Review

Computational Neuroethology: A Call to Action

Sandeep Robert Datta,[5,9,*] David J. Anderson,[1,2,3] Kristin Branson,[4] Pietro Perona,[8] and Andrew Leifer[6,7,*]
[1]Division of Biology and Biological Engineering 156-29, California Institute of Technology, Pasadena, CA 91125, USA
[2]Howard Hughes Medical Institute, Pasadena, CA, 91125, USA
[3]Tianqiao and Chrissy Chen Institute for Neuroscience, California Institute of Technology, Pasadena, CA 91125, USA
[4]Janelia Research Campus, Howard Hughes Medical Institute, Ashburn, VA 20147, USA
[5]Department of Neurobiology, Harvard Medical School, Boston, MA 02115, USA
[6]Department of Physics, Princeton University, Princeton, NJ 08544, USA
[7]Princeton Neuroscience Institute, Princeton University, Princeton, NJ 08544, USA
[8]Division of Engineering & Applied Sciences 136-93, California Institute of Technology, Pasadena, CA 91125, USA
[9]Lead Contact
*Correspondence: srdatta@hms.harvard.edu (S.R.D.), leifer@princeton.edu (A.L.)
https://doi.org/10.1016/j.neuron.2019.09.038

The brain is worthy of study because it is in charge of behavior. A flurry of recent technical advances in measuring and quantifying naturalistic behaviors provide an important opportunity for advancing brain science. However, the problem of understanding unrestrained behavior in the context of neural recordings and manipulations remains unsolved, and developing approaches to addressing this challenge is critical. Here we discuss considerations in computational neuroethology—the science of quantifying naturalistic behaviors for understanding the brain—and propose strategies to evaluate progress. We point to open questions that require resolution and call upon the broader systems neuroscience community to further develop and leverage measures of naturalistic, unrestrained behavior, which will enable us to more effectively probe the richness and complexity of the brain.

Leveraging Naturalistic Behavior to Explore Brain Function

Two distinct traditions have shaped how neuroscientists think about behavior in the lab (Gomez-Marin et al., 2014). Comparative psychology studies the ability of the brain to generate behaviors in response to rewards and punishments (Domjan, 1987). This perspective has led to a large body of work in which animals are trained in the laboratory to respond to specific sensory cues. By combining these behavioral methods with neural recordings and manipulations, modern neuroscience is now addressing fundamental questions about how task-related variables are encoded in the brain and about how neurons and circuits generate task-related behaviors (Jazayeri and Afraz, 2017; Krakauer et al., 2017). Animals are typically trained to produce simple actions (e.g., to lick a port or to reach for a target) that are easy to measure and readily correlated with neural activity patterns. In addition, animals are often (but not always) physically restrained, both to facilitate neural recordings and to avoid spurious movements that complicate inferences about the meaning and purpose of measured patterns of neural activity.

Ethology, on the other hand, has historically focused on natural behavior (Tinbergen, 1951, 1963). The underlying hypothesis of ethology is that exposing the structure of behavior—how behavior in the natural environment is built from components and organized over time in response to ecologically relevant stimuli—will yield insights into how the brain creates behavior (Simmons and Young, 1999; Tinbergen, 1951). However, traditional ethology has focused on observing the behavior of animals without neural recordings or interventions. Exploring neural activity during the expression of "naturalistic" behaviors (which here is taken to mean behaviors that are representative of actions generated during complex real-world tasks, like exploring new environments, obtaining food, finding shelter, and identifying mates, and are therefore largely self-motivated and expressed freely without physical restraint; see Box 1) has the potential to reveal how the brain does much of what the brain evolved to do. Furthermore, ethology has revealed that many naturalistic behaviors are built from sequences of components, a feature that in principle can be used to reveal dependencies between neural activity and actions across multiple timescales and to illuminate how longer-lasting brain states specify the moment-to-moment contents of behavior (Baerends, 1976; Manoli et al., 2006; Tinbergen, 1951).

We argue that understanding the relationship between brain and behavior will require bringing the disciplines of psychology, neurobiology, and ethology together, toward an integrated study of naturalistic behavior that spans a gamut of questions from brain mechanisms to evolution. Despite the compromises imposed by training and/or restraint, this approach to relating neural activity to behavior has yielded, and will continue to yield, key insights into the mechanisms that support perception, govern decision making, and regulate action (Juavinett et al., 2018; Panzeri et al., 2017). Technical advances—ranging from the development of virtual reality-based tasks to the use of touchpads for operant conditioning—are integrating rich contexts and increasingly detailed behavioral measures into training-based experiments (Mar et al., 2013; Minderer et al., 2016). Furthermore, the recent development of deep-learning-based frameworks for tracking, e.g., paws during reaching (see below) is revealing the behavioral richness and variability that underlies even simple behavioral reports like pellet grabs or lever presses (Graving et al., 2019; Guo et al., 2015; Mathis et al., 2018; Nath et al., 2019; Pereira et al., 2019). In contrast, the formidable technical and conceptual challenges of relating

Box 1. Glossary

Behavioral representation—a quantitative distillation of any aspect of the time-varying behavior exhibited by an animal in an experiment. Such representations can vary in form from classical ethograms to low-dimensional plots capturing the trajectory of an animal in space.

Naturalistic—as with "ecologically relevant" (see below), there are many definitions for naturalistic, and indeed most experiments in behavioral neuroscience can justifiably be argued to be naturalistic at some level. Here, we take the word "naturalistic" to mean behaviors that are representative of actions generated during real-world tasks, like exploring new environments, obtaining food, finding shelter, and identifying mates; naturalistic behaviors as referred to herein are also largely self-motivated and expressed freely without physical restraint. This definition is meant to distinguish such behaviors from those that are imposed by researchers on animals through overtraining or those that are more constrained due to, e.g., head fixation (although, as mentioned above, there are contexts in which those types of behaviors are quite reasonably referred to as "naturalistic").

Ecologically relevant—as with "naturalistic" (see above), "ecologically relevant" is an adjective whose meaning is in the eye of the beholder; again, this term can be appropriately applied to many kinds of behavioral experiments, including those in which animals are subject to training and restraint. Here, we take "ecologicallyrelevant" to mean a set of behaviors that support tasks animals have to address as part of the existential challenge of living in a particular environmental or ecological niche.

Behavioral label—a behavioral label is a descriptor applied to an epoch of behavioral data. Behavioral labels can cover descriptions of behavior at many levels of granularity and run the gamut from "a twitch of motor unit 72 in the soleus muscle" to "hibernating."

Behavioral motif—a stereotyped and re-used unit of movement. The terms "motif," "moveme," "module," "primitive," and "syllable" have all been used interchangeably, and none of these terms is linked to a rigorous definition of the spatiotemporal scale at which a unit of behavior is organized. Similarly, action and behavior here and elsewhere are used to refer to collections of units of behavior, but again, there is no rigorous line that separates these or related terms. Perona and Anderson have argued for a taxonomy in which moveme is the simplest movement associated with a behavior, an action is a sequence of movemes, and an activity is a species-characteristic set of movemes and actions (Anderson and Perona, 2014).

Behavioral sequence—an epoch in which more than one behavioral motif is expressed; sequences of motifs can be either deterministic (e.g., motif A always follows motif B) or probabilistic (e.g., motif A follows motif B fifty percent of the time).

Artificial neural network—class of machine learning algorithms that operate by simulating a network of simplified neurons. These often are trained through supervised learning.

Behavioral feature—a relevant attribute of an animal that, when observed over time, helps define behavior, e.g., the location of a paw or the animal's velocity.

Trajectory—the motion of a point in space over time. This can be either a physical object, like an animal's paw through real space, or it can be abstract, like the animal's current behavior state as it travels through behavior state space.

Behavior analysis pipeline—a set of algorithms that, all together, take a raw recording of behavior (usually video) and return high-level representations of the animal's behavior, such as trajectories, motifs, sequences, or behavior labels.

Dimensionality—number of variables required to describe a dataset. For example, the (x,y) position of a mouse paw has dimensionality of 2. A complete description of its pose requires many more variables and thus has higher dimensionality.

Dimensionality reduction—mathematically approximating a dataset using fewer than the original dimensions. For example, Stephens et al., 2008 showed that a worm's centerline originally requiring 100 points could be well approximated by three numbers. Dimensionality reduction usually requires a change in the representation of the data, e.g., as eigenvectors.

Embedding—a type of dimensionality reduction that takes data, which are assumed to exist on a high manifold, and collapses them into a lower-dimensional space where they are more easily visualized. PCA (see below) is a linear embedding method, while tSNE and U-MAP are two examples of non-linear embeddings.

Temporal dynamics—here, how behavior features change over time. These can be mathematically represented in the time domain or in the frequency domain.

Keypoint—a point demarcating a region of interest in an image, such as an animal's joint or appendage.

Supervised learning—a computer algorithm that learns to performs a task, such as identifying an animal's joint in an image, through human-provided training examples.

Unsupervised learning—a computer algorithm that learns to perform a task, such as discriminating between different actions, purely from the structure of the data and without help from labeled training examples.

Behavior state space—a mathematical space (possibly high dimensional) such that a point in this space corresponds to a specific instance of animal behavior.

Behavior map—visualization of a behavior state space; usually refers to a two-dimensional visualization. Convenient for visualizing how an animal's behavior is organized into clusters or how it might differ from another animal. This map can be generated using linear or non-linear embedding methods (see above).

Principal components analysis (PCA)—a mathematical change of basis that is commonly used for dimensionality reduction in many behavioral analysis pipelines. PCA identifies a new basis set in which to represent data. A truncated version of this dataset serves as a useful lower-dimensional approximation of the original data.
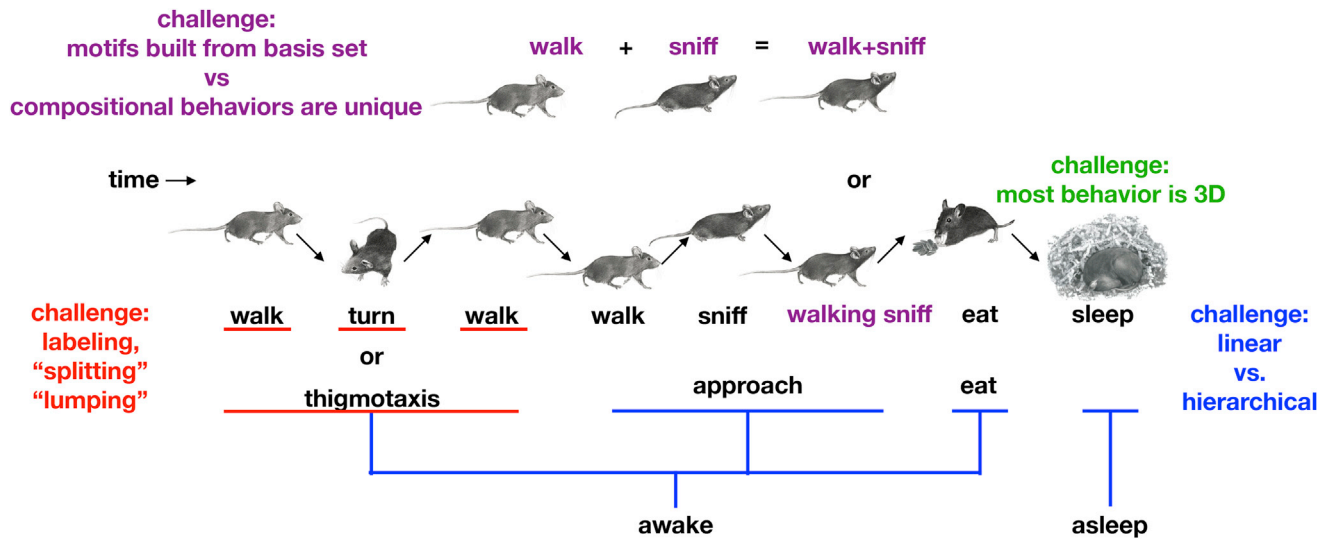
Figure 1. Challenges in Computational Ethology
A characteristic sequence of behaviors exhibited by a mouse in its home cage as it moves around the walls (thigmotaxis), approaches some food, eats, and then sleeps (center, mouse cartoons). Several key challenges facing any segmentation of continuous behavior into components are illustrated. First, the behavior of most model organisms occurs in three dimensions (green). Second, behaviors need to be labeled, which raises the problem of "lumping" versus "splitting" (red); for example, mice engage in thigmotaxis, behaviors in which the mouse exhibits locomotion and turning behaviors that are deterministically sequenced to generate an action where the animal circumnavigates its cage. Is thigmotaxis a singular or unitary behavior (because its elements are deterministically linked during its expression), or it is a sequence of walking, turning, and walking behaviors? Third, should behavior be considered at a single timescale that serially progresses or instead considered a hierarchical process organized at multiple timescales simultaneously (blue)? Fourth, when the mouse is sniffing and running at the same time, is that a compositional behavior whose basis set includes "run" and "sniff," or is "running+sniffing" a fundamentally new behavior (purple)?

naturalistic, unrestrained, and minimally shaped behavior to neural activity are only beginning to be addressed, leaving that area ripe for further development.

In the past decade, a field we now call "computational ethology" has begun to take shape. It involves the use of machine vision and machine learning to measure and analyze the patterns of action generated by animals in contexts designed to evoke ecologically relevant behaviors (Anderson and Perona, 2014). Technical progress in statistical inference and deep learning, the democratization of high-performance computing (due to falling hardware costs and the ability to rent GPUs and CPUs in the cloud), and new and creative ideas about how to apply technology to measuring naturalistic behavior have dramatically accelerated progress in this research area.

Approaches from computational ethology may be particularly important in a future in which we have access to recordings from many thousands of neurons, with the richness of neural codes on full display. Indeed, today nearly all of the neurons in the brains of the worm *C. elegans* and the zebrafish *D. rerio* can be recorded simultaneously, allowing a large fraction of the brain's neural dynamics to be observed (Cong et al., 2017; Kim et al., 2017; Nguyen et al., 2016; Symvoulidis et al., 2017; Venkatachalam et al., 2016). Given that subtle movements can have pervasive effects on neural dynamics, obtaining unbiased and holistic measurements of an animal's behavior - even in restrained animals — will be important to understanding dense neural activity (Musall et al., 2019; Stringer et al., 2019). We further propose that making sense of high-dimensional neural data will ultimately be facilitated by access to behaviors whose complexity and dimensionality is closer to

that of the neural space that is concurrently being surveyed. This perspective makes urgent the need to develop methods that combine analysis of naturalistic, unrestrained behavior with measures of neural activity. Here we review progress toward a science of computational *neuro*ethology—the problem of relating naturalistic behavior to neural recordings or manipulations in order to better understand brain function.

Challenges in Computational Neuroethology
Studying neural activity as animals behave freely has led to some of the most exciting discoveries in brain science over the past 50 years, including place cells (Hartley et al., 2013), grid cells (Rowland et al., 2016), replay (Foster, 2017), mechanisms of non-associative learning (Kandel et al., 2014), and the escape response (Medan and Preuss, 2014). However, to approach these problems without the benefit of restraint (e.g., head fixation), researchers have, by necessity, focused on those behaviors that are simplest to quantify. For example, an animal's head direction can be quantified by a single angle, its location can be quantified by an (x,y) pair of spatial coordinates, and a gill-withdrawal or escape reflex can be quantified by a single binary variable. These behaviors are trivial to estimate by eye, easy to plot on a graph, and unambiguous in their timing.

Unfortunately, most naturalistic behaviors—from exploration of a novel environment to mating rituals—are not well captured by simple parameters like centroid position or head direction (Figure 1). First, naturalistic behaviors involve coordinated movements of limbs, facial features, and other body parts and often include rapid changes in the animal's three-dimensional (3D) pose, i.e., its dynamics. Understanding dynamics requires

simultaneously measuring how the positions of many different body parts evolve over time. Second, although naturalistic behaviors are often built from stereotyped components (referred to variously as behavioral "motifs," "modules," "syllables," "primitives," and "movemes"), labeling behaviors on a moment-to-moment basis is made complicated by spatial and temporal variability in the execution of each behavior (Anderson and Perona, 2014; Berman et al., 2014; Flash and Hochner, 2005; Tinbergen, 1951; Wiltschko et al., 2015). This variability, taken with the observation that many spontaneous behaviors evolve continuously over time, makes it difficult to assign labels and explicit start and stop times to individual actions. Third, because naturalistic behaviors can be described at different levels of granularity, there are many simultaneously valid descriptions of an animal's behavior at any given time point (Dawkins, 1976). For example, replaying a video of a mouse in slow-motion will reveal the kinematics of limb movement as the animal turns its body, but the same movements when played on fast-forward will reveal whether the animal is in the midst of sleep, courtship, or escape behaviors. Finally, identifying actions is complicated by the ability of animals to do more than one thing at once. Some of the authors of this review, for example, claim to be able to walk and chew gum at the same time. This parallelism undermines descriptions of action in which each moment in time is associated with only a single behavior.

It is no coincidence that naturalistic behavior shares many characteristics with neural activity: high dimensionality, time-evolving structure, variability, and organization at multiple temporal and spatial scales (Panzeri et al., 2015). However, efforts to develop methods to measure and represent the spontaneous behavior of unrestrained animals have lagged substantially behind those for neural activity. In part this is because of a long-standing focus on a limited set of assays that provide low-dimensional descriptions of complex patterns of action (e.g., the three-chamber social assay, the open field test, and the tail-suspension test), whose aim is to probe the psychological state of the animal and which therefore have been heavily used in the pharmaceutical industry (Crawley, 2003, 2008). The availability of commercial tracking software has contributed to a perception that automatically measuring and characterizing naturalistic behavior is either a trivial problem or one that has already been solved (Spink et al., 2001; Verbeek, 2005). To the contrary, several outstanding challenges need to be addressed if we seek to leverage the strengths of naturalistic behavior to better understand brain function. Below, we describe the current conceptual and technical framework that supports efforts in computational neuroethology; we later discuss additional progress that remains to be made.

Technology for Quantifying Behavior

Animal behavior inherently evolves over time. Capturing its time-varying structure requires measuring features of an animal's body and pose, tracking those features over time, and then identifying patterns that correspond to different movements, behaviors, or behavioral states (Figure 2). Given the pervasive use of cameras as sensors, here we consider this process largely from the perspective of video data of worms, flies, fish, and mice; however, all of the described steps have been applied to

other animals (e.g., bacteria and birds) and other types of data streams (e.g., accelerometry or ultrasonic vocalizations) (Berg, 1975a; Coffey et al., 2019; Markowitz et al., 2013; Van Segbroeck et al., 2017; Venkatraman et al., 2010). It is also important to note that although video of behaving animals is most often obtained from a single viewpoint using standard video cameras (Drai and Golani, 2001; Spink et al., 2001; Verbeek, 2005), recent innovations (including depth cameras and image fusion approaches) allow researchers to track freely behaving animals via video in 3D (Günel et al., 2019; Hong et al., 2015; Nath et al., 2019; Straw et al., 2011; Wiltschko et al., 2015).

Feature Extraction

A mouse's pose (i.e., its posture), a bird's beak, and the angle of a fly's wing are all features that may be relevant to an analysis of behavior. When studying insect locomotion, for example, one may wish to measure the position of each insect leg in relation to the other legs (Wilson, 1966). Two decades ago, this meant recording video of the insect and manually identifying the location of each of its legs or joints at each point in time (Strauss and Heisenberg, 1990). Early automated techniques required painting the animal's joint or leg, adding a small marker or dye, or using sophisticated imaging modalities like frustrated total internal reflection to highlight points of interest. Image processing algorithms could then automatically extract the location of these points of interests, obviating the need for manual identification (Bender et al., 2010; Kain et al., 2013; Mendes et al., 2013; Petrou and Webb, 2012).

Markerless approaches are an important alternative to these methods, as they enable automatic extraction of specific kinematic features without the use of surrogate markers. For animals whose anatomy is relatively simple, like worms or *Drosophila* larvae, simple computer vision algorithms can automatically identify such features without any human supervision (Broekmans et al., 2016; Gershow et al., 2012; Liu et al., 2018b; Stephens et al., 2008). For animals whose anatomy is more complex, "supervised" machine learning approaches—in which humans identify which features to track and provide labeled examples to train a machine learning-based feature detection algorithm—can be used to facilitate feature identification and tracking from video. Platforms that use this strategy (including CADABRA and JAABA) have been widely used in a variety of contexts (Branson et al., 2009; Dankert et al., 2009; Kabra et al., 2013). Similarly, the LocoMouse platform uses supervised machine learning techniques to recognize the precise position of paws, joints, the snout, and the tail in mice walking on a track (Machado et al., 2015). Recently, there has been an explosion in artificial neural network-based algorithms that can track human-identified anatomical keypoints in video (Graving et al., 2019; Mathis et al., 2018; Pereira et al., 2019). By tracking several of these keypoints in parallel, aspects of an animal's pose can be estimated using a limited amount of training data. These methods are gaining fast adoption for their versatility, ease of use, and robustness (Figure 3).

Any marker or keypoint tracking system can only track points that an experimenter can easily identify, such as joints and paws. In contrast, image properties can also be analyzed without human supervision to identify statistical regularities that recur in the pixel data themselves that can capture or reflect important
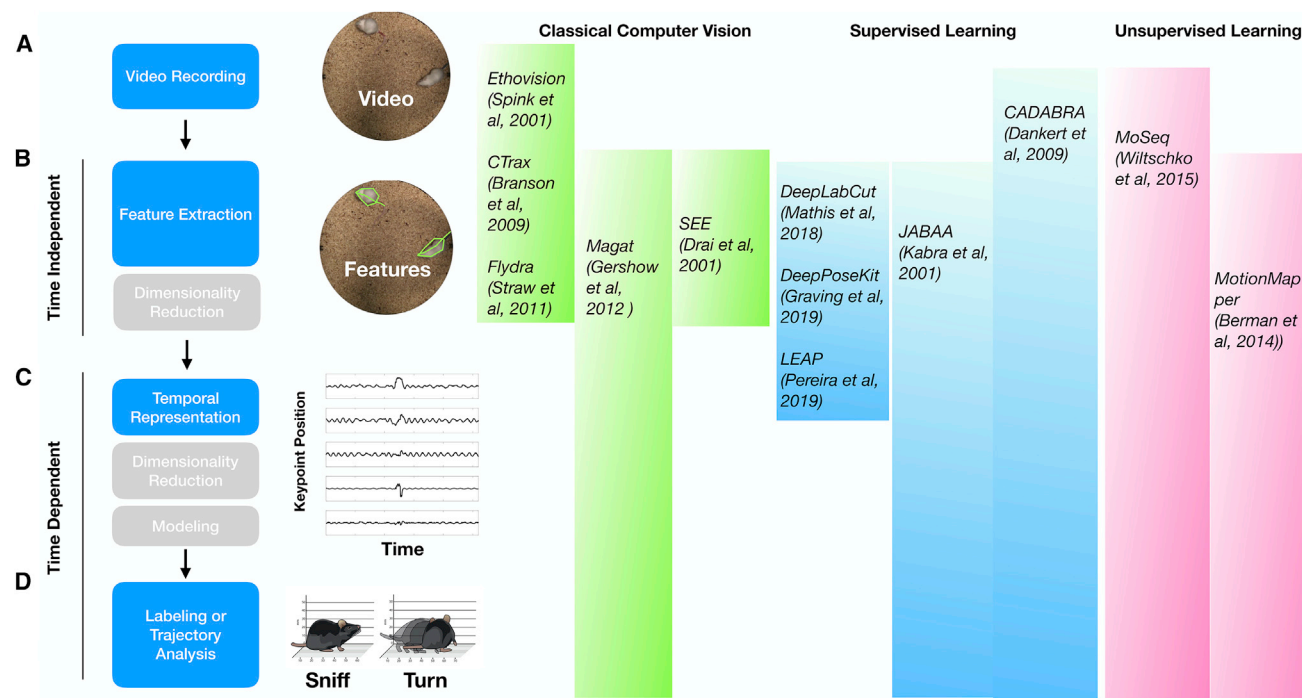
Figure 2. Typical Stages in a Behavioral Analysis Pipeline
Conceptual steps involved in any behavior analysis pipeline (left) with references to published example pipelines that are openly available for use (right). Shaded region indicates the steps that each method implements. In (a), animal behavior is recorded; in (b), features are extracted. Features may be interpretable, e.g., a mouse's paw position, or abstract, e.g., an algorithmically defined, weighted set of pixels. Automatic feature extraction can roughly be divided into algorithms that use classical computer vision feature detection, supervised learning, or unsupervised learning. High-dimensional descriptions of behavior features can optionally be approximated with fewer dimensions via dimensionality reduction. In (c), to leap from features to behavior requires first representing the temporal dynamics of the features. Feature dynamics can be represented in either the time or the frequency domain. Additional dimensionality reduction can optionally be performed at this stage. In (d), the resulting temporal dynamics are organized into behavior that can be either discrete or continuous. For discrete behavior representations, feature dynamics are clustered and labeled (e.g., "sniffing" or "turning") using either supervised or unsupervised machine learning. For continuous descriptions of behavior, trajectories through behavior space are analyzed and interpreted, e.g., using dynamical systems models.

features of an animal's behavior. Such "unsupervised" algorithm-driven approaches can extract features that are recognizable to a human (e.g., the degree to which the left wing is lifted, or a grimace in a mouse face) and can also yield features that a human would be hard-pressed to name (Berman et al., 2014; Musall et al., 2018; Stringer et al., 2019; Wiltschko et al., 2015). Because it is not always clear *a priori* which behavioral features are most relevant or informative in a particular experiment, the ability to identify unexpected features is a potential benefit of taking this sort of approach.

Large numbers of behavioral features are often required to capture behavior within a given experiment. As the high-dimensionality of behavioral data can be cumbersome, dimensionality reduction is commonly used after feature extraction to generate a lower-dimensional dataset that approximates the original feature set. Many such approaches—like principal components analysis (PCA)—are familiar from their use in neural data (Pang et al., 2016). Given a set of data about behavioral features, PCA identifies an ordered set of principal components (PCs), each of which constitutes a different "axis" representing progressively less variance present in the data. A reduced subset of these PCs can be used to approximately reconstruct the underlying features using fewer dimensions than present in the

original data. In *C. elegans*, for example, the animal's centerline captures most of the worm's behavior (Croll, 1975), but reconstructing this sinusoidal centerline requires tens of (x,y) coordinates. When transformed into a new basis set defined by PCA, these centerlines can be represented nearly as well by just three numbers, also known as the "eigenworm" (Stephens et al., 2008), thus providing a more tractable representation for use in subsequent analysis. It is important to note that the choice of how many dimensions to preserve matters, and that the largest sources of behavioral variance may not always reveal those features that most meaningfully describe a particular behavior.

Temporal Dynamics

Imagine a video of a worm crawling on an agar plate. A single frame of video provides no information about the animal's velocity, whether it is accelerating or decelerating, or even whether it is moving forward or backward. Multiple sequential frames, on the other hand, reveal the evolution of the worm's position and pose, which can be used to create a description of behavior. Behavioral representations can incorporate information about time by considering behavioral features in either the time domain or the frequency domain. Most commonly, analysis of behavior takes place exclusively in the time domain—researchers consider how a given behavioral feature (after extraction and
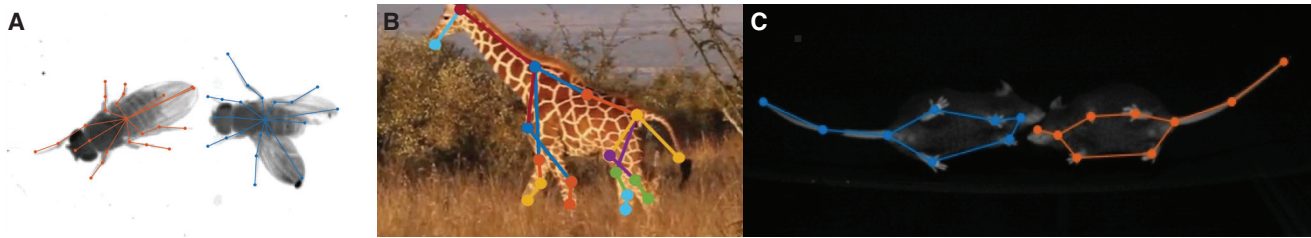
Figure 3. Pose Estimation
Recent advances in artificial neural networks allow identifiable points on an animal's surface (e.g., joints) to be automatically detected from images with minimal human supervision, even when animals interact or are in rich and complex backgrounds (Graving et al., 2019; Mathis et al., 2018; Pereira et al., 2019). Examples of (a) flies, (b) giraffes, and (c) mice are shown from Pereira et al., 2019 and T.D. Pereira, personal communication.

dimensionality reduction) evolves over time and then use that information to characterize behavioral dynamics or to identify behavioral motifs. For example, behavioral motifs in the worm have been identified by using a sliding time window to capture worm postures (Brown et al., 2013).

Alternatively, behavioral features can be first transformed into frequency space before considering how behavior evolves over time. Whereas a time domain analysis would represent a beating wing as the position of the wing over time, a frequency domain analysis would instead represent the same moving wing as its characteristic wing beat frequency. The MotionMapper platform takes this approach to format video data before downstream identification of behavioral motifs (Berman et al., 2016; Klibaite et al., 2017; Liu et al., 2018a; Pereira et al., 2019; Wang et al., 2016). Frequency domain analyses are well suited for representing cyclic motions (e.g., walking gaits, wing flapping, head swinging) and simplify the problem of identifying relationships between behaviors that are similar but out of phase (e.g., a walking bout starting with the right foot and a walking bout starting with the left foot).

Importantly, both time- and frequency-domain approaches require the experimenter to select a relevant timescale at which behavior is thought to be organized. In the time domain, this timescale defines the duration or distribution of durations of each behavioral motif, and in the frequency domain this timescale sets the lowest frequency that can be represented. In both cases the choice of timescale plays an important role in determining whether an animal's action is naturally represented as a single contiguous entity or as separate behavioral motifs.

Organizing Data and Assigning Labels

Behavior can be described as continuous (i.e., following a trajectory through a behavioral space), discrete, or a combination of the two (Fuchs, 1967). For example, the behavior of a worm exploring its environment can be elegantly described using dynamical systems representations as a continuous trajectory through posture space (Stephens et al., 2010). Similarly, brief elements of action (frequently corresponding to semantically low levels of behavior) can follow simple trajectories that are best described as continuous processes (Katsov et al., 2017; Wiltschko et al., 2015). However, animals can also express one or more of a large number of possible discrete behaviors that are stereotyped, distinct from each other, and are organized at a variety of timescales. Behavior therefore often requires labeling to identify which behavioral motifs or states are being expressed

at a particular time point: i.e., was the animal awake or asleep, mating or fighting? Labeling also defines when particular behaviors started and stopped and reveals the sequences of actions taken during an experiment. Traditionally, labeling has been done by hand—ethologists inspected either raw video of behavior or extracted behavioral feature data and then segmented those data using their own implicit criteria to label the types of actions being exhibited by a given subject. This sort of hand-labeling, when used to build transition-probability matrices, leads to the generation of ethograms, the lingua franca of classical ethologists (Baerends, 1976; Tinbergen, 1951). Relatively low-tech heuristic methods have automated some of these human intuitions—for example, when a mouse's nose is high enough to break a laser beam placed in an open field, the mouse is labeled as "rearing" (Crawley, 2003).

Machine learning is now revolutionizing the process of labeling behavioral data and generating ethograms. As with feature extraction, automated labeling often takes advantage of supervised learning approaches, in which machine learning algorithms are trained to identify particular behaviors based upon a set of human-curated examples (Branson et al., 2009; Dankert et al., 2009; Kabra et al., 2013; Machado et al., 2015; Mueller et al., 2019; Ravbar et al., 2019). JAABA, for example, includes an interface that allows researchers to indicate which video snippets correspond to a particular behavior of interest and which can then be used to train a classifier to identify when that behavior occurred (Kabra et al., 2013). Multiple classifiers can be built for different behaviors of interest, allowing researchers to flexibly extract complex information about when different behaviors are expressed (Robie et al., 2017).

Alternatively, unsupervised methods take advantage of statistical regularities in behavioral feature data to identify repeatedly used motifs of action (Berman, 2018; Brown and De Bivort, 2018; Nater et al., 2010). Two exemplar methods (of many) highlight the different paths that can be taken in using the inherent structure of behavioral data to define a description of behavior. MotionMapper takes as its input video data, and then, after pre-processing (which includes a PCA-based dimensionality-reduction step and the reformatting of these dimensionally reduced data into a frequency domain representation), these data are further dimensionally reduced by projecting them into a two-dimensional space through a non-linear method called t-stochastic nearest neighbor (tSNE) embedding (Berman et al., 2016; Klibaite et al., 2017; Liu et al., 2018a; Pereira et al.,

2019; Wang et al., 2016). All unsupervised behavioral methods have to address the problem of "lumping" versus "splitting"—the granularity at which a given behavioral data stream is broken up into parts. MotionMapper addresses this challenge by subjecting the behavioral tSNE embeddings to watershed-based clustering, which identifies reused motifs of action (which appear as peaks in the tSNE embedding) as well as behaviors that are less stereotyped (which appear as valleys). A comparison of clustering-based unsupervised behavioral classification methods can be found in Todd et al., 2017.

Motion Sequencing (MoSeq), on the other hand, takes advantage of a classic method in time-series analysis: the hidden Markov model (HMM) (Eddy, 2004). MoSeq (reviewed in Datta, 2019) takes as its input 3D data obtained from depth cameras and then uses statistical learning techniques to fit a hierarchical variant of the HMM, in which each behavioral component is modeled as a continuous auto-regressive process in pose space, while the components themselves (whose duration distributions are modeled based upon the data) are modeled using an HMM (Johnson et al., 2016; Markowitz et al., 2018; Pisanello et al., 2017; Wiltschko et al., 2015). The fitting procedures used by MoSeq allow it to flexibly learn the identity, number, and ordering of 3D behavioral components (called "syllables") for any given dataset. MoSeq—like all HMMs—is a generative model that after training can generate a synthetic 3D mouse (whose realism can be measured via statistical comparisons to held-out data). The fitting procedure underlying MoSeq explores different descriptions of behavior—"lumping" some movements together and "splitting" others—as it seeks a representation for behavior that best predicts held-out behavioral data.

It should be clear that all behavior pipelines—supervised or unsupervised—require the experimenter to make choices. These choices include which behavioral features to quantify; whether to analyze behavior in the time or frequency domain; which timescales to consider; whether to represent behavior as continuous, discrete, or both; and, if discrete, how to address the problem of granularizing behavior into elements. Optimal choices should reflect the nature of the ethological task the animal is solving and the inherent structure of the data; as our ability to record neural activity improves, these choices should be made with an eye toward maximizing our ability to understand the relationship between brain and behavior.

Emerging Methods in Computational Neuroethology Are Yielding Insight into the Relationship between Brain and Behavior

Methods for monitoring naturalistic behavior in the laboratory are largely in their infancy and yet have already made contributions to understanding the relationship between brain and behavior; below, we review highlights. Note that for brevity here we largely focus on analysis of freely behaving animals during neural recording or neural manipulation, although where relevant we point to examples of interesting analysis in more structured settings.

Forward Screens to Identify Neurons and Circuits for Behavior

In *Drosophila*, modern genetics has yielded collections of driver lines that, either alone or in combination, afford specific access to nearly every neuron in the fly brain (Jenett et al., 2012). The near-simultaneous development of these driver libraries and methods for automated behavioral classification is enabling a new type of forward screen, one that seeks to identify neurons that are necessary or sufficient for particular behaviors or behavioral components. This strategy has been particularly successful at identifying and dissecting neural circuits that underlie fly behaviors (Albin et al., 2015; Hoopfer et al., 2015; von Philipsborn et al., 2011). For example, to identify neurons linked to aggression, researchers expressed neural actuators or inhibitors (such as the thermogenetic activator TrpA1 or the constitutive inhibitor Kir2.1) in specific neural populations and used CADABRA and/or JAABA to quantify the behavioral influence of the targeted neuron (Asahina et al., 2014; Watanabe et al., 2017; Hoopfer et al., 2015). Because these automated methods dramatically reduce the time it takes to score videos, thousands of lines can be quickly analyzed, enabling both screens to focus on likely neurons of interest (i.e., those that express neuromodulators) as well as unbiased screens that survey the entire brain. This work has revealed key roles for tachykinin-expressing neurons, octopamine receptor-expressing neurons, P1 neurons, and fruitless-positive aSP2 neurons in driving or modulating aggression and has identified populations of neurons that control discrete behavioral modules that collectively comprise agonistic behaviors (Asahina et al., 2014; Duistermars et al., 2018; Hoopfer et al., 2015). Within individual modules, e.g., for threat display, quantitative analysis has further decomposed a given behavioral activity into separable motor elements that can be flexibly combined and investigated their neural control (Duistermars et al., 2018). This approach has also revealed epistatic relationships between different identified neural populations, for example, showing that P1 neurons and octopamine receptor-expressing neurons likely functionally converge upon aSP2 neurons (Watanabe et al., 2017). Similar optogenetic-based screens that leverage machine vision at scale have been used to identify neural circuits related to fly feeding and courtship (Albin et al., 2015).

Recent work has built upon this success to characterize the effects of neural activation and silencing on fly behavior more broadly (Robie et al., 2017). In this study, groups of 20 male and female flies were imaged simultaneously; video data were then used to identify a set of 128 hand-engineered features describing the behavior of each fly, which in turn was submitted to a set of supervised classifiers (built using JAABA) to identify specific behaviors (e.g., walking, chasing, mating). The behavioral consequences of thermogenetically activating each of more than 2,000 Gal4 lines (whose anatomy was previously characterized) was assessed using this method. The output of this process was a map linking sub-regions of the fly brain with particular behaviors; this map identified likely relationships (such as that between a series of visual areas and walking behaviors and between fruitless-positive neurons and wing extension) as well as an online resource to mine the data for further hypothesis generation.

These screens demonstrate the power of scalable machine vision-based methods to reveal the neural substrates of behavior. Complementary experiments have also been carried out using unsupervised behavioral classification methods. For example, a variant of hierarchical clustering has been used to

characterize the set of behavioral components and sequences that make up *Drosophila* larvae behavior before and after channelrhodopsin-based actuation of more than 1,000 Gal4 lines (Vogelstein et al., 2014). This experiment identified a set of 29 atomic movements falling into four basic categories (e.g., avoid, escape, backup, turn) and generated a look-up table linking each line to its characteristic behavioral consequence. A related set of experiments has been performed in the adult fly using Motion-Mapper (Berman et al., 2014; Cande et al., 2018). Flies whose descending neurons (which connect the CNS to effector motor centers in the ventral nerve cord) were optogenetically activated exhibited changes in each of five behavioral categories identified by MotionMapper; furthermore, these experiments revealed dependencies between the optogenetically induced behaviors and the behaviors expressed immediately prior (Cande et al., 2018).

Probing Sensorimotor Processing

Rich descriptions of innate animal behavior are proving critical for probing mechanisms of sensorimotor processing. For example, much of our understanding of neural mechanisms underlying sensory-driven navigation in *Drosophila* larvae comes from machine vision-based behavioral quantification. *Drosophila* larvae exhibit a stereotyped head-swing behavior that probes the sensory environment before the animal commits to a new movement direction. Computer vision-based behavior quantification systems such as described in Gershow et al., 2012 automatically detect these head swings and have been used to demonstrate that larvae temporally compare light intensities or chemical concentrations during head swings to chart future movements (Gepner et al., 2015; Gershow et al., 2012; Hernandez-Nunez et al., 2015; Kane et al., 2013; Schulze et al., 2015). Combined behavioral measurements and cell-specific inactivations have also identified specific lateral neurons downstream of the Bolwig's Organ that are crucial for mediating phototaxis (Kane et al., 2013). Similarly, work spanning many labs (reviewed in Calabrese, 2015) has used automated measures of behavior to reveal neural loci where chemosensory signals are integrated with photosensory signals for driving multi-sensory behavioral decisions (Gepner et al., 2015; Hernandez-Nunez et al., 2015; Schulze et al., 2015).

One common thread to these experiments is the combined use of optogenetic stimulation of sensory neurons and simple neural models to link sensory inputs to the animal's head-swings and turns. By automatically detecting a different set of behaviors—"hunching" and "bending"—Jovanic and colleagues conducted a complete neural dissection of the larvae's aversive response to mechanical stimulation, which included functionally and anatomically mapping a set of specific reciprocal inhibitory circuits from sensory input to motor output (Jovanic et al., 2016; Ohyama et al., 2013); this work illustrates how the automated analysis of naturalistic behavior can be used in concert with connectomics, electrophysiology, optogenetics, genetics, and modeling to probe a complete sensorimotor circuit. Similar work in *C. elegans* using a variety of methods (including MotionMapper) has quantified behavioral responses to the optogenetic activation of a nociceptive neuron and characterized the neural computations governing the behavioral responses of thousands of individual worms to optogenetic stimulation of their mechanosensory neurons (Liu et al., 2018a; Schwarz et al., 2015).

Automated behavior measures have also revealed new insights into the sensorimotor transformations driving social behaviors. For example, during courtship, adult *Drosophila* detect the sound and behavior of potential mates to dynamically coordinate their response. Male song production was long thought to be a fixed action pattern, and any variability in the song was considered noise. Careful measures of social behaviors during song production instead showed that the details of the song, such as choice of pulse versus sine, could be quantitatively predicted from the kinematic details of inter-animal behavior (Coen et al., 2014). Further measures of song and inter-animal behavior during calcium imaging revealed new forms of song, detailed neural correlates of male song production, and the neural coding of male song in the female auditory system (Clemens et al., 2015, 2018) as well as internal latent states that are correlated with neural processing (Calhoun et al., 2019). Generalized linear models were used throughout this body of work to mathematically relate inter-animal behavior, song features, and neural coding (Clemens and Murthy, 2017).

Relating Global Brain Dynamics to Naturalistic Behavior

Large-scale recording techniques now allow patterns of neural activity to be measured from hundreds to thousands of neurons at cellular resolution throughout the brain, providing an unprecedented view into neural computations and representational strategies (Williamson et al., 2019). Although such experiments still sub-sample neural activity in rodents or non-human primates, progress is being made in methods to characterize the global pattern of brain dynamics exhibited by a variety of simpler model organisms during naturalistic behavior. This work finds its origins in brain-scale recordings made at cellular resolution via calcium imaging in small transparent organisms like larval zebrafish or *C. elegans* during partial or complete immobilization. Investigations of whole-brain activity during fictive locomotion in zebrafish have been crucial for mapping and identifying the functional roles of new brain areas, including those for motor adaptation learning (Ahrens et al., 2012) and turning behavior (Dunn et al., 2016), and for discerning sensory versus motor areas (Chen et al., 2018). In *C. elegans*, measures of fictive locomotion inferred from interneuron activity have revealed stereotyped low-dimensional neural state-space trajectories that explain over 75% of the variance of brain-wide calcium activity during immobilization (Kato et al., 2015).

In the past few years, such whole-brain imaging at cellular or near cellular resolution has been adapted to freely moving zebrafish and *C. elegans* (Cong et al., 2017; Kim et al., 2017; Nguyen et al., 2016; Symvoulidis et al., 2017; Venkatachalam et al., 2016). Recordings from neurons in the head of larval zebrafish during innate foraging behaviors have identified brain regions related to the animal's swim bout angle and bout speed (Kim et al., 2017); similarly, recordings have been made during prey-capture to relate neural activity to the distance between the zebrafish and its paramecium prey or to features of the animal's eye-convergence angle and head orientation (Cong et al., 2017). And in *C. elegans*, combining whole-brain recordings during unrestrained movements with whole-body posture analysis has revealed new insights into where and how neural activity codes for locomotion (Scholz et al., 2018); in this work, a neural decoder of behavior was used to show that during unrestrained

behavior only a small fraction of the brain's neural dynamics (<5%) are explained by locomotory behavior, suggesting that the rest of the worm brain's activity may be involved in other computations.

Probing the Relationship between Pose Dynamics and Motor Circuits

Animals interact with the world by generating kinematics that support precise task-related movements (e.g., grasping an object) and large-scale changes in pose (e.g., rearing). Computational approaches have played a prominent role in quantifying the detailed kinematics of reaching or grasping behaviors in head-fixed primates and rodents and in relating measured kinematics to neural dynamics (e.g., Churchland et al., 2012; Guo et al., 2015). Improved measures of behavior now allow kinematic measurements during the self-generated locomotory behavior of unrestrained rodents. For example, LocoMouse was recently used to demonstrate that the temporal and spatial aspects of adaptation to a split-belt treadmill are dissociated and that a key locus within cerebellum is required to compensate for lateralized speed perturbations (Darmohray et al., 2019). While LocoMouse estimates gait parameters in a specialized apparatus (in which a camera images the mouse from two orientations), related future experiments exploring locomotion in the open field or other contexts will likely take advantage of deep learning-based point-tracking methods like LEAP, DeepLabCut, and DeepPoseKit (Graving et al., 2019; Mathis et al., 2018; Nath et al., 2019; Pereira et al., 2019).

While point-tracking methods are well suited to monitor the position of easily segmented features like paws or the base of the tail, clearly identifying keypoints can be difficult over much of the surface of many animals (like furry rodents). Alternative methods may therefore play an important complementary role in measuring the global pose dynamics expressed by animals as they generate naturalistic movements (Meyer et al., 2018; Venkatraman et al., 2010). For example, accelerometer data have also been used to parse spontaneous mouse behavior into motifs through the use of unsupervised affinity propagation-based techniques; this behavioral clustering has been combined with miniscope recordings to show that individual behavioral motifs are encoded by ensembles of direct and indirect striatal medium spiny neurons whose variance is systematically related to morphological similarities and differences in 3D behavior (Klaus et al., 2017).

Depth cameras have also been used to directly measure 3D information about an animal's pose dynamics and to use that information to explore brain circuits regulating action. By combining MoSeq with electrophysiological, multi-color photometry, and miniscope methods, neural correlates for 3D behavioral syllables have recently been identified in dorsolateral striatum (DLS) (Markowitz et al., 2018). These experiments reveal a systematic fluctuation in neural activity associated with transitions between behavioral syllables over time and an obligate role for DLS in generating appropriate behavioral sequences during both exploration and odor-guided naturalistic behaviors. Furthermore, MoSeq has been combined with optogenetic stimulation to reveal the differential consequences of activating the motor cortex, the dorsal striatum, and the ventral striatum (Pisanello et al., 2017; Wiltschko et al., 2015). These results are consistent with

similar results recently obtained using marker-based approaches to explore the relationship between 3D posture and activity in posterior parietal cortex (Mimica et al., 2018). Future benchmarking will reveal the trade-offs (if any) between direct measurements of 3D pose with specialized hardware (like depth cameras) and indirect estimates using more accessible hardware (like standard CCDs used to generate 3D keypoint tracking through image fusion or marker-based approaches).

Addressing the Challenges that Remain: A Way Forward for Behavior-Guided Discovery in the Brain

As is made plain by the examples above, although important progress is being made in relating brain activity to naturalistic behaviors, there are a host of conceptual and technical issues that remain. Behavior manifests itself as complex moment-to-moment trajectories, yet it is driven by often long-lasting internal states, e.g., sleep, wakefulness, hunger, and thirst, as well as external states such as the availability of resources. Thus a description of natural behavior must be hierarchically organized in time, and it remains unclear how to best identify behavioral hierarchies in a given dataset (Berman, 2018; Tao et al., 2019). Furthermore, identifying this sort of hierarchical structure requires large-scale data and, in particular, experimental set-ups and analysis pipelines that enable long-term assessment of behavior; this contrasts with most current naturalistic behavioral experiments, in which animal behavior is measured for minutes rather than hours (but see Jhuang et al., 2010; Ohayon et al., 2013).

Second, there is the problem of context: the richness of naturalistic behaviors is most fully observed in complex environments where sensory cues and affordances evoke the complete behavioral repertoire of the animal, and yet naturalistic behaviors in the lab are largely assessed in impoverished arenas. Future improvements in machine vision, virtual reality, and robotics should allow animals to explore increasingly realistic contexts while researchers monitor spontaneous naturalistic behavior (e.g., Del Grosso and Sirota, 2019; Meyer et al., 2018). One humble yet not fully addressed challenge is segmentation: if a lab animal is in an arena with a lot of stuff (as will be the case if, e.g., a mouse is imaged across its lifetime in its home cage during enrichment), it can be difficult to reliably distinguish the surface of the animal from objects around the animal. This is most vexing when considering social behaviors, which requires pose estimation of more than one animal at a time. Recent improvements in deep learning are helping to address this problem, as has the use of 3D cameras, but additional technical work will be required if we are to better understand the behavioral diversity of animals as they interact with realistic environments (which include conspecifics and predators) (Hong et al., 2015; Markowitz et al., 2018; Nath et al., 2019).

Finally, gaining access to multi-scale relationships between brain and behavior requires an understanding of how brain activity itself is organized and evolves over time. While much progress has been made in developing methods to infer structure in high-dimensional neural data, the field is still in flux (Cunningham and Yu, 2014; Williamson et al., 2019). As a consequence, most of the methods currently used to relate brain activity to complex naturalistic behavior (e.g., linear regression) are drawn from a

standard toolkit. An important road forward will be to build methods and model classes in which structure in the neural data and naturalistic behavioral data are jointly inferred (Glaser and Kording, 2016). Ideally, joint inference will allow trial-by-trial variability in neural data to be related to trial-by-trial behavioral performance (e.g., the kinematics that underlie the expression of any given instance of an identified behavioral motif, which in the context of naturalistic behavior can be considered a "trial"). It remains to be seen whether the approaches used to solve the problem of organizing behavioral data will be the same as those used to jointly infer joint structure in neural and behavioral data.

Given that human-imposed design decisions suffuse any quantification of animal behavior, what types of behavioral representations are maximally informative for understanding brain function? Historically, a given behavioral analysis method has been judged to be successful if it can be used to discern a difference after an experimental manipulation. This standard arose from the tradition of behavioral neurogenetics, where observable behavioral differences are used to support forward screens whose goal is to identify individual genes that contribute to the generation of behavior (Benzer, 1971). Importantly, this standard ducks the key conceptual questions in computational neuroethology: which distance metrics tell us whether any two behaviors are similar or different; how do we balance "lumping" versus "splitting"; and, in a given situation, how do we decide whether to characterize behavior as continuous, discrete, or both?

We argue that future development of methods in computational neuroethology should adopt a set of simple design principles. These principles do not address all the key questions posed above but instead are intended to prompt an ongoing conversation about how we measure behavior:

- Timescales. Humans have traditionally applied labels to animal behavior that operate on timescales of seconds or longer—think running, grooming, rearing—because that is the timescale at which perception and language most conveniently intersect. The availability of automated behavioral analysis methods circumvents this limit. Thus behavioral measurements and segmentations should ideally include the timescale at which neural variability is expected to occur. Furthermore, since slow behaviors can reflect fast neural activity and rapid behavioral events can reflect long-term neural dynamics, when possible behavioral descriptions should organize information hierarchically to facilitate multi-scale neurobehavioral discovery. Successfully meeting this imperative for many naturalistic behaviors will require ongoing improvements both in behavioral measurements (e.g., faster and higher-resolution cameras, better tracking, improved data compression) and in analytical methods for characterizing behavior at multiple timescales at once.
- Interpretability. Descriptions of behavior should lend themselves to hypothesis generation. Representations or models whose latent variables can be directly related to neural activity are more useful than those whose underlying variables lack obvious meaning; from this perspective, a behavioral representation can be thought of as "making sense" in light of neural activity (and vice versa), and such representations are useful for generating hypotheses about neural mechanisms underlying action. In a similar vein, representations whose latent variables correspond to a human intuition about behavior are more useful than those that do not; an "eigenworm" is interpretable by a human and therefore useful for articulating hypotheses about how worm behavior might be organized, whereas PCA over arbitrary collections of behavioral parameters may not yield to human intuition. While artificial neural networks are able to generate excellent predictions about behavior, it is often hard for a human to understand what the network actually learned that enabled it to make a prediction or to relate what the network learned back to a latent variable that might be detected in the brain. This does not mean eschewing approaches like deep learning—indeed, such methods are currently playing a central role in detecting behavioral features—but rather deploying them selectively to support hypothesis generation. One such example is the recent use of a variational autoencoder to reformat raw mouse videos before submission to an interpretable generative model (Johnson et al., 2016).
- Prediction. When possible, prediction should be adopted as a standard for judging the quality of behavioral representations. Effective behavioral representations should be able to predict behavior (in those circumstances when behavior is expected to be predictable) or neural activity; conversely, an effective behavioral representation should enable actions to be predicted from neural activity. As neuroscientists, this goal of being able to predict brain from behavior and vice versa strikes to the heart of our motivation for studying behavior. Therefore, prediction quality is a natural arbiter for deciding among model parameters or competing behavioral representations. Prediction also affords the possibility of finding the right balance between parsimony and richness, as testing the predictive performance of different models of behavior potentially offers a solution to the problem of "lumping" versus "splitting." Of course, one fundamental challenge that remains is deciding what sort of predictions are most relevant in a particular experimental context. In the long run, it is likely that multiple types of prediction (including of simultaneously measured neural data, of genomic or transcriptomic data, or of different treatment classes) will be required to fully assess and compare the performance of different behavioral analysis methods.

The Future of Computational Neuroethology

We are now gaining access to powerful tools for recording behavior, including machine learning-based methods to extract key features like joint angles and body postures, and for organizing this information to understand how behavior is structured and evolves over time. These tools will only get better, enabling increasingly dense characterization of action. Thanks to ongoing technical progress in both brain recording and behavioral analysis, we will soon face the challenge of relating high-dimensional time-varying neural data to high-dimensional time-varying behavioral data. A major goal for future research will be to identify behavioral representations that will give us insight into how

neural circuits create behavior. Meeting this goal will require a dedicated effort on the part of neuroscientists to build tools for characterizing naturalistic behavior with the same vigor and creativity that they have thus far largely reserved for measuring and characterizing neural activity.

In looking toward this future, it is helpful to ask what it means to "understand" the relationship between naturalistic behavior and the brain. Psychologists would likely agree that this answer requires a full account of those brain circuits that regulate a given behavior, including testable predictions about how circuit manipulations will affect behavior. Ethologists would wish to understand how behavior helps a given species prosper in its ecological niche, including accounts of how behavior emerges through evolutionary pressures and arises in each individual through the interplay of genetics and learning. These different levels of explanation are interdependent and equally valid (Barlow, 1961).

Yet the conceptual difficulties in reaching this understanding may sometimes appear rather daunting (Gomez-Marin et al., 2014; Jazayeri and Afraz, 2017; Krakauer et al., 2017). How do we search the immense brain for the circuits that are relevant for a given naturalistic behavior? How should we interpret the neural signals that we record, some of which may be as complex as the behavior itself and many of which may be irrelevant? Which is the right representation of behavior, and at what temporal and semantic scale should we look? How do we know which behavioral features are meaningful and which are idiosyncratic? How do we relate the goal we presume the animal is pursuing with the observed behavior? And when we finally describe a circuit, how do we think about the computation that it is intrinsically carrying out to support behavior, independently of the circuit implementation details?

We argue that progress in computational neuroethology will require researchers to propose mechanisms that might allow an animal to reach an ethological goal. Identifying possible mechanisms requires answering three questions that echo Marr's three levels of analysis of perception (Marr, 1982). First, what ethological goal is any given behavior meant to address, and how shall we measure whether a particular action helps an animal reach its goal? Second, given the available sensory inputs and physical constraints, what computational strategies (algorithms) could allow an animal to reach its goal? Third, how should algorithms that support goal-oriented behavior be implemented in the hardware of a nervous system? The book *Vehicles* by Valentino Braitenberg suggests an approach for addressing these questions: testable hypotheses about how the brain creates naturalistic task-driven behaviors can be generated by attempting to design simple automata that can accomplish that same task (Braitenberg, 1986). This perspective has successfully yielded insight into, e.g., the mechanisms and algorithms that enable bacteria to chemotax toward a chemical cue (Berg, 1975b).

Importantly, this process of mapping goals to algorithms produces normative models. These models are valuable in many ways: to generate hypotheses about which neural signals to look for in the brain with relation to behavior, to evaluate whether the observed neural signals are sufficient for a particular task, to assess which behavioral features are noise and which are diagnostic of the main design choices and trade-offs (for example,

sensitivity to input noise versus circuit complexity), and to taxonomize behaviors in the context of a goal at different scales of temporal and semantic resolution. The use of methods in computational neuroethology—whether focused on simple trained behaviors or complex, unrestrained patterns of action, whether done in single animals or at scale—will teach us about the structure of behavior; extracting meaning from these experiments and understanding how behavior meaningfully relates to brain activity will require a notion of the animal's goals in generating a behavior and, in the long run, normative models for how a brain might accomplish that goal. This approach has been very fruitful in understanding perception and learning (Nakayama and Shimojo, 1992; Navalpakkam et al., 2010; Reichardt et al., 1983; Shadlen and Newsome, 1998; Sutton and Barto, 1998). When taken with the conceptual and technical advances in computational neuroethology, we predict this perspective will play an equally powerful role in the study of naturalistic behavior.

DECLARATION OF INTERESTS

S.R.D. is a is a founder of Syllable Life Sciences.

REFERENCES

Ahrens, M.B., Li, J.M., Orger, M.B., Robson, D.N., Schier, A.F., Engert, F., and Portugues, R. (2012). Brain-wide neuronal dynamics during motor adaptation in zebrafish. Nature *485*, 471–477.

Albin, S.D., Kaun, K.R., Knapp, J.-M., Chung, P., Heberlein, U., and Simpson, J.H. (2015). A Subset of Serotonergic Neurons Evokes Hunger in Adult Drosophila. Curr. Biol. *25*, 2435–2440.

Anderson, D.J., and Perona, P. (2014). Toward a science of computational ethology. Neuron *84*, 18–31.

Asahina, K., Watanabe, K., Duistermars, B.J., Hoopfer, E., González, C.R., Eyjólfsdóttir, E.A., Perona, P., and Anderson, D.J. (2014). Tachykinin-expressing neurons control male-specific aggressive arousal in Drosophila. Cell *156*, 221–235.

Baerends, G.P. (1976). The functional organization of behaviour. Anim. Behav. *24*, 726–738.

Barlow, H.B. (1961). Possible Principles Underlying the Transformations of Sensory Messages. In Sensory Communication, W.A. Rosenblith, ed. (MIT Press).

Bender, J.A., Simpson, E.M., and Ritzmann, R.E. (2010). Computer-assisted 3D kinematic analysis of all leg joints in walking insects. PLoS ONE *5*, e13617.

Benzer, S. (1971). From the gene to behavior. JAMA *218*, 1015–1022.

Berg, H.C. (1975a). Chemotaxis in bacteria. Annu. Rev. Biophys. Bioeng. *4*, 119–136.

Berg, H.C. (1975b). Bacterial behaviour Nature. 1975 Apr 3;254(5499):389–392.

Berman, G.J. (2018). Measuring behavior across scales. BMC Biol. *16*, 23.

Berman, G.J., Choi, D.M., Bialek, W., and Shaevitz, J.W. (2014). Mapping the stereotyped behaviour of freely moving fruit flies. J. R. Soc. Interface *11*, 20140672.

Berman, G.J., Bialek, W., and Shaevitz, J.W. (2016). Predictability and hierarchy in Drosophila behavior. Proc. Natl. Acad. Sci. USA *113*, 11943–11948.

Braitenberg, V. (1986). Vehicles: Experiments in synthetic psychology (MIT Press).

Branson, K., Robie, A.A., Bender, J., Perona, P., and Dickinson, M.H. (2009). High-throughput ethomics in large groups of Drosophila. Nat. Methods *6*, 451–457.

Broekmans, O.D., Rodgers, J.B., Ryu, W.S., and Stephens, G.J. (2016). Resolving coiled shapes reveals new reorientation behaviors in C. elegans. eLife *5*, e17227.

Brown, A.E.X., and De Bivort, B. (2018). Ethology as a physical science. Nat. Phys. *14*, 653–657.

Brown, A.E.X., Yemini, E.I., Grundy, L.J., Jucikas, T., and Schafer, W.R. (2013). A dictionary of behavioral motifs reveals clusters of genes affecting Caenorhabditis elegans locomotion. Proc. Natl. Acad. Sci. USA *110*, 791–796.

Calabrese, R.L. (2015). In search of lost scent. eLife *4*, e08715.

Calhoun, A.J., Pillow, J.W., and Murthy, M. (2019). Unsupervised identification of the internal states that shape natural behavior. bioRxiv. https://doi.org/10.1101/691196.

Cande, J., Namiki, S., Qiu, J., Korff, W., Card, G.M., Shaevitz, J.W., Stern, D.L., and Berman, G.J. (2018). Optogenetic dissection of descending behavioral control in *Drosophila*. eLife *7*, 970.

Chen, X., Randi, F., Leifer, A.M., and Bialek, W. (2018). Searching for collective behavior in a small brain. arXiv, 181007623 [cond-mat, physics, qbio].

Churchland, M.M., Cunningham, J.P., Kaufman, M.T., Foster, J.D., Nuyujukian, P., Ryu, S.I., and Shenoy, K.V. (2012). Neural population dynamics during reaching. Nature *487*, 51–56.

Clemens, J., and Murthy, M. (2017). The Use of Computational Modeling to Link Sensory Processing with Behavior in Drosophila. In Decoding Neural Circuit Structure and Function: Cellular Dissection Using Genetic Model Organisms, A. Çelik and M.F. Wernet, eds. (Cham: Springer International Publishing), pp. 241–260.

Clemens, J., Girardin, C.C., Coen, P., Guan, X.-J., Dickson, B.J., and Murthy, M. (2015). Connecting Neural Codes with Behavior in the Auditory System of Drosophila. Neuron *87*, 1332–1343.

Clemens, J., Coen, P., Roemschied, F.A., Pereira, T.D., Mazumder, D., Aldarondo, D.E., Pacheco, D.A., and Murthy, M. (2018). Discovery of a New Song Mode in Drosophila Reveals Hidden Structure in the Sensory and Neural Drivers of Behavior. Curr. Biol. *28*, 2400–2412.e6.

Coen, P., Clemens, J., Weinstein, A.J., Pacheco, D.A., Deng, Y., and Murthy, M. (2014). Dynamic sensory cues shape song structure in Drosophila. Nature *507*, 233–237.

Coffey, K.R., Marx, R.G., and Neumaier, J.F. (2019). DeepSqueak: a deep learning-based system for detection and analysis of ultrasonic vocalizations. Neuropsychopharmacology *44*, 859–868.

Cong, L., Wang, Z., Chai, Y., Hang, W., Shang, C., Yang, W., Bai, L., Du, J., Wang, K., and Wen, Q. (2017). Rapid whole brain imaging of neural activity in freely behaving larval zebrafish (*Danio rerio*). eLife *6*, e28158.

Crawley, J.N. (2003). Behavioral phenotyping of rodents. Comp. Med. *53*, 140–146.

Crawley, J.N. (2008). Behavioral phenotyping strategies for mutant mice. Neuron *57*, 809–818.

Croll, N.A. (1975). Components and patterns in the behavior of the nematode Caenorhabditis elegans. J. Zool. *176*, 159–176.

Cunningham, J.P., and Yu, B.M. (2014). Dimensionality reduction for large-scale neural recordings. Nat. Neurosci. *17*, 1500–1509.

Dankert, H., Wang, L., Hoopfer, E.D., Anderson, D.J., and Perona, P. (2009). Automated monitoring and analysis of social behavior in Drosophila. Nat. Methods *6*, 297–303.

Darmohray, D.M., Jacobs, J.R., Marques, H.G., and Carey, M.R. (2019). Spatial and Temporal Locomotor Learning in Mouse Cerebellum. Neuron *102*, 217–231.e4.

Datta, S.R. (2019). Q&A: Understanding the composition of behavior. BMC Biol. *17*, 44.

Dawkins, R. (1976). Hierarchical organisation: A candidate principle for ethology. In Growing Points in Ethology, P.P.G. Bateson and R.A. Hinde, eds. (Cambridge University Press).

Del Grosso, N.A., and Sirota, A. (2019). Ratcave: A 3D graphics python package for cognitive psychology experiments. Behav. Res. Methods *10*, 433.

Domjan, M. (1987). Comparative psychology and the study of animal learning. J. Comp. Psychol. *101*, 237–241.

Drai, D., and Golani, I. (2001). SEE: a tool for the visualization and analysis of rodent exploratory behavior. Neurosci. Biobehav. Rev. *25*, 409–426.

Duistermars, B.J., Pfeiffer, B.D., Hoopfer, E.D., and Anderson, D.J. (2018). A Brain Module for Scalable Control of Complex, Multi-motor Threat Displays. Neuron *100*, 1474–1490.e4.

Dunn, T.W., Mu, Y., Narayan, S., Randlett, O., Naumann, E.A., Yang, C.-T., Schier, A.F., Freeman, J., Engert, F., and Ahrens, M.B. (2016). Brain-wide mapping of neural activity controlling zebrafish exploratory locomotion. eLife *5*, e12741.

Eddy, S.R. (2004). What is a hidden Markov model? Nat. Biotechnol. *22*, 1315–1316.

Flash, T., and Hochner, B. (2005). Motor primitives in vertebrates and invertebrates. Curr. Opin. Neurobiol. *15*, 660–666.

Foster, D.J. (2017). Replay Comes of Age. Annu. Rev. Neurosci. *40*, 581–602.

Fuchs, A.F. (1967). Saccadic and smooth pursuit eye movements in the monkey. J. Physiol. *191*, 609–631.

Gepner, R., Mihovilovic Skanata, M., Bernat, N.M., Kaplow, M., and Gershow, M. (2015). Computations underlying Drosophila photo-taxis, odor-taxis, and multi-sensory integration. eLife *4*, e06229.

Gershow, M., Berck, M., Mathew, D., Luo, L., Kane, E.A., Carlson, J.R., and Samuel, A.D.T. (2012). Controlling airborne cues to study small animal navigation. Nat. Methods *9*, 290–296.

Glaser, J.I., and Kording, K.P. (2016). The Development and Analysis of Integrated Neuroscience Data. Front. Comput. Neurosci. *10*, 11.

Gomez-Marin, A., Paton, J.J., Kampff, A.R., Costa, R.M., and Mainen, Z.M. (2014). Big behavioral data: psychology, ethology and the foundations of neuroscience. Nat. Neurosci *17*, 1455–1462.

Graving, J.M., Chae, D., Naik, H., Li, L., Koger, B., Costelloe, B.R., and Couzin, I.D. (2019). Fast and robust animal pose estimation. bioRxiv. https://doi.org/10.1101/620245.

Günel, S., Rhodin, H., Morales, D., Campagnolo, J., Ramdya, P., and Fua, P. (2019). DeepFly3D: A deep learning-based approach for 3D limb and appendage tracking in tethered, adult Drosophila. bioRxiv, doi: 0.1101/640375.

Guo, J.-Z., Graves, A.R., Guo, W.W., Zheng, J., Lee, A., Rodríguez-González, J., Li, N., Macklin, J.J., Phillips, J.W., Mensh, B.D., et al. (2015). Cortex commands the performance of skilled movement. eLife *4*, e10774.

Hartley, T., Lever, C., Burgess, N., and O'Keefe, J. (2013). Space in the brain: how the hippocampal formation supports spatial cognition. Philos. Trans. R. Soc. Lond. B Biol. Sci. *369*, 20120510.

Hernandez-Nunez, L., Belina, J., Klein, M., Si, G., Claus, L., Carlson, J.R., and Samuel, A.D. (2015). Reverse-correlation analysis of navigation dynamics in Drosophila larva using optogenetics. eLife. https://doi.org/10.7554/eLife.06225.

Hong, W., Kennedy, A., Burgos-Artizzu, X.P., Zelikowsky, M., Navonne, S.G., Perona, P., and Anderson, D.J. (2015). Automated measurement of mouse social behaviors using depth sensing, video tracking, and machine learning. Proc. Natl. Acad. Sci. USA *112*, E5351–E5360.

Hoopfer, E.D., Jung, Y., Inagaki, H.K., Rubin, G.M., and Anderson, D.J. (2015). P1 interneurons promote a persistent internal state that enhances inter-male aggression in Drosophila. eLife *4*, 2700.

Jazayeri, M., and Afraz, A. (2017). Navigating the Neural Space in Search of the Neural Code. Neuron *93*, 1003–1014.

Jenett, A., Rubin, G.M., Ngo, T.-T.B., Shepherd, D., Murphy, C., Dionne, H., Pfeiffer, B.D., Cavallaro, A., Hall, D., Jeter, J., et al. (2012). A GAL4-driver line resource for Drosophila neurobiology. Cell Rep. *2*, 991–1001.

Jhuang, H., Garrote, E., Mutch, J., Yu, X., Khilnani, V., Poggio, T., Steele, A.D., and Serre, T. (2010). Automated home-cage behavioural phenotyping of mice. Nat. Commun. *1*, 68.

Johnson, M., Duvenaud, D.K., Wiltschko, A., Adams, R.P., and Datta, S.R. (2016). Composing graphical models with neural networks for structured representations and fast inference. Advances in Neural Information Processing Systems *29*, 2946–2954.

Jovanic, T., Schneider-Mizell, C.M., Shao, M., Masson, J.-B., Denisov, G., Fetter, R.D., Mensh, B.D., Truman, J.W., Cardona, A., and Zlatic, M. (2016). Competitive Disinhibition Mediates Behavioral Choice and Sequences in Drosophila. Cell *167*, 858–870.e19.

Juavinett, A.L., Erlich, J.C., and Churchland, A.K. (2018). Decision-making behaviors: weighing ethology, complexity, and sensorimotor compatibility. Curr. Opin. Neurobiol. *49*, 42–50.

Kabra, M., Robie, A.A., Rivera-Alba, M., Branson, S., and Branson, K. (2013). JAABA: interactive machine learning for automatic annotation of animal behavior. Nat. Methods *10*, 64–67.

Kain, J., Stokes, C., Gaudry, Q., Song, X., Foley, J., Wilson, R., and de Bivort, B. (2013). Leg-tracking and automated behavioural classification in *Drosophila*. Nat. Commun. *4*, 1910.

Kandel, E.R., Dudai, Y., and Mayford, M.R. (2014). The molecular and systems biology of memory. Cell *157*, 163–186.

Kane, E.A., Gershow, M., Afonso, B., Larderet, I., Klein, M., Carter, A.R., de Bivort, B.L., Sprecher, S.G., and Samuel, A.D.T. (2013). Sensorimotor structure of Drosophila larva phototaxis. Proc. Natl. Acad. Sci. USA *110*, E3868–E3877.

Kato, S., Kaplan, H.S., Schrödel, T., Skora, S., Lindsay, T.H., Yemini, E., Lockery, S., and Zimmer, M. (2015). Global brain dynamics embed the motor command sequence of Caenorhabditis elegans. Cell *163*, 656–669.

Katsov, A.Y., Freifeld, L., Horowitz, M., Kuehn, S., and Clandinin, T.R. (2017). Dynamic structure of locomotor behavior in walking fruit flies. eLife *6*, e26410.

Kim, D.H., Kim, J., Marques, J.C., Grama, A., Hildebrand, D.G.C., Gu, W., Li, J.M., and Robson, D.N. (2017). Pan-neuronal calcium imaging with cellular resolution in freely swimming zebrafish. Nat. Methods *14*, 1107–1114.

Klaus, A., Martins, G.J., Paixao, V.B., Zhou, P., Paninski, L., and Costa, R.M. (2017). The Spatiotemporal Organization of the Striatum Encodes Action Space. Neuron *95*, 1171–1180.e7.

Klibaite, U., Berman, G.J., Cande, J., Stern, D.L., and Shaevitz, J.W. (2017). An unsupervised method for quantifying the behavior of paired animals. Phys. Biol. *14*, 015006.

Krakauer, J.W., Ghazanfar, A.A., Gomez-Marin, A., MacIver, M.A., and Poeppel, D. (2017). Neuroscience Needs Behavior: Correcting a Reductionist Bias. Neuron *93*, 480–490.

Liu, M., Sharma, A.K., Shaevitz, J.W., and Leifer, A.M. (2018a). Temporal processing and context dependency in *Caenorhabditis elegans* response to mechanosensation. eLife *7*, e36419.

Liu, M., Sharma, A.K., Shaevitz, J.W., and Leifer, A.M. (2018b). Temporal processing and context dependency in C. elegans mechanosensation. arXiv, 180304085 [physics, q-bio].

Machado, A.S., Darmohray, D.M., Fayad, J., Marques, H.G., and Carey, M.R. (2015). A quantitative framework for whole-body coordination reveals specific deficits in freely walking ataxic mice. eLife *4*, 18.

Manoli, D.S., Meissner, G.W., and Baker, B.S. (2006). Blueprints for behavior: genetic specification of neural circuitry for innate behaviors. Trends Neurosci. *29*, 444–451.

Mar, A.C., Horner, A.E., Nilsson, S.R.O., Alsiö, J., Kent, B.A., Kim, C.H., Holmes, A., Saksida, L.M., and Bussey, T.J. (2013). The touchscreen operant platform for assessing executive function in rats and mice. Nat. Protoc. *8*, 1985–2005.

Markowitz, J.E., Ivie, E., Kligler, L., and Gardner, T.J. (2013). Long-range order in canary song. PLoS Comput. Biol. *9*, e1003052.

Markowitz, J.E., Gillis, W.F., Beron, C.C., Neufeld, S.Q., Robertson, K., Bhagat, N.D., Peterson, R.E., Peterson, E., Hyun, M., Linderman, S.W., et al. (2018). The Striatum Organizes 3D Behavior via Moment-to-Moment Action Selection. Cell *174*, 44–58.e17.

Marr, D. (1982). Vision: A computational investigation into the human representation and processing of visual information (Henry Holt and Co., Inc.).

Mathis, A., Mamidanna, P., Cury, K.M., Abe, T., Murthy, V.N., Mathis, M.W., and Bethge, M. (2018). DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. Nat. Neurosci. *21*, 1281–1289.

Medan, V., and Preuss, T. (2014). The Mauthner-cell circuit of fish as a model system for startle plasticity. J. Physiol. Paris *108*, 129–140.

Mendes, C.S., Bartos, I., Akay, T., Márka, S., and Mann, R.S. (2013). Quantification of gait parameters in freely walking wild type and sensory deprived Drosophila melanogaster. eLife *2*, e00231.

Meyer, A.F., Poort, J., O'Keefe, J., Sahani, M., and Linden, J.F. (2018). A Head-Mounted Camera System Integrates Detailed Behavioral Monitoring with Multichannel Electrophysiology in Freely Moving Mice. Neuron *100*, 46–60.e7.

Mimica, B., Dunn, B.A., Tombaz, T., Bojja, V.P.T., and Whitlock, J.R. (2018). Efficient cortical coding of 3D posture in freely behaving rats. Science *362*, 584–589.

Minderer, M., Harvey, C.D., Donato, F., and Moser, E.I. (2016). Neuroscience: Virtual reality explored. Nature *533*, 324–325.

Mueller, J.M., Ravbar, P., Simpson, J.H., and Carlson, J.M. (2019). Drosophila melanogaster grooming possesses syntax with distinct rules at different temporal scales. PLoS Comput. Biol. *15*, e1007105.

Musall, S., Kaufman, M.T., Gluf, S., and Churchland, A. (2018). Movement-related activity dominates cortex during sensory-guided decision making. bioRxiv. https://doi.org/10.1101/308288.

Musall, S., Urai, A., Sussillo, D., and Churchland, A. (2019). Harnessing behavioral diversity to understand circuits for cognition. arXiv, arXiv:1906.09622 [q-bio.NC].

Nakayama, K., and Shimojo, S. (1992). Experiencing and perceiving visual surfaces. Science *257*, 1357–1363.

Nater, F., Grabner, H., and Gool, L.V. (2010). Exploiting simple hierarchies for unsupervised human behavior analysis. In Paper presented at: 2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition.

Nath, T., Mathis, A., Chen, A.C., Patel, A., Bethge, M., and Mathis, M.W. (2019). Using DeepLabCut for 3D markerless pose estimation across species and behaviors. Nat. Protoc. *14*, 2152–2176.

Navalpakkam, V., Koch, C., Rangel, A., and Perona, P. (2010). Optimal reward harvesting in complex perceptual environments. Proc. Natl. Acad. Sci. USA *107*, 5232–5237.

Nguyen, J.P., Shipley, F.B., Linder, A.N., Plummer, G.S., Liu, M., Setru, S.U., Shaevitz, J.W., and Leifer, A.M. (2016). Whole-brain calcium imaging with cellular resolution in freely behaving Caenorhabditis elegans. Proc. Natl. Acad. Sci. USA *113*, E1074–E1081.

Ohayon, S., Avni, O., Taylor, A.L., Perona, P., and Roian Egnor, S.E. (2013). Automated multi-day tracking of marked mice for the analysis of social behaviour. J. Neurosci. Methods *219*, 10–19.

Ohyama, T., Jovanic, T., Denisov, G., Dang, T.C., Hoffmann, D., Kerr, R.A., and Zlatic, M. (2013). High-throughput analysis of stimulus-evoked behaviors in Drosophila larva reveals multiple modality-specific escape strategies. PLoS ONE 8, e71706.

Pang, R., Lansdell, B.J., and Fairhall, A.L. (2016). Dimensionality reduction in neuroscience. Curr. Biol. 26, R656–R660.

Panzeri, S., Macke, J.H., Gross, J., and Kayser, C. (2015). Neural population coding: combining insights from microscopic and mass signals. Trends Cogn. Sci. 19, 162–172.

Panzeri, S., Harvey, C.D., Piasini, E., Latham, P.E., and Fellin, T. (2017). Cracking the Neural Code for Sensory Perception by Combining Statistics, Intervention, and Behavior. Neuron 93, 491–507.

Pereira, T.D., Aldarondo, D.E., Willmore, L., Kislin, M., Wang, S.S.H., Murthy, M., and Shaevitz, J.W. (2019). Fast animal pose estimation using deep neural networks. Nat. Methods 16, 117–125.

Petrou, G., and Webb, B. (2012). Detailed tracking of body and leg movements of a freely walking female cricket during phonotaxis. J. Neurosci. Methods 203, 56–68.

Pisanello, F., Mandelbaum, G., Pisanello, M., Oldenburg, I.A., Sileo, L., Markowitz, J.E., Peterson, R.E., Della Patria, A., Haynes, T.M., Emara, M.S., et al. (2017). Dynamic illumination of spatially restricted or large brain volumes via a single tapered optical fiber. Nat. Neurosci. 20, 1180–1188.

Ravbar, P., Branson, K., and Simpson, J.H. (2019). An automatic behavior recognition system classifies animal behaviors using movements and their temporal context. J. Neurosci. Methods 326, 108352.

Reichardt, W., Poggio, T., and Hausen, K. (1983). Figure-ground discrimination by relative movement in the visual system of the fly. Biol. Cybern. 46, 1–30.

Robie, A.A., Hirokawa, J., Edwards, A.W., Umayam, L.A., Lee, A., Phillips, M.L., Card, G.M., Korff, W., Rubin, G.M., Simpson, J.H., et al. (2017). Mapping the Neural Substrates of Behavior. Cell 170, 393–406.e28.

Rowland, D.C., Roudi, Y., Moser, M.-B., and Moser, E.I. (2016). Ten Years of Grid Cells. Annu. Rev. Neurosci. 39, 19–40.

Scholz, M., Linder, A.N., Randi, F., Sharma, A.K., Yu, X., Shaevitz, J.W., and Leifer, A. (2018). Predicting natural behavior from whole-brain neural dynamics. bioRxiv. https://doi.org/10.1101/445643.

Schulze, A., Gomez-Marin, A., Rajendran, V.G., Lott, G., Musy, M., Ahammad, P., Deogade, A., Sharpe, J., Riedl, J., Jarriault, D., et al. (2015). Dynamical feature extraction at the sensory periphery guides chemotaxis. eLife 4, e06694.

Schwarz, R.F., Branicky, R., Grundy, L.J., Schafer, W.R., and Brown, A.E.X. (2015). Changes in Postural Syntax Characterize Sensory Modulation and Natural Variation of C. elegans Locomotion. PLoS Comput. Biol. 11, e1004322.

Shadlen, M.N., and Newsome, W.T. (1998). The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. J. Neurosci. 18, 3870–3896.

Simmons, P., and Young, D. (1999). Nerve Cells and Animal Behavior (Cambridge University Press).

Spink, A.J., Tegelenbosch, R.A., Buma, M.O., and Noldus, L.P. (2001). The EthoVision video tracking system–a tool for behavioral phenotyping of transgenic mice. Physiol. Behav. 73, 731–744.

Stephens, G.J., Johnson-Kerner, B., Bialek, W., and Ryu, W.S. (2008). Dimensionality and dynamics in the behavior of C. elegans. PLoS Comput. Biol. 4, e1000028.

Stephens, G.J., Johnson-Kerner, B., Bialek, W., and Ryu, W.S. (2010). From modes to movement in the behavior of Caenorhabditis elegans. PLoS ONE 5, e13914.

Strauss, R., and Heisenberg, M. (1990). Coordination of legs during straight walking and turning in Drosophila melanogaster. J. Comp. Physiol. A Neuroethol. Sens. Neural Behav. Physiol. 167, 403–412.

Straw, A.D., Branson, K., Neumann, T.R., and Dickinson, M.H. (2011). Multi-camera real-time three-dimensional tracking of multiple flying animals. J. R. Soc. Interface 8, 395–409.

Stringer, C., Pachitariu, M., Steinmetz, N., Reddy, C.B., Carandini, M., and Harris, K.D. (2019). Spontaneous behaviors drive multidimensional, brainwide activity. Science 364, eaav7893.

Sutton, R.S., and Barto, A.G. (1998). Introduction to reinforcement learning, Vol. 2 (MIT Press).

Symvoulidis, P., Lauri, A., Stefanoiu, A., Cappetta, M., Schneider, S., Jia, H., Stelzl, A., Koch, M., Perez, C.C., Myklatun, A., et al. (2017). NeuBtracker-imaging neurobehavioral dynamics in freely behaving fish. Nat. Methods 14, 1079–1082.

Tao, L., Ozarkar, S., Beck, J.M., and Bhandawat, V. (2019). Statistical structure of locomotion and its modulation by odors. eLife 8, 425.

Tinbergen, N. (1951). The study of instinct (Oxford: Clarendon Press).

Tinbergen, N. (1963). On aims and methods of ethology. Z. Tierpsychol. 20, 410–433.

Todd, J.G., Kain, J.S., and de Bivort, B.L. (2017). Systematic exploration of unsupervised methods for mapping behavior. Phys. Biol. 14, 015002.

Van Segbroeck, M., Knoll, A.T., Levitt, P., and Narayanan, S. (2017). MUPET-Mouse Ultrasonic Profile ExTraction: A Signal Processing Tool for Rapid and Unsupervised Analysis of Ultrasonic Vocalizations. Neuron 94, 465–485.e5.

Venkatachalam, V., Ji, N., Wang, X., Clark, C., Mitchell, J.K., Klein, M., Tabone, C.J., Florman, J., Ji, H., Greenwood, J., et al. (2016). Pan-neuronal imaging in roaming Caenorhabditis elegans. Proc. Natl. Acad. Sci. USA 113, E1082–E1088.

Venkatraman, S., Jin, X., Costa, R.M., and Carmena, J.M. (2010). Investigating neural correlates of behavior in freely behaving rodents using inertial sensors. J. Neurophysiol. 104, 569–575.

Verbeek, J. (2005). Rodent behavior annotation from video. ACM Transactions on Multimedia Computing, Communications, and Applications-TOMCCAP.

Vogelstein, J.T., Park, Y., Ohyama, T., Kerr, R.A., Truman, J.W., Priebe, C.E., and Zlatic, M. (2014). Discovery of brainwide neural-behavioral maps via multiscale unsupervised structure learning. Science 344, 386–392.

von Philipsborn, A.C., Liu, T., Yu, J.Y., Masser, C., Bidaye, S.S., and Dickson, B.J. (2011). Neuronal control of Drosophila courtship song. Neuron 69, 509–522.

Wang, Q., Taliaferro, J.M., Klibaite, U., Hilgers, V., Shaevitz, J.W., and Rio, D.C. (2016). The PSI-U1 snRNP interaction regulates male mating behavior in Drosophila. Proc. Natl. Acad. Sci. USA 113, 5269–5274.

Watanabe, K., Chiu, H., Pfeiffer, B.D., Wong, A.M., Hoopfer, E.D., Rubin, G.M., and Anderson, D.J. (2017). A Circuit Node that Integrates Convergent Input from Neuromodulatory and Social Behavior-Promoting Neurons to Control Aggression in Drosophila. Neuron 95, 1112–1128.e7.

Williamson, R.C., Doiron, B., Smith, M.A., and Yu, B.M. (2019). Bridging large-scale neuronal recordings and large-scale network models using dimensionality reduction. Curr. Opin. Neurobiol. 55, 40–47.

Wilson, D.M. (1966). Insect walking. Annu. Rev. Entomol. 11, 103–122.

Wiltschko, A.B., Johnson, M.J., Iurilli, G., Peterson, R.E., Katon, J.M., Pashkovski, S.L., Abraira, V.E., Adams, R.P., and Datta, S.R. (2015). Mapping Sub-Second Structure in Mouse Behavior. Neuron 88, 1121–1135.