



# The what, how, and why of naturalistic behavior

Ann Kennedy


## Abstract

In the past few years, advances in machine learning have fueled an explosive growth of descriptive and generative models of animal behavior. These new approaches offer higher levels of detail and granularity than has previously been possible, allowing for fine-grained segmentation of animals' actions and precise quantitative mappings between an animal's sensory environment and its behavior. How can these new methods help us understand the governing principles shaping complex and naturalistic behavior? In this review, we will recap ways in which our ability to detect and model behavior have improved in recent years, and consider how these techniques might be used to revisit classical normative theories of behavioral control.

## Addresses

Department of Neuroscience, Northwestern University Feinberg School of Medicine, Chicago, IL 60611, USA

Corresponding author: Kennedy, Ann ([ann.kennedy@northwestern.edu](mailto:ann.kennedy@northwestern.edu))

 (Kennedy A.)

Current Opinion in Neurobiology 2022, 74:102549

This review comes from a themed issue on **Neurobiology of Behavior 2022**

Edited by **Tiago Branco** and **Mala Murthy**

For complete overview of the section, please refer the article collection - [Neurobiology of Behavior 2022](#)

Available online 7 May 2022

<https://doi.org/10.1016/j.conb.2022.102549>

0959-4388/© 2022 Elsevier Ltd. All rights reserved.

## Introduction

Animal behavior is shaped by an interaction of past experience, internal motivational state, and external environmental cues. Behaviors critical for survival, such as feeding and social interaction, are furthermore influenced by the genetically encoded wiring of species-specific neural circuitry. Given this diverse set of influences, the analysis of behavior informs scientific progress in fields spanning genetics, ecology, neuroscience, economics, and robotics. The theories and models of behavior generated by these fields vary in their objective, language, and level of mechanistic detail. They can be broadly divided into categories of **descriptive models** that identify **what** an agent did, **generative models** that tell us **how** an agent can be expected to

behave given a set of conditions, and **normative models** that identify **why** a behavior occurred by defining the underlying principles that inform its structure.

Machine learning tools for animal posture and behavior analysis are becoming increasingly accessible to the research community, fueling new approaches to the quantification, and study of animal behavior. Many of these newly developed methods are descriptive models that improve the throughput and granularity with which we can quantify behavior, bringing its analysis to the realm of big data. An open question is how these methods can give us insight into the organization and control of behavior itself, in the form of generative and normative models. In this review, we will briefly recount scientific applications of descriptive behavioral analysis, before turning to how detailed pose and behavior data might also be used to build generative models and test classical theories about the principles that shape behavioral organization.

## Descriptive models: Turning video files into data

In many applications, the use of computational methods for behavior analysis can be likened to automation of manufacturing: the things produced are not fundamentally different from what was around before, but now they can be produced with greater precision and at scale. Several recent reviews [1–8] have described the use of modern computer vision or machine learning methods to characterize the pose, kinematics, or actions of behaving agents. These methods, particularly the translation from pose to behavior, are areas of active development in machine learning (see **Box 1**). At the heart of many approaches is markerless pose estimation, which characterizes animal posture in terms of a set of experimenter-defined 2D or 3D “key points” [9–12]. Pose estimates or other features extracted from behavior video can be used to quantify postural dynamics of animals [13–15], or paired with supervised (experimenter-defined) behavior detection [16–18] or unsupervised (data-defined) behavior discovery [19–23] methods to segment continuous posture data into actions (**Figure 1a**).

The practical advantage of these methods is twofold. First, they accelerate data processing when applied at scale. By training algorithms to detect behaviors of interest, researchers can distill hundreds of hours of video into precise and detailed records of animals' actions that

**Box 1. Future challenges for descriptive models of behavior.**

**Performance:** A particular challenge facing unsupervised behavior discovery is deciding what constitutes a “good” representation of behavior [74]. Common metrics for evaluating the quality of an unsupervised behavior discovery method include agreement with human annotations, ability to account for variance in neural data, and utility in distinguishing between strains or conditions (Figure 1b), however many articles apply these only to in-house datasets, making it difficult to compare methods. Establishing benchmark datasets for evaluating behavior algorithms may help overcome this issue [75].

**Generalization:** Behavior classifiers and pose estimators trained in one lab rarely perform well out-of-the-box when used by another lab, unless the two groups standardize their data acquisition setup. Creating algorithms that can be tuned to new environments, or that are trained on larger datasets, would help ensure replicability of the results of computational analyses.

**Interpretability:** A common complaint about machine learning methods is that they are a black box. However, automated methods could actually become *more* transparent in their choices than human annotators, who can struggle to communicate their decision process during annotation. Efforts to make human-interpretable behavior classification tools may help scientists have greater confidence in their results [76].

**Long timescales:** Behavior has structure over multiple scales: from the coordination of actions into sequences [77], to effects of priming [78], habituation [79], satiety [80,81], and time of day [82,83]. Long timescales of behavior can be captured intrinsically by sequence learning [84] or hierarchical models [34], or may be recovered post-hoc by coarse-graining of action sequences [85]. But there is still ample space for exploration in this area.

**Level of granularity:** Neural recording datasets introduce another challenge for behavior analysis: not just finding the representation of behavior that best captures observed postural variance, but finding the right level of granularity to account for observed neural activity. The optimal behavior representation in this sense will likely be different for different brain regions, with some regions better explained by fine-grained, sub-second action motifs, and other regions more correlated with an animal’s overall behavioral objective (such as aggression or reproduction.) How best to integrate quantitative behavior analysis with neural imaging datasets is an exciting area of ongoing research.

**Interacting agents:** Multi-agent behavior poses unique challenges for analysis, as these datasets include times when animals are interacting and times when each individual is behaving independently. This affects unsupervised methods in particular, as the relative positioning of agents is sometimes critical for interpreting behavior and sometimes does not matter at all. We hypothesize that one way to overcome this challenge is to base behavior discovery methods on forecasting models, which could be trained to incorporate postural information from an agent’s partner only when that information is relevant to predicting what the agent does next.

are ready for downstream analysis. In addition to saving human effort, this approach can improve annotation quality: whereas human annotators show substantial within-individual and between-individual variability in labeling behavior start and stop times [17], automated algorithms are self-consistent. And second, computational behavior analyses can produce high temporal resolution, highly granular descriptions of animals’ actions, surpassing what is feasible for a human annotator [19–21].

The combined impact of automation and increased granularity is beginning to change the way behavioral experiments are designed and analyzed, with a growing number of labs opting to collect long-term, high-throughput recordings of animals in complex environments. Automating the frame-by-frame annotation of animal behavior has enabled high-throughput screening of hundreds of animals across strains, mutant lines, or experimental perturbations [15,23–30]. Differences between animal cohorts in these assays can most simply be detected by identifying a change in the proportion of time animals express one or more behaviors, or by a change in the transition probability between pairs of behaviors [24,25]. Difference in behavior can also be

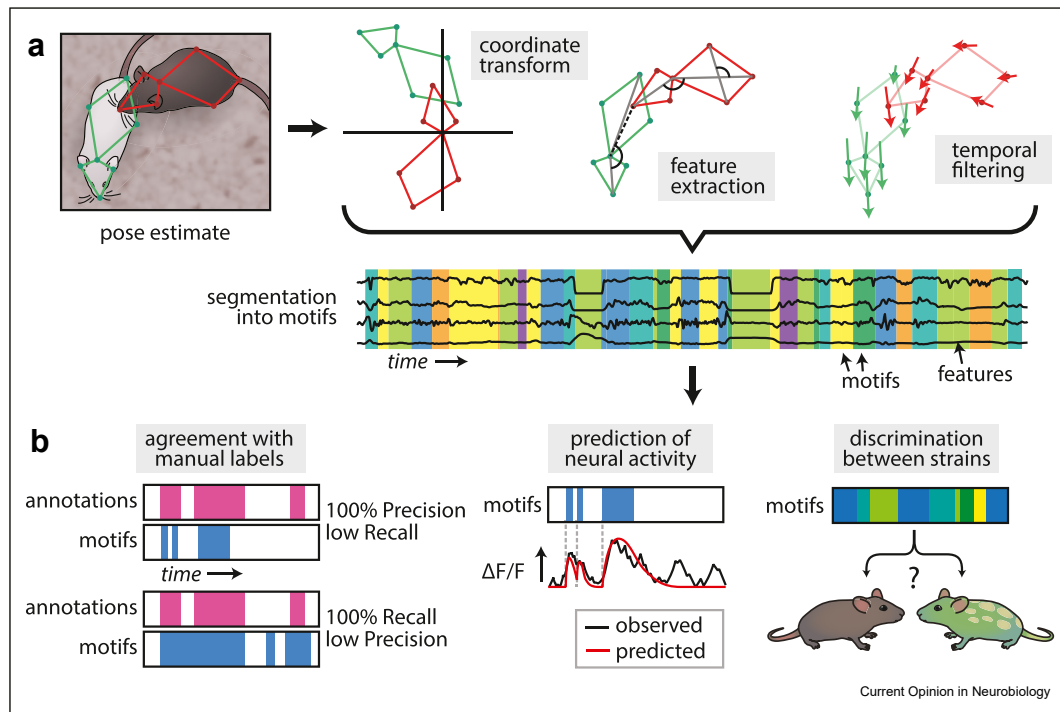
correlated with other within-cohort experimental measures, like gene expression [29].

Beyond the level of individual behaviors, differences between cohorts may be identified by training a classifier to distinguish animal groups given either postural features or a histogram of expressed behaviors [24,25,27]. The magnitude of differences between cohorts can also be quantified and visualized to infer relationships between them [25,27]. Interestingly, some studies have used dimensionality reduction methods (like principal component analysis or Fisher’s linear discriminant analysis) to identify different behavioral modes across individuals, mutant lines, or drug treatments [25,26,29]. These analyses suggest a low-dimensional “behavior manifold” might be learned that can capture the variability in behavior expressions across individuals.

### Generative model and forecasting: The Laplacian demon in the details

If we knew an animal’s entire sensory environment and history, could we predict what it would do next? How far into the future would our predictions hold? Generative models can help us better understand how animals

Figure 1



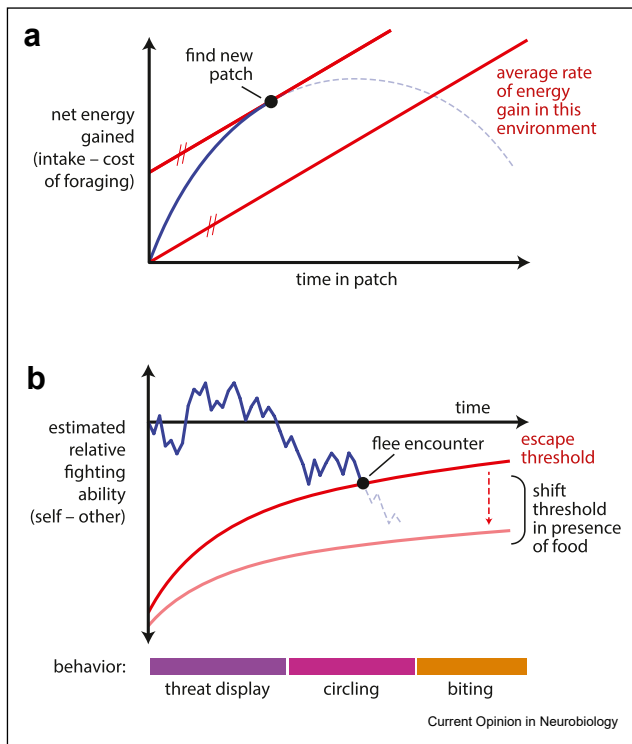
Unsupervised behavior analysis and its evaluation. **(a)** A typical unsupervised behavior analysis pipeline takes in continuous data (e.g. pose estimates) and segments it into motifs. Usually data is first pre-processed to remove sources of variance that are not meaningful for behavior (such as the absolute positions of the animals) and extract some representation of animal movement. These processed features are then segmented into motifs using some form of clustering algorithm. **(b)** Different design choices in an unsupervised analysis produce different motifs: there is therefore a pressing need for ways to evaluate the “usefulness” of unsupervised analysis methods. Three commonly used approaches are comparison of motifs to manual annotations, for example, in terms of precision and recall (left), ability of motifs to predict neural activity (center), and ability to distinguish between animal species, strains, or treatment groups based on the occurrence or patterning of motifs (right).

integrate sensory cues and their own recent history to determine their ongoing behavior. A generative model captures the relationship between two probability distributions, specifically the probability distribution of one signal conditioned on the value of another. In terms of behavior, this might translate to learning a distribution over possible behavioral responses given an animal’s past actions or its sensory environment.

Generative models can be used as a means of categorizing behavior: for example, autoregressive Hidden Markov models capture the statistics of time series evolution, but their inferred hidden states can also be used for unsupervised behavioral segmentation [20]. But generative models have also been used to uncover how animals’ sensory environment, state, and history shape expression of behavior [31–34]. For example, fly courtship song type, previously thought to vary at random, can be predicted by postural cues from the courted female [31], and furthermore this sensory-evoked behavior is better predicted when taking into account an estimate of the singing male’s internal state [33].

Animals often show preparatory movements or gradual escalation of interactions from which we can predict of future movements—for example, mice dart and tail-rattle before initiating an attack. As their name implies, generative models can also be used to generate “realistic-looking” behavioral data, predicting an animal’s actions given initial conditions [32,34]; in machine learning, this area of research is called imitation learning. Model-synthesized trajectories predict how an animal would respond given an environment, provided that environment is close to the conditions the model was trained on: for example, simulated flies follow walls and perform wing extensions when they encounter other simulated flies [34]. Trained models that can forecast behavior are often hierarchical, allowing them to capture structure in behavior on multiple timescales. One exciting promise of these models is that they provide continuous, behavior-related signals of different degrees of granularity, that can be contrasted with recorded neural dynamics to identify neural correlates of behavior or internal states [35].

Figure 2



Normative models make predictions about the temporal organization of behavior. (a) The marginal value theorem predicts that animals foraging in a patchy environment should migrate to a new patch when their net rate of energy accumulation within a patch (blue line) matches the mean rate of energy accumulation across the environment (parallel red lines). Figure adapted from Charnov, 1976. (b) In the resource holding potential model of aggressive escalation, animals integrate evidence to estimate the difference in fitness between fighters, and flee the encounter once their estimate passes a threshold. One possible formulation of this strategy is shown here, adapted from Enquist et al., 1990. The war of attrition model further hypothesizes that animals' motivation to obtain resources can change their willingness to flee, indicated here by the shift of the red line.

### Normative models: What's my motivation?

Why does behavior have the structure it does? A normative model is one that uses some measure of value or utility to derive predictions of how a system "ought to" behave if driven by that utility. At its highest level, we expect animal behavior to either directly or indirectly further the survival of that animal's genetic material. While this maxim is quite broad, principles of survival have informed many more specific theories of naturalistic behavior control. We will review examples from single-agent and dyadic behaviors, and discuss how computational behavior analysis might enrich our experimental investigation of these models. There is also a rich literature on principles of collective behavior [36,37], which will not be covered here.

### Behavior of an agent in an environment

Exploration of new environments is a fundamental behavioral drive, be it to establish territory, identify

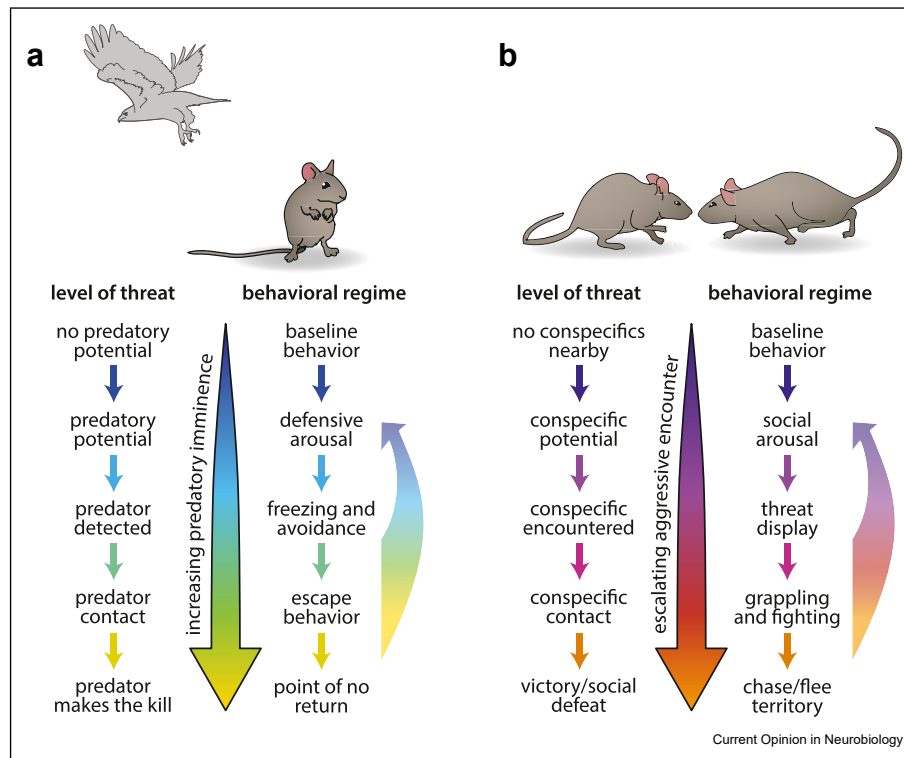
threats, or discover resources. Normative theories predict that exploring animals balance the reward of information gained with the cost of time passage [38] or risk of injury or predation. Foraging behavior, an extension of exploration, also encompasses an assortment of value-based decision-making tasks [39–41]. During foraging, an animal makes dietary choices as to what resources to pursue [42], and also must efficiently use time and energy when collecting resources. In a model environment of food patches, the marginal value theorem determines the precise time to migrate to new food sources, given full knowledge of the underlying average reward rate of the environment [43,44] (Figure 2a). At the opposite extreme, if nothing is known about the structure of the environment, random Lévy flight motion is optimal [45]. Factors such as food source distribution, prey movement, and competition with conspecifics also impact strategy. Foraging under risk of predation introduces the further wrinkle of balancing the energetic cost of monitoring the environment with the survival cost of being caught by a predator [46,47].

### Behavior in pairs of interacting agents

Dyadic behavior can include both symmetric interactions between members of a species and asymmetric interactions such as predator–prey behavior. Predatory imminence theory suggests a topology of defensive behavior, in which the type of behavior a prey species expresses is shaped by the immediacy and magnitude of predator threat [48] (Figure 3a). This escalating organization of defensive behaviors is built around a key conflict: that predators can be evaded either by reducing detectability (through freezing) or by overt escape actions that transiently increase detectability. Because escape behavior draws attention and interrupts ongoing foraging by the prey, an optimal strategy finds a balance between these two defensive behaviors.

Outside of predator-prey conflict, violence in nature occurs predominantly between members of the same species, and rarely results in serious injury. The survival value of intra-species aggression has been questioned since Darwin, who hypothesized that aggression ensures that the fittest members of a species obtain greater territories and resources. Subsequent theories further posed that aggression could help balance distribution of animals across available environment [49]. In social species, establishment of dominance hierarchies is thought to stabilize groups and create networks that shape the flow of information among group members [50,51]. An extensive literature in game theory has explored the structure of aggression, and gave rise to the concept of "evolutionarily stable strategies": behavioral strategies in a species that are stable to small perturbations in the form of mutant strategies [52].

Figure 3



Behavioral escalation in predator-prey and within-species aggression contexts. (a) Predator imminence theory describes an organization of defensive strategies in which level of predatory threat shifts the behavior animals express. Figure adapted from Fanselow & Lester, 1988. (b) Escalation of aggressive encounters between conspecifics shows a similar scaling of behavior, a strategy thought to minimize risk of injury. Escalation could serve to establish animals' relative resource holding potentials, or could be a "war of attrition" process of animals communicating their level of commitment to a contested resource.

Normative models of aggression often focus on the observation that, like defensive behavior, aggressive behaviors vary in their intensity (Figure 3b). Escalation of aggression could serve several functions. In the "resource holding potential" hypothesis, escalation allows animals to exchange information about relative fitness while minimizing risk of injury, by only escalating if the relative fitness of the pair remains unclear (Figure 2b) [53,54]. This hypothesis predicts that animals more closely matched in size escalate aggression further than more asymmetric pairs, a prediction that holds to varying extents across species [55,56]. Alternatively, the "war of attrition" model [57] poses that escalation signals the cost an animal is willing to pay for a resource, explaining how manipulations such as food deprivation might motivate an animal to escalate further or faster [58]. Escalation may also simply reflect the animal's relative levels of aggression and fear [59]. Models of aggressive escalation were extensively investigated in the 1970s, 80s, and 90s, in species ranging from dung beetles to hermit crabs to red deer; this work is excellently reviewed in Ref. [60].

### Going forward

Computational behavior analysis is experiencing a period of explosive growth, fueled and fueled by a push among neuroscientists to study more complex and naturalistic animal behaviors. Yet existing normative models of behavior largely precede the computer vision revolution in animal tracking. New behavior quantification methods will allow more rigorous testing of predictions made by classical normative models, whereas generative modeling provides methods for capturing behavioral control policies based on precise quantification of animals' sensory environment and behavior. An open challenge for modern behavior analysis is to re-evaluate our normative theories of behavior in the light of more plentiful data, to determine how useful they are as descriptors of animals' actions.

Towards this goal, the theoretical framework that has been established around exploration and foraging strategies provides rich ground for experimental exploration. Exploratory behavior and learning in neuroscience are often studied in conceptual tasks that do not require



physical exploration, permitting neural correlates to be measured [61,62]. An exciting promise of computational behavior analysis is that it allows exploration to be studied in complex physical environments such as mazes [63,64], where—in striking contrast to conceptual exploration—mice rapidly learn long sequences of actions, and can show one-shot learning of their home path out of the maze [64]. By tracking exploring animals' posture, researchers can detect deliberative actions such as “vicarious trial and error,” when animals pause to investigate options at a choice point [65,66]. Quantification of deliberative actions may help study another normative theory of exploration, the principle of effort minimization, which predicts that as an environment or scenario becomes familiar animals may swap between decision strategies, balancing low-effort habit and high-effort planning and simulation [67].

Analysis and generative modeling of animal actions and sensory environments during naturalistic foraging could also help determine how animals construct and update their internal model of food availability in their environment. For example, recent work has demonstrated that imitation learning can capture learned foraging behavior in a head-fixed task [62]. Elimination of the behavior annotation bottleneck could also allow for higher volume experiments testing the sensitivity of foraging behavior to combinations of environmental parameters such as degree of patch structure, presence of threat or conspecific cues, choices between food types, or changes in food availability.

Normative theories of dyadic interactions have often focused on the rules governing transitions between behaviors: either the switch from freezing to escape in predator imminence theory, or the gradual escalation of aggression in intra-species encounters. These progressions are challenging to study in the lab, as animals show rapid habituation and priming after a few presentations of a threatening stimulus or conspecific. With automated analyses, classical assays like the resident–intruder paradigm could be expanded to longer timescales in enriched environments, so that animals have time and space for more naturalistic encounters to unfold [68].

Finally, one exciting area for exploration is to develop theories and models for how competing motivational signals should interact to shape behavior. Animal behavior in the wild strikes a balance between competing drives: hunger and thirst, predator defense, and social (reproductive, parenting, and territorial) motivations. These drives can be thought of as low-dimensional intervening variables between sensation and action [69]: any single drive can be affected by multiple sensory cues—for example, both mating experience [70] and social isolation [71] make mice more aggressive, and both dry food and salt make mice thirsty. And drives can be linked to

any learned behavior: both aggression and thirst can be used to motivate nose-poking or lever-pressing in operant tasks [72,73].

We cannot measure an animal's aggressive motivation directly, but we can apply experimental manipulations that we know will alter aggression—and also hunger, thirst, stress, reproductive state, or time of day—and measure resulting changes in behavior. How might these different state manipulations interact? Before we can generate hypotheses of how the brain is controlling behavior, it is good to stop and ask: how complex is behavior in the first place? If the annotation bottleneck can be resolved, we might begin to use high-throughput behavioral studies to ask whether intervening variables of hunger, thirst, fear, or aggression are indeed the one-dimensional signals we call them, or whether behavioral control policies and their underlying neural drives are more complex.

## Conflict of interest statement

Nothing declared.

## Acknowledgements

I am grateful to Jennifer J. Sun, Yisong Yue, Pietro Perona, and David J. Anderson for helpful discussions on the organization of behavior and the application of machine learning to behavior analysis, and to Brady Weissbourd for discussions on theories of social behavior. This manuscript was supported by the National Institute of Mental Health of the National Institutes of Health under Award Number R00MH117264. The content is solely the responsibility of the author and does not necessarily represent the official views of the National Institutes of Health.

## References

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

- \* of special interest
  - \*\* of outstanding interest
1. [Pereira TD, Shaevitz JW, Murthy M: Quantifying behavior to understand the brain. \*Nat Neurosci\* 2020, \*\*23\*\*:1537–1549.](#)
  2. [Anderson DJ, Perona P: Toward a science of computational ethology. \*Neuron\* 2014, \*\*84\*\*:18–31.](#)
  3. [Brown AEX, De Bivort B: Ethology as a physical science. \*Nat Phys\* 2018.](#)
  4. [Datta SR, Anderson DJ, Branson K, Perona P, Leifer A: Computational neuroethology: a call to action. \*Neuron\* 2019, \*\*104\*\*:11–24.](#)
  5. [Mathis MW, Mathis A: Deep learning tools for the measurement of animal behavior in neuroscience. \*Curr Opin Neurobiol\* 2020, \*\*60\*\*:1–11.](#)
  6. [McCullough MH, Goodhill GJ: Unsupervised quantification of naturalistic animal behaviors for gaining insight into the brain. \*Curr Opin Neurobiol\* 2021, \*\*70\*\*:89–100.](#)
  7. [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.](#)
  8. [Calhoun AJ, Murthy M: Quantifying behavior to solve sensorimotor transformations: advances from worms and flies. \*Curr Opin Neurobiol\* 2017, \*\*46\*\*:90–98.](#)
  9. [Mathis A, et al.: DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. \*Nat Neurosci\* 2018, \*\*21\*\*:1281–1289.](#)

10. Pereira TD, *et al.*: **SLEAP: multi-animal pose tracking**. *bioRxiv* 2020, <https://doi.org/10.1101/2020.08.31.276246>. 08.31.276246 (2020).
11. Graving JM, *et al.*: **DeepPoseKit, a software toolkit for fast and robust animal pose estimation using deep learning**. *Elife* 2019, **8**.
12. Karashchuk P, *et al.*: **Anipose: a toolkit for robust markerless 3D pose estimation**. *Cell Rep* 2021, **36**:109730.
13. Machado AS, Darmohray DM, Fayad J, Marques HG, Carey MR: **A quantitative framework for whole-body coordination reveals specific deficits in freely walking ataxic mice**. *Elife* 2015, **4**.
14. Ahamed T, Costa AC, Stephens GJ: **Capturing the continuous complexity of behaviour in *Caenorhabditis elegans***. *Nat Phys* 2020, **17**:275–283.
15. Sheppard K, *et al.*: **Stride-level analysis of mouse open field behavior using deep-learning-based pose estimation**. *Cell Rep* 2022, **38**:110231.
16. Nilsson SRO, *et al.*: **Simple Behavioral Analysis (SimBA) – an open source toolkit for computer classification of complex social behaviors in experimental animals**. *bioRxiv* 2020, <https://doi.org/10.1101/2020.04.19.049452>. 04.19.049452 (2020).
17. Segalin C, *et al.*: **The Mouse Action Recognition System (MARS) software pipeline for automated analysis of social behaviors in mice**. *Elife* 2021, **10**, e64720.
18. Kabra M, Robie AA, Rivera-Alba M, Branson S, Branson K: **JAABA: interactive machine learning for automatic annotation of animal behavior**. *Nat Methods* 2013, **10**:64–67.
19. Berman GJ, Choi DM, Bialek W, Shaevitz JW: **Mapping the stereotyped behaviour of freely moving fruit flies**. *J R Soc Interface* 2014, **11**.
20. Wiltschko AB, *et al.*: **Mapping sub-second structure in mouse behavior**. *Neuron* 2015, **88**:1121–1135.
21. Luxem K, Fuhrmann F, Kürsch J, Remy S, Bauer P: **Identifying behavioral structure from deep variational embeddings of animal motion**. *BioRxiv* 2020.
22. Hsu AI, Yttri EA: **B-SOId, an open-source unsupervised algorithm for identification and fast prediction of behaviors**. *Nat Commun* 2021, **12**:5188.
23. Vogelstein JT, *et al.*: **Discovery of brainwide neural-behavioral maps via multiscale unsupervised structure learning**. *Science* 2014, **344**:386–392.
24. Branson K, Robie AA, Bender J, Perona P, Dickinson MH: **High-throughput ethomics in large groups of *Drosophila***. *Nat Methods* 2009, **6**:451–457.
25. Wiltschko AB, *et al.*: **Revealing the structure of pharmacobehavioral space through motion sequencing**. *Nat Neurosci* 2020, **23**:1433–1443.
- The authors use unsupervised behavioral analysis (MoSeq) to characterize behavior in a cohort of over 700 mice under a variety of pharmacological perturbations. This rich dataset would make an excellent benchmark for evaluating the descriptive power of unsupervised behavioral analyses. The authors also show an interesting and novel application of MoSeq to identify on-target vs off-target effects of drugs in modifying animal behavior.
26. Tang W, *et al.*: **Genetic control of collective behavior in *Zebrafish***. *iScience* 2020, **23**:100942.
27. Hernández DG, *et al.*: **A framework for studying behavioral evolution by reconstructing ancestral repertoires**. *Elife* 2021, **10**.
28. Klibaite U, *et al.*: **Deep phenotyping reveals movement phenotypes in mouse neurodevelopmental models**. *Mol Autism* 2022, **13**:12.
29. Forkosh O, *et al.*: **Identity domains capture individual differences from across the behavioral repertoire**. *Nat Neurosci* 2019, **22**:2023–2028.
- An interesting approach to capturing “personality” of animals, using linear discriminant analysis to identify features of behavior that are stable across time within individual animals, but distinguish those animals from others in their cohort.
30. Geuther BQ, *et al.*: **Action detection using a neural network elucidates the genetics of mouse grooming behavior**. *Elife* 2021, **10**.
31. Coen P, *et al.*: **Dynamic sensory cues shape song structure in *Drosophila***. *Nature* 2014, **507**:233–237.
32. Johnson RE, *et al.*: **Probabilistic models of larval *Zebrafish* behavior reveal structure on many scales**. *Curr Biol* 2020, **30**:70–82.e4.
33. Calhoun AJ, Pillow JW, Murthy M: **Unsupervised identification of the internal states that shape natural behavior**. *Nat Neurosci* 2019, **22**:2040–2049.
- Previous work from the Murthy lab showed that type of fly courtship song can be predicted by sensory cues from the posture of the fly and its potential mate. In this work, Calhoun *et al.* show that model predictions can be improved if combined with a Hidden Markov Model allowing the singing fly to be in different behavioral “states,” which are learned in an unsupervised manner.
34. Eyjolfsson E, Branson K, Yue Y, Perona P: **Learning recurrent representations for hierarchical behavior modeling**. *arXiv [cs.AI]* 2016.
- An exciting application of imitation learning to animal behavior; the authors train recurrent neural networks to predict future actions of interacting flies in a dish, given a simplified representation of the animals’ field of view and posture. The model learns multiple timescales and types of animal behavior.
35. Hattori R, Komiyama T: **Context-dependent persistency as a coding mechanism for robust and widely distributed value coding**. *Neuron* 2021, <https://doi.org/10.1016/j.neuron.2021.11.001>.
36. Sumpter DJT: *Collective animal behavior*. Princeton University Press; 2010.
37. Couzin ID, Krause: **J. Self-organization and collective behavior in vertebrates**. *Adv Stud Behav* 2003.
38. Thrun S: *Efficient exploration in reinforcement learning*. Carnegie Mellon University; 1992. Technical Report.
39. Stephens DW, Krebs: *J. R. Foraging theory*. Princeton University Press; 1987.
40. Calhoun AJ, Hayden BY: **The foraging brain**. *Curr Opin. Behav. Sci.* 2015, **5**:24–31.
41. Mobbs D, Trimmer PC, Blumstein DT, Dayan P: **Foraging for foundations in decision neuroscience: insights from ethology**. *Nat Rev Neurosci* 2018, **19**:419–427.
42. Rangel A: **Regulation of dietary choice by the decision-making circuitry**. *Nat Neurosci* 2013, **16**:1717–1724.
43. Charnov EL: **Optimal foraging, the marginal value theorem**. *Theor Popul Biol* 1976, **9**:129–136.
44. Oaten A: **Optimal foraging in patches: a case for stochasticity**. *Theor Popul Biol* 1977, **12**:263–285.
45. Viswanathan GM, *et al.*: **Optimizing the success of random searches**. *Nature* 1999, **401**:911–914.
46. Lloyd K, Dayan P: **Interrupting behaviour: minimizing decision costs via temporal commitment and low-level interrupts**. *PLoS Comput Biol* 2018, **14**, e1005916.
47. Lima SL, Dill LM: **Behavioral decisions made under the risk of predation: a review and prospectus**. *Can J Zool* 1990, **68**:619–640.
48. Fanselow MS, Lester LS: **A functional behavioristic approach to aversively motivated behavior: predatory imminence as a determinant of the topography of defensive behavior**. In Bolles RC. *Evolution and learning*, vol. 263. Lawrence Erlbaum Associates, Inc, xi; 1988:185–212.
49. Lorenz K: *On aggression*. Routledge; 1966.
50. Holekamp KE, Strauss ED: **Aggression and dominance: an interdisciplinary overview**. *Curr Opin. Behav. Sci.* 2016, **12**:44–51.
51. Grosenick L, Clement TS, Fernald RD: **Fish can infer social rank by observation alone**. *Nature* 2007, **445**:429–432.

52. Smith JM, Price GR: **The logic of animal conflict.** *Nature* 1973, **246**:15.
53. Parker GA: **Assessment strategy and the evolution of fighting behaviour.** *J Theor Biol* 1974, **47**:223–243.
54. Enquist M, Leimar O, Ljungberg T, Mallner Y, Segerdahl N: **A test of the sequential assessment game: fighting in the cichlid fish *Nannacara anomala*.** *Anim Behav* 1990, **40**:1–14.
55. Alward BA, Cathers PH, Blakkan DM, Hoadley AP, Fernald RD: **A behavioral logic underlying aggression in an African cichlid fish.** *Ethology* 2021, **127**:572–581.
56. Tedore C, Johnsen S: **Visual mutual assessment of size in male *Lyssomanes viridis* jumping spider contests.** *Behav Ecol* 2015, **26**:510–518.
57. Parker GA, Thompson EA: **Dung fly struggles: a test of the war of attrition.** *Behav Ecol Sociobiol* 1980, **7**:37–44.
58. Enquist M, Leimar O: **Evolution of fighting behaviour: the effect of variation in resource value.** *J Theor Biol* 1987, **127**:187–205.
59. Hinde RA: **Animal signals: ethological and games-theory approaches are not incompatible.** *Anim Behav* 1981, **29**:535–542.
60. Archer J, Huntingford F: **Game theory models and escalation of animal fights.** In *The dynamics of aggression*. Psychology Press; 2013:21–50.
- A fascinating review of game theory of aggression, including several examples of early efforts to test predictions made by normative models of escalation during aggression.
61. Daw ND, O'Doherty JP, Dayan P, Seymour B, Dolan RJ: **Cortical substrates for exploratory decisions in humans.** *Nature* 2006, **441**:876–879.
62. Hattori R, Danskin B, Babic Z, Mlynaryk N, Komiyama T: **Area-specificity and plasticity of history-dependent value coding during learning.** *Cell* 2019, **177**:1858–1872. e15.
63. Mugan U, MacIver MA: **Spatial planning with long visual range benefits escape from visual predators in complex naturalistic environments.** *Nat Commun* 2020, **11**:3057.
64. Rosenberg M, Zhang T, Perona P, Meister M: **Mice in a labyrinth exhibit rapid learning, sudden insight, and efficient exploration.** *Elife* 2021, **10**.
- While mice are notorious slow learners on many working memory and sensory discrimination tasks, Rosenberg et al. show that when given a more naturalistic task of learning the layout of a complex maze, animals can learn quickly and efficiently.
65. Redish AD: **Vicarious trial and error.** *Nat Rev Neurosci* 2016, **17**:147–159.
66. Zhang T, Rosenberg M, Perona P, Meister M: **Endotaxis: a universal algorithm for mapping, goal-learning, and navigation.** *bioRxiv* 2021, **09.24**:461751, <https://doi.org/10.1101/2021.09.24.461751>. 2021.
67. Keramati M, Smittenaar P, Dolan RJ, Dayan P: **Adaptive integration of habits into depth-limited planning defines a habitual-goal-directed spectrum.** *Proc Natl Acad Sci USA* 2016, **113**:12868–12873.
- Excellent work on the trade-off between habit-driven and goal-driven behavior. While this has been previously explored mostly in virtual tasks, it would be interesting to see a similar approach taken to more naturalistic tasks such as maze navigation.
68. Weissbrod A, et al.: **Automated long-term tracking and social behavioural phenotyping of animal colonies within a semi-natural environment.** *Nat Commun* 2013, **4**:2018.
69. Berridge KC: **Motivation concepts in behavioral neuroscience.** *Physiol Behav* 2004, **81**:179–209.
70. Remedios R, et al.: **Social behaviour shapes hypothalamic neural ensemble representations of conspecific sex.** *Nature* 2017, **550**:388–392.
71. Zelikowsky M, et al.: **The neuropeptide Tac2 controls a distributed brain state induced by chronic social isolation stress.** *Cell* 2018, **173**:1265–1279. e19.
72. Falkner AL, Grosenick L, Davidson TJ, Deisseroth K, Lin D: **Hypothalamic control of male aggression-seeking behavior.** *Nat Neurosci* 2016, **19**:596–604.
73. Golden SA, et al.: **Compulsive addiction-like aggressive behavior in mice.** *Biol Psychiatr* 2017, **82**:239–248.
74. Im DJ, Kwak I, Branson K: **Evaluation metrics for behaviour modeling.** *arXiv [cs.LG]* 2020.
75. Sun JJ, et al.: **The multi-agent behavior dataset: mouse dyadic social interactions.** *arXiv [cs.LG]* 2021.
76. Murdoch WJ, Singh C, Kumbier K, Abbasi-Asl R, Yu B: **Definitions, methods, and applications in interpretable machine learning.** *Proc Natl Acad Sci USA* 2019, **116**:22071–22080.
77. Corver A, Wilkerson N, Miller J, Gordus A: **Distinct movement patterns generate stages of spider web-building.** *bioRxiv* 2021, <https://doi.org/10.1101/2021.05.24.444987>. 05.24.444987 (2021).
78. Potegal M: **Time course of aggressive arousal in female hamsters and male rats.** *Behav Neural Biol* 1992, **58**:120–124.
79. Rankin CH, et al.: **Habituation revisited: an updated and revised description of the behavioral characteristics of habituation.** *Neurobiol Learn Mem* 2009, **92**:135–138.
80. Zhang SX, Rogulja D, Crickmore MA: **Dopaminergic circuitry underlying mating drive.** *Neuron* 2016, **91**:168–181.
81. Zimmerman CA, Knight ZA: **Layers of signals that regulate appetite.** *Curr Opin Neurobiol* 2020, **64**:79–88.
82. Nelson RJ, Bumgarner JR, Walker 2nd WH, DeVries AC: **Time-of-day as a critical biological variable.** *Neurosci Biobehav Rev* 2021, **127**:740–746.
83. Bumgarner JR, Walker 2nd WH, Nelson RJ: **Circadian rhythms and pain.** *Neurosci Biobehav Rev* 2021, **129**:296–306.
84. Whiteway, M. R. et al. **Semi-supervised sequence modeling for improved behavioral segmentation.** doi:10.1101/2021.06.16.448685.bioRxiv
85. Alba V, Berman GJ, Bialek W, Shaevitz JW: **Exploring a strongly non-Markovian animal behavior.** *arXiv [q-bio.NC]* 2020.