Unsupervised quantification of naturalistic animal behaviors for gaining insight into the brain

Michael H. McCullough[1] and Geoffrey J. Goodhill[1,2]

Abstract

Neural computation has evolved to optimize the behaviors that enable our survival. Although much previous work in neuroscience has focused on constrained task behaviors, recent advances in computer vision are fueling a trend toward the study of naturalistic behaviors. Automated tracking of fine-scale behaviors is generating rich datasets for animal models including rodents, fruit flies, zebrafish, and worms. However, extracting meaning from these large and complex data often requires sophisticated computational techniques. Here we review the latest methods and modeling approaches providing new insights into the brain from behavior. We focus on unsupervised methods for identifying stereotyped behaviors and for resolving details of the structure and dynamics of behavioral sequences.

Addresses
[1] Queensland Brain Institute, The University of Queensland, Brisbane, Queensland, 4072, Australia
[2] School of Mathematics and Physics, The University of Queensland, Brisbane, Queensland, 4072, Australia

Corresponding author: Goodhill, Geoffrey J (g.goodhill@uq.edu.au)

Introduction

Neural circuits can only be fully understood in the context of their behavioral outputs [1]. Quantitative models of behavior can therefore serve as a foundation for studying the function, development, and evolution of the brain [2]. Relationships between brain and behavior are often investigated through experiments in which animal behavior is constrained, such as maze navigation or cue-evoked licking of water spouts while head-fixed. These paradigms facilitate well-controlled repeatable trials and relatively straightforward data analysis, but only partly represent the diversity and variability of natural movement, environmental cues, and biological imperatives [1,3]. An important goal for neuroscience is therefore to quantify behavior in more naturalistic settings [4,5].

However, constructing useful models of naturalistic behavior remains a challenging problem both technically and conceptually [1,4,6–10]. For example, an experiment involving one or more animals moving and interacting freely in an open arena is often a process that spans multiple spatial and temporal scales, with many degrees of freedom, and with dynamics that may be driven by unobservable factors [1,6]. Traditional ethological approaches generally depend on human interpretation to identify features of interest in complex behavior. However, manual annotations can be inconsistent between observers [11] and behaviors that may appear random to a human observer can still have underlying structure [12]. Alternatively, mathematical models built from first principles can provide useful and precise abstractions of many complex biological phenomena [12,13], but often rapidly become intractable as dimensionality increases [14]. To achieve consistency, accuracy, and scalability in the era of big behavioral data, neuroscience needs unsupervised and unbiased computational tools for quantifying and modeling behavior based on the data itself [1,4,8,15].

Here we review recent advances in quantitative data-driven approaches for understanding naturalistic behaviors, which draw from a diverse range of fields including machine learning, stochastic modeling, statistical physics, information theory, signal processing, and dynamical systems theory. First, we address unsupervised methods for behavioral mapping: the process of segmenting continuous-time behavioral data into a set of discrete elementary behavioral units. Second, we discuss modeling behavioral sequences and examine frameworks for extracting stereotyped motifs from data over multiple timescales. We review approaches for identifying evidence of underlying hierarchical structures and hidden states modulating patterns of behavior. Finally, we consider continuous models, and emerging frameworks for bridging continuous and discrete behavioral representations. Summaries of some key methods (Box 1), models (Box 2), and terms (Box 3) are included for reference.

Box 1. List of key methods.

Affinity propagation A clustering algorithm based on the concept of message passing [62].

Change-point detection A broad class of methods for identifying epochs in time series data corresponding to different dynamics. Simple change-point detection typically depends on extracting a relevant coarse measure and applying a threshold or identifying peaks, whereas more sophisticated approaches fit statistical models to different temporal segments of the data and identify epochs based on differences in fitted parameters (i.e., locally linear models [66]). Applicable for segmenting behavioral time series data when (assumed) elementary behavioral units have varying duration and form relatively continuous sequences.

Delay embedding A method for reconstructing a representation of phase space (the set of all possible states in a dynamical system) from experimental data where the underlying system variables and mathematical equations are unknown [95−98].

Dictionary-based compression Algorithmic methods which find stereotyped motifs in a sequence by compressing data using information theory. Input must be a sequence of discrete states (e.g., a sequence of elementary behavioral units). These algorithms have been applied to identify and characterize stereotyped patterns of behavior over multiple timescales [30,82,83].

Dynamic time warping An algorithm for measuring the similarity between short time series segments whereby trajectories are first nonlinearly stretched or compressed in time (within some constraints) such they are aligned as best as possible [58]. This method is robust to small changes in alignment, duration, rate and other temporal distortions, and has been applied to compare stereotyped patterns of movement [57].

Eigendecomposition Decomposition of a matrix into its eigenvectors and eigenvalues. Also known as spectral decomposition. In the specific case of a Markov chain transition matrix, the eigenvalues characterize various properties of the model such as the time horizon for predictability [47].

Gaussian mixture models A method for clustering which models data as a set of sub-populations each drawing from a different multivariate Gaussian distribution.

Information bottleneck compression A compression algorithm which can be applied for clustering data [47,86].

Isomap A nonlinear dimensionality reduction technique that aims to preserve local manifold structure by measuring geodesic distances between points on a graph embedding of the data [40].

Minimum description length An approach to model selection from information theory which seeks to find the most compact description of the data [104].

Principal components analysis (PCA) Finds linear projections in multi-dimensional data which capture maximum variance. Commonly used as a dimensionality reduction technique for high-dimensional data.

t-distributed stochastic neighbor embedding (t-SNE) A nonlinear dimensionality reduction technique which aims to preserve local structures [38] (i.e., points that are similar in high-dimensional space should be close in the embedding).

Uniform manifold approximation and projection (UMAP) A nonlinear dimensionality reduction technique which assumes that data lie on a smooth, locally-connected manifold [39].

Wavelet transform (time series analysis) Estimates the frequency content and phase information in a sliding window along a time series using wavelets [51]. If the phase information is discarded, the resulting power spectral density captures linear properties of the time series trajectory in a manner that is robust to small changes in alignment, and is hence advantageous for dealing with the natural variability in timing and duration of stereotyped movements [24].

Mapping discrete stereotyped behaviors in big behavioral data

The first stage of behavioral mapping is feature selection. Video recordings of freely moving behavior typically contain large amounts of redundant information, and the choice to observe specific aspects of kinematics will be guided by the behavioral hypothesis under investigation. If discrete elementary units of behavior exist in the data, then inferring them entails identification of stereotyped patterns occurring over relatively short periods of time that span some subset or combination of the observed variables. The resulting map should be sufficiently discrete to facilitate consistent classification of stereotyped behaviors despite variability between realizations and between animals. Most methods we consider can be broadly categorized as either a sequential pipeline of temporal segmentation followed by low-dimensional embedding and unsupervised clustering, or direct fitting of a generative model simultaneously across time and feature space (Figure 1).

Extracting features from behavioral recordings

A natural choice of measure is change in an animal's posture over time (Figure 1a). Recent developments in computer vision using deep learning now enable automated marker-less posture tracking of animals with complex limb structures and movement patterns [16−20], reviewed in Pereira et al. [11] and Mathis and Mathis [21]. These methods now extend to real-time tracking for integration in closed-loop experimental paradigms [22,23]. A postural representation can also be encoded directly from aligned and cropped images using principal component analysis (PCA) [24,25], or nonlinear techniques for dimensionality reduction such as convolutional autoencoders [26,27]. Fine-scale postural data can be augmented with coarse kinematic variables to produce a combined time series of postural dynamics and whole-body kinematics [28−31]. An additional common behavioral measure is sound, as observed from vocalizations in birdsong [32−34] and wing vibrations in flies [7,35]. More generally, the first stage of behavioral measurement is an explicit choice guided by the behavior

Box 2. List of key models.

Autoencoder A class of neural network architectures for learning efficient representations of data and often used for dimensionality reduction. Autoencoders typically comprise an input (encoding) layer which maps data to an internal representation, and an output (decoding) layer which aims to reconstruct the original data as accurately as possible based on the representation.

Autoregressive process A random process model for time series data where each observation is modeled as a weighted linear sum of previous observations and a stochastic error term.

Autoregressive hidden Markov model (AR-HMM) A class of hierarchical generative models for continuous time series data that assumes a system transitions through a series of discrete states or regimes, each modeling the continuous dynamics as a linear autoregressive process with its own distinct parameters [73]. These models have been applied to investigate the hypothesis that the brain encodes short behavioral patterns as motifs which can be combined together into sequences [25].

Convolutional autoencoder (CAE) An autoencoder with a convolutional input layer that is able to preserve spatial correlations [27]. CAEs are well suited to learning structures in imaging data and have been applied to extract behavioral patterns directly from videos [26].

Hidden Markov model (HMM) A model of an observed variable that is generated based on a set of unobservable discrete states which form a Markov chain. The goal when fitting an HMM is typically to estimate the hidden state of the system for each observation [74].

Ising model A model from the field of statistical physics which describes complex emergent phenomena from a large set of simple binary units which interact locally in a lattice or grid configuration. The model was originally developed to investigate properties of magnetic materials.

Marked renewal process A stochastic model of a sequence of events that occur with varying arrival times. Each event has its own "mark" or type. The type and timing of the next event is conditioned on past events [31]. Can be applied to model complex behavioral sequences which depend on short-time and long-time internal state variables and sensory inputs [31].

Markov chain A discrete state transition model where a system moves to its next state based on fixed probabilities that depend only on the current state (i.e., the system has no memory). This is also known as a first-order Markov chain. The transition probabilities for an n^{th} order Markov chain depend on the history of the last n states visited by the system.

Switching linear dynamical system (SLDS) A class of hierarchical generative models for continuous time series data [77,78] with high-level similarities to AR-HMMs. SLDSs have been extended such that the top level of the hierarchy can model a continuous state variable [72] which has potential for modeling homeostatic processes that modulate behavior.

Variational autoencoder (VAE) A class of generative models for learning latent structure in data [41−44]. Like standard autoencoders, VAEs comprise a pair of encoding and decoding models. However, the mathematical formulation of these models is very different in a VAE, generally imposing stronger assumptions with the aim of learning semantically meaningful representations. VAEs are finding new applications in dimensional reduction for simultaneous segmentation and mapping of behavioral time series [45,46]

of interest—a process of feature engineering that can be aided by machine learning and computer vision.

Segmenting, embedding, and clustering

If we assume that a behavioral sequence corresponds to a series of short nonoverlapping trajectories through the space of the behavioral observations (i.e. trajectories through postural space), one approach for mapping such units is to extract these trajectories by temporal segmentation, and then cluster them into discrete groups based on some measure of similarity (Figure 1b). Given that clustering in high-dimensional space is intrinsically challenging [36,37], the dimensionality of the behavioral data is typically first reduced by PCA or nonlinear embedding algorithms such as t-distributed stochastic neighbor embedding (t-SNE) [38], Uniform manifold approximation and projection (UMAP) [39], or Isomap [40]. Alternatively, variational autoencoders [41−44] are also emerging as a promising new approach for low-dimensional embedding of complex behavioral data [14,45,46]. Despite the conceptual simplicity of the segment-embed-cluster approach, the complexity and variability of natural behavior poses an inherent challenge: different realizations of a stereotyped behavior are seldom identical in the precise shape and timing of their trajectories, nor is there always clear demarcation from one movement pattern to the next. Therefore, methods for segmenting and clustering elementary

units of behavior should ideally be robust to variation between realizations of the same behavior, while simultaneously sensitive to features which discriminate between those that are functionally distinct.

Methods and applications

One approach for unsupervised behavioral mapping is MotionMapper [24,47,48], a framework originally developed to study free walking behaviors of *Drosophila melanogaster* and subsequently extended for the analysis of social behaviors between pairs of animals [49,50]. Rather than attempting to directly segment the postural time series, a continuous wavelet transform [51] is applied along the data to encode short-time trajectories in the frequency domain. By retaining only amplitude information and ignoring phase the encoding is robust against small temporal misalignments. Clustering into discrete behavioral units is achieved by smoothing a 2-dimensional t-SNE embedding and then applying a watershed algorithm. To overcome computational limitations resulting from the large size of their data, the authors implemented a variant of t-SNE which learns a low-dimensional embedding for a representative subset of training data, then maps the remaining points onto this embedding in an additional step. MotionMapper has also been applied for identifying distinct song modes produced by *Drosophila* [52], where short time series

---

Box 3. Glossary of terms.

Attractor A set of states in a dynamical system which trajectories evolve towards or along (e.g., periodic cycle, manifold, fixed point etc.).

Bayesian information criterion A metric for model selection from Bayesian probability which seeks a balance between model performance and model complexity.

Chaos The property of apparent randomness in some deterministic dynamical systems which is caused by a sensitive dependence on initial conditions. In chaotic systems, trajectories which begin infinitesimally close will diverge exponentially fast. Given that the state of a continuous system cannot be measured with infinite precision, chaotic dynamics are unpredictable beyond a finite time horizon (e.g., the motion of a double pendulum).

Chaotic itinerancy A phenomenon in high-dimensional chaotic systems whereby trajectories sequentially jump between a series of meta-stable attractors, and which is hypothesized to explain aspects of neural coding and computation [99].

Generative model A broad class of machine learning models which learn to generate new data similar to training data, based on assumptions about how the original data were created [105]. In this way, they can be used to extract interpretable features from test data during classification or prediction tasks. By contrast, discriminative models only seek to find differences (or similarities) between data points.

Hyperparameter A parameter which is specified rather than learned, inferred or computed during model fitting. For example, many clustering algorithms require the user to specify the number of clusters as an input value, hence this is a hyperparameter.

Latent state Variables in a model which are not directly observed from data and therefore must be inferred via statistical methods.

Lyapunov exponents A measure of the rate of separation of nearby trajectories in a dynamical system. Lyapunov exponents are a characteristic measure for chaos, and quantify the timescale on which a chaotic system becomes unpredictable [96–98].

---

segments were first time aligned by peak energy and normalized over amplitude, rather than performing a wavelet transform. Marshall et al. [53] adapted aspects of MotionMapper into the CAPTURE framework for mapping naturalistic behavior in rodents over week-long periods of observation based on 3-dimensional postural time series. By additionally incorporating supervised feature selection, these authors were able to enhance the discriminability of known behaviors, though at the risk of adding human bias.

Another model system well suited for the study of naturalistic behavior is the larval zebrafish (*Danio rerio*). These swim in discrete bouts [54–56], providing a convenient solution to the problem of temporal segmentation. However, swim bouts vary in duration, precluding direct comparison of postural trajectory segments. Even qualitatively similar bouts are unlikely to have precise phase alignment of tail oscillations. Mearns et al. [57] addressed the problem of temporal alignment using dynamic time warping [58]. Because postural dynamics are often low dimensional [9,15,59–61], Mearns et al. first applied PCA to the tail kinematics of the fish to reduce dimensionality. Pairwise similarities between time-aligned bouts were embedded by a combination of affinity propagation [62] and Isomap embedding [40]. Although acknowledging apparent continuity in the low-dimensional embedding, Mearns et al. applied hierarchical clustering to extract an ethogram of seven characteristic bout types used during hunting. Alternative studies have identified 13 distinct bout types in larval zebrafish behavioral space [28,29] using a novel density-based clustering algorithm [63]. Although the difference in the number of bout types from Mearns et al. [57] might be partly explained by the wider range of stimulus conditions investigated in Marques et al. [28], it highlights more generally that there is rarely one uniquely correct answer for the number of clusters present in complex high-dimensional data.

Unlike larval zebrafish, behavior in worms such as *Caenorhabditis elegans* and *Drosophila* in the larval stage consists largely of continuous patterns of movement, making temporal segmentation challenging. To address this issue, Szigeti et al. [64] devised a change-point detection algorithm based on the local minima of smoothed and summed postural time series. They then applied a probabilistic algorithm for simultaneous temporal alignment and clustering [65] to the resulting trajectory segments, and the Bayesian information criterion to estimate an optimal number of clusters. This method provided not only cluster labels but also probabilities of cluster membership. Based on these probabilities, Szigeti et al. [64] found ambiguity in classification for many movement types, suggesting that worm locomotion may be better described by a continuous representation rather than a discrete behavioral map. In an alternative approach for change-point detection, Costa et al. [66] segmented and mapped elementary behavioral units of varying lengths in *C. elegans* postural time series by sequentially fitting locally linear models to the data and then clustering model parameters. This particular model can also be interpreted as a discrete approximation of a continuous dynamical system. Given the continuous nature of movement patterns in *C. elegans* behavior, methods tailored to this regime can provide complimentary perspectives (see 'Continuous manifolds and chaos' section).

Practical considerations and limitations
The suitability and interpretability of a behavioral mapping pipeline hinges on the combined set of assumptions implicit from every algorithm and

Figure 1
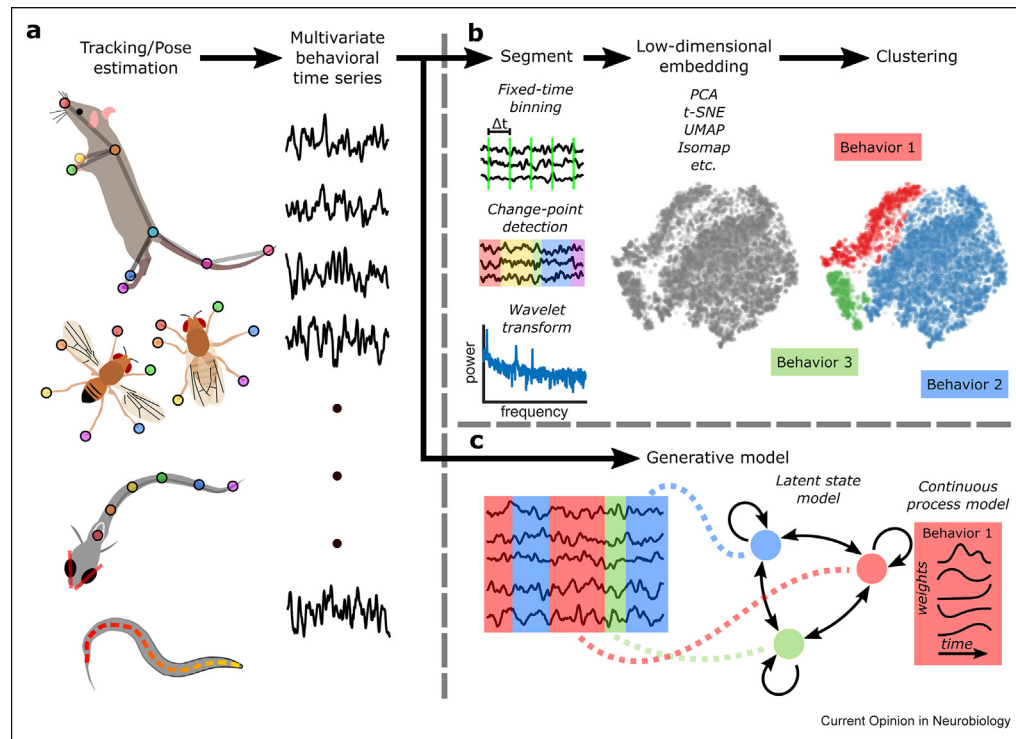
Typical frameworks for unsupervised behavioral mapping from naturalistic behavioral data: (**a**) Computer vision and image processing techniques encode postural information into time series representing continuous recordings of behavior. These data are often augmented with additional features such as coarse kinematics (velocity, acceleration, etc.) or other relevant signals (environmental cues, vibration recordings). (**b**) Elementary units of behavior can be extracted by capturing information about short-time trajectories via temporal segmentation or applying the continuous wavelet transformation, followed by low-dimensional embedding and clustering. (**c**) Alternatively, elementary units of behavior can be extracted by fitting a generative model to the data, which is often hierarchical by construction such that each elementary behavior is modeled by a continuous process and transitions between behaviors are governed by a discrete or continuous latent state model.

hyperparameter choice in the pipeline. Extracting features from fixed-length time bins [53,67] directly imposes a timescale on behavioral units which may not reflect natural variability in the timing of behaviors. Change-point detection [64,66] or the use of features which are robust to small temporal variation [24,57,65] are advantageous approaches in this regard. Where the extracted behavioral features are low dimensional, clustering can be applied directly without the need for dimensionality reduction. For example, Klaus et al. [67] constructed discrete feature histograms from low-dimensional mouse pose data and clustered via affinity propagation to identify behavioral units that correlated with neural coding in spiny projection neurons. When dimensionality reduction is required, as is often the case, it is important to test the effect of different hyperparameters—such as the perplexity parameter for t-SNE or the number of neighbors parameter for Isomap and UMAP—as they can drastically alter results. Furthermore, these algorithms often incorporate stochasticity, meaning they will generate a different embedding each time that the algorithm is run. t-SNE

and UMAP are also sensitive to initialization [68] which should be considered when selecting implementations of these algorithms.

The primary challenges when clustering data are selecting an algorithm, specifying a distance metric, and specifying the number of clusters (or a hyperparameter which indirectly controls this, such as the width of the Gaussian smoothing kernel used in MotionMapper [24]; see Ronan et al. [37] for a practical discussion of clustering for biological data). At a minimum, an appropriately selected clustering validation metric should be applied to assess the quality of the obtained clusters, such as the silhouette score [69], the Calinski-Harabasz criterion [70], or the gap statistic [71]. It is typically more effective to perform ensemble clustering: clustering the data many times while perturbing both hyperparameters and the data (by the addition of noise or bootstrapping, etc.) to identify solutions which are stable and consistent [37]. The dimensionality reduction step can also be incorporated within an ensemble clustering framework to ensure stability for stochastic

algorithms and embedding hyperparameters. Finally, if it is appropriate to use human-labeled data as a ground truth for a given application, then results for a range of algorithms and parameters can be benchmarked against these labels. However, such ground-truth labels can be very costly to generate, and this approach is effectively supervised learning which may impart human bias to the resulting behavioral map.

Generative models

Elementary units of behavior can also be extracted by fitting generative models, which can be configured to describe different levels of complexity in a hierarchical construction (Figure 1c) [25,26,30,35,72]. Generative models enable simultaneous global fitting of the data to automatically identify behavioral units and label the time series accordingly. However, increasing model complexity requires increasing amounts of data for robust fitting of parameters, and necessitates sophisticated approaches for model selection, regularization, and validation to avoid over-fitting.

Wiltschko et al. [25] developed MoSeq, an autoregressive hidden Markov model (AR-HMM) [73], to study free-moving behavior of rodents from depth imaging data. They proposed a two-level hierarchical description of the data, where behavioral units correspond to short-time postural motifs captured by a set of continuous AR processes. An HMM [74] then describes the sequence of transitions between behavioral units. Although the AR-HMM has the flexibility to explain behaviors with different duration, a 'sticky' timescale parameter acts as a lens to focus the model on behaviors of a specific temporal scale. In Wiltschko et al. [25], MoSeq extracted approximately 60 distinct behavioral units from recordings of free-moving mice which explained more than 95% of the data. Furthermore, treating the probability distribution over the set of units as a summary statistic was shown to be significantly more effective than traditional coarse behavioral statistics for discriminating between neuroactive and psychoactive drugs in phenotypic drug discovery experiments [75].

Generative models also offer intriguing possibilities for combined modeling of neural activity and behavior. For example, a convolutional autoencoder can be paired with an AR-HMM to segment behavioral units directly from video data [26]. Coupled with a decoder of simultaneously recorded neural activity from mice, this model could generate videos of estimated behavioral output based on neural activity alone. Switching linear dynamical systems [76–78] have been extended to model data as a three-tier hierarchy comprising a continuous latent state at the top level, which modulates the properties of a discrete state sequence that then controls continuous-time dynamics [72]. This form

of model could have future applications for studying behaviors that are modulated by hidden internal states such as circadian rhythms or hunger, which may be better modeled as continuous variables rather than discrete states.

Characterizing discrete behavioral sequences

Animal behavior is structured over many timescales [1,4,79], a property long theorized to be the result of underlying hierarchical organization guiding patterns of behavior [80]. In this section, we discuss approaches for characterizing temporal dependencies and structures in behavioral sequences which may manifest as motifs (i.e. stereotyped sequences of elementary behavioral units) or reflect internal state changes.

State transition models

Where naturalistic behavioral data can be mapped to a sequence of discrete units then the probabilities of transitioning from one behavior to the next can be estimated. For example, MoSeq behavioral mapping combined with neural activity data showed that lesioning the dorsolateral striatum in mice changed the utilization of different behaviors and increased the complexity of behavioral sequences without altering the underlying set of behavioral units [81]. Markov chains are a useful generalization of this concept for identifying timescales that may be relevant for a behavior [30,47,53,57,82–84]. In an application to larval zebrafish hunting sequences, Mearns et al. [57] showed that second-order Markov chains were only marginally more predictive than first-order Markov chains, and only for a subset of bout types, suggesting a memory-less stimulus–response loop driving the behavior on short timescales. In contrast, eigendecomposition of Markov chains constructed from free-walking *Drosophila* revealed long-timescale behavioral correlations [47]. These authors found evidence of hierarchical structure which may explain this observation, and aligns with prominent theories about the organization of behavior more generally [80,85]. This was achieved by applying information bottleneck compression [86] and using a bespoke 'tree-ness' measure to quantify goodness of fit based on the entropy over all possible forward and backward paths in the hierarchy. These same *Drosophila* data were recently reanalyzed using methods from statistical physics to further investigate temporal structure [84]. Sequences of discrete behavioral units were first mapped to a binary sequence via the information bottleneck algorithm based on the criteria that each state should be maximally predictive of the next. Fitting a one-dimensional Ising model to the binary sequence revealed temporal correlations over long timescales which followed an approximate inverse square law. Although this model does not distinguish between forward- and reverse-time correlations, the result

nonetheless supports the finding that free-walking *Drosophila* behaviors are not well characterized by a low-order Markov chain.

Long-time correlations in behavioral data may be the result of internal states that modulate behavioral organization. HMMs are useful for estimating such states. For example, an HMM could identify discrete finite-memory episodes which explained long-timescale patterns in free-walking *Drosophila* [87]. Furthermore, Calhoun et al. [35] used hierarchical HMMs to infer internal states driving acoustic behavior during *Drosophila* courtship. In this study, a generalized linear model (GLM) described the sensorimotor transformation between social cues and the generation of songs, while the HMM described a discrete latent state which modulated the GLM parameters. The three fitted states of the HMM were interpretable as different stages of social interaction and improved model fit over previous attempts using a GLM alone [52,88]. HMMs have also been applied to free-swimming larval zebrafish to identify internal states representing exploration and exploitation based on dispersion distance—the radius encapsulating movement over a fixed time window [89]. The choice of a two-state model was validated using Gaussian mixture models and the Bayesian information criterion. Simultaneous whole-brain neural imaging during free behavior enabled the authors to identify a distributed network of neurons correlated with internal state transitions. Limitations of the HMMs used in Calhoun et al. [35] and Marques et al. [89] were that the models assumed a discrete internal state, and that the HMM dynamics operated on a fixed timescale. Incorporation of a continuous internal state variable [72] and flexibility in the timescale of the HMM [25,30] are important directions for future study of hierarchical models of behavioral sequences in the context of environmental cues.

Marked renewal processes and dictionary-based approaches

The influence of sensory cues on larval zebrafish hunting sequences has also been investigated in [31] using marked renewal processes [90] with GLMs. The authors found that the previous bout type, previous inter-bout interval, and prey position were most predictive of the next action, and that the time that the fish had spent in the current hunting event was also predictive. A limitation of this modeling approach is that predictive variables are not necessarily causal and may be correlated, hence the model does not necessarily provide an algorithmic explanation for observed behaviors. However, the capacity of the model to identify environmental cues and behavioral states that influence future action selection suggests its potential as a tool for hypothesis generation in the search for corresponding neural circuits and algorithms.

Dictionary-based compression provides an algorithmic approach for identifying stereotyped motifs on multiple timescales [30,82,83]. Repeated compression using the principle of minimum description length [91,92] has been applied to identify arbitrary length motifs from *C. elegans* locomotion, suggesting a hierarchical behavioral structure [82]. Furthermore, the factor by which the algorithm could compress a given behavioral sequence was shown to effectively discriminate between genetic variants. A limitation of this approach was that it identified relatively few stereotyped motifs. This is likely due to the intrinsic rigidity of the lossless compression algorithm, which implicitly accounts for all variations in a sequence of patterns, making it sensitive to natural behavioral variability and noise. A recent study proposed a generative model for dictionary-based representations of behavioral sequences which attempts to account for these issues by modeling variability with random insertion or deletion of behavioral units in a sequence [30]. However, the relationships between motifs extracted by dictionary-based methods, corresponding behavioral functions, and neural coding remain relatively unexplored.

Continuous manifolds and chaos

Certain behaviors may lie on a continuum, including the postural space underlying worm locomotion [59,60,64, 93], *Drosophila* locomotion and gait dynamics [87,94], larval zebrafish bout space [57,61], and mouse vocalizations [33]. Continuous models of naturalistic behaviors can provide complimentary insights into neural algorithms for behavioral control and the relationship between brain and biomechanics. For example, an investigation of inter-limb phase relationships in the gait dynamics of spontaneous walking in *Drosophila* showed a single continuum of coordination patterns [94]. The authors also found that a low-dimensional embedding of *Drosophila* limb kinematics via UMAP [39] produced a continuous bell-shaped manifold underlying gait dynamics. The direction along the bell manifold corresponded simultaneously to limb frequency and gait type, and also captured stimulus-induced perturbations in walking speed. These data were well-fit by a coupled-oscillator model controlled by a single continuous parameter, suggesting that qualitatively distinct gait dynamics may be generated by single continuous neural control variable.

Continuous dynamical systems models of postural dynamics of *C. elegans* worms have shown that phase timing of stimuli can control turn direction [59]. Recently, Ahamed et al. [93] found evidence of continuous chaotic attractors underlying postural dynamics in *C. elegans* during locomotion using phase space reconstruction by delay embedding and estimation of Lyapunov exponents [95–98]. The authors suggested that chaotic dynamics provide a possible framework to

explain how stereotyped yet nonetheless variable behaviors can emerge from a continuous manifold: qualitatively similar trajectories appear on shorter timescales due to the structure of the attractor, and variability naturally arises as trajectories diverge due to chaos. Furthermore, high-dimensional neural network models utilizing chaotic itinerancy [99] can produce both input-driven and spontaneous switching between specified metastable attractors [100]. It was shown that resulting sequences from these models can approximate Markov chains and higher order state transition models. If an elementary behavioral unit can be modeled by a continuous chaotic attractor, then chaotic itinerancy may provide a bridge between the continuous patterns of movement that define behaviors on short timescales, and the qualitatively discrete nature of behavioral sequences on longer timescales.

However, it is not clear the extent to which chaotic behavioral dynamics are the product of neural circuits or biomechanical oscillations [101]. Biomechanical models of larval *Drosophila* can generate oscillatory locomotor dynamics and chaotic turn motions which intrinsically produce an effective exploration strategy, suggesting that exploration can be driven by a diffusion process arising from chaos without an explicit neural control objective [102,103]. Ultimately, behaviors are realized through a coupling of neural dynamics, biomechanics, and the physics of the environment in which an organism finds itself [101], yet how this complex interplay impacts models of behavioral sequences and neural activity across different scales is yet to be fully explored.

Conclusions and outlook

Although there has recently been impressive growth in unsupervised computational methods for investigating naturalistic behavior, many technical and conceptual challenges remain. Unsupervised behavioral mapping techniques can extract detailed biomarkers with improved discrimination over coarse kinematic statistics for genetic and pharmacological phenotyping. However, in isolation these biomarkers provide limited insight into neural circuitry without additional human interpretation of the learned behavioral representation. Furthermore, computational pipelines for behavioral mapping typically apply many transformations to data, each imparting inductive biases which may impact biological interpretation. The risk of erroneous results can be mitigated by using techniques for model validation and selection such as cross-validation, or by testing the stability of solutions over many small random permutations to hyperparameters and data.

Patterns of behavior arise from the complex interaction between an organism's goals, environment, biomechanics, and developmental processes. In the search for algorithms underlying naturalistic behaviors, we should seek models for which there exist plausible neural and biomechanical substrates. The improvement of technologies for real-time pose tracking and neural imaging in free-moving animals will create new opportunities for the measurement of neural and sensory correlates of naturalistic behavior. Continued contributions from the fields of mathematics, statistics, physics, and machine learning will be needed to find meaning in complex behavioral data. Key outstanding problems for the field include resolving behavior across multiple scales, characterizing both discrete and continuous aspects of behavioral dynamics, and integrating sensory information and environmental constraints in behavioral models. Large-scale statistical and machine learning models of behavioral sequences are beginning to enable characterization and prediction from combinations of behavioral history, sensory information, and estimated internal states over multiple timescales. However, so far these typically implement linear transformations or probabilistic maps based on state history and sensory input to model the subsequent action. A promising future direction may be to integrate large-scale models of naturalistic behavioral data with control theory, or similar frameworks for modeling goal-driven behavior with feedback processes. Overall, the rapid development of methods and models for unsupervised quantification of animal behavior promises many important breakthroughs in our understanding of the link between brain and behavior in coming years.

Conflict of interest statement
Nothing declared.

References
Papers of particular interest, published within the period of review, have been highlighted as:

*   of special interest

1. Berman GJ: **Measuring behavior across scales**. *BMC Biol* 2018, **16**:23, https://doi.org/10.1186/s12915-018-0494-7.

2. Krakauer JW, Ghazanfar AA, Gomez-Marin A, MacIver MA, Poeppel D: **Neuroscience needs behavior: correcting a reductionist bias**. *Neuron* 2017, **93**:480−490, https://doi.org/10.1016/j.neuron.2016.12.041.

3. Musall S, Urai AE, Sussillo D, Churchland AK: **Harnessing behavioral diversity to understand neural computations for cognition**. *Curr Opin Neurobiol* 2019, **58**:229−238, https://doi.org/10.1016/j.conb.2019.09.011.

4. Datta SR, Anderson DJ, Branson K, Perona P, Leifer A: **Computational neuroethology: a call to action**. *Neuron* 2019, **104**:11−24, https://doi.org/10.1016/j.neuron.2019.09.038.

5. Gomez-Marin A, Ghazanfar AA: **The life of behavior**. *Neuron* 2019, **104**:25–36, https://doi.org/10.1016/j.neuron.2019.09.017.

6. Hein AM, Altshuler DL, Cade DE, Liao JC, Martin BT, Taylor GK: **An algorithmic approach to natural behavior**. *Curr Biol* 2020, **30**:R663–R675, https://doi.org/10.1016/j.cub.2020.04.018.

7. Calhoun AJ, Murthy M: **Quantifying behavior to solve sensorimotor transformations: advances from worms and flies**. *Curr Opin Neurobiol* 2017, **46**:90–98, https://doi.org/10.1016/j.conb.2017.08.006.

8. Gomez-Marin A, Paton JJ, Kampff AR, Costa RM, Mainen ZF: **Big behavioral data: psychology, ethology and the foundations of neuroscience**. *Nat Neurosci* 2014, **17**:1455–1462, https://doi.org/10.1038/nn.3812.

9. Stephens GJ, Osborne LC, Bialek W: **Searching for simplicity in the analysis of neurons and behavior**. *Proc Natl Acad Sci Unit States Am* 2011, **108**(Supplement 3):15565–15571, https://doi.org/10.1073/pnas.1010868108.

10. Nathan R, Getz WM, Revilla E, Holyoak M, Kadmon R, Saltz D, Smouse PE: **A movement ecology paradigm for unifying organismal movement research**. *Proc Natl Acad Sci Unit States Am* 2008, **105**:19052–19059, https://doi.org/10.1073/pnas.0800375105.

11. Pereira TD, Shaevitz JW, Murthy M: **Quantifying behavior to understand the brain**. *Nat Neurosci* 2020, **23**:1537–1549, https://doi.org/10.1038/s41593-020-00734-z.

12. Leshner A, Pfaff DW: **Quantification of behavior**. *Proc Natl Acad Sci Unit States Am* 2011, **108**(Supplement 3):15537–15541, https://doi.org/10.1073/pnas.1010653108.

13. Egnor SR, Branson K: **Computational analysis of behavior**. *Annu Rev Neurosci* 2016, **39**:217–236, https://doi.org/10.1146/annurev-neuro-070815-013845.

14. Gilpin W, Huang Y, Forger DB: **Learning dynamics from large biological data sets: machine learning meets systems biology**. *Curr Opin Struct Biol* 2020, **22**:1–7, https://doi.org/10.1016/j.coisb.2020.07.009.

15. Bialek W: *What do we mean by the dimensionality of behavior?*. arXiv; 2008:9574. arXiv preprint.

16. Mathis A, Mamidanna P, Cury KM, Abe T, Murthy VN, Mathis MW, Bethge M: **DeepLabCut: markerless pose estimation of user-defined body parts with deep learning**. *Nat Neurosci* 2018, **21**:1281–1289, https://doi.org/10.1038/s41593-018-0209-y.

17. Pereira TD, Aldarondo DE, Willmore L, Kislin M, Wang SS-H, Murthy M, Shaevitz JW: **Fast animal pose estimation using deep neural networks**. *Nat Methods* 2019, **16**:117–125, https://doi.org/10.1038/s41592-018-0234-5.

18. T. Bollu, B. Ito, S. C. Whitehead, B. Kardon, J. Redd, M. H. Liu, J. H. Goldberg, How a mouse licks a spout: cortex-dependent corrections as the tongue reaches for, and misses, targets, bioRxiv doi:10.1101/655852.

19. Hebert L, Ahamed T, Costa AC, O'Shaugnessy L, Stephens GJ: *WormPose: image synthesis and convolutional networks for pose estimation in C. elegans*. bioRxiv; 2020. 2020.07.09.193755, https://doi.org/10.1101/2020.07.09.193755.

20. Dunn TW, Marshall JD, Severson KS, Aldarondo DE, Hildebrand DGC, Chettih SN, Wang WL, Gellis AJ, Carlson DE, Aronov D, Freiwald WA, Wang F, Ölveczky BP: **Geometric deep learning enables 3D kinematic profiling across species and environments**. *Nat Methods* 2021, **18**:564–573, https://doi.org/10.1038/s41592-021-01106-6.

21. Mathis MW, Mathis A: **Deep learning tools for the measurement of animal behavior in neuroscience**. *Curr Opin Neurobiol* 2020, **60**:1–11, https://doi.org/10.1016/j.conb.2019.10.008.

22. Kane GA, Lopes G, Saunders JL, Mathis A, Mathis MW: **Real-time, low-latency closed-loop feedback using markerless posture tracking**. *eLife* 2020, **9**, e61909, https://doi.org/10.7554/eLife.61909.

23. Walter T, Couzin ID: **Trex, a fast multi-animal tracking system with markerless identification, and 2d estimation of posture and visual fields**. *eLife* 2021, **10**, e64000, https://doi.org/10.7554/eLife.64000.

24. Berman GJ, Choi DM, Bialek W, Shaevitz JW: **Mapping the stereotyped behaviour of freely moving fruit flies**. *J R Soc Interface* 2014, **11**:20140672, https://doi.org/10.1098/rsif.2014.0672.

25. Wiltschko A, Johnson M, Iurilli G, Peterson R, Katon J, Pashkovski S, Abraira V, Adams R, Datta S: **Mapping sub-second structure in mouse behavior**. *Neuron* 2015, **88**:1121–1135, https://doi.org/10.1016/j.neuron.2015.11.031.

26. Batty E, Whiteway M, Saxena S, Biderman D, Abe T, Musall S, Gillis W, Markowitz J, Churchland A, Cunningham JP, Datta SR, Linderman S, Paninski L: **BehaveNet: nonlinear embedding and Bayesian neural decoding of behavioral videos**. *Adv Neural Inf Process Syst* 2019, **32**:15706–15717.

27. Masci J, Meier U, Cireşan D, Schmidhuber J: **Stacked convolutional auto-encoders for hierarchical feature extraction**. In *Artificial neural networks and machine learning – ICANN 2011, lecture notes in computer science*. Edited by Honkela T, Duch W, Girolami M, Kaski S, Berlin, Heidelberg: Springer; 2011:52–59, https://doi.org/10.1007/978-3-642-21735-7_7.

28. Marques JC, Lackner S, Félix R, Orger MB: **Structure of the zebrafish locomotor repertoire revealed with unsupervised behavioral clustering**. *Curr Biol* 2018, **28**:181–195. e5, https://doi.org/10.1016/j.cub.2017.12.002.

29. Groneberg AH, Marques JC, Martins AL, Diez del Corral R, de Polavieja GG, Orger MB: **Early-life social experience shapes social avoidance reactions in larval zebrafish**. *Curr Biol* 2020, **30**:4009–4021, https://doi.org/10.1016/j.cub.2020.07.088. e4.

30. Reddy G, Desban L, Tanaka H, Roussel J, Mirat O, Wyart C: *A lexical approach for identifying behavioral action sequences*. bioRxiv; 2020. 2020.08.27.270694, https://doi.org/10.1101/2020.08.27.270694.

31. Johnson RE, Linderman S, Panier T, Wee CL, Song E, Herrera KJ, Miller A, Engert F: **Probabilistic models of larval zebrafish behavior reveal structure on many scales**. *Curr Biol* 2020, **30**:70–82. e4, https://doi.org/10.1016/j.cub.2019.11.026.
The authors fit a large-scale probabilistic model based on a marked renewal process to describe and simulate free swimming and hunting sequences in larval-zebrafish. Fitted model parameters revealed the relative importance of specific environmental cues, behavioral history and hunger state for the prediction of subsequent actions.

32. Mackevicius EL, Bahle AH, Williams AH, Gu S, Denisenko NI, Goldman MS, Fee MS: **Unsupervised discovery of temporal sequences in high-dimensional datasets, with applications to neuroscience**. *eLife* 2019, **8**, e38471, https://doi.org/10.7554/eLife.38471.

33. Goffinet J, Mooney R, Pearson J: *Inferring low-dimensional latent descriptions of animal vocalizations*. bioRxiv; 2019:811661, https://doi.org/10.1101/811661.

34. Cohen Y, Nicholson D, Sanchioni A, Mallaber EK, Skidanova V, Gardner TJ: *TweetyNet: a neural network that enables high-throughput, automated annotation of birdsong*. bioRxiv; 2020. 2020.08.28.272088, https://doi.org/10.1101/2020.08.28.272088.

35. Calhoun AJ, Pillow JW, Murthy M: **Unsupervised identification of the internal states that shape natural behavior**. *Nat Neurosci* 2019, **22**:2040–2049, https://doi.org/10.1038/s41593-019-0533-x.
Using a hierarchical combination of a hidden Markov model and generalized linear models, the authors show how previously undiscovered internal states modulate sensorimotor transformations during *Drosophila* courtship. The trained model was applied in conjunction with optogenetic stimulation to identify a pair of neurons which modulated the inferred internal state.

36. Steinbach M, Ertöz L, Kumar V: *The challenges of clustering high dimensional data, in: new directions in statistical physics*. Springer; 2004:273–309.

37. Ronan T, Qi Z, Naegle KM: **Avoiding common pitfalls when clustering biological data**. re6–re6. arXiv *Sci Signal* 2016, **9**, https://doi.org/10.1126/scisignal.aad1932. URL, https://stke.sciencemag.org/content/9/432/re6.full.pdf. https://stke.sciencemag.org/content/9/432/re6.

38. Maaten Lv d, Hinton G: **Visualizing Data using t-SNE**. *J Mach Learn Res* 2008, **9**:2579–2605.

39. L. McInnes, J. Healy, J. Melville, UMAP: uniform manifold approximation and projection for dimension reduction, arXiv preprint arXiv:1802.03426.

40. Tenenbaum JB, Silva Vd, Langford JC: **A global geometric framework for nonlinear dimensionality reduction**. *Science* 2000, **290**:2319–2323, https://doi.org/10.1126/science.290.5500.2319.

41. Kingma DP, Welling M: *An introduction to variational autoencoders*, **12**. Foundations and Trends® in Machine Learning; 2019: 307–392. arXiv: 1906.02691, https://doi.org/10.1561/2200000056.

42. C. Doersch, Tutorial on variational autoencoders, arXiv: 1606.05908 [cs, stat].

43. D. P. Kingma, M. Welling, Auto-encoding variational bayes, arXiv: 1312.6114 [cs, stat]arXiv: 1312.6114.

44. Rezende DJ, Mohamed S, Wierstra D: **Stochastic backpropagation and approximate inference in deep generative models**. In *International conference on machine learning, PMLR*; 2014:1278–1286.

45. Graving JM, Couzin ID: *VAE-SNE: a deep generative model for simultaneous dimensionality reduction and clustering*. bioRxiv; 2020. 2020.07.17.207993, https://doi.org/10.1101/2020.07.17.207993.

46. Luxem K, Fuhrmann F, Kürsch J, Remy S, Bauer P: *Identifying behavioral structure from deep variational embeddings of animal motion*. bioRxiv; 2020. 2020.05.14.095430. https://doi.org/10.1101/2020.05.14.095430.

47. Berman GJ, Bialek W, Shaevitz JW: **Predictability and hierarchy in Drosophila behavior**. *Proc Natl Acad Sci Unit States Am* 2016, **113**:11943–11948, https://doi.org/10.1073/pnas.1607601113.

48. Hernández DG, Rivera C, Cande J, Zhou B, Stern DL, * Berman GJ: *A framework for studying behavioral evolution by reconstructing ancestral repertoires*. arXiv; 2007:9689 [physics, q-bio], https://doi.org/10.1101/2020.07.17.209361.
MotionMapper [24] was applied to compare free-walking behavior across six closely-related species of *Drosophila*. Based on the results of unsupervised behavioral mapping, the authors proposed a statistical method for investigating how behaviors could co-evolve by using a generalized linear mixed model of phylogenetic and individual variability.

49. Klibaite U, Berman GJ, Cande J, Stern DL, Shaevitz JW: **An unsupervised method for quantifying the behavior of paired animals**. *Phys Biol* 2017, **14**:15006, https://doi.org/10.1088/1478-3975/aa5c50.

50. Klibaite U, Shaevitz JW: **Paired fruit flies synchronize behavior: uncovering social interactions in Drosophila melanogaster**. *PLoS Comput Biol* 2020, **16**, e1008230, https://doi.org/10.1371/journal.pcbi.1008230.

51. Goupillaud P, Grossmann A, Morlet J: **Cycle-octave and related transforms seismic Signal Analysis and Discrimination III in seismic signal analysis**. *Geoexploration* 1984, **23**:85–102. seismic Signal Analysis and Discrimination III. https://doi.org/10.1016/0016-7142(84)90025–5.

52. Clemens J, Coen P, Roemschied FA, Pereira TD, Mazumder D, Aldarondo DE, Pacheco DA, Murthy M: **Discovery of a new song mode in Drosophila reveals hidden structure in the sensory and neural drivers of behavior**. *Curr Biol* 2018, **28**:2400–2412. e6, https://doi.org/10.1016/j.cub.2018.06.011.

53. Marshall JD, Aldarondo DE, Dunn TW, Wang WL, Berman GJ, * Ölveczky BP: **Continuous whole-body 3D kinematic recordings across the rodent behavioral repertoire**. *Neuron* 2021 Feb 3, **109**:420–437, https://doi.org/10.1016/j.neuron.2020.11.016.
Introduced CAPTURE (Continuous Appendicular and Postural Tracking Using Retroreflector Embedding), a postural tracking and naturalistic behavioral analysis framework for rodents. This incorporates high-precision tracking for week-long timescales using deep learning for pose estimation, semi-supervised feature extraction on a fixed timescale, and behavioral mapping via MotionMapper [24]. The authors applied this framework for genetic and pharmacological phenotyping from behavior, revealing that amphetamine increased repetitive behaviors in rats at lower doses than previously thought required, and that *fmr*1-KO mice showed altered grooming behaviors.

54. Budick S, O'Malley D: **Locomotor repertoire of the larval zebrafish: swimming, turning and prey capture**. *J Exp Biol* 2000, **203**:2565–2579.

55. McElligott MB, O'Malley DM: **Prey tracking by larval zebrafish: axial kinematics and visual control**. *Brain Behav Evol* 2005, **66**: 177–196, https://doi.org/10.1159/000087158.

56. Patterson BW, Abraham AO, MacIver MA, McLean DL: **Visually guided gradation of prey capture movements in larval zebrafish**. *J Exp Biol* 2013, **216**:3071–3083, https://doi.org/10.1242/jeb.087742.

57. Mearns DS, Donovan JC, Fernandes AM, Semmelhack JL, * Baier H: **Deconstructing hunting behavior reveals a tightly coupled stimulus-response loop**. *Curr Biol* 2020, **30**:54–69. e9, https://doi.org/10.1016/j.cub.2019.11.022.
The authors constructed a map of larval zebrafish bout space from postural time series using dynamic time warping and Isomap embedding. Fine scale analysis of hunting sequences revealed short-time temporal structure suggestive of a memory-less stimulus-response loop, and uncovered a previously unknown properties of prey-capture swims.

58. Berndt DJ, Clifford J: **Using dynamic time warping to find patterns in time series**. In *KDD workshop*, vol. 10; 1994: 359–370. Seattle, WA, USA.

59. Stephens GJ, Johnson-Kerner B, Bialek W, Ryu WS: **Dimensionality and dynamics in the behavior of C. elegans**. *PLoS Comput Biol* 2008, **4**, e1000028, https://doi.org/10.1371/journal.pcbi.1000028.

60. Stephens GJ, Johnson-Kerner B, Bialek W, Ryu WS: **From modes to movement in the behavior of Caenorhabditis elegans**. *PLoS One* 2010, **5**, e13914, https://doi.org/10.1371/journal.pone.0013914.

61. Girdhar K, Gruebele M, Chemla YR: **The behavioral space of zebrafish locomotion and its neural network analog**. *PLoS One* 2015, **10**, e0128668, https://doi.org/10.1371/journal.pone.0128668.

62. Frey BJ, Dueck D: **Clustering by passing messages between data points**. *Science* 2007, **315**:972–976, https://doi.org/10.1126/science.1136800.

63. Marques JC, Orger MB: **Clusterdv: a simple density-based clustering method that is robust, general and automatic**. *Bioinformatics* 2018, **35**:2125–2132, https://doi.org/10.1093/bioinformatics/bty932.

64. Szigeti B, Deogade A, Webb B: **Searching for motifs in the behaviour of larval Drosophila melanogaster and Caenorhabditis elegans reveals continuity between behavioural states**. *J R Soc Interface* 2015, **12**:20150899, https://doi.org/10.1098/rsif.2015.0899.

65. Gaffney S, Smyth P: **Joint probabilistic curve clustering and alignment**. In *NIPS*; 2004:473–480.

66. Costa AC, Ahamed T, Stephens GJ: **Adaptive, locally linear models of complex dynamics**. *Proc Natl Acad Sci Unit States Am* 2019, **116**:1501–1510, https://doi.org/10.1073/pnas.1813476116.

67. Klaus A, Martins GJ, Paixao VB, Zhou P, Paninski L, Costa RM: **The spatiotemporal organization of the striatum encodes action space**. *Neuron* 2017, **95**:1171–1180. e7, https://doi.org/10.1016/j.neuron.2017.08.015.

68. Karpenko S, Wolf S, Lafaye J, Le Goc G, Panier T, Bormuth V, Candelier R, Debrégeas G: **From behavior to circuit modeling of light-seeking navigation in zebrafish larvae**. *eLife* 2020, **9**, e52882, https://doi.org/10.7554/eLife.52882.

69. Rousseeuw PJ: **Silhouettes: a graphical aid to the interpretation and validation of cluster analysis**. *J Comput Appl Math* 1987, **20**:53–65, https://doi.org/10.1016/0377-0427(87)90125-7.

70. Caliński T, Harabasz J: **A dendrite method for cluster analysis**. *Commun Stat* 1974, **3**:1–27, https://doi.org/10.1080/03610927408827101.

71. Tibshirani R, Walther G, Hastie T: **Estimating the number of clusters in a data set via the gap statistic**. *J Roy Stat Soc B* 2001, **63**:411–423, https://doi.org/10.1111/1467-9868.00293.

72. Linderman S, Johnson M, Miller A, Adams R, Blei D, Paninski L: *Bayesian learning and inference in recurrent switching linear dynamical systems*. In *Artificial intelligence and statistics, PMLR*; 2017:914–922.

73. Fox EB, Sudderth EB, Jordan MI, Willsky AS: *Sharing features among dynamical systems with beta processes*. In *Proceedings of the 22nd international conference on neural information processing systems, NIPS'09*. NY, USA: Curran Associates Inc., Red Hook; 2009:549–557.

74. Rabiner L: **A tutorial on hidden Markov models and selected applications in speech recognition**. *Proc IEEE* 1989, **77**: 257–286, https://doi.org/10.1109/5.18626.

75. Wiltschko AB, Tsukahara T, Zeine A, Anyoha R, Gillis WF,
* Markowitz JE, Peterson RE, Katon J, Johnson MJ, Datta SR: **Revealing the structure of pharmacobehavioral space through motion sequencing**. *Nat Neurosci* 2020, **23**: 1433–1443, https://doi.org/10.1038/s41593-020-00706-3.
The authors applied MoSeq [25] to large-scale pharmacological phenotyping of naturalistic mouse behavior. Results showed that this approach significantly improved discrimination between drugs based on behavioral data when compared with coarse kinematic statistics. Futher, MoSeq identified on-target and off-target effects of drugs for reversing behavioral phenotypes in a model of autism spectrum disorder, thus demonstrating the potential of this approach in phenotypic drug discovery.

76. Ackerson G, Fu K: **On state estimation in switching environments**. *IEEE Trans Automat Contr* 1970, **15**:10–17. conference Name: IEEE Transactions on Automatic Control, https://doi.org/10.1109/TAC.1970.1099359.

77. Hamilton JD: **Analysis of time series subject to changes in regime**. *J Econom* 1990, **45**:39–70, https://doi.org/10.1016/0304-4076(90)90093-9.

78. Fox E, Sudderth E, Jordan M, Willsky A: **Nonparametric bayesian learning of switching linear dynamical systems**. *IFAC Proc Vol* 2009:1591–1596, https://doi.org/10.3182/20090706-3-FR-2004.00264.

79. Brown AEX, de Bivort B: **Ethology as a physical science**. *Nat Phys* 2018, **14**:653–657, https://doi.org/10.1038/s41567-018-0093-0.

80. Tinbergen N: *The study of instinct, the study of instinct*, **xii**. New York, NY, US: Clarendon Press/Oxford University Press; 1951:237.

81. Markowitz JE, Gillis WF, Beron CC, Neufeld SQ, Robertson K,
* Bhagat ND, Peterson RE, Peterson E, Hyun M, Linderman SW, Sabatini BL, Datta SR: **The striatum organizes 3D behavior via moment-to-moment action selection**. *Cell* 2018, **174**:44–58. e17, https://doi.org/10.1016/j.cell.2018.04.019.
MoSeq [25] was applied to segment free-moving mouse behavior into a sequence of stereotyped behavioral units. Combined with simultaneous neural imaging and lesioning experiments, the authors found pathways in the dorsolateral striatum which encoded behavioral sequences on a sub-second timescale. MoSeq also revealed that arena shape can influence behavioral utilization, and that exposure to an aversive odor stimulus up-regulated specific intepretable behaviors.

82. Gomez-Marin A, Stephens GJ, Brown AEX: **Hierarchical compression of Caenorhabditis elegans locomotion reveals phenotypic differences in the organization of behaviour**. *J R Soc Interface* 2016, **13**:20160466, https://doi.org/10.1098/rsif.2016.0466.

83. Gupta S, Gomez-Marin A: *A context-free grammar for Caenorhabditis elegans behavior*. bioRxiv; 2019:708891, https://doi.org/10.1101/708891.

84. Alba V, Berman GJ, Bialek W, Shaevitz JW: *Exploring a strongly non-Markovian animal behavior*. arXiv; 2012:15681 [cond-mat, q-bio].

85. Simon HA: **The architecture of complexity**. *Proc Am Phil Soc* 1962, **106**:467–482.

86. N. Tishby, F. C. Pereira, W. Bialek, The information bottleneck method arXiv:physics/0004057.

87. Katsov AY, Freifeld L, Horowitz M, Kuehn S, Clandinin TR: **Dynamic structure of locomotor behavior in walking fruit flies**. *eLife* 2017, **6**, e26410, https://doi.org/10.7554/eLife.26410.

88. Coen P, Clemens J, Weinstein AJ, Pacheco DA, Deng Y, Murthy M: **Dynamic sensory cues shape song structure in Drosophila**. *Nature* 2014, **507**:233–237, https://doi.org/10.1038/nature13131.

89. Marques JC, Li M, Schaak D, Robson DN, Li JM: **Internal state
* dynamics shape brainwide activity and foraging behaviour**. *Nature* 2020, **577**:239–243, https://doi.org/10.1038/s41586-019-1858-z.
The authors use hidden Markov models to identify exploration and exploitation states in free-swimming larval zebrafish behavior. Notably, the internal state appeared to oscillate independently of hunger or the presence of prey. Using simultaneous imaging of behavior and whole-brain neural activity they found distributed networks of neurons that may be a mechanism for state-dependent neuromodulation of specific sensorimotor circuits.

90. Daley D, Vere-Jones D: **Renewal processes**. In *An introduction to the theory of point processes: volume I: elementary theory and methods, probability and its applications*. Edited by Daley DJ, Vere-Jones D, New York, NY: Springer; 2003:66–110, https://doi.org/10.1007/0-387-21564-6_4.

91. Nevill-Manning CG, Witten IH: **On-line and off-line heuristics for inferring hierarchies of repetitions in sequences**. *Proc IEEE* 2000, **88**:1745–1755, https://doi.org/10.1109/5.892710.

92. Grünwald PD, Myung IJ, Pitt MA: *Advances in minimum description length: theory and applications*. MIT press; 2005.

93. Ahamed T, Costa AC, Stephens GJ: **Capturing the continuous complexity of behaviour in Caenorhabditis elegans**.
* *Nat Phys* 2020:1–9, https://doi.org/10.1038/s41567-020-01036-8.
The authors modeled continuous postural trajectories of *C. elegans* worms during locomotion using a novel method for maximally predictive state space reconstruction. Attractor reconstruction and estimation of topological invariants revealed evidence of continuous chaotic dynamics driving behavior.

94. DeAngelis BD, Zavatone-Veth JA, Clark DA: **The manifold
* structure of limb coordination in walking Drosophila**. *eLife* 2019, **8**, e46409, https://doi.org/10.7554/eLife.46409.
A detailed analysis of inter-limb phase relationships in spontaneous walking behaviors of *Drosophila* using low-dimensional manifold estimation and coupled-oscillator models showed that gait dynamics may lie on a continuum and could be modulated by a single control variable.

95. Takens F: **Detecting strange attractors in turbulence**. In *Dynamical systems and turbulence, Warwick 1980*. Springer; 1981:366–381.

96. Kennel MB, Brown R, Abarbanel HDI: **Determining embedding dimension for phase-space reconstruction using a geometrical construction**. *Phys Rev* 1992, **45**:3403–3411, https://doi.org/10.1103/PhysRevA.45.3403.

97. Kantz H, Schreiber T. *Nonlinear time series analysis*, vol. 7. Cambridge University Press; 2004.

98. Bradley E, Kantz H: **Nonlinear time-series analysis revisited, Chaos: an Interdisciplinary**. *J Nonlinear Sci* 2015, **25**:97610, https://doi.org/10.1063/1.4917289.

99. Tsuda I: **Chaotic itinerancy and its roles in cognitive neurodynamics**. *Curr Opin Neurobiol* 2015, **31**:67–71, https://doi.org/10.1016/j.conb.2014.08.011.

100.Inoue K, Nakajima K, Kuniyoshi Y: **Designing spontaneous behavioral switching via chaotic itinerancy**. *Sci Adv* 2020, **6**, eabb3989, https://doi.org/10.1126/sciadv.abb3989.

101. Tytell E, Holmes P, Cohen A: **Spikes alone do not behavior make: why neuroscience needs biomechanics**. *Curr Opin Neurobiol* 2011, **21**:816–822, https://doi.org/10.1016/j.conb.2011.05.017.

102. Loveless J, Lagogiannis K, Webb B: **Modelling the mechanics of exploration in larval Drosophila**. *PLoS Comput Biol* 2019, **15**, e1006635, https://doi.org/10.1371/journal.pcbi.1006635.

103. Loveless J, Garner A, Issa AR, Webb B, Ohyama T, Alonso CR: *A physical theory of larval Drosophila behaviour*. bioRxiv; 2020: 2020. 08.25.266163, https://doi.org/10.1101/2020.08.25.266163.

104. Hansen MH, Yu B: **Model selection and the principle of minimum description length**. *J Am Stat Assoc* 2001, **96**(454): 746–774, https://doi.org/10.1198/016214501753168398.

105. Bishop CM: *Pattern recognition and machine learning*. Springer; 2006.