computation.

Contents

Abstract **3**

1 Introduction **10**
- 1.1 Background in sensory navigation . . . 11
- 1.2 Background in neural computation . . . 13
- 1.3 Background in *C. elegans* neuroscience . . . 14
- 1.4 Thesis overview . . . 15

2 Controlling and Monitoring Odor Profile to Study Olfactory Navigation in Small Animals **18**
- 2.1 Introduction . . . 19
- 2.2 Results . . . 22
- 2.3 Discussion . . . 42

3 Learning-dependent Navigation Strategies and Neural Computations in *C. elegans* **47**
- 3.1 Introduction . . . 48
- 3.2 Results . . . 50
- 3.3 Discussion . . . 65

4 State-switching Chemotaxis in *C. elegans* **69**
- 4.1 Introduction . . . 70
- 4.2 Results . . . 72
- 4.3 Discussion . . . 87

5 Learning Algorithms and State-dependent Dynamics in Stochastic Neural Networks **91**
- 5.1 Stochastic circuit models for behavioral response in *C. elegans* . . . 92
- 5.2 Measurements and models for whole-brain neural dynamics . . . 102
- 5.3 Learning algorithms for stochastic spiking recurrent networks . . . 110

6 Models for Motion Artifact in Imaging, Predictive Dynamics in Visual Cortex, and Inference for Nonequilibrium Networks **123**
- 6.1 Latent variable model for motion correction in neural imaging . . . 124
- 6.2 Adaptive subunit model for predictive dynamics in mouse visual cortex . . . 133
- 6.3 Coding, dynamics, and inference in nonequilibrium networks . . . 146

7 Conclusion and Outlook **172**
- 7.1 Conclusions . . . 173

7.2 Challenges, ongoing and future work 177

Appendix A Appendix **179**
A.1 Odor-flow chamber project 179
A.2 Olfactory learning project 182
A.3 State-switching chemotaxis project 185
A.4 Alternative formalism for state-space dynamics in RNNs 189
A.5 Repository and resource for data and methods 191

References **212**

To my family, friends, and mentors along the journey.

Acknowledgments

I would never want to leave Princeton. In some sense, I would always want to remain a graduate student. This feeling towards graduation summarizes how much I appreciate and enjoy every minute in the past few years at Princeton. Please don't get me wrong – like any normal senior graduate student, there're plenty tiresome moments and cynicism throughout the process. However, when I zoom out a little bit, what remains is immense gratitude for everyone who supported me along the journey.

First, I want to thank my advisors, Andy Leifer and Jonathan Pillow, for their guidance throughout my PhD. As someone who has never seen the worm or worked with statistical inference initially, Andy and Jonathan truly brought me on an exciting journey of learning about both and combining them. In particular, I want to thank Andy for providing such a resourceful lab environment to explored and conduct hands-on studies. I credit Andy for the strong empiricism training, always emphasizing the importance of making precise measurements and theory-testing perturbations. I also want to thank Jonathan for all the inspiring meetings where he would pull up some Matlab code and teach us how derive or implement math in detail. I credit Jonathan for showing me that professors can still be very hands-on and passionate about science. Last but not least, I am fortunate to have had advisors who cared not only about the science but also about my well-being.

I want to thank my thesis committee members, Mala Murthy and David Tank, for their support both in terms of my thesis studies and career plans. Especially in the early years, their critical feedback made the research direction more exciting and solidified the approach. Mala and David are both my scientific idols; I credit them not only for setting good examples of the scientists I aspire to become but also for their leadership in making PNI such an amazing place. I want to thank Lindy McBride and Gautam Reddy for being my thesis reader. The continuous interaction with Lindy at numerous invertebrate neuroscience journal clubs and strengthened my interest in invertebrate neuroscience. I was fortunate enough to overlap with Gautam in the last year of my thesis studies, learning more about learning algorithms and behavioral models from him. Finally, I want to extend my gratitude to many other faculties who interacted with me during this journey. In particular, I thank Michael Berry for being the faculty member who truly inspired me in the beginning. My early work on the retinal code and my perspective on computational neuroscience have been strongly influenced by him. Additionally, I thank Josh Shaevitz and Bill Bialek for building a wonderful biophysics community at Princeton.

In some sense, peers are the most important component during PhD studies. I am extremely fortunate to be surrounded by brilliant people who are both smart and friendly. I want to thank all the Leifer lab members who overlapped with me: Monika Scholz, Mochi Liu, Xinwei Yu, Milena Chakraverti-Wuerthwein, Francesco Randi, Abdul-Bassit Fijabi, Matthew Creamer, Junang Li, Kelsey Hallinen, Anuj Sharma, Wayan Gauthey, Pearl Thijssen, Emily Osborne, Sophie Dvali,

and Sandeep Kumar. Undergraduate students I've mentored: John Li, Alicia Castillo Bahena, and Anthony Fisher. In additional to vibrant meeting discussions and experimental collaborations, I will miss baking with Kelsey and hunting for swags at conferences with Sophie. I will miss asking Francesco questions and being instantly saved every time. I will miss grabbing coffee with Sandeep, Emily, and Anuj everyday; and would absolutely miss sometimes (with increasing frequency when I'm writing this thesis) getting bubble tea with Emily and Junang.

I want to thank everyone in the Pillow lab: Zeinab Mohammadi, Rich Pang, Dean Pospisil, Alex Riordan, Anuththara Rupasinghe, Victor Geadah, Orren Karniol-Tambour, Aditi Jha, Yoel Sanchez Araujo, Bichan Wu, David Zoltowski, Zoe Ashwood, Iris Stone. These individuals are so intimidatingly smart yet extremely friendly and welcoming. As a half (often more than half) experimentalist, I often worried about not being able to catch up in meetings or discussion, but the Pillow lab somehow provided a warm, supportive environment that allowed me to learn on my own pace. I've enjoyed so many hilarious moments, from nerdy jokes during lab meeting to memes in the slack channel, and I feel so fortunate to be part of this lab. Being a graduate student in the Pillow lab has the benefit of being advised by many brilliant postdocs. Particularly, I want to thank Matt, Rich, Dean, and Brian for many feedback on research, presentations, and career plans.

I would not be able to mention about achievements at Princeton without acknowledging my mentors back in Taiwan. I thank my undergraduate and Master advisors CK Chan, CT Yen, and CC Chen for training me as a scientist. Particularly, I acknowledge CK for truly shaping my approach and taste in science. CK taught me everything from soldering wires to solving nonlinear differential equations. He showed me that it's possible, and extremely rewarding, to work on both experiments and theory. More broadly, as someone who has been passionate about science since elementary school, I am fortunate to have had many inspiring mentors, many of whom may not be listed here. I deeply appreciate all of their patient guidance and support throughout the long academic journey.

As an introvert, Princeton provides an ideal environment with its limited social activities and low social pressure. That being said, the support from friends and occasional non-scientific activities have undeniably played a crucial role in maintaining my sanity. I thank Jimmy and Abe for getting countless bubble teas during graduate school. Many nights talking about Star Wars, violin, and bubble tea made me understand the meaning of life (of course, science and art are also important in the inevitably ending life). Special thanks to Kong Fu Tea and Gong Cha for being the oasis in Princeton whenever I feel homesick and crave for bubble tea. I thank Robin and Mu-Hsi for inviting me to many dinners and art performances. I thank Claire, Shu-Heng, Po-Ju, Yao Lai, Yi-Chin, Ting-An, Jerry, and Sophia for inviting me to many dim-sum gatherings and parties in New York City. I thank Po-Ta, KC, and Ying-Jen for discussing about biophysics with me in Mandarin. I thank the Taiwanese student association for providing support, from me being an incoming student, the chair, to an old student throughout the years. I thank the "study group": Tsung-Lin, Hsiang-Chih, Bor-Ting, Ting-Wei, Hung-Tang, Karen, Mitchell, and many others, for fruitful discussion about complex systems and random trivia over meals. Finally, I thank Chang-Yu and Shang-Fu for being my best friends since college. We've supported each other for more a decade even at long distance.

My journey in Princeton begun with a wonderful cohort: David, Lindsay, Sam, Paula, Tomo, Shruthi, Liz, Sarah, Tim, and Rachel. Particularly, I thank Tomo and Tim going for brunch at GC on many Sundays, Liz and Sarah for hanging out with me at social hours, and Rachel for reluc-

tantly solving many chores for me. I am fortune to have wonderful cohort-roommates: David, Tomo, and Q before and during the pandemic. Returning from military service after the pandemic, Howard was a wonderful roommate with whom I had many late-night philosophical discussions. He taught me how to appreciate TV shows and helped prevent me from working all the time. Lastly, I want to give a shout-out to my outstanding MCN 2024 cohort, with whom I enjoyed an unforgettable, fantastic, nerdy summer at Woods Hole.

I'm a lucky person — not only do I want to stay at Princeton forever, but I'm happy to stay at any point in life. Most importantly, I cannot express my gratitude to my parents in simple words. While they might not fully understand my work (my father still thought I study vision after studying olfaction for 5 years), their genuine support and faith in me really enabled me to pursue my dreams. I deeply appreciate their presence during times of immense pressure, imposture syndrome, or even the pandemic during graduate school. I am especially thankful to Eric, my younger brother who always takes care of me like an older brother. Thanks, Eric for helping me move, visit me over birthdays, and proofread my writing before ChatGPT came along. I thank my cousin Kyle, auntie Ping, and many relatives back in Taiwan, who I see at most once a year, but who always care for me and are ready to feed me sushi and bubble tea whenever I visit home. Last but not least, I want to express gratitude and love for my long distance partner Vicky, who equally resists the concept of growing up as me, but also shows strength and resilience to support me when needed. I cannot be more excited about the future journey together.

Graduation is a milestone, but not the end of education. I hope to continue the quest for being knowledgeable, compassionate, and humble, and to occasionally revisit the wonderful grad school memories, or even this thesis.

1

Introduction

Biological systems are inherently computation — they sense the surrounding environment and produce corresponding reactions[252,80,20,138,157]. In the animal kingdom, the nervous system has evolved to efficiently gather sensory information, incorporate learned experiences, and execute complex behavioral actions[100,80]. Characterizing the information processing from sensory input to behavioral output and the underlying brain computation is a central goal of neuroscience[207,67,139,52,28].

Sensory navigation is one of the most fundamental behaviors observed in species ranging from single cell organism to human[20,10,252]. Animals naturally navigate in the complex chemical environment to seek resource or avoid pathogens and predators, and can rapidly learn new contexts in a few trials[95,87,172]. The tightly coupled sensory information and behavioral actions during nav-

igation constrains our understanding for the underlying brain computation[67,28,10]. To quantitatively study sensory navigation and make predictions about neural computation, the combination of precise experimental measurements and flexible statistical models is crucial.

In the past few decades, the integrative approach with theory-guided experiments has provided fruitful insight into sensory coding, learning, and behavioral control[189,28,52], which are all essential components in sensory navigation. Traditionally, empirical studies were mostly conducted in a relatively constrained setup, with artificial sensory stimuli, restrained animal motion, and structured trials[107,181,17]. More recently, with emerging new experimental tools and theoretical frameworks, the field has moved towards the investigation of complex behavior and neural dynamics during sensory navigation[87,238,83,208,171]. Enabled by techniques for neural measurement and perturbation[8,184,83,156,152,72] and computational models constrained by neural and behavioral measurements[175,230,234], there are substantial opportunities and needs for theory-guided experiments to answer questions such as: What does the animal sense while it navigates? How does the animal's past experience influence its strategies? Is the behavioral strategy static or dynamic? What are the underlying neural substrates and network dynamics for flexible behavior?

1.1 Background in sensory navigation

The classic study of sensory navigation started with chemotaxis behavior in the bacteria *E. coli*[20,18,132,167]. The series of experimental measurements and theoretical analysis revealed the detailed biophysics of chemical sensing and behavioral kinematics. The physical properties in the environment with low Reynolds number constrains both the bacteria motility and chemical sensing mechanisms[19]. In order to climb up chemical gradients, the cell integrates temporal change of chemical cues in order to make reliable behavioral decisions. The molecular sensing process effectively filters noisy chemical signals and adapts to different background concentration. The navigational behavior appears to have two modes: run and tumble, where the persistent runs are disrupted by tumbles with random angular distributions. The bacteria alters its tumbling frequency as a function of the sensed change in concentration, resulting in a process called the bi-

ased random walk.

In another extreme physical realm with high Reynolds number olfactory plumes, recently studies have made progress towards characterizing the sensorimotor transformation insects during plume tracking[68,23]. Animals can utilize precise timing of odor encountering and active sensing mechanisms to navigate in complex plumes. Different from smooth gradients bacteria experience, the turbulent plumes have complex spatiotemporal concentration profiles that provide little information to locally guide locomotion. Short-term memory and evidence accumulation mechanisms are required to navigate with such intermittent signals[33,68,128,251]. Furthermore, the spatiotemporal distribution of odorants and physical source of odor emission alters the ideal strategy animals employ during navigation[33,208].

Together, the study of sensory navigation reveals principles of chemical sensing, including filtering, spatiotemporal measurements, and adaptation[205,10]. In order to transform sensory signals into relevant behavioral actions, the animal has to be able to form memory, maintain internal states, and learn about the environment. Furthermore, the navigation strategy and exact parameters for searching algorithms depend on the specific chemical cue and physical environment.

While there are principles for sensory navigation behavior, a number of gaps remain. The experimental challenges include the precise measurements of sensory cues the animal experiences in the environment. The exact sensory signal obtained by the animal during navigation is crucial for constraining models and analysis for information processing[238,159,269]. In part due to the difficulties to measure airborne concentration, past work mostly approximate olfactory cues with spatial distance or through flow simulation. Another important thread of study is related to how animals integrate prior experience with the current sensory navigation behavior. A large body of studies focus on intrinsic navigation behavior[13,12,51,91,192,241], but less is known about navigation produced through learned experiences[87,95,50,131]. Investigating the integration of memory and internal states is crucial for moving from sensory domains towards decision and control circuits[79]. Lastly, the research community continues to call for theoretical framework to better model and make predictions about sensing mechanisms, behavioral algorithms, and circuit computation for

sensory navigation[205,252,23]. More specifically, the field needs models for navigation that are constrained by observations or directly fitted to experimental measurements.

1.2 Background in neural computation

Many classic studies in neural computation start with sensory encoding[9,232,207]. For decades, neuroscientists presented sensory input to an animal and measure the elicited neural response in peripheral neurons close to the input, thereby building tractable encoding models that transform stimuli to neural activity[193,152,226,107]. In certain cases, the biophysical properties of a single neurons governs the information processing. For instance, slow adaptation to odor concentration in a single sensory neuron AWC in *C. elegans* is required to explain the worm's behavioral decisions[150,50].

As we move towards the inner brain or the motor system, in contrast to single neuron computation, interconnected neural circuits are generally required to perform more elaborate computations[178,72,237,70]. Neural circuits and motifs consist a few number of neurons with specific connectivity[160,250,176]. These neurons can transform input signals into computed output through neural interactions and generate complex temporal patterns[160]. For instance, the interaction between motor-related output and interneurons in *C. elegans* generate state-dependent behavioral output during sensory navigation[123,91].

Modern machine learning architectures, such as deep networks and recurrent neural networks (RNNs), have roots and interaction with computational neuroscience[237,254,100]. Past work trained RNNs on either neural data or the animal behavioral task and find insightful correspondence to representation in the neural state-space and computation through dynamics[237,70,185,202]. In contrast to classic work focusing on single neurons[107,12], the modern description of neural computation has shifted towards dynamics and structure in large neural populations[255,220].

While largely explored in cortical networks and behavioral tasks in vertebrates, the field lacks methods to study neural computation in more compact, noisy circuits[197,84]. Different from large RNNs with quasi-deterministic dynamics that describes fixed behavioral tasks, invertebrate stud-

ies are often in a regime with stochastic, small circuits, studied during complex, natural behavior[171,67]. It is still challenging to connect neural circuit dynamics with the complex behavior it generates. On the experimental side, it is crucial to develop methods to interrogate neural computation in the context of naturalistic behaviors[8,184]. On the theoretical side, the field calls for models that describe stochastic dynamics in compact brains, as well as algorithms to train models and infer parameters from experimental measurements[257,250,165].

1.3 Background in *C. elegans* neuroscience

The nematode worm *C. elegans* is an idea model specie for systems neuroscience research because of its well-characterized neural nervous system and behavioral motifs[260,13,234]. The worm has a compact brain with 302 stereotyped neurons that can be aligned across individuals. The neural wiring diagram — the connectome — of the worm was the first to be completed and repeated[260]. In addition to the structure, physiological measurements and molecular characterization provide insight into signal propagation in the worm brain[39,38,12]. Recent work utilize novel optical and molecular methods to measure and perturb brain-wide neural dynamics in the worm[203,184,8].

Chemical sensing is the major sensory modality in worms. In the compact brain, 32 neurons participate in chemical sensing, with 12 that are specialized with sensory cilia at the nose tip[12,11]. The worm uses these sensors to sense signals such as airborne organic compounds, salt concentration, temperature, and oxygen. Early studies focus on sensory neurons that are tuned to specific molecules and map out the corresponding receptor expression[13,12]. More recent study with systematic neural imaging across the sensory population show collective activation across many odorants[152,198]. This indicates that olfactory signal might be processed as a population code. Many interneurons immediately downstream from then sensory neurons are shared across the sensory population[161,111,110,137]. While certain work pinpoint the crucial sensorimotor transformation in specific interneurons[137,36], the detailed neural computation that transforms sensory input into behavior is still an ongoing research topic in the field.

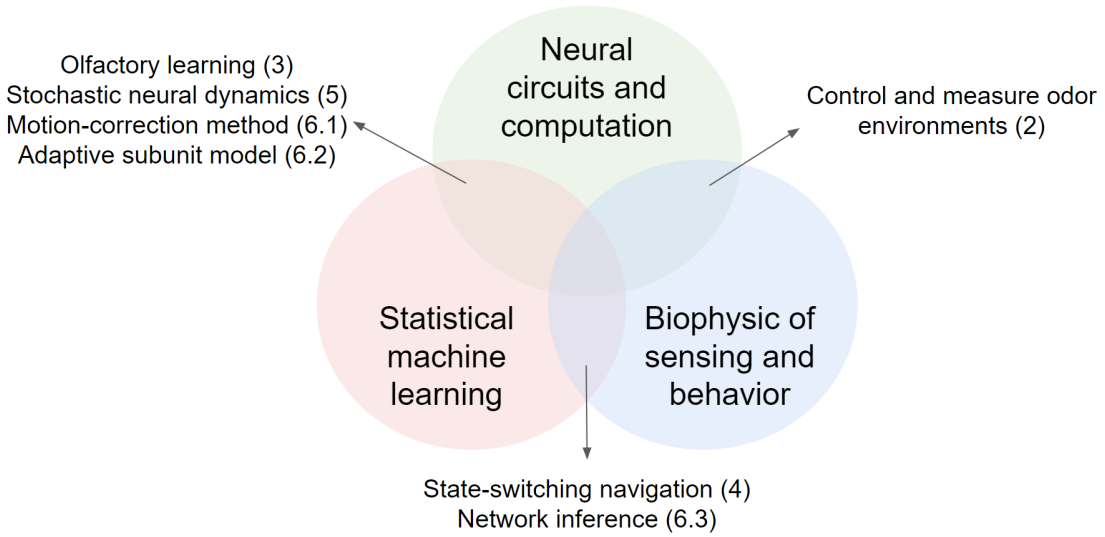

Figure 1.1: **Topic categories and the corresponding chapters of the thesis.**

The worm's behavioral output has been used as a phenotype to readout gene-behavior relations for many decades[39,12,77]. In combination with machine learning methods for behavioral quantification and dynamical systems analysis, recent studies describe the worm's posture dynamics in the behavioral state-space[155,234,2]. Stereotypical behaviors such as turns and runs emerge from the low-dimensional posture space[192,57,233]. More recently, with the development of neural imaging methods in freely moving worms, the neural population code underlying certain posture, movement, and sensory response have been reported[97,8]. While making progress towards the neural computation underlying behavioral actions, many open questions remain. Despite its simplicity, the worm's behavior readily spans across multiple time scales, from sub-second feeding behavior to hour-long learning and memory[129,27,14]. How behavior vary across longer time scales, respond to integrated past experiences, and are driven by time-varying environmental signal are ongoing research topics[79,27,198,40].

1.4 Thesis overview

This thesis aims to address and bridge the gaps highlighted in the aforementioned sections. At the higher level, the work covers topics in neural circuits and computation, the biophysics of sensing and behavior, and methods in statistical machine learning (Fig 1.1). The following 6 chapters

roughly lie in between different intersections. More particularly, chapter 2-4 contains a series of methods, experiments, and modeling for sensory navigation in the worm *C. elegans*. In chapters 5 and 6, we introduced models and algorithms for neural network dynamics. The models range from those constrained by worm behavior or neural recordings to more abstract cases, allowing us to study encoding and inference in networks.

The series of work has been presented at conferences such as the APS physics meeting, CoSyNe, SfN, the Neural Circuit meeting at Cold Spring Harbor, the Navigation Algorithm meeting at Janelia, and locally at Princeton Neuroscience Institute (PNI) and the Center for the Physics of Biological Function (CPBF) communities. Regarding the layout of this thesis, chapters 2-4 encompass the full scope of my first-author projects. In the remaining chapters, the focus is on my key contributions to collaborative works.

For the study of worm navigation behavior, we presented a series of work in the following three chapters. In chapter 2, the novel apparatus is developed in collaboration with Rui Wu and Marc Gershow at NYU, along with Andy Leifer's mentorship. The work has been published in the *eLife* journal[46]. In chapter 3, we presented experiments and models for learned olfactory navigation in worms. An optogenetic strain in this work is from Anuj Sharma, and the project is advised by Andy Leifer and Jonathan Pillow. The work is under revision in *PLOS Biology*[45]. In chapter 4, we introduced additional measurements and a state-space model to capture behavioral dynamics during sensory navigation. This is advised by Andy Leifer and Jonathan, and would be available on arXiv.

For more general models and theories for neural dynamics, we combined unpublished work in chapter 5 and published ones in chapter 6. In chapter 5, we included algorithms and analysis for stochastic network dynamics. The work is largely inspired by experimental observations in chapter 2-4. The work is mainly advised by Jonathan, but mostly unpublished at the moment. Lastly, chapter 6 contains a number of collaborative projects published during the thesis studies. This includes the adaptive model developed with Michael Berry[102], latent dynamic model with Matt Creamer[59], and non-equilibrium network models in collaboration with Ying-Jen Yang at

Stony Brook[47,43].

Towards the final chapter of the thesis, we concluded current progress in the field and some ongoing work. In addition, we briefly describe plans for future studies at the intersection of systems neuroscience, biophysics, and machine learning.

There is only one thing worse than coming home from the lab to a sink full of dirty dishes, and that is not going to the lab at all!

Chien-Shiung Wu

2

Controlling and Monitoring Odor Profile to Study Olfactory Navigation in Small Animals

In this project, we developed a novel tool and resource to create stable odor landscapes (gradients) for the study of chemotaxis in small model organisms. Using a metal oxide-sensor-based approach, odor concentration is measured in space and time and further verified with data-driven simulations of the odor diffusion process and surface interactions. Combined with a system of controlled odorized air flow, odor landscapes are generated to characterize aspects of olfactory navigation in *C. elegans* and the *Drosophila* larva. For presentation in the thesis, we placed a portion of the methodologies in the Appendix section and left out more technical details by ref-

erencing the published paper[46]. As a tools and resource article, the list of equipment, setup, and protocol are fully available. This work is a collaboration with another graduate student, Rui Wu, and professor Marc Gershow at NYU, under the supervision of Andy Leifer. It has been presented at the Annual Physics meeting, Cold Spring Harbor conference, and published in *eLife*[46].

2.1 Introduction

Odor-guided navigation is common across the animal kingdom[10]. Olfactory cues inform an animal of its location in a natural environment[23], and allow it to adjust its locomotion to navigate an odor landscape in a goal directed manner[13,20,1]. Odor guided navigation is an ethologically relevant task that is important for the animal's survival, and it has been a useful framework with which to study genes and circuits underlying sensory-motor transformations[28,52]. We focus in particular on the navigation of small model organisms in continuous, non-turbulent gradients established by the spread of odorants from their sources due to diffusion and drift. How animals interpret these gradients and use them to inform their actions remains an active and productive area of research especially in genetic model systems like *C. elegans* and *Drosophila melanogaster*[13,1,150,170,90,83,244]

A major challenge is to quantitatively relate the animal's behavior to the precise olfactory cue that the animal experiences moment-by-moment. Therefore it is critical to precisely control the odor environment and record the sensory cues experienced by these animals. The need to *control* and *measure* odorants still pose a formidable challenge. While many techniques exist to present odor stimuli[92,128], measure odor concentration[243,33], or develop virtual odor environments in a lab[265,201], it remains challenging for precise control and continuous monitoring of an odor landscape.

All approaches to generate odor landscapes in a lab environment must contend with the odor's diffusivity and interaction with other substrates. Early approaches to *control* odor concentration relied on passive diffusion to construct a quasi-stationary spatial odor gradient, for example by adding a droplet of odorant in a petri dish in a "droplet assay"[177,192,159,110]. Diffusion places se-

vere limits on the space of possible landscapes that can be created and on the timescales over which they are stable, and the created odor profile is sensitive to adsorption of odor to surfaces, absorption into the substrate, temperature gradients, and air currents, all parameters that are difficult to measure, model, or control. Microfluidics allow water-soluble odors to be continuously delivered to a chamber in order to provide spatiotemporal control[51,3,158]. Microfluidics devices, however, are limited in extent, require water-soluble odors and must be tailored-designed to the specific attributes of the animal's size and locomotion. A microfluidic post-array has been shown to support *C. elegans* locomotion, but no similar microfluidic device has been demonstrated to support olfactory navigation of *Drosophila* larvae, for example.

We previously reported a macroscopic gas-based active flow cell that uses parallel flow paths to construct temporally stable odor profiles[85]. That approach allows for finer spatiotemporal control of the odor gradient, is compatible with *Drosophila* larvae, and works with volatile airborne-based odor cues. This device used an array of solenoid valves to generate programmable odor profiles, but perhaps because of its complexity has not been widely adopted.

Most methods to create an odor landscape do not provide a means for knowing or specifying the spatiotemporal odor concentration. In other words while an experimenter may know that some regions of an area have higher odor concentrations, they cannot quantify the animal's behavior given a precise concentration of the odor. This limits the ability to quantitatively characterize sensorimotor processing. To address this shortcoming, various methods have been proposed to *measure* odor concentration across space. For example, gas samples at specific locations have been taken and measured offline[269].

In one of the most comprehensive measurements to date, Louis and colleagues[159,238] used infrared spectroscopy to measure the spatial profile of a droplet based odor gradient. They used this approach to study bilateral olfactory sensing and olfactory encoding in *Drosophila* larvae, clearly demonstrating the value of measuring odor landscapes. That work measured odor concentration at a relatively low temporal resolution ($\sim$ 1 sample / min), only when animals weren't present, and relied on a relatively expensive Fourier-transform infrared spectroscopy machine.

In all of these previous works, measurements were performed offline, not during animal behavior, and the odor concentration was assumed to be the same across repeats of the same experiment, or when animals were present. But even a nominally stable odor landscape is subject to subtle but significant disruptions over time from small changes in airflow, from temperature variation, and from the odor's interaction with the substrate, which can include absorption, adhesion, and reemission[92,269,241,268]. This is challenging to account for and control within a single behavior experiment, and is even more difficult to account for across multiple instances of such experiments. Additional variability may also arise across experiments as a result of the introduction of animals, changes to agar substrates, and alteration in humidity or other environmental conditions. To recover the odor concentration that an animal experienced, there is a need to monitor odor concentration and animals' behaviors concurrently.

Our previously reported flow cell used a photo-ionization detector (PID) sensor moved across the lid before behavioral experiments to measure the odor concentration across space at a single point in time[85]. During experiments, the total concentration of odor in the chamber was monitored concurrently with measurements of behavior. While this provided some assurances that the overall odor concentration was relatively stable, it did not provide any spatial information concurrently with behavior measurements.

Metal-oxide sensors are small and scalable digital gas sensors that have recently been used to monitor odor concentration in studies of olfaction[219,73,243] especially in the context of turbulent odor plumes. Arrays of metal-oxide sensors have been used to measure spatial odor landscapes by sampling across space[26]. Here we leverage these approaches to present a new multi-sensor odor array with metal-oxide sensors combined with a flow chamber in order to measure the odor landscape that an animal crawls in with high spatial and temporal resolution. The array of sensors can be used two ways: the full array can be used to measure the generated gradient throughout the extent of the chamber, or parts of the array can be used on the borders to monitor, *during behavioral experiments*, the odor profile along the boundaries of the chamber. By varying flow rates and the sites of odor introduction, we show a variety of odor landscapes can be generated

and stabilized.

To demonstrate the utility of the apparatus, we applied this instrument to quantitatively characterize the sensorimotor transformation underlying navigational strategies used by *C. elegans* and *D. melanogaster* larva to climb up a butanone odor gradient. Butanone is a water-soluble odorant found naturally in food sources[264] that is often used in odor-guided navigation studies and is of particular scientific interest for studies of olfactory learning[12,150,50,247]. We show that the agar gel used during behavioral experiments greatly disrupts an applied butanone gradient, and we demonstrate a pre-equilibration protocol allowing generation of stable gradients taking into consideration the effects of agar. Moreover we monitor these gradients during ongoing behavior measurements via continuous measurements of the odor profile along the boundaries of the arena.

Using these stable and continuously measured butanone gradients, we measure odor-guided navigation in animals by tracking their posture and locomotion as they navigate the odor landscape. We record chemotaxis behavior and identify navigation strategies in response to the changing odor concentration they experience. In *C. elegans*, we observe navigation to the odorant butanone employs similar strategies to those reported for other sensory-guided navigation conditions, such as salt chemotaxis[110,65,161], namely a biased random turn, known as a pirouette[192], and a gradual veering, known as weathervaning[110,116]. We also measure navigation strategies of a mutant *C. elegans* and implicate the interneuron AIB mutant as being important for olfactory-guided weathervaning. In *Drosophila melanogaster* larvae, we identify runs followed by directed turns[85,159,90]. By using concurrent measurements of behavior and odor gradient we characterize olfactory navigation in these small animals on agar with known butanone odor concentrations, which for *C. elegans* has not been reported before.

2.2 Results

We developed new methods both for generating and measuring odor gradients which we describe here. The systems are modular, scalable, and flexible. The components, which can be used in-

dependently of each other, can be fabricated directly from provided files using online machining services.

2.2.1 Flow chamber for generating spatiotemporal patterns of airborne odors

We first sought to develop a method of creating odor gradients that satisfied the following criteria:

1. The spatial odor profile should be *controllable*. Varying control parameters (e.g. flow rates, tubing connections) should result in predictable changes to the resulting odor landscape.

2. The odor profile should be *stable* and *verifiable*. The same spatial profile should be maintained over the course of an experiment lasting up to an hour, and this should be verifiable via concurrent measurements during behavior experiments.

3. The apparatus should be *straightforward* to construct and to use, and *flexible* to adapt to various experimental configurations, including using with either *C. elegans* or *Drosophila* larva, and with agar arenas of various sizes.

We constructed a flow chamber to control odor air flow across an arena (Fig 2.1a). Odor and humidified air are sourced from two bubblers, one containing pure water and the other an aqueous solution of odorant and water. Flow rates are controlled by separate mass flow controllers (MFCs) upstream of the bubblers. Downstream of the bubblers, the odor and air streams are divided into parallel sections of equal lengthed tubing. Each tube is connected to one input port of the flow chamber. The pattern of connections and the flow rates set by the two MFCs determines the shape of the produced odor profile. For instance, if odor is provided at a single central inlet, the resulting profile is a 'cone' (Fig 2.1d-g) whose peak concentration and divergence are controlled with the MFCs (e.g. speeding up the odor flow while slowing down the air flow broadens the cone). Temporal gradients can be achieved by varying the odor flow in time, subject to constraints imposed by the odor's absorption into the agar gel.

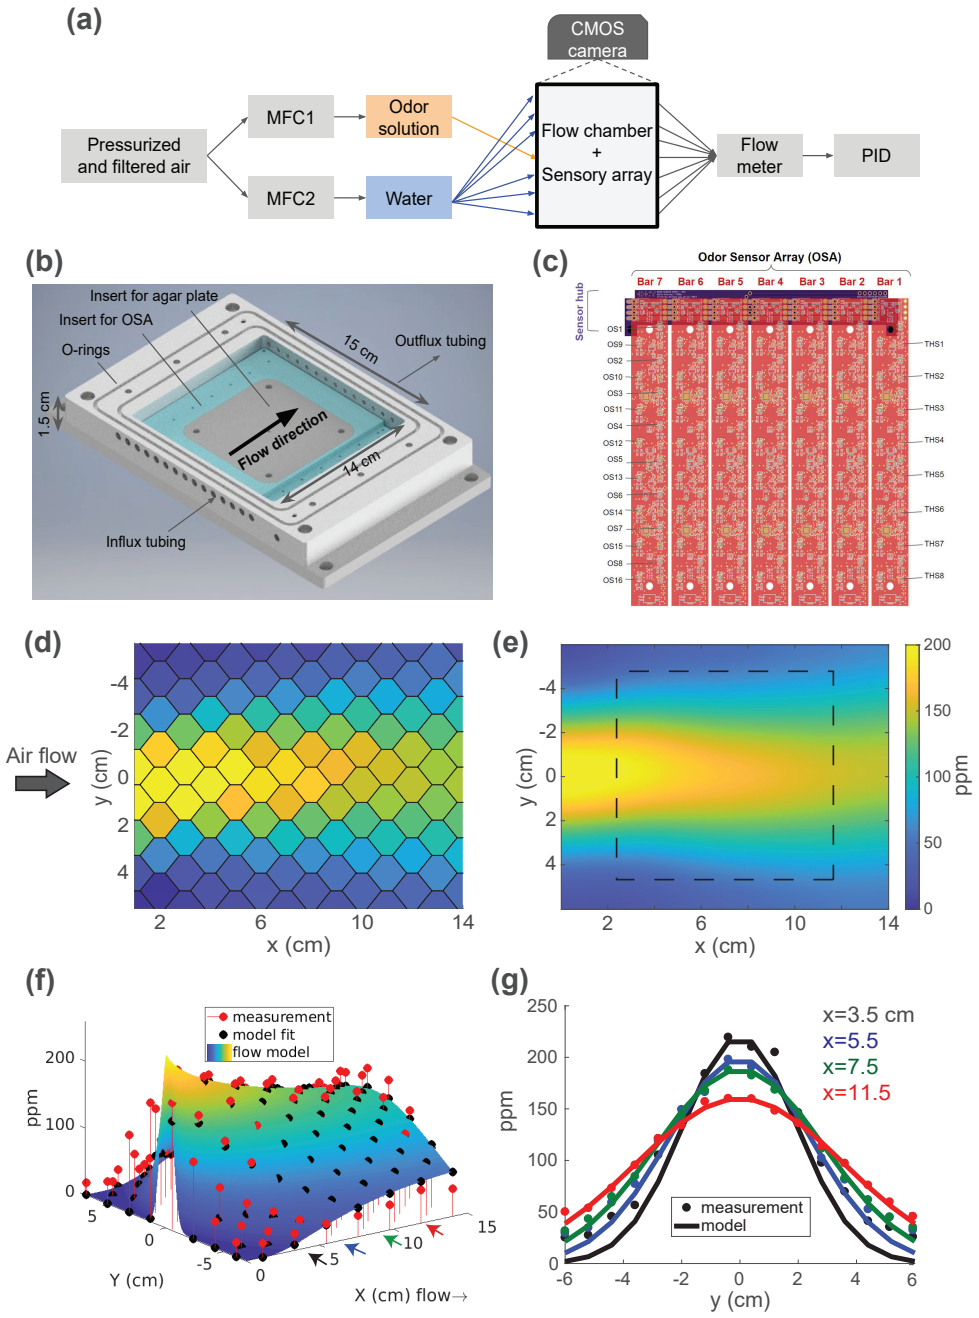

Figure 2.1: **Odor flow chamber with controlled and measured odor concentration. (a)** Schematic of airflow paths. Airflow paths for odor solution and water are controlled separately by mass flow controllers (MFCs) and spatially arranged into the odor chamber. The outflux from the chamber connects to a flow meter and photo-ionization detector (PID). **(b)** Flow chamber design. **(c)** Odor sensory array (OSA). Seven odor sensor bars are connected to a sensor hub. Each bar has 16 odor sensors (OS) and 8 temperature/humidity sensors (THS). Measured odor concentrations from the OSA of a spatially patterned butanone odor concentration shown in **(d)** for each sensor, **(e)** interpolated across the arena with square dashed line indicating the area where agar and animals are placed. **(f)** A two-parameter analytic flow model fit to measurement. **(g)** Cross-sections from (f) at at 4 different x-axis positions (show as colored arrows). Sensor readouts are overlaid points on the smooth curves from a 1D diffusion model.

The outflux from the flow chamber is connected to a flow meter and photo-ionization detector to monitor the overall flow rate and odor concentration respectively. The geometry of this flow chamber is shown in (Fig 2.1b), where parallel tubings are connected from the side and the chamber is vacuum sealed with a piece of acrylic on top during experiments. The chamber is designed for use with ∼100 mm square agar plates. The extra width (2.5 cm on either side) diminishes the influence of the chamber boundary on the odor profile over the arena. Interchangeable inserts allow for different agar substrates (e.g. circular plates) or for full calibration by odor sensor arrays (Fig 2.1c), discussed in the next section. Metal components are designed for low-cost fabrication by automated mechanisms (either laser cut-able or 3-axis CNC machinable). The fabrication plans for the flow chamber and inserts, list of required components, and a detailed tutorial are all publicly available in the paper and repositories (see Appendix).

2.2.2 Measuring the spatiotemporal odor distributions

A central difficulty in measuring animals' responses to olfactory cues is quantifying airborne odor concentrations that vary in space and time[253,243,128,269]. This difficulty is exacerbated in turbulent environments where odor plumes carry abrupt spatial and temporal jumps in concentration far from the source with fundamentally unpredictable dynamics[33,69]. But even in laminar flows, like those considered here, boundary conditions, slight changes in temperature, and the presence of absorbing substrates like agar make this challenging.

There is therefore a need, even for quasi-stationary gradients, to characterize the odor profiles in situ and to monitor these profiles during experiments. Various optical techniques, like laser induced fluorescence or optical absorption[159,238,68], exist to monitor concentration across planar arenas, but these have not been demonstrated for behavior experiments, are expensive to construct, require specially designed arenas, or some combination of these disadvantages. Electronic chemical sensors can reveal the time-varying concentration at a particular point in space. A tiled array of these sensors acts as a 'camera' forming a 2D spatiotemporal reading of the concentration. The gold-standard for measurement of odor concentration is the photo-ionization detector

(PID), but even the smallest versions of these sensors are both too large (∼2 cm in all dimensions) and too expensive (∼ $500 each) to make an array. Metal-oxide odor sensors, designed to be used in commercial air quality sensors, are available in inexpensive and compact integrated circuit packages[219,26]. However, in general, commercial metal-oxide sensors are not designed for precision work - they tend to drift due to variations in heater temperature, humidity, adsorption of chemicals and ageing effects[69]. Most such sensors are designed to detect the presence of gas above a particular concentration but not to precisely measure the absolute concentration. We became aware of a newer metal-oxide sensor, the Sensirion SGP30 that was designed for long-term stability and concentration measurement; we wondered if such a sensor could be calibrated for use in an odor sensor array.

To calibrate the sensor, we created a controllable concentration source by bubbling air through butanone. The odor reservoir contains butanone dissolved in water and is kept below the saturation concentration (11 mM or 110 mM odor sources). We then mixed this odorized air flow into a carrier stream of pure air. We kept the carrier air flow rate constant ($\sim$ 400 mL/min) and varied the flow rate through the odor source ($0 - 50$ mL/min); the odor flow rate was slow enough that the vapor remained saturated, so the concentration of butanone in the mixed stream was proportional to the flow rate through the butanone bubbler, as directly measured with a PID. We typically calibrated concentration with continuously ramped flow rate in triangle waves with a 500 s period for 2-3 cycles. We performed calibrations across a range of odor concentrations and against two PIDs, each with a different dynamic range, and also compared against the PIDs to characterized the metal oxide sensors' temporal response. We found a one-to-one correspondence between the odor sensor reading and the PID reading that persisted over time and showed no obvious hysteresis. These metal oxide sensors' temporal response and stability appear to be well matched for capturing the slow fluctuations in our stable and non-turbulent odor landscapes. We also characterized the heating effects of these sensors in the arena, and conducted measurements to show that powering the sensors does not noticeably change the concentration of odorant in the chamber. We reasoned that after calibrating the metal oxide sensors to a PID, we could use the

array to measure spatiotemporal odor concentration distributions.

We constructed a sensor array from 'odor sensor bars' (OSBs), printed circuit boards each containing 16 sensors in two staggered rows of 8. Each OSB also contained 8 temperature and humidity sensors to allow compensation of the odor sensor readings. The OSBs are mounted orthogonal to the direction of air flow; 7 OSBs fit inside our flow chamber (112 sensors total) allowing a full measurement of the odor profile. Taken together these 112 sensors formed an odor sensor array (OSA), capable of measuring odor concentrations with $\sim$1 cm spatial resolution sampled at 1 Hz. Prior to all experiments we calibrated the OSA in situ by varying the butanone concentration across the entire anticipated range of measurement while simultaneously recording the odor sensor and PID readings.

To verify the ability of the OSA to measure concentration gradients we created an artificially simple steady-state odor landscape and compared it to our predictions based on a convection-diffusion flow model. We flowed in odorized air (~ 30 mL/min) through a central tube and clean air through the others (~ 400 mL/min distributed into 14 surrounding tubings) in an environment without agar and without animals. This results in an air flow velocity ~ 5 mm/s in the flow chamber. The odor concentration profile was established in a few minutes and remained stable over the time scale of 30 minute behavioral experiments. As the flow rates and concentrations are all known and the flow is non-turbulent, and there is no agar or animals present, the concentration across the chamber should match a convection-diffusion model. After establishing the gradient, we recorded from the discrete odor sensors on the array (Fig 2.1d) and estimated the values in between sensors using spline interpolation with length scale equal to the inter-sensor distances (Fig 2.1e). We compared this stationary profile with a two-parameter convection-diffusion flow model (Fig 2.1f,g) fit to the data and described in the Methods section. The measured concentrations in this artificially simplistic odor gradient show good agreement with the fit convection-diffusion model, especially in the central region where experiments are to be conducted, leading us to conclude the OSA can accurately report the odor concentration. Note that the model serves as an independent validation of the sensors, and is not required to perform the measurement. We

proceed to consider more complex odor environments.

We demonstrate several examples of flexible control over the odor profile. In the configuration shown in Fig 2.1 the center-most input provides odorized air and all surrounding inputs provide moisturized clean air to form a cone-shape stationary odor pattern. A narrower cone can be created by increasing the air flow (600 mL/min) of the surrounding inputs relative to the middle odorized flow (Fig 2.2a). The cone can be inverted by placing odorized air in the two most distal inputs, and clean air in all middle inputs (Fig 2.2b). This inverse cone has lower concentration in the middle and higher on the sides. In later animal experiments, we restrict odorized air to one side to form a biased-cone odor landscapes, resulting in a cone with an offset from the middle line of the arena. Many more configurations are possible, demonstrating that the odor flow chamber enables the flexible control of airborne odor landscapes that are much more complex than a single odor point source.

To demonstrate that the flow control and measurement methods are not restricted to any single odor molecule, we create and measure various stationary odor landscapes with benzaldehyde, ethanol, and ehtyl butyrate odorants that, together with butanone, span a range of different vapor pressure and water solubility ((Fig 2.2). For each molecule, the metal-oxide odor sensors are separately calibrated against a PID. To our knowledge, most volatile organic compounds that can be sensed and calibrated by a PID should be compatible with our sensor array system. Depending on the vapor pressure, molecular weight of different molecules, and the geometry of our flow chamber, the cone-shape odor profiles can be skewed or less smooth. The skewed profile is possibly a result of the inhomogeneity in the area created by the odor sensor hub. Nonetheless, we empirically verify that these odor profiles are stationary and reproducible.

2.2.3 Odor-agar interactions dominate classical butanone droplet assays

Classic chemotaxis experiments in small animals commonly construct odor environments with odor droplets in a petri dish, usually with a substrate like agar. Our odor delivery instrument is designed to be compatible with a similar environment. To first better understand classical chemo-

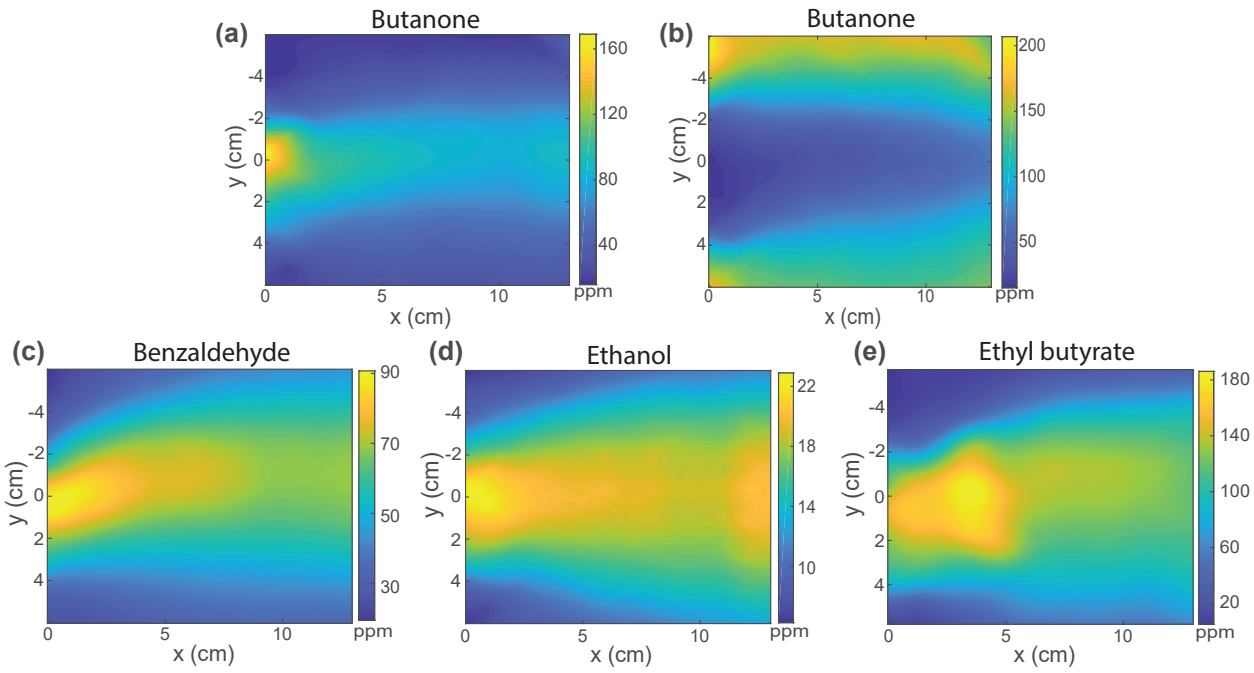

Figure 2.2: **Flexible control of a steady-state odor landscape. (a)** Configuration for a narrow cone with butanone. Colorbar shows interpolated odor concentration in ppm as measured by the odor sensor array. **(b)** An inverse cone landscape that has higher concentration of butanone on both sides and lower in the middle. Stationary cone-shape odor landscape of three different odorants: **(c)** benzaldehyde, **(d)** ethanol, and **(e)** ethyl butyrate.

taxis experiments with butanone, we sought to characterize the spatiotemporal odor profile from an odor droplet point source using our odor sensor array. We first considered the case without agar. In that case the odor concentration should be governed entirely by gas-phase diffusion. We placed a 2μL droplet of 10% butanone in water on the lid of our instrument centered in the arena above the full OSA and without any airflow (Fig 2.3a,b). Butanone was observed to diffuse across the arena in the first three minutes resulting in an equilibrium concentration that was close to uniform across the odor sensors. We note that the final concentration of roughly 100 ppm, and the equilibration timescale both match what we would expect from first principals for $\sim 10^{-6}$ mol of butanone in a ~ 225mL arena that have a diffusion rate of $\sim 0.08\text{cm}^2/\text{s}$ in air. A uniform odor landscape is not helpful for studying odor guided navigation, but most behavioral experiments are not conducted in a bare flow chamber but contain a biologically compatible substrate, such as an agar gel, as is typically used in droplet assays. We therefore sought to investigate the

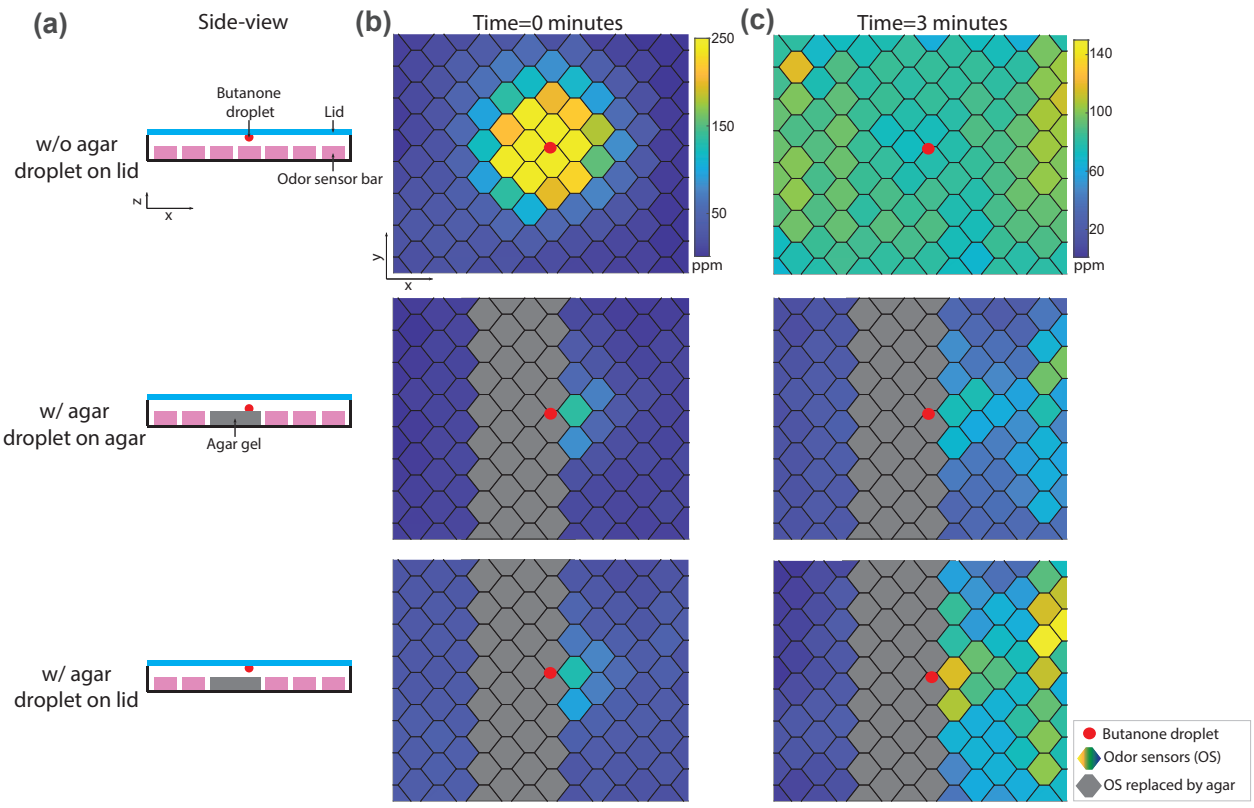

Figure 2.3: **The presence of agar in the droplet assay alters the time-evolution of butanone odor landscape. (a)** Schematic showing configuration of an odor droplet in an arena with or without agar. Red dot indicates the position of butanone droplet ($2\mu L$ of 10% v/v butanone in water) on the lid (blue) or agar (gray). Odor sensor bars are shown in purple. **(b)** Spatial concentration measured by odor sensor array is reported immediately after butanone droplet is introduced into the arena. **(c)** Spatial concentration after 3min.

role that agar plays in sculpting the odor landscape in butanone droplet experiments.

We introduced agar into the droplet assay by removing two sensor bars and replacing them with agar. We placed a butanone droplet directly on the agar or on the lid, as done classically in worm and larval assays respectively, and measured the odor landscape over time (Fig 2.3b,c). The odor concentration measured with agar is dramatically different from that measured without agar. Instead of quickly equilibrating to a uniform concentration, in the presence of agar there was instead a local maximum of butanone surrounding the droplet that persists even after 3 minutes. This difference in airborne odor concentration with and without agar persists after experimental perturbation such as removing and replacing the lid over the chamber. This convinced us of the need to account for odor-agar interactions when working with water soluble odorants such

as butane.

2.2.4 Agar poses challenges for measuring and establishing stable odor landscapes

By placing the agar on the underside of the lid hanging above the odor sensors, we were able to simultaneously record spatial odor concentrations generated by our flow chamber opposite the surface of the agar during animal behavior. Unfortunately, when imaging in this configuration, the odor sensors in the background obstruct the animals from being tracked over the extent of the arena. We therefore sought an alternative strategy: we chose to measure the odor profile along only the boundary of the arena in order to infer the odor concentration inside the arena in such a way as to also take into account the odor-agar-interactions. This approach relies on establishing temporally stable odor-landscapes, even in the presence of agar, which we sought to accomplish by leveraging our ability to control the flow of odor.

We first sought to measure how the presence of agar altered the odor profiles that would have been established by flow in a bare chamber lacking agar (Fig 2.1 and Fig 2.2). As in Fig 2.3c,d, we replaced two odor sensor bars with a rectangular strip of agar gel ('w/ agar') or a metal plate as a control ('w/o agar'), and then measured airborne odor concentration upstream and downstream of the agar under odorized airflow that would normally produce a cone profile (Fig 2.4a-b). While the agar had little effect on the odor landscape upstream of the agar, it drastically altered the downstream odor landscape (Fig 2.4b), suggesting that the agar absorbs the airborne butanone molecules. This finding is consistent with the odor droplet experiments (Fig 2.3), repeated with a full-sized 96 mm square agar plate intended for use with animals (Fig 2.4c), and to be expected since butanone is highly soluble in water (275 g/L). Pulse-chase style experiments confirm that agar does indeed absorb and reemit butanone To accommodate the full sized agar plate we measured only the one dimensional odor profiles upstream and downstream of the agar Fig 2.4d. Taken together, these experiments suggest that agar-butanone interaction presents a challenge for setting up and maintaining stable butanone odor landscapes.

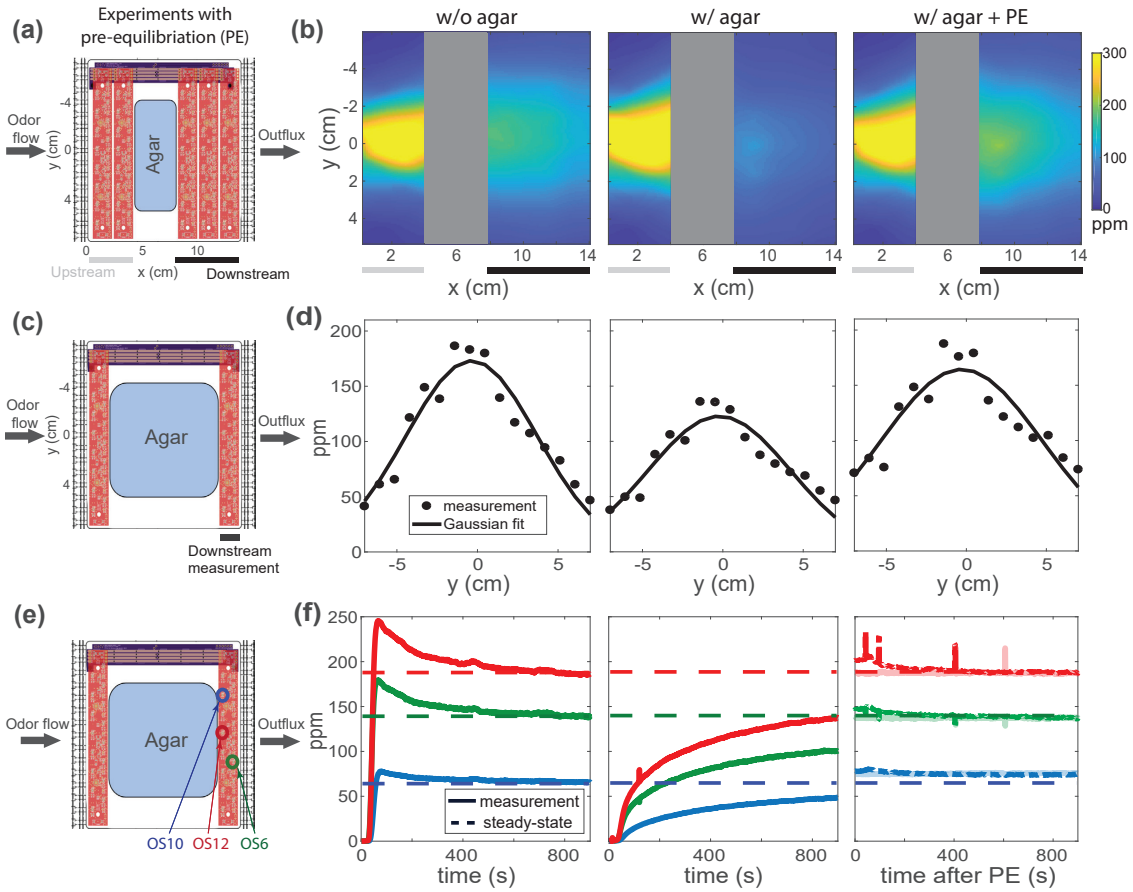

Figure 2.4: **Under flow, agar interacts with butanone odor to disrupt the downstream spatial odor profile, but a pre-equilibration protocol can coax the system into quasi-equilibrium and restore a temporally stable odor profile. (a)** Two odor sensor bars are replaced with agar to observe the effect of introducing agar on the downstream spatial odor profile. **(b)** Measured butanone odor profile is shown upstream and downstream of the removed odor sensor bars in the absence (left) and presence of agar (middle). Transiently delivering a specific higher odor concentration ahead of time via a pre-equilibration (PE) protocol restores the downstream odor profile even in the presence of agar. **(c)** Additional odor sensor bars are removed and replaced with a larger agar, as is typical for animal experiments. **(d)** Measurements from the downstream sensor bar under the same three conditions in (b). The dots are sensor measurements and the smooth curve is a Gaussian fit. **(e)** The same experimental setup in (a), here focusing on time traces of only three downstream odor-sensors (colored circles for selected OS). **(f)** Concentration time series of three sensors color-coded in (e). Traces for three conditions are shown: time aligned to initial flow without agar (left), time aligned to initial flow with agar (middle), and traces after PE (right, with transparent line showing measurements another 20 min after the protocol). The dash-lines indicate the target steady-state concentration for each sensor.

2.2.5 Odor flow compensation overcomes challenges posed by agar

We next sought a method to generate stable odor landscapes even in the presence of agar by applying a constant odor flow. In principle, the disruption caused by agar should be overcome by

the constant flow of a sufficiently long duration, after which the agar and airborne odor would be in quasi-equilibrium at all spatial locations. However, we measured odor concentration downstream from the agar and found that the airborne concentration failed to approach equilibrium on the timescales of single experiments FIGfig4e,f. This suggests that it is not practical to simply wait for the system to reach equilibrium. Instead we developed a pre-equilibration protocol to more efficiently bring the agar and airborne odor into equilibrium before our experiments.

To more rapidly establish a desired airborne odor landscape, we briefly first exposed the agar to an airflow pattern corresponding to higher-than-desired odor concentration, created by replacing the odor reservoir with one containing a higher concentration of butanone. We monitored the odor profile downstream of the agar until it reached the desired concentration and then switched to the original bubbler to maintain that concentration. Using this pre-equilibration protocol, we reached quasi-equilibrium quickly, typically after the order of ten minutes. Note the spatial parameters of the two airflow patterns were the same, only the concentration of the odor source changes. Pre-equilibration allows the generation of airborne odor gradients in the presence of agar that match those in the absence of agar (Fig 2.4b,d,f; right vs. left column), as indicated by odor sensor measurements along the boundary. This demonstrates that temporally stable gradients can be established even in the presence of agar.

Monitoring the boundary determines the odor landscape at quasi-equilibrium

Pre-equilibration allows us to quickly establish a temporally stable odor landscape even in the presence of agar, and stability is indicated by monitoring the odor sensors along the boundary. Now we seek an approach to infer the odor concentration over the agar from the measurements along the boundary. In this section we will show that the odor concentration above the agar after pre-equilibration should be the same as the odor concentration when no agar is present and we will then validate this equivalence empirically (to within 10%). We well then use this knowledge to infer the odor concentration above the agar at quasi-equilibrium using the measurements along the boundary. If the airborne odor concentration along the boundaries upstream and downstream

of the agar matches the profile in the absence of agar, one can infer that the airborne odor concentration landscape above the agar is also the same. We validated this approach three ways: via analytical argument, via numeric simulation, and via direct empirical measurement.

The physical intuition is as follows: in the absence of sources or sinks, the fact that two odor concentration distributions obey the same differential equations and share the same conditions on all boundaries means that the distributions are identical throughout the interior. Agar adds complexity because it can act as a source or sink by absorbing and reemiting the odor, but even then, knowledge that the odor along the boundary is the same with- and without agar places strong constraints on the interior odor landscape. The only way the interior could differ would be if sources and sinks of odor in the agar were precisely arranged so that all excess odor emitted from one point is exactly reabsorbed somewhere else before reaching the downstream boundary. Not only is such an arrangement unlikely, it is inherently temporally unstable and should tend to eliminate the excess sources and sinks over time. A mathematical version of this argument is presented in the Appendix. Therefore when we measure steady odor concentraiton along the boundary with agar that matches the concentration without, we are comfortable concluding that the odor landscape with agar is the same as without.

A numerical reaction-convection-diffusion model also shows that after pre-equilibration, the with-agar odor landscape matches the without-agar odor landscape (Fig 2.5). The model requires reasonable assumptions about the absorption rate, reemission rate, and capacity of the agar. The numerical simulation provide qualitative agreement to our empirical observations in (Fig 2.4).

We experimentally verified that the airborne odor concentration above the agar matched our expectations based on measurements made along the boundary, by performing direct measurements of the airborne odor opposite the agar. We used an upside-down configuration as in control experiments and Fig 2.6: we poured agar onto the lid and measured the odor concentration below it with the full sensor array underneath (Fig 2.6).

After the pre-equilibrium protocol, measured airborne concentration in the arena with agar was similar to measurements made without agar to within an average of less than 10% fractional dif-

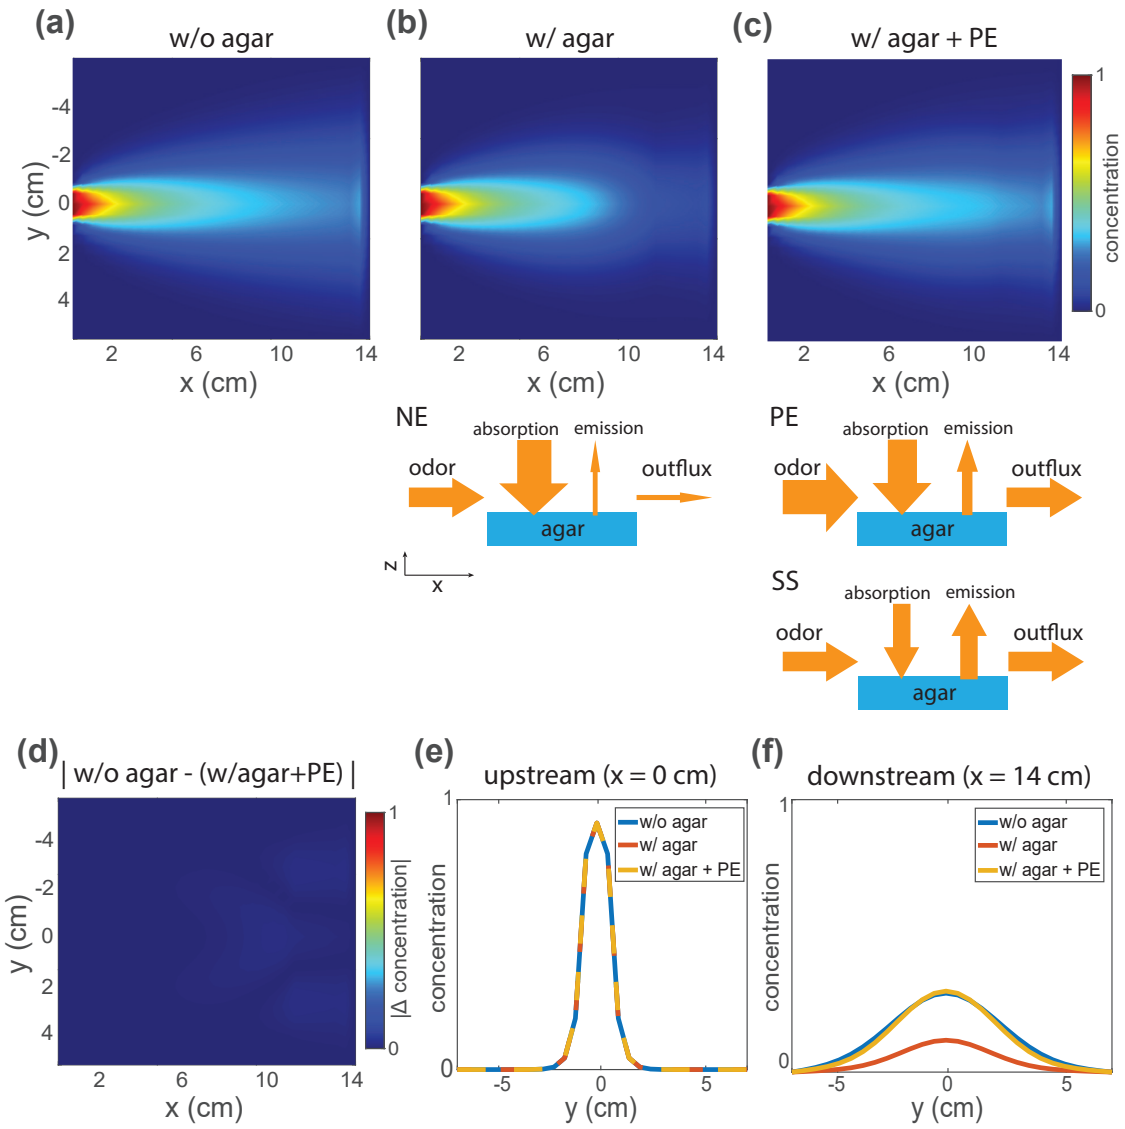

Figure 2.5: **Simulations from a reaction-convection-diffusion model of odor-agar interaction show that at quasi-equilibrium the airborne odor concentration is the same with or without agar. (a)** Simulation results of steady-state odor concentration in air without agar and with flow configured as in Fig 2.1. **(b)** The same simulation condition in (a) but now shortly after agar is introduced and before quasi-equilibrium is reached. A schematic of odor-agar interaction model is shown below. When agar is introduced, it absorbs the odor in air and decreases concentration measured downstream, producing a non-equilibrium (NE) concentration profile. **(c)** Odor concentration profile in air, with agar present, but after the pre-equilibration (PE) protocol brings this system to quasi-equilibrum. The PE protocol is shown in the schematic below, followed with steady-state (SS) with the stable odor concentration profile shown above. **(d)** The absolute difference of concentration profile without agar (a) and with agar after PE (c) is shown. **(e)** Upstream and **(f)** downstream odor concentrations along the agar boundary are shown for all three conditions.

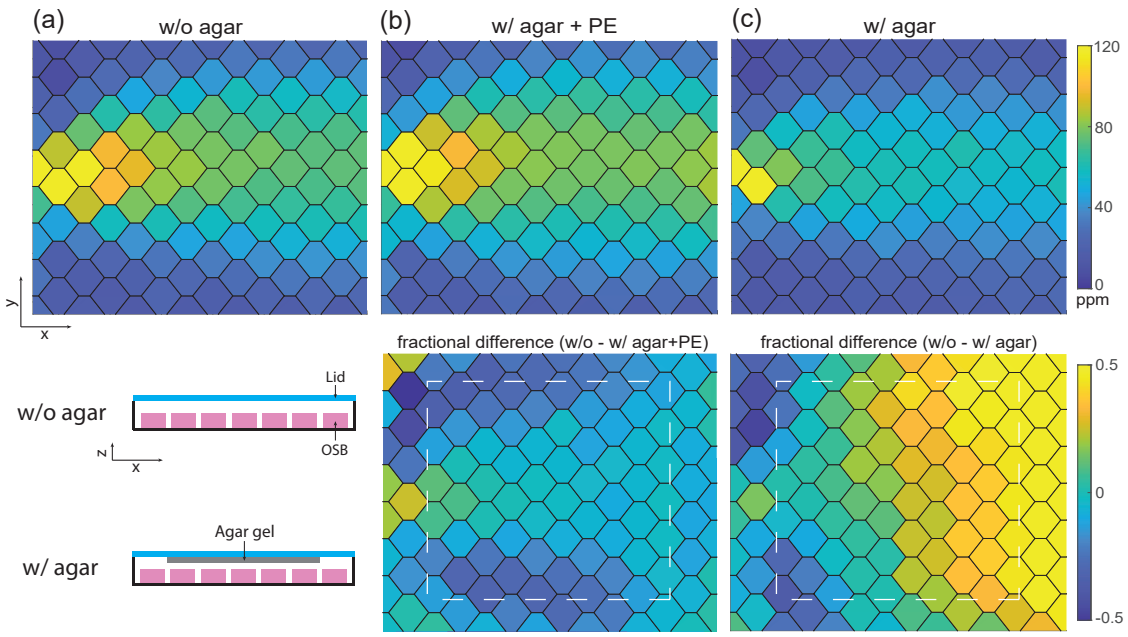

Figure 2.6: **Direct measurement of odor concentration across agar confirms quasi-equilibrium odor profile. (a)** Stationary butanone odor-cone profile without agar. Schematic for configuration with or without agar in this direct measurement shown below. **(b)** Same as (a) but with agar on the lid after the PE protocol. Fractional difference between odor profile with and without agar is shown below and white dash line indicates the region with agar gel. **(c)** Same as (b) but without the PE protocol, 15 minutes into the recording. Fractional difference and agar gel region is shown below.

ference in the area that the animal would crawl, FIGdirect-landscapeb. The agreement of odor landscape measured with and without agar at quasi-equilibrium agrees with our modeling and indicates that the odor landscape is stable and identifiable in our setup by monitoring the odor profile along the boundary. Therefore an experimenter could establish a desired odor gradient without agar, and then, even in the presence of agar, once the boundary measurements matched, the experimenter would be confident that the odor landscapes also matched. Note our hypothetical experimenter would not be required to perform any simulations or calculations.

The presence of worms did not noticeably alter the odor concentration. Our approach relies on reaching quasi-steady state. In measurements made without the pre-equilibration protocol, and therefore presumably before quasi-steady state, the interior odor concentration with agar no longer matches the without agar, as expected. Fig 2.6c.

2.2.6 Butanone chemotaxis in *C. elegans*

We sought to use our experimental approach to directly quantify *C. elegans*' navigation strategies for airborne butanone. *C. elegans* are known to climb gradients towards butanone[12,50,150]. Microfluidic environments suggest that they use a biased random walk strategy to navigate in a liquid butanone environment[150,3]. Worms are also known to use weathervaning to navigate airborne gradients of other odors[110,142] although to our knowledge this has not been specifically investigated for butanone.

Worms were imaged crawling on agar in the flow chamber under an airborne butanone odor landscape illuminated by infrared light. Here 6 recording assays were presented, with approximately $50 - 100$ animals per assay, and two different odor landscapes were used. *C. elegans* navigated up the odor gradient towards higher concentrations of butanone, as expected (Fig 2.7a,b). Importantly, the odor concentration experienced by the animal at every point in time was inferred from concurrent measurements of the odor profile along the boundary of the agar, Fig 2.7c. On average, animals were more likely to travel in a direction up the local gradient than away from the local gradient, as expected for chemo-attraction (Fig 2.7d). We use the term "bearing to local gradient" to describe the animal's direction of travel with respect to the local odor gradient that it experiences.

We find quantitative evidence that the worm exhibits both biased random walk and weathervaning strategies. To investigate biased random walks, we measured the animal's probability of turning (pirouette) depending on its bearing with respect to the local airborne butanone gradient (Fig 2.7e). We find that the animal is least likely to turn when it navigates up the local gradient and most likely to turn when it navigates down the gradient, a key signature of the biased random walk strategy[19,170].

To test for weathervaning, we measured how the curvature of the animal's trajectory depended on its bearing with respect to the local airborne butanone gradient (Fig 2.7f). When the animal's navigation aligns to the axis of the butanone gradient, the distribution of the curvature of its trajectory was roughly symmetric and centered around 0 (straight line trajectory). By contrast, when

the animal navigated more perpendicular to the gradient the distribution of the curvature of its trajectories was skewed. The skew was such that it enriched for cases where the animal curved its trajectories towards the local gradient. We further calculate the average curvature as a function of bearing to local gradients following methods in [110]. The result shows that the worm on average turns towards the gradient direction, which is the key signature of weathervaning strategy. Both the biased random walk and weathervaning behavior was absent in control experiments with flow but not odor, and we observed no evidence of anemotaxis at the ~ 5 mm/s air velocities encountered by the animals. We conclude that *C. elegans* utilize both biased random walk and weathervaning strategies to navigate butanone airborne odor landscapes [110,272]. Both strategies had been observed for salt chemotaxis before, but to our knowledge had not previously been quantitatively characterized for airborne butanone. More broadly, previous measurements of navigation in response to airborne odors for *C. elegans* have not had access to the underlying odor concentration and therefore could only use proxies, such as the bearing angle or distance with respect to a point odor source. Our approach allows a more direct description of the navigation strategy because we characterize turn probabilities and other features with respect to the local gradient directly (Fig 2.7), as opposed to a proxy.

To identify potential neurons involved in mediating the navigation strategy, we measured butanone guided navigation of mutant animals in which the interneuron AIB was chronically inhibited. Disruptions to AIB had previously been shown to either increase weathervaning (AIB ablation) [110] or decrease biased random walks [161] (AIB inhibition) when navigating in a salt gradient. But to our knowledge, AIB's role in influencing the animal's butanone-guided navigation strategy had not previously been measured. We observed that AIB-inhibited animals still climbed butanone gradients, but the worm relied on different behavioral strategies (Fig 2.8). The turn was still biased suggesting that the animal still used a biased random-walk strategy, although turning rate was overall lower. Most dramatically, the animal showed down-regulated weathervaning. In wild-type (N2) animals the distribution of the curvature of the animal's trajectory depends on its orientation to the local odor gradient (Fig 2.7f) indicating weathervaning, but in AIB-inhibited

animals the distribution of trajectory curvature is much less dependant on orientation to the local gradient (Fig 2.8d). Together this suggests that AIB is important for supporting turning rate and the weathervaning strategy, and demonstrates the potential for this approach to characterize mutants or investigate neural circuits underlying navigation.

Having access to quantitative information about the odor the animal experiences provides an opportunity to characterize the animal's overall gradient climbing performance through observations of its local movements on a local gradient[170]. We compute the animal's drift velocity as a function of local gradient steepness (Fig 2.9). This captures the animal's overall gradient climbing performance as a result of all the navigational strategies it uses, including the biased random walk and weathervaning. The analysis provides information relevant to a chemotaxis index, but it has the additional benefit of describing quantitatively how the animal responds to a local gradient and on what timescale. This analysis is only possible with a knowledge of the odor concentration experienced by the animal, and is naturally accessible in our instrument.

2.2.7 Butanone chemotaxis in *Drosophila* larvae

To further evaluate the utility of the flow chamber for the study of small animal navigation, we investigated how larval *Drosophila* navigate butanone. Although butanone is not as commonly used as a stimulus with *Drosophila* as with *C. elegans*, butanone is known to be attractive to larval flies[74,75] and has been variously reported to be attractive[190,127] and aversive[113,147] to adult flies. To investigate the larva's navigational strategy in a butanone gradient, we created a "cone" shaped butanone gradient over the agar substrate using the pre-equilibration protocol, as before, and we confirmed the presence and stability of the gradient by continuously measuring the spatial distribution of butanone upstream and downstream of the agar arena. We monitored the orientation and movement of 59 larvae over 6 separate 10 minute experiments ($\sim$ 10 larvae per experiment) with an average observation time of 7 min per larva (Fig 2.10).

Larvae moved towards higher concentration of butanone(Fig 2.10a). To analyze the strategy by which they achieved this, we first constructed a coordinate system in which 0 degrees was in

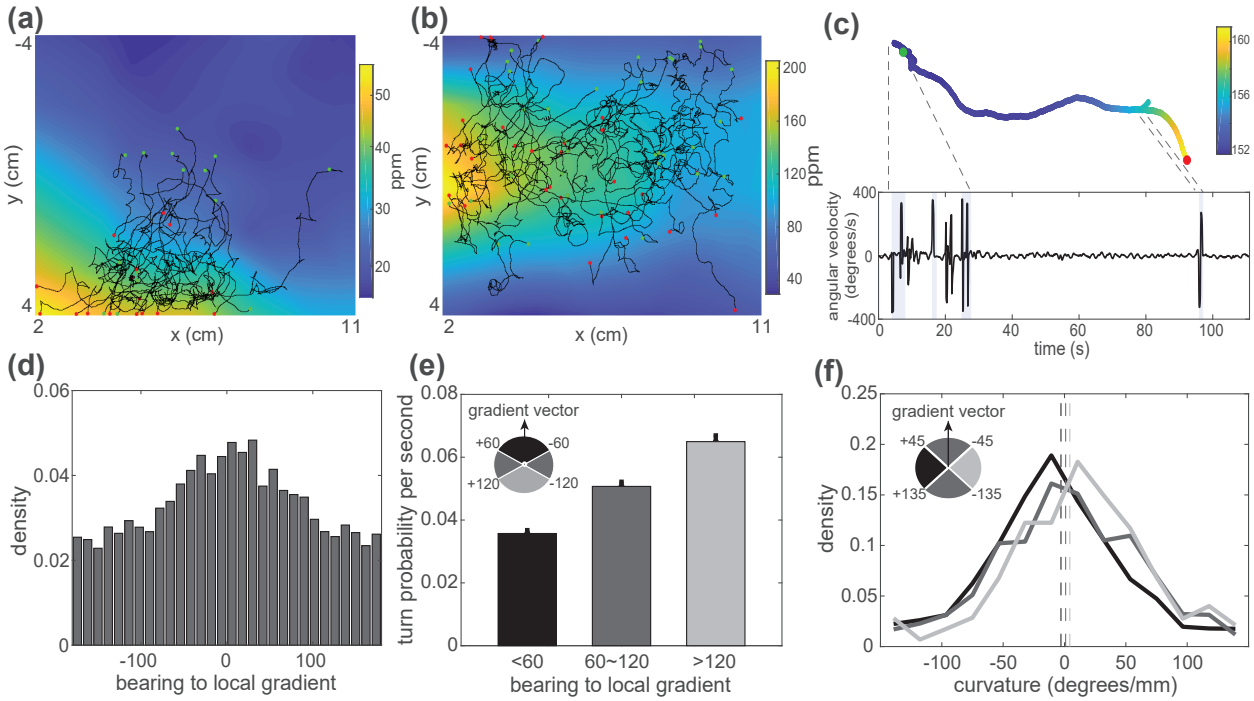

Figure 2.7: ***C. elegans* use both biased random walk and weathervaning to navigate in a butanone odor landscape.** Animals on agar were exposed to butanone in the flow chamber. **(a,b)** Measured animal trajectories are shown overlaid on airborne butanone concentration for different odor landscapes. Green dots are each animal's initial positions and red dots are the endpoints. **(c)** An animal's trajectory is shown colored by the butanone concentration it experiences at each position (top). Its turning behavior is quantified and plotted over time. Turning bouts are highlighted in gray. **(d)** Distribution of the animal's bearing with respect to the local airborne odor gradient is shown. Peak around zero is consistent with chemotaxis. **(e)** Probability of observing a sharp turn per time is shown as a function of the absolute value of the bearing relative to the local gradient. Modulation of turning is a signature of biased random walk. Error bars show error for counting statistics. Data analyzed from 5 experiments, $\sim$100 worms per experiment, resulting in over 90 animal-hours of observation. **(f)** Probability density of the curvature of the animal's trajectory is shown conditioned on bearing with respect to the local gradient. Weathervaning strategy is evident by a skew in the distribution of trajectory curvature when the animal travels more perpendicular to the gradient. Bearing to local gradient defined by gray-scale quadrants. Three distributions are significantly different from each other according to two-sample Kolmogorov-Smirnov test ($p < 0.001$). Means are shown as vertical dashed lines.

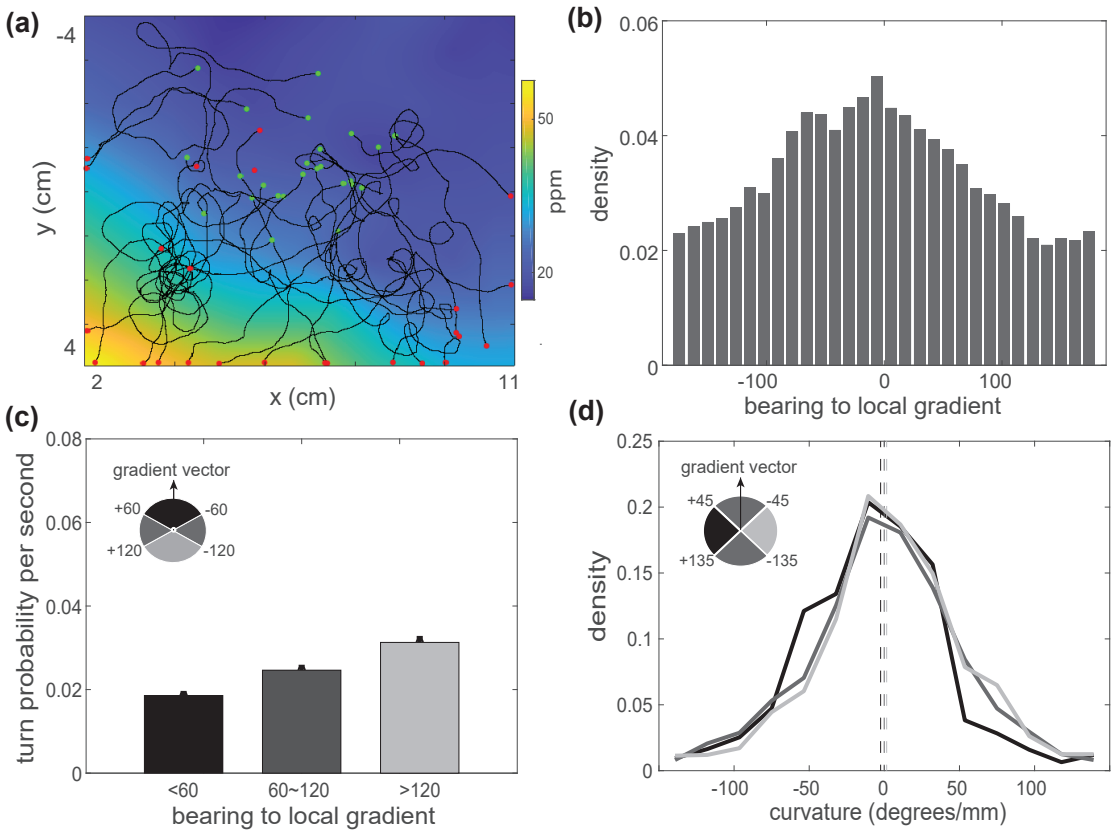

Figure 2.8: **Butanone chemotaxis in worms with down-regulated AIB neuron. (a)** Example behavioral trajectories from worms with down-regulated AIB neuron in the same butanone landscape. Green dots are each animal's initial positions and red dots are the endpoints. **(b)** Distribution of bearing to local gradients. **(c)** Turning probability at different bearing conditions. **(d)** Curvature conditioned on different bearing measurements. Data analyzed from 3 experiments, $\sim$100 worms per experiment, resulting in over 51 animal-hours of observation.

the direction of the odor gradient (towards higher concentration) and 180 degrees was directly down-gradient; angles increased counterclockwise when viewed from above. We found that larvae initiated turns at a higher rate when headed down-gradient ($\pm 180°$ bearing with respect to the local gradient) than up gradient ($0°$) (FIGfig7b). When larvae turned, their reorientations tended to orient up gradient (negative angle changes from $+90°$ bearing with respect to the local gradient, and positive angles changes from $-90°$) (Fig 2.10c). Bias in turning rate and reorientations in larvae are known navigation strategies in larvae[85,90,89], but here it is to our knowledge the first characterization for butanone. With modulation of turning rate and orientation in the same odor environment, we conclude that *Drosophila* larvae use similar navigational strategies to *C. elegans*

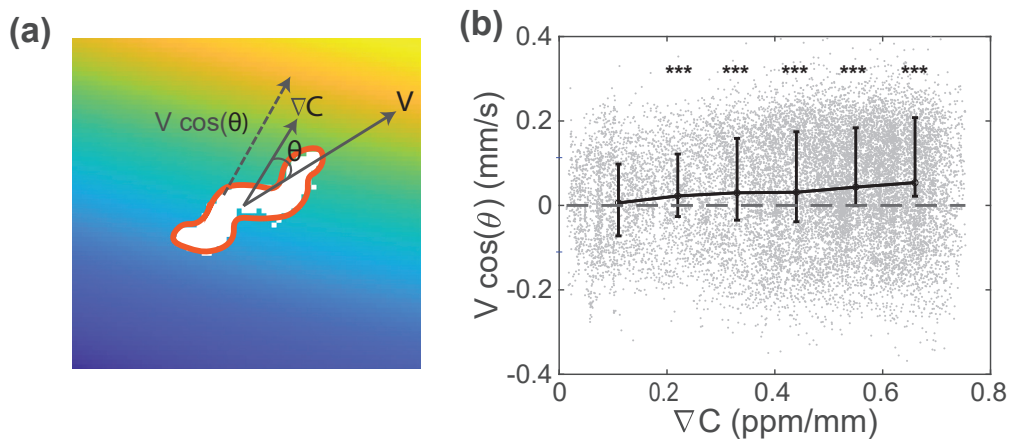

Figure 2.9: **Chemotaxis performance as a function of gradient steepness. (a)** Schematic of a worm tracked in the odor landscape. The crawling velocity vector V, local concentration gradient ∇C, bearing angle θ, and drift velocity $V\cos(\theta)$ are shown. **(b)** Tuning curve shows the drift velocity $V\cos(\theta)$ as a function of the odor concentration gradient. Gray dash line indicates an unbiased performance with zero drift velocity, gray dots are the discrete measurements, and the black line shows the average value within bins. Error bar shows lower and upper quartiles of the measurements. Other than the left-most bin with the lowest gradient, all other bins show significant difference from zero drift velocity using t-test ($p<0.001$, ***). Both two-sample t-test and Kolmogorov-Smirnov test (given the non-Gaussian distribution) show significant differences between bins and the first bin with the lowest concentration gradient ($p<0.001$).

to move towards butanone. We note that the chemotaxis efficiency is much higher in larvae compared to worms given the order of magnitude difference in run velocity ($\sim$ 3 mm/s in larvae and $\sim$ 0.1 mm/s in worms).

2.3 Discussion

We present a custom-designed flow chamber and odor sensor array that enables us to measure navigation strategies of worms and fly larvae within the context of a controlled and measured odor environment. The key features of this odor delivery system are that (1) the odor concentration profile through space is controlled in the flow chamber, (2) the odor sensor array provides a spatial readout to calibrate and measure the profile, and (3) the odor concentration profile is monitored during animal experiments. This last feature, the ability to monitor the spatial profile of odor concentration on the boundary during experiments, to our knowledge sets this method apart from previous approaches.

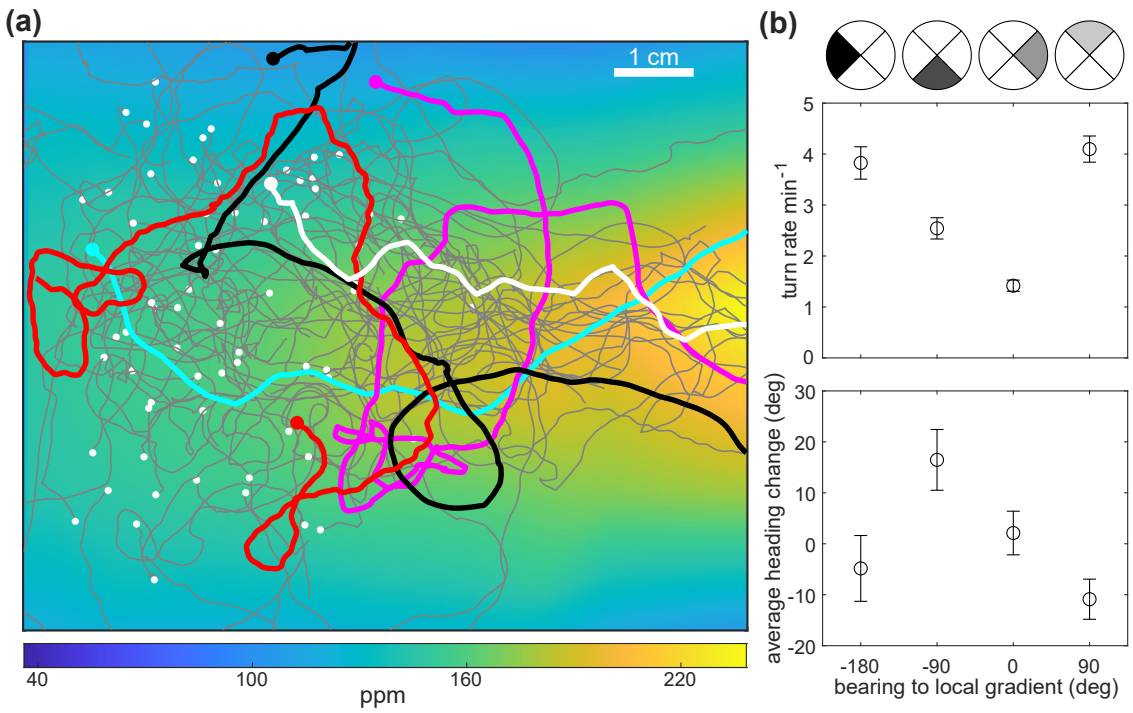

Figure 2.10: ***D. melanogaster*** **larvae chemotaxis in the odor flow chamber. (a)** Trajectories overlaid on the measured butanone odor concentration landscape. Example tracks are highlighted and the initial points are indicated with white dots. **(b)** Top: turn rate versus the bearing, which is the instantaneous heading relative to the gradient defined by quadrants shown on top. Error bar show counting statistics. Bottom: average heading change versus the bearing prior to all turns (re-orientation with at least one head cast). Error bars show standard error of the mean. Data analyzed from 6 experiments, 59 animals, with 620 turns resulting in over 6.8 animal-hours of observation.

The odor landscape experienced by the animal can be recovered one of two ways, with different tradeoffs. Typically, for a given set of flow parameters, we first quantify the odor gradient using the whole sensor array when agar is absent, then during behavioral experiments, we *confirm* the desired odor profile is achieved using sensor bars mounted along the boundary upstream and downstream of agar. During experiments with agar, the real-time monitoring of the odor profile at the edges of the arena guards against a range of possible errors or experimental imperfections. For example, exhaustion of the odor source, blocked connections, imperfect seals, and so on will all be detected by a change in the odor profile along the boundary, thereby avoiding silent changes to the odor landscape that may skew behavioral results. Indeed, it was this real-time monitoring that first alerted us to the fact that absorption of odor into the agar substrate

caused the odor landscape to differ from our naive expectations.

If real-time monitoring of the entire arena during behavioral experiments is desired, the odor array can instead be mounted opposite the agar surface (Fig 2.6). We show that this introduces complications in imaging and makes tracking animals challenging Future work could address these challenges through alternate lighting schemes (e.g. frustrated TIR in[209]) or by masking the sensor array with a darker material. For most applications, however, we expect that confirming the profile via sensors along the boundary will be sufficient because we have demonstrated that the odor landscape established by flow at quasi-equilibrium is stable and repeatable in the absence of experimenter or instrument errors (and detectable when not).

The ability to monitor spatial profile during experiments, along with a quantitative understanding of odor-agar interactions, provides confident knowledge of the odor experienced by the animal over time. This is especially important when we analyze time-varying navigation behavior in an odor environment that potentially changes in time too[269,272,150]. Past work focusing on neural circuit mechanisms for odor navigation in worms have often relied on coarser readout such as the chemotaxis index or using the distance to odor source as a proxy for the inverse of odor concentration. However, it has been shown that the exact odor concentration alters the animal's behavioral response to odor[272,150]. Quantifying tuning curves with respect to more precise odor measurements may better constrain the investigations into neural mechanisms driving the sensorimotor transformations underlying navigation algorithms[52,154].

In contrast to liquid delivery of odor gradients via microfluidic chips[3,145], our method allows worms to crawl freely on an agar surface. This allows our behavior measurements to be directly compared against classical chemotaxis assays[12,159,192]. Additionally, the macroscopic odor airflow chamber makes it straightforward to flexibly adjust the spatial pattern between experiments without the need to redesign the chamber. A detailed tutorial for assembling the instrument and the protocol for experiments is provided in Supplementary File in the paper.

Our setup uses low flow rates corresponding to low wind speed velocities (5 mm/s) to avoid anemotaxis. Larger organisms, including adult flies, navigate towards odor sources by combining

odor and wind flow measurements[252,169]. Fly larvae exhibit negative anemotaxis at wind speeds 200 to 1000 times higher than those used here[126], but previous work showed that they do not exhibit anemotaxis at lower windspeeds like the ones used here[85]. For example, they do not exhibit anemotaxis at 12 mm/s which is still higher than the velocities they experience here[85]. Therefore we do not expect *Drosophila* larvae to exhibit anemotaxis under our flow conditions. *C. elegans* are not thought to respond to airflow. In experiments in aqueous microfluidic chips under flow, *C. elegans* move towards higher concentrations of attractant and do not respond to the flow of the liquid[3]. In agreement, we do not observe any evidence of *C. elegans* anemotaxis in our chamber in response to control experiments without odor and with wind speeds of 5 mm/s.

We focused on the odor butanone because it is important for a prominent associative learning assay[247,131]. Butanone has high vapor pressure and is highly soluble in water and therefore it interacts strongly with agar. In this work we show that this odor-agar interaction makes it difficult to *a priori* infer an odor landscape experienced by the animal when agar is present, but that continuously monitoring the odor profile on the boundary overcomes this challenge.

Odorants with less solubility and lower vapor pressure may not require pre-equilibration to reach quasi-equilibrium quickly, but instead may be able to stabilize across agar on a shorter time scale. For example, ethyl butyrate is less water soluble ($\sim$0.586 g/100 ml in water, compared to 27.5 g/100 ml for butanone), has lower pressure (11.3 mmHg compared to 78 mmHg for butanone) and has a smaller gas phase diffusion coefficient (5.6 x 10^{-3} cm^2/s compared to $\sim$69 cm^2/s for butanone). A droplet of ethyl butyrate has been shown to form a relatively stable odor landscape over a 5 minute timescale in a dish on agar without flow[238,159]. Here we demonstrate stable odor landscapes over longer timescales of at least 30 mins with a more volatile odor, butanone, by using our flow-based assay. A strength of the flow based approach with monitoring is that it should, in principle, allow for the use of a wider range of odors. Crucially, the continuous monitoring makes this approach more robust against other unforeseen odor interactions, such as interactions between an odorant and glass, aluminum or plastic.

Here we have addressed the problem of creating non-turbulent airborne odor landscapes. The

biophysical processes governing odor sensing in small animals such as *C. elegans* are not fully understood. The worm carries a thin layer of moisture around its body as it moves on the agar substrate[11] and it is unclear to what extent the worm pays attention to the concentration of an odorant in the agar below it vs the air above it. Our reaction-convection-diffusion model suggests that at the quasi-equilibrium conditions used in our experiments the odor concentration in agar is related to the airborne odor concentration directly above it up to a scalar that we predict to be constant across the agar. Although we have not measured this empirically, this suggests that even in the extreme case that the animal only senses odor molecules in the agar, the odor concentration experienced by the animal in our experiments should differ by no more than a scaling factor compared to our estimates based on the airborne odor concentration.

Here we also showed the possibility of and challenges associated with directly measuring the full spatial odor profile (instead of using the boundary) during simultaneous behavior tracking. Future work is needed to improve concurrent animal tracking methods, to design more miniaturized and scalable sensors, and to generalize to more complex or three-dimensional environments.

Knowing the concentration experienced by the animal is not only useful for measuring navigational strategies more precisely than in classical assays, like the droplet chemotaxis assays. It will also be crucial for studying *changes* in navigational strategy, such as those in the context of associative learning[50,247], sensory adaptation[150,114], and long time scale behavioral states[27,90,136]. In all those cases, it will be critical to disambiguate slight changes to the odor landscape from gradual changes in the navigational strategies. Continuously monitoring the odor landscape during behavior will remove this ambiguity.

Nothing in neurobiology makes sense – except in the light of behaviour.

Jerry Hirsch

3

Learning-dependent Navigation Strategies and Neural Computations in *C. elegans*

Learning is a feature of species across scales. Understanding how the brain flexibly learns to produce complex behavior is a crucial question in neuroscience. In this project, we focused on how the worm's navigation strategies depend on olfactory learning by utilizing precise measurements and a novel statistical model. The fitted model shows that learning alters two known navigation strategies in worms, and it outperforms a classical metric in decoding the animal's past learning experience. We discovered that learned navigation can express itself in a context-dependent manner.

In addition, through perturbing individual interneurons, we found that most neurons' contribution are distributed across strategies and are deferentially altered by learning. We expect these flexible behavioral algorithms and neural computations to be generalized to other species and behaviors. This work largely leverages the flow chamber apparatus and chemotaxis assays developed in the previous chapter. The optogenetic strain in this work comes from Anuj Sharma. The project is supervised by Jonathan Pillow and Andy Leifer. The work is in revision at *PLOS Biology* at the stage of thesis finalization[45].

3.1 Introduction

Learning is a fundamental property of neural systems that allows an animal to flexibly alter behavior based on experience. Learned olfactory navigation provides an ideal framework for studying learning because it is a naturalistic behavior, common to species across scales[10,195], ethologically relevant for seeking food[247,87] and avoiding pathogens[274,95], and can be rapidly learned in very few trials[87,139,172]. Careful characterization of learned navigation behavior in a naturalistic context can shed light on the flexible computation performed by the brain[52].

To study how animals flexibly alter their olfactory navigation upon learning, we focus on the nematode worm *C. elegans*. This worm has a compact nervous system[260], well-characterized olfactory neural circuits[11,198,56,77,61] and navigation behavior that has been studied in detail[11,192]. Worms learn to associate an odor with the presence or absence of food and then will either navigate towards higher concentrations of the odor or will ignore the odor, respectively[247,131]. However, it is unknown how the worms' navigation strategies are altered by learning to achieve these feats. We aim to quantitatively characterize the learned navigation strategy and their underlying sensory transformations to constrain neural mechanisms of learned navigation.

The worm navigates in sensory environments mainly through two strategies: klinotaxis and klinokinesis[192,110,111]. Klinotaxis is a process in which the worm continuously modulates its heading to align with the local gradient, also known as "weathervaning"[110]. Klinokinesis is a biased random walk[164,132] in which the worm produces sharp turns called "pirouettes" with a probabil-

ity that depends on the animal's estimate of the local gradient of a sensory cue[192,19,164]. In both cases the animal estimates information about the local gradient by comparing measurements of the stimuli across time. Both strategies contribute to sensory-guided navigation in landscapes of temperature[111], salt[110,161,117], and certain odors[110,11,54,241].

Olfactory navigation has significant relevance in studying learning because there are odors for which the animal learns to alter its "valence" (positive or negative) with respect to the odor—similar to associative learning in many other animals. This contrasts with context-dependent salt or thermosensing in *C. elegans* in which the worm instead learns a preferred salt concentration or temperature set point based on past experience[161,111]. For learned olfactory navigation, it is not known how the biased random walk and weathervaning strategies change upon learning, nor how the detailed quantitative transformations between sensory input and motor output are changed.

Quantitative analysis for airborne odor-guided navigation has been limited, in part because of experimental challenges in measuring detailed information about the odor concentration experienced by the animal. While past work estimated odor cues indirectly or with models[87,110], precise concentration measurements are required to characterize sensorimotor transformation in the sensory environment. Recently, however, new experimental methods to control and monitor the odor landscape experienced by small animals, such as the worm,[46,238,85], now make it possible to empirically constrain quantitative models of odor-guided navigation, as we pursue here.

Worms are intrinsically attracted to butanone, a volatile organic compound found in bacterial food in the worm's natural habitat[264]. When worms are exposed to butanone paired with food ("appetitive training"), they increase their attraction and are more likely to navigate towards higher butanone concentrations[247,131,50,198]. In contrast, when butanone is paired with starvation ("aversive training"), worms decrease their tendency to climb up butanone gradients in comparison to worms without exposure to butanone ("naive")[247,50]. The animal's neural and behavioral responses to butanone both change upon learning[50,198]. Sensory neuron AWC^{ON} responds to butanone[50,247], as well as others[152,198]. However, the involvement of downstream interneurons in butanone learning and learned navigation are still unclear.

In this study we seek to answer: (1) How are the worm's navigation strategies altered by olfactory learning? (2) How are sensorimotor transformations altered by learning and how does this vary across a population? (3) What neural substrates may be involved? To answer these questions, we combine precise experimental measurements using our recently developed continuous odor monitoring assay[46] and a novel statistical model to rigorously characterize how butanone associative learning alters odor navigation strategies in worms.

Our measurements and model reveals that the animal's biased random walk is bidirectionally altered by butanone learning and that its weathervaning strategy is down-regulated upon aversive learning. Our approach yields interpretable model parameters that better decode training conditions compared to chemotaxis index, and also predict response to optogenetic perturbation in the sensory neuron AWC^{ON}. We discover that naive worms have higher behavioral variability and demonstrate context dependent behavior. And we provide insights into the role of specific interneurons.

3.2 Results

3.2.1 Learning bidirectionally alters olfactory navigation

We developed a protocol to train worms to associate butanone with either food (appetitive training) or starvation (aversive training) (Fig 3.1a). Our protocol was similar to previously reported training regimens[131,50,198] except that ours exposes the animal to multiple rounds of odor paired with starvation instead of only one, which we found increases consistency in learning. After training, we recorded the movement of populations of worms as they crawled in a defined odor landscape that used metal-oxide sensors to continuously monitor the odor concentration along the boundary[46] (Fig 3.1b). We recorded hundreds of locomotory trajectories per plate in this odor environment after different training conditions, for up to 13 plates per training condition. Animal's locomotory trajectories were qualitatively different depending upon learning, with appetitive trained animals traveling up gradient more often than naive animals, and aversive-trained animals traveling up gradient least of all, broadly consistent with prior reports[50,247,198]

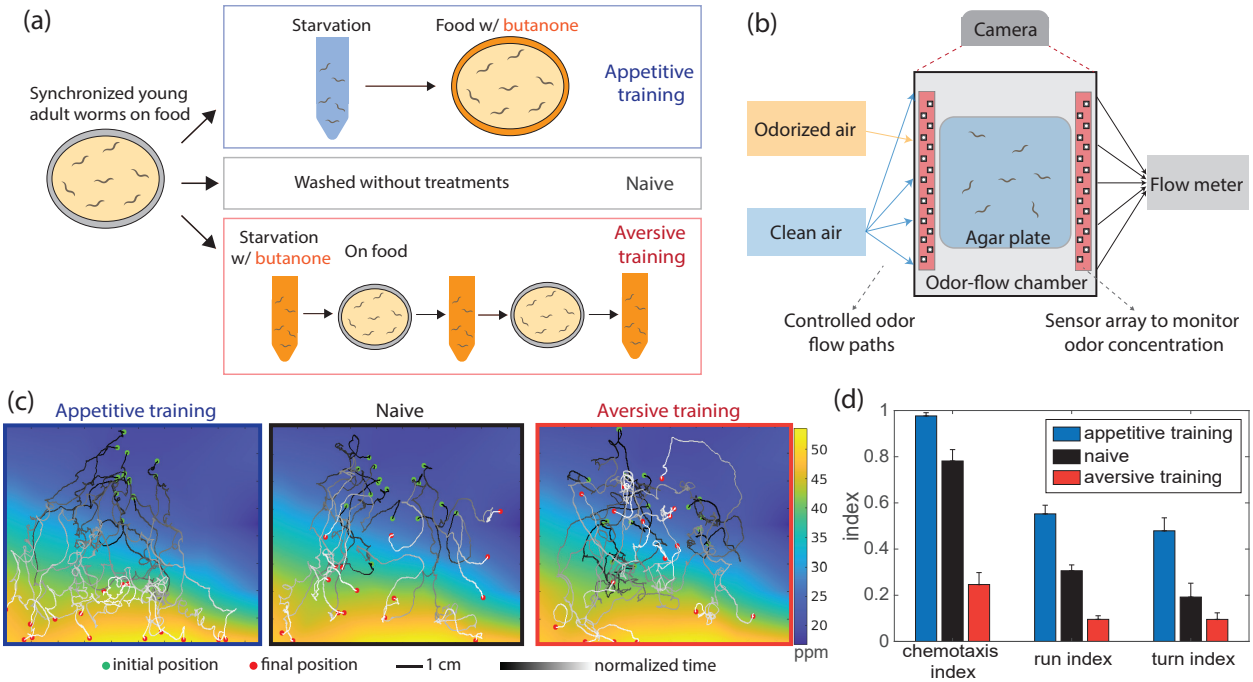

Figure 3.1: **Bidirectional olfactory learning in *C. elegans*.** **(a)** Protocol for butanone associative training in worms. **(b)** After exposure to different training regimens in (a), worms' olfactory navigation is measured in a controlled odor environment. **(c)** Example trajectory after three different training conditions. **(d)** Summary statistics of learning across three conditions. The chemotaxis index is the normalized number of tracks going up gradient N_u versus down N_d: $\frac{N_u - N_d}{N_u + N_d}$, which corresponds to the classic chemotaxis index. The run index is the normalized run length going up gradient r_u versus down r_d: $\frac{r_u - r_d}{r_u + r_d}$, and captures a key feature of a biased random walk. The turn index is the normalized probability turning up gradient p_u versus down p_d: $\frac{p_u - p_d}{p_u + p_d}$, which captures a key feature of weathervaning. Error bar shows standard error of mean across 7-13 plates for each condition. For each index, all pairs of training conditions (appetitive-naive, naive-aversive, and appetitive-aversive) show statistically significant differences (t-test, $p < 0.05$), except for the turn index when comparing naive to aversive conditions. All data underlying this figure can be found at https://doi.org/10.6084/m9.figshare.24764403

(Fig 3.1c). We quantify the performance of traveling up the gradient using a "chemotaxis index" that is calculated by comparing the number of trajectories going up-gradient versus down gradient[50,247] (Fig 3.1d). Performance navigating up gradients is bidirectionally modulated by learning: chemotaxis index increases after appetitive training and decreases after aversive training compared to naive worms that undergo no training. Interestingly, even animals that undergo aversive training do not, on average, navigate down the gradient, suggesting that after learning an association between butanone and starvation, animal are indifferent or possibly still slightly

attracted to butanone.

To explore whether worm navigation superficially resembles a biased random walk, we calculated a "run index" by computing the normalized length of a run moving up-gradient, where the run is defined as the period between pirouettes. To explore whether navigation superficially resembles a weathervaning strategy, we calculated a "turn index" that reports the normalized fraction of turn events that result in heading up-gradient[161]. Our measurements show that both the run index and the turn index are on average bidirectionally modulated by learning with respect to naive animals, suggesting that learning alters both of these navigational strategies.

While the metrics above suggest hypotheses about how navigational strategies change due to learning, they provide little information about the dynamics of navigation, nor the sensorimotor transformations that govern these dynamics. Specifically, the metrics use only binary information about whether the animal is traveling up or down gradient, and ignore details of the sensory landscape like the odor concentration experienced over time. The indices also provide no information about the behavioral variability across the population. To overcome these limitations, we sought a statistical model that captures temporal information, behavioral noise, and explicitly predicts how the animal changes its movement in response to sensory stimuli.

3.2.2 Odor-dependent mixture model of olfactory navigation

To characterize how olfactory sensory inputs are transformed to behavior under different training conditions, we developed a dynamic Pirouette and Weathervaning (dPAW) statistical model of worm olfactory navigation. The dPAW model consists of a mixture of two navigation strategies: a pirouette behavior consisting of an abrupt change in heading angle (a turn) and weathervaning which instead continuously modulates heading angle. The dPAW model describes how the worm implements and balances these two behavioral strategies depending on time-varying sensory inputs (Fig 3.2a). The dPAW model is an extension to a classic biased random walk, where the run intervals are replaced with weathervaning behavior. This framework explicitly models these strategies and is fit to detailed measurements of movement and odor-experience. This is to

our knowledge the first statistical model that explicitly captures the detailed changes to *C. elegans* navigation strategies upon butanone learning.

The model worm samples its heading change $d\theta$ at each time step from one of two distributions: either a "weathervaning distribution" $P_{\text{wv}}(d\theta)$ or a "pirouette distribution" $P_{\text{pr}}(d\theta)$. The weathervaning distribution is narrow, reflecting small changes in heading angle that result from the worm's recent measurements of the concentration gradient. Pirouette behavior, on the other hand, corresponds to large turns that the worm makes when it receives evidence that it is going in the wrong direction. The pirouette distribution is therefore broad, with a peak at $\pm\pi$, indicating a complete reversal in direction. The concentration-dependent "decision" to produce a pirouette in between runs forms a biased random walk[192,19].

The worm's decision to initiate a pirouette or to continue to weathervane on each time step is modeled with a Bernoulli generalized linear model (GLM) that takes the filtered history of the odor concentration $C_{1:t-1}$ and the worm's own movement history $d\theta_{1:t-1}$ as inputs. The output of this GLM is a binary variable β_t that indicates the presence of a pirouette. Thus, the worm samples its heading change from the pirouette distribution $P_{\text{pr}}(d\theta)$ if $\beta_t = 1$ and the weathervaning distribution $P_{\text{wv}}(d\theta)$ if $\beta_t = 0$. The full model can be written:

$$P(d\theta_t|C_{1:t-1}, dC^{\perp}_{1:t-1}, d\theta_{1:t-1}) = \\ P(\beta_t = 1)P_{\text{pr}}(d\theta) + P(\beta_t = 0)P_{\text{wv}}(d\theta \mid dC^{\perp}_{1:t-1}), \quad (3.1)$$

where

$$P(\beta_t = 1) = m + \frac{M - m}{1 + \exp(K_C \cdot \mathbf{C}_{1:t-1} + K_h \cdot |\mathbf{d}\theta_{1:t-1}|)} \quad (3.2)$$

is the mixing probability over the two distributions, with parameters m and M for the minimum and maximum probability of a pirouette on a single time bin, and K_C and K_h corresponding to filters on past odor concentration $\mathbf{C}_{1:t-1}$ and the past absolute angular change $|\mathbf{d}\theta_{1:t-1}|$ vectors,

respectively. The pirouette and weathervaning distributions are in turn given by

$$P_{\mathrm{pr}}(d\theta) = \alpha U[-\pi, \pi] + (1-\alpha) f(\pi, \kappa_{\mathrm{pr}}), \tag{3.3}$$

$$P_{\mathrm{wv}}(d\theta \mid dC^{\perp}_{1:t-1}) = f(-K_{dC^{\perp}} \cdot \mathbf{dC}^{\perp}_{1:t-1}, \kappa_{\mathrm{wv}}) \tag{3.4}$$

where U is uniform in the circular heading and f is a von Mises distribution with mean and precision parameter κ. In the pirouette distribution, scalar $\alpha \in [0, 1]$ is the weight on the uniform distribution and κ_{pr} is the precision parameter that determines the sharpness of the pirouette. In the weathervaning distribution, the mean is altered according to the perpendicular concentration change $dC^{\perp}_{1:t-1}$ and the precision parameter κ_{wv} determines the noise around the head angle. (Note we have made a simplification by allowing the model access to the instantaneous perpendicular odor concentration. In reality, the animal is thought to compute $dC^{\perp}_{1:t-1}$ from sequential measurements in time as it swings its head through space.)

We fit dPAW to measurements of animal movement in the odor arena, including the time varying headings $d\theta_t$, the concentration along the locomotion path C_t, and the concentration perpendicular to the locomotion path $dC^{\perp}_t$. All model parameters in dPAW are jointly inferred through a maximum-likelihood method. To validate that parameters are reliably inferred, we simulated example chemotaxis trajectories from pre-defined parameters, fit them by the model, and confirmed that the model accurately recovered the pre-defined parameters. In the rest of the paper we fit the model to measurements to explore how navigation strategies change with learning.

3.2.3 Model captures navigation altered by olfactory learning

To characterize how olfactory learning alters navigation strategies, we fit the dPAW model to trajectories measured in the butanone odor environment after different training conditions. We first confirmed that the model capture key aspects of the animal's navigation. We confirmed that the fitted model's estimate of pirouette frequency matches that measured empirically[192]. We also used the model to simulate chemotaxis behavior in the odor environment and confirmed that

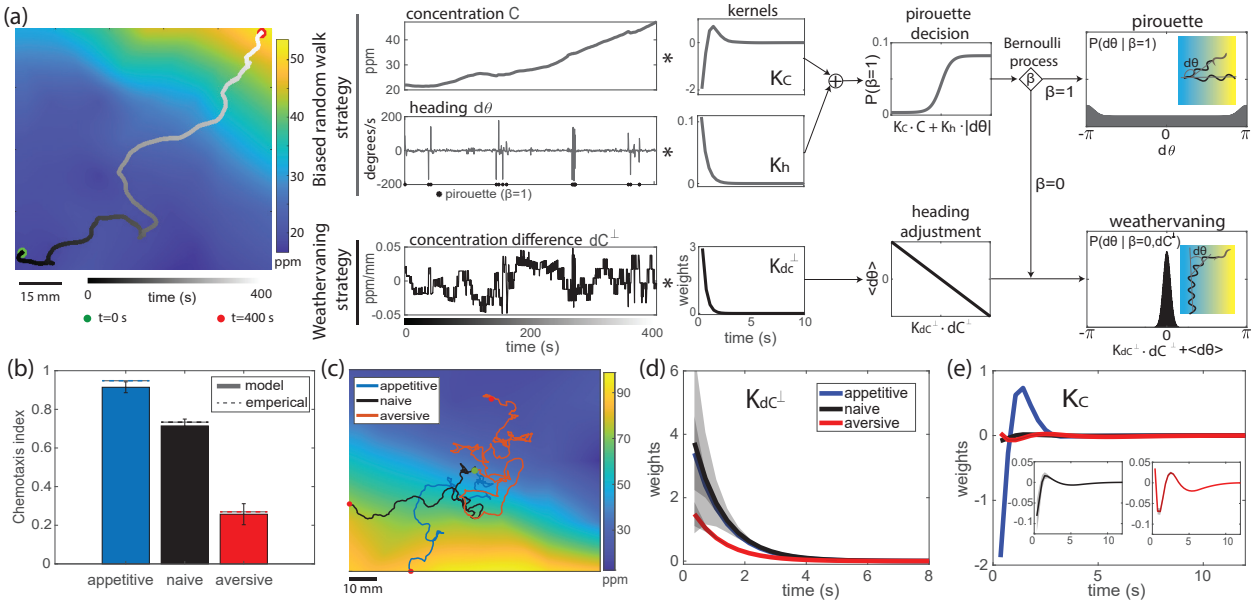

Figure 3.2: **dPAW captures learned olfactory navigation in worms. (a)** Schematic of the dPAW model. An example data trajectory is shown on the far left, providing time series of concentration and angle changes. Response kernels and decision functions are fitted to the time series data. The Bernoulli process β leads to two parallel strategies: biased random walk and weathervaning. **(b)** Chemotaxis index of simulated trajectories with inferred parameters. Error bar shows standard error of mean across 10 repeated simulation, each with 100 trajectories. **(c)** Example trajectories simulate from each training conditions. **(d)** Kernel $K_{dC^\perp}$ and **(e)** Kernel K_C fitted to three training conditions. Shaded area shows standard deviation of the kernel estimate. All data underlying this figure can be found at https://doi.org/10.6084/m9.figshare.24764403

model-generated trajectories have a chemotaxis index that agrees with measurement and is similarly bidirectionally modulated by training conditions (Fig 3.2b). Model-generated trajectories also appear visually similar to experimental observations (Fig 3.2c). And in further agreement, simulated trajectories recapitulate sensory and behavioral statistics measured in experiments, including distributions of the worms: heading angle, pirouette rate, experienced perpendicular odor concentration difference experienced, and tangential odor concentration.

Agreement between model and measurement was not due to chance. The dPAW model on average captures $0.3 - 0.6$ bits/s more information, depending on the animal's training condition, about navigation behavior than a null model that lacks any olfactory sensing mechanisms (Fig 3.3c).

3.2.4 Sensorimotor kernels change upon learning

It is unknown whether sensorimotor kernels, K_C and $K_{dc^\perp}$, should necessarily change upon learning[192,110,50]. For example, an animal could in principle change its chemotaxis index upon learning by altering its overall pirouette rate and the curvature of its runs without having to alter the kernels that govern its sensorimotor response. To test this, we inspected the kernels inferred from dPAW that were fit on animals that underwent different training conditions. Kernels corresponding to both weathervaning and biased random walk strategies are altered by learning (Fig 3.2d,e). The weathervaning kernel $K_{dC^\perp}$ has lower weights after aversive training compared to appetitive trained or naive animals (Fig 3.2d), suggesting that the animal's heading angle is less tightly dependent upon the concentration difference perpendicular to its path. In other words, aversive trained worms don't pay as much attention to perpendicular concentration when choosing their heading angle. There are two extreme ways these worms could not pay attention to the sensory cue: the worm could be uncoordinated and randomly select a heading direction, or alternatively, it could keep its existing heading. We observe the latter. We found that aversive-trained worms alter the behavioral noise and are more likely to preserve their existing heading and exhibit higher persistent length during their runs, as indicated by higher fitted precision parameter of weathervaning κ_{wv} (Fig 3.3b).

The kernels corresponding to pirouette decisions, K_C, change for both aversive and appetitive training compared to naive animals (Fig 3.2e). The amplitude is increased after appetitive training compared to naive animals, suggesting that these animal's pirouette probability more strongly depends on the concentration change along the navigation path than for naive animals. After aversive training, the kernel K_C has a longer time delay and forms a tri-phasic shape, markedly different from the biphasic shape observed in naive animals, suggesting that aversive trained animals may not be responding as much to downward changes in concentration (Fig 3.2e).

Collectively, we show that *C. elegans* alter their chemotaxis upon learning by changing the kernels that govern their sensorimotor response.

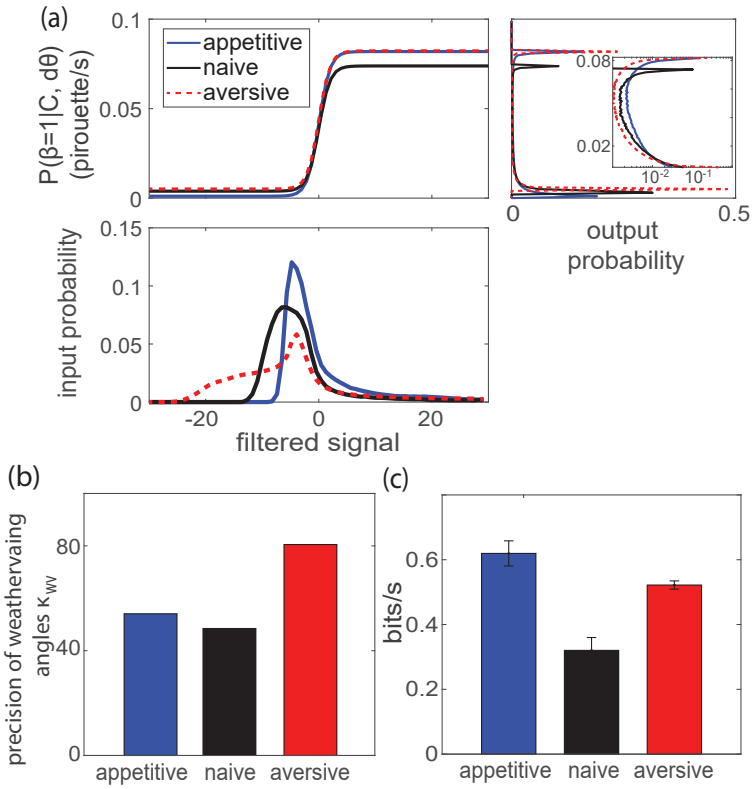

Figure 3.3: **Learning alters pirouette decision and behavioral noise. (a)** The top panel shows decision function $P(\beta = 1|C, d\theta)$ of three training conditions as a function of filtered signal: $K_C \cdot C + K_h \cdot |d\theta|$. The distribution of the filtered signal is shown in the bottom panel. The right panel shows the distribution of the output pirouette rate, with inset showing the same distribution in log scale. **(b)** The precision parameter for weathervaning, κ_{wv} across three training conditions. **(c)** The information rate given fitted model parameters and data across three training conditions. Error bar shows standard error of mean across 10 sampled batches of trajectories. All data underlying this figure can be found at https://doi.org/10.6084/m9.figshare.24764403

3.2.5 Decision function governing pirouettes changes upon learning

To understand how the animal alters its preference for continuing to weathervane versus interrupting weathervaning with pirouettes, we compared the pirouette decision function in equation A.16 inferred from our measurements before and after learning.

We found that learning alters the input-output statistics governing the initiation of a pirouette. This is clear from inspecting the distribution of the filtered signal (Fig 3.3a, bottom, related to odor concentration and past behavior) that serves as input to the decision function. Appetitive training and aversive training pushes the tail of the filtered signal probability distribution in op-

posite directions with respect to naive condition, which changes the statistics of pirouettes. This could reflect either changes to the kernels, or changes to the environment that those animals prefer to explore.

Aversive-trained worms have higher baseline turning rate m than appetitive trained worms, as seen in the output probability (Fig 3.3a, right), which is the result of passing the filtered odor signal (Fig 3.3a, bottom) through a nonlinear decision function (equation A.16 and Fig 3.3a, top). Both aversive- and appetitive-trained worms have higher maximum pirouette rate M compared to naive worms, reflecting a change to the decision function itself.

It is interesting to note that appetitive (but not aversive) trained animals spend more time with their pirouette probabilities in the most sensitive range of the sigmoid (Fig 3.3a right, inset), possibly indicating a more efficient strategy for chemotaxis. Capturing these details is one of the ways that dPAW is able to better detect changes in learning.

3.2.6 Model outperforms other metrics at decoding learned experience

The changes to sensorimotor processing detected by the dPAW model all provide information about the animal's past experience. We therefore wondered whether the model could accurately predict the training (aversive, appetitive or naive) that a population of worms had experienced. To test this we inspected dPAW models fit to measured trajectories from each training conditions $\gamma \in \{\text{appetitive}, \text{naive}, \text{aversive}\}$ corresponding to fitting parameters Θ_γ and made maximum likelihood predictions of the training condition given held-out test trajectories: $\hat{\gamma} = \arg\max_\gamma P(d\vec{\theta}|\vec{C}, d\vec{C}^\perp; \Theta_\gamma)$. On held out data, the fitted model correctly predicted training condition with a performance well above 90%, significantly above chance levels (Fig 3.4a), given sufficiently long recordings.

The inferred kernels and decision function within Θ_γ better reflect the worm's learning than either a classic chemotaxis index that captures the fraction of trajectories that go up-gradient (Fig 3.4a), or the concentration difference along the track. Indeed, dPAW always outperforms the chemotaxis index at decoding past training, for any tested amount of finite data. The model's

predictive power is derived from observing both the odor-history experienced by the animal and the corresponding behavior responses. Models supplied with either odor-history or behavior but not both fail to perform as well.

Some training conditions are more challenging for dPAW to decode then others. Naive worms show the lowest predictive performance with finite data when decoding is performed for each training condition separately (Fig 3.4b). This is consistent with our estimate of lower information exhibited by the trajectories of naive worms than trained worms (Fig 3.3c). Practically, this means more measurements are required to capture navigation behavior in naive worms than in trained worms.

3.2.7 The same computations govern natural and optogenetic odor stimuli

We wondered whether the same underlying computations that governed the animal's response to natural odor stimuli also govern its response to optogenetic-induced sensory stimuli. Optogenetics stimuli are not bound by the same natural statistics that the worm encounters when exploring a physical odor arena. For example, worms in our arena always experience temporally correlated and slowly varying odor responses. This introduces potentially confounding temporal correlations that can be avoided by using optogenetic stimuli[50,83,99]. We therefore investigated response to optogenetic induced odor sensation after learning and adapted our model to incorporate both forms of stimuli.

The animal's behavioral response to optogenetic stimulation was bidirectionally altered by learning (Fig 3.5). We measured the animal's absolute change in heading angle $|d\theta|$ in response to optogenetic stimulation of neuron AWC^{ON} expressing ChR2 (Fig 3.5b). AWC^{ON} is a butanone sensitive neuron known to play an important role in learning[50,198,131]. We measured behavior response to optogenetic stimuli two ways: we delivered pulses of optogenetic stimuli to animals experiencing odor in the arena (Fig 3.5), and we delivered time-varying intensity white noise optogenetic stimuli to animals off odor. In both cases (Fig 3.5e) the animal's turning behavior was more tightly coupled to optogenetic stimuli for appetitive trained animals than naive worms, and

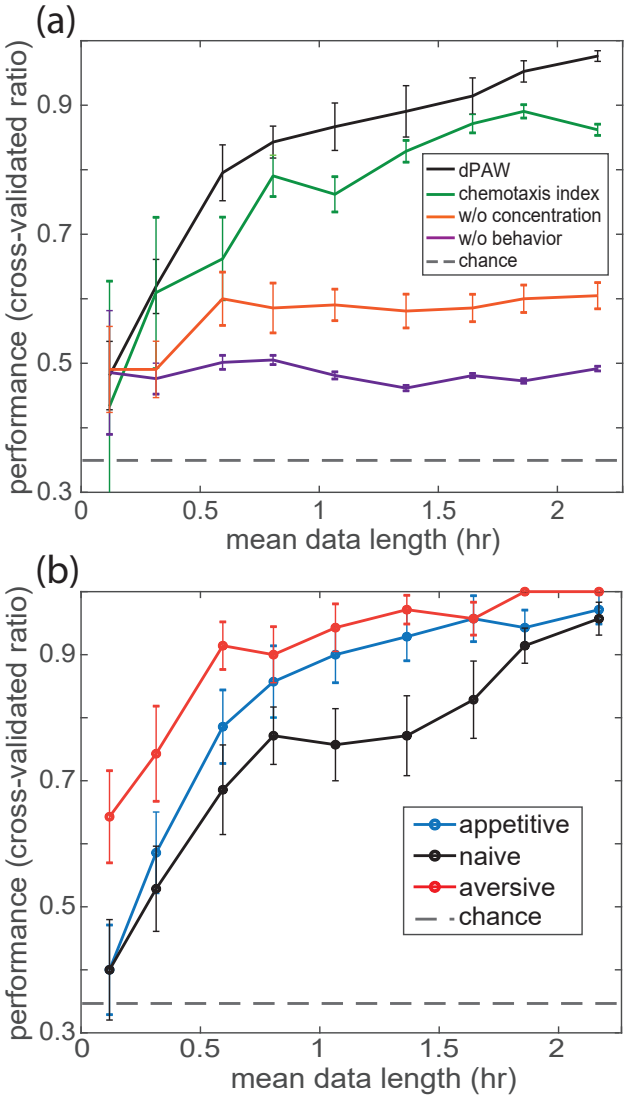

Figure 3.4: **Model-based decoding of learning. (a)** Model performance classifying the prior training of a population of worms (aversive, appetitive or naive) as a function of mean data length. Four models are compared and error bars show standard error of mean across 7-fold cross validation and 10-fold sampling across data length. w/o concentration: model that only captures behavioral statistics and does not account for odor input. w/o behavior: model for concentration difference along each trajectory and does not account for behavioral output. Chance level, 33%, is shown in grey dash line. **(b)** dPAW-based decoding as a function of data length as in (a), but here performance for each training condition is reported separately. All data underlying this figure can be found at https://doi.org/10.6084/m9.figshare.24764403

less so for aversive trained animals. A change in behavior response is consistent with prior reports[50], but note that here we make a more stringent comparison by comparing both appetitive and aversive trained animals against the same naive control condition.

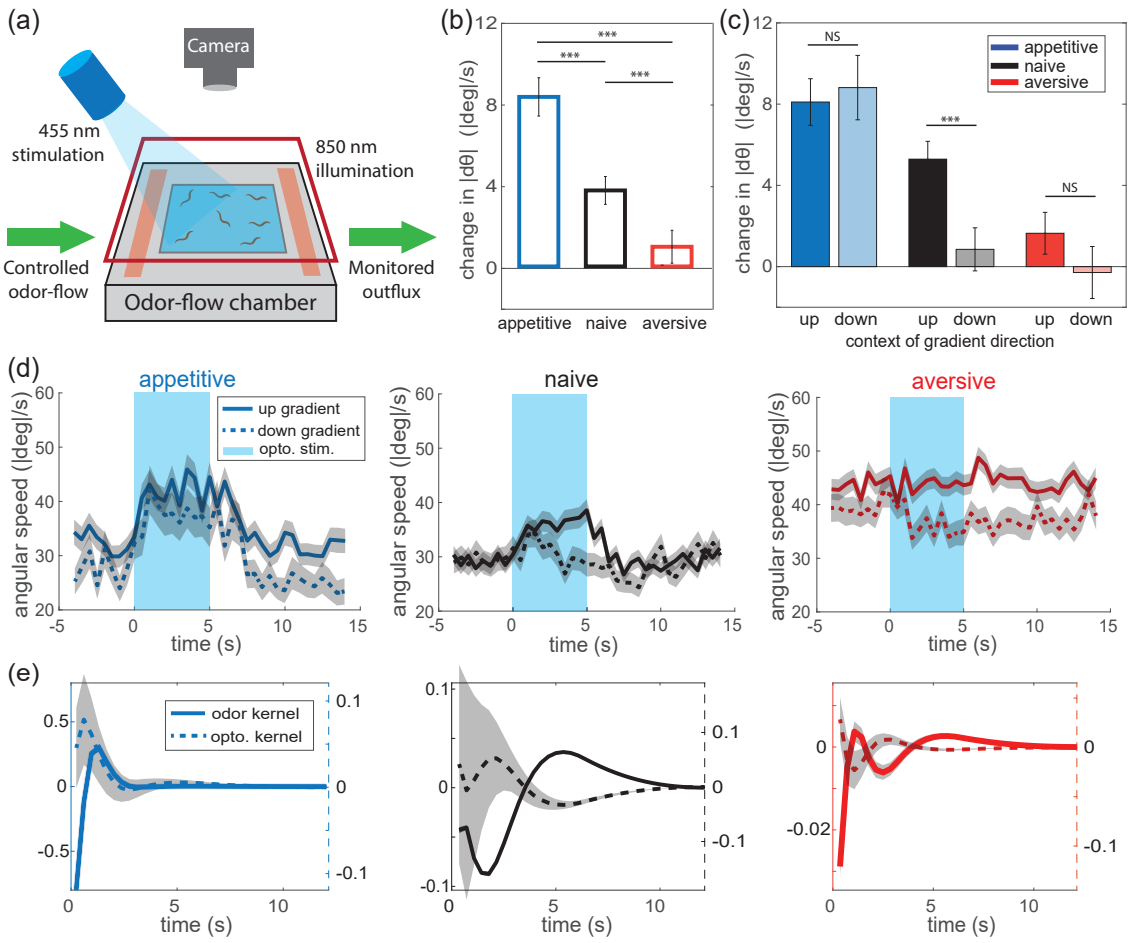

Figure 3.5: **Response to optogenetic induced odor sensation is altered by learning and depends on odor gradient context. (a)** Optogenetic stimulation is delivered to animals expressing ChR2 in AWC^{ON} as they crawl in an odor-flow chamber. **(b)** Change in the absolute heading $|d\theta|$ upon optogenetic impulse across the three training conditions. Change is computed between the average response during the 5s impulse and 5s before it. We record from 4-7 plates per condition. Error bar shows standard error of mean across over 1000 impulses per condition. The results of t-test comparison between training conditions are shown with $***$ for $p < 0.001$. **(c)** Same measurements as (b) upon optogenetic impulse, but with the context of gradient direction. Within each training conditions, measurements are separated into trajectories going up or down gradient. The results of t-test comparison between context are shown with $***$ for $p < 0.001$ and NS for non-significance. **(d)** Time-varying average angular speed aligned to the optogenetic stimuli. Three panels show different training conditions. Solid line and dash line indicate up and sown gradient trajectories. Gray shading shows standard error of mean across traces. **(e)** Temporal kernels for odor K_C and optogenetic K_{opto} fitted to each training conditions. The grey area shows standard deviation around the estimated kernels. All data underlying this figure can be found at https://doi.org/10.6084/m9.figshare.24764403

We extended our model so that light intensity contributes to the pirouette probability during navigation:

$$P(\beta_t = 1) = \frac{M - m}{1 + \exp(K_{odor} \cdot \mathbf{C}_{1:t-1} + K_{opto} \cdot \mathbf{I}_{1:t-1})} + m \tag{3.5}$$

where kernels K_{odor} and K_{opto} are weights on vectors of odor concentration $\mathbf{C}_{1:t-1}$ and light intensity $\mathbf{I}_{1:t-1}$, respectively. Since there is no difference along the perpendicular direction for this five second long spatially uniform optogenetic pulse, this model is simplified with kernels weighting the tangential concentration and neglecting the perpendicular concentration for weathervaning.

Across different training conditions, the optical kernel is close to the mirror image of the odor kernel, most strikingly demonstrated in the naive and aversive conditions (Fig 3.5e). Therefore the sensorimotor computation inferred from freely moving animals is predictive of the response to external perturbations. The inversion is expected and can partly be explained by the biophysics of the AWC^{ON} neuron: it hyperpolarizes when odor is present and depolarizes when odor is removed. This gives us confidence that our findings about sensorimotor processing derived from natural odor stimuli should be relevant to the larger literature based on more artificial stimulation. We therefore will leverage optogenetics to probe behavioral variability during odor navigation.

3.2.8 Learning modulates behavioral variability in response to sensory perturbation

We sought to characterize the variability in learned odor-guided navigation, because variability is a known feature of sensory processing in worms[91]. We observe variability across collections of trajectories in learned navigation. For example, the kernel K_C that governs the timing of pirouettes varies across subsamples of our data, and seems to vary more for aversive-trained than appetitive-trained or naive animals.

Surprisingly, we discovered that animals respond to stimuli differently depending on whether they are traveling up or down an odor gradient and that this contributes to variability (Fig 3.5c,d). Naive worms respond more strongly to optogenetic impulses delivered when traveling up an odor gradient, than when traveling down the gradient (Fig 3.5c). Appetitive trained worms, by contrast, respond consistently to optogenetic impulses regardless of their direction of travel along the

odor gradient. Aversive trained worms show weak response to optogenetic impulse regardless of the gradient context. This indicates that learning alters behavioral variability, and that a response to stimuli can be context-dependent along the navigation trajectory.

3.2.9 Downstream interneurons differentially contribute to learned chemotaxis

We investigated interneurons that may be involved in implementing learned changes to navigation strategy. We focused on interneurons downstream of the odor-sensing neuron AWC^{ON} because optogenetically induced activity in AWC^{ON} is sufficient to recapitulate aspects of learned navigation (Fig 3.5b) and AWC^{ON} is known to play an important role in odor sensing[11], navigation[130,36,150], and butanone learning[247,50]. We selected five interneurons subtypes with direct synaptic inputs (chemical or electrical connections) from AWC^{ON} [56,260]: AIA, AIB, AIZ, AIY, and RIA, several of which are known to be involved in navigation[36,110,98,50,111,161]. For each neuron subtype we compared learned odor-guided navigation in wild type animals to that of mutants for which the neuron subtype was genetically ablated (via miniSOG) or down-regulated (via expressing activated potassium channel), as described in methods, Fig 3.6.

Ablating or down-regulating the downstream interneurons had wide-ranging and statistically significant effects on chemotaxis performance and inferred model parameter, including after learning (Fig 3.6b). Here we characterize chemotaxis performance with a weighted chemotaxis index (wCI) that differs from the chemotaxis index in Fig 3.1d by more heavily weighting trajectories that experience a large change in odor concentration and de-emphasizing tracks that experience little odor change (described in the methods). To capture changes in inferred model parameters we report a stimuli-normalized magnitude of the kernels K_C and $K_{dC^\perp}$ which correspond to the extent to which the animal uses the biased random walk (BRW) or weathervaning strategies (WV), respectively, Fig 3.6b.

Ablation of neuron subtype AIZ was most severe and eliminated gradient climbing behavior across all three conditions (wCI close to zero), consistent with its reported role in salt chemotaxis[110]. In general, though, neuron subtypes contributed differently to chemotaxis performance,

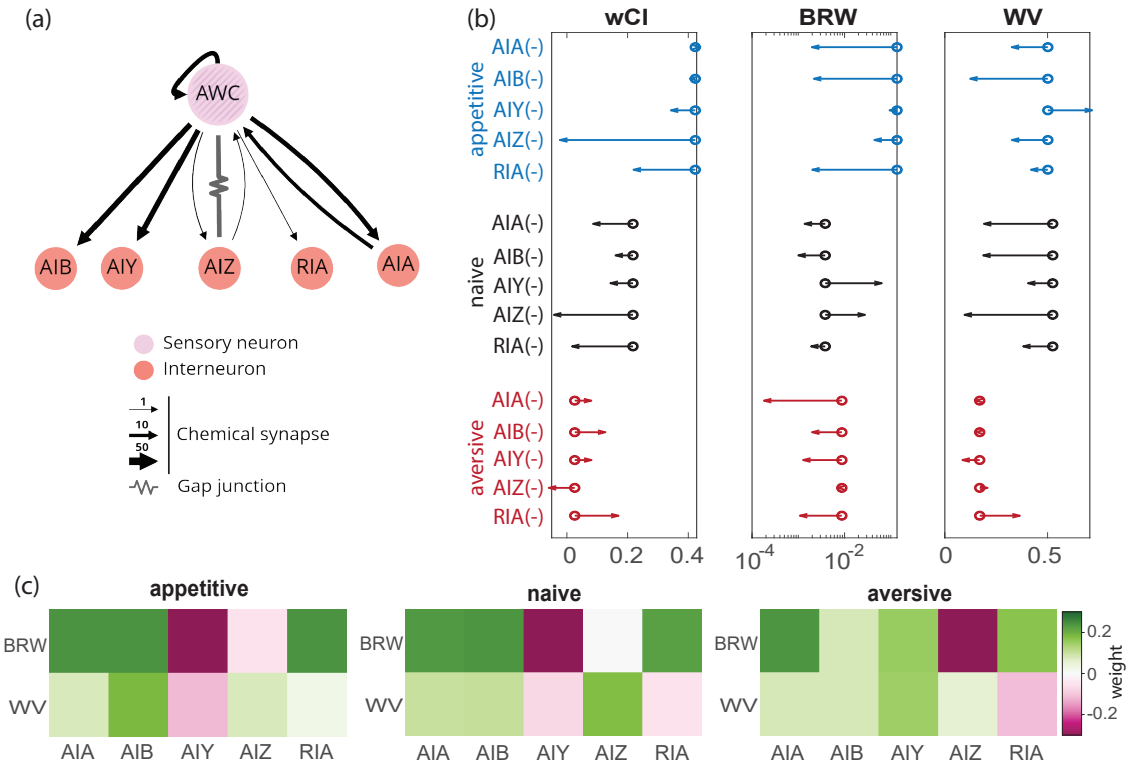

Figure 3.6: **Learned odor navigation in worms with disrupted interneurons. (a)** Five interneurons connected to the AWC sensory neuron. Schematic generated from nemanode.org. **(b)** Weighted chemotaxis index (wCI), biased random walk (BRW), and weathervaning (WV) indices computed from the inferred dPAW for all mutants across three training conditions. BRW and WV correspond to the norm of kernels K_C and $K_{dC^\perp}$ in dPAW. The circles show value from wild-type N2 worms and arrow shows the modulation in mutant worms. We record 3-5 plates for each stain and condition, resulting in 200-500 tracks in each measurement. **(c)** Weights in the neuron-behavior matrix $\mathbf{W}$ across three training conditions. The two rows are for different behavioral strategies and the columns are for different interneurons. All data underlying this figure can be found at https://doi.org/10.6084/m9.figshare.24764403

learning, or had different contributions to the inferred kernels corresponding to each navigation strategies. For instance, AIA, AIB and AIY defective animals showed little difference between naive and aversive training conditions, but did increase chemotaxis performance after appetitive training. The RIA defective animals have similar chemotaxis performance in upon appetitive and aversive learning, suggesting that RIA is involved in fine-tuning heading angle after learning experience. Interestingly, RIA defect results in much lower gradient climbing performance for naive animals. This suggests that RIA may be related to both learning and sensory adaptation in odor navigation.

We systematically quantified each neuron's contribution to the changes in the inferred behavioral strategies with a linear model (Fig 3.6c): $\mathbf{B} = \mathbf{WA}$, where $\mathbf{B}$ is the matrix with behavior readout (BRW and WV indices) along the rows and ablation conditions along the columns, $\mathbf{W}$ is an unknown weight matrix from neuron to behavior, and $\mathbf{A}$ is the neural ablation matrix with binary element with ones on the off diagonal elements. Matrix $\mathbf{A}$ is full-rank and makes the linear model solvable. We normalized the rows in $\mathbf{B}$ to better compare weights across behaviors.

The majority of the weights $\mathbf{W}$ were positive (Fig 3.6c), indicating that a neuron contributes to a given navigation strategy for a given training and that, conversely, disrupting the neuron down-regulates the navigation strategy. Interestingly some neurons, such as RIA, had negative weights. Unlike wild-type animals which decrease their weathervaning upon aversive training, RIA-defective animals fail to adjust the extent of their weathervaning upon aversive training. This suggests that RIA may act in a learning-dependent manner to decrease or increase odor-dependence of head angle, consistent with RIA's known role in controlling head motion[98].

The patterns of neuron weights were similar for naive and appetitive conditions, which both corresponded to wild-type positive chemotaxis. The patterns of neuron weights for aversive behavior, however was quite distinct, suggesting that the neural mechanisms of aversive learning may be quite distinct from that of appetitive learning.

Interestingly, no neurons had non-zero weights exclusively on one strategy (e.g. biased random walk) but not the either (e.g. weathervaning), suggesting that even though these two strategies are mathematically and behaviorally distinct, they do not segregate cleanly into different sub-populations of neurons. Instead the neural basis for learning dependant weathervaning and biased random walks are both distributed across the neural population.

3.3 Discussion

We combined precise experimental measurements and a novel statistical model, dPAW, to characterize learned olfactory navigation in worms. The results show that (1) navigation strategies are bidirectionally altered by learning, (2) the dPAW model decodes the animal's past training con-

ditions based on the observed navigation behavior and outperforms a classic chemotaxis metric, (3) behavior is more variable in naive worms and this variability can in part be explained by a context-dependent response to odor, and (4) interneurons downstream from AWC^{ON} contribute to learned navigation strategies in a distributed manner.

To reach these conclusions, two experimental advances were critical. The first was the measurement of navigation in a precisely defined sensory environment enabled by our recently developed odor delivery system[46]. Previous studies had analyzed navigation based on the proximity to a droplet source[50,247,198], but in those experiments the precise odor concentration experienced by the animal was typically not known. With the odor delivery system, we obtained the odor concentration experienced by the animal and therefore were able to fit more comprehensive models of sensorimotor processing.

The second methodological advance is the development of a more robust training protocol to probe bivalent learning (aversive and appetitive) that shows clearer effects when tested in the same odor concentration range. Prior investigations into butanone learning probed different valence learning (appetitive vs. naive, or aversive vs. naive) in different concentration regimes[131,50,247,40], possibly because the behavior effects of aversive learning are known to be more pronounced at lower odor concentrations of butanone while appetitive is known to be more pronounced at higher concentrations of butanone (see for example Fig 1- figure supplement 2 of[198] and results in[272]). But here we sought to directly compare detailed properties of learning, such as kernels, in the same sensory environment. To do so, we changed the training protocol (more repetitions during aversive than appetitive) in order to achieve a greater difference in learning outcomes. This produced large bivalent training-dependent changes to chemotaxis that were visible even in the same sensory environment (Fig 2.1).

The dPAW model fitted to measurements provides new insight into how the animal responds to time varying sensory stimuli. In particular, we find that learning alters the temporal kernels for odor input. For instance, appetitive training sharpens the tangential concentration kernel (Fig 3.2). By contrast, classical approaches have missed this important change because they im-

plicitly have static sensing kernels and *a priori* fix the time window to calculated gradients[110,161,111,192].

Our measurement and model reveals that learning alters not only the sensing mechanism but also the statistics of behavioral noise (Fig 3.3), which is consistent with recent work showing that starvation and neuromodulation can dramatically alter the statistics of behavior[246,187]. Our finding was possible only because dPAW explicitly includes parameters that control the noise level of the behavioral output, which is another advantage over past work which often excludes noise and relies on predetermined parameters[110,65,192]. Future work is needed to pinpoint the source of behavioral stochasticity, such as noise in the sensory or motor circuits, and its potential functional roles in navigation and exploration[14,91].

A strength of our approach is that it allows us to learn properties of sensorimotor kernels from either artificial optogenetic stimulation or natural odor stimuli. Past work, by contrast, required choosing one approach or another[83,130,50]. Optogenetic stimulation has advantages because it can be finely manipulated to deliver rich informative stimuli, for example white noise stimuli to study sensory encoding[193,226]. But there are challenges to connect optogenetic stimuli to natural stimuli experienced during navigation. Our work shows directly that information from one approach is compatible to the other. We applied statistical inference directly to navigational trajectories in the presence of odor and qualitatively recover similar sensorimotor transformations to those we inferred using optogenetic perturbation (Fig 3.5). This is to our knowledge the first direct comparison between sensorimotor computation during navigation and response upon external perturbation.

An important conclusion from optogenetic perturbations is that the worm's response variability is altered by learning. We found that naive worms are more variable and that learning reduces variability in behavioral strategies across the population of worms. Surprisingly, we further found that the variation across optogenetic responses in naive worms can in part be explained by the context of that worm's odor environment, up or down the gradient (Fig 3.5). Naive worms that travel up gradient pay more attention to both odor and optogenetic stimuli (Fig 3.5c) compared to those that go down gradient. This is broadly consistent with prior literature describing

how the result of associative learning can be heterogeneous across animals or non-stationary in time[149,239,229]. Measurements in response to optogenetically perturbing AWC^{ON} suggest that the source of such context-dependent behavior might be along the AWC^{ON} pathway. Future work is needed to address the source of such variability in the nervous system, as well as its possible functional role for navigation and searching behavior[91,10].

Of the interneurons we investigated, all five classes contribute to both learned navigation strategies (Fig 3.6). Importantly, many connections from AWC^{ON} to the first layer of interneurons are shared with salt sensing neurons ASER/L and thermal sensing neuron AFD[110,111,161]. Some past work hypothesize that experience-dependent changes are localized in specific neurons—- AIB has been shown to play a role in context-dependent thermal navigation[111]. But other work presents evidence that learning effects can be distributed across many neurons[161,198]. For instance, AIB, AIY, and AIZ all seem to be involved in learned salt chemotaxis strategies[161]. Our findings agree with these results in showing that all three interneurons have non-zero and learning-dependent weights for both behavioral strategies.

One puzzling finding is that we sometimes notice changed behavioral strategies but unchanged chemotaxis index. For example, appetitive trained worms with defects in AIB have dramatic reductions in both weathervaning and biased random walk strategies but still show reasonable chemotaxis performance. This may be due to larger model mismatch in mutant worms with strategies deviating from dPAW. Large variability and less consistent behavioral strategies were also observed in earlier ablation studies[110].

In this work we utilize a controlled odor environment and an innovative model to characterize learned odor navigation in worms. The combined approach of precisely delivered sensory stimuli, behavior quantification and navigational modeling continues to be a powerful approach[52,28,67] that can be generalized to other sensory modalities and species to study adaptive sensory navigation. By identifying the specific features of behavior that are altered by learning, our investigation lays the groundwork for followup neural imaging studies and will guide our search for neural representations of learning[184,212].

You have made your way from worm to man, and much within you is still worm.

Frases Nietzsche

4

State-switching Chemotaxis in *C. elegans*

This project focuses on characterizing behavioral dynamics during worm navigation. The study naturally derives from the previous two chapters. After precise measurements of chemotaxis trajectories and model fitting, we began to explore various models and analyze assumptions about worm navigation. While there have been a few decades of work on worm chemotaxis, we revisited assumptions made across various analyses and computational studies. We revealed that worms produce state-switching behavioral strategies to navigate efficiently, rather than a static, parallel combination of strategies. To reach this conclusion, we combined experimental perturbations, novel statistical modeling, and normative frameworks. The project was supervised by Andy Leifer and Jonathan Pillow. This work is in preparation and would be available on arXiv

once it is submitted.

4.1 Introduction

Animals dynamically alter their behavioral response in a sensory environment depending on internal states and behavioral context[7,29,210,25,155]. It is important to characterize state-dependent behavior to better understand behavioral algorithms and the underlying neural computation[211,262,10,123,21,67]. While recent work has discovered state-dependent behavior in fly courtship[29] and rodent sensory decision-making[7,24,109], the role of internal states in goal-directed behavior remains unknown. For instance, it is often unclear how an animal should "choose" between behavioral strategy in order to efficiently navigate a sensory environment[68]. Here we focus on sensory navigation behavior in the nematode worm *C. elegans* and show that—contrary to previous assumptions—its behavior exhibits dependency on discrete states. Moreover, we incorporate neural perturbation and to validate the functional significance of these states.

The *C. elegans* worm is an ideal model organism to study sensory navigation and neural computation because of its well-characterized behavioral dynamics and neural circuits[37,257]. The main navigation strategies include the "pirouette" mechanism that corresponds to the worm making a sharp turn in response to change in sensory cues[192], similar to the biased random walk in bacteria[18,157]; and "weathervaning" mechanism that describes how the worm gradually alters its heading in response to local gradients[110,98], forming a deterministic navigation policy. Past work focuses on the genetic and neural substrates involved in these behavioral strategies. The long-standing assumption has been that the animal uses different strategies in parallel, sampled from fixed frequency for each strategy with no dependency on time or internal states[110,65,111,76]. While the original definition of a "pirouette" is in fact closer to a discrete state with periods of frequent turns (short run intervals)[192], subsequent analyses and computational work turn to describe pirouettes as instantaneous responses without states[65,6,230,210,161,76]. In contrast, there is growing evidence for state-dependency in both behavioral output and neural response to sensory cues[91,155]. The static, state-independent description for sensory processing and employment of strategies

would miss important dynamical features during navigation[241,91,117].

The state-dependent locomotion in worms and many other animals has often been studied in the context of spontaneous behavior without considering sensory signals[14,210,79,233,217,81,27]. These studies characterize how animals explore and adapt to an open environment without specific task structures[79,81,217,210]. The results showed that animals alter their behavioral statistics and transition into different behavioral states, such as local and global search in a longer time scales (order of 10 minutes), and may be related to functions related to foraging in the environment[14,27]. However, the characterization of behavioral states in a shorter time scale (order of seconds) and how it responds to sensory input to support navigation is unclear. A rigorous framework to test for state-dependent behavioral dynamics and incorporate sensory-driven navigation is needed in the field.

In this work, we developed the state-dependent pirouette and weathervane (staPAW) model that describes behavioral strategies that dependent on discrete latent states and fit the model to time-varying trajectories measured from worm navigation. We found that a staPAW model with two alternating states better describes worm's navigation compared to the classical model that does not have state-dependency. The staPAW model flexibly fits to navigation in different sensory environments, such as salt and odor, and provides interpretable parameters describing how the worm evaluates sensory signal to dynamically transition into different behavioral states. With optogenetic perturbation, we recovered similar state-switching strategy with virtual odor signal. In contrast, optogenetically driving interneurons reflect distinct state-transition parameters, indicating that the model fit is non-trivial and may reflect underlying neural computation. Furthermore, we genetically disrupt interneurons to investigate the neural substrates related to state-dependent navigation. Lastly, we found that data-constrained reinforcement learning (RL) model trained on navigation tasks naturally discovers the state-switching solution. These approaches corroborate the description of state-switching strategy in worms that is functionally beneficial for navigation.

4.2 Results

4.2.1 Worms switch between state-dependent navigation strategies

We characterized the worm's navigation behavior by tracking the body centroid location through time. The navigation trajectory in the sensory environment can be characterized with the heading change $d\theta$ and speed dr along the locomotion path. We measured chemotaxis trajectories of worms in either salt or odor gradients (Fig 4.1a) and extract sensory and behavioral time series (Fig 4.1b). The statistics of speed and heading agree with previous reports[192] and show distributions that are multi-modal (Fig 4.1c,d, top). Secondly, the speed autocorrelation shows that there are at least two time scales. Similarly, the turn interval distribution suggests that there is a period of time with more frequent turns (Fig 4.1c,d, bottom). Following analysis in[192], the pirouette is defined as a period of turns that occur within a critical time window, which is computed from a double exponential fit to the turn interval distribution. The correlation in time can be qualitatively observed in exampled navigation trajectories, where speed and turn events seem to co-vary (Fig 4.1e). Lastly, we computed the bearing before and after pirouettes and found that the worm reorients towards the goal direction after the turning bouts (Fig 4.1f). Together, these results indicate that the kinematics are not independent but correlated and persistent in time, and the worm adjusts its navigation path by alternating its kinematics. In the following sections, we seek to develop a rigorous framework to quantify and recapitulate these sensory-driven behavioral dynamics. The approach further makes predictions about the underlying neural substrate and functional benefits in sensory navigation.

Previous work characterized "pirouette" and "weathervaning" strategies along these navigation tracks based on behavioral criteria chosen a priori[192,110]. Inspired by these description of a parallel use of two strategies, we previously constructed a behavioral model with a mixture of strategies known as the dynamical Pirouette and Weathervaning (dPAW) model[45] to flexibly characterize behavior based on the heading statistics and incorporate filters that act on sensory input. We fit the statistical model to observables extracted from the measured navigation (Fig 4.1b). For

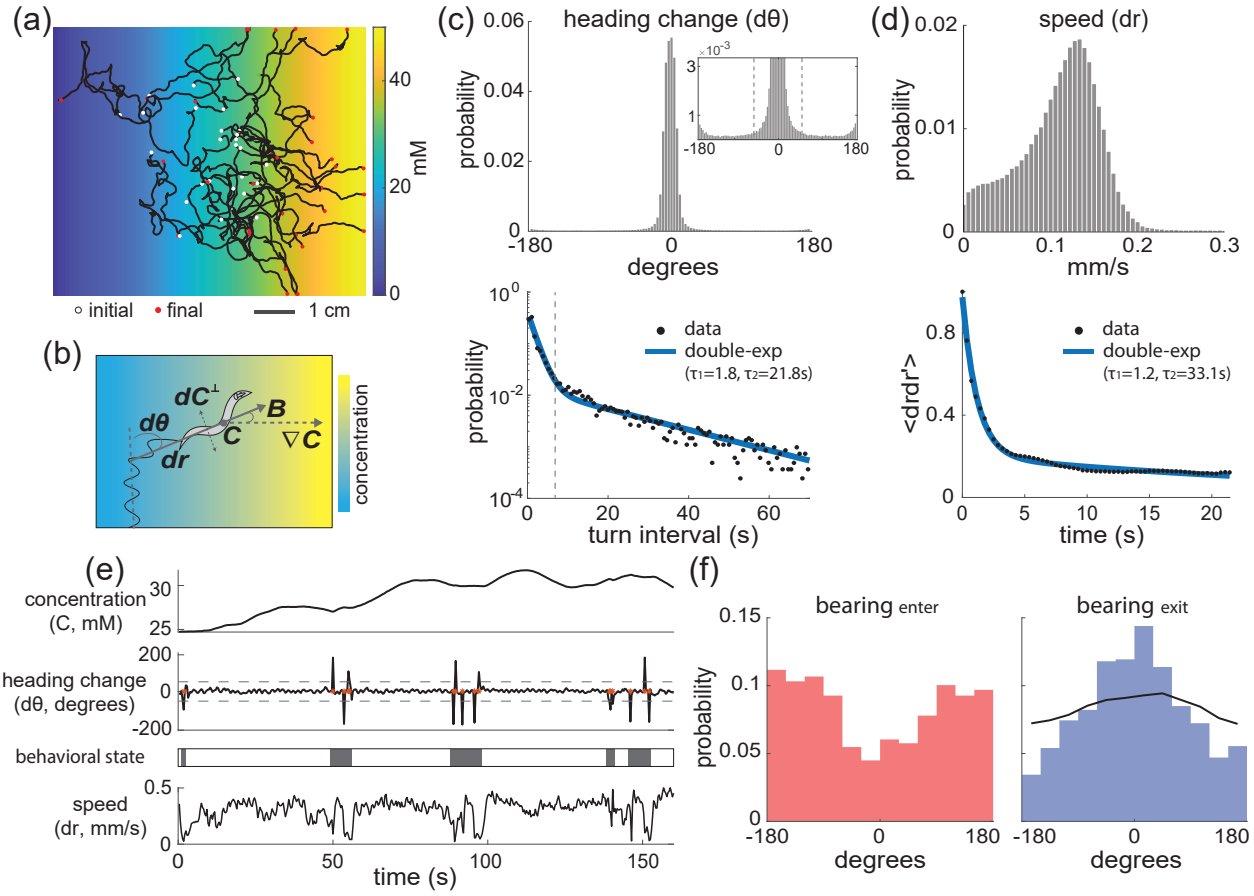

Figure 4.1: **Behavioral dynamics suggest that the worm may be alternating between persistent states. (a)** Example worm chemotaxis trajectories measured in a linear salt gradient. **(b)** Observables are extracted from the chemotaxis experiments. The speed r, heading change $d\theta$, concentration C, perpendicular concentration difference $dC^{\perp}$, and bearing B towards the local gradient ∇C are shown. **(c)** The distribution of heading change $d\theta$ (top) and with inset showing small probability region. The threshold for a turn (± 50 degrees) is indicated with dash lines. The distribution of turn intervals and the double exponential fit in blue are shown (bottom). The critical swim duration ($t_c = 6.8$s) is shown in vertical dash line. **(d)** The distribution of speed dr (top) and its autocorrelation $\langle dr, dr' \rangle$ (bottom). The data points are fitted with a double-exponential curve shown in blue. **(e)** Example navigation trajectory measured from experiment, with turns crossing the threshold (dash lines) shown in red dots. The pirouette state indicated in grey is identified following [192]. **(f)** The bearing before the worm enters a pirouette state (left) and after it exits it (right). The calculation of the exiting bearing with shuffled turning angles shown in black line. Quantification in c, d, and f are for ~ 20 animal hour of navigation trajectories.

dPAW, a binary variable, q, represents the worm's decision to produce a pirouette ($q = 1$) or continue to weathervane ($q = 0$), depending on the filtered sensory input. However, similar to many previous modeling work [65,110,210], dPAW still assumes that the behavioral strategies employed stat-

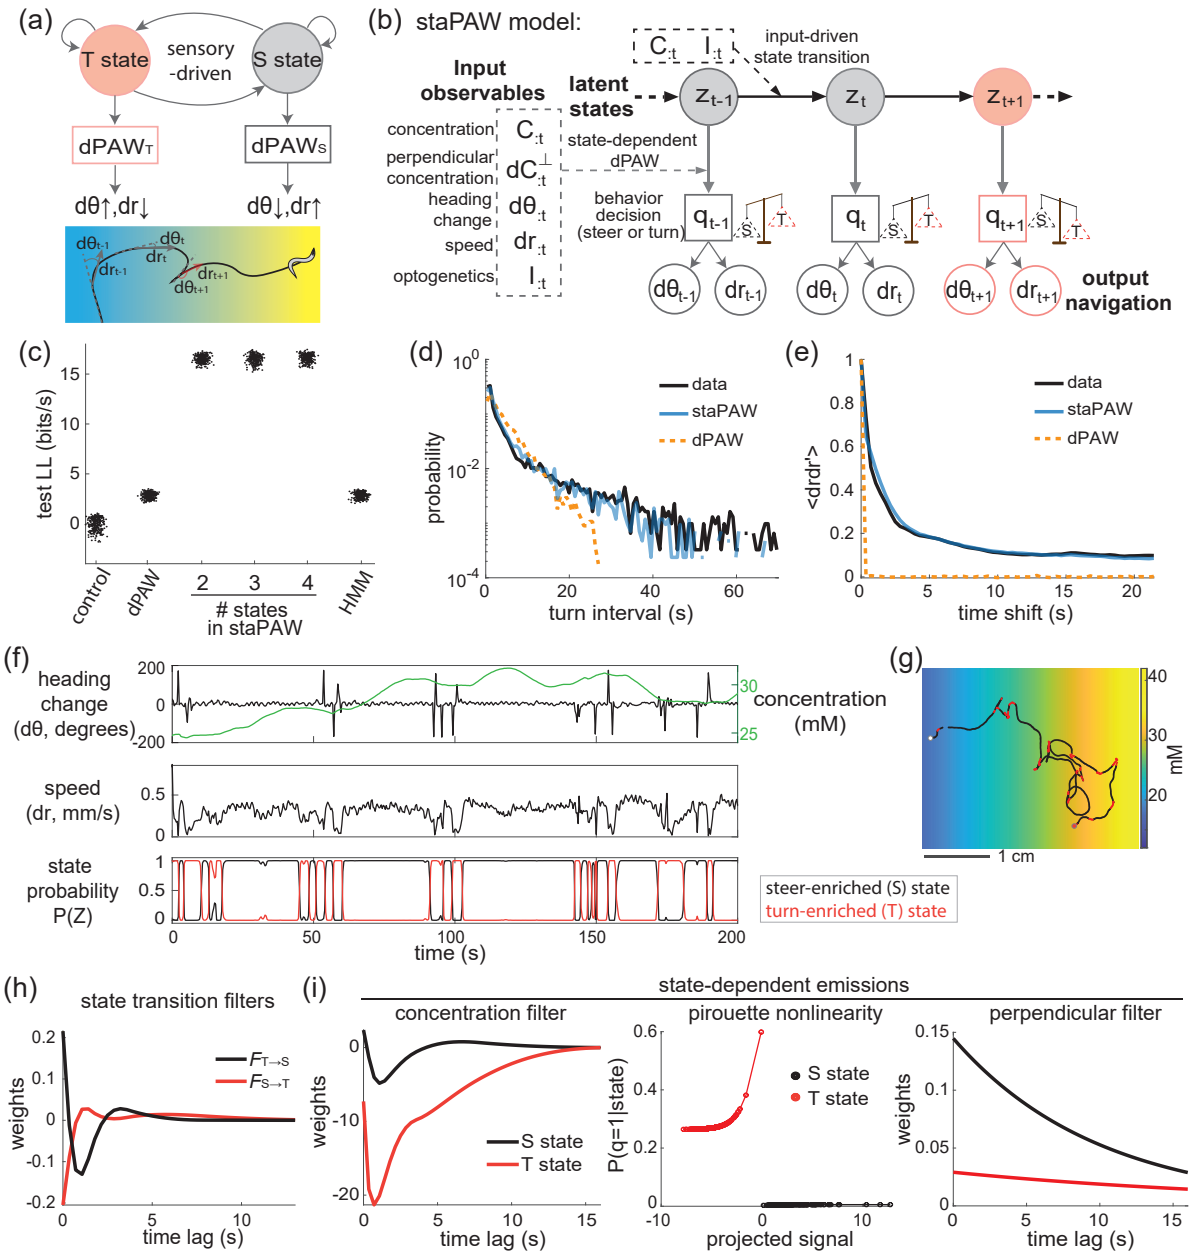

ically across time. More specifically, the pirouette strategy has no persistence and is governed by a Bernoulli process (q) at each time step.

Figure 4.2: **Two states better characterizes navigation behavior than one state. (a)** Schematic of a state-dependent navigation strategy. This is an example with two states, one is the steer-enriched (S) state and another one is the turn-enriched (T) state, each having a dPAW emission for behavioral output. A cartoon time series of worm navigation trajectory color-coded by the behavioral state is shown below. **(b)** Graphical model of staPAW unrolled in time with detailed input-output relation. The input observables are extracted from experimental measurements. The staPAW model assumes a latent discrete state z governing the mapping from input to state-dependent output navigation. **(c)** Test log-likelihood across different statistical models. The difference compared to the mean of control model is plotted on the y-axis. The cross-validation process is 3-fold, each with 5 inference instantiations, and tested with 20 sampled trajectories (with 60 minute total animal time length). **(d)** Turn interval distribution from data and prediction from the staPAW or dPAW models. **(e)** Speed autocorrelation from in data and prediction from the staPAW or dPAW models. **(f)** Example time series of a two-state staPAW fit to data. The posterior probability of being in SE and TE states are shown below. **(g)** Example navigation trajectory plotted in space. Black and red states correspond to the same hidden state in (f). The white circled dot and red circled dot indicate the initial and final location. **(h)** State-transition filters between two latent states. **(i)** State-dependent parameters for emission. Left: concentration filter conditioned on two states. Middle: the nonlinear function that maps projected signal (including history angle change and concentration), conditioned on two states. Right: perpendicular concentration filter conditioned on two states.

To characterize how navigation depends on time-varying states, we extended dPAW with dependence on alternating states (Fig 4.2a). This framework allows different sets of dPAW parameters that persist in time and transition between each other. The novel model is called state-dependent Pirouette and Weathervaning (staPAW, Fig 4.2b). Specifically, the staPAW model is fitted to the extracted time series of heading change ($d\theta$), speed (dr), and sensory experience during worm navigation. The model incorporates discrete latent states, z, that has dynamics following a Markov chain with transition probability driven by sensory input. This resembles a generalized linear model - hidden Markov model (GLM-HMM), where sets of GLM relate sensory input to behavioral output and the HMM introduces state-dependency and transitions between different GLMs. However, the emission of staPAW is not a GLM but an input-dependent mixture model like the dPAW model. Conditioned on each state, the emission includes a binary turn decisions q and samples from behavioral mixtures depending on sensory input. In contrast to past GLM-HMM formalism used to model categorical decisions[29,7,24], the staPAW framework is able to characterize more complex, continuous behavioral kinematics. In addition, staPAW is a generative model that directly produces navigation behavior in the naturalistic context, where the mod-

els behavioral output would feedback and alter the consecutive sensory input in the environment. In the special case of a single state (without dependency on z), the model matches the original dPAW model that includes speed distribution. We developed a null model for dPAW that does not have access to sensory input (control) and another null model for staPAW that can only have Markov transitions independent of sensory input (HMM).

We found that the staPAW model better describes worms' navigation trajectories compared to dPAW model. Through cross-validation, we chose to apply a two-state staPAW to the following analyses (Fig 4.2c). The log-likelihood in held-out test dataset for staPAW is $\sim$ 10 bits/s higher than dPAW model. We used the null model as a baseline to compute model improvement in terms of bit rate. We also compared model fit to a HMM framework that clusters behavioral output but does not relate to sensory input. The HMM outperforms the null model without state-transitions and has a similar fit as dPAW, but has poorer fit compared to staPAW. A direct way to study how staPAW model recapitulates behavioral dynamics is to compute the turn interval distribution and speed autocorrelation. For both metrics, the staPAW model significantly improves the prediction compared to dPAW (Fig4.2d,e). Furthermore, we used the fitted model to predict time-varying behavioral output (see supplement in the paper). While behavioral time series are generated stochastically, we found that staPAW significantly predicts both the turning events and speed correlation in time. Together, this indicates that staPAW captures the time-varying and state-switching sensory navigation strategies in worms.

The fitted staPAW model reveals that worms switch between two behavioral states during navigation (Fig 4.2f,g) — steer-enriched state (S) and turn-enriched state (T). As shown in the time series of salt concentration and behavioral kinematics (Fig 4.2f), the worm intermittently slows down and produces sharp turns as it climbs up the gradient. The turn rate is higher in the T state and results in local searching behavior in space. These results superficially agree with previous report of pirouette-and-weathervaning behavior in worm chemotaxis[192], but note that here staPAW model incorporates the kinematics into a coherent, unsupervised statistical framework with response to sensory cues.

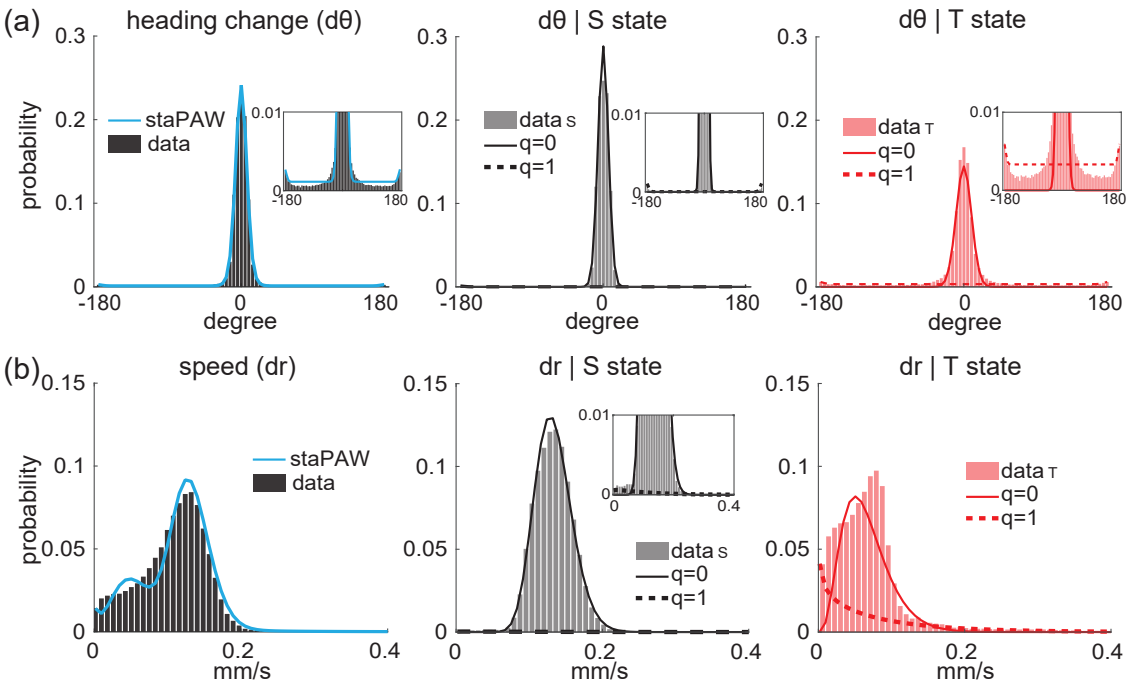

Figure 4.3: **One state is enriched for steering and another state is enriched for turning. (a)** The distribution of heading change ($d\theta$) in data and prediction from staPAW (left). The distribution conditioned on S state (middle) and T state (right) are shown. Within a state, the distribution corresponding to the turn decisions $q = 0$ and $q = 1$ are shown. **(b)** Same representation as in (a) but for the speed (dr) distribution.

4.2.2 Two behavioral states correspond to a turn-enriched and a steer-enriched state

We studied the fitted parameters staPAW to interpret how the transition between latent states are driven by concentration input. We found that the transition filters for two directions (state T to S and S to T) are close to mirror images (flipped-sign) of each other (Fig 4.2h). The downward derivative filter $F_{S\to T}$ is consistent with the biased-random walk description – increase the probability of transitioning into the high-pirouette state when the concentration goes downward, in order to bias the overall movements up the gradient. Interestingly, the mirrored in filter $F_{T\to S}$ indicates that the worm is more likely to exit T state and enter S state when the concentration is rising. This possibly reflects the directed-turn mechanism in worm chemotaxis, where run states are not only biased by down-gradient signals but also promoted when the local gradient is rising [192,117].

The sensory filter that governs pirouette rate differ across behavioral states in staPAW. It is closer

to a derivative form in S state and closer to an integral one in T state (Fig 4.2i, left). The negative shape of these filters reflect that pirouette rate are increased when the concentration is going downwards. The nonlinear function that maps projected signal (including concentration and past behavior) also differs across two behavioral states, with T state mapping to the left-side of the sigmoid curve and S state mostly in the saturated maximum rate regime (Fig 4.2i, middle). This indicates that the worm senses concentration to adjust pirouette rate in T state, but less so in S state. Lastly, the perpendicular filter has large amplitude in S state compared to T state (Fig 4.2i, right), indicating that the worm steers according to perpendicular concentration within the S state. These distinct state-dependent filters reveal how worms navigate differently in the sensory environment depending on its behavioral state.

We investigated the behavioral output in different behavioral states of the fitted staPAW. Consistent with the saturation of test performance in Fig 4.2c, we observed intermediate speed and heading changes in a third state, but the occupancy was variable across data and less interpretable. This result supports the focus on the analyzing two-state staPAW in the following sections. The two-state staPAW model recapitulates the measured distribution of behavioral kinematics (Fig 4.3a,b). We found that the behavioral output significantly vary across two states, with low speed and high turning rate in one state (T state) and high speed and low turning rate in a second state (S state), shown in panels of Fig 4.3a,b. The assignment of T and S states are apparent in the heading change distribution, where the worm rarely produces turns ($q = 1$) and only alters its heading by small changes. In contrast, in the T state the probability of making a turn is much larger, reflecting the broad distribution of heading changes (Fig 4.3a, middle and right). Similarly in speed distribution, we found persistent runs in the S state and slower and variable speed in the T state (Fig 4.3b, middle and right).

We analyzed how temporal patterns in navigation are captured by staPAW. We found that the dwell time of S state is up to a minute and T state mostly within 10 seconds. Note that the dwell-time structure does not follow an exponential form, since the state transitions are driven by non-Markovian sensory input signal. The state occupancy is relatively stationary across time, with

gradual decrease of T states, which may be consistent with adaptation commonly observed in other assays[155,27]. Furthermore, the staPAW model is fitted to data and enables agent-based simulation in the same sensory environment. We simulated navigation trajectories using staPAW with fitted parameters (see supplment in paper). The simulated behavior produces chemotaxis index and marginal behavioral statistics that agree with experimental observations.

4.2.3 Sensory-driven state transitions enables turning towards the goal direction

The mechanism that enables worms to exit turns towards a favorable direction is still less known[192,241]. One possible explanation is that the worm somehow measures the local gradient as it turns and aborts a single turn or reversal so that it end sup in the correct direction. If this were the case, it would suggest that worm nervous system performs a particularly sophisticated estimate of the gradient over very short turning radii or a short duration with 1-2 head swings in time, and places biophysical constraints on its ability to detect concentration changes. If on the other hand, a state switching mechanism would not require measuring the gradient during the single turns but over the persistent state with enrich turns. According to the presented result, the sensory-driven state transition can confer on average the worm's ability to exit the turning state in a favorable direction. An analogy would be a Brownian ratchet, or giving an auditorium of people coins and asking them to flip until they reach heads. Then tabulating the expected value of the coin. A provocative interpretation is that this is a form of "active sensing"[90,10]

To study the contribution of states in the goal of navigation, we analyzed how worms reorient towards the preferred direction when it switches between states. Following analyses in the original pirouette study[192], we computed the bearing after worms exit the T state for staPAW fitted to predict single steps in data and through agent based simulations (Fig4.4a-c). We found that the worms are more likely to orient towards the goal (bearing between -90 and 90 degrees) when it exits the T state. This directed turn can be recapitulated by staPAW fitting and simulation, but missing when there is not state dependency or transition filters (Fig4.4d). Conversely, worms are slightly more likely to enter the T state when it is heading away from the goal, similar to pre-

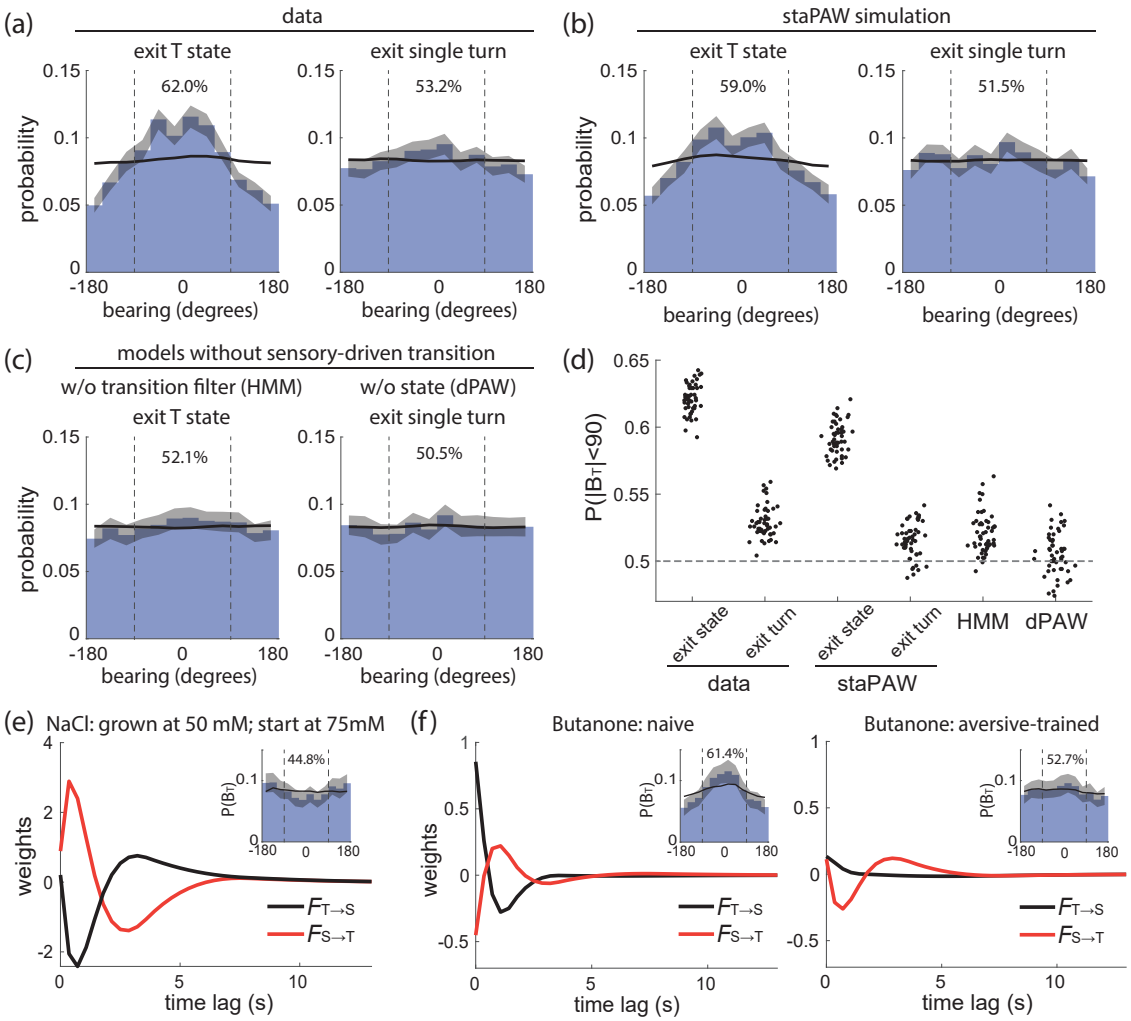

Figure 4.4: **Model state-transition filters govern how worms exit turns towards the goal direction. (a)** Distributions of bearing exiting the T state (left) or a single turn (right) in worm navigation data. **(b)** Same as (a) but for staPAW simulation. **(c)** Bearing distribution from models without sensory-driven states, including HMM and dPAW, **(d)** The probability of aligning to goal, when the bearing exit T state B_T (aligned when $|B_T| < 90$ degrees) for data and different models. The scattered points show 50 repeats samples of 2000 events from the navigation trajectories. **(e)** Worms grown in 50 mM NaCl condition starting at 75 mM along a linear salt gradient produce flipped-sign state-transition filters. The filters have opposite signs compared to worms starting at 25 mM along the linear gradient, as shown in Fig 4.2f. The distribution of B_T and the fraction aligned with gradient are shown in the inset. **(f)** State-transition filters fitted from naive worms (left) and aversive-trained worms (right) navigating in the same butanone environment. 5-9 plates, each with ∼100 worms, are measured per condition.

vious observations[117,192]. Shuffling the different in bearing angle across state transitions does not result in an oriented bearing in the S state, indicating that the correction of heading happens while worms transition between states. These results are consistent with the function of pirouette

in worm chemotaxis, which is to produce directed turns that corrects the animal's course during navigation[192]. Here, the staPAW model not only recapitulate this directed turn property but also provides the underlying computational mechanism with sensory-driven state transitions, which is missing in alternative descriptions without states.

The worm's navigation goal can flexibly depend on the its past experience in specific sensory environments. For instance, salt chemotaxis in worms is governed by the set-point concentration the worm prefers to navigate towards[161,110]. In addition, show that worms can perform associative learning for butanone to alter the attractiveness towards the odor. To test if the state-switching behavior is generic across different navigation goals, we measured salt chemotaxis in a linear gradient with an opposite gradient (100-50 mM) and fitted staPAW to the navigation trajectories. We found similar states as described in Fig 4.2 and focused on the state-transition filters. These filters exhibit opposite results compared to worms in a 0-50 mM linear gradient, which is also reflected in the flipped distribution of bearing after T states (Fig 4.4e). The result indicates that the worms alters its sensory-driven state transitions according to an innate salt concentration preference.

We investigated how worms alter their navigation behavior depending on olfactory learning experiences. Naive worms are attracted by butanone and would climb up the odor gradient, whereas aversive-trained worms down-regulate the tendency perform chemotaxis[50,45]. We applied staPAW to butanone navigation trajectories and found flipped-sign state-transition filters in naive condition, indicating similar gradient-climbing strategies characterized in salt chemotaxis (Fig 4.4f). In comparison, staPAW fitted to aversive-trained worms show delayed $F_{S\to T}$ filter and down-regulated $F_{T\to S}$ filter, which is also reflected in the less aligned bearing after T states. The result indicates that both the biased random walk component and directed-turn mechanisms are down-regulated (less directed) after aversive learning.

Together, the state-switching strategy reflects the navigation goal by tuning the state-transition filters. It appears that state-switching behavior is not specific to climbing salt gradients, but reflects navigation that depend on internal "preference" and reveals similar mechanisms in odor

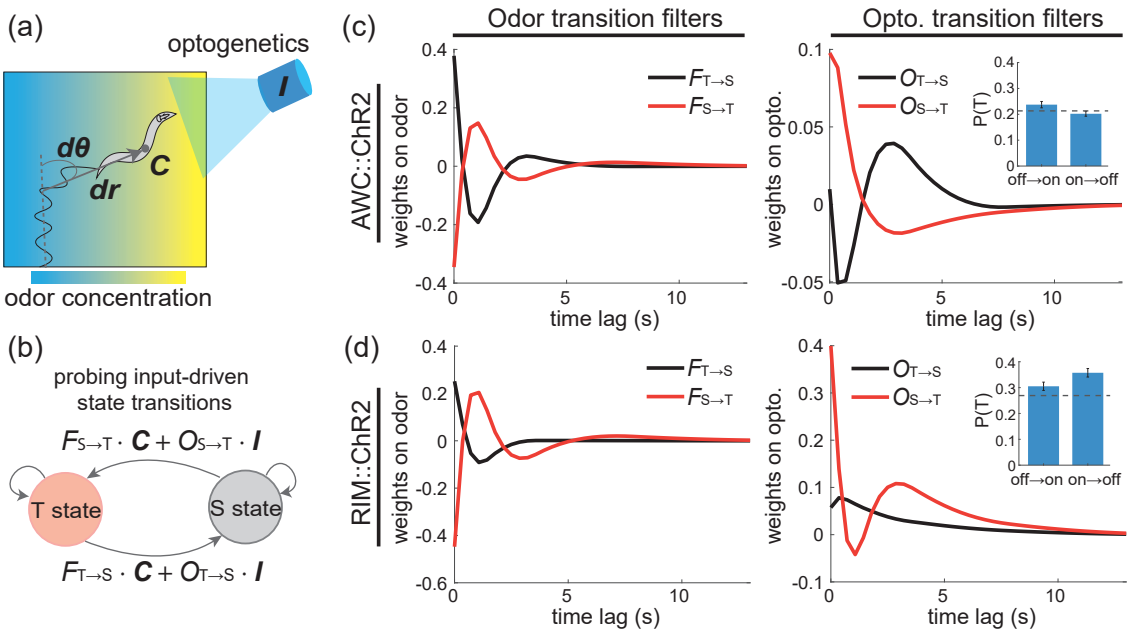

Figure 4.5: **Optogenetic stimulation shows that sensory processing causally drives state transitions. (a)** Experiment with concurrent presentation of odor C and optogenetic input I. **(b)** Schematic for the staPAW model extended to incorporate optogenetic input that drive state transition probability. **(c)** State transition filters for worms expressing ChR2 in AWC sensory neuron. The butanone odor filters $F_{i\to j}$ and optogenetic filters $O_{i\to j}$ are shown. In the optogenetic case, inset shows the probability of being in the T state for the 5s window when light is on versus a 5s window after light is off. Dash line shows the baseline probably before stimuli and error bars show standard deviation across 50 ensembles of 500 stimuli. T-test across all pairs show significant difference ($p < 0.001$). **(d)** Same as (c) but for worms expressing ChR2 in the interneuron RIM. The staPAW model that includes optogenetic input is fitted to data with over 3,000 pulses delivered during worm chemotaxis. In the inset, t-test across all pairs also show significant difference ($p < 0.001$).

chemotaxis under different learning experiences.

4.2.4 State transition filters reflect sensory processing during navigation

To pinpoint study sensory processing and neural circuits involved in state-switching navigation, we investigated introduced optogenetic perturbation and genetic ablation. First, we expressed ChR2 in a sensory neuron AWC that is known to respond to volatile odorants and optogenetically stimulate it during the worm's navigation[50,45,46] (Fig 4.5a). We extended staPAW to include filters that act on optogenetic stimuli in the state-transition probability (Fig 4.5b). The resulting optogenetic filters have flipped signs that match the odor filters (Fig 4.5c). Given the knowledge

of AWC being an OFF cell[50,45], which means that the cell is depolarized when odor is off, we could justify the sign flip in optogenetic filters comparing to odor kernels. This validates that the fitted sensor filters for odor input reflect a similar type of sensory processing as optogenetically induced odor perception.

We applied the same stimulation protocol and model fitting to worms that express ChR2 in an interneurons RIM that is known to drive reversal and turning behavior[257,93]. With the fitted staPAW model, we observed that AWC-ChR2 worms are driven transiently into T state and out in the time window of 5-10 seconds (see supplement in paper). In contrast, RIM-ChR2 worms have slow and prolonged T state after the optogenetic impulse. The fitted optogenetic filter for state transitions from running to pirouette in RIM-ChR2 worms shows a similar shape to those fitted for AWC-ChR2 worms. However, we found that the flipped-sign filter exiting T state and entering S state is no longer preserved (Fig 4.5d). When the optogenetic stimuli is off, in contrast to suppressing the probability of being in T state as observed in the AWC strain, we found that the occupancy increases in the RIM strain (Fig 4.5d, inset). These results indicate that the state-transition filter is not simply a byproduct of worms reversing and turning in the sensory environment. The filters reflect the worm's detailed sensory processing that casually drives the transition between behavioral states.

4.2.5 Interneurons contributes to state-switching behavior

We investigated the contribution of interneurons during the worm's odor navigation by disrupting either neurons AIB, AIY, or AIZ. According to the connectome, these three interneurons are common to both AWC (odor) and ASE (salt) neural circuits and are involved in chemotaxis behavior[37,110,161]. Past work demonstrate that AIB is involved in reversal behavior and AIY is involved in forward movements during navigation[37,93]. The ablation of AIZ ablation significantly down-regulates the overall chemotaxis behavior[110,45]. We applied staPAW model to the transgenic worms and found state-transition structures that are significantly different from wild-type worms (Fig 4.6). For AIB-disrupted worms, the S state is promoted and the T state has less

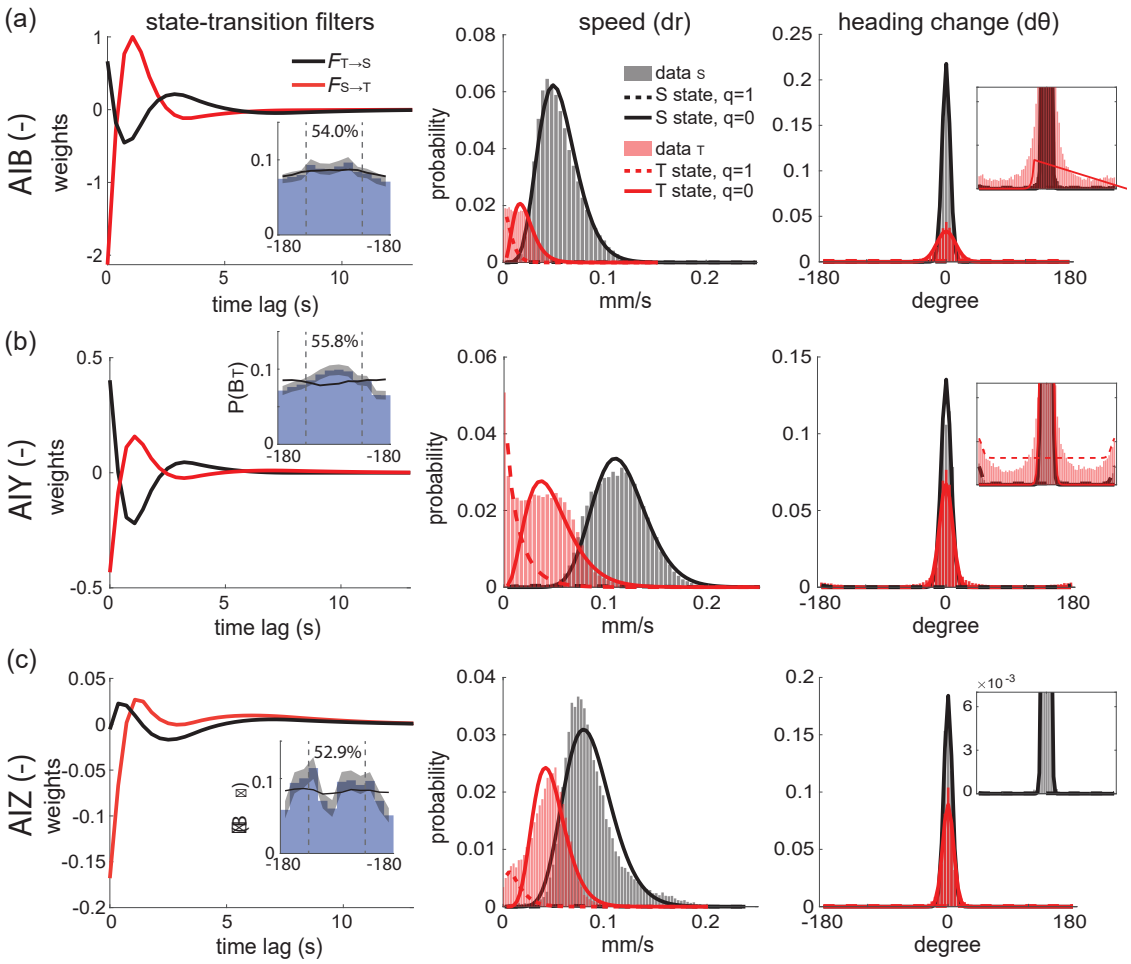

Figure 4.6: **Interneurons have different relative contributions to state occupancy and the sensory-driven state transitions. (a)** State transition filters fitted to worms with disrupted AIB navigating in a butanone gradient (left). The distribution of bearing exiting T state ($P(B_T)$) and percentage of angles aligned to the goal are shown in the inset. The corresponding distribution for speed (middle) and heading change (right) of two states are shown. The enlarged probability for broad heading change is in the inset. **(b)** Same as (a) but for worms with disrupted AIY neuron. **(c)** Same as (a) but for worms with disrupted AIZ neuron. Each transgenic condition includes 5-10 animal hours of navigation trajectories for staPAW fitting.

sharp turns, but the sensor filter for state transitions maintain the same shape as wild type. For AIY-disrupted worms, the S state is down regulated and the T state has enriched sharp turns, and the sensor filter for state transitions also maintain the same shape as wild type. For AIZ-disrupted worms, the S state has slower speed and the T state has little sharp turns and less obvious separation. Importantly, we found that the state transition filters have disrupted shape in AIZ-disrupted worms (Fig 4.6c). Given the fitted staPAW model for each strain, we could sys-

tematically analyze how the bias in behavioral state transitions reflects the overall chemotaxis performances.Together, these analyses showcase how staPAW can flexibly fit to significantly different behavioral kinematics across worm strains. Furthermore, the results indicate that the first layer of interneurons downstream from sensory neurons are readily involved in sensorimotor computation related to the behavioral emissions (AIB and AIY) or the behavioral state transition (AIZ).

4.2.6 Data-constrained reinforcement learning model reveals the functional benefit of the state-switching navigation strategy

Analysis across worm strains showed that the fitted staPAW model predicts a drift velocity going up gradient, which reflects the chemotaxis performance. To quantitatively understand how sensory-driven state transitions benefit chemotaxis performance, we explored the normative aspect of navigation strategies through a data-constrained reinforcement learning (RL) model. For simplicity, we focus on the heading change and compute step-wise concentration differences rather than temporal filters. Importantly, the RL agent can only take actions that are constrained by emissions in the two behavioral states (T and S states). The parameters optimized through RL are weights on concentration difference and a baseline probability that governs state transitions. We first confirmed that the RL model constrained by worm behavior can learn to perform chemotaxis and generate qualitatively similar state-transition behaviors (Fig 4.7a,b,c). Compared to uncontrolled agents (those with shuffled states), the marginal distribution of heading change and occupancy are similar. However, after training, the persistence in inter-turn intervals and the directed turn property emerges through RL iterations (see supplement in paper). We also compared chemotaxis performance with pirouette-only or weathervaning-only strategies. We found that the optimized state-dependent navigation (staPAW) is consistently beneficial across the trained environment and a new environment with shallower gradients. Switching between two strategies is particularly important to generalize across different environments with different sensory statistics (Fig 4.7d). Individual navigation strategies also under-perform the parallel use of both strategies

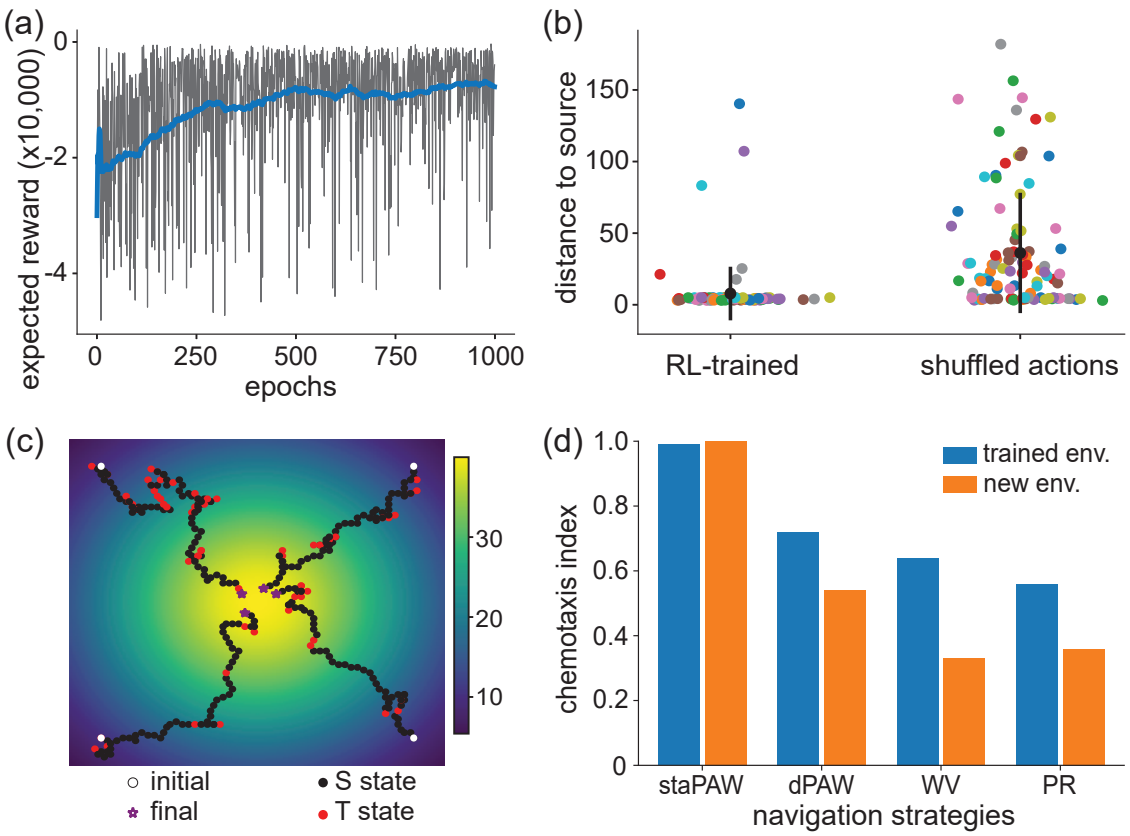

Figure 4.7: **Data-constrained reinforcement learning model suggests that state-switching is beneficial for navigation performance. (a)** Expected reward as a function of epochs during reinforcement learning. The blue line is a running average of every 50 epochs. **(b)** Distance to the point source from 100 navigation trajectories in the RL-trained model and a control comparison with shuffled action states. The mean and standard deviation are shown in black. **(c)** Four example navigation tracks with the RL trained parameters. Two states are colored in black and red. The initial position is in white and the final position is in purple stars. **(d)** Chemotaxis index (CI) across different navigation strategies. CI $= \frac{N_{\text{up}} - N_{\text{down}}}{N_{\text{up}} + N_{\text{down}}}$, which computes the normalized fraction of tracks going up gradient. Results in the trained environment with a Gaussian spatial profile and a new environment with wider Gaussian width (1.7 times) are shown.

(dPAW) and are sensitive to different sensory statistics. The data-constrained RL framework recapitulates the experimental observations of worm navigation strategies and provides a normative explanation for transitions between behavioral states.

The staPAW model and data-constrained RL model are both phenomenological description of the worm's navigation behavior. We further explored a biologically plausible neural circuit model that captures state-switching navigation strategies (see supplement in paper). Our effective network model constitutes two mutually inhibiting rate neurons that both receive sensory input. The

two units correspond to two different navigation states, and the model produces a unit-dependent behavioral strategy according to the unit with larger activity. In result, the circuit model recapitulated similar navigation time series and dwell time of behavioral states. Furthermore, we compared models with and without the mutual inhibiting motif that generates state dependency. The result is consistent with observation in (Fig 4.7d), where models with state-switching has higher chemotaxis compared to models without states. This neural circuit provides a biologically plausible mechanism for state-switching navigation behavior that can be further validated with genetic perturbation and neural imaging in the future.

4.3 Discussion

We investigated how worms produce persistent behavior and directed turns during sensory navigation. To make quantitative predictions about the worm's navigation strategy, we developed a novel state-space model, staPAW, to characterize state-dependent sensorimotor transformation along the navigation path. We found that staPAW better describes experimental measurements compared to models without state-dependency. The fitted staPAW model shows one state that elicits high turning probability (T state) and another state that produces persistent runs (S state). Trough optogenetics perturbation, we verified that the worm's behavior state is causally affected by sensory input. The decision to transition between states depending on local concentration change enables worms to turn towards the goal direction. Furthermore, transgenic studies show that the interneurons have relative contributions to the state transition structure. Lastly, a data-constrained computational model suggests that state-switching strategy can be beneficial for the worm's navigation.

The two states revealed through staPAW model are similar to pirouettes and runs that are reported in the original studies[192]. In fact, the original definition of the pirouette behavior was similar to a discrete state that persists across seconds. While capturing important mechanisms for chemotaxis in worms, the identification of pirouette requires picking a threshold to define turns, measuring inter-turn intervals, then fitting double-exponential curves to the histogram. In each of these

steps, the parameter choice and curve fitting may introduce biases and vary across recordings for different strains. More importantly, the assumptions to define pirouette as a behavioral state was not made explicitly nor compared to other candidate descriptions. Later work conduct analysis or simulations of pirouettes in chemotaxis based on methods ranging from different cutoffs for inter-turn intervals[110,241], a fixed time window to draw sharp turns[230,161], to an instantaneous turning response[65,6,210,76]. More recent work fit mixture models to behavioral kinematics[267] or construct embedding for the worm's posture[57] to more rigorously define pirouette states. However, neither of these frameworks incorporated the time-varying sensory input during navigation. Here, the staPAW model mitigates these gaps by providing a framework to define behavioral states in an unsupervised manner, and characterizes how worms transition between states given sensory input to navigate.

Worms produce behavior across a hierarchy of time scales, from turning and running in seconds to dwelling and roaming in minutes[79,129]. The hierarchical structure of different spontaneous behavioral states have been described in worms and many other animals. However, the functional role of state-switching behavior is rarely studied. Here we utilized the goal-oriented nature of chemotaxis and find that worms produce state-switching navigation strategies during the task. The state-dependency is generalized across different environments and is a better normative solution for worms to climb up chemical gradients. The state-switching navigation strategy captured by staPAW is different from the stochastic models proposed for foraging behavior in worms. Past models for foraging behavior only describe states with different kinematics[79,14], but staPAW explicitly models the sensorimotor transformation during sensory navigation. While we fit staPAW model to navigation trajectories measured during chemotaxis assays in this work, future applications can extend to longer time coarse with different food sources to capture behavioral state in even longer time scales.

The dimensionality of animal behavior is an important aspect to understand behavioral tasks and the underlying neural mechanisms[234,21]. Recently, a number of approaches have been introduced to characterize lower-dimensional structure and dynamics in spontaneous behavior, applying

methods from information theory[21], delay-embedding[2,57], state-space models[153,25], and more complex hierarchical models[205]. The staPAW model is a specific extension of state-space model framework, where the input-output relation within a latent state describes sensorimotor transformation of the worm. An alternative approach to analysis behavioral dimensions is to measure many features and cluster or project down to effective subspace[57]. This non-parametric approach may facilitate discoveries of new behavioral strategies and can be incorporated in future studies. Despite being constrained by more specific parameterization in staPAW, we showed that this framework offers interpretable parameters that can reflect the underlying neural computations. Recent work demonstrate how state-space models enable the characterization of spontaneous behavior structure[262,67], dynamics during perceptual decision making[7], and mating song production[29]. However, the "functional role" of these state-switching behavior are often elusive. For instance, it is unclear how internal states during perceptual decisions in rodents correspond to an exploration strategy in naturalistic tasks[7]. Neither is obvious how internal states during male fly's song production reflect the success rate for mating[29]. The normative framework provides a complimentary approach to evaluate how behavior can be beneficial to achieving certain tasks. For instance, past work numerically train agents to transition between constrained states and perform naturalistic tasks via reinforcement learning (RL)[266]. After training agents to navigate in the sensory environment, certain known strategies and neural computation emerge from the optimal solution[134,208]. Building on these approaches, our work provides connection between the staPAW model that describes statistics of worm navigation behavior and a data-constrained RL model that evaluates navigation performances.

In this work we explore interneuron's contribution to state-switching navigation behavior through optogenetic perturbation and genetically disrupting interneurons. Another approach is to directly measure neural response during naturalistic behaviors in worms. For instance, recent work also show how collateral discharge signal can lead to persistent behavioral states during thermal navigation[123]. Past work has demonstrated lower dimensional state-space neural dynamics in whole-brain measurements and may potentially lead to different behavioral modes[153,25]. Future work is

required for more comprehensive understanding for the neural dynamics underlying behavioral state transitions during navigation[8,184].

As long as our brain is a mystery, the universe, the reflection of the structure of the brain will also be a mystery.

Santiago Ramón y Cajal

5

Learning Algorithms and State-dependent Dynamics in Stochastic Neural Networks

In this chapter, we investigated learning algorithms and dynamics in recurrent neural networks (RNNs) with stochastic properties. The investigation ranges from neural circuits optimized to recapitulate experimentally measured behavioral response in *C. elegan*, stochastic RNN trained to characterize state-dependent phenomena observed in perceptual decision-making, to stochastic spiking dynamics trained to produce complex temporal patterns. While capturing seemingly different target behavior, the models share the common component of being intrinsically stochastic. The stochastic neural response more closely mirrors the biophysical signals measured in the

brain, which exhibit variability at the individual neuron level but still function effectively on the circuit and network levels. Past work mostly focus on modelling the mode of neural dynamics, such as population or trial averages[237,70,202]. More recent work start to reveal the important role of variability and noise in neural computation relevant for naturalistic behavior[171,91,257]. In this chapter, we presented work along this thread, demonstrating how stochastic neural models better describe data and reveal computational principles. In the Leifer lab, we discussed the touch response modeling work with Sandeep Kumar. We collaborated with Emily Osborne on the whole-brain imaging and odor response project. For the development algorithms to train stochastic network, we discussed with Brian DePasquale and the project was advised by Jonathan Pillow. Although the projects in this chapter have not yet been published, they will be continued and expanded upon in future research.

5.1 Stochastic circuit models for behavioral response in *C. elegans*

5.1.1 **Context-dependent sensorimotor transformation**

The touch response in worms has been studied in detail for decades, with particular focus on the molecular machinery underlying signal transduction[39,38]. Research in the mechanical properties of cells and proteins in the worm facilitates our understanding for the detailed biophysics[38]. However, the circuit-level understanding and dynamical picture of how worms respond to touch were less explored. Recent studies in the Leifer lab reveal how worms process temporal patterns of touch stimuli to drive behavior[156,155]. The series of work further discover context-dependent sensory processing: the worm is more likely to reverse in response to touch during runs, but gates the response during turns.

To gain insight into the neural circuits that generate context-dependent responses, we turned to an effective circuit model for worm behavior (Fig 5.1a). We simplified behavior-relevant neurons into single units that are interconnected. The neural dynamics follow:

$$\tau \frac{dv_i}{dt} = -v_i + \Sigma_j J^s_{ij} \phi(v_j) + \Sigma_j J^g_{ij}(v_i - v_j) + I_i + \xi_i \tag{5.1}$$

where the voltage of the ith neuron v_i has time scale τ, receives input I_i, and uncorrelated noise ξ_i. The 3-neuron circuit is connected through chemical synapse J^s_{ij} weighting the nonlinear input $\phi(v_j)$ and gap junctions J^g_{ij}. As justified in previous work[76,157], we did not remove diagonal terms in the weight matrix as it accounts for poly-synaptic connection and possible voltage-dependent current in worm neurons. We used a saturating nonlinearity $\phi(\cdot) = tanh(\cdot)$ to account for stable and state-dependent activity. The three neurons correspond to forward, reversal, and turn behavior. The behavioral readout of this network is the behavior assigned to the unit that has the largest activity at a given time, $B(t) = \text{argmax}_i \mathbf{v}(t)$. We tuned the connectivity J^s_{ij}, J^g_{ij}, and the noise strength D from $\langle \xi(t)\xi(t') \rangle = 2D\delta(t - t')$ to recapitulate context-dependent sensory processing.

Inspired by previous work modeling behavioral sequences in worms[257], we parameterized the input vector I to only drive the reversal neuron. Intuitively, the turning neuron has an inhibitory connection to the reversal neuron to gate the activity and create context dependency. The strong noise input into this circuit is a crucial component to create spontaneous transitions and probabilistic responses. Furthermore, we incorporated short-term depression (STD) in the model to capture depression induced by pre-synaptic activity:

$$\tau_s \frac{dJ^s_{ij}}{dt} = -(J^s_{ij} - J_0) - \gamma[\phi(v_j)]_+ \tag{5.2}$$

where J_0 is the steady-state synaptic connection that the weights relax to when there is little activity. When the pre-synaptic neuron is activated, the connection is weakened by the scalar γ weighting the rectified nonlinear output. Particularly, this STD dynamic is discovered and modelled for the inhibitory projection from the reversal to turn neuron. The motivation for STD is to capture dynamics transitioning from reversal to turn, which is stereotypical in natural worm behavior.

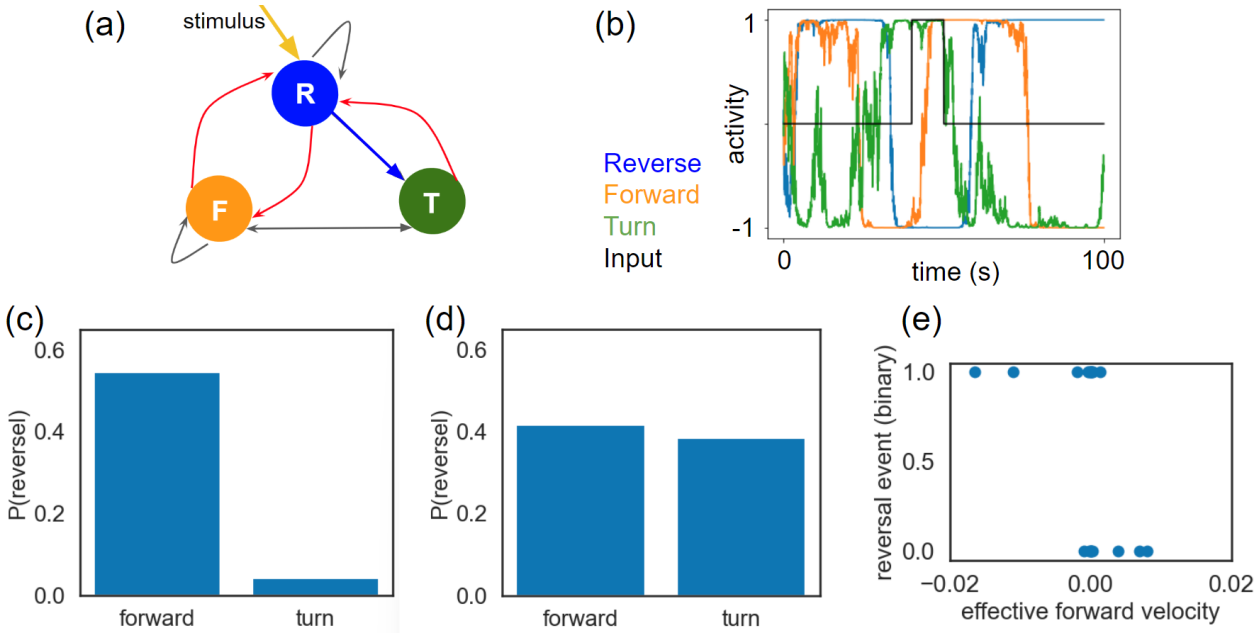

Figure 5.1: **Circuit model for context-dependent touch response**. (a) Schematic of the circuit model with three neurons. The touch stimulus (yellow arrow) is applied to the reversal neuron (R), which has connection to the turn neuron (T) with STD synapse (blue arrow). The static inhibitory connections are in red and excitation are in gray. The forward neuron (F) is also interconnected. (b) Example neural traces simulated from the circuit in response to touch input. (c) The reversal probability depends on the context of this circuit. (d) When the T neuron is inhibited, the touch response has little context dependency. Analysis of (c) and (d) are computed from 200 instantiations. (e) The effective forward velocity measured from the F neuron is negatively related to the binary reversal event.

In result, we found that the activity of three neurons spontaneously transition between different behavioral state when there is no external input (Fig 5.1b). When touch stimuli is presented, the reversal neuron increases in activity and the forward neuron decreases, as expected in the common behavioral response to touch. Importantly, in the context of turning, when the turning neuron happens to have high neural activity upon the stimulus presentation, the model recapitulates the lower probability of entering the reversal state (Fig 5.1c). This demonstrates that a simple circuit with feedback inhibition and STD is sufficient to capture the context-dependent response in worms.

We explored experimental predictions that are generated by the model, including speed tuning and circuit perturbations. In addition to the discrete description of behavioral states, worms have continuously varying forward velocity. We hypothesized that the activity level of the forward

neuron can be an approximation of forward speed, and study how the activity level relates to the triggered reversal event. The result shows that the reversal behavior is negatively related to forward velocity. Worms that have lower speed are more likely to produce a reversal in response to touch (Fig 5.1d). This is qualitatively consistent with experimental observations that are not published in the Leifer lab. In addition, it has been shown in experiments that down-regulation of a population of inhibitory neurons deprive the context dependency[141]. We mimicked the experiments by down-regulating the turning neurons right before the touch stimulation in the model (Fig 5.1e). The simulation consistently shows that the context-dependent behavior is removed – worms have similar reversal probability upon touch stimuli regardless of the forward or turning state.

The circuit modeling approach is closer to a phenomenological one than a detailed biophysical one. Worms use a population of neurons to coordinate behavior, rather than simple units. The single neuron biophysics and synaptic dynamics is also more complicated than the rate neural model presented here[86]. Follow up studies with more details for single neurons and molecular components involved in the process may provided better description. However, a number of past work suggest that single neural response in worms can be approximated with rate models to a large extend[76,157,210,116]. More importantly, on the systems and circuit level, our simple neural circuit model can readily capture experimental observations, make prediction about speed tuning and neural manipulation, and provide intuition for small circuits with gating properties.

5.1.2 **Head-tail sensory integration**

Worms increase forward velocity in response to touch stimuli in the tail and increase reversal probability in response to touch in the head. Such escape response is ethologically relevant to run away from potential threat in the environment. Recent development with genetic manipulation and closed-loop stimulation in the Leifer lab recapitulate the head and tail touch response through high-throughput optogenetics[155]. The worms sprint forward when light activates mechano-sensory neurons in the tail (PVM, PLML/R) and reverses with light activates mechano-sensory

neurons in the head (AVM, ALML/R). Furthermore, worms are biased towards reversal when both head and tail neurons (all six mechano-sensory neurons) are activated. We sought to capture this asymmetric phenomena with a simplified neural circuit model and make predictions about future neural measurements.

The model has a similar form for neural dynamics as the 3-neuron circuit mentioned in the previous section. We formulated a 2-neuron circuit with mutual inhibition (Fig 5.2a). One neuron receives head stimuli and has activity indicating reversal behavior, and another neuron receives tail stimuli and has activity indicating forward behavior. The dynamics follows:

$$\tau \frac{dv_i}{dt} = -v_i + \Sigma_j J_{ij} g(v_j + b_i) + I_i + \xi_i \tag{5.3}$$

with connection J_{ij}, nonlinearity g being a sigmoid $g(\cdot) = \frac{1}{1+\exp(\cdot)}$, input I, bias term b_i, and noise ξ. The time scale of neuron τ is much shorter than behavior, so the behavioral readout B is smoothed from the maximum neural activity v following:

$$\tau_b \frac{dB_i}{dt} = -B_i + \delta(n(t) - i) \tag{5.4}$$

with longer behavioral time scale τ_b and $n(t) = \text{argmax}_i \mathbf{v}$ is the index of the neuron with larger activity. We measured spontaneous and driven behavior across different head-tail stimuli.

In result, we found that the spontaneous behavior produced from this simple circuit has bi-stable transitions between reversal and forward modes (Fig 5.2b). As expected from the circuit design, the model produces reversal when the head neuron is activated, and produces forward when the tail neuron is activated. The modelled mutual inhibition connection are asymmetric, with stronger inhibition from head to tail neuron. This results in the dominance of reversal when both the head and tail are activated. This forms a hypothesis for head-tail sensory integration through inhibitory competition.

The model makes at least two testable predictions: (1) different input strength to the head or tail produces alternative outcome when they are both activated, and (2) the mutual inhibition of for-

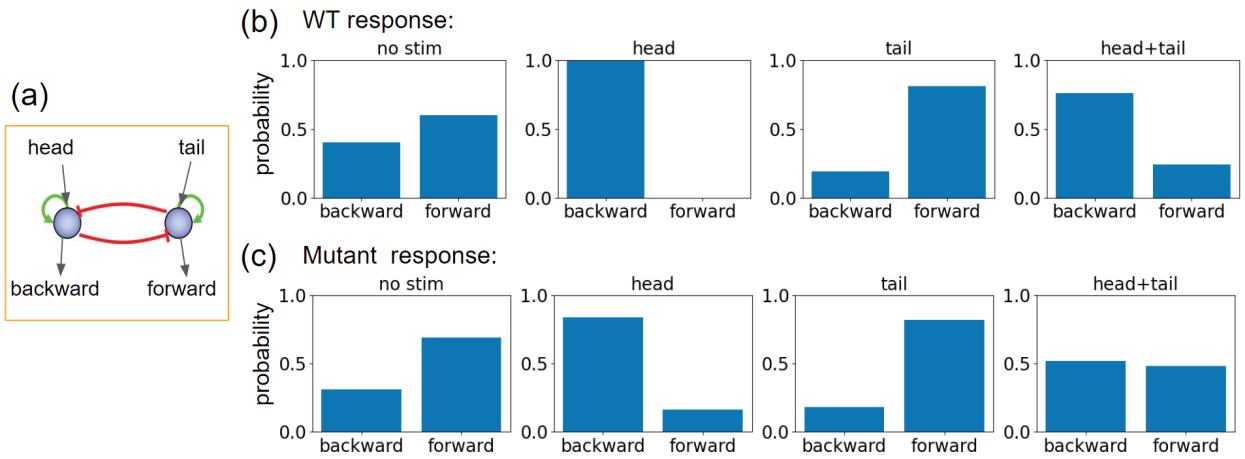

Figure 5.2: **Circuit model for head-tail touch responses**. (a) Schematic of the mutually inhibiting head and tail units. (b) The probability of backward or forward behavior at spontaneous level (no stimuli), upon head stimulation or tail stimulation alone, and when both head and tail are stimulated. This matches the result of wild-type (WT) worms. (c) Same measurements as in (b) but for a transgenic worm that only has input from AVM and PLML/R sensory neurons.

ward and reversal neuron control both the spontaneous transitions and the head-trail integration. The first prediction has already been validated by experiments in the Leifer lab with worm strains that only express opsins in three mechanical sensing neurons, AVM and PLML/R, which leads to less input to the head. The result shows that while spontaneous, head-only, and tail-only stimulus response are relatively intact, there is significantly less reversal when both head and tail stimuli are presented (Fig 5.2c). The second prediction is not tested directly yet. Recent work showing push-pull-like circuits that govern forward and reversal provide candidate interneurons to be measured in the future to understand the integration and decision of mechano-sensory responses[173].

5.1.3 Generating state-dependent behavior in stochastic networks

Naturalistic animal behavior has complex structures, including hierarchical organization, stochasticity, and context dependency[171]. As a model organism, the worm has relatively tractable behavioral dynamics, and its behavior readily demonstrates features across different time scales and probabilistic structure[91,129,2,233]. As described in the context of foraging and chemotaxis, the statistical structure in spontaneous behavior, such as the search area and efficiency, may pro-

vide functional relevance to naturalistic tasks[27,14]. Studies of the structure in the worm's natural behavioral test for normative theories and provide insight into the underlying neural computation[79,130,67].

As described in the previous section, we demonstrated hand-crafted phenomenological models that superficially represent neural dynamics and also account for the worms state-dependent, probabilistic behavior. The models often assume bistable units[235], which are partially supported by sparse neural recordings during chemotaxis[8,91]. These measurements suggest that the neural dynamics in the worm brain can be the internal source of stochastic neural states that alter the sensory processing[91,153]. For instance, while the sensory neuron reliably responds to olfactory input, the information is only transformed into reversal behavioral output when the interneuron is in a certain activity state upon the sensory input[91,150]. We sought to characterize such probabilistic, state-dependent neural dynamics in a generalized framework without hand tuning parameters. Despite success in modern machine learning with large artificial networks, empirically, small models are more difficult to train and the relevant learning algorithms for small systems have been less developed. In fact, smaller biological circuits seem to readily produce complex dynamics and computation. On the theoretical end, it has been shown that small neural circuits can produce rich nonlinear dynamics[15,165]. To optimize parameters for small networks, past work applied brute-force parameter scan[76] or genetic algorithms[115] to optimize circuits for chemotaxis in worms. Inspired by these approaches, we explored an alternative efficient optimization method – the particle swarm optimization (PSO) method with small recurrent neural networks[165,133]. The PSO method is a class of algorithm inspired by collective behavior of swarms and is generically used in engineering optimization problems. The PSO method does not require the objective to be differentiable and has a minimum set of hyper-parameters required for training. While there is no convergence guarantee, the empirical performance and application to "medium" size models such as a small neural network makes it appealing.

We applied PSO to solve for parameters in a 3-neuron circuit with dynamics following:

$$\tau_i \frac{dv_i}{dt} = -v_i + \Sigma_j W_{ij} \phi(v_i + b_i) + I_i(t) + S_i(t) \tag{5.5}$$

where the voltage of the ith neuron v_i is governed by network W_{ij}, nonlinear function ϕ is a sigmoid, offset b_i, time constant τ_i, external input S, and an internal correlated noise signal I_i following:

$$\frac{dI_i}{dt} = \frac{I_i}{\tau_{I_i}} + \xi_i \tag{5.6}$$

which is an Ornstein–Uhlenbeck (OU) process with time scale τ_{I_i} and receiving Gaussian noise $\langle \xi \xi' \rangle = 2D_i \delta(t - t')$. Note that the noise strength and correlation are also private to a single neuron. The behavioral readout from this circuit is simplified as a threshold detection problem. Whenever the first neuron's activity, $\phi(v_i)$, exceeds a threshold θ, then the behavior is assigned into a transitioned state. We sought to fit this abstract representation to the behavioral transitions of forward and reversal states in the worms. The free parameters in the model includes $\{W_{ij}, \tau_i, b_i, D_i, \tau_{I_i}, \theta\}$. The objective is to alter these parameters in order maximize the "fitness" of the swarm of particles. This fitness function is defined by the difference between the simulated and measured behavioral statistics, including the distribution of state occupancy time and the Markov transition probability.

In practice, the PSO algorithm takes the circuit model and an initial guess of the full parameter vector and conducts numerical simulations at each iteration. The particles are sampled from a range of parameters and each particle produces a candidate behavioral output that is then compared to experimental measurements to evaluate fitness. In the parameter space, parameters have position x and velocity v. In each simulation, there would be information about the best parameter for a particle p and for the entire swarm g. For every particle i in every parameter dimension, the motion in parameter space follows:

$$v_i \leftarrow w v_i + \phi_p r_p (p_i - x_i) + \phi_g r_g (g - x_i) \tag{5.7}$$

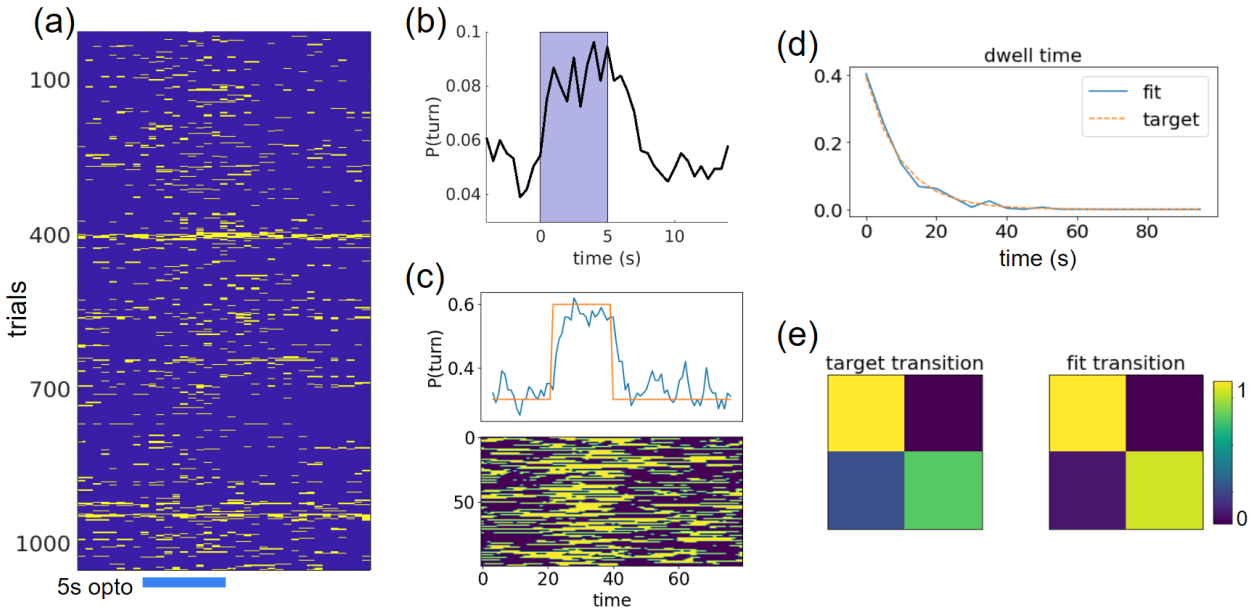

Figure 5.3: **RNN trained to produce stochastic state response**. (a) Experimental measurements of turning behavior (yellow) interleaved with runs (blue) across many trials that are aligned to the 5 second optogenetic input. This is conducted in worm strain AWC::ChR2. Experiments are conducted across 4 plates, each with $\sim$ 100 worms and recorded for 30 minutes. (b) The probability of turn computed from (a). (c) The circuit model trained to produce probabilistic transitions that are input-dependent. The target is in orange and the output from the trained circuit is in blue. Example trails produced by this circuit is shown in the bottom. (d) The dwell time for fitted circuit and the target distribution. (e) The state transition matrix and the transition probability produced by the trained circuit.

where w is a damping factor, ϕ_p and ϕ_g are known as the cognitive and social coefficients, r_p and r_g are separately drawn from uniform distribution $U[0, 1]$. After each trial, the best known parameters p and g are updated according to the fitness function. Empirically, we employed 50 particles and selected parameters $w = 0.5, \phi_p = \phi_g = 2.0$, for stability of the particles and convergence of the solution.

In result, we found that circuits trained through PSO produce behavioral output that agree with experimental measurements in terms of the state transition probability and occupancy (Fig 5.3a-c). Without stimuli, the circuit produces spontaneous state switching, similar to what has been observed experimentally (Fig 5.3d,e). Upon stimulus, the behavioral transition increases the probability of transitioning into reversal state, rather than responding deterministically. The dwell time of the behavioral states can also be matched when it is placed in the fitness function

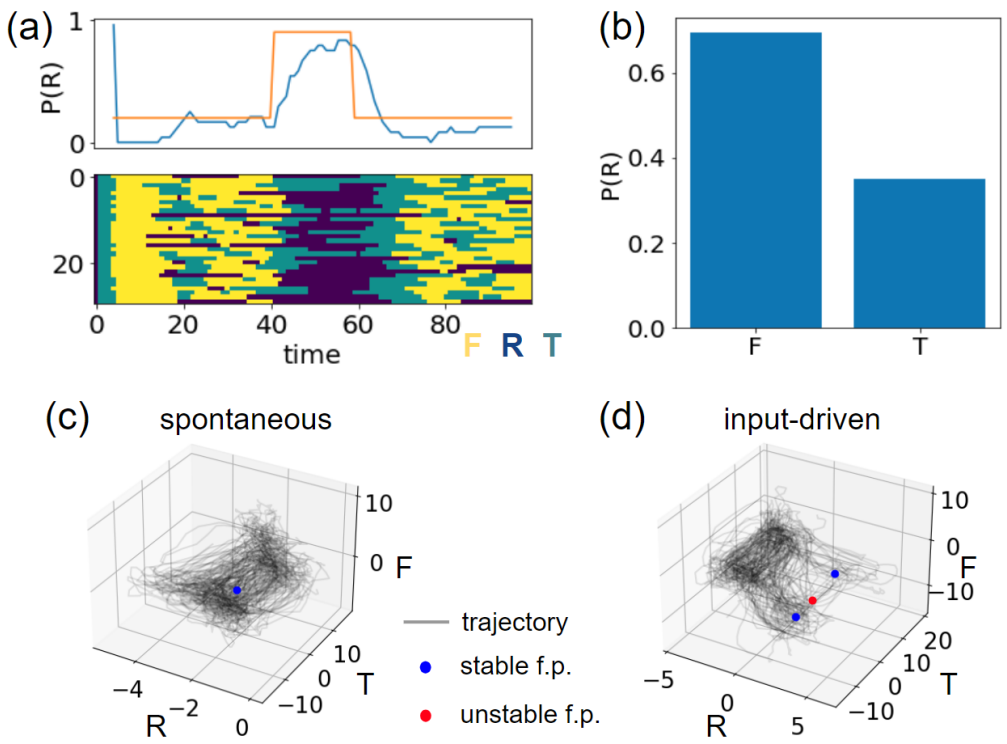

Figure 5.4: **RNN trained to produce state-dependent response**. (a) The probability of reversal (blue) and the target response profile (orange) are shown on top. The time series of three behavioral states (Forward, F; Reversal, R; Turn, T) in response to impulse across trials are shown at the bottom. (b) The reversal probability depends on the state (F or T) right before stimulation. (c) Example of spontaneous neural trajectories in the three-neuron state space, with stable fixed points (f.p.) shown in blue. (d) Same as (c) but for input-driven neural trajectories, where the unstable fixed point is shown in red.

for PSO training. Together, the PSO approach provides solution to a circuit-based model that generates stochastic behavior measured in experiments.

We used this circuit optimization framework to revisit the context-dependent touch response shown in Fig 5.1. Rather than hard-coding the parameters, we explored solutions using PSO and placed in constrains such as the Markov transition matrix across F, R, and T states, as well as the reversal response profile conditioned upon the stimuli (Fig 5.4a). The result shows that the context-dependent touch response can be recapitulated by the optimized circuit, where turning gates the reversal probability (Fig 5.4a,b). We further analyzed an example solution and visualized trajectories in the neural state space (Fig 5.4c,d). During spontaneous motion, there is one stable fixed-points corresponding to dominating forward states and the behavioral output is

only driven by correlated noise. When the stimulus is presented, an unstable fixed point emerges, which separates two fixed points corresponding to reversal and turning states. The learned input weight on the reversal unit is large, indicating a strong drive towards the fixed point with stable reversal. Upon stimulus presentation, however, if the neural trajectory happens to be near the fixed point dominated by turning, it is stable and less likely to cross the unstable fixed point, thereby decreasing the probability to induce a reversal. The context-dependent mechanism therefore comes from the fact that turning is a stable attractor that persists when the reversal state is driven. In contrast to the picture of feed-forward excitation with feedback inhibition for behavioral gating[123], the learned weight matrix W_{ij} is not readily interpretable as we did not see negative weights from T to R. This indicates that there exist more solutions when we considered recurrent, noisy circuits.

Importantly, note that the stochastic neural circuit produces state-switching dynamics that are absent in deterministic models. In a deterministic case, a decision is made after the neural trajectory reaches a fixed point[86,104]. In our stochastic setting, the OU process from I_i alters stability in the phase space and provides flux to cross barrier of unstable fixed point, with escape time proportional to the inverse of noise strength D_i. Indeed, when we checked the learned parameters, the noise strength for the turning unit has a larger strength, which makes the state more short-lived and higher probability to transition out.

Together, the combination of PSO[133] and analysis of neural dynamics[255] inspired by past work[165] demonstrates the flexible solutions in small, noisy, neural circuit models. The result reveals more general computational principles for context-dependent behavior through changing the structure and stability in the neural state space.

5.2 Measurements and models for whole-brain neural dynamics

5.2.1 **Towards whole-brain activity underlying sensory-driven behavior**

To directly investigate the circuit mechanism underlying sensory-driven behavior, we experimentally imaged neural dynamics in freely moving animals upon sensory input. Recent devel-

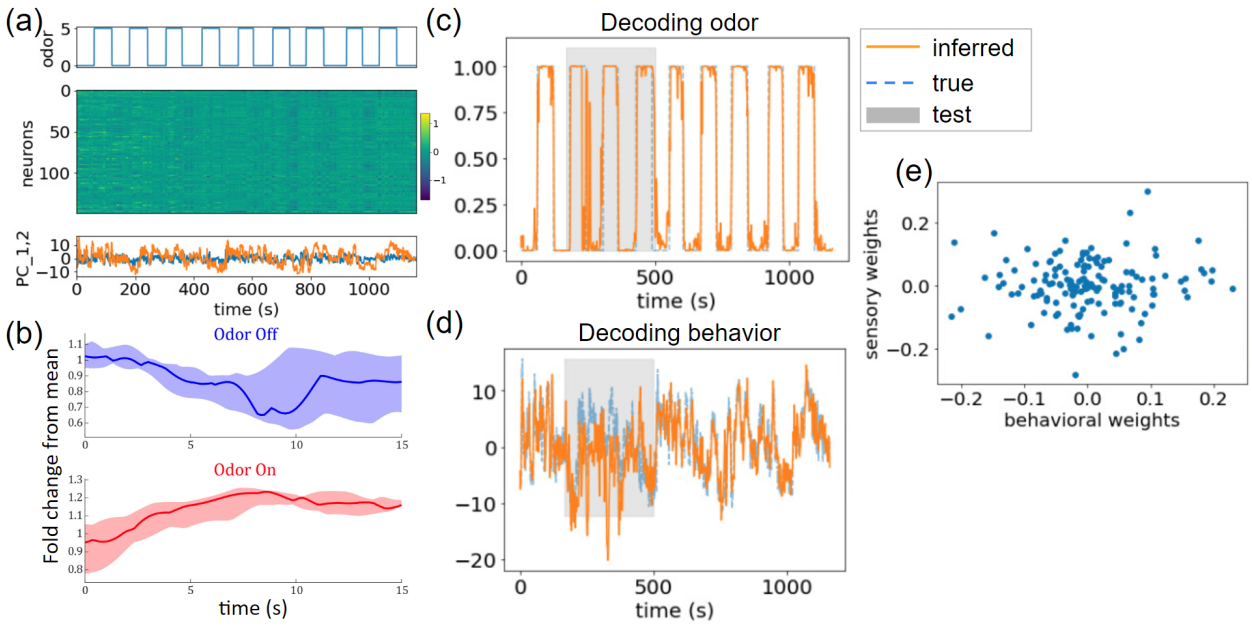

Figure 5.5: **Whole brain imaging during odor presentation in freely moving worms.** (a) Example time series of whole-brain imaging measurements. The odor impulse (top), heat map for neural activity (middle), and PC of the worm posture (bottom) are shown. Odor impulse represent 1.1 mM of butanone in M9 solution, with M9 as background flow then there is no odor. We conducted up to five measurements per training condition; however, only the representative data set is shown in this figure. (b) Two example neurons that are aligned to 9 odor on or off transients. (c) The population decoding performance of odor impulse. The grey region is the testing set and the white is the training set. (d) Same as (c) but for behavioral posture. (e) The scatter plot of neural weights on sensory and behavioral decoding.

opment of optics and closed-loop control in the Leifer lab enables such measurements[184,97]. We developed an experimental assay to conduct whole-brain imaging in worms move in a microfluidic chip while odor flow is presented (Fig 5.5a,b). The experimental setup is inspired by past work[145,198,152] presenting pulsatile odorant to worms in liquid form. We were able to control the odor delivery and monitor the worm's posture while recording from its neural activity. A number of technical challenges remain, including the identification of these measured neurons in the brain and protocols to consistently measure healthy behavior in the setup. At the current stage, we have analyzed preliminary measurements looking into the sensory input and behavioral output relations.

We conducted calcium imaging in a large population of neurons in the worm brain with spinning-disk confocal microscopy. The motorized stage and piezoelectric objective adjust for positions

and focal planes in real time to focus on the worm head and scan through the brain[184]. From the volume of fluorescent signal, we extracted cell-localized signal from individual neurons. This results in a population of $\sim$ 100 neurons that can be correlated to the worm's sensory input and behavioral output. The worm craws in an artificial dirt chip designed to produce fictive crawling through pillars in the PDMS environment[158]. An input port provides constant flow of M9 solution during the experiment, and another channel is used to provide odor flow containing butanone dissolved in M9. With a valve controller (Mux, Multiplex flow matrix, Elveflow) and pressuring source (OB1 - 4 Channel Microfluidic Flow controller, Elveflow) , we periodically swapped between M9 background and odorized flow (1.1 mM butanone in M9 solution) to generate impulse of odor stimuli. The flow rate is held on the order of $\sim$ 10 ml/min, controlled by the valve and driven by pressured source upstream from the odor source. For worm behavior, we fitted centerlines to the worm posture through time and extracted projections onto the eigen-worm basis[234]. The time series of the first two modes are used to characterize worm behavior.

While previous work focus on the neural basis of behavior by regressing behavioral against the population neural dynamics[97,8], the neural representation and transformation of sensory input during behavior is less known. Given the time series measurements, we first identified example neurons that are consistently responding to the onset of offset of the odor impulse. This is in agreement with the sensory population response in literature showing both on and off responses to odor (Fig 5.5b). To understand how the neural population represents odor input, we performed logistic regression on neural activity to study if neural responses can be used to decode the presentation of odor (Fig 5.5c). Similar to behavioral output (Fig 5.5d), the time series sensory input can be decoded from the population of neural activity. We performed the same linear decoding for behavioral time series through linear regression and compare the neural weights for behavioral versus sensory predictions (Fig 5.5e). The result shows little correlation between two weights, indicating that neurons with strong weights to predict sensory input does not necessarily have strong weights to predict odor input.

One assumption for linear or logistic decoding for behavior and sensory signal is that the neu-

ral code has a stable, fixed readout. However, recent work demonstrate that the neural population code can produce significantly different states upon stimuli[8,58]. We hypothesized that there might also be internal-state dependency that are not directly related to sensory input. To study state-dependent neural code, we fitted the neural activity with s switching-linear dynamical system (SLDS) model[153] (Fig 5.6a). This model assumes a discrete latent state that corresponds to a LDS. The state dynamics is a Markov chain with probability altered by sensory input:

$$P(z_{t+1}|z_t) = \frac{\exp(w \cdot O_t + b)}{Z} \tag{5.8}$$

where the discrete latent state evolves in time according to the odor input O_t weighted by w and bias b, with normalization constant Z to form the state transition probability. Given a state, the corresponding LDS follows:

$$x_{t+1} = A^z x_t + b^z + v_t \tag{5.9}$$

$$y_t = C^z x_t + d^z + u_t \tag{5.10}$$

where the superscript indicates LDS parameters that depend on the latent sate z. The matrix A controls linear dynamics of the continuous latent x and the readout matrix C transforms latent to produce the observables y. The bias for latent and observables are b and d. The terms v, u are Gaussian noise following: $v \sim \mathcal{N}(0, Q^z)$ and $u \sim \mathcal{N}(0, R^z)$.

We treated the on-off signal as the odor impulse time series O_t and the neural activity measured through calcium imaging as observable y_t. We fitted measurements to SLDS using the Laplace-expectation-maximization inference proposed previously[153]. In result, we found that SLDS recapitulates reasonable temporal patterns of the population activity. We then hypothesized that these discrete states z not only correspond to different LDS for neural activity and are driven by odor signal but also may correspond to different behavioral readout. To test this hypothesis, we performed linear regression between the population activity and behavioral output conditioned on two different states (Fig 5.6c,d). With the same state occupancy, we shuffled the inferred time series of states and perform the same piece-wise regression as a control analysis. In result, we

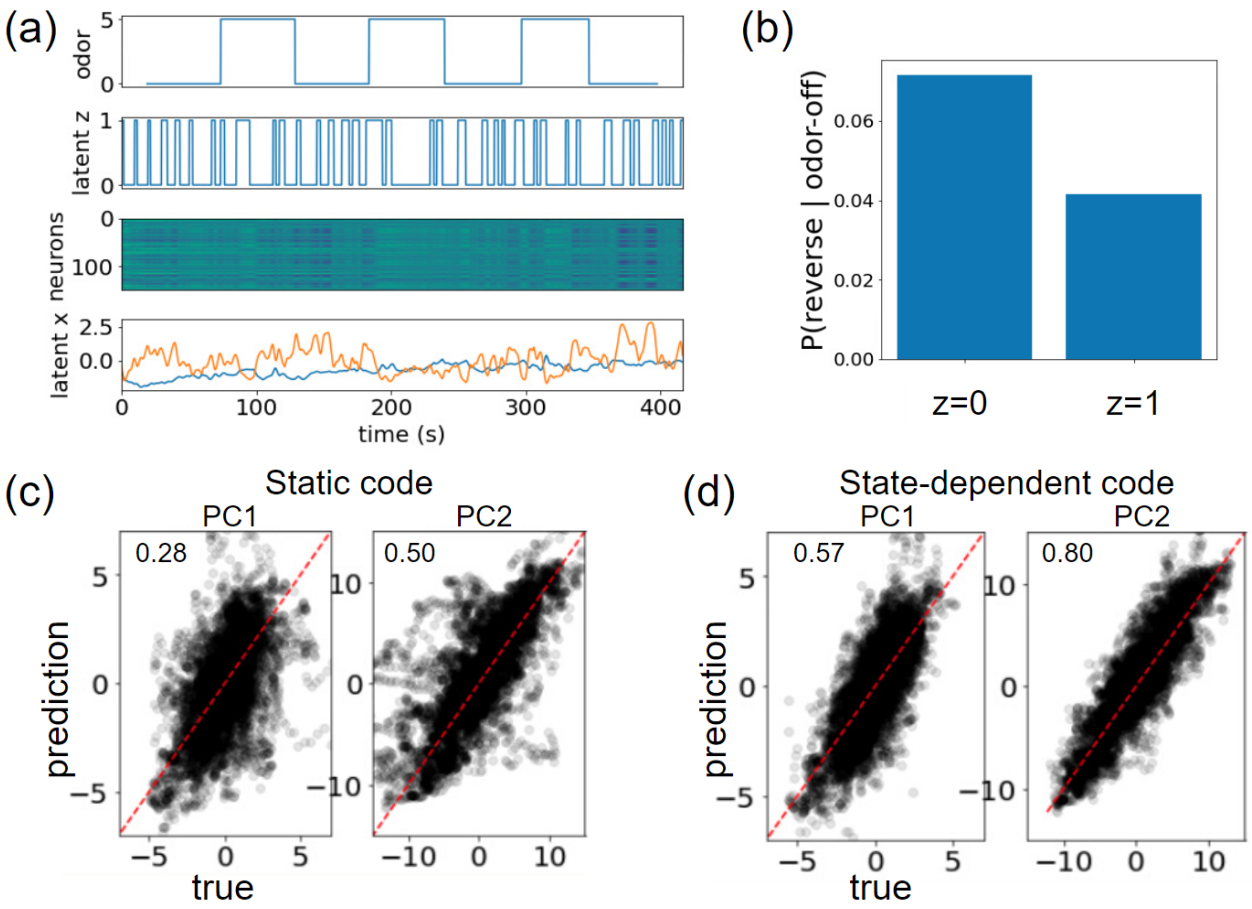

Figure 5.6: **Switching linear dynamical system fit to while brain measurements**. (a) Example traces of SLDS fitted to neural measurements. The odor input, discrete latent state z, reconstructed SLDS neural traces, and the 2D continuous latent x are shown. (b) Condition on different states and triggered on odor off transients, the reversal probability shows state-dependency. (c) The behavioral decoding performance with a fixed weight vector across the neural population. The regression between true and predicted behavior (PC1 on the left and PC2 on the right) are shown, with correlation coefficient indicated in the panels. (d) Same as (c) but with weights that switch depending on the latent state inferred by SLDS.

found that the state-dependent behavioral decoding outperforms a single linear readout or the control analysis. This indicates that the worm's motor program can be better explained by a dynamical readout, rather than a static one.

Together, from the whole-brain imaging setup, odor stimulus presentation, to the decoding analysis through SLDS, we presented progress towards characterizing how stimulus driven and internal state dependent neural activity and behavior are related. We found preliminary evidence of state-dependent response to odor signal by inferring states from neural activity (Fig 5.6b). It

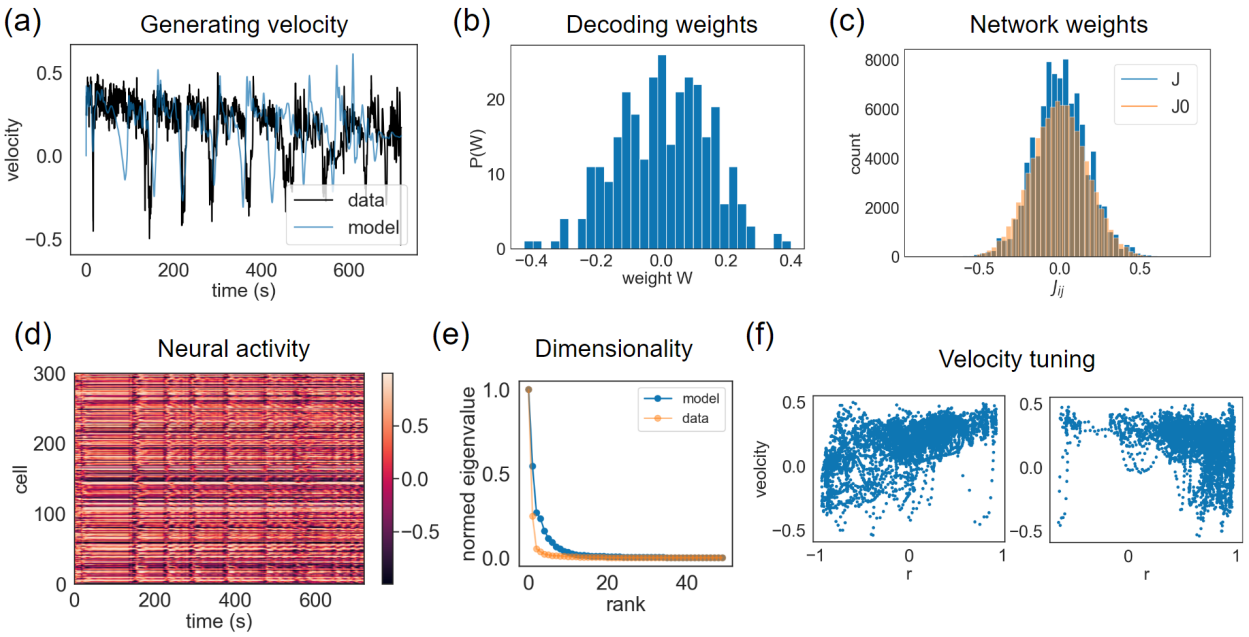

Figure 5.7: **RNN trained with worm behavior**. (a) Experimental data of worm velocity and RNN generated time series. (b) The weights across 300 neurons in the trained RNN to decode behavior. (c) Distribution of the initial J_0 and trained J connectivity. (d) The RNN dynamics during behavior shown in (a). (e) The dimensionality of neural population dynamics in data and model are analyzed with principle component analysis (PCA). (f) Two example neurons from the trained RNN with opposite tuning to velocity output.

remains challenging to connect the neural imaging results with high-throughput behavioral analysis[155,145]. Future study would benefit from the ability to identify neural subtypes to align across individuals for better statistical power, and also have better description for the underlying circuit mechanism for state dependency.

5.2.2 Training neural networks with worm behavior

Studies from the Leifer lab and other recent work show that the worm's behavior can be linearly decoded from the measured population of neural activity[97,8]. The weights on these neural population indicate how each neuron contributes to the behavioral mode or "decision" for specific behavioral action. In artificial neural network, common training procedures of recurrent neural networks (RNN) and reservoir computing also assume linear readout from the population of units[237,70,185]. We explored the connection between behavioral readout and neural network by

training an RNN to produce the worms behavioral dynamics. We then compared the trained neural dynamics and the experimentally measured one.

To train RNNs to produce behavior dynamics, we applied the FORCE algorithm[237]. This algorithm is developed to to study chaotic dynamics and cortical neural activity during monkey's behavioral tasks. In short, FORCE incorporates feedback error signal to make low-rank updates to the network weights during learning. The learning is through iterative least-square approach to update both the readout weights and internal weights. We would revisit this algorithm in detail in the later sections. The algorithm further generalized to full-FORCE that enables full-rank updates[70]. Empirically, we found that full-FORCE produces similar behavior, so here we focus on results from FORCE learning. In practice, we initialized a worm-like RNN with $N = 300$ units and trained it to produce the worm's velocity during whole brain imaging (Fig 5.7). The unites have rate activity and $\tanh$ nonlinearity. The network weight was initialized with $J_{ij} \sim \mathcal{N}(0, \frac{g}{\sqrt{pN}})$, with scaling factor $g = 1.5$ and sparsity $p = 0.3$. This is similar to the standard initialization that poses RNNs at the edge of chaos.

In result, we found that FORCE produces a velocity trace that is qualitatively similar to those measured from worms (Fig 5.7a). The trained RNN alters both the readout and network weights (Fig 5.7b,c) and produces neural activity that has coherent, low-dimensional activity (Fig 5.7d,e). This may be expected as the learning algorithm corrects weights in the RNN according to a behavioral time series that is low dimensional. Across the population, we found neurons that are positively tuned to the trained velocity signal, as well as neurons that have negative tuning or not relation to the behavioral signal (Fig 5.7f). These rich tuning types have also been reported in the worm brain[97].

We further analyzed the structure of the trained population code. By rank-ordering the neurons that contribute to decoding and cumulatively recruit them for population decoding, we recovered a trend that is qualitatively similar to experimental observations[97]. This curve shows that the number of neurons required to decode behavioral rises with the number of neurons and saturates at a size smaller than the full RNN (Fig 5.8a). Furthermore, the population code outperforms de-

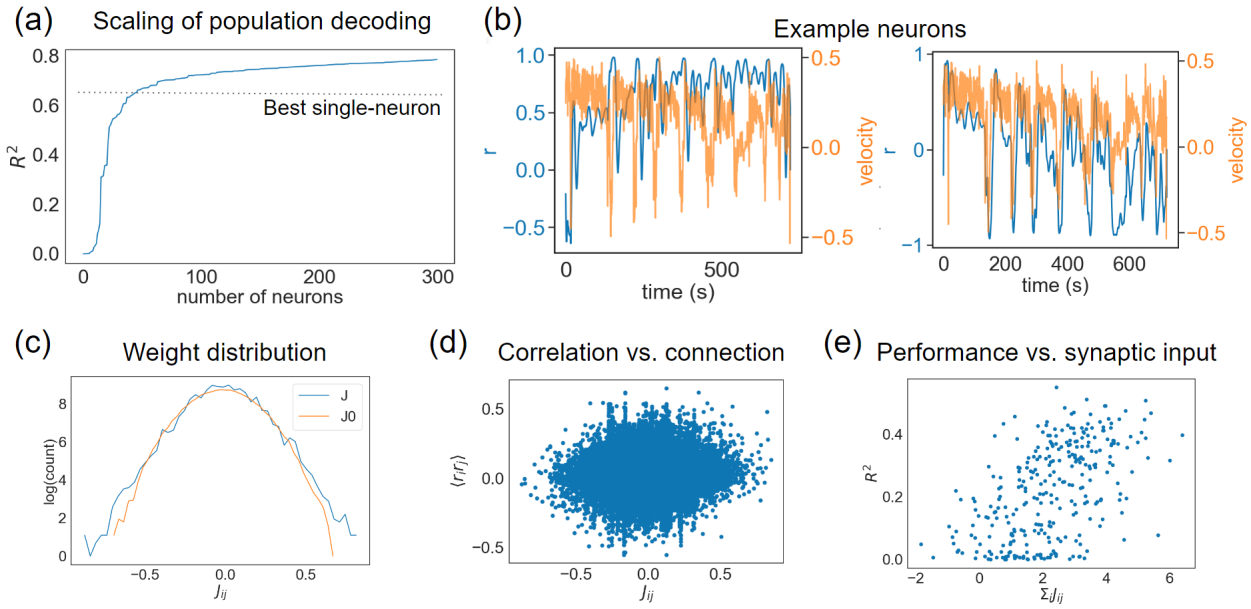

Figure 5.8: **Analysis of network trained with worm behavior**. (a) The behavioral decoding performance (R^2) as a function of the accumulated number of neurons chosen for decoding. The comparison to the performance of a best single neuron is shown in dash line. (b) Two example traces with neural response, r, and the worm velocity are shown. The correlation coefficient between response and behavior is 0.58 (left) and -0.35 (right). (c) The distribution of trained network weights compared to the initial Gaussian distribution, J_0. (d) The scatter plot for the correlation between neuron pairs, $\langle r_i r_j \rangle$, and the trained connectivity, J_{ij}. (e) The scatter plot for decoding performance R^2 of a single neuron and its summed synaptic input, $\Sigma_j J_{ij}$.

coding with the signal best neuron.

The trained RNN offers an opportunity to study how structure relates to neural activity and coding, which is less feasible in biological measurements. We analyzed the weight distribution and found that it deviates with the initialized Gaussian distribution by showing long-tails in the connectivity distribution (Fig 5.8c). This means that there are certain neurons that have over-represented connections to other neurons in the brain, which is the feature observed in the connectome as well[250]. We also found correlation between the cross-correlation of neural signal and the trained coupling strengths, supporting the empirical finding showing that neural dynamics can reflect neural connectivity to some extent (Fig 5.8d). Lastly, we analyzed how the summation of synaptic input correlates to the neuron's weight on behavioral decoding (Fig 5.8e). This analysis is motivated by past work that proposed the degree of connection to be an important indicator for control signal in abstract networks[250]. We found a positive correlation between the synaptic

summation and neural correlation to behavior. This result forms a testable prediction for future experiments, postulating that the candidate neurons with more synaptic connections are more likely to affect behavior through perturbation or ablation. Future studies with neural recordings in combination of the neural identities[8] and connectome-constrained modeling[175] are required to connect network structure with neural dynamics.

5.3 Learning algorithms for stochastic spiking recurrent networks

5.3.1 **Introduction**

An important topic in computational and systems neuroscience is to understand the dynamical properties of recurrent neural networks (RNNs)[237,86,254,255]. However, recent theoretical work has tended to focus on RNNs with deterministic dynamics[237,70,185,202,168]. In experimental studies, stochasticity in single neurons or neural populations are often characterized by statistical models[193,220]. By comparison, relatively little work has focused on the dynamics and network properties of stochastic spiking neural networks. Here, we explored the dynamical repertoire of a RNN with Poisson generalized linear model (GLM) neurons[193]. We measured time correlation and maximum Lyapunov exponent of neural dynamics as a function of connection strength and find that GLM networks exhibit similar range of complex behaviors as rate networks in strong coupling regimes. We proposed a supervised learning algorithm called G-FORCE to modify coupling strengths in GLM network to produce coherent target dynamics. We analyzed how biophysical parameters affect learning and found that heterogeneous filters and synaptic nonlinearity are required. Furthermore, with the tractable statistical properties of GLMs and sparse, low-rank network structure of the trained network, we tested an efficient inference method for recovering the model parameters from spike trains. Together, this work connects theoretical and statistical approaches in stochastic spiking RNNs, extending learning algorithm and inference method that incorporate tractable statistics and theoretical foundations.

5.3.2 The G-FORCE algorithm

Overview for FORCE learning

The original FORCE algorithm trains a vanilla RNN model to autonomously produce dynamics by feeding back the reconstructed signal z to the network[237,185]:

$$\frac{dx_i}{dt} = -x_i + \Sigma_j J_{ij} r_j + u_i z \tag{5.11}$$

where J_{ij} is the network connectivity, $r = tanh(x)$ is the nonlinearity, u is the feedback connection, and $z = \mathbf{W}^T\mathbf{r}$ is the linear reconstruction from the network with weight vector $\mathbf{W}$. Rearranging the network model in vector form gives:

$$\frac{d\mathbf{x}}{dt} = -\mathbf{x} + (\mathbf{J} + \mathbf{u}\mathbf{W}^T)\mathbf{r} \tag{5.12}$$

which shows that the feedback signal is effectively a rank-one update of the connectivity J as we learn the decoding weights W. The learning rule has an objective function to minimize the error between reconstruction and a dynamical target $f(t)$:

$$\mathcal{L}_F = \frac{1}{2}\Sigma_t (W^T r(t) - f(t))^2 \tag{5.13}$$

then with Newton's method for weight learning, the update rule follows recursive least squares algorithm:

$$W(t + \Delta) = W(t) - e(t)P(t)r(t) \tag{5.14}$$

where the error term is $e(t) = W^T r(t) - f(t)$ and the online approximation for inverse covariance $\langle rr^T \rangle_t$ is:

$$P(t) = P(t - \Delta) - \frac{P(t - \Delta)r(t)r(t)^T P(t - \Delta)}{1 + r(t)^T P(t - \Delta)r(t)} \tag{5.15}$$

Together, the FORCE algorithm minimizes the difference between the linear construction and

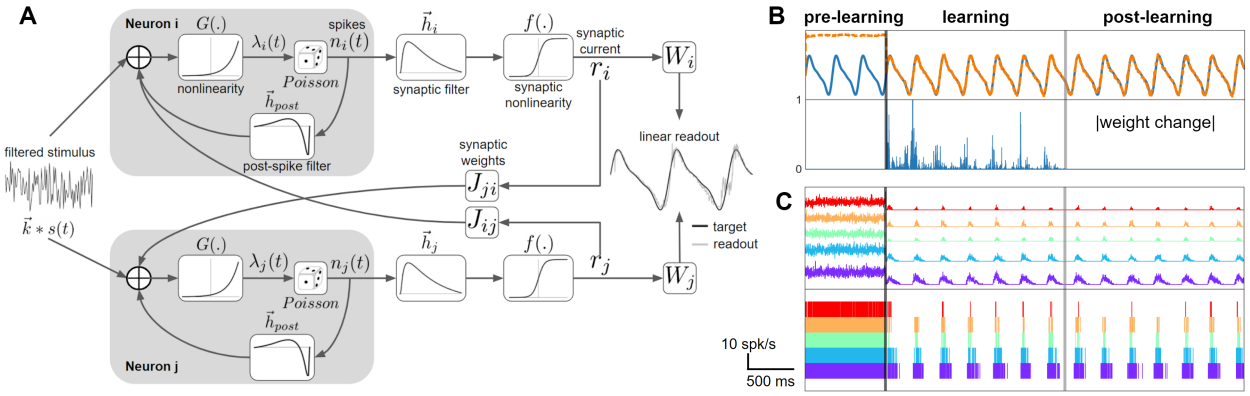

Figure 5.9: **G-FORCE**. (a) Schematic of the GLM network. Each neuron is modeled by Poisson GLM. Connectivity J is trained through G-FORCE learning to generate readout through W. (b) The target (blue solid) and reconstructed time series (orange dash) with change in weight shown below. (c) Time series of 5 example Poisson spiking neurons. Poisson rate shown above and Raster plot for spikes below. Some cells spike precisely and others burst in each cycle. Unless further specified, we use model parameters: $N = 300$, $g = 1.5$, $p_s = 0.8$, kernel length of h is 100 and constructed with 5 basis.

target dynamics. The iterative process alters the decoding weights and makes a rank-one update to the network connections.

G-FORCE

The G-FORCE method extends FORCE algorithm to train a GLM network to produce coherent dynamic output (Fig 5.9). We started by formulating the Poisson-GLM spiking neural model:

$$n_i(t) \sim Poisson(\lambda_i(t)dt) \tag{5.16}$$

with $n_i(t)$ being the spike count in a time bin dt and $\lambda_i(t)$ as the Poisson spike rate of neuron i at time t. The Poisson spike rate is generated from the recurrent neural network:

$$\lambda_i(t) = g(\Sigma_j J_{ij} r_j(t)) \tag{5.17}$$

where g is a nonlinear function, J_{ij} is the recurrent network connectivity with synaptic projection from neurons j to i, and r is the firing rate computed from the generated Poisson spikes in

equation 5.16:

$$r_i(t) = f(\Sigma_j \int h_{ij}(t-\tau)n_j(\tau)d\tau) = f(\Sigma_j(h_{ij} * n_j)(t)) \tag{5.18}$$

where f is another nonlinear function, h_{ij} is the coupling filter for neurons j to i for $i \neq j$ and post-spike filter for $i = j$ (h_{ij} would be a $N \times N \times D$ tensor, with D as the time window for filters). Note that if f is linear, this recovers the original GLM framework in Pillow et al. 2008[193]. Here, we extend the framework to the FORCE algorithm to train spike recurrent neural networks. Following the FORCE algorithm from Sussillo & Abbott 2009[237], we assume linear readout from the network as:

$$\hat{x}(t) = \Sigma_j W_{ij} r_j(t) \tag{5.19}$$

where $\hat{x}$ is the reconstructed dynamical target signal x and W_{ij} is the linear readout weight. During the learning phase, the weight update of row vector w_i follows recursive least square (RLS) algorithm:

$$w_i(t) = w_i(t-\Delta t) + e_i(t)^T P(t) \tilde{r}_j(t) \tag{5.20}$$

where $e_i(t) = x(t) - \Sigma_j W_{ij} r_j(t)$ is the error term, $\tilde{r}_i = \frac{\partial g}{\partial r} r_i$ is the scaled filtered spike train, and P is the inverse of correlation matrix $\langle \tilde{r}(t)\tilde{r}(t)^T \rangle^{-1}$. The matrix P is approximated iterative through:

$$P(t) = P(t-\Delta) - \frac{P(t-\Delta)\tilde{r}(t)\tilde{r}(t)^T P(t-\Delta)}{1 + \tilde{r}(t)^T P(t-\Delta)\tilde{r}(t)} \tag{5.21}$$

Assuming that the weight matrix can be decomposed into static and plastic connections $J = J^0 + \Phi W^T$, where J^0 is static and Φ is an encoding random matrix, the network connectivity matrix can be updated with:

$$J(t) = J(t-\Delta t) + e(t)^T P(t) \tilde{r}(t) \tag{5.22}$$

Together, the G-FORCE algorithm integrates a network of GLM spiking neurons with the FORCE algorithm for supervised learning. We compared this method with known algorithms and further show result of G-FORCE applied to different target dynamics.

Parameter setup for G-FORCE

For the GLM spiking neurons, the nonlinear function $g(\cdot)$ is a rectified linear function $[\cdot]_+$ and $f(\cdot)$ is the sigmoid function $f(x) = \frac{F_{max}}{1+\exp(-x)}$, with F_{max} as the maximum firing rate. Spike rate with $F_{max} = 1$ or small dt turns into a Bernoulli spike train and this did not affect our results. The post-spike and coupling filters h_{ij} are parameterized with raised cosine basis functions[193]:

$$b_n = \frac{1}{2}\cos(a\log[t+c] - \phi_n) + \frac{1}{2} \tag{5.23}$$

where parameters a, c, ϕ are chosen according to the literature. Specifically, weights on each basis function for each h_{ij} is drawn from a Gaussian distribution for $i \neq j$ and a strong negative weight near $t = 0$ kernel is imposed for $i = j$ as self-post-spike signal to enforce refractory periods in its spiking behavior.

For the recurrent network structure, we randomly zeroed out a fraction of coupling kernels to produce sparse connections, given a probability of removal p. Following the original FORCE framework with chaotic recurrent neural networks, the connectivity matrix J_{ij} initially starts with weights drawn from zero-mean Gaussian scaled with its size and sparsity p. We integrate sign and synaptic balancing constraints according to the setup in Nicola & Clopath 2017[185].

Parameters used in the example figures here are: $N = 300$ neurons (note that this is a significantly smaller network compared to the original FORCE framework), $T = 200, dt = 0.1$ for time course of simulation, sparsity for J_{ij} is 0.5 and for h_{ij} is 0.2. The prior term α in FORCE algorithm is 0.1 and we updated weights every two steps. The network J_{ij} is initialized with zero-mean Gaussian distribution re-scaled by $\frac{g}{\sqrt{pN}}$, with $g = 1.5$. For the GLM network, we used five raised cosine basis functions, weights were generated from Gaussian distribution, and with 50 time steps in the temporal window. The diagonal kernels were enforced to be strongly negative post-spike kernels. Target signals were generated from periodic traces of 4 sinusoides and would be further extended to sequences across the population and chaotic dynamics from Lorenz attractor.

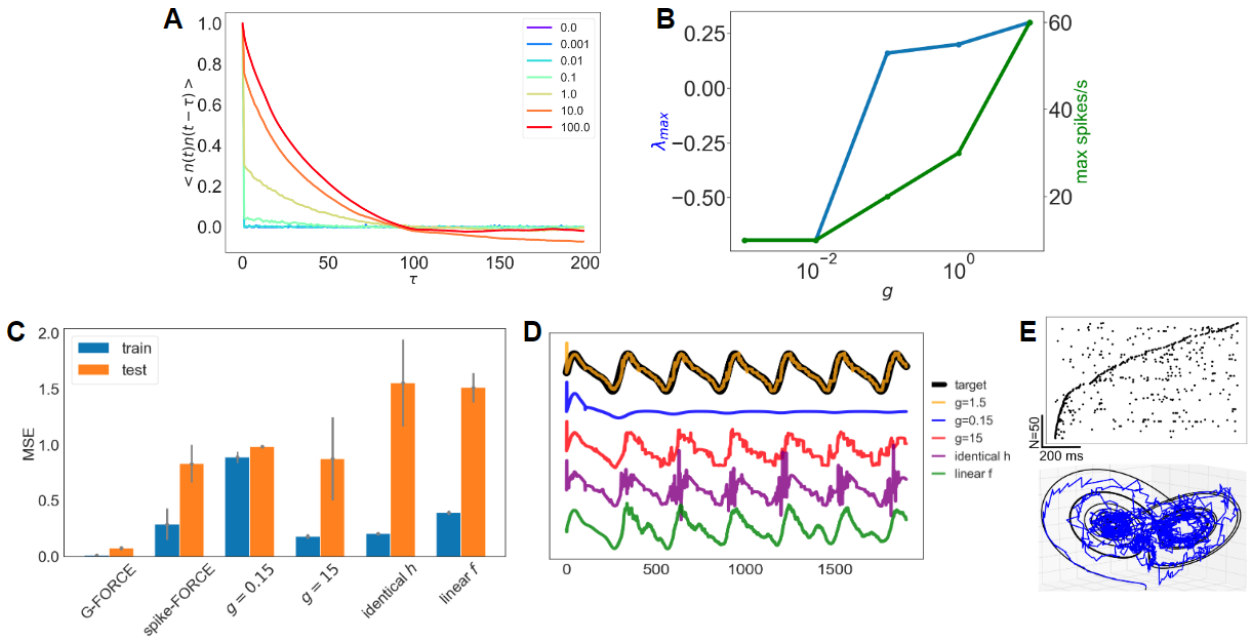

Figure 5.10: **GLM network dynamics and learning**. (a) Autocorrelation of Poisson spike rate with different g color-coded. (b) Largest Lyapunov exponent λ_{max} and maximum Poisson rate as a function of g. (c) Train and test mean squared error (MSE, computed for normalized signals across 20 instantiations) for different parameters. Strength g, kernels h, and function f are parameters in G-FORCE. Spike-FORCE follows[185]. (d) Target time series and trained reconstruction with different conditions. (e) Sorted Raster plot for network trained to produce sequences (top, following conditions in[202]) an Lorenz attractor (bottom, blue for readout and black for target) with G-FORCE.

Results of G-FORCE training

We showed that G-FORCE algorithm can be used to train a network of GLM neurons (Fig 5.9b). The network learns coherent dynamics to reconstruct the target time series. During the training phase, the network weight changes drastically in the first few cycles of the target dynamics and soon reaches a relatively steady value. The resulting population dynamics show coherent spiking patterns (Fig 5.9c). The network successfully reconstructs the target signal autonomously during the testing phase. These results are repeated across different network instantiates.

To gain insights into the stochastic dynamics of GLM network, we first studied the network with random connectivity. The autocorrelation time scale of neural activity elongates beyond the length of single neuron filters when the coupling strength is large (Fig 5.10a). In an intermediate network strength, the maximum firing rate is still biophysically plausible (below 40Hz spike

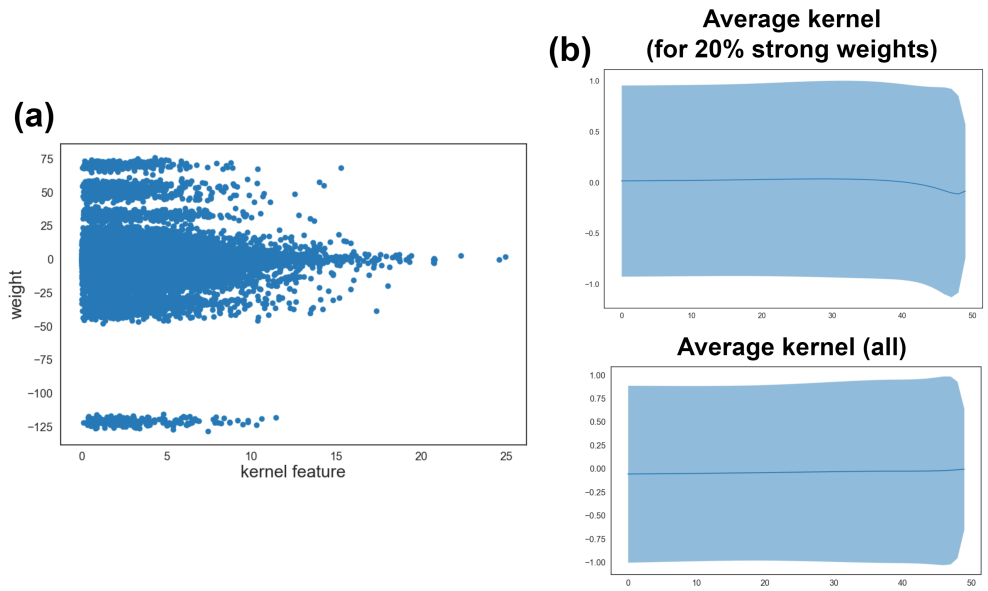

Figure 5.11: **Analysis of the trained GLM network**.(a) The distribution of learned weights in matrix J_{ij} versus the GLM coupling kernels h_{ij}. Kernel feature are computed with its L_2 norm. (b) The average kernel h_{ij} that have strong learned connection in J_{ij} on the top and for all kernels on the bottom.

rate) and exhibits positive Lyapunov exponent (Fig 5.10b), which indicates chaotic spiking dynamics in the network. We then explored how network parameters affect G-FORCE learning and found that an intermediate coupling strength g is required. We found that heterogeneous population of kernel shapes or time scales in h supports stable learning. Synaptic nonlinearity f is also crucial. We observed declined learning performance due to instability when f is a linear function (Fig 5.10c,d). This may reflect that the network learns in a regime with positive Lyapunov exponent and heterogeneous h that provides rich temporal basis, while keeping low firing rate and nonlinear f for stability. G-FORCE with optimal parameters outperforms deterministic spike FORCE method with finite network size ($N = 300$). Finally, we demonstrated that G-FORCE can also be used to learn complex neural dynamics such as sequences and chaotic attractors (Fig 5.10e).

We further analyzed the G-FORCE-trained network parameters. We compared the learned connection strength between neurons against the norm of coupling kernel (Fig 5.11a). It is unclear if the learning procedure would pick up neurons with particular coupling strength in the fast chan-

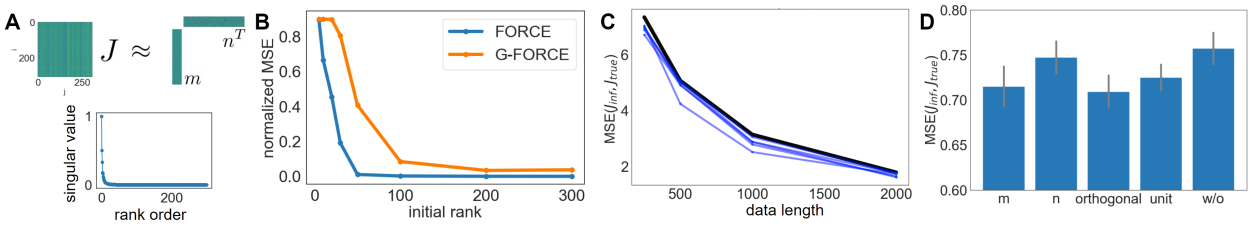

Figure 5.12: **Low-rank structure and network inference.** (a) Network J with low-rank decomposition (top) and singular values (bottom) after G-FORCE learning. (b) Normalized MSE for training with different rank of the static network J_0. (c) MSE between the inferred (J_{inf}) and true (J_{true}) structure converges as a function of data length. Solid dark line is full-rank maximum-likelihood and blue lines are low-rank approximation. (d) MSE(J_{inf},J_{true}) with different projections of the noisy input. Input is projected to m, n, orthogonal ($\{m, n\}^{\perp}$) subspace, unit is an uniform input, and w/o is without input (error bar for 10 repeats).

nel. We found that this is not the case – the network strength has little correlation with the kernel amplitude. We further compared the average kernels that have stronger learned connectivity versus other kernels (Fig 5.11b). We found that those with strong connectivity tend to have slightly negative kernels on average. We also revealed that the constraint on self-post-spike kernel is required to stabilize spiking patterns. Together, these analyses provide insights into training stochastic spiking RNNs that have stable, generic dynamical properties.

5.3.3 Inference with low-rank networks

The Poisson-GLM is a powerful descriptive model that can be fit to experimental measurements and learn about network parameters[193]. We sought to leverage the low-rank structure[168] in networks trained by G-FORCE and rewrite the connectivity as $J = mn^T$, where m and n have dimensions $N \times d$ and $N > d$ (Fig 5.12a). In addition, we found that the static initial connectivity J_0 does not have to be full-rank. In both FORCE and G-FORCE learning, we initialized with a low-rank random matrix and show that the learning performance is not affected if J_0 has sufficient low-rank structure (Fig 5.12b). We developed an efficient inference for low-rank connectivity by parameterizing the joint probability as:

$$P(r, m, n, \gamma, \theta) = P(r|m, n, \theta)P(m|\gamma)P(n|\gamma)P(\gamma) \tag{5.24}$$

where γ is the prior on low-dimensional matrix m and n. Notation $\theta = \{h, k, n, s\}$ includes parameters and data from the GLM network as we focus on inference for connectivity with parameters m, n. We show that the low-rank inference with less parameters outperforms full-rank maximum-likelihood methods (Fig 5.12c).

Many brain-scale experimental measurements seek to characterize neural connectivity from activity induced by stimulation. Here we explore inference in input-driven network with structured input rather than white noise (Fig 5.12d). This setup mimics the non-random external input driving neural dynamics through electrodes or optogenetics. The stimulus $s(t)$ is projected to the subspace of m, n, or an orthogonal subspace. We found that projections onto the right-connectivity vectors n seem to deteriorate network inference, suggesting directions for future experiments with perturbation for network inference.

5.3.4 **Exploring state-dependent neural dynamics**

While FORCE learning works with a variety of different complex target dynamics, a large (~ 1000 units) network size initialized near then dynamical critical point for chaos is required[237,70]. In practice, a number of complex target signals can in fact be decomposed into repeating patterns across alternating states. This includes multiple-period signals, human posture data, and potentially even the Lorenz chaotic attractor. Here, we sought to parameterize the loss-function for RNNs accordingly. In a statistical modelling viewpoint, the original loss function is assuming that the observed decoding signal is conditionally Gaussian with the mean linear readout $P(\hat{f}|r, W) \sim \mathcal{N}(W^T r, \sigma^2)$, with some σ^2 variance around the linear readout. We extended this generative model with a latent state z and state-dependent observation:

$$P(\hat{f}, z, r, W) = \prod_{t=2}^{T} P(\hat{f}_t|r_t, z_t, W) P(z_t|z_{t-1}, r_{t-1}, W) \tag{5.25}$$

where the first probability term on the right hand side is the emission given the trained weights W and latent state z_t. The second term is the state-transition probability. Given specified form of

emission, transition, and initial distribution (at $t = 1$), the state-space model can be solved with EM algorithm. However, the new objective is to train an RNN to represent the state dynamics within the network itself. The method to parameterize the constraints is inspired by recent studies on low-rank RNN dynamics [168,144], where the dynamical systems follows the vanilla RNN form with connectivity decomposed as:

$$J_{ij} = g\chi_{ij} + P_{ij} \tag{5.26}$$

where g is the strength on a zero-mean Gaussian random network χ_{ij} and P_{ij} is an additive term that provides structure to the network. It is constructed with low-rank $P_{ij} = \Sigma_d^D \frac{m_i^d n_j^d}{N}$, where there are D dimensions constructed by a vectors m and n. It has been shown that low-rank network can produce a rich variety of state-space dynamics dependent on the inward and outward vectors m and n added to the random network [144,168]. With intermediate input strength, this rank D structure can result in D stable states in the network dynamics. More importantly, to connect with FORCE learning, we can view the structured mn^T perturbation as the feedback signal uW^T. We hypothesize that this representation can be used for state-dependent learning. A simple parameterization for activity-dependent state-transition (between K states) in the network would be a soft-max probability:

$$P(z_t = k|z_{t-1}, r_t) = \frac{\exp(W_k^T r_t)}{\Sigma_{k'}^K \exp(W_{k'}^T r_t)} \tag{5.27}$$

and we considered $P(\hat{f}_t|r_t, z_t, W)$ conditionally Gaussian for simplicity but be generalized with Poisson emission in the future. The state-space model (SSM) constrained objective can then be formulated as:

$$\mathcal{L}_{ssm} = -\Sigma_t \log P(\hat{f}_t|r_t, z_t; W) P(z_t|z_{t-1}, r_{t-1}; W) \tag{5.28}$$

We computed the first and second derivative with respect to weight W. Now the Newton's method update rule follows:

$$W_k(t + \Delta) = W_k(t) - e_k(t)\bar{P}_k(t)r(t) - \bar{P}_k(t)r(t)(1 - P(z_t = k|z_{t-1}, r_t)) \tag{5.29}$$

where e_k is now the error term for the k-th state governed by the weight W_k, $e_k(t) = W_k^T r(t) - f(t)$. The third term on the right hand side is a new component compared to the original FORCE algorithm. Importantly, the inverse Hessian term $\bar{P}_k \sim \frac{\partial^2 \mathcal{L}_{ssm}}{\partial^2 W_k}^{-1}$ is similar to but not exactly the same as inverse covariance. It follows:

$$\bar{P}_k(t) = \bar{P}_k(t-\Delta) - \frac{\bar{P}_k(t-\Delta)r(t)r(t)^T\bar{P}_k(t-\Delta)}{\lambda_k(t)^{-1} + r(t)^T\bar{P}_k(t-\Delta)r(t)} \tag{5.30}$$

where $\lambda_k(t) = 1 - P(z_t = k)(1 - P(z_t = k))$, corresponding to the calculation of online Hessian calculation for each state k. This is derived from the same matrix inversion lemma. Together, with the new weight update rule, we updated the full network connectivity by the same rule as FORCE. Now each dimension corresponds to the potential latent state:

$$\mathbf{J}(t+\Delta) = \mathbf{J}(t) + \Sigma_k \mathbf{u_k} \mathbf{W_k^T} \tag{5.31}$$

where the second term on the right hand side is a rank-K update to the network connectivity. Note that the constraint with this formalism is that the output weights is now governing the state-transition dynamics as well.

Preliminary results and model comparison

We compared the performance of $\mathcal{L}_{ssm}$ and the original $\mathcal{L}_F$ in FORCE learning. We generated a target that has two latent states periodically switching between low and high frequency (Fig 5.13a). This signal is the target dynamics for a chaotic RNN to be trained on and there are two feedback loops to provide rank-2 update to the trained network. We found that the performance with objective function that includes the latent state switching outperforms those without (Fig 5.13b). Specifically, the trained network generates this state-switching dynamics more reliably given the same network dynamics and size.

Given that the target signal is a probabilistic one, we evaluated the performance by taking the largest value of the time-shifted cross-correlation. The output dynamics of the RNN is often

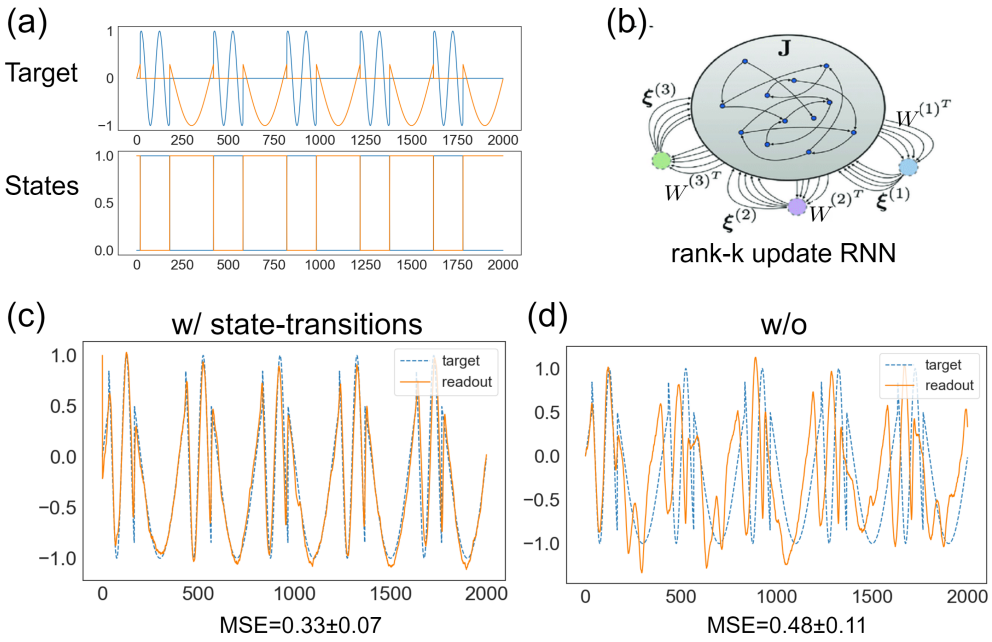

Figure 5.13: **Performance of FORCE with state-dependent target dynamics.** (a) State-switching oscillating target signal with two frequencies. (b) A schematic for rank-K update of the network connectivity. (Figure panel is modified from [144]) (c,d) Target and readout from RNNs trained with and without SSM constraints. The average MSE across 10 instantiaton is shown below.

shifted in phase compare to the target signal, but peak of cross-correlogram takes care of this lag. With the same low-high frequency switching target, we showed that the trained RNN can produce this state-switching signal and the latent states (Fig 5.13c,d).

While we incorporate the objective to learn discrete states, the resulting readout along the W_k projections do not form persistent states. There are different candidate neural implementation to mitigate this issue, including parameterizing another sub-network to learn discrete states [237] or adding winner-take-all (WTA) computation in the network [245]. Future work is required for comparison to other spiking network training algorithms [108,16,135,22,124] that either use approximate gradient or variational methods to train spiking networks. Additional sophisticate network training includes structural constraints from the connectome [175], allows multi-modal observations [140], complex step-wise training and sub-networks [112], or training through variational objectives [175,124]. These generalization of network training would be important for future application to biological measurements. Currently, our ongoing work is to fit GLM networks to spike trains generated

from a state-space model, such as an input-driven GLM-HMM with Poisson emission. Empirically, this method recapitulates state-switching behavioral action observed in experiments[7]. Theoretically, future work is needed to rigorously formalize training with latent state constraints into variational learning algorithms[124]. More detailed derivation, proposed implementation, and ongoing work are shown in the Appendix. Together, we have provided initial formulation of training algorithms and numerical exploration towards the connection of neural dynamics in RNNs and statistical frameworks of state-space models.

Everything should be made as simple as possible,

but not simpler.

Albert Einstein

6

Models for Motion Artifact in Imaging, Predictive Dynamics in Visual Cortex, and Inference for Nonequilibrium Networks

In this chapter, we introduced three theoretical projects that are not directly related to navigation behavior or neural circuits in worms. The projects include: (1) a latent variable model to remove motion artifact in whole-brain imaging in freely-moving worms, (2) an adaptive sub-unit model developed to capture predictive dynamics in mouse visual cortex, and (2) an abstract network model used to study dynamics and inference in the nonequilibrium regime.

The first project has been published in *PLOS Computational Biology*[59], with postdoc Matt Creamer as leading author and advised by Andy Leifer and Jonathan Pillow. I contributed to the inference derivation and implementation of alternative models. The second published in *PNAS*[102], with postdoc Jan Hommann as the leading author and advised by David Tank and Michael Berry. I contributed to the modeling part – developing the adaptive subunit model and fits to measurements. The third one is across arXiv, *Entropy*, and submitted to *PRE*[47,43]. This is a series of theory work that is independent or in collaboration with a postdoc Ying-Jen Yang who works with Ken Dill at Stony Brook. For succinctness, the content presented in this chapter only focuses on main parts I contributed in the projects, rather than the full papers shown in previous chapters. While these projects are independent to one another, they mark as the theory and computational component that are crucial in systems neuroscience and biophysics research.

6.1 Latent variable model for motion correction in neural imaging

6.1.1 Introduction

Recent development in neural imaging enables population-scale measurements of calcium signal in freely moving animals[184,8]. This opens the opportunity to probe the neural basis of behavior. However, animal motion during calcium imaging poses unique challenges both for segmenting and tracking neurons and for accurately extracting calcium traces. Many computational approaches have been proposed to register, segment and track neurons in the presence of motion[273,183,194]. These approaches account for the gross movement of a neuron relative to its neighbors or within the field of view. By contrast, there have been relatively few efforts to account and correct for motion-related changes to the extracted fluorescence intensity time-series itself. Accounting for these latter effects of motion, however, is particularly important because motion related changes to fluorescence intensity can appear similar to behavior-related neural signals of interest. Here, we construct a method called Two-channel Motion Artifact Correction (TMAC) which uses the fluorescence traces from two-channel imaging to correct for shared motion artifacts between the two channels.

Artifact in calcium imaging can result from many sources. For instance, when a neuron deforms or moves relative to an imaging plane, its fluorescence intensity can change unrelated to neural activity (Fig 6.1A). Others include subtleties of image acquisition and segmentation processes. In real measurements, a strategy to remove artifact is to include a calcium-independent indicator, such as RFP, concurrently with the calcium indicator GCaMP. In principle, information in the two-channel setup have shared motion artifact, and the true neural signal can be extracted from the green channel. Following the logic of FRET imaging in single molecular studies, the ratiometric (green over red intensity) approach provides a proxy for neural activity, and proved useful in freely moving animals[53] . However, the underlying assumption of taking the ratio, its limits, and comparison to other metrics are unknown. TMAC explicitly models the statistical distributions of the noise, artifact, and cell activity and uses Bayesian inference to find activity uncontaminated by motion artifact. We use experimental data to compare TMAC with four other approaches to motion artifact removal and demonstrate that it outperforms previous methods.

6.1.2 Statistical model for motion artifact

To account for fluctuations in fluorescence induced by motion artifact, we propose a method called Two-channel Motion Artifact Correction (TMAC) (Fig 6.2A). This method relies on a latent variable model of calcium fluorescence, which we use to infer the latent neural activity of each neuron. Intuitively, inferring the latent activity from the model amounts to subtracting the motion signals present in the red channel from the green channel while accounting for channel independent noise. We first process the fluorescence data by dividing each channel by its time average, such that the processed data has mean 1 and units of fold change from the mean. By specifying our model fluorescence as fold change from the mean the parameters do not depend strongly on experimental conditions such as cell expression, laser intensity, or image acquisition time, allowing us to compare across different neurons and fluorophores (Fig 6.1B,C).
We define the fluorescence from a neural activity-dependent green channel and a neural activity-independent red channel as follows. Let **r** and **g** denote vectors of the preprocessed red chan-

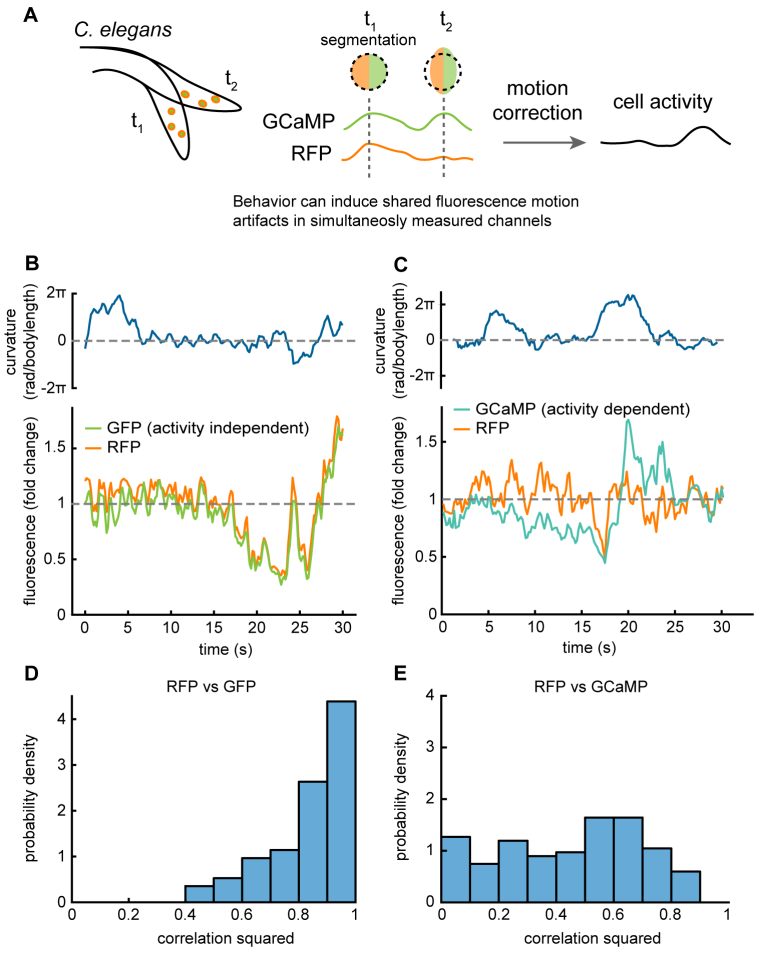

Figure 6.1: **Behavior induces motion artifacts in optical recordings.** (A) Diagram of a moving *C. elegans* undergoing motion and deformation during simultaneous measurements of GCaMP and RFP fluorescence. Deformations of the cell can conspire with imperfect segmentation to create motion artifacts. RFP intensity fluctuations reflect motion or noise. GCaMP intensity fluctuations reflect motion, calcium activity, and noise. (B) Top, Animal body curvature. Bottom fluorescence of a neuron from a whole-brain recording of a moving worm expressing both GFP and RFP. Both fluorophores are activity-independent, yet we observe large highly correlated fluctuations in the two-channel fluorescence. (C) Top, Body curvature. Bottom, fluorescence of a single neuron taken from a whole-brain recording in a moving worm expressing both GCaMP and RFP. The two channels still show correlated fluctuations despite the activity dependence of GCaMP. (D) Histogram of the Pearson correlation coefficient squared between the red and the green channel for a dataset of whole-brain recordings from 10 moving GFP, RFP control worms. (E) Same as D but in 9 GCaMP, RFP worms.

nel and green channel fluorescence from the same neuron. We assume that **r** is the sum of a latent motion artifact time series **m** and additive Gaussian white noise ϵ_r. Similarly, we assume the green channel measurements **g**, arise as the sum of the same latent motion artifact **m**, a la-

tent time series of neural activity **a**, and Gaussian white noise ϵ_g. To obtain a complete generative model, we assume Gaussian process priors over **a** and **m**, which specify that they evolve smoothly over time. Formally the model can be written as:

$$r = 1 + m + \epsilon_r \tag{6.1}$$

$$g = a + m + \epsilon_g \tag{6.2}$$

$$m \sim \mathcal{N}(0, \Sigma_m) \tag{6.3}$$

$$a \sim \mathcal{N}(1, \Sigma_{\mathrm{a}}) \tag{6.4}$$

$$\epsilon_{rg} \sim \mathcal{N}\left(0, \sigma_{r,g}^2 I\right) \tag{6.5}$$

were **1** denotes a vector of 1s, $\mathcal{N}(\cdot, \cdot)$ denotes a multivariate Gaussian distribution, and σ_r^2 and σ_g^2 denote the variance of the additive Gaussian noise in the red and green channels, respectively. The latent time series of the motion artifact m and neural activity a are both governed by Gaussian process priors, with prior covariances Σ_a and Σ_m. The prior covariances are each parameterized by a pair of hyperparameters, (σ_a^2, τ_a) and (σ_m^2, τ_m), where σ^2 denotes the prior variance or amplitude, and τ denotes the temporal length scale, controlling the degree of smoothness. We use a stationary radial basis function (RBF) covariance function, such that the covariance between two time points depends only on the temporal separation between them: $\Sigma = \sigma^2 \exp(-\frac{(t-t')^2}{2\sigma^2})$. Given this model structure, data r and g, parameters a and m, and hyperparameters $\theta(\sigma_a^2, \tau_a, \sigma_m^2, \tau_m \sigma_r^2, \sigma_g^2)$, we fit the model in two steps. We first optimize the hyperparameters θ by maximizing the marginal likelihood $P(r, g|\theta)$. We then compute the maximum a posteriori (MAP) estimates for a and m given the optimized hyperparameters $\hat{\theta}$ by maximizing the posterior $P(a, m|r, g, \hat{\theta})$.

The marginal likelihood is obtained by integrating the joint distribution of data and latents over a and m, which can be computed in closed form:

$$P(\boldsymbol{r}, \boldsymbol{g} \mid \theta) = \int P(\boldsymbol{r}, \boldsymbol{g}, \boldsymbol{a}, \boldsymbol{m} \mid \theta) da dm$$
$$= \mathcal{N}\left(\begin{bmatrix} \mathbf{1} \\ \mathbf{1} \end{bmatrix}, \begin{bmatrix} \boldsymbol{\Sigma}_m + \sigma_r^2 \boldsymbol{I} & \boldsymbol{\Sigma}_m \\ \boldsymbol{\Sigma}_m & \boldsymbol{\Sigma}_a + \boldsymbol{\Sigma}_m + \sigma_g^2 \boldsymbol{I} \end{bmatrix} \right) \tag{6.6}$$

This quantity is the likelihood function for hyperparameters, which can be optimized numerically to provide a maximum (marginal) likelihood estimate: $\hat{\theta} = \text{argmax}_\theta P(\mathbf{r}, \mathbf{g} \mid \theta)$

By Bayes rule, the posterior is proportion to the product of likelihood and prior:

$$P(\boldsymbol{a}, \boldsymbol{m} \mid \boldsymbol{r}, \boldsymbol{g}, \hat{\theta}) \propto P(\boldsymbol{r}, \boldsymbol{g} \mid \boldsymbol{a}, \boldsymbol{m}, \hat{\theta}) P(\boldsymbol{a}, \boldsymbol{m} \mid \hat{\theta}) \tag{6.7}$$
$$= \mathcal{N}\left(\begin{bmatrix} 1 + \boldsymbol{m} \\ \boldsymbol{a} + \boldsymbol{m} \end{bmatrix}, \begin{bmatrix} \sigma_r^2 \boldsymbol{I} & \mathbf{0} \\ \mathbf{0} & \sigma_g^2 \boldsymbol{I} \end{bmatrix} \right) \mathcal{N}\left(\begin{bmatrix} 1 \\ \mathbf{0} \end{bmatrix}, \begin{bmatrix} \boldsymbol{\Sigma}_a & \mathbf{0} \\ \mathbf{0} & \boldsymbol{\Sigma}_m \end{bmatrix} \right) \tag{6.8}$$

For this model, where prior and likelihood are both Gaussian, the posterior distribution is also Gaussian, with mean and covariance that can be expressed in closed form:

$$P(\boldsymbol{a}, \boldsymbol{m} \mid \boldsymbol{r}, \boldsymbol{g}, \hat{\theta}) = \mathcal{N}(\boldsymbol{v}, \boldsymbol{\Lambda}) \tag{6.9}$$

where the covariance $\boldsymbol{\Lambda}$ and mean $\boldsymbol{v}[\hat{\boldsymbol{a}}, \hat{\boldsymbol{m}}]^T$, which corresponds to the MAP estimates for $\hat{\boldsymbol{a}}$ and $\hat{\boldsymbol{m}}$ are given by:

$$\boldsymbol{\Lambda} = \left(\boldsymbol{B}^\top \boldsymbol{C}_{\boldsymbol{n}}^{-1} \boldsymbol{B} + \boldsymbol{C}_{\boldsymbol{x}}^{-1} \right)^{-1}$$
$$\begin{bmatrix} \hat{\boldsymbol{a}} \\ \hat{\boldsymbol{m}} \end{bmatrix} = \boldsymbol{\Lambda} \boldsymbol{B}^\top \boldsymbol{C}_{\boldsymbol{n}}^{-1} \begin{bmatrix} \boldsymbol{r} - 1 \\ \boldsymbol{g} - 1 \end{bmatrix} + \begin{bmatrix} 1 \\ 0 \end{bmatrix} \tag{6.10}$$
$$\text{with } \boldsymbol{B} = \begin{bmatrix} 0 & \boldsymbol{I} \\ \boldsymbol{I} & \boldsymbol{I} \end{bmatrix}, \boldsymbol{C}_n = \begin{bmatrix} \sigma_r^2 \boldsymbol{I} & 0 \\ 0 & \sigma_g^2 \boldsymbol{I} \end{bmatrix}, \boldsymbol{C}_x = \begin{bmatrix} \boldsymbol{\Sigma}_a & 0 \\ 0 & \boldsymbol{\Sigma}_m \end{bmatrix}.$$

Note that the implementation of TMAC relies on a Fourier-domain representation which diagonalizes the prior covariances, substantially reducing computation time[5].

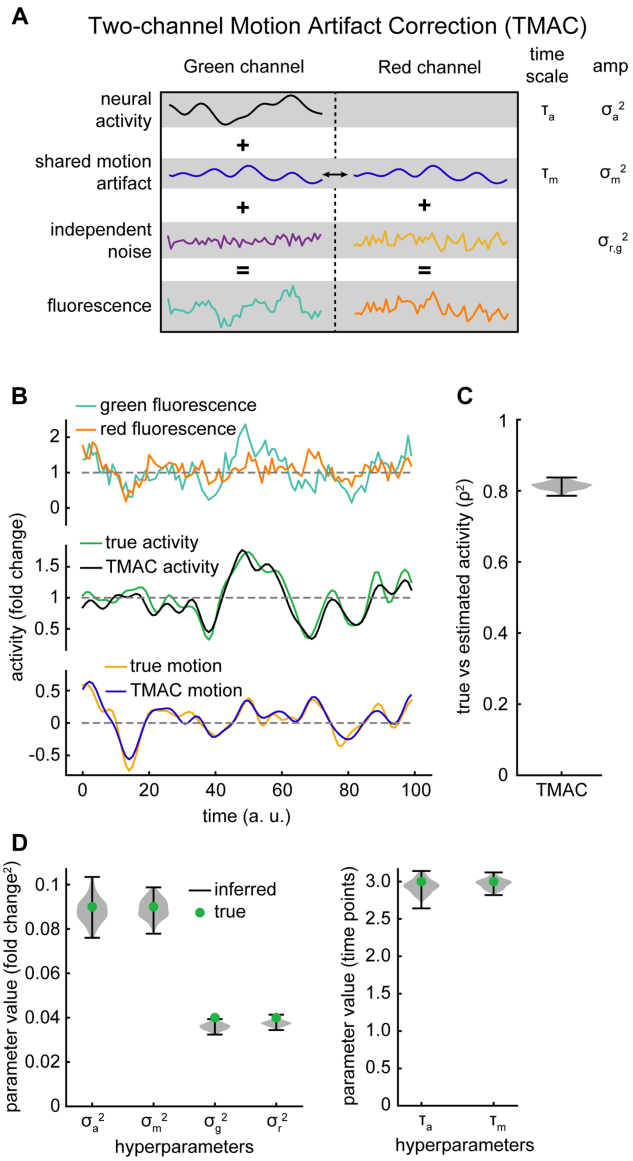

Figure 6.2: **TMAC infers activity and hyperparameters from synthetic data.** (A) Diagram of the structure of TMAC. The green channel is modeled as the sum of the calcium activity, motion artifact, and independent gaussian noise. The red channel is modeled as the sum of the motion artifact and independent gaussian noise. The motion artifact is shared between the two channels. (B) Top: fluorescence from a synthetic green and red channel. Middle and bottom: inferred activity and motion compared with the true activity and motion. (C) Correlation squared between estimated activity from TMAC and true activity over many synthetic datasets. (D) Violin plot of inferred and true parameter values when fitting TMAC

6.1.3 Results for behavioral decoding

To first demonstrate that our model correctly infers ground truth activity and variances, we used TMAC to generate a synthetic dataset (Fig 6.2B top) with 5000 time points and a correlation time

scale which roughly matches experimental datasets from Fig 1, such that 6 time points approximately corresponds to one second. We then used TMAC to infer both the activity and motion artifact from this synthetic dataset (Fig 2B middle and bottom). We find a good correspondence between the inferred activity from TMAC and the true activity (Fig 6.1C). The model also accurately infers the hyperparameters of the model (Fig 6.1D).

We next sought to evaluate TMAC on experimentally acquired neural population calcium imaging recordings in moving worms from[97]. The data was preprocessed to correct for missing values and photobleaching for each channel, following methods of[97]. In result, the hyperparameters learned by TMAC provide a convenient method to characterize the recordings. We identified a putative high signal-to-noise neuron, by finding a neuron that the model associated with a high ratio of activity variance (σ_a^2) to motion and noise variances ($\sigma_{m,r,g}^2$) (Fig 6.3A middle). We also identified a putative low signal-to-noise neuron by finding a neuron that the model associated with high ratio of motion artifact (σ_m^2) to activity and noise ($\sigma_{a,r,g}^2$) (Fig 6.3A bottom). In the case where TMAC estimates high signal-to-noise, TMAC predicts that a smoothed version of the green channel represents true calcium activity (Fig 3A middle). In the case where TMAC estimates that most fluorescent fluctuations are due to motion, TMAC deviates from the measured fluorescent intensities and returns a flatter inferred activity (Fig 3A bottom).

We next wanted to evaluate TMAC's performance on removing motion artifacts. Since we lack access to the ground truth calcium activity, we considered the problem of decoding animal curvature from neural activity (Fig 6.3B). To decode the behavior of each animal we followed the same procedure from Hallinen et al., 2021[97]. Namely, we trained a linear decoder with an L2 penalty on the weights determined using multifold cross validation. Behavior was predicted at each time point from the motion corrected activity, and all reported decoding accuracy scores are from data held out from the center 40% of the recording (Fig 6.3B).

Using decodability as a metric for motion correction performance is challenging because motion artifacts may also provide information about the target behavior. In some regimes, successfully removing motion artifacts could in principle reduce decodability rather than improve it. We

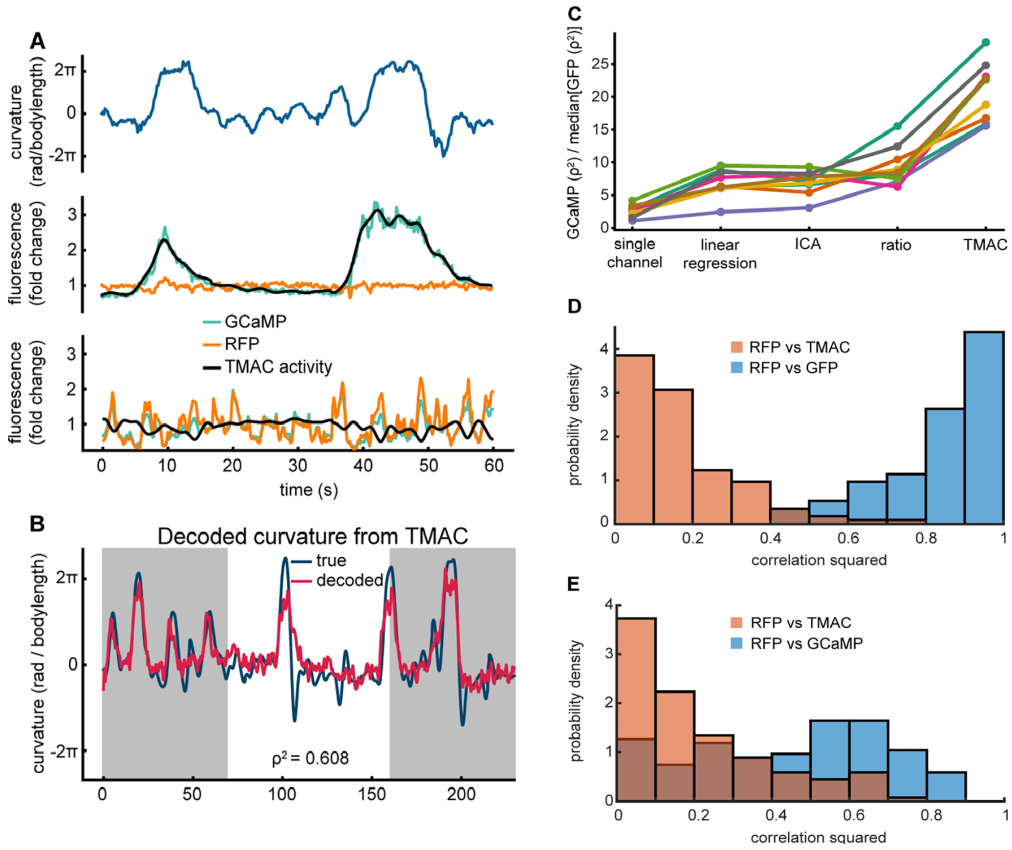

Figure 6.3: **TMAC reduces decodable motion artifacts in experimental data.** (A) Top, animal body curvature over time. Middle, GCaMP and RFP fluorescence from a neuron that TMAC estimates to have high signal to noise, recorded from a behaving worm. Bottom, GCaMP and RFP fluorescence from a different neuron in that same recording that TMAC estimates to have large motion artifacts. (B) Time trace of animal curvature and predicted behavior, decoded from activity inferred by TMAC in a GCaMP worm. Gray shaded regions were used to train the decoder, white region was held out and used to evaluate decoding performance. (C) Ratio of decoding accuracy (ρ^2) when decoding GCaMP divided by the median accuracy for a GFP worm across different models. (D) Histogram over all neurons of correlation squared between RFP and activity inferred by TMAC from a GFP worm. RFP vs GFP data the same as in Fig 6.1D,E. (E) Same as F but in a GCaMP worm.

therefore evaluated motion correction performance by inspecting the decodability of two types of recordings; one that had both activity-dependent and activity-independent fluorophores (GCaMP and RFP), and a control worm that had two activity-independent fluorophores (GFP and RFP), both from Hallinen et al., 2021[97]. Although the GFP control worm contains no neural signal, the worm's behavior can still be decoded, albeit poorly, by relying on behavior-related information in the motion artifact. We reasoned that a successful motion correction algorithm should reduce

the decodability of motion artifacts while retaining the ability to decode from neural signals. We therefore define a new metric that evaluates motion correction performance as the ratio of its decodability of recordings from GCaMP animals to its decodability of recordings from GFP animals (Fig 6.3C).

Using this metric, we further compare five models used for motion correction in two-channel imaging. The models are: single channel GCaMP which is just the green channel fluorescence with no motion correction, the ratio model which is the green channel divided by the red channel, a linear regression model[97], an ICA based approach[221], and TMAC. The linear regression model finds the best linear fit of the red channel to the green channel and then subtracts off that best fit from the green channel. The ICA approach performs ICA using the red and green channel as inputs and returns the independent component least correlated to the red fluorescence.

The activity fit by TMAC is $\sim 20\times$ more effective at predicting curvature from a GCaMP expressing worm than a control worm (Fig 6.3C). We calculated the correlation between the activity inferred by TMAC and the red channel and they did not correlate strongly, suggesting that the model is successful at removing motion artifacts common to the two channels (Fig 6.3D,E). For completeness, we also compare each method on our synthetic dataset.

6.1.4 Discussion

Motion artifacts are prevalent in recordings of behaving animals (Fig 6.1B,C). These artifacts can appear as signals of interest, reducing the interpretability of the data. Here we presented TMAC, a model which infers the latent neural activity without these artifacts by leveraging information in two-channel calcium imaging recordings. We demonstrate that TMAC substantially reduces decodable motion artifacts in experimental data (Fig 6.3A) and outperforms four alternatives (Fig 6.3C).

The model we propose uses an additive interaction between the fluorescence transients caused by motion artifact and the fluorescence transients caused by neural activity. It is unclear in experimental data what the true interaction between artifact and activity is. However, even if the

true interaction is not additive, TMAC will perform well under model mismatch because it can be thought of as a linear approximation of the true interaction between a and m. Consider the case where the activity and motion interact multiplicatively. Multiplicative interactions have linear components when the two variables have nonzero means. The linear component of the multiplicative interaction is clear by writing the equation for the green channel with no noise: $g = (a+1)(m+1) = am+a+m+1$. We can be sure that the a and m have a mean (in this case 1), because fluorescent data is nonnegative. Furthermore, the am term itself can be approximated by $a + m$ as long as a and m are small. If the approximation of a and m as Gaussian is valid, we can be sure that the am and higher order terms are small. For these reasons, while TMAC is a linear approximation to the true interaction between a and m, it is an approximation that is highly accurate.

In this work we demonstrated TMAC's ability to remove decodable motion artifacts from calcium induced fluorescence. However, TMAC will remove motion artifacts in any type of two-channel imaging, so long as one channel is activity-independent and both channels share the same motion artifact component. TMAC could therefore be applied to a wide range of two-channel imaging modalities such as fiber photometry when using an isosbestic wavelength[148] and two-channel two-photon imaging[181].

6.2 Adaptive subunit model for predictive dynamics in mouse visual cortex

6.2.1 Introduction

Rapid detection and processing of stimulus novelty are key elements of adaptive behavior. At the level of human psychophysics, stimulus novelty can enhance salience and capture attention[222], while familiarity can speed visual search[256]. Novelty also affects aversive conditioning and fear conditioning[200]. In human brain imaging, novel stimuli have been shown to generate the mismatch negativity[182] while repeated stimuli lead to repetition suppression[94]. Explicit representation of novelty has been shown at higher stages of the sensory hierarchy, such as in the hippocampus[215] and inferotemporal cortex[174]. Lower in sensory hierarchies, the representation of

novelty can be enhanced by stimulus-specific adaptation[106] as well as by gain control[42]. Novelty signals are also prominently present in midbrain dopamine neurons[223].

Explicit representation of stimulus novelty is also related to theories of predictive coding, in which neural circuits carry out computations that emphasize novel or surprising information. Theories of predictive coding have had a long history, starting with ideas about how the receptive field structure of retinal ganglion cells more efficiently encodes natural visual scenes by removing redundant data[232] and including the idea that active adaptation may aid in this process[9,106,44]. Theories of predictive coding in the neocortex have typically focused on the idea that feedback from higher cortical areas encodes a prediction about lower-level sensory data that is subtracted from the lower-level representation, so that the signals traveling up the cortical hierarchy represent surprise or novelty[204]. However, a recent study failed to find these signatures of predictive coding[231].

Here, we investigate novelty processing in the mouse primary visual cortex. We repeatedly presented a set of images, each composed of a random superposition of Gabor functions, and then occasionally presented novel images drawn from the same ensemble. Using two-photon imaging of the calcium sensor GCaMP6f to measure neural activity in layer 2/3 of awake, head-fixed mice[72], we found that the majority of neurons exhibited excess activity in response to a novel image. This distinction between novel versus familiar images was quickly reached, emerging with a time constant of 2.6 ± 0.9 s. Similarly, when we began presenting a new set of images, a majority of the neurons exhibited elevated firing that relaxed to a steady state with a time constant of 1.4 ± 0.4 s. When we presented novel images within larger image sets, the amplitude of novelty response decreased, defining a capacity of the system to encode $\sim$15 familiar images. All of these findings could be explained qualitatively using an adaptive subunit model in which neurons presynaptic to a recorded neuron have both individual tuning to visual stimuli and adaptive gain control.

6.2.2 Novelty response in mouse visual cortex

Neural imaging and visual stimulation

In order to explore how the primary visual cortex encodes novelty, we used two-photon calcium fluorescence imaging in mice that were awake and head-fixed but free to move on a styrofoam ball placed below them (Fig 6.4A)[72]. Animals were transgenics from a Thy1 line (GP5.3; Janelia) that expressed the protein GCaMP6f[48] in excitatory neurons. In order to locate the primary visual cortex, we first carried out large-scale brain mapping with a one-photon macroscope using drifting bars (Fig 6.4B). We then selected a field of view in V1 and imaged at cellular resolution in layer 2/3. We manually identified regions of interest (ROIs) corresponding to cell bodies having a "halo" pattern of fluorescence indicating expression mostly in the cytoplasm (Fig 6.4C). After averaging across pixels in an ROI to obtain the average fluorescence F, we constructed the time course of the fractional change in fluorescence ($\Delta F/F$). The activity in single ROIs typically exhibited sparse events on a background (Fig 6.4E).

In this two-photon setup a toroidal projection screen was placed around the mouse, which allowed us to display visual stimuli. We chose images that consisted of a random superposition of 100 Gabor functions on a gray background. Those Gabor functions were drawn from a distribution that matched the range of receptive field parameters found in the literature for mice (Fig 6.4D)[186]. Images constructed in this way were statistically similar in global light level, contrast, and spatial scale. We then formed sets of three or four of those images, each displayed for 250 or 300 ms, depending on the specific experiment. In order to create a familiar stimulus, we repeated this image set many times without blanks in between. Occasionally, we substituted one image with a novel image drawn from the same distribution (Fig 6.4D).

In the novelty experiment we repeated a set of four images, each with a duration of 250 ms. A single image set was repeated for $\sim$10 min while novel images were randomly substituted every $\sim$6 s. All novel images were unique images drawn from the same ensemble of random Gabor images that the repeated image set was drawn from. Neurons generally showed a noisy, weakly

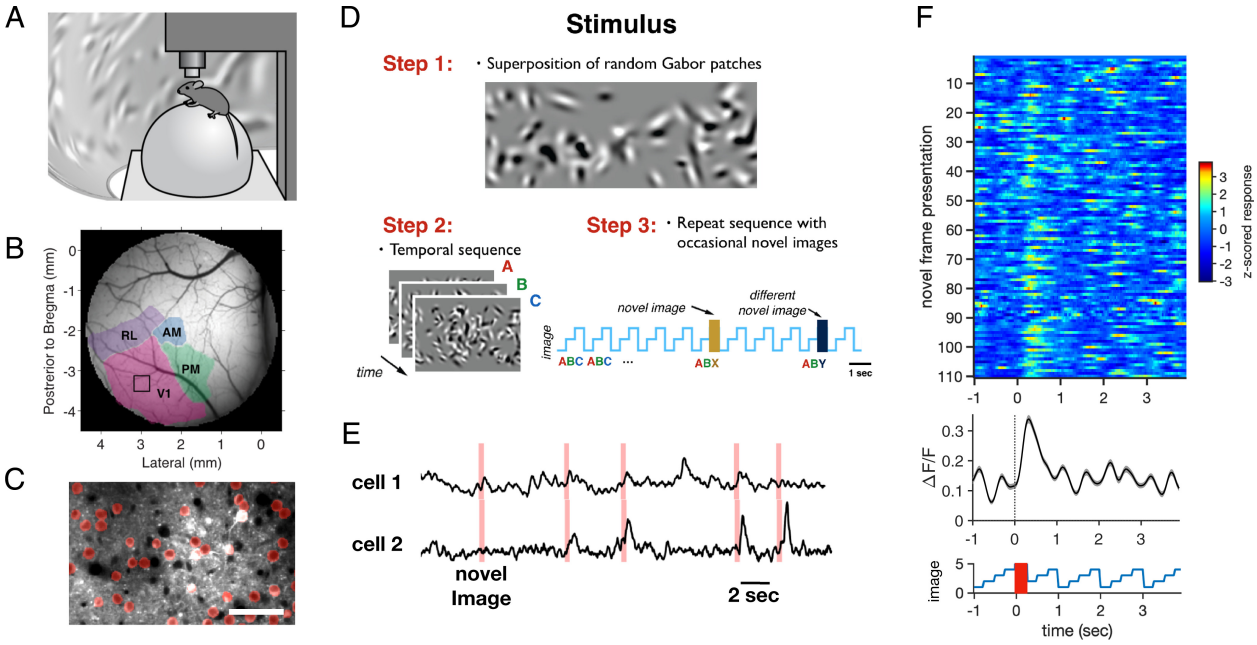

Figure 6.4: **Measuring a novelty response.** (A) Awake mice were head-fixed and placed on an air-suspended styrofoam ball. Visual stimuli were projected on a toroidal screen surrounding the animal. Neural activity was recorded with a two-photon microscope. (B) Wide-field image of visual areas, as determined by one-photon fluorescence measurements (see Materials and Methods). The black square within area V1 shows the size of the field of view for the two-photon microscope. (C) Portion of a field of view taken by the two-photon microscope with ROIs shown in red. (Scale bar, 50 µm.) (D) Stimulus design. Step 1: Images were constructed from a superposition of randomly chosen Gabor functions. Step 2: A set of different images was formed and presented repeatedly in the same order; image sets are represented by plotting the image index versus time. Step 3: Occasionally, an image was substituted by unique novel images drawn from the same image ensemble. (E) Example activity traces with the times of novel image presentations shown in red; all novel images were unique. (F, Top) Matrix of trial-by-trial responses of an example cell to novel images. (F, Middle) Activity averaged across trials. (F, Bottom) Repeated sequence with the time of novel images shown in red.

modulated response to the repeated image set. However, when a novel image was substituted, neurons showed a large, brief response (Fig 6.4E,F). Some neurons responded to a majority of the novel images (example cell in Fig 6.4F), indicating that their novelty response was somewhat unspecific to the particular locations and orientations of the Gabor patches overlapping the cell's receptive field.

Emergence of the Novelty Response.

Across the neural population, we subtracted the activity triggered by the repeated image set from that triggered by the presentation of a novel image. We found that a large majority of the 1,134 neural responses in our sample showed, on average, excess activity to a novel image. Through further control analysis, we verified that the novelty response are not a product of neuropil input from external regions, nor a product of behavioral surprise such as pupil modulation.

Conceptually, a neuron cannot exhibit a differential response to a novel image until the original set of images is repeated at least once. We wanted to see how many repetitions of a sequence were necessary so that neurons would show an elevated response to a novel image. To this end, we designed a variable repetition experiment in which we displayed sets with three images of 300-ms duration. We presented a given image set in a block with either L = 1, 3, 8, 18, or 38 repetitions of the image set before showing a novel image (Fig 6.5A). Each block contained a unique image set and a unique choice of repetitions, L. Blocks followed each other seamlessly with no intertrial period or blank frame.

We found that the novelty response emerged rapidly. Significant excess activity was observed in the population after as few as L = 1 repeats of a new image set (Fig 6.5A, Bottom). The effect increased with more repetitions and saturated at L $\sim$ 20 (Fig 6.5A, Top). Fitting an exponential curve to the effect amplitude versus number of repetitions revealed a time constant of τ_{emerge} = 2.6 $\pm$ 0.9 s, or alternatively 2.9 $\pm$ 1.0 repeats, for the emergence of the novelty response.

The Transient Response

We also noticed that neurons exhibited elevated activity when we began presenting a new image set (Fig 6.5A, elevated activity near t = 0). This transient response adapted strongly as a given image set was repeated, quickly reaching a steady state. This adaptation process had a time constant of τ_{adapt} = 1.4 $\pm$ 0.4 s in the variable repetition experiment. A transient response is also evident after the occurrence of a novel image, because then a new image set was presented (Fig 6.5A, elevated activity near t = 40 s). Similar to the case of the novelty response, a

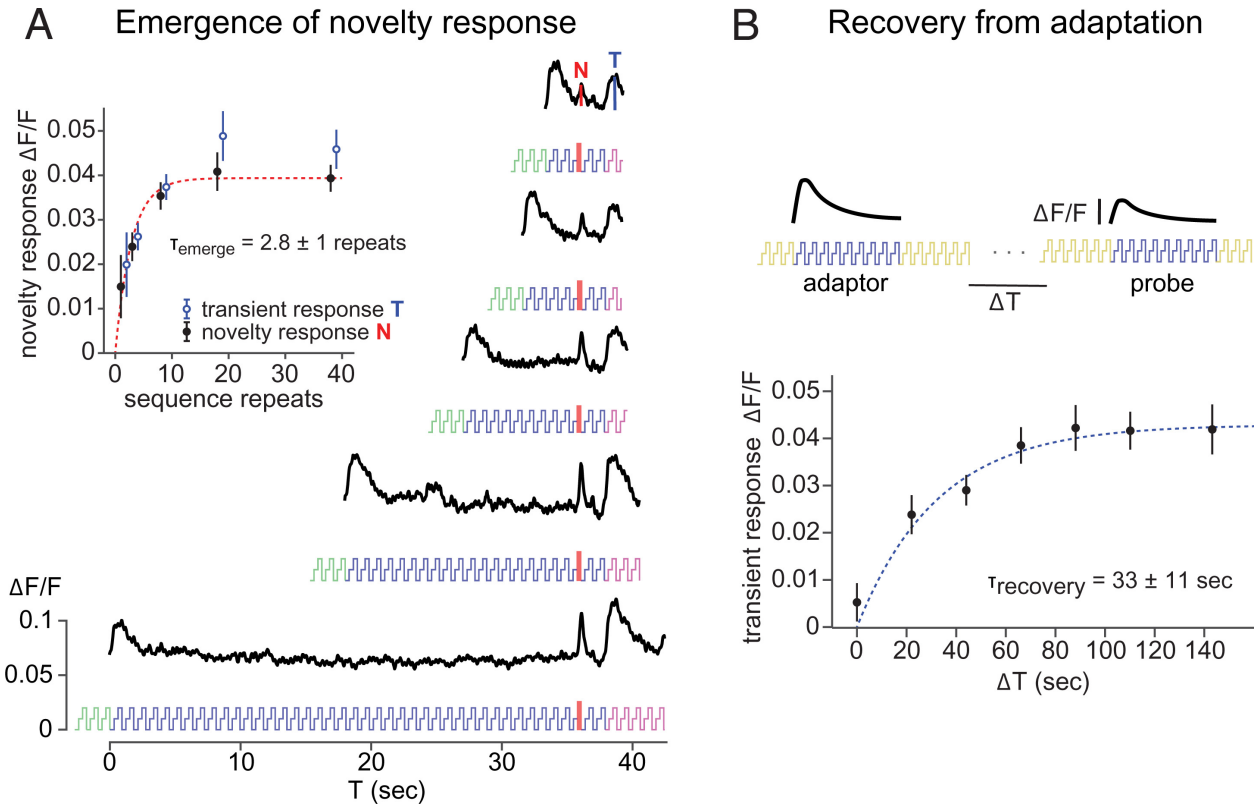

Figure 6.5: **Novelty responses emerge quickly within a repeated sequence.** (A) In the variable repetition experiment, we formed blocks containing a new, randomly chosen image set (shown in blue; previous block shown in green). This unique image set was presented once and then immediately followed by L repetitions. Next, an image set was presented with a unique novel image substituted (shown by red bars). Finally, two more image sets were presented before a new block began (shown in purple). (Bottom) Neural activity averaged over the entire population for different choices of L (curves offset for clarity). (Top) Amplitude of the novelty response versus the number of repetitions, L. Dotted line: Exponential curve fit; error bars are SEM over n = 5 mice. (B, Top) Design of the repeated image set experiment. A given image set was repeated in a block until adaptation reached steady-state (blue). The same image set was presented again after a variable delay during which other image sets were presented (green). (B, Bottom) Transient response amplitude versus delay with an exponential curve fit (red); error bars are SEM over n = 5 mice.

given image set must be repeated before a response can emerge from the transition to a novel sequence. We measured the amplitude of the transient response elicited by the transition to a novel sequence as a function of the number of repeats of the preceding image set, L, and found that the emergence of the transient response closely paralleled that of the novelty response (Fig 6.5A, Top).

How long does adaptation to familiar images persist? All sequences used in the variable repeti-

tion experiment were different. To answer this question, we designed a repeated image set experiment that adapted neurons to a given image set and then, after various lags, presented the same sequence again (Fig 6.5B). After a sufficiently long recovery time, we expect the transient response to recover to its maximum amplitude, similar to what we found for the emergence of the novelty response. To make sure we did not create transition effects due to the lack of neural stimulation during the recovery period, we filled this period with other image sets.

We measured the amplitude of the transient response as the difference between the response to the last few repeats of the previous sequence and the peak calcium fluorescence of the response elicited by the new sequence. The previous sequence was repeated sufficiently often so neurons had reached steady-state activity. When we plotted the amplitude of the transient response as a function of the time since the same image set was presented, we found an exponential rise with a time constant of $\tau_{recovery}$ = 33 ± 11 s (Fig 6.5B). The amplitude at a time interval of zero was set to zero, because in this case there was no pause between the adaptor stimulus and the probe, and thus no transition that could elicit an excess response.

The Novelty Response as a Probe of Capacity

In order for a neural circuit to distinguish novel from familiar images, it must maintain some representation of familiar images. To explore the limitations of this representation, we designed a variable image number experiment, in which we varied the number of images in the set. Similar to the variable repetition experiment, we formed adaptation blocks. In each block, i, we randomly chose the number of images to be S_i = 3, 6, 9, or 12 and then randomly generated a new image set of this size (Fig 6.6A). In all blocks, the given image set was presented for 17 times before a novel image was introduced in the 18th trial. This allowed us to measure the amplitude of the novelty response after adapting to image sets of different size.

We found that the novelty response systematically decreased as the number of images, S, increased (Fig 6.6B,C). This result is consistent with the intuition that as the neural circuit encodes more images, then yet another image will seem less novel. Although the novelty response de-

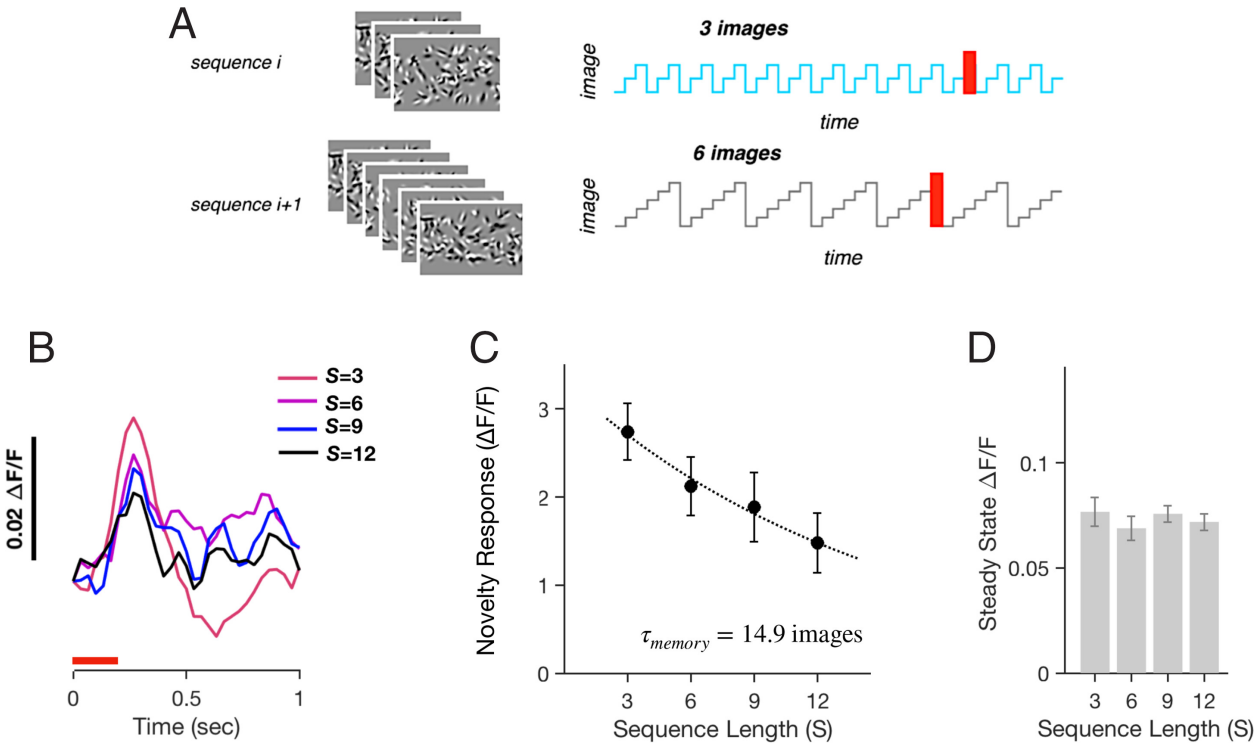

Figure 6.6: **Novelty response from within larger image sets.** (A) Image sets with either S = 3, 6, 9, or 12 images were presented 17 times before one of the images was substituted by a unique novel image. All images were presented for 300 ms each. (B) Trial-averaged response of the neural population to novel images during image sets of different size, S (shown in color). (C) Amplitude of the novelty response versus image set size, S. Dotted line is an exponential curve fit; error bars are SEM for n = 5 mice. (D) Population-averaged steady-state neural activity versus image set size, S; error bars are SEM for n = 5 mice.

creased with larger image sets, it was still present for sets of S = 12 images. This indicates that under our experimental conditions the capacity was at least as large as 12 images. Another way of characterizing the capacity is to examine the trend of novelty response amplitude versus size of the image set, S. Here, we found an exponential decrease with a decay constant of τ_{encode} = 15 images (Fig 6.6C). This decay constant can be seen as a measure of the capacity to store familiar images. One way that the neural circuit could generate a smaller novelty response could be by having a higher steady-state response for larger image sets but always the same response level to novel images. In this case, there would be less dynamic range to generate a large novelty response. However, we found that this was not the case (Fig 6.6D). Because the adaptation process reached the same steady-state activity as a function of the image set size, S, we interpret

the decrease in the novelty response amplitude as arising from limitations in the representation of familiar images. The constancy of steady-state neuron activity across image sets of different sizes is consistent with previous observations that cortical adaptation helps to achieve a form of population homeostasis[17].

6.2.3 The adaptive subunit model

In order to explain and unify our observations, we formulated an adaptive subunit model. The idea is that a given neuron that we recorded from in layer 2/3 received inputs from many other cortical neurons (subunits) that need not have been observed and that each of these presynaptic neurons had its own individual tuning to visual stimuli along with a cell-intrinsic mechanism of adaptation (Fig 6.7A). Because each presynaptic neuron had different tuning to stimuli, there was always a subset of those subunits that had high sensitivity to any novel image. As that image was repeated, those subunits adapted, giving rise to lower activity in the recorded neuron. At the same time, subunits that were poorly stimulated by a given set of images will recover their full sensitivity, so that excess activity was observed in response to a different image.

More formally, in the adaptive subunit model, there was a population of K neurons, denoted by k, that each received input from a set of M subunits, denoted by j, via synapses of strength, q_{jk}. For each synapse, we chose $q_{jk} = q_0$ with probability q and $q_{jk} = 0$ otherwise. Each subunit was activated by images i according to synaptic weights, p_{ij}. We chose these weights from a log-normal distribution having a median of p_0 and a log-standard deviation of σ_p. The choice of such connectivity matches neural activity, synaptic weights, and sparseness observed in cortical neurons. The activity of each subunit, denoted by $x_j(t)$, involved leaky summation of its inputs that was scaled by its own gain factor, $\alpha_j(t)$:

$$\tau_{cell}\frac{dx_j}{dt} = -x_j + \alpha_j \Sigma_i^{images} p_{ij} I_i(t) \tag{6.11}$$

where the input function $I_i(t) = 1$ at time t if image i is present. The gain, α_j, had an equilibrium

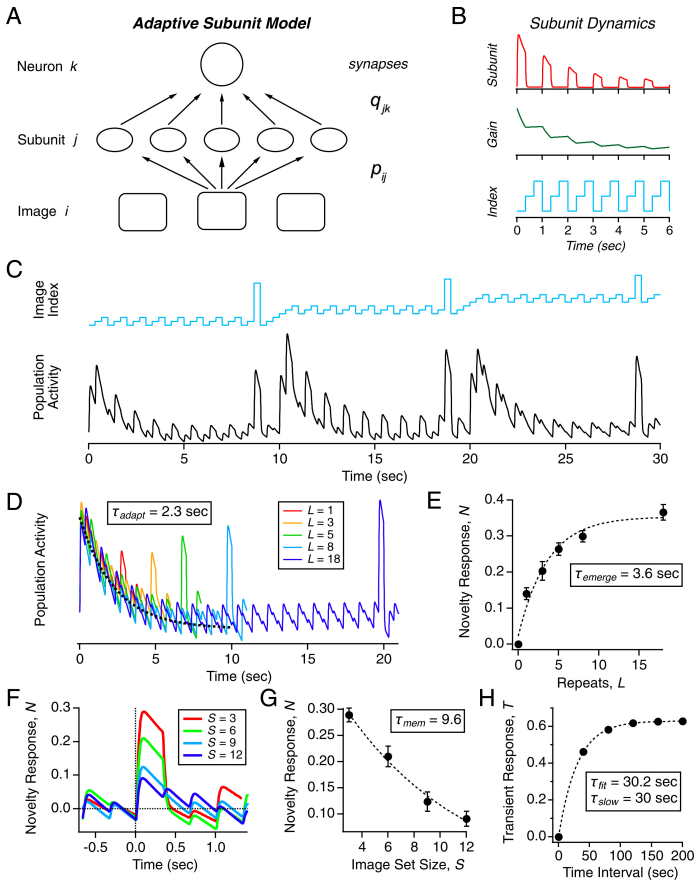

Figure 6.7: **Adaptive subunit model.** (A) Schematic of the model: Images, i, provide input to M subunits, j, via synapses, p_{ij}, and subunits provide input to N neurons, k, via synapses, q_{jk}. (B) Illustration of subunit dynamics: A sequence of images (Bottom) drives sparse responses in a subunit (Top) that decrease over time due to a reduction in gain (Middle) during each image presentation. (C) Population-averaged firing rate (Bottom) in neurons during presentation of three different image sets (Top). (D) Population-averaged firing rate during adaptation blocks with different numbers of repetitions of the same image set before the presentation of a novel image, L; dotted line is an exponential curve fit. (E) Amplitude of the novelty response, N, versus number of repetitions, L; dotted line is an exponential curve fit. (F) Population-averaged firing rate following presentation of a novel image during adaptation blocks with different sizes of image set, S. (G) Amplitude of the novelty response, N, versus size of the image set, S; dotted line is an exponential curve fit. (H) Amplitude of the transient response, T, plotted as a function of the time interval between image set presentations; dotted line is an exponential curve fit that defines the recovery time, τrecovery. All error bars are SEM calculated over n = 4 instantiations.

value of unity, was reduced below this value by the input, and recovered with a time constant given by τ_{slow}:

$$\tau_{slow}\frac{d\alpha_j}{dt} = -(\alpha_j - 1) - \alpha_j \Sigma_i^{images} p_{ij} I_i(t) \tag{6.12}$$

The activity of each neuron, denoted by y_k, was driven by an integration of its inputs, u_k, and had an analogous adaptation process with gain variable, β_k:

$$\begin{aligned} u_k(t) &= \sum_{j=1}^{M} q_{jk} x_j(t) \\ \tau_{\text{cell}} \frac{dy_k(t)}{dt} &= -y_k(t) + \beta_k(t) u_k(t) \\ \tau_{\text{slow}} \frac{d\beta_k(t)}{dt} &= -\left(\beta_k(t) - 1\right) - \beta_k(t) u_k(t) \end{aligned} \tag{6.13}$$

In this model, subunits have a cellular time constant, τ_{cell}, chosen to be 20 ms, along with a time scale of adaptation, τ_{slow}, chosen to be 30 s to match the measured recovery time constant (Fig 6.5B). The most likely mechanism underlying this gain control is an adaptive ion channel[30], such as the sodium-dependent potassium current found in V1 neurons, which has a time constant of 10 to 30 s[218]. Other mechanisms of spike frequency adaptation can have similar time scales and effects[261].

This process of adaptation causes the subunit's gain and hence its response to repeated presentations of the same image to decrease (Fig 6.7B). Notice that the induction of adaptation depends on τ_{slow} as well as both strength parameters, p_0 and q_0. However, because of the sparseness of subunit activation, recovery is determined primarily by τ_{slow}. Recorded cells, k, integrate over all of their subunits with randomly chosen synaptic weights, $q_j k$, and with the same cellular time constant, τ_{cell}. The model contains many recorded cells and averages over their activity to compare with population-averaged experimental data.

Population activity in the model starts out high during the presentation of a new image set and then adapts down to a lower steady state (Fig 6.7C). When a novel image is presented, there is excess activity (Fig 6.7C). This excess activity decays to a steady-state baseline with a time constant, τ_{adapt} = 2.3 s (Fig 6.7D), that agrees with experimental data (Fig 6.5). All new image sets and novel images evoke similar levels of increased activity. When we increase the number of image-set repetitions, L, before the presentation of a novel image, the amplitude of the novelty response, N, increases (Fig 6.7D) with a roughly exponential dependence on L (Fig 6.7E). This

trend is dominated by the fact that the transient response had not fully adapted away for small L. In addition, there is a modest recovery of sensitivity of the subunits most strongly driven by the novel stimulus that leads to a small increase in the absolute amplitude of the novelty response. This recovery occurs because most of the subunits driven by the novel image were not activated by the image set, allowing their gain to recover.

When we increase the size of the image set, S, the amplitude of the novelty response, N, decreases (Fig 6.7F) with a roughly exponential dependence (Fig 6.7G). This phenomenon results from cross-adaptation. Subunits tended to have a strong response to a small fraction of all the images and weak responses to the other images. With a larger image set, the weak responses to these other images induces increasingly more adaptation. Because of this accumulated weak adaptation, subunits had a smaller gain when the novel image is presented. Finally, the transient response for a given image set, T, recovers with a time constant of τ_{fit} = 30.2 s, consistent with our measurements (Fig 6.5B). Taken together, the adaptive subunit model provides a qualitative match for all of our experimental observations with 6 free parameters. Only three of these parameters were fully adjusted to match experimental data (p_0, q_0, and τ_{slow}); the others were constrained to reasonable values.

Discussion

In this study we found that the excitatory neurons in layer 2/3 of the mouse primary visual cortex exhibited a distinctive pattern of excess activity when a novel image was presented among a set of familiar images (Fig 6.4). This excess activity was present in a large fraction of neurons and across different sets of images. The observation did not result from changes in the animals running, pupil, or eye movements.

The process that determined the novelty of an image was rapid, as the differential response between novel and familiar images could be observed after only one repetition of an image set (Fig 6.5A). The characteristic time scale over which the novelty response emerged was approximately three repeats (Fig 6.5A, Inset). By presenting novel images in the image sets of different

durations, S, we found that the novelty response decayed for larger S (Fig 6.6). This defined a capacity for the system to encode ∼15 familiar images, under our experimental conditions. This representation of familiarity decayed relatively slowly, with a recovery from adaptation on a time scale of ∼30 s (Fig 6.5B).

In order to gain insight into the circuit mechanisms responsible for novelty responses, we formulated an adaptive subunit model (Fig 6.7). In this model, each observed neuron received rectified input from many subunits that each had different tuning for stimuli—a form similar to Hubel and Wiesel's model of the complex cell[107] as well as models of cascaded adaptation[71]. In addition, each subunit had its own gain control. This model could account for all of the main qualitative features of our data, including: 1) excess activity evoked by all novel images, 2) rapid emergence of the novelty response among a new set of familiar images, 3) rapid decay of excess activity as novel images were repeated, 4) decay of the novelty response among larger sets of familiar images, and 5) recovery of full activity with a longer time scale.

While mechanisms of short-term synaptic plasticity, like depression, appear to be nearly ubiquitous in early sensory pathways[31], we did not explicitly include such terms. This was for two reasons: 1) The dominant time scale of synaptic depression in layer 2/3 of V1 is ∼400 ms[249], which is too fast to account for our observations, and 2) the gain control in each subunit adjusts the strength of presynaptic inputs in a manner similar to depression, obviating the need for another term. In our model, the gain control variable in subunits accounted for all adaptation in subcortical pathways as well as adaptation in layer 4 of V1. As such, we used a stronger gain control parameter ($p_0 > q_0$). Breaking up all of these processes into multiple biophysical mechanisms would presumably have allowed for a closer quantitative match to our observations but would have been achieved at the cost of a more complicated model. One strength of our resulting model is that it captures such a wide range of phenomena with a relatively simple structure.

The novelty response can also be thought of as a form of predictive coding, which emphasizes the representation of novel or surprising information. While many theories of predictive coding in the neocortex focus on the role of descending feedback[204], mechanisms of adaptation can also

contribute[258]. The fact that a simple and biophysically common mechanism can give rise to the novelty response suggests that similar processing may be present at many stages of the sensory processing hierarchy. If this operation were repeated at multiple stages, then adaptation would cascade and potentially get stronger, faster, or longer-lasting. If a mechanism like the adaptive subunit model existed at multiple cortical stages, then novelty responses could be present for increasingly complex stimulus features at higher stages of the hierarchy, without needing to invoke further changes, like adaptive synapses[71]. Thus, this simple mechanism could lead to powerful adaptive processing across the cortical hierarchy.

6.3 Coding, dynamics, and inference in nonequilibrium networks

In this section, we included two projects related to nonequilibrium networks. In section 6.3.1, we introduced a framework to analyze thermodynamical properties of an input-driven network. In section 6.3.2, we presented an inference method to recover causal network interaction from spiking networks. Lastly, we discussed relation to the literature and future work in 6.3.3.

6.3.1 Nonequilibrium physics of input-driven networks

Introduction

The Hopfield model provides a framework to describe neural dynamics in terms of an energy function or Lyapunov function for dynamical systems[103,104]. The resulting neural dynamics asymptotically reaches a stationary state known as the attractor in dynamical systems. In terms of thermodynamics, this is the result when the system reaches thermal equilibrium and the probability of activity patterns follows Boltzmann distribution[100,101]. Specifically, these equilibrium solutions are learned through synaptic learning rules and the recall process of learned patterns captures the property of associative memory[100]. In more biological settings, the neural network receives time-varying external input and the synaptic weights may change on the time-scale that is not separable from neural dynamics. This leads to nonequilibrium conditions as the system is constantly driven away from the equilibrium neural state. Stability analysis and Boltzmann distri-

bution for equilibrium systems may not be relevant for analyzing dynamics far-from-equilibrium. Nonequilibrium thermodynamics studies state variables that are away from equilibrium and how macroscopic variables used in equilibrium thermodynamics can be described under such conditions[118,119]. In the past few decades, the relation between nonequilibrium work distribution and equilibrium free energy difference has been theorized[62,118] and supported experimentally[55,35,224]. More specifically, Crooks theory relates the ratio between forward and backward processes with dissipated work, as well as providing a nonequilibrium method for measuring free energy difference between states[62,63]. Fluctuation theory enables thermodynamic characterization of systems that are stochastic and driven by external input or time-varying parameters, such as biological neural networks. Characterizing neural activity in terms of free energy is of interest as it informs kinetics, relates to thermodynamics properties, and further tests for theories of brain computations[199,100,80].

Recently, nonequilibrium statistical mechanics has been applied to input-driven continuous attractors[275,270] and unsupervised learning of restricted Boltzmann machine[216]. The former derives limits of attractor dynamics and memory capacity for a bump attractor model under dynamical input, whereas the later characterizes the learning process of energy models with thermodynamic variables. For supervised learning process, the thermodynamic efficiency of different learning rules have been computed in a stochastic thermodynamics framework[88]. This thread of theoretical work provides better understanding for the nonequilibrium nature of neural dynamics and learning processes. In this project, we extended the picture by applying generalized fluctuation theorem to energy based neural dynamics driven by external input. We demonstrated that the system satisfies Crooks fluctuation theorem (CFT) under dynamic input and heat exchange fluctuation theorem (XFT) in contact with heat reservoir with another temperature.

Settings

We introduced nonequilibrium thermodynamical measurements through fluctuation theorems. We formulated the toy example of input-driven stochastic Hopfield network and specify ther-

modynamics quantities in this model. We measured work performed during a set of stimulation protocol and further extend the method with other parameter choices.

Fluctuation theorems Considering an isothermal process moving from an initial equilibrium state to a final equilibrium state, the second law of thermodynamic suggests that the work performed during this process W and free energy difference between two states ΔF follows $W \geq \Delta F$. Free energy is defined by $F = \langle E \rangle - TS$, where E is the internal energy, T is temperature, and S is the entropy at equilibrium. When the process between two states is nonequilibrium, Jarzynski equality states the average work performed is still related to the equilibrium free energy difference[118]:

$$\langle \exp(-\beta W) \rangle = \exp(-\beta \Delta F), \tag{6.14}$$

where $\beta = 1/k_B T$ is the inverse temperature T and k_B is the Boltzmann constant. This leads to a similar form in statistical mechanics: $\langle W \rangle \geq \Delta F$, according to Jensen's inequality. For reversible process, the equality $\langle W \rangle = \Delta F$ holds, so one can calculate the free energy difference from the average of nonequilibrium work trajectories. However, for irreversible processes with dissipated work $W_{diss} = \langle W \rangle - \Delta F \neq 0$, this method of estimating free energy difference leads to error when the system is far from equilibrium. With work parameters that evolve slowly, it is shown that the work distribution can be approximated as Gaussian distribution. This leads to $\Delta F = \langle W \rangle - \beta \frac{\sigma_W^2}{2}$, where σ_W^2 is the standard deviation of work distribution[119]. The result agrees with fluctuation-dissipation theorem with $W_{diss} = \beta \frac{\sigma_W^2}{2}$.

Crooks fluctuation theorem is a generalized form that relates forward and reverse work processes with free energy[63]. Given a controllable external input parameter I that controls the switching between two states in a system, the work W distribution for forward P_f and reverse P_r is related to the equilibrium free energy difference:

$$\frac{P_f(W)}{P_r(-W)} = \exp(\beta(W - \Delta F)) \tag{6.15}$$

This fluctuation theorem implies that under certain control protocol I, one can analyze fluctuation of nonequilibrium work trajectories and identify W^* so $P_f(W^*) = P_r(W^*)$, leading to $W^* = \Delta F$. This is the nonequilibrium work method of measuring equilibrium free energy difference between two states.

Heat exchange between two systems without any work performed follows the heat exchange fluctuation theorem[120]. After a system reaches equilibrium at temperature T_1, it is placed in thermal contact with another heat reservoir at a different temperature T_2. Heat exchanged between two system follows a similar form:

$$\frac{P(+\Delta Q)}{P(-\Delta Q)} = \exp\left(\Delta\beta\Delta Q\right) \tag{6.16}$$

where $+\Delta Q$ and $-\Delta Q$ are heat transferred in two directions and $\Delta\beta = 1/T_1 - 1/T_2$. With the definition of internal energy $E = W + Q$, we can replace ΔQ variables with ΔE in XFT under the condition without work performed $W = 0$.

Input driven stochastic Hopfield network model Neural activities for each neuron i in the Hopfield network are binary spins $V_i = \{+1, -1\}$ and neuron j to neurons i is connected with synaptic weight W_{ij}[104]. The energy function can be written as a function of neural activation patterns and external input I_i:

$$E(V, I) = -\frac{1}{2}\Sigma_i\Sigma_j W_{ij}V_iV_j - \Sigma_j V_iI_i + \Sigma_i U_iV_i \tag{6.17}$$

where U_i is the value of neuron i that sets an activation value for the net input H_i:

$$H_i = \Sigma_j W_{ij}V_j + I_i \tag{6.18}$$

which leads to neural dynamics through time t with discrete time steps:

$$V_i(t+1) = \sigma(H_i(t) - U_i) \tag{6.19}$$

where σ is the a nonlinear function that maps to a probability of flipping V_i to $+1$ or suppressing at -1. We used the sigmoid function for be consistent with the mean field approximation of spin glass models:

$$\sigma(x) = \frac{1}{1 + \exp(-\beta x)} \tag{6.20}$$

Note that this form of neural dynamics in equation 6.18-6.20 agrees with the minimization of energy function in equation 6.17. Simulation of such energy-based neural dynamics is known is Glauber dynamics. At thermal equilibrium, the probability of finding neural activity given a static external stimulus follows the Boltzmann distribution:

$$P(V) = \frac{1}{Z_I} \exp(-\beta E(V, I)) \tag{6.21}$$

where the partition function $Z_I = \Sigma_V \exp(-\beta E(V, I))$. Associative memory patterns should be retrieved at thermal equilibrium. To build in memory of patterns in the Hopfield network, namely the local minimum in the energy landscape, we applied Hebbian learning rule to the connectivity matrix W_{ij} [103]:

$$W_{ij} = \frac{1}{M} \Sigma_{\mu}^{M} \epsilon_i^{\mu} \epsilon_j^{\mu} \tag{6.22}$$

where M patterns are learned and each pattern ϵ^{μ} is a binary vector with the length of N neurons in the network. Note that this learning rule produces a symmetric matrix $W_{ij} = W_{ji}$ and we removed self-connections $W_{ii} = 0$, as done in the original Hopfield model.

Importantly, we can then define work W performed by controlling the external input I and heat production through relaxation of V dynamics.

$$W = \Sigma_{t=0}^{\tau-1} E(V_t, I_{t+1}) - E(V_t, I_t) \tag{6.23}$$

$$Q = \Sigma_{t=1}^{\tau} E(V_t, I_t) - E(V_{t-1}, I_t) \tag{6.24}$$

$$\Delta E = W + Q \tag{6.25}$$

where the time evolution from 0 to τ follows the time-varying stimuli protocol $I(t)$. Note that we integrated work performed through the process for multiple instantiations to compute the distribution of work W. With equation 6.17 and 6.23, we derived that $W = \Sigma_t^\tau \delta w_t$, where $\delta w_t = \delta I V(t)$ with δI being the change in dynamic input signal. Note that this is consistent with electrical circuits, where power is defined by the injected current multiplying voltage $P = IV$.

Model settings and stimulus protocol We constructed a simple Hopfield network with two memory patterns. The factor γ that controls the weighting of memory is introduced to provide asymmetry between to stable states ϵ^1 and ϵ^2:

$$W_{ij} = \gamma \epsilon_i^1 \epsilon_j^1 + (1-\gamma)\epsilon_i^2 \epsilon_j^2 + \alpha \epsilon_i^1 \epsilon_j^2 \tag{6.26}$$

where the $0 \leq \gamma \leq 1$ and $0 \leq \alpha \leq 1$. Note that γ affects the energy function, thereby changing the probability of recovering two patterns at equilibrium according to equation 6.21. The parameter α weights the overlapping between two patterns and when this it is non-zero the weight matrix can be asymmetric. This is the original Hopfield model with two learned patterns with $\gamma = 0.5$ and when $\alpha = 0$. We designed a stimulation protocol that drives the system away from equilibrium. This time-varying input switches between two states. The initial state is defined as ϵ^1 and the other as ϵ^2. The forward process is $\lambda_{12} : \epsilon^1 \rightarrow \epsilon^2$ and reverse is $\lambda_{21} : \epsilon^2 \rightarrow \epsilon^1$. The input protocol follows $\frac{st}{t_s}\epsilon^2$ for the forward process and $\frac{st}{t_s}\epsilon^1$ for the reverse one. Here s is a factor that controls the rate of input strength increasing through time t and t_s is the time course of stimuli to normalize the time length of process. We computed the work performed along two processes with equation 6.23 and compare with the equilibrium free energy difference by $\Delta F = \Delta\langle E\rangle - T\Delta S$, where the internal energy is $\langle E\rangle = \Sigma_V P(V)E(V)$ and entropy is $S = -\Sigma_V P(V)\log P(V)$ with $P(V)$ given by equation 6.21. For ground truth of ΔF between states, we numerically computed the value through $F = -\beta^{-1}\log Z$, where Z is the partition function in equation 6.21.

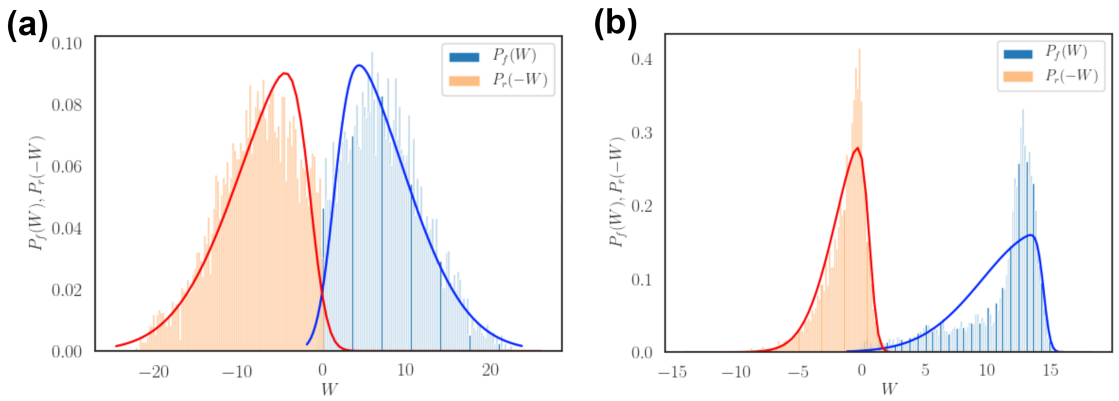

Figure 6.8: **Forward and backward work distributions.** Work distribution under networks with (a) $\Delta F = 0.1$ and (b) $\Delta F = 1.2$. Curves are skewed Gaussian fit to distributions.

For the following numerical analyses, we used the same set of parameters, if not further noted. Network parameters: $N = 15$, $k_B T = 15$, $U = 0$, $\alpha = 0$, $\gamma = 0.2$; Input protocol: $s = 0.005$ with maximum iteration steps equals 1000. Stored patterns ϵ^1 and ϵ^2 are randomly generated for each instantiation and the protocols are repeated for more than 5000 repetitions for statistical measurements.

Results

Work protocol and distribution We measured work performed along trajectories simulated across different random instantiates with the same network parameter and nonequilibrium protocol. The distribution of work is non-Gaussian and has skewed tails or multi-modal, depending on the network and work parameters. The multi-modal cases correspond to patterns with "glassy" energy landscape and have multiple local minimum, resulting in different modes of work required. The shape of distribution would not affect the measurement of equilibrium free energy difference according to equation 6.15. We can measure the value of W^* at $P_f(W^*) = P_r(W^*)$ by searching for the crossing point of two distributions (Fig 6.8). The value roughly agrees with true equilibrium ΔF calculated numerically. In addition to directly plotting the intersection of two distributions (Fig 6.8a,c), we applied sampling method based on Bennett's acceptance ratio to solve for ΔF (Fig 6.8b,d). This method does not rely on histogram binning and minimizes sam-

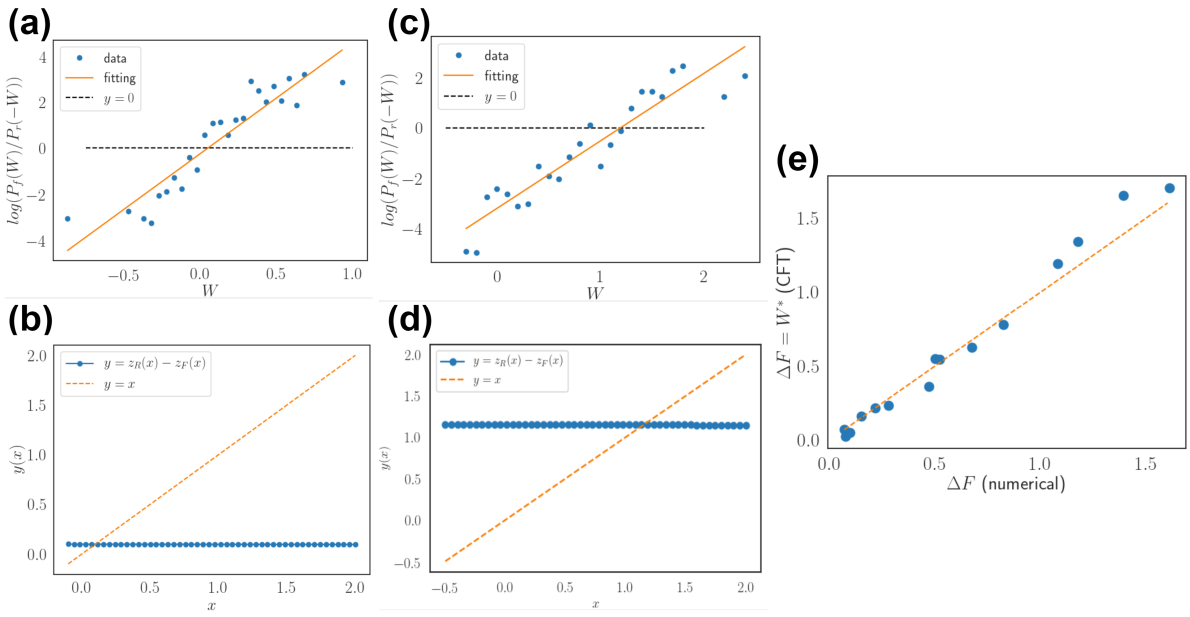

Figure 6.9: **Verifying CFT.** (a,c) Measuring the crossing point of two distributions $P_f(W)$ and $P_r(-W)$ as a function of work value W. Another method is through BAR (b,d), where we sampled with a function $f_x(W) = (1 + \exp(\beta(x - W))^{-1}$. The optimal $W^* = \Delta F$ value is then the crossing of function $y(x) = z_R(x) - z_F(x)$ and $x = x$, where $z_F = \log(\langle f_x(W) \exp(-\beta W)\rangle_F$ and $z_R = \log(\langle f_x(W)\rangle_R$. The true network $\Delta F = 0.49$ for (a,b) and $\Delta F = 1.25$ for (c,d). A systematic evaluation by constructing networks with different ΔF through tuning $\gamma = 0.1 - 0.9$ is shown in (e). Dash-line is when the numerical results agrees with our measurement through CFT.

pling variance[55,35]. We showed that the result agrees with histogram overlapping method and the ground truth free energy value.

The result is not sensitive to the choice of work parameter within a range nor the memory patterns stored in the network. Note that the rate of forward and reverse work protocols would change the width of work distributions. When the changing rate is sufficiently slow, it is near equilibrium and the distribution concentrates near the equilibrium free energy. Similarly, the temperature of the system also governs the work distribution. Higher temperature corresponds to noisier trajectories, resulting in broader distributions or longer tails.

Verifying fluctuation theorem We systematically varied the parameter γ in equation 6.26 to change the equilibrium free energy difference ΔF and estimate it through nonequilibrium work protocols shown in the former section. The results show that the nonequilibrium method recovers

the value of ΔF within some range, supporting the original prediction from CFT (Fig 6.9e). This range of parameters suitable for CFT has to suffice certain conditions: (i) tails of two distribution should overlap and provide unambiguous crossing point, (ii) fluctuation of the finite measurements has to be smaller than the scale of true equilibrium free energy difference, and (iii) the work protocol should not act too fast and dominating neural dynamics. Following these criteria, this method fails when the temperature is too high, the network is "glassy" and stored with multiple patterns, or when the work distribution is not sufficiently sampled.

Another variation of the fluctuation theorem focuses on the exchange of heat between two systems. We extended our thermodynamics analysis for neural networks to a condition not driven by external input, but connected to another temperature after equilibrium, similar to a global modulation in the neural network (see figure 5 in paper[43]). This corresponds to XFT that has a similar form as the original fluctuation theorem: $\frac{P(+\Sigma)}{P(-\Sigma)} = \Delta\beta\Sigma$, where Σ is the measurement of entropy production. XFT replaces Σ with heat flow Q defined by equation 6.24 and can simply be replaced by energy E when no work is performed. We showed that the Hopfield network produces heat distribution following XFT.

Adding biological constraints The network with Hebbian learning has connectivity matrix W_{ij} that changes according to the ongoing input pattern $I_i(t)$. This process effectively accelerates the rate of convergence towards an equilibrium point in the energy landscape during our work protocol, since we designed the input pattern to be a target fixed point. A changing energy landscape makes the definition of the true equilibrium free energy landscape ill-posed[270], but we argue that the nonequilibrium measurement holds as long as the learning dynamics is much slower than work protocols.

We explored two more constraints for biophysically plausible recurrent neural networks. One is the threshold of activation? (see figure 7 in paper[43]) and the other is asymmetric connections[270] (Fig 6.10). Different from the Hopfield model, biological neural networks have directed synaptic connections and form asymmetric connectivity. The parameter α in equation 6.26 controls

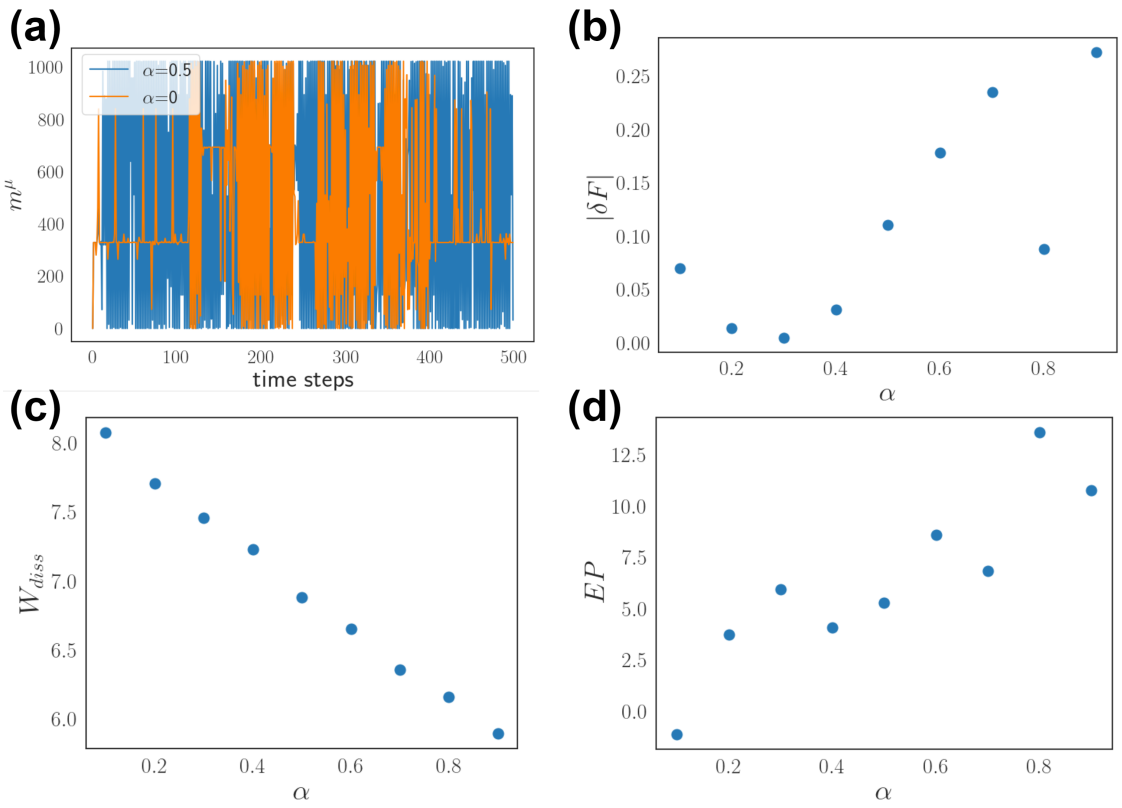

Figure 6.10: **Effects of asymmetric connections.** (a) Time series of neural states m^μ in neural networks with different asymmetric strengths α. (b) The absolute error of free energy estimation $|\delta F|$ as a function of parameter α. (c) The dissipated work W_{diss} as a function of parameter α. $\gamma = 0.2$ in all simulations, resulting with true ΔF at the scale of 1 for all instantiates. (d) Entropy production (EP) of networks as a function of asymmetric parameter α.

the weight of the cross over between two memory states. This can be thought as building an intermediate state between two attractors in the energy landscape. An important outcome of this term is that the connectivity matrix W_{ij} may not be symmetric as in the original Hopfield model. The energy function would no longer be well-defined, since the derivative of energy functions would not guarantee to be pure gradient force (matrix W_{ij} is not symmetric and may no longer positive semi-definite)[270]. However, at weak asymmetry we can be agnostic to this mismatch and still proceed to analyze the nonequilibrium work distributions of the system (Fig 6.10). The results hold for networks with asymmetric connections as long as the distribution of work protocols overlap and are sufficient for statistical sampling. When α is comparable to γ or $1-\gamma$ that weights one of the attractor sates, the error of estimating free energy difference is small

(Fig 6.10b). As α increases, the dissipated work W_{diss} for forward trajectories decreases and the entropy production from "housekeeping" heat Q_{hk} increase (Fig 6.10c,d). Here we quantified this part of entropy production with:

$$Q_{hk} = \log\left(\frac{P_{ss}(\epsilon_1)M(\epsilon_2|\epsilon_1)}{P_{ss}(\epsilon_2)M(\epsilon_1|\epsilon_2)}\right) \tag{6.27}$$

where P_{ss} is the steady-state probability of a given neural state and M is the transition probability between states measured from steady-state[206]. With the original symmetric network $\alpha = 0$, detailed balance holds and this term would be zero. In other words, the entropy production here captures how irreversible the system is. This result is expected as the increase in asymmetric weight α on the crossover term results in a more irreversible process and the flux between state reduces dissipated work required. We noticed that the notion of equilibrium free energy is inaccurate with calculation of equation 6.17 when connectivity W_{ij} is asymmetric. However, results of recovering these free energy values still work in a reasonable range of asymmetry[60]. Further correction can be done by using steady-state values to compute non-equilibrium free energy[191,270].

6.3.2 Max Cal inference for spiking networks

Introduction

Inferring network properties from neural activity is an important goal of modern neuroscience. The spatiotemporal statistical properties of the spike trains are determined by the synaptic coupling between neurons and by how neurons response to stimuli they received from other neurons and/or from the environment[86,138]. Spike trains contain information about the underlying neuronal network, and can thus be used to infer network properties of interest[193,213]. However, questions remain about how much information about the network is encoded in simple spiking statistics and what constitutes the minimal dynamical model based on these statistics.

The connection between statistical observables and network properties can be established through

building the minimal, unbiased model based solely on the observables. Such minimal model is derived by the Maximum Entropy Principle (Max Ent) as a generic model inference principle[121,227,248], or its dynamics generalization known as the Maximum Caliber Principle (Max Cal)[122,64,196,49,271]. For example, Schneidman et al. showed that the Max Ent model built from the same-time pairwise correlations between neurons can capture the distribution of network spike configurations[220]; Mora et al. demonstrated that the Max Cal model generated from the autocorrelation of total spike numbers in neighboring time windows can reproduce the size and duration of the spiking-activity avalanche in the network[179]. Max Ent or Max Cal transforms statistics from data to minimal models, from which network properties can be extrapolated and compared to other properties of the system to reveal connections not known before.

Here, we are interested in the Max Cal inference of elemental network properties —primarily the synaptic coupling strength, as well as the neuronal response functions and the inter-spike interval distribution. In the present work, we utilized a more fundamental set of statistical observables to infer effective synaptic coupling and neuronal responses with Max Cal. Our method is simple, and yet it correctly identifies the direction and sign of synaptic coupling, as well as the heterogeneity in neuronal responses. We also showed that the minimal dynamic model from Max Cal can capture the inter-spike interval distributions to the leading order. Inference performance for synaptic coupling is tested across different neural motifs and subsets of neurons in a network. With such validation from simulated spike trains from biophysical models, we applied our method to salamander retina data[189,163] to identify motifs.

Neural states and Max Cal approach

Effective neural coupling and responses A network of N spiking neurons with binary states can have 2^N possible configurations — we call them *network states*. The time scale relevant for the biophysical properties of synaptic transmission is around 10 ms. We therefore used a sliding window of width in the order of 10 ms to construct a continuous-time jump process in the network state space from the network spike trains, as shown in Fig 6.11. To our knowledge, this

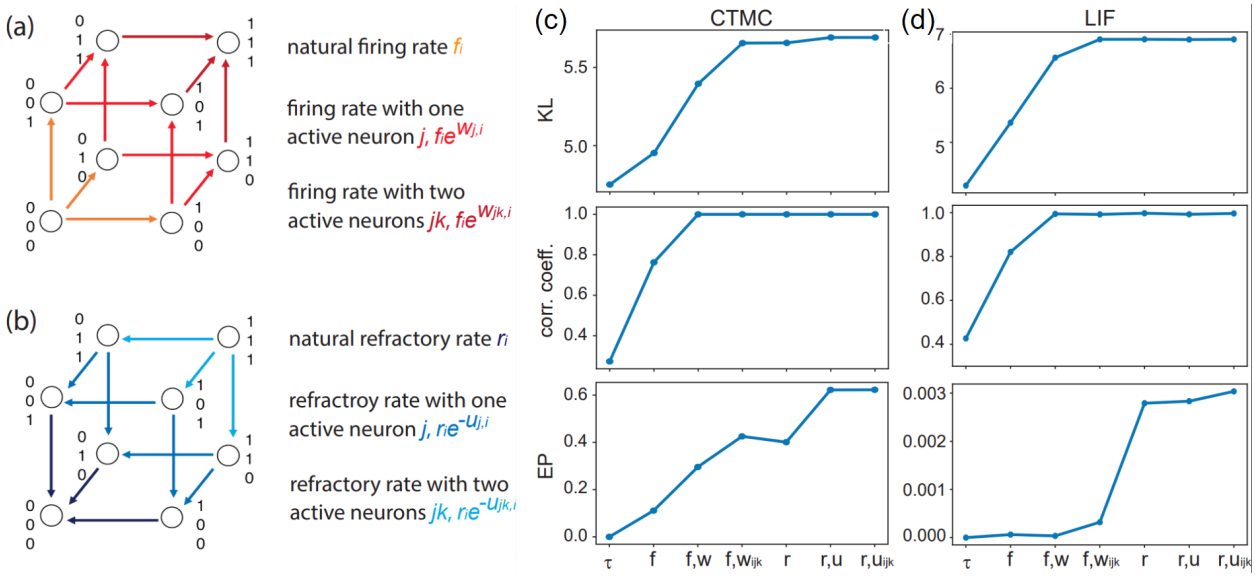

Figure 6.11: **Network states and inference learning curve.** (a) Firing rate and coupling parameters governing network state transitions to generate spikes. (b) Refractory rate and coupling parameters governing network state transition to silence. (c) Inference for continuous-time Markov chain (CTMC) across accumulation of different groups of observables in the x-axis. The optimized KL term (top) Pearson correlation coefficient against ground-truth weights (middle), and the notion of time irreversibility called "entropy production" in stochastic thermodynamics (EP, bottom) are shown. (d) Same as (c) but for a motif of three leaky-integrate-and-fire (LIF) neurons.

sliding-window construction was first proposed by Lynn et al.[163]. With this procedure, the network spike trains of N neurons is converted into a continuous-time jump process in a state space with 2^N possible network states, as illustrated in Fig 6.11a,b.

A continuous-time process has vanishing probability to have two jumps occur at the same instance. As each of the 2^N has N possible ways to jump out, the network state space is thus a hypercube with $N2^N$ edges indicating the possible jumps, leading to a *multi-partite* process[105,263,163]. Given long spike trains, we measured the total occupancy time in each network state $\boldsymbol{x}$

$$\tau_{\boldsymbol{x}} = \text{total time in state } \boldsymbol{x} \tag{6.28}$$

and the counts of transitions from one network state $\boldsymbol{x}$ to another $\boldsymbol{y}$,

$$C_{\boldsymbol{x},\boldsymbol{y}} = \text{number of } \boldsymbol{x} \mapsto \boldsymbol{y}. \tag{6.29}$$

From them, one can get the empirical distribution $\rho_{\boldsymbol{x}}$ of each network state $\boldsymbol{x}$

$$\rho_{\boldsymbol{x}} = \frac{\tau_{\boldsymbol{x}}}{\text{total time } T} \tag{6.30}$$

and the empirical flux $\rho_{\boldsymbol{x},\boldsymbol{y}}$ from state $\boldsymbol{x}$ to $\boldsymbol{y}$

$$\rho_{\boldsymbol{x},\boldsymbol{y}} = \frac{C_{\boldsymbol{x},\boldsymbol{y}}}{\text{total time } T}. \tag{6.31}$$

Their ratio, in the infinite sampling idealization $T \to \infty$, gives the rate/tendency of making the jump $\boldsymbol{x} \mapsto \boldsymbol{y}$ when the network is in the state $\boldsymbol{x}$:

$$R_{\boldsymbol{x},\boldsymbol{y}} = \lim_{T\to\infty} \frac{\rho_{\boldsymbol{x},\boldsymbol{y}}}{\rho_{\boldsymbol{x}}} = \lim_{T\to\infty} \frac{C_{\boldsymbol{x},\boldsymbol{y}}}{\tau_{\boldsymbol{x}}}. \tag{6.32}$$

These rates can help us reveal the synaptic coupling type and effective strength as well as the effective neuronal response functions.

With our interpretation of the spike rate and coupling parameters along the hypercube of network states, we can back out the parameters such as firing rate f_i for individual neuron i, coupling $w_{j,i}$ from j to i, higher order coupling $w_{jk,i}$; as well as the refractory counterparts r_i, $u_{j,i}$, and $u_{jk,i}$. Furthermore, by using $\sum_{j(\neq i)} w_{j,i}$ as the effective total "synaptic current" received by neuron i, denoted as I_i, and firing rate R_i as the i-th neuron firing rate under different network firing patterns, we can also construct effective response function of neuron i, denoted as $R_i = \varphi(I_i)$, based on the network state jump rates $R_{\boldsymbol{\xi},\boldsymbol{\eta}}$.

Maximum Caliber infers Continuous-time Markov chain as the minimal model We have shown above that the network state jump rates $R_{\boldsymbol{\xi},\boldsymbol{\eta}}$ defined from network state occupancy time τ and transition counts C can be used to reveal effective network properties, such as the coupling strength and the neuronal response functions. How many other network properties could be captured by the information encoded in these rates $R_{\boldsymbol{\xi},\boldsymbol{\eta}}$? To see this, we used the dynamics ex-

tension of the Maximum Entropy principle — Maximum Caliber principle (Max Cal) — to build the minimal dynamical model based on the observables τ and C, or equivalently $(\rho_{\boldsymbol{\xi}}, \rho_{\boldsymbol{\xi},\boldsymbol{\eta}})$, that define the rates $R_{\boldsymbol{\xi},\boldsymbol{\eta}}$. Then, by comparing the minimal model and the data, we can verify how many properties can be explained by the transition rates. In short, Max Cal will tell us that the minimum model based on (τ, C) observables is a Markov jump process.

The Maximum Entropy principle (Max Ent), originated from equilibrium statistical physics, has been generalized as an inference principle to generate an unique, unbiased model with minimal update to a prior model[121,227,248,196,125,32]. In a nutshell, given a prior probability distribution model $P_0(z)$ and the set of candidate posteriors $\mathcal{P}_c$ satisfying a set of constraints c, Max Ent finds the unique posterior $P^*(z)$ that is Kullback-Leibler (KL) closest to the prior P_0:

$$\begin{aligned} P^* &= \arg\min_{P\in\mathcal{P}_c} D(P||P_0) \\ &= \arg\min_{P\in\mathcal{P}_c} \sum_z P(z)\log\frac{P(z)}{P_0(z)}. \end{aligned} \tag{6.33}$$

As explained intuitively by Jaynes[121], since Kullback-Leibler (KL) divergence $D(P||P_0)$ is the unique measure of information gain[225], the Max Ent posterior P^* is thus unbiased in that it is the minimal model with no other information added other than the prior P_0 as the reference model and the constraints c. Jaynes' intuitive explanation has been made mathematically precise recently through the axiomatization of Max Ent as a statistical model inference[227,125,32]. When applying Max Ent to build minimal model to real data, Max Ent transforms constraints (our knowledg/data about the system) into a minimal model P^* that can be used to find out what network properties the constraints are enough to predict[220].

With z taken as the system's state, Max Ent only infers probabilistic models for static state distributions, not models about dynamics. To apply Max Ent to dynamics, Jaynes further proposed taking z as paths $\{x(t), 0 \le t \le T\}$, and named such Max Ent of dynamics the principle of *Maximum Caliber* (Max Cal)[122]. For recent reviews, see, *e.g.*,[196,188]. When considering only the averages of time-additive state and transition observables, $\langle\sum_t F_{x(t)}\rangle$ and $\langle\sum_t G_{x(t-\mathrm{d}t),x(t)}\rangle$, it has

been shown that Max Cal produces Markov dynamics as the minimal dynamical model[64,82,146]. For examples, the counting observables we considered above — distribution $\rho_{\boldsymbol{\xi}}$ and flux $\rho_{\boldsymbol{\xi},\boldsymbol{\eta}}$— correspond to choosing delta functions for F and G in continuous time with delta functions[82,271]:

$$\rho_{\boldsymbol{\xi}}[\boldsymbol{x}(t)] = \int_0^T \frac{\delta_{\boldsymbol{\xi},\boldsymbol{x}(t)}}{T} \mathrm{d}t \tag{6.34}$$

$$\rho_{\boldsymbol{\xi},\boldsymbol{\eta}}[\boldsymbol{x}(t)] = \sum_{t_m} \frac{\delta_{\boldsymbol{\xi},\boldsymbol{x}(t_{m-1})}\delta_{\boldsymbol{\eta},\boldsymbol{x}(t_m)}}{T} \tag{6.35}$$

where t_m are jump times and $\boldsymbol{x}(t_m)$ is the state immediately after the jump. Given these counting observables $\rho_{\boldsymbol{\xi}}$ and $\rho_{\boldsymbol{\xi},\boldsymbol{\eta}}$, Max Cal produces effective Markov models.

If possible, we would measure these counting observables $\rho_{\boldsymbol{\xi}}$ and $\rho_{\boldsymbol{\xi},\boldsymbol{\eta}}$ for arbitrary long paths — ideally for infinitely long paths so that their statistical values have converged to their asymptotic, steady-state values as probability distribution $\pi_{\boldsymbol{\xi}}$ and flux $p_{\boldsymbol{\xi},\boldsymbol{\eta}}$. Max Ent, as a theory, is built for the idealized infinite sampling that gives us the expected values (π, p)[271]. When we have finite but large data, we would treat our measured $\rho_{\boldsymbol{\xi}}$ and $\rho_{\boldsymbol{\xi},\boldsymbol{\eta}}$ as surrogates to $\pi_{\boldsymbol{\xi}}$ and $p_{\boldsymbol{\xi},\boldsymbol{\eta}}$, as if the data were infinitely long, an idea from thermodynamics of building a theory for infinitely large system and applying it to finite system as a leading order approximation[4,271].

Given a set of steady-state statistics $(\pi_{\boldsymbol{\xi}}, p_{\boldsymbol{\xi},\boldsymbol{\eta}})$ from idealized long-term data, the Max Cal of time homogeneous processes minimizes the steady-state KL rate that dominates the KL divergence $D(R||R^0)$ in the long-term limit:

$$d(R||R^0) = \lim_{T\to\infty} \frac{D(R||R^0)}{T} \tag{6.36a}$$

$$= \sum_{\boldsymbol{\xi}\neq\boldsymbol{\eta}} \pi_{\boldsymbol{\xi}} R_{\boldsymbol{\xi},\boldsymbol{\eta}} \log \frac{R_{\boldsymbol{\xi},\boldsymbol{\eta}}}{R^0_{\boldsymbol{\xi},\boldsymbol{\eta}}} + \pi_{\boldsymbol{\xi}} R_{\boldsymbol{\xi},\boldsymbol{\eta}} - \pi_{\boldsymbol{\xi}} R^0_{\boldsymbol{\xi},\boldsymbol{\eta}} \tag{6.36b}$$

where the steady-state distribution $\pi_i(R_{\boldsymbol{\xi},\boldsymbol{\eta}})$ is a function of the transition rates R_{ij}. Max Cal finds the continuous-time Markov chain (CTMC) with the posterior transition rates R^*_{ij} that mini-

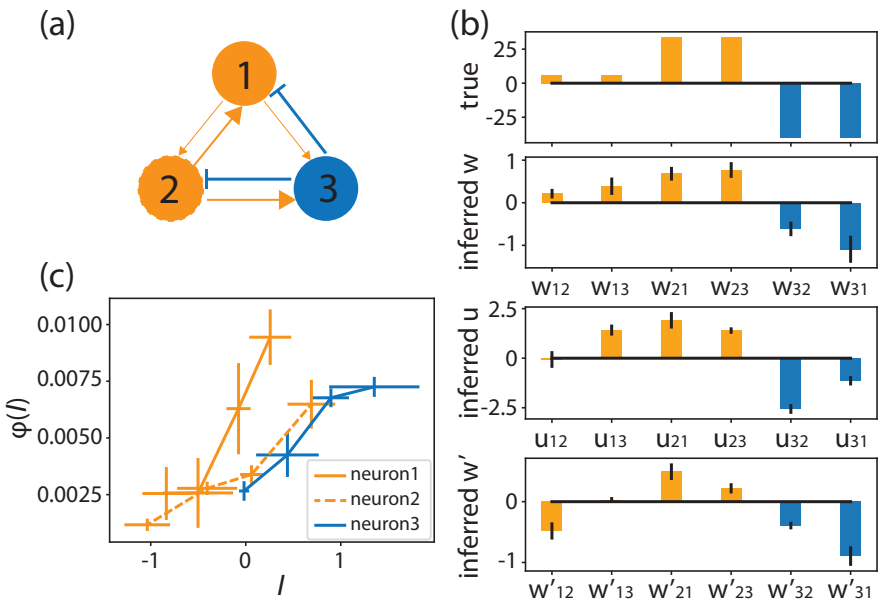

Figure 6.12: **Inferring coupling and response in LIF circuits.** (a) Three-neuron circuit with balanced E-I. Neuron 1 and 2 are excitatory (orange) whereas neuron 3 is inhibitory (blue). (b) The true (top) and inferred weights w_{ij}, refractoriness u_{ij}, and effective coupling w'_{ij}. Error bars show standard deviation across 10 instantiations. (c) The inferred response function φ for the same three neurons shown in (b). Neuron 2 has a lower firing threshold compared to neuron 1 and 3.

mizes d and satisfies the constraints:

$$R^*_{\boldsymbol{\xi},\boldsymbol{\eta}} = \arg\min_{R\in\mathcal{R}_c} d(R||R^0) \tag{6.37}$$

where $\mathcal{R}_c$ collects all possible processes $R_{\boldsymbol{\xi},\boldsymbol{\eta}}$ that satisfies a given set of constraints c, *e.g.* a bunch of distribution and flux measurements, no need to be the full set.

These posterior rates $R^*_{\boldsymbol{\xi},\boldsymbol{\eta}}$ can then provide us the inferred effective synaptic couplings and neuronal responses based on our set up in the last subsection. And the posterior KL-rate $d^* = d(R^*||R^0)$ represents the *information gain* from the set of constraints c. With more and more constraints feed into Max Cal, we gain more information with d^* increasing monotonically as shown in the first row of Fig. 6.11c,d.

Results

Inferring neural response and circuit motif We first tested inference for network coupling and response from spike trains simulated with a 3-neuron circuit (Fig 6.12a). We modelled a fully connected excitatory-inhibitory (E-I) balanced motif, with two excitatory neurons and an inhibitory neuron with stronger projection. The neurons are LIF units that integrate current input and generate discrete spikes when the voltage exceeds a threshold. The synaptic coupling is modelled through a weighted and filtered spike impulse from the pre-synaptic neuron. In this biophysical neural model, we tuned spiking threshold and synaptic weights to investigated inference for neuronal response and network coupling, respectively.

The inferred weights w and refractory coupling u have relatively consistent trends as the true synaptic weights underlying LIF circuit, but not the effective coupling w' (Fig. 6.12b). Particularly, w correctly reflects the underlying sign and relative amplitude of synaptic connection. The computed w' is with larger variations, potentially due to the limited occurrence for the state transitions between more synchronized states such as (011) to (111), but other than w'_{12} the excitation and inhibition inference are consistent with the true network. Additionally, we explored the effects of time window size and signal to noise ratio in the spiking network and found consistent results within a range of hyperparameters.

In the simulation, one of the two excitatory neuron (neuron 2) has a lower spiking threshold (-60 mV, compared to the default value at -50 mV). Such neuronal heterogeneity can be inferred from the response curves ϕ computed and shown in Fig 6.12c. Neuron 2 has a higher response curve that reflects its lower spiking threshold; neuron 1 and 3 have essentially the same response curve within the error bars. Together, these results show that the simple CTMC framework captures the biophysics of a generic spiking model. This not only validates the inference method but also suggests application to biological data, which we will visit in last section.

We further explored the flexibility of the inference method in identifying different motifs by simulating spike trains from three LIF motifs: (1) common-input, (2) chain, (3) cyclic motifs (Fig 6.13). Across different motifs, we found that the inferred weights correctly identify the true

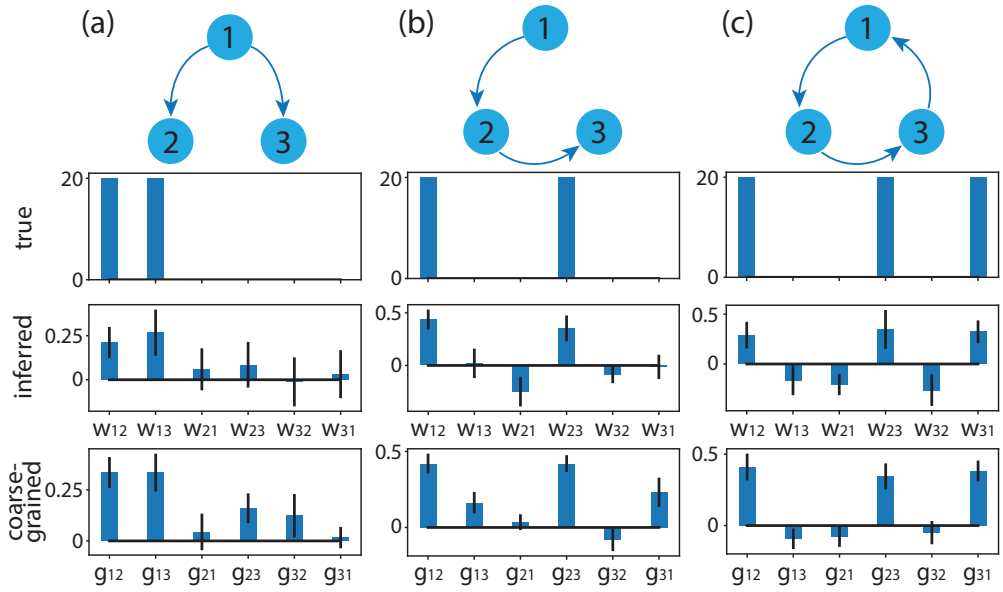

Figure 6.13: **Inferring different circuit motifs.** (a) Motif with common input shown on top. The true connection and inferred weights w_{ij} shown in the middle. The coupling inferred through coarse-graining g_{ij} is shown on the third row. Error bar shows standard deviation across 10 instantiations. (b) Same as (a) but for a chain motif. (c) Same as (a) but for a cyclic motif.

connections. We note that the successful inference indicates that the methods identifies causal coupling, rather than correlation due to common input. Furthermore, the inference results hold even when we take a coarse-graining approach to analyze only pair-wise neural interaction and marginalize over the third one. This coarse-graining analysis of network generalizes to larger network when we apply to subsamples of a larger network with 5 neurons (see figure 6 in paper[47]). This result indicates that one can confidently identify motifs from sub-sampled experimental measurements, under the mild assumption that the hidden neurons coupled to the observed units with slightly weaker weights.

Max Cal reconstruction of spike trains The Max Cal framework enables us to generate posterior spike trains with the CTMC posterior model. We characterized the network spike train generated from CTMC with its inter-spike interval (ISI) distribution for the whole network (the sum of each neuron's ISI distribution). We explored how well CTMC recapitulates the ISI of LIF simulations across different aforementioned network motifs and input noise strength (Fig. 6.14). The Max Cal reconstruction is evaluated with the KL divergence between the ISI distribution

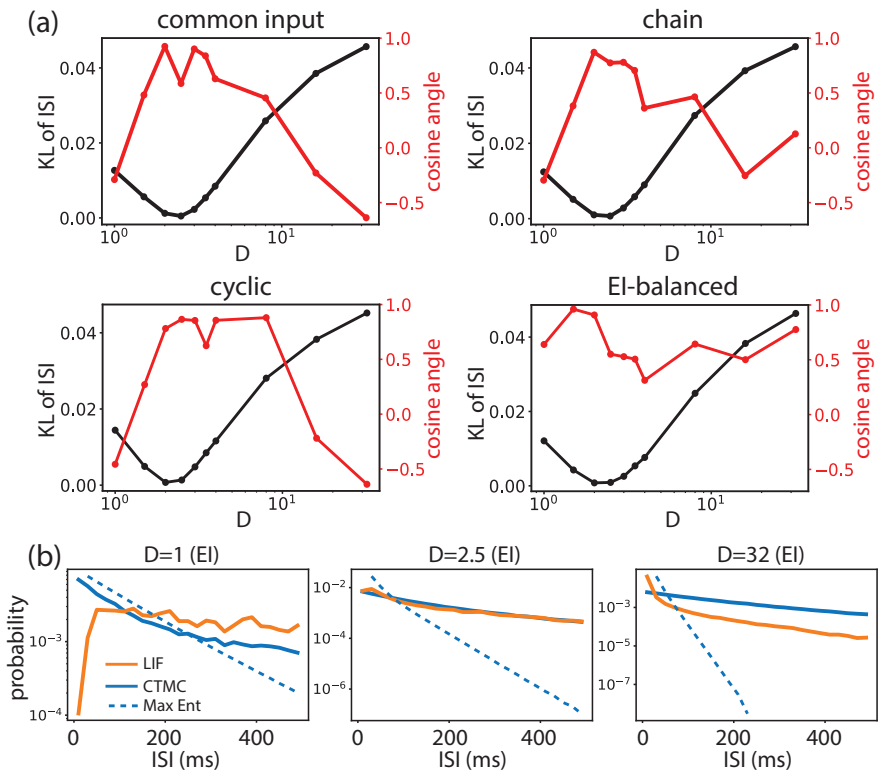

Figure 6.14: **Spike train reconstruction and network inference.** (a) The KL divergence between ISI of LIF simulation and CTMC reconstruction as a function of noise strength, D, is shown in black. The cosine angle between LIF network weights and inferred w_{ij} is shown in red. The same analysis is done for four different network motifs. (b) Example ISI distribution of the E-I balanced circuit in (a), with three different noise strength D selected. The empirical measurement from LIF spike train, CTMC reconstruction, and Max Ent prediction are shown.

measured from LIF spike trains and the CTMC reconstruction. We found consistent U-shape tuning to noise of this reconstruction across motifs, meaning that the network is best described by CTMC at a certain noise strength. Furthermore, the noise regime that generates spike trains well-approximated by CTMC coincides with better weight inference among different motifs (Fig. 6.14a). We also observed that coupling strength inference is more robust than ISI under varying noise levels across different motifs, showing that coupling strength inference is an task easier than ISI.

We studied the reconstruction of ISI distribution across three noise levels in the E-I circuit (Fig. 6.14b). To understand how much temporal structure is captured in the CTMC minimal model from Max Cal, we constructed a Max Ent model with discrete time bins for comparison[220].

We used discrete time windows to measure the long-term network-state distribution and after which we roll a 2^N-face dice repeatedly to reconstruct spike trains from Max Ent. The resulting "process" would be independent and identically-distributed (i.i.d.) sampling of the measured distribution, with no temporal correlation. The ISI distributions inferred by Max Ent are depicted by dashed lines in Fig. 6.14b.

One sees clearly that the CTMC representation is a correction to the empirical ISI distribution generated by Max Ent and matches better to that of LIF circuits across all noise. Consistent with the U-shape KL divergence as a function of noise strength in Fig. 6.14a, we found that ISI distribution is best capture by CTMC at an intermediate noise level. Particularly, the CTMC misses the drop of distribution in low noise, and it over estimates long ISI in high noise cases.

Neurons are known to have refractory periods with much lower firing probability immediately after they fired[86,138,193]. Such refractory period is required in models to describe a number of neural dynamical properties[259]. Here, the refractory manifests in the low occurrence of small ISI in the LIF simulated data in Fig. 6.14b. The simple two-state coarse-grained neurons in our present Max Cal CTMC model does not have such refractory mechanisms, which is the primary reason of the mismatch between LIF and CTMC there. With membrane potentials $V(t)$ data simulated from the LIF model, one could coarse grain a single neuron's states into three by the firing threshold θ and the resting potential V_r. One could define the firing state 1 as $V \geq \theta$, the refractory state -1 as $V \leq V_r$, and the resting state 0 as $V_r < V < \theta$. The resulting network state space would be more complicated — there will be 3^N number of network states in a network of N neurons instead of 2^N. The exact same operations can be conducted to define effective synaptic couplings and neuronal response functions. We leave such generalization to future works.

Application to retinal spike train We explored network inference for retinal spike trains measured from Palmer *et al.*[189]. The population spikes are measured from a group of salamander retinal ganglion cells driven by complex visual stimuli. We processed the spike train into network states with a 20 ms sliding window. In an example triplet of retinal ganglion cells, we found that

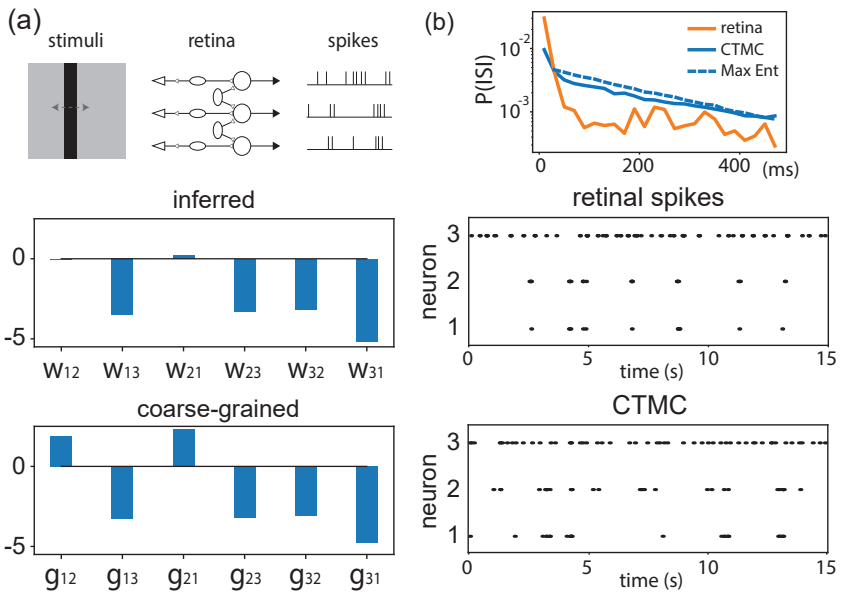

Figure 6.15: **Network inference with retinal spike trains.** (a) Schematic of the retinal spike trains measured under a stochastic moving-bar stimuli (top). The inferred weights w_{ij} (middle) and coarse-grain coupling g_{ij} (bottom) of three selected retinal ganglion cells are shown. The inferred $w_{12} = -0.05$ is weak and inconsistent with g_{12}. (b) The ISI distribution (top) of retinal spike trains from the three neurons in (a). The measurement is compare to CTMC reconstruction and Max Ent. The measured (middle) and CTMC generated (bottom) network spike trains are shown.

a pair of neurons that may be mutually exciting and inhibiting the third one (Fig. 6.15a). The result is consistent with previous results showing positive and negative effective coupling across retinal spike trains [193,220].

Similar to reconstruction of LIF spike trains in Fig. 6.14, we reconstructed the ISI distribution using the inferred CTMC representation and the Max Ent fit. We consistently found that CTMC is slightly closer to the empirical retinal spike train distribution. While the reconstructed spike train appear qualitatively similar (Fig. 6.15b), however, the mismatch of ISI across intermediate time scale ($\sim$100 ms) is still prevalent. This is different from the mismatch trends in noise-driven cases in Fig. 6.12b. One hypothesis is that the retinal spike train is a result of both the dynamical visual stimuli and more intricate pre-synaptic computation. The idea is supported by the fact that we observed extra effective coupling structure beyond the time scale of single neurons (see Appendix [47]). While future work is needed to incorporate more elaborate encoding model in the

framework, the consistent weight inference validates that the method can be generalize to biological measurements.

6.3.3 Discussion

Non-equilibrium dynamics in driven networks

Biological networks are often driven by time-varying external input while maintaining internal structure such as memory patterns stored in neural networks[242,78]. The resulting far-from-equilibrium trajectories of these driven systems can not be analyzed by methods based on equilibrium statistics[34,270,275]. In this project, we introduced nonequilibrium methods of recovering equilibrium free energy difference in recurrent neural networks with a defined energy function. We showed that the fluctuation theorem holds in our network model and further discuss how biophysical parameter alter the results.

The non-equilibrium analysis of neural network dynamics has been explored recently[270,275,214,178]. In a stochastic thermodynamics setting, the change in weights of an perceptron during supervised learning can be defined as a forcing term and the objective function being an effective energy function[88]. The thermodynamic efficiency during learning processes can then be defined in terms of the heat production due to weight updates and mutual information between the target signal and predictive output. In a nonequilibrium thermodynamics approach, another work[216] analyzed the unsupervised learning process of a restricted Boltzmann machine (RBM). The results show that fluctuation theorem holds in this abstract model and the learning process is related to the reduction of dissipated work. Interestingly, if we inspect the Hopfield's original simplified model with dimension analysis, one would conclude that the "energy" term has a unit that is closer to entropy production rate in nonequilibrium systems[41]. Lastly, a recent study derived the nonequilibrium Green's function of neural networks and attempts to characterize functional connectivity with its time-dependent response to impulse[203]. It is known that glassy systems violate fluctuation dissipation theorem, thereby diminishing linear response theorem applied near equilibrium[166]. More methods incorporating stochastic and nonequilibrium thermodynamics are

required in the future.

The information measurement plays a role in the fluctuation theorem when the stimulus is not independent to neural activity. According to the thermodynamics of information, we can rewrite the Jarzynski equality as $\langle \exp(-\beta W) \rangle = \exp(-\beta \Delta F - \Delta MI)$, where ΔMI is the mutual information measured between neural activity and input protocol $MI(V, I)$[191]. Mutual information between two random variables X and Y is calculated from the difference in entropy once an variable is observed: $MI(X, Y) = S(X) - S(X|Y)$, where S is entropy defined in the former section. Since our stimulation protocol is independent to the neural activity in our simulations, $MI(V, I)$ is negligible (on the order of statistical error when we numerically computed the value). When this information term dominates, generalized fluctuation theorem should include the mutual information term and one can verify the theory through analyzing the distribution of forward and backward work distribution along with the information measurement. Note that this quantity of $MI(V, I)$ serves as a lower bound of dissipated work and a designed protocol with higher efficiency minimizes dissipated work by maximizing the predictive information during the work protocol[236,191]. In addition, mutual information between stimulus and neural response in the network has been used to quantify neural encoding in neuroscience[240,193]. Whether the thermodynamic efficiency relates to the encoding performance is the goal for future directions.

Max Cal inference for spiking network

In this work, we applied Maximum Caliber (Max Cal) inference to the jump-process representation of network spike trains to recover network and response properties, using the inferred Continuous-Time Markov Chain (CTMC) as a minimal model. For numerical simulations of leaky integrate-and-fire (LIF) neurons, we were able to infer synaptic couplings and neural responses. The inference is applicable to different neural motifs and generalizes to cases with unobserved neurons. Our method demonstrated the ability to extract causality, not correlation, and to identify direct couplings instead of indirect effective couplings. Additionally, the interspike interval (ISI) distribution can be reconstructed through Max Cal over a range of input noise

strength to the network. Lastly, we applied our method to spike trains measured from the retina. While the ISI reconstruction showed a larger discrepancy, potentially due to strong stimuli applied to the system, we were able to characterize the retinal motif using this method.

The CTMC representation of spike trains grows exponentially with the number of observed neurons. This is a similar challenge for past work building state space for neural measurements[163,213]. While mainly studying various 3-neuron motifs in this work, we have shown that this 3-neuron spike trains could be a subset of the total network of 5 neurons, given that the ignored couplings are slightly weaker. According to the setup in Fig. 6.11, the degree of freedom (dof.) for 5 neurons grows to 160, in comparison to 32 dof for 3 neurons. As dof grows with the size of the hypercube, the demand of data for sufficient statistics becomes large, and should be an important future direction. One possible alternative is to construct lower dimensional state-space that may not scale with the number of measured units.

An interesting benefit of considering larger network is the richer variety of motifs and topologies. As shown in previous work[96], the number of possible motifs with directed connections is on the order of thousands for just 4 neurons. Furthermore, following our network-state jump-process framework, different cycles and flux across motifs can be studied, allowing for further exploration on how "entropy production" is related to the function of these motifs[162,163].

We focused on small network motifs[176,160] for starters to test our method, as couplings in small motifs with indirect/cyclic couplings and heterogeneous neurons could already be hard to infer[197,84,96]. Prior work applied Max Ent models to characterize functional coupling that are related to correlation of the spike trains[220,213]. In protein folding studies, it has been shown that Max Ent can distinguish between direct and indirect coupling[180], but the dynamics generalization to coupling with time delays is not known, which our present method paves the way to solve.

The causal coupling has been explored in pair-wise interaction frameworks such as the generalized linear model (GLM) or Granger causality[193,143,84]. Empirically, it has been shown that GLMs can successfully reconstruct network connectivity[84,66,143]. However, common across all methods and our results, the inference is particularly challenging in strongly interacting networks

and when units are unobserved[66,151]. Our framework is similar to GLM in characterizing causal interactions, but generalizes from pair-wise interaction to network states that include higher-order interactions. Our approach is similar to Max Ent in finding minimal models given data statistics, but Max Cal model offers a novel, principled approach to infer network properties and capture dynamics in spike trains. The present paper aims to fully explore our framework as a new way to infer network properties. We will leave detailed methods performance comparison to future works.

While we analyzed CTMC from network spike trains, this method does not require the underlying neural dynamics to be Markovian. Our approach captures the first-order temporal interaction, and we showed that many of these descriptions are consistent with the underlying biophysical parameters in simulation. To extend the current framework, additional observables such as state dwell time moments and distribution[228] could be incorporated into Max Cal to enhance inference accuracy. Sensory neural spike trains enables computation such as encoding readout[193,207]. Another promising future direction is to study how these network states are utilized in information processing[162].

7

Conclusion and Outlook

In this thesis, we presented experimental measurements and computational models for sensory navigation in worms and neural network dynamics. The projects in chapter 2-4 include aspects ranging from instrument construction, experimental measurements, to model fitting. We measured the worm's navigation with the apparatus developed in chapter 2. We then introduced a statistical model to characterize navigation and learning effects. This framework was generalized in chapter 3 into a state-space model to study behavioral dynamics. The cross-talk between experimental and theoretical work is evident – we made precise measurements to constrain models, then modified models and designed perturbation experiments according to the empirical measurements. For instance, the statistical model was fitted to precise measurements from our novel flow

chamber; and only through such modeling framework can one quantitatively map out the relation between optogenetics and odor response.

In addition to experiment-oriented work, we presented theory work in chapter 5-6, which spanned from follow-up models for worm behavior to more abstract neural dynamic models. The computational framework can assist better experimental analysis and predictions. For instance, in chapter 6, the latent variable model corrects for motion artifact in neural imaging signals and better decodes behavior in freely moving worms. Theory work can also open the avenue for more quantitative investigation or formal framework. For instance, in chapter 5, we investigated the connection between RNNs used in mechanistic modeling and state-space models used for descriptive models.

7.1 Conclusions

In chapter 2, we introduced a novel experimental apparatus for controlling and monitoring an odor environment, specifically designed to study navigation in small animal. We built a scalable odor sensory array within an odor flow chamber to capture spatiotemporal patterns of the odor field. By working with steady-state flow patterns and accounting for non-equilibrium odor-agar interaction, we could determine the odor landscape through experimental validations and theoretical arguments. The novel method enabled measurements of detailed navigation strategies in animals, quantitatively showing that worms employ both biased random walk and weathervaning during odor navigation. The setup also allows flexible use to study other species, such as the fly larva that navigates in the same environment. This is to our knowledge the first methodology to precisely control and measure airborne odor cues in the context of worm chemotaxis. There are a number of scientific and methodological future work following this project. Scientifically, how does the navigation strategy vary across different chemical cues and statistics of the sensory environment[192,110,241]? How does the navigation strategy in airborne odorant differ from non-volatile cues such as salt and temperature? More specifically, what is the biophysics of molecular sensing at the interface of gas and the moisture layer on agar? In terms of the instrument, a num-

ber of extension and improvement can be improved, including the new electronic components and geometry of the flow chamber that allows simultaneous odor and animal measurements. For potential future applications, better understanding of the active sensing elements in nature – navigating small animals – would support the design of artificial molecular sensors, and vice versa for advanced artificial detectors to study animal sensing.

In chapter 3, we examined how associative learning changes sensory navigation and neural computation in worms. We utilized the experimental setup in chapter 2 to measure odor navigation in worms after different training protocols. We developed a novel statistical model, dynamic Pirouette and Weathervaning (dPAW), that fits to time-varying navigation measurements. By explicitly incorporate biased random walk and weathervaning strategies in the model, dPAW characterizes how the worm weights sensory signals and transforms it into behavioral actions. We found that olfactory learning alters not only how the worm processes odor signals but also the precise sensorimotor transformation in two strategies. We probed neural mechanisms with optogenetics perturbation and genetic disruption. The result shows that while the sensory neuron AWC is crucial in generating learning-dependent sensorimotor transformation, the effects of learning on behavioral strategies may be distributed across the interneuron circuit in worms. The result gave rise to many follow-up questions. For instance, how is the learned olfactory memory stored and retrieved in the network to produce relevant navigation behavior[95,50]? What is the temporal dynamic of learning and can we pinpoint molecular or synaptic level learning processes? Are there additional behavioral strategies not captured by dPAW that can be recruited for learned navigation[111]? Future work with whole-brain imaging during odor navigation and more sophisticate behavioral models would be required to tackle these questions.

In chapter 4, we revisited a fundamental assumption of navigation strategies in worms, which describes a parallel use of navigation strategies without temporal dynamics. We tested this assumption explicitly by extending the dPAW model to a state-dependent Pirouette and Weathervaning (staPAW) model that includes the dynamics of latent states. We found that staPAW better describes navigation data in worms by capturing the state-switching behavioral strategies. By intro-

ducing optogenetic perturbation, we confirmed that the behavioral state transitions are non-trivial and related to the detail computation of local concentration gradients. To climb up the gradient, worm enters a state with slow speed and large angle change when the concentration drops, and transitions into a state with high speed and small angle change when the concentration rises. This reveals a novel navigation strategy that enables directed reorientation. We showed evidence that this phenomena is generic across odor and salt navigation. Future work is needed to explore how the navigation strategy vary across different sensory environment, and what neural mechanisms are involved in the state-dependent behavioral dynamics. While we only analyzed and model the navigation tracks, it is possible to extract more information from the posture dynamics in worms by combing theoretical frameworks developed recently[57,2,266]. Lastly, we proposed hypothesis of functional benefits for using state-switching strategies through numerical experiments and reinforcement learning models. It would be interesting to design experiments with different sensory statistics to test these normative predictions.

In chapter 5, we introduced different algorithms to train stochastic neural networks. We presented both fine-tuned models and flexible models optimized with a particle swarm algorithm. The model structure is a small stochastic RNN that captures the worm's behavioral dynamics, such as context-dependent response, head-tail integration of touch stimuli, and stochastic state transitions. We demonstrated that such stochastic neural circuit models can recapitulate the worm's behavioral responses. Recent development in neuro-AI discipline strengthens the connection between computation neuroscience and machine learning. However, most of the common architectures in machine learning use large networks, deterministic dynamics, and non-spiking signals. This is in contrast to invertebrate nervous system that is relatively compact and stochastic. In addition to small, stochastic rate circuits, we developed the G-FORCE algorithm and analyzed its ability to train stochastic spiking RNNs to produce complex temporal patterns. The chapter puts emphasis on biologically relevant architecture and target outputs. Future work can further constrain models with the detailed connectome in worms[250,260]. One can also utilize the whole-brain imaging measurements in worms to constrain dynamics in the artificial units[184,8]. This approach

would not only strengthen our understanding of neural computation but also facilitate the development of alternative architecture and algorithms.

Finally, in chapter 6, we included a series of theory work, ranging from a latent variable model known as two-channel motion artifact correction (TMAC) to remove motion artifact in neural imaging, an adaptive subunit model for predictive dynamics in the mouse visual cortex, to the nonequilibrium analysis of spiking networks. In the first project, the TMAC approach significantly improves neural decoding for the worm's behavior. The model structure is a simple approximation for how motion alters the fluorescence intensity during imaging. Future work is needed to test for generalization of this method to other imaging system, or to incorporate more sophisticate physical models for experimental artifacts[194]. The second project demonstrates how a simple phenomenological model can capture predictive visual response observed in experiments. The model makes predictions about neural adaptation to sequential input that can be tested with experiments in the future. The model demonstrates that simple adaptation can explain seemingly complex temporal patterns, but whether or not it matches the precise biophysical mechanism is still unknown[138,30,218,261]. Lastly, the theory project with nonequilibrium network demonstrates inference and dynamics for spiking neurons. Different from traditional physics approach that deal with symmetric, large networks, we demonstrated that modern nonequilibrium physics methods can be applied to perform better inference on parameters and analyze network computation. There are still a number of remaining gaps, including how to generalize to more complex neural network architectures[250] and methods to apply to measurements in larger brains[203].

Together, this thesis integrates topics from worm navigation behavior to stochastic neural dynamics. The research approach span the intersections of neuroscience, biophysics, and machine learning. Overall, we learned that more careful measurements and model designed to fit to data would facilitate our understanding for sensory navigation, learning, and internal states in animals. On the methodological development side, the research community has a growing set of methods to characterize and model naturalistic animal behavior[21,171,156]. The development of advance

sensors that are compatible in biological experiments would assist more precise experimental measurements. On the scientific side, modern neuroscience has underscored the importance of studying the brain in the context of animal behavior and its ethological relevance[139,52,28]. More specific to the worm, we have revealed how seemingly simple navigation behavior can be affected by experiences and demonstrate precise state-dependent dynamics. Building on top of these frameworks, we expect future work to generalize to more complex species, behavior, and neural computation.

7.2 Challenges, ongoing and future work

As neural recording techniques advance, measurements of neural activity with higher dimension and resolution become available. Together with theory approaches and machine learning methods, modern neuroscience is equipped to tackle complex aspects of the brain. However, certain challenges remain at the current state, including (1) approach to understand ethologically relevant behavior and the underlying neural computation, (2) methods in tackling non-stationarity in time and variability across individuals, and (3) framework to bridge detailed neural recordings with high-throughput behavior analysis.

As an preliminary attempt to address these challenges, in this thesis work, we built measurement systems (chapter 2) and experimental protocols (chapter 3) to study the worms' navigation behavior and learning, which constrains our understanding for flexible neural computation underlying natural behavior. We introduced state-space models (chapter 4) to tackle behavioral dynamics and capture the variable behavior during the worm's navigation. With ongoing and preliminary studies in chapter 5, we sought to bridge the level of neural dynamics with behavior. While advance neural imaging techniques enable whole-brain measurements during animal behavior, the throughput is order of magnitudes lower compare to behavioral analysis. The challenge lies in both continuous advance in measurements and also conceptual frameworks to seek relevant behavioral characteristics to relate with neural dynamics. While large-scale neural and behavioral measurements become accessible, many of these challenges suggest the importance to formal-

ize hypothesis and develop theory-testing experiments to understand the neural basis of complex behavior.

My higher-level goals for future research include two main aspects that build upon and extend from this thesis: (1) To study seemingly simple, small species capable of generating complex behavior, such as information processing and learning. (2) To comprehend intricate network dynamics through tractable computational models, including statistical and mechanistic descriptions. I aim to continue integrating experimental and theoretical approaches in both aspects, anticipating new observations, principles, and challenges alongside the research community.

A

Appendix

In this appendix, we included detailed derivation and explanation for contents corresponding to each chapter. Due to length considerations, we left out the detailed figures for projects that have been published, and directly refer to the appendix or supplementary information in the corresponding article. Lastly, we provided links to open source repository for code, data, and figures.

A.1 Odor-flow chamber project

To model odor flow landscapes, we used two models in our work: (1) a convection-diffusion model that captures quasi-steady state odor concentration profile measured without agar used for the fits in Fig 2.1f,g and (2) a reaction-convection-diffusion model for odor-agar interaction

shown in Fig 2.5. A version of this second model is also used to justify the pre-equilibration protocol, as discussed in the Appendix in[46].

A.1.1 Convection diffusion model for odor flow without agar

To model odor flow without agar, for example for the fits in Fig 2.1f,g, we use a two-dimensional convection-diffusion model:

$$\frac{\partial C(x,y,t)}{\partial t} = -v\nabla C + D\nabla^2 C \tag{A.1}$$

where the concentration across space and time is $C(x,y,t)$, flow velocity is v, and the diffusion coefficient of our odor is D. In our chamber, at steady state ($\frac{\partial C}{\partial t} = 0$) we have:

$$v\frac{\partial C}{\partial x} = D\frac{\partial^2 C}{\partial y^2} \tag{A.2}$$

because with our configuration flow along the x axis is dominated by convection while flow along the y axis is dominated by diffusion, and therefore $\frac{\partial^2 C}{\partial x^2} \ll \frac{\partial^2 C}{\partial y^2}$.
The fit in Fig 2.1f is the solution to equation A.2:

$$C(x,y) = \frac{C_o}{2}(1 - \text{erf}(\frac{x}{2\sqrt{D\frac{x}{v}}}))\exp(-\frac{y^2}{4D\frac{x}{v}}), \tag{A.3}$$

where erf is the error function and C_o is the odor source concentration measured in air.
In Fig 2.1g we show a fit for a one dimensional slice along y at various positions along x_c, for the situation in which there is an odor-source at $(y = 0, x = 0)$:

$$C(y) = \frac{C(x_c, y = 0)}{\sqrt{4\pi D\frac{x_c}{v}}}\exp(-\frac{y^2}{4D\frac{x_c}{v}}) \tag{A.4}$$

where $\frac{x_c}{v}$ is an analogy of time in non-stationary diffusion process at the cross-section at x_c.
For the fits in Fig 2.1 the air flow velocity is set to be $v \sim 0.5$ cm/s based on the flow rate and geometry of the chamber (15 parallel tubes provide around 450 mL/min of flow into a ~255 mL

chamber with $\sim$15 cm^2 cross section). The diffusion coefficient D is left as a free parameter and the value that minimizes the mean-squared error between the model and the empirical measurement is used. We chose to leave the diffusion coefficient as a free parameter instead of using butanone's nominal diffusion constant of $D \sim 0.08$ cm^2/s, because we expect butanone's effective diffusion coefficient to be different in a confined chamber with background flow. We note that the fitted profile shown in Fig 2.1f,g and the fitted value agrees with what is expected in a stable convection-diffusion process (Peclet number ~ 80).

A.1.2 Reaction-convection-diffusion model for odor-agar interaction

To justify the pre-equilibration protocol of Fig 2.4 and to show that measurements of odor concentration along the agar's boundary allows us to infer the concentration on the agar, we propose a reaction-convection-diffusion model. This phenomenological model forms the basis of Fig 2.5 and agrees with experimental observations Fig 2.6. Compared to the convection-diffusion model, we include the "reaction" term to account for odor-agar interactions.

The model used is a 2D generalization of this non-spatial model:

$$\frac{dC}{dt} = -\frac{1}{\tau}(C - C_o) - w\frac{dA}{dt} \tag{A.5}$$

$$\frac{dA}{dt} = k_a C(1 - \frac{A}{M}) - k_d A \tag{A.6}$$

where C is a downstream concentration readout after the airflow has surface interaction with the agar gel. The influx odor concentration is C_o and the odor concentration in agar is A. Without agar interaction, the flow chamber has its own timescale τ and the molecular flux into the agar is weighted by a scalar w (so $w = 0$ when there's no agar in the chamber). The association and dissociation constants are k_a and k_d and the maximum capacity of odor concentration that can be absorbed is M. This model is similar to the description of odorant pulse kinetics shown in[92].

In Fig 2.5, we use the 2D generalization:

$$\frac{\partial}{\partial t}C(x,y) = \mathcal{L}C(x,y) - w\frac{\partial}{\partial t}A(x,y) \tag{A.7}$$

$$\frac{\partial}{\partial t}A(x,y) = k_a C(x,y)(1 - \frac{A(x,y)}{M(x,y)}) - k_d A(x,y) \tag{A.8}$$

where $\mathcal{L} = -v\nabla + D\nabla^2$ (Equation A.1) is a linear operator for the convection-diffusion process and the odor influx is at the boundary $C(x = 0, y = 0) = C_o$. We perform numerical analysis on the set of 2D equations and permit A to be non-zero only in the region where agar is present. We use a target concentration C_o that is lower than M and $k_a \gg k_d$ to capture odor absorption into agar. In the simulated pre-equilibration protocol we temporarily increase C_o above M then switch back to the target concentration to efficiently reach a steady-state.

A slightly simplified version of this model forms the basis of the arguments in the Appendix of the paper[46].

A.2 Olfactory learning project

Behavioral analysis and dPAW inference

We tracked the location of the worm's centroid and fit a centerline to its posture as described in[46]. We removed trajectories that are shorted than 1 minute or have displacement less than 3 mm across the recordings. In addition, trajectories that started above 70% of the maximum odor concentration were also removed to prevent double-counting worms that may have already started or traveled up-gradient. This results in 270-1,140 animal hours of chemotaxis trajectory per training condition for model fitting. For each processed navigation trajectory, we computed the displacement vectors every 5 time bins (5/14 seconds) and compute the angle between consecutive vectors to obtain $d\theta_t$. We computed the odor concentration it passes through using the two-dimensional odor landscape measured in the flow chamber C_t. The perpendicular concentration difference $dC^{\perp}$ is calculated with unit vectors that are orthogonal to each displacement vector.

We also recorded the length of each displacement vector to form the empirical speed distribution. We fit dPAW to the ensemble of trajectories by maximizing the log-likelihood:

$$\text{argmax}\Sigma_n^N \Sigma_t^T \log P(d\theta_{n,t+1}|C_{n,:t}, dC_{n,:t}^{\perp}, d\theta_{n,:t}) - \lambda|K_C|^2 \tag{A.9}$$

where N is the number of trajectories, T is the time steps, and λ is the regularization term. To impose smoothness on kernels, K_C is parameterized with 4 raised-cosine basis function[193], $K_{dC^{\perp}}$ and K_h are parameterized with an exponential form. Optimization was performed with constrained optimization in Matlab, where the constrains are positivity of the precision parameters and sigmoid probabilities. Uncertainty about the inferred parameters are characterized by numerically computing the Hessian of the log-likelihood function around the maximum likelihood estimation. For model-based decoding, we perform 7-fold cross validation with all measured trajectories. To test with finite data, we subsample 10 ensembles of trajectories to test for performance in Fig 3.4.

To validate the accuracy of maximum likelihood inference for dPAW model parameters, we simulated behavioral time series from dPAW with a fixed set of ground-truth parameters. The model generated simulated trajectories using Gaussian random white noise for concentration time series C_t and $dC_t^{\perp}$. The time series has 50,000 data points, which is in the same scale of our data length ($\sim$ 300 animal hours of recording at $5/14$ Hz sampling frequency). We confirmed that the inference procedure works with simulated data and recovers the ground truth parameters in the supplementary information of[45].

To generate chemotaxis trajectories from inferred parameters shown in Fig 3.2, we conduct agent-based simulation by measuring concentration in the same odor landscape and drawing an-

gular change $d\theta_t$ from dPAW. We simulate two-dimensional navigation trajectories with:

$$x_{t+1} = x_t + v_t \cos(\theta_t) \tag{A.10}$$

$$y_{t+1} = y_t + v_t \sin(\theta_t) \tag{A.11}$$

$$\theta_{t+1} = \theta_t + d\theta_t \tag{A.12}$$

where v_t is speed drawn from Gaussian fit to the empirical distribution.

For information rate (Fig 3.3) and model comparison (Fig 3.4), we construct null models to compare with dPAW. The random walk model has similar statistical structure for behavior but is independent to odor concentration:

$$P(d\theta_t) = P(\beta = 1)P_{\mathrm{pr}}(d\theta) + P(\beta = 0)P_{run}(d\theta) \tag{A.13}$$

Note that in this null model, the turning probability is time-independent, so β does not have a time subscript t. The pirouette behavior has the similar sharp angle as the full model, but now in the null model the weathervaning is changed to "runs" that have zero-mean and do not take perpendicular concentration change into account. This model was fitted to the same ensemble of trajectories, and the log-likelihood difference between the full dPAW and this null model is normalized by $\log(2)$ per time to compute the bit rate. We also used this model to compute behavior-only model prediction in Fig 3.4. The chemotaxis model is formulated with a binomial distribution with expected fraction of track going up gradient $\hat{p}$. The estimation for chemotaxis index is then $2\hat{p} - 1$. Lastly, the concentration change model takes $C_{final} - C_{initial}$ for all tracks and uses naive Bayes classifier for prediction.

Statistical analysis for chemotaxis in mutant worms

We computed the concentration weighted chemotaxis index:

$$CI = \frac{(N_{up}\Delta C_{up} - N_{down}\Delta C_{down})}{(N_{up}\Delta C_{up} + N_{down}\Delta C_{down})} \tag{A.14}$$

where N is the number of tracks going up or down gradient and ΔC is computed for every track. We find the maximum and minimum concentration across the full observed odor landscape, then define $\Delta C_{up} = C_{max} - C_{final}$ for tracks going up gradient and $\Delta C_{down} = C_{final} - C_{min}$ for tracks going down gradient, with ending concentration C_{final}.

To conduct statistical tests between indices in Fig 3.6, we re-sampled chemotaxis trajectories in each conditions.For the weighted chemotaxis index, we sampled 50 tracks for 20 times to compute the standard deviation from mean. For the biased random walk strategy, we computed $|K_C|/std(C_t)$ to quantify the weights on concentration C and normalized by the standard deviation of the time series itself to compare across strains that experience different concentration inputs. For the weathervaning strategy, we computed $\mathrm{std}(K_{dC^\perp} * dC^\perp_{:t})$ to characterize how much the worm weights and corrects for the heading in response to $dC^\perp$. To conduct statistical tests on the behavioral strategies, we sampled 100 times from the posteriors of these kernels, with Gaussian approximation around the maximum likelihood estimate shown in Fig 3.2 b,c. For each transgenic strain and training condition, we computed the standard deviation of the metrics and conducted t-test between the transgenic strain and wild type worms, as shown in detail in supplement of[45].

A.3 State-switching chemotaxis project

A.3.1 The staPAW model and inference

The staPAW model describes how the consecutive step speed dr_{t+1} and heading change $d\theta_{t+1}$ are governed by experimental observations including concentration C_t, perpendicular concentration difference $dC^\perp_t$, past behavior ($d\theta_t$ but not dr_t), and optogenetic input I_t presented in the perturbation experiments. The emission of speed and heading change also depends on a latent state z_t that is modeled as an input-driven Markov chain. Within a state, it is similar to dPAW model that makes the Bernoulli decision q_t to perform a pirouette ($q = 1$) or to weathervane ($q = 0$). Similar to the proposed form in[45], the staPAW model conditioned on the latent state z_t is written as:

$$P(d\theta_t, dr_t | C_{1:t-1}, dC^{\perp}_{1:t-1}, d\theta_{1:t-1}; z_t) = \\ P(q_t = 1|z_t) P_{\text{pr}}(d\theta, dr|z_t) + P(q_t = 0|z_t) P_{\text{wv}}(d\theta, dr \mid dC^{\perp}_{1:t-1}|z_t) \tag{A.15}$$

where

$$P(q_t = 1|z_t) = m^z + \frac{M^z - m^z}{1 + \exp(K^z_C \cdot \mathbf{C}_{1:t-1} + K^z_h \cdot |\mathbf{d}\theta_{1:t-1}|)} \tag{A.16}$$

is the mixing probability of pirouette and weathervaning behavioral strategies. The state-dependent parameters are denoted with superscript z, with the maximum M, minimum pirouette probability m, concentration filter K_C, and history behavioral filter K_h. We assumed that the conditional probability of speed and heading change to be independent: $P(dr, d\theta) = P(dr)P(d\theta)$, for simplicity. The speed is modeled with Gamma distribution $P(dr|a, b)$ with shape a and scale b parameters conditioning on pirouette and weathervaning strategies q. The heading change densities follow:

$$P_{\text{pr}}(d\theta) = \alpha^z U[-\pi, \pi] + (1 - \alpha^z) f(\pi, \kappa^z_{\text{pr}}), \tag{A.17}$$

$$P_{\text{wv}}(d\theta \mid dC^{\perp}_{1:t-1}) = f(-K^z_{dC^{\perp}} \cdot \mathbf{dC}^{\perp}_{1:t-1}, \kappa^z_{\text{wv}}) \tag{A.18}$$

where α is the mixture weight on uniform distribution U, f is the von Mises distribution with mean and variance κ parameters. The filter $K_{dC^{\perp}}$ acts on the perpendicular concentration change to steer the heading during weathervaning.

In addition to the emissions that are concentration dependent (equation A.15), the state transitions of z is also input-driven, with transition probability following a soft-max form:

$$P(z_{t+1} = j|z_t = i, C_{1:t-1}) = \frac{\exp(F_{ij} \cdot \mathbf{C}_{1:t-1} + b_i)}{Z} \tag{A.19}$$

where the filter acting on concentration time series that drive transitions from state i to j is denoted as F_{ij}, the baseline probability of self-transition is b_i, and Z is the normalization constant. For optogenetic experiments, we added in terms $O_{ij} \cdot \mathbf{I}_{1:t-1}$ in the state-transition probability (equation A.19) and $K_O \cdot \mathbf{I}_{1:t-1}$ in the pirouette decision function (equation A.16). The state-transition filters and pirouette filters for optogenetic input are denoted as O and K_O, respectively. For inference, the experimental data for fitting is $D = \{C, dC^{\perp}, d\theta, dr, I\}$ and the history components with time lag is denoted H. The objective for fitting staPAW is the log-likelihood that sums over latent states. The joint distribution of data and latent is the complete-data log-likelihood with all parameters Θ:

$$\text{ECLL}(\Theta) = \Sigma_z P(z|D, \Theta^{old}) \log(P(dr, d\theta, z|H, \Theta)) \tag{A.20}$$

We optimized this objective through expectation-maximization (EM) algorithm. The E-step applies forward-backward algorithm to estimate the marginal and joint posterior state probabilities and the M-step uses the state estimates to compute and maximize the ECLL. In practice, we initialized parameters in the EM runs around the maximum likelihood estimate of a single state (dPAW model). We repeated different insinuations for 5 times for each of the 3-fold cross validation training set and choose the parameter with the highest training likelihood for testing in Fig 4.2. Due to the inclusion of strictly-positive parameters, such as pirouette probabilities, variance, and shape parameters, we performed constrained optimization during the M-steps. We empirically found that EM converges around 20-30 iterations.

A.3.2 Reinforcement learning model

The RL model constrained by worm navigation takes actions that correspond to the behavioral states fitted in staPAW. Here we simplify the model by using constant speed and state-dependent heading changes. The policy takes in a local concentration difference measurement $\Delta C(s)$ at location s across a single time step and makes the decision following:

$$\pi_\theta(a = 1|\Delta C(s)) = \frac{1}{1 + \exp(K\Delta C(s) + b)} \tag{A.21}$$

where K is the scalar weight on concentration difference $\Delta C(s)$ and b is a baseline. These are the model parameters θ. This policy is similar to the soft-max state transition in staPAW (equation A.19) but simplified for concentration difference in a single step. Given the action that corresponds to state-dependent kinematics (speed and heading change), the worm navigates to the next location in the environment $P(s'|s, a, \Delta C)$. Note that this model setting forms a partially observed Markov decision process (POMDP), since the policy only takes in a local concentration difference in the environment, rather than the actual states in the environment. The POMDP setting is similar to formulation in previous navigation RL models[251,208].

The reward is a function of the state and action, where state is the coordinates in the environment s and action is the discrete output from equation A.21. The function is a combination of the negative distant to source s^* and a filtered action $\phi(a)$ weighted by scalar λ:

$$r(s, a) = -d(s, s^*) - \lambda\phi(a) \tag{A.22}$$

where the function ϕ is an exponential filter on the discrete actions to penalize behavioral switching. This term is added to qualitatively recapitulate the experimentally observed state occupancy of behavioral states.

The parameters in the policy is optimized through policy gradient method (Fig 4.7):

$$\theta' \leftarrow \theta + \eta\nabla J, \tag{A.23}$$

$$J(\theta) = \mathbb{E}_{\pi_\theta}[\Sigma_{t=1}^{T}\gamma^t r^t] \tag{A.24}$$

where η is the learning rate and γ is the temporal discount factor used to compute the expected total reward across time J. In practice, we used Adam optimizer to adjust the learning rate during gradient descent and run for 1000 epochs and observed convergence of the expected reward. In

the same model structure, we shuffled actions to make the dPAW model, and fixed the action at 0 or 1 for weathervaning or pirouette-only models, respectively.

A.3.3 Neural circuit model

The effective circuit model has neural activity following:

$$\tau \frac{dx_i}{dt} = -x_i + \Sigma_j J_{ij} \phi(x_i) + B_i I_i(t) + \xi_i \tag{A.25}$$

where x_i is the voltage of neuron i, with time constant τ, circuit connection J_{ij}, input weight B_i, and time-varying white noise ξ_i. The input $I_i(t) = F(t) - S(t)$ is from the fast and slow signal processing concentration input $C(t)$ of the sensory environment experienced by the agent:

$$\frac{dF}{dt} = \alpha C - \beta F, \tag{A.26}$$

$$\frac{dS}{dt} = \gamma (F - S) \tag{A.27}$$

with coefficients α, β, and γ governing the concentration processing. In the effective circuit model, there are two units $i \in \{S, T\}$ corresponding to S and T states.Both units have the same parameters and only differ in the coupling strength J_{ij}. The agent employs behavioral strategy from the state that has a higher activity at any given time point. In practice, we simulated neural dynamics and navigation behavior with Euler method with $dt = 0.1$ s and measured dwell time and chemotaxis index from 200 trajectories.

A.4 Alternative formalism for state-space dynamics in RNNs

We have proposed algorithms to enforce state-switching behavior in RNNs in chapter 5. Here we derive another objective by incorporating a winner-take-all (WTA) readout for the state, with online update rule:

$$z_k(t+\Delta) = z_k(t) - \eta_z[e_k^2(t) - z_k(t) + T\log(z_k(t)) + n(t)] \tag{A.28}$$

where the evolution of latent state variable of state k is z_k that evolves with rate η_z, with state-dependent error e_k, and normalization term n_t:

$$n(t+\Delta) = n(t) - \eta_n(1 - \Sigma_k z_k(t)) \tag{A.29}$$

This form was derived from the Chklovskii group modeling biological circuits for WTA computation[245]. With this setup, the objective function can be written as:

$$\mathcal{L}_{wta} = \Sigma_t \Sigma_k z_k(t)(W_k^T r(t) - f(t))^2 \tag{A.30}$$

where z_k can now be the one-hot encoding for states $z_k = \delta_{kk^*(t)}$, with $k^* = \text{argmax}_k \mathcal{L}_{wta}$. Together, this would be an alternative way of training the network with state-dependent linear readout weights W_k.

Another thread of investigation would be to change the Gaussian emission class, for instance to Poisson target spiking patterns Y_t. In this case, the objective naturally extends as:

$$\mathcal{L}_{Poisson-ssm} = -\Sigma_t \Sigma_k [Y_t(W_k^T r_t) - \exp(W_k^T r_t) + \log P(z_t = k | z_{t-1}, r_t; W)] \tag{A.31}$$

now the corresponding update rule becomes:

$$W_k(t+\Delta) = W_k(t) - \hat{P}_k(t) r(t)[Y(t) - \exp(W^T r(t)) + 1 - P(z(t)|z(t-1), r(t))] \tag{A.32}$$

and the online approximation for inverse Hessian is:

$$\hat{P}_k(t+\Delta) = \hat{P}_k(t) - \frac{\hat{P}_k(t-\Delta) r(t) r(t)^T \hat{P}_k(t-\Delta)}{\gamma_k(t)^{-1} + r(t)^T \hat{P}_k(t-\Delta) r(t)} \tag{A.33}$$

with the diagonal term:

$$\gamma_k(t) = -\exp(W_k^T r(t)) - P(z_t = k)(1 - P(z_t = k)) \tag{A.34}$$

In addition, depending on the generative model, one can change the latent state dynamics to conditionally Gaussian form, forming a linear dynamical system. The objective function would become:

$$\mathcal{L}_{LDS} = \Sigma_t \frac{1}{2\sigma_o^2}(W^T r(t) - f(t))^2 + \frac{1}{2\sigma_l^2}(r(t) - J^T r(t))^2 \tag{A.35}$$

where we assume that the linear dynamics is approximated by the network itself, and here it is simplified with noise strength σ_o and σ_l for the observation and latent. Altogether, we propose alternative candidate framework to flexibly train chaotic RNNs to produce target patterns with state-space dynamics. More work on numerically investigating the methods and applying them to data would be promising future directions.

A.5 Repository and resource for data and methods

For chapter 2, the code, measurements, and user manual for the odor flow chamber are published publicly:

- Acquisition software that interfaces with the odor sensors:https://github.com/GershowLab/OdorSensorArraycopyarchivedatGershowLab
- Demonstration code that regenerates many of the figures based on the data in 10.6084/m9.figshare.21737303
- Acquisition software for C. elegans imaging and associated analysis: https://github.com/Kevin-Sean-Chen/leifer-Behavior-Triggered-Averaging-Tracker-new
- Acquisition software for Drosophila larvae: https://github.com/GershowLab/Chen-Wu-eLife-Drosophila-Image-Capture

For chapter 3, the code and data for the olfactory learning project can be found:

- Data to generate all figures in this work: 10.6084/m9.figshare.24764403
- Code used for analysis and to generate figures: https://github.com/Kevin-Sean-Chen/Chemotaxis_function

For chapter 4, the code and data for the state-switching navigation project can be found:

- Code used for analysis and to generate figures: https://github.com/Kevin-Sean-Chen/Chemotaxis_function
- Code for reinforcement learning model: https://github.com/Kevin-Sean-Chen/wormRL
- Data to generate all figures in this work: 10.6084/m9.figshare.25996498

For chapter 5 and 6, a number of independent projects have their own repositories.

- Code for circuit models for worm behavior: https://github.com/Kevin-Sean-Chen/CTRNN_PSO
- Code for G-FORCE, stochastic RNN training, and fitting to state-space models: https://github.com/Kevin-Sean-Chen/GLM_RNN
- Code for TMAC motion correction: https://github.com/Nondairy-Creamer/tmac
- Code for Max Cal inference: https://github.com/Kevin-Sean-Chen/MaxCal_network

References

[1] Aceves-Piña, E. O. & Quinn, W. G. (1979). Learning in normal and mutant drosophila larvae. *Science*, 206(4414), 93–96.

[2] Ahamed, T., Costa, A. C., & Stephens, G. J. (2021). Capturing the continuous complexity of behaviour in caenorhabditis elegans. *Nat. Phys.*

[3] Albrecht, D. R. & Bargmann, C. I. (2011). High-content behavioral analysis of caenorhabditis elegans in precise spatiotemporal chemical environments. *Nat. Methods*, 8(7), 599–605.

[4] Anderson, P. W. (1972). More Is Different. *Science*, 177(4047), 393–396.

[5] Aoi, M. C. & Pillow, J. W. (2017). Scalable bayesian inference for high-dimensional neural receptive fields. *bioRxiv*, (pp. 212217).

[6] Appleby, P. A. (2012). A model of chemotaxis and associative learning in c. elegans. *Biological cybernetics*, 106, 373–387.

[7] Ashwood, Z. C., Roy, N. A., Stone, I. R., International Brain Laboratory, Urai, A. E., Churchland, A. K., Pouget, A., & Pillow, J. W. (2022). Mice alternate between discrete strategies during perceptual decision-making. *Nat. Neurosci.*, 25(2), 201–212.

[8] Atanas, A. A., Kim, J., Wang, Z., Bueno, E., Becker, M., Kang, D., Park, J., Kramer, T. S., Wan, F. K., Baskoylu, S., Dag, U., Kalogeropoulou, E., Gomes, M. A., Estrem, C., Cohen, N., Mansinghka, V. K., & Flavell, S. W. (2023). Brain-wide representations of behavior spanning multiple timescales and states in c. elegans. *Cell*, 186(19), 4134–4151.e31.

[9] Atick, J. J. & Redlich, A. N. (1992). What does the retina know about natural scenes? *Neural computation*, 4(2), 196–210.

[10] Baker, K. L., Dickinson, M., Findley, T. M., Gire, D. H., Louis, M., Suver, M. P., Verhagen, J. V., Nagel, K. I., & Smear, M. C. (2018). Algorithms for olfactory search across species. *Journal of Neuroscience*, 38(44), 9383–9389.

[11] Bargmann, C. I. (2006). Chemosensation in c. elegans. *WormBook*, (pp. 1–29).

[12] Bargmann, C. I., Hartwieg, E., & Horvitz, H. R. (1993). Odorant-selective genes and neurons mediate olfaction in c. elegans. *Cell*, 74(3), 515–527.

[13] Bargmann, C. I. & Horvitz, H. R. (1991). Chemosensory neurons with overlapping functions direct chemotaxis to multiple chemicals in c. elegans. *Neuron*, 7(5), 729–742.

[14] Bartumeus, F., Campos, D., Ryu, W. S., Lloret-Cabot, R., Méndez, V., & Catalan, J. (2016). Foraging success under uncertainty: search tradeoffs and optimal space use. *Ecol. Lett.*, 19(11), 1299–1313.

[15] Beer, R. D. (1995). On the dynamics of small continuous-time recurrent neural networks. *Adaptive Behavior*, 3(4), 469–509.

[16] Bellec, G., Wang, S., Modirshanechi, A., Brea, J., & Gerstner, W. (2021). Fitting summary statistics of neural data with a differentiable spiking network simulator. *Advances in Neural Information Processing Systems*, 34, 18552–18563.

[17] Benucci, A., Saleem, A. B., & Carandini, M. (2013). Adaptation maintains population homeostasis in primary visual cortex. *Nature neuroscience*, 16(6), 724–729.

[18] Berg, H. C. (1975). Chemotaxis in bacteria. *Annual review of biophysics and bioengineering*, 4(1), 119–136.

[19] Berg, H. C. (1993). *Random walks in biology*. Princeton University Press.

[20] Berg, H. C. & Brown, D. A. (1972). Chemotaxis in escherichia coli analysed by three-dimensional tracking. *nature*, 239(5374), 500–504.

[21] Berman, G. J., Bialek, W., & Shaevitz, J. W. (2016). Predictability and hierarchy in drosophila behavior. *Proc. Natl. Acad. Sci. U. S. A.*, 113(42), 11943–11948.

[22] Boerlin, M., Machens, C. K., & Denève, S. (2013). Predictive coding of dynamical variables in balanced spiking networks. *PLoS computational biology*, 9(11), e1003258.

[23] Boie, S. D., Connor, E. G., McHugh, M., Nagel, K. I., Ermentrout, G. B., Crimaldi, J. P., & Victor, J. D. (2018). Information-theoretic analysis of realistic odor plumes: What cues are useful for determining location? *PLoS Comput. Biol.*, 14(7), e1006275.

[24] Bolkan, S. S., Stone, I. R., Pinto, L., Ashwood, Z. C., Iravedra Garcia, J. M., Herman, A. L., Singh, P., Bandi, A., Cox, J., Zimmerman, C. A., et al. (2022). Opponent control of behavior by dorsomedial striatal pathways depends on task demands and internal state. *Nature neuroscience*, 25(3), 345–357.

[25] Buchanan, E. K., Lipschitz, A., Linderman, S. W., & Paninski, L. (2017). Quantifying the behavioral dynamics of c. elegans with autoregressive hidden markov models. In *Workshop on Worm's neural information processing at the 31st conference on neural information processing systems*.

[26] Burgués, J., Hernández, V., Lilienthal, A. J., & Marco, S. (2020). Gas distribution mapping and source localization using a 3d grid of metal oxide semiconductor sensors. *Sensors and Actuators B: Chemical*, 304, 127309.

[27] Calhoun, A. J., Chalasani, S. H., & Sharpee, T. O. (2014). Maximally informative foraging by caenorhabditis elegans. *Elife*, 3, e04220.

[28] Calhoun, A. J. & Murthy, M. (2017). Quantifying behavior to solve sensorimotor transformations: advances from worms and flies. *Current opinion in neurobiology*, 46, 90–98.

[29] Calhoun, A. J., Pillow, J. W., & Murthy, M. (2019). Unsupervised identification of the internal states that shape natural behavior. *Nat. Neurosci.*, 22(12), 2040–2049.

[30] Carandini, M. & Ferster, D. (1997). A tonic hyperpolarization underlying contrast adaptation in cat visual cortex. *Science*, 276(5314), 949–952.

[31] Castro-Alamancos, M. A. & Connors, B. W. (1996). Short-term synaptic enhancement and long-term potentiation in neocortex. *Proceedings of the National Academy of Sciences*, 93(3), 1335–1339.

[32] Caticha, A. (2021). Entropy, Information, and the Updating of Probabilities. *Entropy*, 23(7), 895.

[33] Celani, A., Villermaux, E., & Vergassola, M. (2014). Odor landscapes in turbulent environments. *Physical Review X*, 4(4), 041015.

[34] Cessac, B. (2019). Linear response in neuronal networks: from neurons dynamics to collective response.

[35] Çetiner, U., Raz, O., Sukharev, S., & Jarzynski, C. (2020). Recovery of equilibrium free energy from nonequilibrium thermodynamics with mechanosensitive ion channels in e. coli. *Phys. Rev. Lett.*, 124(22), 228101.

[36] Chalasani, S. H., Chronis, N., Tsunozaki, M., Gray, J. M., Ramot, D., Goodman, M. B., & Bargmann, C. I. (2007). Dissecting a circuit for olfactory behaviour in caenorhabditis elegans. *Nature*, 450(7166), 63–70.

[37] Chalasani, S. H., Kato, S., Albrecht, D. R., Nakagawa, T., Abbott, L. F., & Bargmann, C. I. (2010). Neuropeptide feedback modifies odor-evoked dynamics in caenorhabditis elegans olfactory neurons. *Nat. Neurosci.*, 13(5), 615–621.

[38] Chalfie, M. & Au, M. (1989). Genetic control of differentiation of the caenorhabditis elegans touch receptor neurons. *Science*, 243(4894), 1027–1033.

[39] Chalfie, M., Sulston, J. E., White, J. G., Southgate, E., Thomson, J. N., & Brenner, S. (1985). The neural circuit for touch sensitivity in caenorhabditis elegans. *Journal of Neuroscience*, 5(4), 956–964.

[40] Chandra, R., Farah, F., Muñoz-Lobato, F., Bokka, A., Benedetti, K. L., Brueggemann, C., Saifuddin, M. F. A., Miller, J. M., Li, J., Chang, E., Varshney, A., Jimenez, V., Baradwaj, A., Nassif, C., Alladin, S., Andersen, K., Garcia, A. J., Bi, V., Nordquist, S. K., Dunn, R. L., Garcia, V., Tokalenko, K., Soohoo, E., Briseno, F., Kaur, S., Harris, M., Guillen, H., Byrd, D., Fung, B., Bykov, A. E., Odisho, E., Tsujimoto, B., Tran, A., Duong, A., Daigle,

K. C., Paisner, R., Zuazo, C. E., Lin, C., Asundi, A., Churgin, M. A., Fang-Yen, C., Bremer, M., Kato, S., VanHoven, M. K., & L'Étoile, N. D. (2023). Sleep is required to consolidate odor memory and remodel olfactory synapses. *Cell*, 186(13), 2911–2928.e20.

[41] Chang, H.-J., Huang, K.-S., & Huang, K.-T. (1989). 'entropy production rate'and'entropy'for neural networks. In *[1989] Proceedings of the Thirteenth Annual International Computer Software & Applications Conference* (pp. 658–663).: ieeexplore.ieee.org.

[42] Chen, E. Y., Chou, J., Park, J., Schwartz, G., & Berry, M. J. (2014). The neural circuit mechanisms underlying the retinal response to motion reversal. *Journal of Neuroscience*, 34(47), 15557–15575.

[43] Chen, K. S. (2020). Nonequilibrium thermodynamics of input-driven networks. *arXiv preprint arXiv:2012.13252*.

[44] Chen, K. S., Chen, C.-C., & Chan, C. (2017). Characterization of predictive behavior of a retina by mutual information. *Frontiers in computational neuroscience*, 11, 66.

[45] Chen, K. S., Sharma, A. K., Pillow, J. W., & Leifer, A. M. (2024). Olfactory learning alters navigation strategies and behavioral variability in c. elegans. *ArXiv*.

[46] Chen, K. S., Wu, R., Gershow, M. H., & Leifer, A. M. (2023). Continuous odor profile monitoring to study olfactory navigation in small animals. *Elife*, 12.

[47] Chen, K. S. & Yang, Y.-J. (2024). Maximum caliber infers effective coupling and response from spiking networks.

[48] Chen, T.-W., Wardill, T. J., Sun, Y., Pulver, S. R., Renninger, S. L., Baohan, A., Schreiter, E. R., Kerr, R. A., Orger, M. B., Jayaraman, V., et al. (2013). Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature*, 499(7458), 295–300.

[49] Chetrite, R. & Touchette, H. (2015). Variational and optimal control representations of conditioned and driven processes. *J. Stat. Mech.*, 2015(12), P12001.

[50] Cho, C. E., Brueggemann, C., L'Etoile, N. D., & Bargmann, C. I. (2016). Parallel encoding of sensory history and behavioral preference during caenorhabditis elegans olfactory learning. *eLife Sciences*, 5, e14000.

[51] Chronis, N., Zimmer, M., & Bargmann, C. I. (2007). Microfluidics for in vivo imaging of neuronal and behavioral activity in caenorhabditis elegans. *Nature methods*, 4(9), 727–731.

[52] Clark, D. A., Freifeld, L., & Clandinin, T. R. (2013). Mapping and cracking sensorimotor circuits in genetic model organisms. *Neuron*, 78(4), 583–595.

[53] Clark, D. A., Gabel, C. V., Gabel, H., & Samuel, A. D. (2007). Temporal activity patterns in thermosensory neurons of freely moving caenorhabditis elegans encode spatial thermal gradients. *Journal of Neuroscience*, 27(23), 6083–6090.

[54] Colbert, H. A. & Bargmann, C. I. (1995). Odorant-specific adaptation pathways generate olfactory plasticity in c. elegans. *Neuron*, 14(4), 803–812.

[55] Collin, D., Ritort, F., Jarzynski, C., Smith, S. B., Tinoco, Jr, I., & Bustamante, C. (2005). Verification of the crooks fluctuation theorem and recovery of RNA folding free energies. *Nature*, 437(7056), 231–234.

[56] Cook, S. J., Jarrell, T. A., Brittin, C. A., Wang, Y., Bloniarz, A. E., Yakovlev, M. A., Nguyen, K. C. Q., Tang, L. T.-H., Bayer, E. A., Duerr, J. S., Bülow, H. E., Hobert, O., Hall, D. H., & Emmons, S. W. (2019). Whole-animal connectomes of both caenorhabditis elegans sexes. *Nature*, 571(7763), 63–71.

[57] Costa, A. C., Ahamed, T., Jordan, D., & Stephens, G. J. (2023). A markovian dynamics for c. elegans behavior across scales. *bioRxiv*.

[58] Costa, A. C., Ahamed, T., & Stephens, G. J. (2019). Adaptive, locally linear models of complex dynamics. *Proceedings of the National Academy of Sciences*, 116(5), 1501–1510.

[59] Creamer, M. S., Chen, K. S., Leifer, A. M., & Pillow, J. W. (2022). Correcting motion induced fluorescence artifacts in two-channel neural imaging. *PLoS computational biology*, 18(9), e1010421.

[60] Crisanti, A. & Sompolinsky, H. (1987). Dynamics of spin systems with randomly asymmetric bonds: Langevin dynamics and a spherical model. *Phys. Rev. A Gen. Phys.*, 36(10), 4922–4939.

[61] Croll, N. A. (1977). Sensory mechanisms in nematodes. *Annual Review of Phytopathology*, 15(1), 75–89.

[62] Crooks, G. E. (1998). Nonequilibrium measurements of free energy differences for microscopically reversible markovian systems. *J. Stat. Phys.*

[63] Crooks, G. E. (1999). Entropy production fluctuation theorem and the nonequilibrium work relation for free energy differences. *Phys. Rev. E Stat. Phys. Plasmas Fluids Relat. Interdiscip. Topics*, 60(3), 2721–2726.

[64] Csiszar, I., Cover, T., & Byoung-Seon Choi (1987). Conditional limit theorems under Markov conditioning. *IEEE Trans. Inform. Theory*, 33(6), 788–801.

[65] Dahlberg, B. A. & Izquierdo, E. J. (2020). Contributions from parallel strategies for spatial orientation in c. elegans. *ALIFE 2020: The 2020 Conference on*.

[66] Das, A. & Fiete, I. R. (2020). Systematic errors in connectivity inferred from activity in strongly recurrent networks. *Nature Neuroscience*, 23(10), 1286–1296.

[67] Datta, S. R., Anderson, D. J., Branson, K., Perona, P., & Leifer, A. (2019). Computational neuroethology: a call to action. *Neuron*, 104(1), 11–24.

[68] Demir, M., Kadakia, N., Anderson, H. D., Clark, D. A., & Emonet, T. (2020). Walking drosophila navigate complex plumes using stochastic decisions biased by the timing of odor encounters. *Elife*, 9, e57524.

[69] Dennler, N., Rastogi, S., Fonollosa, J., Van Schaik, A., & Schmuker, M. (2022). Drift in a popular metal oxide sensor dataset reveals limitations for gas classification benchmarks. *Sensors and Actuators B: Chemical*, 361, 131668.

[70] DePasquale, B., Cueva, C. J., Rajan, K., Escola, G. S., & Abbott, L. (2018). full-force: A target-based method for training recurrent networks. *PloS one*, 13(2), e0191527.

[71] Dhruv, N. T. & Carandini, M. (2014). Cascaded effects of spatial adaptation in the early visual system. *Neuron*, 81(3), 529–535.

[72] Dombeck, D. A., Harvey, C. D., Tian, L., Looger, L. L., & Tank, D. W. (2010). Functional imaging of hippocampal place cells at cellular resolution during virtual navigation. *Nature neuroscience*, 13(11), 1433–1440.

[73] Drix, D. & Schmuker, M. (2021). Resolving fast gas transients with metal oxide sensors. *ACS sensors*, 6(3), 688–692.

[74] Dubin, A. E., Heald, N. L., Cleveland, B., Carlson, J. R., & Harris, G. L. (1995). Scutoid mutation of drosophila melanogaster specifically decreases olfactory responses to short-chain acetate esters and ketones. *Journal of neurobiology*, 28(2), 214–233.

[75] Dubin, A. E., Liles, M. M., Seligman, F., Le, T., Tolli, J., & Harris, G. L. (1998). Involvement of genes encoding a k+ channel (ether a go-go) and a na+ channel (smellblind) in drosophila olfaction a. *Annals of the New York Academy of Sciences*, 855(1), 212–222.

[76] Dunn, N. A., Conery, J. S., & Lockery, S. R. (2007). Circuit motifs for spatial orientation behaviors identified by neural network optimization. *Journal of neurophysiology*, 98(2), 888–897.

[77] Dusenbery, D. B., Sheridan, R. E., & Russell, R. L. (1975). Chemotaxis-defective mutants of the nematode caenorhabditis elegans. *Genetics*, 80(2), 297–309.

[78] Ferrari, U., Deny, S., Chalk, M., Tkacik, G., Marre, O., & Mora, T. (2018). Separating intrinsic interactions from extrinsic correlations in a network of sensory neurons.

[79] Flavell, S. W., Pokala, N., Macosko, E. Z., Albrecht, D. R., Larsch, J., & Bargmann, C. I. (2013). Serotonin and the neuropeptide PDF initiate and extend opposing behavioral states in c. elegans. *Cell*, 154(5), 1023–1035.

[80] Friston, K. (2010). The free-energy principle: a unified brain theory? *Nat. Rev. Neurosci.*, 11, 127.

[81] Gallagher, T., Bjorness, T., Greene, R., You, Y.-J., & Avery, L. (2013). The geometry of locomotive behavioral states in c. elegans. *PloS one*, 8(3), e59865.

[82] Ge, H., Pressé, S., Ghosh, K., & Dill, K. A. (2012). Markov processes follow from the principle of maximum caliber. *The Journal of Chemical Physics*, 136(6), 064108.

[83] Gepner, R., Mihovilovic Skanata, M., Bernat, N. M., Kaplow, M., & Gershow, M. (2015). Computations underlying drosophila photo-taxis, odor-taxis, and multi-sensory integration. *Elife*, 4.

[84] Gerhard, F., Kispersky, T., Gutierrez, G. J., Marder, E., Kramer, M., & Eden, U. (2013). Successful reconstruction of a physiological circuit with known connectivity from spiking activity alone. *PLoS computational biology*, 9(7), e1003138.

[85] Gershow, M., Berck, M., Mathew, D., Luo, L., Kane, E. A., Carlson, J. R., & Samuel, A. D. (2012). Controlling airborne cues to study small animal navigation. *Nature methods*, 9(3), 290–296.

[86] Gerstner, W., Kistler, W. M., Naud, R., & Paninski, L. (2014). *Neuronal Dynamics: From Single Neurons to Networks and Models of Cognition*. Cambridge University Press, 1st edition edition.

[87] Gire, D. H., Kapoor, V., Arrighi-Allisan, A., Seminara, A., & Murthy, V. N. (2016). Mice develop efficient strategies for foraging and navigation using complex natural stimuli. *Current Biology*, 26(10), 1261–1273.

[88] Goldt, S. & Seifert, U. (2017). Stochastic thermodynamics of learning. *Phys. Rev. Lett.*, 118(1), 010601.

[89] Gomez-Marin, A. & Louis, M. (2014). Multilevel control of run orientation in drosophila larval chemotaxis. *Frontiers in behavioral neuroscience*, 8, 38.

[90] Gomez-Marin, A., Stephens, G. J., & Louis, M. (2011). Active sampling and decision making in drosophila chemotaxis. *Nat. Commun.*, 2, 441.

[91] Gordus, A., Pokala, N., Levy, S., Flavell, S. W., & Bargmann, C. I. (2015). Feedback from network states generates variability in a probabilistic olfactory circuit. *Cell*, 161(2), 215–227.

[92] Gorur-Shandilya, S., Martelli, C., Demir, M., & Emonet, T. (2019). Controlling and measuring dynamic odorant stimuli in the laboratory. *J. Exp. Biol.*, 222(Pt 23).

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

[94] Grill-Spector, K., Henson, R., & Martin, A. (2006). Repetition and the brain: neural models of stimulus-specific effects. *Trends in cognitive sciences*, 10(1), 14–23.

[95] Ha, H.-I., Hendricks, M., Shen, Y., Gabel, C. V., Fang-Yen, C., Qin, Y., Colón-Ramos, D., Shen, K., Samuel, A. D. T., & Zhang, Y. (2010). Functional organization of a neural network for aversive olfactory learning in caenorhabditis elegans. *Neuron*, 68(6), 1173–1186.

[96] Haber, A. & Schneidman, E. (2022). Learning the architectural features that predict functional similarity of neural networks. *Physical Review X*, 12(2), 021051.

[97] Hallinen, K. M., Dempsey, R., Scholz, M., Yu, X., Linder, A., Randi, F., Sharma, A. K., Shaevitz, J. W., & Leifer, A. M. (2021). Decoding locomotion from population neural activity in moving c. elegans. *Elife*, 10, e66135.

[98] Hendricks, M., Ha, H., Maffey, N., & Zhang, Y. (2012). Compartmentalized calcium dynamics in a c. elegans interneuron encode head movement. *Nature*, 487(7405), 99–103.

[99] Hernandez-Nunez, L., Belina, J., Klein, M., Si, G., Claus, L., Carlson, J. R., & Samuel, A. D. (2015). Reverse-correlation analysis of navigation dynamics in drosophila larva using optogenetics. *Elife*, 4, e06225.

[100] Hertz, J. A. (2018). *Introduction to the theory of neural computation*.

[101] Hinton, G. E. (2012). A practical guide to training restricted boltzmann machines. *Neural networks: Tricks of the trade*.

[102] Homann, J., Koay, S. A., Chen, K. S., Tank, D. W., & Berry, M. J. (2022). Novel stimuli evoke excess activity in the mouse primary visual cortex. *Proceedings of the National Academy of Sciences*, 119(5), e2108882119.

[103] Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci. U. S. A.*, 79(8), 2554–2558.

[104] Hopfield, J. J. (1984). Neurons with graded response have collective computational properties like those of two-state neurons. *Proc. Natl. Acad. Sci. U. S. A.*, 81(10), 3088–3092.

[105] Horowitz, J. M. & Esposito, M. (2014). Thermodynamics with Continuous Information Flow. *Phys. Rev. X*, 4(3), 031015.

[106] Hosoya, T., Baccus, S. A., & Meister, M. (2005). Dynamic predictive coding by the retina. *Nature*, 436(7047), 71–77.

[107] Hubel, D. H. & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of physiology*, 160(1), 106.

[108] Huh, D. & Sejnowski, T. J. (2018). Gradient descent for spiking neural networks. *Advances in neural information processing systems*, 31.

[109] Hulsey, D., Zumwalt, K., Mazzucato, L., McCormick, D. A., & Jaramillo, S. (2024). Decision-making dynamics are predicted by arousal and uninstructed movements. *Cell Reports*, 43(2).

[110] Iino, Y. & Yoshida, K. (2009). Parallel use of two behavioral mechanisms for chemotaxis in caenorhabditis elegans. *J. Neurosci.*, 29(17), 5370–5380.

[111] Ikeda, M., Nakano, S., Giles, A. C., Xu, L., Costa, W. S., Gottschalk, A., & Mori, I. (2020). Context-dependent operation of neural circuits underlies a navigation behavior in caenorhabditis elegans. *Proceedings of the National Academy of Sciences*, 117(11), 6178–6188.

[112] Inoue, K., Nakajima, K., & Kuniyoshi, Y. (2020). Designing spontaneous behavioral switching via chaotic itinerancy. *Science advances*, 6(46), eabb3989.

[113] Israel, S., Rozenfeld, E., Weber, D., Huetteroth, W., & Parnas, M. (2022). Olfactory stimuli and moonwalker sez neurons can drive backward locomotion in drosophila. *Current Biology*, 32(5), 1131–1149.

[114] Itskovits, E., Ruach, R., Kazakov, A., & Zaslaver, A. (2018). Concerted pulsatile and graded neural dynamics enables efficient chemotaxis in c. elegans. *Nature communications*, 9(1), 1–11.

[115] Izquierdo, E. J. & Lockery, S. R. (2010). Evolution and analysis of minimal neural circuits for klinotaxis in caenorhabditis elegans. *Journal of Neuroscience*, 30(39), 12908–12917.

[116] Izquierdo, E. J., Williams, P. L., & Beer, R. D. (2015). Information flow through a model of the c. elegans klinotaxis circuit. *PLoS One*, 10(10), e0140397.

[117] Jang, M. S., Toyoshima, Y., Tomioka, M., Kunitomo, H., & Iino, Y. (2019). Multiple sensory neurons mediate starvation-dependent aversive navigation in caenorhabditis elegans. *Proc. Natl. Acad. Sci. U. S. A.*

[118] Jarzynski, C. (1997). Nonequilibrium equality for free energy differences. *Phys. Rev. Lett.*

[119] Jarzynski, C. (2011). Equalities and inequalities: Irreversibility and the second law of thermodynamics at the nanoscale. *Annu. Rev. Condens. Matter Phys.*

[120] Jarzynski, C. & Wójcik, D. K. (2004). Classical and quantum fluctuation theorems for heat exchange. *Phys. Rev. Lett.*, 92(23), 230602.

[121] Jaynes, E. T. (1957). Information Theory and Statistical Mechanics. *Phys. Rev.*, 106(4), 620–630.

[122] Jaynes, E. T. (1980). The Minimum Entropy Production Principle. *Annual Review of Physical Chemistry*, 31(1), 579–601.

[123] Ji, N., Venkatachalam, V., Rodgers, H. D., Hung, W., Kawano, T., Clark, C. M., Lim, M., Alkema, M. J., Zhen, M., & Samuel, A. D. (2021). Corollary discharge promotes a sustained motor state in a neural circuit for navigation. *Elife*, 10.

[124] Jimenez Rezende, D. & Gerstner, W. (2014). Stochastic variational learning in recurrent spiking networks. *Frontiers in computational neuroscience*, 8, 38.

[125] Jizba, P. & Korbel, J. (2019). Maximum Entropy Principle in Statistical Inference: Case for Non-Shannonian Entropies. *Phys. Rev. Lett.*, 122(12), 120601.

[126] Jovanic, T., Winding, M., Cardona, A., Truman, J. W., Gershow, M., & Zlatic, M. (2019). Neural Substrates of Drosophila Larval Anemotaxis. *Current biology: CB*, 29(4), 554–566.e4.

[127] Jung, S.-H., Hueston, C., & Bhandawat, V. (2015). Odor-identity dependent motor programs underlie behavioral responses to odors. *Elife*, 4, e11092.

[128] Kadakia, N., Demir, M., Michaelis, B. T., DeAngelis, B. D., Reidenbach, M. A., Clark, D. A., & Emonet, T. (2022). Odour motion sensing enhances navigation of complex plumes. *Nature*, (pp. 1–8).

[129] Kaplan, H. S., Thula, O. S., Khoss, N., & Zimmer, M. (2020). Nested neuronal dynamics orchestrate a behavioral hierarchy across timescales. *Neuron*, 105(3), 562–576.

[130] Kato, S., Xu, Y., Cho, C. E., Abbott, L. F., & Bargmann, C. I. (2014). Temporal responses of c. elegans chemosensory neurons are preserved in behavioral dynamics. *Neuron*, 81(3), 616–628.

[131] Kauffman, A., Parsons, L., Stein, G., Wills, A., Kaletsky, R., & Murphy, C. (2011). C. elegans positive butanone learning, short-term, and long-term associative memory assays. *J. Vis. Exp.*, (49).

[132] Keller, E. F. & Segel, L. A. (1971). Model for chemotaxis. *Journal of theoretical biology*, 30(2), 225–234.

[133] Kennedy, J. & Eberhart, R. (1995). Particle swarm optimization. In *Proceedings of ICNN'95-international conference on neural networks*, volume 4 (pp. 1942–1948).: ieee.

[134] Kim, J. & Shlizerman, E. (2020). Deep reinforcement learning for neural control.

[135] Kim, R., Li, Y., & Sejnowski, T. J. (2019). Simple framework for constructing functional spiking recurrent neural networks. *Proceedings of the national academy of sciences*, 116(45), 22811–22820.

[136] Klein, M., Krivov, S. V., Ferrer, A. J., Luo, L., Samuel, A. D., & Karplus, M. (2017). Exploratory search during directed navigation in c. elegans and drosophila larva. *Elife*, 6, e30503.

[137] Kocabas, A., Shen, C.-H., Guo, Z. V., & Ramanathan, S. (2012). Controlling interneuron activity in caenorhabditis elegans to evoke chemotactic behaviour. *Nature*, 490(7419), 273–277.

[138] Koch, C. (2004). *Biophysics of computation: information processing in single neurons*. Oxford university press.

[139] Krakauer, J. W., Ghazanfar, A. A., Gomez-Marin, A., MacIver, M. A., & Poeppel, D. (2017). Neuroscience needs behavior: correcting a reductionist bias. *Neuron*, 93(3), 480–490.

[140] Kramer, D., Bommer, P. L., Tombolini, C., Koppe, G., & Durstewitz, D. (2021). Reconstructing nonlinear dynamical systems from multi-modal time series. *arXiv preprint arXiv:2111.02922*.

[141] Kumar, S., Sharma, A. K., & Leifer, A. (2023). Inhibitory motor signals gate mechanosensory processing in c. eleagns. *arXiv preprint arXiv:2301.02709*.

[142] Kunitomo, H., Sato, H., Iwata, R., Satoh, Y., Ohno, H., Yamada, K., & Iino, Y. (2013). Concentration memory-dependent synaptic plasticity of a taste circuit regulates salt concentration chemotaxis in caenorhabditis elegans. *Nature communications*, 4(1), 1–11.

[143] Ladenbauer, J., McKenzie, S., English, D. F., Hagens, O., & Ostojic, S. (2019). Inferring and validating mechanistic models of neural microcircuits based on spike-train data. *Nature communications*, 10(1), 4933.

[144] Landau, I. D. & Sompolinsky, H. (2018). Coherent chaos in a recurrent neural network with structured connectivity. *PLoS computational biology*, 14(12), e1006309.

[145] Larsch, J., Flavell, S. W., Liu, Q., Gordus, A., Albrecht, D. R., & Bargmann, C. I. (2015). A circuit for gradient climbing in c. elegans chemotaxis. *Cell Rep.*, 12(11), 1748–1760.

[146] Lee, J. & Pressé, S. (2012). A derivation of the master equation from path entropy maximization. *The Journal of Chemical Physics*, 137(7), 074103.

[147] Lerner, H., Rozenfeld, E., Rozenman, B., Huetteroth, W., & Parnas, M. (2020). Differential role for a defined lateral horn neuron subset in naïve odor valence in drosophila. *Scientific reports*, 10(1), 1–17.

[148] Lerner, T. N., Shilyansky, C., Davidson, T. J., Evans, K. E., Beier, K. T., Zalocusky, K. A., Crow, A. K., Malenka, R. C., Luo, L., Tomer, R., et al. (2015). Intact-brain analyses reveal distinct information carried by snc dopamine subcircuits. *Cell*, 162(3), 635–647.

[149] Lesar, A., Tahir, J., Wolk, J., & Gershow, M. (2021). Switch-like and persistent memory formation in individual drosophila larvae. *Elife*, 10.

[150] Levy, S. & Bargmann, C. I. (2020). An Adaptive-Threshold mechanism for odor sensation and animal navigation. *Neuron*, 105(3), 534–548.e13.

[151] Liang, T. & Brinkman, B. A. (2024). Statistically inferred neuronal connections in subsampled neural networks strongly correlate with spike train covariances. *Physical Review E*, 109(4), 044404.

[152] Lin, A., Qin, S., Casademunt, H., Wu, M., Hung, W., Cain, G., Tan, N. Z., Valenzuela, R., Lesanpezeshki, L., Venkatachalam, V., Pehlevan, C., Zhen, M., & Samuel, A. D. T. (2023). Functional imaging and quantification of multineuronal olfactory responses in c. elegans. *Sci Adv*, 9(9), eade1249.

[153] Linderman, S. W., Nichols, A. L. A., Blei, D. M., Zimmer, M., & Paninski, L. (2019). Hierarchical recurrent state space models reveal discrete and continuous dynamics of neural activity in c. elegans.

[154] Liu, A., Papale, A. E., Hengenius, J., Patel, K., Ermentrout, B., & Urban, N. N. (2020). Mouse navigation strategies for odor source localization. *Frontiers in Neuroscience*, 14, 218.

[155] Liu, M., Kumar, S., Sharma, A. K., & Leifer, A. M. (2022). A high-throughput method to deliver targeted optogenetic stimulation to moving c. elegans populations. *PLoS biology*, 20(1), e3001524.

[156] Liu, M., Sharma, A. K., Shaevitz, J. W., & Leifer, A. M. (2018). Temporal processing and context dependency in caenorhabditis elegans response to mechanosensation. *Elife*, 7.

[157] Lockery, S. R. (2011). The computational worm: spatial orientation and its neuronal basis in c. elegans. *Current opinion in neurobiology*, 21(5), 782–790.

[158] Lockery, S. R., Lawton, K. J., Doll, J. C., Faumont, S., Coulthard, S. M., Thiele, T. R., Chronis, N., McCormick, K. E., Goodman, M. B., & Pruitt, B. L. (2008). Artificial dirt: microfluidic substrates for nematode neurobiology and behavior. *Journal of neurophysiology*, 99(6), 3136–3143.

[159] Louis, M., Huber, T., Benton, R., Sakmar, T. P., & Vosshall, L. B. (2008). Bilateral olfactory sensory input enhances chemotaxis behavior. *Nat. Neurosci.*, 11(2), 187–199.

[160] Luo, L. (2021). Architectures of neuronal circuits. *Science*, 373(6559), eabg7285.

[161] Luo, L., Wen, Q., Ren, J., Hendricks, M., Gershow, M., Qin, Y., Greenwood, J., Soucy, E. R., Klein, M., Smith-Parker, H. K., Calvo, A. C., Colón-Ramos, D. A., Samuel, A. D. T., & Zhang, Y. (2014). Dynamic encoding of perception, memory, and movement in a c. elegans chemotaxis circuit. *Neuron*, 82(5), 1115–1128.

[162] Lynn, C. W., Cornblath, E. J., Papadopoulos, L., Bertolero, M. A., & Bassett, D. S. (2021). Broken detailed balance and entropy production in the human brain. *Proceedings of the National Academy of Sciences*, 118(47), e2109889118.

[163] Lynn, C. W., Holmes, C. M., Bialek, W., & Schwab, D. J. (2022). Decomposing the Local Arrow of Time in Interacting Systems. *Phys. Rev. Lett.*, 129(11), 118101.

[164] Macnab, R. M. & Koshland Jr, D. E. (1972). The gradient-sensing mechanism in bacterial chemotaxis. *Proceedings of the National Academy of Sciences*, 69(9), 2509–2512.

[165] Maesani, A., Ramdya, P., Cruchet, S., Gustafson, K., Benton, R., & Floreano, D. (2015). Fluctuation-driven neural dynamics reproduce drosophila locomotor patterns. *PLoS computational biology*, 11(11), e1004577.

[166] Marinari, E., Parisi, G., Ricci-Tersenghi, F., & Ruiz-Lorenzo, J. J. (1998). Violation of the fluctuation-dissipation theorem in finite-dimensional spin glasses. *J. Phys. A Math. Gen.*, 31(11), 2611.

[167] Masson, J.-B., Voisinne, G., Wong-Ng, J., Celani, A., & Vergassola, M. (2012). Noninvasive inference of the molecular chemotactic response using bacterial trajectories. *Proceedings of the National Academy of Sciences*, 109(5), 1802–1807.

[168] Mastrogiuseppe, F. & Ostojic, S. (2018). Linking connectivity, dynamics, and computations in low-rank recurrent neural networks. *Neuron*, 99(3), 609–623.

[169] Matheson, A. M. M., Lanz, A. J., Medina, A. M., Licata, A. M., Currier, T. A., Syed, M. H., & Nagel, K. I. (2022). A neural circuit for wind-guided olfactory navigation. *Nature Communications*, 13(1), 4613. Number: 1 Publisher: Nature Publishing Group.

[170] Mattingly, H., Kamino, K., Machta, B., & Emonet, T. (2021). Escherichia coli chemotaxis is information limited. *Nature Physics*, 17(12), 1426–1431.

[171] Mazzucato, L. (2022). Neural mechanisms underlying the temporal organization of naturalistic animal behavior. *Elife*, 11, e76577.

[172] Meister, M. (2022). Learning, fast and slow. *Current opinion in neurobiology*, 75, 102555.

[173] Meng, J., Ahamed, T., Yu, B., Hung, W., EI Mouridi, S., Wang, Z., Zhang, Y., Wen, Q., Boulin, T., Gao, S., et al. (2024). A tonically active master neuron modulates mutually exclusive motor states at two timescales. *Science Advances*, 10(15), eadk0002.

[174] Meyer, T. & Rust, N. C. (2018). Single-exposure visual memory judgments are reflected in inferotemporal cortex. *elife*, 7, e32259.

[175] Mi, L., Xu, R., Prakhya, S., Lin, A., Shavit, N., Samuel, A., & Turaga, S. C. (2021). Connectome-constrained latent variable model of whole-brain neural activity. In *International Conference on Learning Representations*.

[176] Milo, R., Shen-Orr, S., Itzkovitz, S., Kashtan, N., Chklovskii, D., & Alon, U. (2002). Network motifs: simple building blocks of complex networks. *Science*, 298(5594), 824–827.

[177] Monte, P., Woodard, C., Ayer, R., Lilly, M., Sun, H., & Carlson, J. (1989). Characterization of the larval olfactory response indrosophila and its genetic basis. *Behavior genetics*, 19(2), 267–283.

[178] Monteforte, M. & Wolf, F. (2012). Dynamic flux tubes form reservoirs of stability in neuronal circuits. *Phys. Rev. X*, 2(4), 041007.

[179] Mora, T., Deny, S., & Marre, O. (2015). Dynamical Criticality in the Collective Activity of a Population of Retinal Neurons. *Phys. Rev. Lett.*, 114(7), 078105.

[180] Morcos, F., Pagnani, A., Lunt, B., Bertolino, A., Marks, D. S., Sander, C., Zecchina, R., Onuchic, J. N., Hwa, T., & Weigt, M. (2011). Direct-coupling analysis of residue coevolution captures native contacts across many protein families. *Proceedings of the National Academy of Sciences*, 108(49), E1293–E1301.

[181] Musall, S., Kaufman, M. T., Juavinett, A. L., Gluf, S., & Churchland, A. K. (2019). Single-trial neural dynamics are dominated by richly varied movements. *Nature neuroscience*, 22(10), 1677–1686.

[182] Näätänen, R., Paavilainen, P., Rinne, T., & Alho, K. (2007). The mismatch negativity (mmn) in basic research of central auditory processing: a review. *Clinical neurophysiology*, 118(12), 2544–2590.

[183] Nguyen, J. P., Linder, A. N., Plummer, G. S., Shaevitz, J. W., & Leifer, A. M. (2017). Automatically tracking neurons in a moving and deforming brain. *PLoS computational biology*, 13(5), e1005517.

[184] Nguyen, J. P., Shipley, F. B., Linder, A. N., Plummer, G. S., Liu, M., Setru, S. U., Shaevitz, J. W., & Leifer, A. M. (2016). Whole-brain calcium imaging with cellular resolution in freely behaving caenorhabditis elegans. *Proceedings of the National Academy of Sciences*, 113(8), E1074–E1081.

[185] Nicola, W. & Clopath, C. (2017). Supervised learning in spiking neural networks with force training. *Nature communications*, 8(1), 2208.

[186] Niell, C. M. & Stryker, M. P. (2008). Highly selective receptive fields in mouse visual cortex. *Journal of Neuroscience*, 28(30), 7520–7536.

[187] Omura, D. T., Clark, D. A., Samuel, A. D. T., & Horvitz, H. R. (2012). Dopamine signaling is essential for precise rates of locomotion by c. elegans. *PLoS One*, 7(6), e38649.

[188] Pachter, J. A., Yang, Y.-J., & Dill, K. A. (2024). Entropy, irreversibility and inference at the foundations of statistical physics. *Nat Rev Phys*, (pp. 1–12). Publisher: Nature Publishing Group.

[189] Palmer, S. E., Marre, O., Berry, M. J., & Bialek, W. (2015). Predictive information in a sensory population. *Proceedings of the National Academy of Sciences*, 112(22), 6908–6913.

[190] Park, S.-K., Shanbhag, S., Dubin, A., De Bruyne, M., Wang, Q., Yu, P., Shimoni, N., D'Mello, S., Carlson, J. R., Harris, G., et al. (2002). Inactivation of olfactory sensilla of a single morphological type differentially affects the response of drosophila to odors. *Journal of neurobiology*, 51(3), 248–260.

[191] Parrondo, J. M. R., Horowitz, J. M., & Sagawa, T. (2015). Thermodynamics of information. *Nat. Phys.*, 11, 131.

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

[193] Pillow, J. W., Shlens, J., Paninski, L., Sher, A., Litke, A. M., Chichilnisky, E., & Simoncelli, E. P. (2008). Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature*, 454(7207), 995–999.

[194] Pnevmatikakis, E. A. & Giovannucci, A. (2017). Normcorre: An online algorithm for piecewise rigid motion correction of calcium imaging data. *Journal of neuroscience methods*, 291, 83–94.

[195] Porter, J., Craven, B., Khan, R. M., Chang, S.-J., Kang, I., Judkewitz, B., Volpe, J., Settles, G., & Sobel, N. (2007). Mechanisms of scent-tracking in humans. *Nature neuroscience*, 10(1), 27–29.

[196] Pressé, S., Ghosh, K., Lee, J., & Dill, K. A. (2013). Principles of maximum entropy and maximum caliber in statistical physics. *Rev. Mod. Phys.*, 85(3), 1115–1141.

[197] Prinz, A. A., Bucher, D., & Marder, E. (2004). Similar network activity from disparate circuit parameters. *Nature neuroscience*, 7(12), 1345–1352.

[198] Pritz, C., Itskovits, E., Bokman, E., Ruach, R., Gritsenko, V., Nelken, T., Menasherof, M., Azulay, A., & Zaslaver, A. (2023). Principles for coding associative memories in a compact neural network. *Elife*, 12.

[199] Qian, H. & Reluga, T. C. (2005). Nonequilibrium thermodynamics and nonlinear kinetics in a cellular signaling switch. *Phys. Rev. Lett.*, 94(2), 028101.

[200] Radulovic, J., Kammermeier, J., & Spiess, J. (1998). Relationship between fos production and classical fear conditioning: effects of novelty, latent inhibition, and unconditioned stimulus preexposure. *Journal of Neuroscience*, 18(18), 7452–7461.

[201] Radvansky, B. A. & Dombeck, D. A. (2018). An olfactory virtual reality system for mice. *Nature communications*, 9(1), 839.

[202] Rajan, K., Harvey, C. D., & Tank, D. W. (2016). Recurrent network models of sequence generation and memory. *Neuron*, 90(1), 128–142.

[203] Randi, F. & Leifer, A. M. (2020). Nonequilibrium green's functions for functional connectivity in the brain.

[204] Rao, R. P. & Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature neuroscience*, 2(1), 79–87.

[205] Reddy, G., Desban, L., Tanaka, H., Roussel, J., Mirat, O., & Wyart, C. (2022). A lexical approach for identifying behavioural action sequences. *PLoS Comput. Biol.*, 18(1), e1009672.

[206] Riechers, P. M. & Crutchfield, J. P. (2016). Fluctuations when driving between nonequilibrium steady states.

[207] Rieke, F., Warland, D., Van Steveninck, R. d. R., & Bialek, W. (1999). *Spikes: exploring the neural code*.

[208] Rigolli, N., Reddy, G., Seminara, A., & Vergassola, M. (2022). Alternation emerges as a multi-modal strategy for turbulent odor navigation. *Elife*, 11.

[209] Risse, B., Thomas, S., Otto, N., Löpmeier, T., Valkov, D., Jiang, X., & Klämbt, C. (2013). Fim, a novel ftir-based imaging method for high throughput locomotion analysis. *PloS one*, 8(1), e53963.

[210] Roberts, W. M., Augustine, S. B., Lawton, K. J., Lindsay, T. H., Thiele, T. R., Izquierdo, E. J., Faumont, S., Lindsay, R. A., Britton, M. C., Pokala, N., Bargmann, C. I., & Lockery, S. R. (2016). A stochastic neuronal model predicts random search behaviors at multiple spatial scales in c. elegans. *eLife Sciences*, 5, e12572.

[211] Roemschied, F. A., Pacheco, D. A., Aragon, M. J., Ireland, E. C., Li, X., Thieringer, K., Pang, R., & Murthy, M. (2023). Flexible circuit mechanisms for context-dependent song sequencing. *Nature*, 622(7984), 794–801.

[212] Roman, A., Palanski, K., Nemenman, I., & Ryu, W. S. (2023). A dynamical model of c. elegans thermal preference reveals independent excitatory and inhibitory learning pathways. *Proc. Natl. Acad. Sci. U. S. A.*, 120(13), e2215191120.

[213] Roudi, Y., Aurell, E., & Hertz, J. A. (2009). Statistical physics of pairwise probability models. *Frontiers in computational neuroscience*, 3, 652.

[214] Roudi, Y. & Hertz, J. (2011). Mean field theory for nonequilibrium network reconstruction. *Phys. Rev. Lett.*, 106(4), 048702.

[215] Rutishauser, U., Mamelak, A. N., & Schuman, E. M. (2006). Single-trial learning of novel stimuli by individual neurons of the human hippocampus-amygdala complex. *Neuron*, 49(6), 805–813.

[216] Salazar, D. S. P. (2017). Nonequilibrium thermodynamics of restricted boltzmann machines. *Phys Rev E*, 96(2-1), 022131.

[217] Salvador, L. C., Bartumeus, F., Levin, S. A., & Ryu, W. S. (2014). Mechanistic analysis of the search behaviour of caenorhabditis elegans. *Journal of The Royal Society Interface*, 11(92), 20131092.

[218] Sanchez-Vives, M. V., Nowak, L. G., & McCormick, D. A. (2000). Membrane mechanisms underlying contrast adaptation in cat area 17in vivo. *Journal of Neuroscience*, 20(11), 4267–4285.

[219] Schmuker, M., Bahr, V., & Huerta, R. (2016). Exploiting plume structure to decode gas source distance using metal-oxide gas sensors. *Sensors and Actuators B: Chemical*, 235, 636–646.

[220] Schneidman, E., Berry, M. J., Segev, R., & Bialek, W. (2006). Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, 440(7087), 1007–1012.

[221] Scholz, M., Linder, A. N., Randi, F., Sharma, A. K., Yu, X., Shaevitz, J. W., & Leifer, A. (2018). Predicting natural behavior from whole-brain neural dynamics. *BioRxiv*, 445643.

[222] Schomaker, J. & Meeter, M. (2012). Novelty enhances visual perception. *PloS one*, 7(12), e50599.

[223] Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, 275(5306), 1593–1599.

[224] Seif, A., Hafezi, M., & Jarzynski, C. (2020). Machine learning the thermodynamic arrow of time. *Nat. Phys.*

[225] Shannon, C. E. (1948). A Mathematical Theory of Communication. *Bell Syst. Tech.*, 27(3), 379–423.

[226] Sharpee, T., Rust, N. C., & Bialek, W. (2004). Analyzing neural responses to natural signals: maximally informative dimensions. *Neural Comput.*, 16(2), 223–250.

[227] Shore, J. & Johnson, R. (1980). Axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy. *IEEE Trans. Inf. Theor.*, 26(1), 26–37.

[228] Skinner, D. J. & Dunkel, J. (2021). Estimating Entropy Production from Waiting Time Distributions. *Phys. Rev. Lett.*, 127(19), 198101. Publisher: American Physical Society.

[229] Smith, M. A.-Y., Honegger, K. S., Turner, G., & de Bivort, B. (2022). Idiosyncratic learning performance in flies. *Biol. Lett.*, 18(2), 20210424.

[230] Soh, Z., Sakamoto, K., Suzuki, M., Iino, Y., & Tsuji, T. (2018). A computational model of internal representations of chemical gradients in environments for chemotaxis of caenorhabditis elegans. *Scientific Reports*, 8(1), 17190.

[231] Solomon, S. S., Tang, H., Sussman, E., & Kohn, A. (2021). Limited evidence for sensory prediction error responses in visual cortex of macaques and humans. *Cerebral Cortex*, 31(6), 3136–3152.

[232] Srinivasan, M. V., Laughlin, S. B., & Dubs, A. (1982). Predictive coding: a fresh view of inhibition in the retina. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, 216(1205), 427–459.

[233] Stephens, G. J., Bueno de Mesquita, M., Ryu, W. S., & Bialek, W. (2011). Emergence of long timescales and stereotyped behaviors in caenorhabditis elegans. *Proc. Natl. Acad. Sci. U. S. A.*, 108(18), 7286–7289.

[234] 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.

[235] Stern, M., Sompolinsky, H., & Abbott, L. F. (2014). Dynamics of random neural networks with bistable units. *Physical Review E*, 90(6), 062710.

[236] Still, S., Sivak, D. A., Bell, A. J., & Crooks, G. E. (2012). The thermodynamics of prediction.

[237] Sussillo, D. (2009). Generating coherent patterns of activity from chaotic neural networks. *Neuron*, 63(4), 544–557.

[238] Tadres, D., Wong, P. H., To, T., Moehlis, J., & Louis, M. (2022). Depolarization block in olfactory sensory neurons expands the dimensionality of odor encoding. *Science Advances*, 8(50), eade7209.

[239] Tait, C., Mattise-Lorenzen, A., Lark, A., & Naug, D. (2019). Interindividual variation in learning ability in honeybees. *Behav. Processes*, 167, 103918.

[240] Takagi, K. (2019). Principles of mutual information maximization and energy minimization affect the activation patterns of large scale networks in the brain. *Front. Comput. Neurosci.*, 13, 86.

[241] Tanimoto, Y., Yamazoe-Umemoto, A., Fujita, K., Kawazoe, Y., Miyanishi, Y., Yamazaki, S. J., Fei, X., Busch, K. E., Gengyo-Ando, K., Nakai, J., Iino, Y., Iwasaki, Y., Hashimoto, K., & Kimura, K. D. (2017). Calcium dynamics regulating the timing of decision-making in c. elegans. *Elife*, 6.

[242] Tank, D. W. & Hopfield, J. J. (1987). Neural computation by concentrating information in time. *Proc. Natl. Acad. Sci. U. S. A.*, 84(7), 1896–1900.

[243] Tariq, M. F., Lewis, S. M., Lowell, A., Moore, S., Miles, J. T., Perkel, D. J., & Gire, D. H. (2021). Using head-mounted ethanol sensors to monitor olfactory information and determine behavioral changes associated with ethanol-plume contact during mouse odor-guided navigation. *Eneuro*, 8(1).

[244] Tastekin, I., Khandelwal, A., Tadres, D., Fessner, N. D., Truman, J. W., Zlatic, M., Cardona, A., & Louis, M. (2018). Sensorimotor pathway controlling stopping behavior during chemotaxis in the drosophila melanogaster larva. *Elife*, 7, e38740.

[245] Teşileanu, T., Golkar, S., Nasiri, S., Sengupta, A. M., & Chklovskii, D. B. (2022). Neural circuits for dynamics-based segmentation of time series. *Neural computation*, 34(4), 891–938.

[246] Thapliyal, S., Beets, I., & Glauser, D. A. (2023). Multisite regulation integrates multimodal context in sensory circuits to control persistent behavioral states in c. elegans. *Nat. Commun.*, 14(1), 3052.

[247] Torayama, I., Ishihara, T., & Katsura, I. (2007). Caenorhabditis elegans integrates the signals of butanone and food to enhance chemotaxis to butanone. *J. Neurosci.*, 27(4), 741–750.

[248] van Campenhout, J. & Cover, T. (1981). Maximum entropy and conditional probability. *IEEE Trans. Inform. Theory*, 27(4), 483–489.

[249] Varela, J. A., Sen, K., Gibson, J., Fost, J., Abbott, L., & Nelson, S. B. (1997). A quantitative description of short-term plasticity at excitatory synapses in layer 2/3 of rat primary visual cortex. *Journal of Neuroscience*, 17(20), 7926–7940.

[250] Varshney, L. R., Chen, B. L., Paniagua, E., Hall, D. H., & Chklovskii, D. B. (2011). Structural properties of the caenorhabditis elegans neuronal network. *PLoS computational biology*, 7(2), e1001066.

[251] Verano, K. V. B., Panizon, E., & Celani, A. (2023). Olfactory search with finite-state controllers. *Proceedings of the National Academy of Sciences*, 120(34), e2304230120.

[252] Vergassola, M., Villermaux, E., & Shraiman, B. I. (2007). 'Infotaxis' as a strategy for searching without gradients. *Nature*, 445(7126), 406–409. Number: 7126 Publisher: Nature Publishing Group.

[253] Vickers, N. J. & Baker, T. C. (1994). Reiterative responses to single strands of odor promote sustained upwind flight and odor source location by moths. *Proceedings of the National Academy of Sciences*, 91(13), 5756–5760.

[254] Vogels, T. P., Rajan, K., & Abbott, L. F. (2005). Neural network dynamics. *Annu. Rev. Neurosci.*

[255] Vyas, S., Golub, M. D., Sussillo, D., & Shenoy, K. V. (2020). Computation through neural population dynamics. *Annual review of neuroscience*, 43, 249–275.

[256] Wang, Q., Cavanagh, P., & Green, M. (1994). Familiarity and pop-out in visual search. *Perception & psychophysics*, 56, 495–500.

[257] Wang, Y., Zhang, X., Xin, Q., Hung, W., Florman, J., Huo, J., Xu, T., Xie, Y., Alkema, M. J., Zhen, M., & Wen, Q. (2020). Flexible motor sequence generation during stereotyped escape responses. *Elife*, 9.

[258] Weber, A. I., Krishnamurthy, K., & Fairhall, A. L. (2019). Coding principles in adaptation. *Annual review of vision science*, 5, 427–449.

[259] Weistuch, C., Mujica-Parodi, L. R., & Dill, K. (2021). The Refractory Period Matters: Unifying Mechanisms of Macroscopic Brain Waves. *Neural Computation*, 33(5), 1145–1163.

[260] White, J., Southgate, E., Thomson, J., & Brenner, S. (1986). The structure of the nervous system of the nematode caenorhabditis elegans. *Philosophical Transactions of the Royal Society of London Series B*, 314(1165), 1–340.

[261] Whitmire, C. J. & Stanley, G. B. (2016). Rapid sensory adaptation redux: a circuit perspective. *Neuron*, 92(2), 298–315.

[262] 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.

[263] Wolpert, D. H. (2020). Minimal entropy production rate of interacting systems. *New J. Phys.*, 22(11), 113013. Publisher: IOP Publishing.

[264] Worthy, S. E., Haynes, L., Chambers, M., Bethune, D., Kan, E., Chung, K., Ota, R., Taylor, C. J., & Glater, E. E. (2018). Identification of attractive odorants released by preferred bacterial food found in the natural habitats of c. elegans. *PloS one*, 13(7), e0201158.

[265] Yamada, M., Ohashi, H., Hosoda, K., Kurabayashi, D., & Shigaki, S. (2021). Multisensory-motor integration in olfactory navigation of silkmoth, bombyx mori, using virtual reality system. *Elife*, 10, e72001.

[266] Yamaguchi, S., Naoki, H., Ikeda, M., Tsukada, Y., Nakano, S., Mori, I., & Ishii, S. (2018). Identification of animal behavioral strategies by inverse reinforcement learning. *PLoS Comput. Biol.*, 14(5), e1006122.

[267] Yamazaki, S. J., Ohara, K., Ito, K., Kokubun, N., Kitanishi, T., Takaichi, D., Yamada, Y., Ikejiri, Y., Hiramatsu, F., Fujita, K., et al. (2019). Steftr: A hybrid versatile method for state estimation and feature extraction from the trajectory of animal behavior. *Frontiers in Neuroscience*, 13, 626.

[268] Yamazoe-Umemoto, A., Fujita, K., Iino, Y., Iwasaki, Y., & Kimura, K. D. (2015). Modulation of different behavioral components by neuropeptide and dopamine signalings in non-associative odor learning of caenorhabditis elegans. *Neurosci. Res.*, 99, 22–33.

[269] Yamazoe-Umemoto, A., Iwasaki, Y., & Kimura, K. D. (2018). Measuring spatiotemporal dynamics of odor gradient for small animals by gas chromatography. *Bio Protoc*, 8(7), e2797.

[270] Yan, H., Zhao, L., Hu, L., Wang, X., Wang, E., & Wang, J. (2013). Nonequilibrium landscape theory of neural networks. *Proc. Natl. Acad. Sci. U. S. A.*, 110(45), E4185–94.

[271] Yang, Y.-J. & Qian, H. (2023). Statistical Uncertainty Principle in Stochastic Dynamics. arXiv:2205.09321 [cond-mat].

[272] Yoshida, K., Hirotsu, T., Tagawa, T., Oda, S., Wakabayashi, T., Iino, Y., & Ishihara, T. (2012). Odour concentration-dependent olfactory preference change in c. elegans. *Nature communications*, 3(1), 739.

[273] Yu, X., Creamer, M. S., Randi, F., Sharma, A. K., Linderman, S. W., & Leifer, A. M. (2021). Fast deep neural correspondence for tracking and identifying neurons in c. elegans using semi-synthetic training. *Elife*, 10, e66410.

[274] Zhang, Y., Lu, H., & Bargmann, C. I. (2005). Pathogenic bacteria induce aversive olfactory learning in caenorhabditis elegans. *Nature*, 438(7065), 179–184.

[275] Zhong, W., Lu, Z., Schwab, D. J., & Murugan, A. (2018). Non-equilibrium statistical mechanics of continuous attractors.