

Rapid modulation of dynamics and computation in neural systems

Rich Pang

A dissertation

submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2019

Reading Committee:

Adrienne Fairhall, Chair

Eric Shea-Brown

Jeff Riffell

Program Authorized to Offer Degree:

Neuroscience

©Copyright 2019

Rich Pang

University of Washington

Abstract

Rich Pang

Chair of the Supervisory Committee:

Adrienne Fairhall

Department of Physiology and Biophysics

A central goal in theoretical neuroscience is to understand how neural systems perform computations over the continuum of timescales that underlie behavior. In particular, what are the algorithms and mechanisms enabling single-neuron membrane voltage fluctuations, which occur over milliseconds, to produce the dynamics and information processing in behavior that unfold over hours to years? Notably, while the core ionic processes of membrane voltage fluctuations have been largely elucidated and while extensive theories and evidence exist to explain how slow modulation of neural network structures might underlie learning, almost nothing is known about the liminal regime of seconds to minutes that bridges these two timescales. In the work that follows I address three questions in three different systems, each of which centers around neural computations occurring over the timescales of seconds to minutes. I first investigate the navigational decisions made by flying insects during odor tracking, where I show that fruit flies and mosquitoes exhibit a history dependence in their odor-triggered turning responses that is qualitatively similar to an information-maximizing tracking strategy, but not to others. Next, in collaboration with Ari Zolin, Raphael Cohn, and Vanessa Ruta, I analyze the dynamics of dopaminergic neuromodulation of a short-term memory circuit in the fruit fly mushroom body, where we suggest that the fly dopamine system encodes multiplexed representations of a wide diversity of sensory, motor, and valence signals, some of which predict behavior several seconds in the future. Third, I develop a spiking neural network model capable of storing and replaying sequential activity patterns using a heterosynaptic and fast-acting biological plasticity rule, and which reconstructs sequences through the existing recurrent network structure. Collectively, these results elucidate the computational capacities of three distinct systems and shed new light on short-term information processing in neural computations from three novel angles.

Finally, in collaboration with Sid Henriksen and Mark Wronkiewicz, I describe a simple network-growth model reproducing several statistical features of mouse brain network connectivity at the mesoscale; while this work does not explicitly address short-term computations, simplified statistical network models will be crucial to eventually understanding how such computations occur within large scale distributed brain networks.

To my family.

Table of contents

Acknowledgments	10
Introduction	12
i. Insect odor tracking	13
ii. Dopamine signals in the mushroom body	14
iii. Sequences in model neural networks	15
iv. Mesoscale brain connectivity	17
1. History dependence in insect flight decisions during odor tracking	23
1.1 Abstract	23
1.2 Summary	23
1.3 Introduction	24
1.4 Results	26
1.5 Discussion	37
1.5.1 Measuring and controlling turbulent signals	38
1.5.2 Informative features of plume encounters	39
1.5.3 Identifying evidence of theoretical search strategies in data	41
1.5.4 On the role of infotaxis in biological source localization	43
1.5.5 The role of vision in search and infotaxis	44
1.6 Materials and Methods	45
1.6.1 Dataset	45
1.6.2 Identifying plume crossings	46
1.6.3 Data exclusion criteria	47
1.6.4 Partial correlation analysis	47
1.6.5 Threshold and threshold-linear model fitting	48
1.6.6 History dependence analysis	48
1.6.7 Base model for surge-cast and centerline-inferring tracking algorithms	49
1.6.8 Surge-cast model	49
1.6.9 Centerline-inferring agent simulation	50
1.6.10 Infotaxis simulation	51
1.7 Supporting information	52

2. Multiplexed dopamine signals in the fruit fly mushroom body	62
2.1 Abstract	62
2.2 Introduction	63
2.2.1 The role of dopamine in computation and behavior	63
2.2.2 Dopaminergic control of mushroom body computations	64
2.3 Results	66
2.3.1 Absent external stimuli, DAN activity represents moment-to-moment behaviors	66
2.3.2 In sensory-motivated context DANs represent stimulus features	71
2.3.3 Interhemispheric DAN differences encode air stream angular movements	74
2.3.4 DANs recruit extended sensorimotor representations in the presence of odor	78
2.4 Discussion	84
2.4.1 Moment-to-moment sensorimotor representations in DAN inputs to the mushroom body	86
2.4.2 Multiplexed dopamine signals for memory and learning	88
2.4.3 On the relationship between spontaneous and evoked neural activity	91
2.5 Additional Materials and Methods	91
2.5.1 Data Collection	92
2.5.2 Data normalization	92
2.5.3 Filter fitting	92
2.5.4 Open loop sensory paradigm	92
2.5.5 Nested model analysis	93
2.6 Supplementary Figures	93
3. Fast and flexible sequence induction in spiking neural networks via rapid excitability changes	104
3.1 Abstract	104
3.2 Introduction	104
3.3 Results	107
3.3.1 LTP-IE increases membrane voltages and spike rates under random gating inputs	108
3.3.2 Spike sequences propagate along LTP-IE-defined paths through a network	110
3.3.3 Dependence of LTP-IE-based sequence propagation on network parameters	112
3.3.4 LTP-IE-based sequences can encode temporary stimulus-response mappings	115
3.4 Discussion	117

3.4.1 Comparison to existing sequence-generation models	118
3.4.2 Model limitations	120
3.4.3 Experimental predictions	121
3.4.4 Computational significance	122
3.5 Methods	123
3.5.1 Membrane and synaptic dynamics	123
3.5.2 Neuron tuning and network structure	124
3.5.3 Long-term potentiation of intrinsic excitability (LTP-IE) of gating inputs	125
3.5.4 Sequence replay	126
3.5.5 Simulations	126
3.5.5.1 Single-neuron LTP-IE simulation (Fig 1)	126
3.5.5.2 Replay in networks of neurons (Fig 2-3)	126
3.5.5.3 Replay dynamics parameter sweep (Fig 3)	126
3.5.5.4 LTP-IE-based stimulus-response mappings (Fig 4)	127
3.5.5.5 Other simulation parameters	128
3.6 Supplementary figures	128
 4. A simple generative model of the mouse mesoscale connectome	 139
4.1 Abstract	139
4.2 Introduction	139
4.3 Results	142
4.3.1 Comparison of the connectome to undirected graph models	142
4.3.2 The directed perspective and novel random graphs	144
4.3.3 Comparison of the connectome to the direct graph models	145
4.3.4 Higher-order connectivity statistics	149
4.4 Discussion	152
4.4.1 The importance of directed analysis	152
4.4.2 Indications of functional segregation and integration in the connectome	152
4.4.3 Organizational principles of the mouse connectome	153
4.4.4 A biophysical interpretation of source growth	154
4.4.5 A non-power law explanation of the connectome's undirected degree distribution	155
4.4.6 Limitations of the SGPA model	156

4.5 Conclusion	157
4.6 Materials and Methods	157
4.6.1 Empirical connectivity graph	158
4.6.2 Random graphs	158
4.6.3 Standard random graphs	158
4.6.4 Custom random graphs	159
4.6.5 Purely geometric models	159
4.6.6 Target-attraction-proximal-attachment (TAPA) random graph	159
4.6.7 Source-growth-proximal-attachment (SGPA) random graph	160
4.6.8 Total-degree SGPA random graphs	160
4.6.9 Node-by-node SGPA random graph	160
4.6.10 Undirected metrics and lesioning of directed graphs	161
4.7 Acknowledgments	161
4.8 Supplementary information	161
5. Conclusions and future directions	172
5.1 Insect odor tracking	172
5.2 Dopaminergic modulation of mushroom body computations	175
5.3 Rapid excitability-based sequence induction	177
5.4 Mesoscale connectivity in mammalian brain networks	179
5.5 Final remarks	182

Acknowledgments

My graduate studies at the University of Washington have proved one of the most exciting, challenging, and rewarding chapters of my life. None of this could have been undertaken alone, however, and it is overarchingly because of the friendship, support, inspiration, and wisdom—both scientific and personal—of the many people I’ve had the pleasure to interact with these last years that I’ve managed to arrive at where I am today.

I of course owe an immense debt of gratitude to my PhD advisor, Adrienne Fairhall, for showing me the world of computational neuroscience and helping me carve my own way through its many domains. Above all, she was willing to take a chance on my often ill-posed and circuitous meandering through our field, always helping me find value and meaning in my work along the way. Beyond fueling myriad fantastic scientific conversations, providing endless encouragement and support through the many trials of grad school and research, and connecting me with many wonderful opportunities and people both at UW and across the world, she has also been tremendously inspirational both in her ability to keep track of and direct our field as a whole and in her many efforts to make computational neuroscience more exciting and welcoming for all newcomers.

Getting to know and work with the many members of the Fairhall Lab, both past and present, has also been a true honor and delight. From those from the very beginning—Alison Duffy, Sharri Zamore, Heather Barnett, and Sangwook Lee—to those who passed through or joined the lab over the many years that followed—Yoni Browning, Ben Lansdell, Stefano Vrizzi, Chaoqing Wang, Richard Decal, Anatoly Buchin, Tom Richner, Kenneth Latimer, Guillaume Lajoie, Stefano Recanatesi, Hengji Wang, Argha Mondal—everyone in the lab has taught me an incredible amount, and they have been a crucial part of my looking forward to coming into work each day. I would also be remiss not to thank Eric Shea-Brown and his lab—in particular Ali Weber, Kam Decker Harris, Hannah Choi, and Matthew Farrell—who has been our sister lab since Eric and Adrienne came to UW. It was due largely to Eric’s lectures and journal club and the many happy hours with his lab members that I found my interest in biological neural networks, which has now become a key part of my research.

The majority of the work that follows would not exist without the energy and efforts of my collaborators. I’d particularly like to thank Mark Wronkiewicz and Sid Henriksen, with whom we managed to coerce a small summer school endeavor into an official research project; Floris van Breugel, Clement Vinauger, Jeff Riffell, and Michael Dickinson, for welcoming me into the world of insect neuroscience;

and Ari Zolin, Raffi Cohn, and Vanessa Ruta for showing me the intrigue and subtlety of the fruit fly dopamine system. I'd also like to thank my other committee members, who besides Adrienne, Eric, and Jeff, include Greg Horwitz, who gave me my first computational project when I rotated with him and with whom I've shared many scientific coffee hours, and Raj Rao, who very kindly allowed me to TA multiple classes he taught. I'd additionally like to thank the organizers of the two summer schools I was able to participate in—Christof Koch and Adrienne Fairhall of the Summer Workshop on the Dynamic Brain, and Joe Raimondo, Tim Vogels, Peter Latham, and Alex Antrobus, of the 2019 Computational Neuroscience Imbizo. These summer schools have created some of the most jam-packed and fondest memories of my PhD and have been highly influential in my deciding to continue on in the field.

Equally important to the growth occurring in the lab and over scientific luncheons is that which takes place over happy hours, around campfires, on the beach, and at dim sum tables with hot tea and dumplings. The friends I've shared these experiences and more with—with their good humor and laughter, unconditional encouragement and support, and assistance in keeping track of the bigger picture parts of life—have been absolutely paramount to my making it to the end of grad school. Besides everyone mentioned above I'd like to particularly thank Phil Mardoun, Matt Soleiman, Eric Thomas, Bethany Kondiles, Clare Gamlin, Abhishek De, Dean Pospisil, Pam Baker, Zach Small, Tanya Novotnak, and Peter Novotnak, for amplifying the good times and making the trying ones not seem so trying. Finally, I owe immense gratitude to my family, who despite not knowing exactly what I was doing in grad school have been unconditionally supportive as well. My mom has taught me the crucial skill of having faith in myself, my father the perspective of not doing things just because other people do them, and my sister Sarah the inspiring attitude of always seeking out the brighter things in life. Thank you for your encouragement and wisdom.

Introduction

A central question in neuroscience is how the millisecond timescales of single action potentials unfold into the lifelong timescales over which memories are consolidated and behaviors evolve and crystallize. On one end of the spectrum, extensive work has refined our understanding of how rapid membrane potential fluctuations shape computation at the level of single neurons. Since the first measurement of the action potential [1] through its mathematical modeling [2] through modern analysis of the stimuli that drive a neuron to spike [3, 4, 5, 6, 7, 8], we have developed a thorough understanding of how the electronics of lipid bilayers and ion channels give rise to specific information-processing algorithms operating over milliseconds to tens of milliseconds [9, 10]. On the other end of the spectrum, we have begun to reach a broad consensus about how memories are stored over the long term, namely via gradual modifications of synaptic connectivity patterns [11, 12, 13, 14, 15] and neuronal excitabilities [16, 17, 18], which take place over hours to days to years to reshape how our brains integrate information to drive behavior. One of the least understood but perhaps most interesting realms of temporal information processing, however, is that of timescales slower than the rapid electrical signaling of single neurons but faster than the long-term reconfiguration of brain connectivity. In this case, information must linger within a brain circuit far beyond the timescale at which membrane potentials decay, and must furthermore have some capacity to influence long-term network operations. Such rapid behavioral timescales of seconds to minutes underlie manifold processes, such as olfactory search, decision-making, navigation, conversation, and diverse motor sequences. More generally, neural computation at this timescale enables a temporary bridge or buffer for rapid electrical signals to influence long-term network reorganization. Elucidating the dynamics and function of neuronal processing over this timescale will thus be crucial to a fuller understanding of brain function and how nervous systems interact with the continuum of timescales that underlie the natural world.

What neural computations unfold over this liminal timescale of seconds to minutes to affect behavior? What neural and network algorithms underlie these computations and what purpose do they serve? How do such algorithms emerge from the collective interactions of cellular and synaptic mechanisms? In the work that follows I'll apply a variety of computational analyses and modeling techniques to quantitatively explore four different neural systems. I'll first analyze and model behavioral patterns in the odor-tracking flights of fruit flies to show that they integrate information over multiple seconds and use this to reshape odor responses. Next, I'll analyze dynamic recordings of dopaminergic inputs to the fruit fly mushroom body, an insect neuropil implicated in olfactory processing and memory.

Third, I'll model a generic network of spiking neurons to show how a novel, rapid empirical plasticity rule might subserve the stable encoding of flexible dynamic patterns that could be used for cognitive computations. Finally, I'll present a collaborative modeling project elucidating structural patterns in the mammalian brain network, the statistics of which will be necessary for future full-brain models of network dynamics and memory processing. While each of these projects is self-contained, the first three each shed new light on neural algorithms unfolding over seconds to minutes, and the fourth, while not explicitly addressing this timescale, provides useful groundwork for the future modeling of dynamics and computations supported by large-scale distributed brain networks.

I'll briefly outline the four specific prongs of my investigation and how they connect to our overarching question of intermediate neuronal processing timescales, followed by a thorough account of the motivation, methods, results, and conclusions for each project. Finally I'll bring together the key findings of all four arcs to elucidate potential principles linking them together and to suggest future research directions that would further refine our understanding of these systems.

i. Insect odor tracking

We begin by considering the problem of odor-tracking, a compelling natural decision-making behavior involving extended response sequences to a complex stimulus, and which pervades the animal kingdom [19]. Many animals follow odor plumes to locate food sources or mates, but natural plumes are broken up by turbulent airflow patterns into discrete odor puffs and filaments [20]. As a consequence, animals must use sparse, stochastic encounters with broken odor signals to localize the source of the plume, a highly nontrivial problem. Whereas in the well characterized *chemotaxis* behavior, organisms can simply orient their locomotion along chemical gradients to localize a source [21, 22], when odors are disturbed by turbulence the instantaneous gradients in the broken up signal tend to point toward the center of the puffs or filaments, rather than the source. What information about the plume source can an animal extract from instantaneous encounters with puffs and filaments using only biological sensors? Moreover, how should it integrate this information to drive its search trajectory for reliable and efficient source localization?

We address this problem by characterizing the flight patterns of fruit flies and mosquitoes as they interact with a wind-borne odor plume. By analyzing reconstructed trajectories of these insects flying freely in a wind tunnel containing a stationary appetitive odor plume, we ask which observables predict their flight responses to odor encounters, and whether the variability in these responses is consistent

with a battery of theoretical tracking algorithms. Since theory suggests that optimal tracking of natural plumes requires active integration of information about the source over a full multi-second flight trajectory [23], our core question is whether the recorded trajectories exhibit any evidence of such information integration, or whether purely “reactive”, or memoryless, responses to odor suffice for successful source localization.

To this end, we first deconstruct the substantial variability in odor-triggered flight responses across hundreds of plume crossings and correlate its key feature—the odor-triggered change in upwind orientation—both with features of the individual encounter and with features dependent on the prior history of the trajectory. Our results first reveal the concentration experienced during an encounter as a positive predictor of upwind turning, suggesting odor packet/filament concentration as a usable plume feature for source localization. Second, we find that the number of prior odor encounters in the trajectory weakens upwind turning, suggesting insects indeed integrate information and modulate their behavior over the multi-second timescales of a single trajectory.

Second, we compare the sign of our observed history dependence with that predicted by multiple normative tracking algorithms to understand how information accumulation is incorporated into the neural computations underlying effective tracking. We find that the observed history dependence is qualitatively consistent with that predicted by the theoretical tracking strategy “infotaxis” [23], but not with a simpler strategy asserting that insects should fly more upwind as they refine their estimate of the plume centerline. This suggests that odor-tracking insects may at least be partially updating an estimate over source location as they gather information about plume structure, and it sets the stage for unraveling the neural mechanisms for storing information over the seconds-to-minutes trajectory timescales and using it to shape behavior.

ii. Dopamine signals in the mushroom body

Where in the insect nervous system might recently accumulated information about plume structure be temporarily stored, and how might this actively reshape behavioral responses to odor stimuli in natural contexts? A compelling possibility only recently beginning to be elucidated is the dopaminergic innervation of the insect mushroom body [24], a third-order sensory neuropil noted for its involvement in associative learning [25] and sparse representation of odor identities [26]. Recent work, however, has revealed that the dopaminergic inputs to the mushroom body can rapidly, selectively, and reversibly gate information flow from the mushroom body’s input neurons (Kenyon cells) to the mushroom body

output neurons [?], suggesting the dopamine system’s potential as a substrate for the short-term information storage and signal-flow modulation needed for neural computations underlying effective odor tracking and other behaviors.

Through collaboration with Vanessa Ruta’s lab at Rockefeller University, and in particular with experimentalists Ari Zolin and Raphael Cohn, we have sought to disentangle what information is represented in the dopaminergic inputs to the mushroom body in order to gain insight into whether and how different signals ultimately reshape sensorimotor transformations in a behaviorally relevant manner. To this end, we used tethered fruit flies to analyze simultaneous recordings of dopaminergic activity in the mushroom body and walking behaviors in a simplified virtual reality environment. Our analysis reveals that beyond modulation by reward- and punishment-associated stimuli like sucrose or electric shock, respectively [?], dopamine signals are substantially modulated by moment-to-moment variations in behavior and external sensory stimulation. In particular, different dopamine cells innervating spatially localized compartments in the mushroom body gamma lobe [24] multiplex differential representations of (1) walking state, (2) walking speed, (3) turning velocity, (4) wind direction, and (5) odor onset. Moreover, our followup analyses reveal that sensorimotor representations in these dopamine signals may themselves be modulated by the presence of external odors. For instance, in the presence of odor, several dopaminergic signals appear to recruit a representation not only of the fly’s recent heading relative to upwind, but additionally the upwind turns it is likely to make several seconds in the future. These results shed new light on a specific neural substrate for short-term information processing in the insect olfactory system, in the context of the naturalistic behavior of odor tracking. Moreover, these findings set the stage for asking how extensively multiplexed signals representing sensory stimuli, motor actions, and states encoding intrinsic valence, may be integrated together to subserve a larger framework for information processing, such as reinforcement learning, and how they may suggest common functions of dopamine systems preserved across a diversity of species, from insects to mammals.

iii. Sequences in model neural networks

Having elucidated a novel feature of information-accumulation dynamics in a particular natural behavior and having shed light on the neural processes that might underlie the relevant timescales, we now turn to the question of more generic rapid modulation of information processing in neural networks. Specifically, while it is well known that the distributed architecture of neural networks affords

a rich flexibility of computations that can be learned through modifications of synaptic connections, it is far more poorly understood how more limited but still substantial flexibility over seconds to minutes might arise through purely biophysical mechanisms. For example, if instructed to raise our left hand in response to one word and our right hand in response to another, we can perform the task immediately and with little error. The plasticity mechanisms typically associated with reconfiguring network function, however, for instance spike-timing-dependent plasticity (STDP), generally act over slow timescales [27, 28] and so are unlikely to play a substantial role in rapid flexibility. What other cellular and biophysical mechanisms might enable such flexibility? Moreover, how do they depend on the existing network structures that have formed over the course of long-term learning, for example, through mechanisms like STDP?

While substantial work has shown how single or sets of stimuli can be stored in a network through realistic fast-acting biological mechanisms, such as reverberating activity or short-term synaptic potentiation, an open question is how such mechanisms might store relationships among stimuli, like sequences or associations, a core component of highly flexible representations. Motivated by the empirical phenomenon of hippocampal “replay” [29, 30, 31, 32, 33], we addressed this question by considering the storage and recall of novel sequences corresponding to trajectories through space. In particular, it has been repeatedly observed in rodent hippocampal regions CA3 and CA1 that neuronal spike sequences initially activated by and encoding a recent locomotor trajectory through the environment, subsequently replay when the animal pauses to rest. These sequences replay in either the forward or reverse order in which they were originally evoked, and at a highly compressed timescale relative to the behavioral sequence.

Using these observations as an empirical entry point into rapid cognitive flexibility, we modeled networks of position-tuned spiking neurons to ask how a specific recently discovered fast-acting cellular mechanism might give rise to a simplified version of replay, yet which nonetheless encodes arbitrary trajectories through an environment. Of note, biophysical mechanisms associated with sequences are typically assumed to be homosynaptic, e.g. in STDP a synapse from cell A to cell B facilitates increases only if A and B fire in rapid succession, thus reinforcing causal signal propagation at the synapse [12, 27]. Here, however, we consider an alternative, heterosynaptic mechanism, long-term plasticity of intrinsic excitability (LTP-IE) [34, 35], which effectively serves to increase the excitability of a recently active cell, regardless of which input caused the cell to spike. Our modeling work reveals that given a simple recurrent network architecture, modulation of excitability alone, which might occur much more

rapidly than recurrent synaptic modulation, is indeed sufficient to store and replay sequences corresponding to recent trajectories through space, and that LTP-IE might further be used to rapidly and flexibly store novel stimulus-response mappings in an otherwise untrained recurrent network. Besides suggesting a specific mechanism underlying replay, this sheds novel light on the “neurocentric” theory of learning [18], in which memories and computations are encoded predominantly as sets of neurons, rather than the connections between them. This helps frame the more overarching theoretical question of which recurrent network architectures best support the induction of flexible computations through excitability changes alone. Untangling this idea within the neurocentric framework of computation may reveal many novel insights into cognitive flexibility, both at rapid and extended timescales.

iv. Mesoscale brain connectivity

One of the core features of neural computation distinguishing it from conventional computing paradigms is its highly distributed and parallelized nature. Indeed, multi-region and whole-brain recording studies suggest that neural representations of information, both perceptual and over the short-term memory timescales we have discussed so far, are highly distributed, recruiting diverse brain areas to encode stimuli [36, 37]. To understand how real network architectures in the brain support flexible dynamics and computations, it is therefore necessary to develop a sufficient understanding of the core connectivity patterns underlying empirically observed brain networks, such that one can disentangle the computations they allow within parsimonious models. Accordingly, recent high-throughput technologies and industrial scientific endeavors have allowed for the rapid collection of comprehensive connectivity datasets, enabling a thorough, data-driven analysis of the potential patterns that underlie them. One prominent example of such a dataset is the Allen Institute for Brain Science’s (AIBS) mesoscale mouse connectome, a complete map of all directed interconnections among several hundred identified brain areas, reconstructed through high-throughput tracing and imaging [38], providing a holistic entry point to understanding connectivity patterns that wire together the entire brain.

To this end, in collaboration with theoreticians Mark Wronkiewicz and Sid Henriksen, we characterized several core statistical features of AIBS’ mesoscale connectome and developed a parsimonious network-growth model capable of reproducing these features *in silico* [39], in order to understand whether mesoscale connectivity could be explained by simple underlying principles. In particular, whereas most prior analyses of brain connectivity have relied on “undirected” models, in which connection directionality is ignored, the tracing procedure used in the AIBS connectome allowed the

experimenters to retain directional information such that the resulting connectome could be represented as a directed graph among 413 regions in the mouse brain. Using this information, we found that when incoming vs outgoing connectivity patterns were separated prior to analyses, these two types of connections exhibited distinct degree distributions, with in-degree following an approximately Gaussian distribution and out-degree following an approximately exponential one. Crucially, we found nearly the exact same distributions could emerge from a network-growth model based on two simple rules. First, outgoing connections should obey a rich-get-richer rule, such that brain regions are more likely to generate more outgoing connections in proportion to how many outgoing connections they already have. Second, connections should preferentially target nearby brain regions but should otherwise be random. Beyond yielding the empirically observed in- and out-degree distributions, these rules also recapitulated the inverse relationship between clustering coefficient and total degree that we observed in the connectome.

These results suggest that rather than requiring complicated preprogrammed instructions, large-scale statistics of mesoscale connectivity might arise from parsimonious principles. This may simplify future modeling studies in which microscale connectivity structures are embedded in larger, mesoscale connection structures, in order to disentangle multiscale effects of network structure on cognitive flexibility and the dynamic modulation of information flow and computation in biologically inspired networks more generally.

References

- [1] Stephen M Schuetze. The discovery of the action potential. *Trends in Neurosciences*, 6:164–168, 1983.
- [2] Alan L Hodgkin and Andrew F Huxley. A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of physiology*, 117(4):500–544, 1952.
- [3] Hugh L Bryant and José P Segundo. Spike initiation by transmembrane current: a white-noise analysis. *The Journal of physiology*, 260(2):279–314, 1976.
- [4] Zachary F Mainen and Terrence J Sejnowski. Reliability of spike timing in neocortical neurons. *Science*, 268(5216):1503–1506, 1995.
- [5] Prakash Kara, Pamela Reinagel, and R Clay Reid. Low response variability in simultaneously recorded retinal, thalamic, and cortical neurons. *Neuron*, 27(3):635–646, 2000.
- [6] Blaise Agüera y Arcas and Adrienne L Fairhall. What causes a neuron to spike? *Neural Computation*, 15(8):1789–1807, 2003.
- [7] Blaise Agüera y Arcas, Adrienne L Fairhall, and William Bialek. Computation in a single neuron: Hodgkin and huxley revisited. *Neural Computation*, 15(8):1715–1749, 2003.
- [8] Adrienne L Fairhall, C Andrew Burlingame, Ramesh Narasimhan, Robert A Harris, Jason L Puchalla, Michael J Berry, et al. Selectivity for multiple stimulus features in retinal ganglion cells. *Journal of neurophysiology*, 96(5):2724–2738, 2006.
- [9] Peter Dayan and Laurence F Abbott. Theoretical neuroscience: computational and mathematical modeling of neural systems. 2001.
- [10] Wulfram Gerstner, Werner M Kistler, Richard Naud, and Liam Paninski. *Neuronal dynamics: From single neurons to networks and models of cognition*. Cambridge University Press, 2014.
- [11] Tim VP Bliss and Graham L Collingridge. A synaptic model of memory: long-term potentiation in the hippocampus. *Nature*, 361(6407):31, 1993.
- [12] Henry Markram, Joachim Lübke, Michael Frotscher, and Bert Sakmann. Regulation of synaptic efficacy by coincidence of postsynaptic aps and epsps. *Science*, 275(5297):213–215, 1997.

- [13] Stephen J Martin, Paul D Greenwood, and Richard GM Morris. Synaptic plasticity and memory: an evaluation of the hypothesis. *Annual review of neuroscience*, 23(1):649–711, 2000.
- [14] Guo-qiang Bi and Mu-ming Poo. Synaptic modification by correlated activity: Hebb’s postulate revisited. *Annual review of neuroscience*, 24(1):139–166, 2001.
- [15] SJ Martin and RGM Morris. New life in an old idea: the synaptic plasticity and memory hypothesis revisited. *Hippocampus*, 12(5):609–636, 2002.
- [16] Wei Zhang and David J Linden. The other side of the engram: experience-driven changes in neuronal intrinsic excitability. *Nature Reviews Neuroscience*, 4(11):885, 2003.
- [17] Riccardo Mozzachiodi, Fred D Lorenzetti, Douglas A Baxter, and John H Byrne. Changes in neuronal excitability serve as a mechanism of long-term memory for operant conditioning. *Nature neuroscience*, 11(10):1146, 2008.
- [18] Heather K Titley, Nicolas Brunel, and Christian Hansel. Toward a neurocentric view of learning. *Neuron*, 95(1):19–32, 2017.
- [19] Neil J Vickers. Mechanisms of animal navigation in odor plumes. *The Biological Bulletin*, 198(2):203–212, 2000.
- [20] Boris I Shraiman and Eric D Siggia. Scalar turbulence. *Nature*, 405(6787):639, 2000.
- [21] Julius Adler. Chemotaxis in bacteria. *Science*, 153(3737):708–716, 1966.
- [22] Evelyn F Keller and Lee A Segel. Model for chemotaxis. *Journal of theoretical biology*, 30(2):225–234, 1971.
- [23] Massimo Vergassola, Emmanuel Villermaux, and Boris I Shraiman. ‘infotaxis’ as a strategy for searching without gradients. *Nature*, 445(7126):406, 2007.
- [24] Yoshinori Aso, Daisuke Hattori, Yang Yu, Rebecca M Johnston, Nirmala A Iyer, Teri-TB Ngo, Heather Dionne, LF Abbott, Richard Axel, Hiromu Tanimoto, et al. The neuronal architecture of the mushroom body provides a logic for associative learning. *Elife*, 3:e04577, 2014.
- [25] J Steven de Belle and Martin Heisenberg. Associative odor learning in drosophila abolished by chemical ablation of mushroom bodies. *Science*, 263(5147):692–695, 1994.

- [26] Andrew C Lin, Alexei M Bygrave, Alix De Calignon, Tzumin Lee, and Gero Miesenböck. Sparse, decorrelated odor coding in the mushroom body enhances learned odor discrimination. *Nature neuroscience*, 17(4):559, 2014.
- [27] Guo-qiang Bi and Mu-ming Poo. Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *Journal of neuroscience*, 18(24):10464–10472, 1998.
- [28] Sen Song and Larry F Abbott. Cortical development and remapping through spike timing-dependent plasticity. *Neuron*, 32(2):339–350, 2001.
- [29] David J Foster and Matthew A Wilson. Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature*, 440(7084):680, 2006.
- [30] Thomas J Davidson, Fabian Kloosterman, and Matthew A Wilson. Hippocampal replay of extended experience. *Neuron*, 63(4):497–507, 2009.
- [31] Anoopum S Gupta, Matthijs AA van der Meer, David S Touretzky, and A David Redish. Hippocampal replay is not a simple function of experience. *Neuron*, 65(5):695–705, 2010.
- [32] Margaret F Carr, Shantanu P Jadhav, and Loren M Frank. Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval. *Nature neuroscience*, 14(2):147, 2011.
- [33] Shantanu P Jadhav, Caleb Kemere, P Walter German, and Loren M Frank. Awake hippocampal sharp-wave ripples support spatial memory. *Science*, 336(6087):1454–1458, 2012.
- [34] Jung Ho Hyun, Kisang Eom, Kyu-Hee Lee, Won-Kyung Ho, and Suk-Ho Lee. Activity-dependent downregulation of d-type k⁺ channel subunit kv1.2 in rat hippocampal ca3 pyramidal neurons. *The Journal of physiology*, 591(22):5525–5540, 2013.
- [35] Jung Ho Hyun, Kisang Eom, Kyu-Hee Lee, Jin Young Bae, Yong Chul Bae, Myoung-Hwan Kim, Sooyun Kim, Won-Kyung Ho, and Suk-Ho Lee. Kv1.2 mediates heterosynaptic modulation of direct cortical synaptic inputs in ca3 pyramidal cells. *The Journal of physiology*, 593(16):3617–3643, 2015.
- [36] Thomas Naselaris, Kendrick N Kay, Shinji Nishimoto, and Jack L Gallant. Encoding and decoding in fmri. *Neuroimage*, 56(2):400–410, 2011.

- [37] Kartik K Sreenivasan, Clayton E Curtis, and Mark D'Esposito. Revisiting the role of persistent neural activity during working memory. *Trends in cognitive sciences*, 18(2):82–89, 2014.
- [38] Seung Wook Oh, Julie A Harris, Lydia Ng, Brent Winslow, Nicholas Cain, Stefan Mihalas, Quanxin Wang, Chris Lau, Leonard Kuan, Alex M Henry, et al. A mesoscale connectome of the mouse brain. *Nature*, 508(7495):207, 2014.
- [39] Sid Henriksen, Rich Pang, and Mark Wronkiewicz. A simple generative model of the mouse mesoscale connectome. *Elife*, 5:e12366, 2016.

1. History dependence in insect flight decisions during odor tracking

This chapter has been previously published under the Creative Commons Attribution (CC BY) license and is available via the following citation:

Pang, Rich, et al. “History dependence in insect flight decisions during odor tracking.” *PLoS Computational Biology* 14.2 (2018): e1005969. (DOI: 10.1371/journal.pcbi.1005969).

1.1 Abstract

Natural decision-making often involves extended decision sequences in response to variable stimuli with complex structure. As an example, many animals follow odor plumes to locate food sources or mates, but turbulence breaks up the advected odor signal into intermittent filaments and puffs. This scenario provides an opportunity to ask how animals use sparse, instantaneous, and stochastic signal encounters to generate goal-oriented behavioral sequences. Here we examined the trajectories of flying fruit flies (*Drosophila melanogaster*) and mosquitoes (*Aedes aegypti*) navigating in controlled plumes of attractive odorants. While it is known that mean odor-triggered flight responses are dominated by upwind turns, individual responses are highly variable. We asked whether deviations from mean responses depended on specific features of odor encounters, and found that odor-triggered turns were slightly but significantly modulated by two features of odor encounters. First, encounters with higher concentrations triggered stronger upwind turns. Second, encounters occurring later in a sequence triggered weaker upwind turns. To contextualize the latter history dependence theoretically, we examined trajectories simulated from three normative tracking strategies. We found that neither a purely reactive strategy nor a strategy in which the tracker learned the plume centerline over time captured the observed history dependence. In contrast, “infotaxis”, in which flight decisions maximized expected information gain about source location, exhibited a history dependence aligned in sign with the data, though much larger in magnitude. These findings suggest that while true plume tracking is dominated by a reactive odor response it might also involve a history-dependent modulation of responses consistent with the accumulation of information about a source over multi-encounter timescales. This suggests that short-term memory processes modulating decision sequences may play a role in natural plume tracking.

1.2 Summary

Many important behaviors require animals to make extended decision sequences in response to complex stimuli. In this study we investigated the sequences of navigational decisions made by fruit flies and mosquitoes while tracking odor plumes, which are generally subject to turbulence. By examining video-taped 3D trajectories of these insects flying in a wind tunnel containing an attractive odor plume, we asked what features of encounters with the plume influenced odor-triggered flight decisions. Most notably, we found that although the average response was dominated by a reflexive upwind turn, its strength was modulated by the history of prior plume encounters throughout the trajectory. While no theoretical strategy we simulated captured all aspects of the data, the algebraic sign of the history dependence was only recapitulated in a model where a simulated tracking agent maximizes information about the position of the plume source. This suggests that real odor tracking may involve short-term memory processes over multi-encounter timescales that are consistent with the accumulation of information about source location.

1.3 Introduction

In contrast to two-alternative laboratory decision-making paradigms, an animal making decisions in the natural world typically executes extended sequences of choices in response to stimuli with complex spatiotemporal statistics. One such natural decision process is the tracking of odor plumes to locate food sources and mates [1]. Natural odor plumes have complex distributions due to turbulence, which breaks plumes into intermittent filaments of odor interrupted by large areas of clean air [2, 3]. Thus, an insect’s navigational decisions during plume tracking must be made in response to sparsely distributed odor encounters that provide incomplete information about the source location [4]. At present, it is not well understood which features of in-flight encounters with a turbulent plume contain the most useful information about source position, nor how insects integrate these features for successful localization. Investigating this problem in flying insects will help illuminate this decision process and identify biologically relevant statistics of natural windborne signals. Ultimately, this may help in the development of strategies to combat insect-borne illness as well as inspire the design of robots used to track signals advected by fluids [5]. More generally, we hope our work will inform theories of the neuroscience underlying the accumulation of sensory evidence and how this evidence is used in active sensing and goal-directed decision making.

Because turbulent stimuli are difficult to both measure and control, experiments investigating search strategies during flight can only approximate naturalistic environments. This is frequently done

using small plume-containing wind tunnels, in which the trajectories of flying insects are recorded with video cameras. Such studies showed, for example, that plume-tracking behaviors are highly influenced by the large-scale statistics of plume structure [1]. When tracking CO₂, for instance, mosquitoes exhibited increased upwind flight in a fluctuating, as compared to a homogeneous, plume [6]; further, continuous ribbon plumes elicited less upwind flight in moths than rapidly fluctuating plumes [7]. These both suggest that at least the coarse spatiotemporal structure of a plume affects tracking flight, although the mechanisms underlying this behavior remain unknown. It has also been suggested that the duration, concentration, and concentration gradient of odor packets may be informative plume features [8, 9, 10, 11, 12, 13], but it is not known how these properties influence fine-scale navigational decisions or how these fine-scale decisions combine to form a global search strategy.

Fine-scale aspects of plume-tracking strategies have been primarily investigated in laminar flow (i.e., non-turbulent) wind tunnels with stationary plumes. Here, knowledge of the plume’s stationary profile allows one to infer the approximate plume concentration experienced by an insect based on its position at every moment in its flight trajectory. Subsequently, one can ask what behavioral responses are triggered by encounters with the plume. The current model for how flying insects track a source is that they follow a sequence of reflexes, in which they “surge” upwind upon entering the plume and “cast” across the wind upon leaving it. Some evidence of the surge-cast reflex strategy has been found in moths [14], in fruit flies tracking an ethanol plume [15], and in mosquitoes tracking CO₂ [16]. A related but distinct “counterturning” (zigzagging) behavior, triggered by plume-crossing, has been observed in moths tracking pheromone plumes [7, 14, 17, 18].

While the surge-cast strategy is consistent with some aspects of plume tracking and provides an intuitive mechanism that can lead a simulated tracking agent to the source of a laminar plume [15], its most important shortcoming is that while it explains average behavioral responses to plume encounters, it does not explain the diversity of individual navigational decisions. Indeed, though the mean behavioral response of a fruit fly crossing an ethanol plume is to turn briefly upwind, individual responses vary over a range of nearly 100 degrees of turning angle (Fig 1D; the corresponding odor concentration timecourses for example crossings are shown in Fig 1E). In a similar analysis of moth flight trajectories, for example, investigating variability around the mean led to novel observations about moment-to-moment behaviors, such as ground speed variations in turning vs straight flight and decreases in ground speed upon approach to the pheromone source, which were not apparent from the mean responses alone [19]. What is not known is whether the diversity in fruit fly odor-

triggered responses arises purely from unobservable variations in small air currents or the insect’s internal state, or whether some of it can be predicted from measurable stimulus features. In the context of courtship, for instance, an analysis of behavioral diversity indeed showed that much variability in fruit fly singing previously assumed to be random could in fact be explained by variations in the fly’s recent sensory experience [20], which allowed the authors to better hypothesize about the underlying neural mechanisms.

Here, we explore the sources of diversity in plume-tracking-related flight maneuvers by analyzing a dataset of fruit fly (*Drosophila melanogaster*) and mosquito (*Aedes aegypti*) flight trajectories through a wind tunnel containing a stationary odor plume, and by asking which detectable features of odor encounters systematically modulate flight maneuvers. We find that fruit flies, and possibly mosquitoes as well, turn more strongly upwind in response to higher experienced concentrations, and that later encounters in a sequence of encounters along a trajectory trigger weaker upwind turns. To our knowledge, the latter result provides the first evidence for information accumulation over multiple plume encounters in in-flight plume-tracking in fruit flies and mosquitoes. This suggests that purely reactive models of insect plume-tracking might be significantly augmented by a history dependence term in which upwind flight slows as more encounters occur.

1.4 Results

We analyzed datasets of fruit fly and mosquito flight trajectories recorded with a camera-based tracking system as they interacted with either an ethanol (fruit flies) or carbon dioxide (mosquitoes) plume. The data were originally described in van Breugel and Dickinson [15] (fruit flies) and van Breugel et al. [21] (mosquitoes). Briefly, a low velocity non-turbulent wind flowed along the long axis of a wind tunnel in the center of which was a stationary laminar odor plume (Fig 1A and 1B). The plume had an approximately Gaussian cross-sectional profile, with maximum concentration along the centerline, and the floor of the wind tunnel contained a checkerboard pattern to provide the insects with a visual cue by which to orient. Insects were allowed to explore the wind tunnel for several hours, and whenever they took off, their trajectories were captured by the tracking system until they landed. The 3D position sequence of the trajectory was reconstructed in real time from the videos. Because the plume was stationary with a previously measured concentration profile, the dynamically varying odor concentration experienced by the insect at each time point could be estimated from its position. To understand how specific plume crossing features influenced decisions, we investigated the time-varying




Fig 1: Example flight patterns and illustration of variability. *A. Example fruit fly trajectory in a wind tunnel containing a stationary ethanol plume. The plume outline is shown in red, with the shell indicating 2 standard deviations from the plume's centerline. The trajectory color shows the instantaneous odor concentration experienced by the fly, and the green star marks where the fly took off. B. Example plume-crossings from several trajectories. Blue arrows on a subset of crossings indicate the direction of flight. C. Heading is defined as the angle between the upwind direction and the fly's*

Continued from Fig 1: *velocity vector. D. Heading time courses surrounding example crossings (gray), as well as mean (black solid) and standard deviation (black dashed), relative to the crossing time. E. Odor concentration time courses surrounding example crossings.*

heading (where heading is defined as the angle between the fly’s velocity vector and upwind [Fig 1C]), time-locked to the time of the peak concentration experienced during the plume crossing (Fig 1B, 1D and 1E). We defined a plume crossing as the portion of a trajectory when the odor experienced by the fly exceeded a minimum concentration threshold. Our criteria for choosing this threshold are discussed in Methods, but the conclusions that follow are not strongly affected by its precise value.

Because natural flight trajectories typically cross the odor plume repeatedly, each time yielding a potentially unique sensory experience, we first asked whether the crossing-triggered flight maneuvers depended on the peak concentration experienced during the crossing. We need, however, to control for effects of the animal’s initial heading and location in the tunnel. To do this, we calculated the partial correlation between peak experienced odor concentration and the insect’s heading at different points in the near future, conditional on x_0 , the insect’s position along the long axis of the wind tunnel, and h_0 , its initial heading at the time of the crossing. This conditioning allowed us to subtract out any linear dependence on these purely geometric factors. To further isolate the concentration-dependent effects from those deriving from initial location or heading angle, in all subsequent analysis we include only crossings occurring in the middle portion of the wind tunnel (i.e., not in the upwind- or downwind-most 30 cm) and only crossings for which the insect crossed the plume such that h_0 was between 60 and 120. This selection of the crosswind heading range was motivated by the canonical view that plume crossings on average trigger upwind turns [15, 16], which would be more apparent if the insect were not headed upwind already. As shown in Fig 2A and 2B for fruit flies we found a significant negative partial correlation between the peak concentration experienced in the crossing and their flight heading 300 ms later, suggesting that higher concentrations elicit stronger upwind turns.

A correlation between peak concentration and subsequent heading could also arise, however, if the response were purely binary but noisy—as has been implicitly assumed in previous odor-tracking investigations [15] as well as in theoretical models for tracking turbulent signals [22]. To test whether, in contrast, heading may depend continuously on odor-encounter concentration, we compared two simple models predicting h_{300} , (the heading 300 ms post-crossing, or 500 ms [h500] for mosquitoes), from c_{peak} , the peak concentration experienced in the crossing preceding the heading measurement (Fig 2C). In the binary model, h_{300} was assumed to take one mean value when c_{peak} was below




Fig 2: Dependence of heading on peak concentration. *A. Partial correlation of peak odor concentration experienced by the animal and subsequent heading at various times past the odor peak for fruit flies and mosquitoes. Shading indicates 2.5–97.5% confidence interval. B. Heading at 300 ms since odor concentration peak as a function of peak odor concentration during crossing. Each point represents a single plume-crossing. Though a significant correlation exists, the joint distribution is dominated by variability. C. Threshold and threshold-linear models for identifying non-binary dependence of h_{300} (the heading 300 ms post-crossing) on concentration. For mosquitoes, h_{500} (the heading 500 ms post-crossing) was used. P-values indicating the probability that the threshold-linear model would have fit better by chance (F-test) are shown in the gray box for four different experiments (fruit flies following an ethanol plume in 0.3, 0.4, or 0.6 m/s winds, and mosquitoes following a CO_2 plume in 0.4 m/s wind).*

a threshold c_{th} and another when c_{peak} was above c_{th} , with the linear effects of h_0 and x_0 again subtracted out to minimize geometric confounds. We compared this with a threshold-linear model, in which h_{300} was assumed to be a linear function of c_{peak} when c_{peak} was greater than c_{th} . The threshold-linear model allowed us to retain the possibility that sufficiently weak concentrations are not detectable, without having to assume that this would be accounted for by a linear prediction. We fit both models by minimizing the squared prediction error of h_{300} . We found that for fruit flies the threshold-linear model fit significantly better than the binary model (F-test, with p-values shown in Fig 2C), corroborating our hypothesis that behaviorally relevant concentrations are not processed in a simple binary way during in-flight plume tracking. For mosquitoes, there was not enough data to see a significant difference between the fits ($p = 0.175$) (Fig 2C).

Beyond purely reflexive flight maneuvers, recent theoretical results suggest that more efficient source localization can result from integrating information from multiple plume encounters over an extended period of time [22, 23]. Such information integration should lead to a dependence of crossing-triggered flight maneuvers on trajectory history. To explore this, we asked whether it is possible to distinguish flight maneuvers that are triggered by later compared with earlier plume crossings in an extended sequence of crossings.

When we separated plume crossings into early (the first or second in a trajectory) vs. late (the third or later in a trajectory) crossings, we found that while in both cases, the average flight is strongly dominated by a brief upwind turn, late crossings tended to trigger weaker upwind turns than early crossings (Fig 3A–3D). Fig 3E–3J show the equivalent crossing distributions for three different models, one reactive (surge-cast [Fig 3E]) and two whose behavior depends on information about the plume accumulated over time (centerline-inferring [Fig 3F] and infotaxis [Fig 3G–3J]), which are discussed in more detail later in the text. There may be geometric confounds to identifying history dependence in the data: later crossings may generally occur in the more upwind section of the wind tunnel since the insect has made more upwind progress, which might result in more crosswind turns simply because the insect has a smaller area to turn into. However, we found that the position along the long axis of the wind tunnel at the time of crossing, x_0 , differed by less than 3 centimeters between early and late crossings (about 2% of the total wind tunnel length) (S2 Fig).

To treat potential confounds, not only geometric but also those possibly due to fatigue from extended flying (e.g. flies that have been flying longer may choose more crosswind paths to save energy), we defined a new heading change variable $h^*(t)$ at each time point post-crossing, computed as the




Fig 3: History dependence of crossing-triggered turns in data and models. *A-D.* Crossing-triggered heading time courses for early (blue) or late (green) crossings for fruit flies tracking an ethanol plume in three different wind speeds, or mosquitoes flying in a wind tunnel with a 0.4 m/s wind. Thick lines indicate means, with shading indicating standard error of the mean. *E.* Same as *A* but for trajectories generated using the surge-cast model. *F.* Same as *A* but for trajectories generated using the centerline-inferring model. *G-J.* Same as *A-D* but for trajectories generated using the infotaxis algorithm, with wind speeds and plume profiles matched to each of the four experiments.

original $\Delta h(t)$ with the best-fit linear prediction from x_0 and T , the total flight time since take-off, subtracted out, thus removing any linear dependence on these variables of whether the crossing was early vs. late. As shown in Fig 4, the early vs. late time courses of $h^*(t)$ were significantly different from about 300 ms post-crossing onwards for all of the fruit fly experiments, with late heading changes being more crosswind than early heading changes. We also calculated the partial correlation between the heading change, time-averaged from 350 ms to 450 ms post-crossing, and the crossing number, conditioned on x_0 and T . Note: these plots appear symmetrical about $h^*(t) = 0$ because the mean $\Delta h(t)$ across all crossings is recomputed and subtracted out during the linear fitting at each timepoint. Thus, $h^*(t)$ averages to 0 by construction. Since there are similar numbers of crossings in each group, and since the early crossings have stronger upwind components ($h^*(t) \downarrow 0$), whereas the late crossings have stronger crosswind components ($h^*(t) \downarrow 0$), the average within-group $h^*(t)$'s end up about equidistant from 0. The heading change and crossing number were significantly correlated even after removing all correlations through x_0 and T (S4 Fig). These two analyses provide further evidence that the small history dependence in the fly's trajectories was due neither to the geometry of the wind tunnel nor to fatigue of the insects, but was instead related to their extended interactions with the plume. Additionally, while increased wind speed across the fruit fly experiment tended to yield stronger upwind turns overall, as shown in Fig 3A–3C and Fig 4A–4C it did not appear to have any consistent effect on the history dependence.

How does the history dependence in plume-crossing-triggered flight responses compare with that predicted by theoretical algorithms? To address this we considered three plume-tracking strategies capable of generating complete flight trajectories, in a simulated wind tunnel containing a plume with the same geometry as that in the experiments. We first considered a purely reflexive “surge-cast” algorithm, based on the model proposed by van Breugel and Dickinson [15]. Here, however, we initially modeled the insect to have baseline flight dynamics given by a correlated random 3D walk with a moderate bias towards casting (crosswind flight), with parameters fit to match the insects' empirical flight speed, angular velocity, and crosswind position distribution as best as possible (S6 Fig). In addition to the random walk dynamics, whenever the insect crossed the plume, it was subjected to a brief upwind “surge” force, modeled as an alpha function with a time constant of 0.07 s. An example empirically observed trajectory is shown in Fig 5A, with a trajectory generated by the surge-cast model shown in Fig 5B. Fig 3E shows that the surge-cast model recapitulates the dominant upwind component of crossing-triggered turns but does not yield an observable distinction between early and




Fig 4: History dependence of post-crossing turns after accounting for position and flight time. Each panel shows the mean and standard error of the mean (shading) of $h^*(t)$ across different time points post-crossing for early (blue) and late (green) crossings. The dashed line indicates the probability that such a difference in means would have arisen from chance, with the gray line marking 0.05. A. Fruit fly experiment in 0.3 m/s wind (783 early crossings, 766 late crossings). B. Fruit fly experiment in 0.4 m/s wind (1158 early, 828 late). C. Fruit fly experiment in 0.6 m/s wind (394 early, 294 late). D. Mosquito experiment in 0.4 m/s wind (125 early, 143 late).

late crossings. This shows that the reflexive maneuvers generated by this model are not sufficient to yield the history dependence we observed in the data.

An alternative model suggests that as an insect should bias its flight more strongly upwind as it gains information about the location of the plume centerline, that is, it should slowly switch from an explorative crosswind to an exploitative upwind strategy as it learns where the plume is located [23]. This model allows for the probability distribution over the source’s centerline to be updated in a Bayesian way with only simple arithmetic computations. Instead of using Grunbaum’s model directly, in order to study it we applied its general “centerline-inferring” principle to our baseline correlated random 3D walking model: instead of reflexively surging upon crossing the plume (as in the surge-cast model), whenever this modeled insect crossed the plume at a certain location, it updated the mean and uncertainty of its estimate of the plume’s centerline position according to Bayes’ rule. An additional biasing term then drove the insect upwind, with a magnitude that increased as the uncertainty of the centerline estimate decreased. We refer to this as the “centerline-inferring model”, an example trajectory of which is shown in Fig 5C. We additionally included a finite memory timescale, allowing the estimate over the plume centerline to decay back to its baseline over the course of about 10 seconds.

The history dependence of trajectories generated by the centerline-inferring model was reversed compared to the empirical data (Fig 3F). Specifically, late plume crossings yielded stronger upwind turns than early plume crossings, suggesting that this model likely cannot account for the sign of the history dependence observed in empirical trajectories. Note that neither the surge-cast nor centerline-inferring models used wind speed as a parameter, since flight velocities were calculated relative to the environment. While it would have been possible to include such a parameter in the models so that we could reproduce, say, the stronger upwind turns in higher winds (Fig 3A–3C), our goal here was to study history dependence, which was not systematically affected by wind speed, and so we excluded it to focus and simplify our model.

The final plume-tracking strategy we investigated, which is more theoretically principled but also more cognitively demanding, was “infotaxis”, an algorithm proposing that a tracking agent should move to maximally decrease the uncertainty over possible source locations in 3D space, with the specific form and parameters of the distribution derived from the statistics of turbulence [22]. An example trajectory is shown in Fig 5D. Though the “true” plume used in our infotaxis simulations was cylindrical, as in the experiments, it is important to note that the tracking agent’s internal model is of a turbulent one. That is, we simulated how an agent used to tracking natural, turbulent plumes, would




Fig 5: Example empirical and simulated trajectories. *Each of the simulated trajectories*

Continued from Fig 5: (B-D) began at the same location in the wind tunnel (green star) as the empirical trajectory (A) and was run for the same length of time as the empirical trajectory.

respond when introduced to a laminar one, a paradigm we thought best reflected the situation in the experiments. As with the centerline-infering model, we compared infotaxis to our data by examining the history dependence in early vs. late crossing-triggered heading time-series in a set of infotaxis trajectories generated within a simulated plume-containing wind tunnel. The wind speed parameter, which influences the agent’s internal plume model, we set equal to that from each experiment, assuming that an ideal agent would be able to measure the wind speed. We chose the coefficient of turbulent diffusivity, the principal parameter in infotaxis governing how far a tracking agent thinks the plume would naturally spread from a point source, so as to roughly reflect the length scale of the wind tunnel and the assumption that fluctuations in a non-laminar flow would be at about the same scale as the mean wind speed itself.

The infotaxis trajectories resembled the empirical data in the sign of the history dependence observed: late plume crossings in infotaxis trajectories tended to yield turns with a weaker or completely absent upwind component than early plume crossings (Fig 3G–3J), as in the real data (Fig 3A–3D). The magnitude of the odor-triggered turns, however, was quite different, suggesting that infotaxis does not fully capture all aspects of the plume crossing distribution. Nevertheless, this suggests that the true plume-tracking algorithm used by fruit flies shares certain features with infotaxis that cannot be captured by other normative tracking models.

To quantify how much a simpler model, such as the surge-cast algorithm, would have to be modified by an infotaxis-like model in order to recapitulate the magnitude of the history dependence observed in the data, we also investigated a hybrid model of crossing-triggered heading responses. Here, a new set of crossings was built using weighted sums of the surge-cast and infotaxis crossings. Specifically, for each infotaxis crossing we selected a random surge-cast crossing with the same crossing number (i.e. we matched a third crossing from an infotaxis trajectory with a third crossing from a surge-cast trajectory), and then generated a new hybrid crossing composed of $P\%$ of the surge-cast crossing and $(1-P)\%$ of the infotaxis crossing, i.e., $\Delta h_{hybrid}(t) = P/100 \times \Delta h_{surge-cast}(t) + (1 - P)/100 \times \Delta h_{infotaxis}(t)$. As with the other models, we then grouped the hybrid crossings into early vs. late categories and compared them quantitatively (Fig 6A). When we optimized the mixing percentage P to match the data, we found that the early vs. late crossing-triggered heading response could be best reconstructed with P




Fig 6: Analysis of hybrid surge-cast-infotaxis crossing model. *A. Hybrid surge-cast-infotaxis crossings, as described in text, separated into early vs. late crossings with a crossing model mixing 65% surge-cast with 35% infotaxis. B. The late vs. early difference in mean headings of the hybrid crossings (time-averaged from 0.3 to 0.8 s post-crossing) as a function of percentage of surge-cast making up the model (red), overlaid with that calculated from the 0.3 m/s wind speed fruit fly experiment (black). C. The distribution of data-matched surge-cast % over multiple random pairings of surge-cast and infotaxis crossings.*

equal to approximately 70% (Fig 6B and 6C), suggesting that an infotaxis-like history dependence gives about a 30% contribution to explaining the full crossing-triggered heading response.

To show that infotaxis reproduced the sign of the empirically observed history dependence but not necessarily other qualities of the data, we compared the distribution of positions within the wind tunnel generated by either the real trajectories or those simulated by infotaxis, as well as by the two other models. Notably, the position distribution of the empirical trajectories exhibited a salient peak at the upwind end of the wind tunnel (Fig 7A, reproduced roughly from [15]). This was true for the surge-cast and centerline-inferring simulations as well (Fig 7B and 7C). Notably, while the infotaxis simulation also yielded a salient peak in the upwind end of the wind tunnel, it exhibited an additional peak at the very downwind end of the wind tunnel (Fig 7D). This is likely due to the fact that in infotaxis it is usually preferable to initially travel to the downwind end of the wind tunnel in order to increase the probability of first encountering an odor hit. However, we note that in the infotaxis simulations corresponding to the 0.4 m/s and 0.6 m/s wind speeds, the strong peak in the upwind end of the wind tunnel was much less pronounced (S8 Fig). This, combined with the lack of a quantitative match between the crossing-triggered heading time courses in the data vs. in infotaxis, suggests that although infotaxis is able to capture certain features of the empirical data there are other key features that it fails to reproduce.

1.5 Discussion




Fig 7: Trajectory heatmaps for data and models. *Each distribution was calculated by binning all timepoints generated from all trajectories, empirical or modeled, used in our analyses. In this figure we have not excluded any timepoints based on position in the wind tunnel or heading. A. Empirical trajectories from fruit flies in a wind tunnel with an ethanol plume and 0.3 m/s wind. B-D. Equivalent trajectories generated by the surge-cast (B), centerline-inferring (C), or infotaxis (D) algorithms.*

To explore how animals enact sequences of decisions in response to complex stimuli in a natural context, we examined flying insects tracking odor plumes and asked whether deviations from the mean odor-triggered responses could be predicted from specific features of the odor encounters. We found that higher concentrations of odor led to stronger upwind turns, and that encounters occurring later in a sequence triggered weaker upwind turns. While a reactive surge-cast strategy predicted the mean odor-triggered upwind turn well, only trajectories generated by the infotaxis strategy [22] showed consistency with sign of the observed history dependence, though the magnitude of the history dependence in infotaxis was greatly exaggerated relative to that of the data. This suggests that sequential decision processes in real plume tracking may be dominated by reactive responses but may also be modulated by information about a source location accumulated over the course of a flight trajectory.

1.5.1 Measuring and controlling turbulent signals

Although it is relatively straightforward to reconstruct natural flight trajectories using videos, precise spatiotemporal patterns of a naturalistic odor signal are challenging to measure, since odor is usually advected by turbulent flow [24]. Instead one often measures the coarse-grained statistics of a turbulent signal, such as its time-averaged concentration, intermittency, or plume envelope, foregoing the ability to infer the odor concentration at precise timepoints, but allowing one to explore how trajectory features depend on these statistics [6, 7]. Alternatively one can present a stationary plume with a pre-measured profile, which allows one to infer the concentration experienced by a flying insect on a

moment-by-moment basis, but in which case the true signal is less naturalistic. In this analysis we have examined flight trajectories generated in the latter context, since it allowed us to infer the moment-by-moment sensory experience of each insect and to look for correlations with subsequent behaviors. Further, from the perspective of the fly, which follows a highly variable path through the wind tunnel, its encounters with the plume may be a sufficient approximation to the encounters it would have with a turbulent plume [15]. Importantly, although the plume itself is stationary, the variation in each trajectory yields variation in the sensory experience of each fly. This was key in allowing us to explore how different features of plume encounters elicited different navigational decisions.

1.5.2 Informative features of plume encounters

Our first result that flies respond to differences in the concentration experienced during an in-flight odor encounter suggests that this feature may contain usable information about the fly’s current position relative to the plume source. Indeed, it has been shown theoretically that the distribution of nonzero instantaneous concentrations depends on displacement from the plume source [13], and our results propose that flies detect and act on this information. This suggests that theoretical trajectory-generating plume-tracking strategies, such as infotaxis [22], should incorporate non-binary odor detections in order to better recapitulate observed animal trajectories; whether such information significantly improves tracking efficiency, however, remains an open question. For example, it may be the case that higher instantaneous concentrations indicate to the fly that it is closer to the plume centerline, such that it would behoove it to make a stronger upwind turn, as we observe in the data. Our result also corroborates the fact that flies have neural mechanisms capable of detecting graded levels of odor concentration, which have been observed at the level of ORNs [25, 26], but which may also be present in the antennal lobe, one synapse downstream from the detector neurons.

Our second result, that the history prior to an odor encounter influences a fly’s response to that encounter, suggests that the number of encounters throughout a naturalistic trajectory also contains information about the fly’s displacement from the source. For example, our observations are consistent with the possibility that increased encounters might indicate that the fly has moved closer to the source, such that it should diminish its upwind turns so as not to overshoot it. Indeed, such a strategy would be consistent with moth plume tracking behavior, in which it was seen that moths tended to slow their flight as they approached a pheromone source [19]. Alternatively, as the flies accumulate odor encounters they might be more inclined to investigate the visual features of the surrounding wind

tunnel, such as the floor pattern [15], which might slow their upwind progress.

More generally, the dependence of the flight trajectories on odor history is consistent with the flies' acting on information obtained from multiple, temporally separated samples of the plume, an idea fundamental to strategies based on extended accumulation of information [22, 23]. The neural mechanisms enabling history-dependent behaviors in plume tracking are unknown, but a simple possibility is that they result from adaptation of the peripheral olfactory neurons (ORNs) in the insects' antenna and maxillary palp, which typically decrease their sensitivity after extended odor exposure [27, 28]. ORN adaptation is typically studied, however, with much longer pulses than the sparse, intermittent puffs expected in a turbulent plume, so it is unclear how strongly this mechanism would influence natural tracking. An alternative is that history dependence arises from state changes in the antennal lobe (AL), the insect olfactory network immediately downstream of the initial receptors, analogous to the mammalian olfactory bulb. Indeed, many local neurons within the *Drosophila* AL substantially change their odor-evoked responses over the course of an approximately 10 second train of short pulses, with some exhibiting decreased, and others increased responses to later pulses in the train [29]. Interestingly, these AL response dynamics occur on about the same timescale as the history dependence we have observed from in the behavioral analysis. Determining the precise neural mechanisms underlying our results will be an exciting future task, and this suggests that exploring antennal lobe mechanisms may be a promising avenue for investigation.

An important factor in experiments involving flight trajectories through wind tunnels is the visual experience of the insect. Since wind tunnels are typically small due to technical constraints, i.e., between 1–2 m long and between 0.3–1 m wide, such that the insects can likely see the boundaries of the environment, it is likely the case that flight decisions are driven by a combination of both visual and olfactory stimuli. Indeed, the optomotor response of flying insects requires visual input in order for the flies to determine the wind direction relative to the ground [1], so that experiments cannot be performed in the absence of visual input. In our analyses we have attempted to control for visual effects in several ways. Since the main effect we expect is for the fly's distance to the upwind or downwind walls to influence their propensity to turn upwind (for example, they may be less inclined to turn upwind if there is little space for them to turn into), in our analyses of the empirical data we excluded plume crossings occurring in the upwind-most and downwind-most 30 cm of the wind tunnel. As an additional control, however, in each of our data analyses we subtracted out the best linear prediction of the quantity of interest given the position x_0 of the fly along the upwind-downwind axis of the wind

tunnel. Specifically, in quantifying the concentration dependence of plume-crossing-triggered heading, we first subtracted the best prediction of heading from x_0 and then calculated the correlation of the residual heading variation with plume-crossing concentration, with a similar analysis applied to the dependence of plume-crossing-triggered heading on the number of prior crossings in the trajectory. In the future, it will be useful to conduct plume-tracking experiments in larger wind tunnels, such that the visual experience does not substantially vary across different locations in the wind tunnel, but doing so at present remains a significant technical challenge.

1.5.3 Identifying evidence of theoretical search strategies in data

There has been much interest in elucidating theoretical principles of how insects track turbulent plumes [2]. However, due to the highly dynamical nature of flight trajectories, identifying evidence of specific search strategies in data is not straightforward. In particular, because the dynamical odor experience of each insect is controlled not by the experimenter but rather by the insect’s own movements, one cannot compare theoretical vs. empirically observed trajectories on a one-to-one basis. More precisely, one cannot force a simulated trajectory to have a data-matched odor experience without causing it to deviate potentially quite strongly from the algorithm that was supposed to have generated it.

Given these limitations, it is preferable to instead ask how the distributions over specific trajectory features vary between simulations and data. Since a key theoretical question asks about the extent to which insects accumulate and use information about the plume’s source location over extended time periods, we chose to ask how such information accumulation might affect the history dependence of navigational decisions observed in real trajectories vs. those simulated by theoretical algorithms. Consequently, we examined the joint distributions of plume-crossing-triggered heading responses and a variable indicating whether each plume crossing occurred early vs. late in the total sequence. Our findings were that although plume-crossing-triggered flight maneuvers were dominated by a brief upwind turn (Fig 1D), they indeed differed early or late in the trajectory (Fig 3A–3D, Fig 4). This suggests that although the real plume-tracking strategy used by the insects we explored likely involves a significant reactive component, it is also modulated by the insect’s history with the plume since it took to the air. This rules out plume-tracking strategies that are 100% reactive, such as the pure surge-cast strategy depicted in Fig 3E, and it suggests that including such history-dependent modulation may increase the efficiency of successful source localization.

Since the purely reactive surge-cast model (Fig 3E) did not reproduce the early-late distinction we observed in the empirical crossing distribution, we examined the history dependence predicted by two information-accumulation-based theoretical strategies: the centerline-inferring model, based on the model by Grunbaum and Willis [23], and infotaxis, in which movement of a tracking agent explicitly maximizes expected information about the location of the source in 3D space [22]. We found that the centerline-inferring model predicted a history dependence opposite to that seen in the data: later crossings generated stronger upwind turns. This inconsistency suggests that accumulated information about a plume’s centerline may not be directly converted into increased upwind advancement along the inferred centerline.

The failure of the centerline-inferring model motivates the consideration of infotaxis, an algorithm in which the tracking agent uses internal knowledge of turbulent plume statistics to interpret odor encounters and guide its search. Though the true plume used both in the original experiments and in our infotaxis simulations was cylindrical and contained within a laminar flow, we nonetheless allowed the simulated agent to assume an internal model of a more natural, plume governed by turbulent diffusion. Essentially, we wanted to ask how a search agent used to turbulence would behave when the plume was actually laminar and cylindrical; this, we felt, best approximated the experimental situation in which flies evolved to track natural plumes were introduced to an artificial one. Additionally, though it is possible the flies in the wind tunnel noticed and adapted to the real plume’s non-turbulent structure, given the high variability of their trajectories and the intermittency of their odor encounters, their dynamic flight and odor experience was likely reasonably similar to that one would expect in more naturalistic plume tracking. Performing similar experiments in the presence of turbulent, yet measurable plumes (measured, for example, using Schlieren imaging, a technique that could detect concentration changes through the optical distortions they cause over a patterned background [30, 31] will be an important step to confirming this hypothesis.

In infotaxis, though the magnitude substantially differed, the sign of the observed history dependence was consistent with infotaxis. One possible reason that infotaxis, and perhaps the empirical trajectories, exhibit weaker upwind turns for later encounters is because the physical scale at which information is gained is larger for earlier encounters and smaller for later encounters. That is, the first encounters in a trajectory confine the possible source location to a relatively large region of space upwind of where the encounters occurred, so that it advantages the tracking agent to fly upwind into this space to increase its chances of encountering the source or additional odor puffs. Later encounters,

however, provide primarily fine-scale information about the source, since the large-scale location may already be known. Thus, it may behoove the tracking agent to pay more attention to the specific details of where it thinks the source might be, as opposed to flying indiscriminately upwind.

In the experiments we analyzed, the fruit flies tracked ethanol, an attractive odorant that emanates from fruit, and the mosquitoes tracked carbon dioxide, which is produced by mammalian hosts. While we did not investigate plume tracking history dependence in the presence of other odors, if the behaviors we observed were indeed evidence of efficient source localization strategies we would expect to see the same history dependence when flies tracked other odors arising from sources with positive valence. Indeed, for other attractive odors, such as Vector 960 (an odorant used in some fruit fly traps), the reflexive odor-triggered turns are quite similar to those triggered by ethanol [15]. Investigating more detailed plume tracking features as a function of odor identity would be a fruitful topic of future research. For aversive odors, on the other hand, we would not expect the same behavioral responses, as the insects’ goal would presumably no longer be to localize and approach the odor source, and so one would expect different navigation strategies.

1.5.4 On the role of infotaxis in biological source localization

Although infotaxis was proposed in the context of biological plume tracking, few studies have compared theory to data. In fact, to our knowledge the only study that has directly compared infotaxis to animal behavior was that of Calhoun, Chalasani, and Sharpee [32], in which the authors showed that the foraging behavior of *C. elegans* in the absence of a food source exhibited transitions from local to global search strategies similar to those predicted by infotaxis. In addition to the already discussed challenge of comparing theoretical to empirical flight trajectories, infotaxis trajectories bear the challenge of depending on the parameters the tracking agent uses in its turbulent plume model, such as the source emission rate or turbulent diffusivity coefficient. While our qualitative results were relatively invariant to small changes in these parameters, the observed infotaxis history dependence changed sign when the source emission rate assumed by the insect decreased too sharply (S3 Fig). D , the turbulent diffusivity coefficient assumed by an infotaxis agent in its internal plume model, also strongly affects the source position information gained from each odor detection: with a very small D , odor detection implies the source is almost directly upwind, since little spreading should have occurred and which should strongly bias movement in that direction; with a larger D , however, odors can be detected from sources with more variable positions. When we re-ran our infotaxis simulations we saw

this was indeed the case: smaller D led to the predicted stronger upwind turns (S3 Fig). The sign of the history dependence, however, which was the key focus of our analysis, did not depend on D ; for both small and large D early plume-crossing-triggered turns had a stronger upwind component than late turns (S3 Fig). This shows that while trajectories differ in some ways when the internal plume model changes, certain coarse-resolution features remain the same.

Since an infotaxis agent uses these parameters to infer the source location given a sequence of odor encounters, a key question is how such an agent would learn these parameters in the first place, which can vary over several orders of magnitude in natural environments [33], and incorrectly estimated values of which can affect the efficiency of its search trajectory [34]. One possibility is that by incorporating a reactive component to their behaviors, as observed in the data we analyzed, a tracking strategy might gain robustness to variations in the plume statistics that would otherwise confound a pure infotaxis strategy. That is, a moderately strong reactive component might stabilize behavioral responses that would otherwise be sensitive to the parameters of the plume. Indeed, when plume-tracking robots were programmed with different search strategies, under conditions of high odor concentration, reactive strategies outperformed infotaxis [35], perhaps because of mismatches in the true plume parameters vs. those assumed by the robot. Future work could test this hypothesis explicitly by building a full trajectory-generating model that was a hybrid of infotaxis and a reactive strategy, and calculating source-finding efficiency under varying plume statistics.

1.5.5 The role of vision in search and infotaxis

An important component of an insect’s search strategy is to explore visual features after encountering attractive odors [15]. Their choice to explore a visual feature likely depends on many factors, which may include the concentration and frequency and number of odor encounters. Thus, it is conceivable that our observed difference in behavior for early and late plume crossings is influenced by the insect’s choice to explore visual features (such as the checkerboard pattern on the floor of the wind tunnel) rather than solely being a function of plume interactions. These same visual features, however, could also help to guide the animal’s surge and cast behavior by providing helpful landmarks.

Of the models we tested, only the infotaxis model captured the correct qualitative change in behavior between early and late plume encounters. The infotaxis model, however, requires that the animal has access to its exact position in world coordinates, something that is quite challenging for an insect flying in the wind over large distances (separating distance, ground speed, and wind speed, from

measurements of optic flow and air speed is not trivial [15]. Instead, it is possible that insects could use visual features as relative landmarks in an infotaxis-like strategy. We hope that our results will inspire the development of new biologically plausible search theories similar to infotaxis, but that take into account the limited information available to insects, which includes visual landmarks, but not exact location information.

1.6 Materials and Methods

All analyses and simulations were written in the Python programming language. All code is available at https://github.com/rkp8000/wind_tunnel.

1.6.1 Dataset

We used datasets previously described by van Breugel and Dickinson [15] and van Breugel et al. [21]. See the referenced works for more details. Briefly, groups of either fruit flies or mosquitoes were released in a 1.5 m x 0.3 m x 0.3 m wind tunnel (with a 1.3 m x 0.3 m x 0.3 m section available for flight) containing a gentle, laminar wind flowing along the long axis (30–60 cm/s wind speed). When any insect took off, a camera-based tracking system began recording it, and the insect’s 3D trajectory was later reconstructed from the video. In the experimental conditions we have focused on, a laminar odor plume was also present in the wind tunnel (ethanol for fruit fly experiments and CO₂ for mosquito experiments). The plume retained a stationary profile due to the laminar wind flow, and was approximately cylindrical in shape, with the highest concentration along the center of the cylinder. The odor concentration as a function of 3D position in the wind tunnel was measured prior to releasing any insects. For the ethanol plume in the fruit fly experiments, there was no discernible widening of the plume from the upwind to the downwind end of the wind tunnel, such that the plume’s cross section did not change along the wind axis. For the CO₂ plume some diffusion occurred, with the plume widening slightly with increasing x , i.e. towards the downwind end of the wind tunnel. However, since our history dependence analysis accounted for x before comparing early vs. late plume crossings, it is unlikely that the trend observed in mosquitoes tracking CO₂ was due to perceived changes in the plume cross section. With both plume types, however, the plume’s stationarity allowed the approximate concentration experienced by the insect to be inferred from its position. The resulting dataset that we used as our starting point thus contained registered time-series of experienced odor

concentration (up to a multiplicative factor), 3D position and 3D velocity (Kalman filtered with a constant velocity Kalman filter) at every time point in every trajectory of every insect.

1.6.2 Identifying plume crossings

We defined a plume crossing as the interval during which the odor concentration experienced by the insect rose above a certain threshold. To determine an optimal threshold for each species we asked which threshold best separated potential plume crossings (in which the insect entered the vicinity of the plume but may or may not have crossed it) into below- and above-threshold crossing groups. Specifically, for a range of candidate thresholds, we asked when the difference between the post-crossing headings in the above-threshold group (averaged over crossings and time) was maximally different from that in the below-threshold group. We reasoned that if our chosen threshold were below its true value, the above-threshold group would get diluted with non-crossings, making it more closely resemble the below-threshold group; if our chosen threshold were above the true value, however, then the below-threshold group would get diluted by true crossings, making it more closely resemble the above-threshold group. An optimal threshold that ideally separates the below- and above-threshold groups should exist somewhere in between. However, upon performing this analysis a large uncertainty arose surrounding the calculated difference between the two groups' heading time series (S1 Fig), such that the optimality of the chosen threshold could not be guaranteed. To keep our analysis straightforward, we chose an ethanol threshold approximately corresponding to the "elbows" in S1A Fig (for fruit flies) and S1D Fig (for mosquitoes), corresponding to 0.0009% ethanol and 430 ppm CO₂, respectively, to define plume crossings.

Importantly, however, none of our results depend on the precise threshold used to define a plume crossing, since an additional threshold was either kept as a free parameter (Fig 2) or the chief comparison was made between groups defined by the same plume-crossing threshold but which differed along other dimensions, such as crossing history. To provide further evidence that the precise crossing threshold value did not affect our qualitative results, we re-ran the history dependence analysis from Fig 4 using the 0.3 m/s wind speed fly-tracking-ethanol experiments, but extracting crossing sets according to a range of thresholds that varied over almost two orders of magnitude (S5 Fig). We found, however, that regardless of which threshold we chose, the results followed the same trend: late crossings yielded turns with a weaker upwind component than earlier crossings, all else held constant. Additionally, we re-ran our hybrid model analysis using a wide range of thresholds (S7 Fig) but also

found that this parameter did not substantially affect our results.

In all of our analyses we chose $t = 0$ to correspond to the peak odor concentration experienced during the plume crossing (e.g., as in Fig 1D). We defined the end of each plume crossing as the time when the insect either reentered the plume (the experienced concentration rose above threshold again) or landed.

1.6.3 Data exclusion criteria

In order to limit the influence of geometrical confounds on the results of our analysis, unless otherwise noted we excluded the following sets of plume crossings from all empirical and model-generated datasets when performing our analysis. First, we excluded all crossings occurring in the upwind- or downwind-most 30 cm of the wind tunnel. This was motivated by the possibility that flies already in the upwind end of the wind tunnel may exhibit more crosswind-oriented flight simply because they have less of an area to turn into upwind of them. Second, we excluded all crossings that did not occur at a crosswind angle, i.e., in which the initial crossing heading was not between 60 degrees and 120 degrees from upwind. We made this exclusion simply because we were seeking out upwind directed changes in heading, which would not be visible if flies were already heading upwind. Relaxing these criteria did not change the qualitative results of our analysis of the dataset, but they did introduce confounding factors into our simulated trajectories. For example, when we included the upwind-most 30 cm of the wind tunnel in the surge-cast model analysis (Fig 3E), late crossings indeed appeared more crosswind, simply because many late crossings occurred in this upwind-most portion, where the simulated agents' contacting of the upwind wall often prevented them from turning as upwind as in the earlier crossings that occurred in the central portion of the wind tunnel.

1.6.4 Partial correlation analysis

To calculate the partial correlation between peak concentration experienced during a plume crossing and the fly's heading at time t following the crossing (Fig 2), we first fit a linear model predicting the heading $h(t)$ from the initial heading $h_0 = h(0)$ and the position x_0 along the long axis of the wind tunnel. We then subtracted this prediction from $h(t)$ to yield a residual heading $h'(t)$, and we subsequently calculated the correlation between peak concentration c_{peak} and $h'(t)$. Given this framework, any correlations between c_{peak} and $h'(t)$ cannot be through the confounding factors of initial heading or position in the wind tunnel.

1.6.5 Threshold and threshold-linear model fitting

We defined the threshold model as:

$$h(t)a_h h_0 + a_x x_0 + h_{<} + h_{>} \Theta(c_{peak} - c_{th})$$

where $h(t)$ is the heading t seconds after the plume crossing; h_0 is the heading at the time of the plume crossing, scaled by parameter a_h ; x_0 is the position at the time of the plume crossing along the upwind/downwind axis, scaled by parameter a_x ; $h_{<}$ is the heading predicted when the peak crossing concentration c_{peak} is below the threshold c_{th} ; $h_{>}$ is the additional heading change predicted when the concentration was above c_{th} , and Θ is the Heaviside step function, which is 0 below 0 and 1 above 0.

We defined the threshold-linear model as:

$$h(t)a_h h_0 + a_x x_0 + h_{<} + (h_{>} + a_c c_{peak}) \Theta(c_{peak} - c_{th})$$

where a_c is an additional parameter specifying the linear dependence of heading on concentration when c_{peak} is greater than the threshold c_{th} .

In both models we fit a_h , a_x , $h_{<}$, $h_{>}$, and c_{th} . In the threshold-linear model we additionally fit a_c . The threshold model is thus “nested” within the threshold-linear model. For both models we chose as the best fit parameters those which minimized the squared difference between the true and the model-predicted headings. To determine whether the threshold-linear model fit the data significantly better, we used an F-test. When calculating the p-value, however, we replaced the number of crossings with the number of unique trajectories, since each trajectory yielded multiple crossings, thus providing a more conservative estimate of significance by not treating all crossings as independent occurrences.

1.6.6 History dependence analysis

To prevent geometrical factors from confounding the results of our comparison of early vs. late plume crossings, we first constructed a new variable $h^*(t)$ representing the change in heading at time t that remained after linear predictions of $\Delta h(t)$ from x_0 and t_{flight} (the time spent flying since takeoff) were subtracted from $\Delta h(t)$. In particular, we fit the predictor coefficients and calculated the predictions of $\Delta h(t)$ using all non-excluded plume crossings, regardless of crossing number. Following the calculation of $h^*(t)$ for each crossing at each time point, we grouped crossings into early and late classes as before

and compared $h^*(t)$ between the two groups using a t-test. Finally, since there were occasionally multiple non-excluded crossings from one trajectory in the same class, when calculating the p-values shown in Fig 4 we used as the number of data points in each class the number of unique trajectories, as opposed to the number of unique crossings.

In an additional related analysis, also aimed at uncovering a relationship between crossing number and post-crossing heading, instead of splitting crossings into early vs. late crossings we simply calculated the partial correlation between crossing number and Δh time-averaged from 350–450 ms post-crossing. This partial correlation was calculated by first subtracting out the best linear predictions of the time-averaged Δh from x_0 and t_{flight} (S4 Fig).

1.6.7 Base model for surge-cast and centerline-inferring tracking algorithms

We based both our surge-cast and centerline-inferring algorithms on a correlated random walker, described by:

$$\tau \dot{\mathbf{v}} = -\mathbf{v} + \boldsymbol{\eta} + b\mathbf{w}_C$$

where v is the tracking agent’s velocity, with the dot indicating its time derivative, τ is the timescale of turning, $\boldsymbol{\eta}$ is a Gaussian white noise variable sampled independently at each timestep with zero-mean and diagonal covariance matrix with diagonal entries η , w_c is a unit vector in the crosswind plane pointing toward the centerline of the wind tunnel, and b is a positive scalar multiplying w_c . Before augmenting the base model to include either surging or centerline-inferring features, we fit τ , η , and b to approximately match the speed, angular velocity, and crosswind position distributions of a simulated agent to those of the empirical data (S6 Fig). This yielded values of $\tau = 0.42$ s, $\eta = 1.9$ m/s and $b = 0.25$ m/s. Trajectories generated by the base model alone thus followed a correlated random walk structure with a 0.42 s timescale and with a moderate “casting” bias of oscillating around the wind tunnel centerline.

In both the surge-cast and centerline inferring models, we modeled the wind tunnel to have the same geometry and the plume to have the same concentration profile as in the empirical data. We then analyzed the resulting plume crossings using the same exclusion criteria as we did with the real data.

1.6.8 Surge-cast model

To add a surging component to the base model just described we augmented it with an additional upwind bias term such that

$$\tau \dot{\mathbf{v}} = -\mathbf{v} + \boldsymbol{\eta} + b\mathbf{w}_c + a\mathbf{w}_u$$

where w_u is a unit vector pointing upwind, and $a(t)$ is the convolution of an alpha function with timescale 0.07 s and maximum amplitude 0.8 m/s and a train of delta functions placed at the time points corresponding to the peak odor experienced in each crossing. The timescale and amplitude of the alpha function were chosen so that the simulated tracking agent's plume-crossing triggered responses exhibited approximately the same dynamics as observed in the empirical plume-crossing data.

1.6.9 Centerline-inferring agent simulation

To generate trajectories for the centerline-inferring algorithm we augmented the base model with a different upwind bias term:

$$\tau \dot{\mathbf{v}} = \mathbf{v} + \boldsymbol{\eta} + b\mathbf{w}_c + k\mathbf{w}_u$$

Here k is a positive-valued scalar that increases as the uncertainty of the plume's centerline location decreases. In particular, we let $k = k^* / |K|$, where k^* is the upwind bias and K is the covariance matrix of a posterior belief probability over the 2D centerline location (y^*, z^*) , whose determination is given shortly. Furthermore, in the centerline-inferring agent, we let w_c be a vector that points not to the centerline of the wind tunnel, but rather to the maximum a posteriori estimate of the centerline of the plume. As with the surge-cast model, we chose the scaling factor for the upwind bias such that the plume-crossing triggered responses showed approximately the same dynamics as in the data.

The 2D belief distribution over the plume centerline location (y^*, z^*) was updated in a Bayesian way each time the simulated insect crossed the plume at location (x_n, y_n, z_n) (with n indexing the plume crossing number):

$$P(y^*, z^* | y_n, z_n, y_{n-1}, z_{n-1}, \dots) \propto P(y_n, z_n | y^*, z^*, y_{n-1}, z_{n-1}) P(y^*, z^* | y_{n-1}, z_{n-1}, \dots)$$

We assumed that the likelihood and prior (the two terms on the right-hand side of the equation,

respectively) were both Gaussian with diagonal covariances. Thus, the posterior (the term on the left-hand side of the equation), the likelihood, and the prior could each be parameterized by a mean and covariance matrix. Note that in this model the posterior after one crossing acts as the prior at the next crossing. We chose the likelihood covariance K_s and the covariance K_0 of the prior over (y^*, z^*) before any crossings had occurred to have diagonal entries $(2m)^2$ and $(5m)^2$, respectively. Finally, to model short-term memory decay over several seconds, we allowed the mean of the posterior to decay exponentially to $(0, 0)$ over $\tau_m = 10$ seconds, with the posterior covariance K decaying to K_0 over the same timescale.

1.6.10 Infotaxis simulation

The infotaxis algorithm we used was developed by and detailed in [22]. In this algorithm a tracking agent moves along a rectangular lattice in order to localize the source of an odor plume. To model turbulence, the odor plume is not assumed to have a concentration gradient, but rather to yield random “hits”, with hit probability maximized immediately downwind of the source. The tracking agent maintains a belief distribution over the possible source locations that is updated at each time step, depending on whether the agent receives a hit or a miss. This update is Bayesian and thus depends on an explicit formula for the probability of receiving a hit a function of displacement from the source, that is, the likelihood. The principle difference between the Bayesian update in the centerline-infering model vs. in infotaxis is that in the centerline-infering model the belief distribution is a 2D Gaussian over the centerline location, whereas in infotaxis the distribution is over the 3D source position coordinate and has no parametric form. Since hit probability depends on the agent’s position, and hits and misses update the source position distribution in different ways, the tracking agent moves so as to maximize the expected decrease in uncertainty of the source position distribution, as measured by its Shannon entropy.

The likelihood formula is derived in [22] from the time-averaged hit rate one would expect at different displacements from a source in a turbulent medium. Importantly, this hit rate depends on D , the effective diffusivity of the turbulent medium [24]. When D is small, the plume is confined to a narrow band downwind of the source, and when D is large the plume becomes much more spread out, and can even yield hits upwind of the source.

We ran our simulations on a 3D lattice that simulated positions accessible within the wind tunnel, with lattice points 2 cm apart from one another. A discretized version of the non-turbulent cylindrical




Fig S1: Determination of concentration thresholds for each experiment. *The thick line shows the difference between the mean plume-crossing-triggered heading time-series for crossings above the threshold and the mean plume-crossing-triggered heading time-series for crossings below the threshold, time-averaged over the first one second following the plume crossing. Shading represents uncertainty, calculated by propagating the standard errors of the means of each group through the difference calculation. Each panel corresponds to a different experiment, labeled by the insect and the wind speed.*

plume used in the experiments was simulated in the center of the lattice, and the tracking agent received a hit at every time step it was inside the plume. Note, however, that even though the simulated plume was cylindrical and non-turbulent, the actions of the tracking agent in our simulation were driven by its assumption that it was interacting with a naturalistic, turbulent plume. For each experimental flight trajectory we generated one corresponding infotaxis trajectory that started at the same (but discretized) location in the wind tunnel as the experimental trajectory and which lasted for the same number of timesteps as the number of lattice points that the experimental trajectory traversed.

1.7 Supporting information




Fig S2: Positional differences between early and late plume crossings. *Each panel shows the distribution of the upwind/downwind components of plume-crossing positions for either the early or late groups shown in Fig 3, measured at the time of the plume crossing. As in Fig 3, we have excluded all crossings occurring in the most upwind or most downwind 30 cm of the wind tunnel (leaving 70 cm of valid flight space for our analysis) and of the remaining crossings we have included only those in which the heading at the time of the crossing was between 60 and 120 degrees. Each panel corresponds to one insect/wind-speed, with the difference in the mean plume-crossing upwind/downwind position (x) shown in meters in the title.*




Fig S3: Infotaxis history dependence for varied plume estimate parameters. *Same layout as in Fig 3E–3H, but with different infotaxis parameters. In A–B, the source emission rate R is 1000Hz, and the turbulent diffusivity coefficient D is varied, as indicated in the panel titles. In C–E, the turbulent diffusivity coefficient D is 0.09 m²/s, with R varied, as indicated in the panel titles.*




Fig S4: Partial correlation between heading change and crossing number, conditioned on x_0 and t_{flight} . Each panel corresponds to a different experiment. The x -axis shows x_0 the x -position at the time of the crossing, and the y -axis shows the change in heading, time-averaged from 350 to 450 ms post-crossing. Each point corresponds to one crossing, with the color of the points denoting the crossing number (c_n ; crossing numbers of 1 or 2 correspond to early crossings and 3, 4, 5 to late crossings; we did not include the small number of crossings with crossing number ≥ 5). The partial correlation coefficients between crossing number and the $\Delta\text{heading}$, conditioned on x_0 , and the corresponding p -values, are shown in the panel titles. When calculating p -values, we used the number of unique trajectories, rather than the number of crossings, since one trajectory frequently contained multiple crossings.




Fig S5: Effects of varying odor detection threshold on history dependence analysis. *In each panel the analysis from Fig 4 is performed on a set of plume crossings where crossings are defined as trajectory portions in which the odor concentration rises at least once above a minimum detection threshold (varying by panel and listed in panel titles). All panels show crossings calculated from trajectories in which flies were tracking ethanol in a 0.3 m/s wind speed. As crossing-detection threshold varies over more than an order of magnitude, the key history dependent features remain qualitatively constant.*




Fig S6: Distributions of kinematic quantities in base model (before surge-cast or centerline-inferring features were added) vs. empirical data. Each panel shows the distribution of speeds (A), angular velocities (B), or crosswind positions (C) (calculated across all time point) of the empirical trajectories (black) vs. the trajectories generated by best-fit base model for the surge-cast and centerline-inferring models. Y-axis units are arbitrary and denote relative proportions of time points.




Fig S7: Effects of varying odor detection threshold on hybrid model analysis. Here we show the optimal surge-cast percentage in the surge-cast-infotaxis hybrid crossing model analysis introduced in Fig 6 as a function of the odor detection threshold used to generate and extract crossings from the trajectories. The points and error bars show the mean and standard deviation of the threshold-dependent distributions equivalent to Fig 6C.




Fig S8: Position distributions for data and infotaxis results for different wind speeds. Equivalent to Fig 6A and 6D. A, B show position distributions for empirical data and infotaxis simulations in a wind tunnel with 0.4 m/s wind speeds, respectively. C, D show the same for a 0.6 m/s wind speed.

References

- [1] John S. Kennedy and D. Marsh. Pheromone-regulated anemotaxis in flying moths. *Science*, 184(4140):999–1001, 1974.
- [2] John Murlis, Joseph S. Elkington, and Ring T. Carde. Odor plumes and how insects use them. *Annual review of entomology*, 37(1):505–532, 1992.
- [3] D. R. Webster and M. J. Weissburg. Chemosensory guidance cues in a turbulent chemical odor plume. *Limnology and oceanography*, 46(5):1034–1047, 2001.
- [4] Ring T. Carde and Mark A. Willis. Navigational strategies used by insects to find distant, wind-borne sources of odor. *Journal of Chemical Ecology*, 34(7):854–866, 2008.
- [5] Hiroshi Ishida, Takamichi Nakamoto, Toyosaka Moriizumi, Timo Kikas, and Jiri Janata. Plume-tracking robots: A new application of chemical sensors. *The Biological Bulletin*, 200(2):222–226, 2001.
- [6] MARTIN Geier, OLIVER JÖRG Bosch, and JÜRGEN Boeckh. Influence of odour plume structure on upwind flight of mosquitoes towards hosts. *Journal of Experimental Biology*, 202(12):1639–1648, 1999.
- [7] Agenor Mafra-Neto and Ring T Cardé. Fine-scale structure of pheromone plumes modulates upwind orientation of flying moths. *Nature*, 369(6476):142–144, 1994.
- [8] Eugene Yee, DJ Wilson, and BW Zelt. Probability distributions of concentration fluctuations of a weakly diffusive passive plume in a turbulent boundary layer. *Boundary-Layer Meteorology*, 64(4):321–354, 1993.
- [9] Christopher M Finelli, N Dean Pentcheff, Richard K Zimmer-Faust, and David S Wetthey. Odor transport in turbulent flows: constraints on animal navigation. *Limnology and Oceanography*, 44(4):1056–1071, 1999.
- [10] Paul A Moore, Marc J Weissburg, J Michael Parrish, Richard K Zimmer-Faust, and Greg A Gerhardt. Spatial distribution of odors in simulated benthic boundary layer flows. *Journal of Chemical Ecology*, 20(2):255–279, 1994.
- [11] Jelle Atema. Eddy chemotaxis and odor landscapes: exploration of nature with animal sensors. *The Biological Bulletin*, 191(1):129–138, 1996.

- [12] Ring T Cardé, Gabriella Gibson, et al. Host finding by female mosquitoes: mechanisms of orientation to host odours and other cues. *Olfaction in vector-host interactions*, 2010:115–142, 2010.
- [13] Antonio Celani, Emmanuel Villermaux, and Massimo Vergassola. Odor landscapes in turbulent environments. *Physical Review X*, 4(4):041015, 2014.
- [14] Neil J Vickers and Thomas C Baker. 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, 1994.
- [15] Floris van Breugel and Michael H Dickinson. Plume-tracking behavior of flying drosophila emerges from a set of distinct sensory-motor reflexes. *Current Biology*, 24(3):274–286, 2014.
- [16] Teun Dekker and Ring T Cardé. Moment-to-moment flight manoeuvres of the female yellow fever mosquito (*aedes aegypti* l.) in response to plumes of carbon dioxide and human skin odour. *Journal of Experimental Biology*, 214(20):3480–3494, 2011.
- [17] Neil J Vickers. Mechanisms of animal navigation in odor plumes. *The Biological Bulletin*, 198(2):203–212, 2000.
- [18] Neil J Vickers. Winging it: moth flight behavior and responses of olfactory neurons are shaped by pheromone plume dynamics. *Chemical senses*, 31(2):155–166, 2006.
- [19] MA Willis and EA Arbas. Variability in odor-modulated flight by moths. *Journal of Comparative Physiology A*, 182(2):191–202, 1998.
- [20] Philip Coen, Jan Clemens, Andrew J Weinstein, Diego A Pacheco, Yi Deng, and Mala Murthy. Dynamic sensory cues shape song structure in drosophila. *Nature*, 507(7491):233, 2014.
- [21] Floris Van Breugel, Jeff Riffell, Adrienne Fairhall, and Michael H Dickinson. Mosquitoes use vision to associate odor plumes with thermal targets. *Current Biology*, 25(16):2123–2129, 2015.
- [22] Massimo Vergassola, Emmanuel Villermaux, and Boris I Shraiman. ‘infotaxis’ as a strategy for searching without gradients. *Nature*, 445(7126):406, 2007.
- [23] Daniel Grünbaum and Mark A Willis. Spatial memory-based behaviors for locating sources of odor plumes. *Movement ecology*, 3(1):11, 2015.
- [24] Boris I Shraiman and Eric D Siggia. Scalar turbulence. *Nature*, 405(6787):639, 2000.

- [25] AJ Grant, JG Aghajanian, RJ O'Connell, and BE Wigton. Electrophysiological responses of receptor neurons in mosquito maxillary palp sensilla to carbon dioxide. *Journal of Comparative Physiology A*, 177(4):389–396, 1995.
- [26] Vikas Bhandawat, Gaby Maimon, Michael H Dickinson, and Rachel I Wilson. Olfactory modulation of flight in drosophila is sensitive, selective and rapid. *Journal of Experimental Biology*, 213(21):3625–3635, 2010.
- [27] Marien De Bruyne, Peter J Clyne, and John R Carlson. Odor coding in a model olfactory organ: The drosophila maxillary palp. *Journal of Neuroscience*, 19(11):4520–4532, 1999.
- [28] Katherine I Nagel and Rachel I Wilson. Biophysical mechanisms underlying olfactory receptor neuron dynamics. *Nature neuroscience*, 14(2):208, 2011.
- [29] Katherine I Nagel and Rachel I Wilson. Mechanisms underlying population response dynamics in inhibitory interneurons of the drosophila antennal lobe. *Journal of Neuroscience*, 36(15):4325–4338, 2016.
- [30] Hubert Schardin. Die schlierenverfahren und ihre anwendungen. In *Ergebnisse der exakten Naturwissenschaften*, pages 303–439. Springer, 1942.
- [31] Michael John Hargather and Gary S Settles. Natural-background-oriented schlieren imaging. *Experiments in fluids*, 48(1):59–68, 2010.
- [32] Adam J Calhoun, Sreekanth H Chalasani, and Tatyana O Sharpee. Maximally informative foraging by *Caenorhabditis elegans*. *Elife*, 3:e04220, 2014.
- [33] Luther H Allen Jr. Turbulence and wind speed spectra within a Japanese larch plantation. *Journal of Applied Meteorology*, 7(1):73–78, 1968.
- [34] Juan Duque Rodríguez, David Gómez-Ullate, and Carlos Mejía-Monasterio. On the performance of blind-infotaxis under inaccurate modeling of the environment. *The European Physical Journal Special Topics*, 226(10):2407–2420, 2017.
- [35] Nicole Voges, Antoine Chaffiol, Philippe Lucas, and Dominique Martinez. Reactive searching and infotaxis in odor source localization. *PLoS computational biology*, 10(10):e1003861, 2014.

2. Multiplexed dopamine signals in the fruit fly mushroom body

The following work was undertaken in collaboration with experimental neuroscientists Ari Zolin, Raphael Cohn, and Vanessa Ruta at Rockefeller University. Unless otherwise noted, all experiments were designed and conducted by Zolin et al., and I largely led the formal analytical work detailed in this chapter. My primary experimental contribution was designing the stimulation sequence for the open loop sensory paradigm (Fig 6). All figures and writing are my own.

2.1 Abstract My preceding results suggested fruit flies can accumulate information over the course of an odor-tracking flight and use it to reshape subsequent odor responses. What mechanisms might support the rapid reconfiguration of neural computations required for this and other behaviors that unfold over multiple seconds? One compelling component of the insect brain that might subserve this is the mushroom body (MB), a third-order sensory and memory area whose internal information flow is under rapid control by dopaminergic inputs (DANs) from upstream areas, as shown by Cohn et al. [1]. Whereas dopaminergic circuitry is typically associated with external reward and punishment, Cohn et al. (2015) noticed that DAN activity in the MB tracked not only rewarding (sugar) and aversive stimuli (electric shock), but also value-neutral behaviors like walking initiation, suggesting modulation of MB information flow may be ongoing and driven by more than valence signals alone. The precise relationship between DAN activity and observable sensorimotor variables, however, remains unknown. In collaboration with Ari Zolin, Raphael Cohn, and Vanessa Ruta at Rockefeller University, I've addressed this problem by analyzing simultaneous time-series recordings of DAN activity, sensory stimulation, and behavior in the fruit fly to characterize the DAN-sensorimotor relationship. Our results indicate that (1) DAN activity contains multiplexed representations of several sensorimotor variables, including wind direction, walking state, walking speed, turning velocity, and odor; and (2) DAN representations of sensorimotor variables change under the presence of appetitive odors, potentially recruiting anticipatory representations of behaviors the fly will undertake several seconds in the future. This suggests that DAN-based control of information flow through the mushroom body may depend on external sensory variables, internally generated motor commands, and reward—all key variables in reinforcement learning; further, in addition to flexibly reconfiguring MB information flow, the DAN representations may exhibit a state-dependent flexibility themselves. These results shed new light on how dopaminergic signalling may shape neural processing and behavior over rapid timescales.

2.2 Introduction

A hallmark of even the simplest nervous systems is their remarkable functional flexibility. Despite having only about 100,000 neurons, for instance, the fruit fly is capable of a diverse repertoire of behaviors spanning multiple timescales, from rapid escape reflexes [2] to odor-tracking [3, 4] and courtship [5, 6, 7, 8] to long-term associative learning [9, 10, 11], among many others. How could relatively simple circuitry underlie such manifold behavioral functions? One possible mechanism conserved across many animals is the reshaping of neuronal processing via neuromodulatory signalling molecules, such as serotonin, norepinephrine, octopamine, and dopamine [12]. By temporarily altering cellular or synaptic function, these neuromodulators can endow a biological neural network with more varied dynamics and computational properties than its baseline circuitry might suggest. Interestingly, while neuromodulatory activity and the circuit dynamics it shapes have traditionally been associated with changes in global behavioral states like quiescence or arousal [13, 14, 15, 16], exciting new lines of work are emerging suggesting that neuromodulatory signals might be multidimensional and exert much more fine-grained control over neural processing than previously thought [1, 17, 18]. Unraveling the dynamics and function of such fine-grained neuromodulatory control may be crucial to our understanding of how neural circuitry supports the flexible computations needed to navigate an ever-changing world.

2.2.1 The role of dopamine in computation and behavior

Perhaps having received the most scrutiny for its computational properties is the neurotransmitter dopamine, which is present in both vertebrate [19, 20, 17] and invertebrate nervous systems [21, 22, 23] and appears to have at least a partially conserved role across species. Disrupted in several disorders, such as addiction [24, 25, 26, 27] or schizophrenia [28, 29, 30], dopamine has been primarily investigated for its role in signalling the rewards and punishments experienced by an animal. In particular, the foundational finding by Schultz et al. [31] that firing rates of dopaminergic neurons in the mammalian midbrain tracked not reward itself, but rather the difference between expected and received reward—the so-called *reward prediction error*—sparked a rich line of inquiry into dopamine’s functional role in the nervous system. In particular, dopamine has been strongly implicated as a potential key player in a biological implementation of reinforcement learning (RL) [32, 33], in which an agent uses trial-and-error learning to develop a decision-making policy by assigning a scalar value to external states and internal actions [34]. This framework has been used, for instance, to develop computational models

of addiction in an effort to gain a more principled understanding of the affliction [35]. Dopaminergic responses to reward and punishment [36], and potentially reward prediction error [37] have also been observed in the fruit fly, demonstrating the potential conservation of dopaminergic RL computations across many species.

A parallel line of research into dopaminergic function suggests that beyond its role in RL, dopamine may be more generally involved in the representation and control of an animal’s motor and sensory experience. An impairment of the nigrostriatal dopamine system, for example, is one of the key biomarkers of Parkinson’s Disease, whose sufferers exhibit a substantially disrupted capacity for movement initiation [38, 39, 40, 41]. In the sensory domain, it was shown that burst firing of dopamine neurons in cat ventral tegmental area (VTA) could be evoked by unconditioned visual and auditory stimulation [42], with similar firing patterns observed in dopamine neurons in rat VTA in response to vibrissal stimulation [43]. The tuning of dopaminergic responses to sensory stimuli has also recently been observed in zebra fish [17], suggesting this sensorimotor role of dopamine may be conserved across species as well. Beyond this, recent work has begun to tease apart potential dopaminergic responses to valence-neutral motor states in fruit flies [1]. Tying together the diverse roles of dopamine both within and across species will be a central step in gaining a full understanding of how this neuromodulator interacts with and shapes behaviorally relevant neural computations.

2.2.2 Dopaminergic control of mushroom body computations

One dopaminergic system that is being increasingly investigated for its rich role during behavior is the dopaminergic innervation of the fruit fly mushroom body (MB), a third-order insect sensory area. Lying immediately downstream of the antennal lobe (the insect homolog to the mammalian olfactory bulb [44, 45]), the fruit fly MB consists of approximately 2000 excitatory *Kenyon cells* (KCs) that synapse onto 34 *output neurons* (MBONs) according to a stereotyped, compartmentalized organization [46]. Fruit fly KCs are divided into three functionally independent classes, α/β , α'/β' , and γ , each of which arborize along individual lobes which are themselves spatially and functionally subdivided into multiple compartments [47, 48, 46]. The γ lobe, for instance, is subdivided into five visually and genetically distinguishable compartments G1-5, which may have different roles in mushroom body olfactory processing [36]. KC-to-MBON synapses are themselves subject to stereotyped, compartmentalized innervation by dopaminergic inputs (DANs) from upstream areas [49, 23, 36], suggesting they may be a prime target of neuromodulatory control. Functionally, the mushroom body has been implicated in the

representation of olfactory and gustatory stimuli by sparse KC activations [50, 51], and it appears to be a neuropil critical for associative learning. For example, when their mushroom bodies were chemically lesioned, flies were unable to learn associations between odor cues and an electric shock, despite being able to do so in the absence of the lesion [10]. Finally, whereas some studies have suggested relative homogeneity in the functions of dopamine neurons in a mammalian midbrain area VTA [52], DAN activity in different MB compartments appears to convey differential signals, with distinct DAN clusters activating in response to either reward or punishment, rather than a single DAN signal encoding a scalar level of valence [36]. Overall, the mushroom body’s involvement in odor representation and learning, its multidimensional dopaminergic neuromodulation, and its relatively simple organization makes it an ideal and potentially tractable system for comprehensively investigating dopamine’s role in shaping neural computation and behavior.

One of the most recent and exciting studies on dopaminergic modulation of MB information processing comes from [1], who found that DAN activation could yield large, rapid, reversible, and compartmentally localized effects on KC-to-MBON synaptic transmission. In particular, they found that an isolated *unpaired* one-second stimulation of the DANs innervating γ -lobe compartment G4 caused KC-to-MBON synapses in G4 to become and remain strongly potentiated; however, when DAN stimulation was *paired* with simultaneous activation of a subset of KCs, this caused the synapses of the activated KCs onto their downstream MBONs to become strongly depressed. Furthermore, both potentiation and depression of KC-to-MBON synapses could be subsequently reversed simply by applying either paired or unpaired stimulation to the preparation, respectively. As a result, when a second pulse of an odor initially paired with G4 DAN activation was presented, it evoked a weaker G4 MBON response due to the depressed synaptic transmission from KCs responsive to the odor. This observation of rapid DAN-based synaptic modification thus suggests a mechanistic underpinning of the hypothesized preferential involvement of the MB γ -lobe in short-term memory [53]. In addition to observing the expected compartmentally specific DAN responses to reward vs punishment, Cohn et al. also observed that DAN responses appeared to track spontaneous behavioral state not explicitly associated with reward or punishment. Indeed, when flies “flailed” within the tethered preparation, the G2 and G3 DANs showed substantive increases in activity, the same DANs activated during electric shock; when flies were still, the G4 and G5 DANs exhibited increased activation, the same DANs activated by the presentation of a sucrose reward. This suggests that DAN activity in the mushroom body and its effect on olfactory computations and behavior may be modulated not only by external positive or

negative stimuli but by internally generated motor activity as well, suggesting a multi-dimensional and nuanced function of the insect dopamine system.

The above results beg two questions. First, what combination of stimulus, motor, and reward events drive or are driven by different activity dynamics in the different DAN compartments in the MB? Second, how do these shape information processing through the KC-to-MBON synapses to eventually modify behavior in natural contexts? In collaboration with experimentalists Ari Zolin, Raphael Cohn, and Vanessa Ruta, we have attempted to characterize the DAN responses to a variety of sensorimotor observables and to postulate their involvement in the representation and transformation of behavior. To do so we studied flies’ locomotor patterns as they walked on a trackball with simultaneous recordings made of γ -lobe DAN activations in either or both MBs under a variety of conditions. Our results suggest DANs impinging on the MB γ lobe may represent a wide range of multiplexed signals, and that they may potentially be involved in predicting flies’ tracking behaviors over a timescale of several seconds into the future. We discuss the computational implications of these findings in the Discussion.

2.3 Results

2.3.1 Absent external stimuli, DAN activity represents moment-to-moment behaviors

We first explored DAN activity while flies walked on a trackball without additional sensory stimulation. DAN activity in the different gamma lobe compartments was measured via a synaptically tagged calcium indicator, while the fly’s walking velocity was reconstructed from a simultaneous video recording of the trackball (Fig 1A). Prior to analysis, the video of the DAN activity from one hemisphere was segmented by hand into four regions of interest (ROIs), corresponding to the gamma lobe regions $\gamma 2-5$ (G2-5) (Fig 1B). This yielded time-locked recordings at 10 Hz of multiple walking motion variables and 4 DAN time-series (Fig 1C).

When we overlaid the fly’s total walking motion (the net rotational velocity of the trackball) with each of the 4 DAN traces, we observed a strong correlation between the DAN fluctuations and the fly’s walking state (walking or paused) throughout the trial (Fig 4D). In particular, inspection of example traces revealed that while G2 and G5 exhibited much variability, G3 and G4 tended to exhibit a strong, consistent relationship with walking state, often with an apparent lag of less than one second. Further, additional experiments by Zolin, Cohn, and Ruta (not shown) suggest that the DAN fluctuations accompanying walking state are on the same order of magnitude as DAN responses to sucrose reward.

This demonstrates that the dopaminergic projections to the mushroom body gamma lobe appear to encode the fly’s moment-to-moment behavioral state, absent any external stimulation or reward.

To quantify the relationship between DAN activity and behavioral state we fit a simple logistic-regression classifier predicting the fly’s current state from the DAN activity alone. For all time points that could clearly be assigned a walking or paused label, we trained the classifier to predict which of these two states the fly was in from a symmetric 1-second window of either one single DAN or all DANs together. For each trial we trained the classifier on 80% of the time points and tested it on the 20% of held-out time points, where time points were grouped in extended “chunks” prior to the training-test split. In doing so we found G3 to be the best predictor of walking state, with an average of 98% accuracy (Fig 1E), and which dropped only to 97% when we included only “fitful” trials (trials where the fly spent at least 20% of the time paused) (Fig 1F-G). Other DANs contained some predictive power of walking state, but given the strong predictive power of G3, including them as additional predictors along with G3 only marginally improved predictive accuracy (Fig 1G). This shows that the moment-to-moment walking state fluctuations can be almost entirely predicted by the moment-to-moment fluctuations in G3 alone.

We next asked whether the relationship between walking state and DANs other than G3 might be explained simply by correlations of those other DANs with G3, as would arise if there were internal network interactions among the DANs. To this end, we computed cross-correlation functions between all pairs of DANs (Fig 2). Consistent with the observed G2-G3 and G4-G5 correlations observed when flies are presented with stimuli of positive and negative valence (sugar reward, electric shock, respectively) [1], we observed strong positive G2-G3 correlations and strong positive G4-G5 correlations in the asensory context described in Fig 1, in which behavior and DAN activity were purely spontaneous (Fig 2B,E,L,O). Further, consistent with the G2-G3 vs G4-G5 opponency observed in the presence of positive- or negative-valence stimuli, we observed a strong negative correlation between G2 and G4. Compared to other pairs of compartments, there was a relatively weak correlation between G3 and G4, suggesting that the positive relationship of both of these with walking state (Fig 1D,G) is probably driven by both DAN populations independently encoding walking, as opposed to, say, G3 following walking, and G4 being simply driven by G3. We note, however, the potentially bimodal distribution of G3-G4 correlations (Fig 2G, J), suggesting they may also be dependent on an unobserved internal state that varies trial-by-trial. These results suggest that either behaviorally tied or purely internal network dynamics among the four DAN populations in question drive spontaneous correlation structures that




Fig 1: Asensory experiments and initial observations. *A.* Flies walked on a patterned trackball whose rotation was videotaped to reconstruct a fictive trajectory. Mushroom body activity in one hemisphere was simultaneously imaged. *B.* Schematic of mushroom body architecture. Kenyon Cells (KCs) synapse onto mushroom body output neurons (MBONs) in four distinct anatomically and functionally segregated "gamma" compartments (G2-5). Dopaminergic inputs (DANs) to each compartment synapse onto KC → MBON synapses themselves, allowing selective modulation of synaptic transmission. In experiments by Zolin et al., DAN synapses were tagged with a calcium indicator,

Continued from Fig 1: *such that activity could be reconstructed from observed fluorescence. C. Example recordings of basic behavioral variables (Motion [net ball rotational speed], ForVel [forward rotation], AngVel [angular/rotational velocity around vertical axis]) and DANs across the four gamma compartments G2-5. All traces have been normalized to have same scale on plot. D. Example overlay of each DAN activation trace (colors) with total ball motion (black). All trials were 300 seconds long. E. Accuracy of walking vs paused classification using any individual DAN vs all DANs over a one-second window centered at the time point being classified. Error bars show standard deviations in accuracy. F. Distribution across trials of fraction of time spent walking. Trials with less than 80% of time spent walking were deemed “fitful” trials. G. Same as E but for fitful trials only.*

reflect the correlation structure observed in the reward/punishment context [1], but that multi-DAN representations of walking may be partially independent.

Motivated by G3’s strong predictions for walking state, we formally quantified the temporal relationship between each DAN compartment and multiple features of the fly’s walking behavior. For each trial, we fit linear filters predicting either total motion, forward velocity, (signed) angular velocity, or absolute (unsigned) angular velocity on the trackball, using as predictors an 8 second window of DAN activity centered at the time point of the behavior prediction (Fig 3). (Example overlays of each DAN with each of the behavioral variables we examined are shown in Fig 3S1.) Our results show that for the most part, DANs exhibited a lagged encoding of behavioral variables, with most DANs responding to a weighted time average of any given behavioral variable over approximately 1-3 seconds in the past. Most filters were predominantly monophasic, indicative of a simple moving average, e.g. G4 and G5 from forward velocity (Fig 3G-H) or G2 from absolute angular velocity (Fig 3M). Interestingly, however, a couple of filters took on unexpected shapes. G2, for example, exhibited a salient biphasic filter on forward velocity (Fig 3E), suggesting that it might be encoding a time derivative, i.e. forward acceleration. The extension of the positive lobe of this filter beyond $\Delta t = 0$ suggests G2 may show some “anticipatory” coding of movement, rather than simply following the fly’s most recent motion; that is, for G2 to encode the time-derivative of forward velocity at $\Delta t = 0$, it would require some knowledge of the fly’s forward velocity 500-1000 ms in the future, suggesting it may be involved not just in reading out the most recent forward velocity but in shaping it as well. Note, however, that due to the 500 ms lag in the calcium indicator signaling, it is difficult to draw strong conclusions about anticipatory coding at this timescale. G4 also exhibited an interesting filter shape on absolute angular velocity, which, though weak, also appeared to have a secondary lobe extending into the future and suggesting a possible anticipatory code of this quantity. G5’s filters appeared to simply reflect those observed for G4, suggesting that its computational distinction from G5 might only be evident




Fig 2: Cross covariance functions between distinct DANs. *Each timepoint shows the cross covariance between two DANs at a given time lag. All cross covariances were normalized by the product of the standard deviations of the signals to yield a maximum value of 1. Light gray traces are individual trials, and thick red lines show the average across trials.*

in alternate behavioral contexts. These results thus demonstrate that G2-4 encode different features of behavior in distinct ways, some potentially anticipatory, and combined with the result of Fig 2 suggest the DAN system is a complex of internal network interactions superimposed with behavior-driven and/or behavior-driving activity patterns. Thus, beyond simply coding rewarding and punishing stimuli, the DAN system in the mushroom body additionally exhibits a rich and diverse representation of moment-to-moment behaviors.

2.3.2 In sensory-motivated context DANs represent stimulus features

To examine DAN representations in a more naturalistic setting, we turned to data collected by Zolin et al. in a “sensory-motivated” context. Here, flies walking on a track ball were additionally subject to a steady air stream controlled by the fly’s own turning/angular velocity on the track ball (Fig 4A). We defined the fly’s heading as the angle that the air stream was offset relative to the direction the fly was facing (Fig 4B). The setup was calibrated in a “virtual reality” sense, so that when the fly turned right, the air stream moved left, and vice versa (Fig 4B-C). Additionally, an appetitive odor (apple cider vinegar [ACV]) could be injected into the air stream at computer-controlled times and in such a way that the total air flow remained approximately constant during and outside of odor presentation.

Flies generally kept the air stream in front of them and showed salient behavioral and DAN responses to odor onset. Heading over time within and outside of the odor for an example trial is shown in Fig 4D, where following odor onset relatively little variation in heading is observed beyond an approximately 90 degree range centered in front of the fly, possibly indicating a tracking behavior reflective of a search for the appetitive odor source. As in their representations of behavior, individual DANs exhibited a diversity of odor responses. In particular, the odor-triggered G2 average showed a suppressed response to odor onset (Fig 4E), G3 showed a moderate average increase to odor with substantial variability (Fig 4F), and G4 and G5 showing strong positive responses to odor onset (Fig 4G-H). Additionally, G3 exhibited a salient offset response to the odor, suggesting G3 (Fig 4K) might signal unsigned changes in odor presence. Neither G2, G4, nor G5 showed substantial responses to odor offset (Fig I-L). Finally, while substantial variability remained, ACV onset on average tended to trigger increases in forward walking velocity, which was additionally reflected in odor-triggered increases in total motion (Fig 4M-N). Increases in forward velocity appeared to last through the duration of the odor presence, with no discernable decreases following odor offset (Fig 4Q-R). Neither absolute angular velocity nor absolute heading, however, showed any overt responses to odor onset or offset




Fig 3: Linear filters predicting different DANs from different behavioral variables. *Linear filters were fit using standard least-squares regression. Negative time lags correspond to the behavioral time point preceding the DAN time point being predicted. Filters were fit predicting DANs G2-5 using either total ball motion (Motion), forward walking velocity (ForVel), angular/rotational velocity (AngVel), or absolute (unsigned) angular velocity (AbsAngVel). Gray lines indicate individual trials and red lines show cross-trial averages.*




Fig 4: “Closed loop” sensory setup and odor-triggered responses. *A. The sensory setup from Fig 1 was augmented by introducing an air stream that was controlled by the fly’s angular velocity so as to imitate a virtual reality environment. Experiments were conducted in darkness so the fly could not see the air tube. Additionally, mushroom body gamma lobes in both hemispheres were imaged in the sensory setup. B. Heading was defined as the angle between the direction the fly was facing*

Continued from Fig 4: *and the air tube. C. Example overlays of AngVel and Heading in the closed loop setup. D. Example heading trace over a full trial, with a 60-second pulse of ACV shown in red. E-H. Odor-onset-triggered DAN responses, averaged across hemispheres (e.g. $G2A = (G2L + G2R)/2$). Gray lines are individual trials and black lines cross-trial averages. I-L. As in E-H but for odor-offset-triggered responses. M-T. As in E-L but for behavioral quantities, including absolute (unsigned) heading (AbsHeading).*

(Fig 4O-P, S-T).

Overall, these results show that in addition to representing moment-to-moment behaviors the DANs may carry instantaneous sensory information as well; indeed, if one reasonably imagines ACV as being predictive of a future food source, this is consistent with the canonical understanding of the mammalian dopamine system [54] in which dopaminergic activity increases not only in response to reward, but in response to reward predictors. More generally, our results suggest G2-4 likely play different roles in the representation and processing of olfactory signals at the level of dopaminergic inputs to the mushroom body. Furthermore, the opposite G2 and G4 responses to odor appear to reflect the negative correlations between G2 and G4, suggesting that DAN representations of sensory information may be shaped by the same inter-DAN dynamics that drive spontaneous activity in the sensory environment.

2.3.3 Interhemispheric DAN differences encode air stream angular movements

We next characterized the relationship between the DAN activations and various sensorimotor observables in the sensory-motivated context. As before, we fit filters linearly predicting DAN activity from an 8-second window of a given sensorimotor observable centered at the time of the DAN activity being predicted (Fig 5). As in the sensory context we observed salient filters on motion, forward velocity, and absolute angular velocity that predicted G2-4 in distinct ways, with the G5 filters largely resembling the G4 filters (Fig 5A-H). Of note, in the sensory-motivated context, G4 showed a substantially stronger filter on forward velocity than G3 (Fig 5F-G), as compared to the sensory context, in which G3 and G4 filters on forward velocity were of approximately the same magnitude (Fig 3F-G). Flies in the sensory-motivated context spent much less time not walking, consequently undergoing many fewer walking state transitions than in the sensory case. Our results in Fig 5 are thus consistent with the hypothesis that G3 largely represents walking state whereas G4 may more richly represent the specific value of forward walking velocity, giving rise to a filter that remains salient even when the fly walks continuously.


DAN-predicting filters on absolute angular velocity in the sensory-motivated context also had

shapes resembling those in the asensory context (Fig 5I-J). However, in the sensory-motivated context G2 and G4 filters on absolute angular velocity were on a similar order of magnitude as the G3 filter, whereas in the asensory context, the G3 filter was much larger than the G2 and G4 filters. This may be due once again to the more continuous walking in the sensory-motivated context; it could also potentially arise as a consequence of sensory feedback that allows the fly to more precisely compute its angular velocity relative to the external environment. When we fit DAN-predicting filters on absolute (unsigned) heading (Fig 5M-P), a quantity with no direct meaning in the asensory context, we observed strongly salient biphasic filter shapes predicting G2 and G4, extending for 2-4 seconds into the past, with a small 1s filter predicting G3. This suggests that in a sensory-motivated context, the fly’s orientation relative to a continuous air stream is more strongly encoded in G2 and G4 than in G3, and that G2 and G4 are predicted by distinct biphasic features of the recent sequence of headings experienced by the fly. The filtering properties of the various DANs on behavior did not exhibit any strong differences when fit using only time points within vs outside of the odor epoch (not shown).

Finally, we took advantage of the fact that in the sensory-motivated context Zolin et al. were able to record from both hemispheres simultaneously, to ask whether signed sensorimotor variables might be encoded in the interhemispheric differences between DAN signals. Indeed, when our predicted variable was the difference, G4D, between the left and right G4 activations, we observed salient, stereotyped filters on both signed angular velocity (Fig 5Q) and signed heading (Fig 5R). These filters were largely monophasic and extended for 1-2 s into the past, relative to the time point of the predicted G4D signal. This suggests that, “riding on top” of the sensorimotor-driven fluctuations in the DAN signals averaged across hemispheres, fluctuations in the left vs right DAN signals relative to the hemisphere-averaged signals may encode sensorimotor information useful for wind-guided navigation (anemotaxis). No DAN compartments besides G4 exhibited salient representations in their interhemispheric differences of the sensorimotor variables we considered in this analysis.

We next asked whether the heading- and angular-velocity-dependent fluctuations in G4D (the interhemispheric difference in G4 DAN activities) were representing sensory or motor information, or a combination of the two. Since the original “closed loop” sensory-motivated experiment paired the fly’s angular velocity with its heading (the angle of the air stream), in these experiments it was impossible to distinguish which of the two G4D might be representing. In coordination with Zolin et al., we therefore designed a novel experimental paradigm in which closed loop stimulation (necessary to motivate the fly to walk consistently) was interleaved with 60-second periods of “open loop” stimulation (Fig 6A),

Filters predicting DANs averaged across hemispheres



Filters predicting difference between interhemispheric G4 DANs




Fig 5: Linear filters predicting DAN activity from behavior in closed loop sensory setup. A-P. As in Fig 3A-P but in the closed loop sensory condition. Q-R. Best fit linear filters predicting the difference between G_4 levels in the two hemispheres ($G_4D = G_4L - G_4R$), using either (signed) *AngVel* or (signed) *heading* as a predictor.

in which the air stream angle was driven by an Ornstein-Uhlenbeck process (a parsimonious version of temporally filtered white noise) [55] instead of the fly’s own motion. Our hope was that by dissociating the sensory and motor observables during the experiment we could assess their independent and joint relationships with G4D.

To identify whether G4D represented sensory or motor information or both, we fit linear filters predicting G4D during the open-loop epochs of the experiments. In particular, we attempted to predict G4D using filters fit only to the fly’s angular velocity or only to the air stream’s angular. In doing so we found that while the angular velocity and air stream angular velocity filters both had salient integrator shapes over 500ms - 2s in scale (Fig 6B-C, Red), the filtered version of the fly’s angular velocity explained between 10-20% of the G4D variance, the filtered version of the air stream’s angular velocity was able to explain 30-60% of the G4D variance (Fig 6D). This suggests G4D is better predicted by the candidate sensory feature (air stream angular velocity) than by the motor feature (the fly’s angular velocity).

To account for the possibility that the fly’s angular velocity varied over a different timescale than the air stream angular velocity, potentially distorting the results of linear filter fitting and consequent variance explained, we additionally fit a model predicting G4D from both filtered signals. Indeed, when we allowed our model to use either the fly’s or the air stream’s angular velocity as predictors, or both together, we found that nonzero coefficients were only assigned to the air stream’s angular velocity (Fig 6B-C, Black). Furthermore, when both filtered variables were allowed to be predictors, the G4D variance explained only increased by a few percent relative to that explained from the air stream’s angular velocity (Fig 6D). These results together suggest that G4D, the interhemispheric difference between the G4 DAN activities, may be driven largely by the fly’s sensory experience rather than its angular motor outputs. To further confirm this finding, we note that in example traces, the G4D signal predicted purely from the air stream angular velocity often closely follows the true G4D signal, even when the fly is not walking at all, lending further evidence to G4D’s encoding of predominantly sensory information.

Finally, we note that in the sensory-motivated context, inter-compartmental DAN correlations (Fig 6S1) appear largely identical to those computed in the asensory context (Fig 2). This suggests that although the DANs recruit novel representations of sensorimotor variables in the sensory-motivated context, their internal network dynamics may be similar between contexts; however, there may still be context-dependent differences in higher-order correlation structures. Beyond intercompartmental

correlations, interhemispheric correlations between equivalent DANs across hemispheres appear largely symmetric, with G2 almost perfectly correlated between hemispheres, and with G3 and G4 exhibiting strong but still somewhat variable correlations, in line with our observation that G4D fluctuations may encode additional information useful to the fly. In sum, in combination with the results we have discussed until this point, it appears that the DAN activations encode not only valence-tied stimuli, but also both moment-to-moment behavioral variables and moment-to-moment sensory variables, and that these representations are often multiplexed within the same DAN compartment. In other words, modulation of olfactory information flow and transformations through the mushroom body may be controlled by a variety of sensorimotor and valence variables multiplexed into a relatively low-dimensional modulatory signals.

2.3.4 DANs recruit extended sensorimotor representations in the presence of odor

Given mammalian dopamine’s canonical association with rewarding stimuli [54], we asked how reward-associated stimuli for the fruit fly might alter DAN function relative to other sensorimotor variables. To this end, Zolin et al. modified the odor stimulation paradigm such that multiple 10-second odor pulses were presented during trials, interleaved with 30-second periods of clean air (Fig 7A). Additionally, to increase behavioral variation they decreased the air flow rate, as the higher air flow of Fig 4-6 tended to entrain flies into keeping the air stream directly in front of them. Interestingly, this changed the baseline relationships between odor, DANs, and behavioral variables. In particular, in this lower air flow condition, odor pulses did not trigger an average increase in forward velocity (Fig 7B-C), indicating that behavioral responses became less stereotyped. Also, as in previous experiments, odor onset did not trigger any salient average change in absolute heading (Fig 7D-E). Furthermore, whereas the odor-triggered DAN responses were not as sharp as in the higher flow conditions, they had roughly the same shapes, with G2 showing a negative and G4 a positive response to odor onset, and G3 exhibiting a weak offset response. Notably, in the lower air flow condition, the once prominent absolute heading filters predicting G2-G4 became far less salient, with only a barely discernible positive and negative lobe just prior to the time point of the prediction in G2 and G4, respectively, with no apparent structure in the G3 or G5 filters (Fig 7J-M). The filters predicting G2-G5 from forward velocity, however, were largely unchanged in the lower vs higher air flow conditions. This suggests the lower air flow condition with pulsed odors leads to less stereotyped odor responses but DAN-behavioral relationships similar to in the higher flow condition.




Fig 6: Distinguishing sensory vs motor representations in G4D. *A. Example of a 60-second epoch of open-loop movement of the air stream angle. The y-axis is in units of degrees per second, and G4D is in arbitrary units and has been scaled for visualization. B, C. Linear filters predicting G4D from either the fly's angular velocity or the air stream angular velocity, respectively (red), or using both signals independently (black). The x-axis is relative to the G4D time point being predicted. D. G4D variance explained using only the fly's angular velocity, only the air stream's angular velocity, or both. E. True G4D (dark green) and G4D predicted from the air stream angular velocity (light green), with the fly's scaled angular velocity shown below, with three non-walking periods indicated in magenta.*




Fig 7: Multi-odor pulse closed loop sensory paradigm. *A. Demonstration of fly's motion and the 4 DAN values (averaged across hemispheres), with two 10-second odor pulses shown in red. Missing values are due to the recording system occasionally dropping signals. B-E. Odor-triggered averages for 4 behavioral variables (B-E) and the 4 DAN variables. Gray lines are for individual pulses, with dark black indicating the average over pulses. J-Q. Linear filters predicting each of the 4 DAN values from*

Continued from Fig 7: *an 8-second window surrounding the time point of the DAN prediction, using either absolute (unsigned) heading or forward velocity as predictors. Gray lines are individual trials, and red are averages.*

Taking advantage of the behavioral variability in the lower air flow condition, we subsequently asked whether the way in which DANs represent external variables might be modulated by the presence of an odor. To address this question, we focused in particular on the relationship between different DANs and the representation of absolute heading, a core variable needed for the fly to generate behaviors according to its orientation relative to the wind. Of note, when we fit linear filters predicting different DANs from heading using epochs either inside or outside odor pulses, there was no noticeable difference in filter shapes (not shown). As an alternative approach, we isolated then aggregated all odor pulses across flies and trials and computed correlations between DAN levels and absolute heading at different offsets relative to the odor pulse onset (Fig 8A-D, E-F). This analysis revealed that there were indeed changes in the correlations between specific DAN activity levels and absolute heading in response to the onset of the appetitive odor. For example, the G2-absolute-heading and G3-absolute-heading correlations increased following odor onset and decreased again after odor offset (Fig 8A,B). In particular, this correlation increased across multiple seconds, indicating that during the odor pulse, G2 and G3 both increased their correlations not only with present absolute heading, but also with heading multiple seconds in the past and future. Thus, it would appear that the G2 or G3 DAN responses immediately after an odor contain a predictive power of future heading that we had not observed in prior analyses. G4 exhibited an opposite but similarly strong response, with prominent negative correlations with a wide temporal range of headings following odor onset (Fig 8C). Overall, this suggests that at a correlational level odor modulates DAN representations of at least one crucial navigational variable.


As an aside, we note that due to the frequent extended odor-triggered G2 and G4 DAN responses, these DANs exhibited an increase in their own autocorrelation timescale following odor onset (Fig 8S1); thus, one should question whether the increased DAN-behavior correlation timescale is simply reflective of the increased autocorrelation timescale in the DANs. While we have not fully explored this possibility, consideration of the following null hypothesis suggests this is probably not the case. Specifically, if we assume relatively short autocorrelation time in heading, as suggested by Fig 8S1D, then there should be zero correlation between absolute heading at one time $|h_t|$ and absolute heading a few seconds later $|h_{t+\Delta t}|$. Given this, if an arbitrary variable X were to increase its correlation strength with $|h_t|$, it would still remain uncorrelated with $|h_{t+\Delta t}|$ unless the correlation existed through a dif-

ferent pathway. Therefore, if we consider the null hypothesis that the increased DAN autocorrelation timescale is due to a random increase in DAN activity (or even a $|h_t|$ -dependent increase) that endures throughout the odor pulse, this will indeed increase the DAN autocorrelation time; however, for the above reason its correlations with $|h_{t+\Delta t}|$ should remain zero unless either (1) the heading autocorrelation time similarly increases during odor, or (2) the DAN- $|h_{t+\Delta t}|$ correlations arise through a different pathway. Since (1) does not appear to be the case, i.e. the absolute-heading autocorrelation timescale does not appear to increase substantially in the presence of odor (Fig 8S1D), it would seem that only (2) remains a likely option. Thus, we believe it is likely that odor recruits increased DAN correlations with absolute heading that in fact reflect a change in the relationship between the two variables.

To ensure our findings were indeed a product of the odor pulse and not a byproduct of our specific analysis, we repeated the analysis on “false odor pulses” offset by a certain number of seconds relative to the actual odor pulses. As shown in Fig 8I-L,M-P, we did not see the same prominent change in DAN-heading correlations when we triggered our analysis on false odor pulses. Moreover, since the false odor pulse shown in Fig 8I-L were offset 10 seconds prior to the real odor pulses, the initial results using real odor pulses are visible in the lower right-hand corner of the correlation matrices, corresponding to the expected 10 s offset. Thus, our results were due to the presence of the odor and not the structure of our analyses. Notably, some of these correlations correspond to anticipatory representations of the fly’s future behavior; for example, the correlation between G4 2-seconds-post-odor-onset and heading 10-seconds-post-odor-onset suggests that in the presence of odor G4 may be predictive of the navigational decisions the fly may make several seconds in the future. Overall, these combined results suggest that the DAN representations of absolute heading indeed strengthen in the presence of an appetitive odor, with these correlations stretching over an extended, multi-second timescale that is not seen outside of odor.

Finally, we asked whether the DANs’ anticipatory representations of future absolute heading could arise from the relationships we have already observed in previous analyses. For example, could it be the case that odor-triggered G4 simply reflects the odor-triggered change in forward walking velocity, which in turn is predictive of the fly’s change in heading a short time later? To address this question we defined several quantities chosen to assess the statistical structure of each DAN value, forward velocity, and absolute heading at a coarse temporal resolution. To this end we defined early and late epochs as 1-4 seconds and 7-10 seconds, respectively, following the onset of the 10-second odor pulse (Fig 9A-F), and used these epochs to compute 3-second time averages over different variables. For each DAN we

Odor-pulse triggered correlations



False-odor-pulse triggered correlations




Fig 8: Odor-triggered DAN-heading correlation matrices.. *A.* Correlation values (green = positive, pink = negative) between normalized G2 fluorescence and the fly's heading, offset at different times relative to the ($t=0$) odor pulse. For example, the point at $t = 2$ s on the x-axis and $t = 10$ s on the y-axis shows the correlation (across all pulses) between G2 at 2 seconds post odor onset and absolute heading 10 seconds post odor onset. *B-D.* Same as *A* but for G3-5. *E-H.* Corresponding P-values for correlation matrices in *A-D*, with blue indicating $P < 0.05$. *I-P.* Same as *A-H* except for false/control odor pulses offset 10 seconds prior to true odor pulses.

defined a *reduced model* as a linear model predicting the early DAN level, e.g. $\Delta G_{4,4}$ (time-averaged from 1-4 s post odor onset) from early heading $\Delta h_{1,4}$ and early forward velocity $\Delta v_{1,4}$. To account for the possibility that the odor-triggered DAN response was affected by the fly’s initial orientation when the odor came on, we also included the initial heading $h_{-10,0}$, time-averaged for 10 seconds prior to odor onset, in our reduced model. In other words, the reduced model predicts the early odor-triggered DAN response following odor onset using the early odor-triggered heading and forward velocity responses. Our reasoning was then the following: if the early DAN response contained more information about late heading $h_{7,10}$ than that contained in early heading $h_{1,4}$ and early forward velocity $v_{1,4}$, i.e. if it contained a true, independent anticipatory representation of heading, then augmenting the reduced model by adding in $h_{7,10}$ as an additional predictor should significantly increase the model’s predictive capacity of the early DAN value.

Indeed, we found that augmenting our reduced model by adding late heading $\Delta h_{7,10}$ as an additional predictor significantly increased the model’s prediction of early G2, G3, and G4 ($\Delta G_{2,4}$, $\Delta G_{3,4}$, and $\Delta G_{4,4}$, respectively) (Fig 9G-J, Red). Precisely, including late heading led to a 10% increase in explained variance for $\Delta G_{2,4}$ and $\Delta G_{3,4}$, and a 5% increase for $\Delta G_{4,4}$ ($P < 0.01$ for $\Delta G_{2,4}$, $\Delta G_{3,4}$, and $\Delta G_{4,4}$; F-test). Importantly, these increases were only observed during odor pulses, as when we ran the same analyses on false (control/sham) odor pulses offset either 10, 20, or 30 seconds prior to the true odor pulse, adding late heading as a predictor yielded no substantial improvement in explaining DAN variance (Fig 9G-J, Blue, Dark Blue, Black). (We note, however, that as expected, building up the reduced model by adding the reduced predictors $h_{-10,0}$, $\Delta h_{1,4}$, $\Delta G_{4,4}$ one-by-one frequently improved the model’s prediction of the various DANs, often regardless of whether or not the odor pulse was present.) This result suggests that at least at a coarse resolution, the early odor-triggered G2-4 responses may themselves be predictive not only of the odor’s presence and the fly’s instantaneous sensorimotor experience, as was suggested by our prior analyses, but additionally of how the fly will orient relative to the wind several seconds in the future. Thus, the presence of an appetitive odor for the fly may induce the recruitment of DAN activity patterns that anticipate subsequent behaviors, suggesting that the DANs may additionally be involved in context-dependent computations spanning multi-second timescales. An open question is how such representations would shape plasticity and information flow through the mushroom body to potentially give rise to or modulate these behaviors.

2.4 Discussion




Fig 9: Explaining odor-triggered DAN variance with future heading changes. *A-F. Depictions of variables used in building the predictive model. Gray shading shows the odor pulse; in A-D blue shading indicates the “early” epoch (1-4 s post odor onset) used to compute the DAN value; in E green shading shows the early epoch used to compute the forward velocity value; and in F green shading shows the early and late epochs used to compute the heading value. G. Fraction of early G2 variance explained as a function of which predictors were included in the model, for true odor*

Continued from Fig 9: *pulses (red) and false/control odor pulses at several offsets relative to the true odor pulse (blue, dark blue, black). Single stars indicate $P < 0.05$ and double stars indicate $P < 0.01$ (F-test; nested model analysis). H-J. Same as G but predicting early G3-5 instead of G2.*

Although the computational role of dopamine may be better characterized than that of any other neurotransmitter [31, 46, 1, 20, 17, 18], the molecule’s nuanced role in shaping neural function and behavior outside explicit reward contexts remains largely uncharted territory. By characterizing a battery of responses of dopamine neurons (DANs) innervating the fruit fly mushroom body and examining how these responses change under different conditions, we have attempted to shed light on how dopamine may be involved in the reconfiguration of information flow through a highly organized and potentially tractable neural system. While much work remains to fully contextualize prior dopamine studies and our own observations within a common framework, we hope our results provide a useful stepping stone for illuminating the multiplexed roles of dopaminergic computation during natural behaviors and for directing future experiments to further treat this question. Below we describe possible implications of our results and expound on how they may relate to specific learning frameworks.

2.4.1 Moment-to-moment sensorimotor representations in DAN inputs to the mushroom body

The γ -lobe of the fruit fly mushroom body has been shown to be involved in short-term memory [53] and to induce rapid, reversible, and compartmentalized changes in the efficacy of synaptic transmission from Kenyon cells (KCs) to mushroom body output neurons (MBONs) [1]. Our core novel results were the characterization of what information about a fly’s behavioral experience was represented by DAN projections to γ -lobe compartments G2-G5. In doing so we hoped to shed light on what features of a fly’s behavioral experience might drive KC-to-MBON plasticity, and how this might shape future behavior. We found that in addition to representing rewarding and punishing experiences [36], the DANs tracked a diverse array of sensorimotor variables on a moment-to-moment timescale, including walking state/speed, forward and rotational components of the fly’s self-driven locomotion, odor onsets and offsets, and both signed and unsigned heading/heading changes relative to angle of an upwind air stream. While some responses, such as the G4D response to heading, were relatively small compared to the reward-driven DAN fluctuations [1], others, such as the G3 and G4 response to walking state, were on a similar order of magnitude. This suggests that dopaminergic control of MB olfactory information processing may be driven not only by events with positive or negative valence but by

internally generated behavior as well, and perhaps by valence-neutral sensory stimuli.

While future experiments are needed to identify the precise magnitudes and timecourse of KC and DAN activations to induce KC-to-MBON synaptic plasticity, our results allow us to speculate as to how different behavioral experiences for the fly might modulate odor representations and responses over short timescales. First, in G4, when an odor is paired with reward the KCs representing the odor should weaken their MBON projections such that future encounters with that odor yield weaker MBON responses, as observed by [1]. Such a finding is in line with the canonical neural phenomenon of *adaptation* observed across a diversity of organisms and neural systems, in which neural responses to repeated stimuli gradually and/or rapidly decrease [56, 57, 58, 59]. Generally, adaptation is thought to subserve a favored representation of novel over familiar stimuli, for example to conserve metabolic resources by allocating spikes primarily to changes in the environment. The results of Cohn et al. [2015] suggest that in MBONs this may occur at the presynaptic level under the control of dopamine. Indeed, the DAN-triggered potentiation of non-activated KC-to-MBON synapses predicts that absent additional feedback processes, odors that were not paired with DAN activity should evoke much stronger MBON responses in G4.

Our current results add further nuance to this idea, suggesting that odor responses in G4 MBONs might adapt not only when paired with reward, but also with other behavioral events, such as odor onset itself or locomotion, which activate the MB-projecting DANs in a similar way as explicit reward. That an appetitive odor like apple cider vinegar (ACV) should evoke dopaminergic responses is perhaps not surprising given the well established observation that mammalian dopaminergic neurons learn to represent reward predictors following an associative conditioning paradigm [31, 54], and here it suggests that MBON responses to ACV should decrease regardless of whether ACV was paired with reward. DAN-based adaptation of odor response to locomotion, on the other hand, could point to a potentially novel role of dopamine in modulating neural information processing. In particular, if locomotion-induced DAN activity had similar effects on rapid KC-to-MBON plasticity as reward, this would imply that G4 MBON responses to odors presented during locomotion would adapt more strongly than responses to odors presented while the fly was still, potentially reflecting a novel form of behavior-dependent response adaptation. A remaining open question is whether smaller and noisier fluctuations in DANs that track other variables like heading would have an appreciable effect on KC-to-MBON transmission, and if not, what role they might play in behavior. More generally, our results beg the questions of to what extent the continuous DAN dynamics transform into continuous

modulation of KC-MBON synaptic strengths, how such a continuously modulated circuit might be stabilized, what effects these have on external and internal state representations in the MBONs, and how these shape behavior. Designing experiments to disentangle these questions will be an exciting future line of inquiry.

2.4.2 Multiplexed dopamine signals for memory and learning

Could dopamine, canonically associated with reinforcement learning (RL) in mammalian systems [31, 33], play a similar computational role in the fruit fly mushroom body? We first note that as the goal of RL is to learn a decision-making policy based on estimates of rewards that different stimulus-response pairings are expected to produce, a biological implementation of RL should require convergent representations of relevant sensory information, relevant motor information, and information about the current reward or its difference from the expected reward [33, 60]. The striatum, for instance, one of the chief hypothesized locations of RL in mammalian nervous systems [60], receives sensory and motor signals from cortex [61] and dopaminergic inputs from VTA and the substantia nigra pars compacta [62, 61]. In a way the mushroom body exhibits a similar convergence of inputs necessary for RL. KCs represent odor identity [50] and the DANs appear to provide information about additional sensorimotor, as well as rewarding or punishing events [1]. The key difference, however, is that whereas in striatum all sensorimotor information is thought to originate from cortex, with dopaminergic inputs signalling a more global reward prediction error signal [60], our results indicate that in the mushroom body much sensorimotor information is absorbed into the DAN signals themselves. This suggests that if the MB is implementing a similar RL-like computation as has been hypothesized in mammalian dopamine systems, it may be doing it in a different way. Given the relatively small nervous system underlying fly behavior, it is perhaps not surprising that a small number of cells would carry multiple different decodable signals, as has been observed for instance, in the fruit fly courtship circuit [8], but it remains an open question whether the multiplexing we observe is sufficient for RL and if so, whether it provides any computational advantage. As we refine our experimental understanding of the MB, such a question could be addressed through computational modeling studies.

One feature of our findings and others' that might shed light on the DAN-based control of MB information processing is the limited but still multidimensional DAN activity patterns. Indeed, whereas dopamine has canonically been thought to provide a global scalar signal representing an animal's overall estimate of the difference between the received and expected reward [31, 33, 60, 52], in the DAN data

collected by Zolin and Cohn one observes differential encodings of various sensorimotor variables across different DANs. This in itself suggests the γ -lobe DANs likely serve an additional purpose beyond conveying reward information, but a few recent studies attempting to disentangle dopamine function in mammals may shed additional light on the possible utility of multi-dimensional DAN signalling. First, it has been found that there may actually exist substantial heterogeneity in mammalian dopamine neurons, with some neurons found to respond, for instance, to only reward or only punishment and others to both reward and punishment [63], the latter which might be analogous to our observation that in the higher air flow closed loop condition, G3 appeared to respond to both the onset and offset of an extended ACV pulse (Fig 4F,J). Additionally, midbrain dopamine neurons may convey different information and subserve different behavioral function depending on their downstream targets [64], lending further evidence of generic heterogeneity in dopaminergic systems. From a computational perspective one function of multidimensional responses may be to reflect the different dimensions of the reward stimulus itself. Indeed, reward is rarely accurately represented by a single number but instead often has multiple dimensions, such as its overall value, risk and uncertainty, and type of reward [65]. While [65] found that midbrain dopamine neurons in primates appeared to integrate these different *objective* reward dimensions into a single scalar-valued *subjective* representation of reward prediction error, the possibility remains that it might prove useful in other contexts to retain the separation of different reward dimensions at the level of dopaminergic activity. Takahashi, et al. [66], for example, recently found midbrain dopamine neurons that signaled not only a prediction error of a scalar reward level, but also signaled prediction errors relative to sensory features of the predicted reward that had no correlation with reward value. Drawing parallels between dopaminergic computations in the mushroom body and the mammalian midbrain and striatum within a common computational framework through future experimental and modeling studies will likely be key to gaining insight into what purpose multidimensional dopaminergic and reward representations could serve with regard to behaviorally relevant computations.

An outstanding question about our particular results is what role the MB γ lobe plays in short vs. long-term memory. Indeed, the γ lobe has traditionally been implicated in short-term memory, relative to the α/β and α'/β' lobes, which are generally thought to be required long-term and medium-term memory, respectively [53, 67, 68, 69]. The rapid plasticity results of Cohn et al. [1] that motivated this investigation additionally corroborate the γ lobe’s role in fast, short-term storage of information in the mushroom body that alters MBON odor representations. Although the goal of RL is to shape decision-

policies over the long-term, it also requires short-term memory in order to solve the temporal credit assignment problem, in which the animal must construct associations between its own sensorimotor activity and rewards that arrive after a delay [?]. To do so, the animal must keep a temporary record of its recent actions so that it can discern which ones were involved in eliciting the delayed reward. One of the core RL mechanisms for storing such a record is the *eligibility trace*, in which different states or actions are “tagged” so as to indicate their eligibility for reinforcement, should a reward arrive a short time later [70, 71, 54, 32, 72]. The immediate responses of DAN activity to sensorimotor state and the rapid plasticity the DANs are capable of inducing in the γ lobe indeed poise them as a potential mechanism for storing such an eligibility trace. For example, by our reasoning in the previous section the arrival of an odor during active locomotion could tag a set of KC-to-MBON synapses encoding that odor, such that if a reward were encountered shortly after, the γ lobe would contain information about the odor identity or identities that had recently been encountered, which could theoretically be used to update the fly’s decision-making policy. The cellular and network mechanisms that could transform such information into future decisions remain to be elucidated, but the γ lobe’s role in long-term memory could be addressed experimentally by varying the experiments performed by Zolin et al. for this analysis. In particular, one would allow flies to learn to pair neutral odors with reward or punishment over the course of multiple trials and examine how the KC and DAN representations of the initially neutral stimulus changed over learning. Conversely, recordings such as the ones described in this chapter might be similarly performed in the α/β and α'/β' lobes to identify whether these MB neuropils exhibited any capacity for short-term memory.

Given the extensive role of odor-tracking in the life of an insect, one might also consider whether the γ -lobe DAN computations elucidated here might be used to shape this crucial behavior. Indeed, it was recently shown that flies exhibit some evidence of short-term memory and information accumulation over the course of single odor-tracking flights [4], and it is possible that the γ lobe might partially subserve such short-term memories. Unknown is whether and how the short-term memory of the γ lobe might be integrated into potential spatial computations that could optimize odor-tracking [73, 74, 4]. Interestingly, one modeling study seeking to understand the rich spatial behaviors of ants found that mushroom body circuitry could be co-opted to encode navigational routes using sparse codes for rapid learning of many different visual images [75]. Furthermore, recent investigations into biologically inspired machine learning algorithms have shown that the canonical properties of MB computations might be usable for more general purpose tasks such as image classification and one-shot

learning [76]. Understanding how the different MB lobes and their DAN-based neuromodulation might be used for spatial and general purpose computations remains an exciting line of future inquiry.

2.4.3 On the relationship between spontaneous and evoked neural activity

Lastly, our characterization of DAN activity in the mushroom body introduces new evidence regarding our understanding of the relationship between internally generated, “spontaneous” neural activity, and that evoked by external sensory stimulation. Indeed, our observations of the positive G2-G3 and G4-G5 correlations and the negative G2-G4 correlations in the DANs during the asensory condition (Fig 2) strongly reflect the correlational structure of DAN activity seen when flies are presented sugar rewards or electric shock [1]. Further, when we examined fluctuations of G2 and G4 DANs riding on top of their modulations due to the fly’s walking patterns (not shown), we found that these fluctuations were negatively correlated as well, lending further evidence that spontaneous dynamics in the DANs reflect those driven by behaviorally salient stimuli. Further, the G4 and G5 DANs appear to either directly or indirectly inhibit the G2 DANs [1], suggesting one should likely see such G2-G4 opponency during internally generated network activity. These results are in line with several other results suggesting that over development or learning, intrinsic brain activity in the absence of an explicit stimulus comes to reflect patterns of activity evoked by sensorimotor activity [77, 78, 79], and that furthermore such intrinsic or spontaneous activity might play an active role in shaping how future sensory are represented by the brain [80, 81, 82]. Indeed, at first blush at least, such observations may reflect the process by which our nervous systems acquire information about the world through short- or long-term memory and use this information to reshape how we process future sensory stimulation, for instance to improve our processing efficiency. That such a relationship between evoked and spontaneous activity is present in dopaminergic networks in addition to cortical and hippocampal networks lends evidence to the possibility that this may be a universal feature of biological neural networks that places both neuromodulatory circuitry and the networks they modulate on common footing.

2.5 Additional Materials and Methods

Unless otherwise noted, all experiments were designed and conducted by Zolin et al., and I largely led the formal analytical work. My primary experimental contribution was designing the stimulation sequence for the open loop sensory paradigm (Fig 6).

2.5.1 Data Collection

To measure behavior flies were tethered below a two-photon microscope but otherwise unconstrained. A small patterned trackball was placed beneath the flies and movement data extracted at 60 Hz using the FicTrac recording system [83] and custom MATLAB scripts by Zolin et al. DAN synaptic activations in the γ lobe were measured using the calcium indicator GCaMP6s attached to the C-terminus of the synaptic protein synaptotagmin, which was selectively expressed in DAN synapses using a combination of the tyrosine hydroxylase and dopa-decarboxylase promoters [1]. Activations of the G2-G5 DANs were imaged in a single imaging plane at 10 frames per second, and univariate activations in each of the 4 compartments (or 8 when both hemispheres were recorded in the sensory setup) were extracted by spatial averaging over hand-drawn regions of interest corresponding to each of the 4 visually identifiable compartments. Final synaptic activations were then computed by dividing the GCaMP6s fluorescence in each ROI by a red labeling protein whose fluorescence did not alter in response to DAN activations, to control for possible total variations in luminance. All behavioral variables were binned into 100 ms windows and aligned with the DAN recordings prior to analysis. For the sensory setup the air tube position was driven either in response to the fly's own motion on the trackball (closed loop) or via a predetermined filtered white noise sequence (open loop). All trials used in these analyses were 300 seconds long.

2.5.2 Data normalization

Prior to analysis all DAN activations were either z-scored across time (Fig 1-6) or computed relative to the time-averaged DAN activity surrounding a single odor pulse (Figs 7-9).

2.5.3 Filter fitting

All linear filters were fit using standard linear regression in which the least squared error between the true and predicted output was minimized to identify best-fit filters.

2.5.4 Open loop sensory paradigm

For the open loop sensory paradigm 5 trials were used across 3 flies. Closed loop movement of the air stream occurred for epochs 0-60s, 120-180s, and 240-300s, with open loop movement used for epochs

60-120s and 180-240s. Open loop air stream positions were driven using an O-U process with a standard deviation of 45 degrees and a timescale extracted by fitting a two-sided decaying exponential function to the average autocorrelation function of the flies' angular velocity in the closed loop paradigm.

2.5.5 Nested model analysis The nested model analysis of Fig 9 used linear models, with the significance levels for the increase in variance explained as stepwise predictors were added computed using an F-test.

2.6 Supplementary Figures




Fig 3S1: Example DAN overlays with three behavioral variables. *A. Forward walking velocity. B. The fly's angular/rotational velocity. C. The fly's absolute (unsigned) angular/rotational velocity. All variables are in arbitrary units and were rescaled for visual alignment.*




Fig 6S1: Inter-DAN cross covariance functions in the closed loop sensory condition. A-P. As in Fig 2 but for the closed loop trials, where each DAN signal used (G2A - G5A) was now the average of the corresponding DAN values across hemispheres. Q-T. Cross covariance functions between corresponding DANs in different hemispheres.




Fig 8S1: Autocorrelation functions of G2-4 and absolute heading. As in Fig 8A-D except showing correlation functions of three DAN values and one behavioral variable with themselves.

References

- [1] Raphael Cohn, Ianessa Morantte, and Vanessa Ruta. Coordinated and compartmentalized neuromodulation shapes sensory processing in drosophila. *Cell*, 163(7):1742–1755, 2015.
- [2] Sarah Hammond and Michael O’Shea. Escape flight initiation in the fly. *Journal of Comparative Physiology A*, 193(4):471–476, 2007.
- [3] Floris van Breugel and Michael H Dickinson. Plume-tracking behavior of flying drosophila emerges from a set of distinct sensory-motor reflexes. *Current Biology*, 24(3):274–286, 2014.
- [4] Rich Pang, Floris van Breugel, Michael Dickinson, Jeffrey A Riffell, and Adrienne Fairhall. History dependence in insect flight decisions during odor tracking. *PLoS computational biology*, 14(2):e1005969, 2018.
- [5] Therese A Markow and Stephen J Hanson. Multivariate analysis of drosophila courtship. *Proceedings of the National Academy of Sciences*, 78(1):430–434, 1981.
- [6] Arthur W Ewing. Functional aspects of drosophila courtship. *Biological Reviews*, 58(2):275–292.
- [7] Y Yu Jai, Makoto I Kanai, Ebru Demir, Gregory SXE Jefferis, and Barry J Dickson. Cellular organization of the neural circuit that drives drosophila courtship behavior. *Current Biology*, 20(18):1602–1614, 2010.
- [8] Philip Coen, Marjorie Xie, Jan Clemens, and Mala Murthy. Sensorimotor transformations underlying variability in song intensity during drosophila courtship. *Neuron*, 89(3):629–644, 2016.
- [9] Ronald L Davis. Mushroom bodies and drosophila learning. *Neuron*, 11(1):1–14, 1993.

- [10] J Steven de Belle and Martin Heisenberg. Associative odor learning in drosophila abolished by chemical ablation of mushroom bodies. *Science*, 263(5147):692–695, 1994.
- [11] Alberto Pascual and Thomas Pr  at. Localization of long-term memory within the drosophila mushroom body. *Science*, 294(5544):1115–1117, 2001.
- [12] Eve Marder. Neuromodulation of neuronal circuits: back to the future. *Neuron*, 76(1):1–11, 2012.
- [13] Craig W Berridge and Barry D Waterhouse. The locus coeruleus–noradrenergic system: modulation of behavioral state and state-dependent cognitive processes. *Brain research reviews*, 42(1):33–84, 2003.
- [14] Jessica A Cardin and Marc F Schmidt. Noradrenergic inputs mediate state dependence of auditory responses in the avian song system. *Journal of Neuroscience*, 24(35):7745–7753, 2004.
- [15] Seung-Hee Lee and Yang Dan. Neuromodulation of brain states. *Neuron*, 76(1):209–222, 2012.
- [16] Eve Marder, Timothy O’Leary, and Sonal Shruti. Neuromodulation of circuits with variable parameters: single neurons and small circuits reveal principles of state-dependent and robust neuromodulation. *Annual review of neuroscience*, 37:329–346, 2014.
- [17] Sebastian Reinig, Wolfgang Driever, and Aristides B Arrenberg. The descending diencephalic dopamine system is tuned to sensory stimuli. *Current Biology*, 27(3):318–333, 2017.
- [18] Yves Kremer, J  r  me Flakowski, Cl  ment Rohner, and Christian L  scher. VTA dopamine neurons multiplex external with internal representations of goal-directed action. *bioRxiv*, page 408062, 2018.
- [19] Roy A Wise. Dopamine, learning and motivation. *Nature reviews neuroscience*, 5(6):483, 2004.
- [20] Vikram Gadagkar, Pavel A Puzerey, Ruidong Chen, Eliza Baird-Daniel, Alexander R Farhang, and Jesse H Goldberg. Dopamine neurons encode performance error in singing birds. *Science*, 354(6317):1278–1282, 2016.
- [21] Kyung-An Han, Neil S Millar, Michael S Grotewiel, and Ronald L Davis. Damb, a novel dopamine receptor expressed specifically in drosophila mushroom bodies. *Neuron*, 16(6):1127–1135, 1996.
- [22] Susan E Fahrbach. Structure of the mushroom bodies of the insect brain. *Annu. Rev. Entomol.*, 51:209–232, 2006.

- [23] Yoshinori Aso, Andrea Herb, Maite Ogueta, Igor Siwanowicz, Thomas Templier, Anja B Friedrich, Kei Ito, Henrike Scholz, and Hiromu Tanimoto. Three dopamine pathways induce aversive odor memories with different stability. *PLoS genetics*, 8(7):e1002768, 2012.
- [24] Joshua D Berke and Steven E Hyman. Addiction, dopamine, and the molecular mechanisms of memory. *Neuron*, 25(3):515–532, 2000.
- [25] EP Noble. Addiction and its reward process through polymorphisms of the d2 dopamine receptor gene: a review. *European Psychiatry*, 15(2):79–89, 2000.
- [26] Nora D Volkow, Joanna S Fowler, Gene-Jack Wang, and James M Swanson. Dopamine in drug abuse and addiction: results from imaging studies and treatment implications. *Molecular psychiatry*, 9(6):557, 2004.
- [27] David J Nutt, Anne Lingford-Hughes, David Erritzoe, and Paul RA Stokes. The dopamine theory of addiction: 40 years of highs and lows. *Nature Reviews Neuroscience*, 16(5):305, 2015.
- [28] Jonathan D Cohen and David Servan-Schreiber. Context, cortex, and dopamine: a connectionist approach to behavior and biology in schizophrenia. *Psychological review*, 99(1):45, 1992.
- [29] Oliver D Howes and Shitij Kapur. The dopamine hypothesis of schizophrenia: version iii—the final common pathway. *Schizophrenia bulletin*, 35(3):549–562, 2009.
- [30] Anthony A Grace. Dysregulation of the dopamine system in the pathophysiology of schizophrenia and depression. *Nature Reviews Neuroscience*, 17(8):524, 2016.
- [31] Wolfram Schultz, Peter Dayan, and P Read Montague. A neural substrate of prediction and reward. *Science*, 275(5306):1593–1599, 1997.
- [32] Richard S Sutton and Andrew G Barto. *Introduction to reinforcement learning*, volume 135. MIT press Cambridge, 1998.
- [33] Peter Dayan. Dopamine, reinforcement learning, and addiction. *Pharmacopsychiatry*, 42(S 01):S56–S65, 2009.
- [34] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.

- [35] A David Redish. Addiction as a computational process gone awry. *Science*, 306(5703):1944–1947, 2004.
- [36] Scott Waddell. Reinforcement signalling in drosophila; dopamine does it all after all. *Current opinion in neurobiology*, 23(3):324–329, 2013.
- [37] Kanta Terao and Makoto Mizunami. Roles of dopamine neurons in mediating the prediction error in aversive learning in insects. *Scientific reports*, 7(1):14694, 2017.
- [38] Stephen J Kish, Kathleen Shannak, and Oleh Hornykiewicz. Uneven pattern of dopamine loss in the striatum of patients with idiopathic parkinson’s disease. *New England Journal of Medicine*, 318(14):876–880, 1988.
- [39] P Damier, EC Hirsch, Y Agid, and AM Graybiel. The substantia nigra of the human brain: Ii. patterns of loss of dopamine-containing neurons in parkinson’s disease. *Brain*, 122(8):1437–1448, 1999.
- [40] Julie Lotharius and Patrik Brundin. Pathogenesis of parkinson’s disease: dopamine, vesicles and α -synuclein. *Nature Reviews Neuroscience*, 3(12):932, 2002.
- [41] Tamara Pringsheim, Nathalie Jette, Alexandra Frolkis, and Thomas DL Steeves. The prevalence of parkinson’s disease: A systematic review and meta-analysis. *Movement disorders*, 29(13):1583–1590, 2014.
- [42] Jon C Horvitz, Tripp Stewart, and Barry L Jacobs. Burst activity of ventral tegmental dopamine neurons is elicited by sensory stimuli in the awake cat. *Brain research*, 759(2):251–258, 1997.
- [43] Arthur S Freeman and Benjamin S Bunney. Activity of a9 and a10 dopaminergic neurons in unrestrained rats: further characterization and effects of apomorphine and cholecystokinin. *Brain research*, 405(1):46–55, 1987.
- [44] J Boeckh, P Distler, KD Ernst, M Hösl, and D Malun. Olfactory bulb and antennal lobe. In *Chemosensory information processing*, pages 201–227. Springer, 1990.
- [45] Leslie B Vosshall, Allan M Wong, and Richard Axel. An olfactory sensory map in the fly brain. *Cell*, 102(2):147–159, 2000.

- [46] Yoshinori Aso, Daisuke Hattori, Yang Yu, Rebecca M Johnston, Nirmala A Iyer, Teri-TB Ngo, Heather Dionne, LF Abbott, Richard Axel, Hiromu Tanimoto, et al. The neuronal architecture of the mushroom body provides a logic for associative learning. *Elife*, 3:e04577, 2014.
- [47] Tzumin Lee, Arthur Lee, and Liqun Luo. Development of the drosophila mushroom bodies: sequential generation of three distinct types of neurons from a neuroblast. *Development*, 126(18):4065–4076, 1999.
- [48] David-Benjamin G Akalal, Curtis F Wilson, Lin Zong, Nobuaki K Tanaka, Kei Ito, and Ronald L Davis. Roles for drosophila mushroom body neurons in olfactory learning and memory. *Learning & memory*, 13(5):659–668, 2006.
- [49] Zhengmei Mao and Ronald L Davis. Eight different types of dopaminergic neurons innervate the drosophila mushroom body neuropil: anatomical and physiological heterogeneity. *Frontiers in neural circuits*, 3:5, 2009.
- [50] Iori Ito, Rose Chik-ying Ong, Baranidharan Raman, and Mark Stopfer. Sparse odor representation and olfactory learning. *Nature neuroscience*, 11(10):1177, 2008.
- [51] Colleen Kirkhart and Kristin Scott. Gustatory learning and processing in the drosophila mushroom bodies. *Journal of Neuroscience*, 35(15):5950–5958, 2015.
- [52] Neir Eshel, Ju Tian, Michael Bukwich, and Naoshige Uchida. Dopamine neurons share common response function for reward prediction error. *Nature neuroscience*, 19(3):479, 2016.
- [53] T Zars, M Fischer, R Schulz, and M Heisenberg. Localization of a short-term memory in drosophila. *Science*, 288(5466):672–675, 2000.
- [54] Wolfram Schultz. Predictive reward signal of dopamine neurons. *Journal of neurophysiology*, 80(1):1–27, 1998.
- [55] Daniel T Gillespie. Exact numerical simulation of the ornstein-uhlenbeck process and its integral. *Physical review E*, 54(2):2084, 1996.
- [56] Ying-Hui Liu and Xiao-Jing Wang. Spike-frequency adaptation of a generalized leaky integrate-and-fire model neuron. *Journal of computational neuroscience*, 10(1):25–45, 2001.
- [57] Jan Benda and Andreas VM Herz. A universal model for spike-frequency adaptation. *Neural computation*, 15(11):2523–2564, 2003.

- [58] Barry Wark, Brian Nils Lundstrom, and Adrienne Fairhall. Sensory adaptation. *Current opinion in neurobiology*, 17(4):423–429, 2007.
- [59] Katrien Segaert, Kirsten Weber, Floris P de Lange, Karl Magnus Petersson, and Peter Hagoort. The suppression of repetition enhancement: a review of fmri studies. *Neuropsychologia*, 51(1):59–66, 2013.
- [60] Paul W Glimcher. Understanding dopamine and reinforcement learning: the dopamine reward prediction error hypothesis. *Proceedings of the National Academy of Sciences*, 108(Supplement 3):15647–15654, 2011.
- [61] Barbara J Hunnicutt, Bart C Jongbloets, William T Birdsong, Katrina J Gertz, Haining Zhong, and Tianyi Mao. A comprehensive excitatory input map of the striatum reveals novel functional organization. *Elife*, 5:e19103, 2016.
- [62] Martin Darvas and Richard D Palmiter. Restriction of dopamine signaling to the dorsolateral striatum is sufficient for many cognitive behaviors. *Proceedings of the National Academy of Sciences*, 106(34):14664–14669, 2009.
- [63] Masayuki Matsumoto and Okihide Hikosaka. Two types of dopamine neuron distinctly convey positive and negative motivational signals. *Nature*, 459(7248):837, 2009.
- [64] Stephan Lammel, Byung Kook Lim, and Robert C Malenka. Reward and aversion in a heterogeneous midbrain dopamine system. *Neuropharmacology*, 76:351–359, 2014.
- [65] Armin Lak, William R Stauffer, and Wolfram Schultz. Dopamine prediction error responses integrate subjective value from different reward dimensions. *Proceedings of the National Academy of Sciences*, 111(6):2343–2348, 2014.
- [66] Yuji K Takahashi, Hannah M Batchelor, Bing Liu, Akash Khanna, Marisela Morales, and Geoffrey Schoenbaum. Dopamine neurons respond to errors in the prediction of sensory features of expected rewards. *Neuron*, 95(6):1395–1405, 2017.
- [67] Sean E McGuire, Phuong T Le, and Ronald L Davis. The role of drosophila mushroom body signaling in olfactory memory. *Science*, 293(5533):1330–1333, 2001.

- [68] Allison L Blum, Wanhe Li, Mike Cressy, and Josh Dubnau. Short-and long-term memory in drosophila require camp signaling in distinct neuron types. *Current Biology*, 19(16):1341–1350, 2009.
- [69] Séverine Trannoy, Christelle Redt-Clouet, Jean-Maurice Dura, and Thomas Preat. Parallel processing of appetitive short-and long-term memories in drosophila. *Current Biology*, 21(19):1647–1653, 2011.
- [70] A Harry Klopff. Brain function and adaptive systems: a heterostatic theory. Technical report, AIR FORCE CAMBRIDGE RESEARCH LABS HANSCOM AFB MA, 1972.
- [71] Satinder P Singh and Richard S Sutton. Reinforcement learning with replacing eligibility traces. *Machine learning*, 22(1-3):123–158, 1996.
- [72] Wulfram Gerstner, Marco Lehmann, Vasiliki Liakoni, Dane Corneil, and Johanni Brea. Eligibility traces and plasticity on behavioral time scales: Experimental support of neohebbian three-factor learning rules. *arXiv preprint arXiv:1801.05219*, 2018.
- [73] Massimo Vergassola, Emmanuel Villermanx, and Boris I Shraiman. ‘infotaxis’ as a strategy for searching without gradients. *Nature*, 445(7126):406, 2007.
- [74] Jean-Baptiste Masson. Olfactory searches with limited space perception. *Proceedings of the National Academy of Sciences*, 110(28):11261–11266, 2013.
- [75] Paul Ardin, Fei Peng, Michael Mangan, Konstantinos Lagogiannis, and Barbara Webb. Using an insect mushroom body circuit to encode route memory in complex natural environments. *PLoS computational biology*, 12(2):e1004683, 2016.
- [76] Charles B Delahunt, Jeffrey A Riffell, and J Nathan Kutz. Biological mechanisms for learning: A computational model of olfactory learning in the manduca sexta moth, with applications to neural nets. *arXiv preprint arXiv:1802.02678*, 2018.
- [77] Pietro Berkes, Gergő Orbán, Máté Lengyel, and József Fiser. Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science*, 331(6013):83–87, 2011.
- [78] Margaret F Carr, Shantanu P Jadhav, and Loren M Frank. Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval. *Nature neuroscience*, 14(2):147, 2011.

- [79] Shengjin Xu, Wanchen Jiang, Mu-ming Poo, and Yang Dan. Activity recall in a visual cortical ensemble. *Nature neuroscience*, 15(3):449, 2012.
- [80] Artur Luczak, Peter Barthó, and Kenneth D Harris. Spontaneous events outline the realm of possible sensory responses in neocortical populations. *Neuron*, 62(3):413–425, 2009.
- [81] Artur Luczak, Peter Bartho, and Kenneth D Harris. Gating of sensory input by spontaneous cortical activity. *Journal of Neuroscience*, 33(4):1684–1695, 2013.
- [82] Luis Carrillo-Reid, Jae-eun Kang Miller, Jordan P Hamm, Jesse Jackson, and Rafael Yuste. Endogenous sequential cortical activity evoked by visual stimuli. *Journal of Neuroscience*, 35(23):8813–8828, 2015.
- [83] Richard JD Moore, Gavin J Taylor, Angelique C Paulk, Thomas Pearson, Bruno van Swinderen, and Mandyam V Srinivasan. Fictrac: a visual method for tracking spherical motion and generating fictive animal paths. *Journal of neuroscience methods*, 225:106–119, 2014.

3. Fast and flexible sequence induction in spiking neural networks via rapid excitability changes

The following work is additionally available as a bioRxiv preprint via the following citation:

Pang, Rich, and Adrienne L. Fairhall. "Fast and Flexible Sequence Induction In Spiking Neural Networks Via Rapid Excitability Changes." *bioRxiv* (2018): 494310. (DOI: 10.1101/494310)

3.1 Abstract

Cognitive flexibility, the adaptation of mental processing to changes in task demands, is thought to depend on biological neural networks' ability to rapidly modulate the dynamics governing how they process information. While extensive work has elucidated how network dynamics can be reshaped by slowly occurring structural changes, e.g. the gradual modification of recurrent synaptic patterns, much less is known about how dynamics might be reconfigured over faster timescales of seconds. One compelling example of rapid and selective modulation of network dynamics potentially involved in cognitive flexibility is observed in rodent hippocampus, where short bouts of exploratory behavior cause new activity sequences to preferentially "replay" during subsequent awake rest periods without continued sensory input. Fast mechanisms for selectively biasing sequential activity through networks, however, remain unknown. Using a spiking neural network model, we asked whether a simplified version of sequence replay could arise from three biophysically plausible components: recurrent, spatially organized connectivity; homogeneous, stochastic "gating" inputs; and rapid, activity-dependent scaling of gating input strengths, based on a phenomenon known as long-term potentiation of intrinsic excitability (LTP-IE). Indeed, these enabled both forward and reverse replay of flexible sequences reflecting recent behavior, despite unchanged recurrent weights. Specifically, activation-triggered LTP-IE "tags" neurons in the recurrent network by increasing their spiking probability when gating input is applied, and the sequential ordering of spikes is reconstructed by the existing recurrent connectivity. In a proof-of-concept demonstration, we also show how LTP-IE-based sequences can implement temporary stimulus-response mappings in a straightforward manner. These results elucidate a simple yet previously unexplored combination of biological mechanisms that converge in hippocampus and suffice for fast and flexible reconfiguration of sequential network dynamics, suggesting their potential role in cognitive flexibility over rapid timescales.

3.2 Introduction

We can rapidly and flexibly adapt how we process incoming information from the environment, a mental faculty known as cognitive flexibility. For example, after being instructed to raise our left hand when one word is heard and our right hand when another is heard, we perform the task with little error. How brain networks quickly induce novel stimulus-response mappings such as this into their underlying neural dynamics, however, remains mysterious. In particular, while extensive prior work has elucidated how stimulus-response mappings might be implemented biologically, the mechanisms for inducing these mappings typically require slow, gradual modifications to network structure e.g. by incrementally training connection weights to minimize errors between correct and predicted responses ([1, 2, 3, 4, 5]) or by allowing local plasticity to reshape network dynamics in response to internal activity [6, 7, 8, 9, 10, 11, 12]. Biologically observed spike-timing-dependent plasticity (STDP) mechanisms, however, typically increase post-synaptic potentials (PSPs) by at most a few percent [13, 14, 15, 16, 17]. Furthermore, due to the precise timing requirements of canonical plasticity mechanisms (e.g. in STDP spike pairs must occur in the correct order within tens of milliseconds) and low firing rates of cortex and hippocampus (typically less than a few tens of Hz [18, 19, 20]), STDP-triggering spike patterns may occur relatively rarely, especially given the asynchronous nature of cortical firing patterns in awake animals [21]. As a result, computational models for shaping dynamics that modify network structures via synaptic plasticity typically rely on at least dozens of learning trials over extended time periods [6, 22, 10], challenging their suitability for rapid reconfiguration of network dynamics. Computationally, modifying connections to change network function might also interfere with long-term memories or computations already stored in the network’s existing connectivity patterns, leading, for instance, to catastrophic forgetting [23]. Consequently, it is unclear (1) how biologically observed plasticity mechanisms could reshape network function over the timescales of seconds required for rapid cognitive flexibility, and (2) how such restructuring of synaptic connectivity could occur over the short term without degrading existing long-term memories.

One feature of neural dynamics potentially reflecting processes of cognitive flexibility is stereotyped sequential firing patterns [24, 25, 26, 27, 28, 29, 30]. Functionally, firing sequences are thought to be involved in various cognitive processes, from short-term memory [28, 29] to decision-making [30]. More generally, one can imagine sequential activity as reflecting information propagating from one subnetwork to another, e.g. stimulus S evoking a cascade of activity that eventually triggers motor output M . A compelling empirical example of memory-related firing sequences that arise in a neural network almost immediately after a sensorimotor event, apparently without requiring repeated

experience or long-term learning, occurs in so-called awake hippocampal replay. Here, sequences of spikes in hippocampal regions CA1 and CA3 originally evoked by a rodent traversing its environment along a specific trajectory subsequently replay when the rodent pauses to rest [31, 28, 32, 33]. Such replay events occur at compressed timescales relative to the original trajectory and often in reverse order. Replay has also been observed in primate cortical area V4, where firing sequences evoked by a short movie were immediately reactivated by a cue indicating the movie was about to start again, but without showing the movie [34]. The functional role of replay has been implicated in memory, planning, and learning [35, 33, 34, 36, 37], but little is known about the mechanisms enabling the underlying sequential activity to become induced in the network dynamics in the first place. Elucidating biological mechanisms for rapidly inducing sequential firing patterns in network dynamics may not only illuminate the processes enabling replay but may also shed light on principles for fast and flexible reconfiguration of computations and information flow in neural networks more generally.

An intriguing fast-acting cellular mechanism whose role in shaping network dynamics has not been investigated is the rapid, outsize, and activity-dependent modulation of cortical inputs onto pyramidal cells (PCs) in hippocampal region CA3 [38, 39, 40]. Specifically, following a 1-2 s train of 20 action potentials in a given CA3 PC, excitatory postsynaptic potentials (EPSPs) from medial entorhinal cortex (MEC) more than doubled in magnitude within eight seconds (the first time point in the experiment). Thought to arise through inactivation of K^+ -channels colocalized with MEC projections onto CA3 PC dendrites and deemed “long-term potentiation of intrinsic excitability” (LTP-IE), potentiation occurred regardless of whether the CA3 PC spikes were evoked via current injection or by upstream physiological inputs, indicating its heterosynaptic nature, since only MEC EPSPs, and not others, exhibited potentiation [38, 39]. Notably, the 10-20 Hz spike rate required for potentiation matches the range of in vivo spike rates in hippocampal “place cells” when rodents pass through specific locations [41, 42], suggesting LTP-IE may occur in natural contexts. Furthermore, EPSP potentiation persisted throughout the multi-minute course of the experiment [39], suggesting that in addition to fast onset, the modulation could extend significantly into the future. Although inducing sequential activation patterns in neural networks is typically associated with homosynaptic plasticity (e.g. in STDP a postsynaptic following a presynaptic spike strengthens the activated synapse, thereby incrementally increasing the probability of subsequent presynaptic spikes triggering postsynaptic spikes), the rapidity, strength, and duration of this heterosynaptic potentiation mechanism suggest it might significantly modulate network dynamics in natural conditions, warranting further investigation within

the context of neural sequences. Moreover, while this mechanism has so far been observed in hippocampus only, an intriguing possibility is that functionally similar mechanisms exist in cortex also, enabling rapid effective changes in excitability of recently active cells, with potentially similar computational consequences (see Discussion).

To explore this idea we first develop a computational model for the effect of LTP-IE on EPSPs from upstream inputs as a function of their spiking responses to physiological inputs. Next, we demonstrate how LTP-IE combined with recurrent PC connectivity in a spiking network can yield spike sequences reflecting recent sensorimotor sequences that replay in both forward and reverse and which are gated by an upstream gating signal. We subsequently identify parameter regimes allowing and prohibiting LTP-IE-based sequence propagation and examine the effect of specific parameters on replay speed, spontaneous activity, and information transmission. Finally, we show how LTP-IE-based sequences can be used to induce temporary stimulus-response mappings in an otherwise untrained recurrent network. We discuss implications for cognitive flexibility and rapid, non-stereotyped memory storage that does not require modification of recurrent network weights.

3.3 Results

To investigate its consequences on neuronal spiking dynamics we implemented LTP-IE in a network of leaky integrate-and-fire (LIF) neurons [43]. Overall neurons received three types of inputs: sensorimotor inputs S , carrying tuned “external” information; recurrent input R from other cells in the network (developed further, Fig 2); and “gating” inputs G , which are assumed to be random, with a homogeneous rate across the network, but independent to each cell. We first provide an overview and intuition of the model that follows.

First, we demonstrate our implementation of LTP-IE in spiking neurons: LTP-IE in a given neuron is activated by that neuron’s firing within a physiologically plausible range (20 Hz) for 1-2 seconds. LTP-IE specifically augments EPSPs from gating (G) inputs (analogous to MEC inputs); consequently, recently activated cells, “tagged” by LTP-IE, will show stronger positive voltage deflections in response to G inputs. Thus, when continuous G inputs are present, recently active, LTP-IE-tagged cells will, on average, sit closer to spiking threshold. Unless otherwise specified, in all simulations of reactivation dynamics that follow we assume the presence of G inputs, as LTP-IE cannot affect membrane voltages without them.

Second, in analogy to hippocampal “place cells” [41, 42] we consider the case where different

neurons maximally activate when a simulated animal is in different positions in the environment. Consequently, when the simulated animal navigates along a trajectory, LTP-IE tags the cells with place fields along the trajectory, thus storing the trajectory memory as a set of positions, without order or timing information. When G inputs are turned on post-navigation (representing an awake resting state), these recently active LTP-IE-tagged cells sit closer to spiking threshold than cells that were not activated by the trajectory.

Finally, inspired by excitatory recurrence in CA3 [44] we assume recurrent excitatory connections in the network, which allows activity to propagate through the network over time. We assume that cell pairs with more similar tuning (nearby “place fields”) have stronger connections [44, 45, 46]. Thus, activity from cell A is more likely to propagate to cell B if (1) cell B has been LTP-IE-tagged after being activated by the trajectory, and (2) cell B has a nearby place field to cell A. Consequently, given the right model parameters, activity that begins at one point in the network (e.g. in cells tuned to the animal’s resting position) should propagate primarily among LTP-IE-tagged cells, in an order potentially reflecting the original trajectory. We discuss additional model components we excluded for the sake parsimony, such as inhibition, structured G inputs, additional tuning parameters, and LTP-IE extinction, in the Discussion.

Note that our model accords two meanings of “excitability” to the LTP-IE acronym. As coined by [38, 39], LTP-IE’s “excitability” originally refers to the augmented EPSPs of G inputs (MEC inputs, in [38, 39]) arriving to distal dendrites. Here, however, “excitability” also refers to the increased voltage of LTP-IE-tagged cells, which makes them more likely to spike in response to inputs.

3.3.1 LTP-IE increases membrane voltages and spike rates under random gating inputs

We first model LTP-IE mathematically and demonstrate its effect on neuronal membrane voltages and spiking. To model LTP-IE, when a naive cell spikes at a sufficiently high rate (10 Hz) over 1 second, we increase the strength of that cell’s G inputs by a factor γ , in accordance with its spike rate, such that future G inputs yield augmented EPSPs (Fig 1A, B). Consequently, cells that have recently emitted several spikes in quick succession end up with high LTP-IE levels, which causes them to sit at higher average voltages given a steady random stream of G inputs (Fig 1C-E); additionally, cells with higher LTP-IE exhibit increased variability in their membrane voltages (Fig 1D, E).

Due to their increased membrane voltage, cells with high LTP-IE also exhibit higher spiking probability, both spontaneously (Fig 1S1) and in response to depolarizing current inputs (Fig 1F,G), although




Fig 1: Mechanism and consequences of LTP-IE. *A.* Demonstration using leaky integrate-and-fire (LIF) neuron model of fast activity-driven LTP-IE, which doubles membrane voltage responses to gate inputs as described in [38, 39]. A spike in an upstream gate neuron (G) first elicits a small EPSP in the pyramidal cell (PC); a 1-second spike train (dots) at approximately 20 Hz is evoked by strong stimulation of sensory inputs S; when G spikes again, the EPSP has doubled in size. “RCR” refers to recurrent inputs from other PCs (not used in this figure). *B.* Shifted logistic function for LTP-IE strength (effective weight scaling factor) σ vs. PC firing rate over 1 second. *C.* Example voltage responses (V_m) of PCs with different receiving statistically identical stochastic gating input spikes. *D.* Distribution of PC membrane voltage for σ values shown in C. *E.* Mean (thick) and standard deviation (shading) of V_m as a function of σ for three gate firing rates. Black: $r_G = 75$ Hz; dark red: $r_G = 125$ Hz; red: $r_G = 175$ Hz. *F.* Example differential sensitivities of PC spike responses to injected current (blue) for two different σ . Dashed lines show leak and spike threshold potential; dots indicate spikes (which only occur for the $\sigma = 2$ case [cyan]). *G.* Difference between current-evoked spike probability and spontaneous spike probability as a function of σ for four initial gate input weights (unitless). Color code, in order of increasing lightness: $w^{PG,G} = .004, .008, .012, .016$.

substantial variability remains, and even with high LTP-IE, current-evoked spiking is not assured (Fig 1F,G). Nonetheless, for moderate values of the synaptic strength of initial G inputs, LTP-IE substantially increases the probability of transforming a depolarizing input into an output spike from near zero to approximately 0.5 (Fig 1G). (For overly weak initial G inputs, LTP-IE does not substantially facilitate spiking; for overly strong initial G inputs, spontaneous spike rates are already high and increased spiking leads to more frequent voltage resetting [Fig 1G, 1S1].) Thus, given moderate initial G input strengths, when a cell undergoes LTP-IE, its chance of spiking in response to future inputs (i.e. its excitability) substantially increases, although variability remains. Since LTP-IE-triggering spike rates are only around 10 Hz, this suggests LTP-IE may play an active role in modulating firing properties of recently active cells in physiological conditions.

3.3.2 Spike sequences propagate along LTP-IE-defined paths through a network

We next asked how LTP-IE would shape activity in a recurrent spiking network of 1000 excitatory neurons. Individual neurons were inspired by hippocampal place cells [47, 41], with firing rates tuned to specific positions (“place fields”) within an environment (Fig 2A-C). We furthermore assumed preferential connections between neurons with nearby place fields (Fig 2D and Methods). Since our primary goal was to evaluate whether LTP-IE alone was sufficient to induce activity sequences we did not include inhibitory neurons, but potential roles for inhibition are described in the Discussion.

To investigate network activity depending on recent sensorimotor sequences, we considered the scenario in which a rodent has just completed a short trajectory through its environment, in analogy with experimental setups used to measure neuronal replay in hippocampus [31, 28, 33]. Since our goal was to understand sequential reactivation during the awake, quiescent period following the trajectory, we only used the trajectory geometry to predict the LTP-IE levels one would expect following its termination. To do so, we computed each neuron’s maximum expected spike rate during the trajectory as a function of the distance from its place field to the nearest point on the trajectory (Fig 2C). We then passed the result through the soft-thresholding LTP-IE activation function (Fig 1B) to compute the final expected LTP-IE level (Fig 2C). This allowed us to model the expected LTP-IE profile over the network of neurons as a function of the recent trajectory (Fig 2E). As per our design, the LTP-IE profile stores which locations were covered by the original trajectory but bears no explicit information about its speed or direction.

Can replay-like sequences that recapitulate the original trajectory structure emerge from the LTP-




Fig 2: Demonstration of history-dependent sequence replay via LTP-IE. *A. Network architecture used in simulations. B. Squared exponential position tuning for an example neuron (20 Hz max. firing rate, 0.15 m length scale). C. Resultant LTP-IE as a function of distance between a cell's maximum tuning and the closest point on the trajectory through the environment (red), computed as a sigmoidal function (Fig 1B) of position-dependent firing rate (black). D. Recurrent weights as squared exponential function of distance between peak tuning positions of two cells. E. LTP-IE profile induced by example trajectory (inset). Cells are positioned according to peak tuning. F. Cells activated during a 100 ms replay epoch, colored by the ordering of the first spikes emitted during replay. Black cells did not activate during replay. G. As in F, but for a different trajectory (inset). H. Raster plot depicting spike times for select cells across several replay epochs. Current-pulse triggers to cells in upper-left or lower-right portion of trajectory are shown by triangles and inverted triangles, respectively. I-K. Zoom in on two triggered replay events (I, J) and one spontaneous event (K). L. Trajectories decoded from replay spikes when LTP-IE was induced by trajectories in F or G.*

IE profile stored in the network? Indeed, when we triggered replay with a depolarizing current pulse into cells at the lower-right corner of the trajectory, an approximately 100 ms sequence of activity cascaded along the path through the network aligned with the LTP-IE profile (Fig 2F). Due to the refractory period of 8 ms (see Methods), activity propagated in one direction without reversing. When an alternative trajectory was used to induce the LTP-IE profile, triggered replay recapitulated that trajectory instead (Fig 2G), indicating the replay of a specific trajectory was not an a priori consequence of the recurrent connectivity; indeed, this is to be expected from the spatial uniformity of the recurrence. Additionally, the replay sequence was able to turn corners, suggesting robustness to multidimensional trajectory geometries. Moreover, replay could occur in either direction along the LTP-IE profile depending on which end of the trajectory received the initial trigger (Fig 2H-L), in line with empirical observations in which replay often occurs in the reverse direction as the original trajectory [31, 28, 33]. Replay events also arose spontaneously, in either direction (Fig 2H, K), and sometimes—due to the lack of inhibition—arose from the middle of the trajectory, propagating outwards in both directions (not shown). Thus, despite sizable variability in membrane voltages arising from the random gating inputs G , LTP-IE was able to induce activity sequences into the network that recapitulated recent sensorimotor experiences.

3.3.3 Dependence of LTP-IE-based sequence propagation on network parameters

What conditions must hold for LTP-IE to induce successful sequences through the network? To address this we explored how the network activity evoked by a trigger at the lower-right corner of the Z-shaped LTP-IE profile from Fig 2E changed as we varied different pairs of network parameters. For each parameter combination we classified triggered activity into one of four classes: in fadeout (Fig 3A), activity fails to propagate along the full LTP-IE profile; in blowup (Fig 3B), activity propagates through whole network instead of being constrained to the LTP-IE profile; in disordered replay (Fig 3C), activity remains constrained to the LTP-IE profile but does not propagate strictly from one end to the other, as may occur if a spontaneous event arises at the same time as a triggered one. If neither fadeout, blowup, nor disordered replay occurred, we classified the response as successful replay. In general, successful replay was sensitive to certain parameter variations but robust to others (Fig 3D-F). For instance, whereas a fine balance was required between the recurrent connection strength and the recurrent connectivity length scale (Fig 3D), replay was more robust to variations in the gating input rate r_G and the maximum LTP-IE strength $\max$ (Fig 3E). (We note, however, that sensitivity to

recurrent connection strength and length scale could be improved by increasing the density of cells in the network [Fig 3S1].) In general, excessive decreases in any of these four parameters led to fadeout whereas excessive increases led to blowup, with disordered replay occurring when either max or rG was too large while other parameters retained moderate values (Fig 3E-F).

Within parameter regimes enabling successful replay, precise parameter values modulated specific activity dynamics. For instance, spontaneous spiking (typically dominated by spontaneous replay events) increased with the recurrent connection strength, maximum LTP-IE strength, and gating input rate (Fig 3G-I). Furthermore, the “replay speed” (the speed of movement along the trajectory decoded during the replay event) showed the greatest positive relationship with the recurrence length scale, but also exhibited some dependence on other parameters (Fig 3J-K). This result corroborates the hypothesis that the speed of replay events emerges from internal network structure, as opposed to the temporal structure of the animal’s original trajectory [28]. Note that empirical replay speeds have been measured near 8 m/s [28] whereas ours were in the range of 20-50 m/s; this range is largely determined by the density of our network, however. When we doubled the density of neurons (2000 neurons with place fields spanning the same environment), allowing us to decrease the recurrence length scale, replay became more reliable, and replay speed decreased substantially across all parameter sets (compare Fig 3D with Fig 3S1A and Fig 3J-K with Fig 3S1B-C).

A key component of brain dynamics underlying cognitive flexibility may be the ability to selectively route information from one neural subnetwork to another [48, 49, 50]. One substrate for such selective information propagation may be flexible sequential activity patterns that coordinate activity between disparate areas within the network [51, 52, 4], which our results suggest LTP-IE may facilitate. For instance, if activating one end of the LTP-IE profile reliably causes activation at the other, this could be used to quickly carve out a temporary communication channel within the network. To this end, we measured the mutual information (MI) between the presence/absence of a triggering pulse at the lower right end of the trajectory, and subsequent elevated spiking activity at the opposing end in the upper left (where perfect transmission is 1 bit). When we computed MI along the LTP-IE profile as a function of gating input rate r_G and maximum LTP-IE level σ_{max} , which both influence the spontaneous spiking that is most likely to interrupt MI, we found that MI was maximized for a high r_G and moderate max, although MI remained below 1 bit for all parameter values, indicating that transmission was never 100% reliable (Fig 3L). This suggests that for LTP-IE-induced sequences, information transmission through these sequences may be maximized for moderate values of σ_{max} , so stronger differentiation




Fig 3: Parameter dependence of LTP-IE-based sequential replay. *A. Demonstration of “fade-out”, in which sequential activity initiated at lower right fails to propagate along full trajectory defined by LTP-IE profile. Same format as Fig 2F-G, with colors indicating spiking order (with only first spikes used for neurons spiking multiple times) during replay epoch. B. Demonstration of “blowup”, in which all neurons activate during replay epoch. C. Demonstration of “disordered” replay, in which LTP-IE-tagged neurons activate during replay epoch but not in either a fully forward or fully reverse order. In this example, activity was triggered by a current injection in the middle of the trajectory, which caused propagation in both directions. D. Profile of replay dynamics regimes as a function of recurrent connectivity length scale $\lambda^{PC,PC}$ and scale factor $w^{PC,PC}$ (see methods). Colors corresponding to Fig 3F colorbar indicate number of successful replay events out of ten trials, for parameter regimes that yielded successful replay. If fewer than 2/10 successful replay events occurred, colors indicate most common dynamics (fadeout: blue; blowup: cyan; disordered: green). E. Same as D but showing dynamics as function of gate input rate r_G and max LTP-IE value σ_{max} . F. Same as E, F but showing dynamics as function of $w^{PC,PC}$ and r_G . G. Spontaneous spike rate of neurons with at least 50% of max LTP-IE value, calculated over 10-second time period absent external stimulation, as function of*

Continued from Fig 3: $\lambda^{PC,PC}$ and $w^{PC,PC}$. Spontaneous spike rates were calculated for all non-fadeout/non-blowup parameter regimes, i.e. where ordered or disordered replay occurred. H. Same as G but as function of r_G and σ_{max} . I. Same as G, H except as function of $w^{PC,PC}$ and r_G . J. “Virtual” replay speed of triggered sequential activity propagation for all parameter regimes in G. K. Same as J but for parameter regimes in H. L. Mutual information between trigger presence and spike detection at upper left between 50-250 ms following trigger, calculated for all non-fadeout and non-blowup parameter regimes. Perfect information transmission is 1 bit.

between LTP-IE-tagged and non-LTP-IE-tagged cells does not necessarily imply improved information transmission. We note further that we could increase MI to 1 by increasing the density of neurons in the network and shortening their recurrence length scale (Fig 3S1C), with information transmission likely improving due to decreased variability in the average signal propagation (since more neurons are involved). Thus, experience-dependent LTP-IE at biologically realistic levels can rapidly and flexibly modulate sequence-based information propagation through a network of spiking neurons.

3.3.4 LTP-IE-based sequences can encode temporary stimulus-response mappings

A popular hypothesis in cognitive neuroscience is the multiplexing and interaction of mnemonic and spatial representations, contributing, for example, to the formation of “memory maps”. Indeed, hippocampus’ experimentally observed roles in both memory and spatial navigation lends compelling evidence to this notion [53]. The potential neural mechanisms underlying this phenomena, however, remain poorly understood. To demonstrate how LTP-IE could shape temporary memory storage and computation more generally, beyond simply replaying recent sensorimotor sequences, we considered the simple task of requiring a network to represent one of two possible stimulus-response mappings, and then asked how LTP-IE could induce these mappings in the network (Fig 4A). We assume such an induction could be generated by appropriately transformed sensorimotor inputs from upstream areas, but we did not model this explicitly, as our goal was to demonstrate the final results.

LTP-IE was able to induce multiple stimulus-response mappings in the network, each in multiple ways. As an example we let two regions in the network represent two stimuli S1 and S2, and two regions represent motor outputs M1 and M2, with the average activity in each region read out by a corresponding downstream readout unit (Fig 4B). We then introduced non-intersecting LTP-IE paths into the network connecting each stimulus region to its associated motor output region, and recorded the ability of a trigger in each stimulus region to subsequently activate the correct motor output (Fig 4C-N). (We discuss intersections in the Discussion.) We first verified our idea with a simple mapping




Fig 4: Using LTP-IE-based sequential activity propagation to maintain and decode pairwise associations. *A. Example of two different mappings between a pair of stimuli ($S1$, $S2$) and a pair of responses ($M1$, $M2$). In Mapping 1 (solid), activating $S1$ should activate $M1$, and activating $S2$ should activate $M2$. In Mapping 2 (dashed) activating $S1$ should activate $M2$, and activating $S2$ should activate $M1$. **B.** Depiction of simulation structure. Either $S1$ or $S2$ is activated with a recall cue (injected current input into colored cells). Readout units average all units from either $S1$, $S2$, $M1$, or $M2$, indicated by colors. **C.** Example LTP-IE-based encoding of Mapping 1 with neurons laid out in 2-D coordinate space. **D.** Time-dependent readout responses (normalized to maximum readout response over the 3 s), with colors corresponding to readout units depicted in **B**. $S1$ and $S2$ were each alternately activated 3 times by appropriate recall cue (direct current injection into the relevant neurons; inverted colored triangles), and all readout responses were plotted. In this example, activating $S1$ (magenta) causes $M1$ (red) to activate, and activating $S2$ (cyan) causes $M2$ (blue) to activate, due to spike propagation along paths defined by LTP-IE profile in **C**. **E.** Order of first spikes of all neurons that spiked during shaded time period in **D**. **F.** Example LTP-IE-based encoding of Mapping 2 with neurons laid*

Continued from Fig 4: *out in same 2-D coordinate space as in C. G. Readout responses for same recall cue pattern as D. H. First spike order of all neurons that spiked during time-period shaded in G. I-K. Same as F-H but for an LTP-IE profile encoding Mapping 2 that is distinct from that in F. L-N. Same as F-H, I-K but for a third alternative LTP-IE profile encoding Mapping 2.*

in which each stimulus was connected via a one-dimensional LTP-IE profile to its appropriate motor output (Fig 4C-E). Indeed, triggering each stimulus region with a depolarizing current input led to subsequent activation of its corresponding motor output after a short delay (Fig 4D-E), although there were occasionally low levels of spontaneous activity. We next explored a mapping requiring at least one LTP-IE path to take a roundabout course through the network (Fig 4F-H). As before, triggering each stimulus region led to activation of its corresponding motor output, with the longer LTP-IE path reflected in a longer stimulus-response delay (Fig 4G-H). Responses, however, appeared less reliable than when shorter LTP-IE paths were used, likely due to the increased opportunities for propagation to fail due to variability in the baseline G inputs. Notably, the same stimulus-response mapping could be implemented via LTP-IE in several different ways (Fig 4F, I, L), with the different LTP-IE paths reflected in the different stimulus-response delays. This contrasts with alternative models of temporarily binding together distinct components of a neural network, which typically suppose a single stereotyped structure of the mapping/binding representation [54, 55, 56]. Thus, LTP-IE combined with recurrent connectivity might serve as a biophysically plausible, yet highly flexible substrate for inducing temporary stimulus-response mappings in a recurrent spiking network, a potential key property enabling rapid and flexible induction of temporary information-processing patterns in brain networks underlying cognitive flexibility.

3.4 Discussion

Structured sequential spiking activity is a key feature of recurrent neural network dynamics, potentially reflecting information flow and computations within the network. One compelling example of spike sequences observed is the so-called “replay” of neural sequences representing recent sensorimotor sequences through the environment. This has been observed in vivo in both hippocampus and cortex during awake quiescent periods in both rodents and primates [31, 28, 57, 32, 33, 34] and is thought to be involved in navigational planning [58, 59] and memory consolidation [33, 36, 60, 61]. Such replay serves as a compelling entry point for uncovering the biophysical mechanisms that support rapid modulation of information processing. Accordingly, here we take the perspective that sequential

spike patterns reflect information transmission between distinct subnetworks. While substantial work has shown how modifications to the structure of recurrent network connections can induce sequences into a network’s dynamical repertoire [2, 10, 3, 4], little was known about alternative mechanisms acting on faster timescales. Here we have demonstrated the sufficiency of an empirically observed, strong, and fast-acting heterosynaptic plasticity mechanism known as long-term potentiation of intrinsic excitability (LTP-IE) [38, 39], applied to a spatially organized recurrent spiking network, to coerce a network into generating selective sequences. These are capable of recapitulating recent sensorimotor sequences, implicating their potential role in memory storage, as well as encoding simple stimulus-response mappings, suggesting their suitability for organizing flexible computations.

Additionally, as observed *in vivo* and as predicted by our model, the sequences that replay during awake quiescence are significantly compressed relative to the initial behavioral timescales, suggesting replayed spikes may occur within the appropriate time windows for long-term, spike-timing-dependent plasticity mechanisms (STDP) [13, 14, 62]. For example, it was shown in CA1 place cells that replaying place-tuned firing sequences *in vitro* that were previously recorded *in vivo* could induce long-term connectivity changes between cells with overlapping place fields [63]. Our model reveals a mechanism by which spatial sequences could be transiently stored in the first place, in order to be later embedded into long-term memories.

Finally, although our work was inspired primarily by hippocampus and the notion of LTP-IE-based short-term sequences, the core concept of excitability-based network dynamics might naturally extend to cortical areas, and to more diverse timescales [64, 65]. For example, it was recently found that small neuronal ensembles could be “imprinted” into a mouse cortical network and subsequently reactivated *in vivo* via repeated optogenetic stimulation [66]. While this result was taken as evidence of Hebbian (associative) synaptic plasticity, our work suggests the same result might also arise simply via increasing the excitability of the neurons in the ensemble. Distinguishing these two mechanisms in cortical circuits and understanding how they interact presents an exciting avenue for further investigation.

3.4.1 Comparison to existing sequence-generation models

While we do not aim to provide a complete account of hippocampal replay, but rather to demonstrate how specific biophysical mechanisms can rapidly generate selective network sequences, we believe it is still worthwhile to compare our model to existing models for rapidly storing and recalling sequences in neural systems, both in terms of biological plausibility and computational robustness. One family

of models for encoding and decoding sequences in hippocampus supposes that sequential information is stored in the timing of spikes relative to theta (5-10 Hz) and gamma (40 Hz) oscillatory cycles in the hippocampus [67, 44]. Briefly, gamma cycles are “nested” within theta cycles, and a sequence of stimuli is stored by stimulus-specific neurons spiking at unique gamma cycles within the encompassing theta cycle. While these models provide a persuasive computational account of how oscillations might be used to store temporary information, it is not clear how the sensory input would be appropriately transformed so as to be “entered” into the oscillation cycle at the correct time, how stable these mechanisms are in the face of noise, nor how multiple sequences might be stored simultaneously in the network; in particular, sequence information could be lost if the oscillation were disrupted, and one would require an independent nested oscillation for a second sequence. In our LTP-IE-based model, however, memories are directly entered into the network by physiological spiking patterns (Fig 1), are stored in effective synaptic weight changes likely driven by ion channel inactivation (as opposed to persistent spiking activity) and might therefore be more robust to noise, and multiple sequences can theoretically be stored in the same network (Fig 4).

The second main family of replay models either explicitly [68, 69, 70] or implicitly [71] assume that the sequences produced by the network are initially induced in the network via modification of recurrent network weights, for example through homosynaptic mechanisms like spike-timing-dependent plasticity (STDP) [13, 14, 7]. While such models can indeed selectively bias network sequences when implemented computationally, given the hypothesized weak magnitude of STDP, it is unclear whether such a mechanism could account for the awake trajectory replay observed in rodents, which can sometimes occur even when the animal has experienced the original trajectory only once [31], nor whether it could be used to immediately induce novel computations that the animal cannot afford to learn over extensive repetitions. In our model we rely on a strong, fast-acting mechanism, LTP-IE, which although heterosynaptic in nature (unlike STDP, LTP-IE follows spiking regardless of which input elicited the spikes [38, 39]) suffices in biasing a network toward producing highly specific sequences: LTP-IE “tags” cells in the recurrent network with a higher probability of spiking in response to input (Fig 1), and sequential ordering is reconstructed by the existing but unmodified recurrent connectivity. While recent research has begun to investigate how small, biophysically plausible recurrent weight changes might successfully encode new memories [72, 73], to our knowledge this has not been applied in the context of sequences. Overall, it is likely that memory and computational flexibility over short timescales relies on a host of plasticity and dynamics mechanisms, but we propose that LTP-IE may

play a significant role.

3.4.2 Model limitations

As a proof-of-concept account demonstrating the sufficiency of a small number of mechanisms to quickly bias selective sequence production in a network, we did not explicitly model several features in order to keep our model parsimonious. First, we did not include inhibition, as we discovered that LTP-IE-based sequences could arise through excitatory connectivity alone. Among a number of possible roles, inhibition in hippocampal and cortical networks supports network oscillations [74, 36] and balances excitatory activity to stabilize dynamics [75, 71]. Indeed, empirical hippocampal sequence replay co-occurs with brief oscillatory events called sharp-wave ripples [Carr 2011, Nat. Neurosci.], suggesting that hippocampal inhibitory circuitry may contribute to oscillatory replay patterns, perhaps serving as a clock-like signal for controlling signal propagation. An additional role for inhibition in sequence replay may be to induce competition among reactivating representations, for example so that multiple sequences would not replay simultaneously, or so activity traveling along a branching LTP-IE-path would follow only one branch at a time. Inhibitory balance might also serve to decrease the parameter sensitivity of our model to small variations in $w^{PC,PC}$ or PC,PC (Fig 3D) by providing a counterbalance to over-excitation, but precisely how the structure and dynamics of inhibitory interactions with our network would modulate LTP-IE-based sequences is beyond the scope of this work.

A further issue with LTP-IE-based sequences surrounds trajectories that intersect themselves. At an intersection point, our current model would typically lead to sequential reactivation that spread in multiple directions, thus failing to recapitulate the original trajectory. This, however, could be resolved by the embedding of trajectories in a higher-dimensional space beyond just the $x - y$ plane. Indeed, neuronal firing is typically modulated by multiple stimulus features; for example, hippocampal neurons are also tuned to head-direction (θ) [76] and running velocity [77]. Thus, if excitatory connectivity between neurons with similar x - and y -tuning were only present if neurons had similar θ -tuning also, the trajectory representation would expand into a third θ dimension, such that it might no longer intersect with itself and replay would reflect the actual trajectory travelled. Consequently, extending our assumption of preferential connectivity among similarly tuned neurons to additional input dimensions could increase the capacity of our network to store non-interfering sequences.

Additionally, while we have assumed the upstream “gating” inputs in our model are stochastic

with homogeneous firing rates, cortical activity likely contains higher order structure beyond mean firing rates, such as the grid-like representations of space observed in neural firing patterns in entorhinal cortex [78, 79], the hippocampal input that inspired our gating inputs. Our model demonstrates that replay can occur absent additional mnemonic information in such, but one would imagine such additional information would only increase the capacity and robustness for memory storage and computational flexibility, since this would act as an additional memory buffer.

Finally, our model begs the question of how information would be erased from the network or prevented from entering it in the first place. Indeed, with no forgetting or selective LTP-IE activation, one would expect all neurons in our model to undergo LTP-IE after the animal had fully explored the environment, thus preventing sequential reactivation along specific trajectories through the network. While this problem would be partially solved by considering additional input dimensions as described above (i.e. so that even if the $x - y$ space were fully explored, only a fraction of the $x - y - \theta$ space would have been explored), one might also imagine additional neuromodulatory inputs serving to control the strength of LTP-IE during exploration. For example, the neurotransmitter acetylcholine is known to interact with the voltage-gated potassium channel Kv1.2 [80], the ion channel hypothesized to underlie LTP-IE [39], and hippocampus receives state-dependent cholinergic inputs from medial septum [81, 82]. Thus, sufficient machinery for erasing or selectively modulating the overall effects of LTP-IE, say, as a function of arousal or motivation state, may converge in the brain region that has served as the inspiration for our model.

3.4.3 Experimental predictions

Despite its parsimony, our model of LTP-IE-based network sequences makes specific predictions that could be tested in hippocampal replay experiments. First, pharmacological blocking of LTP-IE during exploration should prevent the encoding of new trajectory information into CA3, consequently preventing subsequent replay. This could be achieved by in vivo application of nimodipine or PP2, which block intracellular calcium increase, protein tyrosine kinase activation, and Kv1.2 channel endocytosis, all potential mechanisms of LTP-IE, and which have been shown to block LTP-IE in vitro [39]. Second, inactivating medial entorhinal layer II (MEC II), which projects to CA3 and which we associate with gating inputs in our analogy with hippocampal circuitry, should inhibit CA3 sequential replay events. Indeed, inactivation of MEC III inputs has been shown to inhibit forms of replay in CA1 [79]; we propose that a parallel effect would be seen in MEC II inputs to CA3. Finally, the core

prediction of our sequence induction model is that the neurons do not have to be initially activated in the order (or the reverse order) in which they will later reactivate. As the only information explicitly added to the network at the time of encoding is the set of neurons activated, speed or ordering information must be reconstructed from existing connectivity. While identifying the position tunings of individual neurons in vivo and subsequently artificially activating them out of order currently presents a substantial technical challenge, one could potentially achieve a similar effect using a virtual reality environment with place-specific sensory cues, and in which specific places along a virtual trajectory were presented out of order. Our model predicts that as long as the animal had sufficient knowledge of which sensory cues corresponded to which virtual environment locations, replay would occur in an order corresponding to a path through the environment, rather than the order in which the virtual places were presented to the animal.

3.4.4 Computational significance

Besides revealing a potential core mechanism for sequence induction, our model lends two key insights into the structure and control of replay in biological neural networks. First, it suggests a link between isotropy in pre-existing spatial representations in the network (i.e. the spatially organized connectivity) and the experimental observation that replay sometimes occurs in the forward, and sometimes in the reverse order relative to the neurons’ initial activation during the trajectory [31, 28, 33]. In particular, since no directionality information is stored by LTP-IE (which only tags which neurons were recently active), replay should occur with equal chance in forward or reverse. (Note however, that as observed in [28], replay events may have an increased probability of originating at cells tuned to the animal’s current location, as a result of place-tuned inputs to the network during awake quiescence.) One would therefore not expect to observe reverse replay for neural sequences that do not have a natural spatial embedding, e.g. spoken sentences in humans. Second, the upstream gating signal in our model might serve as a substrate for the short-term state-dependence of replay, e.g. with arousal or motivation controlling when the gating signal is present or absent. This could ensure that replay occurs only at appropriate times (e.g. quiescent rest) without interfering with other computations being performed by the network, such as pattern separation in external sensory inputs [83, 84].

Further, while our network was embedded in 2-D space, our results frame the more general question of how a network embedded on a higher dimensional manifold, or with no natural embedding, supports excitability- based sequence induction and information flow. For example, one could imagine the

rapid induction of spike sequences corresponding to movement along 3-dimensional trajectories if the recurrent architecture reflected a 3-dimensional space, as might occur within the three-dimensional hippocampal place-cell network of free-flying bats [85]. More generally it remains an open and intriguing question as to which network structures are most suited to rapid excitability-based flexibility among sequences and behaviorally relevant computations. One example of a network model with no obvious spatial embedding, and which could potentially be recast in the framework of excitability-based information flow, is the “binding pool” model of associative memory proposed by Swan Wyble [56]. Here, associations are encoded in a set of steady, persistently active “binding” units that allow preferential signal transmission among specific sets of units representing object features, to which the binding units are randomly connected. While the role of such persistent spiking in coding mnemonic information is presently undergoing substantial reevaluation [86], one could theoretically replace it with excitability increases, with an equivalent consequence for information transmission. Whether this could be done via biological mechanisms remains an open question, but our present work suggests its plausibility.

Finally, one can draw an equivalence between the modulation of cellular properties to shape sequences and information flow and selective variation in upstream modulatory inputs to a cell. For example, a neuron receiving strong inhibition might be less likely to participate in a computation than one receiving excitation, since the latter would be more likely to spike in response to inputs, similar to if one had increased its excitability through intrinsic cellular mechanisms. Indeed, it was recently shown that external modulatory inputs to a recurrent network could control the speed at which a computation performed by the network unfolded through time, closely recapitulating experimental results [87]. We propose that the control of information flow and computation in recurrent networks via effective excitability alone, without structural connectivity changes, might be a rich framework for investigating rapid cognitive flexibility and short-term information storage. This may lead to significant insights into how we adapt our mental processes and behaviors in the face of ever-changing environments and task demands.

3.5 Methods

3.5.1 Membrane and synaptic dynamics

We modeled single neurons as leaky integrate-and-fire neurons with conductance-based synapses according to the equation

$$\tau_m \frac{dv^i}{dt} = -(v^i - E_{leak}) - g_E^i(t)(v^i - E_E) + I_{ext}^i(t)$$

where v_i is the membrane voltage of neuron i , $\tau_m = 50$ ms is the membrane time constant, $E_{leak} = -68$ mV is the leak/resting potential, g_E is the time-varying conductance of excitatory synaptic inputs, $E_E = 0$ mV is the excitatory synaptic reversal potential and I_{ext} is externally injected currents (expressed as voltage deflections). If v exceeded a threshold $v_{th} = -36$ mV at any time step, a spike was emitted and v was reset to E_{leak} for a refractory period $\tau_r = 8$ ms. This outsize refractory period ensured that the sequential events in our simulations propagated only in one direction and is consistent with refractory periods measured in hippocampal region CA3 [88].

Synaptic dynamics were conductance-based and were modeled as exponentially filtered trains of weighted delta functions representing input spikes arriving from different upstream neurons. Specifically, for an individual neuron i

$$g_E^i(t) = \sum_j \sum_k w_{ij} h(t) * \delta(t - t_k^j)$$

where j indexes neurons and k indexes spikes, such that w_{ij} is the synaptic weight from neuron j onto neuron i and t_{kj} is the k -th spike of the j -th neuron; and $h(t)$ was a one-sided exponential with time constant $\tau_E = 2$ ms. In our simulations we assumed all synapses were excitatory and did not explicitly model effects of inhibition. Weights were unitless since they acted only to scale conductance changes.

3.5.2 Neuron tuning and network structure

In network simulations, we assigned each neuron a position (x_i, y_i) corresponding to the peak of its tuning curve, the (x, y) position eliciting maximal firing. Positions were assigned so as to approximately tile a grid spanning -1 to 1 m in both dimensions (Fig 2, 3) or -2 to 2 m (Fig 4), using 1000 or 4000 neurons, respectively. We used the following position-dependent firing-rate equation to calculate the maximal firing rate we expected to be evoked by the trajectory through the simulated environment (the firing rate evoked by the closest point on the trajectory to the neuron's place field center):

$$r^i(x, y) = r_{max} \exp \left[-((x - x^i)^2 + (y - y^i)^2) / (2\lambda_{PL}^2) \right]$$

where $r_{max} = 20$ Hz is the maximum firing rate (i.e. when the trajectory passes directly through the neuron’s peak tuning position) and $\lambda_{PL} = 0.15$ m is the length constant determining how close a neuron’s peak tuning must be to the trajectory to significantly activate. All neurons had the same r_{max} and λ_{PL} .

We explicitly modeled two sources of synaptic input to the neurons in our recurrent network: recurrent inputs and gating inputs. Recurrent synaptic weights were assigned such that neurons with similar position tuning had stronger connectivity, with the connectivity between two neurons i and j with peak tunings separated by d_{ij} given by

$$w_{ij} = w^{PC,PC} \exp \left[-d_{ij}^2 / (2(\lambda^{PC,PC})^2) \right]$$

This was motivated by past modeling studies [45, 89] and by analyses of correlated hippocampal activity in the absence of sensory input [46]. Note that in our network architecture, all recurrent connections between position-tuned cells are reciprocal, due to the symmetry of the distance function d_{ij} . All w_{ij} below a minimum value $w_{min}^{PC,PC} = 0.001$ were set to 0.

Synaptic weights w_i^G on gating inputs G were initially uniform across all position-tuned neurons, and each neuron received an independent instantiation of an upstream G spike train, generated from a Poisson-distributed point process with constant rate r_G . While effective w_i^G varied as a function of each neuron’s trajectory-dependent LTP-IE, r_G was identical across all neurons regardless of tuning or LTP-IE status.

3.5.3 Long-term potentiation of intrinsic excitability (LTP-IE) of gating inputs

Although LTP-IE is thought to result from changes in dendritic conductances [39], we modeled it as an effective synaptic weight change (since conductance changes can be absorbed into connection weights according to the equations underlying conductance-based synaptic dynamics), such that initial weights w_i^G were scaled by a factor σ_i as a function of neuron i ’s position along the initial trajectory. For computational efficiency we did not model network activity during the initial trajectory, but instead directly calculated the σ_i expected to result from the maximum firing rate $r_{max}^i = \max(r_i(x, y) | (x, y) \in trajectory)$. From r_{max}^i we computed σ_i according to the following equation (Fig 1B):

$$\sigma_i = 1 + \frac{\sigma_{max} - 1}{1 + \exp[-\beta_\sigma(r_{max}^i - r_\sigma)]}$$

with $r_\sigma = 10$ Hz and $\beta_\sigma = 1$. Thus, neurons with position-tuning peaks close to the original trajectory had a i near max whereas neurons far away from the trajectory had σ_i near unity.

3.5.4 Sequence replay

We fully simulated the replay epoch, during which all neurons in the recurrent network received independent stochastic gating input G at constant rate r_G . Replay was triggered by applying a short depolarizing current pulse to a set of neurons with peak position-tuning within 0.4 m (Fig 2-3) or 0.3 m (Fig 4) of a central trigger location (e.g. $(x = 1m, y = -0.75m)$ in Fig 2F) for 3 ms, of amplitude 8 mV (Fig 2-3) or 7 mV (Fig 4); exact trigger parameters were unimportant, as long as they elicited significant spiking activity in the neurons at the start of the sequence. In Fig 2L, replayed positions were decoded during by taking the median peak-tuning of neurons that spiked within 10 ms windows spanning the replay event.

3.5.5 Simulations

3.5.5.1 Single-neuron LTP-IE simulation (Fig 1)

LTP-IE-level-dependent voltage distributions (Fig 1D, E) were computed using 15-second simulations. In measuring Fig 1F-G, current pulses were 10 ms in duration, and spiking probability was measured within this window. Current-evoked spike probabilities (Fig 1G) were computed using 125-second simulations, with pulses presented every 250 ms.

3.5.5.2 Replay in networks of neurons (Fig 2-3)

In replay simulations, we used 1000 neurons, with place fields distributed on an approximate lattice over a 2 m x 2 m environment. Replay was triggered by a short, depolarizing current pulse into all cells with place field centers within a 0.4 m radius of either $(-1\text{ m}, 0.75\text{ m})$ or $(1\text{ m}, -0.75\text{ m})$, although our results were not sensitive to these precise numbers, so long as sufficiently many LTP-IE-tagged cells at one end of the trajectory were activated by the trigger.

3.5.5.3 Replay dynamics parameter sweep (Fig 3)

For each parameter set we ran ten 800 ms simulations starting with different random number generator seeds (which controlled the instantiation of gating inputs, which is the only stochasticity in our simulation). Each simulation contained one replay trigger at 500 ms, and we counted how many of the ten triggers resulted in either fadeout, blowup, disordered replay, or successful replay. If successful replay occurred 2 or more times out of the 10 simulations, we colored the point in parameter space by how many successful replay events occurred; otherwise we colored it by the most common unsuccessful event type (fadeout: blue; blowup: cyan; disordered: green).

We first varied the magnitude factor $w^{PC,PC}$ and the length scale $\lambda^{PC,PC}$ of the recurrent excitatory weight profile (Fig 2D), while holding $\sigma_{max} = 2$ and $r_G = 125$ Hz. We next varied σ_{max} and r_G while holding $w^{PC,PC} = 0.03$ and $\lambda^{PC,PC} = 0.083$ m, (Fig 3E). Finally, we varied $w^{PC,PC}$ and r_G while fixing $\lambda^{PC,PC} = 0.083$ m and $\sigma_{max} = 2$ (Fig 3F).

Spontaneous activity was determined by computing spike rates over 10 seconds without no additional current inputs. Replay speed was calculated as the slope of a line fit to the neuronal positions along the trajectory profile and the times at which they activated during replay. Note that for high levels of spontaneous activity but without blowup, two replay events sometimes occurred in this window, distorting the speed measure. Since these distorted measurements consistently yielded computed replay speeds between 0 and 20 m/s, whereas non-distorted replay speeds were grouped in an approximately Gaussian-shaped cluster well over 20 m/s, we excluded all measured replay speeds ≤ 20 m/s from our analysis and plots.

To measure mutual information between a replay trigger at one end of the trajectory and elevated spiking activity at the opposing end, we defined the trigger as an injected current into cells with place fields within 0.4 m of (1 m, -0.75 m) and measured elevated spiking activity of neurons within 0.4 m of (-1 m, 0.75 m) between 50 and 250 ms following the trigger. The binary trigger variable corresponded to the presence or absence of the trigger, and the binary opposing-end-spiking variable was positive if fewer than 90% of the opposing-end neurons spiked and otherwise negative.

3.5.5.4 LTP-IE-based stimulus-response mappings (Fig 4)

For this simulation, in order to construct sufficiently non-interfering LTP-IE-defined paths, we constructed a model network of 4000 neurons with neural place fields ranging from -2 to +2 m in both the x - and y - dimensions and let two clusters of radius of 0.25 m, centered at (-1, 1) and (-1, -1), represent two stimuli S1 and S2, respectively, and two additional clusters, centered at (1, 1) and (1,

Symbol	Definition	Value
τ_m	Membrane time constant	50 ms
E_{leak}	Leak potential	-68 mV
v_{th}	Spike threshold	-36 mV
τ_r	Refractory period	8 ms
E_E	Excitatory synaptic reversal potential	0 mV
τ_E	Excitatory synaptic time constant	2 ms
r_{max}	Max position-driven firing rate	20 Hz
λ_{PL}	Position-tuning length constant	0.15 m
r^G	Gating input firing rate	125 Hz (Fig 2-3), 105 Hz (Fig 4)
N	Number of neurons in recurrent network	1000 (Fig 2-3), 4000 (Fig 4)
w^G	Initial gating input weight	0.008216
$w^{PC,PC}$	Max recurrent weight	0.03 (Fig 2-3), 0.034 (Fig 4)
$w_{min}^{PC,PC}$	Min nonzero recurrent weight	0.001
$\lambda^{PC,PC}$	Recurrent connectivity length scale	0.083 m (Fig 2-3), 0.0835 m (Fig 4)
σ_{max}	Max LTP-IE level	2
r_σ	Threshold firing rate for LTP-IE	10 Hz
β_σ	Scale factor for LTP-IE onset	1
ΔT	Simulation time step	0.5 ms

-1), encode motor outputs M1 and M2, respectively. Readout units summed activity in any of the four clusters at every timestep. Sequence reactivation was triggered by depolarizing current injections into cells contained in the S1 or S2 clusters.

3.5.5.5 Other simulation parameters

Unless otherwise noted we used the following values for all simulation parameters:

All code for this work was written in Python 3 and is available at https://github.com/rkp8000/seq_speak.

3.6 Supplementary figures




Fig 1S1: Effects of current pulses on spiking under LTP-IE. *A. Probability of current-evoked spiking for different $w^{PC,G}$ and σ . B. Spontaneous spiking probability for different $w^{PC,G}$ and σ . C. Example voltage traces for neurons with two different σ given large $w^{PC,G}$. Same format as Fig 1F.*

References

- [1] Ronald J Williams and David Zipser. A learning algorithm for continually running fully recurrent neural networks. *Neural computation*, 1(2):270–280, 1989.
- [2] David Sussillo and Larry F Abbott. Generating coherent patterns of activity from chaotic neural networks. *Neuron*, 63(4):544–557, 2009.
- [3] Rodrigo Laje and Dean V Buonomano. Robust timing and motor patterns by taming chaos in recurrent neural networks. *Nature neuroscience*, 16(7):925, 2013.
- [4] Kanaka Rajan, Christopher D Harvey, and David W Tank. Recurrent network models of sequence generation and memory. *Neuron*, 90(1):128–142, 2016.
- [5] Wilten Nicola and Claudia Clopath. Supervised learning in spiking neural networks with force training. *Nature communications*, 8(1):2208, 2017.
- [6] Sen Song and Larry F Abbott. Cortical development and remapping through spike timing-dependent plasticity. *Neuron*, 32(2):339–350, 2001.




Fig 3S1: Replay features for denser networks. *A.* As in Fig 3D but for a “dense” network of 2000 neurons. *B.* As in Fig 3J but for dense network. *C.* As in Fig 3K but for dense network. *D.* As in Fig 3L but for dense network.




Fig 3S2: Mutual information calculation. *Plots of conditional probabilities used to calculate mutual information in Fig 3L. A. Probability of no elevated spiking at opposing end of trajectory (“final = F”) given no initiating trigger (“trigger = F”). B. Probability of elevated spiking at opposing end of trajectory (“final = T”) given no initiating trigger. C. Probability of no elevated spiking at opposing end of trajectory given initiating trigger (“trigger = T”). D. Probability of elevated spiking at opposing end of trajectory given initiating trigger.*

- [7] Ila R Fiete, Walter Senn, Claude ZH Wang, and Richard HR Hahnloser. Spike-time-dependent plasticity and heterosynaptic competition organize networks to produce long scale-free sequences of neural activity. *Neuron*, 65(4):563–576, 2010.
- [8] Matthieu Gilson, Anthony Burkitt, and Leo J Van Hemmen. Stdp in recurrent neuronal networks. *Frontiers in computational neuroscience*, 4:23, 2010.
- [9] Tyler P Lee and Dean V Buonomano. Unsupervised formation of vocalization-sensitive neurons: a cortical model based on short-term and homeostatic plasticity. *Neural computation*, 24(10):2579–2603, 2012.
- [10] Stefan Klampfl and Wolfgang Maass. Emergence of dynamic memory traces in cortical microcircuit models through stdp. *Journal of Neuroscience*, 33(28):11515–11529, 2013.
- [11] Danilo J Rezende, Daan Wierstra, and Wulfram Gerstner. Variational learning for recurrent spiking networks. In *Advances in Neural Information Processing Systems*, pages 136–144, 2011.
- [12] Peter U Diehl and Matthew Cook. Unsupervised learning of digit recognition using spike-timing-dependent plasticity. *Frontiers in computational neuroscience*, 9:99, 2015.
- [13] Henry Markram, Joachim Lübke, Michael Frotscher, and Bert Sakmann. Regulation of synaptic efficacy by coincidence of postsynaptic aps and epsps. *Science*, 275(5297):213–215, 1997.
- [14] Guo-qiang Bi and Mu-ming Poo. Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *Journal of neuroscience*, 18(24):10464–10472, 1998.
- [15] Per Jesper Sjöström, Gina G Turrigiano, and Sacha B Nelson. Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron*, 32(6):1149–1164, 2001.
- [16] Huai-Xing Wang, Richard C Gerkin, David W Nauen, and Guo-Qiang Bi. Coactivation and timing-dependent integration of synaptic potentiation and depression. *Nature neuroscience*, 8(2):187, 2005.
- [17] Natalia Caporale and Yang Dan. Spike timing-dependent plasticity: a hebbian learning rule. *Annu. Rev. Neurosci.*, 31:25–46, 2008.
- [18] JS Griffith and G Horn. An analysis of spontaneous impulse activity of units in the striate cortex of unrestrained cats. *The Journal of Physiology*, 186(3):516–534, 1966.

- [19] Alexei A Koulakov, Tomáš Hromádka, and Anthony M Zador. Correlated connectivity and the distribution of firing rates in the neocortex. *Journal of Neuroscience*, 29(12):3685–3694, 2009.
- [20] Kenji Mizuseki and György Buzsáki. Preconfigured, skewed distribution of firing rates in the hippocampus and entorhinal cortex. *Cell reports*, 4(5):1010–1021, 2013.
- [21] Alfonso Renart, Jaime De La Rocha, Peter Bartho, Liad Hollender, Néstor Parga, Alex Reyes, and Kenneth D Harris. The asynchronous state in cortical circuits. *science*, 327(5965):587–590, 2010.
- [22] Timothée Masquelier, Rudy Guyonneau, and Simon J Thorpe. Spike timing dependent plasticity finds the start of repeating patterns in continuous spike trains. *PloS one*, 3(1):e1377, 2008.
- [23] Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, pages 109–165. Elsevier, 1989.
- [24] Richard HR Hahnloser, Alexay A Kozhevnikov, and Michale S Fee. An ultra-sparse code underlies the generation of neural sequences in a songbird. *Nature*, 419(6902):65, 2002.
- [25] Yuji Ikegaya, Gloster Aaron, Rosa Cossart, Dmitriy Aronov, Ilan Lampl, David Ferster, and Rafael Yuste. Synfire chains and cortical songs: temporal modules of cortical activity. *Science*, 304(5670):559–564, 2004.
- [26] Artur Luczak, Peter Barthó, Stephan L Marguet, György Buzsáki, and Kenneth D Harris. Sequential structure of neocortical spontaneous activity in vivo. *Proceedings of the National Academy of Sciences*, 104(1):347–352, 2007.
- [27] Eva Pastalkova, Vladimir Itskov, Asohan Amarasingham, and György Buzsáki. Internally generated cell assembly sequences in the rat hippocampus. *Science*, 321(5894):1322–1327, 2008.
- [28] Thomas J Davidson, Fabian Kloosterman, and Matthew A Wilson. Hippocampal replay of extended experience. *Neuron*, 63(4):497–507, 2009.
- [29] David A Crowe, Bruno B Averbeck, and Matthew V Chafee. Rapid sequences of population activity patterns dynamically encode task-critical spatial information in parietal cortex. *Journal of Neuroscience*, 30(35):11640–11653, 2010.

- [30] Christopher D Harvey, Philip Coen, and David W Tank. Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature*, 484(7392):62, 2012.
- [31] David J Foster and Matthew A Wilson. Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature*, 440(7084):680, 2006.
- [32] Anoopum S Gupta, Matthijs AA van der Meer, David S Touretzky, and A David Redish. Hippocampal replay is not a simple function of experience. *Neuron*, 65(5):695–705, 2010.
- [33] Margaret F Carr, Shantanu P Jadhav, and Loren M Frank. Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval. *Nature neuroscience*, 14(2):147, 2011.
- [34] Sarah L Eagleman and Valentin Dragoi. Image sequence reactivation in awake v4 networks. *Proceedings of the National Academy of Sciences*, 109(47):19450–19455, 2012.
- [35] Valérie Ego-Stengel and Matthew A Wilson. Disruption of ripple-associated hippocampal activity during rest impairs spatial learning in the rat. *Hippocampus*, 20(1):1–10, 2010.
- [36] Shantanu P Jadhav, Caleb Kemere, P Walter German, and Loren M Frank. Awake hippocampal sharp-wave ripples support spatial memory. *Science*, 336(6087):1454–1458, 2012.
- [37] H Freyja Ólafsdóttir, Daniel Bush, and Caswell Barry. The role of hippocampal replay in memory and planning. *Current Biology*, 28(1):R37–R50, 2018.
- [38] Jung Ho Hyun, Kisang Eom, Kyu-Hee Lee, Won-Kyung Ho, and Suk-Ho Lee. Activity-dependent downregulation of d-type k⁺ channel subunit kv1.2 in rat hippocampal ca3 pyramidal neurons. *The Journal of physiology*, 591(22):5525–5540, 2013.
- [39] Jung Ho Hyun, Kisang Eom, Kyu-Hee Lee, Jin Young Bae, Yong Chul Bae, Myoung-Hwan Kim, Sooyun Kim, Won-Kyung Ho, and Suk-Ho Lee. Kv1.2 mediates heterosynaptic modulation of direct cortical synaptic inputs in ca3 pyramidal cells. *The Journal of physiology*, 593(16):3617–3643, 2015.
- [40] Nelson Rebola, Mario Carta, and Christophe Mulle. Operation and plasticity of hippocampal ca3 circuits: implications for memory encoding. *Nature Reviews Neuroscience*, 18(4):208, 2017.
- [41] Edvard I Moser, Emilio Kropff, and May-Britt Moser. Place cells, grid cells, and the brain’s spatial representation system. *Annu. Rev. Neurosci.*, 31:69–89, 2008.

- [42] Kenji Mizuseki, Sebastien Royer, Kamran Diba, and György Buzsáki. Activity dynamics and behavioral correlates of ca3 and ca1 hippocampal pyramidal neurons. *Hippocampus*, 22(8):1659–1680, 2012.
- [43] Wulfram Gerstner, Werner M Kistler, Richard Naud, and Liam Paninski. *Neuronal dynamics: From single neurons to networks and models of cognition*. Cambridge University Press, 2014.
- [44] John E Lisman. Relating hippocampal circuitry to function: recall of memory sequences by reciprocal dentate–ca3 interactions. *Neuron*, 22(2):233–242, 1999.
- [45] Szabolcs Káli and Peter Dayan. The involvement of recurrent connections in area ca3 in establishing the properties of place fields: a model. *Journal of Neuroscience*, 20(19):7463–7477, 2000.
- [46] Chad Giusti, Eva Pastalkova, Carina Curto, and Vladimir Itskov. Clique topology reveals intrinsic geometric structure in neural correlations. *Proceedings of the National Academy of Sciences*, 112(44):13455–13460, 2015.
- [47] John O’Keefe. A review of the hippocampal place cells. *Progress in neurobiology*, 13(4):419–439, 1979.
- [48] Thomas Akam and Dimitri M Kullmann. Oscillations and filtering networks support flexible routing of information. *Neuron*, 67(2):308–320, 2010.
- [49] Andrea Stocco, Christian Lebiere, and John R Anderson. Conditional routing of information to the cortex: A model of the basal ganglia’s role in cognitive coordination. *Psychological review*, 117(2):541, 2010.
- [50] Christoph Kirst, Marc Timme, and Demian Battaglia. Dynamic information routing in complex networks. *Nature communications*, 7:11061, 2016.
- [51] Giovanni Pezzulo, Matthijs AA van der Meer, Carien S Lansink, and Cyriel MA Pennartz. Internally generated sequences in learning and executing goal-directed behavior. *Trends in cognitive sciences*, 18(12):647–657, 2014.
- [52] Shantanu P Jadhav, Gideon Rothschild, Demetris K Roumis, and Loren M Frank. Coordinated excitation and inhibition of prefrontal ensembles during awake hippocampal sharp-wave ripple events. *Neuron*, 90(1):113–127, 2016.

- [53] Daniela Schiller, Howard Eichenbaum, Elizabeth A Buffalo, Lila Davachi, David J Foster, Stefan Leutgeb, and Charan Ranganath. Memory and space: towards an understanding of the cognitive map. *Journal of Neuroscience*, 35(41):13904–13911, 2015.
- [54] Antonino Raffone and Gezinus Wolters. A cortical mechanism for binding in visual working memory. *Journal of Cognitive Neuroscience*, 13(6):766–785, 2001.
- [55] Matthew Botvinick and Takamitsu Watanabe. From numerosity to ordinal rank: a gain-field model of serial order representation in cortical working memory. *Journal of Neuroscience*, 27(32):8636–8642, 2007.
- [56] Garrett Swan and Brad Wyble. The binding pool: A model of shared neural resources for distinct items in visual working memory. *Attention, Perception, & Psychophysics*, 76(7):2136–2157, 2014.
- [57] Mattias P Karlsson and Loren M Frank. Awake replay of remote experiences in the hippocampus. *Nature neuroscience*, 12(7):913, 2009.
- [58] David J Foster and James J Knierim. Sequence learning and the role of the hippocampus in rodent navigation. *Current opinion in neurobiology*, 22(2):294–300, 2012.
- [59] Brad E Pfeiffer and David J Foster. Hippocampal place-cell sequences depict future paths to remembered goals. *Nature*, 497(7447):74, 2013.
- [60] H Freyja Ólafsdóttir, Francis Carpenter, and Caswell Barry. Task demands predict a dynamic switch in the content of awake hippocampal replay. *Neuron*, 96(4):925–935, 2017.
- [61] Mark C Zielinski, Wenbo Tang, and Shantanu P Jadhav. The role of replay and theta sequences in mediating hippocampal-prefrontal interactions for memory and cognition. *Hippocampus*, 2017.
- [62] Yang Dan and Mu-ming Poo. Spike timing-dependent plasticity of neural circuits. *Neuron*, 44(1):23–30, 2004.
- [63] John TR Isaac, Katherine A Buchanan, Robert U Muller, and Jack R Mellor. Hippocampal place cell firing patterns can induce long-term synaptic plasticity in vitro. *Journal of Neuroscience*, 29(21):6840–6850, 2009.
- [64] Wei Zhang and David J Linden. The other side of the engram: experience-driven changes in neuronal intrinsic excitability. *Nature Reviews Neuroscience*, 4(11):885, 2003.

- [65] Heather K Tittley, Nicolas Brunel, and Christian Hansel. Toward a neurocentric view of learning. *Neuron*, 95(1):19–32, 2017.
- [66] Luis Carrillo-Reid, Weijian Yang, Yuki Bando, Darcy S Peterka, and Rafael Yuste. Imprinting and recalling cortical ensembles. *Science*, 353(6300):691–694, 2016.
- [67] John E Lisman and Marco A Idiart. Storage of 7 ± 2 short-term memories in oscillatory subcycles. *Science*, 267(5203):1512–1515, 1995.
- [68] Colin Molter, Naoyuki Sato, and Yoko Yamaguchi. Reactivation of behavioral activity during sharp waves: a computational model for two stage hippocampal dynamics. *Hippocampus*, 17(3):201–209, 2007.
- [69] Alan Veliz-Cuba, Harel Z Shouval, Krešimir Josić, and Zachary P Kilpatrick. Networks that learn the precise timing of event sequences. *Journal of computational neuroscience*, 39(3):235–254, 2015.
- [70] Tatsuya Haga and Tomoki Fukai. Recurrent network model for learning goal-directed sequences through reverse replay. *Elife*, 7:e34171, 2018.
- [71] Nikolay Cherkov, Henning Sprekeler, and Richard Kempter. Memory replay in balanced recurrent networks. *PLoS computational biology*, 13(1):e1005359, 2017.
- [72] Carina Curto, Anda Degeratu, and Vladimir Itskov. Flexible memory networks. *Bulletin of mathematical biology*, 74(3):590–614, 2012.
- [73] Pierre Yger, Marcel Stimberg, and Romain Brette. Fast learning with weak synaptic plasticity. *Journal of Neuroscience*, 35(39):13351–13362, 2015.
- [74] Xiao-Jing Wang and György Buzsáki. Gamma oscillation by synaptic inhibition in a hippocampal interneuronal network model. *Journal of neuroscience*, 16(20):6402–6413, 1996.
- [75] Nicolas Brunel. Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *Journal of computational neuroscience*, 8(3):183–208, 2000.
- [76] S Leutgeb, KE Ragozzino, and SJY Mizumori. Convergence of head direction and place information in the ca1 region of hippocampus. *Neuroscience*, 100(1):11–19, 2000.
- [77] BL McNaughton, Carol A Barnes, and J O’keefe. The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. *Experimental Brain Research*, 52(1):41–49, 1983.

- [78] John O'keefe and Neil Burgess. Dual phase and rate coding in hippocampal place cells: theoretical significance and relationship to entorhinal grid cells. *Hippocampus*, 15(7):853–866, 2005.
- [79] Jun Yamamoto and Susumu Tonegawa. Direct medial entorhinal cortex input to hippocampal ca1 is crucial for extended quiet awake replay. *Neuron*, 96(1):217–227, 2017.
- [80] David Hattan, Edmund Nesti, Teresa G Cachero, and Anthony D Morielli. Tyrosine phosphorylation of kv1. 2 modulates its interaction with the actin-binding protein cortactin. *Journal of Biological Chemistry*, 277(41):38596–38606, 2002.
- [81] Ryan M Yoder and Kevin CH Pang. Involvement of gabaergic and cholinergic medial septal neurons in hippocampal theta rhythm. *Hippocampus*, 15(3):381–392, 2005.
- [82] Omar Mamad, Harold M McNamara, Richard B Reilly, and Marian Tsanov. Medial septum regulates the hippocampal spatial representation. *Frontiers in behavioral neuroscience*, 9:166, 2015.
- [83] Jill K Leutgeb, Stefan Leutgeb, May-Britt Moser, and Edvard I Moser. Pattern separation in the dentate gyrus and ca3 of the hippocampus. *science*, 315(5814):961–966, 2007.
- [84] Arnold Bakker, C Brock Kirwan, Michael Miller, and Craig EL Stark. Pattern separation in the human hippocampal ca3 and dentate gyrus. *Science*, 319(5870):1640–1642, 2008.
- [85] Michael M Yartsev and Nachum Ulanovsky. Representation of three-dimensional space in the hippocampus of flying bats. *Science*, 340(6130):367–372, 2013.
- [86] Kartik K Sreenivasan, Clayton E Curtis, and Mark D'Esposito. Revisiting the role of persistent neural activity during working memory. *Trends in cognitive sciences*, 18(2):82–89, 2014.
- [87] Evan D Remington, Devika Narain, Eghbal A Hosseini, and Mehrdad Jazayeri. Flexible sensorimotor computations through rapid reconfiguration of cortical dynamics. *Neuron*, 98(5):1005–1019, 2018.
- [88] Morten Raastad and Gordon MG Shepherd. Single-axon action potentials in the rat hippocampal cortex. *The Journal of physiology*, 548(3):745–752, 2003.
- [89] Trygve Solstad, Hosam N Yousif, and Terrence J Sejnowski. Place cell rate remapping by ca3 recurrent collaterals. *PLoS computational biology*, 10(6):e1003648, 2014.

4. A simple generative model of the mouse mesoscale connectome

This chapter has been previously published under the Creative Commons Attribution (CC BY) license and is available via the following citation:

Henriksen, Sid, Rich Pang, and Mark Wronkiewicz. “A simple generative model of the mouse mesoscale connectome.” *eLife* 5 (2016): e12366. (DOI: 10.7554/eLife.12366)

This project was conceived, executed, and the final report written collaboratively with Mark Wronkiewicz and Sid Henriksen. I developed our model’s core “source-growth” rule, and all other project components were divided equally among us.

4.1 Abstract

Recent technological advances now allow for the collection of vast data sets detailing the intricate neural connectivity patterns of various organisms. Oh et al. (2014) recently published the most complete description of the mouse mesoscale connectome acquired to date. Here we give an in-depth characterization of this connectome and propose a generative network model which utilizes two elemental organizational principles: proximal attachment – outgoing connections are more likely to attach to nearby nodes than to distant ones, and source growth – nodes with many outgoing connections are likely to form new outgoing connections. We show that this model captures essential principles governing network organization at the mesoscale level in the mouse brain and is consistent with biologically plausible developmental processes.

4.2 Introduction

The network of physical connections among neurons in the brain provides the medium for neural communication. Investigations of these anatomical networks are typically categorized as macro-, meso-, or microscale, depending on the spatial resolution of the techniques used. The mesoscale, which describes the connectivity among local populations (hundreds to thousands) of neurons, is an attractive intermediate between the two more extreme scales: it has higher granularity than macroscale data, which details connectivity between large anatomically defined brain areas, but it has a broader lens than the microscale, which is concerned with synaptic level connections, often in relatively small volumes

of tissue (see [1] for review). These properties make the mesoscale tractable enough for whole-brain (i.e., connectomic) studies with current technological and analytical tools.

The Allen Institute for Brain Science recently constructed a mesoscale connectome for the mouse (the Allen Mouse Brain Connectivity Atlas), which was the first complete connectivity dataset of a mammalian brain at the mesoscale [2]. Using injections of an anterograde fluorescent viral tracer and serial two-photon microscopy, Oh et al. [2] comprehensively mapped both intra- and interhemispheric axonal tracts and estimated the directed connectivity structure among 213 non-overlapping anatomical regions. The authors also conducted a preliminary graph theoretic analysis and showed that basic network properties of the mouse connectome could not be explained by any one standard network model [2]. Rubinov et al. [3] extended this analysis by identifying small-worldness, a hierarchical modular structure, and non-optimal wiring in the connectome [3]. As the mouse is one of the most pervasive model organisms in biomedical science, a deeper characterization of its connectome is likely to provide pertinent groundwork for future studies of brain development and function as well as yield insights into the broader organizational principles of the mammalian brain.

Graph theory provides a mathematical framework for investigating the organization of networks and has been increasingly applied in neuroscience over the last 15 years. Graphs are mathematical objects that consist of nodes and connections between the nodes, called edges [4, 5]. Edges can be either directed or undirected, as well as binary or weighted. In the Allen mouse connectome, nodes and edges correspond to brain regions and axonal tracts, respectively. Conventionally, real world networks (such as the World Wide Web or social networks) are compared to binary undirected graphs such as small-world [6] and scale-free graphs [7]. These models have been critical for understanding the conditions under which various network properties, such as small-worldness (high clustering among nodes combined with short average path lengths between node pairs) and scale-freeness (defined by power-law degree distributions), emerge in the brain and other systems.

Modeling brain networks with these standard graphs, however, requires some limiting simplifications. First, because they are not embedded in physical space, these networks ignore the biological costs of constructing physical fiber tracts, as well as spatial constraints imposed by the surrounding tissue. Second, small-world and scale-free models ignore directionality, intrinsically discarding information about differences in incoming and outgoing connection patterns. Both of these properties are crucial considerations when modeling the brain [8, 9, 10, 11]. Third, these models were either not developed with neural data [7] or used data from simple model organisms, such as *C. elegans*, at a

particular spatial scale [6], which constrains their relevance to understanding connectomes of other organisms. For these reasons, the field has recently turned to exploring non-standard network models to elucidate generative principles of real brain networks.

Recent work in developing plausible generative network models for the brain has primarily addressed spatial embedding [9, 10, 11, 12, 13]. While the exact approaches differ in their implementation and scale of the networks being modeled, a common theme is that a node is more likely to connect to nearby nodes than distal ones. This organizational principle has been able capture a range of properties observed in cortex, including the distribution of connection lengths [9, 10, 11], the inverse relationship between degree and clustering coefficient [6,9,14,15], and the relative frequency of three-node motifs [11]. Additional generative rules have been explored by Klimm et al. [12], and while the resulting models have captured many properties of cortical networks, the authors note that these rules likely do not reflect the underlying generative principles of cortical networks. Similarly, Betzel et al., [14], recently reported that individual human macroscale connectomes are well-fitted by generative network models which use both spatial proximity and homophilic attraction (i.e. nodes with similar graph theoretic properties are more likely to form connections). However, the homophilic rules employed by Betzel et al. also do not lend themselves to straightforward biophysical interpretations. Indeed, the difficulty of developing biophysically interpretable rules is a recurring challenge in generative network models [6, 9, 12, 14, 15].

Here, we provide an in-depth analysis of the mouse connectome’s properties and use the findings to develop a generative network model of the mesoscale connectome. We characterized the directed and undirected degree distributions, clustering coefficient distribution, reciprocity, global efficiency, physical edge length distribution, nodal efficiency, and the characteristic path length of the connectome (see Table 2 for definitions or [4, 5] for review). Informed by these data, we developed a spatially embedded directed network model. This model uses two simple generative principles: proximal attachment – outgoing connections are more likely to attach to nearby nodes than distal ones, and source growth – nodes with many outgoing connections are more likely to develop new outgoing connections. We show that this simple model, parameterized only by a length constant and the number of nodes and edges, can capture directed, undirected, and spatial properties of the mouse connectome. This work supports the existing literature on the importance of spatial embedding and provides strong evidence that source growth is a major phenomenological rule that shapes connectivity patterns in the mouse brain. Lastly, we propose biological mechanisms that might account for these two generative principles.

4.3 Results

We analyzed the Allen Mouse Connectivity Atlas [2], which is the most comprehensive mesoscale connectome collected to date. We used the linear model from Oh et al. [2] to build an adjacency matrix containing connections between 213 symmetric pairs of nodes (426 total) and 8820 directed edges (7804 undirected).

4.3.1 Comparison of the connectome to undirected graph models

We first compared the undirected structure of the mouse connectome with that of three well-characterized standard graphs commonly used in the literature: a degree-controlled random network [16], a small world network [6], and a scale-free network [7]. The mouse connectome is characterized by a degree distribution with many low-degree nodes and a long tail of high-degree nodes (Figure 1a; [2]). The degree distribution was not well replicated by any standard graph (Figure 1a and b), nor was the clustering coefficient distribution – a finding also shown in Oh et al. [2]. Although the scale-free network’s degree distribution most closely resembles that of the connectome, by construction, it cannot capture the distribution of low degree nodes (as all nodes are instantiated with a minimum number of edges).

The structural discrepancy between the mouse connectome and standard networks was further exposed by examining the relationship between degree and clustering coefficient (Figure 1c-f). In the connectome, nodes with lower degree tend to be more clustered. This relationship was well-fitted with a power law: $C_i \propto k_i^{-0.44}$, where C_i and k_i denote clustering coefficient and degree, respectively, and with the best-fit = -0.44. Previous studies have shown that such a relationship (with varying < 0) is common to many real-world networks [17], including the human connectome [12]. While the small-world network exhibits a clustering coefficient distribution similar to that of the connectome [2], as well as an inverse relationship between clustering coefficient and degree, these similarities are superficial: inspecting Figure 1e reveals that the small-world network’s degree distribution is much more homogeneous. While the scale-free network partly captures the connectome’s degree distribution, its nodes have a much lower clustering coefficient, and there is a weaker relation between these two metrics (Figure 1f). The degree-controlled random model shows that shuffling the connectome’s edges yields a graph resembling the scale-free network (Figure 1d). Thus, these graphs fail to capture key aspects of the connectome’s undirected structure.




Fig 1: Standard graph models vary in their ability to recreate the mouse connectome's degree distribution and relationship between degree and clustering coefficient. *a-b)* Probability density of degree distributions for the mouse connectome, and an average over 100 repeats of scale-free, small-world, and (degree-controlled) random networks plotted with *a)* linear and *b)* log-linear scales. *c-f)* Clustering coefficient as a function of degree for each node in *c)* the mouse connectome, *d)* a degree-controlled random network, *e)* a small-world network, and *f)* a scale-free network. Each plot shows data from 426 nodes and the best fitting power law function (dashed line). *a)* is similar to Figure 3c from Oh et al. [2].

Previous modeling efforts have shown that several properties of brain networks can be captured by simple models which assume that spatially nearby nodes are more likely to connect than distal ones [9, 10, 11, 13]. Therefore, we explored an undirected spatially embedded network model. We first assigned to each node a spatial position randomly sampled from a 7mm x 7mm x 7mm cube, which gave an inter-nodal distance distribution similar to the connectome (Figure 5-figure supplement 1). Edges were then added between pairs of nodes i and j by choosing node i at random and node j with probability $P_{ij} \exp(-d_{ij}/L)$. That is, the probability of choosing a target node j decayed with the Euclidean distance d_{ij} between nodes i and j according to length constant L (see Methods). We call this rule proximal attachment (PA).

Relative to the connectome, the purely geometric model exhibited relatively narrow Gaussian-like degree distributions across several values of L (Figure 2a). Thus, the PA rule fails to generate low- and high-degree nodes. This geometric model does, however, exhibit an inverse relationship between degree and clustering (Figure 2b) as in the connectome. However, Figure 2b demonstrates that there is no value of L for which this model adequately captures the joint degree-clustering distribution seen in the connectome; the PA model misses both low-degree nodes with high clustering and high-degree nodes with low clustering. This suggests that the purely geometric rules explored in previous studies [9, 10, 11, 13] are not sufficient to recreate the mouse connectome’s properties. Therefore, we explored the hypothesis that topological rules, which are based on properties of individual nodes, play an important role in forming the connectivity patterns of the mouse connectome.

4.3.2 The directed perspective and novel random graphs

Conventionally, topological rules in generative network models have been applied to undirected networks. However, the mouse connectome contains directed edges, allowing us to probe whether the directionality of connections in the mouse connectome plays a role in shaping connectivity patterns. In a directed graph, edges point from source nodes to target nodes, so one can consider each node’s in-degree and out-degree (i.e., the total number of incoming and outgoing connections for a node, respectively). We found a surprisingly asymmetric relationship between these distributions in the connectome – while the in-degree of the network was approximately normally distributed, the out-degree distribution exhibited a peak near zero and a long tail (Figure 3a), much like an exponential distribution. There was no significant correlation between in- and out-degree ($r = 0.127$, $P = 0.065$, Spearman rank correlation). Figure 3b shows the proportion of incoming edges as a function of total




Fig 2: Example networks generated using the (purely geometric) proximal attachment (PA) rule where connection probability between two nodes decreased with distance. *a) The degree distribution for networks grown with three values of L (in mm), each averaged across 100 repeats. The mouse connectome is shown for comparison. b) Clustering coefficient as a function of degree for representative networks grown with the same values of L used in a).*

degree for all nodes in the mouse connectome. As the total degree of a node increases, the proportion of incoming edges decreases.

To model the directed mouse connectome, we extended the purely geometric PA model by considering two mathematically symmetrical, yet phenomenologically distinct frameworks: target attraction and source growth (TAPA and SGPA). When adding an edge in the TAPA model, a target node is chosen with a probability proportional to its in-degree, while the source node is chosen with a probability that decreases with its distance from the target node (Figure 4, top). In the SGPA model, a source node is chosen with a probability proportional to its out-degree, while the target node is chosen with a probability that decreases with its distance from the source node (Figure 4, bottom). The TAPA and SGPA models can be considered directed spatial variants of the preferential attachment algorithm introduced by Barabasi Albert (1999) and lead to “rich-get-richer” patterns of either incoming or outgoing connection formation. In both cases, all nodes were initialized upon network instantiation (unlike in the preferential attachment algorithm where nodes are added iteratively), and directed edges were added iteratively until the number of edges matched the mouse connectome. Since these models also incorporated the geometric PA property, both models were parameterized by a single free parameter L exactly as in the PA model.

4.3.3 Comparison of the connectome to the direct graph models




Fig 3: Directed analysis of the mouse connectome reveals different distributions for in- and out-degree. a) Out-degree as a function of in-degree for all nodes in the mouse connectome. Margins show in- and out-degree distributions with the blue lines and axis labels corresponding to a logarithmic scale. In-degree is approximately normally distributed while the out-degree approximately follows an exponential distribution. b) Proportion in-degree as a function of total degree (i.e. in-degree divided by the sum of incoming and outgoing edges). Low degree nodes tend to have mostly incoming edges, while high degree nodes are characterized by mostly outgoing edges.

The SGPA model exhibits an in-degree distribution that is approximately normal and an out-degree distribution that is approximately exponential (Figure 5a, cyan). This matches the mouse connectome (Figure 3a) but is exactly opposite for the TAPA model (Figure 5a, red). Figure 5b shows proportion in-degree as a function of total degree (in-degree + out-degree) for both network models. In the SGPA model (and the mouse connectome; Figure 3b) high-degree nodes tend to have a large proportion of outgoing connections. Again, this is opposite of the TAPA model, where high-degree nodes tend to have a large proportion of incoming connections. These results suggest that the directed graph theoretic properties of the mouse connectome are best captured by a source growth model of network generation.

Previous work has shown that both the reciprocity coefficient and average clustering coefficient of the brain are well above chance [2, 5, 18, 19]. We found that the mouse connectome has a reciprocity coefficient of 0.13, which is high compared to that expected from chance (0.03). Reciprocal edges in the connectome are on average shorter than nonreciprocal edges (Figure 5c; $t(8818) = 13.25$, $P < 10^{-10}$, independent samples t-test; Cohen's $d = 0.47$). As shown in Figure 5e, the reciprocity coefficient of our model networks decreases with increasing length constant, approximately matching that of the mouse connectome (0.13) when $L = 0.725$ mm for both the SGPA and TAPA models. However,




Fig 4: Target attraction and source growth network generation algorithms. The numbers inside each node indicate the probability of growth or attachment. For illustration, the most probable node (dashed) is selected in both diagrams (T and S, corresponding to target and source, respectively). Top row: target attraction proximal attachment (TAPA) model. A target node is selected with a probability proportional to its in-degree (left), while the source node is chosen with a probability that decreases exponentially with the node's Euclidean distance from the target (right). Two nodes have zero probability of forming an edge since the target is already receiving projections from these nodes. The dashed red line shows the most probable edge. Bottom row: source growth proximal attachment (SGPA) model. A source node is selected with a probability proportional to its out-degree (left), while the target node is chosen with a probability that decreases exponentially with the node's Euclidean distance from the source (right). The dashed cyan line shows the most probable edge. In both algorithms, we assume that each node begins with a self-connection (corresponding to an outgoing and incoming edge) to avoid zero-valued probabilities, though self-connections are not shown here or included when calculating any metrics.




Fig 5: Directed analysis of single representative TAPA and SGPA network models and reciprocity comparison with the connectome. *a)* Out-degree as a function of in-degree for both algorithms with $L = 0.725$ mm, which was chosen to match the connectome's reciprocity – see e). Margins show in- and out-degree distributions. *b)* Proportion in-degree as a function of total degree for both algorithms. The SGPA model qualitatively captures the connectome's directed degree distributions and proportion in-degree (cf. Figure 2). *c)* Edge length distribution for the connectome, shown for both reciprocal (blue) and non-reciprocal edges (black). *d)* Same as c) but for the SGPA model. *e)* Reciprocity coefficient as a function of the length parameter for both TAPA (red) and SGPA (cyan), with shading indicating standard deviation over 100 repeats. Both models intersect the connectome at $L = 0.725$ mm. For reference, the reciprocity coefficient for the connectome (magenta) and a corresponding degree-controlled random graph (gold) are also shown. SGPA and TAPA models overlap.

fitting the length constant to reciprocity (e.g., in the SGPA model) underestimates the physical edge lengths of the connectome (compare Figure 5c to 5d). Using a larger length constant improved the fit to the reciprocal and nonreciprocal edge length distributions (not shown), but reduced the model’s reciprocity.

Incorporating source growth into the PA model increases the width of the degree distribution, allowing the model to also capture the joint clustering-degree distribution seen in the connectome (Figure 6a). A network grown solely with the PA rule and a random degree-controlled network both showed a joint clustering-degree distribution that differed from the connectome’s, again suggesting that both topological and geometric rules are important when modeling the connectome. The inset in Figure 6a quantifies this by showing that the SGPA model’s fitted power law relationship between clustering coefficient and degree is the most similar to the connectome. Additionally, this increased clustering is spatially localized: In the connectome, there is a negative correlation between a node’s clustering coefficient and its mean edge length. That is, nodes with short average edge lengths tend to be highly clustered ($=-0.238$, $P < 10^{-6}$, Spearman rank correlation). A similar, but stronger relationship occurs in the SGPA model (median = -0.482 , $P < 10^{-10}$, Spearman rank correlation).

4.3.4 Higher-order connectivity statistics

Our previous analyses of degree, in-degree, out-degree, and clustering coefficient distributions describe the connectivity patterns of a node in the context of its immediate neighbors. To examine the role played by each node in the context of the entire network, we calculated distributions of nodal efficiency. A node’s nodal efficiency is defined as the mean inverse directed shortest path length between itself and all other nodes in the network and quantifies the ease with which that node can theoretically transmit information to all other nodes (see Methods). As shown in Figure 6b, there was a close match between the nodal efficiency distributions for the connectome and the SGPA model (with $L = 0.725$ mm): both exhibited an approximately normal distribution, save for a small selection of nodes with zero-valued nodal efficiency (corresponding to nodes with no outgoing connections). A degree-controlled random graph also showed approximately normally distributed nodal efficiency distributions, but the mean of the distribution was slightly higher than either the SGPA model or the mouse connectome. In contrast, the purely geometric PA model showed a sharply peaked nodal efficiency distribution, similar to what one would expect for a directed Erdos-Renyi graph (Figure 5-figure supplement 1, Figure 6-figure supplement 1). This analysis suggests that the SGPA model




Fig 6: Clustering and nodal efficiency for the connectome and other directed models. *a)* Clustering-degree joint distribution for connectome and one representative instantiation of each model graph. Inset: best-fit power-law exponent (see Methods) for the connectome ($\gamma = -0.44$, $R^2=0.58$) and random (mean std. = -0.17 ± 0.02 , median $R^2=0.16$), SGPA (mean std. = -0.48 ± 0.04 , median $R^2=0.59$), and geometric PA (mean std. = -0.80 ± 0.06 , median $R^2=0.42$) models. *b)* Distributions of nodal efficiencies (see Methods) with mean standard deviation (line and shaded regions, respectively) for model networks. Mean standard deviation of mean nodal efficiency (averaged over nodes) is 0.409 ± 0.001 for the random model, 0.393 ± 0.004 for the SGPA model, and 0.426 ± 0.001 for the pure geometric model. The connectome's average nodal efficiency is 0.375 . The PA model's histogram peaks at 279 at a nodal efficiency of 0.44 . *b)* and the inset in *a)* both used 100 sample instantiations of each model.

captures the statistical connectivity patterns of individual nodes in relation to the whole network.

Prior work has shown that a network's resilience to the removal of nodes can provide insight into its structural composition [20, 21]. These studies typically explore the structure of undirected networks, so for our analyses we converted the directed SGPA model to an undirected one by ignoring directionality of edges (see Methods). We then simulated a lesioning (or percolation) process to compare the undirected SGPA model and other standard undirected models to the mouse connectome. When nodes (and the edges connected to them) were removed in order of decreasing degree (i.e., targeted attack), the SGPA model's global efficiency (akin to nodal efficiency; see Methods) and largest (giant) component size both decreased in a manner more similar to the mouse connectome than any standard graph (Figure 7). However, Figure 7a shows that the mouse connectome disintegrates the fastest, and Figure 7b shows that the global efficiency falls more rapidly in the connectome than any model. Thus, the connectome appears more vulnerable to targeted attack than the model networks explored here.

The density of edges can influence some metrics (like global efficiency). To test if our lesioning results were a product of connection density (which is itself determined by degree distribution), rather than connectivity patterns specific to each model, we also investigated the connection density throughout the lesioning process (Figure 7-figure supplement 1a). We found that, throughout the lesioning




Fig 7: Response of undirected networks to targeted lesions where nodes are removed in order of highest degree. Mean values (lines) standard deviations (shaded regions) are plotted for each model after 100 repeats. a) Size of the largest (giant) connected component in response to targeted attack. The randomly shuffled connectome (Random) is obscured by the scale-free graph. b) Global efficiency for the mouse connectome and network models in response to targeted attack.

Model	Clustering coefficient	Characteristic Path Length	Global Efficiency
Connectome	0.361	2.226	0.492
Random	0.140 ± 0.0013	2.002 ± 0.0020	0.528 ± 0.0003
Small-world	0.359 ± 0.0058	2.129 ± 0.0058	0.507 ± 0.0010
Scale-free	0.159 ± 0.0037	1.998 ± 0.0026	0.527 ± 0.0004
SGPA	0.343 ± 0.0094	2.166 ± 0.0159	0.501 ± 0.0027

Table 1: Average undirected metrics for the connectome, standard, and model networks. For the random, small-world, scale-free, and SGPA models, the standard deviation for 100 repeats is also shown. See Table 2 for metric definitions. Note that properties for the original (single-hemisphere) connectome are presented in Oh, et al. 2014.

process, the connectome’s connection density was similar to that of all models except the small world network (and necessarily matched the degree-controlled random graph). We also found that while global efficiency and connection density were correlated, their specific relationship depended on the model (Figure 7-figure supplement 1b). Interestingly, the way these two variables were related in the connectome most resembled how they were related in the SGPA model, relative to all the others. This shows that the results in Figure 7 are not simply due to differences in degree distribution and the dependence of global efficiency on connection density. Therefore, we propose that these lesioning results expose differences in higher-order connectivity structure between the connectome and the examined models.

4.4 Discussion

The Allen Mouse Brain Connectivity Atlas [2] provides a unique view into the mesoscale structure of the mammalian brain. Through anatomical tracing experiments using genetically identical mice registered to a common reference frame, this dataset permitted a graph theoretic analysis of whole-brain connectivity. We first showed that the connectome’s lower degree nodes are more clustered – a relationship absent in the standard random graphs used in the literature. We then demonstrated that in-degree is approximately normally distributed and out-degree is approximately exponentially distributed in the connectome; because in- and out-degree were uncorrelated, this meant that nodes with more total connections tended to have a higher percentage of outgoing connections. Additionally, we found that reciprocal edges were shorter than nonreciprocal edges on average and that the proportion of reciprocal edges was substantially higher in the connectome than that expected by chance. Finally, we developed a directed generative network model based on two simple rules: source growth and proximal attachment. This directed model captures many (but not all) spatial and graph theoretic properties of the mouse connectome and provides important biological insights into the organizational principles governing neural connectivity at the mesoscale, which we discuss later.

4.4.1 The importance of directed analysis

The fact that in- and out-degree were differentially distributed in the connectome highlights an important limitation of undirected graphs: they do not discriminate between in- and out-degree distributions and may therefore fail to reveal key connectivity properties arising from such directed structure. Interestingly, the in- and out-degree relationships observed at other scales and in other organisms differ from our findings here. For instance, the *C. elegans* connectome (for the full nervous system) [22] has nearly identical in- and out-degree (exponential) distributions [23]. While numerous studies have found common properties across scales (e.g., small-worldness), more pronounced differences in network structure, such as different in- and out-degree distributions, suggest that a single set of generative rules cannot capture the brain’s network structure across different scales and/or organisms.

4.4.2 Indications of functional segregation and integration in the connectome

Functional segregation and integration comprise two components of a framework commonly used to interpret brain network architecture [5, 24]. Functional segregation asserts that the brain carries

out specialized computations in anatomically localized and highly interconnected regions. As high clustering coefficient is thought to indicate potential for participation in this sort of computational unit [5], our finding that nodes with a small number of neighbors tend to be highly clustered suggests that these nodes are candidates for specialized processing. However, functional data would be required to validate this conjecture since high clustering coefficient does not necessarily imply grouping of nodes into specialized clusters. Functional integration predicts that some high-degree “hub” brain regions consolidate the results of these specialized computations for higher function (e.g., as in behavior and perception) [25]. Surprisingly, connections associated with such hub nodes were primarily outgoing in the directed connectome (Figure 3b). Assuming again that anatomical connectivity is indicative of functional relationships, our results suggest that hubs play a stronger role in distribution, rather than integration, of information. However, these ideas are not necessarily contradictory – it is possible that many brain regions are involved in both integration and distribution to varying extents. In fact, the proportion in-degree distribution roughly resembles a uniform distribution (Figure 3-figure supplement 1) suggesting a continuum of these properties rather than discrete categories.

Above-chance reciprocity, as we found in the mouse connectome, is also hypothesized to be important for functional integration and segregation [24]. Previous work has suggested that bidirectional connections between brain regions constrain the dynamics in such a way that a balance arises between the two [26, 27]. Considering the propensity for feedback within the brain, however, perhaps the relatively high reciprocity is not surprising. Like the asymmetric degree distributions, the connectome’s high reciprocity also highlights the importance of using directed graph models (as opposed to undirected ones) when analyzing brain networks.

4.4.3 Organizational principles of the mouse connectome

The generative rules explored here follow from a series of recent publications attempting to develop generative network models of animal connectomes [9, 10, 11, 12, 14, 28]. Several papers have documented the importance of proximal attachment in generative network models of animal connectomes [9, 10, 11, 12, 13]. For example, Kaiser, Hilgetag, van Ooyen [10] noted that if axonal outgrowth occurs in a straight line and the axon attaches to the first node it physically encounters, then the probability of connecting to a target neuron depends exponentially on the target’s distance from the source neuron [10]. Previous research has also shown that some chemical gradients responsible for axon guidance decay exponentially with distance [?, 29] and have been modeled as such [30]. Simi-

larly, Rubinov et al. [3], recently reported that distance-dependent inter-areal connection strengths in the mouse connectome were best captured by a power law. We explored a purely geometric model based on proximal attachment and found that while this model induces clustering between low-degree nodes, it fails to account for the broad degree distribution of the connectome. This motivated the exploration of two mathematically symmetric generative models that included an additional topological rule: target attraction proximal attachment (TAPA) and source growth proximal attachment (SGPA).

Only when incorporating the source-growth rule could we capture the in- and out-degree distributions of the mouse connectome. Note that these distributions arise independently of the PA rule, but other characteristics require it, as discussed previously (see Figure 5-figure supplement 2 for the in- and out-degree distributions of a graph with source-growth but not proximal attachment). To explore whether source selection based specifically on out-degree was a necessary property, we also examined a network grown with source selection probability proportional to the total degree raised to a power. While this model also exhibited many of the connectome’s properties, it makes predictions that are not observed in the connectome (e.g. it predicts that the number of incoming and outgoing connections should be correlated; Figure 5-figure supplement 5); the model itself also requires an additional parameter. This indicates that a source growth rule specifically depending on out-degree is a better candidate mechanism for generating the properties of mouse connectome.

Our results indicate that much of the organizational structure of the mouse connectome is captured by the geometric and topological generative rules employed by the SGPA model. The ability of these simple rules to closely capture the mouse connectome’s network structure raises the possibility that brain organization at the mesoscopic scale does not require precise specification of connectivity (e.g., through genetic or transcriptional factors), but might instead be largely based on a set of relatively simple instructions. Many generative schemes proposed previously, such as the homophilic rules by Betzel et al. ([14]; nodes with similar graph theoretic properties are more likely to connect), or the minimal wiring length networks by Klimm et al. [12], are also based on simple connectivity rules, but these rules are not readily interpretable in terms of biophysical processes. We previously discussed possible biological underpinnings of the geometric PA rule, but the topological SG rule also lends itself to speculation on a set of possible biological mechanisms.

4.4.4 A biophysical interpretation of source growth

The source growth rule might be realized during brain development by the actions of a family of proteins known as neurotrophins, which play a major role in promoting the survival of innervating neurons [31, 32] and the growth and branching of their axons [33]. Found throughout both the central and peripheral nervous systems, these proteins are typically secreted by a target (postsynaptic) neuron at functional synapses, endocytosed by a source (presynaptic) neuron, and retrogradely transported to the source neuron’s soma, where they trigger the above-mentioned processes [31, 34]. Because neurotrophins are available in limited quantities, they are hypothesized to cause competitive interactions among growing neuronal populations [35]. Indeed, such an interaction resembles the way in which a single source node is probabilistically selected for each edge addition in the SGPA model (as opposed to allowing all nodes to generate outgoing connections simultaneously). Most importantly, since populations that maintain many functional outgoing connections will tend to receive more neurotrophins, they should on average be more fit for survival and new connection generation, and their existing connections should be more likely to branch. This would in turn increase their probability of establishing new connections to novel targets. Such a biological mechanism would lead to the “rich-get-richer” phenomenon for outgoing connections that was fundamental in modeling the network properties of the mouse connectome.

4.4.5 A non-power law explanation of the connectome’s undirected degree distribution

A final insight provided by the SGPA model concerns the undirected degree distribution in the mouse connectome. Degree distributions in the brain have often been characterized by power laws, where $p(k) \propto k^{-\gamma}$ for some $\gamma < 0$ [4, 21]. The SGPA model provides an alternative because the in- and out-degree distributions are driven by different rules. In-degree is approximately Gaussian because of the central limit theorem (connections are random in the sense that they only depend on distance and not degree or other topological properties), and out-degree is approximately exponential, owing to the source growth process. If in- and out-degree are independent, as they are in our model and seem to be in the connectome, the distribution of their sum (i.e., of total degree) will be the convolution of the two individual distributions. Thus, we would expect the distribution over undirected degree (which is approximately the sum of in- and out-degree, the approximation arising since we do not count reciprocal edges twice) to approach the convolution of a Gaussian and an exponential, i.e., an exponentially modified Gaussian (Figure 6-figure supplement 2). Additionally, the SGPA model is able to capture the low-degree portion of the distribution (Figure 1b) unlike the scale-free graph. This

serves as a more biologically motivated alternative to the power-law degree distribution often used to describe the brain’s network properties.

4.4.6 Limitations of the SGPA model

While the SGPA model captured a number of the connectome’s characteristics, there are also limitations worth highlighting. Our choice to instantiate all the nodes at the start of the network’s generation is somewhat biologically implausible. This is different from the Barabasi-Albert algorithm used to generate scale-free graphs where, aside from an initial group of nodes, all the nodes are added sequentially over time. However, a version of our SPGA model in which nodes are added one-by-one (see Methods) yielded a qualitatively similar network (see Figure 5-figure supplement 3 and Figure 6-figure supplement 3). The only notable differences were that the youngest nodes had an extremely low degree and a clustering coefficient of either zero or one, and that there was a positive correlation between in- and out-degree. This deviation from the properties observed in the mouse connectome could potentially be corrected by “pruning” the final network or by modifying the probability of generating new connections as the final nodes are added.

In our simulations, we chose to randomly draw node positions from a 7 mm cube to spatially embed our model. The dimensions of this cube were chosen to match the inter-nodal distance distribution of the connectome (Figure 5-figure supplement 4). All the simulations we explored, however, yielded comparable results when run with node coordinates determined by the centroids of brain regions in the connectome. One interesting feature of the SGPA model is that it predicts that edge lengths should be significantly shorter than those observed in the connectome. This is a consequence of fitting the length constant to the connectome’s reciprocity. A more satisfactory fit to the edge length distribution can be obtained by using a larger length constant, but this substantially reduces the reciprocity and degrades the inverse relationship observed between clustering and degree (Figure 6a). Equivalently, we can scale the dimensions of the model, which allows us to match the reciprocity and the edge length distribution but not the inter-nodal distance distribution. Thus, our model is unable to account for both the edge length distribution and the reciprocity of the connectome while also maintaining the appropriate dimensionality.

Within the lesioning analysis, we found that the connectome is more susceptible to targeted attack than our generative SGPA model network (or any other model). This suggests that the mouse connectome is less resilient than the synthetic networks explored here. More broadly, addressing whether the

SGPA model can reproduce the macroscale hierarchical modularity recently reported by Rubinov et al. [3] is also a target for future research. Previous work on targeted attacks on the macaque and cat macroscale connectome has shown different patterns of results [21], but it is difficult to assess whether this is due to differences in the organism (mouse vs. monkey or cat), scale (meso vs. macro), or both.

Lastly, we treated the entire connectome as a network grown with homogeneous growth rules. The success of our approach lends some merit to this assumption, but it is nevertheless likely that cortico-subcortical connections follow different generative rules than cortico-cortical ones, for example. A similar point was raised by Kaiser Varier [19] and Oh et al. [2] – both noted that cortico-cortical projections have higher reciprocity than cortico-subcortical ones in macaque macroscale and mouse mesoscale connectomes, respectively. Uncovering the differences in generative rules employed by sub-networks in the brain, as well as those at different scales or time points, is a target for future research. Dynamically changing the physical scale of the model (e.g., to simulate physical growth) is another promising avenue for future work which some have begun to explore [36], and may account for our inability to capture the edge length distribution in the connectome. Here, our network was embedded within a physical space that maintained a constant size. We did find that a node’s age affected its final properties (Figure 5-figure supplement 3, Figure 6-figure supplement 3), but the growth of neural tissue typically occurs in an expanding physical space which may stretch or alter the oldest connections as development proceeds. Continued use of real data to both inspire and evaluate network models will be crucial for elucidating the principles that govern the network organization of connectomes.

4.5 Conclusion

We have characterized the network properties of the mouse mesoscale connectome, a system that highlights the importance of using spatial, directed graphs when modeling brain networks. A model that uses two simple organizational principles – source growth and proximal attachment – can capture a large number of directed, undirected and spatial network properties of the mouse connectome. Importantly, these rules have biologically plausible connections to developmental mechanisms and wiring properties in real brains. This model not only serves as a simple mathematical tool that can be used to model and understand mesoscopic brain organization, but also provides a parsimonious framework for informing future investigations of brain network formation.

4.6 Materials and Methods

All code used in this analysis is available at <https://github.com/neofunkatron/neofunkatron>.

4.6.1 Empirical connectivity graph

To generate a connectivity graph from empirical data, we used the mouse connectome developed by the Allen Brain Institute. To compile this dataset, anterograde tracers were injected in a single hemisphere and the authors traced projections into both hemispheres. Thus, the connectome probed connections originating in the right hemisphere and terminating in either the right or left hemisphere. The original work computed projection strengths and associated p-values (the probability that the observed projection would arise by chance) between 213 pairs of ipsi- and contralateral anatomically defined brain regions [2]. To generate a binary adjacency matrix, we set all connection strengths with $p < 0.01$ to one and all other elements to zero. This yielded a graph with a reasonable connection density for analysis. We then assumed that all connectivity projections were bilaterally symmetric in order to construct a 426×426 binary adjacency matrix, which completely defines a (single component) graph. The results remained qualitatively the same for alternative connection thresholds ($p < 0.05$ and $p < 0.001$).

4.6.2 Random graphs

The following sections describe the random graphs we included in our analysis. For each random graph, the number of nodes was set to 426, and the number of edges to 7804 (for undirected models) or 8820 (for directed models).

4.6.3 Standard random graphs

Standard random graphs were generated in Python using the NetworkX module (<https://networkx.github.io>) [37]. The small-world graph [6] is parameterized by the number of nodes n , the number of initial connections for each node k_{SW} , and the probability of rewiring, p . Here we used $k_{SW} = 18$ to match the mean degree of the connectome, and $p = 0.23$ to approximately match the mean clustering coefficient of the connectome. The scale-free graph [7] is parameterized by the number of nodes n and the number of connections k_{SF} formed by each node as it is added to the network. We used $k_{SF} = 18$ to match the number of edges in the connectome. We generated the undirected degree-controlled random network by shuffling the mouse connectome's edges while holding the degree distribution constant (similar to

Maslov-Sneppen, 2002 [16]); this generates a control graph with random connectivity but identical degree distribution. The directed degree-controlled random was generated in a similar way, except both the in- and out-degree distributions were held nearly fixed (“nearly” because the algorithm often converts a small number of edges to self- or double-connections, which are ignored in our analysis; however, these only represent about 5% of all connections). These standard graphs are similar to those used in Oh et al. [2]. The directed Erdos-Renyi graph was generated using NetworkX, with an edge probability of 0.0487, which on average yielded the number of directed edges present in the mouse connectome.

4.6.4 Custom random graphs

The following sections describe the algorithms used to generate the custom random graphs examined in this study.

4.6.5 Purely geometric models

For the purely geometric proximal attachment (PA) graph, 426 nodes were randomly assigned centroids within a 7mm x 7mm x 7mm cube, with no edges connecting the nodes. We then generated each undirected edge by first selecting a source node i at random and subsequently selecting a target node j with probability $P_{ij} \propto \exp(-d_{ij}/L)$, where d_{ij} denotes the distance between node i and node j . If an undirected edge already existed between the source and target, the probability of selecting that target was set to zero. Edges were added until the number of (undirected) edges matched that of the connectome ($K=7804$). The directed PA algorithm was identical, except that directionality of edges was retained. For the graphs incorporating topological rules (i.e. SGPA and TAPA), the growth algorithm was initialized by instantiating n nodes with no edges connecting them, except for one self-connection per node to prevent zero-valued connection probabilities for nodes with no outgoing or incoming edges (for SGPA and TAPA, respectively). However, these self-connections were ignored when calculating all metrics for the final graph. In all of the graphs we generated (except for node-by-node SGPA) we matched n and K to the number of nodes and directed edges, respectively, in the empirical connectome ($n = 426$, $K = 8820$).

4.6.6 Target-attraction-proximal-attachment (TAPA) random graph

At each step in the growth process an edge is added by (1) selecting the target node j from all nodes without maximum in-degree according to the in-degree of j : $P(\text{target} = j) \propto k_j^{\text{in}}$, and (2) selecting the source node i with a probability that decreases with distance from the target: $P(\text{source} = i | \text{target} = j) \propto \exp(-d_{ij}/L)$. If a connection already existed from node i to j , this probability was set to zero to avoid a duplicate edge. d_{ij} is the Euclidean distance between nodes i and j , and L is a parameter governing the strength of distance-dependence. For the nonspatial TA graph (Figure 5-figure supplement 2), the algorithm was as the TAPA algorithm, except the source was chosen with a uniform probability for all nodes (i.e. no distance-dependence).

4.6.7 Source-growth-proximal-attachment (SGPA) random graph

At each step in the growth process an edge is added by (1) selecting the source node i according to its out-degree as $P(\text{source}=i) \propto k_i^{\text{out}}$ from all nodes that do not have maximal out-degree, and (2) selecting the target node j with a probability that decreases with distance from the source: $P(\text{target} = j | \text{source} = i) \propto \exp(-d_{ij}/L)$. If a connection already existed from node i to j , this probability was set to zero to avoid a duplicate edge. For the nonspatial SG graph (Figure 5-figure supplement 2), the algorithm was as the SGPA algorithm, except the target was chosen with a uniform probability for all nodes (i.e. no distance-dependence). Both the target attraction and source growth rules can be conceptualized as directed spatial variants of the preferential attachment rule introduced by Barabasi and Albert [7].

4.6.8 Total-degree SGPA random graphs

This graph was identical to the SGPA graph mentioned above, except that probability of selecting the source node i was proportional to its total degree (in-degree + out-degree) raised to a power : $P(\text{source} = i) \propto (k_i^{\text{out}} + k_i^{\text{in}})$. Target selection was as in the SGPA model.

4.6.9 Node-by-node SGPA random graph

In this algorithm, nodes are added one at a time with positions sampled uniformly from within a 7mm x 7mm x 7mm cube until the graph contains 426 nodes. Upon each node addition (starting with the second node addition), each existing edge “branches” with a probability $p = 0.016$ (which was chosen so as to yield approximately the same number of edges as in the connectome). When an edge

“branches” a new edge is formed that has the same source node and whose target node is chosen with probability $P(\textit{target} = j | \textit{source} = i) \propto \exp(-d_{ij}/L)$ (i.e., in the same way as in the standard SGPA model). If an edge already exists between the source node and the selected target node, no new edge is added. When all the edges and nodes have been added, any nodes that are not connected to the largest (giant) graph component (typically only three or four, representing $< 1\%$ of the network) are removed from the graph.

4.6.10 Undirected metrics and lesioning of directed graphs

We calculated undirected metrics (e.g., clustering coefficients and undirected degree) and carried out a lesioning study for directed graphs by first casting the directed graph to an undirected one. This yielded a graph which contained an undirected edge between every pair of nodes that had been connected by at least one directed edge in the directed graph.

4.7 Acknowledgments

We thank the Allen Institute for Brain Science for generously providing the dataset used in this analysis and for funding the Summer Workshop on the Dynamic Brain – the course in which this project was conceived. We also thank the directors of this course, Christof Koch and Adrienne Fairhall. Additionally, we are grateful to Marcus Kaiser, Martha Bosma, and the reviewers for helpful comments and suggestions.

4.8 Supplementary information




Fig 3S1: *Distribution of proportion in-degree in the mouse connectome.*




Fig 5S1: *Directed degree distributions and proportion in-degree for a directed Erdos-Renyi graph. a) In- and out-degree are independent and both are approximately normally distributed. b) Proportion in-degree is independent of total degree.*




Fig 5S2: Directed degree distributions and proportion in-degree for a purely topological source-growth or target-attraction directed graph (SG only, TA only, respectively), with no proximal attachment (equivalent to $L = \infty$ in SGPA or TAPA). The qualitative properties of the in- and out-degree distributions are the same as in SGPA and TAPA with $L = 0.725$.




Fig 5S3: Directed degree distributions and proportion in-degree for an SGPA model where the network is grown one node at a time. Hotter colors indicate more recently added nodes. a) Out- and in-degree relationships for one representative model instantiation with $L = 0.725$ mm. b) Proportion in-degree as a function of total degree.




Fig 5S4: *Inter-nodal distance distribution for the mouse brain (magenta bars) and a 7mm³ cube (black line). The use of a 7mm x 7mm x 7mm cube was justified on the basis that its inter-nodal distance distribution closely mimicked that of the connectome. The results did not change considerably when coordinates from the mouse brain were used instead of the 7mm³ cube.*




Fig 5S5: Directed degree distributions and proportion in-degree for a graph in which source selection was proportional to total degree (in-degree + out-degree) in a) and b), and in which source selection was proportional to total degree raised to the power = 1.67 in c) and d). In both a) and c), there are significant correlations between the in- and out-degrees ($=0.45$, $P < 10^{-10}$ and $=0.60$, $P < 10^{-10}$ for a and b, respectively; Spearman rank correlation). In all graphs, $L = 0.725$ mm.




Fig 6S1: Clustering and nodal efficiency for Erdos-Renyi (ER) and TAPA models. a) Clustering-degree joint distribution for connectome and one representative instantiation of each model graph. The inset shows best-fit power-law exponents, as in Figure 6 (ER: mean std. = -0.01 0.04, median $R^2 = 0.00$; TAPA: mean std. = -0.47 0.035, median $R^2 = 0.60$). b) Distributions of nodal efficiencies (see Methods). Solid lines are means and shaded regions are standard deviations. b) and the inset in a) both used 100 sample instantiations of each model. Mean standard deviation (taken over 100 graph instantiations) of mean nodal efficiency (averaged over nodes) is 0.467 0.002 for ER, 0.395 0.003 for TAPA.




Fig 6S2: Undirected degree distribution for the SGPA model and the mouse connectome in a) linear and b) logarithmic scales. For the SGPA model, the line and shaded region represent the mean and standard deviation across 100 graph instantiations.




Fig 6S3: Clustering vs. degree a) and nodal efficiency b) for the node-by-node SGPA network ($L = 0.725$) used in Figure 5-figure supplement 3. Hotter colors indicate more recently added nodes.




Fig 7S1: Network connectivity patterns (and not just connection density) affect global efficiency. a) Connection density as a function of nodes removed (for comparison with Figure 7). Except for the small-world network, the connection density was changed similarly throughout the lesioning process. Mean values (lines) standard deviations (shaded regions) are plotted for each model after 100 repeats. b) Global efficiency and connection density at each point during the lesioning process for representative networks grown in the same manner as Figure 7. Note that global efficiency varies across graphs even at similar connection densities.

Metric	Interpretation	Formal definition
Degree	Number of edges connected to a node.	$k_i = \sum_{j \in N} a_{ij}$, where N is the set of all nodes and a_{ij} is 1 if an edge from node i to j exists and 0 otherwise.
In-degree	Number of incoming edges to a node.	$k_i^{in} = \sum_{j \in N} a_{ji}$
Out-degree	Number of outgoing edges from a node.	$k_i^{out} = \sum_{j \in N} a_{ij}$
Clustering coefficient [6]	Level of connectivity among nearest neighbors of node i .	$\frac{2t_i}{k_i(k_i - 1)}$, where t_i is the number of triangles including node i
Characteristic path length [6]	Mean shortest undirected path length over all pairs of nodes.	$\frac{1}{n(n-1)} \sum_{i,j \in N, i \neq j} d_{ij}$, where n is the number of nodes and d_{ij} is the shortest undirected path from i to j
Global efficiency [40]	Mean inverse shortest path <i>undirected</i> path length over all pairs of nodes.	$\frac{1}{n(n-1)} \sum_{i,j \in N, i \neq j} 1/d_{ij}$.
Nodal efficiency (generalized from [41])	Mean inverse shortest path <i>directed</i> path length over all pairs of nodes.	$\frac{1}{n(n-1)} \sum_{i,j \in N, i \neq j} 1/d_{ij}$, where d_{ij} is the shortest directed path from i to j .
Reciprocity coefficient	Proportion of edges from node i to node j that have a reciprocal connection from node j to node $i \neq j$	$\frac{1}{N_e} \sum_{i,j \in N, j \neq i} a_{ij}a_{ji}$, where N_e is the total number of undirected edges.

Table 2: Graph theoretical metrics used for analysis.

References

- [1] Olaf Sporns, Giulio Tononi, and Rolf Kötter. The human connectome: A structural description of the human brain. *PLoS Computational Biology*, 1(4):e42, 2005.
- [2] Seung Wook Oh, Julie A Harris, Lydia Ng, Brent Winslow, Nicholas Cain, Stefan Mihalas, Quanxin Wang, Chris Lau, Leonard Kuan, Alex M Henry, et al. A mesoscale connectome of the mouse brain. *Nature*, 508(7495):207, 2014.
- [3] Mikail Rubinov, Rolf JF Ypma, Charles Watson, and Edward T Bullmore. Wiring cost and topological participation of the mouse brain connectome. *Proceedings of the National Academy of Sciences*, 112(32):10032–10037, 2015.

- [4] Ed Bullmore and Olaf Sporns. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature Reviews Neuroscience*, 10(3):186, 2009.
- [5] Mikail Rubinov and Olaf Sporns. Complex network measures of brain connectivity: uses and interpretations. *Neuroimage*, 52(3):1059–1069, 2010.
- [6] Duncan J. Watts and Steven H. Strogatz. Collective dynamics of ‘small-world’ networks. *Nature*, 393(6684):440–442, jun 1998.
- [7] Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *science*, 286(5439):509–512, 1999.
- [8] Simon B Laughlin and Terrence J Sejnowski. Communication in neuronal networks. *Science*, 301(5641):1870–1874, 2003.
- [9] H. Francis Song, Henry Kennedy, and Xiao-Jing Wang. Spatial embedding of structural similarity in the cerebral cortex. *Proceedings of the National Academy of Sciences*, 111(46):16580–16585, nov 2014.
- [10] Marcus Kaiser, Claus C Hilgetag, and Arjen Van Ooyen. Random axon outgrowth and synaptic competition generate realistic connection lengths and filling fractions. *BMC Neuroscience*, 10(1):P58, 2009.
- [11] Mária Ercsey-Ravasz, Nikola T Markov, Camille Lamy, David C Van Essen, Kenneth Knoblauch, Zoltán Toroczkai, and Henry Kennedy. A predictive network model of cerebral cortical connectivity based on a distance rule. *Neuron*, 80(1):184–197, 2013.
- [12] Florian Klimm, Danielle S Bassett, Jean M Carlson, and Peter J Mucha. Resolving structural variability in network models and the brain. *PLoS computational biology*, 10(3):e1003491, 2014.
- [13] Marcus Kaiser and Claus C Hilgetag. Spatial growth of real-world networks. *Physical Review E*, 69(3):036103, 2004.
- [14] Richard F Betzel, Andrea Avena-Koenigsberger, Joaquín Goñi, Ye He, Marcel A De Reus, Alessandra Griffa, Petra E Vértés, Bratislav Mišić, Jean-Philippe Thiran, Patric Hagmann, et al. Generative models of the human connectome. *Neuroimage*, 124:1054–1064, 2016.

- [15] P. E. Vertes, A. F. Alexander-Bloch, N. Gogtay, J. N. Giedd, J. L. Rapoport, and E. T. Bullmore. Simple models of human brain functional networks. *Proceedings of the National Academy of Sciences*, 109(15):5868–5873, mar 2012.
- [16] Sergei Maslov and Kim Sneppen. Specificity and stability in topology of protein networks. *Science*, 296(5569):910–913, 2002.
- [17] Erzsébet Ravasz and Albert-László Barabási. Hierarchical organization in complex networks. *Physical Review E*, 67(2):026112, 2003.
- [18] Daniel J Felleman and DC Essen Van. Distributed hierarchical processing in the primate cerebral cortex. *Cerebral cortex (New York, NY: 1991)*, 1(1):1–47, 1991.
- [19] Marcus Kaiser and Sreedevi Varier. Evolution and development of brain networks: from caenorhabditis elegans to homo sapiens. *Network: Computation in Neural Systems*, 22(1-4):143–147, 2011.
- [20] Mark Newman. *Networks*. Oxford university press, 2018.
- [21] Marcus Kaiser, Robert Martin, Peter Andras, and Malcolm P Young. Simulation of robustness against lesions of cortical networks. *European Journal of Neuroscience*, 25(10):3185–3192, 2007.
- [22] J.G. White. Neuronal connectivity in caenorhabditis elegans. *Trends in Neurosciences*, 8:277–283, jan 1985.
- [23] Luis A Nunes Amaral, Antonio Scala, Marc Barthelemy, and H Eugene Stanley. Classes of small-world networks. *Proceedings of the national academy of sciences*, 97(21):11149–11152, 2000.
- [24] G. Tononi, O. Sporns, and G. M. Edelman. A measure for brain complexity: relating functional segregation and integration in the nervous system. *Proceedings of the National Academy of Sciences*, 91(11):5033–5037, may 1994.
- [25] Karl Friston. Functional integration and inference in the brain. *Progress in neurobiology*, 68(2):113–143, 2002.
- [26] Giulio Tononi, Olaf Sporns, and Gerald M. Edelman. Reentry and the problem of integrating multiple cortical areas: Simulation of dynamic integration in the visual system. *Cerebral Cortex*, 2(4):310–335, 1992.

- [27] O. Sporns, G. Tononi, and G. M. Edelman. Modeling perceptual grouping and figure-ground segregation by means of active reentrant connections. *Proceedings of the National Academy of Sciences*, 88(1):129–133, jan 1991.
- [28] Sol Lim and Marcus Kaiser. Developmental time windows for axon growth influence neuronal network topology. *Biological cybernetics*, 109(2):275–286, 2015.
- [29] Carolyn M Isbister, Paul J Mackenzie, Kenneth CW To, and Timothy P O’Connor. Gradient steepness influences the pathfinding decisions of neuronal growth cones in vivo. *Journal of Neuroscience*, 23(1):193–202, 2003.
- [30] Duncan Mortimer, Julia Feldner, Timothy Vaughan, Irina Vetter, Zac Pujic, William J Rosoff, Kevin Burrage, Peter Dayan, Linda J Richards, and Geoffrey J Goodhill. A bayesian model predicts the response of axons to molecular gradients. *Proceedings of the National Academy of Sciences*, 106(25):10296–10301, 2009.
- [31] Eric J Huang and Louis F Reichardt. Neurotrophins: roles in neuronal development and function. *Annual review of neuroscience*, 24(1):677–736, 2001.
- [32] Moses V Chao. Neurotrophins and their receptors: a convergence point for many signalling pathways. *Nature Reviews Neuroscience*, 4(4):299, 2003.
- [33] M. Tessier-Lavigne and C. S. Goodman. The molecular biology of axon guidance. *Science*, 274(5290):1123–1133, nov 1996.
- [34] Sigrun Korsching. The neurotrophic factor concept: a reexamination. *Journal of Neuroscience*, 13(7):2739–2748, 1993.
- [35] Arjen van Ooyen and David J. Willshaw. Competition for neurotrophic factor in the development of nerve connections. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 266(1422):883–892, may 1999.
- [36] Jonathan Ozik, Brian R Hunt, and Edward Ott. Growing networks with geographical attachment preference: Emergence of small worlds. *Physical Review E*, 69(2):026108, 2004.
- [37] Aric Hagberg, Pieter Swart, and Daniel S Chult. Exploring network structure, dynamics, and function using networkx. Technical report, Los Alamos National Lab.(LANL), Los Alamos, NM (United States), 2008.

5. Conclusions and future directions

In the preceding work we have shed new light on neural systems from four different angles. The first three addressed the problem of short-term modulation of dynamics and information processing, and the fourth investigated structural models of mesoscale brain networks. We began within the context of a natural behavior—odor-tracking—utilized by a plethora of organisms to locate food and mates; we statistically analyzed the odor-tracking flight patterns of fruit flies and mosquitoes and compared the identified history dependence in their flight patterns with theoretical models. Next, we investigated a candidate mechanism potentially enabling rapid modulation of neural processing that might underlie history-dependent tracking behaviors; to this end we dissected components of the neuromodulatory dopamine system in the fruit fly mushroom body through experimental collaboration with Vanessa Ruta’s lab, and have discovered that within a handful of dopamine signals there may exist a multiplexed representation of diverse features of the fly’s current sensorimotor and valence experience. To generalize our investigation beyond a specific computation we next considered biophysical mechanisms that might subserve the rapid embedding of complex memories into neural circuits responsible for arbitrary changes in computation and behavior; we showed that a specific outsize plasticity mechanism operating at a cellular, rather a synaptic level, could indeed enable flexible induction of sequences and mappings in a simple recurrent network on a rapid timescale. Finally, to address the highly distributed nature of neural systems we analyzed the structural connectivity patterns of the mouse brain connectome at the mesoscale, in order to refine our understanding of the structural principles that must be taken into account when modeling computations that occur within full-scale brain networks; we found that multiple core mesoscale connectivity statistics could be recreated via two simple network growth rules, suggesting that multi-scale network models could potentially be implemented atop parsimonious architectures, greatly reducing complexity and potentially facilitating mathematical analysis. Below we briefly summarize our key findings from each branch of this overarching investigation, and point to next questions that could form the seed for future research arcs.

5.1 Insect odor tracking

Here we considered the problem of how a search agent should integrate the sparse, stochastic encounters with an odor plume broken up by realistic turbulence to reliably generate an efficient trajectory toward the plume source. We addressed this by analyzing odor-modulated flight trajectories and

comparing our results to theoretical models. We found that beyond known odor-encounter-triggered upwind turns, encounter-by-encounter variability within trajectories also depended on both the concentration experienced during each encounter, as well as the history of encounters that preceded it. Higher concentrations triggered stronger upwind turns, and later encounters triggered weaker upwind turns than earlier encounters. To shed light on what normative strategy might be reflected in this history dependence, we asked whether several previously identified theoretical algorithms produced the same history dependence, finding that only “infotaxis” [1] captured the sign of the history dependence, although it differed substantially in magnitude. Thus, although none of the algorithms we explored precisely explained the variation observed in real flight trajectories, these trajectories were more consistent with the model predicting search based on accumulation of information about a specific source position than with other models.

These results motivate and contextualize several follow-up investigations, both experimental and theoretical. First, whereas our analyses were based on flight trajectories within a stationary plume, under the assumption that the stochasticity of insect flight would yield variation in odor encounters approximately reflecting those within a turbulent plume, new technologies may soon enable simultaneous recording of flight trajectories and odor encounters in a truly turbulent plume. Schlieren imaging, for instance, an optical technique for inferring variations in chemical concentrations in fluids based on changes in refractive indices [2, 3], may prove a useful tool. Indeed, with such recordings made simultaneously with recordings of an insect flight trajectories, one could identify the precise concentration experienced by the insect at each point in its trajectory, thereby providing the accurate estimate of the moment-to-moment olfactory stimulus needed to dissect possible tracking strategies within a more natural scenario. An additional use of Schlieren imaging could be to infer local wind direction, which though generally steady on average in turbulent plumes, exhibits large variation across space due to the higher-order airflow statistics like vortices introduced by turbulence, and which may affect flight and odor processing. An alternative approach, which we have begun to investigate in collaboration with Thierry Emonet’s lab at Yale University (not shown) is to generate chemical plumes that are both attractive to insects and which can be imaged directly, for example a smoke plume containing an attractive odorant. Indeed, new datasets from the Emonet lab contain smoke plumes that drive upwind tracking imaged simultaneously with the walking trajectories of fruit flies in a 2D environment, serving as a natural first step to the more complex case of three-dimensional flight.

A final problem with the above-mentioned experimental setups that has yet to be fully addressed

but which will be crucial in accurately characterizing odor tracking is disentangling odor-driven from visual responses, the latter which are known to heavily influence flight control [4, 5, 6] and which likely play a role in odor tracking [7, 8, 9]. Unfortunately, since vision is required for flying insects to orient relative to the wind, one cannot remove the visual environment from an experiment. One workaround is to perform experiments in much larger wind tunnels with uniform visual statistics (e.g. so the upwind and downwind ends of the wind tunnel appear identical); this will likely require more advanced optical technologies to track small insects like flies and mosquitoes in a large area, but may eventually be feasible. An alternative solution would be to model visual responses directly, e.g. informed by studies on the neural circuitry and computations of fruit fly vision [10, 11, 12, 13, 14, 15, 16, 17], and subsequently use this to project out predictions of flight patterns from vision in order to isolate odor-driven flight dynamics.

Finally, while the theoretical plume-tracking theory “infotaxis” [1] has spawned numerous variations and become incorporated into studies of plume-tracking robotics [18, 19, 20, 21], its specification poses key problems to seeking evidence of the theory in empirical behavioral data. Indeed, to the best of our knowledge ours was the first study to quantitatively seek evidence of infotaxis in plume-tracking data, with a single additional study comparing the predictions of infotaxis to the movement patterns of *C. elegans* [22]. The first core issue is that existing implementations depend heavily on a number of parameters presumably not known *a priori* by the biological tracking agents. For example, “infotaxis” and its variations presuppose, either explicitly or implicitly, an internal model of the plume statistics, which depends on the source emission rate and crucially on the turbulent diffusivity coefficient [1, 23] (this latter governs how much the plume is expected to spread out through time and space) [24]. Besides posing the key scientific question of how an insect could learn these parameters over time, or respond to their natural variation (which can sometimes span orders of magnitude), this poses a key practical challenge. How should one parameterize an infotaxis model when generating specific predictions about empirical trajectory structure? Alternatively, are there predictions generated by the core infotaxis algorithm independent of the parameters used to tune it? Refining such predictions will be crucial to seeking out further evidence of this theory, and in making precise exactly where it converges and diverges with empirical observation.

The second core issue is that “infotaxis” and its variations are active sensing algorithms, which specify full time-varying interactions of a search agent with an intermittent odor source, with behavior and sensation reciprocally influencing one another. Thus, it is not straightforward to directly compare indi-

vidual empirically measured trajectories to those generated by infotaxis, since the latter will generally contain a completely different odor sequence such that moment-to-moment comparisons of trajectory features have little meaning. Instead, it will be crucial in future investigations of odor-tracking theory to identify robust predictions of infotaxis that distinguish it on a multi-trial level from other tracking algorithms, and which don't depend on responses to specific sequences of odor encounters and behaviors. While this observation was what motivated our analysis of the odor-encounter-triggered history dependence, it will be a useful future investigation to specify additional observables that can be identified or falsified by experimental observation.

5.2 Dopaminergic modulation of mushroom body computations

We next considered a candidate neural substrate potentially involved in odor tracking and other natural behaviors dependent on short-term memory: dopaminergic innervation of the insect mushroom body, which controls information propagation from the mushroom body's inputs to its output neurons [25]. To this end, we analyzed recordings of dopamine neuron (DAN) activations made simultaneously with video recordings of a fly's walking behavior in a minimal virtual reality environment. In addition to the expected encoding of positive and negative valences, which has received much study in both mammalian and invertebrate nervous systems, we found that these dopamine neurons additionally tracked moment-to-moment variations in both motor variables, such as walking state and forward walking velocity; and sensory variables, such as wind direction and odor. This suggests that the neuromodulatory signals controlling information flow through the mushroom body, despite being confined to only tens of cells, may themselves be governed by a multiplexed representation of diverse sensorimotor and valence variables, potentially reflecting a highly integrative computation that shapes behavior. Beyond this, our work has revealed a potential context dependence of how the DANs represent behavior, with appetitive odors appearing to cause the DANs to recruit representations not only of current locomotor patterns, but of tracking movements several seconds in the future. This result thus exposes an additional reflection of the how network interactions in the fly olfactory circuit might extend representations and computation beyond the fast timescales of fluctuating neuronal membrane voltages to span the seconds long timescales underlying odor tracking and other behaviors requiring short-term memory.

Our results in the DAN inputs to the mushroom body pose several questions regarding their role within larger, behaviorally relevant computations. First, given what's known about how the DANs gate

synaptic transmission between mushroom body Kenyon cells (KCs) and output neurons (MBONs), how would the naturalistic, sensorimotor-dependent sequences of DAN activations we’ve elucidated in these results in turn shape KC-to-MBON synaptic transmission? Importantly, given that DANs can either strongly potentiate or depress KC-to-MBON synapses, dependent on coincident KC and DAN activation [26], what would be required for the average gating dynamics to remain stable over time? Furthermore, how would natural behaviors modulate the representations of odor identity known to be encoded in KCs [27, 28, 29] as they leave the mushroom body through the MBONs? While further refinement of the DAN-modulated KC-to-MBON synaptic dynamics would be required for a full analysis, in particular a detailed profiling of how the time course and magnitude of DAN and KC signals shape effective synapse strength, our existing results provide a first step for addressing this through modeling. For instance, one could model (1) KC population responses to odors based on literature [27, 28, 29], (2) synaptically connect the KC population to a modeled MBON population, and those synapses themselves to the DAN population, (3) develop a simple model of KC and DAN activations’ core effects on synaptic modulation, and (4) drive the KCs and DANs with the sensory and sensorimotor signals, respectively, corresponding a recorded fly’s walking trajectory through its virtual reality environment. Such a modeling analysis would lend insight into how the representations of appetitive odors that drive behavior might change over the course of a realistic search trajectory.

A further next step in this line of inquiry is to consider the role DAN modulation of mushroom body signaling might play within the well studied computational framework of reinforcement learning (RL). In this framework, a formalization of trial-and-error learning, an agent learns an optimal decision policy by gradually inferring causal connections between behavioral sequences and external rewards [30]. Furthermore, extensive study in mammalian systems has identified dopamine as a potential core actor in biological RL implementations [31, 32]. As a key component of RL is the interaction of short-term behavioral sequences and long-term learning, one might naturally ask whether and how the behaviorally modulated DAN signals could subserve the long-term associative learning processes traditionally associated with mushroom body function [33, 25]. This could be addressed through an alternate network model incorporating long-term mechanisms, such as spike-timing-dependent plasticity [34, 35], and examining whether dynamic DAN signals enable accelerated learning. Such an approach might also reveal any computational advantage gained by multiplexing sensorimotor and valence representations in the same neuromodulatory controls, and the RL framework could set the stage for more formal comparison of invertebrate vs mammalian dopaminergic computations.

Finally, it is worth acknowledging the mushroom body’s recent popularity as inspiration for novel machine learning algorithms for more general tasks [36, 37]. For instance, one recent study has suggested that mushroom-body-inspired network architecture, and in particular dopamine-like modulation, might hasten learning for classification tasks like hand-written digit recognition [36]. Continuing to unravel the nuances of mushroom body computations should therefore be fruitful not only in our understanding of neural processes underlying key behaviors, but potentially for engineered network computations also.

5.3 Rapid excitability-based sequence induction

Our work investigating rapid sequence induction provides a first treatment of the potential for ample flexibility of short-term information maintenance in neural circuits. Here we identified a fast, heterosynaptic plasticity mechanism capable of inducing novel sequences in model networks that replay both in response to external triggers and spontaneously. More precisely, the heterosynaptic plasticity mechanism, long-term potentiation of intrinsic excitability (LTP-IE) [38, 39], acted to “tag” recently active neurons by increasing their spiking probability when homogeneous “gating” inputs were applied to the network, such that these neurons were more likely to be recruited into population activity patterns. Since the recurrent connectivity was spatially organized, these population activity patterns typically corresponded to continuous movement through the space represented by the network, in particular following the LTP-IE profile induced by the preceding sensorimotor sequence. As a consequence, following LTP-IE and the introduction of homogeneous gating inputs, the network was highly biased towards replaying temporally compressed versions of recent sensorimotor-driven neural sequences, both in forward and reverse, recapitulating core features of empirically observed hippocampal replay [40, 41, 42, 43]. Beyond this, LTP-IE could be used to more generally establish temporary communication channels between different areas in the network, for instance enabling fast, flexible storage of sensory-to-motor mappings that had never been previously experienced by the network. This work demonstrated the capacity of physiological excitability changes to underlie a flexible and stable short-term memory buffer, indicating they might play a significant role in biological neural network models and in rapid cognitive flexibility in general.

While this work was not intended as a complete account of hippocampal replay, it nonetheless reveals a potential core mechanism behind the phenomenon and sets the stage for more thorough models. One canonical feature of hippocampal activity we did not model, for instance, is the observation that

replay events tend to co-occur with so-called sharp-wave ripples (SWRs), salient 100-200 ms oscillatory events in the surrounding local field potential, with oscillations around 150-200 Hz [41, 42, 43]. Furthermore, SWRs are hypothesized to play a role in coordinating information transfer from hippocampus to other cortical regions in the brain, which can exhibit replay events simultaneous with hippocampus during sleep [44], and disrupting SWRs has been shown to inhibit memory consolidation [45]. While we did not model SWRs, (recall that we did not include inhibition, a likely source of oscillations in general networks [46, 47, 48]), incorporating SWRs into a future LTP-IE-based model might serve to impose stricter regulation of the timing of replay, for instance with the SWR times controlling exactly when replay events arise, and with the oscillations acting as a potential clock-like signal to control replay speed. A further phenomenon that must be addressed in future replay models is the conversion of short-term into long-term memories, in particular through long-term mechanisms like spike-timing-dependent plasticity (STDP), which was indeed originally identified in cultured hippocampal neurons [34]. Since replay events tend to compress spike sequences into intervals aligned with the established 10s-of-ms STDP window, and since disruption of replay impairs memory consolidation, it appears likely that STDP may play a role in consolidating memories stored and replayed in the short-term memory buffer. Understanding this short-term and long-term interaction in both hippocampus as well as the cortical networks it may help coordinate will be a paramount future research direction, and we believe LTP-IE could play a prominent role in the short-term storage mechanism.

From a computational perspective, this work speaks more generally to the recently popular “neurocentric” theory of learning, in which information is predominantly stored in cellular, rather than synaptic changes. Mathematically, in a network of N neurons this difference corresponds to storing memories in N values (the number of neurons) vs $O(N^2)$ values (the number of synapses). Of note, the former additionally encapsulates certain forms of synaptic memories that are dependent on neuronal (as opposed to synaptic) activity alone, e.g. calcium buildup in all of a neuron’s synapses following spiking [49]; such synaptic memories are therefore also N - rather than $O(N^2)$ -dimensional.

While at first blush one might assume substantial computational advantage in storing memories in $O(N^2)$ values, there are several reasons to consider using only N . First, theoretically a system of N variables that can each take on M states can specify M^N different memories, which may still be quite sufficient for biological purposes; indeed, recent work has shown that such exponential memories could indeed exist stably side-by-side within a network’s dynamic landscape [50, 51]. Second, many identified mechanisms with rapid onset but enduring effects, like LTP-IE [38, 39] and the just-mentioned calcium

buildup [49], as well as effects caused by, say, changes in mean or fluctuations of baseline noise inputs, occur at the N -dimensional neuron level. In contrast, plasticity mechanisms occurring at the $O(N^2)$ -dimensional synapse level, the most notable example being STDP [34], likely occur far too gradually to significantly modify network dynamics over the short term. Third, theoretical work in the field of reservoir computing [52, 53, 54] has shown that given a sufficiently connected network of N neurons, universal time-series transformations can be specified via only N parameters, in this case the output weights of the reservoir network; suggesting N dimensions may be sufficient for specifying highly flexible computations. Finally, if the brain stores long-term memories in the $O(N^2)$ -dimensional space of synaptic strengths, it might make sense to store short-term memories in a different substrate, such that their interference with and/or integration into long-term memories would not be by default, but rather potentially controllable, e.g. via reactivation that occurs as a function of external reward [55]. Given this line of reasoning, we therefore believe it may be fruitful to consider how flexible short-term memories might be stored without requiring synapse-by-synapse storage mechanisms.

The above consideration frames several questions addressing the rapid emergence and retention of temporary information structures in a neural network. Overarchingly, given a (N -dimensional) neuron-level memory buffer, how does the existing synaptic connectivity structure shape what memories can be stored and recalled? In other words, how could the existing representations in the network, presumably learned through long-term experience and gradual plasticity mechanisms like STDP, be recruited to stabilize short-term representations? Are there specific network structures that maximize the set of stable, recallable, short-term memories that can be stored by modifying neuron-level variables? What connectivity statistics would such networks obey? These are problems all approachable through mathematical modeling, yet which could nonetheless reveal substantial insight into a brain-like computational paradigm. Furthermore, by revealing computationally useful connectivity structures and the neuronal dynamics they support, such models could yield direct predictions that could be tested via experiment. Indeed, such an approach could provide a concrete neural underpinning and mathematical entry point into comprehending and contextualizing the phenomenology of human short-term and working memory faculties collected over the last century [56, 57]. While memory in neuronal systems likely depends on a diverse host of biophysical mechanisms, we believe a such neurocentric research arc may be very promising for shedding new light on this fascinating system.

5.4 Mesoscale connectivity in mammalian brain networks

Finally, we investigated mesoscale structural patterns in mammalian brain networks, which will be needed to model memory processes and neuronal dynamics more generally within architectures inspired by full brain networks. Remarkably, in examining the mesoscale mouse connectome collected by the Allen Institute for Brain Science [58], we found that multiple core connectivity statistics could be reproduced with a parsimonious model. This model relied on two simple rules for network growth: “source growth”, a rich-get-richer rule where brain regions with many existing outgoing connections were more likely to grow more outgoing connections; and “proximal attachment”, in which outgoing connections were more likely to land on nearby regions than on distant ones, but were otherwise random. In particular, these rules recapitulated the exponential out-degree and Gaussian in-degree distributions observed in the mesoscale connectome, as well as the inverse relationship between clustering coefficient and total degree. This revealed that models simulating brain networks can substantially benefit from treating incoming and outgoing connectivity patterns independently from the start, rather than beginning with “undirected” models and later adding perturbations to account for directed patterns.

While we took advantage of what was to our knowledge the first complete mammalian connectome [58], work is underway to characterize many other connectomes at different scales, from the drosophila larval and adult connectomes [59] to marmoset [60] and human [61, 62] connectomes. Indeed, novel techniques such as MAPseq [63], which enable tracing of many axonal tracts simultaneously in a single brain, may rapidly accelerate experimental connectomics to produce extremely rich datasets. A natural question will be to ask whether such connectomes beyond the mouse brain can similarly be explained by simple growth rules, and if so, how these rules compare to mesoscale mouse brain’s. Such an investigation could provide a useful characterization of connectivity structures across scales and organisms and will be crucial to identifying (1) whether there exist any universal principles of brain connectivity and (2) what core dimensions best account for the structural variability observed both within and across systems. This work also makes predictions and sets the stage for future developmental studies, potentially providing a mathematically tractable anchorpoint for interpreting experimental observations of how connections within real brain networks develop and refine throughout development. A further step to refining our understanding of connectivity, presently underway at the Allen Institute for Brain Science [64], is to additionally characterize cell-type specific connectivity, in particular connectivity patterns that depend on whether each cell type is excitatory or inhibitory. Such an analysis will be crucial to running dynamical models atop whole brain connectivity, since in such models the sign of interaction between brain regions is of paramount importance in shaping dynamics

and computation.

This work suggests that in developing future models of full brain networks, the above two rules might largely suffice to specify connectivity. This could greatly simplify modeling efforts and make results more interpretable and generalizable, potentially easing the challenge of separating core structural properties from otherwise diverse higher-order connectivity patterns. Indeed, in combination with recent work revealing that both large-scale structural and dynamical features of microscale brain networks might be predominantly specified by local connection motifs [65], this suggests the space of network structures we must explore to accurately model brain connectivity might be smaller than it initially appears.

A compelling next step combining our own work with that of others will be to model the activity dynamics such network models with parsimonious, biologically inspired connectivity rules might support. Indeed, while the connections within a network highly constrain and enable the dynamics it can generate [66], even with simple connectivity patterns it can be extremely challenging to predict which dynamics will unfold. In particular, if we are to one day understand memory and neuronal dynamics as they are distributed across the whole brain at multiple scales, a key step will be developing models that combine mesoscale with microscale connectivity. For instance, one might use mesoscale connectivity as a “scaffold” for microscale connectivity, e.g. implementing a coarse-resolution interareal architecture according to our identified source-growth and proximal-attachment rules, then modeling each area itself as a network of locally connected neurons, with long-range connections to other areas determined by the mesoscale architecture. We can then formally address several questions treating multiscale structure and dynamics. For instance, if individual areas are implemented as balanced excitatory-inhibitory networks [67], does the mesoscale architecture disrupt or stabilize this balance, and what parameters control this? Moreover, how does mesoscale architecture shape information flow through the full brain network? This latter question returns us to our overarching theme of flexible short-term memory representations and dynamics. In particular, could the same excitability-based specification of neuronal sequences and information-routing play out at the mesoscale in addition to the microscale, as we explored in the previous section on sequence induction in model networks? If so, similar mechanisms in both micro- and mesoscale dynamics might suggest a novel, scale-free principal of brain computation. That all of these questions may be addressable within parsimonious network architectures suggests that gaining a comprehensive yet principled understanding of the interaction between structure and dynamics in the brain may in fact be tractable.

5.5 Final remarks

These four research projects have exposed several different components likely involved in short-term memory maintenance and rapid reshaping of computation and behavior, a computational faculty that appears to be involved in a plethora of neural systems. Our full understanding of memory and computational flexibility at this timescale will therefore be a crucial component to a comprehensive and principled understanding of how our nervous systems function in the natural world. One aspect of this work we have not much addressed is the comparison between invertebrate and mammalian nervous systems. Whereas mammalian cortical and hippocampal network structures arise from the interaction within large collections of ostensibly quite similar and interchangeable cell identities, invertebrate nervous systems are strongly characterized by their reliance on *identified neurons*, which are stereotyped across individuals on a cell-by-cell basis, and which often appear to exhibit highly nuanced and multiplexed functionality. The compartmentalized dopaminergic innervation of the mushroom body gamma lobe, for instance, which we have examined in detail, is preserved across individuals. Thoroughly characterizing the differences between identified-neuron computations and those arising from more homogeneous cortical/hippocampal networks will likely prove indispensable to how we think about the immense space of nervous systems that exist across the animal kingdom, and moreover the information processing strategies these nervous systems support. A key question within such a space of nervous systems is whether any computational principles are universally preserved. For instance, in spite of the marked difference between invertebrate and mammalian nervous systems, the insect antennal lobe and mammalian olfactory bulb share a striking resemblance in both structure and function within olfactory processing, suggesting some universal principles may indeed span quite distinct systems. Moreover, refining this understanding may yield insights into developing artificial neural systems that capitalize on biomimetic principles to perform flexible, efficient computations. Overall, the work presented in this dissertation refines our understanding of short-term memory and rapid reconfiguration of neural computations from multiple perspectives, and we hope it sets the stage for several exciting future lines of inquiry.

References

- [1] Massimo Vergassola, Emmanuel Villermaux, and Boris I Shraiman. ‘infotaxis’ as a strategy for searching without gradients. *Nature*, 445(7126):406, 2007.
- [2] Hubert Schardin. Die schlierenverfahren und ihre anwendungen. In *Ergebnisse der exakten Naturwissenschaften*, pages 303–439. Springer, 1942.
- [3] Michael John Hargather and Gary S Settles. Natural-background-oriented schlieren imaging. *Experiments in fluids*, 48(1):59–68, 2010.
- [4] Mark A Frye and Michael H Dickinson. Closing the loop between neurobiology and flight behavior in drosophila. *Current opinion in neurobiology*, 14(6):729–736, 2004.
- [5] Lance F Tammero, Mark A Frye, and Michael H Dickinson. Spatial organization of visuomotor reflexes in drosophila. *Journal of Experimental Biology*, 207(1):113–122, 2004.
- [6] Floris Van Breugel and Michael H Dickinson. The visual control of landing and obstacle avoidance in the fruit fly drosophila melanogaster. *Journal of Experimental Biology*, 215(11):1783–1798, 2012.
- [7] Floris van Breugel and Michael H Dickinson. Plume-tracking behavior of flying drosophila emerges from a set of distinct sensory-motor reflexes. *Current Biology*, 24(3):274–286, 2014.
- [8] Conor J McMeniman, Román A Corfas, Benjamin J Matthews, Scott A Ritchie, and Leslie B Vosshall. Multimodal integration of carbon dioxide and other sensory cues drives mosquito attraction to humans. *Cell*, 156(5):1060–1071, 2014.
- [9] Floris van Breugel, Jeff Riffell, Adrienne Fairhall, and Michael H Dickinson. Mosquitoes use vision to associate odor plumes with thermal targets. *Current Biology*, 25(16):2123–2129, 2015.
- [10] Klaus Hausen and Martin Egelhaaf. Neural mechanisms of visual course control in insects. In *Facets of vision*, pages 391–424. Springer, 1989.
- [11] Shawn R Olsen and Rachel I Wilson. Cracking neural circuits in a tiny brain: new approaches for understanding the neural circuitry of drosophila. *Trends in neurosciences*, 31(10):512–520, 2008.
- [12] Haleh Fotowat, Amir Fayyazuddin, Hugo J Bellen, and Fabrizio Gabbiani. A novel neuronal pathway for visually guided escape in drosophila melanogaster. *Journal of neurophysiology*, 2009.

- [13] John C Tuthill, M Eugenia Chiappe, and Michael B Reiser. Neural correlates of illusory motion perception in drosophila. *Proceedings of the National Academy of Sciences*, 108(23):9685–9690, 2011.
- [14] Saskia EJ De Vries and Thomas R Clandinin. Loom-sensitive neurons link computation to action in the drosophila visual system. *Current Biology*, 22(5):353–362, 2012.
- [15] Shin-ya Takemura, Arjun Bharioke, Zhiyuan Lu, Aljoscha Nern, Shiv Vitaladevuni, Patricia K Rivlin, William T Katz, Donald J Olbris, Stephen M Plaza, Philip Winston, et al. A visual motion detection circuit suggested by drosophila connectomics. *Nature*, 500(7461):175, 2013.
- [16] James A Strother, Aljoscha Nern, and Michael B Reiser. Direct observation of on and off pathways in the drosophila visual system. *Current Biology*, 24(9):976–983, 2014.
- [17] James A Strother, Shiuan-Tze Wu, Allan M Wong, Aljoscha Nern, Edward M Rogers, Jasmine Q Le, Gerald M Rubin, and Michael B Reiser. The emergence of directional selectivity in the visual motion pathway of drosophila. *Neuron*, 94(1):168–182, 2017.
- [18] Eduardo Martin Moraud and Dominique Martinez. Effectiveness and robustness of robot infotaxis for searching in dilute conditions. *Frontiers in neurorobotics*, 4:1, 2010.
- [19] Alejandro R Garcia Ramirez, Amarilys Lima Lopez, Andy Blanco Rodríguez, Alejandro Durán C de Albornoz, and Edson Roberto De Pieri. An infotaxis based odor navigation approach. In *ISSNIP Biosignals and Biorobotics Conference 2011*, pages 1–6. IEEE, 2011.
- [20] Victor Hernandez Bennetts, Achim Josef Lilienthal, Patrick Neumann, and Marco Trincavelli. Mobile robots for localizing gas emission sources on landfill sites: is bio-inspiration the way to go? *Frontiers in neuroengineering*, 4:20, 2012.
- [21] Nicole Voges, Antoine Chaffiol, Philippe Lucas, and Dominique Martinez. Reactive searching and infotaxis in odor source localization. *PLoS computational biology*, 10(10):e1003861, 2014.
- [22] Adam J Calhoun, Sreekanth H Chalasani, and Tatyana O Sharpee. Maximally informative foraging by caenorhabditis elegans. *Elife*, 3:e04220, 2014.
- [23] Juan Duque Rodríguez, David Gómez-Ullate, and Carlos Mejía-Monasterio. On the performance of blind-infotaxis under inaccurate modeling of the environment. *The European Physical Journal Special Topics*, 226(10):2407–2420, 2017.

- [24] Boris I Shraiman and Eric D Siggia. Scalar turbulence. *Nature*, 405(6787):639, 2000.
- [25] Yoshinori Aso, Daisuke Hattori, Yang Yu, Rebecca M Johnston, Nirmala A Iyer, Teri-TB Ngo, Heather Dionne, LF Abbott, Richard Axel, Hiromu Tanimoto, et al. The neuronal architecture of the mushroom body provides a logic for associative learning. *Elife*, 3:e04577, 2014.
- [26] Raphael Cohn, Ianessa Morantte, and Vanessa Ruta. Coordinated and compartmentalized neuromodulation shapes sensory processing in drosophila. *Cell*, 163(7):1742–1755, 2015.
- [27] Glenn C Turner, Maxim Bazhenov, and Gilles Laurent. Olfactory representations by drosophila mushroom body neurons. *Journal of neurophysiology*, 2008.
- [28] Kyle S Honegger, Robert AA Campbell, and Glenn C Turner. Cellular-resolution population imaging reveals robust sparse coding in the drosophila mushroom body. *Journal of Neuroscience*, 31(33):11772–11785, 2011.
- [29] Andrew C Lin, Alexei M Bygrave, Alix De Calignon, Tzumin Lee, and Gero Miesenböck. Sparse, decorrelated odor coding in the mushroom body enhances learned odor discrimination. *Nature neuroscience*, 17(4):559, 2014.
- [30] Richard S Sutton, Andrew G Barto, et al. *Introduction to reinforcement learning*, volume 135. MIT press Cambridge, 1998.
- [31] Wolfram Schultz, Peter Dayan, and P Read Montague. A neural substrate of prediction and reward. *Science*, 275(5306):1593–1599, 1997.
- [32] Paul W Glimcher. Understanding dopamine and reinforcement learning: the dopamine reward prediction error hypothesis. *Proceedings of the National Academy of Sciences*, 108(Supplement 3):15647–15654, 2011.
- [33] J Steven de Belle and Martin Heisenberg. Associative odor learning in drosophila abolished by chemical ablation of mushroom bodies. *Science*, 263(5147):692–695, 1994.
- [34] Guo-qiang Bi and Mu-ming Poo. Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *Journal of neuroscience*, 18(24):10464–10472, 1998.
- [35] Stijn Cassenaer and Gilles Laurent. Conditional modulation of spike-timing-dependent plasticity for olfactory learning. *Nature*, 482(7383):47, 2012.

- [36] Charles B Delahunt and J Nathan Kutz. Putting a bug in ml: The moth olfactory network learns to read mnist. *arXiv preprint arXiv:1802.05405*, 2018.
- [37] Constructing models of heterogeneous dopamine activity. *Cosyne Abstracts 2019, Lisbon, PT*, 2019.
- [38] Jung Ho Hyun, Kisang Eom, Kyu-Hee Lee, Won-Kyung Ho, and Suk-Ho Lee. Activity-dependent downregulation of d-type k⁺ channel subunit kv1. 2 in rat hippocampal ca3 pyramidal neurons. *The Journal of physiology*, 591(22):5525–5540, 2013.
- [39] Jung Ho Hyun, Kisang Eom, Kyu-Hee Lee, Jin Young Bae, Yong Chul Bae, Myoung-Hwan Kim, Sooyun Kim, Won-Kyung Ho, and Suk-Ho Lee. Kv1. 2 mediates heterosynaptic modulation of direct cortical synaptic inputs in ca3 pyramidal cells. *The Journal of physiology*, 593(16):3617–3643, 2015.
- [40] David J Foster and Matthew A Wilson. Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature*, 440(7084):680, 2006.
- [41] Kamran Diba and György Buzsáki. Forward and reverse hippocampal place-cell sequences during ripples. *Nature neuroscience*, 10(10):1241, 2007.
- [42] Thomas J Davidson, Fabian Kloosterman, and Matthew A Wilson. Hippocampal replay of extended experience. *Neuron*, 63(4):497–507, 2009.
- [43] Margaret F Carr, Shantanu P Jadhav, and Loren M Frank. Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval. *Nature neuroscience*, 14(2):147, 2011.
- [44] Daoyun Ji and Matthew A Wilson. Coordinated memory replay in the visual cortex and hippocampus during sleep. *Nature neuroscience*, 10(1):100, 2007.
- [45] Shantanu P Jadhav, Caleb Kemere, P Walter German, and Loren M Frank. Awake hippocampal sharp-wave ripples support spatial memory. *Science*, 336(6087):1454–1458, 2012.
- [46] Xiao-Jing Wang and György Buzsáki. Gamma oscillation by synaptic inhibition in a hippocampal interneuronal network model. *Journal of neuroscience*, 16(20):6402–6413, 1996.
- [47] Maxim Bazhenov, Mark Stopfer, Mikhail Rabinovich, Ramon Huerta, Henry DI Abarbanel, Terrence J Sejnowski, and Gilles Laurent. Model of transient oscillatory synchronization in the locust antennal lobe. *Neuron*, 30(2):553–567, 2001.

- [48] Nobuaki K Tanaka, Kei Ito, and Mark Stopfer. Odor-evoked neural oscillations in drosophila are mediated by widely branching interneurons. *Journal of Neuroscience*, 29(26):8595–8603, 2009.
- [49] Gianluigi Mongillo, Omri Barak, and Misha Tsodyks. Synaptic theory of working memory. *Science*, 319(5869):1543–1546, 2008.
- [50] Gilles Wainrib and Jonathan Touboul. Topological and dynamical complexity of random neural networks. *Physical review letters*, 110(11):118101, 2013.
- [51] Rishidev Chaudhuri and Ila Fiete. Associative content-addressable networks with exponentially many robust stable states. *arXiv preprint arXiv:1704.02019*, 2017.
- [52] Herbert Jaeger. The echo state approach to analysing and training recurrent neural networks-with an erratum note. *Bonn, Germany: German National Research Center for Information Technology GMD Technical Report*, 148(34):13, 2001.
- [53] Wolfgang Maass, Thomas Natschläger, and Henry Markram. Real-time computing without stable states: A new framework for neural computation based on perturbations. *Neural computation*, 14(11):2531–2560, 2002.
- [54] Mantas Lukoševičius and Herbert Jaeger. Reservoir computing approaches to recurrent neural network training. *Computer Science Review*, 3(3):127–149, 2009.
- [55] R Ellen Ambrose, Brad E Pfeiffer, and David J Foster. Reverse replay of hippocampal place cells is uniquely modulated by changing reward. *Neuron*, 91(5):1124–1136, 2016.
- [56] Alan Baddeley. Working memory. *Science*, 255(5044):556–559, 1992.
- [57] Alan Baddeley. Working memory: looking back and looking forward. *Nature reviews neuroscience*, 4(10):829, 2003.
- [58] Seung Wook Oh, Julie A Harris, Lydia Ng, Brent Winslow, Nicholas Cain, Stefan Mihalas, Quanxin Wang, Chris Lau, Leonard Kuan, Alex M Henry, et al. A mesoscale connectome of the mouse brain. *Nature*, 508(7495):207, 2014.
- [59] Katharina Eichler, Feng Li, Ashok Litwin-Kumar, Youngser Park, Ingrid Andrade, Casey M Schneider-Mizell, Timo Saumweber, Annina Huser, Claire Eschbach, Bertram Gerber, et al. The complete connectome of a learning and memory centre in an insect brain. *Nature*, 548(7666):175, 2017.

- [60] Hideyuki Okano and Partha Mitra. Brain-mapping projects using the common marmoset. *Neuroscience research*, 93:3–7, 2015.
- [61] Olaf Sporns, Giulio Tononi, and Rolf Kötter. The human connectome: a structural description of the human brain. *PLoS computational biology*, 1(4):e42, 2005.
- [62] Klaus H Maier-Hein, Peter F Neher, Jean-Christophe Houde, Marc-Alexandre Côté, Eleftherios Garyfallidis, Jidan Zhong, Maxime Chamberland, Fang-Cheng Yeh, Ying-Chia Lin, Qing Ji, et al. The challenge of mapping the human connectome based on diffusion tractography. *Nature communications*, 8(1):1349, 2017.
- [63] Justus M Kebschull, Pedro Garcia da Silva, Ashlan P Reid, Ian D Peikon, Dinu F Albeanu, and Anthony M Zador. High-throughput mapping of single-neuron projections by sequencing of barcoded rna. *Neuron*, 91(5):975–987, 2016.
- [64] Julie A Harris, Stefan Mihalas, Karla E Hirokawa, Jennifer D Whitesell, Joseph Knox, Amy Bernard, Phillip Bohn, Shiella Caldejon, Linzy Casal, Andrew Cho, et al. The organization of intracortical connections by layer and cell class in the mouse brain. *BioRxiv*, page 292961, 2018.
- [65] Yu Hu, James Trousdale, Kršimir Josić, and Eric Shea-Brown. Motif statistics and spike correlations in neuronal networks. *BMC neuroscience*, 13(1):P43, 2012.
- [66] Hannah Choi and Stefan Mihalas. Synchronization dependent on spatial structures of a mesoscopic whole-brain network. *bioRxiv*, page 319830, 2018.
- [67] Nicolas Brunel. Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *Journal of computational neuroscience*, 8(3):183–208, 2000.