

UC Santa Barbara

UC Santa Barbara Electronic Theses and Dissertations

Title

Patterns of variability in brain and behavior across spatiotemporal scales

Permalink

<https://escholarship.org/uc/item/3nc5j0s6>

Author

Mueller, Joshua

Publication Date

2020

Peer reviewed|Thesis/dissertation

University of California
Santa Barbara

**Patterns of variability in brain and behavior across
spatiotemporal scales**

A dissertation submitted in partial satisfaction
of the requirements for the degree

Doctor of Philosophy
in
Dynamical Neuroscience

by

Joshua M. Mueller

Committee in charge:

Professor Jean M. Carlson, Chair
Professor Julie H. Simpson
Professor Jeff Moehlis
Professor Michael Goard

September 2020

The Dissertation of Joshua M. Mueller is approved.

Professor Julie H. Simpson

Professor Jeff Moehlis

Professor Michael Goard

Professor Jean M. Carlson, Committee Chair

August 2020

Patterns of variability in brain and behavior across spatiotemporal scales

Copyright © 2020

by

Joshua M. Mueller

Acknowledgements

Thanks to Jean Carlson, Julie Simpson, Emily Jacobs, Eric Jones, Primoz Ravbar, Laura Pritschet, Neil Zhang, Tyler Santander, Caitlin Taylor, Scott Grafton, Kimberly Schlesinger, Zipeng Wang, Li Guo, and Jordan Zawaydeh for their contributions to this work. The research presented here would not have been possible without their knowledge, patience, guidance, advice, encouragement, hard work, and generosity.

Thanks to Sam Verbanic, Will Summers, Eric Jones, Mark Hirsch, Elizabeth Forbes, Owen Campbell, Caitlin McAvoy, Sarah Halperin, Erin O'Reilly, Laura Gray, and Sölen Kiratli for providing balance and perspective. I arrived in Santa Barbara a stranger to all of you; now I hear your voices when I look at the world.

Finally, thanks to my family for a lifetime of support. Accomplishments do not occur in a vacuum; I am grateful always for the unearned opportunities that have allowed me to follow my interests. Wherever I land, I know that my journey began with yours.

This work was supported by UCSB through the Chancellor's Fellowship, the Towbes Graduate Research Fellowship, and the Worster Summer Research Fellowship. This work was also supported the David and Lucile Packard Foundation and the Institute for Collaborative Biotechnologies through contract no. W911NF-09-D-0001 from the U.S. Army Research Office.

The content of the information in this dissertation does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Thanks to Nicholas M. Timme, Najja Marshall, Nicholas Bennett, Monica Ripp, Edward Lautzenhiser, and John M. Beggs for permission to use the data sets analyzed in Chapter 2.

Curriculum Vitæ

Joshua M. Mueller

Education

- 2020 Ph.D. in Dynamical Neuroscience (Expected), University of California, Santa Barbara.
- 2015 B.A. in Biochemistry, University of Pennsylvania.

Publications

Mueller JM, Pritschet L, Santander T, Taylor CM, Grafton ST, Jacobs EG, Carlson JM (2020) Dynamic community detection reveals transient reorganization of functional brain networks across a female menstrual cycle. In revision, bioRxiv preprint available at doi: <https://doi.org/10.1101/2020.06.29>.

Wang Z, Jones EW, **Mueller JM**, Carlson JM (2020) Control of ecological outcomes through deliberate parameter changes in a model of the gut microbiome. *Physical Review E* 101(5): 052402.

Mueller JM, Ravbar P, Simpson JH, Carlson JM (2019) *Drosophila melanogaster* grooming possesses syntax with distinct rules at different temporal scales. *PLoS Comput. Biol.* 15(6): e1007105.

Chen HS, Martin KA, Lu F, Lupey LN, **Mueller JM**, Lieberman PM, Tempera I (2014) Epigenetic deregulation of the LMP1/LMP2 locus of Epstein-Barr virus by mutation of a single CTCF-cohesin binding site. *J. Virology* 88(3): 1703-1713.

Mueller JM, Carlson JM (2020) Dissociated hippocampal cultures spontaneously develop low-dimensional spatiotemporal firing patterns. In submission.

Abstract

Patterns of variability in brain and behavior across spatiotemporal scales

by

Joshua M. Mueller

Brain systems span a wide range of spatial and temporal scales, from individual neurons firing every few milliseconds to whole-brain connectivity changing over many years. Network models offer a mathematical framework for describing, predicting, and controlling a diverse array of brain-related phenomena. Regardless of the particular system of study, it is possible to characterize the range of observed variability in system behavior over time by calculating topological and statistical metrics from inferred networks. In this thesis, tools from network science, computational neuroscience, and statistics are applied to data sets spanning several systems and spatiotemporal scales to characterize patterns of variability in brain activity and behavior within and between individuals over time.

First, electrophysiological activity of dissociated rat hippocampal neurons *in vitro* is analyzed to reveal the emergence and persistence of stereotyped activity patterns over time in the absence of external stimuli. These discoveries motivate further use of *in vitro* hippocampal cultures as model systems for studying developmental phenomena related to learning and memory.

Next, fruit fly behavior is modeled using tools from computational ethology and statistics to define and compare the syntax of grooming, a common behavior in *Drosophila*. We identify duration dependence in syntax for the first time, suggesting that central pattern generators contribute to organizing behavior in a partially sensory-independent manner.

We also identify inter- and intra-species variation in grooming syntax, but do not observe a significant correspondence between genetic heterogeneity and grooming variability. We also examine optogenetically-stimulated flies and find variability in grooming responses despite careful control of sensory presentation. Taken together, these results demonstrate that genetic differences can produce variation in grooming behavior, but that sensory experiences and stochastic effects contribute more strongly than genetics to the natural range of expressed grooming phenotypes.

Finally, we characterize the stability and changes in human brain network connectivity over the course of a complete menstrual cycle in a “dense sampling” experimental paradigm. We identify stable communities that persist for the majority of the menstrual cycle. However, we also discover a correspondence between sex hormones and transient large-scale brain network reorganization, implicating hormones in shaping network topology.

Despite the differences in systems and phenomena of interest in this thesis, the work contained within is unified by a concern for variability in brain activity and behavior. Across contexts, we demonstrate the utility of network models in describing changes in complex systems over time.

Contents

Curriculum Vitae	v
Abstract	vi
1 Introduction	1
1.1 Networks, brains, and behavior	1
1.2 Variability	8
1.3 Permissions and Attributions	13
2 Spontaneous Development of Low-dimensional Spatiotemporal Activity Patterns in Neuronal Cultures <i>In Vitro</i>	15
2.1 Introduction	15
2.2 Methods	19
2.3 Results	26
2.4 Discussion	30
3 Discovery and Characterization of Duration Dependence in Fruit Fly Grooming Syntax	36
3.1 Introduction	36
3.2 Methods and Models	43
3.3 Results	50
3.4 Discussion	61
3.5 Conclusion	68
4 Variability and Variation in <i>Drosophila</i> Grooming Syntax	70
4.1 Introduction	70
4.2 Results	73
4.3 Discussion	81
4.4 Materials and methods	85

5	Brain Network Reorganization Over A Single Menstrual Cycle	88
5.1	Introduction	88
5.2	Results	92
5.3	Discussion	102
5.4	Conclusion	108
5.5	Methods	109
6	Concluding Remarks	114
A		116
A.1	Supporting Figures	116
B		119
B.1	Supporting Figures	119
C		125
C.1	Supporting Methods	125
C.2	Supporting Figures	128
D		139
D.1	Supporting Figures	139
	Bibliography	148

Chapter 1

Introduction

1.1 Networks, brains, and behavior

Across systems and scales, complex dynamics emerge from the interactions between many individual elements. The growth of microbial populations, the patterns formed by flocking birds, and even shifts in voting preferences have all been described by network models[1, 2, 3], which decompose these systems into component parts and the interactions between them. These models are complementary to, and sometimes reducible to, analytical approaches, such as mean-field theories, that describe bulk or statistical behavior of large systems.

In this formalism, components are referred to as *nodes* and connections between them as *edges*, which together make up a *graph* (graph and network may be used interchangeably here)[4]. Network models possess several useful features not present in other formalisms. For one, networks possess observable *topology* that can be formally quantified using mathematical metrics. More simply, with network models it is possible to inspect which components are connected to one another, how strongly they are connected, and whether geometric patterns (such as loops) occur. By analyzing metrics such as these,

it is possible to characterize more abstract properties of the network, such as the ability to “traverse” from one node to another along edges.

In perhaps the most famous example of this approach, the layout of the city of Königsberg, Prussia was formulated as a network by Leonhard Euler in 1736. As shown in Figure 1.1 below, Euler conceptualized the land masses of the city that were connected by seven bridges as a mathematical object to prove that one could not walk the entire city in a loop while crossing each bridge only once, laying the foundations for modern graph theory.




Figure 1.1: **The seven bridges of Königsberg.** On the left is a map schematic of the city of Königsberg, which consisted of four landmasses connected by seven bridges. On the right is a simple graph representation of the city, with numbered nodes corresponding to landmasses and edges indicating the presence of a bridge between landmasses. Since graphs are abstract representations, their particular spatial layout does not matter; identifying the connections between nodes fully defines the graph topology.

1.1.1 Structural and functional brain networks

The seven bridges of Königsberg network is an example of a *structural* network, which provides information only about whether components are connected to one another. Additionally, this network is referred to as *undirected* and *unweighted*, as edges

are reciprocal and use only binary values to indicate the presence or absence of connections. These types of networks can be extended, however, to consider situations in which connections may be *directed* or *weighted*. Mathematically, these graph objects may be represented simply as an *adjacency matrix*, A , where the matrix entry A_{ij} indicates the strength of the connection from node i to node j .




Figure 1.2: **Common graph representations.** On the left is a weighted, directed network graph schematic. Edge thicknesses indicate weight and arrows indicate directionality of connections. On the right is a heatmap representation of the same graph; these objects are mathematically identical. In this representation, the color of entry (i,j) indicates the strength of the edge from node i to node j . Due to the abstract nature of graphs, these objects can represent a wide variety of systems and processes.

Though the brain is a biological organ that seemingly shares few, if any, similarities with urban geography, it is fundamentally an object composed of spatially segregated subunits (neurons) connected by long-range projections (axons and dendrites). This observation has motivated researchers in the field of *connectomics* to map the physical connections between neurons, circuits, and brain regions in organisms ranging in scale from *C. elegans* to humans [5]. By mapping these features, researchers have discovered general topological properties of brain organization, such as “small-world structure”, that may provide insight into how the brain coordinates activity across large spatial scales while maintaining coherent function at the level of microcircuits [6, 7].

In the human brain, *gray matter* regions, which are composed of cell bodies and

form the outer cortical surfaces of the brain, form connections with one another via bundles of axons (*white matter*) that pass through deeper layers of the brain. Using magnetic resonance imaging (MRI) to measure diffusion in the brain, several techniques enable mapping the white matter connections between gray matter regions [8]. Structural networks generated from this imaging exhibit a wide range of topological features that correspond to cognitive task performance and disease state [9, 10], indicating that brain architecture is an important determinant of brain function.

Graphs may also be used to represent *functional* networks, where edges indicate interactions between nodes. Here, the usage of “interactions” is loose and allows for many definitions. For example, correlation metrics are commonly used to define edges in functional networks when nodes represent objects that possess time series data, providing a proxy for potential interactions between elements without attributing causality.

In human brain studies, several techniques can measure brain activity directly or indirectly to generate time series data of activity in different regions at rest or during task performance. Functional MRI (fMRI), which detects blood-oxygen-level-dependent (BOLD) signal as a proxy for brain activity, is commonly used to generate functional brain networks. As stated previously, despite the nomenclature, the resulting networks are typically correlational, indicating coupling between regions without ascribing causal influence. Even so, differences in functional brain networks have been associated with cognitive task performance and disease state as well [11], affirming the principle that cognition is often dependent on coordinated activity between spatially disparate brain regions.

1.1.2 Dynamics of networks, dynamics on networks

Though the structural and functional networks described above may use time series data to define the network structure, they are still intrinsically *static* networks in the sense that relationships between nodes are fixed. It is possible, however, to consider *dynamics of networks* if there is reason to believe that edges change over time. One way to represent dynamic networks is using *multilayer* (or *multislice*) networks, which can extend structural and functional networks to include a temporal dimension by incorporating many “snapshots” of the network over time. Formally, these networks can be defined using a multilayer adjacency tensor similar to an adjacency matrix, but where entry A_{ijt} represents the connection between nodes i and j at time (layer/slice) t . Multilayer networks can be extended further to include links between layers, so that A_{ijst} denotes the connection between nodes i and j between layers s and t . Figure 1.3 provides a schematic overview of a multilayer network.




Figure 1.3: **Multilayer networks capture changes in network structure over time.** Shown here is a schematic of a multilayer network. At each time step, the graph connectivity undergoes changes in topology, but the identity of the nodes stays the same. Tools from network science can be used to characterize these changes, providing insight about the stability of network structure over time.

As tools for measuring, describing, and analyzing brain networks have advanced, efforts to link structure and function have also proliferated. In particular, patterns of human brain activity coordination at rest and during tasks have been found to strongly

correspond to existing inter-regional structural connections [12]. However, this correspondence is not one-to-one, illustrating that *dynamics on networks* must be accounted for to generate the observed functional data given structural constraints. That is to say, both the physical connections and the rules by which brain regions influence one another (e.g. activation or inhibition) are needed to fully understand brain functioning; studying dynamics of networks (as opposed to dynamics on networks) only yields partial information by describing changes in topology over time without necessarily explaining the underlying mechanistic sources of observed time series data.

Both inferring dynamics on networks in a data-driven fashion and generating mechanistic models of brain dynamics from first principles are difficult, as illustrated by the well-known example of *C. elegans* connectivity. *C. elegans* is a relatively simple organism containing only 302 neurons whose nervous system wiring diagram has been fully mapped [13]. However, a model that fully recapitulates *C. elegans* behavior by implementing dynamics on the structural network has yet to be described, though efforts at generating individual behaviors, such as crawling and bending, in a biologically-plausible fashion have been successful [14]. Given the challenges of understanding even relatively simple nervous systems, using model systems, such as fruit flies or mice, to understand the neuroscientific underpinnings of behavior is a necessary and powerful complement to investigations of the human brain.

Finally, beyond characterizing and reproducing brain activity from structural and functional measurements, networks also provide a basis for *graphical models*. Broadly speaking, these models represent probabilistic relationships between variables represented by nodes and are typically explicitly non-biological, as they abstract brain dynamics even further to a purely statistical process.

One of the most commonly used graphical models is a Markov chain, in which edges represent transition probabilities between states (nodes). Just like nodes, these states

are abstract and can represent a wide variety of objects, such as behaviors. A simple example of a Markov chain is a first-order process, in which the probability of the system entering a state is dependent only on the current state of the system. Mathematically, the transition probabilities are given by the transition matrix P , where P_{ij} provides the probability of entering state j at time $t + 1$, given that the system is in state i at time t . This graphical model possesses identical structure to a weighted, directed graph, so similar topological analysis is possible, but these models are often used as an alternative to biologically-derived network models.

For example, Markov chains may be used generatively to produce time series data related to brain activity or behavior. In particular, Markov models have been used to describe observed fluctuations in resting-state fMRI connectivity [15] and spontaneous behavior in fruit flies [16]. While such models do not directly correspond to reductionist biological processes, they still serve as a useful guide by connecting phenomenological scales; successful approximate models provide measurable benchmarks for future efforts to discover mechanistic models that reproduce the behavior of their coarse-grain counterparts. Additionally, by identifying relevant dimensions of a phenomenon or behavior, these graphical models may suggest targets for further biological investigation even if the model components do not correspond directly to anatomical elements.

Due to their flexibility and interpretability, network models offer a mathematical framework for describing, predicting, and controlling a diverse array of system dynamics. When these models possess dynamics, it is possible to characterize the range of observed variability in system behavior using topological and statistical metrics. In this thesis, tools from network science, computational neuroscience, and statistics are applied to data sets spanning several systems and spatiotemporal scales to characterize patterns of variability in brain activity and behavior within and between individuals over time. Figure 1.4 provides a schematic of the systems analyzed here.

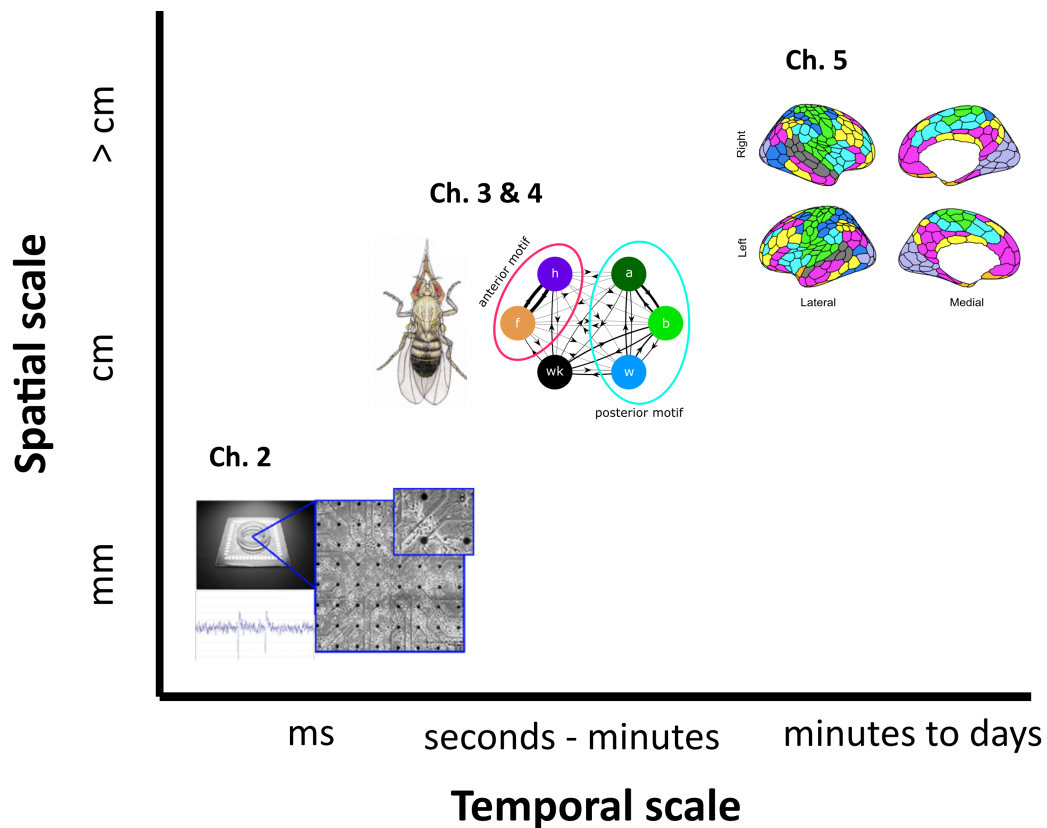


Figure 1.4: **Analysis of variability in brain activity and behavior across temporal and spatial scales.** In this thesis, three different neuroscientific systems are characterized using tools from network science. In Ch. 2, electrophysiological activity of rat hippocampal neurons *in vitro* is analyzed to reveal the emergence of stereotyped activity patterns over time in the absence of external stimulus. This system occupies the smallest spatial scale and fastest temporal scale. In Chs. 3 and 4, fruit fly behavior is modeled to define and compare the syntax of grooming. This behavior occurs at the spatiotemporal mesoscale. In Ch. 5, human brain imaging data is studied to uncover reorganization of large-scale brain networks over the course of a complete menstrual cycle. This system exhibits the slowest temporal dynamics and a range of spatial scales.

1.2 Variability

Despite the differences in systems and phenomena of interest in this thesis, the work contained within is unified by a concern for *variability* in brain activity and behavior. Here, variability is defined as the natural range of phenotypic expression within a ge-

netically identical or highly similar population, following Ayroles et al [17]. Techniques from statistics and machine learning, such as dimensionality reduction and supervised classification, can be used to quantify variability. In this thesis, these techniques are applied to various systems, as described below.

1.2.1 *In vitro* neuronal cultures

Recent advances in neuroscience propose that a variety of cognitive deficits may arise from disruptions in function at the level of the neuronal microcircuit. An effort has been made to characterize the organization of neurons *in vitro*, meaning existing in an artificial setting outside of a living organism. In one experimental platform, neurons can be grown on a multi-electrode array (MEA), which possesses electrodes that detect changes in extracellular voltage due to action potentials. Such a platform is advantageous for studying the function of neuronal microcircuitry because it allows for the collection of firing activity from individual neurons with better than millisecond resolution while also admitting the possibility of various manipulations, such as electrical or chemical intervention. In addition, neuronal cultures grown *in vitro* permit pre-experimental genetic manipulations and post-experimental screenings such as single cell RNA-seq.

Due to their accessibility and relative simplicity, rodent neuronal cultures *in vitro* are a common system for studying functional properties of networks of neurons. Neuronal cultures *in vitro* exhibit spontaneous firing activity that begins within days of seeding and that persists for several weeks. Phenomena such as coordinated bursting events have been shown to develop in many types of *in vitro* preparations [18]. Several types of bursts have been identified as well, suggesting that *in vitro* cultures possess a repertoire of possible activity patterns [19]. Although cultures of dissociated rat hippocampal neurons exhibit synchronous bursts of activity reminiscent of sharp wave ripples observed in

the mammalian hippocampus, little is known about the developmental origins of these bursting events.

Chapter 2 characterizes the range of patterns that develop within dissociated rat hippocampal cultures *in vitro* over the course of approximately 30 days. In this case, variability refers to the size of the observed firing repertoire, as the neuronal populations are derived from a common progenitor and are not subjected to genetic manipulation over the course of the experiment.

We demonstrate the emergence of several distinct neural activity trajectories in a low-dimensional space over time in this system. Clustering of longitudinal spike train data collected from rat hippocampal dissociated cultures on multi-electrode arrays identifies repeated, separable patterns of neural activity. The observed activity patterns in specific subpopulations become both more frequent and distinct over time, suggesting that *in vitro* neuronal cultures may share functional properties of neuronal populations *in vivo* such as distributed population codes. These findings encourage the use of dissociated cultures as model systems for probing the development of coding schemes in hippocampal networks.

1.2.2 Fruit fly grooming as a model of sequential behavior

Fruit flies perform grooming actions after exposure to irritants, which is thought to maintain sensory fidelity and prevent buildup of pathogens [20]. These grooming actions are stereotyped at the level of movement primitives, but the order in which the grooming movements are arranged varies within and between individuals [21]. Experiments demonstrate that, following stimulation by an irritant, grooming progresses gradually from an early phase dominated by anterior cleaning to a later phase with increased walking and posterior cleaning. Mathematical modeling of behavioral sequences yields insight into

the rules and mechanisms underlying sequence generation, which can illuminate general sensorimotor principles and provide targets for future investigations.

In Chapter 3, we develop a series of data-driven Markov models that isolate and identify the behavioral features governing transitions between individual grooming bouts. We identify action order as the primary driver of probabilistic, but non-random, syntax structure, as has previously been identified. Subsequent models incorporate grooming bout duration, which also contributes significantly to sequence structure. Our results show that, surprisingly, the syntactic rules underlying probabilistic grooming transitions possess action duration-dependent structure, suggesting that sensory input-independent mechanisms guide grooming behavior at short time scales. Finally, the inclusion of a simple rule that modifies grooming transition probabilities over time yields a generative model that recapitulates the key features of observed grooming sequences at several time scales. These discoveries suggest that sensory input guides action selection by modulating internally generated dynamics. Additionally, the discovery of these principles governing grooming in *D. melanogaster* demonstrates the utility of incorporating temporal information when characterizing the syntax of behavioral sequences.

Having established the structure of grooming syntax, in Chapter 4 we investigate genetic differences as a possible source of variation and variability in grooming behavior. Grooming features of five drosophilid species exposed to an irritant were analyzed. We found that components of grooming behavior were conserved between species. However, significant differences in grooming syntax were identified between and within drosophilid species, corresponding largely to differences in transition probabilities between grooming actions involving the posterior set of legs. These findings indicate that genetic differences may contribute to variation in grooming behavior, but the degree to which genetics dictate the natural range of observed grooming phenotypes within a homogeneous population is not yet understood.

To address this, *melanogaster* flies bred for maximal (minimal) genetic heterogeneity were examined to determine whether and how genetic heterogeneity within a population corresponds to the observed range of phenotypic variability. Genetic heterogeneity was not found to be related to grooming variability, as *melanogaster* flies interbred to increase (decrease) genetic heterogeneity did not exhibit increased (decreased) variability in grooming syntax. These results suggest that differences in sensory and life experience outweigh genetic contributions to grooming variability.

To assess the contribution of sensory experience to grooming variability, optogenetically-stimulated flies were examined. Despite tighter standardization of sensory presentation, these flies exhibited a similar range of variability as dusted flies. Taken together, these results indicate that further stochastic sources of variability, such as intrinsic circuit noise and developmental history, also contribute to grooming and require further investigation.

1.2.3 Longitudinal human brain functional organization

The human brain exhibits variable activity patterns on a broad range of scales. Even at rest, brain activity fluctuates noticeably on time scales as short as hundreds of milliseconds and on spatial scales as large as entire functional-anatomical networks. Sex steroid hormones have been shown to alter regional brain activity, but the extent to which they modulate connectivity within and between large-scale functional brain networks over time has yet to be characterized. In Chapter 5, we examine functional connectivity at rest in a single subject over the course of a complete menstrual cycle. Here, variability refers to changes in coupling strengths between regions over time in concert with naturally-fluctuating hormone levels characteristic of the menstrual cycle.

We applied dynamic community detection techniques to data from a highly sampled female with 30 consecutive days of brain imaging and venipuncture measurements to

characterize changes in resting-state community structure across the menstrual cycle. Four stable functional communities were identified consisting of nodes from visual, default mode, frontal control, and somatomotor networks. Limbic, subcortical, and attention networks exhibited higher than expected levels of nodal flexibility, a hallmark of between-network integration and transient functional reorganization.

The most striking reorganization occurred in a default mode subnetwork localized to regions of the prefrontal cortex, coincident with peaks in serum levels of estradiol, luteinizing hormone, and follicle stimulating hormone. Nodes from these regions exhibited strong intra-network increases in functional connectivity, leading to a split in the stable default mode core community and the transient formation of a new functional community. Probing the spatiotemporal basis of human brain-hormone interactions with dynamic community detection suggests that ovulation results in a temporary, localized patterns of brain network reorganization.

1.3 Permissions and Attributions

1. The content of Chapter 2 and Appendix A is the result of research performed under Dr. Jean M. Carlson and is currently in submission.
2. The content of Chapter 3 and Appendix B is the result of a collaboration with Dr. Julie H. Simpson and has previously appeared in *PLoS Computational Biology* (Mueller JM et al. *Drosophila melanogaster* grooming possesses syntax with distinct rules at different temporal scales. *PLoS Computational Biology* 15(6): e1007105). It is reproduced here with permission under the Creative Commons Attribution (CC BY) license.
3. The content of Chapter 4 and Appendix C is the result of a collaboration with

- Dr. Julie H. Simpson and is currently available as a pre-print via BioRxiv (Mueller JM et al. Variation and variability in *Drosophila* grooming behavior. MS ID#: BIORXIV/2020/252627).
4. The content of Chapter 5 and Appendix D is the result of a collaboration with Dr. Emily G. Jacobs and is currently available as a pre-print via BioRxiv (Mueller JM et al. Dynamic community detection reveals transient reorganization of functional brain networks across a female menstrual cycle. MS ID#: BIORXIV/2020/178152).
 5. The images in Figure 1.1 are reproduced from Creative Commons with permission under the CC0 1.0 license.

Chapter 2

Spontaneous Development of Low-dimensional Spatiotemporal Activity Patterns in Neuronal Cultures *In Vitro*

2.1 Introduction

Technology for recording electrophysiological activity from neuronal populations has advanced sufficiently to allow researchers to generate large-scale data sets (> 100 neurons observed for several hours) of neuronal firing activity [22, 23]. *In vitro* neuronal cultures serve as convenient model systems due to their scale, cost, ease of recording access, and potential for multimodal manipulation. Though they lack physiological context, recordings of hundreds of neurons spanning several hours can capture many aspects of the functional repertoire of neuronal ensembles, such as coordinated bursts. Since the mammalian hippocampus exhibits complex population-wide firing patterns during

important functions such as sleep and navigation but is experimentally difficult to access due to its subcortical location, it is a candidate for *in vitro* studies, which may uncover fundamental principles underlying circuit activity.

The role of the hippocampus in navigation and memory formation is well-documented [24, 25], as are some of the mechanisms underlying these functions [26]. Researchers have experimentally observed propagating waves of neuronal activity in the hippocampus across species, preparations, and recording modalities, suggesting that this brain region utilizes population-wide coordinated responses as a coding strategy [27, 28]. Several *in vivo* studies suggest that oscillatory activity in the hippocampus plays a role in episodic replay and memory consolidation [29], connecting dynamics to high-level function. However, the complex interplay between intrinsic developmental phenomena and extrinsic feedback to the hippocampus has made it difficult to fully understand how such functionality emerges.

To investigate this further, researchers have utilized *ex vivo* and *in vitro* hippocampal preparations. In these contexts, hippocampal preparations display a wide repertoire of spontaneous oscillatory phenomena and culture-wide synchrony [22, 30], motivating further study of spontaneous and evoked activity patterns [31] in the hopes that new results will shed light on the developmental and computational mechanisms responsible for hippocampal function. To determine the appropriateness of generalizing from this model system, further specific quantitative comparisons are necessary.

Several *in vivo* studies have demonstrated that dimensionality reduction is a helpful mathematical tool for characterizing complex neural activity. For example, projection of high-dimensional neural firing activity into a low-dimensional space can reveal the existence of neural activity trajectories which correspond to behavioral states [32, 33]. This phenomenon is in line with population coding theories, which propose that neural circuits carry out their functions via sparse, sequential firing patterns implemented by

multiple neurons. Not all brain networks are thought to use population coding schemes, so any assumptions about coding strategies *in vitro* should correspond to the observed phenomena in related regions *in vivo*. Specifically, rate codes and temporal codes have been identified in several cortical regions, such as V1 and M1, respectively [34], suggesting that analyzing activity for population codes in dissociated cultures derived from those regions may be misleading.

Recent *in silico* work complements biological experiments by identifying fundamental dynamics that allow for the spontaneous formation of complex neural trajectories in simulated neuronal populations. Specifically, a few simple local synaptic weight update rules can entrain a simulated neuronal network in an unsupervised manner to exhibit repeated, complex, spatiotemporally constrained neural trajectories [35, 36]. Moreover, this work identifies patterns consistent with a population coding scheme, as activity exists in a low-dimensional space with far fewer dimensions than the number of neurons. It is an obvious extension of this work, then, to search for such low-dimensional activity in dissociated cultures.

Here, we combine dimensionality reduction techniques with clustering analysis to demonstrate that spontaneous activity in dissociated rat hippocampal cultures can be interpreted as a low-dimensional population code which develops over time to become both more tightly stereotyped and frequent (Fig 2.1). Specifically, we find that linear combinations of neuronal firing, discovered using principal components analysis (PCA), repeatedly trace out activity “trajectories” with distinct spatiotemporal signatures. These trajectories begin to appear at approximately 8 days *in vitro* (DIV) and persist for several weeks - some remain until 35 DIV.

Additionally, we characterize sub-population activity and find that relatively few neurons consistently contribute to repeated trajectories. This is suggestive of a sparse population coding scheme, which has been documented in the mammalian hippocampus

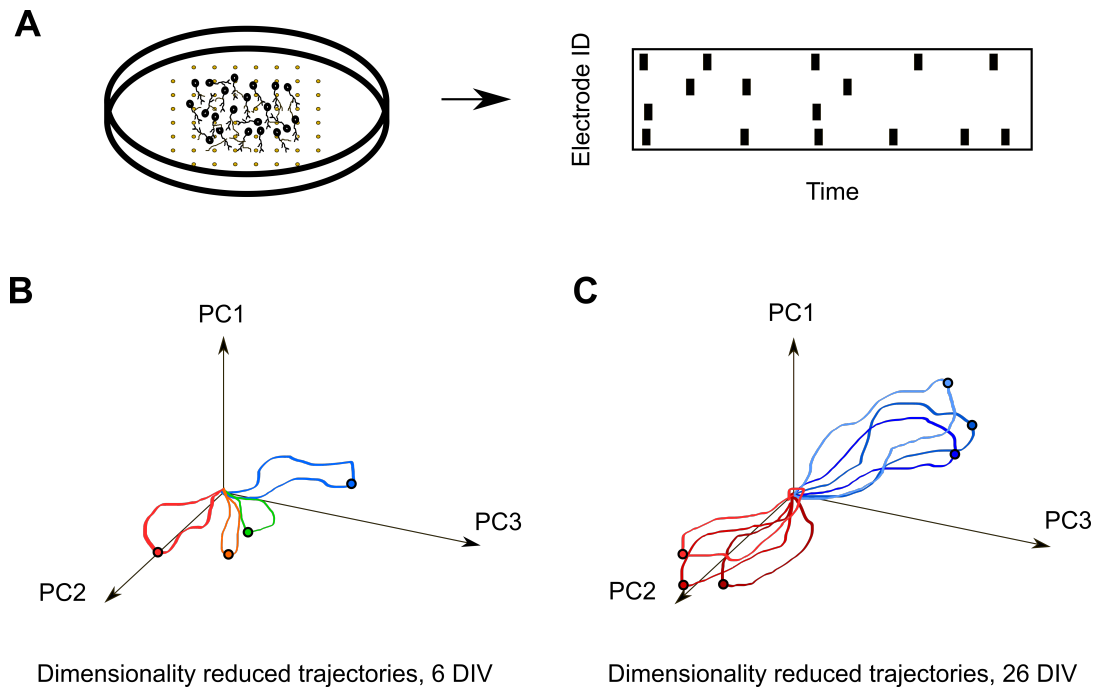


Figure 2.1: Rat dissociated hippocampal neurons develop stereotyped activity patterns over the course of development. **(A)** The electrical activity of rat dissociated hippocampal neurons ($N = 424$ recordings, 1 hr. each) was recorded on a multi-electrode array (MEA) by Timme et al. [37]. Spike sorting was applied to yield binary spike train data (right). **(B)** Applying dimensionality reduction techniques, namely principal component analysis (PCA), to smoothed spike data yielded low-dimensional activity trajectories. The “maximum excursion point”, which is the point on a trajectory in activity space furthest from the origin, is shown for each trajectory. Early in development, these trajectories are random and disorganized. **(C)** After several weeks of development, some stereotyped trajectories are repeated and can be clustered into groups (indicated by trajectory color). This discovery illustrates that dissociated cultures retain some fundamental features of the hippocampus, as this phenomenon resembles memory formation and consolidation.

[38, 39]. However, we also find that, over their entire life spans, cultures exhibit many population-wide synchronized bursts, as has been previously reported [19]. These events recruit many neurons and are distinct from the trajectories we identify, though several classes of bursts in dissociated cultures have also been shown to develop and persist over time [40].

Surprisingly, we find that some sub-populations responsible for the observed low-

dimensional trajectories are highly stable over the course of several days but that others change dramatically between recordings. This variability may result from the lack of external feedback in this system, which may be necessary to fully consolidate activity patterns over long periods of time.

These results illustrate the development of coherent, complex spatiotemporal activity patterns in dissociated hippocampal cultures for the first time. The existence of these patterns suggests that this system is potentially useful as a substrate for studying the development of memory-like functions in artificial biological settings. Since these cultures display a capacity to acquire and retain nontrivial firing sequences (i.e. sequences that do not simply recruit every neuron in an fully synchronized “avalanche”), external electrical or chemical feedback could be used in future experiments to promote and maintain the development of activity trajectories. The development of such feedback systems would provide significant opportunities to further study aberrant hippocampal dynamics, which are believed to underlie Alzheimer’s disease and other types of dementia.

2.2 Methods

Data

The data used in this study were provided courtesy of Nicholas M. Timme, Najja Marshall, Nicholas Bennett, Monica Ripp, Edward Lautzenhiser, and John M. Beggs [37]. The data consist of spiking activity of rat hippocampal dissociated cultures ($N = 28$) observed at semi-regular intervals from days *in vitro* (DIV) 6 through 35, producing 424 recordings in total. The raw electrophysiological data was collected using a Multichannel Systems multi-electrode array with 60 flat electrodes in a square lattice with $30 \mu\text{m}$ diameter and $200 \mu\text{m}$ spacing between electrodes. Recordings were sampled at 20 kHz

for approximately 1 hour. The data was then spike sorted using `wave_Clus`, as described in [41], resulting in binary neuronal spike trains. Fig 2.2 shows summary statistics for this data set.




Figure 2.2: Summary statistics of neuronal recordings are relatively stable over the course of development in dissociated rat hippocampal neuronal cultures. Here, box plots illustrate summary statistics, with the circle representing the mean value across cultures and color indicating day *in vitro* (DIV). Shaded rectangles indicate the boundaries of the top and bottom quartiles. Whiskers indicate the 95% bound. Outliers are shown as colored points. The number of active channels detected by the MEA (left) is relatively consistent over time (mean = 53.9 channels, 1.7 neurons per channel), though early recordings possess slightly more active channels. Likewise, the distribution of firing rates (middle) is stable (mean = 1.9 Hz). However, the number of trajectories detected (Methods) changes over the course of development (right). Between approximately 20 and 30 DIV, cultures exhibit a greater number of trajectories than younger or older cultures. This suggests that cultures possess the most dynamic, biologically relevant behavior during this period.

The data described above were then further modified from their original format before analysis. Spike trains of neurons recorded by the same electrode were collapsed into a single unit, resulting in binary electrode spike trains. This step obscures information about neuron identity, but has been used successfully before to follow the development of functional networks *in vitro* [42], so we adapt it here. In doing so, we follow the activity of what we assume to be the same neurons over many days, as it is unlikely that neurons drift far enough in space to be recorded on a different electrode across different days. Additionally, this step eliminates the potential ambiguities introduced by attempting to

map neuron identities across days using spike waveform profiles.

Electrode spike trains were then smoothed using a Gaussian filter with standard deviation $\sigma = 150$ ms. This value is in line with standard smoothing kernels used for decoding in brain-machine interfaces [43] and is smaller than the average inter-spike interval in this data set, ensuring that it is a reasonable choice of parameter. Filters with $\sigma = 50$ and 100 ms were also used but did not affect the results.

Data Partitioning

Recordings were placed in one of three classes, as determined by their age: young (6-14 DIV), middle (15-25 DIV), and old (26-35 DIV). These classes contained $N = 139$, 147, and 138 recordings, respectively. These boundaries were chosen so that each category would have as close to the same number of recordings as possible. Changing the boundaries by one day in either direction did not change our results.

Trajectory Identification

Neural activity trajectories were identified in activity space using a thresholding method (2.3). A trajectory was identified if the culture-wide smoothed firing activity exceeded 20% of the maximum possible value for the recording, given by $0.2nf_{max}$, where n is the number of neurons identified in the recording and f_{max} is the maximum value of the filter used to smooth the spike train data. We also examined threshold values ranging between $0.2nf_{max}$ and $0.4nf_{max}$, at which point we began to observe a change in our results (Appendix A.1), so we used $0.2nf_{max}$ to be more permissive in our definition of neural trajectories and to examine trajectories that do not necessarily fit the description of synchronous burst or neural “avalanche”.

Then, looking backward in time from when activity exceeded this threshold, a tra-

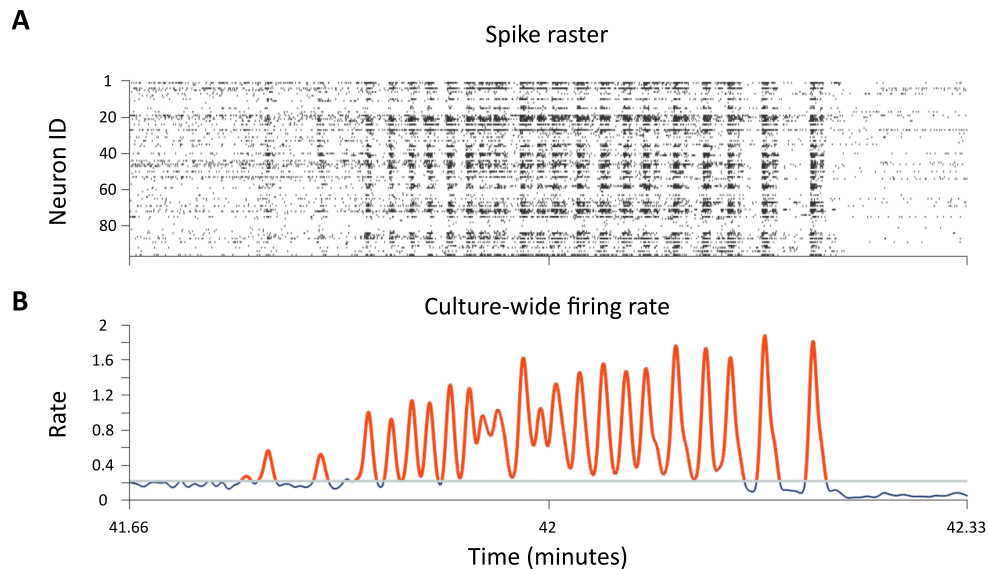


Figure 2.3: Example of trajectory identification. **(A)** Example raster plot of spiking activity recorded from a single culture. Black bars indicate a spike at a given time point. **(B)** Illustration of trajectory identification protocol. Curve shows the instantaneous overall firing rate after smoothing the spike train using a Gaussian kernel with $\sigma = 150$ ms. The grey line indicates the threshold for trajectory identification, which was set at 20% of the maximum firing rate for the recording. Red portions of the firing rate curve represent individual trajectories and periods of below threshold activity are shown in blue.

jectory began at the latest time point where culture-wide activity was below the median value for the recording. Similarly, a trajectory ended when activity first decreased back to the median value after having crossed the threshold. As a result, a trajectory as defined here consists of the time series of firing activity across all electrodes between its start and end time.

For the subsequent analysis, trajectories are referred to in a condensed format. Specifically, each trajectory is represented by the single point at which the sum of culture-wide activity was the largest, which we refer to as the “maximum excursion point”.

Dimensionality Reduction

Trajectories for each recording were transformed with principal component analysis (PCA) using the MATLAB *pca* function. PCA captures the covariance structure of high-dimensional data, yielding a new set of orthogonal basis vectors that maximize the amount of variance captured when the data is projected onto those vectors. Fig 2.4 shows an example of PCA-transformed trajectories projected onto the first 3 principal components. Since principal components are a linear combination of electrode firing activities, each dimension in principal component space can be interpreted as a weighted sum of neuronal activities. Though PCA is simpler than many other dimensionality reduction approaches, the results of this method are easily interpretable and therefore useful for understanding the dynamics of neural population activity as a sum of individual neuronal activities [44, 45].

The composition of an example principal component is shown in Fig 2.4. In subsequent analysis, data were projected onto only the first 5 principal components. This number of components is roughly 10% of the original dimensionality, providing significant compression, and accounts for 80.7% of the variance on average, indicating that we are not discarding an excessive amount of information about trajectories.

Clustering

After transformation by PCA, maximum excursion points of trajectories were initially clustered with density-based spatial clustering of applications with noise (DBSCAN) [46] using the MATLAB 2019a function *dbscan*. Maximum excursion points were chosen as clustering objects because they represent the largest deviation from baseline activity and can serve as a representative “signature” of a trajectory without including potentially noisy activity near the baseline. The DBSCAN algorithm has two free parameters: ϵ ,

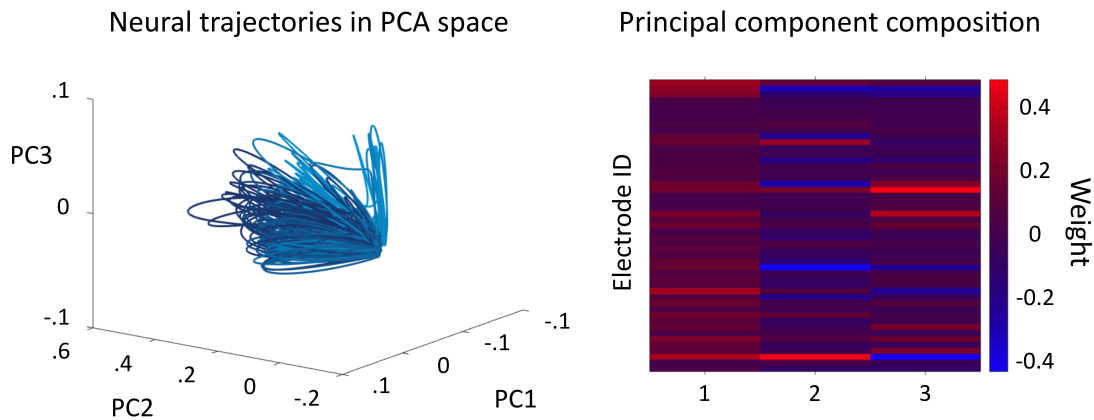


Figure 2.4: Shown on the left are 100 activity trajectories from culture 12 on day 13 *in vitro* projected onto the first three principal components. Each trajectory is indicated by a different color and originates and terminates near the origin. From inspection, it is apparent that many trajectories follow a similar path through PCA space, indicating similar spatiotemporal dynamics at the population level. Notably, this culture appears to exhibit two distinct trajectory classes which are nearly orthogonal to one another. Shown on the right is a heat map of the principal component composition for this set of trajectories. Principal components are column vectors composed of a linear combination of weighted activity values, providing an interpretable illustration of sub-populations which co-vary in their firing behavior. Colors indicate the weight each electrode contributes to a principal component. Positive weights indicate correlations between electrodes while negative weights indicate anti-correlated activity. Here, only three principal components are shown for the sake of illustration, but subsequent clustering analysis uses five principal components.

a sensitivity parameter, and m , the minimum number of points in a cluster. Here, we restrict ϵ to values between 0.01 and 0.02, as smaller values result in nearly as many clusters as points and larger values result in all points being assigned to a single cluster.

Null Model Validation

To validate the significance of the clustering results, we generate permuted versions, or null models, of the observed trajectories for comparison (Fig 2.5). We consider two null models. In the first, spikes are randomly reassigned across electrodes within a time bin, preserving instantaneous firing rate. In the second, spikes are randomly reassigned

across time bins within an electrode, preserving individual electrode firing rate statistics. We apply the methodology as previously described to ten instantiations of these permuted trajectories to determine whether perturbing various aspects of trajectory structure changes the results of clustering.

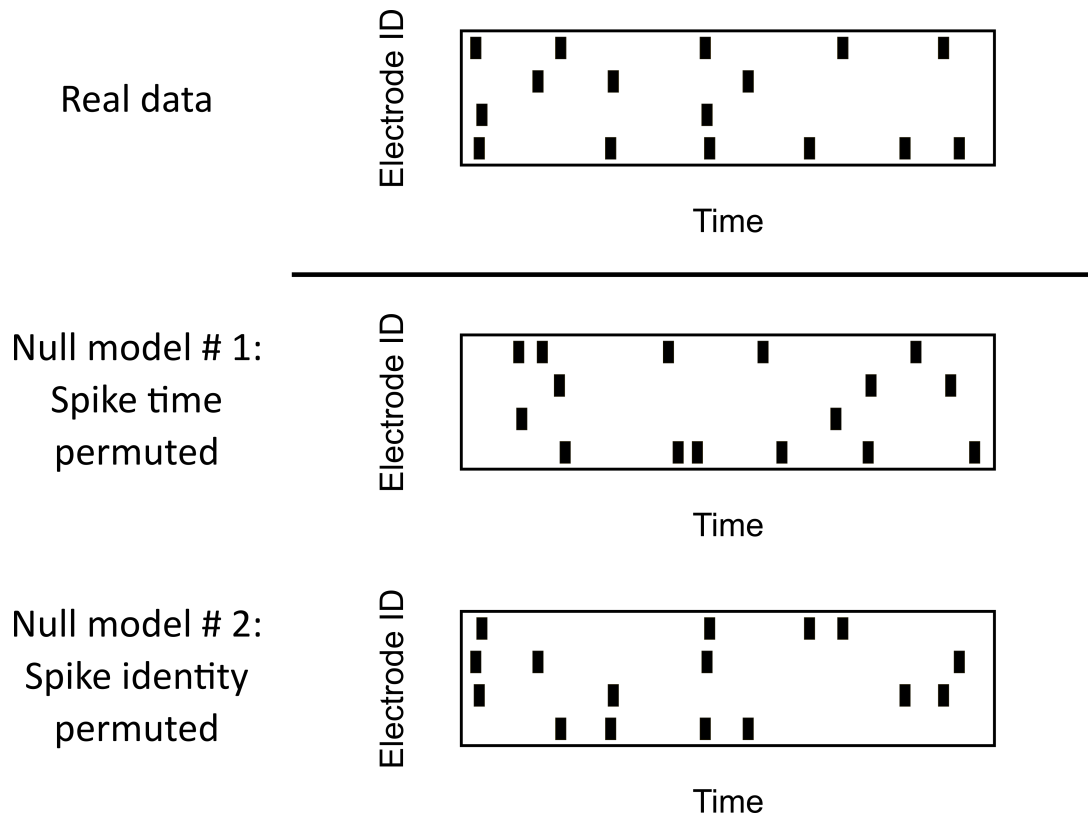


Figure 2.5: Shown here are the null model permutations used to identify neural activity trajectory features. Actual trajectories (top) are shuffled in one of two ways: spike times for an electrode are randomly distributed (middle) or the identity of a spike at a given time is randomly assigned (bottom). When spike times are permuted, the firing rate of each electrode is preserved but culture-wide instantaneous firing rates are disrupted. In contrast, when spike assignments are permuted, culture-wide instantaneous firing rates are preserved, but individual electrode firing rates are not.

2.3 Results

Distinct neural activity trajectories emerge in rat dissociated hippocampal cultures over the course of development

Collecting trajectory clustering statistics reveals that the number of clusters identified varies as cultures age. Young, middle, and old cultures contained 3.64 ± 4.66 , 5.88 ± 5.04 , and 4.14 ± 3.96 trajectory clusters, respectively (mean \pm standard deviation). Differences between distributions for young, middle-aged, and old cultures were evaluated using a two-sided Kruskal-Wallis test. Fig 2.6 illustrates these distributions and indicates that the observed clustering distributions are statistically significantly different between each age condition. In contrast, null models exhibit far fewer clusters, indicating that the permutations used here disrupt the ability of the clustering procedure to discriminate between trajectories.

Trajectories become increasingly separable during development

We find that, over the course of development, trajectory clusters move further apart in activity space. Panel B of Fig 2.6 illustrates this phenomenon; young trajectories (left) are tightly packed, whereas middle-aged trajectories (middle) visit distinct regions of activity space. We define the distance between clusters as the Euclidean distance between cluster centroids in three-dimensional PCA space. The density plots of these between-cluster distances are shown in Fig 2.6.

Though we find that the median between-cluster distances are nearly identical across young, middle, and old categories (0.107, 0.132, and 0.145, respectively), the tails of the distribution differ significantly. This is illustrated by the fact that a two-sided Kruskal-Wallis test differentiates between each pairwise comparison of distributions with $p \ll$

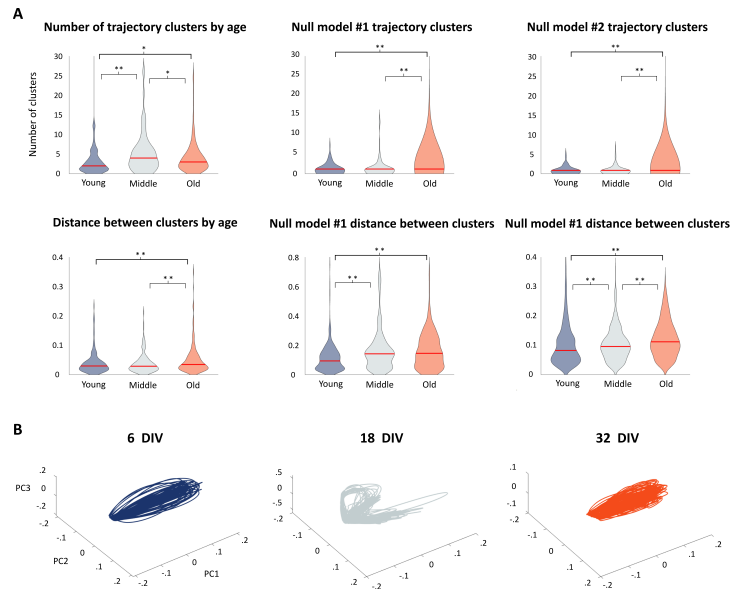


Figure 2.6: Trajectory clusters emerge and consolidate over the course of development **(A)** As dissociated hippocampal cultures age, they exhibit stereotyped firing activity patterns that can be classified into clusters. Shown in the top row are density plots of the number of identified clusters across 424 hour-long recordings, with the median value shown in red. In the observed data (left), middle-aged cultures (15-25 DIV, grey) exhibit the highest number of distinct clusters (median = 4). In contrast, young (< 15 DIV, blue) and old (> 25 DIV, red) cultures possess fewer clusters (median values = 2 and 3, respectively), indicating that the repertoire of firing patterns develops over the first few weeks *in vitro* but eventually decays without external feedback. Across time periods, the distribution of the number of clusters are statistically significantly different (Kruskal-Wallis test, * $p < .05$, ** $p < .01$). The middle and right panels display the clustering statistics for two null model permutations of the data. In each condition, null model permutations of the data have fewer clusters (median = 1), indicating that there is structure in the data related to neuron identity and firing rate. Shown in the bottom row are the density plots of the corresponding distances between cluster centroids. Distances between clusters are determined by calculating the Euclidean distance between cluster centroids in three-dimensional PCA space. Over the course of development, trajectory clusters become better separated compared to those exhibited by young cultures. Notably, null model clusters exhibit a more pronounced trend in separability, suggesting that true firing patterns are restricted to certain regions of activity space, constraining their separability. **(B)** Example of trajectory development over time. Young cultures (left) typically exhibit few activity patterns, whereas middle-aged cultures (middle) possess several distinct patterns. Old cultures (right) display a decrease in the number of activity patterns.

.01.

Notably, null model clusters exhibit similar properties, though at a more pronounced scale. This effect is potentially driven by the fact that actual neuronal firing patterns are constrained by the underlying connectivity of the culture, restricting patterns to a few regions of activity space. In contrast, null model permutations of trajectories may consist of arbitrary activity patterns, allowing them to sample a larger volume of activity space than real cultures.

Co-varying sub-populations that participate in neural activity trajectories are stable over time

To assess whether neural sub-populations become functionally coupled and remain so, we compared the first three principal component eigenvectors over the course of each culture's development using the standard dot product as the measure of similarity, as shown in Fig 2.7. Across cultures, we find that the eigenvector corresponding to the first principal component is remarkably stable, as the dot product of the first eigenvector of temporally adjacent recordings (i.e. the first and second recording sessions, second and third sessions, etc.) has a value of 0.63 ± 0.03 (mean \pm standard deviation). Additionally, the second PC was significantly less stable, as demonstrated by a dot product value of 0.37 ± 0.05 . The third eigenvector was even less self-similar over time, with a dot product value of 0.26 ± 0.03 .

In contrast, both null models exhibit disrupted principal component stability. The first null model, which permutes spike timing, possesses second and third principal components that are almost orthogonal in consecutive recordings, as indicated by a near-zero dot product (Fig 2.7, middle). In addition, the first principal component is markedly less stable than observed in real data, with a wide range in stability over null model runs.

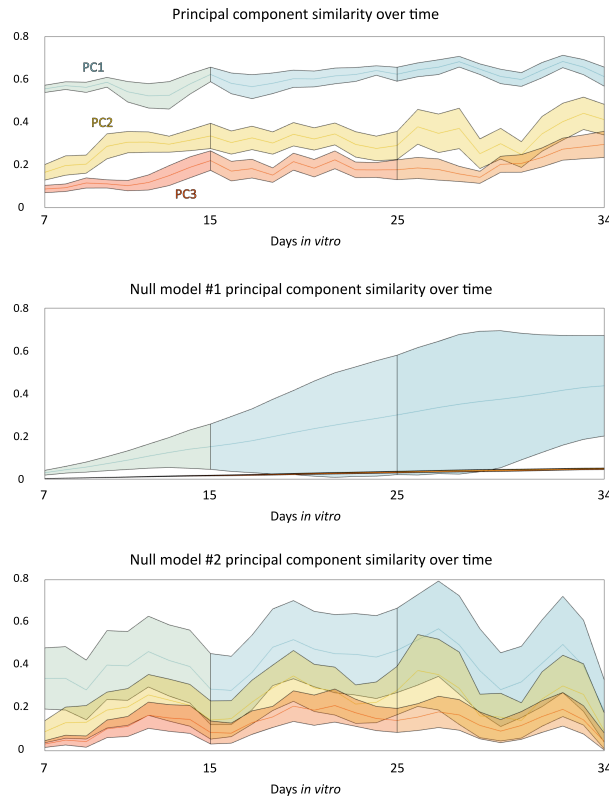


Figure 2.7: Co-active neurons consolidate into stable sub-populations over several weeks *in vitro*. Principal component eigenvectors, which represent a weighted sum of neuronal firing activity, can be compared using the dot product as a measure of similarity. This similarity value ranges from 0, indicating fully orthogonal vectors, to 1, indicating identical vectors. Shown here are the three-day moving average dot products between the first three principal components and their variances (shaded region) on consecutive recordings for the observed data (top) and null models (middle and bottom). In each condition, the first principal component (blue shading) is the most stable over time. In the real data, the first principal component has a mean similarity of 0.63. The second and third components are less consistent but gradually increase in stability. The second principal component has a mean dot product of 0.33 from 7-15 DIV which increases to 0.40 from 25-34 DIV. Similarly, the values for the third component increase from 0.21 to 0.28 over that same period. Notably, the second null model, which permutes spike timing but maintains the firing rate within an electrode during a trajectory, possesses a similar stability profile, potentially indicating that trajectories represent a rate coding scheme. In contrast, the first principal component of the first null model is both less stable and has a higher variance across cultures and the second and third components are almost completely orthogonal across consecutive recordings. This indicates that trajectories possess characteristic spatial structure and do not rely purely on bulk culture-wide firing rate.

The second null model, which permutes spike identity, exhibits stability profiles which are more similar to the real data, but mean stabilities are still lower and the variance is significantly wider than is observed in real data. Taken together, these results suggest that the covarying activity structure observed in the data persists over time more than would be expected if spike timing and identity did not matter.

2.4 Discussion

Low-dimensional neural trajectories in dissociated cultures *in vitro* suggest a self-organizing population code

Dissociated neuronal cultures consist of cells that have been removed from a living brain and homogenized, severing any physical connections between cells that may have formed during development. These systems serve as a simple alternative to organotypic slices removed from developed animal brains, often referred to as *ex vivo* preparations. In either case, these model systems are useful to neuroscientists for many reasons; importantly, they are more accessible to measurement and manipulation than *in vivo* systems, they are simpler than fully connected *in vivo* systems, and they provide realistic biological context that may be missing from *in silico* models.

In contrast to *ex vivo* systems, dissociated neuronal cultures have no developmentally-guided network structure and therefore are subject only to intrinsic cellular mechanisms as they age. Consequently, while slices exhibit complex firing patterns, there is no expectation dissociated cultures should as well. In order to explain the development of coherent spatiotemporal firing patterns, we must consider plausible cellular mechanisms that could result in this phenomenon. Perhaps the most well-studied mechanism relating to network formation is Hebbian learning - often sloganized as “cells that fire together,

wire together.” Hebbian learning via spike timing-dependent plasticity (STDP) has been implicated in the formation of associative memory and navigation reward learning in hippocampus [47, 48], so it is a natural candidate for the trajectory consolidation and maintenance observed here. Additionally, recent *in silico* work shows that simple STDP-like rules are sufficient to entrain an initially randomly connected network of neurons to learn and maintain a limited repertoire of specific low-dimensional spatiotemporal firing patterns [49, 35].

Several *in vitro* studies describe the existence of repeated spatiotemporal firing patterns in dissociated neuronal cultures [50, 19, 51, 52, 40, 53, 54, 18]. However, no work to date has demonstrated that dissociated hippocampal cultures, which have been established as model systems of culture-wide synchronous activity [31, 30], exhibit this type of structured spontaneous activity, which is typically identified in neocortex. Taken together, these lines of evidence suggest that, absent external guidance, isolated hippocampal networks can develop a repertoire of sub-bursting firing patterns in a self-organizing manner. Moreover, this phenomenon may require only a few simple rules, making it an attractive target for study via genetic and electrochemical manipulation.

Dissociated cultures as a model system

Many proposed population coding schemes require covariant activity in subsets of neurons involved in a functional circuit [29, 55, 56, 57]. Sparse, distributed codes provide robustness to both external signal noise and internal transmission noise while also providing flexibility, as individual neurons may serve different roles in multiple subnetworks. The functional roles of coupled subnetworks *in vivo* are well-studied at multiple spatiotemporal scales, providing mechanistic insight into how brains execute complex behavior [33, 32]. However, less is known about the processes which govern the devel-

opment of functional circuitry. Dissociated cultures *in vitro* provide excellent models for studying this problem, as they presumably develop connections in a purely unsupervised manner, disentangling the roles of genetics and intrinsic activity during maturation from externally-provided feedback and manipulation. However, their usefulness is contingent upon whether or not they exhibit activity which could plausibly serve as an analog to *in vivo* dynamics.

While general properties of *in vitro* cultures, such as the development of synchronous, culture-wide bursting activity after an initial period of sparse, asynchronous firing, are well-documented across cell types and recording modalities [22, 23, 19], little is understood about to what extent *in vitro* dissociated hippocampal cultures serve as useful models for spontaneous activity in microcircuits *in vivo*, which exhibit complex, repeated, multi-neuron firing patterns. Attractor-like patterns, replay activity, and chaotic trajectories have all been proposed to play roles in brain function [58, 56, 59, 60]; if similar dynamics exist in *in vitro* cultures, researchers would be able to probe fundamental properties of complex neural activity in these simplified systems.

The results presented here provide crucial context for future *in vitro* studies, as they demonstrate that such systems are at least partially appropriate models for developing circuits *in vivo*. This remains an unresolved issue and is a weakness of *in vitro* studies which attempt to generalize claims about network activity properties observed across contexts. As with other studies, however, there are limitations to how broadly applicable the results presented here may be. First, changes in preprocessing methodology may introduce slight differences to downstream analysis. In this case, spikes were assigned to an electrode rather than individually identified neurons. This makes longitudinal comparisons of principal component eigenvectors possible but obscures information about cell-level activity. In addition, the number of active electrodes can vary across recordings from dissociated cultures. This illustrates the inherent stochastic nature of development

in cultures of this type, suggesting that researchers produce many technical replicates to ensure a representative sample of outcomes.

This work suggests several interesting avenues of investigation. First, it would be a significant advance in the field of *in vitro* research if the types of trajectories identified here could be trained, initiated, or ablated by pharmacological, genetic, or electrical manipulation, as such functionality would make these systems more valuable as model microcircuits. In addition, if these trajectories can be modified through intervention, they could be used as appropriate model systems for hippocampal replay, memory formation and consolidation, and potentially even brain-machine interfaces.

Stability of neuronal activity patterns

Here, we find that co-varying sub-populations do exist within developing *in vitro* cultures but that these sub-populations exhibit differences in composition and stability over time. Principal components analysis reliably identifies a small number of co-varying sub-populations which explain a significant amount of the variance observed in hour-long recordings. This demonstrates that dimensionality reduction successfully identifies redundancies in the culture-wide activity in the form of co-varying activity, allowing for compression of the activity into a simplified, low-dimensional form. Importantly, since each principal component eigenvector is a linear combination of electrode activities, we can retain information about the contribution of specific electrodes to sub-population dynamics. It should be noted that, while many nonlinear dimensionality reduction techniques exist and have been used to uncover the dynamics of behaviorally relevant neuronal firing patterns [44, 61], linear methods such as principal components analysis are often sufficient to reveal complex co-activation patterns [45, 32]. Since population codes are common across brain regions and model systems, PCA is a reasonable first choice of di-

mensionality reduction technique, though other techniques may reveal further properties of *in vitro* dissociated cultures.

Due to the abundance of evidence that coordinated population activity plays a role in the execution of cognitive and motor tasks *in vivo*, the presence of non-trivial population dynamics *in vitro* is a prerequisite for treating these systems as appropriate models. Many computational studies treat connectivity in dissociated cultures as random, which may limit the range of activity patterns the systems can produce. It is also possible that the character of coordinated population activity changes over the course of development in dissociated cultures, making it necessary to understand not only if it is present but also when and to what degree.

From a computational standpoint, questions still remain about the robustness and flexibility of population codes as they are actually implemented in biological systems. Specifically: how stable are firing patterns over time? How much variability exists within firing patterns corresponding to the same function? How quickly can this pattern-function mapping be altered? In this vein, several studies have characterized the functional repertoire of neuronal populations by performing similarity analysis on sequences of spike activity [19, 51, 40, 53, 54, 18]. These studies illustrate that several types of repeated spiking patterns emerge in dissociated cortical cultures over approximately two weeks of development. One set of studies relies on template matching algorithms which stipulate that only the order of spikes in activity patterns must be consistent in order to be considered a repeated sequence, largely disregarding the temporal dynamics and distribution of spikes across the culture. Implicit in such analysis is the proposition that these systems utilize a temporal coding strategy that cares only about spike order. Another set considers spike timing more explicitly, assuming a temporal coding scheme. Finally, another set restricts analysis to bursting activity, using the total culture-wide firing rate as the signature of complex, distributed activity. In order to generalize these findings to

in vivo systems, this type of analysis assumes a rate coding scheme.

The observation that our null model which permutes spike timing while retaining firing rate statistics within electrodes possesses a similar sub-population stability profile as the observed data suggests that the trajectories we observe *in vitro* are a type of distributed rate coding scheme (Fig 2.7). Additionally, the fact that the null model which permutes spikes across electrodes while preserving instantaneous culture-wide firing rates does not possess the same profile indicates that such trajectories are neuron-specific. This is a useful property for model systems to possess, as it is a close analog to *in vivo* systems of interest [33, 34]. Our results indicate that dissociated cultures can serve as a valuable model system, provided that adequate technology for recording and manipulation are accessible.

Chapter 3

Discovery and Characterization of Duration Dependence in Fruit Fly Grooming Syntax

3.1 Introduction

Sequential animal behaviors are often composed of repetitions of simple subroutines. In driving motor action execution, nervous systems must integrate sensory information with internal priorities and dynamics. Both external and internal conditions contribute to sequence organization, but their respective weights are unknown. At one extreme, purely sensory-driven, reflexive processes produce actions, such as larval escape sequences, solely based on recent sensory input [62]. At the other, internally generated and maintained nervous system dynamics, such as those found in the crustacean stomatogastric ganglion, may produce behaviors which proceed irrespective of changing external conditions [63, 64].

Grooming, a common behavior across species, confers social and survival benefits [65,

66] and provides a rich source of data for discovering rules that organisms use to produce behavioral sequences. The vocabulary of grooming consists of the possible actions that may be executed and has been cataloged in mice and several species of flies [67, 68, 69]. However, the rules for grooming action sequence organization, or syntax, remain poorly understood. In *Drosophila melanogaster*, or fruit flies, these sequences are variable between individuals and within individuals across grooming sessions, suggesting that flies use non-deterministic rules to make grooming decisions. Because of the complex structure of these sequences, analysis of *D. melanogaster* grooming can reveal distinct rules of sequence generation at different temporal scales: the long time scale of bulk behavioral progression, the intermediate scale of grooming motifs, and the short scale of individual grooming bouts.

On long time scales, *D. melanogaster* grooming in response to exposure to an irritant is well-described as a process that typically progresses from an early phase characterized by preferential grooming of anterior body parts, such as eyes, to a later phase which exhibits heightened proportions of walking, abdomen cleaning, and wing cleaning. This progression occurs gradually over the course of many minutes. Previously published computational models which utilize either hierarchically structured suppression or graded sensory gain exhibit gross features of this progression, namely a gradual transition from early anterior-heavy grooming to a later quasi-steady state featuring increased posterior grooming and walking levels [70].

At an intermediate temporal scale, grooming is organized in units that we refer to as motifs. We identify two classes of motifs (anterior and posterior), which are named for the set of legs used to execute them (front and back, respectively). Motifs consist of consecutive alternations between body-directed grooming bouts and leg rubbing bouts. Bouts are defined as sustained periods of a single grooming action (e.g. head cleaning or wing cleaning). Grooming bouts occur at the shortest time scale we consider here, as

individual bouts typically last somewhere between 150 ms and 2 s. Flies use the same pair of legs to execute bouts within a motif, allowing them to transition between within-motif actions easily. Specifically, an anterior grooming motif consists of consecutive alternations between bouts of head cleaning and front leg rubbing. A posterior motif consists of alternating bouts of abdomen cleaning, wing cleaning, and back leg rubbing.

Analysis of freely-behaving flies reveals behavioral structure at multiple time scales, with coarse-grained descriptions being sufficient to describe long-term trends and higher resolution descriptions providing increased predictive power at shorter time scales [71]. This finding suggests that there is value in examining sequential data at several levels of temporal resolution. Additionally, it suggests that treating behaviors as both continuously varying and discretely separated in the same analysis can yield more insight than considering either purely continuous or discrete models. Currently, we do not know how flies integrate sensory information with internal states in order to produce multi-scale grooming behavior. It is possible that long time scale grooming trends are governed principally by changing sensory conditions, as grooming results in the removal of irritant over time. Several studies also provide evidence that sensory input is sufficient to drive grooming behavior in a reflex-like fashion at short time scales [72, 73]. Here, we use statistical models to characterize grooming syntax at each of the temporal scales mentioned above.

Several classes of Markov models, in which the probability of an event occurring is contingent upon previous events, have been used to describe factors involved in non-deterministic decision-making in various animals. Different Markov models vary in their structure and number of parameters but they each require a well-defined state space. Navigation and foraging have been described using Markov models, as these sequential behaviors can be decomposed into subroutines which can be categorized using easily observable dimensions such as direction and velocity [74, 75, 76]. These models are ap-

plicable to common behaviors exhibited by *D. melanogaster* as well, such as courtship and locomotion [77, 78, 79]. Recent work from Tao et al. uses a Hierarchical Hidden Markov Model (HHMM) to analyze fruit fly locomotor behavior in the presence of odors [80]. They find that, although locomotor behavior is non-stereotyped, it can be decomposed into stereotyped units, making it suitable for analysis within a Markovian framework. Berman et al. [81] use high-order Markov models to describe the behavior of freely-roaming *D. melanogaster*, suggesting the applicability of such models to behavioral subsets identified in their analysis, including grooming behavior.

It should be noted that non-Markovian dynamics have also been identified in animal vocal sequence production, suggesting that other classes of models may be useful for describing behavioral generation [82, 83, 84]. In fact, Berman et al. [81] report that, although a Markovian framework is useful for illustrating some features of behavioral transitions, they observe long time dependencies in their data that cannot be captured by their Markov models.

Here, we use our in-house Automatic Behavior Recognition System (ABRS, [85]), which can classify different grooming movements from videos of flies covered in dust, to generate a large data set of ethograms - records of cleaning actions over time - from wild-type flies removing dust. We analyze more than 40 total hours of video from over 90 flies (Fig 3.1).

First, this data set allows us to quantify behavioral trends at several temporal scales on a larger data set than has been previously described. Fig 3.2 provides a schematic overview of the temporal scales we analyze here. Notably, we observe the progression from anterior to posterior grooming that has been identified by Seeds et al. [70]. We also observe a strong correlation between body-directed and leg-directed cleaning actions within motifs. That is, the amount of head cleaning and front leg rubbing are extremely strongly correlated over the entire course of grooming, as are the amount of back leg

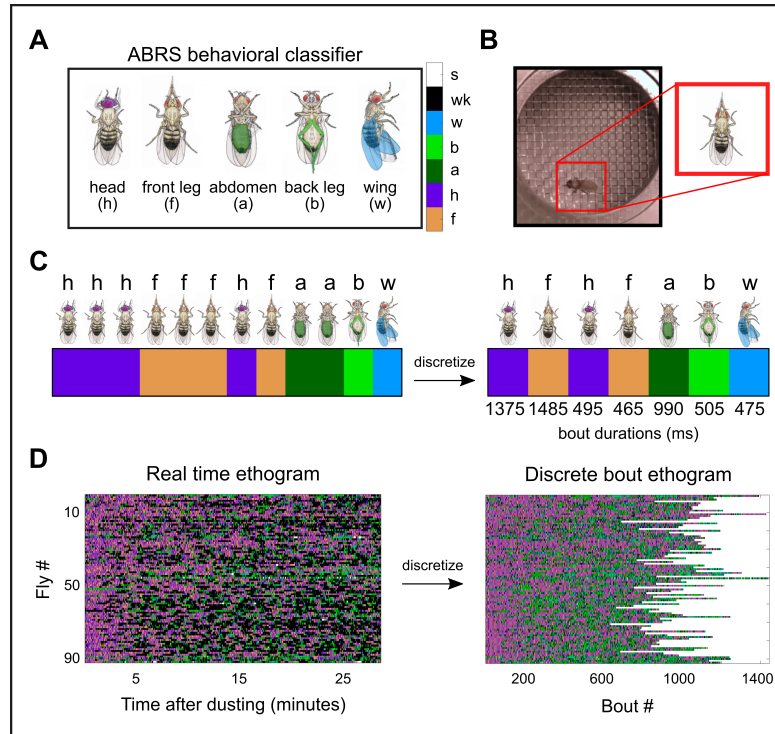


Figure 3.1: Applying the ABRs classifier to video data generates large-scale ethogram data sets. A: The ABRs classifier is trained to annotate video data with labels corresponding to one of five grooming actions observed in *D. melanogaster* or two non-grooming actions. Color bar (right) displays the corresponding color code used for visualization of grooming action sequences. B: Flies are covered with dust, placed in a chamber, and recorded from above. Video data is then passed through the ABRs classifier, which locates the fly in each frame and applies a behavioral label. Here, the fly is shown performing a front leg rub. C: After classification, video data can be represented by a vector. Each vector element represents one frame and contains a behavioral label, visualized here using a color code. Without loss of information, the real time ethogram (left) is recast as a discrete bout ethogram, defined by discrete transitions between individual bouts, which are labeled by both action and duration. This results in a vector (right) containing both an ordered list of grooming action labels (color) and grooming action durations (number). For illustrative purposes, we choose bout durations, shown below the discrete bout ethogram, that are representative of the observed distributions. We refer to instances of sustained identical grooming actions as bouts. D: Real time ethograms for 92 dusted flies (left) are shown as a matrix containing 92 rows and 50,000 columns. After discretization, ethograms contain variable numbers of bouts but retain information about bout order and bout duration (right).

rubbing and the sum of abdomen and wing cleaning. Finally, we observe a relaxation into steady state-like behavior after approximately 13 minutes of grooming.




Figure 3.2: **Schematic overview of the multiple time scales present in *Drosophila melanogaster* grooming ethograms.** Grooming ethograms possess at least three distinct time scales. The grooming progression (top), which takes many minutes to observe in its entirety, occurs at the longest time scale we consider here. We refer to this phenomenon as a progression because we observe that, on average, flies begin grooming with anterior-oriented actions and end with posterior-oriented actions and more walking. At an intermediate time scale, ethograms consist of grooming motifs (middle). Motifs are composed of consecutive grooming bouts which use the same pair of legs (either anterior or posterior). At the shortest time scale, we observe individual grooming bouts (bottom), which consist of sustained grooming movements that range from several hundreds of milliseconds to a few seconds in duration.

Next, we use a set of Markov models to characterize probabilistic rules governing grooming action transitions. We use a Markovian framework as a first approximation in order to discover relevant features of grooming, since Markov models are simple and carry few assumptions. As such, these models serve as a tool for grooming sequence feature identification rather than as a full explication of grooming decision-making factors. Previous work has indicated that such analysis can reveal temporal relationships in sequential behavioral data, indicating that, even if the behavior being analyzed is not fully Markovian in reality, Markov models can still be useful as a descriptive exploratory tool [77, 78, 79, 81, 80].

The simplest of our models, a first order discrete time Markov chain, highlights the presence of grooming motif structure. Subsequent models incorporate temporal informa-

tion by considering grooming bout durations. By comparing these models to statistical null model hypotheses, which shuffle bout order or bout duration, we discover the contribution of bout duration to grooming sequence structure at short time scales.

Though we rely on a Markovian framework to identify features of grooming syntax, our analysis does not preclude the possibility of non-Markovian dynamics in fruit fly grooming. Indeed, our data exhibits two notable nonstationarities. First, we observe an overall trend in grooming proportions, as has been reported previously [70], wherein flies favor anterior grooming immediately after being dusted but favor posterior grooming and walking after approximately 10-15 minutes. Second, grooming transition probabilities change slightly over the course of grooming (though they are unexpectedly stationary over time), indicating that a purely stationary model is an oversimplification. To account for this, we extend our analysis to include a simple nonstationarity which captures sequence progression more closely than a first order Markov chain.

We find that a time-varying Markov renewal process (MRP) which incorporates bout duration dependence recapitulates the observed grooming structure at long, intermediate, and short time scales. This model is partially Markovian, as transition probability matrices dictate state transition dynamics. However, the model also includes renewal process dynamics, in which the duration of the ensuing bout is determined using a random draw from the empirically observed bout duration distribution. This model is nonstationary due to the fact that Markov transition probabilities change over time. This model recapitulates both the observed bulk grooming progression statistics and the fine-grain temporal structure of grooming motifs, suggesting that sensory drive can modulate internal programs in a simple manner to prioritize different grooming strategies in different contexts.

Our models clarify the syntax of grooming, demonstrating that, at the scale of individual bouts, transitions are dictated by both the identity and duration of the pre-

ceding bout. While previous studies have established the non-random nature of fruit fly grooming [67, 70, 77, 78, 79, 81], this work reveals the previously unknown contribution of grooming bout duration to sequence structure, suggesting a role for sensory input-independent decision-making. Finally, a model that includes duration dependence at the scale of individual grooming decisions and a linear progression to mimic dust removal is able to generate synthetic ethograms that closely resemble observed data at the intermediate scale of motifs and the long scale of the grooming progression.

The discovery of duration dependence in grooming sequence structure suggests the existence of neural circuits in *D. melanogaster* that can generate patterned behavior irrespective of sensory drive, as transition rules and bout durations remain nearly constant even as the dust distribution varies over the course of grooming. Additionally, the tight correlation between body and leg-directed actions within motifs demonstrates for the first time the existence of duration dependence at an intermediate temporal scale. Finally, the model improvements provided by the inclusion of a simple nonstationarity suggest that sensory drive plays a role in guiding the implementation of internal programs over long time scales rather than serving a purely reflexive action selection function at short time scales. These discoveries provide the groundwork for future experiments and analysis of the neural underpinnings of duration-dependent action selection and sensory-driven modulation of decision-making in grooming and related behaviors.

3.2 Methods and Models

Ethogram generation

Drosophila melanogaster (N = 92) were recorded at 30 Hz for 27.8 minutes after being uniformly covered in Reactive Yellow 86 dust (Organic Dyestuffs Corporation, Concord,

North Carolina). To generate labeled ethograms, we apply the ABRS classifier to video (<https://github.com/AutomaticBehaviorRecognitionSystem/ABRS>). Following this, real time ethograms are converted to discrete bout ethograms and de-noised (Fig 3.1).

Briefly, the classifier is trained on grooming data using a combination of supervised and unsupervised protocols in several steps. First, a Fourier transform is performed on pixel intensity data across 17-frame sliding windows (~ 0.5 s), resulting in two-dimensional spectra referred to as space-time images (ST-images). ST-images are subsequently modified by Radon transformation, which can be used to represent 2D shapes in a rotation and translation-invariant manner. Fast Fourier transformation is then applied to the Radon transformed ST-images, resulting in position and orientation-invariant ST-images [86].

ST-images are then decomposed using singular value decomposition (SVD) and clustered using linear discrimination analysis (LDA) with human labeled behavioral categories. Clusters identified using LDA correspond to five predefined grooming behaviors (f = front leg rubbing, h = head cleaning, a = abdomen cleaning, b = back leg rubbing, w = wing cleaning) and two non-grooming behaviors (wk = walking, s = standing). After classification, each video is represented by a 50,000-element vector, with each element containing a label denoting the behavior identified in the corresponding frame. Classifier performance and validation are shown in Appendix B.1.

To generate discrete bout ethograms, each 50,000-element vector is transformed into a two-dimensional vector in a series of steps. The first row of this discrete bout ethogram contains the bout identity, or the label corresponding to one of the seven behaviors described above. The second row contains the bout durations, or the amount of time spent in the corresponding behavior before transitioning to another behavior.

Consecutive frames containing identical actions are consolidated into one bout identity entry. Then, the corresponding bout duration entry is calculated by counting the

number of consecutive identical frames and converted into seconds by dividing by the video frame rate (30 Hz). As a result, the number of entries in discrete bout ethograms corresponds to the number of observed bouts, which varies between flies. However, the sum of bout durations for each ethogram is identical (27.8 minutes).

After ethogram generation, behavioral bouts persisting for less than 167 ms are deemed artifacts and deleted. This threshold is chosen because manual inspection of video data suggests that individual leg sweeps occur at approximately this time scale. This is also the approximate average transition time between behavioral states described by Berman et al. [71]. Deleted bouts account for 1.03% of total recording time. In subsequent analysis, we also neglect standing bouts, as they account for $< 1\%$ of total recording time.

Data binning

In order to include bout duration as a data dimension, we introduce a binning scheme. To implement this change, the continuous-valued entries in the second row of discrete bout ethograms are collapsed into duration category labels. Grooming bouts are classified into three categories (short, medium, and long) based on the duration of the associated bout. Although our results are consistent even when using more than three duration categories (Appendix B.2), we limit analysis to three categories for the sake of simplicity.

Bin edges for these categories are determined independently for each grooming action (f, h, a, b, w) and walking (wk), as each action possesses a unique duration distribution. Bin edges are chosen such that each bin is populated by an equal number of samples (i.e. the total number of actions labeled “short” is equal to the number of “long” actions) and are depicted in Fig 3.3.

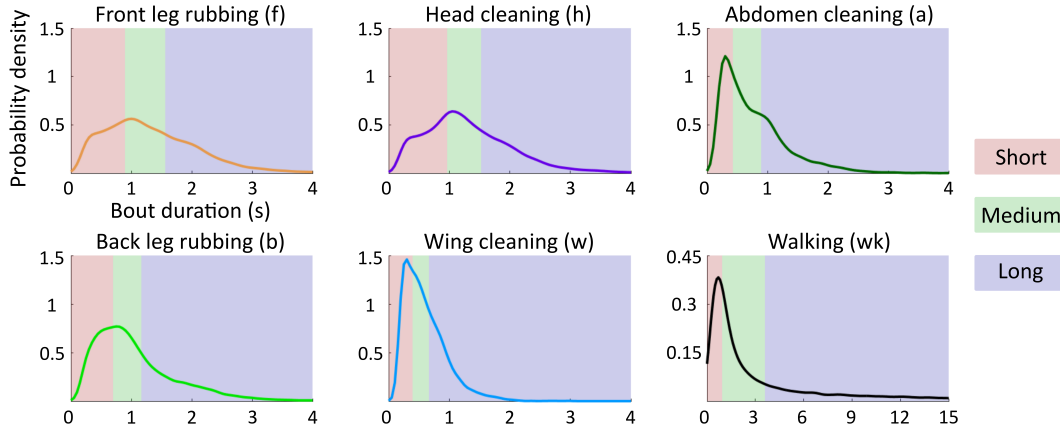


Figure 3.3: **Binning grooming bouts by duration allows for analysis of temporal dependence.** Ethograms are modified by introducing a binning scheme based on grooming bout duration. Here, we show the binning strategy used to create three duration categories for grooming actions. Grooming actions are classified as short, medium, or long bouts, indicated by the red, green, and blue regions of the probability density functions, respectively. Region boundaries are chosen such that each duration category contains an equivalent number of samples. Notably, anterior motif actions (f, h) possess strikingly similar duration distributions, suggesting coupling between these actions.

Markov chain analysis

Markov processes are a class of probabilistic models that describe transitions between states of a system in terms of previous states. The simplest version obeys the Markov property, which declares that the next transition between states depends only on the current state. As such, long-term history dependence can be neglected when considering Markov processes, greatly simplifying the models. In discrete time, the dynamics of a first order Markov chain are described by the equation

$$\mathbf{x}_{t+1} = \mathbf{x}_t \mathbf{M}, \quad (3.1)$$

where \mathbf{x}_t denotes the state probability vector at time t . The entries of \mathbf{x}_t represent the probability of the system being in the state corresponding to that entry at time t . \mathbf{M} is the matrix of state transition probabilities such that entry \mathbf{M}_{ij} is the probability of

transitioning from state i to state j . The vector \mathbf{x}_t will have dimensionality corresponding to the number of possible states in the system. In this case, there are either six, twelve, or eighteen grooming states, depending on whether grooming bout duration is incorporated into the model. Given sufficient data about state transitions in a real system, it is possible to generate a maximum likelihood estimate (MLE) transition probability matrix, denoted by $\widehat{\mathbf{M}}$, which provides the best estimate of the transition probabilities governing the observed transitions. The entries of $\widehat{\mathbf{M}}$ are given by

$$\widehat{\mathbf{M}}_{ij} = \frac{n_{ij}}{\sum_{u=1}^v n_{iu}}, \quad (3.2)$$

where n_{ij} is the total number of observed transitions from state i to state j . The summation over n_{iu} is a normalization factor which counts the total number of transitions from state i to any other state, indexed by u , and v denotes the total number of possible states.

First order Markov chain maximum likelihood transition probabilities are determined for each ethogram as described in Eq 3.2. Probabilities are determined for ethograms with either one, two, or three bout duration bins. The resulting matrices represent population-wide average transition probabilities. We also calculated transition matrices for individual flies, examples of which are shown in Appendix B.3.

Null model hypothesis generation

Statistical null models, which randomize isolated features of data while maintaining overall statistical features, provide a basis for hypothesis testing to determine the significance of results obtained from experimental data. To assess the contributions of bout identity and bout duration to sequence structure, we generate null hypotheses which independently randomize these features of the ethogram data while preserving bout du-

ration frequency distributions within grooming action types. Null hypothesis transition probabilities are analytically tractable in the limit of infinite permutations, so we use exact formulas to calculate matrix entries (Appendix Table B.1). A representative example of each permutation on grooming bout order and duration is shown schematically in Fig 3.4. Our null hypotheses are as follows:

1. Duration permuted: This null hypothesis preserves action order information, but transition matrix entries are independent of duration.
2. Order permuted: This null hypothesis preserves action duration information, but transition matrix entries are independent of action order.




Figure 3.4: **Statistical null model hypotheses (schematic)**. By shuffling bout duration or bout order and comparing the resultant null hypothesis transition matrices to the maximum likelihood estimate, we identify which features are statistically significant in the data and contribute to sequence structure. Shown here are schematics of the permutations we perform on the observed ethograms to generate null hypothesis transition matrices. In the duration permuted hypothesis, bout order is preserved but duration categories are randomly permuted (h = head, wk = walking, f = front leg, s = short, m = medium, l = long). In the order permuted hypothesis, bout order is permuted but the durations associated with each bout remain the same.

The resulting model fit is determined using the Bayesian Information Criterion (BIC),

defined as

$$\text{BIC} = \log(n)k - 2\log(L), \quad (3.3)$$

where n is the number of data points observed, k is the number of free model parameters, and $\log(L)$ is the maximized log-likelihood value. For first order Markov chains,

$$\log(L) = \sum_{t=1}^{T-1} P(\mathbf{x}_{t+1} | \mathbf{x}_t), \quad (3.4)$$

and

$$P(\mathbf{x}_{t+1} | \mathbf{x}_t) = \widehat{\mathbf{M}}_{ij}, \quad (3.5)$$

where i is the row index for the behavior observed at time t and j is the column index for the behavior observed at time $t+1$.

The BIC provides a metric for evaluating model goodness-of-fit by penalizing models which use a large number of free parameters (the first term in Eq 3.3) and rewarding models which accurately predict unobserved data (the second term in Eq 3.3). Thus, models with low BIC values possess more explanatory power than models with higher BIC values but still avoid overfitting. BIC is similar to a related metric, Akaike's Information Criterion (AIC), but penalizes models with more free parameters more severely than AIC in cases where the number of data points is much larger than the number of free parameters.

Synthetic ethogram generation

To evaluate the ability of our models to generate sequences that are similar to our observed data, we use them to create synthetic ethograms. Synthetic sequences are generated using Monte Carlo methods within a Markov renewal process (MRP) framework. In an MRP, state transitions are governed in a manner identical to Markov chains (i.e. using a transition probability matrix). However, the amount of time spent in each state is determined in parallel by a renewal process, in which action durations are drawn from a known distribution. Here, we use the empirically observed duration distributions shown in Fig 3.3.

3.3 Results

Drosophila melanogaster grooming behavior progresses from anterior to posterior motifs on long time scales

After exposure to an irritant, *D. melanogaster* grooming begins with an initial phase of mostly front leg rubbing and head cleaning (f, h). We refer to these behaviors as anterior motif grooming actions. After roughly 13 minutes, flies reach an approximate steady state consisting of heightened amounts of abdomen cleaning, back leg rubbing, and wing cleaning (a, b, w). Together, these actions constitute the posterior motif. Flies also exhibit an increased proportion of walking (wk) during the late phase of grooming. Our result, visualized in Fig 3.5 corroborates a similar finding reported by Seeds et al. [70] but utilizes a much larger, automatically annotated data set.

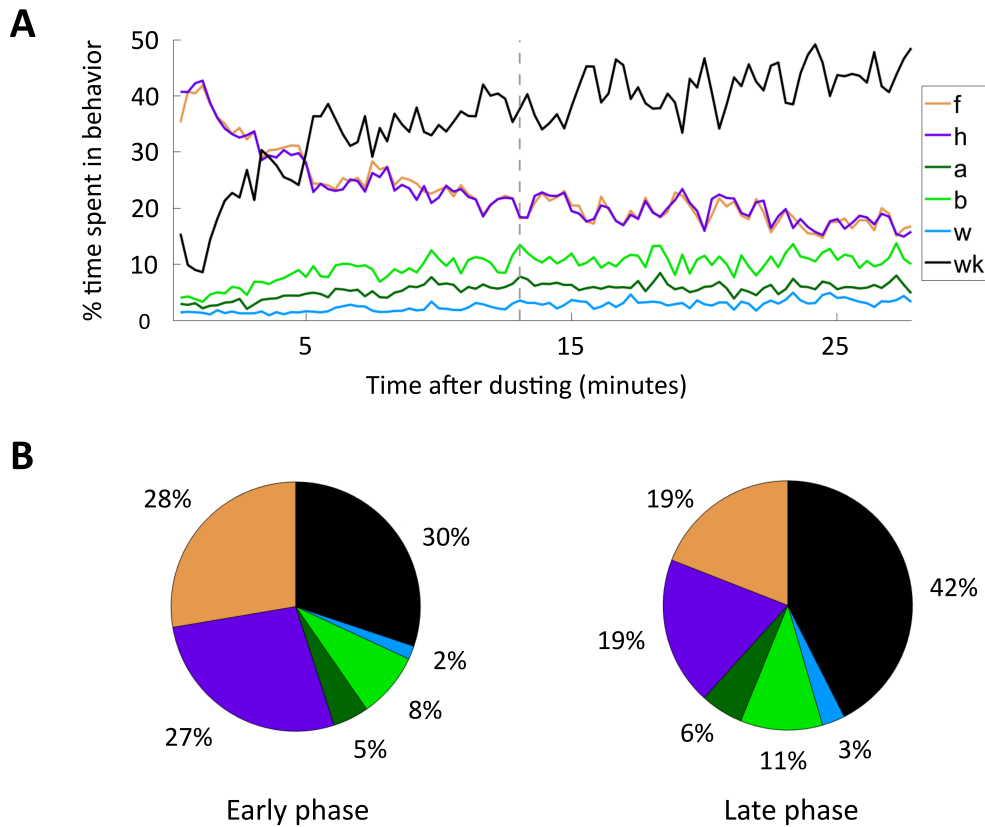


Figure 3.5: *Drosophila melanogaster* grooming progresses from anterior motif to posterior motif actions. **A**: Average grooming progression of *D. melanogaster* population ($N = 92$) after exposure to an irritant. Colored lines show the percentage of time spent in each action (f = front leg, h = head, a = abdomen, b = back leg, w = wing, wk = walking) across a 16.7 second sliding window (500 frames). Front leg rubbing and head cleaning proportions track closely with one another for the entire course of grooming. Additionally, back leg rubbing tracks closely with the sum of abdomen and wing cleaning. Dashed gray vertical line indicates the boundary between early and late grooming phases, which occurs after approximately thirteen minutes, or nearly half of the recording duration. **B**: Proportion of time spent in each action during the early phase (left) and late phase (right) of grooming. During the early phase, *D. melanogaster* spend the majority of time performing anterior grooming movements. In the late phase, they spend a relatively larger proportion of time engaged in posterior grooming movements and walking.

Grooming bouts possess characteristic duration structure

We find that grooming bouts possess action-specific duration distributions. Fig 3.3 depicts this result. We find that anterior motif grooming actions possess remarkably

similarly shaped action duration profiles, with peaks in their probability functions near 1 s. In contrast, posterior motif actions tend to be shorter. Most notably, abdomen and wing cleaning exhibit sharp distribution with peaks near 250 ms, whereas back leg rubbing exhibits a smoother distribution which more closely resembles those of anterior motif actions.

Interestingly, bouts durations are not normally distributed, suggesting a mechanism other than random generation of durations around a mean value. Instead, anterior motif duration distributions possess “shoulders” on each side of the central peak, suggesting that dividing actions into bout duration categories is not an unnatural distinction. Moreover, these “shoulders” are relatively broad due to the fact that grooming bouts consist of many quanta of individual leg sweeps (in the case of body-directed actions) or rubs (in the case of leg-directed actions), which are cyclical in nature and last approximately 150 ms per cycle. Due to video frame rate resolution limitations, we do not consider this extremely fast behavioral scale here. In this analysis, we highlight results which divide grooming actions into three duration categories in order to provide an intuitive demarcation between “short”, “medium”, and “long” actions. However, it is important to note that our conclusions are not strictly dependent on this choice of categorization scheme, as our results hold for higher numbers of duration categories as well (Appendix B.2).

Intra-motif grooming actions are tightly coupled

We find that the vast majority of transitions occur within grooming motifs. Approximately 88% of transitions from anterior motif grooming bouts are to another anterior motif behavior, while about 70% of transitions from posterior motif bouts are to another posterior motif behavior. This tight intra-motif coupling is illustrated in Fig 3.5, as the

amount of front leg rubbing tracks the amount of head cleaning extremely closely. The amount of back leg rubbing also tracks the sum of abdomen and wing cleaning very closely. This result is not affected by the choice of sliding window, as we observe a similarly strong correlation even when using a 1.67 s (50 frame) window instead of the 16.7 s (500 frame) window used to illustrate the progression shown in Fig 3.5.

We find that this strong coupling emerges from the structure of individual grooming motifs. Fig 3.6 provides an illustrated example of the structure of an anterior motif. We find that, across flies and independent of the time after grooming, anterior motifs contain nearly perfectly correlated amounts of body and leg-directed grooming. This strong correlation is somewhat surprising given that consecutive anterior grooming bouts exhibit a non-linear relationship (Appendix B.4).

This tight intra-motif coupling is also reflected in the structure of Markov transition matrices fit to our behavioral data (ethograms). We fit a first order discrete time Markov chain to our data and the resulting matrix, denoted by $\widehat{\mathbf{M}}$ (Eq 3.2), contains the maximum likelihood estimates of the probabilities of transitioning from one grooming action to another. Fig 3.7 depicts the structure of these transitions in both network and matrix form, illustrating the transition probabilities used to define the Markov chain in Eq 3.1. In Fig 3.7, anterior and posterior motif grooming actions are indicated by light red and blue outlines, respectively. Note that self-transitions are explicitly prohibited as a result of the ethogram discretization method. Consequently, the network diagram lacks self-loops and all diagonal elements in transition matrices are blacked out, signifying zero value entries.

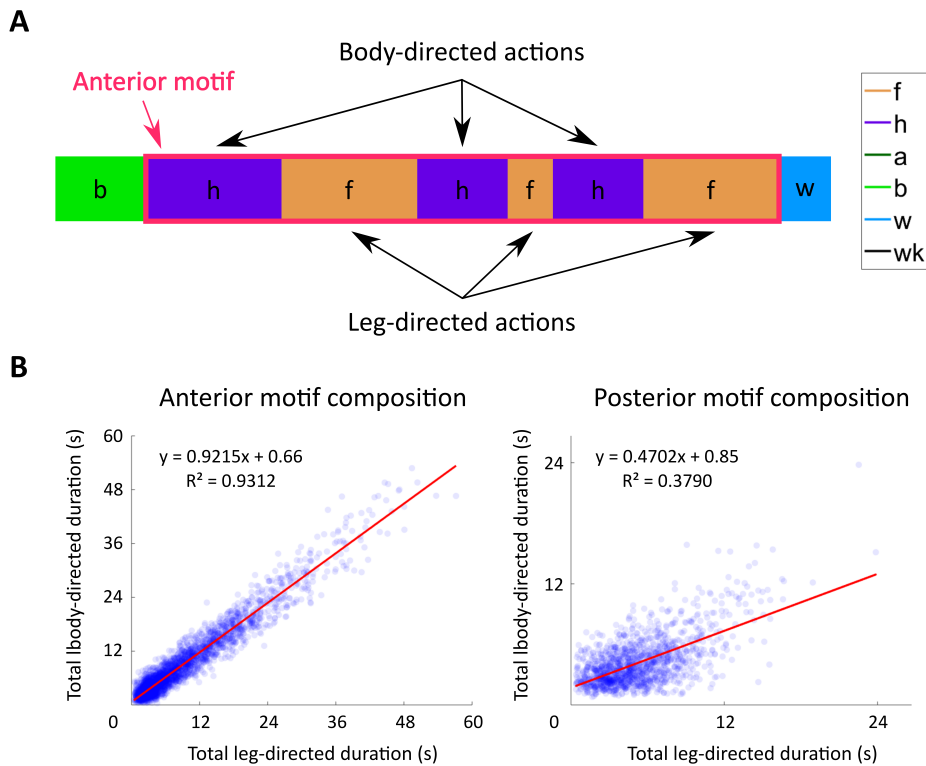


Figure 3.6: **Anterior motifs are composed of highly correlated amounts of body-directed and leg-directed actions.** A: Schematic of an anterior motif. Anterior motifs are defined as continuous consecutive bouts of head grooming (h) or front leg rubbing (f) flanked by non-anterior motif actions on either side. Posterior motifs are defined similarly, but consist of abdomen grooming, back leg rubbing, and wing grooming (a, b, and w). Both motifs contain body-directed actions (h, a, and w), which clear irritant from the body. Leg-directed actions (f and b) clear irritant from the legs, which collect irritant during body-directed grooming actions. In this example, the anterior motif consists of six grooming bouts of varying duration. B: Anterior motifs exhibit a strong linear relationship between the total amount of time spent performing body and leg-directed actions (left). Each point corresponds to an observed motif consisting of four or more actions. Posterior motifs (right) display a weaker linear trend with greater amounts of leg-directed grooming.

Grooming bout duration contributes to sequence structure at short time scales

In Fig 3.7, we illustrate the fine-grain temporal structure present within the identified motifs. Transition matrices with either two (left) or three (right) duration bins reveal the

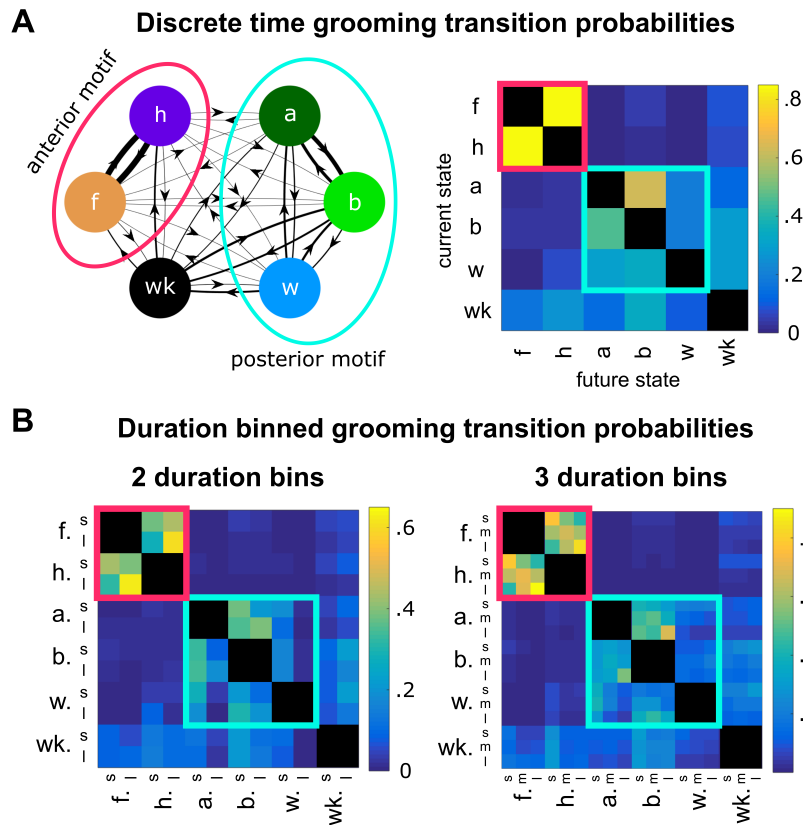


Figure 3.7: **Within-motif transitions dominate grooming syntax.** A: The maximum likelihood transition matrix, $\widehat{\mathbf{M}}$ (Eq 3.2), provides the best estimate of the Markov chain dynamics observed in grooming behavior (Eq 3.1). The network representation of $\widehat{\mathbf{M}}$ (left) illustrates transition probabilities using edge thickness (thicker edges indicate higher probabilities). Shown on the right is the matrix representation of discrete time transition probabilities, with probability magnitude indicated by color. Here, the anterior motif is indicated by the red circle (left) and red square (right). The posterior motif is delineated by the blue circle (left) and blue square (right). B: Maximum likelihood transition probabilities for ethograms binned according to the schema shown in Fig 3.3. Matrices fit to data with two (left) and three (right) duration categories show similar within-motif structure with increased resolution (s = short, m = medium, l = long). For example, both matrices illustrate that anterior motif transitions between long bouts dominate syntax.

presence of certain high-probability transitions. Most notably, we observe that transitions between long anterior grooming actions dominate to an extent not found in the null hypotheses we use for comparison. For example, when using three duration categories, the probability of transitioning from a long front leg rub to a long head cleaning bout

is $\sim 39\%$. This is significantly higher than the corresponding transition probabilities in the duration permuted and order permuted null hypotheses, which are $\sim 28\%$ and $\sim 13\%$ respectively. These differences are depicted in greater detail in Fig 3.8.

Additionally, we observe that, regardless of the number of duration bins, the anterior motif transition probability structure is nearly mirror symmetrical along the diagonal (i.e. transitions between front leg rubbing and head cleaning bouts have the same probabilities). In contrast, posterior motif transitions are less strongly symmetrical, with transitions between abdomen cleaning and back leg rubbing exhibiting the highest probabilities (Fig 3.7).

We find that the probabilistic rules describing transitions between consecutive bouts depend on two features of immediate behavioral history: previous bout identity and previous bout duration. To validate this observation, we evaluate the quality of the maximum likelihood model using the Bayesian Information Criterion (BIC). This metric, which is described mathematically by Eq 3.3, rewards models that produce accurate predictions and penalizes those that contain more parameters. The model with the lowest BIC score is considered the best fit to the observed data. Using BIC, we show that, compared to statistical null model hypotheses, the maximum likelihood transition probability best describes the data without overfitting. BIC values and null hypothesis matrices used for comparison are shown in Fig 3.8. Since each null hypothesis matrix contains fewer free parameters than the maximum likelihood model, the advantage of the maximum likelihood matrix in BIC is due entirely to its superior explanatory power.

Inspection of BIC values indicates that bout order is the strongest individual determinant of grooming syntax and that bout duration provides an additional contribution. Fig 3.8 shows the BIC values for each matrix as determined using data binned into three duration categories. The order permuted null hypothesis possesses a significantly larger BIC value than the duration permuted null hypothesis, indicating that disrupting bout

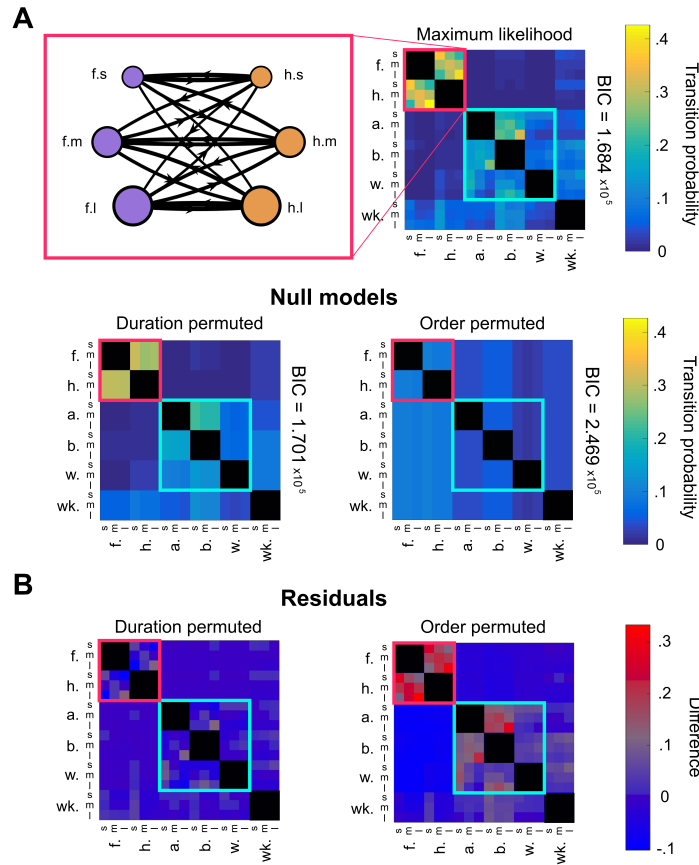


Figure 3.8: **Grooming bout duration contributes to syntax at the scale of individual transitions.** A: Probabilities from the maximum likelihood transition matrix (top) differ significantly from null hypothesis transition matrices (bottom) to varying degrees, indicating that both grooming action order and bout duration contribute to grooming syntax. Shown on the top left is the anterior motif transition network, with edge weights proportional to transition probabilities. Transitions between actions belonging to the same duration category possess high probabilities and appear nearly symmetric, suggesting a coupling mechanism between anterior motif actions. BIC values (right of matrices) provide validation that the maximum likelihood model captures statistically significant features of the data. The maximum likelihood transition matrix has a lower BIC value than the null hypotheses used for comparison, indicating that both bout order and duration contribute to sequence syntax. B: The residual values illustrate that specific transitions differ from what the null hypotheses would predict. The duration permuted null model matrix exhibits block structure but fails to capture the temporal relationship between bouts, as illustrated by the red values in several “long-to-long” transitions, for example. The order permuted null model differs even more severely, indicating that, while duration dependence plays a role in sequence structure, action order is still the primary determinant.

order degrades predictive power most significantly.

By construction, each null hypothesis transition matrix contains identical row structure within sets of rows corresponding to the same grooming action. However, only the duration permuted transition matrix possesses block structure similar to that found in the maximum likelihood matrix, as seen in Fig 3.8. This feature also reinforces the claim that action order provides a relatively larger contribution to syntax, since preserving action order results in the preservation of transition matrix block structure.

Bout-to-bout transitions are largely stationary despite changing sensory conditions

Fig 3.9 shows the average transition probability matrices with three duration categories for the early and late phases of grooming, as separated by the dashed line in Fig 3.5. Data were separated into early and late phase data for each fly and maximum likelihood transition probabilities were determined as described in Methods and Models. Here, we determine the boundary between phases to be the average time elapsed after flies have performed half of their discrete grooming actions. We also conducted a similar analysis in which we fit transition matrices only to the first and last third of data so as to avoid using the middle portion of the progression. This did not alter our results (Appendix B.5).

Qualitatively, these matrices exhibit strong similarities, despite the differences in grooming proportions, as shown in Fig 3.5. In particular, transitions between anterior motif grooming movements are similar across phases of the recording. Additionally, we find that 253 of the 266 (95.1%) non-zero entries change by less than 5% and 168 (63.2%) entries change by less than 1% (Appendix B.6). This suggests that the rules for sequence generation are nearly fixed, in spite of changing sensory input. Most notably, transitions

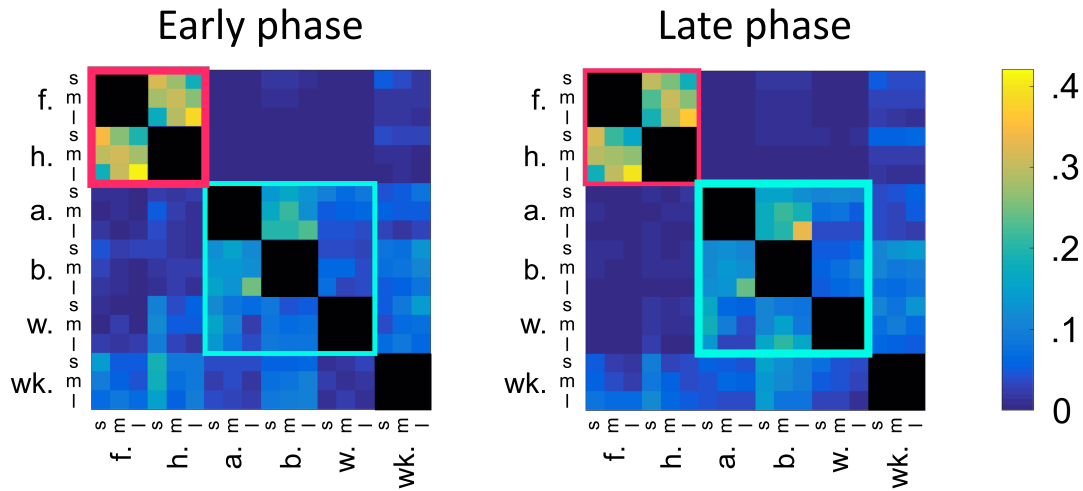


Figure 3.9: **Transition probabilities are largely stationary across the entire grooming progression despite changing sensory conditions.** Shown here are the population average transition probabilities for the early phase (left) and late phase (right) of grooming. The border between these phases is indicated by the dashed gray line in Fig 3.5. In the early phase, flies prioritize anterior grooming motifs, indicated by the thick red outline. In the late phase, flies perform more posterior grooming, indicated by the thick blue outline. We observe that transitions between front leg rubbing and head cleaning bouts exhibit consistent duration dependence regardless of when they occur in the sequence. Posterior motif transitions display similarities as well, but transitions between long abdomen grooming bouts and long back leg rubbing bouts are significantly more likely late in grooming. Overall, the relative stationarity of these transition probabilities despite changing sensory conditions suggests the existence of an internal mechanism that dictates bout durations. However, sensory stimuli also appear to play a role in modulating grooming transitions on long time scales, as transitions from posterior motif actions to anterior motif actions become less likely in the late phase.

between long abdomen grooming bouts and long back leg rubbing bouts are enriched in the late phase of grooming, illustrating that grooming rules, while similar, are not completely stationary over time.

A slowly varying model captures grooming syntax at long, intermediate, and short time scales

Markov models are steady state models and, as such, only exhibit equilibrium dynamics when used in a generative manner to create synthetic data. The first order maximum likelihood Markov model presented above is best suited for describing the late stages of grooming, which resembles steady state behavior. This is an obvious shortcoming of using a Markovian framework for analyzing the ethograms presented here, as flies display a dynamic progression (Fig 3.5), likely due to changing amounts of irritant over the course of grooming.

In order to more explicitly model these changing sensory dynamics, we utilize a non-stationary, or time-varying, Markov renewal process (MRP). Here, we introduce a non-stationarity in the form of a time-varying transition probability matrix. In order to maintain as parsimonious a model as possible, we define the transition probability matrix as a time-dependent convex combination of an early phase matrix and a late phase matrix (Eq 3.6). The early phase matrix is determined using maximum likelihood fitting on only the first 200 actions of each fly. Likewise, the late matrix uses each fly's final 200 actions. This ensures that early and late phases are well-separated, as each ethogram averages approximately 1000 bouts.

The time-varying transition probability matrix $\mathbf{M}(t)$ used in the MRP is defined as

$$\mathbf{M}(t) = \begin{cases} \frac{13-t}{13}\mathbf{M}_{early} + \frac{t}{13}\mathbf{M}_{late}, & t \leq 13, \\ \mathbf{M}_{late}, & t > 13 \end{cases} \quad (3.6)$$

where t is in minutes. After 13 minutes, the late phase transition matrix is used due to the observation that flies exhibit steady state behavior after approximately that amount

of time. Bout duration distributions used in the MRP do not change over time, as we did not find any significant difference between early and late bout durations (Appendix B.7).

Surprisingly, using this simple linear interpolation as a proxy for changing sensory conditions yields a model which produces ethograms that closely resemble our observed data (Fig 3.10). This interpolation serves as a first-order approximation to changing sensory drive, as flies constantly remove dust at an unknown rate over the course of grooming. Synthetic flies prioritize anterior grooming early on and gradually transition into a late phase which prioritizes posterior grooming and walking. This synthetic progression closely matches the actual progression shown in Fig 3.5. The synthetic ethograms also exhibit a tight coupling between the amount of anterior grooming actions over time. Moreover, synthetic motifs possess similar structure to observed motifs, as shown in panel C of Fig 3.10. Finally, the use of empirical transition matrices and bout duration distributions in our MRP guarantees that synthetic ethograms possess the same duration dependence at short time scales as actual data.

3.4 Discussion

***Drosophila melanogaster* respond to irritant by executing structured, modular sequences of cleaning movements**

After exposure to an irritant, *D. melanogaster* engage in a series of grooming actions in order to clear the irritant from their body. These movements are highly non-random, as indicated by their structure at three separate temporal scales. On the longest time scale, grooming progresses over several minutes from an early anterior-heavy phase to a later phase which contains elevated levels of posterior-directed actions and walking as

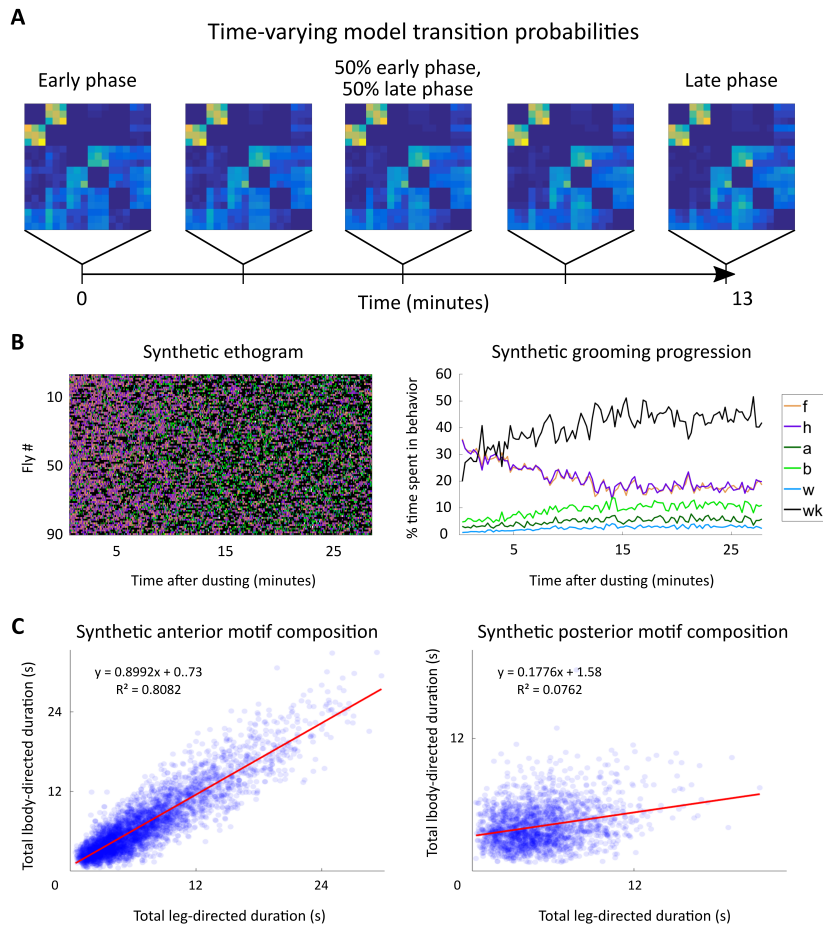


Figure 3.10: **A nonstationary Markov renewal process recapitulates grooming progression and bout structure.** A: Illustration of the time-varying transition matrix, $\mathbf{M}(t)$, used for generating synthetic ethograms. To approximate changing sensory conditions in the simplest possible manner, two transition matrices, \mathbf{M}_{early} and \mathbf{M}_{late} , are fit to the first and last 200 actions in the discrete time ethogram with 3 duration categories, respectively. As time evolves in the synthetic ethogram simulation, $\mathbf{M}(t)$ changes as described in Eq 3.6. At the beginning of the simulation, $\mathbf{M}(t)$ is identical to \mathbf{M}_{early} and after 13 minutes, it is identical to \mathbf{M}_{late} . Between those times, it is a linear combination of the two matrices. B: Synthetic ethograms display the characteristic progression from anterior to posterior grooming, as seen from comparison with Fig 3.5. C: Synthetic ethograms reproduce observed anterior motif composition. Synthetic anterior motifs (left) exhibit a similar, though slightly weaker, trend as observed in our data (Fig 3.6). Posterior motifs are less similar, indicating that other factors may be necessary to explain posterior motif structure.

illustrated in Fig 3.5. At an intermediate time scale, anterior and posterior grooming motifs exhibit strong correlations between the amount of body and leg-directed move-

ments present in each motif, as shown in Fig 3.6. Finally, grooming syntax exhibits non-random structure at the short time scale of individual grooming bout decisions, as illustrated by the non-random structure of transition matrices shown in Fig 3.7. Together, these findings suggest that there are distinct syntactic rules that govern grooming sequence structure at different temporal scales.

Here, we isolate short time scale grooming syntax by using a binning scheme to represent grooming bouts as semi-discrete actions (e.g. short, medium, or long bouts). Across different binning schemes, within-motif transitions (i.e. anterior-to-anterior or posterior-to-posterior) dominate syntax at the scale of consecutive bouts, as indicated by the block-like structure of transition matrices in Fig 3.7. The likely functional explanation for this tight correlation is that *D. melanogaster* engage in leg-centric grooming as a way to clear irritant that has collected on their limbs after body-centric grooming. Past work on blowflies also suggests that this coupling may be due to postural considerations, since fewer motor actions are required to transition between movements which use the same set of legs [67].

Moreover, the amount of front leg rubbing over the course of the grooming progression corresponds nearly exactly to the amount of head cleaning performed by the flies. The amount of back leg rubbing is also directly proportional to the amount of abdomen and wing cleaning (Fig 3.5). Though the total amount of time spent performing actions within motifs is tightly coupled, consecutive bout durations do not exhibit a simple linear relationship (Appendix B.4). Instead, transitions between body-centric and leg-centric actions of certain durations dominate. In particular, anterior motif movements exhibit nearly symmetric dynamics, with a preference for long-to-long transitions. In contrast, posterior motif transitions are asymmetric - this can be seen in Fig 3.7, as transitions from body-directed actions to leg-directed actions have different values than transitions in the opposite direction. This suggests a mechanism whereby internally generated dynamics

that are specific to each motif guide intra-motif transitions.

Symmetric anterior motif transition rules suggest higher-level transition frequency control

Anterior motif transitions display symmetric transition rules, as shown in Fig 3.7. The transition from head cleaning to leg rubbing may exhibit a strong correlation in duration for the following reason - the longer the head cleaning bout, the more particulate the legs accumulate, which subsequently requires more time to clear. However, it is more surprising that after performing a bout of leg rubbing the animal should display a similar duration preference upon transitioning back to head cleaning; once the front legs are sufficiently clean, we would not expect the animal to exhibit any duration preference on the subsequent bout. Additionally, the bout duration distributions for grooming actions remain strikingly similar across the entire recording, suggesting that sensory considerations do not account for the duration dependence we observe (Appendix B.7).

These observations support the hypothesis that alternations between anterior grooming bouts are regulated by a common source or common dynamics that dictate intra-motif transition frequency. If this is the case, future investigation of neural circuitry could focus on identifying neural activity that oscillates in phase with grooming bout alternations. Under this model, flies would then need to combine discrete decisions about which motif to perform with continuous decisions about how long to maintain a motif. This provides other targets for future experiments, as we could look for neural activity corresponding to transitions between motifs, explore manipulations that induce transitions between motifs, and search for factors that affect the duration of motifs.

Modeling suggests that sensory input modifies internal grooming programs

Although sensory input is sufficient to initiate individual grooming bouts, it has been difficult to assess the role of sensory input in grooming on longer time scales, since direct quantification of irritant relies on invasive protocols [87]. To date, the contribution of internal dynamics to grooming sequence generation also remains unknown. The discovery of duration dependence in grooming suggests that flies do not rely exclusively on sensory information to make grooming decisions, leaving open the possibility that sensory input modifies pre-existing autonomous grooming programs.

Here, we abstract the contribution of sensory input, incorporating it as a parameter in a time-varying Markov renewal process. Specifically, we approximate changing sensory conditions as a linear phenomenon which dictates the degree to which flies favor early versus late grooming rules (Fig 3.10). This choice of nonstationarity preserves the parsimony of our MRP, as it precludes the necessity of re-calculating transition probabilities over time. Additionally, the similarity between early and late intra-motif transition probabilities (Fig 3.9) provides some leeway for the choice of nonstationarity, bolstering our confidence that a simple linear rule is sufficient to approximate true sensory conditions.

Since our MRP reliably reproduces true ethogram statistics without explicitly modeling sensory input, we propose that, rather than acting in a purely reflexive manner, *D. melanogaster* instead use sensory information to modulate internal grooming programs. This type of slowly varying process has been observed in other organisms, suggesting that it is a useful framework for describing sequential behaviors [88]. Additionally, this hypothesis can be tested in future experiments which manipulate sensory conditions using irritants with different properties or optogenetic stimulation of sensory receptors.

A role for sensory information-independent transition rules in grooming

From a computational perspective, sensing is costly when it requires high-frequency updates and can provide readouts at a high resolution, as specialized neural circuits must be formed, maintained, and integrated in order to provide accurate, real-time readouts. This can lead to increased metabolic costs on several fronts, as increased neural activity and behaviors intended to protect and maintain sensory machinery can also require energetic resources [89]. In contrast, autonomous internal dynamics require only sparse sensory updates and may not require the same level of neural activity as is needed for sensory processing. If sensory drive alone, triggered by the presence of an irritant, dictates sequence generation in a purely reflexive manner, competing priorities may lead to rapid alternations between actions without any distinctive duration structure, consistent with neural circuitry implementing competitive inhibition with a transition cost [70].

During the execution of a grooming motif, frequent, high-resolution sensory updates may not be necessary. Instead, the moment-to-moment decisions about execution of grooming bouts can be automated by utilizing patterned internal dynamics. Several models of behavioral generation in which internal triggers drive sequence generation have been proposed [90, 91]. Since we show that intra-motif transitions exhibit bout duration-dependent structure, we suspect that internal triggers may provide a basis for transitions on short time scales after a discrete grooming motif decision is made. We consider this a strong possibility due to the fact that anterior motif transition probabilities are nearly stationary across grooming even though the dust distribution on the body is constantly changing. The near stationarity of these rules in spite of a changing stimulus suggests to us that behavior possesses an internally generated, fixed component. The discovery that grooming progresses gradually from anterior to posterior movements suggests that, on

long time scales, *D. melanogaster* utilize sensory information to dictate decisions about which grooming motif to perform, since the dust distribution changes over the period of cleaning.

In discussing temporal correlations between neural activity across spatial scales, Berman et al. [71] note that “(a)lthough no such correlation has been specifically found in *Drosophila*, our results suggest that such neuronal patterns may exist: perhaps by combining descending commands from the brain with local circuitry within and emerging from the ventral nerve cord.” This observation is consistent with our interpretation of our results, as local circuitry in the ventral nerve cord could plausibly generate internal dynamics which introduce duration dependence to action selection.

By combining sensory input and internal dynamics, *Drosophila melanogaster* nervous systems may utilize multi-level control algorithms which make discrete, “ballistic” decisions about the onset and type of continuous behaviors, allowing them to update infrequently at low sensory resolution. Once the decision is made, the execution of the behavior can be guided primarily by autonomous internal dynamics. In flies, these internal dynamics may not be unique to grooming, as freely-behaving *D. melanogaster* exhibit hierarchical Markov-like behavior in the absence of external sensory stimuli [79, 71]. Additionally, the identification of stereotyped subroutines in *D. melanogaster* locomotion further supports the idea that specialized neural circuits can automate and reproduce portions of behavior, reducing the need for constant calibration via sensory feedback [80]. We can further explore this idea by studying whether spontaneously grooming flies or those stimulated using optogenetic methods exhibit similar syntactic rules.

In larval and mature *D. melanogaster*, specialized, isolated neural circuits known as central pattern generators (CPGs), use rhythmic activity to guide execution of locomotion and courtship song sequences [92, 93]. In humans, CPGs are believed to facilitate chewing and breathing subroutines, such as individual jaw or pharyngeal movements [94, 95],

indicating that such neural circuit elements are common and flexible enough to carry out a wide variety of functions.

Since CPGs can modulate motor programs independently of sensory input, they are a natural candidate to execute grooming behavior on the scale of individual leg sweeps and rubs within a grooming bout. We do not explicitly model behavior at this temporal scale, but the ABRS classifier applied to video at the frame rate used here identifies actions nearly at the temporal resolution of individual leg movements. We plan to extend analysis to this time scale in future work.

Additionally, robotic systems that utilize multi-layer CPG architectures can carry out locomotor tasks, indicating that hierarchical CPG structures can generate sequential behaviors from smaller elements [96]. Recent work indicates that circadian rhythms and genetic factors contribute to *D. melanogaster* grooming behavior as well, leaving open a role for internal regulation of sequence generation at longer time scales [97].

3.5 Conclusion

When describing complex behavioral phenomena, formal mathematical models can bridge the gap between phenomenology and biological mechanisms by providing parameters that may correspond to underlying neuronal activity. In many cases, statistical model parameters do not directly describe underlying biological structures, but data analysis can suggest hypotheses for future exploration.

Here, we analyze large-scale automatically annotated data sets to reveal the syntax of *D. melanogaster* grooming. Using Markov models to detect features of grooming sequences, we find that grooming actions exhibit duration dependence at the scale of individual bouts and grooming motifs. Then, we produce realistic synthetic ethograms by introducing a slow modulation of grooming rules meant to abstract the contribution of

sensory input. Together, these findings suggest that internal programs dictate grooming decision on fast time scales but are modulated by sensory input over longer time scales.

As the tools for computational analysis of behavioral data continue to develop, interdisciplinary approaches that use mathematical tools to illuminate biological mechanisms will yield further insight into the generation of sequential behavior. Future experiments involving the manipulation of sensory experiences using optogenetic stimulation or other methods will also allow researchers to test hypotheses with unprecedented precision and scope. Together, these advances promise to improve our understanding of sequence generation by connecting mathematical descriptions of behavior with the underlying neural circuitry, as we are now motivated to search for circuits controlling long, intermediate, and short time scale grooming rules and determine how they interact with sensory circuits. Finally, our work suggests that temporal dynamics can and should be included when using statistical models to assess complex behaviors.

Chapter 4

Variability and Variation in *Drosophila* Grooming Syntax

4.1 Introduction

For over a century, researchers have understood that phenotypic variation arises from genotypic differences between organisms [98]. Animal behavior is a phenotype partially under genetic control, so a sizable research effort has been dedicated to uncovering specific genes associated with observable differences in behavior between and within species [99]. Different mouse species exhibit variation in monogamous behavior and parental care and different fly species show variation in courtship song [100, 101]. Within a species, natural variability produces individual mice that differ in aggression and flies that implement different foraging strategies [102, 103]. Mutant screens have also uncovered genes associated with differences in locomotion, courtship routines, and sleep patterns, among other complex behaviors [104, 105, 99, 17].

Fruit flies live in dirty environments, from laboratory vials to rotting fruit, and perform grooming actions to remove accumulated particulate. This behavior has been ob-

served in several drosophilid species and is important for social behavior and survival [20, 106, 107]. Past work has demonstrated that the leg movements used in grooming are stereotyped, but that sequences of actions are probabilistic and variable as opposed to fixed. However, the rules, or syntax, underlying grooming exhibit observable structure [21, 108]. Different sensory experiences and life histories may influence probabilistic grooming behavior, leading us to ask: *how much variability in fruit fly grooming is due to genetics?*

To address this question, $N = 372$ flies were covered in dust and their grooming behavior was recorded for approximately 30 minutes each (Figure 4.1A). We analyzed flies from five drosophilid species (*melanogaster*, *santomea*, *sechellia*, *simulans*, and *erecta*), which are genetically distinct and inhabit different ecological niches. We also examined four common *melanogaster* stocks (Canton-S, Oregon-R, Berlin-K, and w1118), and several isogenic lines derived from parent stocks in the laboratory. To analyze this large data set, we used tools from computational ethology [109]. An automated behavioral recognition system (ABRS, [85]) was used to classify fly behavior into one of five grooming actions (front leg cleaning, head grooming, abdomen grooming, back leg cleaning, wing grooming) and two non-grooming actions (walking and standing). After generating ethograms (behavioral time series records) for each fly, several grooming features were extracted (Figure 4.1B). Measures of central tendency and spread in the distributions of these features were calculated for each group and classification analysis was performed to determine the degree of separability between groups.

We found that several features of grooming are conserved across different drosophilid species. All species performed stereotyped movements in a probabilistic sequence and tended to groom anterior body parts before posterior portions.

We evaluated different grooming features by comparing their average values between groups and their range within groups, using the concepts of *variation* and *variability*

as defined by Ayroles et al. [17]. Variation refers to differences in phenotype between genetically distinct populations, such as drosophilid species. For example, inter-species variation in drosophilid larval digging behavior has been identified based on the duration distributions of “dives” into an agarose substrate [110].

Here, inter-species comparisons were used to look for variation in grooming features (such as syntax and action duration distributions) between flies with large genetic differences. Several significant differences in syntax were identified, allowing for highly accurate species-specific classification.

Next, we looked for variation within *melanogaster* stock lines using similar analysis. Although grooming features were more similar between *melanogaster* stocks than between species, some lines did exhibit detectable differences in grooming, largely related to overall grooming activity levels (e.g. after dusting, some lines spend a large proportion of time grooming, while others walk or stand more).

In contrast, variability describes how a trait differs within a genetically similar population. Intra-species variability has been characterized in the locomotor behavior of fruit flies, as some populations exhibit a wide range of locomotor behaviors, whereas others exhibit a narrower range [17].

We hypothesized that if a heterogeneous gene pool contributes to variability in grooming behavior, it should be possible to observe trends in variability between *melanogaster* populations with high and low genetic diversity. Starting from different commonly used stocks, flies were interbred for several generations to produce flies with high genetic variability and, in parallel, individuals were isogenized to minimize genetic variability. All three groups exhibited similar variability in measured grooming features.

Furthermore, flies tested on sequential days revealed that between-fly differences in syntax were similar to within-fly differences over several sessions. Finally, flies stimulated using optogenetic manipulation to induce grooming exhibited within-individual variabil-

ity in grooming syntax comparable to dusted flies. These results suggest that genetic diversity is not a primary contributor to the variability in grooming action sequences. Future work that controls for sensory experience and life history could shed light on the degree to which stochastic effects influence behavioral variability.

4.2 Results

Drosophilids exhibit a robust grooming response but different syntax after irritant exposure

Regardless of genotype, all flies possessed several common grooming features. Here, five grooming actions were observed consistently across all five species of drosophilids (*erecta*, *sechellia*, *santomea*, *simulans*, *melanogaster*), indicating a conserved behavioral response. In this analysis, *melanogaster* flies interbred between stock lines for maximum genetic heterogeneity (MaxVar) were used as the representative *melanogaster* group to avoid any stock-specific trends.

To quantify the behavioral response to dusting, the proportion of time spent grooming (as opposed to walking and standing, the two non-grooming actions identified here) was calculated for each fly (Figure 4.2A). Overall, the distributions were statistically indistinguishable with one exception: *drosophila erecta* groomed significantly less than every other species (Wilcoxon rank-sum test, $p < .05$ after Bonferroni correction).

Grooming proportions were then examined in more detail by considering the progression of all seven actions over time, as in Figure 4.1B. All species exhibited a qualitatively similar grooming progression characterized by initial anterior grooming followed by increased posterior grooming, walking, and standing, but the relative proportions of behaviors over time differed between several populations (Appendix C.4).

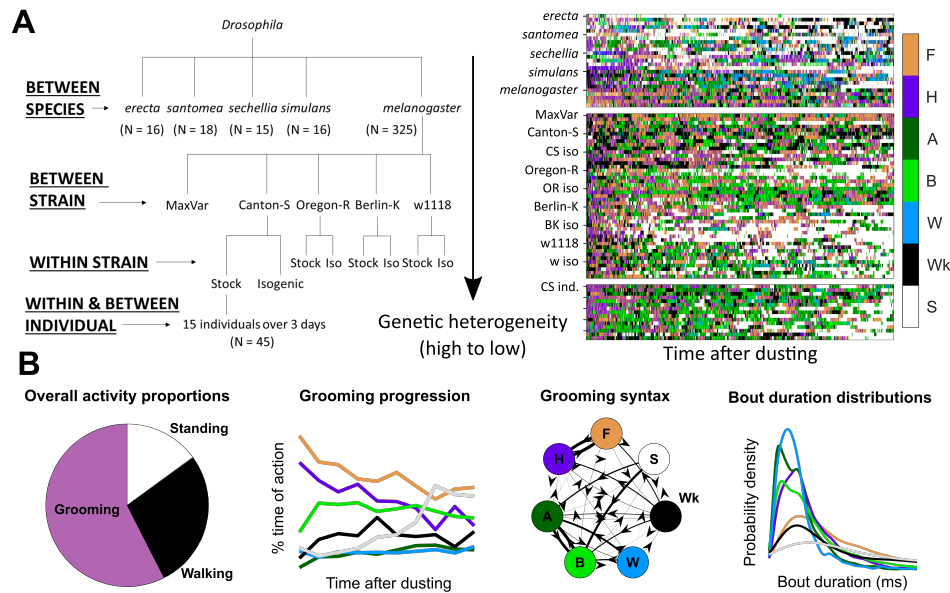


Figure 4.1: **Grooming variability dataset and analysis overview.** A. In total, $N = 372$ flies were dusted and their grooming activity was recorded for approximately half an hour each. Five drosophilid species, four *melanogaster* stock lines, one interbred *melanogaster* line, and six isogenic *melanogaster* lines were analyzed for similarity, differences, and variability in grooming behavior. On the left is a schematic of the different drosophilid groups included in this analysis. Higher levels of the tree indicate higher levels of genetic diversity (scale is relative, not absolute). On the right is a sample of ethograms generated by automated annotation of grooming video. Color indicates the occurrence of the five grooming actions (front leg cleaning, head grooming, abdomen grooming, back leg cleaning, wing grooming), walking, and standing. B. Features scored from ethograms provide summary representations of grooming behavior. The proportion of time spent grooming, overall or over time, provides the coarsest description of the behavioral response to a dust stimulus. Regardless of genotype, all flies exhibit probabilistic (not fixed) action sequences consisting of the same set of five grooming actions, walking, and standing after exposure to irritant. Additionally, most flies follow a typical grooming progression pattern: initial anterior grooming followed by increased posterior grooming. The amount and timing of walking and standing, however, can vary significantly between flies. Grooming transitions (syntax) and action (bout) duration distributions provide representations of individual grooming decisions within a recording session. Shown here are example features scored from Canton-S flies.

Grooming syntax (the transition probabilities between discrete behaviors) was calculated from the ethogram of each dusted fly ($N = 372$). With seven behavioral states, 42 transitions were possible, excluding self-transitions. Thus, syntax was represented as a

42-dimensional vector for subsequent classification analysis and visualization. Grooming syntax across all flies was probabilistic with high transition probabilities within anterior grooming actions (front leg cleaning, head grooming) and posterior grooming actions (abdomen grooming, back leg cleaning, wing grooming). The average syntax for each species is illustrated as a weighted, directed graph in Figure 4.2.

Finally, action duration distributions were calculated from ethograms. Distributions of grooming actions were similar across species and had probability peaks between 500 and 750 ms (Appendix C.6).

Several significant differences in grooming syntax between *melanogaster* and non-*melanogaster* species were identified. To compare syntax, transition probability distributions for each action transition were compared between species in a pairwise manner. 38 of 42 unique syntax elements (90.5%) were significantly different between at least two species (Wilcoxon rank-sum test, $p < 0.05$ after Bonferroni correction). Overall, 148 of 420 (35.2%) of pairwise syntax comparisons revealed differences between species.

Of these syntactic differences, 100 (67.6%) occurred between *melanogaster* and non-*melanogaster* species. In particular, posterior motif grooming transitions (i.e. transitions between abdomen grooming, back leg rubbing, and wing grooming) were consistently significantly different, on average, as were transitions between back leg rubbing, standing, and walking. Figure 4.2B illustrates these syntactic differences.

Figure 4.2C depicts a low-dimensional embedding of species syntax using t-SNE. This visualization suggests that different species possess distinguishable syntax, as points are aggregated by species. Classification and clustering analysis was applied to species syntax to probe this interpretation. Logistic regression classified flies by species according to syntax with >90% accuracy in 9 of 10 pairwise comparisons (Figure 4.2D). K-medoids clustering reinforced this finding, as these groups were classified with >90% accuracy using fewer than 20 syntax elements (Appendix C.12). Classification according to progres-

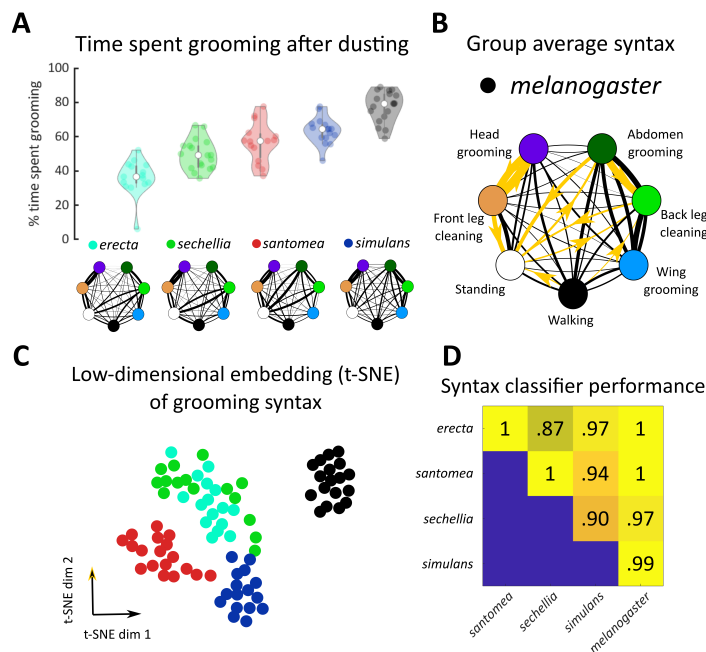


Figure 4.2: *Drosophila* species share grooming features but exhibit between-species variation in grooming syntax. A. Dusting elicits a conserved grooming response across drosophilids. All species produce a probabilistic grooming sequence (as shown in Figure 4.1A), which can be characterized by the transition probabilities (syntax) between grooming actions (as represented in Figure 4.1B, calculated as in [108]). Shown are violin plots of grooming proportions for each species examined here. Colored points represent a single fly, with color indicating species (as indicated by points next to species names in A and B). White points indicate the median value and plot with indicates density. B. Several syntax elements differ between drosophilid species. The mean syntax for each species is depicted as a graph, with nodes representing grooming actions and edges indicating transition probability. Thicker edges indicate higher probabilities. On the *melanogaster* syntax graph, the 10 action transitions exhibiting the largest magnitude differences between *melanogaster* and non-*melanogaster* species are highlighted in red. These differences are largely found in posterior motif transitions, which use the back legs to perform grooming actions. C. Each fly’s 42-dimensional syntax vector was plotted in two dimensions after dimensionality reduction using t-SNE. t-SNE preserves local distance structure, indicating that tightly grouped clusters of points are similar to one another. In this case, dimensionality reduction reveals that drosophilid species exhibit significant differences in grooming syntax, as syntax vectors congregate by color. D. Classification analysis confirms the qualitative clustering observed in C. Shown is a heat maps of accuracy rates resulting from pairwise logistic regression classifiers trained on syntax vectors. Samples used for classification were chosen to be the same size, so classification at chance levels would be 0.5. Consistent classification accuracy values $> .9$ indicate that species are highly separable by syntax.

sion and grooming proportions achieved similar performance, but using bout durations performed the worst (Appendix C.10).

***Melanogaster* strains exhibit variation in grooming behavior**

Standard lab strains (Canton-S, Berlin-K, Oregon-R, w1118) were analyzed for differences in grooming proportions, progression, syntax, and bout duration distributions. Grooming syntax differed between stocks, allowing for classification above chance levels. Grooming progressions were also different between stocks. Overall grooming proportions and bout durations achieved the worst classification performance, indicating that these features are more highly conserved.

The grooming syntax of *melanogaster* stocks is depicted in Figure 4.3A. Similar to the species analysis, all action transition probability distributions were compared in a pairwise manner to look for variation in syntax. Only 37 of 252 syntax element comparisons (14.7%) differed significantly between any two stocks and, of these, only two within-motif transitions (posterior motif) differed significantly. Within-motif syntax elements are of particular interest because they represent the most common, most highly stereotyped action transitions observed across flies of all genotypes (see Figure 4.2A and B for visualizations of these transitions). However, these differences allowed for classification of stocks above chance, as illustrated in Figure 4.3A.

Examination of other grooming features revealed that overall grooming proportions could also account for differences observed between stocks. Figure 4.3B shows a ternary plot of action proportions, showing that Canton-S flies spend more time walking than other stocks. Relatedly, the syntax element exhibiting the greatest stastically significant difference was the wing grooming to walking transition, shown in the middle of Figure 4.3A. Classification accuracy was poor when using grooming proportions but comparing

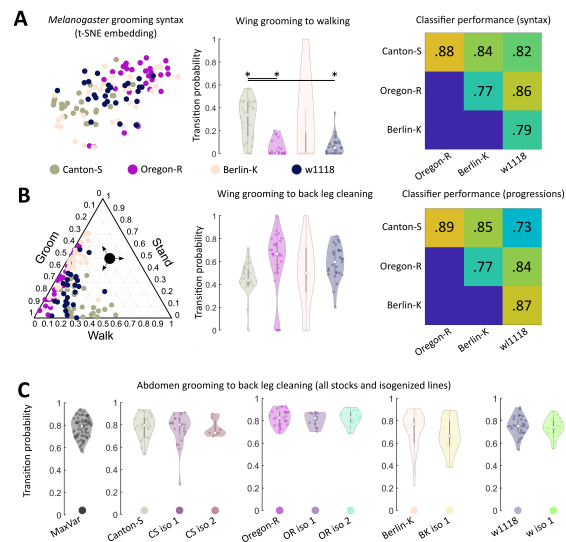


Figure 4.3: **Within *melanogaster*, different stocks exhibit similar syntax but different activity levels. Genetic homogeneity does not correspond to behavioral stereotypy.** A. Compared to inter-species differences in grooming syntax (Figure 4.2), *melanogaster* stocks did not exhibit strong variation in grooming features. This is indicated by the high degree of overlap in the t-SNE plot of grooming syntax vectors (left). The middle panel displays distributions of one representative strong grooming transition; 215/252 (85.3%) of transition distributions were not significantly different from one another. Of the 17 unique transitions that were significantly different, only two were within-motif transitions; both were posterior motif transitions. Due to the strong similarity of syntax between stocks, classifier analysis was not able to distinguish between stocks with greater than 70% accuracy (right). B. Most syntax elements are similar between *melanogaster* stocks, but Canton-S flies walked more than other stocks. On the left is a ternary plot of grooming, walking, and standing proportions for each species. Colored points represent individual flies. The large black point with arrows indicates how to read activity proportions; the example point corresponds to 20% grooming, 30% walking, and 50% standing. Due to differences in activity levels, some walking-related syntax elements differed between Canton-S flies and other stocks. The middle plot illustrates the most significantly different syntax element between stocks (Wilcoxon rank-sum test, $p < .05$ corrected for multiple comparisons). Classifier performance on action proportions (right) is similar to performance on syntax, indicating that syntax differences are not a more informative source of variation between stocks than activity levels. C. Genetic homogeneity did not correspond to behavioral variability. Shown are the distributions of wing to back leg transitions for a line bred for maximum genetic heterogeneity (MaxVar), each stock line, and lines derived from stocks to minimize genetic heterogeneity (iso). MaxVar flies did not consistently exhibit a higher degree of variability than stock lines. Isogenized lines did not exhibit a lower degree of variability than their parent stocks.

grooming progressions, which are similar to proportions but account for changes in grooming behavior over time, yielded significant differences. Canton-S flies higher propensity to walk after grooming their wings is reflected both in their syntax and grooming proportions, resulting in similar classification accuracy (Figure 4.3B, right) whether using syntax or progressions.

Figure 4.3B also illustrates the high degree of variability in *melanogaster* syntax. The wing grooming to back leg cleaning transition exhibited the largest difference between median values of any syntax element (comparison of Canton-S and Oregon-R yielded this difference), but none of these distributions possessed detectable statistical differences due to their concomitantly large variances.

Grooming behavioral variability is similar across *melanogaster* genotypes

To examine the relationship between genetic heterogeneity and behavioral variability, each *melanogaster* lab strain was compared to lines bred to minimize genetic heterogeneity (“isogenic” lines). To quantify this, the variances of grooming transition probability distributions were calculated and compared.

Only 6/252 (2.4%) syntax element distributions possessed statistically significantly different variances between MaxVar, Canton-S, and the isogenic lines out of all possible pairwise comparisons (Levene’s test, $p < .05$ after Bonferroni correction). Moreover, none of these differences corresponded to within-motif transitions, indicating that variability of common transitions is similar regardless of genetic heterogeneity in a population. These findings were similar for Oregon-R (8/252), Berlin-K (18/108), and w1118 (2/108) stock and isogenic comparisons.

Figure 4.3C provides the transition probability distributions for the most common

posterior motif transition (abdomen grooming to back leg cleaning) for all stocks and stock-derived isogenic lines. This transition exhibits wide variability in many populations and even populations with smaller variability (CS iso 2) are not different enough to achieve statistical significance after accounting for multiple hypothesis testing.

Standardizing sensory experience does not abolish grooming behavioral variability

Finally, we analyzed 15 Canton-S flies that were recorded over three consecutive days after irritant exposure. In this experiment we aimed to isolate the magnitude of grooming variability that is due to differences in sensory experience (since the dusting protocol does not allow for perfect replication of sensory experience) and life history (since flies will have been exposed to the same irritant several times by the end of the experiment). Ethograms from three example flies are provided in Figure 4.4A.

Flies did not exhibit strong longitudinal grooming trends. Observing individuals over several sessions provided no evidence of stable between-individual variation (individuality) in grooming proportions, progression, syntax, and bout duration distributions. Intra-individual variability across three sessions was statistically similar to inter-individual variation on any given day, suggesting that non-genetic factors account for a significant proportion of grooming variability (Appendix C.13).

To further probe the sensory basis of within-individual variability, 10 Bristle-spGAL4-1 flies (Bloomington Stock Center, RRID: BDSC_71032) were stimulated using optogenetic manipulation to induce anterior grooming (as in [111]). Figure 4.4B provides ethograms from this experiment, with red bars indicating the two stimulation windows. Even when sensory experience was controlled using optogenetic stimulation, flies exhibited variability in their grooming response. Grooming stereotypy was quantified using the

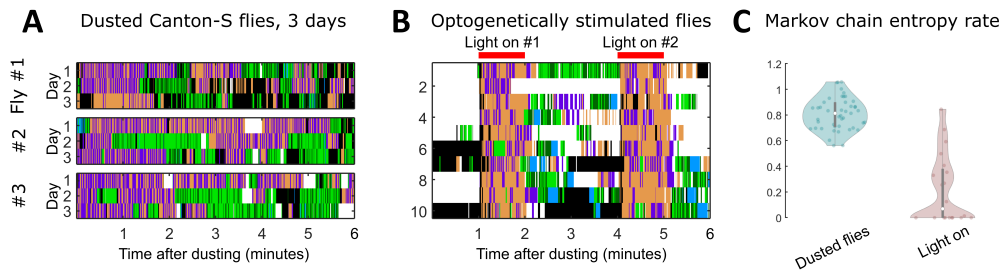


Figure 4.4: **Within-individual grooming variability suggests non-genetic factors account for a significant portion of differences in behavior.** A. Shown are portions of ethograms from three Canton-S flies observed on consecutive days after dust irritant exposure. The qualitative differences in ethograms on consecutive days indicate that non-genetic factors must account for some amount of within-individual grooming variability. B. Shown are ethograms of 10 Bristle-spGAL4-1 flies ([111]). Flies were optogenetically stimulated to induce anterior grooming in two separate one-minute windows, indicated by the red bars. Between these windows, flies still exhibit within-individual grooming variability even though the sensory experience is presumed to be more uniform than repeated dust exposure. C. Markov chain entropy, a measure of grooming stereotypy, was calculated from syntax. Optogenetically stimulated flies (right) exhibited lower entropies, corresponding to a higher degree of stereotypy, than dusted flies (left). However, optogenetically stimulated flies exhibited differences in stereotypy between stimulation windows, implicating sources of grooming variability beyond genetic and sensory influences.

entropy rate of the grooming syntax, represented as a first-order Markov chain. While the entropy rate for optogenetically stimulated flies was lower than for dusted flies, indicating a higher degree of stereotypy, we nonetheless observed within-individual variability in grooming stereotypy between stimulation windows. Figure 4.4C illustrates this finding, as the differences in entropy between sessions for three example flies are highlighted using red lines.

4.3 Discussion

Here, we analyzed the grooming behavior of 372 flies, including five different drosophilid species and four common *melanogaster* stocks, to investigate the relationship between genetic heterogeneity, variation, and variability in grooming behavior. Large genetic

differences (species-level) correspond to identifiable differences in several grooming features. Within *melanogaster* flies, stocks also exhibited variation in grooming features. However, increased genetic heterogeneity did not correspond to increased behavioral variability. Analysis of 15 Canton-S flies recorded over consecutive days found no evidence for individuality in grooming behavior, as all flies possessed highly similar grooming features across all days.

Genetic influences of behavioral variation and variability

In a classic example of behavioral variation, fruit fly larvae expressing the *rover* and *sitter* alleles of the *foraging* gene exhibit, respectively, greater or lesser propensities to leave patches of food [104, 112, 103]. Several genes that affect fruit fly grooming behavior have been identified in mutant screens [113], but the degree of normal variability due to genetic factors has not been characterized.

Here, we identified significant inter-species variation in grooming syntax, suggesting a genetic basis for group differences in grooming behavior. Some species differ from *melanogaster* in their propensity to perform posterior grooming actions, providing an avenue for future research directions aimed at exploring differences in posterior neuronal circuitry and sensory physiology further.

We also identified differences in grooming behavior between commonly used *melanogaster* stock lines. Other work has identified differences in locomotor behavior between ostensibly “standardized” *melanogaster* sub-strains [114], suggesting that within-species genetic differences, perhaps introduced by founder effects or laboratory selection, may impact behavioral phenotypes. Such considerations are important in designing and interpreting experiments aimed at relating genetics to behavior [115].

We also examined whether greater genetic heterogeneity within a population corre-

sponded to greater behavioral variability. The grooming syntax variability in a population of *melanogaster* flies bred to maximize genetic diversity was statistically equivalent to the variability of parent stocks and isogenic lines, suggesting that non-genetic factors contribute significantly to grooming variability.

Variability also encompasses *individuality* in animal behavior, typically defined as a trait-like feature that persists stably over several observations. Individuality has been identified in fruit fly turning decisions [116], mouse roaming behavior [117], and bumblebee foraging behavior [118], among others [119, 120], indicating the importance of considering *within-individual variability*, or the degree to which an individual’s behavior varies in repeated trials.

Our analysis found no evidence for individuality in grooming, as Canton-S flies exhibited similar degrees of within-individual variability and between-individual variation. One explanation for this observation is that, within a species, genetic influences on grooming behavior are outweighed by differences in sensory experience, which is controlled here to the greatest extent possible using standardized experimental protocols but some differences in sensory experience between trials is assumed to exist. Grooming variability was also assessed in optogenetically-stimulated flies, which allowed for greater stimulus standardization. However, optogenetically-stimulated flies exhibited non-trivial grooming variability across trials, suggesting that other factors also contribute to grooming variability.

Environmental and stochastic influences of behavioral variation and variability

Environmental (or contextual) differences have long been considered a source of behavioral variation, visible most notably perhaps in the canonical “nature vs. nurture”

paradigm of phenotypic differentiation. In *Drosophila*, environmental enrichment has been shown to affect variability in several behaviors, including locomotion and phototactic behaviors [121]. Limiting genetic diversity within a population, then, allows for characterization of environmental and stochastic effects on behavior.

Comparing several grooming features between populations of *melanogaster* flies, increasing genetic homogeneity did not decrease behavioral variability. This suggests that grooming variability may be weakly under genetic control, but the influence of genotype is likely outweighed by environmental and stochastic factors. A recent study of *melanogaster* visual object orientation demonstrated that stochastic, developmental asymmetries in dorsal cluster neuron projections to the medulla yield stable “personalities” in this paradigm [120].

Differences in sensory experience and life history likely contributed to the variable nature of grooming sequences within and between flies observed here. Even if sensory experience were perfectly controlled, however, neural circuits in genetically identical organisms are not guaranteed to be identical due to developmental differences arising from epigenetic factors, transcriptional variability, or post-translational regulation. Recently, these stochastic effects, defined as non-heritable effects that cannot be predicted from measurable variables, have been suggested to be important drivers of variability [122]. Studies of phototactic and thermotactic responses in *Drosophila* provide evidence that a single genotype is subject to enough developmental stochasticity to produce different, stable, non-heritable behavioral phenotypes [123].

4.4 Materials and methods

Genetic stocks

Canton-S, Oregon-R, Berlin-K, and w1118 stocks were obtained from the Bloomington Stock Center. Isogenic (more accurately, reduced genetic variability) stocks were made by crossing single males to double-balanced stocks and then back-crossing males to the double balancer stock to isolate single second and third chromosomes. Single pairs were mated to reduce variability of X and IV. ~ 2 independent isogenic lines from each *melanogaster* stock were generated; note that many attempts to isogenize result in lethality, as anecdotally reported by colleagues. Maximum Variability stocks were obtained by crossing each *melanogaster* strain to double balancers and then crossing the progeny together and selecting against the balancers. This allowed combination of chromosomes for all four strains. The progeny were allowed to interbreed for several generations to enable recombination in the females. Full cross schemes available on request. Drosophilid species stocks were obtained from Tom Turner, UCSB, and the National Drosophila Species Stock Center (<https://www.drosophilaspecies.com/>).

Data collection and processing

Grooming was induced and analyzed as described in Zhang et al. 2020 and Seeds et al. 2014 [111, 21]. Three chambers were used in fly dusting assay: dusting chamber (24 well Corning tissue culture plate #3524), transfer chamber and recording chamber. Recording chambers were coated with Insect-a-slip (BioQuip Products Cat#2871A) to discourage wall-climbing and cleaned daily. Dust-induced grooming assays were performed in 21-23. 4-7 day old male flies were anesthetized on ice and transferred to the middle four wells of the transfer chamber. Flies were left in the transfer chamber for 15 minutes to

recover. Approximately 5 mg Reactive Yellow 86 dust (Organic Dyestuffs Corporation CAS 61951-86-8) was added into each of the 4 middle wells of dusting chamber. For fly dusting, the transfer chamber was aligned with the dusting chamber. Flies were tapped into the dusting chamber and shaken 10 times. After dusting, flies and dust were transferred back into the transfer chamber. Transfer chamber was tapped against an empty pipette tip box to remove extra dust. Dusted flies were then immediately tapped into recording chamber for video recording. The entire dusting process was performed in a WS-6 downflow hood. Approximately 10 individuals were recorded for each genotype. 30 Hz videos were recorded for 50,000 frames (27.78 min) with a DALSA Falcon2 color 4M camera. A white LED ring light was used for illumination.

Optogenetic stimulation protocol is replicated from [111]. Further details can be found in Appendix C.

Videos were processed through the Automated Behavior Recognition System (ABRS, [85]), trained on a classifier using Canton-S *melanogaster* flies and validated for accuracy on other species, to generate ethograms. Grooming actions were described previously ([124]). ABRS was used to generate ethograms. Briefly, the raw video frames were pre-processed to generate 3-channel spatial-temporal images (ST images). Features were extracted in three timescales and saved into different channels of ST images: 1. raw frame; 2. difference between two frames; 3. spectral features extracted from a 0.5 sec window. A convolutional network trained by ST images under different light conditions was then used to label the behavior identified in each frame. A different network was trained for classification of each species due to differences in body size and light conditions. All networks achieved >70% validation accuracy, which is similar to the agreement between human-annotated ethograms.

Data analysis

All ethogram features were extracted using custom-written code in MATLAB 2019a. Grooming progression vectors were generated for each fly by calculating the proportion of each action in 10 non-overlapping windows (2.78 minutes each), yielding a 70-dimensional vector for each fly (10 windows with 7 behavioral proportions). Grooming syntax was defined as the first-order transition probabilities between actions. Syntax for each fly was calculated as described in [108]. Bout duration distributions were generated as described in [108], using a normalized histogram with 20 bins of equal width for each behavior. Bin width was determined independently for grooming and non-grooming actions, as standing and walking exhibit longer tailed distributions than grooming actions. Thus, duration distribution vectors were 140-dimensional for each fly. Examples of progression, syntax, and duration distribution vectors can be found in Appendix C.

Statistics for comparisons between grooming features were calculated using built-in MATLAB functions. K-medoids clustering, t-SNE, and logistic regression classification analysis were performed using built-in MATLAB functions (Appendix C).

Chapter 5

Brain Network Reorganization Over A Single Menstrual Cycle

5.1 Introduction

The application of network science techniques to the study of the human brain has revealed a set of large-scale functional brain networks that meaningfully reorganize both intrinsically and in response to external task demands [125]. One technique, dynamic community detection (DCD), has emerged as a powerful tool for conceptualizing and quantifying changes in mesoscale brain network connectivity patterns by identifying sets of nodes (communities) with strong intra-community connections [126] to enable identification of communities that persist or change over time. DCD complements other statistical approaches used in fMRI data analysis by identifying when functionally coupled brain regions undergo sufficiently large changes in connectivity to warrant re-assignment to separate functional communities. Additionally, this method provides an interpretable summary of whether strongly connected sets of brain regions undergo transient, but significant, changes that could be missed when time-averaging data within and between

sessions.

This method is particularly suited for examining relationships between brain dynamics and physiological variables that vary over relatively short time scales, such as sex hormone fluctuations over the human menstrual cycle. A typical cycle, occurring every 25-30 days, is characterized by significant rises in estradiol (~ 12 -fold) and progesterone (~ 800 -fold), both of which are powerful neuromodulators that have a widespread influence on the central nervous system [127]. Converging evidence from animal studies has established sex hormones' influence on regions supporting higher-order cognition, including the prefrontal cortex (PFC) and hippocampus [128, 129]. Within these regions, fluctuations in estradiol enhance spinogenesis and synaptic plasticity while progesterone largely abolishes this effect [130, 131]. Importantly, sex hormones are expressed broadly throughout the cerebellum and cerebrum, suggesting that whole-brain effects might be observed beyond the regions targeted in these studies.

Human neuroimaging studies have demonstrated that sex hormones influence brain activity across broad regions of cortex [132, 133]. Additionally, a handful of studies have demonstrated that menstrual cycle stage uniquely alters resting-state functional connectivity (rs-fc) [134, 135, 136, 137]. However, these studies typically involve group-based or sparse-sampling (2-4 time points) designs that are unable to capture transient day-to-day relationships between sex hormones and functional brain dynamics, and this relatively low temporal resolution has led to inconsistencies in the literature [138]. Therefore, new approaches are needed that can address these spatial and temporal limitations, as doing so will provide novel perspectives on human brain-hormone interactions.

Recently, Pritchet et al. applied a “dense sampling” approach [139, 140] to a naturally-cycling female who underwent 30 consecutive days of brain imaging and venipuncture to capture rs-fc variability over a complete menstrual cycle (Figure 5.1). The authors found day-to-day fluctuations in estradiol to be associated with widespread increases in

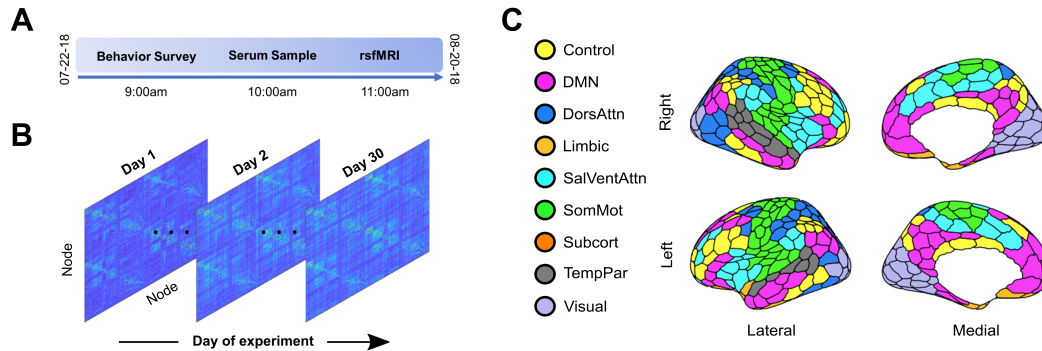


Figure 5.1: **28andMe dataset.** **A.** Subject LP (naturally cycling female, age 23) participated in a month-long “dense sampling” experimental protocol to provide a multimodal, longitudinal dataset referred to as 28andMe [143]. The subject completed daily assessments of diet, mood, and sleep, provided blood for assessment of serum hormone concentrations, and underwent a 10 minute resting-state fMRI scan. **B.** For each resting-state scan, functional connectivity matrices were constructed by calculating the pairwise mean magnitude-squared coherence between each region from the entire 10-minute scan. The result is a $415 \times 415 \times 30$ data structure, in which each entry indicates the coherence between two nodes on a given day. **C.** The brain was parcellated into 415 regions and regions were assigned to one of nine networks based on previously identified anatomical and functional associations [144]. Colors indicate regional network membership.

rs-fc across the whole brain, with progesterone showing an opposite, negative relationship. Using time series modeling and graph theoretical analysis, they also found that estradiol drives variation in topological network states, specifically within-network connectivity (global efficiency) of default mode and dorsal attention networks that encompass regions rich with estrogen receptors (ER). These findings have important implications for the field of network neuroscience where dense-sampling, deep-phenotyping approaches have emerged to aid in understanding sources of intra/inter-individual variability in functional brain networks over days, weeks, months, and years [140, 141, 142].

Pritchet and colleagues’ approach identified node-averaged trends in rs-fc changes within canonical functional networks across the cycle, but questions remain regarding whether and where functional reorganization takes place between large-scale networks. As changes in edge weight can result in the formation of functional “communities” not

captured by traditional rs-fc methods, complementary approaches are needed to characterize trends in brain connectivity at intermediate spatial and temporal scales. Examining mesoscale networks has further revealed fundamental principles of functional brain networks, such as the modular, integrated architecture underpinning flexible task performance [145, 146]. Additionally, a better understanding of mesoscale connectivity may provide an avenue for improving personalized medicine by increasing the efficacy of targeted therapeutic interventions [147].

Here, we applied DCD to examine whole-brain dynamics in relation to sex hormone fluctuations across a menstrual cycle. Our results reveal that a stable set of “core” communities persist over the course of a menstrual cycle, primarily consisting of nodes belonging to distinct *a priori* defined functional-anatomical networks, namely visual, somatomotor, attention, default mode, and control networks. Though these core communities were largely stable, nodes from limbic, subcortical, attention, and control networks changed community affiliation (referred to as flexibility) at higher rates than expected compared to a null hypothesis. DCD also identified a transient split of the DMN core into two smaller subcommunities concurrent with peaks in estradiol, luteinizing hormone (LH), and follicle stimulating hormone (FSH) levels defining the ovulatory window. This community split was driven by strong increases of within-network integration between prefrontal nodes of the DMN, which subsided immediately after the ovulatory window. The default mode, temporoparietal, limbic, and subcortical networks also exhibited significantly increased flexibility during ovulation, suggesting a role for estradiol, LH, and FSH in regulating localized, temporary changes in regional connectivity patterns. Taken together, while a large degree of functional brain network stability was observed across the menstrual cycle, peaks in sex hormones over the ovulatory window resulted in temporary brain network reorganization, suggesting sex hormones may have the ability to rapidly modulate rs-fc on shorter time scales than previously documented.

5.2 Results

A single female underwent brain imaging and venipuncture for 30 consecutive days. For each session, the brain was parcellated into 400 cortical regions from the Schaefer atlas and 15 subcortical regions from the Harvard-Oxford atlas (Figure 5.1C) and 415 x 415 functional association matrices were constructed via magnitude-squared coherence [144]. Dynamic community detection was applied to this data, revealing a stable set of communities that persist over the course of a menstrual cycle. However, significant transient changes in community structure occurred within the default mode network during the ovulatory window concomitant with peaks in estradiol, luteinizing hormone, and follicle stimulating hormone.

Stable functional cores persisted over the course of one menstrual cycle

The degree to which functional brain network connectivity changes over the course of a human menstrual cycle has yet to be fully characterized. Here, dynamic community detection (also referred to as multislice or multilayer modularity maximization [148]) consistently identified four functional communities that were largely stable in a naturally cycling female over 30 consecutive days. In this context, “community” refers to a set of nodes whose intra-set connections are significantly stronger than would be expected when compared to an appropriate null model. A representative example of this consensus temporal community structure (the community designation that best matches the output of 50 runs of the non-deterministic community detection algorithm) is shown in Figure 5.2C. This structure was conserved over a range of community detection parameter values that, roughly speaking, must be defined to set the “spatial” and “temporal” resolutions of community identification (see Methods for detailed description). Across all temporal

resolutions considered here, consensus community partitions with a spatial resolution parameter $0.975 \leq \gamma \leq 1.01$ possessed exactly four communities.

For the standard parameter choice (temporal and spatial resolution parameters both set to 1), the four identified communities had distinct compositional characteristics. These communities were largely bilaterally symmetric, with analogous brain regions in each hemisphere assigned to the same community 71% of the time. The four communities correspond roughly to a visual core, a somatomotor-attention core, a default mode core, and a control core. The compositions of these four communities are shown in Figure 5.3A. The composition value was calculated by summing the total number of instances in which a node belonging to an *a priori* functional-anatomical network [144] also belonged to the community identified in the consensus community partition.

The core communities identified here were named based on the highest representation of nodes belonging to *a priori* functional networks. The visual core was 81% composed of visual network nodes and had an average size of 51.5 nodes per day. The somato-attention core was composed of 52% somatomotor, 27% salience-ventral attention, and 13% dorsal attention network nodes and had an average size of 134.1 nodes per day. The default mode core consisted of 56% DMN nodes and approximately 10% of each control, limbic, and temporoparietal network nodes and contained 132.9 nodes on average per day. Finally, the control core consisted of 48% control and 28% dorsal attention network nodes and contained 96.5 nodes on average per day. Importantly, for all parameter combinations in which four communities were detected, the composition of these communities was consistent (Appendix D.2). These community partitions were also stable across the entire menstrual cycle. Specifically, 315 of the 415 nodes (75.9%) did not change community affiliation across the 30-day experiment.

Taken together, these results suggest the presence of a stable solution to the dynamic community detection algorithm and a reliable coarse-grained community architecture

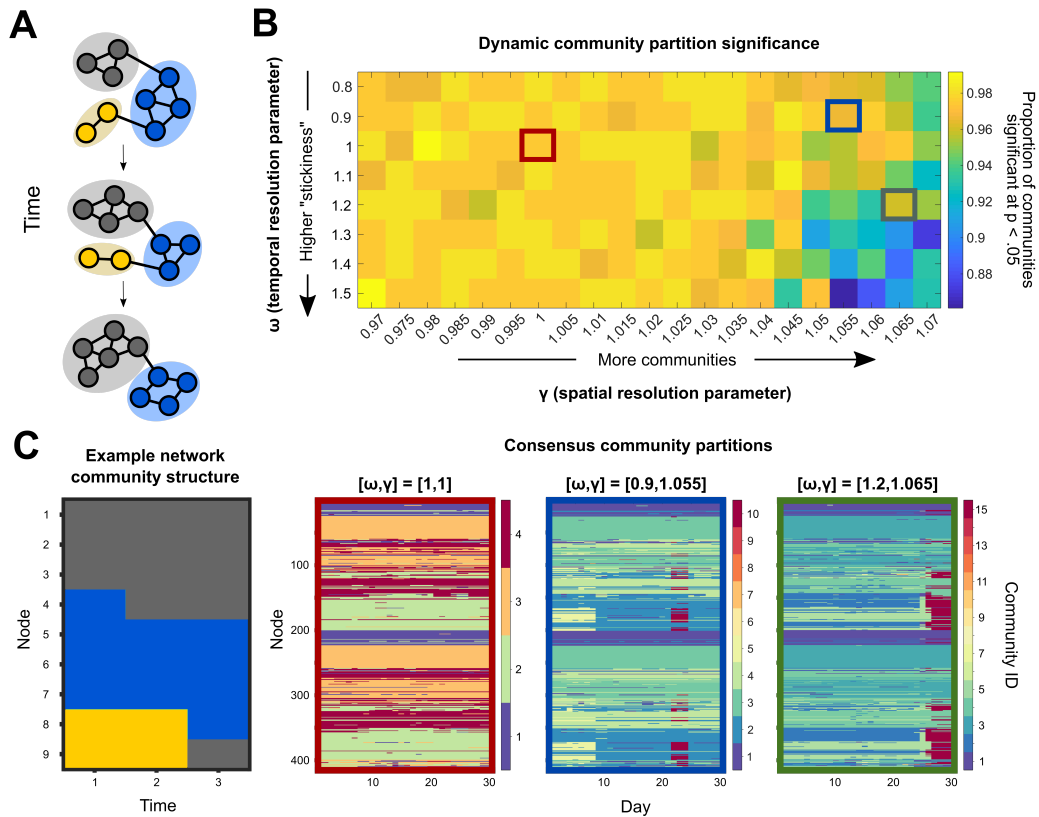


Figure 5.2: **Dynamic community detection identified changing modular structure over time at multiple scales.** **A.** A toy network example illustrates the dynamic community detection algorithm. For each time point, every node is assigned to a community so as to maximize the strength of intra-community connections relative to inter-community links while also taking community assignments over time into account (Eq. 1). In this case, three communities are identified and denoted by color. **B.** To assess temporal structure in the 28andMe resting-state fMRI data, community assignments were calculated for a range of parameter values. In this procedure, two parameters, ω and γ , specify the temporal and spatial scales of analysis, respectively. After performing 50 runs of the community detection algorithm for each parameter combination, the statistical significance of each community partition relative to a random null model was calculated. The color for each entry in the heat map indicates the proportion of communities at that parameter combination which are significant at the $p < .05$ level. **C.** Consensus partition structure varied according to the choice of resolution parameters. The example network community structure (left) changes at each time point, with node community assignment given by color on the y-axis and time indicated on the x-axis. For three different parameter combinations (outlined in red, blue, and green, respectively), the consensus partitions varied in the total number of communities identified, ranging from four to fifteen, with more communities identified when the temporal resolution was low and the spatial resolution was high.

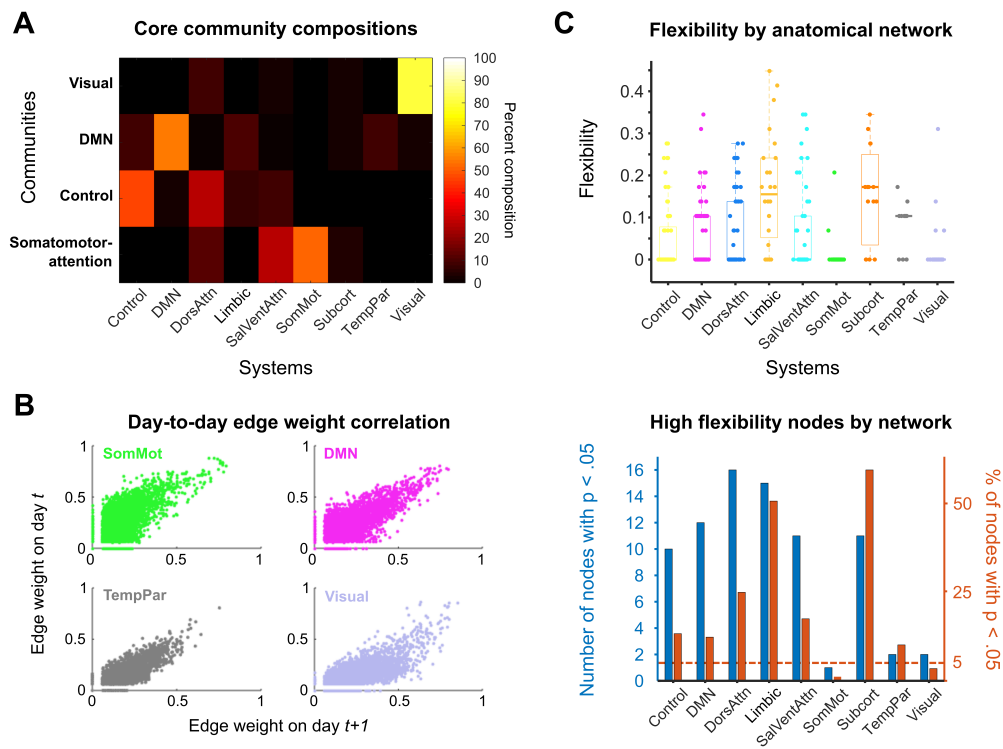


Figure 5.3: **Dynamic community detection uncovered stable cores across a complete menstrual cycle.** **A.** Four core communities (y-axis) were consistently identified in the 28andMe dataset across spatial and temporal resolution parameter values. For these parameter combinations, the compositions of the visual, default mode, control, and somatomotor-attention network cores are shown as a heat map, with color corresponding to the percentage of nodes in a community belonging to a functional-anatomical network. **B.** The four networks that constituted the hubs of the core communities possessed stable pairwise connectivity between nodes across days. Scatter plots show the day-to-day correspondence between edge weights for all of the nodes of the somatomotor, default mode, temporoparietal, and visual networks on days t and $t + 1$. These network edges had Pearson correlation coefficients of 0.379, 0.573, 0.590, and 0.538, respectively. **C.** The subcortical, limbic, and dorsal attention networks exhibited the highest median node flexibility. Top: Normalized flexibility values for each node over the entire cycle are plotted as points, with color indicating network affiliation. Thick horizontal lines on box plots indicate median values. A flexibility value of 1 indicates that a node changes community assignment at each possible time point, whereas a value of 0 indicates that the node never changes community assignment. Bottom: A 95% cutoff value is calculated using the flexibility values for each node over all 50 community detection runs. For each functional-anatomical network, the blue bar indicates the number of nodes belonging to that network which have flexibility values above the cutoff threshold. The red bars indicate the proportion of nodes in each network that surpass the cutoff value (i.e. the value for each blue bar is normalized by the number of nodes in the network). Once again, limbic, subcortical, dorsal attention, and control networks contained the highest proportion of highly flexible nodes.

present in the data. In several functional-anatomical networks, there was little to no modification of network architecture over time; for instance, greater than 85% of nodes in each of the somatomotor, default mode, temporoparietal, and visual networks did not change community affiliation over the entire menstrual cycle. The strong day-to-day correlations between edge weights in these networks (Figure 5.3B) reinforce the existence of these stable cores.

Functional-anatomical networks exhibited distinct patterns of flexibility

Though network community structure was stable over a complete menstrual cycle when classifying nodes into four communities, specific nodes did change community affiliation at levels above chance when modifying the sensitivity of the community detection algorithm. Specifically, when γ , the spatial resolution parameter, was increased, the dynamic community detection algorithm subdivided the four core communities into smaller communities, providing a finer-grained classification of subnetwork structure. At an intermediate parameter combination ($\omega = 1$, $\gamma = 1.05$), nine communities significant at the $p < .05$ level were identified over the course of the experiment, as visualized in Figure 5.2C (blue outlines). The subsequent analysis uses community partitions at this parameter combination, but the results were consistent across a range of neighboring parameter values (Appendix D.3).

This “higher-resolution” partition revealed trends in functional organization over time that were not observable with coarser partitions. First, inspecting the median flexibility value, or the proportion of times a node changed community affiliation out of the total possible number of changes, demonstrates that functional-anatomical networks possessed distinct flexibility distributions (Figure 5.3C, top). The limbic, subcortical, dorsal atten-

tion, and control networks were overrepresented in terms of highly flexible nodes relative to a null hypothesis (Figure 5.3C, bottom).

Fine-scale community reorganization occurred on experiment day 23 and persisted until day 25, as illustrated in Figure 5.4A. Across these days, 62 nodes belonging to the default mode core community split from the default mode core community to transiently form a small, strongly connected community. This was the only large-scale reorganization event detected during the experiment, as indicated by the nodal flexibility values illustrated in Figure 5.4B. Global flexibility was significantly higher (Wilcoxon rank-sum test, $p < .05$) during ovulation (days 23-25) than during follicular or luteal phases (days 11-22 and 26-10, respectively). Specifically, global mean flexibility during the ovulatory window was 0.142, whereas flexibility during follicular and luteal phases was 0.049 and 0.050, respectively.

Notably, 31 (50%) of the nodes in the community that emerged during the ovulatory window belonged to the DMN, 12 nodes (19%) belonged to the temporoparietal network, and 7 (11%) were subcortical regions (as defined by functional - anatomical atlases [144, 149], Figure 5.5A). The functional-anatomical network memberships of the node-node pairs exhibiting the strongest increases in coherence (top 5%) indicated that enhanced connectivity between DMN nodes drove this community split, as opposed to DMN nodes being “converted” to a new community via increased connectivity to non-DMN regions (Appendix D.8). More specifically, nodes within prefrontal regions belonging to DMN subnetwork B drove this reorganization event, as 118 of the 371 (32%) strongest increases in coherence occurred between nodes in this subnetwork.




Figure 5.4: **Fine-grain community partitioning revealed a bifurcation in the default mode core during ovulation.** **A.** When the spatial resolution parameter (which alters the size of communities identified by dynamic community detection) was increased from the standard value, the four core communities identified previously were subdivided into smaller subcommunities (reproduced from Figure 5.2C). Here, a split in the default mode core community (light blue) appeared at day 22 (red-orange), concomitant with ovulation and a spike in sex hormones. This community (red) rejoined the default mode core on day 25. For illustrative purposes, only the consensus partition for one parameter value is shown, but this trend was consistent across nearby parameter combinations (Appendix D.3). **B.** Shown are flexibility values for each node by menstrual cycle phase. Color in each region indicates flexibility value, with hotter colors indicating higher values. The following days of the experiment corresponded to the phases of the menstrual cycle: follicular, days 11-22; ovulatory, days 23-25; luteal, days 1-10 and 26-30. Flexibility values are noticeably higher in many regions from the temporoparietal, limbic, subcortical, and default mode networks during the ovulatory phase compared to the follicular and luteal phases.




Figure 5.5: Nodes in a default mode subnetwork drove community bifurcation via strong increases in coherence. A. The newly formed functional community on day 23 and 24 contained 62 nodes that belonged to the community on both days. The functional-anatomical network and subnetwork affiliations of these nodes are shown on the left and right, respectively. The new community contained 31 DMN nodes (50%), 12 temporoparietal nodes (19%), and 7 subcortical nodes (11%). **B.** The edges that exhibited large weight changes from day 22 to day 23 (top 5% of changes, left) were predominantly within-network connections between DMN network nodes (118/371). Examining subnetwork structure reveals that all of the strongly enhanced connections between nodes in the DMN belonged to subnetwork B, indicating that this subnetwork, which consists of regions in prefrontal cortex, drove the default mode core community bifurcation at ovulation.

Network reorganization timing coincided with peaks in hormone levels during ovulation

Mean flexibility of each network over a 5-day sliding window is depicted in Figure 5.6A. The DMN, temporoparietal, subcortical, and limbic networks exhibited peaks in flexibility at days 23 and 24 of the experiment, coincident with the peaks in estradiol, LH, and FSH which are a hallmark signals of the ovulatory window (Figure 5.6B). To determine whether the bifurcation of the default mode core community was significantly associated with sex hormones, we compared functional-anatomical network flexibility values to serum hormone levels.

To assess the temporal relationship between network flexibility values and sex hormones, cross-covariance structure between each time series was calculated. The control, default mode, limbic, salience/ventral attention, subcortical, and temporoparietal networks had maximum cross-variance values greater than 0.6 (where maximum value of 1 indicates fully shared covariance structure and 0 indicates no covariance) with estradiol, which were significant when compared to cross-covariance values for a null model of time-permuted estradiol levels (Bonferroni-corrected at $p < .05$). Each network except for the control and attention networks had maximum cross-variance values greater than 0.6 with LH as well (permutation test, $p < .05$ after Bonferroni correction). In each case, maximum cross-covariance values occurred at lags less than 2 days and no other significant cross-covariance structure existed, indicating that most functional communities exhibited changes in composition concurrent with significant rises in estradiol and LH levels.

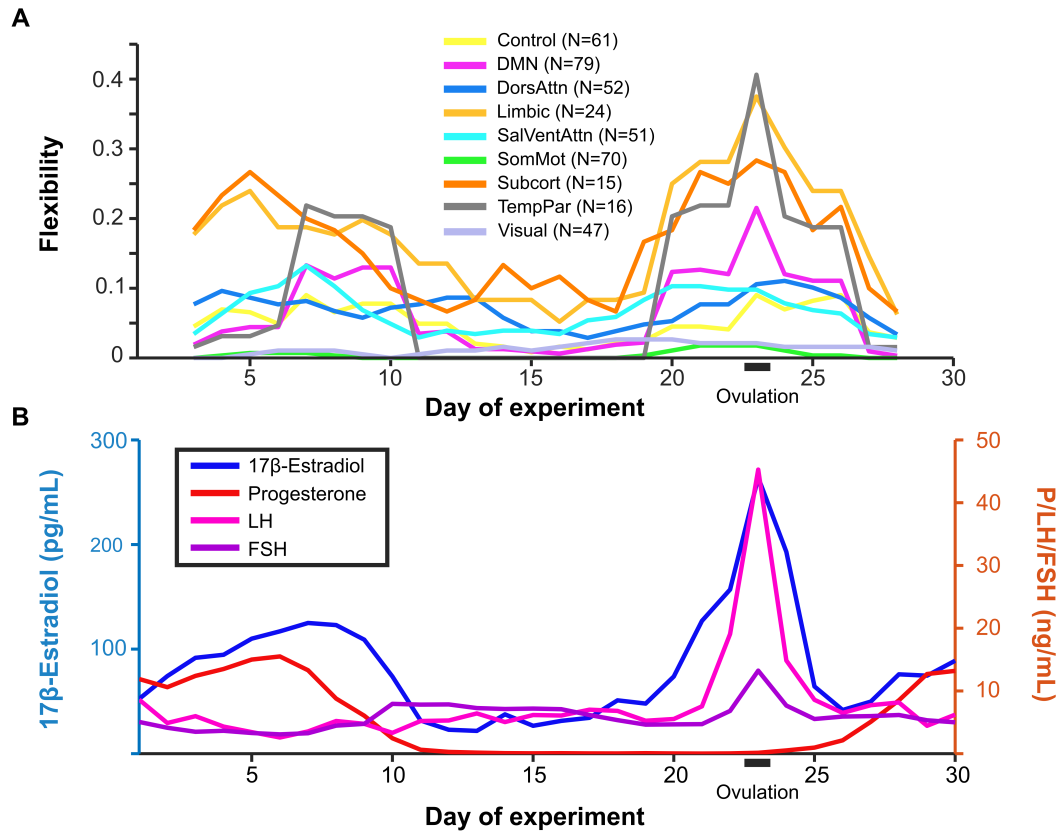


Figure 5.6: **Community reorganization was temporally localized to ovulation.** Changes in community assignment (**A**) were coordinated and closely tracked the timing of spikes in serum hormone concentrations (**B**). Prior to day 20 of the experiment, all networks except for the subcortical network exhibited low baseline rates of flexibility (mean = 0.04). However, several networks exhibited sharp increases in flexibility between days 20 and 26, indicating brain-wide functional reorganization during the ovulatory window. The pattern of flexibility shown here corresponds to the network reorganization observed for dynamic community detection performed with the parameter combination $\omega = 1$, $\gamma = 1.05$ (blue outline in Figure 5.2). To note, flexibility is calculated over a 5 day sliding window.

5.3 Discussion

In this study, we applied DCD to data from a densely sampled female who underwent 30 consecutive days of brain imaging and venipuncture to investigate the extent of intrinsic spatiotemporal functional reorganization over a menstrual cycle. We identified four stable community cores across the cycle, represented here as visual, somatomotor, default mode, and control network cores; interestingly, the exception to this stability occurred simultaneously with peaks in estradiol, LH and FSH. During this event, we observed a transient reorganization of the default mode core into a newly formed community, as well as increases in nodal flexibility among prefrontal, limbic, and subcortical nodes. Taken together, our results suggest that the interplay between the nervous and endocrine systems over a menstrual cycle result in temporary, localized patterns of brain network reorganization occurring during ovulation. These results highlight DCD as a new avenue for investigating the intricate relationship between sex hormones and human brain dynamics.

Dynamic community detection characterizes network-specific functional stability across a menstrual cycle

Dense-sampling, deep-phenotyping studies offer new ways to investigate intra/inter-individual variability in functional brain networks by identifying features of rs-fc that are stable traits within an individual or change in conjunction with biological factors and state-dependent variables [140, 142]. Recent dense-sampling studies have shown that frontoparietal regions/networks exhibit high degrees of intra-individual rs-fc stability while also being characteristically unique across individuals, suggesting that these higher-order regions may be especially critical for uncovering individual differences in brain function and improving applications into personalized medicine [150, 142]. Our

findings provide new insight towards the ongoing explorations into stability within functional brain networks. In this dataset, frontal control and DMN nodes exhibited high day-to-day connection weight correlations and low propensity to change functional community membership over the experiment (Figure 5.3), while, on average, somatomotor, temporoparietal, visual, and salience/ventral attention networks were also largely stable. Therefore, our results align with previous research suggesting both a high degree of network stability in resting-state networks over a menstrual cycle[138] and in individuals over time [140, 142, 150].

In conjunction with this observed stability, network-specific changes in functional community organization were also identified. Control subnetwork C, encompassing posterior cingulate cortex/precuneus regions, was the most flexible functional subnetwork identified, with 10 of the 12 nodes exhibiting significantly higher than expected flexibility (i.e. how often a node switches community affiliation, see Appendix D.7). Limbic and subcortical networks displayed intermediate levels of flexibility. Regions from these systems receive input from and project to many cortical areas and are implicated in functions such as sensorimotor integration via the cortico-basal ganglia-thalamo-cortical loop [151]; therefore, the high degree of flexibility observed here may reflect the tendency of these systems to serve as relays between functionally segregated communities.

Particular changes in rs-fc were significantly related to the sharp rises in sex hormones seen across the ovulatory window. During this time, we observed a spatially-specific transient reorganization of the DMN, during which nodes from the temporoparietal, limbic, subcortical, and default mode networks split from the default mode core to form a short-lived community (2 days) before rejoining the original core community. Using time-lagged analyses, Pritschet and colleagues previously reported that within-network connectivity of the DMN was regulated by previous states of estradiol [143]. Here, we expand on this finding and identify a subnetwork of the DMN that is likely driving this

reorganization. Notably, regions constituting this new community are located in PFC, an area exquisitely sensitive to sex steroid hormones [152] where, for instance, nearly 50% of pyramidal neurons in the dorsolateral PFC (dlPFC) express ER-alpha [128]. Together, this presents the possibility that endocrine signaling may, in part, regulate intrinsic brain dynamics within the frontal cortex.

Neurobiological interpretations of sex hormones on PFC function

Cross-species investigations have established estrogen's ability to shape the PFC [153, 154, 152, 133, 155]. In rodents, estradiol increases fast-spiking interneuron excitability in deep cortical layers [156]; in non-human primates, estradiol treatment increases dendritic spine density in dlPFC neurons [157] and this potentiation is observed only if the treatment is administered in the typical cyclical pattern observed across a menstrual cycle. In parallel, human brain imaging studies have implicated estradiol in enhancing the efficiency of PFC-based circuits. In cycling women performing a working memory task, PFC activity is exaggerated under low estradiol conditions and reduced under high estradiol conditions [133]. Similarly, when estradiol declines across the menopausal transition, working-memory related PFC activity becomes more exaggerated despite no differences in task performance [155]. Examining rs-fc across the cycle, Petersen and colleagues found that women in the late follicular stage (near ovulation) showed increased coherence within the default mode and executive control networks compared to those in luteal stages [134]. Our findings extend this body of work by demonstrating that dlPFC nodal flexibility tracks significantly with sharp increases in estradiol and LH across the cycle, which may support the brain's ability to reorganize at the mesoscale level.

This tight temporal coupling highlights the potential for a mechanistic link between endocrine signaling and large-scale network reorganization. While future multimodal

brain imaging studies are needed to establish this link, one possible neurobiological mechanism of action may be through estradiol's interaction with the dopaminergic system. For instance, the PFC is innervated by midbrain dopaminergic neurons that enhance the signal-to-noise ratio of PFC pyramidal neurons and drives cortical efficiency [158]. In turn, estradiol enhances dopamine release and modifies the basal firing rate of dopaminergic neurons, providing one explanation for how alterations in estradiol could impact cortical efficiency.

Implications for cognition and disease

Several studies have begun utilizing DCD to relate “task-free” and “task-based” functional network reorganization to cognitive performance. High levels of nodal flexibility have been associated with enhanced performance on working memory tasks [159], improved learning of a motor task [160], and visual cue learning [161]. In each study, flexibility was associated with performance in regions known to underlie each task, implicating frontal, motor, and visual cortical cortices and subcortical structures such as thalamus and striatum. Notably, similar associations were not observable when analyzing these experiments through region-based activation patterns alone, indicating that temporal organization of brain-wide functional activity (e.g. dynamic community structure) may provide important information related to cognitive functioning that might be missed with traditional analyses.

Indeed, Mattar et al. used DCD to characterize cognitive systems like those defined here [144] in a 64-task battery, demonstrating that functional networks fluidly reconfigure to form new cohesive communities under different task settings [162]. Similar work has revealed that primary motor, visual, and auditory regions typically participate in a single or a small number of functional networks during various tasks, whereas “hub” regions in

frontal cortex, including precuneus and posterior cingulate gyrus participate in multiple functional networks [163]. Together, these studies indicate that network-specific temporal reconfiguration of functional connectivity has implications for a wide variety of cognitive functions. While whole-brain activity patterns during task-free states differ from that of goal-directed cognitive states, the capacity for the brain to fluctuate between integrated and segregated (modular) states at rest allows for rapid and efficient transitions to various task states [164, 165, 166]. Here, we leverage these techniques to characterize the brain’s response to both subtle and pronounced hormonal changes typical of a menstrual cycle.

Highly flexible nodes were identified in precuneus and posterior cingulate gyrus, with changes in community affiliation occurring simultaneously with sharp peaks in estradiol and LH levels, raising the possibility that hormonal fluctuations could also be associated with task-based network reorganization. For instance, if high levels of estradiol increase nodal flexibility among hub regions in the PFC, one would predict that performance on PFC-dependent tasks will improve. Further, pregnancy—a period of profound hormonal change—leads to long-lasting gray matter reductions in regions within the default mode network [167]. Therefore, future work examining whether task-based functional brain networks undergo transient changes in flexibility and community structure both across the menstrual cycle and during other hormonal transition periods, and whether this impacts cognitive performance, will be imperative.

Examining how large-scale brain networks are disrupted between healthy and patient populations may enhance our understanding of neurological conditions [168]. Notable intrinsic connectivity differences within the DMN are observed among individuals with depression [169] and Alzheimer’s disease [170] — two conditions that display a sex-skewed prevalence towards women [171]. Recent studies have applied DCD methods to characterize functional brain network reconfigurations in different disease states: region-specific flexibility at rest has been linked to symptom severity in autism spectrum disorder [172]

and a recent investigation used DCD to associate pronounced community reorganization during seizures with poorer surgical outcomes [173]. Here, using similar methods, we demonstrate that high estradiol days are associated with significant reorganization of the default mode network and increased flexibility of several brain networks. Understanding the relationship between brain network reconfiguration (time-varying communities) and the endocrine system (dynamic fluctuations in sex hormones) may offer new ways to understand complex neurological conditions, especially those with pronounced sex differences in disease prevalence.

Limitations and future directions

The following limitations should be taken into consideration. First, this study involved densely sampling a single female over one complete menstrual cycle, hindering our ability to generalize these findings to other individuals. Therefore, it is critical for this approach to be extended to a larger and more diverse set of women to establish the consistency of these results while taking individual differences into consideration. Second, we used a well-established group-based atlas to mitigate the limitations inherent to a single-subject design and improve generalizability [144]. However, recent work has demonstrated that group-based atlases can lead to loss in individual-level specificity and overlook meaningful spatial reconfigurations in parcellations themselves [174]. Future work using an individual-derived atlas is needed to confirm whether these results are stable across various parcellation applications. Furthermore, an ongoing debate in network neuroscience surrounds test-retest reliability and what constitutes a “substantial” amount of data per individual. While some studies suggest that large amounts of data (> 20 minutes) is needed [142], others contend that shorter durations (5-15 minutes) of sampling is sufficient to achieve reliability [175, 141]. Repeating this experiment under

longer scanning durations (>10 minutes per day) will be critical for exploring the degree of network stability across the menstrual cycle. Finally, this work only considers network reconfiguration across multiple sessions. Due to the aforementioned concerns about test-retest reliability when using short rs-fc windows for analysis and our sample size of $N = 1$, we opted not to examine within-session reorganization. However, previous studies have used DCD methods to examine flexibility on shorter time scales [159, 176], so extending this experimental protocol to a larger sample could allow for future analysis of within-session network reorganization.

5.4 Conclusion

In sum, we demonstrate that resting-state functional connectivity is largely stable within an individual over the course of a complete menstrual cycle. The exception to this stability occurs around the ovulatory window, during which peaks in sex hormones result in temporary patterns of brain network reorganization largely localized within areas of the default mode network. Historically, brain-level phenomena resulting from hormone fluctuations have been treated as an unwanted source of variance in population studies and, consequently, studies of this relationship are sparse and underpowered. This work demonstrates that dynamic network methods can reveal important, transient effects of sex hormones that may be overlooked by traditional approaches and provides a novel template for examining the nature of human brain-endocrine relationships.

5.5 Methods

28andMe experimental protocol

Data was collected and preprocessed as reported in [143]; methods briefly reproduced here. The participant was a right-handed Caucasian female, aged 23 years for duration of the study. The participant had no history of neuropsychiatric diagnosis, endocrine disorders, or prior head trauma. She had a history of regular menstrual cycles (no missed periods, cycle occurring every 26-28 days) and had not taken hormone-based medication in the 12 months prior to the study. The participant gave written informed consent and the study was approved by the University of California, Santa Barbara Human Subjects Committee.

The participant underwent daily testing for 30 consecutive days, with the first test session determined independently of cycle stage for maximal blindness to hormone status. The participant began each test session with a daily questionnaire (9:00am) followed by a time-locked blood sample collection 10:00am (± 30 min). Endocrine samples were collected, at minimum, after two hours of no food or drink consumption (excluding water). This was followed by a one-hour MRI session (11:00am) consisting of structural and functional MRI sequences. To note, the participant refrained from consuming caffeinated beverages before each test session.

A licensed phlebotomist inserted a saline-lock intravenous line into the dominant or non-dominant hand or forearm daily to evaluate hypothalamic-pituitary-gonadal axis hormones, including serum levels of gonadal hormones (17β -estradiol, progesterone and testosterone) and the pituitary gonadotropins luteinizing hormone (LH) and follicle stimulating hormone (FSH). One 10cc mL blood sample was collected in a vacutainer SST (BD Diagnostic Systems) each session. The sample clotted at room temperature for 45 min. until centrifugation (2,000 g for 10 minutes) and was then aliquoted into three 1

ml microtubes. Serum samples were stored at -20 C until assayed. Serum concentrations were determined via liquid chromatography-mass spectrometry (for all steroid hormones) and immunoassay (for all gonadotropins) at the Brigham and Women-s Hospital Research Assay Core.

fMRI data acquisition and preprocessing

The participant underwent a daily magnetic resonance imaging scan on a Siemens 3T Prisma scanner equipped with a 64-channel phased-array head coil. First, high-resolution anatomical scans were acquired using a T1-weighted magnetization prepared rapid gradient echo (MPRAGE) sequence (TR = 2500 ms, TE = 2.31 ms, TI = 934 ms, flip angle = 7°; 0.8 mm thickness) followed by a gradient echo fieldmap (TR = 758 ms, TE1 = 4.92 ms, TE2 = 7.38 ms, flip angle = 60°). Next, the participant completed a 10-minute resting-state fMRI scan using a T2 -weighted multiband echo-planar imaging (EPI) sequence sensitive 468 to the blood oxygenation level-dependent (BOLD) contrast (TR = 720 ms, TE = 37 ms, flip angle = 56°, multiband factor = 8; 72 oblique slices, voxel size = 2 mm). In an effort to minimize motion, the head was secured with a custom, 3D-printed foam head case (<https://caseforge.co/>) (days 8-30 of Study 1). Overall motion (mean framewise displacement) was negligible, with fewer than 130 microns of motion on average each day.

Initial preprocessing was performed using the Statistical Parametric Mapping 12 software (SPM12, Wellcome Trust Centre for Neuroimaging, London) in MATLAB. Functional data were realigned and unwarped to correct for head motion and the mean motion-corrected image was coregistered to the high-resolution anatomical image. All scans were then registered to a subject-specific anatomical template created using Advanced Normalization Tools (ANTs) multivariate template construction. A 5 mm full-width at

half-maximum (FWHM) isotropic Gaussian kernel was subsequently applied to smooth the functional data. Further preparation for resting-state functional connectivity was implemented using in-house MATLAB scripts. Global signal scaling (median = 1,000) was applied to account for fluctuations in signal intensity across space and time, and voxelwise timeseries were linearly detrended. Residual BOLD signal from each voxel was extracted after removing the effects of head motion and five physiological noise components (CSF + white matter signal). Motion was modeled using a Volterra expansion of translational/rotational motion parameters, accounting for autoregressive and nonlinear effects of head motion on the BOLD signal. All nuisance regressors were detrended to match the BOLD timeseries.

Functional network nodes were defined based on a 400-region cortical parcellation and 15 regions from the Harvard-Oxford subcortical atlas. For each day, a summary time-course was extracted per node by taking the first eigenvariate across functional volumes. These regional timeseries were then decomposed into several frequency bands using a maximal overlap discrete wavelet transform. Low-frequency fluctuations in wavelets 3-6 (0.01-0.17 Hz) were selected for subsequent connectivity analyses. Finally, we estimated the spectral association between regional timeseries using magnitude-squared coherence: this yielded a 415 x 415 functional association matrix each day, whose elements indicated the strength of functional connectivity between all pairs of nodes (FDR-thresholded at $q < .05$).

Dynamic community detection and analysis

A multi-layer tensor ($415 \times 415 \times 30$) was constructed from the association matrices described above for network analysis. Each layer corresponded to the strictly positive, weighted, FDR-thresholded rs-fc association matrix for the corresponding day of the

experiment. Interlayer links were added only between adjacent layers. Communities in resting-state connectivity were identified by maximizing multislice modularity, given by

$$Q = \frac{1}{2\mu} \sum_{ijlr} \left((A_{ijl} - \gamma_l \frac{k_{il}k_{jl}}{2m_l} \delta_{lr} + \delta_{ij} \omega_{jlr}) \delta(g_{il}, g_{jr}) \right), \quad (5.1)$$

where μ is the total edge weight in the network, i and j index nodes in slices l and r , A is the adjacency matrix containing edge weights between nodes and slices, γ is the structural resolution parameter, k_{il} is the strength of node i in slice l , m_l is the total nodal strength in slice l , δ is the Kronecker delta, ω is the temporal resolution parameter, and g is the community assignment index [148].

Community assignments that maximize modularity were determined 150 times over a grid of parameter values $(\gamma, \omega) = [.97, 1.07] \times [.8, 1.5]$ using the `genlouvain` function from Jeub et al. in MATLAB 2019a [177]. From these community assignments, the consensus partition for each parameter combination was determined using the `consensus_similarity` function from the Network Connectivity Toolbox (NCT, <http://commdetect.weebly.com/>).

Node flexibility is defined as the proportion of times a node changes community assignment out of all possible opportunities to change its assignment. Thus, a flexibility value of 1 indicates that a node changes community membership at every time step and a value of 0 indicates that it never changes communities. Partition significance, node flexibility, and persistence were also calculated using functions from the NCT [160]. Cross-covariance values were calculated and statistical tests were performed using built-in MATLAB functions.

Head motion was low (< 220 microns), not significantly associated with hormone concentrations (all pairwise Pearson correlations $> .05$ after Bonferroni correction), and was nearly identical between days 20 and 24 of the experiment (when reorganization

occurred), suggesting that head motion is not a confounding factor when considering community reconfiguration. On the day of the experiment with the fewest connections (day 26), the network had an edge density of 0.9317 (i.e. 93.17% of possible edges have non-zero values) and the median density was 0.9713. This represents a 4% difference in density and density was not significantly correlated with hormone levels, so we do not believe the community detection algorithm was biased by disparities in edge density.

Chapter 6

Concluding Remarks

Here, a diverse set of neuroscientific systems were analyzed using tools from network science, statistics, and computational ethology. In Chapter 2, neuronal cultures *in vitro* were demonstrated to develop stereotyped activity patterns in the absence of stimulation or reinforcement. Several advances in technology and experimental protocols would provide avenues for building upon this work. First, recording apparatuses that allow for greater spatial resolution (such as CMOS arrays or optical recording methods) will improve the reliability and granularity of the results reported here, as one major concern with this work is that the currently used MEA technology severely undersamples the total number of neurons present in the experiment. Next, application of spike sorting techniques would likely improve these results by characterizing trajectories in terms of individual neuronal firing rates rather than summed electrode activity. As spatial resolution of recording techniques is increased, this step will be increasingly necessary, since densely packed electrodes are more likely to pick up signal from multiple neurons concurrently. Finally, development and implementation of technologies that can stimulate or ablate the activity of sets of neurons would provide finer control over neuronal firing activity, allowing researchers to attempt to construct, modify, or “delete” activity patterns

from the repertoire. Such developments would have implications for models of learning and memory and further illuminate the capacity of *in vitro* cultures to serve as useful model systems.

In Chapters 3 and 4, *Drosophila* grooming behavior was characterized, revealing both conserved syntactic structure and significant phenotypic differences within and between species. Going forward, careful genomic and transcriptomic assessment of *Drosophila* populations could identify particular sets of genes implicated in behavioral differences; this work uses only approximate genetic distance as a measure of genotypic distance, so more rigorous quantification would enhance the results presented here. In a similar vein, cross-species circuit tracing comparisons could reveal the structural basis for the observed differences in grooming behavior by identifying differences in sensory circuit architecture, for example. More realistically, large-scale experiments utilizing targeted optogenetic stimulation could solidify these results by reducing sensory noise, providing a better estimate of the contribution of non-sensory factors to grooming behavior.

In Chapter 5, transient human functional brain network reorganization was linked to hormonal fluctuations during the menstrual cycle in a single subject. Most obviously, this analysis would be bolstered by increased sample size, sensitivity analysis across brain parcellation schemes, and extension to multiple menstrual cycles within the same subject. Additionally, analysis of task-based fMRI would provide insight into whether the observed resting-state changes are present in other contexts across the menstrual cycle. In general, widespread adoption of dense-sampling methodology in conjunction with dynamic network methods promises to illuminate meaningful relationships between the brain and other physiological systems.

Appendix A

This appendix contains additional figures to support the work presented in Chapter 2.

A.1 Supporting Figures

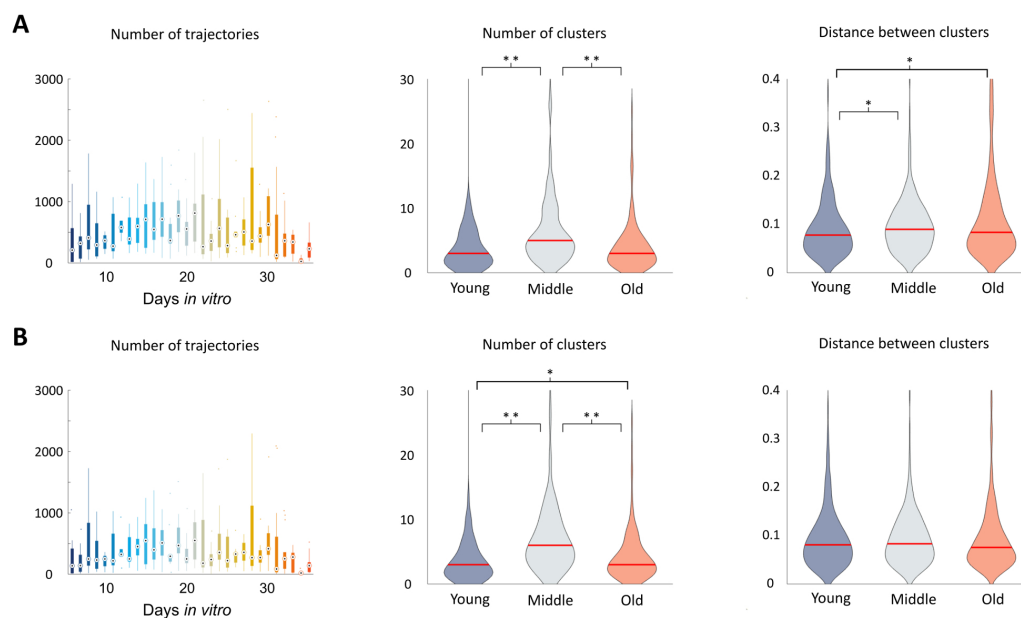


Figure A.1: **Clustering results are robust to choice of trajectory identification threshold.** A. Changing the trajectory identification threshold to $0.3n f_{max}$ reduces the number of trajectories identified overall but preserves the main results of clustering. B. Changing the trajectory identification threshold to $0.4n f_{max}$, which is twice the limit used in the manuscript, preserves the result related to the number of clusters in the data but does not preserve the result related to the separability of clusters. This suggests that trajectories with extremely high firing rates may resemble culture-wide synchronous bursts, or “avalanches”, and therefore lack the complex spatiotemporal structure identified at lower threshold values.




Figure A.2: **PCA results are robust to choice of trajectory identification threshold.** A. Changing the trajectory identification threshold to $0.3n_{f_{max}}$ preserves the similarity of principal components over time. B. Changing the trajectory identification threshold to $0.4n_{f_{max}}$ slightly changes the similarity values, but preserves the overall trend.

Appendix B

This appendix contains additional figures to support the work presented in Chapter 3.

B.1 Supporting Figures

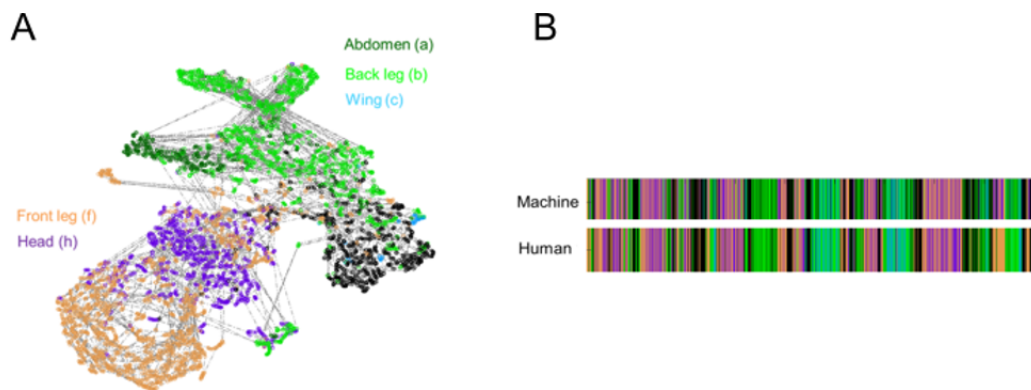


Figure B.1: **Machine classifier performance is similar to human classification accuracy.** A: Here, grooming behaviors are visualized in low-dimensional space using t-SNE, which preserves local distances between nearby points. Color indicates behavior type. Behaviors are well-separable by type. B: Example classifications of ethograms of grooming behaviors are shown. Human-algorithm agreement is about 75% (in terms of number of matching frames).

Duration permuted	$\mathbf{M}_{ij} = P(x_{action,j} x_{action,i})P(x_{duration,j})$
Order permuted	$\mathbf{M}_{ij} = P(x_j)$

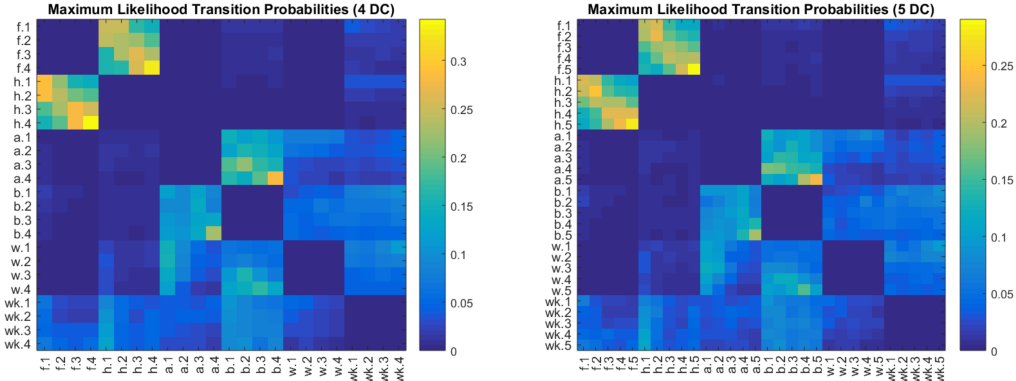


Figure B.2: **Transition probability matrices exhibit duration-dependent structure across many choices of binning schemes.** When transition probabilities are calculated from ethograms with 4 (left) or 5 (right) duration categories, fine-grain temporal structure is still present. The existence of this trend across binning schemes ensures that duration dependence is a phenomenon present in the data and not an artifact of data processing.

Table 1. Analytic formulas for null model transition matrices. Given in B.1 are the analytic formulas for transition matrix entries, \mathbf{M}_{ij} , for each null hypothesis transition matrix. Here, i denotes the row and j denotes the column of \mathbf{M} . Note that, in the three duration category case, each state has both an action and duration identifier (e.g. row h.s. corresponds to a head cleaning action in the short duration category). $P(x_{action,j}|x_{action,i})$ is the absolute probability of transitioning from the grooming action represented by row i to the action represented by row j , regardless of duration. $P(x_{duration,j})$ is the absolute probability of the grooming action represented by row j belonging to the duration category represented by row j . $P(x_j)$ is the absolute probability of state j .

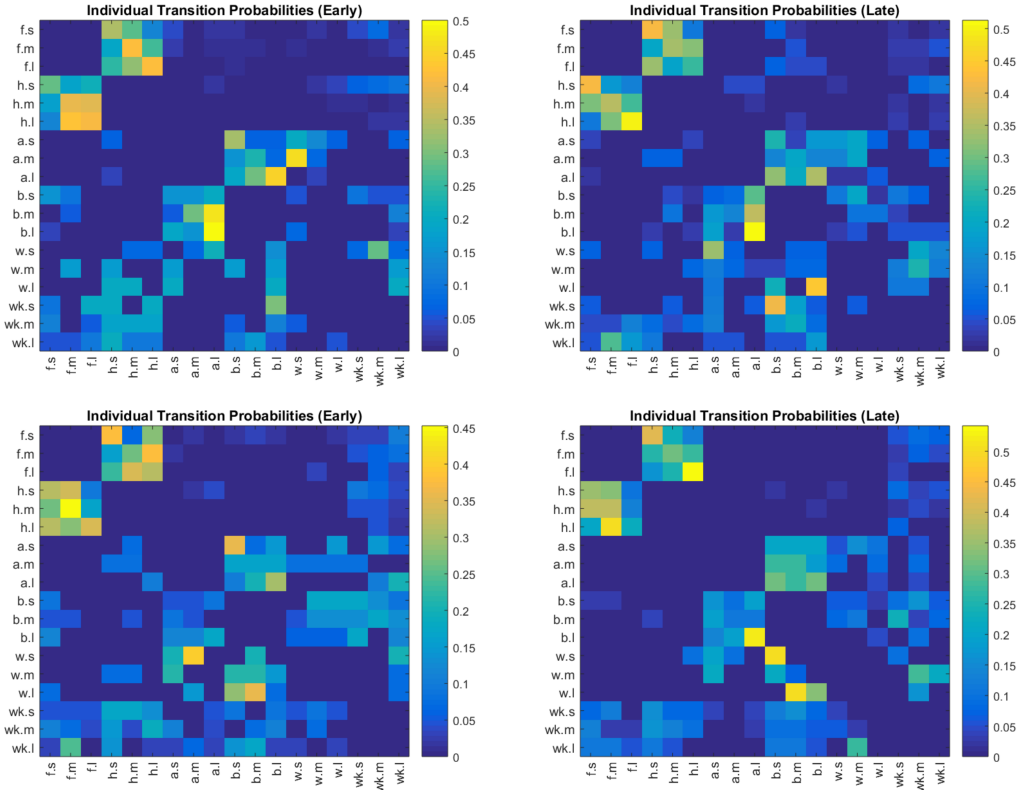


Figure B.3: **Individual flies possess stable syntax that resembles the group average.** Two examples of individual fly transition probabilities illustrate that, although syntax varies between individuals and over time, individuals are relatively stable and resemble the group average. Though we did not fully characterize individual differences here, future studies will aim to examine individuals more closely.

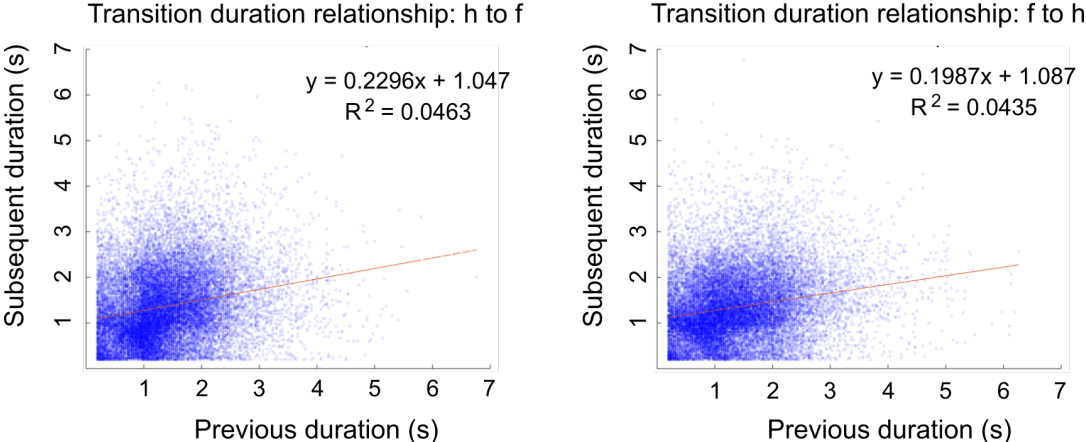


Figure B.4: **Consecutive anterior bout durations do not exhibit a linear relationship.** Plotted here is the relationship between consecutive anterior grooming action durations. Linear regression produces a poor fit, indicating that duration dependence is not linear.




Figure B.5: **Sampling early and late phases differently does not alter the observed syntax.** To ensure that our choice of boundary between early and late phases does not alter our results, we partition our data in a stricter manner so that the transition period which occurs during the middle of the grooming progression is excluded. Partitioning the data into thirds produces transition matrices with structure and stationarity similar to those shown the main text results.




Figure B.6: **Small residuals between early and late transition probabilities indicate that most grooming rules are stationary over time.** We find that 267 of the 270 (98.9%) non-zero entries change by less than 5% and 167 (61.8%) change by less than 1%. This was a very surprising finding and suggests that, although some non-stationarity does exist, it is much smaller than we had expected given that the average time spent performing different behaviors varies between the early and late phases.

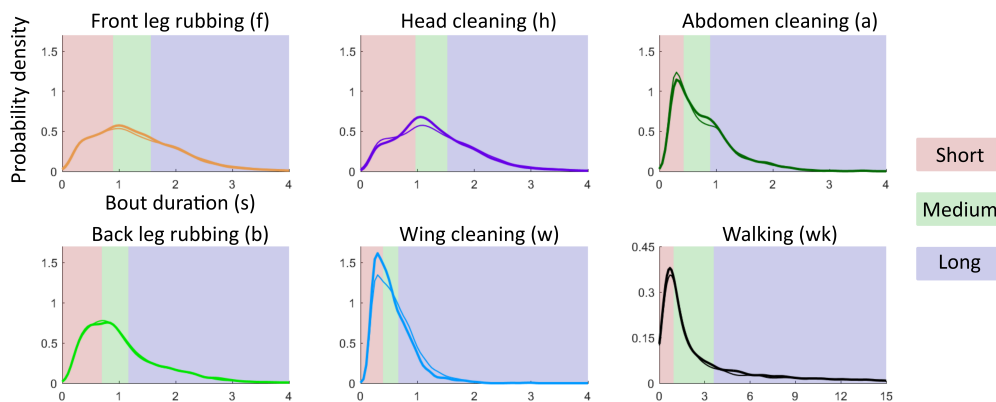


Figure B.7: **Bout durations do not change between early and late phases.** Probability density plots of bout duration distributions illustrate that bout durations remain consistent between early and late phases of grooming. Early phase data is plotted with a thick line. Late phase data is plotted with a thin line.

Appendix C

This appendix contains additional details about methodology and figures to support the work presented in Chapter 4.

C.1 Supporting Methods

Classification analysis

To determine whether flies were distinguishable on the basis of several grooming features, logistic regression classification and k-medoids clustering were applied. Logistic regression classifiers were trained using the MATLAB function `fitglm` with a logit link function. Models were validated using 5-fold cross-validation to prevent overfitting. For each grooming feature, PCA was first applied to reduce the dimensionality of feature vectors to 10 dimensions to ensure that the number of feature dimensions was smaller than the number of samples.

Separate classifiers were trained on each grooming feature (syntax, progression, behavioral proportions, and bout duration distributions) for pairwise comparisons within the following multi-class datasets: *erecta*, *santomea*, *sechellia*, *simulans*, and MaxVar; Canton-S, Berlin-K, Oregon-R, and w1118; Canton-S, Canton-S isoJ3, Canton-S isoJ1; Berlin-K, Berlin-K isoJ13; Oregon-R, Oregon-R isoJ5, Oregon-R isoJ8; w1118, w1118

isoJ5. Strains with “iso” designations refer to isogenized lines generated as described in the Methods section of the main text. When performing pairwise classification, sample sizes were matched so that chance classification accuracy would be 50%.

Finally, k-medoids clustering was applied 100 times to the same data used in the above pairwise syntax comparison using the MATLAB function `kmedoids`. For $k = 2$, the median accuracy of each classification is shown in Appendix C.12, with normalized mutual information between true and predicted labels serving as the accuracy measure. Normalized mutual information ranges from 0 to 1, with 0 indicating classification at chance and 1 indicating perfect classification.

Visualization

To visualize high-dimensional grooming features in low-dimensional space, t-Distributed Stochastic Neighbor Embedding (tSNE) was used [178]. Briefly, tSNE performs dimensionality reduction by converting high-dimensional pairwise Euclidean distances into conditional probabilities (or similarities). These probabilities are given by

$$p_{j|i} = \frac{e^{-\|x_i - x_j\|^2 / 2\sigma_i^2}}{\sum_{k \neq i} e^{-\|x_i - x_k\|^2 / 2\sigma_i^2}} \quad (\text{C.1})$$

where σ_i is the variance of a Gaussian distribution centered at x_i . These probabilities represent the likelihood that any two points would be nearest neighbors in low-dimensional space, assuming a Student-t distribution of pairwise distances between points. The low-dimensional counterpart joint probabilities, then, are given by

$$q_{ij} = \frac{(1 + \|y_i - y_j\|^2)^{-1}}{\sum_{k \neq l} (1 + \|y_k - y_l\|^2)^{-1}} \quad (\text{C.2})$$

where y_i is the low-dimensional embedding of x_i . To create a map between high and

low dimensions, the Kullback-Leibler divergence between P , the high-dimensional joint probability distribution, and Q , the low-dimensional distribution, are minimized using the cost function

$$\sum_i KL(P_i||Q_i) = \sum_i \sum_j p_{j|i} \log \frac{p_{j|i}}{q_{j_i}} \quad (\text{C.3})$$

where P_i is the high-dimensional conditional probability distribution of the point x_i and Q_i is its low-dimensional counterpart. A non-deterministic solution to this mapping cost function is found using gradient descent. Here, the `tsne` function in MATLAB was used to perform this computation. 5000 gradient descent steps were used, as solutions were found to converge sufficiently using that number of steps. The perplexity value, a hyperparameter that approximates the effective number of neighbors of a point, was set between 15 and 30. Note that typical perplexity values range between 5 and 50.

Markov chain entropy rate

Using first-order transition dynamics to define syntax, the Markov chain entropy rate was calculated as a measure of stereotypy in grooming. An entropy rate of 0 indicates that grooming transitions are perfectly predictable, while higher values indicate decreased stereotypy. Entropy was defined as

$$H = - \sum_{ij} \mu_i P_{ij} \ln(P_{ij}) \quad (\text{C.4})$$

where P_{ij} is the transition probability from action i to action j and μ is the steady-state action distribution.

Optogenetic stimulation

Protocols for optogenetic stimulation of grooming are reproduced here in brief from Zhang et al. (2020).

After cold anesthesia, flies were left to recover in recording a chamber for at least 20 min. Custom-made LED panels (LXM2-PD01-0050, 625nm) were used for light activation from below. 20 Hz 20% light duty cycle was used in all experiments. Light intensity was measured by Thorlabs S130VC power sensor coupled with PM100D console. The light intensity used for Bristle-spGAL4-1 flies was 0.84 mW/cm². The Bristle-spGal4-1 line is available from the Bloomington Stock Center (RRID: BDSC_71032).

C.2 Supporting Figures




Figure C.1: **Ethogram for non-*melanogaster* drosophilids.** Shown are the ethograms for *erecta*, *santomea*, *sechellia*, and *simulans* flies. Rows are individual ethograms, columns indicate frame. Behaviors are coded by color.




Figure C.2: **Ethogram for *melanogaster* flies.** Shown are the ethograms for each group of *melanogaster* flies analyzed here. Rows are individual ethograms, columns indicate frame. Behaviors are coded by color.




Figure C.3: **Ethogram for Canton-S individuals recorded across multiple sessions.** Shown are the ethograms for 15 Canton-S flies recorded for 3 sessions each over 3 consecutive days. Row labels indicate the individual fly, with rows after each label providing ethograms for sequential days (i.e. session 1, session 2, session 3). Columns indicate frame. Behaviors are coded by color.




Figure C.4: Different drosophilid species exhibit robust grooming responses in this assay. The amount of time spent grooming varies within and between species. However, regardless of genetic background, drosophilids spend several minutes performing grooming actions after exposure to a dust stimulus. Shown are violin plots of grooming proportion distributions for four non-*melanogaster* species (left) and five different *melanogaster* groups (right). White dots indicate median values. For all groups except *erecta* and Berlin-K, the median proportion of time spent grooming in a 27.8 minute recording after dusting is between 50 and 60% (13.9 to 16.7 minutes), indicating that nearly all drosophilids exhibit a robust grooming response.




Figure C.5: Grooming progressions are similar across drosophilid species. Shown are the group mean grooming progressions for non-*melanogaster* populations. Grooming proportions were calculated across 10 non-overlapping windows (approximately 3 min. each). Y-axis indicates the proportion of time spent in an action. X-axis indicates the window. Note that all populations exhibit many similar trends, including an initial high proportion of head grooming and front leg rubbing followed by an increase in other grooming actions. Many populations differ significantly in their propensity to walk or stand rather than groom. Though group means are shown here, individual vectors identical in format were used for classification analysis and visualization.

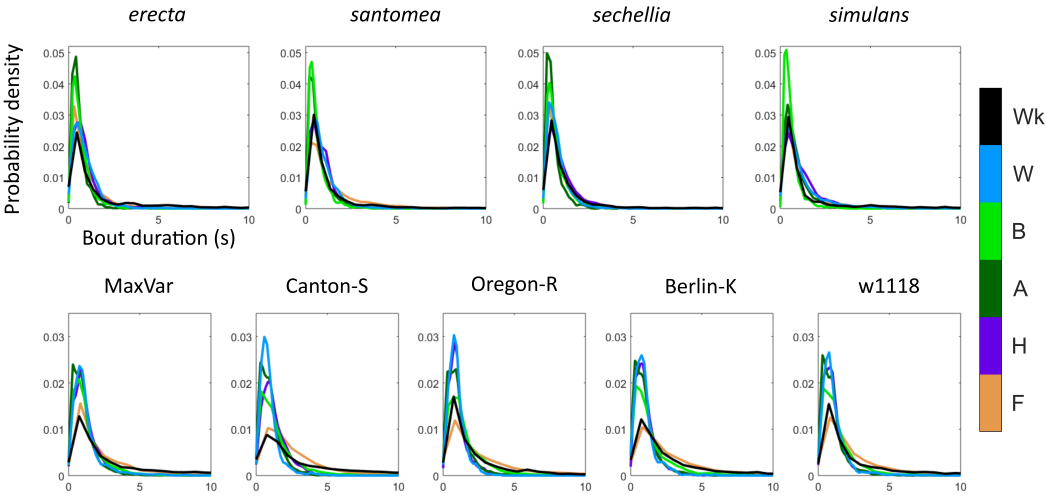


Figure C.6: **Bout duration distribution vectors of drosophilid species, *melanogaster* stocks, and Canton-S.** Shown here are the group mean bout duration distribution vectors for several populations. Y-axis indicates density. X-axis indicates bin. 20 bins of equal width were used, with the same bin width for all grooming actions. Non-grooming actions had larger bin widths due to longer average action durations for these actions. Note that non-*melanogaster* species perform a higher proportion of short actions, while *melanogaster* populations possess similar distributions to one another.




Figure C.7: **Grooming syntax separates species most strongly.** When comparing species, non-syntactic features of grooming do not exhibit the same degree of separability as grooming syntax. Shown here are t-SNE plots of each grooming figure, where points represent the grooming feature of an individual fly. Color indicates species. Qualitatively, a high degree of overlap between points indicates that grooming features are similar. Notably, grooming syntax provides the most distinct separation between species. As such, we focused on syntax analysis to a greater extent than other grooming features. However, bout duration distributions do separate *melanogaster* from non-*melanogaster* species, suggesting that action persistence is also an important syntactic feature.




Figure C.8: *Melanogaster* progressions differ in the amount of grooming and non-grooming actions. Shown are the group mean grooming progressions for *melanogaster* populations. Grooming proportions were calculated across 10 non-overlapping windows (approximately 3 min. each). Y-axis indicates the proportion of time spent in an action. X-axis indicates the window. As with non-*melanogaster* species, all populations exhibited a general anterior to posterior grooming progression, with increased amounts of standing later in the progression. Canton-S exhibited a higher to walk than other stock lines, allowing for accurate classification when using grooming progressions as a feature.

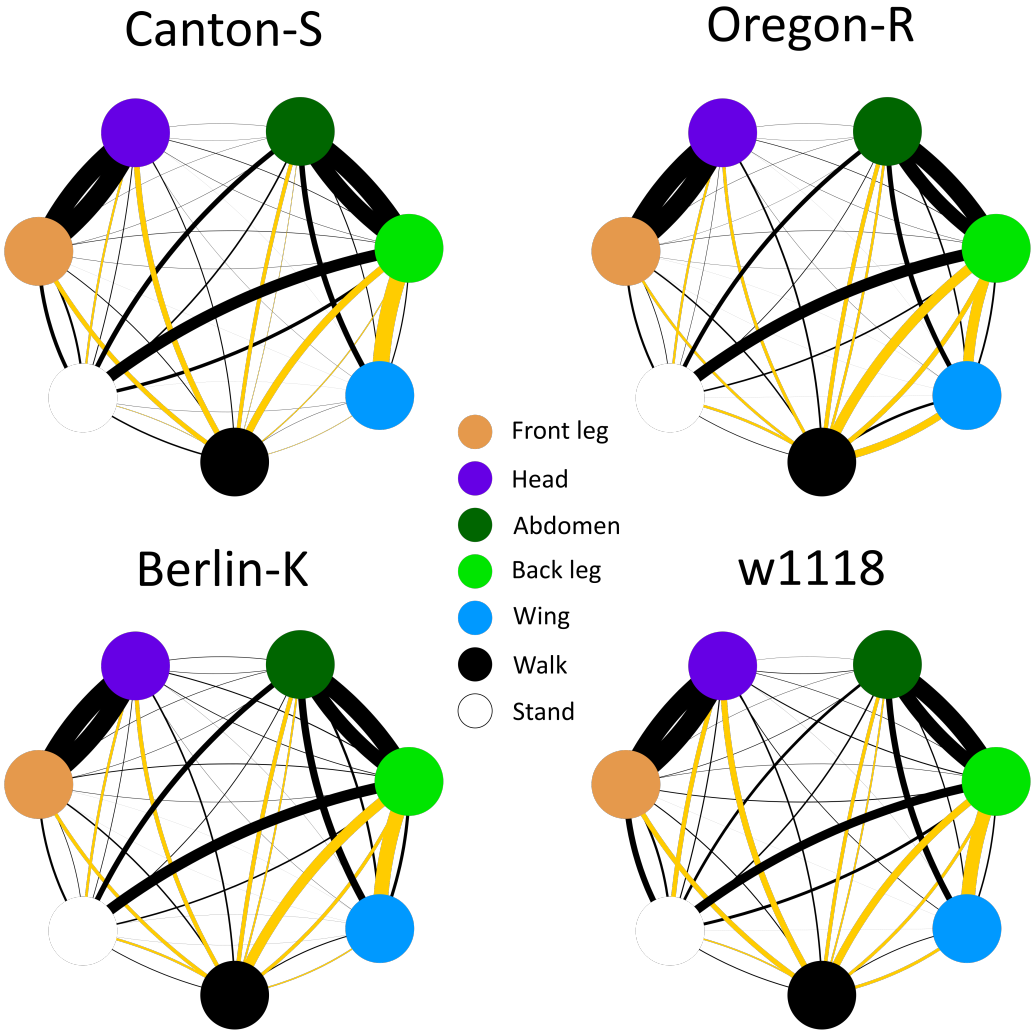


Figure C.9: **Grooming syntax of *melanogaster* stocks is highly conserved.** Shown are graphs of the group mean grooming syntax for each *melanogaster* stock. Nodes denote actions and edge thicknesses indicate transition probabilities between actions. Highlighted in red are transitions within motifs (i.e. head-front leg or abdomen-back leg-wing) that differ significantly (Wilcoxon rank-sum test, $p < .05$ after Bonferroni correction) between stocks. Of 48 within-motif transition comparisons, only 2 possessed statistically different group means, indicating that syntax is highly conserved within *melanogaster* stock lines. In addition, none of these transition distributions possessed statistically distinguishable variances, demonstrating that variability within stocks is also similar.




Figure C.10: **Drosophilids differ in several grooming features.** Shown here are heatmaps of the classification performances of pairwise logistic regression classifiers for between-species comparisons of four grooming features. Color indicates classification accuracy, with 0.5 indicating chance performance. Pairwise between-species classification is generally highly accurate, as syntax and progressions separate *melanogaster* from all non-*melanogaster* species with $> 90\%$ accuracy.




Figure C.11: **Some *melanogaster* stocks differ in syntax and grooming progressions.** Shown here are heatmaps of the classification performances of pairwise logistic regression classifiers for within-*melanogaster* comparisons of four grooming features. Color indicates classification accuracy, with 0.5 indicating chance performance. Pairwise between-species classification performs above chance for many syntax and progressions comparisons, indicating that several stocks exhibit distinct grooming behavior. Closer inspection of syntax and progression elements that differ indicates that the amount of non-grooming actions (walking and standing) differentiates Canton-S and Berlin-K from other stock lines.

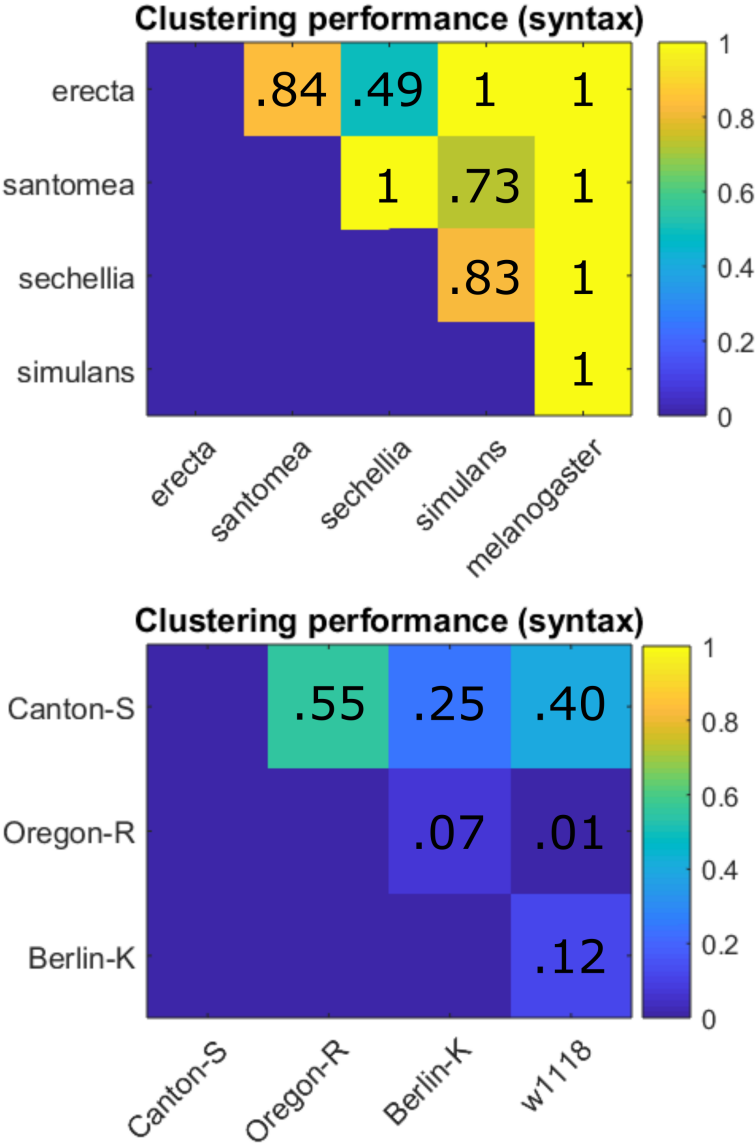


Figure C.12: **K-medoids clustering also separates populations by grooming syntax.** Shown here are heatmaps of the classification performances of k-medoids clustering on between-species and within-species syntax comparisons. Normalized mutual information (NMI) is used to evaluate the accuracy of results from k-medoids clustering using 2 clusters and the known labels, with color indicating the value of NMI. This metric ranges from 0 to 1, with 1 indicating perfect classification. Pairwise between-species classification (top) is generally highly accurate, while within-species classification (bottom) is less reliable.




Figure C.13: **Canton-S individuals do not possess distinct grooming syntax.** 15 Canton-S flies were observed longitudinally three consecutive days (three recordings). Syntax for each fly was calculated for each session and pairwise Euclidean distances were calculated between syntax vectors. Shown here are violin plots of the distances between syntax vectors for inter-individual comparisons (left) and intra-individual comparisons (right). The 15 most similar inter-individual comparisons are more similar than all of the intra-individual comparisons, indicating that flies do not possess strongly idiosyncratic grooming syntax between individuals.




Figure C.14: **Optogenetically stimulated flies exhibit similar variability to dusted flies.** Shown are grooming proportions and transition probabilities of anterior motif transitions for undusted CS flies, undusted CS flies stimulated by optogenetic manipulation, and dusted CS flies. Optogenetic stimulation results in high levels of anterior grooming but does not abolish grooming variability.

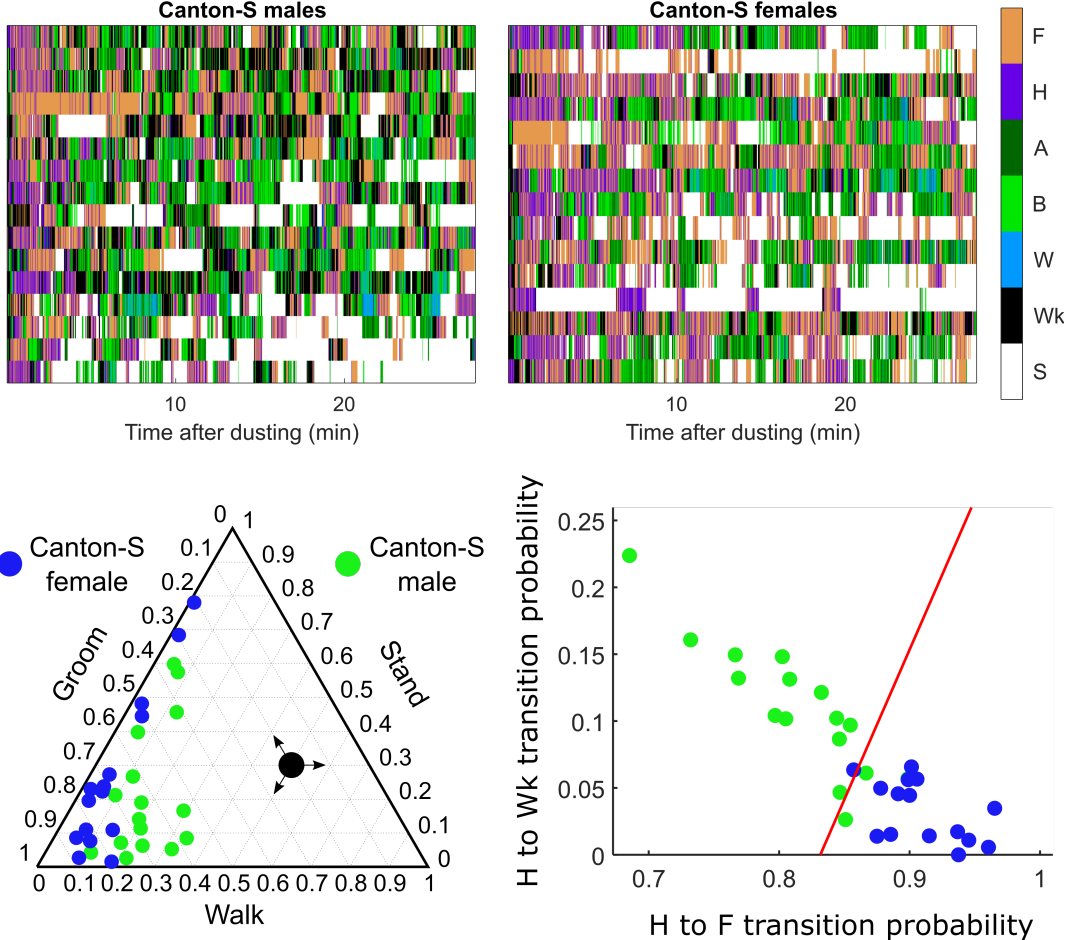


Figure C.15: **Male and female Canton-S flies possess distinct grooming features.** Shown are the ternary plots of the proportion of time spent grooming, walking, and standing for male and female flies. Female flies groom more than male flies, while male flies walk more than female flies. LDA classifies male and female flies with high accuracy using only two syntax elements (H to Wk and H to F transitions). Using only two predictors, LDA classifies male and female flies with an error rate < 10%. Similar error rates are achieved for progressions and grooming proportions, indicating that male and female flies possess identifiable differences in grooming behavior, largely driven by differences in locomotor behavior.

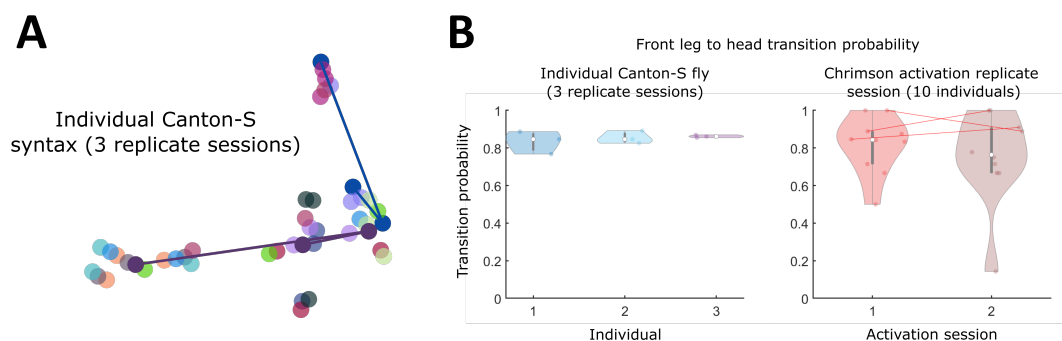


Figure C.16: **Within-individual grooming variability across multiple sessions is comparable to between-individual variation.** A. Shown here is a t-SNE embedding of syntax from 15 Canton-S flies recorded three times on consecutive days. Two example individuals are highlighted, with colored lines connecting their syntax vectors from each session. The high degree of overlap between points and lack of aggregation by color indicates that individual flies do not exhibit strong individuality in their grooming phenotype. This was confirmed quantitatively, as classification analysis performs poorly at discriminating between flies (Appendix D). B. Upon repeated trials, even highly stereotyped anterior transitions exhibit within-individual variability. On the left are front leg to head transition probabilities for three flies, each recorded three times. This anterior motif transition is consistently strong, but intra-individual variability is on the same scale as between-individual variation, indicating that these flies were not strongly idiosyncratic groomers but rather distributed around a group mean. On the right is the same transition probability for ten flies from two sessions of sensory stimulation via optogenetic manipulation (see Methods and Appendix D for more detail). Even when sensory input is carefully controlled, flies exhibit both between-individual variation and within-individual variability, indicating that several factors may account for stochasticity in grooming behavior.

Appendix D

This appendix contains additional figures to support the work presented in Chapter 5.

D.1 Supporting Figures




Figure D.1: **Resting-state networks are modular over a wide range of community detection parameter choices.** Shown is a heat map of the mean modularity quality value (referred to as Q in the main Methods text) of 150 runs of the dynamic community detection algorithm for each parameter combination considered here. Color indicates the value of Q , which ranges from -1 to 1, with a value of 0 indicating a lack of modular structure. While no cutoff exists to identify “true” modular structure per se, higher non-zero positive values of Q indicate more pronounced modularity. Across the parameters considered here, the value of Q ranges from approximately 0.11 to 0.15, indicating weak, but present, modular structure. Consensus partitions are similar across these parameter choices and the networks analyzed here were weighted, undirected, and dense; these factors explain the consistency of Q values.

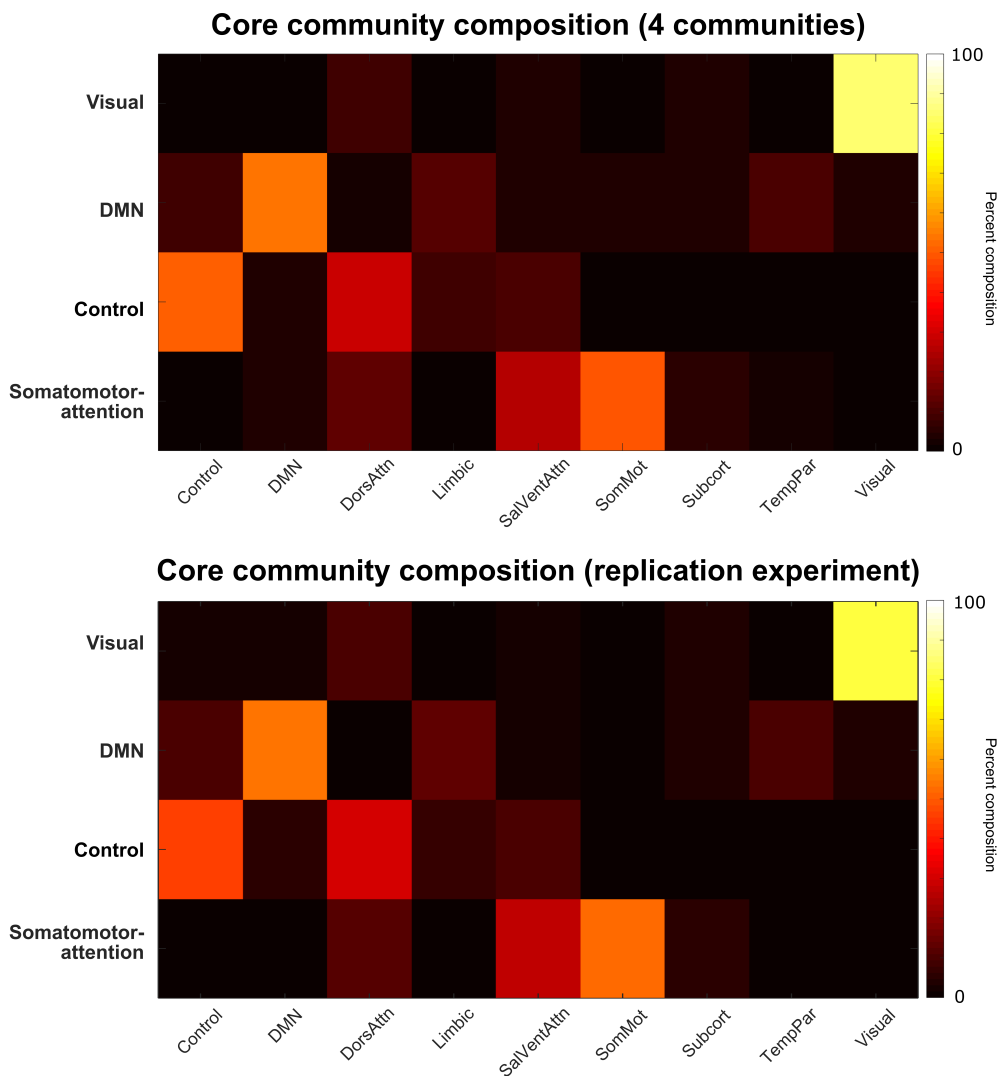


Figure D.2: **Core community compositions are consistent across dynamic community detection parameter choices.** Shown are heat maps of the core community compositions for the original experiment (top) and hormone suppression replication experiment (bottom). In each case, community compositions are averaged over the set of parameter combinations which result in four communities being identified ($.8 < \omega < 2$, $.995 < \gamma < 1.01$). The compositions shown here are nearly identical to that shown in the main text results, which is the core community composition identified at the standard parameter combination for the dynamic community detection algorithm ($\omega = \gamma = 1$).

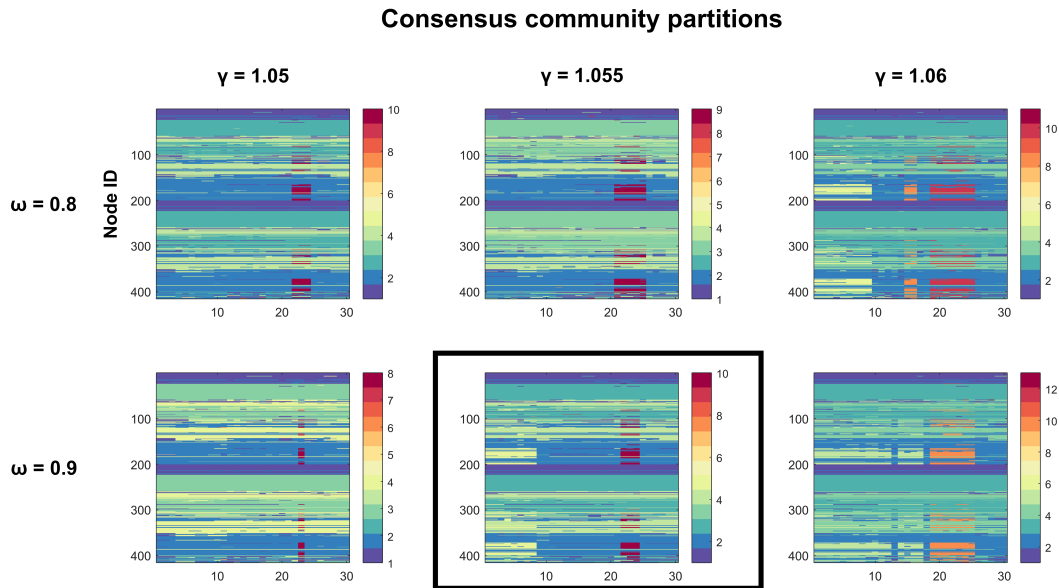


Figure D.3: **Functional communities exhibit transient reorganization around the ovulatory window over a range of dynamic community detection parameter values.** Shown are the consensus community partitions for a range of dynamic community detection algorithm parameters ($.8 \leq \omega \leq .9$, $1.05 \leq \gamma \leq 1.06$). Y-axis values indicate node identity, x-axis values indicate the day of experiment, and color indicates community membership. The partition outlined in black (bottom middle) is the basis of the quantitative analysis in the main text. Ovulation occurs on day 23 (black underline in bottom middle panel). In each case displayed here, a new subcommunity with a consistent composition splits from the default mode core during the ovulatory window but rejoins after day 25 at the latest, concurrent with a sharp decline in estradiol, LH, and FSH. This indicates the presence of a stable, reliable solution to the dynamic community detection algorithm in this parameter range.




Figure D.4: **Solutions to the dynamic community detection algorithm are similar and reliable.** Since the Louvain algorithm is non-deterministic, it is important to characterize the stability/similarity of community detection solutions. To compare the consensus partition used for analysis in the main text to results from each run of the dynamic community detection algorithm, we used normalized mutual information (NMI) as a measure of partition similarity and normalized variation of information (NVI) as a measure of partition distance, per [179]. NMI ranges from 0 to 1, with a value of 0 indicating similarity between partitions at no better than chance and 1 indicating a perfect match. NVI ranges from 0 to 1 as well, but with an inverse interpretation (i.e. smaller distances indicate higher similarity). Pairwise NMI and NVI values were calculated between the consensus partition and each output of the 150 runs of the dynamic community detection algorithm. Pairwise NMI and NVI values were also calculated between 100 repetitions of a null model consensus partition, which permuted node identities, and each DCD output. This null model permutation preserves the flexibility trends and community sizes across the experiment, but shuffles node identity to provide a baseline hypothesis against which to test the significance of the true community detection results. The NMI and NVI distributions are significantly higher and lower, respectively, than the null model comparison, indicating that the consensus partition used in subsequent analysis is much better than what would be arrived at by chance. The median NMI between the consensus partition and each DCD algorithm output is 0.72 and the median NVI is 0.09, indicating that the consensus partition is highly representative of each individual output partition from the 150 runs. Therefore, the community detection solution is highly stable and solutions are highly similar.




Figure D.5: **Network flexibility is significantly correlated with estradiol.** Shown on the left is a heat map of the Spearman correlation values between the flexibility and hormone curves shown in the main text. Color indicates correlation values and red outlines indicate statistical significance at the $p < .05$ level after Bonferroni correction. All networks except the control, dorsal attention, and visual networks have correlation values > 0.6 with estradiol, indicating that these variables are tightly coupled. The visual network also had a statistically significant negative correlation with progesterone, but this network was minimally flexible so the effect size was negligible. No other hormone-flexibility relationships were significant, suggesting that estradiol is the primary driver of flexibility. On the right, a heat map illustrates the lack of significant relationships between hormones and mood variables, which were surveyed each day of the experiment. Only depression correlated significantly with luteinizing hormone.

Node Label	Network	Subnetwork	xMNI	yMNI	zMNI
17Networks_LH_DefaultB_Temp_2	DMN	DMN B	-54	-2	-30
17Networks_LH_DefaultB_Temp_3	DMN	DMN B	-62	-18	-20
17Networks_LH_DefaultB_Temp_4	DMN	DMN B	-56	-8	-14
17Networks_LH_DefaultB_Temp_5	DMN	DMN B	-60	-34	-4
17Networks_LH_DefaultB_Temp_6	DMN	DMN B	-52	-22	-6
17Networks_LH_DefaultB_IPL_1	DMN	DMN B	-46	-58	20
17Networks_LH_DefaultB_IPL_2	DMN	DMN B	-56	-54	30
17Networks_LH_DefaultB_PFCd_1	DMN	DMN B	-4	52	28
17Networks_LH_DefaultB_PFCd_2	DMN	DMN B	-14	58	30
17Networks_LH_DefaultB_PFCd_3	DMN	DMN B	-22	50	32
17Networks_LH_DefaultB_PFCd_4	DMN	DMN B	-8	42	52
17Networks_LH_DefaultB_PFCd_5	DMN	DMN B	-12	24	60
17Networks_LH_DefaultB_PFCd_6	DMN	DMN B	-6	10	64
17Networks_LH_DefaultB_PFCi_1	DMN	DMN B	-40	20	48
17Networks_LH_DefaultB_PFCi_2	DMN	DMN B	-42	8	48
17Networks_LH_DefaultB_PFCv_1	DMN	DMN B	-36	22	-16
17Networks_LH_DefaultB_PFCv_2	DMN	DMN B	-36	36	-12
17Networks_LH_DefaultB_PFCv_3	DMN	DMN B	-46	32	-10
17Networks_LH_DefaultB_PFCv_4	DMN	DMN B	-48	28	0
17Networks_LH_DefaultB_PFCv_5	DMN	DMN B	-54	20	12
17Networks_RH_DefaultB_Temp_1	DMN	DMN B	64	-24	-8
17Networks_RH_DefaultB_Temp_2	DMN	DMN B	64	-38	0
17Networks_RH_DefaultB_AntTemp_1	DMN	DMN B	50	8	-32
17Networks_RH_DefaultB_PFCd_1	DMN	DMN B	6	58	30
17Networks_RH_DefaultB_PFCd_2	DMN	DMN B	16	52	36
17Networks_RH_DefaultB_PFCd_3	DMN	DMN B	4	44	40
17Networks_RH_DefaultB_PFCd_4	DMN	DMN B	14	38	52
17Networks_RH_DefaultB_PFCd_5	DMN	DMN B	12	20	62
17Networks_RH_DefaultB_PFCv_1	DMN	DMN B	34	22	-18
17Networks_RH_DefaultB_PFCv_2	DMN	DMN B	48	32	-8
17Networks_RH_DefaultB_PFCv_3	DMN	DMN B	54	24	6
17Networks_LH_ContC_pCun_1	Control	Control C	-10	-70	32
17Networks_LH_ContC_pCun_2	Control	Control C	-10	-78	46
17Networks_LH_ContC_pCun_3	Control	Control C	-4	-64	52
17Networks_LH_ContC_Cingp_1	Control	Control C	-6	-40	24
17Networks_LH_ContC_Cingp_2	Control	Control C	-4	-22	30
17Networks_RH_ContC_pCun_1	Control	Control C	16	-64	28
17Networks_RH_ContC_pCun_2	Control	Control C	14	-72	40
17Networks_RH_ContC_pCun_3	Control	Control C	6	-64	44
17Networks_RH_ContC_pCun_4	Control	Control C	8	-50	44
17Networks_RH_ContC_pCun_5	Control	Control C	8	-72	52
17Networks_RH_ContC_Cingp_1	Control	Control C	8	-44	20
17Networks_RH_ContC_Cingp_2	Control	Control C	6	-28	28

Figure D.6: **Node identities from the Schaefer functional-anatomical atlas.** Shown here are the node identities for regions in the Control C and DMN B subnetworks. Control C nodes are identified as being highly flexible over the entire course of the menstrual cycle. Within-network connectivity between nodes in DMN B increases around the ovulatory window, resulting in a transient bifurcation of the default mode core community.




Figure D.7: **Functional-anatomic subnetworks have distinct flexibility profiles.** Within functional-anatomical networks, nodes belonging to different subnetworks exhibit different flexibility trends. Specifically, Control C subnetwork nodes are the most likely to be highly flexible within the Control network, suggesting a specific “integrator” role for this subnetwork.




Figure D.8: **Strong changes in functional connectivity during the ovulatory window are localized to a default mode subnetwork.** Shown here are the differences in edge weights (magnitude-squared coherence values) between nodes on days 22 and 23, with color indicating value. Positive values indicate an increase in edge weight from day 22 to 23. On the right, only the top 5% of changes by magnitude are displayed. As demonstrated by the extent of edge weights present in the rows and columns labeled DMN B, it can be seen that within-subnetwork changes in edge strength are responsible for the observed community reorganization during the ovulatory window.

Bibliography

- [1] M. Layeghifard, D. Hwang, and D. Guttman, *Disentangling interactions in the microbiome: A network perspective*, *Trends in Microbiology* **25** (2017), no. 3 217–228.
- [2] S. Rosenthal, C. Twomey, A. Hartnett, H. Wu, and I. Couzin, *Revealing the hidden networks of interaction in mobile animal groups allows prediction of complex behavioral contagion*, *PNAS* **112** (2015), no. 15 4690–4695.
- [3] R. Durrett, J. Gleeson, A. Lloyd, P. Mucha, F. Shi, D. Sivakof, J. Socolar, and C. Varghese, *Graph fission in an evolving voter model*, *PNAS* **109** (2012), no. 10 3682–3687.
- [4] M. Newman, *Networks: an introduction*. Oxford University Press, 2010.
- [5] O. Sporns and D. Bassett, *Editorial: New trends in connectomics*, *Network Neuroscience* **2** (2018), no. 25 125–127.
- [6] M. Rubinov and O. Sporns, *Complex network measures of brain connectivity: Uses and interpretations*, *NeuroImage* **52** (2010) 1059–1069.
- [7] S. Muldoon, E. Bridgeford, and D. Bassett, *Small-world propensity and weighted brain networks*, *Sci. Rep.* **6** (2016) 22057.
- [8] P. Hagmann, L. Jonasson, P. Maeder, J. Thiran, V. Wedeen, and R. Meuli, *Understanding diffusion mr imaging techniques: From scalar diffusion-weighted imaging to diffusion tensor imaging and beyond*, *RadioGraphics* **26** (2006) S205–S223.
- [9] S. Davis, J. Kragel, D. Madden, and R. Cabeza, *The architecture of cross-hemispheric communication in the aging brain: Linking behavior to functional and structural connectivity*, *Cerebral Cortex* **22** (2012), no. 1 232–242.
- [10] A. Fornito and E. Bullmore, *Reconciling abnormalities of brain network structure and function in schizophrenia*, *Curr. Opin. Neurobiology* **30** (2015) 44–50.
- [11] J. Medaglia, M. Lyndall, and D. Bassett, *Cognitive network neuroscience*, *J. Cogn. Neurosci.* **27** (2015), no. 8 1471–1491.

- [12] A. Hermundstad, D. Bassett, K. Brown, E. Aminoff, D. Clewett, S. Freeman, A. Frithsen, A. Johnson, C. Tipper, M. Miller, S. Grafton, and J. Carlson, *Structural foundations of resting-state and task-based functional connectivity in the human brain*, *PNAS* **110** (2013), no. 15 6169–6174.
- [13] E. Towlson, P. Vértes, S. Ahnert, W. Schafer, and E. Bullmore, *The rich club of the C. elegans neuronal connectome*, *J. Neurosci.* **33** (2013), no. 15 6380–6387.
- [14] G. Yan, P. Vértes, E. Towlson, Y. Chew, D. Walker, W. Schafer, and A. Barabási, *Network control principles predict neuron function in the caenorhabditis elegans connectome*, *Nature* **550** (2017) 519–523.
- [15] R. Liégeois, T. Laumann, A. Snyder, J. Zhou, and B. Yeo, *Interpreting temporal fluctuations in resting-state functional connectivity mri*, *NeuroImage* **163** (2017) 437–455.
- [16] G. Berman, W. Bialek, and J. Shaevitz, *Predictability and hierarchy in drosophila behavior*, *PNAS* **113** (2016), no. 42 11943–11948.
- [17] J. Ayroles, S. Buchanan, C. O’Leary, K. Skutt-Kakaria, J. Grenier, A. Clark, D. Hartl, and B. de Bivort, *Behavioral idiosyncrasy reveals genetic control of phenotypic variability*, *PNAS* **112** (2015), no. 21 6706–6711.
- [18] M. Gandolfo, A. Maccione, M. Tedesco, S. Martinoia, and L. Berdondini, *Tracking burst patterns in hippocampal cultures with high-density cmos-meas*, *J. Neur. Eng.* **7** (2010), no. 5 doi:056001.
- [19] D. Wagenaar, J. Pine, and S. Potter, *An extremely rich repertoire of bursting patterns during the development of cortical cultures*, *BMC Neuroscience* **7** (2006), no. 11 doi:10.1186/1471-2202-7-11.
- [20] A. Szébenyi, *Cleaning behaviour in drosophila melanogaster*, *Anim. Behav.* **17** (1969) 641–651.
- [21] A. Seeds, P. Ravbar, P. Chung, S. Hampel, F. Midgley Jr., B. Mensh, and J. Simpson, *A suppression hierarchy among competing motor programs drives sequential grooming in drosophila*, *eLife* **3** (2014) e02951. doi: 10.7554/eLife.02951.
- [22] L. Millett and M. Gillette, *Over a century of neuron culture: From the hanging drop to microfluidic devices*, *Yale J. Biol. and Med.* **85** (2012) 501–521.
- [23] S. Marom and G. Shahaf, *Development, learning and memory in large random networks of cortical neurons: lessons beyond anatomy*, *Q. Rev. Biophys.* **35** (2002), no. 1 63–87.
- [24] L. Squire, *Memory and the hippocampus: A synthesis from findings with rats, monkeys, and humans*, *Psych. Rev.* **99** (1992), no. 2 195–231.

- [25] R. Morris, P. Garrud, J. Rawlins, and J. O’Keefe, *Place navigation impaired in rats with hippocampal lesions*, *Nature* **297** (1982) 681–683.
- [26] A. Treves and E. Rolls, *Computational analysis of the role of the hippocampus in memory*, *Hippocampus* **4** (1994), no. 3 374–391.
- [27] G. Buszáki, *Hippocampal sharp waves: Their origin and significance*, *Brain Res.* **398** (1986), no. 2 242–252.
- [28] G. Buszáki, Z. Horváth, R. Urioste, J. Hetke, and K. Wise, *High-frequency network oscillation in the hippocampus*, *Science* **256** (1992) 1025–1027.
- [29] A. Malvache, S. Reichinnek, V. Villette, C. Haimerl, and R. Cossart, *Awake hippocampal reactivations project onto orthogonal neuronal assemblies*, *Psych. Rev.* **99** (1992), no. 2 195–231.
- [30] J. Butler and O. Paulsen, *Hippocampal network oscillations - recent insights from in vitro experiments*, *Curr. Opin. Neurobiol.* **31** (2015) 40–44.
- [31] A. Draguhn, R. Traub, D. Schmitz, and J. Jefferys, *Electrical coupling underlies high-frequency oscillations in the hippocampus in vitro*, *Nature* **394** (1998) 189–192.
- [32] M. Churchland, J. Cunningham, M. Kaufman, J. Foster, P. Nuyujukian, S. Ryu, and K. Shenoy, *Neural population dynamics during reaching*, *Nature* **487** (2012) 51–56.
- [33] M. Goard, G. Pho, J. Woodson, and M. Sur, *Distinct roles of visual, parietal, and frontal motor cortices in memory-guided sensorimotor decisions*, *eLife* **5** (2016) e13764.
- [34] K. Shenoy, M. Sahani, and M. Churchland, *Cortical control of arm movements: A dynamical systems perspective*, *Ann. Rev. of Neuro.* **36** (2013) 337–359.
- [35] J. Liu and D. Buonomano, *Embedding multiple trajectories in simulated recurrent neural networks in a self-organizing manner*, *J. Neuro.* **29** (2009), no. 42 13172–13181.
- [36] E. Izhikevich, *Polychronization: Computation with spikes*, *Neural Comp.* **18** (2006) 245–282.
- [37] N. Timme, N. Marshall, N. Bennett, M. Ripp, E. Lautzenhiser, and J. Beggs, *Criticality maximizes complexity in neural tissue*, *Frontiers Physiol.* **7** (2016) 425.
- [38] J. Wixted, S. Goldinger, L. Squire, J. Kuhn, M. Papesh, K. Smith, D. Treiman, and P. Steinmetz, *Coding of episodic memory in the human hippocampus*, *PNAS* **115** (2018), no. 5 1093–1098.

- [39] M. Beyeler, E. Rounds, K. Carlson, N. Dutt, and J. Krichmar, *Neural correlates of sparse coding and dimensionality reduction*, *PLoS Comput. Biol.* **15** (2019), no. 6 e1006908.
- [40] N. Raichman and E. Ben-Jacob, *Identifying repeated motifs in the activation of synchronized bursts in cultured neuronal networks*, *J. Neuro. Meth.* **170** (2008), no. 1 96–110.
- [41] R. Quiroga, Z. Nadasdy, and Y. Ben-Shaul, *Unsupervised spike detection and sorting with wavelets and superparamagnetic clustering*, *Neur. Comp.* **16** (2004) 1661–1687.
- [42] M. Schroeter, P. Charlesworth, M. Kitzbichler, O. Paulsen, and E. Bullmore, *Emergence of rich-club topology and coordinated dynamics in development of hippocampal functional networks in vitro*, *J. Neuroscience* **35** (2015), no. 5459-5470.
- [43] J. Cunningham, V. Gilja, S. Ryu, and K. Shenoy, *Methods for estimating neural firing rates, and their application to brain-machine interfaces*, *Neural Networks* **22** (2009), no. 9 1235–1246.
- [44] J. Cunningham and B. Yu, *Dimensionality reduction for large-scale neural recordings*, *Nat. Neuro.* **17** (2014), no. 11 1500–1509.
- [45] A. Peyrache, K. Benchenane, M. Khamassi, S. Wiener, and F. Battaglia, *Principal component analysis of ensemble recordings reveals cell assemblies at high temporal resolution*, *J. Comp. Neuro.* **29** (2010), no. 1-2 309–325.
- [46] M. Ester, H. Kriegel, J. Sander, and X. Xu, *A density-based algorithm for discovering clusters in large spatial databases with noise*, *KDD Proceedings* (1996) 226–231.
- [47] P. Sjöström, E. Rancz, A. Roth, and M. Häusser, *Dendritic excitability and synaptic plasticity*, *Physiol. Rev.* **88** (2008) 769–840.
- [48] Z. Brzosko, S. Zannone, W. Schultz, C. Clopath, and O. Paulsen, *Sequential neuromodulation of hebbian plasticity offers mechanism for effective reward-based navigation*, *eLife* **6** (2017) e27756.
- [49] D. Buonomano, *Harnessing chaos in recurrent neural networks*, *Neuron* **63** (2009) 423–425.
- [50] D. Poli, V. Pastore, and P. Massobrio, *Functional connectivity in in vitro neuronal assemblies*, *Frontiers Neur. Circ.* **9** (2015), no. 57 doi:10.3389/fncir.2015.00057.
- [51] J. Sun, W. Kilb, and H. Luhmann, *Self-organization of repetitive spike patterns in developing neuronal networks in vitro*, *Euro. J. Neuro.* **32** (2010) 1289–1299.

- [52] D. Wagenaar, Z. Nadasdy, and S. Potter, *Persistent dynamic attractors in activity patterns of cultured neuronal networks*, *Phys. Rev. E. Stat. Nonlin. Soft Matter Phys.* **73** (2006), no. 5.1 051907.
- [53] J. Rolston, D. Wagenaar, and S. Potter, *Precisely timed spatiotemporal patterns of neural activity in dissociated cortical cultures*, *Neuroscience* **148** (2007), no. 1 294–303.
- [54] R. Madhavan, Z. Chao, and S. Potter, *Plasticity of recurring spatiotemporal activity patterns in cortical networks*, *Phys. Biol.* **4** (2008), no. 3 381–193.
- [55] T. Isomura, K. Kotani, and Y. Jimbo, *Cultured cortical neurons can perform blind source separation according to the free-energy principle*, *PLoS Comp. Biol.* **11** (2015), no. 12 e1004643.
- [56] S. Jahnke, M. Timme, and R. Memmesheimer, *A unified dynamic model for learning, replay, and sharp-wave/ripples*, *Journal of Neuroscience* **35** (2015), no. 49 16236–16258.
- [57] V. Cutsuridis, S. Cobb, and B. Graham, *Encoding and retrieval in a model of the hippocampal ca1 microcircuit*, *Hippocampus* **20** (2009), no. 3 423–446.
- [58] C. Rennó-Costa, J. Lisman, and P. Verschure, *A signature of attractor dynamics in the ca3 region of the hippocampus*, *PLoS Comput. Biol.* **10** (2014), no. 5 e1003641.
- [59] D. Sussillo and L. Abbott, *Generating coherent patterns of activity from chaotic neural networks*, *Neuron* **63** (2009), no. 4 544–557.
- [60] H. Johnson, A. Goel, and D. Buonomano, *Neural dynamics of in vitro cortical networks reflects experienced temporal patterns*, *Nat. Neuro.* **13** (2010), no. 8 917–919.
- [61] J. Gallego, M. Perich, L. Miller, and S. Solla, *Neural manifolds for the control of movement*, *Neuron* **94** (2017) 978–984.
- [62] C. Hu, M. Petersen, N. Hoyer, B. Spitzweck, F. Tenedini, D. Wang, A. Gruschka, L. Burchardt, E. Szpotowicz, M. Schweizer, A. Guntur, C. Yang, and P. Soba, *Sensory integration and neuromodulatory feedback facilitate drosophila mechanonociceptive behavior*, *Nat. Neurosci.* **20** (2017), no. 8 1085–1095.
- [63] E. Marder and D. Bucher, *Understanding circuit dynamics using the stomatogastric nervous system of lobsters and crabs*, *Annu. Rev. Physiol.* **69** (2007) 291–316.

- [64] A. Selverston, R. Elson, M. Rabinovich, R. Huerta, and H. Abarbanel, *Basic principles for generating motor output in the stomatogastric ganglion*, *Ann. NY Acad. Sci.* **16** (1998), no. 860 35–50.
- [65] B. Spruijt, J. van Hooff, and W. Gispen, *Ethology and neurobiology of grooming behavior*, *Physiological Rev.* **72** (1992), no. 3 825–852.
- [66] M. Zhukovskaya, A. Yanagawa, and B. Forschler, *Grooming behavior as a mechanism of insect disease defense*, *Insects* **4** (2013) 609–630.
- [67] R. Dawkins and M. Dawkins, *Hierarchical organization and postural facilitation: rules for grooming in flies*, *Anim. Behav.* **24** (1976) 739–755.
- [68] J. Fentress and F. Stilwell, *Grammar of a movement sequence in inbred mice*, *Nature* **244** (1973) 52–53.
- [69] A. Szebenyi, *Cleaning behaviour in drosophila melanogaster*, *Anim. Behav.* **17** (1969) 641–651.
- [70] A. Seeds, P. Ravbar, P. Chung, S. Hampel, F. Midgley Jr., B. Mensh, and J. Simpson, *A suppression hierarchy among competing motor programs drives sequential grooming in drosophila*, *eLife* **3** (2014) e02951. doi: 10.7554/eLife.02951.
- [71] G. Berman, W. Bialek, and J. Shaevits, *Predictability and hierarchy in drosophila behavior*, *PNAS* **113** (2016), no. 42 11943–11948.
- [72] H. Raad, J. Ferveur, N. Ledger, M. Capovilla, and A. Robichon, *Functional gustatory role of chemoreceptors in drosophila wings*, *Cell Rep.* **15** (2016), no. 7 1442–1454.
- [73] A. Yanagawa, A. Guigue, and F. Marion-Poll, *Hygienic grooming is induced by contact chemicals in drosophila melanogaster*, *Front. Beh. Neurosci.* **8** (2014) 254.
- [74] E. Agarwala, H. Chiel, and P. Thomas, *Pursuit of food versus pursuit of information in a markovian perception-action loop model of foraging*, *J. Theor. Biol.* **304** (2012) 235–272.
- [75] S. Lee and S. Kang, *Characterization of the crawling activity of caenorhabditis elegans using a hidden markov model*, *Theory Biosci.* **134** (2015), no. 3-4 117–125.
- [76] S. Sim, S. Kang, and S. Lee, *Using hidden markov models to characterize termite traveling behavior in tunnels with different curvatures*, *Behav. Processes* **111** (2015) 101–108.
- [77] V. Cane, *Behaviour sequences as semi-markov chains*, *J. Royal. Stat. Society Series B (Methodological)* **21** (1959), no. 1 36–58.

- [78] T. Markow, *Behavioral and sensory basis of courtship success in drosophila melanogaster*, *PNAS* **84** (1987) 6200–6204.
- [79] A. Katsov, L. Freifeld, M. Horowitz, S. Kuehn, and T. Clandinin, *Dynamic structure of locomotor behavior in walking fruit flies*, *eLife* **6** (2017) e24610. doi:10.7554/eLife.26410.
- [80] L. Tao, S. Ozarkar, J. Beck, and V. Bhandawat, *Statistical structure of locomotion and its modulation by odors*, *eLife* **8** (2019) e41235.
- [81] G. Berman, D. Choi, W. Bialek, and J. Shaevits, *Mapping the stereotyped behavior of freely moving flies*, *J. R. Soc. Interface* **11** (2014) 20140672.
- [82] A. Kershenbaum, A. Bowles, T. Freeberg, D. Jin, A. Lameira, and K. Bohn, *Animal vocal sequences: not the markov chains we thought they were*, *Proc. R. Soc. B* **281** (2014) 20141370.
- [83] D. Jin, *Generating variable birdsong syllable sequences with branching chain networks in avian premotor nucleus hvc*, *Phys. Rev. E* **80** (2009) 051902.
- [84] K. Katahira, K. Suzuki, K. Okanoya, and M. Okada, *Complex sequencing rules of birdsong can be explained by simple hidden markov processes*, *PLoS One* **6** (2011), no. 9 e24516.
- [85] P. Ravbar, K. Branson, and J. Simpson, *An automatic behavior recognition system classifies animal behaviors using movements and their temporal context*, *J. Neuro. Methods* **326** (2019), no. 1 108352.
- [86] M. van Ginke, C. Luengo Hendriks, and L. van Vliet, *A short introduction to the radon and hough transforms and how they relate to each other*, *Delft University of Technology* (2004).
- [87] F. Barradale, K. Sinha, and T. Lebestky, *Quantification of drosophila grooming behavior*, *J. Vis. Exp.* **125** (2017) e55231, doi:10.3791/55231.
- [88] W. Heiligenberg, *Random processes describing the occurrence of behavioural patterns in a cichlid fish*, *Anim. Behav.* **21** (1973) 169–182.
- [89] J. Niven and S. Laughlin, *Energy limitation as a selective pressure on the evolution of sensory systems*, *J. Exp. Biol.* **211** (2008) 1792–1804.
- [90] W. Kristan, *Behavioral sequencing: Competitive queuing in the fly cns*, *Curr. Biol.* **24** (2014) R743–746.
- [91] M. Long, D. Jin, and M. Fee, *Support for a synaptic chain model of neuronal sequence generation*, *Nature* **468** (2010), no. 7322 394–399.

- [92] M. Suster and M. Bate, *Embryonic assembly of a central pattern generator without sensory input*, *Nature* **416** (2002), no. 6877 174–178.
- [93] A. von Philipsborn, T. Liu, J. Yu, C. Masser, S. Bidaye, and B. Dickson, *Neuronal control of drosophila courtship song*, *Neuron* **69** (2011), no. 3 509–522.
- [94] T. Bautista, Q. Sun, and P. Pilowsky, *The generation of pharyngeal phase of swallow and its coordination with breathing: interaction between the swallow and respiratory central pattern generators*, *Prog. Brain. Res.* **212** (2014) 253–275.
- [95] K. Westberg and A. Kolta, *The trigeminal circuits responsible for chewing*, *Int. Rev. Neurobiol.* **97** (2011) 77–98.
- [96] J. Nassour, P. Henaff, F. Benouezdo, and G. Cheng, *Multi-layered multi-pattern cpg for adaptive locomotion of humanoid robots*, *Biol. Cybern.* **108** (2014), no. 3 291–303.
- [97] B. Qiao, C. Li, V. Allen, M. Shirasu-Hiza, and S. Syed, *Automated analysis of long-term grooming behavior in drosophila using a k-nearest neighbors classifier*, *eLife* **7** (2015) e34497.
- [98] W. Johanssen, *The genotype conception of heredity*, *Int. J. Epidemiol.* **43** (2014), no. 4 989–1000.
- [99] B. Baker, B. Taylor, and J. Hall, *Are complex behaviors specified by dedicated regulatory genes? reasoning from drosophila*, *Genetics* **105** (2001), no. 1 13–24.
- [100] A. Bendesky, Y. Kwon, J. Lassance, C. Lewarch, S. Yao, B. Peterson, M. He, C. Dulac, and H. Hoekstra, *The genetic basis of parental care evolution in monogamous mice*, *Nature* **544** (2017), no. 7651 434–439.
- [101] Y. Ding, J. Lillvis, J. Cande, G. Berman, B. Arthur, X. Long, M. Xu, B. Dickson, and D. Stern, *Neural evolution of context-dependent fly song*, *Curr. Biol.* **29** (2019), no. 7 1089–1099.
- [102] D. Anderson, *Circuit modules linking internal states and social behaviour in flies and mice*, *Nature Rev. Neurosci.* **17** (2016) 692–704.
- [103] A. Allen, I. Anreiter, M. Neville, and M. Sokolowski, *Feeding-related traits are affected by dosage of the foraging gene in drosophila melanogaster*, *Genetics* **205** (2017), no. 2 761–773.
- [104] M. Sokolowski, *Drosophila: Genetics meets behavior*, *Nat. Rev. Genetics* **2** (2001) 879–890.
- [105] B. Gaertner, E. Ruedi, L. McCoy, J. Moore, M. Wolfner, and T. Mackay, *Heritable variation in courtship patterns in drosophila melanogaster*, *G3* **5** (2015) 531–539.

- [106] B. Spruijt, J. van Hooff, and W. Gispen, *Ethology and neurobiology of grooming behavior*, *Physiological Rev.* **72** (1992), no. 3 825–852.
- [107] M. Zhukovskaya, A. Yanagawa, and B. Forschler, *Grooming behavior as a mechanism of insect disease defense*, *Insects* **4** (2013) 609–630.
- [108] J. Mueller, P. Ravbar, J. Simpson, and J. Carlson, *Drosophila melanogaster grooming possesses syntax with distinct rules at different temporal scales*, *PLoS Comp. Biol.* **15** (2019), no. 6 e1007105.
- [109] S. Datta, D. Anderson, K. Branson, P. Perona, and A. Leifer, *Computational neuroethology: A call to action*, *Neuron* **104** (2019) 11–24.
- [110] D. Kim, M. Alvarez, L. Lechuga, and M. Louis, *Species-specific modulation of food-search behavior by respiration and chemosensation in drosophila larvae*, *eLife* **6** (2017) e27057.
- [111] N. Zhang, L. Guo, and J. Simpson, *Spatial comparisons of mechanosensory information govern the grooming sequence in drosophila*, *Curr. Biol.* **30** (2020), no. 6 988–1001.
- [112] J. de Belle, A. Hilliker, and M. Sokolowski, *Genetic localization of foraging (for): A major gene for larval behavior in drosophila melanogaster*, *Genetics* **123** (1989), no. 1 157–163.
- [113] R. Phillis, A. Bramlage, C. Wotus, A. Whittaker, L. Gramates, D. Seppala, F. Farahanchi, P. Caruccio, and R. Murphey, *Isolation of mutations affecting neural circuitry required for grooming behavior in drosophila melanogaster*, *Genetics* **133** (1993), no. 3 581–592.
- [114] J. Colom and B. Brembs, *Sub-strains of drosophila canton-s differ markedly in their locomotor behavior*, *F1000Research* **3** (2015) 176.
- [115] B. Voelkl et al., *Reproducibility of animal research in light of biological variation*, *Nat. Rev. Neuro.* **21** (2020) 384–393.
- [116] S. Buchanan, J. Kain, and B. de Bivort, *Neuronal control of locomotor handedness in drosophila*, *PNAS* **112** (2015), no. 21 6700–6705.
- [117] J. Freund, A. Brandmaier, L. Lewejohann, I. Kirste, M. Kritzler, A. Kruger, N. Sachser, U. Lindenberger, and G. Kempermann, *Emergence of individuality in genetically identical mice*, *Science* **340** (2013), no. 6133 756–759.
- [118] S. Klein, C. Pasquaretta, A. Barron, J. Devaud, and M. Lihoreau, *Inter-individual variability in the foraging behaviour of traplining bumblebees*, *Sci. Reports* **7** (2017), no. 4561 DOI:10.1038/s41598-017-04919-8.

- [119] S. Takagi and R. Benton, *Animal behavior: A neural basis of individuality*, *Curr. Biol.* **30** (2020) R698–R720.
- [120] G. Linneweber, M. Andriatsilavo, S. Dutta, M. Bengochea, L. Hellbruegge, G. Liu, R. Ejsmont, A. Straw, M. Wernet, and P. Hiesinger et al., *A neurodevelopmental origin of behavioral individuality in the Drosophila visual system*, *Science* **367** (2020) 1112–1119.
- [121] J. Akhund-Zade, S. Ho, C. O’Leary, and B. de Bivort, *The effect of environmental enrichment on behavioral variability depends on genotype, behavior, and type of enrichment*, *J. Exp. Biol.* **222** (2019) jeb202234. doi:10.1242/jeb.202234–R12.
- [122] K. Honegger and B. de Bivort, *Stochasticity, individuality and behavior*, *Curr. Biol.* **28** (2015), no. 1 R8–R12.
- [123] J. Kain, S. Zhang, J. Akhund-Zade, A. Samuel, M. Klein, and B. de Bivort, *Variability in thermal and phototactic preferences in drosophila may reflect an adaptive bet-hedging strategy*, *Evolution* **69** (2015), no. 12
<https://doi.org/10.1111/evo.12813>.
- [124] S. Hampel, R. Franconville, J. Simpson, and A. Seeds, *A neural command circuit for grooming movement control*, *eLife* **7** (2015), no. 4 e08758.
- [125] D. Bassett and O. Sporns, *Network neuroscience.*, *Nat. Neurosci.* **20** (2017), no. 3 353–364.
- [126] M. Newman, *Modularity and community structure in networks.*, *PNAS* **103** (2006), no. 23 8577–8582.
- [127] L. Galea, K. Frick, E. Hampson, F. Sohrabji, and E. Choleris, *Why estrogens matter for behavior and brain health*, *Neurosci. Biobehav. Rev* **76** (2017) 363–379.
- [128] A. Wang, Y. Hara, W. Janssen, P. Rapp, and J. Morrison, *Synaptic estrogen receptor-alpha levels in prefrontal cortex in female rhesus monkeys and their correlation with cognitive performance.*, *J. Neurosci.* **30** (2010), no. 38 12770–12776.
- [129] K. Frick, *Molecular mechanisms underlying the memory-enhancing effects of estradiol.*, *Hormones and Behavior* **74** (2015) 4–18.
- [130] C. Woolley and B. McEwen, *Roles of estradiol and progesterone in regulation of hippocampal dendritic spine density during the estrous cycle in the rat.*, *J. Comp. Neurology* **336** (1993), no. 2 293–306.
- [131] Y. Hara, E. Waters, B. McEwen, and J. Morrisoh, *Estrogen effects on cognitive and synaptic health over the lifecourse*, *Physiol. Rev.* **95** (2015) 472–493.

- [132] K. Berman, P. Schmidt, D. Rubinow, M. Danaceau, J. Van Horn, G. Esposito, J. Ostrem, and D. Weinberger, *Modulation of cognition-specific cortical activity by gonadal steroids: a positron-emission tomography study in women.*, *PNAS* **94** (1997), no. 16 8836–8841.
- [133] E. Jacobs and M. D’Esposito, *Estrogen shapes dopamine-dependent cognitive processes: implications for women’s health.*, *J. Neurosci.* **13** (2011), no. 14 5286–5293.
- [134] N. Petersen, L. Kilpatrick, A. Goharзад, and L. Cahill, *Oral contraceptive pill use and menstrual cycle phase are associated with altered resting state functional connectivity.*, *NeuroImage* **90** (2014) 24–32.
- [135] S. Weis, S. Hodgetts, and M. Hausmann, *Sex differences and menstrual cycle effects in cognitive and sensory resting state networks*, *Brain and Cognition* **131** (2019) 66–73.
- [136] K. Arélin, K. Mueller, C. Barth, P. Rekkas, J. Kratzsch, I. Burmann, A. Villringer, and J. Sacher, *Progesterone mediates brain functional connectivity changes during the menstrual cycle - a pilot resting state mri study*, *Frontiers in Neuroscience* **9** (2015), no. 44 doi: 10.3389/fnins.2015.00044.
- [137] N. Lisofsky, J. Martensson, A. Eckert, U. Lindenberger, J. Gallinat, and S. Kuhn, *Hippocampal volume and functional connectivity changes during the female menstrual cycle.*, *NeuroImage* **118** (2015) 154–162.
- [138] H. Hjelmervik, M. Hausmann, B. Osnes, R. Westerhausen, and K. Specht, *Resting states are resting traits - an fmri study of sex differences and menstrual cycle effects in resting state cognitive control network*, *PLoS One* **9** (2014), no. 7 e103492.
- [139] T. Laumann et al., *Functional system and areal organization of a highly sampled individual human brain*, *Neuron* **87** (2015), no. 3 657–670.
- [140] R. Poldrack, T. Laumann, and O. Koyejo et al, *Long-term neural and physiological phenotyping of a single human*, *Nat Commun* **6** (2015), no. 8885 doi:10.1038/ncomms9885.
- [141] B. Chen, T. Xu, C. Zhou, L. Wang, N. Yang, and Z. Wang et al., *Individual variability and test-retest reliability revealed by ten repeated resting-state brain scans over one month*, *PLoS ONE* **10** (2015), no. 12 e0144963.
- [142] C. Gratton et al., *Functional brain networks are dominated by stable group and individual factors, not cognitive or daily variation*, *Neuron* **98** (2018), no. 2 439–452.

- [143] L. Pritschet, T. Santander, E. Layher, C. Taylor, S. Yu, M. Miller, S. Grafton, and E. Jacobs, *Functional reorganization of brain networks across the human menstrual cycle*, *bioRxiv* (2019) doi: <https://doi.org/10.1101/866913>.
- [144] A. Schaefer, R. Kong, E. Gordon, T. Laumann, X. Zuo, A. Holmes, S. Eickhoff, and B. Yeo, *Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity mri.*, *Cereb. Cortex* **28** (2018), no. 9 3095–3114.
- [145] M. Bertolero, B. Yeo, and M. D’Esposito, *The modular and integrative functional architecture of the human brain.*, *PNAS* **112** (2015), no. 49 E6798–E6807.
- [146] A. Khambhati, A. Sizemore, R. Betzel, and D. Bassett, *Modeling and interpreting mesoscale network dynamics.*, *NeuroImage* **180** (2018) 337–349.
- [147] S. Gu, F. Pasqualetti, M. Cieslak, Q. Telesford, A. Yu, A. Kahn, J. Medaglia, J. Vettel, M. Miller, S. Grafton, and D. Bassett, *Controllability of structural brain networks.*, *Nature Comm.* **6** (2015) 8414.
- [148] D. Bassett, M. Porter, N. Wymbs, S. Grafton, J. Carlson, and P. Mucha, *Robust detection of dynamic community structure in networks*, *Chaos* **23** (2013) 01342.
- [149] M. Jenkinson, C. Beckmann, T. Behrens, M. Woolrich, and S. Smith, *Fsl.*, *Neuroimage* **62** (2012), no. 2 782–790.
- [150] C. Horien, X. Shen, D. Scheinost, and R. Constable, *The individual functional connectome is unique and stable over months to years*, *NeuroImage* **189** (2019) 676–687.
- [151] P. Bell and J. Shine, *Subcortical contributions to large-scale network communication.*, *Neuro. & Biobehav. Rev.* **71** (2016) 313–322.
- [152] S. Shanmugan and C. Epperson, *Estrogen and the prefrontal cortex: Towards a new understanding of estrogen’s effects on executive functions in the menopause transition.*, *Human Brain Mapping* **35** (2014), no. 3 847–865.
- [153] C. Galvin and N. I., *Regulation of the mouse medial prefrontal cortical synapses by endogenous estradiol.*, *Neuropsychopharm.* **39** (2014), no. 9 2086–2094.
- [154] Y. Hara, F. Yuk, R. Puri, W. Janssen, P. Rapp, and J. Morrison, *Estrogen restores multisynaptic boutons in the dorsolateral prefrontal cortex while promoting working memory in aged rhesus monkeys.*, *Journal of Neuroscience* **36** (2016), no. 3 901–910.
- [155] E. Jacobs, B. Weiss, N. Makris, S. Whitfield-Gabrieli, S. Buka, A. Klibanski, and J. Goldstein, *Reorganization of functional networks in verbal working memory circuitry in early midlife: The impact of sex and menopausal status.*, *Cerebral Cortex* **25** (2017), no. 5 2857–2870.

- [156] A. Clemens, C. Lenschow, P. Beed, L. Li, R. Sammons, R. Naumann, H. Wang, D. Schmitz, and M. Brecht, *Estrus-cycle regulation of cortical inhibition.*, *Current Biology* **29** (2019), no. 4 605–615.
- [157] J. Hao, P. Rapp, A. Leffler, S. Leffler, W. Janssen, W. Lou, H. McKay, J. Roberts, S. Wearne, P. Hof, and J. Morrison, *Estrogen alters spine number and morphology in prefrontal cortex of aged female rhesus monkeys.*, *J. Neurosci.* **26** (2006), no. 9 2571–2578.
- [158] G. Williams and P. Goldman-Rakic, *Modulation of memory fields by dopamine d1 receptors in prefrontal cortex.*, *Nature* **376** (1995), no. 6541 572–575.
- [159] U. Braun et al., *Dynamic reconfiguration of frontal brain networks during executive cognition in humans*, *PNAS* **112** (2015), no. 37 11678–11683.
- [160] D. Bassett, N. Wymbs, M. Porter, P. Mucha, J. Carlson, and S. Grafton, *Dynamic reconfiguration of human brain networks during learning*, *PNAS* **108** (2011), no. 18 7641–7646.
- [161] R. Gerraty, J. Davidow, K. Foerde, A. Galvan, D. Bassett, and D. Shohamy, *Dynamic flexibility in striatal-cortical circuits supports reinforcement learning.*, *J. Neurosci.* **38** (2018), no. 10 2442–2453.
- [162] M. Mattar, M. Cole, S. Thompson-Schill, and D. Bassett, *A functional cartography of cognitive systems.*, *PLoS Comp. Biol.* **11** (2015), no. 12 e1004533. doi:10.1371/journal.pcbi.1004533.
- [163] M. van den Heuvel and O. Sporns, *Network hubs in the human brain.*, *Trends in Cogn. Sci.* **17** (2013), no. 12 683–696.
- [164] A. Kabbara, W. Falou, M. Khalil, K. Wendling, and M. Hassan, *The dynamic functional core network of the human brain at rest.*, *Sci. Repor.* **7** (2017) 2936.
- [165] J. Shine, P. Bissett, P. Bell, O. Koyejo, J. Balsters, K. Gorgolewski, C. Moodie, and R. Poldrack, *Dynamics of functional brain networks: Integrated network states during cognitive task performance.*, *Neuron* **92** (2016), no. 2 544–554.
- [166] A. Zalesky, A. Fornito, L. Cocchi, L. Gollo, and M. Breakspear, *Time-resolved resting-state brain networks.*, *Proc. Natl. Acad. Sci.* **111** (2014), no. 28 10341–10346.
- [167] E. Hoekzema, E. Barba-Müller, C. Pozzobon, M. Picado, F. Lucco, D. García-García, and A. Ballesteros, *Pregnancy leads to long-lasting changes in human brain structure.*, *Nat. Neuro.* **20** (2017), no. 2 287–296.

- [168] M. Hallquist and F. Hillary, *Graph theory approaches to functional network organization in brain disorders: A critique for a brave new small-world.*, *Netw. Neurosci.* **3** (2019), no. 1 1–26.
- [169] M. Greicius, B. Flores, V. Menon, G. Glover, H. Solvason, H. Kenna, A. Reiss, and A. Schatzberg, *Resting-state functional connectivity in major depression: Abnormally increased contributions from subgenual cingulate cortex and thalamus.*, *Biol. Psychiatry* **62** (2007), no. 5 429–437.
- [170] R. Buckner, J. Sepulcre, T. Talukdar, F. Krienen, H. Liu, T. Hedden, J. Andrews-Hanna, R. Sperling, and K. Johnson, *Cortical hubs revealed by intrinsic functional connectivity: mapping, assessment of stability, and relation to alzheimer’s disease.*, *J. Neurosci.* **29** (2009), no. 6 1860–1873.
- [171] R. Nebel, N. Aggarwal, L. Barnes, A. Gallagher, J. Goldstein, K. Kantarci, M. Mallampalli, E. Mormino, L. Scott, W. Yu, P. Maki, and M. Mielke, *Understanding the impact of sex and gender in alzheimer’s disease: a call to action.*, *Alzheimer’s & Dementia* **14** (2018), no. 9 1171–1183.
- [172] U. Braun, A. Schäfer, D. Bassett, F. Rausch, J. Schweiger, E. Bileka, S. Erk, N. Romanczuk-Seiferth, O. Grimm, L. Geiger, L. Haddad, K. Otto, S. Mohnke, A. Heinz, M. Zink, H. Walter, E. Schwarz, A. Meyer-Lindenberg, and H. Tost, *Dynamic brain network reconfiguration as a potential schizophrenia genetic risk mechanism modulated by nmda receptor function.*, *PNAS* **113** (2016), no. 44 12568–12573.
- [173] L. Martinet, M. Kramer, W. Viles, L. Perkins, E. Spencer, C. Chu, S. Cash, and E. Kolaczyk, *Robust dynamic community detection with applications to human brain functional networks.*, *Nature Comm.* **11** (2020) 2785.
- [174] M. Salehi, A. Greene, A. Karbasi, X. Shen, D. Scheinost, and R. Constable, *There is no single functional atlas even for a single individual: Functional parcel definitions change with task.*, *NeuroImage* **208** (2020) 116366.
- [175] R. Birn, E. Molloy, R. Patriat, T. Parker, T. Meier, G. Kirk, V. Nair, M. Meyerand, and V. Prabhakaran, *The effect of scan length on the reliability of resting-state fmri connectivity estimates.*, *Neuroimage* **83** (2013) 550–558.
- [176] Q. Telesford, M. Lynall, J. Vettel, M. Miller, S. Grafton, and D. Bassett, *Detection of functional brain network reconfiguration during task-driven cognitive states.*, *NeuroImage* **142** (2016) 198–210.
- [177] L. Jeub, M. Bazzi, I. Jutla, and P. Mucha, “A generalized louvain method for community detection implemented in matlab.”
<http://netwiki.amath.unc.edu/GenLouvain>,
<https://github.com/GenLouvain/GenLouvain>, 2011-2019.

- [178] L. van der Maaten and G. Hinton, *Visualizing data using t-sne*, *J. of Machine Learning Res.* **9** (2008) 2579–2608.
- [179] M. Meila, *Comparing clusterings-an information based distance*, *J. of Multivar. Anal.* **97** (2007), no. 5 873–895.