



Published in final edited form as:

Neuron. 2021 July 21; 109(14): 2224–2238. doi:10.1016/j.neuron.2021.05.021.

Promises and Challenges of Human Computational Ethology

Dean Mobbs^{1,2,*}, Toby Wise^{1,3,4,*}, Nanthia Suthana^{5,6}, Noah Guzmán², Nikolaus Kriegeskorte^{7,8}, Joel Z. Leibo⁹

¹Department of Humanities and Social Sciences and Computation, 1200 E California Blvd, HSS 228–77, Pasadena, California 91125, USA

²Neural Systems Program at the California Institute of Technology, 1200 E California Blvd, HSS 228–77, Pasadena, California 91125, USA

³Wellcome Centre for Human Neuroimaging, University College London, London, United Kingdom

⁴Max Planck UCL Centre for Computational Psychiatry and Ageing Research, University College London, London, United Kingdom

⁵Department of Psychiatry and Biobehavioral Sciences, Jane and Terry Semel Institute for Neuroscience and Human Behavior, University of California Los Angeles, Los Angeles, USA

⁶Departments of Neurosurgery, Psychology, and Bioengineering, University of California Los Angeles, Los Angeles, USA

⁷Department of Psychology, Columbia University, New York, USA

⁸Department of Neuroscience, Zuckerman Mind Brain Behavior Institute, Columbia University, New York, USA

⁹Google DeepMind, London, United Kingdom

SUMMARY

The movements an organism makes provide insights into its internal states and motives. This principle is the foundation of the new field of computational ethology, which links rich automatic measurements of natural behaviors to motivational states and neural activity. Computational ethology has proven transformative for animal behavioral neuroscience. This success raises the question whether rich automatic measurements of behavior can similarly drive progress in human neuroscience and psychology? New technologies for capturing and analyzing complex behaviors in real and virtual environments enable us to probe the human brain during naturalistic dynamic interactions with the environment that were beyond experimental investigation so far. Inspired by nonhuman computational ethology, we explore how these new tools can be used to test important questions in human neuroscience. We argue that the application of this methodology will help human neuroscience and psychology extend on limited behavioral measurements such as

*Equal contribution

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

reaction time and accuracy, permit novel insights into how the human brain produces behavior, and ultimately reduce the growing measurement gap between human and animal neuroscience.

In Brief.

Computational ethology has revolutionized comparative neuroscience through automated measurement of naturalistic behavior. Mobbs, Wise et al., formalize how such approaches can benefit human neuroscience, providing richer behavioral assays and narrowing the measurement gap between human and animal studies.

INTRODUCTION

In the natural world, animals modify their behavior in response to changes in their environment, such as predation and competition, as well as changes in their internal metabolic drives (e.g., hunger and thirst; LeDoux, 2012; Mobbs et al., 2018). These observable behaviors can range from the deliberately controlled to impulsive or reactive, yet are consistent in that they provide information about the animal's latent motivational state and reflect strategic responses to a variety of survival demands. Measuring these unconstrained and seemingly random behaviors in extensive datasets has been a major challenge for behavioral neuroscientists. However, in non-human studies, machine learning methods that automate the registration and analysis of locomotor activity on the basis of video have given us much richer measurements of behavior. By better characterizing the behavioral motifs (See Glossary) of animals, a deeper understanding of the neural circuits involved in a rich variety of survival behaviors can be achieved (Anderson and Perona, 2014; Datta et al., 2019). This computational enhancement of behavioral observation has created the new field of computational ethology. Its methods appear equally pertinent to humans as to animals. However, these methods have yet to develop their full impact on human psychology and human neuroscience.

Here, we explore the promises and challenges of applying the methods of computational ethology to human neuroscience. We focus on the fields of human decision, social, and affective neuroscience and discuss ways in which experimental paradigms can be designed to evoke a wide range of natural defensive, appetitive and social behaviors. We detail both current and potential future methods, focusing on increased use of virtual ecologies to probe behavioral motifs and their underlying computations. This potential shift in approach to human neuroscience parallels recent calls from the field of comparative neuroscience where there is a need for more effective probing of behavior (Anderson and Perona, 2014; Datta et al., 2019). In turn, this will provide better models of how the brain produces behaviors (Babayan and Konen, 2019; Niv, 2020; Krakauer et al., 2017). As we argue throughout this article, if our goal is ultimately to explain real-world behavior, we will need to study naturalistic behavior and its neural underpinnings. Such an approach will 1) allow the detection of naturalistic behavioral patterns that are hidden in traditional, constrained experimental paradigms, 2) characterize neural systems supporting spontaneous, naturalistic behavior, and 3) determine how cognitive processes (e.g., decision-making) unfold in complex, naturalistic scenarios. So far, computational ethology has led to a range of novel computationally identified behaviors, and their associated neural basis, being

catalogued in rodents and *Drosophila*. These include behaviors that reflect appetitive, social and defensive behaviors (see Glossary; Anderson and Perona 2014), which up until now were frequently passed off as noise, being too fast and too stochastic to measure. But neither these challenges, nor their emerging solutions, are unique to animal research. We contend that computational ethology methods will also prove critical for the progress of human neuroscience.

NON-HUMAN COMPUTATIONAL ETHOLOGY

Description and analysis of animal behavior has traditionally relied on human observation and recording. In recent years however, the development of modern recording and analysis methods has facilitated the emergence of computational ethology (Anderson and Perona, 2014, Datta et al., 2019), which uses machine learning methods to automatically identify and quantify behavior, obviating the need for human observers. In lab settings, these approaches typically take data acquired from video cameras positioned around one or more animals and output a continuous representation of the animal's location or pose, for example recording limb or head position. This circumvents the subjectivity of human observations and promises observations higher in precision and quality, being unaffected by human visual and attentional capacity. Perhaps most importantly, computational ethology provides a dramatic increase in throughput, the benefits of which have been felt most strongly in fields that depend on high-frequency observations from large numbers of animals, such as those examining the roles of specific neural circuits in *Drosophila* (Hoopfer et al., 2015).

Although a new field, computational ethology has already demonstrated its utility across a range of studies. In *drosophila*, these methods have allowed the identification of neural circuits underlying distinct sensorimotor states (Calhoun et al., 2019), while combining computational ethology with optogenetics has enabled causal links between neural circuits and specific behaviors to be tested (Jovanic et al., 2016). Recently, the combination of automated classification of behavior with high-throughput neural recordings from rodents has revealed distributed patterns of neural activity previously regarded as noise to be associated with specific behavioral patterns (Musall et al., 2019; Stringer et al., 2019), a discovery with obvious relevance to “noisy” human neuroimaging.

Computational ethology has been a major beneficiary of developments in machine vision, which allows streams of video data to be automatically mined for behaviors of interest. This is not a trivial task: First, it is essential to continuously detect and monitor unique animals without confusing separate individuals. Second, animal features that constitute specific behaviors must be accurately extracted and tracked. Although simply tracking animals' location and direction of movement will be sufficient to answer many questions, a behavioral phenotype often depends on more complex features of the animal's actions. These may be subtle, for example based on pose or specific limb movements. Finally, classification of behaviors based on these features must be accurate. This process is approached either in a supervised way (Graving et al., 2019; Mathis et al., 2018; Pereira et al., 2019), where an experimenter manually identifies the features of the animal that should be tracked, or an unsupervised way (Berman et al., 2014; Wiltischko et al., 2015), where features of interest are identified without human intervention.

The number of tools developed for this purpose has increased dramatically over the past decade, boosted by the growth of deep learning. Experimental setups used in computational ethology represent the kind of nonlinear many-to-many classification problem (where many features must be mapped to many categories) that neural networks excel at, and their use has enabled the automatic identification and classification of behavioral features previously limited to human observation. These methods have evolved from algorithms designed for the detection of human poses (Insafutdinov et al., 2016), and have been made possible through transfer learning (Donahue et al., 2013), which takes advantage of pre-trained neural networks to facilitate performance without the need for extensive training data from the task at hand. Notable examples here are those that have been able to identify animal body and limb positions (Graving et al., 2019; Mathis et al., 2018), enabling automatic evaluation of movements and pose and classification of behaviors based on these. The result of this is that continuous video recordings of naturalistic animal behavior can be automatically processed, producing detailed ethograms representing how behavior unfolds over time.

In both human and non-human research, we face the challenge of understanding how complex high-dimensional behaviors are generated by neural systems. Computational ethology lends itself to this problem naturally. First, computational methods allow fine-grained assessment of behavior that accounts for its temporal dynamics. Importantly, this permits the dynamics of behavior to be linked to unfolding neural activity, potentially elucidating time-dependent neural processes underlying behavior. Second, computational ethology allows detailed quantification of free, naturalistic movements, making it possible to link neural processes to behavior without the need for highly controlled, unrealistic tasks.

HUMAN COMPUTATIONAL ETHOLOGY

It has recently been argued (e.g., Babayan and Konen, 2019; Balleine, 2019; Niv 2020) that behavior is essential to understanding the animal, including the human, brain. The continuous measurement of behavior that characterizes computational ethology may add to existing measures of behavior, including categorical decisions that are common to most experiments in human cognitive neuroscience (Fig. 1A–J). It is not, however, obvious how to best integrate computational ethological approaches into traditional paradigms in human neuroscience. Consider, for example, the study of human fear and anxiety. Traditional approaches include fear conditioning and presentation of visually aversive stimuli (e.g., fearful faces). These studies provide no clear path to the approaches advocated by computational ethology as these paradigms minimize behavioral dynamics (e.g., binary button presses), in contrast with the rich behavioral outputs that are central to animal computational ethology. However, new experimental paradigms have begun to engage subjects in dynamic interactions, typically in virtual environments. One of the first studies to move beyond classic fear conditioning paradigms used virtual predators to create an Active Escape Task. Although simple, the task involved subjects actively escaping from an attacking virtual predator in a 2D maze which allowed the subjects to visually keep track of the distance to the predator, and providing richer behavior than is typical of these tasks (Mobbs et al., 2007). Distance, therefore, could be parametrically coupled with BOLD-signal measurements taken from fMRI, permitting the identification of neural circuits involved in processing proximal and distal threats. This approach was later used and

extended by showing that panic-related motor errors correlated with brain areas commonly implicated in human and animal models of panic (Fig. 1B; Glossary; wrong button presses resulting in collisions with the virtual walls of the maze; Mobbs et al. 2009). Despite the promise of these early studies, more causal research is needed and computational ethology may be one direction that can address the shortfalls of previous empirical work.

These experiments were followed by several studies using similar paradigms (Bach et al., 2014; Gold et al., 2015; Meyer et al., 2019). For example, Bach and colleagues examined the role of human hippocampus in arbitrating approach-avoidance conflict under different levels of potential threat (Bach et al., 2014). Subjects were instructed to move a green triangle around a 2D gridded environment to collect tokens (exchanged for money) where one of three differently dangerous predators was located in a corner of the grid. At any time, the predator can begin to chase the subject, yet the subject can also choose to hide in the safe place (black box in the corner of the grid). Once caught, the subject loses all their tokens and the epoch is over. This task was able to measure several variables including time spent in the safe place, time spent close to the walls, and distance from the threat (Bach et al. 2014; Figure 1E). In another study, Gold et al., created a task where subjects were asked to capture prey and evade predators in a 2D maze similar to previous studies (Bach et al., 2014; Mobbs et al., 2007, 2009). The main result showed that when a threat was unpredictable, there was increased connectivity between the amygdala and vmPFC (Gold et al., 2014).

Although these studies did not take full advantage of the movement trajectories of the virtual environment, or apply unsupervised machine learning, they do provide a template for how to apply computational ethology to human neuroscience. They also highlight the advantages of using less restricted behavioral measures than is common in human neuroscience to reveal behavioral and neural patterns that would not otherwise be observable. Creation of such virtual ecologies enables experimentalists to measure less constrained types of behavior (Fig. 1).

EXAMPLES OF NOVEL BEHAVIORAL ASSAYS IN HUMANS

The standard behavioral measures used in laboratory tasks are decision accuracy and reaction time (RT). Reaction time (RT) is often used as a marker of decision confidence, deliberation and learning. However, while not undermining the importance of RT, it can be problematic as latencies in RT can arise due to many factors, including those of little interest to the experimenter such as tiredness and distraction. The additional measures of behavioral motifs and sequences help minimize these potentially confounding effects. In many tasks, decision accuracy is also used as a measure of learning (see special issue on behavior in *Neuron* Oct 2019), which encourages decision making paradigms to depend on infrequent decisions between a limited set of options in order to provide clearly delineated opportunities for learning to occur. Dynamic measurement of complex behaviors replaces the RTs and accuracies of a discrete sequence of actions by an essentially continuous stream of motor control signals and performance measures. This in turn, provides an excellent way to examine between-subject differences in behavior. Moving beyond the standard protocols, paradigms have been developed based on simple 2D environments or virtual ecologies that can capture multiple measures of threat anticipation, escape and conflict (Fig. 1 A).

Virtual ecologies can move one step closer to the real world by including levels of threat imminence, visually clear vs. opaque environments (e.g., forest vs open field), and changes in competition density (Mobbs et al., 2013; Silston et al., 2020). This approach provides the experimentalist not only with tools to question how the environment changes decision processes, but also how it affects locomotor activity. Indeed, in the real world, behaviors can be fluid, stilted, fleeting and subtle. Below, we give examples of several types of behavior that can be measured using virtual ecologies.

Anticipation of Danger

Anticipation of danger is a critical part of anxiety, which is typically defined as a future-oriented emotional state associated with ‘potential’ and ‘uncertain’ threats (Grupe and Nitschke, 2013). One classic measure observed by ethologists and behavioral neuroscientists is thigmotaxis (Glossary), which is an index of anxiety typically associated with the animal moving to the peripheral area of an open field. Thigmotaxis is observed in rodents, fish and humans (Walz et al., 2016). Other anxiety like behaviors include intermittent locomotion (i.e., movement pauses) when a threat is anticipated. Place aversion, which is a form of Pavlovian conditioning, is also possible in virtual ecologies, where avoidance of particular areas of the environment demonstrates aversion (e.g., a predator was encountered in the location). Another approach taken from the field of behavioral ecology is the concept of margin of safety, or when prey adopt choices that prevent deadly outcomes from occurring by keeping close proximity to a safety refuge and increasing the success of escape (Cooper and Blumstein, 2015; Qi et al., 2020). Finally, potential threats lead to vigilance behaviors, including orienting toward and attending to threat, both in non-human species and in humans (Mobbs and Kim, 2015; Wise et al., 2019).

We do not wish to argue that the results of traditional, non-naturalistic tasks are incorrect or entirely invalid. Their ability to break down complex behaviors into their constituent parts has provided substantial insight into the basic processes governing behavior. However, we do argue that reductionist approach limits their ability to explain naturalistic behavior as a whole. Taking fear conditioning as an example, we now have a rich understanding of the specifics of how humans acquire and lose fears of stimuli linked to aversive outcomes in the lab. However, typical fear conditioning experiments involve repeated, contiguous pairing of unconditioned and conditioned stimuli to engender learning. In the real world, such clean learning experiences are the exception rather than the norm. For example, a student may become fearful of exams after learning receiving a poor grade on a single exam taken weeks previously, without needing repeated experiences of taking exams and receiving immediate negative feedback. In terms of behavioral output, traditional tasks fail to capture the complexity of human behavior. Fear conditioning studies may require subjects to provide an expectancy rating, or in some tasks may require a binary stimulus selection. In the real world, our behaviors in response to feared stimuli are far more varied and complex, but traditional fear-conditioning paradigms tell us little about how acquired fears influence these behaviors. This gap between lab-based fear conditioning and real-world fear has been previously described as a barrier to successful treatment of pathological fears (Scheveneels et al., 2016).

Escape Behaviors

Escape is associated with fear and is elicited during predatory attack (Mobbs et al., 2020). Escape differs from avoidance in that escape is driven by the moment-to-moment adjusted movements of the attacking predator. Behaviorally, escape is associated with ballistic movements and increased vigor and is less coordinated than avoidance. This will also result in protean escape which is driven by the trajectory of the predator's attack and often results in unpredictable flight (e.g., zigzagging, spinning) (Humphries and Driver, 1970). The first human studies of escape and its neural correlates used a virtual predator that chase subjects in a 2D maze and examined the shift of activity in the brain as the threat came closer or farther away (Mobbs et al., 2007, 2009). More recent work has used flight initiation distance (FID), or the distance at which the subject flees from the approaching threat. FID is both a spatiotemporal measure of threat sensitivity and economic decision-making (Ydenberg and Dill, 1986). In a recent study, Qi and colleagues measured subject's volitional fleeing distance when they encounter a virtual predator. This study was the first to examine escape decisions in humans and importantly, showed that different part of the defensive circuits were engaged for fast and slow attacking threats (Fung et al., 2019; Qi et al., 2018).

Appetitive Behaviors. Pursuit and Hunting

Several experiments have used virtual ecologies to measure reward activity. These include foraging for rewards (Bach et al. 2014; Gold et al., 2014), and chasing prey for reward (Mobbs et al., 2009). In nature, appetitive behaviors take several forms including approach, increased vigor, exploration, stealth and sit and wait behaviors associated with surprise attack. Other examples include movement strategies when pursuing a virtual prey including angle of attack (e.g. cutting off corners to reduce escape time and feinting; Fig. 1), place preference and impulsive errors such as over-shooting the prey's anticipated movements (Mobbs et al., 2009). Using similar 2D environments to human studies (Mobbs et al., 2007; Bach et al., 2014) a recent study in non-human primates showed how the dorsal anterior cingulate cortex is involved in pursuit predictions including velocity, prey position and acceleration (Yoo et al., 2020). Finally, some studies have taken advantage of VR to explore human place preference in 3D environments. This has been demonstrated across both primary (Astur et al., 2014) and secondary (Astur et al., 2016; Molet et al., 2013) reinforcers, allowing simple conditioning paradigms to be extended to more realistic environments.

Dynamic Switching, Arbitration and Conflict between circuits

Virtual ecologies allow for dynamic switches in behavior. This of course can be measured in conventional task designs (e.g., task switching), however, in virtual ecologies, the switches can be either reactive, volitional or ramped up providing a unique way to study the human brain. One example is the Active Escape Task where results showed that when the artificial predator is distant, increased activity is observed in the ventromedial prefrontal cortex (vmPFC). However, as the artificial predator moves closer, a switch to enhanced activation in the midbrain periaqueductal gray (PAG) is observed (Mobbs et al., 2007). Arbitration between approach and avoidance has also been measured using more dynamic paradigms. Bach and colleagues aimed to support well established animal models of how approach-

avoidance conflict drive anxiety by showing that subjects exhibited passive avoidance behavior to threats when foraging for money in a 2D maze. FMRI results implicated the ventral hippocampus in this passive avoidance behavior, with lesions resulting in reduced avoidance (Bach et al., 2014). This was later extended on by showing that amygdala lesion patients (i.e., two Urbach-Wiethe syndrome patients) and healthy subjects administered Lorazepam showed reduced avoidance of threat (Korn et al., 2017).

Social Behaviors

Human social interaction features rich temporal and spatial dynamics. In the case of cooperative behaviors, most prior experimental and computational research has treated the choice to cooperate or defect as atomic. For instance, Camerer (2003) established a canon of rigorously controlled experimental paradigms where participants make atomic decisions such as whether to cooperate or defect in a Prisoner's Dilemma game. However, studies that abstract over the substructure of group behavior obscure its multi-scale dynamics. More recently, there has been a trend toward more complex paradigms using computer game-like virtual ecologies (Janssen, 2010; Mobbs et al., 2013). In this setting, self-interested individuals cooperate or defect through emergent policies that sequence lower-level actions. That is, participants must sequence primitive actions like move forward, turn left, etc. to implement their higher-level strategic decisions, say to cooperate or defect. Thus, the fine structure of an ecologically valid collective action problem is determined by coupled interactions between natural properties of the environment and the actions of other group members.

AUTOMATED CLASSIFICATION OF BEHAVIOR

Non-human work in computational ethology has availed itself of advances in deep learning to extract ethograms from high-dimensional behavioral data. In humans, similarly, complex behaviors can be captured with wearable behavioral sensors and video (Carreira and Zisserman, 2017; Topalovic et al., 2020) combined with machine learning, and this technology looks set to transform this field. More dynamic virtual ecologies provide a middle ground, with joystick or mouse input providing relatively rich behavioral data. Such data would lend themselves to analyses with machine learning methods, which can discover useful latent variables that more concisely capture consistent structure in dynamic behaviors. For example, detailed measures of a subject's position and velocity relative to rewarding stimuli could be used to identify particular reward-guided behaviors (Fig. 1). Moving beyond purely visual worlds, virtual environments enabling full bodily movement could allow the use of similar video-based techniques to those used in animal work. The use of unsupervised methods could also be particularly interesting when applied to human data, allowing the identification of behavioral patterns that are not easily detectable by human observers. Furthermore, methods linking behavioral data to other variables of interest (for example physiological measures or subjective state) could facilitate the identification of particular behaviors that have relevance to broader constructs such as anxiety.

A major disadvantage with the collection of large datasets is that much of the data is meaningless. However, when used in combination with dimensionality reduction methods

(for example clustering approaches, as discussed elsewhere in this article), signal can be separated from the noise to an extent. The approach we advocate provides two advantages over previous behavioral measures. First, the measurement of targeted behaviors such as thigmotaxis or pauses, and second, the ability to discover new behaviors that might be indicative of a decision or emotional state. The former uses naturalistic environments and rich data to confirm theory-driven predictions, while the latter uses a data-driven approach that can be used to inform new theory.

However, while automated dimensionality reduction can help greatly in the face of high-dimensional, unconstrained data, this does not eliminate the need for theory entirely. First, the choice of dimensionality reduction technique will be guided by theoretically motivated questions (e.g., what is the dimensionality of the data? Are we seeking to cluster brief behavioral motifs, or trajectories through an environment?). Second, it will be necessary to validate extracted behavioral patterns based on existing theory (e.g., do these newly-detected behaviors have meaningful neural correlates?). Finally, theory can be used to constrain the inferences that can be drawn from new observations, and help identify areas where new theory is needed (e.g., are automatically-identified threat-related behavioral patterns in line with theories about avoidance behavior?).

There are three examples that we believe illustrate the value of human computational ethology. First, Rosenberg et al (2021) show that more naturalistic environments can produce surprising insights into behavior even without introducing complex, multidimensional behavioral measures. In this study, mice were allowed to freely roam through a complex maze in search of reward, and the authors found that learning about the location of rewards was approximately 1000 times faster than in a standard two alternative forced choice task as typically used to study learning and decision-making. This clearly shows that behavior in standard, constrained task is not necessarily reflective of real-world behavior, which is ultimately what we are trying to explain. Second, Calhoun et al. (2019) used a data-driven modelling approach to identify three discrete behavioral states in *Drosophila* during courtship and were able to then identify neural systems supporting behavioral state switching. These behavioral states were not *a priori* hypothesized and would not have been visible in more constrained conditions. This demonstrates that using naturalistic behavior can identify behaviors that would simply not be identified otherwise. Finally, Stringer et al. (2019) combined data-driven parsing of high-dimensional natural behaviors and neural activity in mice to demonstrate that a large proportion of neural activity across the cortex is linked to behavioral patterns. This shows that neural signals that may otherwise be considered noise can in fact be clearly linked to behavior, a finding that emerges by virtue of a data-driven approach using multidimensional, naturalistic behavior.

We believe these studies show that 1) naturalistic behavior can be qualitatively different to that seen in constrained tasks, 2) data-driven analysis of naturalistic behavior can identify novel behavioral states, and 3) combining measures of behavior and neural activity in naturalistic environments can provide insights into the role of neural systems that would not otherwise be seen. Drawing parallels with human avoidance, 1) it is possible that naturalistic avoidance decisions are qualitatively different from those seen in constrained avoidance tasks, 2) escape may involve key behavioral states that would only be visible through

data-driven analysis of high-dimensional behavioral data, and 3) variability in neural activity may be linked to these behavioral states during escape.

HYBRID APPROACHES – PREPROGRAMMED AND AUTOMATIC CLASSIFICATION

Validating Machine Learning Methods

Despite the upsides of automated methods, even in non-human computational ethology, some analyses such as syllable identification and segmentation in bird song remain difficult to fully automate, and often hand coding and simple preprogrammed approaches remain the gold standard (Mets and Brainard, 2018). The results of automated machine learning models can be validated using traditional methods, but automated methods such as unsupervised learning can also help investigators identify features of behavior they may have missed using traditional methods. In cases such as these, a fruitful approach is to open a dialogue between traditional and automated methods, termed “human-in-the-loop” or “interactive” machine learning (Holzinger et al., 2019). In this framework, input from a human user is utilized to select or provide feedback to aspects of a model or learning algorithm, resulting in performance that adheres better to domain-specific expertise. The results produced by these models can give the human user clues regarding features they were not aware of, expanding their domain-specific expertise which can then be fed back into the model or algorithm. Notably, this approach has been advocated as a training method for human grandmasters in games such as Chess and Go (Kasparov, 2018).

Using Deep Neural Networks to Develop Better Virtual Ecologies

A limitation of the current virtual ecologies is that they consist of preprogrammed environments whose realism is questionable. For example, some virtual ecologies require subjects to interact with virtual agents whose behavior does not necessarily reflect known strategies used by real agents (Yoo et al., 2020). Hybrid environments are possible where real agents interact with each other, but the environments in which they interact are impoverished with respect to the information they would encounter and use to guide behavior in more naturalistic scenarios (Tsutsui et al., 2019).

One issue with more naturalistic task environments is the lack of control over properties of the stimuli and information in these environments. Indeed, in naturalistic settings, it is not always clear what the relevant properties are in terms of guiding behavior (Hamilton and Huth, 2020; Patterson, 1974; Sonkusare et al., 2019) since naturalistic stimuli are nonparametric and complex (Geisler, 2008). By pairing insights from ecological psychology with the relatively new tools provided by deep neural networks, investigators can discover the properties of naturalistic stimuli relevant for behavior as well as parameterize them, enabling the construction of more realistic virtual ecologies whose properties can be precisely controlled. Recent work on automated, unsupervised environment design for reinforcement learning agents may allow virtual ecologies to be defined without experimenter input (Dennis et al., 2020).

Feature learning methods can help in identifying such variables by extracting higher-order features of environments that reliably differ between task conditions using artificial neural networks (ANNs; Dosovitskiy et al., 2015). Recent work has argued that ANNs, rather than explicitly representing features in their environment, implicitly learn the structure of their environment that corresponds to task-appropriate actions (Hasson et al., 2020); this work lends further credibility to the use of ANNs for identifying higher-order latent variables in naturalistic stimuli. Deep generative models such as generative adversarial networks (GANs; Goodfellow et al., 2014) and variational autoencoders (VAEs; Doersch, 2016) can be used to construct generative models of naturalistic stimuli such as images, audio, videos, and even task-specific video game environments (Li et al., 2019; Yan et al., 2016) which can then be sampled from. These tools are already being used for the generation of naturalistic audio stimuli in the animal vocalization community (Sainburg et al., 2019). New methods for feature specific guidance of the output of these deep generative models can provide investigators with precise control over the statistics and dynamics of the higher-order latent variables relevant for behavior (Brookes et al., 2020; Lee and Seok, 2019).

LINKING BEHAVIORAL ETHOGRAMS TO NEURAL CIRCUITS

Computational ethology provides tools to generate ethograms (representations of different types of behavior over time; Fig. 2) with great accuracy and ease. Ethograms provide a detailed representation of the frequency of different behaviors over time; for example, in an avoidance task we may wish to identify patterns of danger anticipation and escape behavior in response to environmental threat. With computational methods, we can automatically generate an ethogram describing how different behaviors emerge over the course of the task. The temporal information embedded within ethograms makes them ideal candidates for linking to unfolding neural events. This presents a new challenge, however: how can we optimally map these detailed behavioral observations on to high-dimensional data provided by neuroimaging?

Multivariate decoders (Haxby et al., 2001; Kriegeskorte and Douglas, 2019) have great potential if we wish to identify distributed patterns of brain activity or connectivity associated with specific behaviors. Given behavioral labels taken from ethograms, multivariate classifiers may be trained to identify neural patterns associated with distinct behaviors. This would then permit the identification of distinct, distributed patterns of activity associated with specific behavioral patterns emerging from naturalistic behavior, in a similar way to previous work relating activation measured using fMRI to naturalistic movie stimuli (Spiers and Maguire, 2007) or continuous speech (Willems et al., 2016). Alternatively, for more hypothesis-driven work, joint brain-behavior modelling approaches may be effective (Turner et al., 2019). These methods rely on a single pre-specified model that accounts for both behavioral and neural data, accounting for their covariance through a hierarchical parameter structure, with shared parameters at the top level constraining the two modalities.

Encoding models (Kay et al., 2008) predict each channel of measured brain activity from external variables. They, too, could provide an effective method for identifying how brain activity corresponds to complex behavior. Alternatively, given a rich characterization of

behavioral patterns, decoders could be trained on behavior to predict associated neural states, in line with work in non-human animals (Clemens et al., 2015).

A complementary multivariate method that could link ethograms to neural data is representational similarity analysis (RSA; Kriegeskorte et al., 2008). RSA characterizes the representation in each brain region by means of a representational dissimilarity matrix that reveals how dissimilar the activity patterns are for each pair of experimental conditions. RSA could help establish relationships between complex behavioral descriptions and high-dimensional response patterns with minimal need for fitting of parameters that define the relationship between each channel of measured brain activity and each behavior as needed when using encoding and decoding models (Kriegeskorte and Douglas, 2019).

A promising avenue toward reducing the number of states to be considered is clustering of behavioral and neural data. Tools from computational ethology already constitute a form of clustering, typically identifying combinations of actions that constitute specific behaviors, resulting in detailed ethograms. In a similar vein, clustering methods have been applied to fMRI (Allen et al., 2014; and MEG data (Baker et al., 2014) to identify recurrent patterns of functional connectivity or activity, demonstrating that the brain cycles through a limited number of neural states. Exploring relationships between recurrent behavioral and neural states could aid in the identification of robust associations between brain and behavior.

LINKING ETHOGRAMS AND NEURAL CIRCUITS TO COMPUTATIONAL MODELS

Ultimately, we would like to have computational models that explain the information processing performed by brains and predict neural and behavioral activity (Kriegeskorte & Douglas 2018). Reinforcement learning algorithms are commonly divided into two categories: model-free (MF) and model-based (MB). MF learning gradually updates cached value estimates retrospectively from experience, and MF control uses those value estimates for decision-making. MF learning is associated with algorithms like temporal-difference learning (Sutton and Barto, 1998) and Q-learning (Watkins and Dayan, 1992). By contrast, MB algorithms calculate prospectively, for instance by simulating possible future states (as in replay and preplay; (Mattar and Daw, 2018; Pfeiffer and Foster, 2013; Wise et al., 2020)).

Applying computational models to dynamic and unconstrained behaviors presents a new challenge: in terms of decision making, we are now faced with a series of complex decisions, unconstrained in time and reflected in behavior more elaborate than a button press. For example, human tasks that depend on multi-step decision trees typically focus on a single initial decision point at the start of the tree, with a decision made once per trial (Daw et al., 2011; Momennejad et al., 2017). In generating richer behavioral datasets, virtual ecologies make such models more difficult to test in relation to behavior. However, recent advances in deep RL algorithms have created the opportunity for modeling complex behavior in virtual ecologies (Mnih et al., 2015). This is especially true in the model-free case, though now showing promise for incorporating model-based algorithms as well (Schrittwieser et al., 2020).

In one notable example, artificial agents learned through reinforcement learning to play a first-person shooter computer game with realistic physics and complex objectives involving both competition and teamwork (capture the flag) (Jaderberg et al., 2019). This was a computer game played by simulated agents, so in principle the researcher could have full experimental access to any internal variable of the system. However, in practice, the long timescale and high dimensionality of the generated dataset meant computational ethology methods were still necessary to analyze the resulting agent behavior. In particular, the authors employed an unsupervised computational ethology analysis inspired by Wiltischko et al. (2015). Results showed that internal representations of important game events like teammate following and home-base defense emerged as a result of reinforcement learning in this environment, suggesting the significant extent to which the agents come to “understand” these game-related concepts.

Multi-agent deep reinforcement learning algorithms have also been applied to model the cooperation behaviors of groups in mixed-motivation settings (Foerster et al., 2018; Leibo et al., 2017; Lerer and Peysakhovich, 2018; Perolat et al., 2017). This line of work extends classical game-theoretic models based on matrix game formulations (Camerer (2003) to capture complex spatiotemporally extended virtual ecologies. “Rational” (selfish) agent models are no better at cooperating in virtual ecologies than they are in matrix games. For example, Perolat et al. (2017) studied a virtual ecology designed to model common-pool resource appropriation (Janssen et al., 2010). As in the human studies where individuals could not communicate (Janssen et al., 2014), they found that failures of cooperation yield a tragedy of the commons, where overuse of resources degrades the environment to mutual detriment. In other circumstances, they found the spontaneous emergence of exclusion behaviors and inequality (Perolat et al., 2017). Another study measured proxemics (i.e., distance to conspecifics) and preferences to different individuals that emerged from reinforcement learning in mixed motivation scenarios (McKee et al., 2020). This approach, underpinned by models derived from multi-agent reinforcement learning research, holds promise to uncover interactions between the fine spatiotemporal structure of behavior and its strategic content that are not easily seen in traditional paradigms.

The use of complex behavioral measures will require a move away from models of simple choice likelihood based on the value of individual options, as is common in standard decision-making tasks, toward models that make predictions about more complex ongoing aspects of locomotor activity. Modelling approaches to more ecologically realistic behaviors have been developed, for example using computational models of reward-guided place preference-like behavior (Wu et al., 2018) or threat-guided place aversion (Wise and Dolan, 2020) in 2D environments, however these have focused on relatively coarse-grained trial-by-trial behavioral measures of location. Fully explaining behavior in these environments will require modelling not only position on a 2D grid, but also motion measures such as velocity and acceleration. In essence, the action space for modelling becomes larger and more complex. Additionally, it will be necessary to determine the optimal level of granularity for behavioral outcome measures. However, despite this added complexity, these rich measures may improve our behavioral models greatly. Additionally, unconstrained virtual ecologies naturally result in richer and more varied behaviors, a characteristic that

provides more flexibility in designing experiments to elicit behavioral patterns that will differentiate candidate models (Palminteri et al., 2017).

THE USES AND FUTURE OF VIRTUAL REALITY (VR)

2D environments provide simple and clear ways to provide the subjects with task-relevant information that may not be visible from a first-person point of view. On the other hand, in some circumstances the perceptual uncertainty provided by a 3D environment could be useful. For example, given that 3D environments most accurately represent our perception of the real world, they increase the ecological validity of the task, while also allowing identification of behavioral dynamics including intermittent locomotion and eye-movements. To strengthen connections between behavioral research with human participants and its counterpart with artificial agent participants, it is sometimes even helpful to simulate within a 3D environment a scenario where the agent stands in front of a flat screen to perform a task. This allows for virtual “eye movements” on the 2D environment projected within the simulated 3D environment (Leibo et al., 2018). Furthermore, 3D environments where aspects of the world are obscured from view encourage the subject to build and use an internal model of their environment, rather than relying on what is directly in front of them (Wayne et al., 2018).

Immersive VR technology has significantly moved forward in the last decade. Its use in human neuroimaging and psychological experiments has been revolutionary as it provides a more enriched and naturalistic approach to computerized environments (Bohil et al., 2011; Reggente et al., 2018; see Fig. 3)). As with simple 2D environments, defensive behaviors can be measured including thigmotaxis, place aversion and escape. As VR becomes more realistic and feasible, for example with smaller headsets equipped with improved eye-tracking and pupillometry capability, and integrated with mobile intracranial EEG (iEEG; Topalovic et al., 2020), MEG and fMRI hardware (Fig. 3), there will be a need to understand how to use this technology best. Limitations on the immersive experience in the MRI scanner, however, provide a major challenge given the absence of body-based cues related to vestibular, motor, and somatosensory input. Several creative ways around these limitations have begun to merge including the use of 3D glasses and VR training outside the scanner. For example, Huffman and Ekstrom, (Huffman and Ekstrom, 2019) used VR outside the MRI scanner to train people in enriched (on a treadmill), limited (using a joystick and head-mounted display) and impoverished (joystick only) environments, where their goal was to spatially navigate a virtual large-scale environment. After training, subjects were placed in an MRI scanner where they performed a judgement of relative direction task, showing that body-based cues influenced spatial navigation. Full immersion in a virtual environment did not result in any behavioral or neural differences between conditions and proved the hypothesis that body-based cues are not necessary for retrieval of spatial information related to large-scale environments, a hypothesis that would be difficult to test without VR. Numerous other studies have shown VR using desktop computers, head mounted displays and other technologies can be used to study several aspects of human cognition and related neural mechanisms (Ekstrom et al., 2003; Jacobs et al., 2013; Chrastil et al., 2015; Diersch and Wolbers, 2019; Hartley et al., 2003). While several aspects of real-world human behavior seem to be modeled effectively in VR (Huffman and Ekstrom,

2021; Chrastil and Warren, 2015), there may be certain cognitive abilities that do not transfer as well in VR. For example the type of learning strategy used to accomplish a spatial navigation task and the transferring of this knowledge to novel situations may be altered in VR compared to the real-world (Clemenson et al., 2020; Hejtmanek et al., 2020). It will thus be important for future studies to determine the boundary-conditions in which reality can and cannot be effectively modeled with immersive VR technologies.

Evidence suggests that some of these challenges may become less prominent, however, as the level of immersion continues to improve and full vestibular, motor, and somatosensory inputs are present (Huffman and Ekstrom, 2021; Hejtmanek et al., 2020). Recently, Topalovic et al., (2020) combined fully immersive VR technologies with full body- and eye-tracking, as well as biometrics (e.g., heart rate, respiration, and galvanic skin response) in participants implanted with chronic deep brain devices capable of recording iEEG activity that is unsusceptible to motion-related artifacts. One recent study is using combined immersive VR and iEEG recordings in moving subjects to understand the neural representations of actual physical space during memory formation and retrieval (Aghajan et al., 2019). Future research studies of a similar nature can use this technology to record synchronized behavioral and neural data from a wide range of brain structures (e.g., amygdala, hippocampus, vmPFC) in naturally behaving humans.

This technology can also be integrated with augmented reality (AR) headsets that allow for objects/events/agents to be superimposed onto the real-world (Topalovic et al., 2020). One recent study used full-body motion capture combined with on-body world-view cameras and eye-tracking in an environment shared with others to investigate social neural mechanisms of location-encoding (Stangl et al., 2021; Fig. 3 A–C). Altogether, these studies open up exciting opportunities for applying computational ethological methods to high-resolution behavioral data captured during naturalistic experiences in social scenarios within real, virtual, or augmented environments. The use of iEEG in participants with temporary deep brain electrodes (e.g., in the epilepsy monitoring unit) combined with biometric recordings is also primed for VR use (Yilmaz Balban et al., 2020). Further, with the creation of moveable optically pumped magnetometer magnetoencephalography (OPM-MEG; (Boto et al., 2018)), providing electrophysiological measurements at the millisecond resolution, there is promise in combining VR with spatially and temporally high-resolution brain imaging on a wider population of subjects not limited to only those who have implanted brain electrodes.

A small but accumulating set of VR studies are beginning to demonstrate how VR can be used to study fear and anxiety. Traditionally, it has been difficult to expose subjects to realistic threats in the lab, and instead studies have typically relied on painful stimuli such as electric shocks. In contrast, VR has permitted the assessment of common fears, such as height and public speaking, in the lab (Gromer et al., 2019; Stupar-Rutenfrans et al., 2017). Virtual versions of the Elevated Plus Maze (EPM) have been used in humans, and like rodents, high anxiety individuals show increased avoidance of open arms (Biedermann et al., 2017). Other have shown that VR can be used to elicit anxiety in flight phobics (Mühlberger et al., 2001). Interestingly, VR can have therapeutic effects in flight and spiders phobics (Mühlberger et al., 2003; Shibani et al., 2015). Similarly, exposure therapy in VR

has begun to see use in treating PTSD in war veterans. Rizzo et al. (Rizzo et al., 2010) developed an exposure therapy system for veterans of Iraq and Afghanistan combining realistic 3D environments in VR with physiological measurements such as the galvanic skin response. The same team later employed this system in conjunction with fMRI to monitor improvements in cerebral function in veterans undergoing treatment for PTSD (Roy et al., 2010). Finally, Yilmaz Balban et al., (Yilmaz Balban et al., 2020) have shown that exposure to virtual threats such as scary heights can elicit increases in autonomic arousal. Further, using iEEG, the authors show higher gamma activity in the insula for virtual heights compared to no heights control conditions. Together, these studies show how VR can elicit autonomic and neural responses to threat and show promise for implementing the behavioral measures advocated by computational ethology.

CONCLUSIONS

Introducing methods from computational ethology to human neuroscience promises to help us uncover novel behavioral assays and better understand the dynamic nature of the human brain and how it might function in the real world. Further, through the extraction of individualized ethograms from tasks using virtual ecologies, we can link patterns of behavior in naturalistic environments to brain states, potentially revealing links between neural circuits and behavior that are not observable with current methods. Current experimental paradigms are restricted to specific processes thought to be important by the researcher, and therefore miss behavioral characteristics (in both healthy function and psychiatric disorders) that might be essential for understanding brain function in naturalistic environments. The approaches laid out in this paper enable the quantification of behavior and its disruption in a variety of psychiatric conditions at both a behavioral and neural level. Unsupervised methods of behavior classification may identify novel patterns of behavior that are diagnostic of particular symptom clusters (see Box 1). While there are challenges to overcome, the use of approaches advocated by computational ethology (i.e., unsupervised quantification of behavior) provides an exciting and powerful approach to engaging the dynamics of natural cognition in human neuroscience.

Acknowledgements.

This work was supported by the US National Institute of Mental Health grant no. 2P50MH094258, a Chen Institute Award (grant no. P2026052) and a Templeton Foundation grant TWCF0366 (All to D.M.). T.W. is supported by a Wellcome Trust Sir Henry Wellcome Fellowship (grant no.206460/17/Z). This work is also supported by the National Institutes of Health (NIH) National Institute of Neurological Disorders and Stroke (NINDS; grants NS103802 and NS117838), the McKnight Foundation (Technological Innovations Award in Neuroscience to N.S.) and a Keck Junior Faculty Award (to N.S.). We thank Matthew Botvinick for his feedback on an earlier version of this paper.

GLOSSARY

Behavioral motif

A unit of organized movement often interchangeably used with the terms “motif,” “moveme,” “module,” “primitive,” and “syllable” (see Datta et al., 2019; Anderson and Perona, 2014).

Dimensionality

The number of variables that are present in a dataset (Also see Dimensionality reduction; Datta et al., 2019).

Ethogram

A repeatable and predefined set of movements that are either learned or hard-wired. These include such things as thigmotaxis (see below), approach, and pauses.

Machine learning (ML)

Where computers are programmed to learn without explicit instructions, providing accurate predictions based on either labelled or unlabeled training data. Using a variety of mathematical models including support vector machines or deep neural networks, a data set is first used to train the algorithms. Once the algorithm is trained, it is then tested on a test data set. In the case of human behavior, ML can be used to detect and categorize human and animal behavioral motifs using what has been called a “action classifier.”

Model-based (MB) inference and decision making

Inferences and decisions based on an internal model of the world. MB methods exploit a (possibly learned) model of the environment to calculate prospectively the likely consequences of actions, for instance by simulating possible future states.

Model-free (MF) inference and decision making

The agent learns what to do to maximize long-run return or learns value estimates of those long-run returns. MF methods acquire values by a bootstrapping process of enforcing consistency between successive estimates.

Protean escape

Unpredictable escape trajectories, such as zig-zagging, that prevent a predator anticipating the future position or actions of its prey.

Temporal dynamics

How behavior features change over time.

Trajectory

The movement of the agent through time and space.

Thigmotaxis

A measure of anxiety where animals stay close to walls rather than maneuvering in open spaces.

Virtual ecology

Self-contained virtual environments where subjects can freely move throughout the environment.

For more definitions, see Anderson and Perona, 2014; and Datta et al., 2019.

REFERENCES

- Aghajian Z, Schuette P, Fields TA, Tran ME, Siddiqui SM, Hasulak NR, Tchong TK, Eliashiv D, Mankin EA, Stern J, et al. (2017). Theta Oscillations in the Human Medial Temporal Lobe during Real-World Ambulatory Movement. *Curr Biol* 27, 3743–3751.e3. [PubMed: 29199073]
- Allen EA, Damaraju E, Plis SM, Erhardt EB, Eichele T, and Calhoun VD (2014). Tracking Whole-Brain Connectivity Dynamics in the Resting State. *Cereb. Cortex* 24, 663–676. [PubMed: 23146964]
- Anderson DJ, and Perona P (2014). Toward a Science of Computational Ethology. *Neuron* 84, 18–31. [PubMed: 25277452]
- Astur RS, Carew AW, and Deaton BE (2014). Conditioned place preferences in humans using virtual reality. *Behavioural Brain Research* 267, 173–177. [PubMed: 24657735]
- Astur RS, Purton AJ, Zaniewski MJ, Cimadevilla J, and Markus EJ (2016). Human sex differences in solving a virtual navigation problem. *Behav Brain Res* 308, 236–243. [PubMed: 27108050]
- Babayan BM, and Konen CS (2019). Behavior Matters. *Neuron* 104, 1. [PubMed: 31600506]
- Bach DR, Guitart-Masip M, Packard PA, Miró J, Falip M, Fuentemilla L, and Dolan RJ (2014). Human hippocampus arbitrates approach-avoidance conflict. *Curr. Biol* 24, 541–547. [PubMed: 24560572]
- Baker AP, Brookes MJ, Rezek IA, Smith SM, Behrens T, Probert Smith PJ, and Woolrich M (2014). Fast transient networks in spontaneous human brain activity. *ELife* 3, e01867. [PubMed: 24668169]
- Balleine BW (2019). The Meaning of Behavior: Discriminating Reflex and Volition in the Brain. *Neuron* 104, 47–62. [PubMed: 31600515]
- Beattie C, Köppe T, Duéñez-Guzmán EA, and Leibo JZ (2020). DeepMind Lab2D. ArXiv:2011.07027 [Cs]
- Berman GJ, Choi DM, Bialek W, and Shaevitz JW (2014). Mapping the stereotyped behaviour of freely moving fruit flies. *Journal of The Royal Society Interface* 11, 20140672.
- Biedermann SV, Biedermann DG, Wenzlaff F, Kurjak T, Nouri S, Auer MK, Wiedemann K, Briken P, Haaker J, Lonsdorf TB, et al. (2017). An elevated plus-maze in mixed reality for studying human anxiety-related behavior. *BMC Biol* 15, 125. [PubMed: 29268740]
- Bohil CJ, Alicea B, and Biocca FA (2011). Virtual reality in neuroscience research and therapy. *Nat Rev Neurosci* 12, 752–762. [PubMed: 22048061]
- Boto E, Holmes N, Leggett J, Roberts G, Shah V, Meyer SS, Muñoz LD, Mullinger KJ, Tierney TM, Bestmann S, et al. (2018). Moving magnetoencephalography towards real-world applications with a wearable system. *Nature* 555, 657–661. [PubMed: 29562238]
- Brady LS, Potter WZ, and Gordon JA (2019). Redirecting the revolution: new developments in drug development for psychiatry. *Expert Opinion on Drug Discovery* 14, 1213–1219. [PubMed: 31545131]
- Brookes DH, Park H, and Listgarten J (2020). Conditioning by adaptive sampling for robust design. ArXiv:1901.10060 [Cs, Stat]
- Burock MA, Buckner RL, Woldorff MG, Rosen BR, and Dale AM (1998). Randomized event-related experimental designs allow for extremely rapid presentation rates using functional MRI. *Neuroreport* 9, 3735–3739. [PubMed: 9858388]
- Calhoun AJ, Pillow JW, and Murthy M (2019). Unsupervised identification of the internal states that shape natural behavior. *BioRxiv* 691196.
- Camerer CF (2003). Behavioural studies of strategic thinking in games. *Trends in Cognitive Sciences*. 7, (5), 225–231. [PubMed: 12757825]
- Carl E, Stein AT, Levihn-Coon A, Pogue JR, Rothbaum B, Emmelkamp P, Asmundson GJG, Carlbring P, and Powers MB (2019). Virtual reality exposure therapy for anxiety and related disorders: A meta-analysis of randomized controlled trials. *J Anxiety Disord* 61, 27–36. [PubMed: 30287083]
- Carreira J, and Zisserman A (2017). Quo Vadis, Action Recognition? A New Model and the Kinetics Dataset. In 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 4724–4733.

- de Chaumont F, Coura RD-S, Serreau P, Cressant A, Chabout J, Granon S, and Olivo-Marin J-C (2012). Computerized video analysis of social interactions in mice. *Nature Methods* 9, 410–417. [PubMed: 22388289]
- Clemens J, Girardin CC, Coen P, Guan X-J, Dickson BJ, and Murthy M (2015). Connecting Neural Codes with Behavior in the Auditory System of *Drosophila*. *Neuron* 87, 1332–1343. [PubMed: 26365767]
- Cook JL, Blakemore S-J, and Press C (2013). Atypical basic movement kinematics in autism spectrum conditions. *Brain* 136, 2816–2824. [PubMed: 23983031]
- Cooper J, William E, and Blumstein DT (2015). *Escaping From Predators: An Integrative View of Escape Decisions* (Cambridge: Cambridge University Press).
- Chrastil ER, Sherrill KR, Hasselmo ME, and Stern CE (2015). There and back again: hippocampus and retrosplenial cortex track homing distance during human path integration. *J. Neurosci.* 35, 15442–15452. doi: 10.1523/JNEUROSCI.1209-15.2015 [PubMed: 26586830]
- Clemenson Gregory D, Wang Lulian, Zeqian Mao, Stark Shauna M, Stark Craig EL Exploring the Spatial Relationships Between Real and Virtual Experiences: What Transfers and What Doesn't. *Frontiers in Virtual Reality* 2020. Doc: 10.3389/frvir.2020.572122.
- Datta SR, Anderson DJ, Branson K, Perona P, and Leifer A (2019). Computational Neuroethology: A Call to Action. *Neuron* 104, 11–24. [PubMed: 31600508]
- Davis T, Larocque KF, Mumford JA, Norman KA, Wagner AD & Poldrack RA 2014. What do differences between multi-voxel and univariate analysis mean? How subject-, voxel-, and trial-level variance impact fMRI analysis. *Neuroimage*, 97, 271–283. [PubMed: 24768930]
- Daw ND, Gershman SJ, Seymour B, Dayan P, and Dolan RJ (2011). Model-Based Influences on Humans' Choices and Striatal Prediction Errors. *Neuron* 69, 1204–1215. [PubMed: 21435563]
- Dennis M, Jaques N, Vinitzky E, Bayen A, Russell S, Critch A, and Levine S (2020). Emergent Complexity and Zero-shot Transfer via Unsupervised Environment Design. *Advances in Neural Information Processing Systems* 33.
- Diersch N, and Wolbers T (2019). The potential of virtual reality for spatial navigation research across the adult lifespan. *J. Exp. Biol* 222:jeb.187252. doi: 10.1242/jeb.187252
- Doersch C (2016). Tutorial on Variational Autoencoders. ArXiv:1606.05908 [Cs, Stat]
- Donahue J, Jia Y, Vinyals O, Hoffman J, Zhang N, Tzeng E, and Darrell T (2013). DeCAF: A Deep Convolutional Activation Feature for Generic Visual Recognition. ArXiv:1310.1531 [Cs]
- Dosovitskiy A, Fischer P, Springenberg JT, Riedmiller M, and Brox T (2015). Discriminative Unsupervised Feature Learning with Exemplar Convolutional Neural Networks. ArXiv:1406.6909 [Cs]
- Ekstrom AD, Kahana MJ, Caplan JB, Fields TA, Isham EA, Newman EL, et al. (2003). Cellular networks underlying human spatial navigation. *Nature* 425, 184–188. doi: 10.1038/nature01964 [PubMed: 12968182]
- Fajen BR (2007). Affordance-Based Control of Visually Guided Action. *Ecological Psychology* 19, 383–410.
- Fajen BR, and Warren WH (2003). Behavioral dynamics of steering, obstacle avoidance, and route selection. *J Exp Psychol Hum Percept Perform* 29, 343–362. [PubMed: 12760620]
- Fajen BR, and Warren WH (2007). Behavioral dynamics of intercepting a moving target. *Exp Brain Res* 180, 303–319. [PubMed: 17273872]
- Foerster JN, Chen RY, Al-Shedivat M, Whiteson S, Abbeel P, and Mordatch I (2018). Learning with Opponent-Learning Awareness. ArXiv:1709.04326 [Cs]
- Fung BJ, Qi S, Hassabis D, Daw N, and Mobbs D (2019). Slow escape decisions are swayed by trait anxiety. *Nature Human Behaviour* 1.
- Geisler WS (2008). Visual perception and the statistical properties of natural scenes. *Annu Rev Psychol* 59, 167–192. [PubMed: 17705683]
- Gillan CM, Kosinski M, Whelan R, Phelps EA, and Daw ND (2016). Characterizing a psychiatric symptom dimension related to deficits in goal-directed control. *ELife* 5, e11305. [PubMed: 26928075]

- Gold AL, Morey RA, and McCarthy G (2015). Amygdala–Prefrontal Cortex Functional Connectivity During Threat-Induced Anxiety and Goal Distraction. *Biol Psychiatry* 77, 394–403. [PubMed: 24882566]
- Goodfellow IJ, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A, and Bengio Y (2014). Generative Adversarial Networks. *ArXiv:1406.2661 [Cs, Stat]*
- Grabb MC, Cross AJ, Potter WZ, and McCracken JT (2016). De-risking Psychiatric Drug Development: The NIMH’s Fast Fail Program, a Novel Precompetitive Model. *J Clin Psychopharmacol* 36, 419–421. [PubMed: 27404513]
- Graving JM, Chae D, Naik H, Li L, Koger B, Costelloe BR, and Couzin ID (2019). DeepPoseKit, a software toolkit for fast and robust animal pose estimation using deep learning. *ELife* 8, e47994. [PubMed: 31570119]
- Gromer D, Reinke M, Christner I, and Pauli P (2019). Causal Interactive Links Between Presence and Fear in Virtual Reality Height Exposure. *Front Psychol* 10.
- Grupe DW, and Nitschke JB (2013). Uncertainty and anticipation in anxiety: an integrated neurobiological and psychological perspective. *Nat Rev Neurosci* 14, 488–501. [PubMed: 23783199]
- Hamilton WD (1971). Geometry for the selfish herd. *Journal of Theoretical Biology* 31, 295–311. [PubMed: 5104951]
- Hamilton LS, and Huth AG (2020). The revolution will not be controlled: natural stimuli in speech neuroscience. *Language, Cognition and Neuroscience* 35, 573–582.
- Hartley T, Maguire EA, Spiers HJ, and Burgess N (2003). The well-worn route and the path less traveled: distinct neural bases of route following and wayfinding in humans. *Neuron* 37, 877–888. doi: 10.1016/S0896-6273(03)00095-3. [PubMed: 12628177]
- Hasson U, Nastase SA, and Goldstein A (2020). Direct Fit to Nature: An Evolutionary Perspective on Biological and Artificial Neural Networks. *Neuron* 105, 416–434. [PubMed: 32027833]
- Haxby JV, Gobbini MI, Furey ML, Ishai A, Schouten JL, and Pietrini P (2001). Distributed and Overlapping Representations of Faces and Objects in Ventral Temporal Cortex. *Science* 293, 2425–2430. [PubMed: 11577229]
- Hejtmanek L, Starrett M, Ferrer E, Ekstrom AD. How Much of What We Learn in Virtual Reality Transfers to Real-World Navigation? *Multisens Res* 2020 Mar 17;33(4–5):479–503. doi: 10.1163/22134808-20201445. [PubMed: 31972540]
- Holzinger A, Plass M, Kickmeier-Rust M, Holzinger K, Cri an GC, Pintea C-M, and Palade V (2019). Interactive machine learning: experimental evidence for the human in the algorithmic loop. *Appl Intell* 49, 2401–2414.
- Hoopfer ED, Jung Y, Inagaki HK, Rubin GM, and Anderson DJ (2015). P1 interneurons promote a persistent internal state that enhances inter-male aggression in *Drosophila*. *ELife* 4, e11346. [PubMed: 26714106]
- Huffman DJ, and Ekstrom AD (2019). A Modality-Independent Network Underlies the Retrieval of Large-Scale Spatial Environments in the Human Brain. *Neuron* 104, 611–622.e7. [PubMed: 31540825]
- Huffman DJ, Ekstrom AD. An Important Step toward Understanding the Role of Body-based Cues on Human Spatial Memory for Large-Scale Environments. *J Cogn Neurosci* 2021 Feb;33(2):167–179. [PubMed: 33226317]
- Humphries DA, and Driver PM (1970). Protean defence by prey animals. *Oecologia* 5, 285–302. [PubMed: 28309783]
- Ilg E, Mayer N, Saikia T, Keuper M, Dosovitskiy A, and Brox T (2016). FlowNet 2.0: Evolution of Optical Flow Estimation with Deep Networks. *ArXiv:1612.01925 [Cs]*
- Insafutdinov E, Pishchulin L, Andres B, Andriluka M, and Schiele B (2016). DeeperCut: A Deeper, Stronger, and Faster Multi-Person Pose Estimation Model. *ArXiv:1605.03170 [Cs]*
- Insel T, Cuthbert B, Garvey M, Heinssen R, Pine DS, Quinn K, Sanislow C, and Wang P (2010). Research Domain Criteria (RDoC): Toward a New Classification Framework for Research on Mental Disorders. *AJP* 167, 748–751.

- Jacobs J, Weidemann CT, Miller JF, Solway A, Burke JF, Wei XX, et al. (2013). Direct recordings of grid-like neuronal activity in human spatial navigation. *Nat. Neurosci* 16, 1188–1190. [PubMed: 23912946]
- Jaderberg M, Czarnecki WM, Dunning I, Marris L, Lever G, Castañeda AG, Beattie C, Rabinowitz NC, Morcos AS, Ruderman A, et al. (2019). Human-level performance in 3D multiplayer games with population-based reinforcement learning. *Science* 364, 859–865. [PubMed: 31147514]
- Janssen M (2010). Introducing Ecological Dynamics into Common-Pool Resource Experiments. *Ecology and Society* 15.
- Janssen M, Tyson M, and Lee A (2014). The effect of constrained communication and limited information in governing a common resource. *International Journal of the Commons* 8.
- Jovanic T, Schneider-Mizell CM, Shao M, Masson J-B, Denisov G, Fetter RD, Mensh BD, Truman JW, Cardona A, and Zlatić M (2016). Competitive Disinhibition Mediates Behavioral Choice and Sequences in *Drosophila*. *Cell* 167, 858–870.e19. [PubMed: 27720450]
- Kaiser RH, Andrews-Hanna JR, Wager TD, and Pizzagalli DA (2015). Large-Scale Network Dysfunction in Major Depressive Disorder: A Meta-analysis of Resting-State Functional Connectivity. *JAMA Psychiatry* 72, 603–611. [PubMed: 25785575]
- Kasparov G (2018). Chess, a *Drosophila* of reasoning. *Science* 362, 1087–1087. [PubMed: 30523085]
- Kay KN, Naselaris T, Prenger RJ, and Gallant JL (2008). Identifying natural images from human brain activity. *Nature* 452, 352–355. [PubMed: 18322462]
- Korn CW, Vunder J, Miró J, Fuentemilla L, Hurlmann R, and Bach DR (2017). Amygdala Lesions Reduce Anxiety-like Behavior in a Human Benzodiazepine-Sensitive Approach-Avoidance Conflict Test. *Biol Psychiatry* 82, 522–531. [PubMed: 28364943]
- Kriegeskorte N, and Douglas PK (2019). Interpreting encoding and decoding models. *Current Opinion in Neurobiology* 55, 167–179. [PubMed: 31039527]
- Kriegeskorte N, Mur M, and Bandettini PA (2008). Representational similarity analysis-connecting the branches of systems neuroscience. *Front. Syst. Neurosci* 2.
- LeDoux J (2012). Rethinking the Emotional Brain. *Neuron* 73, 653–676. [PubMed: 22365542]
- Lee M, and Seok J (2019). Controllable Generative Adversarial Network. *IEEE Access* 7, 28158–8169.
- Leibo JZ, Zambaldi V, Lanctot M, Marecki J, and Graepel T (2017). Multi-agent Reinforcement Learning in Sequential Social Dilemmas. ArXiv:1702.03037 [Cs]
- Leibo JZ, d'Áutume C. de M., Zoran D, Amos D, Beattie C, Anderson K, Castañeda AG, Sanchez M, Green S, Gruslyś A, et al. (2018). Psychlab: A Psychology Laboratory for Deep Reinforcement Learning Agents. ArXiv:1801.08116 [Cs, q-Bio]
- Lerer A, and Peysakhovich A (2018). Maintaining cooperation in complex social dilemmas using deep reinforcement learning. ArXiv:1707.01068 [Cs]
- Li J, Ma H, and Tomizuka M (2019). Conditional Generative Neural System for Probabilistic Trajectory Prediction. ArXiv:1905.01631 [Cs, Stat]
- Mathis A, Mamidanna P, Cury KM, Abe T, Murthy VN, Mathis MW, and Bethge M (2018). DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. *Nat Neurosci* 21, 1281–1289. [PubMed: 30127430]
- Mattar MG, and Daw ND (2018). Prioritized memory access explains planning and hippocampal replay. *Nature Neuroscience* 1.
- McKee KR, Gemp I, McWilliams B, Duéñez-Guzmán EA, Hughes E, and Leibo JZ (2020). Social diversity and social preferences in mixed-motive reinforcement learning. ArXiv:2002.02325 [Cs]
- Mets DG, and Brainard MS (2018). An automated approach to the quantitation of vocalizations and vocal learning in the songbird. *PLOS Computational Biology* 14, e1006437. [PubMed: 30169523]
- Meyer C, Padmala S, and Pessoa L (2019). Dynamic Threat Processing. *J Cogn Neurosci* 31, 522–542. [PubMed: 30513044]
- Mnih V, Kavukcuoglu K, Silver D, Rusu AA, Veness J, Bellemare MG, Graves A, Riedmiller M, Fidjeland AK, Ostrovski G, et al. (2015). Human-level control through deep reinforcement learning. *Nature* 518, 529–533. [PubMed: 25719670]

- Mnih V, Badia AP, Mirza M, Graves A, Lillicrap TP, Harley T, Silver D, and Kavukcuoglu K (2016). Asynchronous Methods for Deep Reinforcement Learning. ArXiv:1602.01783 [Cs]
- Mobbs D, Petrovic P, Marchant JL, Hassabis D, Weiskopf N, Seymour B, Dolan RJ, and Frith CD (2007). When Fear Is Near: Threat Imminence Elicits Prefrontal-Periaqueductal Gray Shifts in Humans. *Science* 317, 1079–1083. [PubMed: 17717184]
- Mobbs D, Marchant JL, Hassabis D, Seymour B, Tan G, Gray M, Petrovic P, Dolan RJ, and Frith CD (2009). From Threat to Fear: The Neural Organization of Defensive Fear Systems in Humans. *J. Neurosci* 29, 12236–12243. [PubMed: 19793982]
- Mobbs D, Hassabis D, Yu R, Chu C, Rushworth M, Boorman E, and Dalgleish T (2013). Foraging under Competition: The Neural Basis of Input-Matching in Humans. *J. Neurosci* 33, 9866–9872. [PubMed: 23739983]
- Mobbs D, Trimmer PC, Blumstein DT, and Dayan P (2018). Foraging for foundations in decision neuroscience: insights from ethology. *Nature Reviews Neuroscience* 1.
- Mobbs D, Headley DB, Ding W, and Dayan P (2020). Space, Time, and Fear: Survival Computations along Defensive Circuits. *Trends in Cognitive Sciences* 24, 228–241. [PubMed: 32029360]
- Molet M, Billiet G, and Bardo MT (2013). Conditioned place preference and aversion for music in a virtual reality environment. *Behavioural Processes* 92, 31–35. [PubMed: 23089383]
- Momennejad I, Russek EM, Cheong JH, Botvinick MM, Daw ND, and Gershman SJ (2017). The successor representation in human reinforcement learning. *Nature Human Behaviour* 1, 680.
- Mühlberger A, Herrmann MJ, Wiedemann G, Ellgring H, and Pauli P (2001). Repeated exposure of flight phobics to flights in virtual reality. *Behaviour Research and Therapy* 39, 1033–1050. [PubMed: 11520010]
- Mühlberger A, Wiedemann G, and Pauli P (2003). Efficacy of a one-session virtual reality exposure treatment for fear of flying. *Psychother Res* 13, 323–336. [PubMed: 21827246]
- Musall S, Kaufman MT, Juavinett AL, Gluf S, and Churchland AK (2019). Single-trial neural dynamics are dominated by richly varied movements. *Nature Neuroscience* 22, 1677–1686. [PubMed: 31551604]
- Niv Y (2020). The primacy of behavioral research for understanding the brain
- Palminteri S, Wyart V, and Koechlin E (2017). The Importance of Falsification in Computational Cognitive Modeling. *Trends in Cognitive Sciences* 0.
- Patterson GR (1974). A Basis for Identifying Stimuli Which Control Behaviors in Natural Settings. *Child Development* 45, 900–911. [PubMed: 4143880]
- Pereira TD, Aldarondo DE, Willmore L, Kislin M, Wang SS-H, Murthy M, and Shaevitz JW (2019). Fast animal pose estimation using deep neural networks. *Nat Methods* 16, 117–125. [PubMed: 30573820]
- Perolat J, Leibo JZ, Zambaldi V, Beattie C, Tuyls K, and Graepel T (2017). A multi-agent reinforcement learning model of common-pool resource appropriation. ArXiv:1707.06600 [Cs, q-Bio]
- Pfeiffer BE, and Foster DJ (2013). Hippocampal place-cell sequences depict future paths to remembered goals. *Nature* 497, 74–79. [PubMed: 23594744]
- Qi S, Hassabis D, Sun J, Guo F, Daw N, and Mobbs D (2018). How cognitive and reactive fear circuits optimize escape decisions in humans. *PNAS* 115, 3186–3191. [PubMed: 29507207]
- Qi S, Cross L, Wise T, Sui X, O’Doherty J, and Mobbs D (2020). The Role of the Medial Prefrontal Cortex in Spatial Margin of Safety Calculations. *BioRxiv* 2020.06.05.137075
- Reggente N, Essoe JK-Y, Aghajani ZM, Tavakoli AV, McGuire JF, Suthana NA, and Rissman J (2018). Enhancing the Ecological Validity of fMRI Memory Research Using Virtual Reality. *Front. Neurosci* 12.
- Rizzo AS, Difede J, Rothbaum BO, Reger G, Spitalnick J, Cukor J, and McLay R (2010). Development and early evaluation of the Virtual Iraq/Afghanistan exposure therapy system for combat-related PTSD. *Ann N Y Acad Sci* 1208, 114–125. [PubMed: 20955333]
- Robie AA, Hirokawa J, Edwards AW, Umayam LA, Lee A, Phillips ML, Card GM, Korff W, Rubin GM, Simpson JH, et al. (2017). Mapping the Neural Substrates of Behavior. *Cell* 170, 393–406.e28. [PubMed: 28709004]

- Roy M, Harvey P-O, Berlim MT, Mamdani F, Beaulieu M-M, Turecki G, and Lepage M (2010). Medial prefrontal cortex activity during memory encoding of pictures and its relation to symptomatic improvement after citalopram treatment in patients with major depression. *J Psychiatry Neurosci* 35, 152–162. [PubMed: 20420766]
- Sainburg T, Thielk M, and Gentner TQ (2019). Latent space visualization, characterization, and generation of diverse vocal communication signals. *BioRxiv* 870311
- Schrittwieser J, Antonoglou I, Hubert T, Simonyan K, Sifre L, Schmitt S, Guez A, Lockhart E, Hassabis D, Graepel T, et al. (2020). Mastering Atari, Go, Chess and Shogi by Planning with a Learned Model. *ArXiv:1911.08265 [Cs, Stat]*
- Sharp PB, Dolan RJ, and Eldar E (2020). Cognitive map learning is disrupted in compulsivity and anxious arousal (*PsyArXiv*)
- Shiban Y, Brütting J, Pauli P, and Mühlberger A (2015). Fear reactivation prior to exposure therapy: does it facilitate the effects of VR exposure in a randomized clinical sample? *J Behav Ther Exp Psychiatry* 46, 133–140. [PubMed: 25460259]
- Shine JM, and Poldrack RA (2018). Principles of dynamic network reconfiguration across diverse brain states. *NeuroImage* 180, 396–405. [PubMed: 28782684]
- Shine JM, Breakspear M, Bell PT, Martens KAE, Shine R, Koyejo O, Sporns O, and Poldrack RA (2019). Human cognition involves the dynamic integration of neural activity and neuromodulatory systems. *Nat Neurosci* 22, 289–296. [PubMed: 30664771]
- Silston B, Wise T, Qi S, Sui X, Dayan P, and Mobbs D (2020). Neural encoding of socially adjusted value during competitive and hazardous foraging. *BioRxiv* 2020.09.11.294058
- Silver D, and Veness J (2010). Monte-Carlo Planning in Large POMDPs. In *Advances in Neural Information Processing Systems 23*, Lafferty JD, Williams CKI, Shawe-Taylor J, Zemel RS, and Culotta A, eds. (Curran Associates, Inc.), pp. 2164–2172.
- Sonkusare S, Breakspear M, and Guo C (2019). Naturalistic Stimuli in Neuroscience: Critically Acclaimed. *Trends in Cognitive Sciences* 23, 699–714. [PubMed: 31257145]
- Spiers HJ, and Maguire EA (2007). Decoding human brain activity during real-world experiences. *Trends in Cognitive Sciences* 11, 356–365. [PubMed: 17618161]
- Stangl M, Topalovic U, Hiller S, Inman C, Hasulak N, Rao V, Halpern C, Fried I, and Suthana N (2021). Boundary-anchored neural mechanisms of location-encoding for self and others. *Nature* 589, 420–425. [PubMed: 33361808]
- Stringer C, Pachitariu M, Steinmetz N, Reddy CB, Carandini M, and Harris KD (2019). Spontaneous behaviors drive multidimensional, brainwide activity. *Science* 364.
- Stupar-Rutenfrans S, Ketelaars LEH, and van Gisbergen MS (2017). Beat the Fear of Public Speaking: Mobile 360° Video Virtual Reality Exposure Training in Home Environment Reduces Public Speaking Anxiety. *Cyberpsychol Behav Soc Netw* 20, 624–633. [PubMed: 29039704]
- Sutton RS, and Barto AG (1998). *Introduction to Reinforcement Learning* (Cambridge, MA, USA: MIT Press).
- Topalovic U, Aghajani ZM, Villaroman D, Hiller S, Christov-Moore L, Wishard TJ, Stangl M, Hasulak NR, Inman CS, Fields TA, et al. (2020). Wireless Programmable Recording and Stimulation of Deep Brain Activity in Freely Moving Humans. *Neuron* 108, 322–334.e9. [PubMed: 32946744]
- Turner BM, Palestro JJ, Mileti S, and Forstmann BU (2019). Advances in techniques for imposing reciprocity in brain-behavior relations. *Neuroscience & Biobehavioral Reviews* 102, 327–336. [PubMed: 31128445]
- Vidaurre D, Smith SM, and Woolrich MW (2017). Brain network dynamics are hierarchically organized in time. *Proc. Natl. Acad. Sci. U.S.A* 114, 12827–12832. [PubMed: 29087305]
- Vidaurre D, Abeyesuriya R, Becker R, Quinn AJ, Alfaro-Almagro F, Smith SM, and Woolrich MW (2018). Discovering dynamic brain networks from big data in rest and task. *NeuroImage* 180, 646–656. [PubMed: 28669905]
- Vinyals O, Babuschkin I, Czarnecki WM, Mathieu M, Dudzik A, Chung J, Choi DH, Powell R, Ewalds T, Georgiev P, et al. (2019). Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature* 575, 350–354. [PubMed: 31666705]
- Walz N, Mühlberger A, and Pauli P (2016). A Human Open Field Test Reveals Thigmotaxis Related to Agoraphobic Fear. *Biol Psychiatry* 80, 390–397. [PubMed: 26876946]

- Warren WH (2018). Collective Motion in Human Crowds. *Curr Dir Psychol Sci* 27, 232–240. [PubMed: 30197472]
- Watkins CJCH, and Dayan P (1992). Q-learning. *Mach Learn* 8, 279–292.
- Wayne G, Hung C-C, Amos D, Mirza M, Ahuja A, Grabska-Barwinska A, Rae J, Mirowski P, Leibo JZ, Santoro A, et al. (2018). Unsupervised Predictive Memory in a Goal-Directed Agent. ArXiv:1803.10760 [Cs, Stat]
- Willems RM, Frank SL, Nijhof AD, Hagoort P, and van den Bosch A (2016). Prediction During Natural Language Comprehension. *Cereb Cortex* 26, 2506–2516. [PubMed: 25903464]
- Wiltchko AB, Johnson MJ, Iurilli G, Peterson RE, Katon JM, Pashkovski SL, Abreira VE, Adams RP, and Datta SR (2015). Mapping Sub-Second Structure in Mouse Behavior. *Neuron* 88, 1121–1135. [PubMed: 26687221]
- Wise T, and Dolan RJ (2020). Associations between aversive learning processes and transdiagnostic psychiatric symptoms in a general population sample. *Nature Communications* 11, 4179.
- Wise T, Michely J, Dayan P, and Dolan RJ (2019). A computational account of threat-relateds attentional bias. *PLOS Computational Biology* 15, e1007341. [PubMed: 31600187]
- Wise T, Liu Y, Chowdhury F, and Dolan RJ (2020). Model-based aversive learning in humans is supported by preferential task state reactivation. *BioRxiv* 2020.11.30.404491
- Woo CW, Krishnan A & Wager TD 2014. Cluster-extent based thresholding in fMRI analyses: pitfalls and recommendations. *Neuroimage*, 91, 412–9. [PubMed: 24412399]
- Wu CM, Schulz E, Speekenbrink M, Nelson JD, and Meder B (2018). Generalization guides human exploration in vast decision spaces. *Nature Human Behaviour* 2, 915.
- Yan X, Yang J, Sohn K, and Lee H (2016). Attribute2Image: Conditional Image Generation from Visual Attributes. ArXiv:1512.00570 [Cs]
- Ydenberg RC, and Dill LM (1986). The Economics of Fleeing from Predators. In *Advances in the Study of Behavior*, Rosenblatt JS, Beer C, Busnel M-C, and Slater PJB, eds. (Academic Press), pp. 229–249.
- Yilmaz Balban M, Cafaro E, Saue-Fletcher L, Washington MJ, Bijanzadeh M, Lee AM, Chang EF, and Huberman AD (2020). Human Responses to Visually Evoked Threat. *Curr Biol*
- Yoo SBM, Tu JC, Piantadosi ST, and Hayden BY (2020). The neural basis of predictive pursuit. *Nature Neuroscience* 23, 252–259. [PubMed: 31907436]

Box 1.**Implications for psychiatric populations and RDoC****Movement in psychiatric disorders.**

The emergence of machine learning techniques provides a new avenue from which to study a variety of complex movements that capture behaviors that have, until now, been difficult to measure. This could in turn present unique opportunities for identifying novel markers of mental health problems. Similar approaches have already shown promise in prior studies, demonstrating subtle behavioral signatures of mental health problems. For example, inspired by animal models, researchers have shown that thigmotaxis is higher in social phobia patients compared to healthy controls (Walz et al., 2016). Movement kinetics have also been used to detect altered movement patterns in autism (Cook et al., 2013) and may be used to detect prodromal markers of psychiatric conditions. These early studies demonstrate that taking a more naturalistic approach to the study of behavior, with rich indices of the movements human subjects make, can facilitate the identification of markers of disorder that could not otherwise be seen.

New transdiagnostic models of psychiatric disorders.

The benefits of computational ethology go beyond movement-based markers of psychiatric disorder, however. In recent years, computational psychiatry has begun to demonstrate how dysfunction in learning and decision-making processes can result in symptoms of mental health problems. However, these studies have relied on highly constrained, artificial tasks with limited behavioral measures. As detailed in other sections, computational ethology has the potential to bring new insights into learning and decision-making through richer and more natural behavioral measures. This deeper understanding of how humans learn about and act within their environment will naturally provide further targets for studies of how these processes go awry in psychiatric disorders. As an example, prior work has considered the importance of model-based planning in compulsive symptoms, showing both that individuals high in these traits have difficulty both in learning a model of the world and in using this learned model to guide behavior (Gillan et al., 2016; Sharp et al., 2020). However, tasks used to assess these processes rely on simplistic task structures, with only a few task states. Using methods from computational ethology could encourage the development of new models that are able to explain the relationship between model-based and model-free control in more complex and naturalistic environments, which would in turn provide new targets for studies investigating dysfunction in these processes and its association with symptom dimensions such as compulsivity. This may take inspiration from artificial intelligence, for example, where planning in complex environments has received a great deal of attention (Schrittwieser et al., 2020; Silver and Veness, 2010). As a result, new models of dysfunction could be developed that account for behavior that only emerges in these more naturalistic virtual environments.

Research Domain Criteria (RDoC).

RDoC is a framework within which to investigate and classify mental disorders, focusing on systems that span traditional diagnostic categories (Insel et al., 2010). The ability to detect and measure new behaviors will also advance the objectives of the RDoC, where one goal is to measure a full range of behaviors and link them to health and disorder. For example, while the RDoC negative valence systems matrix suggest that researchers should measure freezing, risk assessment, approach, avoidance and escape, there are few existing paradigms that can evoke these behaviors in the truest sense. It follows that gaining methods to measure these behaviors in human subjects will aid translation of animal models to humans and the identification of human behaviors that can also be evoked in animals, which could have benefits for the development of new pharmacological therapies. Drug development in psychiatry has largely stalled (Brady et al., 2019), and there have been numerous examples of drug candidates that were apparently efficacious in animals failing to show benefits in human trials, likely as a result of animal models of disease not truly representing the conditions they intend to (Grabb et al., 2016). Computational ethology, and its focus on naturalistic behaviors, will allow the proper measurement of key behaviors, highlighted in the RDoC, across both humans and animals and could in turn facilitate drug development.

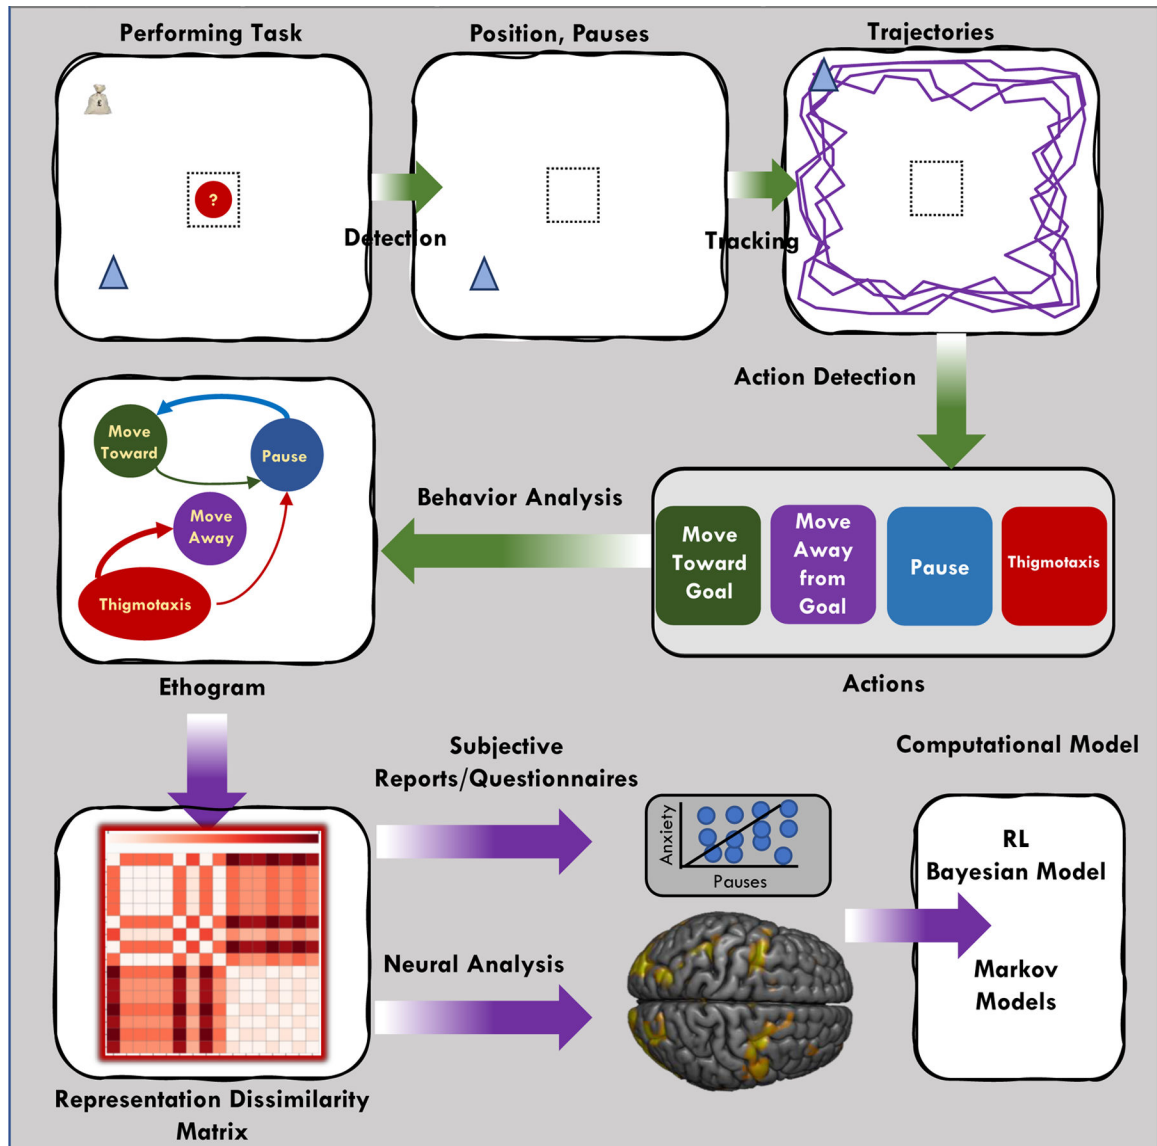


Figure 2. Steps in the automated analysis and modeling of natural behavior.

An example of how software can measure human behavior in 2D or 3D environment. This occurs in several stages including tracking the movements, action classification and behavior analysis (Anderson and Peroni, 2014). Starting from the top left: (A) the subject performs a task where they learn about safe patches and where rewards of high or low value will appear. Such a task should result in place preference or aversion. The software occurs in several stages including the (B)detection (C) tracking of the movements, (D) action classification and (E) behavior analysis (Anderson and Peroni, 2014; Dankert et al., 2009). This will result in an (F) ethogram that will illustrate the different behaviors of pausing, thigmotaxis, movement away and so forth. (G) The behaviors can then be used for correlation with neural activity or subjective reports and questionnaire data. Finally, these data can be used to inform or create (H) computational accounts of the behavior.

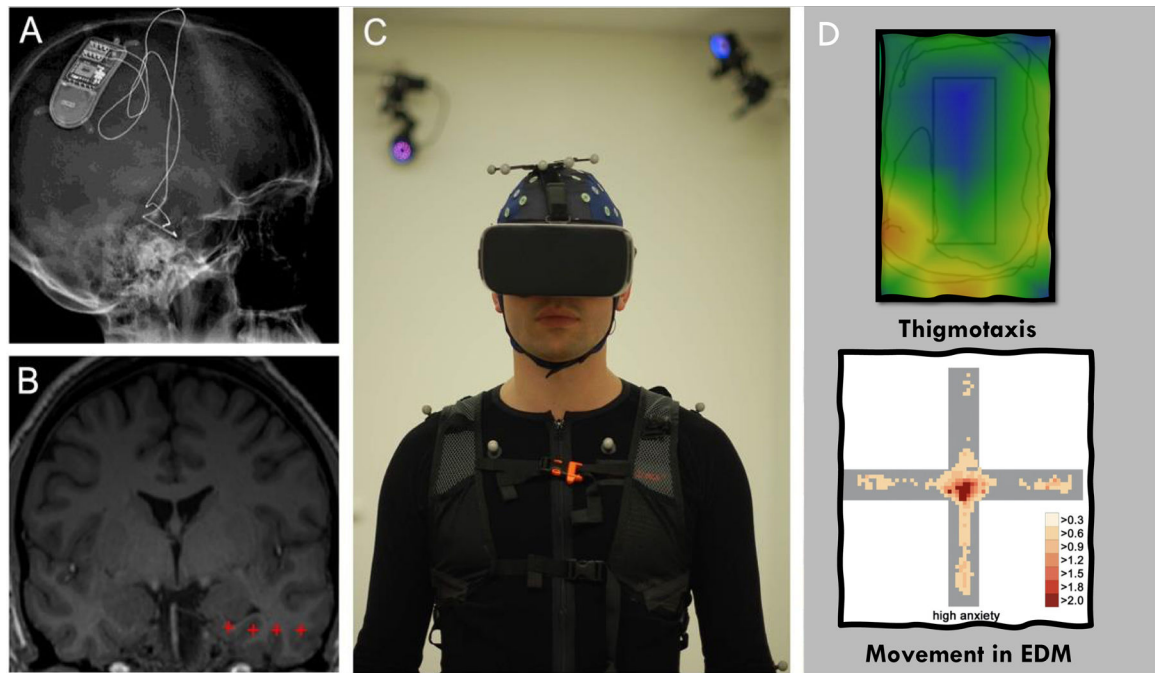


Fig. 3. Example of how VR can be used to create virtual ecologies that measure naturalistic behaviors in participants with chronically implanted electrodes.

A) X-ray image and B) MRI of an example participant with a chronically implanted electrode with four contacts (red crosses) in the temporal lobe for iEEG recording, during which ambulatory VR and full-body motion capture (C) can be integrated (Stangl et al., 2021). Example of how VR can be used to create virtual ecologies that measure behaviors similar to those observed in rodents (e.g., thigmotaxis ((Walz et al., 2016), and movement in the EDM (Biedermann et al., 2017).