

Neural mechanisms underlying the temporal organization of naturalistic animal behavior

Luca Mazzucato*

Institute of Neuroscience, Departments of Biology, Mathematics and Physics,
University of Oregon, Eugene, United States

Abstract Naturalistic animal behavior exhibits a strikingly complex organization in the temporal domain, with variability arising from at least three sources: hierarchical, contextual, and stochastic. What neural mechanisms and computational principles underlie such intricate temporal features? In this review, we provide a critical assessment of the existing behavioral and neurophysiological evidence for these sources of temporal variability in naturalistic behavior. Recent research converges on an emergent mechanistic theory of temporal variability based on attractor neural networks and metastable dynamics, arising via coordinated interactions between mesoscopic neural circuits. We highlight the crucial role played by structural heterogeneities as well as noise from mesoscopic feedback loops in regulating flexible behavior. We assess the shortcomings and missing links in the current theoretical and experimental literature and propose new directions of investigation to fill these gaps.

Introduction

Naturalistic animal behavior exhibits a striking amount of variability in the temporal domain (**Figure 1A**). An individual animal's behavioral variability can be decomposed across at least three axes: hierarchical, contextual, and stochastic. The first source of variability originates from the vast hierarchy of timescales underlying self-initiated, spontaneous behavior ranging from milliseconds to minutes in animals (and to years for humans). On the sub-second timescale, animals perform fast movements varying from tens to hundreds of milliseconds. In rodents, these movements include whisking, sniffing, and moving their limbs. On the slower timescales of seconds, animals concatenate these fast movements into behavioral sequences of self-initiated actions, such as exploratory sequences (moving around an object while sniffing, whisking, and wagging their noses) or locomotion sequences (coordinating limb and head movements to reach a landmark). These sequences follow specific syntax rules (**Berridge et al., 1987**) and can last several seconds. On the timescales of minutes or longer, mice may repeat the 'walk and explore' behavioral sequence multiple times, when engaged in some specific activity, such as foraging, persisting toward their goal for long periods. In these simple examples, a freely moving mouse exhibits behavior whose temporal organization varies over several orders of magnitudes simultaneously, ranging from the sub-second scale (actions), to several seconds (behavioral sequences), to minutes (goals to attain). A leading theory to explain the temporal organization of naturalistic behavior is that behavioral action sequences arrange in a hierarchical structure (**Tinbergen, 1951; Dawkins, 1976; Simon, 1962**), where actions are nested into behavioral sequences that are then grouped into activities. This hierarchy of timescales is ubiquitously observed across species during naturalistic behavior.

The second source of temporal variability stems from the different contexts in which a specific behavior can be performed. For example, reaction times to a delivered stimulus can be consistently faster when an animal is expecting it compared to slower reaction times when the stimulus

*For correspondence:
lmazzuca@uoregon.edu

Competing interest: The author declares that no competing interests exist.

Funding: See page 25

Received: 11 January 2022

Accepted: 07 June 2022

Published: 06 July 2022

Reviewing Editor: Gordon J Berman, Emory University, United States

© Copyright Mazzucato. This article is distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use and redistribution provided that the original author and source are credited.

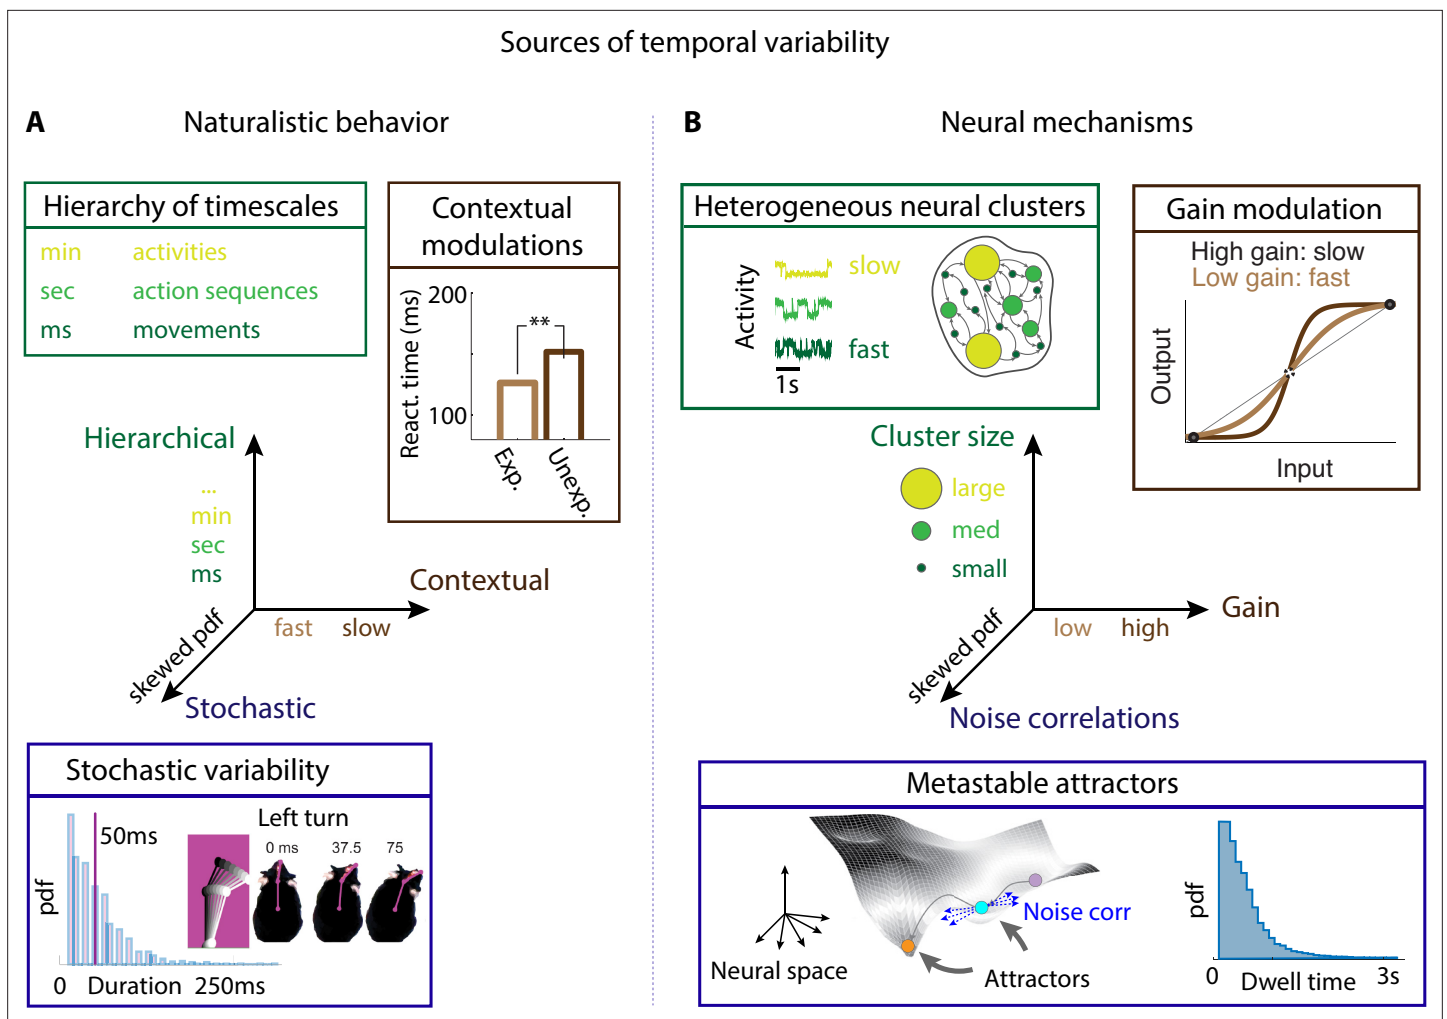


Figure 1. Neural mechanisms underlying the temporal organization of naturalistic animal behavior. **(A)** Three sources of temporal variability in naturalistic behavior: hierarchical (from fast movements, to behavioral action sequences, to slow activities and long-term goals), contextual (reaction times are faster when stimuli are expected), and stochastic (the distribution of ‘turn right’ action in freely moving rats is right-skewed). **(B)** Neural mechanisms underlying each source of temporal variability: hierarchical variability may arise from recurrent networks with a heterogeneous distribution of neural cluster sizes; contextual modulations from neuronal gain modulation; stochastic variability from metastable attractor dynamics where transitions between attractors are driven by low-dimensional noise, leading to right-skewed distributions of attractor dwell times. Panel **(A)** adapted from Figure 2 of *Jaramillo and Zador, 2011*. Panel **(A, bottom)** reproduced from Figure 6 of *Findley et al., 2021*.

is unexpected (*Jaramillo and Zador, 2011*). Contextual sources can be either internally generated (modulations of brain state, arousal, expectation) or externally driven (pharmacological, optogenetic or genetic manipulations; or changes in task difficulty or environment).

After controlling for all known contextual effects, an individual’s behavior still exhibits a large amount of residual temporal variability across repetitions of the same behavioral unit. This residual variability across repetitions can be quantified as a property of the distribution of movement durations, which is typically very skewed with an exponential tail (*Wiltchko et al., 2015; Findley et al., 2021; Recanatani et al., 2022*), suggesting a stochastic origin. We define as stochastic variability the fraction of variability in the expression of a particular behavioral unit, which cannot be explained by any readily measurable variables.

These three aspects of temporal variability (hierarchical, contextual, stochastic) can all be observed in the naturalistic behavior of a single individual during their lifetime. Other sources of behavioral variability have also been investigated. Variability in movement execution and body posture (*Dhawale et al., 2017*) may originate from the motor periphery, such as noise in force production within muscles (*van Beers et al., 2004*), or from the deterministic chaotic dynamics arising from neural activity (*Costa*

et al., 2021) or the biomechanical structure of an animal's body (*Loveless et al., 2019*). One additional source is individual phenotypic variability across different animals, an important aspect of behavior in ethological and evolutionary light (*Honegger and de Bivort, 2018*). Although we will briefly discuss some features of variability in movement execution and phenotypic individuality, the main focus of this review will be on the three main axes of variability described above, which pertain to the behavior of a single animal.

From this overview, we conclude that the temporal structure of naturalistic, self-initiated behavior can be decomposed along at least three axes of temporal variability (**Figure 1A**): a large hierarchy of simultaneous timescales at which behavior unfolds, ranging from milliseconds to minutes; contextual modulations affecting the expression of behavior at each level of this hierarchy; and stochasticity in the variability of the same behavioral unit across repetition in the same context. These three axes may or may not interact depending on the scenarios and the species. Here, we will review recent theoretical and computational results establishing the foundations of a mechanistic theory that explains how these three sources of temporal variability can arise from biologically plausible computational motifs. More specifically, we will address the following questions:

1. *Neural representation*: How are self-initiated actions represented in the brain and how are they concatenated into behavioral sequences?
2. *Stochasticity*: What neural mechanisms underlie the temporal variability observed in behavioral units across multiple repetitions within the same context?
3. *Context*: How do contextual modulations affect the temporal variability in behavior, enabling flexibility in action timing and behavioral sequence structure?
4. *Hierarchy*: How do neural circuits generate the vast hierarchy of timescales from milliseconds to minutes, hallmark of naturalistic behavior?

The mechanistic approach we will review is based on the theory of metastable attractors (**Figure 1B**), which is emerging as a unifying principle expounding many different aspects of the dynamics and function of neural circuits (*La Camera et al., 2019*). We will first establish a precise correspondence between behavioral units and neural attractors at the level of self-initiated actions (the lowest level of the temporal hierarchy). Then we will show how the emergence of behavioral sequences originates from sequences of metastable attractors. Our starting point is the observation that transitions between metastable attractors can be driven by the neural variability internally generated within a local recurrent circuit. This mechanism can naturally explain the action timing variability of stochastic origin. We will examine biologically plausible mesoscopic circuits that can learn to flexibly execute complex behavioral sequences. We will then review the neural mechanisms underlying contextual modulations of behavioral variability. We will show that the average transition time between metastable attractors can be regulated by changes in single-cell gain. Gain modulation is a principled neural mechanism mediating the effects of context, which can be induced by either internal or external perturbations and supported by different neuromodulatory and cortico-cortical pathways; or by external pharmacological or experimental interventions. Finally, we will review how a large hierarchy of timescales can naturally and robustly emerge from heterogeneities in a circuit's structural connectivity motifs, such as neural clusters with heterogeneous sizes. Although most of the review is focused on the behavior of

Table 1. Definitions of behavioral units (see Appendix 1 for how to extract these features from videos).

Action	The simplest building blocks of behavior are at the lowest level of the hierarchy. This definition depends on the level of granularity (spatiotemporal resolution) required and the scientific questions to be investigated (see Appendix 1). A widely used definition of action is a short stereotypical trajectory in posture space (synonyms include movemes, syllables [<i>Anderson and Perona, 2014; Brown and de Bivort, 2018</i>]), operationally defined as a discrete latent trajectory of an autoregressive state–space model fit to pose-tracking time-series data (Figure 2 ; <i>Wiltchko et al., 2015; Findley et al., 2021</i>). An alternative definition is in terms of short spectrotemporal representations from a time-frequency analysis of videos (<i>Berman et al., 2016; Marshall et al., 2021</i>). Examples include poking in or poking out of a nose port, waiting at a port, pressing a lever.
Behavioral action sequence	A combination of actions concatenated in a meaningful yet stereotyped way, lasting up to a few seconds. A sequence can occur during trial-based experimental protocols (e.g., a short sequence of actions aimed at obtaining a reward in an operant conditioning task [<i>Geddes et al., 2018; Murakami et al., 2014</i>]; running between opposite ends of a linear track [<i>Maboudi et al., 2018</i>]) or during spontaneous periods (e.g., repeatedly scratching own head; picking up and manipulating an object).
Activity	A concatenation of multiple behavioral sequences, often repeated and of variable duration, typically aimed at obtaining a goal and lasting up to minutes or even hours. Examples include grooming, foraging, mating, feeding, and exploration. Activities typically unfold in naturalistic freely moving settings devoid of experimenter-controlled trial structure.

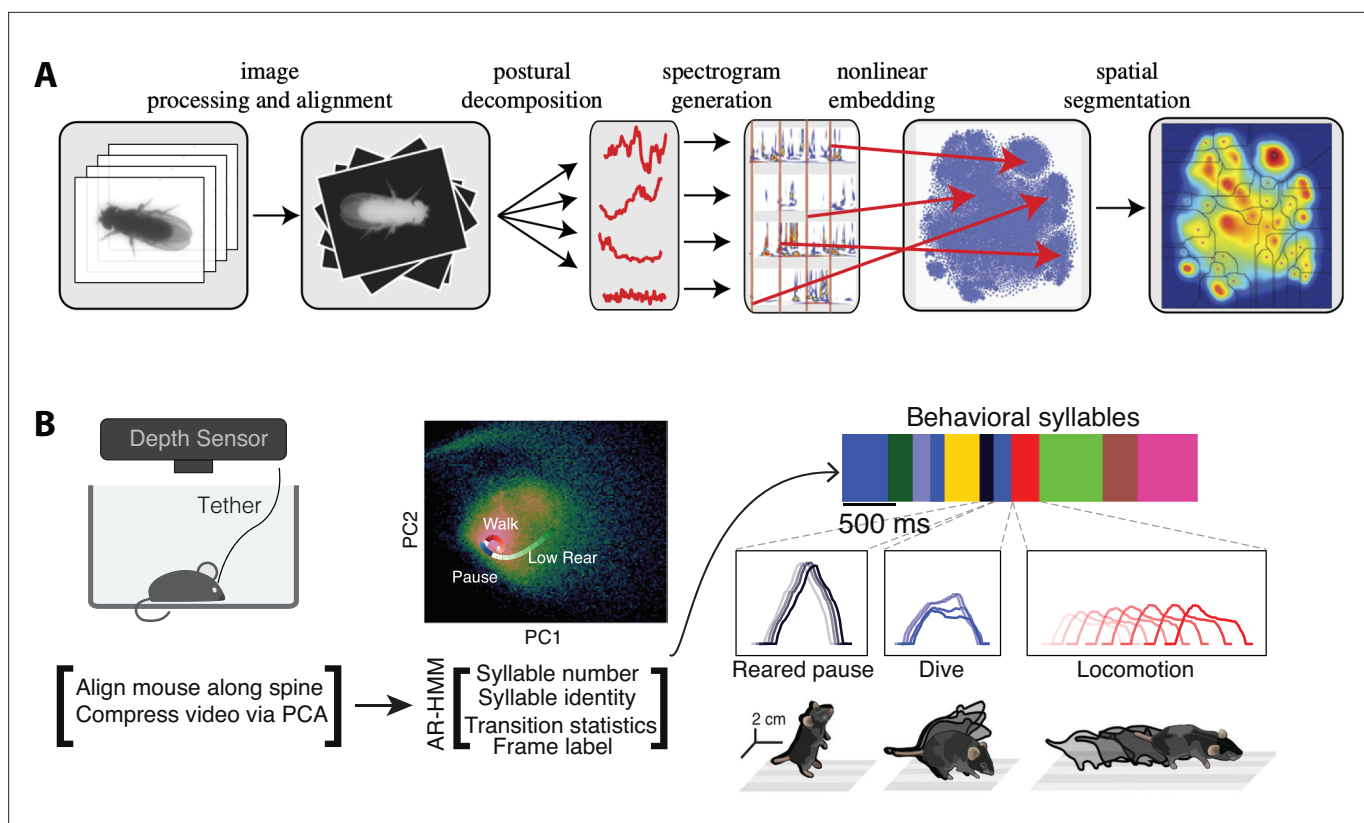


Figure 2. The complex spatiotemporal structure in naturalistic behavior. **(A)** Identification of behavioral syllables in the fruit fly. From left to right: raw images of *Drosophila melanogaster* (1) are segmented from background, rescaled, and aligned (2), then decomposed via principal component analysis (PCA) into a low-dimensional time series (3). A Morlet wavelet transform yields a spectrogram for each postural mode (4), mapped into a two-dimensional plane via t-distributed stochastic neighbor embedding (t-SNE) (5). A watershed transform identifies individual peaks from one another (6). **(B)** Identification of behavioral action sequences from 3D videos with MoSeq. An autoregressive hidden Markov model (AR-HMM, right) fit to PCA-based video compression (center) identifies hidden states representing actions (color bars, right, top: color-coded intervals where each HMM state is detected). Panel **(A)** reproduced from Figures 2 and 5 of *Berman et al., 2014*. Panel **(B)** adapted from Figure 1 of *Wiltchko et al., 2015*.

individual animals, we discuss how recent results on multianimal interactions and social behavior may challenge existing theories of naturalistic behavior and brain function. Appendices 1–3 provide guides to computational methods for behavioral video analyses and modeling; theoretical and experimental aspects of attractors dynamics; and biologically plausible models of metastable attractors.

The stochastic nature of naturalistic behavior

The study of naturalistic behavior based on animal videos has recently undergone a revolution due to the spectacular accuracy and efficiency of computational methods for animal pose tracking (*Stephens et al., 2008; Berman et al., 2014; Hong et al., 2015; Mathis et al., 2018; Pereira et al., 2019; Nourizonoz et al., 2020; Segalin et al., 2020*); see Appendix 1 and Table 1 for details). These new methods work across species and conditions, ushering a new era for computational neuroethology (*Anderson and Perona, 2014; Brown and de Bivort, 2018; Datta et al., 2019*). They have led to uncovering a quantitative classification of self-initiated behavior revealing a stunning amount of variability both in its lexical features (which actions to choose, in which order) and in its temporal dimension (when to act) (*Berman et al., 2016; Wiltchko et al., 2015; Markowitz et al., 2018; Marshall et al., 2021; Recanatesi et al., 2022*). Different ways to characterize the behavioral repertoire on short timescales have been developed, depending on the underlying assumptions of whether the basic units of behavior are discrete or continuous (Appendix 1). One can define stereotypical behaviors or postures by clustering probability density maps of spectrotemporal features extracted from behavioral videos, as in the example of fruit flies (*Berman et al., 2014; Berman et al., 2016; (Figure 2A)* or rats (*Marshall et al., 2021; Figure 7B*). Alternatively, one can capture actions or postures as discrete

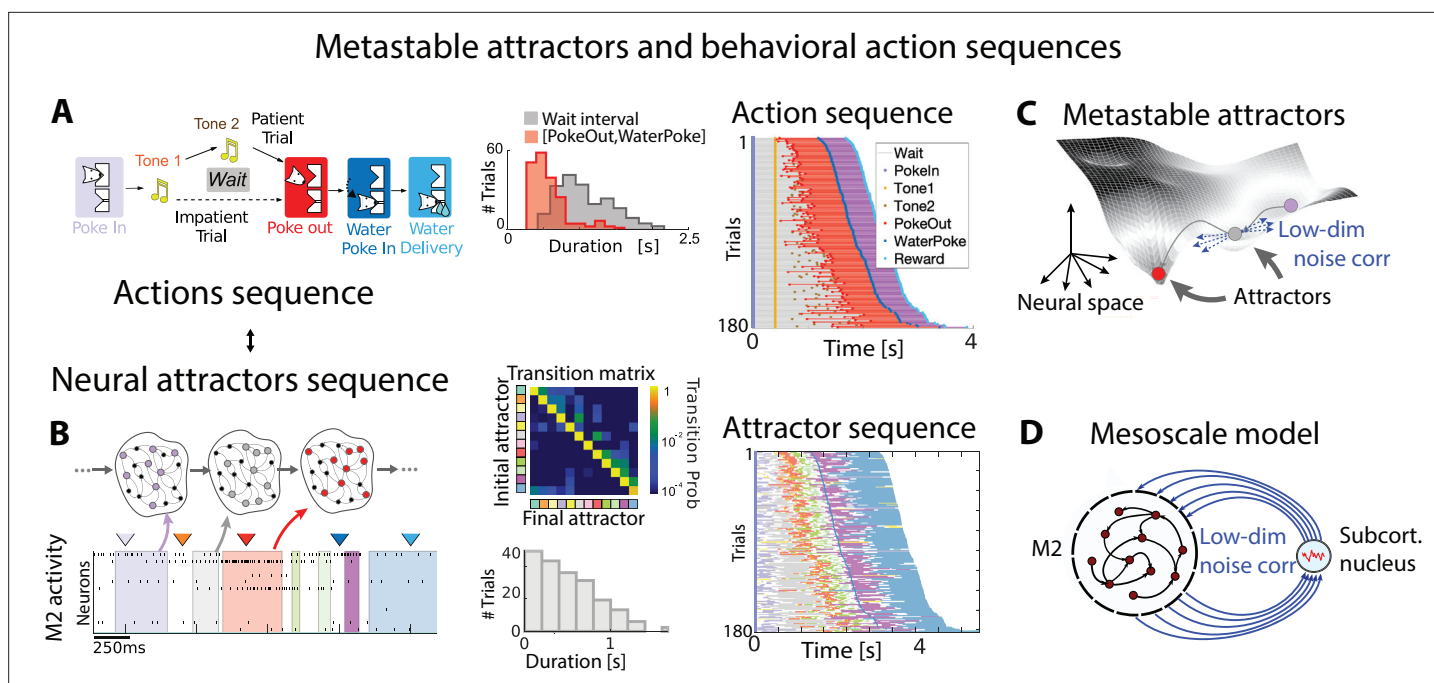


Figure 3. Metastable attractors in secondary motor cortex can account for the stochasticity in action timing. **(A)** In the Marshmallow experiment, a freely moving rat poked into a wait port, where tone 1 signaled the trial start. The rat could poke out at any time (after tone 2, played at random times, in patient trials; or before tone 2, in impatient trials) and move to the reward port to receive a water reward (large and small for patient and impatient trials, respectively). Center: distributions of wait times (gray) and intervals between Poke Out and Water Poke (red) reveal a large temporal variability across trials. Right: behavioral sequences (sorted from shortest to longest) in a representative session. **(B)** Left: representative neural ensemble activity from M2 during an impatient trial (tick marks indicate spikes; colored arrows indicate the rat's actions, same colors as in **A**) with overlaid hidden Markov model (HMM) states, interpreted as neural attractors each represented as a set of coactivated neurons within a network (colored intervals indicate HMM states detected with probability above 80%). Top center: transition probability matrix between HMM states. Bottom center: the distribution of state durations (representative gray pattern from left plot) is right-skewed, suggesting a stochastic origin of state transitions. Right: sequence of color-coded HMM states from all trials in the representative session of panel **(A)**. **(C)** Schematic of an attractor landscape: attractors representing HMM states in panel **(C)** are shown as potential wells. Transitions between consecutive attractors are driven by low-dimensional correlated noise. **(D)** Schematic of a mesoscale neural circuit generating stable attractor sequences with variable transition times comprising a feedback loop between M2 and a subcortical nucleus. Panels **(A)** and **(B)** adapted from Figures 1 and 5 of *Recanatesi et al., 2022*.

states of a state–space model based on Markovian dynamics, each state represented as a latent state autoregressive trajectory accounting for stochastic movement variability (*Figure 2*). In the case of discretization of the behavioral repertoire, the definition of the smallest building block of behavior depends on the level of granularity that best suits the scientific question being investigated. Within this discrete framework, state–space models were applied successfully to *Caenorhabditis elegans* (*Linderman et al., 2019*), *Drosophila* (*Tao et al., 2019*), zebrafish (*Johnson et al., 2020*), and rodents (*Wiltschko et al., 2015; Markowitz et al., 2018*). At the timescale of actions and postures, transition times between consecutive actions are well described by a Poisson process (*Killeen and Fetterman, 1988*), characterized by a right-skewed distribution of inter-action intervals (*Figure 3*). In this and the next section, we will focus on short timescales (up to a few seconds) where state–space models can provide parsimonious accounts of behavior. However, these models fail to account for longer timescale and non-Markovian structure in behavior, and in later sections we will move to a more data-driven approach to investigate these aspects of behavior (see Appendix 1).

Self-initiated actions and ensemble activity patterns in premotor areas

We begin the investigation of naturalistic behavior starting at the lowest level of the hierarchy and examining the neural mechanisms underlying self-initiated actions. How does the brain encode the intention to perform an action? How are consecutive actions concatenated into a behavioral sequence? Unrestrained naturalistic behavior, as it has been recently characterized in the foundational studies mentioned above, revealed a vast repertoire of dozens to hundreds of actions (*Berman et al., 2014*;

Wiltchko et al., 2015; Markowitz et al., 2018; Marques et al., 2018; Costa et al., 2019; Schwarz et al., 2015) (although the repertoire depends on the coarse-graining scale of the behavioral analysis), leading to a combinatorial explosion in the number of possible action sequences. To tame the curse of dimensionality, typical of unconstrained naturalistic behavior, a promising approach is to control for the lexical variability in behavioral sequences and design a naturalistic task where an animal performs a single behavioral sequence of self-initiated actions, where each action retains its unrestrained range of temporal variability. *Murakami et al., 2014, Murakami et al., 2017, and Recanatesi et al., 2022* adopted this strategy to train freely moving rats to perform a self-initiated task (the rodent version of the 'Marshmallow task'; *Mischel and Ebbesen, 1970, Figure 3A*), where a specific set of actions had to be performed in a fixed order to obtain a reward. The many repetitions of the same behavioral sequence yielded a large sample size to elucidate the source of temporal variability across trials. Action timing retained the temporal variability hallmark of naturalistic behavior, characterized by skewed distributions of action durations (*Figure 3A*).

What is the neural mechanism generating the large variability in action timing? A large number of studies implicated the secondary motor cortex (M2) in rodents as part of a distributed network involved in motor planning in head-fixed mice (*Li et al., 2016; Barthas and Kwan, 2017*) and controlling the timing of self-initiated actions in freely moving rats (*Murakami et al., 2014; Murakami et al., 2017*). Ensemble activity recorded in M2 during the Marshmallow task unfolded via sequences of multineuron-firing patterns, each one lasting hundreds of milliseconds to a few seconds; within each pattern, neurons fired at an approximately constant rate (*Figure 3B*). Such long dwell times, much longer than typical single-neuron time constants, suggest that the observed metastable patterns may be an emergent property of the collective circuit dynamics within M2 and reciprocally connected brain regions. Crucially, both neural and behavioral sequences were highly reliable yet temporally variable, and the distribution durations of action and neural pattern durations were characterized by a right-skewed distribution. This temporal heterogeneity suggests that a stochastic mechanism, such as that found in noise-driven transitions between metastable states, could contribute to driving transitions between consecutive patterns within a sequence (see below). A dictionary between actions and neural pattern could be established, revealing that the onset of specific patterns reliably preceded upcoming self-initiated actions (*Figure 3A and B*; e.g., the onset of the red pattern reliably precedes the poke out movement). The dictionary trained on the rewarded sequence generalized to epochs where the animal performed erratic nonrewarded behavior, where pattern onset predicted upcoming actions as well. The use of state-space models with underlying Markovian dynamics as generative models, capturing both naturalistic behavior (*Wiltchko et al., 2015; Batty et al., 2019; Johnson et al., 2020; Parker et al., 2021; Findley et al., 2021*) and the underlying neural pattern sequences (*Maboudi et al., 2018; Linderman et al., 2019; Recanatesi et al., 2022*), is a powerful tool to bridge the first two levels of the temporal hierarchy in naturalistic behavior: a link from actions to behavioral sequences. This generative framework further revealed fundamental aspects of neural coding in M2 ensembles, such as their distributed representations, dense coding, and single-cell multistability (*Recanatesi et al., 2022; Mazzucato et al., 2015*).

Action timing variability from metastable attractors

Neural patterns in M2 may represent attractors of the underlying recurrent circuit dynamics. Attractors are a concept in the theory of dynamical systems, defined as the set of states a system evolves to starting from a large set of initial conditions. The central feature of an attractor is that, if the system is perturbed slightly away from an attractor, it tends to return to it. Within the context of neural activity, an attractor is a persisting pattern of population activity where neurons fire at an approximately constant rate for an extended period of time (see Appendix 1 for models of attractor dynamics). Foundational work in head-fixed mice showed that licking preparatory activity in M2 during a delayed response task is encoded by choice-specific discrete attractors (*Inagaki et al., 2019*). Attractors can be stable (*Amit and Brunel, 1997*), as observed in monkey IT cortex in working memory tasks (*Fuster and Jervey, 1981; Miyashita and Chang, 1988*). Attractors can also be metastable (see Appendix 3), when they typically last for hundreds of milliseconds and noise fluctuations spontaneously trigger transitions to a different attractor (*Deco and Hugues, 2012; Litwin-Kumar and Doiron, 2012*).

Metastable attractors can be concatenated into sequences, which can either be random, as observed during ongoing periods in sensory cortex (*Mazzucato et al., 2015*), or highly reliable,

encoding the evoked response to specific sensory stimuli (*Jones et al., 2007; Miller and Katz, 2010; Mazzucato et al., 2015*), or underlying freely moving behavior (*Maboudi et al., 2018; Recanatesi et al., 2022*). In particular, M2 ensemble activity in the Marshmallow experiment was consistent with the activity generated by a specific sequence of metastable attractors (*Figure 3*). The main hypothesis underlying this model is that the onset of an attractor drives the initiation of a specific action as determined by the action/pattern dictionary (*Figure 3A and B*) and the dwell time in a given attractor sets the inter-action-interval. The dynamics of the relevant motor output and the details of variability in movement execution are generated downstream of this attractor circuit (see *Figure 5* and 'Discussion' section). The main features observed in the M2 ensemble dynamics during the Marshmallow task (i.e., long-lived patterns with a right-skewed dwell-time distribution, concatenated into highly reliable pattern sequences) can be explained by a two-area mesoscopic network where a large recurrent circuit (representing M2) is reciprocally connected to a small circuit lacking recurrent couplings (a subcortical area likely representing the thalamus [*Guo et al., 2018; Guo et al., 2017; Inagaki et al., 2022*], see *Figure 5C*). In this biologically plausible model, metastable attractors are encoded in the M2 recurrent couplings, and transitions between consecutive attractors are driven by low-dimensional noise fluctuations arising in the feedback projections from the subcortical nucleus back to M2. As a consequence of the stochastic origin of the transitions, the distribution of dwell times for each attractor is right-skewed, closely matching the empirical data. This model's prediction was confirmed in the data, where the ensemble fluctuations around each pattern were found to be low-dimensional and oriented in the direction of the next pattern in the sequence (*Figure 3*). This model presents a new interpretation for low-dimensional (differential) correlations: although their presence in sensory cortex may be detrimental to sensory encoding (*Moreno-Bote et al., 2014*), their presence in motor circuits seems to be essential for motor generation during naturalistic behavior (*Recanatesi et al., 2022*).

Open issues

The hypothesis that preparatory activity for upcoming actions is encoded in discrete attractors in M2 has been convincingly demonstrated using causal manipulations in head-fixed preparations in mice (see *Appendix 1*), although a causal test of this hypothesis in freely moving animals is currently lacking. A shortcoming of the metastable attractor model of action timing in *Recanatesi et al., 2022* is the unidentified subcortical structure where the low-dimensional variability is originating. Thalamus and basal ganglia are both likely candidates as part of a large reciprocally connected mesoscopic circuit underlying action selection and execution (see *Figure 5B*), and more work is needed to precisely identify the origin of the low-dimensional variability.

The metastable attractor model assumes the existence of discrete units of behavior at the level of actions, although large variability in movement execution originating downstream of cortical areas may blur the distinction between the discrete behavioral units. The extent to which behavior can be interpreted as a sequence of discrete behavioral units or, rather, a superposition of continuously varying poses (see *Appendix 1* for an in-depth discussion of this issue) is currently open for debate.

At higher levels of the behavioral hierarchy, repetitions of the same behavioral sequence (such as a jump attempt, an olfactory search trial, or a waiting trial, *Figures 2 and 3*) exhibit large temporal variability as well, characterized by right-skewed distributions (*Lottem et al., 2018*). It remains to be examined whether temporal variability in sequence duration may originate from a hierarchical model where sequences themselves are encoded in slow-switching metastable attractors in a higher cortical area or a distributed mesoscopic circuit (see *Figure 5C* and *Figure 9*).

A theory of metastable dynamics in biologically plausible models

A variety of neural circuit models have been proposed to generate metastable dynamics (see *Appendix 3*). However, a full quantitative understanding of the metastable regime is currently lacking. Such theory is within reach in the case of recurrent networks of continuous rate units. In circuits where metastable dynamics arises from low-dimensional correlated variability (*Recanatesi et al., 2022*), dynamic mean field methods could be deployed to predict the statistics of switch times from underlying biological parameters. In biologically plausible models based on spiking circuits, it is not known how to quantitatively predict switching times from underlying network parameters. Phenomenological birth–death processes fit to spiking network simulations can give a qualitative understanding

of the on–off cluster dynamics (*Huang and Doiron, 2017; Shi et al., 2022*), and it would be interesting to derive these models from first principles. Mean field methods for leaky-integrate-and-fire networks can give a qualitative prediction of the effects of external perturbations on metastable dynamics, explaining how changes in an animal's internal state can affect circuit dynamics (*Mazzucato et al., 2019; Wyrick and Mazzucato, 2021*). These qualitative approaches should be extended to fully quantitative ones. A promising method, deployed in random neural networks, is based on the universal colored-noise approximation to the Fokker–Planck equation, where switch times between metastable states can be predicted from microscopic network parameters such as neural cluster size (*Stern et al., 2021*). Finally, a crucial direction for future investigation is to improve the biological plausibility of metastable attractor models to incorporate different inhibitory cell types. Progress along this line will open the way to quantitative experimental tests of the metastable attractor hypothesis using powerful optogenetic tools.

Stochastic individuality

One important source of behavioral variability is phenotypic variability across different individuals with identical genetic profile, an important aspect of behavior in ethological and evolutionary light (*Honegger and de Bivort, 2018*). Stochastic individuality is defined as the part of the phenotypic variability in nonheritable effects that cannot be predicted from measurable variables such as learning or other developmental conditions – such as behavioral differences in identical twins reared in the same environment. Signatures of stochastic individuality have been found in rodents (*Kawai et al., 2015; Marshall et al., 2021*) and flies (*Tao et al., 2019*). Existing theoretical models have not examined which neural mechanisms may underlie this individuality. This variability may arise from differences in developmental wiring of brain circuits related to axonal growth. The metastable attractor model in *Recanatesi et al., 2022* may naturally accommodate some stochastic individuality. The location of each attractor in firing rate space is drawn from a random distribution, so that across-animal variability may stem from different random realization of the attractor landscape with the same underlying hyperparameters (i.e., mean and covariance of the neural patterns).

Variability in movements and body posture

A large source of variability in behavior originates from variability in movement execution (*Dhawale et al., 2017*). In a delayed response task in overtrained primates, preparatory neural activity in premotor areas could only account for half of the trial-to-trial variability in movement execution (*Churchland et al., 2006*). This fraction of motor variability originating from preparatory neural activity can be potentially explained by the mesoscopic attractor model in *Figure 3D*. The remaining fraction may originate from the motor periphery, such as noise in force production within muscles (*van Beers et al., 2004*), which is not accounted for by our attractor models. In humans, trial-to-trial motor variability can be interpreted as a means to update control policies and motor output within a reinforcement learning paradigm (*Wu et al., 2014*). It is challenging to dissect this finer movement variability from the state–space model approach to behavioral classification as it aims at capturing discrete stereotypical movement features. Variability in postural behavior can also stem from purely deterministic dynamics, which have been modeled using state–space reconstruction based on delay embedding of multivariate time series (*Costa et al., 2019; Costa et al., 2021; Costa et al., 2021*). These methods are fundamentally different from state–space models: while the latter assume an underlying probabilistic origin of variability, the former assume that variability arises from deterministic chaos. In order to capture the finer scale of movements and body posture in *C. elegans* (*Schwarz et al., 2015; Szigeti et al., 2015*), a data-driven state–space reconstruction revealed that locomotion behavior exhibits signature of chaos and can be explained in terms of unstable periodic orbits (*Costa et al., 2021*). A completely different approach posits that variability in movements emerges from an animal's body mechanics in the absence of any neural control. Strikingly, in a mechanical model of the *Drosophila* larvae body, the animal's complex exploratory behaviors were shown to emerge solely based on the features of the body biomechanics in the absence of any neural dynamics (*Loveless et al., 2019*). In particular, behavioral variability stems from an anomalous diffusion process arising from the deterministic chaotic dynamics of the body.




Figure 4. Contextual modulations mediated by changes in internal states. **(A)** Expectation modulates reaction times. Top: freely moving rats were trained to initiate a trial by choosing either side port depending on the frequency features of a presented stimulus (a target frequency embedded in a train of distractors) to collect reward. Bottom: reaction time for targets that were expected (light brown) was faster than unexpected (dark brown). **(B)** Expectation induces faster stimulus coding. Contextual modulations that accelerate stimulus coding and reaction times may operate via a decrease in single-cell intrinsic gain (left), lowering the energy barrier separating noncoding attractors to the stimulus-coding attractor (center). Lower barriers allow for faster transitions into the stimulus coding attractor, mediating faster encoding of sensory stimuli in the expected condition compared to the unexpected condition. Right: representative ensemble activity from rats gustatory cortex in two trials where the taste delivery was expected (bottom) or unexpected (top). The onset of the taste-coding attractor (blue) occurs earlier when the taste delivery is expected. Panel **(A)** adapted from Figures 1 and 2 of *Jaramillo and Zador, 2011*. Panel **(B)** adapted from Figure 5 of *Mazzucato et al., 2019*.

Contextual modulation of temporal variability

The second source of temporal variability in naturalistic behavior arises from contextual modulations, which can be *internally* or *externally* driven. When internally driven, they may arise from changes in brain or behavioral state such as arousal, expectation, or task engagement. When externally driven, they may arise from changes in environmental variables, from the experimenter's imposed task conditions, or from manipulations such as pharmacological, optogenetic, or genetic ones. Contextual modulations can affect several qualitatively different aspects of behavioral units at each level of the hierarchy (actions, sequences, activities): average duration; usage frequency; and transition probabilities between units. Moreover, context may also change the motor execution of a behavioral unit, for example, by improving the vigor of a certain movement upon learning or motivation. For each type of modulation, we will give several examples and review computational mechanisms that may explain them.

Action timing

The distribution of self-initiated action durations typically exhibits large variability, whose characteristic timescale can be extracted from their average duration (*Figure 3*). The average timing of an action is strongly modulated by contextual factors, both internally and externally driven. Examples of internally generated contextual factors include expectation and history effects. When events occur at predictable instants, anticipation improves performance such as reaction times. This classic effect of expectation was documented in an auditory two-alternative choice task (*Jaramillo and Zador, 2011*), where freely moving rats were rewarded for correctly discriminating the carrier frequency of a frequency-modulated target sound immersed in pure-tone distractors (*Figure 4A*). The target could occur early or late within each sound presentation, and temporal expectations on target timing were modulated by changing the ratio of trials with early or late targets within each block. When manipulating the expectations about sound timing, valid expectations accelerated reaction times and improved detection accuracy, showing enhanced perception. The auditory cortex is necessary to perform this task, and firing rates in auditory cortex populations are modulated by temporal expectations.

Contextual changes in brain state may also be induced by varying levels of neuromodulators. In the self-initiated waiting task of **Figure 3A**, optogenetic activation of serotonergic neurons in the dorsal raphe nucleus selectively prolonged the waiting period, leading to a more patient and less impulsive behavior, but did not affect the timing of other self-initiated actions (**Fonseca et al., 2015**). Internally generated contextual modulations include history effects, which can affect the timing of self-initiated actions. In an operant conditioning task, where freely moving mice learned to press a lever for a minimum duration to earn a reward (**Fan et al., 2012**), the distribution of action timing showed dependence on the outcome of the preceding trial. After a rewarded trial, mice exhibited longer latency to initiate the next trial, but shorter press durations; after a failure, the opposite behavior occurred, with shorter latency to engage and longer press durations. Trial history effects are complex and action-specific and depend on several other factors, including prior movements, and wane with increasing inter-trial intervals. Inactivation experiments showed that these effects rely on frontal areas such as medial prefrontal and secondary motor cortices (**Murakami et al., 2014; Murakami et al., 2017; Schreiner et al., 2021**).

Although the contextual modulations considered so far occur on a fast timescale of a few trials, they may also be the consequence of associative learning. In a lever press task, mice learned to adjust the average duration of the lever press to different criteria in three different conditions where lever presses were always rewarded regardless of duration or only rewarded if longer than 800 or 1600 ms. Within each of the three conditions, the distribution of action timing exhibited large temporal variability, yet the average duration was starkly different between the three criteria as the mice learned the different criteria (**Fan et al., 2012; Schreiner et al., 2021**). Reaction times to sensory stimuli and self-initiated waiting behavior, in the form of long lever presses or nose-pokes, have emerged as a fruitful approach to test hypotheses on contextual modulations and decision-making in naturalistic scenarios.

Controlling action timing via neuronal gain modulation

The paramount role of contextual modulations in regulating action timing during naturalistic scenarios has been well documented. However, the neural mechanisms underlying these effects remain elusive. Results from head-fixed preparations revealed some possible explanations, which have the potential to generalize to the freely moving case. Within the paradigm of metastable attractors (see Appendix 3), the speed at which cortical activity encodes incoming stimuli can be flexibly controlled in a state-dependent manner by transiently changing the baseline level of afferent input currents to a local cortical circuit. These baseline changes may be driven by top-down projections from higher cortical areas or by neuromodulators. In a recurrent circuit exhibiting attractor dynamics, changes in baseline levels modulate the average transition times between metastable attractors (**Mazzucato et al., 2019; Wyrick and Mazzucato, 2021**). In these models, attractors are represented by potential wells in the network's energy landscape, and the height of the barrier separating two nearby wells determines the probability of transition between the two corresponding attractors (lower barriers are easier to cross and lead to faster transitions, **Figure 4B**). Changes in input baseline that decrease (increase) the barrier height lead to faster (slower) transitions to the coding attractor, in turn modulating reaction times.

Although it is not possible to measure potential wells directly in the brain, using mean field theory one can show that the height of these potential wells is directly proportional to the neuronal gain as measured by single-cell transfer functions (for an explanation, see Appendix 3). In particular, a decrease (increase) in pyramidal cells gain can lead to faster (slower) average action timing. This biologically plausible computational mechanism was proposed to explain the acceleration of sensory coding observed when gustatory stimuli are expected compared to when they are delivered as a surprise (**Samuelsen et al., 2012; Mazzucato et al., 2019**); and the faster encoding of visual stimuli observed in V1 populations during locomotion periods compared to when the mouse sits still (**Wyrick and Mazzucato, 2021**). Transition rates between attractors may also be modulated by varying the amplitude and color of the fluctuations in the synaptic inputs (see Appendix 3), while keeping the barrier heights fixed. Changes in E/I background synaptic inputs were also shown to induce gain modulation and changes in integration time (**Chance et al., 2002**).

How can a neural circuit learn to flexibly adjust its responses to stimuli or the timing of self-initiated actions? Theoretical work has established gain modulation as a general mechanism to flexibly control

network activity in recurrent network models of motor cortex (*Stroud et al., 2018*). Individual modulation of each neuron's gain can allow a recurrent network to learn a variety of target outputs through reward-based training and combine learned modulatory gain patterns to generate new movements. After learning, cortical circuits can control the speed of an intended movement through gain modulation and affect the shape or the speed of a movement. Although the model in *Stroud et al., 2018* could not account for the across-repetition temporal variability in action timing, it is tempting to speculate that a generalization of this learning framework to incorporate the metastable attractor models of *Mazzucato et al., 2019* could allow a recurrent circuit to learn flexible gain modulation via biologically plausible synaptic plasticity mechanisms (*Litwin-Kumar and Doiron, 2014*). This hypothetical model could potentially explain the contextual effects of learning on action timing observed in *Fan et al., 2012*; *Schreiner et al., 2021*, and the acceleration of reaction times in the presence of auditory expectations (*Figure 4A*; *Jaramillo and Zador, 2011*). Although this class of models has not been directly tested in freely moving assays, we believe that pursuing this promising direction could lead to important insights.

Behavioral sequences

Contextual modulations of temporal variability may affect other aspects of natural behavior beyond average action timing, such as the frequency of occurrence of an action ('state usage') or the transition probabilities between actions within a behavioral sequence (*Figure 5A*). These effects can be uncovered by analyzing freely moving behavioral videos with state-space models, such as the autoregressive hidden Markov model (HMM) (*Wiltchko et al., 2015*; *Figures 2B and 5A* and Appendix 1). Differences in state usage and transition probabilities between conditions may shed light on the computational strategies animals deploy to solve complex ethological tasks, such as those involving sensorimotor integration. For example, in a distance estimation task where mice were trained to jump across a variable gap, a comparison of monocular and binocular mice revealed the different visually guided strategies mice may use to perform a successful jump. Mice performed more vertical head movements under monocular conditions compared to control (*Figure 5A*, states 2 and 3 occur more frequently in the monocular condition), revealing a reliance on motion parallax cues (*Parker et al., 2021*).

During ongoing periods, in the absence of a task, animal behavior features a large variety of actions and behavioral sequences (*Figure 2*). Experimentally controlled manipulations can lead to strong changes in ongoing behavior reflected both in changes of state usage and of transition probabilities between actions, resulting in different repertoires of behavioral sequences. Examples of manipulations include exposing mice to innately aversive odors and other changes in their surrounding environment; optogenetic activation of corticostriatal pathways (*Wiltchko et al., 2015*); and pharmacological treatment (*Wiltchko et al., 2020*). In the latter study, a classification analysis predicted with high accuracy which drug and specific dose was administered to the mice from a large panel of compounds at multiple doses. Comparison of state usage and transition rates can also reveal subtle phenotypical changes in the structure of ongoing behavior in genetically modified mice compared to wildtype ones. This phenotypic fingerprinting has led to insights into the behavior of mouse models of autism spectrum disorder (*Wiltchko et al., 2015*; *Wiltchko et al., 2020*; *Klibaite et al., 2021*).

Computational mechanisms underlying flexible behavioral sequences

Recent studies have begun to shed light on the rules that may control how animals learn and execute behavioral sequences. These studies revealed various types of contextual modulations such as changes in the occurrence of single actions or in the transition probabilities between pairs of actions and proposed potential mechanisms underlying these effects. Biologically plausible models of mesoscopic neural circuits can generate complex sequential activity (*Logiaco et al., 2021*; *Murray and Escola, 2017*). In a recent model of sequence generation (*Logiaco and Escola, 2020*; *Logiaco et al., 2021*), an extensive library of behavioral motifs and their flexible rearrangement into arbitrary sequences relied on the interaction between motor cortex, basal ganglia, and thalamus. In this model (*Figure 5B*), the basal ganglia sequentially disinhibit motif-specific thalamic units, which in turn trigger motif preparation and execution via a thalamocortical loop with the primary motor cortex (M1). Afferent inputs to M1 set the initial conditions for motif execution. This model represents a biologically plausible neural implementation of a switching linear dynamical system (LDS) (*Linderman et al., 2017*; *Nassar et al.,*

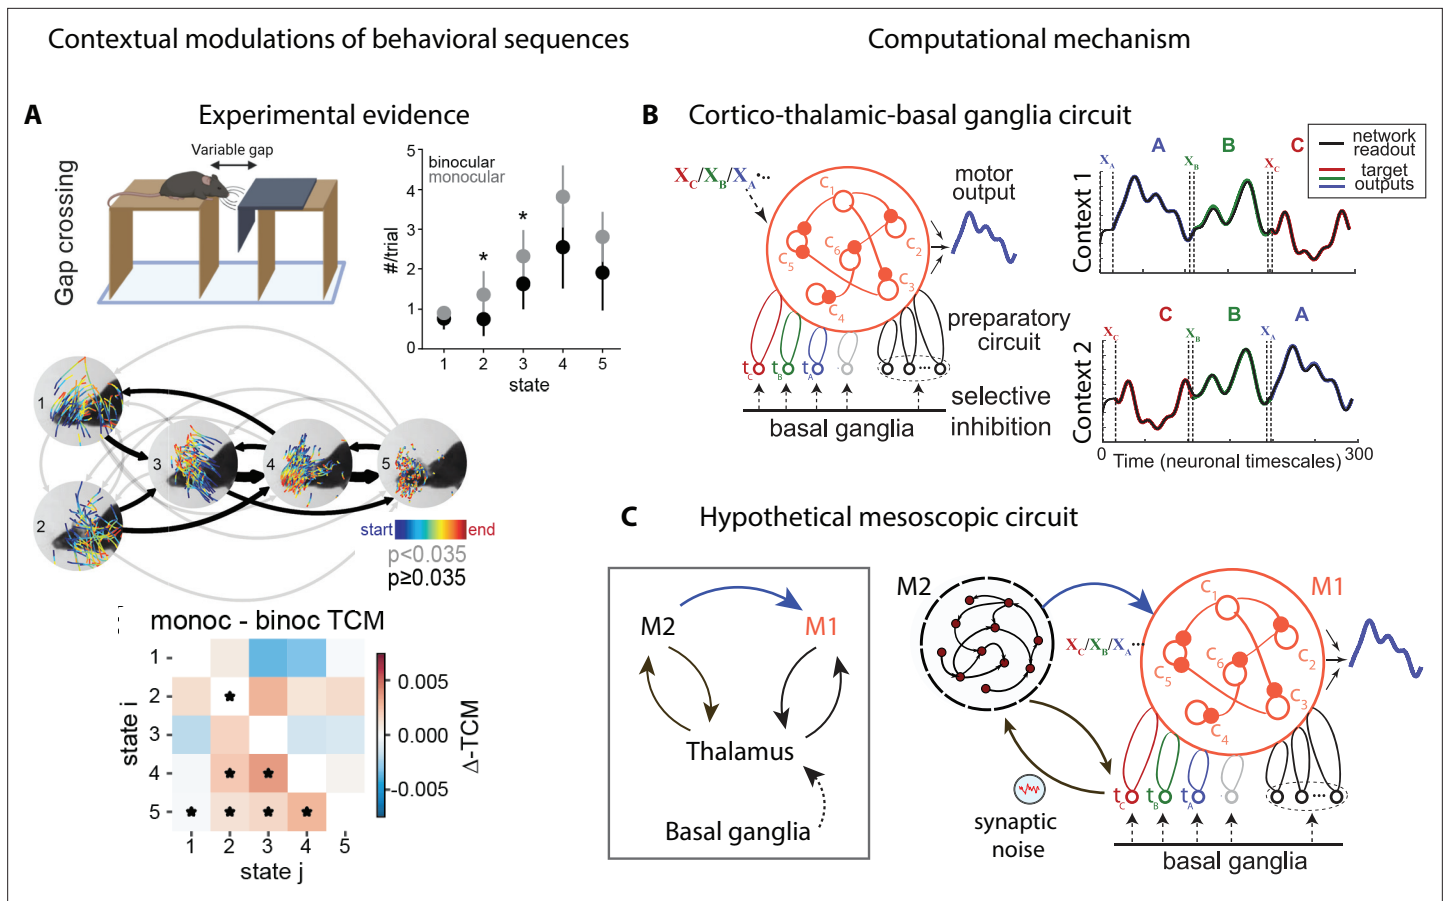


Figure 5. Contextual modulations of behavioral sequences. **(A)** Freely moving mice were rewarded for successfully jumping across a variable gap. Center: example traces of eye position from five movement states labeled with autoregressive hidden Markov model (AR-HMM) of DeepLabCut-tracked points during the decision period (progressing blue to red in time) in average temporal order (arrow line widths are proportional to transition probabilities between states; gray $< 0.035 \leq$ black; states 2–3 represent vertical head movements). Top right: frequency of occurrence of each state for binocular (black) and monocular (gray) conditions. Bottom: difference between monocular and binocular transition count matrices; red transitions are more frequent in monocular, blue in binocular ($*p < 0.01$). **(B)** Left: cortex–thalamus–basal ganglia circuit for behavioral sequence generation. The basal ganglia projections select thalamic units ($t_{A,B,C}$) needed for either motif execution or preparation. During preparation, the cortical population c_i also receives an input x_m specific to the upcoming motif m . Right: generation of sequences of arbitrary orders, using preparatory periods (between vertical dashed lines) before executing each motif. **(C)** Hypothetical cortex–thalamus–basal ganglia circuit for behavioral sequence generation combining the metastable attractor model (**Figure 2C and D**) with the model of panel **(B)**. Secondary motor cortex (M2) provides the input to primary motor cortex (M1), setting the initial conditions $x_{A,B,C}$ for each motif A, B, C . A thalamus–M2 feedback loop sustains the metastable attractors in M2, and synaptic noise in these thalamus-to-M2 projections generates temporal variability in action timing.

© 2021, Parker et al. Panel A reproduced from Figure 3 of **Parker et al., 2021** published under the [CC BY-NC-ND 4.0 license](https://creativecommons.org/licenses/by-nc-nd/4.0/). Further reproduction of the panel must follow the terms of this license.

© 2021, Logiaco et al. Panels B and C are adapted from Figure 4 of **Logiaco et al., 2021**. They are not covered by the CC-BY 4.0 license and further reproduction of these panels would need permission from the copyright holder.

2018), a class of generative models whose statistical structure can capture the spatiotemporal variability in naturalistic behavior (see Appendix 1).

How do animals learn context-dependent behavioral sequences? Within the framework of corticostriatal circuits, sequential activity patterns can be learned in an all-inhibitory circuit representing the striatum (**Murray and Escola, 2017**). Learning in this model is based on biologically plausible synaptic plasticity rules, consistent with the decoupling of learning and execution suggested by lesion studies showing that cortical circuits are necessary for learning, but that subcortical circuits are sufficient for expression of learned behaviors (**Kawai et al., 2015**). This model can achieve contextual control over temporal rescaling of the sequence speed and facilitate flexible expression of distinct sequences via selective activation and concatenation of specific subsequences. Subsequent work

uncovered a new computational mechanism underlying how motor cortex, thalamus, and the striatum coordinate their activity and plasticity to learn complex behavior on multiple timescales (*Murray and Escola, 2020*). The combination of fast cortical learning and slow subcortical learning may give rise to a covert learning process through which control of behavior is gradually transferred from cortical to subcortical circuits, while protecting learned behaviors that are practiced repeatedly against overwriting by future learning.

Open issues

It remains to clarify the exact extent to which the stochastic and contextual variability are independent sources. In principle, one could hypothesize that detailed knowledge and control of all experimental variables and behavioral state and their history might explain part of the trial-to-trial variability as contextual variability conditioned on these variables. For example, in the Marshmallow task about 10% of the temporal variability in waiting times could be attributed to a trial history effect, which relied on an intact medial prefrontal cortex and was abolished with its inactivation (*Murakami et al., 2017*). Moreover, in a subset of sessions this fraction of variability could be predicted by the activity of transient neurons before trial onset (*Murakami et al., 2014; Murakami et al., 2017*). The remaining 90% of the unexplained across-repetition variability was attributed to a stochastic mechanism, likely originating from metastable dynamics (*Recanatesi et al., 2022*). In general, it is an open question to investigate whether what we think of as noise driving trial-to-trial variability could just be another name for a contextual variable that we have not yet quantified. Alternatively, stochastic variability could be genuinely different from contextual variability and originate from noise inherent in neural spiking or other activity-dependent mechanisms.

A mesoscopic circuit for flexible action sequences

The computational mechanisms discussed so far can separately account for some specific features of contextual modulations, but none of the existing models can account for all of them. Metastable attractor models, where an attractor/action dictionary can be established, explain how the large variability in action timing may arise from correlated neural noise driving transitions between attractors (*Figure 3C*). These models can explain how contextual modulations of action timing may arise from neuronal gain modulation (*Figure 4B*). However, it is not known whether such models can explain the flexible rearrangement of actions within a behavioral sequence (*Figure 5A*). Conversely, models of flexible sequence execution and learning (*Logiaco et al., 2021; Murray and Escola, 2017; Stroud et al., 2018*) can explain the latter effect, but do not incorporate temporal variability in action timing across repetition. We would like to propose a hypothetical circuit model that combines the cortex–basal ganglia–thalamic model of *Figure 5B* together with the metastable attractor model of preparatory activity of *Figure 3C and D* to provide a tentative unified model of temporally variable yet flexible behavioral sequences. In the original model of *Figure 5B*, the movement motif *A* in M1 is executed following a specific initial condition x_A , representing an external input to M1. In the larger mesoscopic circuit of *Figure 5C*, we can interpret x_A as the preparatory activity for the motif *A*, dynamically set by an M2 attractor, which encodes the upcoming motif, and mediated by an M2-to-M1 projection. In the model of *Figure 5B*, the motif *A* in M1 is selected by the activation of a thalamic nucleus t_A , which is gated in by basal ganglia-to-thalamus projections. In the larger mesoscopic circuit of *Figure 5C*, the thalamic nuclei are also part of a feedback loop with M2 (already present in the model of *Figure 3D*) responsible for sustaining the metastable attractors encoding the preparatory activity for motif *A*. Presynaptic noise in the thalamus-to-M2 projections implement the variability in action timing via noise-driven transitions between M2 attractors. Another prominent source of behavioral variability, the one in movement execution, can potentially be incorporated into this model. Specifically, fluctuations in neural preparatory activity (which account for up to half of the total variability in movement execution; *Churchland et al., 2006*) can be naturally incorporated as trial-to-trial variability in the initial conditions $x_{A,B,C}$ for motifs in M1, originating from firing rate variability in M2 attractors upstream and in the M2-to-M1 projections. This model represents a direct circuit implementation of the state–space models of behavior based on Markovian dynamics: discrete states representing actions/motifs correspond to discrete M2 attractors, whereas the continuous latent trajectories underlying movement execution correspond to low-dimensional trajectories of M1 populations driving muscle movements. Moreover, this model can provide a natural and parsimonious

neural implementation for the contextual modulations of behavioral sequences (**Figure 5A**). Two different neural mechanisms can drive this behavioral flexibility: either a variation in the depth of the M2 attractor potential wells mediated by gain modulations (**Figure 4B**) or a change in the noise amplitude originating at thalamus-to-M2 projections. These contextual modulations induce changes in transition times between consecutive actions/attractors. As a consequence, the temporal dynamics are only Markovian when conditioned on each context and the across-context dynamics exhibits non-Markovian features such as context-specific transition rates. This class of models can serve as a useful testing ground for generating mechanistic hypotheses and guide future experimental design.

Flexible sequences in birdsongs

A promising model system to investigate the neural mechanisms of flexible motor sequences is songbirds. Canary songs consist of repeated syllables called phrases, and the ordering of these phrases follows long-range rules in which the choice of what to sing depends on the song structure many seconds prior (**Markowitz et al., 2013**). These long-range contextual modulations are correlated to the activity of neurons in the high vocal center (HVC, a premotor area) (**Cohen et al., 2020**). Seminal work has established that motor variability can be actively generated and controlled by the brain for the purpose of learning (**Olveczky et al., 2005; Kao et al., 2005**). Variability in song production is not simply due to intrinsic noise in motor pathways but is introduced into robust nucleus of the arcopallium (RA, a motor cortex analog) by a dedicated upstream area lateral magnocellular nucleus of the anterior nidopallium (LMAN, analog to a premotor area), which is required for song learning. Theoretical modeling further revealed that this motor variability could rely on topographically organized projections from LMAN to RA for amplifying correlated neural fluctuations (**Darshan et al., 2017**). Strikingly, this mechanism provides universal predictions for the statistics of babbling shared by songbirds and human infants. Timing variability may allow animals to explore the temporal aspects of a given sequence of behavior independently of the choices of actions. Animals could learn proper timing of an action sequence by a search in timing space independent of action selection and vice versa. Future work should explore the advantages of temporal variability in driving learning of precise timing.

Contextual modulations of behavior may occur on longer timescales as well, including those induced by circadian rhythms, hormones, and other changes in internal states; and they may span the whole lifetime of an individual, such as homeostasis and development. Social behavior provides another source of complex contextual modulations. We briefly address these different aspects of contextual modulations below.

Variability on longer timescales

Internal states, such as hunger (**Corrales-Carvajal et al., 2016**), and circadian rhythms (**Patke et al., 2020**) induce daily modulations of several aspects of behavior. In the fruit fly, different clusters of clock neurons are implicated in regulating rhythmic behaviors, including wake–sleep cycles, locomotion, feeding, mating, courtship, and metabolism. Activation of circadian clock neurons in different phases of the cycle drives expression of specific behavioral sequences, and targeted manipulations of particular clusters of clock neurons are sufficient to recapitulate those sequences artificially (**Dissel et al., 2014; Yao and Shafer, 2014**). In mammals, the circadian pacemaker is located in the hypothalamus, interacting with a complex network of neuronal peripheral signals downstream of it (**Hastings et al., 2018; Schibler et al., 2015**). Other sources of contextual variability include variations in the levels of hormones and neuromodulators, which modulate behavioral sequences on longer timescales (**Nelson, 2005**). Although hormones and neuromodulators may not directly drive expression of specific behaviors, they typically prime an animal to elicit hormone-specific responses to particular stimuli in the appropriate behavioral context, as observed prominently during social behavior such as aggression and mating (**Schibler et al., 2015**). Statistical models based on Markovian dynamics, such as state–space models, whose building blocks are sub-second movements syllables, are challenged by contextual modulations occurring on very long timescales. Introducing by hand a hierarchical structure in these models might mitigate this issue (**Tao et al., 2019**). However, approaches based on discrete states, such as HMMs, are not well suited to capture long-term modulation of continuous nature as in the case of circadian rhythms. Fitting a discrete model to slow–continuous dynamics would lead to the proliferation of a large number of discrete states tiling the continuous trajectory (**Alba et al., 2020; Recanatesi et al., 2022**). Data-driven models may provide a more promising way

to reveal slow behavioral modulations, although they have not been tested in this realm yet (*Marshall et al., 2021*). Alternative methods based on continuous latents, such as Gaussian processes (*Yu et al., 2008*), might also provide useful alternative approaches. In trial-based experiments, tensor component analysis was shown to capture slow drifts in latent variables as well (*Williams et al., 2018*). Although the behavioral effects of these long-term sources of contextual modulations have not been included in current models of neural circuits, the theoretical framework based on attractor networks in **Figure 5C** could be augmented to account for them. An afferent higher-order cortical area, recurrently connected to M2, may toggle between different behavioral sequences and control long-term variations in their expression, for example, via gain modulation. A natural candidate for this controller area is the medial prefrontal cortex, which is necessary to express long-term biases in the waiting task (**Figure 3**; *Murakami et al., 2017*). In this augmented model, a top-down modulatory input to the higher-order controller, representing afferent inputs from circadian clock neurons or neuromodulation, may modulate the expression of different behavioral sequences implementing these long-term modulations.

Aging

Other sources of contextual modulation of behavior may act on even longer timescales spanning the whole lifetime of an individual (*Churgin et al., 2017*), such as homeostasis and development. The interplay between these two mechanisms may explain how the motor output of some neural circuits maintains remarkable stability in the face of the large variability in neural activity observed across a population or in the same individual during development (*Prinz et al., 2004*; *Bucher et al., 2005*). Across the entire life span, different brain areas develop, mature, and decline at different moments and to different degrees (*Sowell et al., 2004*; *Yeatman et al., 2014*). Connectivity between areas develops at variable rates as well (*Yeatman et al., 2014*). It would be interesting to investigate whether these different mechanisms, unfolding over the life span of an individual, can be accounted for within the framework of multiarea attractor networks presented above.

Social behavior and contextual modulations

Although most of the experiments discussed in this review entail the behavior of individual animals, contextual modulations of behavior are prominently observed in naturalistic assays comprising the interaction between pairs of animals, such as hunting and social behaviors. In the prey–capture paradigm, a mouse pursues, captures, and consumes live insect prey (*Hoy et al., 2016*; *Michaël et al., 2020*). Prey–capture behavior was found to strongly depend on context. Experimental control of the surrounding environment revealed that mice rely on vision for efficient prey–capture. In the dark, the hunting behavior is severely impaired: only at close range to the insect is the mouse able to navigate via auditory cues. Another remarkable example of context-dependent social behavior was demonstrated during male–female fruit fly mating behavior (*Calhoun et al., 2019*). During courtship bouts, male flies modulate their songs using specific feedback cues from their female partner such as their relative position and orientation. A simple way to model the relationship between sensory cues and the choice of a specific song in terms of linear ‘filters’ (*Coen et al., 2014*), where a common assumption is that the sensorimotor map is fixed. Relaxing this assumption and allowing for moment-to-moment transition between more than one sensorimotor map via a hidden Markov model with generalized linear model emissions (GLM-HMM), *Calhoun et al., 2019* uncovered latent states underlying the mating behavior corresponding to different sensorimotor strategies, each strategy featuring a specific mapping from feedback cues to song modes. Combining this insight with targeted optogenetic manipulation revealed that neurons previously thought to be command neurons for song production are instead responsible for controlling the switch between different internal states, thus regulating the courtship strategies. Finally, a tenet of naturalistic behavior is vocal communication, which combines aspects of sensory processing and motor generation in the realm of complex social interactions. A particularly exciting model system is the marmoset, where new techniques to record neural activity in freely moving animals during social behavior and vocalization (*Nummela et al., 2017*) together with newly developed optogenetic tools (*MacDougall et al., 2016*) and multi-animal pose tracking algorithms (*Pereira et al., 2020b*; *Lauer et al., 2021*) hold the promise to push the field into entirely new domains (*Eliades and Miller, 2017*).

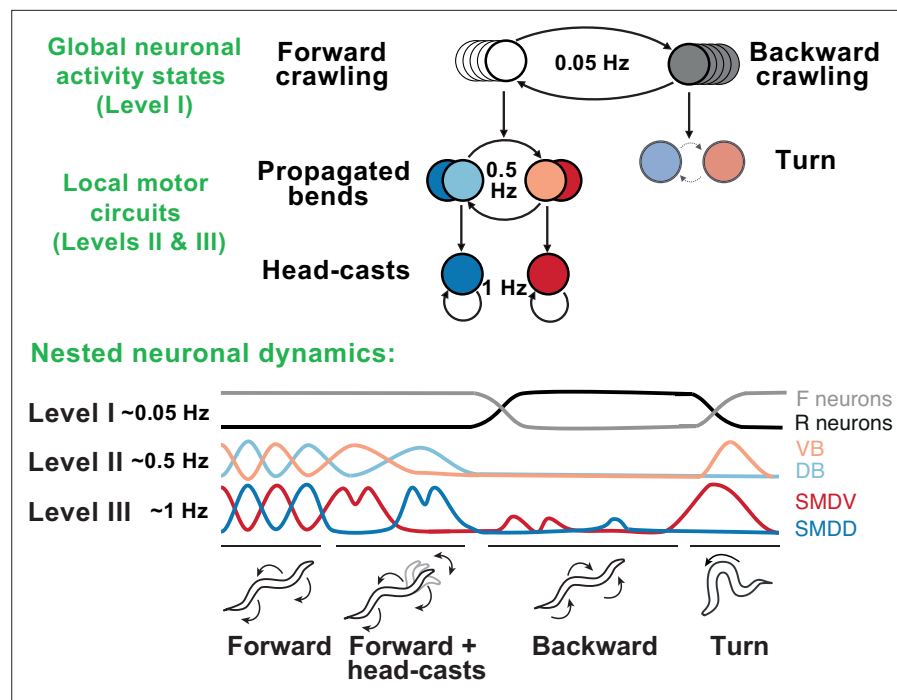


Figure 6. Behavioral and neural hierarchies in *C. elegans*. Top: behavioral hierarchy in *C. elegans*. A 0.05Hz cycle drives switches between forward and reverse crawling states, with intermediate level 0.5Hz crawling undulations, and lower level 1Hz head-casts. Bottom: slow dynamics across whole-brain circuits reflect upper-hierarchy motor activity; fast dynamics in motor circuits drive lower-hierarchy movements. Slower dynamics tightly constrain the state and function of faster ones. Adapted from Figure 8 of *Kaplan et al., 2020*.

Hierarchical structure

The temporal organization of naturalistic behavior exhibits a clear hierarchical structure (*Tinbergen, 1951; Dawkins, 1976; Simon, 1962*), where actions are nested into behavioral sequences that are then grouped into activities. Higher levels in the hierarchy emerge at longer timescales: actions/movements occur on a sub-second scale, behavioral sequences span at most a few seconds and activities last for longer periods of several seconds to minutes. The crucial aspect of this behavioral hierarchy is its complexity: an animal's behavior unfolds along all timescales simultaneously. What is the organization of this vast spatiotemporal hierarchy? What are the neural mechanisms supporting and generating this nested temporal structure?

A case study: *C. elegans* locomotor behavior

Remarkable progress in both behavioral and neurophysiological aspects of the hierarchy has been made in the worm *C. elegans* (*Figure 6*). The worm locomotor behavior exhibits a clear hierarchical organization with three distinct timescales: sub-second body-bends and head-casts, second-long crawling undulations, and slow reverse-forward cycles (*Gomez-Marin et al., 2016*). At the neuronal level, this hierarchy is generated by nested neuronal dynamics where upper-level motor programs are supported by slow activity spread across many neurons, while lower-level behaviors are represented by fast local dynamics in small multifunctional populations (*Kaplan et al., 2020*). Persistent activity driving higher-level behaviors gates faster activity driving lower-level behaviors, such that, at lower levels, neurons show dynamics spanning multiple timescales simultaneously (*Hallinen et al., 2021*). Specific lower-level behaviors may only be accessed via switches at upper levels, generating a nonoverlapping, tree-like hierarchy, in which no lower-level state is connected to multiple upper-level states.

At the top of this motor hierarchy, we find a much longer-lasting organization of states in terms of exploration, exploitation, and quiescence. In contrast to the strict, tree-like structure observed in the motor hierarchy, lower-level motor features are shared across these states, albeit with different

frequency of occurrence. Whereas the motor hierarchy is directly generated by neuronal activity, this state-level hierarchy may rely on neuromodulation (Ben Arous et al., 2009; Flavell et al., 2013)

Behavioral hierarchies in flies and rodents

How much of this tight correspondence between behavioral and neural hierarchies generalizes from worms to insects and mammals? Are action sequences organized as a chain? Alternatively, is there a hierarchical structure where individual actions, intermediate subsequences, and overall sequences can be flexibly combined? A chain-like organization would require a single controller concatenating actions, but could be prone to error or disruption. A hierarchical structure could be error-tolerant and flexible at the cost of requiring controllers at different levels of the hierarchy (Geddes et al., 2018). We will first examine compelling evidence from the fruit fly and then discuss emerging results in rodent experiments.

By analyzing videos of ground-based fruit fly during long sessions of spontaneous behavior (Figure 2B), Berman et al., 2014 and Berman et al., 2016 mapped the entire behavioral space of

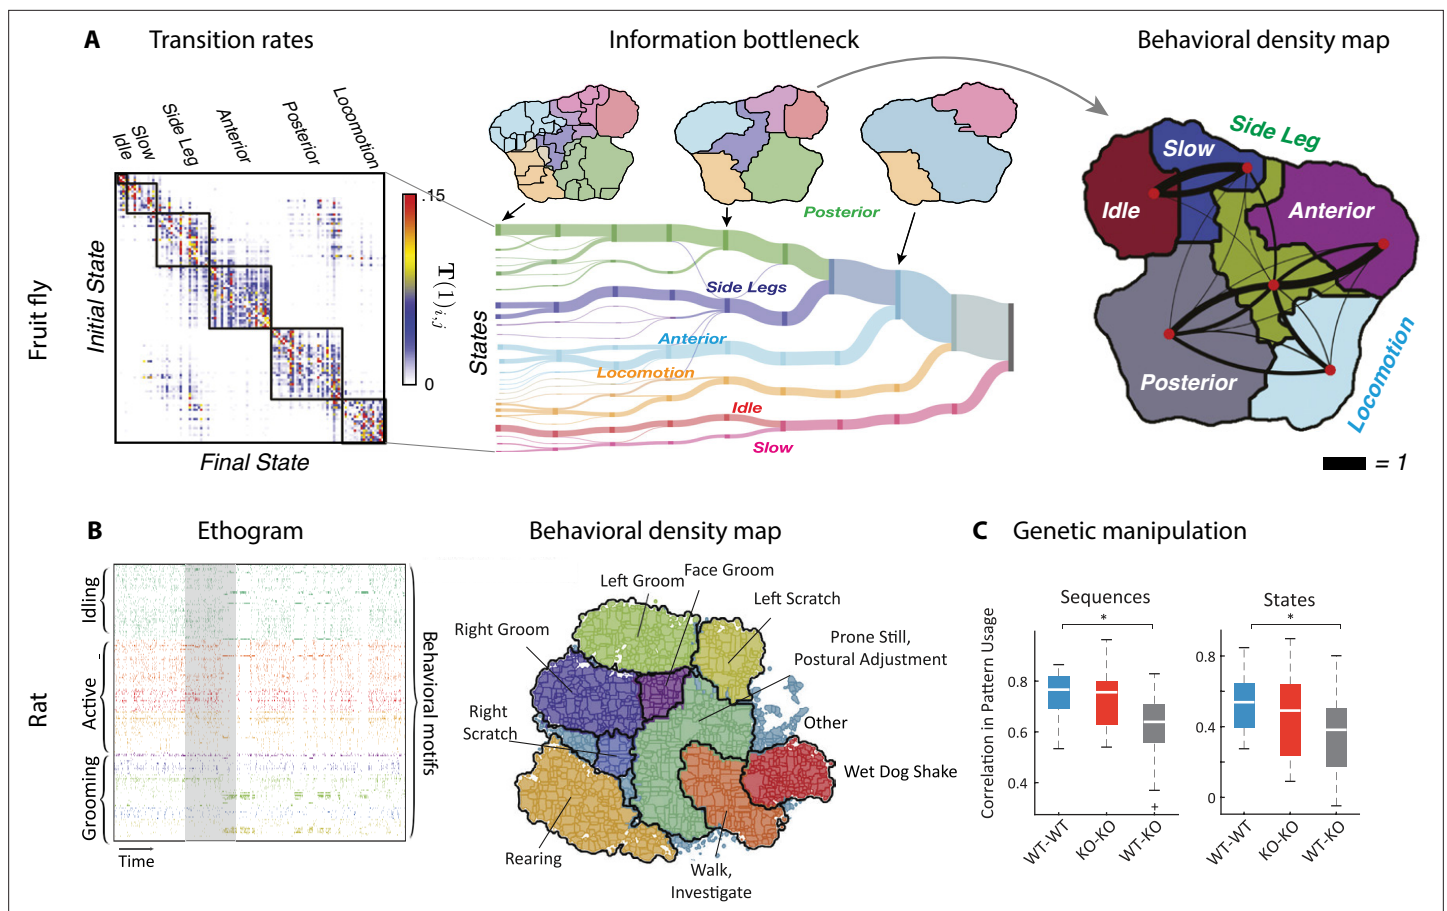


Figure 7. Hierarchical structure in freely moving fly and mouse behavior. (A) Hierarchical variability in the fruit fly behavior. The Markov transition probability matrix (left) between postures reveals a clustered structure upon applying the predictive information bottleneck with six clusters (black outline on the left plot). Labels (e.g., ‘anterior,’ ‘side leg’) refer to movements involving specific body parts. Center: hierarchical organization for optimal solutions of the information bottleneck for predicting behavior on increasingly slower timescales (varying clusters from 25 to 1, left to right; colored vertical bars are proportional to the percentage of time a fly spends in each cluster). Right: behavioral clusters are contiguous in behavioral space (same clusters as in the transition matrix in the left panel; black lines represent transitions probabilities between states, the thicker the more likely). (B) Left: a temporal pattern matching algorithm detected repeated behavioral patterns (rows, sorted in grooming, active, idling classes; color-coded clusters explained in panel B) in freely moving rats behavioral recordings. Right: data from 16 rats co-embedded in a two-dimensional t-distributed stochastic neighbor embedding (t-SNE) behavioral map was clustered with a watershed transform, revealing behavioral clusters segregated to different regions of the map. The color-coded ethogram on the left is annotated from these behavioral clusters. (C) Sequence and state usage probabilities for wild-type and Fmr1-KO rats show a significantly decreased correlation between different genotypes. Panel (A) adapted from Figures 1 and 5 of Berman et al., 2016. Panels (B) and (C) adapted from Figures 3, 4, and 6 of Marshall et al., 2021.

freely moving flies, identifying a hundred stereotyped, frequently reoccurring actions, interspersed with bouts of nonstereotyped behaviors. Recurring behavioral categories emerged as peaks in the behavioral space probability landscape, labeled as walking, running, head grooming, and wing grooming (**Figure 7A**). In order to uncover the organization of behavior at different timescales, the authors estimated which behavioral representations (movements, sequences, or activities) could optimally predict the fly's future behavior on different temporal horizons (from 50 ms to minutes), applying the information bottleneck method (**Tishby et al., 2000**). This predictive algorithm revealed multiple timescales in the fly behavior, organized into a hierarchical structure reflecting the underlying behavioral programs and internal states. The near future could be optimally predicted by segmenting behavior according to actions at the fastest level of the hierarchy. The optimal representation of behavior that could optimally predict the distant future up to minutes away was based on slower, coarser groups of actions grouped into activities. These longer timescales manifest as nested blocks in the transition probability matrix, implying that a strongly non-Markovian structure emerges on longer timescales (**Alba et al., 2020; Figure 7A**).

A hallmark of the fly behavior emerging from this analysis is that the different branches of this hierarchy are segregated into a tree-like, nonoverlapping structure: for example, actions occurring during grooming do not occur during locomotion. This multiscale representation was leveraged to dissect the descending motor pathways in the fly. Optogenetic activation of single neurons during spontaneous behavior revealed that most of the descending neurons drove stereotyped behaviors that were shared by multiple neurons and often depended on the behavioral state prior to activation (**Cande et al., 2018**). An alternative statistical approach based on a hierarchical HMM revealed that, although all flies use the same set of low-level locomotor features, individual flies vary considerably in the high-level temporal structure of locomotion, and in how this usage is modulated by different odors (**Tao et al., 2019**). This behavioral idiosyncrasy of individual-to-individual phenotypic variability has been traced back to specific genes (**Ayroles et al., 2014**) regulating neural activity (**Buchanan et al., 2015**) in the central complex of the fly brain. This series of studies suggest that the fly behavior is organized hierarchically in a tree-like structure from actions to sequences to activities. This hierarchy seems to lack flexibility, such that behavioral units at lower levels are segregated into different hierarchical branches and cannot be rearranged at higher levels. Although a more flexible structure seems to emerge during social interactions (see above and **Calhoun et al., 2019**).

What is the structure of the behavioral hierarchy in mammals? We will start by considering the first two levels of the hierarchy, namely, how actions concatenate into sequences. A recent experimental *tour de force* demonstrated the existence of a hierarchical structure in the learning and execution of heterogeneous sequences in a freely moving operant task (**Geddes et al., 2018**). Mice learned to perform a 'penguin dance' consisting of a sequence of two or three consecutive left lever presses (LL or LLL) followed by a sequence of two or more right lever presses (RR or RRR) to obtain a reward. Mice acquired the sequence hierarchically rather than sequentially, a learning scheme that is inconsistent with the classic reinforcement learning paradigm, which predicts sequence learning occurs in the reverse order of execution (**Sutton and Barto, 2018**). Using closed-loop optogenetic stimulation, the authors revealed the differential roles played by striatal direct and indirect pathways in controlling, respectively, the expression of a single action (either L or R), or a fast switch from one subsequence to the next (from the LL block to RR, and from RR to the reward approach).

What is the structure of the rodent behavioral hierarchy at slower timescales? A multiscale analysis of ongoing behavior was performed using rat body piercing, which allowed for tracking of the three-dimensional pose with high accuracy (**Figure 7B; Marshall et al., 2021**). A watershed algorithm applied to behavioral probability density maps revealed the emergence of hierarchical behavioral categories at different temporal scales. Examining the behavioral transition matrix at different timescales, signature of non-Markovian dynamics peaked at 10–100 s timescales. On 15 s timescales, pattern sequences emerged featuring sequentially ordered actions, such as grooming sequences of the face followed by the body. On minute-long timescales, states of varying arousal or task engagement emerged, lacking stereotyped sequential ordering. Although the results of **Marshall et al., 2021** seem to suggest the presence of a flexible, overlapping structure in the rat hierarchy, consistent with **Geddes et al., 2018** and different from the tree-like hierarchy in flies (**Berman et al., 2016**), the analysis of **Marshall et al., 2021** lacked the information bottleneck step, which was crucial to capture the

fly behavior on long timescales. Future work may shed light on the nature of behavioral hierarchies in mammals using the information bottleneck.

Manipulation experiments could help clarify the structure of the rodent behavioral hierarchy and whether it allows for flexible behavioral sequences. Earlier studies of contextual modulations of rodent freely moving behavior by pharmacological or genetic manipulations were mostly confined to comparison of usage statistics of single actions (Wiltschko et al., 2015). However, more recent studies (Marshall et al., 2021; Klibaite et al., 2021), which mapped the rodent behavioral hierarchy, allowed a multiscale assessment of contextual effects. In a rat model of the fragile X syndrome, while mutant and wildtype rats had similar locomotor behavior, the former showed abnormally long grooming epochs, characterized by different behavioral sequences and states compared to wildtypes (Figure 7C). This study highlights the advantages of multiscale comparative taxonomy of naturalistic behavior to investigate behavioral manifestations of complex conditions such as autism spectrum disorder and classify the behavioral effects of different drugs. Another benefit of this multiscale approach is to allow the investigation of the individuality of rodent behavior, revealing that, although all animals drew from a common set of temporal patterns, differences between individual rats were more pronounced at longer timescales compared to short ones (Marshall et al., 2021). It would be interesting to further explore the features of individuality across species as a way of testing ecological ideas such as ‘bet-hedging’ (Honegger and de Bivort, 2018).

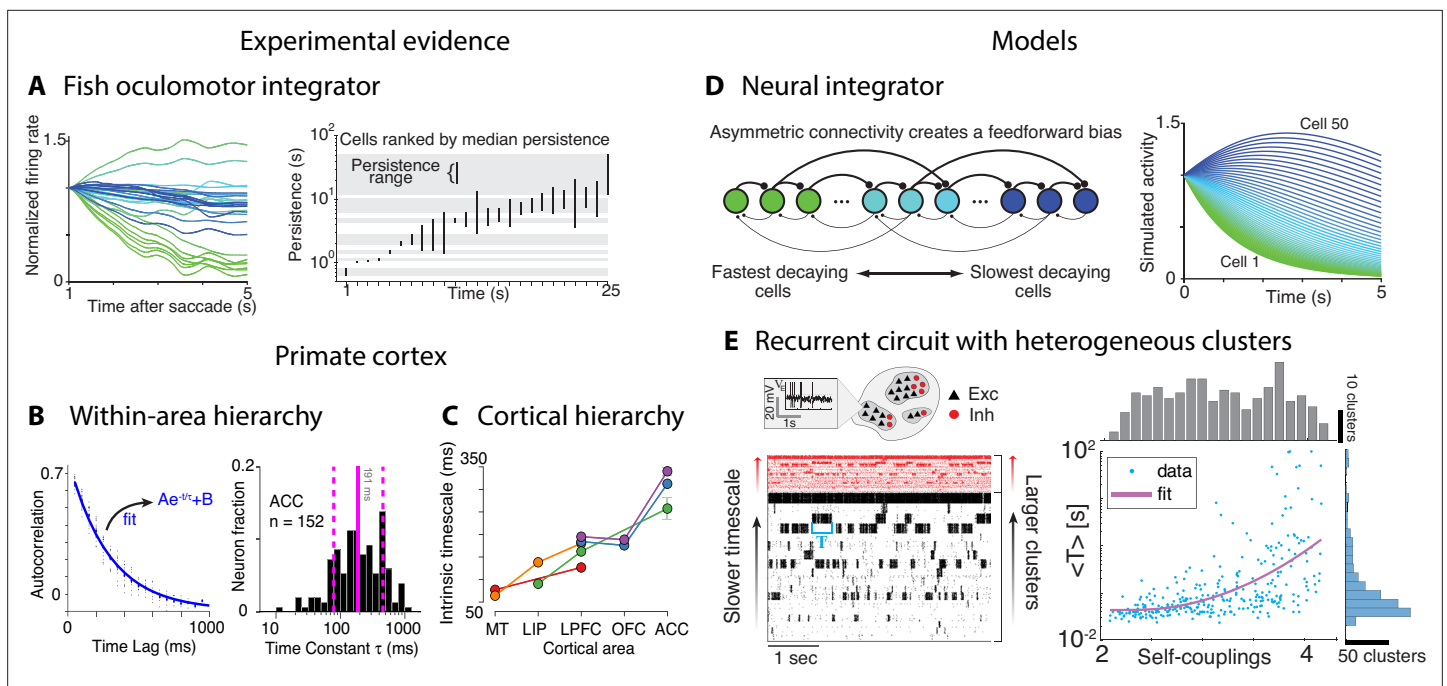


Figure 8. Computational principles underlying the heterogeneity of timescales. (A) Firing rate of neurons imaged in zebrafish larvae (left, colored according to the rostrocaudal and dorsoventral neuron location) reveals the correspondence between persistence time (right, each bar represents the persistence time range for each cell) and location along these dimensions. (B) Left: autocorrelation of an example neuron from orbitofrontal cortex (OFC) in awake monkey (blue: exponential fit). Right: histogram of the time constants reveals a large variability across OFC neurons (solid and dashed vertical lines represent mean and SD). (C) Intrinsic timescales across the visual-prefrontal hierarchy in five datasets estimated as the average population autocorrelation. (D) Persistent activity with heterogeneous timescales in the oculomotor system can be explained by a progressive filtering of activity propagating down a circuit including a mixture of feedback and functionally feedforward interactions, realizing a uniformly detuned line attractor. (E) A heterogeneous distribution of timescales naturally emerges in a recurrent network of excitatory (black) and inhibitory (red) spiking neurons arranged in clusters, generating time-varying activity unfolding via sequences of metastable attractors (left: representative trial, clusters activate and deactivate at random times; neurons are sorted according to cluster membership). Larger clusters (at the top) activate for longer intervals. Right: the distribution of cluster activation timescales $<T>$, proportional to a cluster size, exhibits a large range from 20 ms to 100s. Panel (A) adapted from Figures 4 and 8 of Miri et al., 2011. Panels (C) and (D) adapted from Figure 1 of Murray et al., 2014 and Figure 8 of Miri et al., 2011. Panel (B) adapted from Figure 2 of Cavanagh et al., 2016. Panel (E) adapted from Figure 5 of Stern et al., 2021.

Multiple timescales of neural activity

What are the neural mechanisms generating the hierarchy of timescales observed during naturalistic behavior? Is there evidence that neural activity is simultaneously varying over multiple timescales? Although no studies directly addressed these questions yet, a number of experimental and theoretical approaches have provided evidence for multiple timescales of neural activity. Some evidence for temporal heterogeneity in neural activity was reported in restrained animals during stereotyped behavioral assays. A heterogeneous distribution of timescales of neural activity was found in the brainstem of the restrained larval zebrafish by measuring the decay time constant of persistent firing across a population of neurons comprising the oculomotor velocity-to-position neural integrator (**Figure 8A**; **Miri et al., 2011**). The decay times varied over a vast range 0.5–50 s across cells in individual larvae. This heterogeneous distribution of timescales was later confirmed in the primate oculomotor brainstem (**Joshua et al., 2013**), suggesting that timescale heterogeneity is a common feature of oculomotor integrators conserved across phylogeny. Single-neuron activity may also encode a long memory of task-related variables. In head-fixed monkeys performing a competitive game task, temporal traces of the reward delivered in previous trials were encoded in single-neuron spiking activity in frontal areas over a wide range of timescales, obeying a power law distribution up to 10 consecutive trials (**Bernacchia et al., 2011**).

The timescale of intrinsic fluctuations in spiking activity can be also estimated from single-neuron spike autocorrelation functions (**Figure 8B**). In awake head-fixed primates, during periods of ongoing activity, the distribution of intrinsic autocorrelation timescales within the same cortical circuit was found to be right-skewed and approximately lognormal (**Figure 8B**), in the range from 10 ms to 1 s (**Cavanagh et al., 2016**). Moreover, comparison of the population-averaged autocorrelations during ongoing periods revealed a hierarchical structure across cortical areas, varying from 50 to 350 ms along the occipital-to-frontal axis (**Murray et al., 2014**; **Figure 8C**).

While all these results were obtained in restrained animals, it is an open question whether neural activity during naturalistic behavior exhibits temporal hierarchies similar to those observed in behavior. Current evidence from freely moving rodents engaged in waiting tasks (**Murakami et al., 2017**; **Schreiner et al., 2021**; **Recanatesi et al., 2022**) revealed the presence of multiple timescales in single-neuron activity from secondary motor and prefrontal cortices. These timescales range from the sub-second scale (single attractors), to a few seconds (attractor sequences), to tens of seconds or minutes (trial history dependence). It is tantalizing to speculate that even longer timescales may be present for neural activity to be able to generate the vast hierarchy of timescales observed in naturalistic behavior (**Figure 7**). Evidence from associative learning tasks in rodents found that population activity in hippocampal CA1 and CA3 encodes multiplexed information about several aspects of the task occurring on multiple trials such as context, place, value, and objects (**McKenzie et al., 2014**). Although an explicit analysis of temporal correlations was not carried out in this study, these results suggest that a hierarchy of timescales may be present in the hippocampus and emerge during learning.

Computational mechanisms underlying heterogeneous distributions of timescales

What are the computational mechanisms underlying this hierarchical temporal structure featuring a wide distribution of timescales? Single-cell biophysical properties such as differences in membrane time constant across cell types or synaptic time constants across receptors (e.g., faster AMPA and GABA vs. slower NMDA receptors) may generate multiple timescales in sub-second range, from a few milliseconds to a few hundred milliseconds (**Gjorgjieva et al., 2016**). Although this class of single-cell mechanisms might explain the temporal hierarchy observed in the posterior–anterior cortical axis (**Murray et al., 2014**) by relying on systematic differences in cell-type-specific features, they are unlikely to explain the much wider temporal hierarchy observed across neurons within the same area (**Figure 8A and B**), or the even larger hierarchy observed in behavior (**Figure 7**).

Theoretical studies highlighted the central role played by recurrent synaptic couplings within a local circuit for generating long timescales emerging from the recurrent dynamics. Recurrent networks with random recurrent couplings can generate time-varying neural activity whose timescale may be tuned to very slow values at a critical point (**Toyoizumi and Abbott, 2011**; **Magnasco et al., 2009**). Long timescales may also emerge in random neural networks with an excess of symmetric couplings

(Martí et al., 2018). In these examples, all cells share the same timescale. Although this common timescale can be tuned to arbitrarily long values, these circuits are incapable of giving rise to temporal heterogeneity either at the single cell or at the population level.

Two biologically plausible ingredients were shown to be sufficient to generate long timescales with temporal heterogeneity in their distributions: recurrent couplings realizing local functional neural clusters and heterogeneity in synaptic couplings. The first requirement of functional assemblies is a connectivity motif ubiquitously observed in biological circuits, namely, the fact that the strength of recurrent couplings decays with spatial distance between pairs of neurons, leading to the emergence of local functional assemblies of strongly coupled neurons (Song et al., 2005; Perin et al., 2011; Lee et al., 2016; Kiani et al., 2015; Miller et al., 2014). In recurrent linear network models, functional assemblies can generate neural activity exhibiting slow relaxation times following stimulation. In order to generate a heterogeneity in the distribution of relaxation times, functional assemblies can be coupled by heterogeneous long-range connections, arranged along a spatial feedforward gradient (Figure 8D; Miri et al., 2011; Joshua et al., 2013; Chaudhuri et al., 2014). Linear networks featuring both local functional assemblies and spatial heterogeneity in long-range couplings were able to reproduce the heterogeneous decay times found in the brainstem oculomotor integrator circuit in zebrafish and primates (Miri et al., 2011; Joshua et al., 2013) and in the reward integration times in the macaque cortex (Bernacchia et al., 2011). It remains an open question whether these integrator models could be generalized to explain hierarchical activity in motor generation, perhaps stacking multiple layers of them. A shortcoming of these models is the fact that they require fine-tuning of synaptic couplings and do not generate attractor dynamics, which recent experimental evidence suggests is the basic building block of preparatory motor activity (Inagaki et al., 2019; Recanatesi et al., 2022). Other network models generating multiple timescales of activity fluctuations were proposed based on self-tuned criticality with anti-Hebbian plasticity (Magnasco et al., 2009) or multiple block-structured connectivity (Aljadeff et al., 2015).

An alternative robust and biologically plausible way to generate a vast hierarchy of timescales was proposed based on the ingredients of recurrent functional clusters and heterogeneity in synaptic couplings (Figure 8E; Stern et al., 2021), two common features observed across cortical circuits (Marshall et al., 2019; Kiani et al., 2015; Miller et al., 2014; Song et al., 2005; Perin et al., 2011; Lee et al., 2016), supported by theoretical evidence (Litwin-Kumar and Doiron, 2012; Wyrick and Mazzucato, 2021). In this model, excitatory and inhibitory neurons are arranged in clusters of heterogeneous sizes, generating metastable activity whose typical timescale is measured by the on-off cluster switching time (see Appendix 3). In this model, a cluster's timescale is proportional to its size and larger clusters exhibit longer timescales, yielding a heterogeneous distribution of timescales in the range observed in cortex (Figure 8A–C). An appealing feature of this model is that it could be generalized to other domains that exhibit fluctuations simultaneously varying over a large range of timescales, such as spin glasses (Bouchaud, 1992), metabolic networks of *Escherichia coli* (Almaas et al., 2004; Emmerling et al., 2002), and yeast cultures (Roussel and Lloyd, 2007; Aon et al., 2008). It remains an open question whether the relationship between a neural cluster's size and its intrinsic timescale is realized in cortical circuits and whether it can explain the origin of the hierarchical variability in naturalistic behavior.

Future directions: Neural mechanisms generating temporal hierarchies

Future studies should address how the hierarchy of timescales found in behavior may emerge from computational mechanisms. There are several missing links along this path. First, although temporally heterogeneous neural activity was found in several brain areas and species under different experimental conditions, no study to date investigated the presence of temporal hierarchies in neural activity during naturalistic behavior. All the necessary tools are available: recent advances in neurotechnology demonstrate the feasibility of chronic recordings of large neural populations during freely moving behavior (Juavinett et al., 2019; van Daal et al., 2021); simultaneously performing behavioral classification analyses from pose tracking software (Mathis et al., 2018; Marshall et al., 2021; Pereira et al., 2019). This first piece of the puzzle is thus within reach.

On the theory side, although recurrent dynamics can generate heterogeneous distributions of timescales (Figure 8E), this model needs to be extended to a fully realistic framework for explaining nested temporal hierarchies in naturalistic behavior such as the proposed mesoscopic circuit in

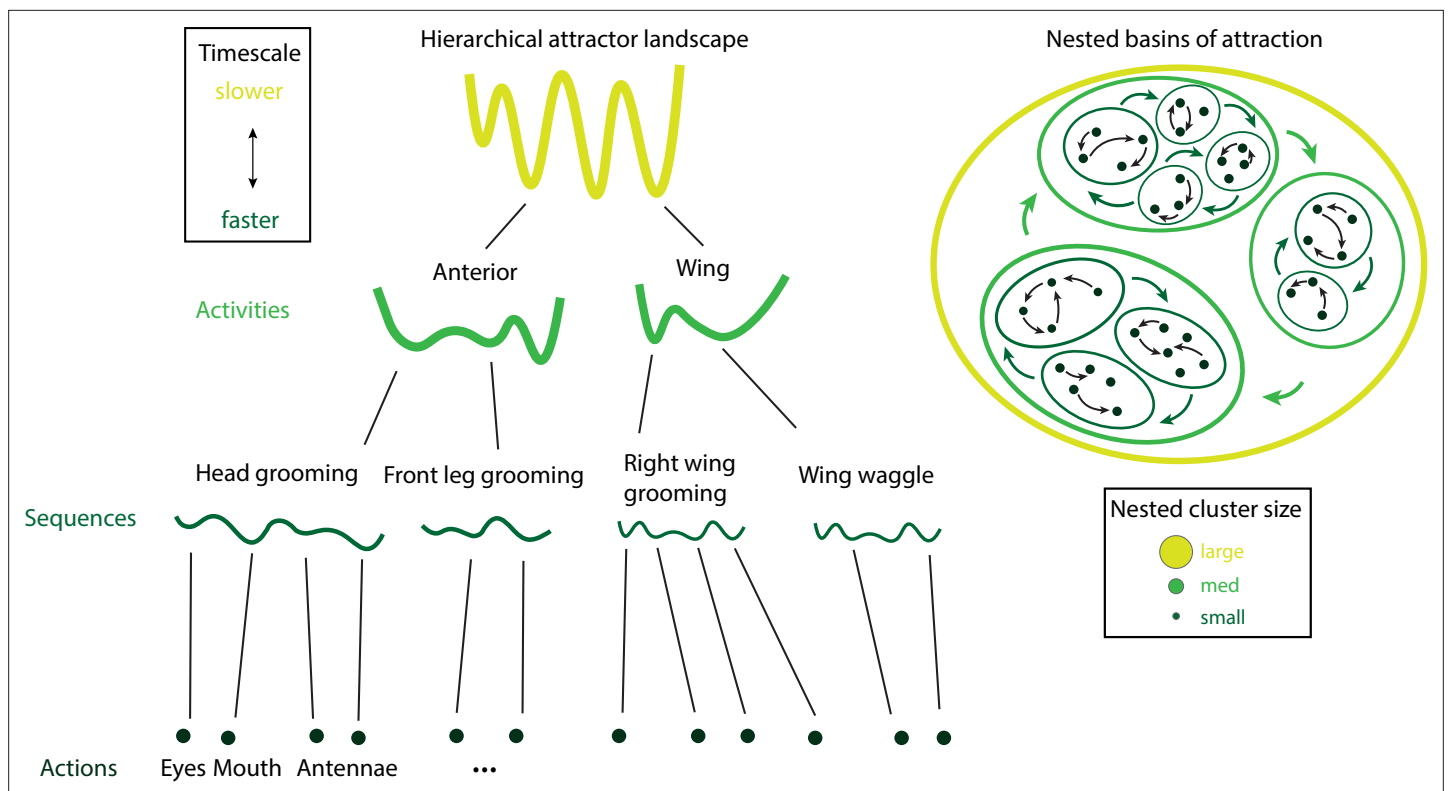


Figure 9. Hierarchical attractor networks can explain a tree-like behavioral hierarchy. A computational framework to explain the hierarchical structure of behavior was discovered in the fruit fly and rodents (**Figure 7**). Attractors encoding for actions, sequences, and activities (left) are supported by nested basins of attractions (right), where the size of a basin determines the intrinsic timescale of the corresponding activity. Actions have small basins, corresponding to fast timescales, while activities have large basins, corresponding to slow timescales. Attractors are nonoverlapping, consistent with the tree-like structure of the behavioral hierarchy.

Figure 5C. Here, we propose a simple roadmap to bridge this gap, building on some of the theoretical ideas we reviewed above. This theory is based on a hierarchical structure that, for lack of a better term, we will denote ‘*Attractors all the way down*’ (**Figure 9**). We start from the observation that self-initiated actions within a behavioral sequence are represented as metastable attractors in secondary motor cortex (*Inagaki et al., 2019; Recanatesi et al., 2022; Figure 3*). We then extend this observation to a general principle positing that behavioral units at each level of the temporal hierarchy (not only actions, but also sequences, activities) are represented by metastable attractors. The stochasticity in behavioral unit duration can be achieved by generating transitions between attractors at each level via low-dimensional variability, arising from mesoscopic feedback loops involving cortex and subcortical areas (*Recanatesi et al., 2022; Figure 3D*). The average transition time between behavioral units at each level will then depend on the barrier height separating the corresponding attractors (*Cao et al., 2016; Mazzucato et al., 2019; Wyrick and Mazzucato, 2021*). As one moves up the hierarchy from actions (fast transitions = lower barriers) to sequences and activities (slow transitions = higher barriers), the potential wells separating the corresponding attractors become deeper and the basins of attraction wider. This increase in timescales can be achieved in a biological manner by assuming that increasingly larger neural populations encode for slower behavioral features (*Stern et al., 2021; see Figure 8E*), which can be realized in a biologically plausible way as a hierarchy of clusters within clusters (*Schaub et al., 2015*). Contextual modulations in average duration of a behavioral unit at each level of the hierarchy can be implemented by top-down changes of the barrier heights between attractors via gain modulation (**Figure 4C**). The new theoretical ingredient required for this theory to work is the presence of a nested structure in the attractor landscape, whereas the basin of attraction of lower-level attractors (actions) is contained within the basins of higher level attractors (sequences) all the way up to the largest basins representing long-term activities. The architecture of this model provides some immediate predictions for neural activity. First, neurons exhibit conjunctive selectivity

to multiple variables at different levels of the hierarchy, which has been observed, for example, in rodent hippocampus during freely moving tasks ([McKenzie et al., 2014](#)) and in primate prefrontal cortex during complex decision-making tasks ([Rigotti et al., 2013](#)). Second, single-cell activity is modulated by many attractors, in agreement with the multistable activity observed in sensory ([Mazzucato et al., 2015](#)) and frontal areas ([Recanatesi et al., 2022](#)) where neuronal representations of attractors are dense rather than sparse. Third, fluctuations in the activity of single neurons should vary over multiple timescales *simultaneously*, encoding information about multiple levels of the hierarchy; this is the structure that was observed in *C. elegans* where neurons representing movements exhibit both fast and for slow fluctuations correlated to multiple levels of the behavioral hierarchy ([Kaplan et al., 2020; Figure 6](#)).

Which network architectures can implement a hierarchical attractor landscape? The behavioral hierarchy discovered in flies ([Berman et al., 2016](#)) can be approximated as a tree: behavioral units at a lower level of the hierarchy only occur during a specific unit at a higher level of the hierarchy ([Figure 7](#)). For example, movements occurring during the fruit fly locomotion activities never occur during idle activities. This structure is typical of phylogenetic classification in taxonomy, where the distance from the root to the leaves is the same for every leaf and the tree is called ultrametric ([Rammal et al., 1986](#)). It is not known whether behavioral trees are ultrametric; in order to clarify this structure, a notion of distance in behavioral space will have to be introduced and examined. A classic model of hierarchical attractors, which could potentially capture the tree-based hierarchy, is an extension of the Hopfield network where the stored pattern is hierarchically arranged into a tree ([Parga and Virasoro, 1986; Figure 9](#)). This class of models was originally proposed to explain the effect of word semantic categorization during memory encoding, storage, and retrieval. A contamination between the domains of naturalistic behavior and natural language processing has been proposed early on ([Dawkins, 1976](#)) and recently applied to describe *C. elegans* ([Gomez-Marin et al., 2016; Gupta and Gomez-Marin, 2019](#)) and larval zebrafish behavior ([Reddy et al., 2022](#)). It is still an open question to determine the precise structure of the behavioral hierarchy in mammals. The results of [Marshall et al., 2021](#) seem to suggest that the hierarchy in rats may not be tree-like as the one observed in flies ([Berman et al., 2016](#)). However, the analysis of the rodent behavior in [Marshall et al., 2021](#) lacked the information bottleneck step, which was crucial to infer the structure of the fly hierarchy at long timescale. On general grounds, it would be interesting to investigate whether the mammalian hierarchical motor system allows for flexible and adaptable behaviors, by utilizing the same motor elements across different sequences in a hierarchical manner, as represented in the mesoscopic circuit models of [Murray and Escola, 2017; Logiaco and Escola, 2020; Logiaco et al., 2021](#) (see [Figure 5](#)) and suggested by experimental evidence in [Geddes et al., 2018](#). This flexible hierarchy could be realized by overlapping branches and would be different from the tree-like structure found in the fly.

Are the computational mechanisms generating the tree-like hierarchy sufficient to generate flexible hierarchies? Recent theoretical work highlighted the importance of a neural circuit architecture that segregates the motor preparation and execution in different areas. In the cortico–thalamic–basal ganglia model of [Figure 5B](#), flexible sequences can be generated by rearranging existing motifs as well as by learning new motifs without interfering with previous ones ([Logiaco et al., 2021](#)). Moreover, flexible sequences can be learned using biologically plausible Hebbian plasticity in the striatum ([Murray and Escola, 2017](#)). These theoretical models provide an exciting blueprint for further investigation into learning and expression of flexible behavioral hierarchies.

What are potential neural implementations of this network architecture? In the case of mammals, at the lowest level of the hierarchy, neural populations in the rodent secondary motor cortex (M2) generate metastable attractors representing upcoming actions ([Inagaki et al., 2019; Recanatesi et al., 2022](#)). Other subcortical areas, such as the basal ganglia or the thalamus, are likely involved in the generation of complex behavioral sequences ([Murray and Escola, 2017; Logiaco et al., 2021; Recanatesi et al., 2022](#)). How are higher-order behavioral units, such as sequences and activities, encoded? One possibility is that they could also be encoded in M2 attractors as recent evidence suggests M2 populations encode for a wide range of behavioral variables ([Cazettes et al., 2021](#)). We hypothesize an architecture featuring nested assemblies where larger populations, whose activity varies over progressively slower timescales ([Stern et al., 2021; Schaub et al., 2015](#)), hierarchically encode for slower features of behavior (from actions to sequences to activities). Experimental evidence shows that single neurons in M2 are selective to multiple actions ([Recanatesi et al., 2022](#)),

suggesting that mixed selectivity (*Rigotti et al., 2013*) could play an important role in generating the high-dimensional hierarchical attractor landscape necessary to capture the complexity of naturalistic behavior. An alternative architecture might involve a distributed cortical circuit where the neural representations of behavioral units at different levels of the hierarchy are encoded in multiple frontal cortical areas, such as action sequences in M2, and activities in prefrontal cortex, which is known to control behavior on longer timescales such as trial history effects (*Murakami et al., 2014; Murakami et al., 2017; Schreiner et al., 2021*). In both scenarios, we hypothesize that while behavioral units are encoded as cortical attractors, transitions between attractor rely on feedback loops involving cortico–subcortical circuits (*Murray and Escola, 2017; Logiaco et al., 2021; Recanatesi et al., 2022*).

Behavioral energy landscape

The structure of naturalistic behavior emerging on long timescales in the fruit fly is a hierarchical tree-like organization (*Figure 7A*), which is consistent with an underlying neural circuit architecture based on hierarchical basins of attractions (*Figure 9*). Can we derive a representation of behavior as a complex energy landscape directly from the behavioral data itself? A promising approach is given by the probability density map approach (*Berman et al., 2016*). One could define the log-probability of the density map as an energy potential, where the transition rates between behavioral features yield a probability flux along this potential landscape. It is tantalizing to speculate that using the combination of energy gradient and probability flux one could derive a data-driven nonlinear dynamical system describing behavior (*Wang, 2015*). Other approaches to infer a potential energy landscape directly from data have been successfully applied to spike trains (*Genkin and Engel, 2020; Duncker et al., 2019*), although in a regime where the energy landscape has only a few minima. Alternatively, in chemical kinetics or fluid flow, methods based on the transfer operator formalism have been applied to find effective free energy landscapes and metastable states from experimental data or simulations (*Bowman et al., 2013*), and have been recently applied to animal behavior (*Costa et al., 2021*).

Neural mechanisms of flexible behavioral hierarchies

The model of nested basins of attractions proposed in *Figure 9* can explain a tree-like behavioral hierarchy, where behavioral units at lower levels (e.g., actions) are not shared by different units at a higher level (e.g., sequences), as observed in fruit flies (*Berman et al., 2016*). Although the tree-like hierarchical organization emerged from assays where individual animals were monitored in isolation, recent studies of social behavior seem to challenge this structure. A large variety of qualitatively new behaviors arise from social interactions, including fighting, mating, and others. While some of these behaviors involve specialized behavioral units, others involve simultaneous execution of multiple behavioral units leading to multitasking. For example, during courtship, a male fly can simultaneously approach a female fly (locomotion) and sing, two behaviors that would be mutually exclusive in the absence of a female fly. As a consequence, the simple tree-like hierarchical organization of behavior observed in isolated individuals (*Figure 8*) might break down during social interactions and lead to a flexible organization where actions are shared between multiple sequences and activities.

Acknowledgements

The author thanks Zach Mainen, James Murray, Cris Niell, Matt Smear, Osama Ahmed, Gordon Berman, the participants of the Computational Neuroethology Workshop 2021 in Jackson, and the members of the Mazzucato and Murray labs for many discussions and suggestions. The author acknowledges the reviewers for providing critical feedback and suggestions. LM was supported by the National Institute of Neurological Disorders and Stroke grant R01-NS118461 and the National Institute on Drug Abuse grant R01-DA055439 (CRCNS).

Additional information

Funding

Funder	Grant reference number	Author
National Institute of Neurological Disorders and Stroke	R01-NS118461	Luca Mazzucato
National Institute on Drug Abuse	R01-DA055439	Luca Mazzucato

The funders had no role in study design, data collection and interpretation, or the decision to submit the work for publication.

Author contributions

Luca Mazzucato, Conceptualization, Writing - original draft, Writing - review and editing

Author ORCIDs

Luca Mazzucato  <http://orcid.org/0000-0002-8525-7539>

References

- Ahamed T, Costa AC, Stephens GJ. 2021. Capturing the continuous complexity of behaviour in *Caenorhabditis elegans*. *Nature Physics* **17**:275–283. DOI: <https://doi.org/10.1038/s41567-020-01036-8>
- Alba V, Berman GJ, Bialek W, Shaevitz JW. 2020. Exploring a Strongly Non-Markovian Animal Behavior. *ArXiv Preprint*. <https://arxiv.org/abs/2012.15681>
- Aljadeff J, Stern M, Sharpee T. 2015. Transition to chaos in random networks with cell-type-specific connectivity. *Physical Review Letters* **114**:088101. DOI: <https://doi.org/10.1103/PhysRevLett.114.088101>, PMID: 25768781
- Almaas E, Kovács B, Vicsek T, Oltvai ZN, Barabási AL. 2004. Global organization of metabolic fluxes in the bacterium *Escherichia coli*. *Nature* **427**:839–843. DOI: <https://doi.org/10.1038/nature02289>, PMID: 14985762
- Amit DJ, Gutfreund H, Sompolinsky H. 1985. Spin-glass models of neural networks. *Physical Review. A, General Physics* **32**:1007–1018. DOI: <https://doi.org/10.1103/physreva.32.1007>, PMID: 9896156
- Amit DJ, Brunel N. 1997. Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cerebral Cortex (New York, N.Y)* **7**:237–252. DOI: <https://doi.org/10.1093/cercor/7.3.237>, PMID: 9143444
- Anderson DJ, Perona P. 2014. Toward a science of computational ethology. *Neuron* **84**:18–31. DOI: <https://doi.org/10.1016/j.neuron.2014.09.005>, PMID: 25277452
- Aon MA, Roussel MR, Cortassa S, O'Rourke B, Murray DB, Beckmann M, Lloyd D. 2008. The scale-free dynamics of eukaryotic cells. *PLOS ONE* **3**:e3624. DOI: <https://doi.org/10.1371/journal.pone.0003624>, PMID: 18982073
- Ayroles JF, Buchanan SM, Jenney C, Skutt-Kakaria K, Grenier J, Clark AG, Hartl DL, de Bivort BL. 2014. Behavioral Individuality Reveals Genetic Control of Phenotypic Variability. *bioRxiv*. DOI: <https://doi.org/10.1101/009027>, PMID: 24773151
- Barthas F, Kwan AC. 2017. Secondary Motor Cortex: Where 'Sensory' Meets 'Motor' in the Rodent Frontal Cortex. *Trends in Neurosciences* **40**:181–193. DOI: <https://doi.org/10.1016/j.tins.2016.11.006>, PMID: 28012708
- Bateson PPG, Hinde RA. 1976. *Growing Points Ethology*. Cambridge University Press.
- Batty E, Whiteway M, Saxena S, Biderman D, Abe T, Musall S, Gillis W, Markowitz J, Churchland A, Cunningham JP. 2019. BehaveNet: nonlinear embedding and Bayesian neural decoding of behavioral videos. *NeurIPS*.
- Ben Arous J, Laffont S, Chatenay D. 2009. Molecular and sensory basis of a food related two-state behavior in *C. elegans*. *PLOS ONE* **4**:e7584. DOI: <https://doi.org/10.1371/journal.pone.0007584>, PMID: 19851507
- Berman GJ, Choi DM, Bialek W, Shaevitz JW. 2014. Mapping the stereotyped behaviour of freely moving fruit flies. *Journal of The Royal Society Interface* **11**:20140672. DOI: <https://doi.org/10.1098/rsif.2014.0672>, PMID: 25142523
- Berman GJ, Bialek W, Shaevitz JW. 2016. Predictability and hierarchy in *Drosophila* behavior. *PNAS* **113**:11943–11948. DOI: <https://doi.org/10.1073/pnas.1607601113>
- Berman GJ. 2018. Measuring behavior across scales. *BMC Biology* **16**:1–11. DOI: <https://doi.org/10.1186/s12915-018-0494-7>, PMID: 29475451
- Bernacchia A, Seo H, Lee D, Wang XJ. 2011. A reservoir of time constants for memory traces in cortical neurons. *Nature Neuroscience* **14**:366–372. DOI: <https://doi.org/10.1038/nn.2752>, PMID: 21317906
- Berridge KC, Fentress JC, Parr H. 1987. Natural syntax rules control action sequence of rats. *Behavioural Brain Research* **23**:59–68. DOI: [https://doi.org/10.1016/0166-4328\(87\)90242-7](https://doi.org/10.1016/0166-4328(87)90242-7), PMID: 3828046
- Bouchaud JP. 1992. Weak ergodicity breaking and aging in disordered systems. *Journal de Physique I* **2**:1705–1713. DOI: <https://doi.org/10.1051/jp1:1992238>

- Bowman GR**, Pande VS, Noé F. 2013. An Introduction to Markov State Models and Their Application to Long Timescale Molecular Simulation, *Advances in Experimental Medicine and Biology*. Springer Science & Business Media. DOI: <https://doi.org/10.1007/978-94-007-7606-7>
- Brinkman BAW**, Yan H, Maffei A, Park IM, Fontanini A, Wang J, La Camera G. 2022. Metastable dynamics of neural circuits and networks. *Applied Physics Reviews* 9:011313. DOI: <https://doi.org/10.1063/5.0062603>, PMID: 35284030
- Brown AEX**, Yemini EI, Grundy LJ, Jucikas T, Schafer WR. 2013. A dictionary of behavioral motifs reveals clusters of genes affecting *Caenorhabditis elegans* locomotion. *PNAS* 110:791–796. DOI: <https://doi.org/10.1073/pnas.1211447110>, PMID: 23267063
- Brown AEX**, de Bivort B. 2018. Ethology as a physical science. *Nature Physics* 14:653–657. DOI: <https://doi.org/10.1038/s41567-018-0093-0>
- Brunel N**. 2000. Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *Journal of Computational Neuroscience* 8:183–208. DOI: <https://doi.org/10.1023/a:1008925309027>, PMID: 10809012
- Buchanan SM**, Kain JS, de Bivort BL. 2015. Neuronal control of locomotor handedness in *Drosophila*. *PNAS* 112:6700–6705. DOI: <https://doi.org/10.1073/pnas.1500804112>, PMID: 25953337
- Bucher D**, Prinz AA, Marder E. 2005. Animal-to-Animal Variability in Motor Pattern Production in Adults and during Growth. *Journal of Neuroscience* 25:1611–1619. DOI: <https://doi.org/10.1523/JNEUROSCI.3679-04.2005>, PMID: 15716396
- Calhoun AJ**, Pillow JW, Murthy M. 2019. Unsupervised identification of the internal states that shape natural behavior. *Nature Neuroscience* 22:2040–2049. DOI: <https://doi.org/10.1038/s41593-019-0533-x>, PMID: 31768056
- Cande J**, Namiki S, Qiu J, Korff W, Card GM, Shaevitz JW, Stern DL, Berman GJ. 2018. Optogenetic dissection of descending behavioral control in *Drosophila*. *eLife* 7:e34275. DOI: <https://doi.org/10.7554/eLife.34275>, PMID: 29943729
- Cao R**, Pastukhov A, Mattia M, Braun J. 2016. Collective Activity of Many Bistable Assemblies Reproduces Characteristic Dynamics of Multistable Perception. *Journal of Neuroscience* 36:6957–6972. DOI: <https://doi.org/10.1523/JNEUROSCI.4626-15.2016>, PMID: 27358454
- Cavanagh SE**, Wallis JD, Kennerley SW, Hunt LT. 2016. Autocorrelation structure at rest predicts value correlates of single neurons during reward-guided choice. *eLife* 5:e18937. DOI: <https://doi.org/10.7554/eLife.18937>, PMID: 27705742
- Cazettes F**, Murakami M, Renart A, Mainen Z. 2021. Reservoir of Decision Strategies in the Mouse Brain. *bioRxiv*. DOI: <https://doi.org/10.1101/2021.04.01.438090>
- Chance FS**, Abbott LF, Reyes AD. 2002. Gain modulation from background synaptic input. *Neuron* 35:773–782. DOI: [https://doi.org/10.1016/s0896-6273\(02\)00820-6](https://doi.org/10.1016/s0896-6273(02)00820-6), PMID: 12194875
- Chaudhuri R**, Bernacchia A, Wang XJ. 2014. A diversity of localized timescales in network activity. *eLife* 3:e01239. DOI: <https://doi.org/10.7554/eLife.01239>, PMID: 24448407
- Churchland MM**, Afshar A, Shenoy KV. 2006. A central source of movement variability. *Neuron* 52:1085–1096. DOI: <https://doi.org/10.1016/j.neuron.2006.10.034>, PMID: 17178410
- Churgin MA**, Jung SK, Yu CC, Chen X, Raizen DM, Fang-Yen C. 2017. Longitudinal imaging of *Caenorhabditis elegans* in a microfabricated device reveals variation in behavioral decline during aging. *eLife* 6:e26652. DOI: <https://doi.org/10.7554/eLife.26652>, PMID: 28537553
- Coen P**, Clemens J, Weinstein AJ, Pacheco DA, Deng Y, Murthy M. 2014. Dynamic sensory cues shape song structure in *Drosophila*. *Nature* 507:233–237. DOI: <https://doi.org/10.1038/nature13131>, PMID: 24598544
- Cohen Y**, Shen J, Semu D, Leman DP, Liberti WA, Perkins LN, Liberti DC, Kotton DN, Gardner TJ. 2020. Hidden neural states underlie canary song syntax. *Nature* 582:539–544. DOI: <https://doi.org/10.1038/s41586-020-2397-3>, PMID: 32555461
- Corrales-Carvajal VM**, Faisal AA, Ribeiro C. 2016. Internal states drive nutrient homeostasis by modulating exploration-exploitation trade-off. *eLife* 5:e19920. DOI: <https://doi.org/10.7554/eLife.19920>, PMID: 27770569
- Costa AC**, Ahamed T, Stephens GJ. 2019. Adaptive, locally linear models of complex dynamics. *PNAS* 116:1501–1510. DOI: <https://doi.org/10.1073/pnas.1813476116>, PMID: 30655347
- Costa AC**, Ahamed T, Jordan D, Stephens GJ. 2021. Maximally Predictive Ensemble Dynamics from Data. *Biophysics*. DOI: <https://doi.org/10.1101/2021.05.26.445816>
- Darshan R**, Wood WE, Peters S, Leblois A, Hansel D. 2017. A canonical neural mechanism for behavioral variability. *Nature Communications* 8:1–13. DOI: <https://doi.org/10.1038/ncomms15415>, PMID: 28530225
- Datta SR**, Anderson DJ, Branson K, Perona P, Leifer A. 2019. Computational Neuroethology: A Call to Action. *Neuron* 104:11–24. DOI: <https://doi.org/10.1016/j.neuron.2019.09.038>, PMID: 31600508
- Dawkins R**. 1976. Hierarchical organisation: A candidate principle for ethology. *Growing Points in Ethology* 7:54.
- Deco G**, Hugues E. 2012. Neural network mechanisms underlying stimulus driven variability reduction. *PLOS Computational Biology* 8:e1002395. DOI: <https://doi.org/10.1371/journal.pcbi.1002395>, PMID: 22479168
- Deutsch D**, Pacheco D, Encarnacion-Rivera L, Pereira T, Fathy R, Clemens J, Girardin C, Calhoun A, Ireland E, Burke A, Dorkenwald S, McKellar C, Macrina T, Lu R, Lee K, Kemnitz N, Ih D, Castro M, Halageri A, Jordan C, et al. 2020. The neural basis for a persistent internal state in *Drosophila* females. *eLife* 9:e59502. DOI: <https://doi.org/10.7554/eLife.59502>, PMID: 33225998
- Dhawale AK**, Smith MA, Ölveczky BP. 2017. The Role of Variability in Motor Learning. *Annual Review of Neuroscience* 40:479–498. DOI: <https://doi.org/10.1146/annurev-neuro-072116-031548>, PMID: 28489490
- Dissel S**, Hansen CN, Özkaya Ö, Hemsley M, Kyriacou CP, Rosato E. 2014. The logic of circadian organization in *Drosophila*. *Current Biology* 24:2257–2266. DOI: <https://doi.org/10.1016/j.cub.2014.08.023>, PMID: 25220056

- Duncker L**, Bohner G, Boussard J, Sahani M. 2019. Learning interpretable continuous-time models of latent stochastic dynamical systems. *International Conference on Machine Learning PMLR*. 1726–1734.
- Eliades SJ**, Miller CT. 2017. Marmoset vocal communication: Behavior and neurobiology. *Developmental Neurobiology* **77**:286–299. DOI: <https://doi.org/10.1002/dneu.22464>, PMID: 27739195
- Emmerling M**, Dauner M, Ponti A, Fiaux J, Hochuli M, Szyperski T, Wüthrich K, Bailey JE, Sauer U. 2002. Metabolic Flux Responses to Pyruvate Kinase Knockout in *Escherichia coli*. *Journal of Bacteriology* **184**:152–164. DOI: <https://doi.org/10.1128/JB.184.1.152-164.2002>, PMID: 11741855
- Fan D**, Rossi MA, Yin HH. 2012. Mechanisms of Action Selection and Timing in Substantia Nigra Neurons. *Journal of Neuroscience* **32**:5534–5548. DOI: <https://doi.org/10.1523/JNEUROSCI.5924-11.2012>, PMID: 22514315
- Findley TM**, Wyrick DG, Cramer JL, Brown MA, Holcomb B, Attey R, Yeh D, Monasevitch E, Nouboussi N, Cullen I, Songco JO, King JF, Ahmadian Y, Smear MC. 2021. Sniff-synchronized, gradient-guided olfactory search by freely moving mice. *eLife* **10**:e58523. DOI: <https://doi.org/10.7554/eLife.58523>, PMID: 33942713
- Finkelstein A**, Fontolan L, Economo MN, Li N, Romani S, Svoboda K. 2021. Attractor dynamics gate cortical information flow during decision-making. *Nature Neuroscience* **24**:843–850. DOI: <https://doi.org/10.1038/s41593-021-00840-6>, PMID: 33875892
- Flavell SW**, Pokala N, Macosko EZ, Albrecht DR, Larsch J, Bargmann CI. 2013. Serotonin and the neuropeptide PDF initiate and extend opposing behavioral states in *C. elegans*. *Cell* **154**:1023–1035. DOI: <https://doi.org/10.1016/j.cell.2013.08.001>, PMID: 23972393
- Fonseca MS**, Murakami M, Mainen ZF. 2015. Activation of dorsal raphe serotonergic neurons promotes waiting but is not reinforcing. *Current Biology* **25**:306–315. DOI: <https://doi.org/10.1016/j.cub.2014.12.002>, PMID: 25601545
- Fourcaud N**, Brunel N. 2002. Dynamics of the firing probability of noisy integrate-and-fire neurons. *Neural Computation* **14**:2057–2110. DOI: <https://doi.org/10.1162/089976602320264015>, PMID: 12184844
- Funahashi S**, Bruce CJ, Goldman-Rakic PS. 1989. Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *Journal of Neurophysiology* **61**:331–349. DOI: <https://doi.org/10.1152/jn.1989.61.2.331>, PMID: 2918358
- Fuster JM**, Jervey JP. 1981. Inferotemporal neurons distinguish and retain behaviorally relevant features of visual stimuli. *Science (New York, N.Y.)* **212**:952–955. DOI: <https://doi.org/10.1126/science.7233192>, PMID: 7233192
- Geddes CE**, Li H, Jin X. 2018. Optogenetic Editing Reveals the Hierarchical Organization of Learned Action Sequences. *Cell* **174**:32–43. DOI: <https://doi.org/10.1016/j.cell.2018.06.012>, PMID: 29958111
- Genkin M**, Engel TA. 2020. Moving beyond generalization to accurate interpretation of flexible models. *Nature Machine Intelligence* **2**:674–683. DOI: <https://doi.org/10.1038/s42256-020-00242-6>
- Gjorgjieva J**, Drion G, Marder E. 2016. Computational implications of biophysical diversity and multiple timescales in neurons and synapses for circuit performance. *Current Opinion in Neurobiology* **37**:44–52. DOI: <https://doi.org/10.1016/j.conb.2015.12.008>, PMID: 26774694
- Gomez-Marin A**, Stephens GJ, Brown AEX. 2016. Hierarchical compression of *Caenorhabditis elegans* locomotion reveals phenotypic differences in the organization of behaviour. *Journal of the Royal Society, Interface* **13**:20160466. DOI: <https://doi.org/10.1098/rsif.2016.0466>, PMID: 27581484
- Guo ZV**, Inagaki HK, Daie K, Druckmann S, Gerfen CR, Svoboda K. 2017. Maintenance of persistent activity in a frontal thalamocortical loop. *Nature* **545**:181–186. DOI: <https://doi.org/10.1038/nature22324>, PMID: 28467817
- Guo K**, Yamawaki N, Svoboda K, Shepherd GMG. 2018. Anterolateral Motor Cortex Connects with a Medial Subdivision of Ventromedial Thalamus through Cell Type-Specific Circuits, Forming an Excitatory Thalamo-Cortico-Thalamic Loop via Layer 1 Apical Tuft Dendrites of Layer 5B Pyramidal Tract Type Neurons. *The Journal of Neuroscience* **38**:8787–8797. DOI: <https://doi.org/10.1523/JNEUROSCI.1333-18.2018>, PMID: 30143573
- Gupta S**, Gomez-Marin A. 2019. A Context-Free Grammar for *Caenorhabditis Elegans* Behavior. *bioRxiv*. DOI: <https://doi.org/10.1101/708891>
- Hallinen KM**, Dempsey R, Scholz M, Yu X, Linder A, Randi F, Sharma AK, Shaevitz JW, Leifer AM. 2021. Decoding locomotion from population neural activity in moving *C. elegans*. *eLife* **10**:e66135. DOI: <https://doi.org/10.7554/eLife.66135>, PMID: 34323218
- Hänggi P**, Talkner P, Borkovec M. 1990. Reaction-rate theory: fifty years after Kramers. *Reviews of Modern Physics* **62**:251–341. DOI: <https://doi.org/10.1103/RevModPhys.62.251>
- Hastings MH**, Maywood ES, Brancaccio M. 2018. Generation of circadian rhythms in the suprachiasmatic nucleus. *Nature Reviews. Neuroscience* **19**:453–469. DOI: <https://doi.org/10.1038/s41583-018-0026-z>, PMID: 29934559
- Honegger K**, de Bivort B. 2018. Stochasticity, individuality and behavior. *Current Biology* **28**:R8–R12. DOI: <https://doi.org/10.1016/j.cub.2017.11.058>, PMID: 29316423
- Hong W**, Kennedy A, Burgos-Artizzu XP, Zelikowsky M, Navonne SG, Perona P, Anderson DJ. 2015. Automated measurement of mouse social behaviors using depth sensing, video tracking, and machine learning. *PNAS* **112**:E5351–E5360. DOI: <https://doi.org/10.1073/pnas.1515982112>, PMID: 26354123
- Hopfield JJ**. 1982. Neural networks and physical systems with emergent collective computational abilities. *PNAS* **79**:2554–2558. DOI: <https://doi.org/10.1073/pnas.79.8.2554>
- Hopfield JJ**. 1984. Neurons with graded response have collective computational properties like those of two-state neurons. *PNAS* **81**:3088–3092. DOI: <https://doi.org/10.1073/pnas.81.10.3088>
- Hoy JL**, Yavorska I, Wehr M, Niell CM. 2016. Vision Drives Accurate Approach Behavior during Prey Capture in Laboratory Mice. *Current Biology* **26**:3046–3052. DOI: <https://doi.org/10.1016/j.cub.2016.09.009>, PMID: 27773567

- Huang C, Doiron B. 2017. Once upon a (slow) time in the land of recurrent neuronal networks. *Current Opinion in Neurobiology* **46**:31–38. DOI: <https://doi.org/10.1016/j.conb.2017.07.003>, PMID: 28756341
- Inagaki HK, Fontolan L, Romani S, Svoboda K. 2019. Discrete attractor dynamics underlies persistent activity in the frontal cortex. *Nature* **566**:212–217. DOI: <https://doi.org/10.1038/s41586-019-0919-7>, PMID: 30728503
- Inagaki HK, Chen S, Ridder MC, Sah P, Li N, Yang Z, Hasanbegovic H, Gao Z, Gerfen CR, Svoboda K. 2022. A midbrain-thalamus-cortex circuit reorganizes cortical dynamics to initiate movement. *Cell* **185**:1065–1081. DOI: <https://doi.org/10.1016/j.cell.2022.02.006>, PMID: 35245431
- Jaramillo S, Zador AM. 2011. The auditory cortex mediates the perceptual effects of acoustic temporal expectation. *Nature Neuroscience* **14**:246–251. DOI: <https://doi.org/10.1038/nn.2688>, PMID: 21170056
- Johnson RE, Linderman S, Panier T, Wee CL, Song E, Herrera KJ, Miller A, Engert F. 2020. Probabilistic Models of Larval Zebrafish Behavior Reveal Structure on Many Scales. *Current Biology* **30**:70–82. DOI: <https://doi.org/10.1016/j.cub.2019.11.026>, PMID: 31866367
- Jones LM, Fontanini A, Sadacca BF, Miller P, Katz DB. 2007. Natural stimuli evoke dynamic sequences of states in sensory cortical ensembles. *PNAS* **104**:18772–18777. DOI: <https://doi.org/10.1073/pnas.0705546104>, PMID: 18000059
- Joshua M, Medina JF, Lisberger SG. 2013. Diversity of neural responses in the brainstem during smooth pursuit eye movements constrains the circuit mechanisms of neural integration. *The Journal of Neuroscience* **33**:6633–6647. DOI: <https://doi.org/10.1523/JNEUROSCI.3732-12.2013>, PMID: 23575860
- Juavinett AL, Bekheet G, Churchland AK. 2019. Chronically implanted Neuropixels probes enable high-yield recordings in freely moving mice. *eLife* **8**:e47188. DOI: <https://doi.org/10.7554/eLife.47188>, PMID: 31411559
- Kao MH, Doupe AJ, Brainard MS. 2005. Contributions of an avian basal ganglia-forebrain circuit to real-time modulation of song. *Nature* **433**:638–643. DOI: <https://doi.org/10.1038/nature03127>, PMID: 15703748
- Kaplan HS, Salazar Thula O, Khoss N, Zimmer M. 2020. Nested Neuronal Dynamics Orchestrate a Behavioral Hierarchy across Timescales. *Neuron* **105**:562–576. DOI: <https://doi.org/10.1016/j.neuron.2019.10.037>, PMID: 31786012
- Kawai R, Markman T, Poddar R, Ko R, Fantana AL, Dhawale AK, Kampff AR, Ölveczky BP. 2015. Motor cortex is required for learning but not for executing a motor skill. *Neuron* **86**:800–812. DOI: <https://doi.org/10.1016/j.neuron.2015.03.024>, PMID: 25892304
- Kiani R, Cueva CJ, Reppas JB, Peixoto D, Ryu SI, Newsome WT. 2015. Natural Grouping of Neural Responses Reveals Spatially Segregated Clusters in Prearcuate Cortex. *Neuron* **85**:1359–1373. DOI: <https://doi.org/10.1016/j.neuron.2015.02.014>, PMID: 25728571
- Killeen PR, Fetterman JG. 1988. A behavioral theory of timing. *Psychological Review* **95**:274–295. DOI: <https://doi.org/10.1037/0033-295x.95.2.274>, PMID: 3375401
- Klibaite U, Kislin M, Verpeut JL, Sun X, Shaevitz JW, Wang SSH. 2021. Deep Phenotyping Reveals Movement Phenotypes in Mouse Neurodevelopmental Models. *In Review*. DOI: <https://doi.org/10.21203/rs.3.rs-798847/v1>
- Koch KW, Fuster JM. 1989. Unit activity in monkey parietal cortex related to haptic perception and temporary memory. *Experimental Brain Research* **76**:292–306. DOI: <https://doi.org/10.1007/BF00247889>, PMID: 2767186
- La Camera G, Fontanini A, Mazzucato L. 2019. Cortical computations via metastable activity. *Current Opinion in Neurobiology* **58**:37–45. DOI: <https://doi.org/10.1016/j.conb.2019.06.007>, PMID: 31326722
- Lauer J, Zhou M, Ye S, Menegas W, Nath T, Rahman MM, Di Santo V, Soberanes D, Feng G, Murthy VN, Lauder G, Dulac C, Mathis MW, Mathis A. 2021. Multi-Animal Pose Estimation and Tracking with DeepLabCut. *Animal Behavior and Cognition*. DOI: <https://doi.org/10.1101/2021.04.30.442096>
- Lee W-CA, Bonin V, Reed M, Graham BJ, Hood G, Glattfelder K, Reid RC. 2016. Anatomy and function of an excitatory network in the visual cortex. *Nature* **532**:370–374. DOI: <https://doi.org/10.1038/nature17192>
- Li N, Daie K, Svoboda K, Druckmann S. 2016. Robust neuronal dynamics in premotor cortex during motor planning. *Nature* **532**:459–464. DOI: <https://doi.org/10.1038/nature17643>, PMID: 27074502
- Lim S, McKee JL, Woloszyn L, Amit Y, Freedman DJ, Sheinberg DL, Brunel N. 2015. Inferring learning rules from distributions of firing rates in cortical neurons. *Nature Neuroscience* **18**:1804–1810. DOI: <https://doi.org/10.1038/nn.4158>, PMID: 26523643
- Linderman S, Johnson M, Miller A, Adams R, Blei D, Paninski L. 2017. Bayesian learning and inference in recurrent switching linear dynamical systems. Singh A, Zhu J (Eds). *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*. PMLR. p. 914–922. DOI: <https://doi.org/10.1109/CAMSAP.2017.8313132>
- Linderman S, Nichols A, Blei D, Zimmer M, Paninski L. 2019. Hierarchical Recurrent State Space Models Reveal Discrete and Continuous Dynamics of Neural Activity in *C. Elegans*. *Neuroscience*:621540. DOI: <https://doi.org/10.1101/621540>
- Litwin-Kumar A, Doiron B. 2012. Slow dynamics and high variability in balanced cortical networks with clustered connections. *Nature Neuroscience* **15**:1498–1505. DOI: <https://doi.org/10.1038/nn.3220>, PMID: 23001062
- Litwin-Kumar A, Doiron B. 2014. Formation and maintenance of neuronal assemblies through synaptic plasticity. *Nature Communications* **5**:1–12. DOI: <https://doi.org/10.1038/ncomms6319>, PMID: 25395015
- Loggiaco L, Escola GS. 2020. Thalamocortical Motor Circuit Insights for More Robust Hierarchical Control of Complex Sequences. *ArXiv Preprint*. <https://arxiv.org/abs/2006.13332>
- Loggiaco L, Abbott LF, Escola S. 2021. Thalamic control of cortical dynamics in a model of flexible motor sequencing. *Cell Reports* **35**:109090. DOI: <https://doi.org/10.1016/j.celrep.2021.109090>, PMID: 34077721

- Lottem E**, Banerjee D, Vertechi P, Sarra D, Lohuis MO, Mainen ZF. 2018. Activation of serotonin neurons promotes active persistence in a probabilistic foraging task. *Nature Communications* **9**:1–12. DOI: <https://doi.org/10.1038/s41467-018-03438-y>, PMID: 29520000
- Loveless J**, Lagogiannis K, Webb B. 2019. Modelling the mechanics of exploration in larval *Drosophila*. *PLOS Computational Biology* **15**:e1006635. DOI: <https://doi.org/10.1371/journal.pcbi.1006635>, PMID: 31276489
- Maboudi K**, Ackermann E, de Jong LW, Pfeiffer BE, Foster D, Diba K, Kemere C. 2018. Uncovering temporal structure in hippocampal output patterns. *eLife* **7**:e34467. DOI: <https://doi.org/10.7554/eLife.34467>, PMID: 29869611
- MacDougall M**, Nummela SU, Coop S, Disney A, Mitchell JF, Miller CT. 2016. Optogenetic manipulation of neural circuits in awake marmosets. *Journal of Neurophysiology* **116**:1286–1294. DOI: <https://doi.org/10.1152/jn.00197.2016>, PMID: 27334951
- Magnasco MO**, Piro O, Cecchi GA. 2009. Self-tuned critical anti-Hebbian networks. *Physical Review Letters* **102**:258102. DOI: <https://doi.org/10.1103/PhysRevLett.102.258102>, PMID: 19659122
- Markowitz JE**, Ivie E, Kligler L, Gardner TJ. 2013. Long-range order in canary song. *PLOS Computational Biology* **9**:e1003052. DOI: <https://doi.org/10.1371/journal.pcbi.1003052>, PMID: 23658509
- Markowitz JE**, Gillis WF, Beron CC, Neufeld SQ, Robertson K, Bhagat ND, Peterson RE, Peterson E, Hyun M, Linderman SW, Sabatini BL, Datta SR. 2018. The Striatum Organizes 3D Behavior via Moment-to-Moment Action Selection. *Cell* **174**:44–58. DOI: <https://doi.org/10.1016/j.cell.2018.04.019>, PMID: 29779950
- Marques JC**, Lackner S, Félix R, Orger MB. 2018. Structure of the Zebrafish Locomotor Repertoire Revealed with Unsupervised Behavioral Clustering. *Current Biology* **28**:181–195. DOI: <https://doi.org/10.1016/j.cub.2017.12.002>, PMID: 29307558
- Marshall JD**, Aldarondo DE, Dunn TW, Wang WL, Berman GJ, Ölveczky BP. 2021. Continuous Whole-Body 3D Kinematic Recordings across the Rodent Behavioral Repertoire. *Neuron* **109**:420–437. DOI: <https://doi.org/10.1016/j.neuron.2020.11.016>, PMID: 33340448
- Marshall JH**, Kim YS, Machado TA, Quirin S, Benson B, Kadmon J, Raja C, Chibukhchyan A, Ramakrishnan C, Inoue M, Shane JC, McKnight DJ, Yoshizawa S, Kato HE, Ganguli S, Deisseroth K. 2019. Cortical layer-specific critical dynamics triggering perception. *Science (New York, N.Y.)* **365**:eaaw5202. DOI: <https://doi.org/10.1126/science.aaw5202>, PMID: 31320556
- Martí D**, Brunel N, Ostojic S. 2018. Correlations between synapses in pairs of neurons slow down dynamics in randomly connected neural networks. *Physical Review. E* **97**:062314. DOI: <https://doi.org/10.1103/PhysRevE.97.062314>, PMID: 30011528
- Mascaro M**, Amit DJ. 1999. Effective neural response function for collective population states. *Network (Bristol, England)* **10**:351–373. PMID: 10695764.
- Mathis A**, Mamidanna P, Cury KM, Abe T, Murthy VN, Mathis MW, Bethge M. 2018. DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. *Nature Neuroscience* **21**:1281–1289. DOI: <https://doi.org/10.1038/s41593-018-0209-y>, PMID: 30127430
- Mattia M**, Sanchez-Vives MV. 2012. Exploring the spectrum of dynamical regimes and timescales in spontaneous cortical activity. *Cognitive Neurodynamics* **6**:239–250. DOI: <https://doi.org/10.1007/s11571-011-9179-4>, PMID: 23730355
- Mattia M**, Pani P, Mirabella G, Costa S, Del Giudice P, Ferraina S. 2013. Heterogeneous attractor cell assemblies for motor planning in premotor cortex. *The Journal of Neuroscience* **33**:11155–11168. DOI: <https://doi.org/10.1523/JNEUROSCI.4664-12.2013>, PMID: 23825419
- Mazzucato L**, Fontanini A, La Camera G. 2015. Dynamics of multistable states during ongoing and evoked cortical activity. *The Journal of Neuroscience* **35**:8214–8231. DOI: <https://doi.org/10.1523/JNEUROSCI.4819-14.2015>, PMID: 26019337
- Mazzucato L**, La Camera G, Fontanini A. 2019. Expectation-induced modulation of metastable activity underlies faster coding of sensory stimuli. *Nature Neuroscience* **22**:787–796. DOI: <https://doi.org/10.1038/s41593-019-0364-9>, PMID: 30936557
- McKenzie S**, Frank AJ, Kinsky NR, Porter B, Rivière PD, Eichenbaum H. 2014. Hippocampal representation of related and opposing memories develop within distinct, hierarchically organized neural schemas. *Neuron* **83**:202–215. DOI: <https://doi.org/10.1016/j.neuron.2014.05.019>, PMID: 24910078
- Michaël AM**, Abe ET, Niell CM. 2020. Dynamics of gaze control during prey capture in freely moving mice. *eLife* **9**:e57458. DOI: <https://doi.org/10.7554/eLife.57458>, PMID: 32706335
- Miller P**, Katz DB. 2010. Stochastic transitions between neural states in taste processing and decision-making. *The Journal of Neuroscience* **30**:2559–2570. DOI: <https://doi.org/10.1523/JNEUROSCI.3047-09.2010>, PMID: 20164341
- Miller JK**, Ayzenshtat I, Carrillo-Reid L, Yuste R. 2014. Visual stimuli recruit intrinsically generated cortical ensembles. *PNAS* **111**:E4053–E4061. DOI: <https://doi.org/10.1073/pnas.1406077111>, PMID: 25201983
- Miri A**, Daie K, Arrenberg AB, Baier H, Aksay E, Tank DW. 2011. Spatial gradients and multidimensional dynamics in a neural integrator circuit. *Nature Neuroscience* **14**:1150–1159. DOI: <https://doi.org/10.1038/nn.2888>, PMID: 21857656
- Mischel W**, Ebbesen EB. 1970. Attention in delay of gratification. *Journal of Personality and Social Psychology* **16**:329–337. DOI: <https://doi.org/10.1037/h0029815>
- Miyashita Y**, Chang HS. 1988. Neuronal correlate of pictorial short-term memory in the primate temporal cortex. *Nature* **331**:68–70. DOI: <https://doi.org/10.1038/331068a0>, PMID: 3340148
- Moreno-Bote R**, Beck J, Kanitscheider I, Pitkow X, Latham P, Pouget A. 2014. Information-limiting correlations. *Nature Neuroscience* **17**:1410–1417. DOI: <https://doi.org/10.1038/nn.3807>, PMID: 25195105

- Murakami M**, Vicente MI, Costa GM, Mainen ZF. 2014. Neural antecedents of self-initiated actions in secondary motor cortex. *Nature Neuroscience* **17**:1574–1582. DOI: <https://doi.org/10.1038/nn.3826>, PMID: 25262496
- Murakami M**, Shteingart H, Loewenstein Y, Mainen ZF. 2017. Distinct Sources of Deterministic and Stochastic Components of Action Timing Decisions in Rodent Frontal Cortex. *Neuron* **94**:908–919. DOI: <https://doi.org/10.1016/j.neuron.2017.04.040>, PMID: 28521140
- Murray JD**, Bernacchia A, Freedman DJ, Romo R, Wallis JD, Cai X, Padoa-Schioppa C, Pasternak T, Seo H, Lee D, Wang X-J. 2014. A hierarchy of intrinsic timescales across primate cortex. *Nature Neuroscience* **17**:1661–1663. DOI: <https://doi.org/10.1038/nn.3862>, PMID: 25383900
- Murray JM**, Escola GS. 2017. Learning multiple variable-speed sequences in striatum via cortical tutoring. *eLife* **6**:e26084. DOI: <https://doi.org/10.7554/eLife.26084>, PMID: 28481200
- Murray J.M**, Escola GS. 2020. Remembrance of things practiced with fast and slow learning in cortical and subcortical pathways. *Nature Communications* **11**:1–12. DOI: <https://doi.org/10.1038/s41467-020-19788-5>, PMID: 33361766
- Nassar J**, Linderman SW, Bugallo M, Park IM. 2018. Tree-Structured Recurrent Switching Linear Dynamical Systems For Multi-Scale Modeling. *arXiv*. <https://arxiv.org/abs/1811.12386>
- Nelson RJ**. An introduction to behavioral endocrinology. Sinauer Associates; 2005.
- Nourizonoz A**, Zimmermann R, Ho CLA, Pellat S, Ormen Y, Prévost-Solié C, Reymond G, Pifferi F, Aujard F, Herrel A, Huber D. 2020. EthoLoop: automated closed-loop neuroethology in naturalistic environments. *Nature Methods* **17**:1052–1059. DOI: <https://doi.org/10.1038/s41592-020-0961-2>, PMID: 32994566
- Nummela SU**, Jovanovic V, de la Mothe L, Miller CT. 2017. Social Context-Dependent Activity in Marmoset Frontal Cortex Populations during Natural Conversations. *The Journal of Neuroscience* **37**:7036–7047. DOI: <https://doi.org/10.1523/JNEUROSCI.0702-17.2017>, PMID: 28630255
- Olveczky BP**, Andalman AS, Fee MS, Schultz W. 2005. Vocal Experimentation in the Juvenile Songbird Requires a Basal Ganglia Circuit. *PLOS Biology* **3**:e153. DOI: <https://doi.org/10.1371/journal.pbio.0030153>, PMID: 15826219
- Parga N**, Virasoro MA. 1986. The ultrametric organization of memories in a neural network. *Journal de Physique* **47**:1857–1864. DOI: <https://doi.org/10.1051/jphys:0198600470110185700>
- Parker PRL**, Abe ETT, Beatie NT, Leonard ESP, Martins DM, Sharp SL, Wyrick DG, Mazzucato L, Niell CM. 2021. Distance Estimation from Monocular Cues in an Ethological Visuomotor Task. *bioRxiv*. DOI: <https://doi.org/10.1101/2021.09.29.462468>
- Patke A**, Young MW, Axelrod S. 2020. Molecular mechanisms and physiological importance of circadian rhythms. *Nature Reviews Molecular Cell Biology* **21**:67–84. DOI: <https://doi.org/10.1038/s41580-019-0179-2>, PMID: 31768006
- Pereira TD**, Aldarondo DE, Willmore L, Kislin M, Wang SSH, Murthy M, Shaevitz JW. 2019. Fast animal pose estimation using deep neural networks. *Nature Methods* **16**:117–125. DOI: <https://doi.org/10.1038/s41592-018-0234-5>, PMID: 30573820
- Pereira TD**, Shaevitz JW, Murthy M. 2020a. Quantifying behavior to understand the brain. *Nature Neuroscience* **23**:1537–1549. DOI: <https://doi.org/10.1038/s41593-020-00734-z>, PMID: 33169033
- Pereira TD**, Tabris N, Li J, Ravindranath S, Papadoyannis ES, Wang ZY, Turner DM, McKenzie-Smith G, Kocher SD, Falkner AL, Shaevitz JW, Murthy M. 2020b. SLEAP: Multi-Animal Pose Tracking. *Animal Behavior and Cognition*. DOI: <https://doi.org/10.1101/2020.08.31.276246>
- Perin R**, Berger TK, Markram H. 2011. A synaptic organizing principle for cortical neuronal groups. *PNAS* **108**:5419–5424. DOI: <https://doi.org/10.1073/pnas.1016051108>, PMID: 21383177
- Prinz AA**, Bucher D, Marder E. 2004. Similar network activity from disparate circuit parameters. *Nature Neuroscience* **7**:1345–1352. DOI: <https://doi.org/10.1038/nn1352>, PMID: 15558066
- Rammal R**, Toulouse G, Virasoro MA. 1986. Ultrametricity for physicists. *Reviews of Modern Physics* **58**:765–788. DOI: <https://doi.org/10.1103/RevModPhys.58.765>
- Recanatesi S**, Pereira-Obilinovic U, Murakami M, Mainen Z, Mazzucato L. 2022. Metastable attractors explain the variable timing of stable behavioral action sequences. *Neuron* **110**:139–153. DOI: <https://doi.org/10.1016/j.neuron.2021.10.011>, PMID: 34717794
- Reddy G**, Desban L, Tanaka H, Roussel J, Mirat O, Wyart C, Faisal AA. 2022. A lexical approach for identifying behavioural action sequences. *PLOS Computational Biology* **18**:e1009672. DOI: <https://doi.org/10.1371/journal.pcbi.1009672>, PMID: 35007275
- Rigotti M**, Barak O, Warden MR, Wang XJ, Daw ND, Miller EK, Fusi S. 2013. The importance of mixed selectivity in complex cognitive tasks. *Nature* **497**:585–590. DOI: <https://doi.org/10.1038/nature12160>, PMID: 23685452
- Rostami V**, Rost T, Riehle A, Albada SJ, Nawrot MP. 2020. Spiking Neural Network Model of Motor Cortex with Joint Excitatory and Inhibitory Clusters Reflects Task Uncertainty, Reaction Times, and Variability Dynamics. *bioRxiv*. DOI: <https://doi.org/10.1101/2020.02.27.968339>
- Roussel MR**, Lloyd D. 2007. Observation of a chaotic multioscillatory metabolic attractor by real-time monitoring of a yeast continuous culture. *The FEBS Journal* **274**:1011–1018. DOI: <https://doi.org/10.1111/j.1742-4658.2007.05651.x>, PMID: 17250739
- Samuelson CL**, Gardner MPH, Fontanini A. 2012. Effects of cue-triggered expectation on cortical processing of taste. *Neuron* **74**:410–422. DOI: <https://doi.org/10.1016/j.neuron.2012.02.031>, PMID: 22542192
- Schaub MT**, Billeh YN, Anastassiou CA, Koch C, Barahona M. 2015. Emergence of Slow-Switching Assemblies in Structured Neuronal Networks. *PLOS Computational Biology* **11**:e1004196. DOI: <https://doi.org/10.1371/journal.pcbi.1004196>, PMID: 26176664

- Schibler U**, Gotic I, Saini C, Gos P, Curie T, Emmenegger Y, Sinturel F, Gosselin P, Gerber A, Fleury-Olela F, Rando G, Demarque M, Franken P. 2015. Clock-Talk: Interactions between Central and Peripheral Circadian Oscillators in Mammals. *Cold Spring Harbor Symposia on Quantitative Biology* **80**:223–232. DOI: <https://doi.org/10.1101/sqb.2015.80.027490>, PMID: 26683231
- Schreiner DC**, Cazares C, Renteria R, Gremel CM. 2021. Mice Are Not Automatons; Subjective Experience in Premotor Circuits Guides Behavior. *Neuroscience*. DOI: <https://doi.org/10.1101/2021.06.23.449617>
- Schwarz RF**, Branicky R, Grundy LJ, Schafer WR, Brown AEX. 2015. Changes in Postural Syntax Characterize Sensory Modulation and Natural Variation of *C. elegans* Locomotion. *PLOS Computational Biology* **11**:e1004322. DOI: <https://doi.org/10.1371/journal.pcbi.1004322>, PMID: 26295152
- Segalin C**, Williams J, Karigo T, Hui M, Zelikowsky M, Sun JJ, Perona P, Anderson DJ, Kennedy A. 2020. The Mouse Action Recognition System (MARS): A Software Pipeline for Automated Analysis of Social Behaviors in Mice. *Neuroscience*. DOI: <https://doi.org/10.1101/2020.07.26.222299>
- Shadlen MN**, Newsome WT. 1998. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *The Journal of Neuroscience* **18**:3870–3896 PMID: 9570816.,
- Shi YL**, Steinmetz NA, Moore T, Boahen K, Engel TA. 2022. Cortical state dynamics and selective attention define the spatial pattern of correlated variability in neocortex. *Nature Communications* **13**:1–15. DOI: <https://doi.org/10.1038/s41467-021-27724-4>
- Simon HA**. 1962. The architecture of complexity. *PNAS* **106**:467.
- Softky WR**, Koch C. 1993. The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *The Journal of Neuroscience* **13**:334–350 PMID: 8423479.,
- Song S**, Sjöström PJ, Reigl M, Nelson S, Chklovskii DB, Friston KJ. 2005. Highly Nonrandom Features of Synaptic Connectivity in Local Cortical Circuits. *PLOS Biology* **3**:e68. DOI: <https://doi.org/10.1371/journal.pbio.0030068>
- Sowell ER**, Thompson PM, Toga AW. 2004. Mapping changes in the human cortex throughout the span of life. *The Neuroscientist* **10**:372–392. DOI: <https://doi.org/10.1177/1073858404263960>, PMID: 15271264
- Stephens GJ**, Johnson-Kerner B, Bialek W, Ryu WS. 2008. Dimensionality and dynamics in the behavior of *C. elegans*. *PLOS Computational Biology* **4**:e1000028. DOI: <https://doi.org/10.1371/journal.pcbi.1000028>, PMID: 18389066
- Stern M**, Istrate N, Mazzucato L. 2021. A Reservoir of Timescales in Random Neural Networks. *bioRxiv*. DOI: <https://doi.org/10.1101/2021.10.11.463861>
- Stroud JP**, Porter MA, Hennequin G, Vogels TP. 2018. Motor primitives in space and time via targeted gain modulation in cortical networks. *Nature Neuroscience* **21**:1774–1783. DOI: <https://doi.org/10.1038/s41593-018-0276-0>, PMID: 30482949
- Sutton RS**, Barto AG. Reinforcement learning: An introduction. MIT press; 2018.
- Szigeti B**, Deogade A, Webb B. 2015. Searching for motifs in the behaviour of larval *Drosophila melanogaster* and *Caenorhabditis elegans* reveals continuity between behavioural states. *Journal of the Royal Society, Interface* **12**:20150899. DOI: <https://doi.org/10.1098/rsif.2015.0899>, PMID: 26609067
- Tao L**, Ozarkar S, Beck JM, Bhandawat V. 2019. Statistical structure of locomotion and its modulation by odors. *eLife* **8**:e41235. DOI: <https://doi.org/10.7554/eLife.41235>, PMID: 30620334
- Tinbergen N**. 1951. The Study of Instinct. Plunkett Lake Press.
- Tishby N**, Pereira FC, Bialek W. 2000. The Information Bottleneck Method. *arXiv*. <https://arxiv.org/abs/physics/0004057>
- Toyoizumi T**, Abbott LF. 2011. Beyond the edge of chaos: amplification and temporal integration by recurrent networks in the chaotic regime. *Physical Review E, Statistical, Nonlinear, and Soft Matter Physics* **84**:051908. DOI: <https://doi.org/10.1103/PhysRevE.84.051908>, PMID: 22181445
- van Beers RJ**, Haggard P, Wolpert DM. 2004. The role of execution noise in movement variability. *Journal of Neurophysiology* **91**:1050–1063. DOI: <https://doi.org/10.1152/jn.00652.2003>, PMID: 14561687
- van Daal RJJ**, Aydin Ç, Michon F, Aarts AAA, Kraft M, Kloosterman F, Haesler S. 2021. Implantation of Neuropixels probes for chronic recording of neuronal activity in freely behaving mice and rats. *Nature Protocols* **16**:3322–3347. DOI: <https://doi.org/10.1038/s41596-021-00539-9>, PMID: 34108732
- van Vreeswijk C**, Sompolinsky H. 1996. Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science (New York, N.Y.)* **274**:1724–1726. DOI: <https://doi.org/10.1126/science.274.5293.1724>, PMID: 8939866
- Vogelstein JT**, Park Y, Ohyama T, Kerr RA, Truman JW, Priebe CE, Zlatic M. 2014. Discovery of brainwide neural-behavioral maps via multiscale unsupervised structure learning. *Science (New York, N.Y.)* **344**:386–392. DOI: <https://doi.org/10.1126/science.1250298>, PMID: 24674869
- Wang J**. 2015. Landscape and flux theory of non-equilibrium dynamical systems with application to biology. *Advances in Physics* **64**:1–137. DOI: <https://doi.org/10.1080/00018732.2015.1037068>
- Williams AH**, Kim TH, Wang F, Vyas S, Ryu SI, Shenoy KV, Schnitzer M, Kolda TG, Ganguli S. 2018. Unsupervised Discovery of Demixed, Low-Dimensional Neural Dynamics across Multiple Timescales through Tensor Component Analysis. *Neuron* **98**:1099–1115. DOI: <https://doi.org/10.1016/j.neuron.2018.05.015>, PMID: 29887338
- Wilson FAW**, Scialdhe SPÓ, Goldman-Rakic PS. 1993. Dissociation of Object and Spatial Processing Domains in Primate Prefrontal Cortex. *Science* **260**:1955–1958. DOI: <https://doi.org/10.1126/science.8316836>, PMID: 8316836
- Witschko AB**, Johnson MJ, Iurilli G, Peterson RE, Katon JM, Pashkovski SL, Abaira VE, Adams RP, Datta SR. 2015. Mapping Sub-Second Structure in Mouse Behavior. *Neuron* **88**:1121–1135. DOI: <https://doi.org/10.1016/j.neuron.2015.11.031>, PMID: 26687221

- Wiltschko AB**, Tsukahara T, Zeine A, Anyoha R, Gillis WF, Markowitz JE, Peterson RE, Katon J, Johnson MJ, Datta SR. 2020. Revealing the structure of pharmacobehavioral space through motion sequencing. *Nature Neuroscience* **23**:1433–1443. DOI: <https://doi.org/10.1038/s41593-020-00706-3>, PMID: 32958923
- Wu HG**, Miyamoto YR, Gonzalez Castro LN, Ölveczky BP, Smith MA. 2014. Temporal structure of motor variability is dynamically regulated and predicts motor learning ability. *Nature Neuroscience* **17**:312–321. DOI: <https://doi.org/10.1038/nn.3616>, PMID: 24413700
- Wyrick D**, Mazzucato L. 2021. State-Dependent Regulation of Cortical Processing Speed via Gain Modulation. *The Journal of Neuroscience* **41**:3988–4005. DOI: <https://doi.org/10.1523/JNEUROSCI.1895-20.2021>, PMID: 33858943
- Yao Z**, Shafer OT. 2014. The *Drosophila* circadian clock is a variably coupled network of multiple peptidergic units. *Science (New York, N.Y.)* **343**:1516–1520. DOI: <https://doi.org/10.1126/science.1251285>, PMID: 24675961
- Yeatman JD**, Wandell BA, Mezer AA. 2014. Lifespan maturation and degeneration of human brain white matter. *Nature Communications* **5**:1–12. DOI: <https://doi.org/10.1038/ncomms5932>
- Yu BM**, Cunningham JP, Santhanam G, Ryu S, Shenoy KV, Sahani M. 2008. Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *Advances in Neural Information Processing Systems* **102**:614–635. DOI: <https://doi.org/10.1152/jn.90941.2008>

Appendix 1

From behavioral videos to behavioral features

Several video analysis tools are available to extract behavioral features from videos, ranging from standard computer vision algorithms to deep learning-based methods. These algorithms are fast and reliable and allow to track animal movements at different levels of spatiotemporal resolution, typically yielding three different kinds of time-series outputs. Deep learning algorithms with human-annotated video frames yield a set of coordinates in ambient space representing tracked points on an animal's body (Mathis et al., 2018; Segalin et al., 2020; Pereira et al., 2019). Unsupervised algorithms yield a set of low-dimensional variables obtained from dimensionally reducing the image pixel space via principal component analysis or other nonlinear compression methods (Wiltschko et al., 2015; Batty et al., 2019). Spectral decomposition of the videos obtained from a time-frequency analysis yield probability density maps (Berman et al., 2014; Szigeti et al., 2015). Translating these low-dimensional observables into a meaningful representation of behavior for subsequent analyses requires further assumptions on the nature of feature representations, leading to difficult choices that should be carefully evaluated case by case.

Discrete vs. continuous features

The first choice regards the definition of what constitutes a fundamental unit of behavior for a particular model system and whether this unit is discrete or continuous. Is a 'slow locomotion' bout distinct from a 'fast locomotion' bout? Are they just different manifestation within the variability range of the same 'locomotion' motif? This choice is related to the timescale at which behavior is analyzed to test a particular hypothesis and is based on the level of perceived stereotypy present in the observed behavior (Brown and de Bivort, 2018; Berman, 2018; Pereira et al., 2020a). At the sub-second timescale, movements can be represented as continuous low-dimensional trajectories of body parts. A continuous representation of behavior assumes that behavioral time series can be represented as a superposition of continuous behavioral motifs or 'eigenshapes' and has been fruitful in quantifying the behavior of the worm *C. elegans* and the *Drosophila* larva (Stephens et al., 2008; Szigeti et al., 2015) (although discrete representations have been applied to the same systems; Brown et al., 2013; Vogelstein et al., 2014; Schwarz et al., 2015). Continuous trajectories of movements can be interpreted and investigated as complex dynamical systems as well (Ahamed et al., 2021). In the zebrafish swimming, clear peaks in distribution of kinematic parameters suggest categorical distinctions (Marques et al., 2018), although for some movements the kinematics vary more continuously, allowing the fish to perform a smooth range of swim types. The approach where units of behavior are discrete and can be separated in time leads to representations based on ethograms, namely, graphs where nodes are discrete actions and edges transition rates between them (Bateson and Hinde, 1976). Ethograms implicitly assume that variability across repetitions of the same action ('stereotypy') is smaller than the variability between different actions. In order to bridge the continuous nature of movements with its inherent variability to a discrete interpretation in terms of actions, two successful approaches have been used. When using state-space models (Figures 2A and 5B), the fundamental units of behavior can be captured by stereotyped, low-dimensional, autoregressive trajectories or LDS. The behavioral time series unfolds as a sequence of piecewise LDS trajectories, where each trajectory represents a discrete action, yet it accommodates for large variability in its continuous expression. This approach has been successfully applied both in rodents (with autoregressive HMMs; Wiltschko et al., 2015) and worms (with switching LDS; Linderman et al., 2019). Two of its advantages are the fact that it provides a generative model that can be used to simulate synthetic behavior, and the fully Bayesian implementation of inference. One limitation is that the number of behavioral units must be specified as a hyperparameter, raising the thorny issue of model selection. In freely moving rodents, up to a dozen states were sufficient to capture animal behavior when performing a task (Findley et al., 2021; Parker et al., 2021), but many more states (up to a hundred) were found during spontaneous behavior (Wiltschko et al., 2015; Figure 2A). Sample size and the bias-variance tradeoff are important issues to consider in this class of models. An alternative data-driven approach to segmentation of behavior avoids having to choose the number of units and their dynamical features a priori, but rather lets the structure emerge from the data at different timescales. Using density-based clustering to obtain a distribution of time points in feature space, one can identify peaks in the density map via watershed transforms (Figure 2B). This approach has been successfully used in the adult *Drosophila* (Berman et al., 2014) and rodents (Marshall et al., 2021). The advantage of the watershed transform is that the segmentation into discrete units depends on the level of coarse-graining or the timescale at which the behavior is analyzed, allowing greater flexibility compared to state-space models. Both

approaches combine continuous representations of movements with discrete segmentation into behavioral units. A different data-driven approach used locally linear models to extract behavioral states from posture dynamics (Costa et al., 2019; Costa et al., 2021). Here, behavioral states emerge at different coarse-graining levels through a hierarchical clustering approach combining short-time movements into slower interpretable behaviors, thus surpassing some of the limitations of state–space models, since the number of states is not imposed as a hyperparameter but emerges from the data.

Non-Markovianity and long timescales

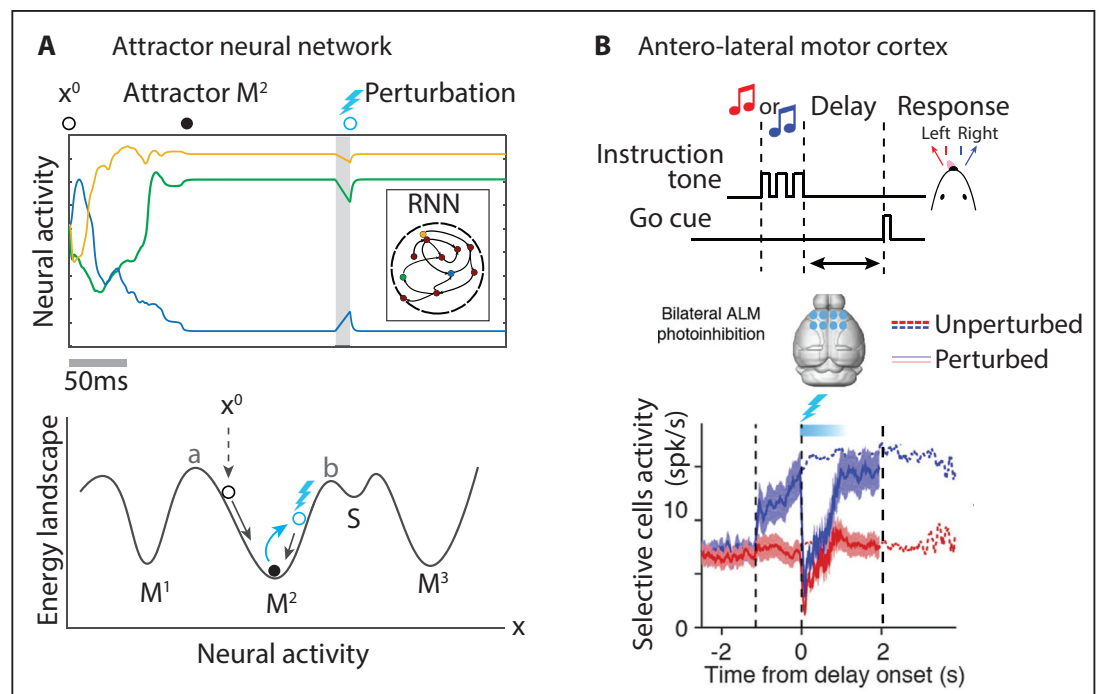
Discrete behavioral representations based on state–space models (such as HMM or SLDS) are convenient and effective for capturing naturalistic behavior with a limited number of actions, represented by discrete hidden states. These models are based on the Markovian assumption that the state of the system at time t only depends on the state at time $t - 1$. A large number of studies used these models to capture a single sub-second timescale related to fast movements emerging from the data, as in the examples of **Figures 2B, 3B and 5A**; while a fast and a slow timescale were found in an olfactory search task, the sub-second duration of single movements and the second-long duration of the investigation/approach sequences (Findley et al., 2021). In principle, a Markov model with N states can capture up to $N - 1$ different timescales, represented by the nontrivial eigenvalues of the transition probability matrix. In these models, transition probabilities between states over longer intervals are generated by powers of the elementary transition probability matrix. However, naturalistic behavior typically exhibits strong deviations from this Markovian structure (Tinbergen, 1951; Dawkins, 1976; Simon, 1962; Marshall et al., 2021), whereby at long intervals much slower timescales emerge, compared to the ones predicted by the Markovian dynamics. A simple workaround is to introduce an ad hoc hierarchical structure in the state–space model (capturing timescale of a few seconds; Tao et al., 2019), although the hierarchical architecture has to be imposed a priori and may reflect human prejudices. An alternative data-driven approach is provided by the information bottleneck, whereby a non-Markovian structure at long timescales can be inferred directly from the data. Using this approach in the fruit fly, a hierarchical representation of actions naturally emerged when aiming to optimally predict the future behavioral state of the fly using the information bottleneck (Berman et al., 2016). An efficient way to model the slow non-Markovian structure is by compressing the behavioral repertoire to just two states and using maximum entropy models to capture the interactions between states. This statistical physics approach revealed that interactions between states at time τ decay as slow as $1/\tau^2$. If one wanted to fit an HMM to data with such slowly decaying interactions, an infinite number of discrete states would be required (Alba et al., 2020). On the other hand, it would be interesting to fit Markov models (such as HMM or SLDS) to behavior for classifying the fast movement motifs and then apply the information bottleneck for compressing them and revealing the underlying slow timescales. This non-Markovian structure in behavior poses an interesting challenge to existing frameworks based on discrete states and calls for new theoretical models to explain the dynamical evolution of naturalistic behavior over long timescales.

Appendix 2

Attractors dynamics in neural circuits

Attractors are a concept in the theory of dynamical systems. An attractor is defined as the set of states a system evolves to starting from a large set of initial conditions, defined as its basin of attraction. The main feature of an attractor is that if the system is perturbed slightly, it tends to return to it. In the context of recurrent neural networks (RNN, **Appendix 2—figure 1A**), attractors are configurations of the network activity that exhibit these features and they can be either time-independent (point attractors, **Appendix 2—figure 1A**), time-varying but periodic (limit cycles), or time-varying but nonperiodic (chaotic attractors). The classic example of an attractor neural network exhibiting point attractors is the Hopfield network (**Hopfield, 1982; Amit et al., 1985**), which was introduced as a model of auto-associative memory. The Hopfield network represents an RNN with N units x_i (for $i = 1, \dots, N$) evolving according to the dynamics $\tau \dot{x}_i = -x_i + \sum_{j=1}^N w_{ij} \phi(x_j)$, where τ and $\phi(x)$ are single-cell time constant and transfer function, the latter typically chosen with a saturating nonlinearity in neuroscience models (**Hopfield, 1984**). In a Hopfield network, the synaptic couplings $w_{ij} = \sum_{\mu=1}^p M_i^\mu M_j^\mu$ encode the network attractors $\{M_i^\mu\}_{\mu=1}^p$. After initializing the network activity at a random value x_i^0 (empty circle in **Appendix 2—figure 1A**), the network temporal dynamics quickly converge to the attractor M^μ , which is closest to x^0 (black circle in **Appendix 2—figure 1A**). The set of initial conditions that converge to a particular attractor defines its basin of attraction, represented as potential wells in an energy landscape (in **Appendix 2—figure 1A**, the basin of the attractor of M^2 is represented by the values of x such that $a < x < b$). If we interpret this model as an auto-associative memory circuit, then each attractor M^μ represents a stored memory; when presenting the network with a hint (x^0) that is close to the memory M^2 , the network will quickly retrieve the full memory. A crucial feature of attractor dynamics is that when we slightly perturb the network activity away from the attractor (blue lightning in **Appendix 2—figure 1A**), then if the new configuration (empty blue circle) still lies within the basin of attraction of M^2 , the network will quickly converge back to the attractor.

It is possible to generalize these Hopfield-like stable attractors to biologically plausible networks of spiking neurons with cell-type-specific connectivity (**Amit and Brunel, 1997**). These models explain the features of neural activity observed in several areas of the association cortex (inferotemporal cortex [**Fuster and Jervey, 1981; Miyashita and Chang, 1988**]; prefrontal cortex [**Funahashi et al., 1989; Wilson et al., 1993**]; posterior parietal cortex [**Koch and Fuster, 1989**]). In these experiments featuring a delay–response task, a stimulus is presented and then withdrawn, followed by a delay period at the end of which the monkey performs a stimulus-specific saccade to obtain a reward. Neural activity in association areas exhibits stimulus-selective persistent activity during the delay period following the removal of the stimulus, which can be interpreted as originating from attractors in the underlying neural circuit. Such self-sustained persistent activity preserves an active memory of a visual stimulus after it is removed. Long-lasting persistent activity was also reported in the fruit fly following optogenetic stimulation and connectome reconstruction, suggesting that this activity may be supported by strong recurrent couplings between persistently active neurons (**Deutsch et al., 2020**). Recent advances with optogenetics in rodents further confirmed that delay persistent activity is stable against perturbations, a crucial feature of attractor dynamics. In a delay–response task (**Inagaki et al., 2019; Finkelstein et al., 2021; Appendix 2—figure 1B**), neural activity recorded in the anterolateral motor cortex (a part of the secondary motor cortex) showed persistent choice-selective activity during the delay period (dashed lines in **Appendix 2—figure 1B**). In trials where bilateral optogenetic perturbations transiently silenced the local circuit at the beginning of the delay epoch (full lines in **Appendix 2—figure 1B**), stimulus-selective population activity quickly recovered to the unperturbed values (full lines in **Appendix 2—figure 1B**), suggesting that delay activity in this area can be interpreted as originating from attractors.



Appendix 2—figure 1. Neural attractors. **(A)** In an attractor neural network, the activity of three representative units starting from initial values x^0 rapidly converges to the closest attractor (M^2). A small perturbation displaces the activity away from the attractor (blue circle) and the dynamics quickly returns to the attractor. Bottom: attractors are represented as potential wells in an energy landscape. Initial conditions or perturbations within the basin of attraction of the attractor M^2 (defined by $a < x < b$) quickly converge back to the attractor. **(B)** In a head-fixed delay–response task (top), one of two tones was presented during the sample epoch and the mouse reported its decision after the delay epoch by directional licking (randomized delay duration). The anterolateral motor cortex was photoinhibited bilaterally during the first 0.6s of the delay epoch (cyan bar) and quickly returned to its unperturbed level. Mean spike rate of lick-right preferring neurons (blue and red thick lines: unperturbed lick-right and -left trials; dashed lines: perturbed trials). Panel **(B)** adapted from Figures 1, 6, and S6 of *Inagaki et al., 2019*.

Appendix 3

Metastable attractors

Although in the models of memory and delay persistent activity discussed so far attractors are stable, and noise originating from some stochastic process may generate spontaneous transitions between them. The hallmark of these noise-driven metastable dynamics is a right-skewed distribution of attractor dwell times with an exponential tail. The source of this stochasticity can be either external (i.e., incoming Poisson spike trains relaying neural activity from other brain areas) or intrinsically generated from within the local recurrent circuit. In the case of recurrent networks of rate units (**Figure 3D**), metastable dynamics can be triggered by low-dimensional correlated variability arising from synaptic dynamics in a feedback loop between a cortical and subcortical area (**Recanatesi et al., 2022**). In the case of a spiking circuit, intrinsic variability may arise from the asynchronous irregular spiking activity ubiquitously observed across many brain areas (**Softky and Koch, 1993; Shadlen and Newsome, 1998**). The asynchronous irregular regime is a hallmark of neural circuits operating in a fluctuation-driven regime (**Amit and Brunel, 1997**) and balanced networks (**van Vreeswijk and Sompolinsky, 1996**). The stochasticity arising from irregular spiking activity can generate noise-driven transitions between stable attractors and represents a robust and parsimonious mechanism to explain the origin of metastable attractors (**Deco and Hugues, 2012; Litwin-Kumar and Doiron, 2012; Mazzucato et al., 2015; Wyrick and Mazzucato, 2021**). This class of biologically plausible models is based on an architecture whereby excitatory and inhibitory neurons are arranged in neural clusters, or functional assemblies, where neurons belonging to the same cluster have potentiated synaptic couplings compared to neurons belonging to different clusters (**Appendix 3—figure 1A**). During ongoing periods, in the absence of external stimulation, the network activity unfolds via a sequence of metastable attractors, where each attractor is defined by a subset of activated clusters (**Appendix 3—figure 1B**). Using mean field theory, one can represent attractors as potential wells in an attractor landscape (**Mascaro and Amit, 1999; Litwin-Kumar and Doiron, 2012; Mattia and Sanchez-Vives, 2012; Wyrick and Mazzucato, 2021**). The attractor landscape can be calculated analytically using a Lyapunov function (in networks with symmetric synaptic couplings the only minima are point attractors **Appendix 3—figure 1A**), although in general more complex minima may arise (**Brinkman et al., 2022**). In a simplified network with two clusters (**Appendix 3—figure 1C**), each well corresponds to a configuration where one cluster is active and the other is inactive. When the network activity dwells in the attractor represented by the left potential well, it may escape to the right potential well because of internally generated variability. This process will occur with a probability determined by the height Δ of the barrier separating the two wells: the higher the barrier, the less likely the transition (**Hänggi et al., 1990**). The mean first passage time for this process is captured by the Arrhenius law (**Huang and Doiron, 2017**), which can be generalized to the more realistic case of transitions driven by colored noise (**Stern et al., 2021**).

Fast and slow fluctuations

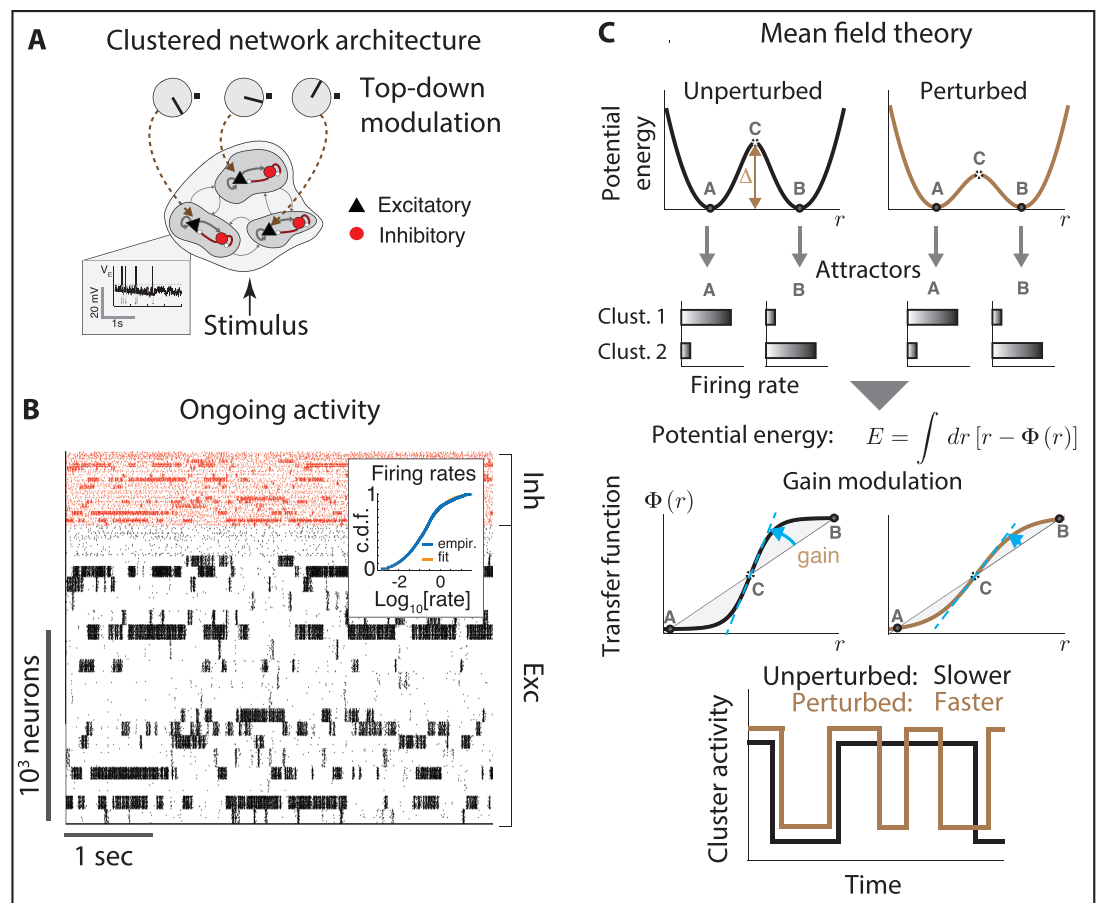
Given a multistable landscape with multiple fixed points, such as the one generated by clustered networks, there are specific requirements on the statistics of noise fluctuations in order to generate metastable dynamics. Within a mean field approach, we can approximate the on-off switches of a single cluster as driven by escape noise from the energy potential (**Appendix 3—figure 1C**). There are at least three sources of noise fluctuations in this spiking circuit, contributing to the input current to the representative cluster: the fluctuations due to the finite size of neural clusters, the shot noise generated by the incoming Poisson spike trains from other recurrently coupled neurons, and the slow rate fluctuations due to the metastable dynamics of the other clusters in the network. Finite size effects can be modeled as white noise, whose amplitude is inversely proportional to the cluster size (**Brunel, 2000; Huang and Doiron, 2017**). This fast noise source is sufficient to account for the emergence of a slow metastable dynamics of the cluster's on-off switches, captured by Kramers theory of noise-induced escape in a bistable system (**Hänggi et al., 1990; Huang and Doiron, 2017; Litwin-Kumar and Doiron, 2012**). In this simple picture, increasing the cluster size has the dual effect of both reducing the noise amplitude and also increasing the potential barriers. However, the picture significantly changes if we account for the other two sources of noise, which are both colored. In the diffusion approximation for the spiking network, an incoming spike train filtered by exponential synapses gives rise to an Ornstein-Uhlenbeck process, whose color is set by the synaptic time constant and whose amplitude is proportional to recurrent coupling strength (**Fourcaud and Brunel, 2002**). Finally, the self-consistent rate fluctuations originating from the metastable dynamics of the other network clusters give rise to a colored noise whose slow timescale is proportional to the average switching time (**Stern et al., 2021**). In order to deal with escape colored noise, we

can replace the Kramers theory (*Hänggi et al., 1990; Huang and Doiron, 2017; Litwin-Kumar and Doiron, 2012*) with the universal colored noise approximation (UCNA) to the Fokker–Planck equation, which provides a good approximation of the escape process in this regime (*Stern et al., 2021*). Here, both the strength of the noise amplitude and the barrier height are proportional to a product of the noise color and the recurrent coupling strength. This theory shows that the overall effect of colored noise is to slow down the average transition rate. Although UCNA has been applied to networks of rate units, a full account of the spiking network dynamics is still an open problem. In summary, in models with metastable attractors, we can find a separation of timescales between the fast fluctuations around a particular attractor and the slow fluctuations involving switching between different attractors. Moreover, network simulations suggest that the details of the network architecture might play an important role in generating metastable dynamics over a large range of parameters. Whereas in networks with only excitatory but not inhibitory clusters the metastable regime might require fine-tuning (*Litwin-Kumar and Doiron, 2012; Mazzucato et al., 2019*), an architecture with coupled excitatory and inhibitory clusters allows metastable dynamics to persist over a wide range of synaptic coupling strength and cluster sizes (*Rostami et al., 2020; Wyrick and Mazzucato, 2021*). This E/I clustered architecture can lead to the emergence of a heterogeneous distribution of timescales within the same circuit, as explained in *Figure 8E* (*Stern et al., 2021*).

Top-down modulation of network metastable dynamics

A central challenge in experimentally testing computational models of metastable attractors is that the reconstruction of the attractor landscape and the height of its potential wells relies on the knowledge of the network's structural connectivity including its coupling strength and connection probability, which is out of reach of current neurotechnology. However, one can use mean field theory methods to obtain an alternative formulation of the attractor landscape only involving quantities directly accessible to experimental observation (*Mascaro and Amit, 1999; Mattia et al., 2013; Mazzucato et al., 2019; Wyrick and Mazzucato, 2021*). The double potential well representing the two attractors can be directly mapped to the effective transfer function of a neural population representing a single cluster (*Appendix 3—figure 1C*). Using this correspondence, one can map changes in the barrier height Δ separating metastable attractors to changes in the slope (or 'gain') of the intrinsic transfer function estimated during ongoing periods. This map provides a direct relationship between changes in cluster activation timescale and single-cell gain modulation, which can be inferred from either intracellular or extracellular recordings (*Lim et al., 2015; Wyrick and Mazzucato, 2021*). These computational tools can be used to test mechanistic hypothesis on the effect that circuit perturbations and manipulations exert on a network metastable dynamics (*Mazzucato et al., 2019; Wyrick and Mazzucato, 2021*). In particular, perturbations inducing steeper gain increase well depths and barrier heights, and thus increase the cluster timescale, and vice versa. This approach assumes that fluctuations in synaptic inputs are constant and additive, such that we can interpret the changes in first passage times as coming from changes in the transfer function. Transition rates between attractors may also be modulated by varying the levels of amplitude and color of the synaptic inputs, which can arise either from recurrent dynamics or from external projections. An increase in the noise amplitude can naturally arise in balanced E/I networks by simultaneously increasing both the E and I activity or E and I external inputs, leading to noise-induced gain modulation and a faster integration time in conductance-based models (*Chance et al., 2002*). Variations in multiplicative synaptic noise within thalamocortical loops (e.g., in the model of *Recanatesi et al., 2022*) may also modulate average transition times between metastable attractors, although this mechanism has not been explored yet. Variations in the noise color were also shown to affect the transition rate between potential wells (*Stern et al., 2021*), where a slower noise timescale led to slower transition times compared to white noise.

An exciting new direction aims to infer the dynamical system underlying the observed spiking activity directly from the data. Assuming that high-dimensional neural activity is generated by a low-dimensional set of latent variables evolving according to Langevin dynamics, one can directly infer the effective energy potential underlying their dynamics from the data (*Genkin and Engel, 2020; Duncker et al., 2019*). A potential limitation of these new methods may be their reliance on low-dimensional latent manifolds, whereas the structure of neural patterns in motor areas during freely moving behavior seems to be high dimensional (*Recanatesi et al., 2022*).



Appendix 3—figure 1. Gain modulation regulates metastable attractor dynamics. **(A)** Biologically plausible model of a cortical circuit generating metastable attractor dynamics. A recurrent network of E (black triangles) and I (red circles) spiking neurons arranged in clusters is delivered sensory stimuli and top-down modulation representing changes in cortical states. Inset: membrane potential trace from a representative E neuron. **(B)** Representative neural activity during ongoing periods (tick marks represent spike times of E [black] or I [red] neurons). The network activity unfolds through metastable attractors, each attractor corresponding to subsets of transiently active cluster. Inset: the cumulative distributions of single-cell firing rates is lognormal. **(C)** The effect of perturbations on network dynamics in a two-cluster network can be captured by a double-well potential (top). Potential wells represent two attractors where either cluster is active (A, B), separated by a barrier with height. Mean field theory links perturbation effects on the barrier height (top right, lower barrier) to changes in the intrinsic neuronal gain (center right, lower gain). Bottom: a perturbation that decreases barrier heights and gain leads to faster transitions between attractors. All panels adapted from Figures 2 and 3 of *Wyryck and Mazzucato, 2021*.