Measuring and modeling the motor system with machine learning

Sébastien B. Hausmann[1,†], Alessandro Marin Vargas[1,†], Alexander Mathis[1,*], & Mackenzie W. Mathis[1,*]

[1]EPFL, Swiss Federal Institute of Technology, Lausanne, Switzerland | †co-first *co-senior

arXiv:2103.11775v1 [q-bio.QM] 22 Mar 2021

The utility of machine learning in understanding the motor system is promising a revolution in how to collect, measure, and analyze data. The field of movement science already elegantly incorporates theory and engineering principles to guide experimental work, and in this review we discuss the growing use of machine learning: from pose estimation, kinematic analyses, dimensionality reduction, and closed-loop feedback, to its use in understanding neural correlates and untangling sensorimotor systems. We also give our perspective on new avenues where markerless motion capture combined with biomechanical modeling and neural networks could be a new platform for hypothesis-driven research.

Highlights:

1. Deep learning-based tools allow for robust automation of movement capture and analysis
2. New approaches to modeling the sensorimotor system enable new hypotheses to be generated
3. These tools are poised to transform our ability to study the motor system

Introduction

Investigations of the motor system have greatly benefited from theory and technology. From neurons to muscles and whole-body kinematics, the study of one of the most complicated biological systems has a rich history of innovation. In their quest to understand the visual system, David Marr and H. Keith Nishihara described how static shape powerfully conveys meaning (information), and gave two examples of how "skeletonized" data and 3D shapes, such as a series of cylinders can easily represent animals and humans (Figure 1A,B; 1). These series of shapes evolving across time constitute actions, motion, and what we perceive as the primary output of the motor system. Now with deep learning, the ability to measure these postures has been massively facilitated. But the measurement of movement is not the sole realm where machine learning has increased our ability to understand the motor system. Modeling the motor system is also leveraging new machine learning tools. This opens up new avenues where neural, behavioral, and muscle-activity data can be measured, manipulated, and modeled with ever-increasing precision and scale.

Measuring motor behavior

Motor behavior remains the ultimate readout of internal intentions. It emerges from a complex hierarchy of motor control circuits in both the brain and spinal cord, with the ultimate output being muscles. Where these motor commands arise and how they are shaped and adapted over time still remains poorly understood. For instance, locomotion is produced by central pattern generators in the spinal cord and is influenced by goal-directed descending command neurons from the brainstem and higher motor areas to adopt specific gaits and speeds—the more complex and goal-directed the movement, the larger number of neural circuits recruited (5, 6). Of course, tremendous advances have already been made in untangling the functions of motor circuits since the times of Sherrington (7), but new advances in machine learning for modeling and measuring the motor system will allow for better automation and precision, and new approaches to modeling, as we highlight below.

In the past few years modern deep learning techniques have enabled scientists to closely investigate behavioral changes in an increasingly high-throughput and accurate fashion. Often eliminating hours of subjective and time-consuming human annotation of posture, "pose estimation"—the geometric configuration of keypoints tracked across time—has become an attractive workhorse within the neuroscientific community (8–11). While there are still computer vision challenges in both human and animal pose estimation, these new tools perform accurately within the domain they are trained on (12, 13). Naturally, the question arises: now that pose estimation is accelerating the pace and ease at which behavioral data can be collected, how can these extracted poses help explain the underlying neural computations in the brain?

Meaningful data from pose estimation.

Advances in pose estimation have opened new research possibilities. Software packages that tackle the needs of animal pose estimation with deep neural networks (DNNs) include DeepLabCut (4, 14), LEAP (15), DeepBehavior (16), DeepPoseKit (17), DeepFly3D (18), and DeepGraphPose (19). These packages provide customized tracking for the detailed study of behavior, and come with their own pitfalls and perspectives (20). Pose estimation was reviewed elsewhere (9–11, 20), which is why we focus on how such data can be used to understand the motor system.

Several groups have delved into the challenge of deriving additional metrics from deep learning-based pose estimation data. There are (at least) two common paths: (1) derive kinematic variables, (2) derive semantically-meaningful behavioral actions. Both have unique challenges and potential uses.

* alexander.mathis@epfl.ch, mackenzie@post.harvard.edu

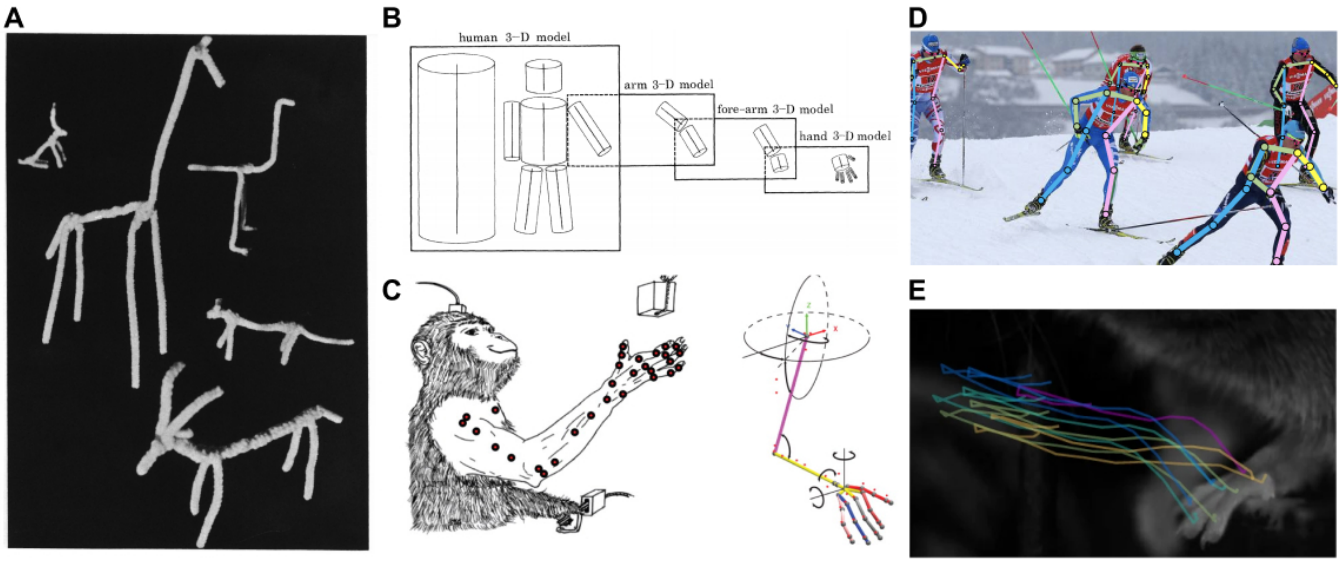

Figure. 1. Measuring movement: Keypoint-based "skeletons" provide a reduced representation of posture. From Marr (A, B, adapted from 1) to now, classical and modern computer vision has greatly impacted our ability to measure movement (C adapted from 2, D adapted from 3, and E adapted from 4).

Kinematic analysis.

Whether marker-based or marker-free, tracked keypoints can be viewed as the raw building blocks of joint coordinates and subsequent measurements. Namely, key points can form the basis of movement analysis. Kinematics—the analysis of motion, and in biomechanics typically the study of position, velocity, and acceleration of joints—is essential to describe the motion of systems.

How has kinematic data been used to understand the motor system? In elegant work, Vargas-Irwin and colleagues demonstrated how highly detailed pose estimation data can be used to uncover neural representations in motor cortex during skilled reaching. Using a marker-based approach, they showed that a small number of primary motor cortex (M1) neurons can be used to accurately decode both proximal and distal joint locations during reaching and grasping (2) (Figure 1C). Others have powerfully used marker-based tracking to quantify recovery in spinal cord injuries in mice, rats and macaques (21–23).

Markerless pose estimation paired with kinematic analysis is now being used for a broad range of applications. Human pose estimation tools, such as state-of-the-art (on COCO) HRNet (3) (Figure 1D) or DeepLabCut (Figure 1E), have been used in applications such as sports biomechanics (24, 25), locomotion (Figure 2) and clinical trials[1]. For example, Williams et al. recently showed that finger tap bradykinesia could be objectively measured in people with Parkinson's disease (26). They used deep learning-based pose estimation methods and smartphone videos of finger tapping kinematics (speed, amplitude and rhythm) and found correlations with clinical ratings made by multiple neurologists. With a broader adoption of markerless approaches in both the biomechanical and neuroscience community, we foresee a growing use of machine learning for developing clinically-relevant biomarkers, and automating the process of clinical scoring of disease states (i.e., automating scoring Parkinson's disease).

Reducing dimensionality to derive actions.

Analyzing and interpreting behavior can be challenging as data are complex and high-dimensional. While pose estimation already reduces the dimensionality of the problem significantly, there are other processing steps that can be used to transform video into "behavior actions". Specifically, dimensionality reduction methods can transform the data into a low-dimensional space, enabling a better understanding and/or visualization of the initial data.

Dimensionality reduction tools (e.g., PCA, t-SNE (31), UMAP (32)) are commonly used methods to cluster data in an unsupervised manner. Principal Component Analysis (PCA) aims to find the components that maximize the variance in the data and the principal components rely on orthogonal linear transformations, and has been important for estimating the dimensionality of movements, for example (2, 33). t-distributed Stochastic Neighbor Embedding (t-SNE) is suitable for visualizing non-linear data and seeks to project data into a lower dimensional space such that the clustering in the high dimensional space is preserved. However, it *typically* does not preserve the global data structure (i.e., only within-cluster distances are meaningful; but cf. 34). In conjunction with an impressive system to track hunting zebrafish, t-SNE was used to visualize the hunting states (Figure 2). Lastly, Uniform Manifold Approximation and Projection for Dimension Reduction (UMAP) typically preserves the data's global structure (cf. 35). Many tools allow for seamless passing of pose estimation output to perform such clustering, such as MotionMapper (36) and B-SOiD (37).

The utility of unsupervised clustering for measuring behav-

[1] https://clinicaltrials.gov/ct2/show/NCT04074772

ior is well illustrated by Bala et al. 38. They explored motor behavior in freely moving macaques in large unconstrained environments in the laboratory setting. In combination with 62 cameras—a remarkable feat on its own—OpenMonkeyStudio uses 13 body landmarks and trains a generalizable view-invariant 2D pose detector to then compute 3D postures via triangulation. Coherent clusters extracted with UMAP correlated with semantic actions such as sitting, climbing, and climbing upside down in monkeys (38). Importantly, in this case, a 3D pose estimation compared to 2D pose estimation yielded more meaningful clusters. Together with the 3D location of a second macaque, social interactions were derived from co-occurrence of actions.

Dimensionality reduction can also be applied to kinematic features (such as joint angles). For example, DeAngelis et al. used markerless pose estimation to extract gait features from freely moving hexapods (39). Then, UMAP was applied to generate a low-dimensional embedding of the tracked limbs while preserving the local and global topology of the high-dimensional data. UMAP revealed a vase-like manifold parametrized by the coordination patterns of the limbs. In conjunction with modulation by means of optogenetic or visual perturbations, different gaits and kinematic modalities were evoked and directly interpretable in the UMAP embedding. In a similar approach, using the clustering algorithm *clusterdv* (40), different types of swimming bouts were iden-

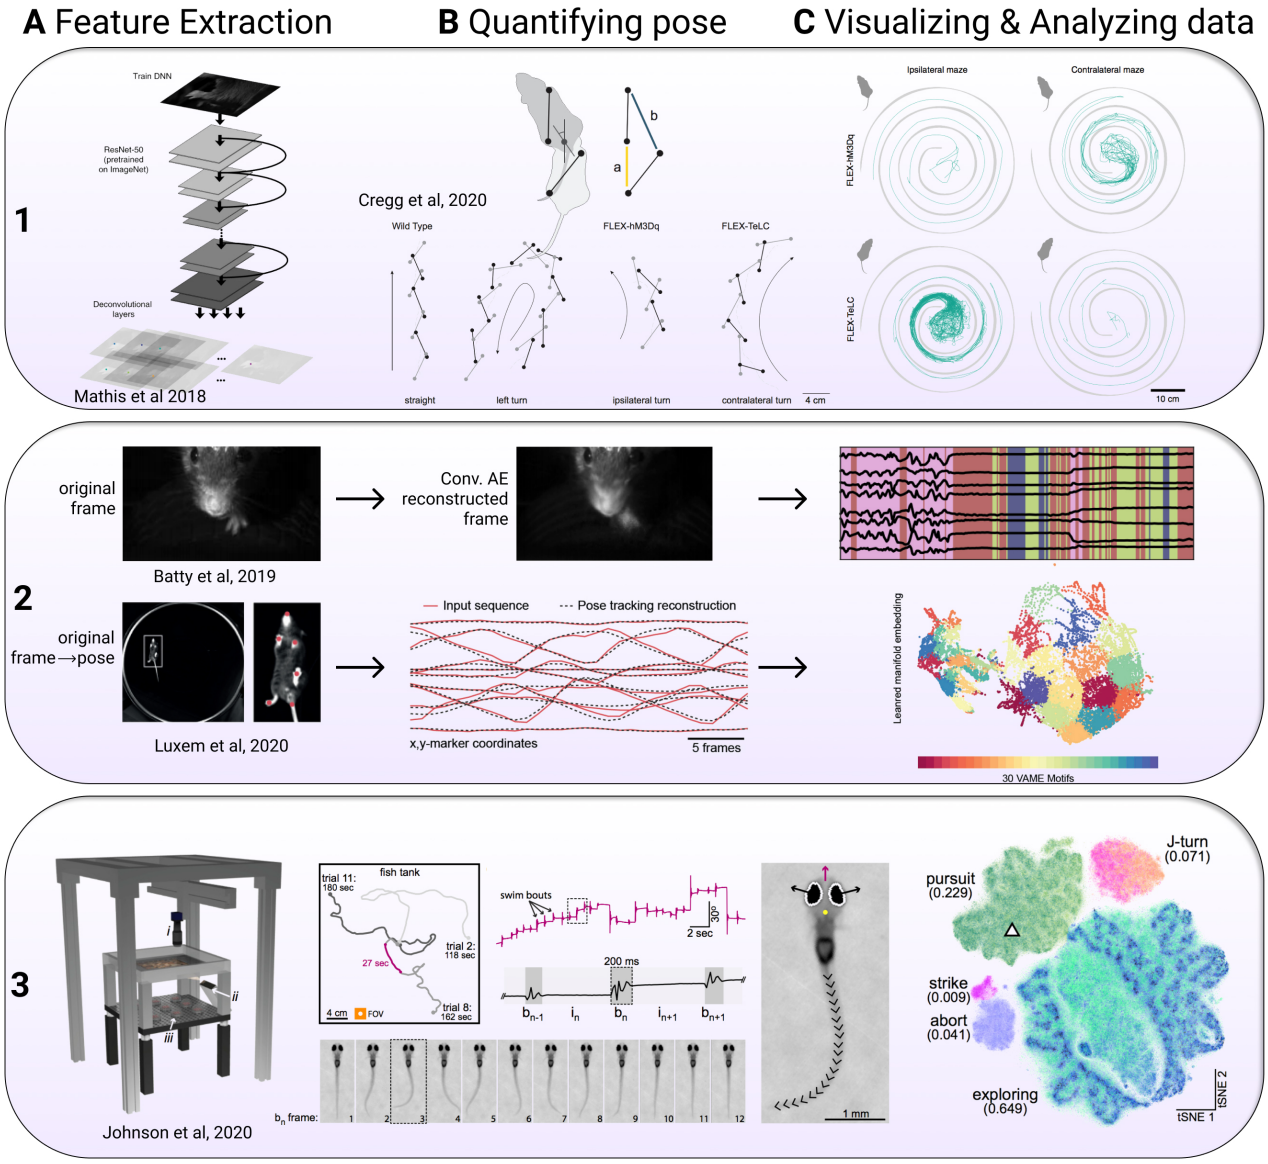

Figure. 2. From Video to Behavioral Metrics: How deep learning tools are shaping new studies in neural control of locomotion, analyzing natural behavior, and kinematic studies. The general workflow consists in (A) extracting features (pose estimation), (B) quantifying pose and (C) visualizing and analyzing data. (1) Limb tracking to analyze gait patterns in mice whose locomotion has been altered by brainstem optical stimulation (adapted from 4, 27). (2) Autoencoders: Top row: The use of a convolutional autoencoder from video to inferred latents and states, adopted from 28. Bottom row: Architecture of VAME (adapted from 29) constituted by an encoder (input is key points), which learns a representation of the hidden states. The original sequence can be reconstructed by a decoder or the evolution of the time series can be predicted. UMAP embedding of the pose estimation. (3) Behavioral data acquisition enabling precise tracking of zebrafish larvae and behavioral clustering of swimming bouts (adapted from 30).

tified from zebrafish larvae movements (41). Similar methods enabled precise visualization of hunting strategies that were dependent on visual cues (42).

Autoencoders are powerful tools for nonlinear dimensionality reduction via generative models (43). They have found wide use in sequence modeling for language, text, and music (44, 45). They can also be used for unsupervised discovery of behavioral motifs by using manually selected features (i.e., the keypoints defined for pose estimation) as input, or by inputting video (Figure 2. For instance, BehaveNet uses video directly to model the latent dynamics (28). Variational Animal Motion Embedding (29) (Figure 2) used keypoints as the input to a recurrent variational auto encoder to learn a complex distribution of the data. Variational autoencoder stochastic neighbor embedding (VAE-SNE) combines a VAE with t-SNE to compress high dimensionality data and automatically learn a distribution of clusters within the data (46).

In addition to unsupervised analysis, supervised methods are powerful ways to use the outputs of pose estimation to derive semantically meaningful actions (47–50). For instance, the Mouse Action Recognition System (MARS), a pipeline for pose estimation and behavior quantification in pairs of freely behaving (differently colored) mice, works in conjunction with BENTO, which allows users to annotate and analyze behavior states, pose estimates, and neural recording data simultaneously through a graphical interface (50). This framework enabled the authors to use supervised machine learning to identify a subset of 28 neurons whose activity was modulated by mounting behavior. These methods build on open source tools, which provided many supervised and unsupervised learning tools in a highly customize way (47).

Another approach is to use video directly, instead of applying pose estimation first (28, 36, 51). Here, the pixels themselves can be used for action recognition, which can be highly useful when the behaviors of interest do not involve kinematics, such as blushing in humans, or freezing in mice. In computer science the rise of deep learning has gone hand-in-hand with the development of larger datasets. For example, pre-training models on the Kinetic Human Action Video dataset drastically improves action recognition performance (52). Others pushed video recognition further by using a *multi-fiber* architecture, where sparse connections are introduced inside each residual block, thereby reducing computations (53). Several groups have leveraged this approach for facial expression (54, 55), pharmacological behavior-modulation (56) or for measuring posture and behavior representations (57, 58).

Closed-loop feedback based on behavior.

Closed-loop feedback based on behavioral measurements can be informative for causal testing of learning algorithms (such as reinforcement learning) and for probing the causal role of neural circuits in behavior (59, 60). Recent efforts to translate offline pose estimation and analysis to real time have enabled new systems, such as in EthoLoop (60) and DeepLabCut-live! (61). EthoLoop is a multi-camera, closed-loop tracking system using a two-step process: a marker-based camera tracking system, followed by DeepLabCut-based pose estimation analysis. This system is capable of providing close-up views and analyses of the ethology of tracked freely moving primates. The closed-loop aspect enabled real-time wireless neuronal recordings and optical stimulation of individuals striking specific poses. We expect that real-time (and predictive) low-latency pose estimation for closed-loop systems will certainly play a crucial role in the years to come (61–63). These tools are bound to include more customized behavior-dependent real-time feedback options. While the delays can be minimal for such computations, time-delayed (hardware) systems are not real-time controllable. Thus, to instantly provide feedback we added a forward prediction mode to DeepLabCut-Live! (61), which we believe will be crucial as these tools grow in complexity. Overall, the above discussed methods are highly effective in disentangling behavioral data, finding patterns, and capturing actions within behavioral data.

Neural correlates of behavior.

At the heart of neuroscience research is the goal to causally understand how neurons (from synapses to networks) relate to behavior. With the rise of new tools to measure behavior, there is a homecoming to the quest of relating movement to neural activity. As described above, both kinematic features or lower dimensional embeddings of behavior can be used to regress against neural activity, or using real-time feedback, to causally probe their relationship to actions.

In recent years, many groups have utilized such tools to uncover new knowledge of the motor system. It remains debated what the role of motor cortex is, yet new tools are enabling careful and precise studies of the system in goal-directed actions. For example, Ebina et al. used markerless tracking and optogenetics to show that in marmosets stimulation with varying spatial or temporal patterns in motor cortex could elicit simple or even direction-specific forelimb movements (64). Sauerbrei et al. also used markerless tracking to show that motor cortex inactivation halted movements, but this was a result of disrupted inputs (from thalamus), thus revealing that multiple interacting brain regions were responsible for dexterous limb control (65). Moreover, others have used these tools to show that brainstem neurons highly correlate and causally drive locomotor behaviors (27, 66).

Furthermore, it is becoming increasingly recognized that brain-wide (or minimally, cortex-wide) neural correlates of movement are ubiquitous in animals performing both spontaneous or goal-directed actions (67–69). Stringer et al. (67) showed spontaneous movements constituted much of the neural variance ("noise") in visual cortex, and Musall, Kaufmann et al. (68) report similar findings across many brain regions. It was also previously shown that neuronal activity in the posterior parietal cortex and the pre-motor cortex (M2) of rats accurately correlate with animal's head posture (70), and of course even in sensory areas such as visual cortex encode movement (67, 71, 72).

Relating neural activity to not only kinematic or behavioral features, but to muscle output is also highly important for correlating neurons to movement. In a recent study, human neonatal motoneuron activity was characterized using non-invasive neural interface and joint tracking. Using markerless pose estimation, they found that fast leg movements in neonates are mediated by high motoneuron synchronization, and not simply due to an increase in discharge rate as previously observed in adults (73, 74). Another example of combining EMG (electromyography) and motion capture, Herent et al. used peripheral recordings via a diaphragmatic EMG to reveal an absence of synchronization (e.g., temporal correlation) of breaths to strides in mice at various displacement speeds during locomotion (75). This rich data is crucial for understanding the neural code, and it is shaping efforts to model the system at an even finer resolution.

Being able to not just capture movement, but model it has a rich history in movement neuroscience (76–78). In the next sections we will discuss how machine learning has also influenced modeling the motor system.

Neural networks as sensorimotor system models

The brain excels at orchestrating adaptive animal behavior, striving for robustness across environments (7, 78). Thereby the brain, hidden in the skull, takes advantage of multiple sensory streams in order to act as a closed-loop controller of the body. How can we elucidate the function of this complex system? The solution to this problem is not trivial due to multiple challenges such as high-dimensionality, redundancy, noise, uncertainty, non-linearity and non-stationarity of the system (78). Since DNNs excel at learning complex input-output mappings (79, 80), they are well positioned for modeling motor, sensory and also sensorimotor circuits. Moreover, unlike in the biological brain, DNNs are fully observable such that one can easily "record" from all the neurons in the system and measure the connectome. Therefore, in the next sections we discuss how modeling the motor, sensory and combine sensorimotor systems may lead to new principles of motor control.

Modeling the motor system.

How do neural circuits produce adaptive behavior, and how can deep learning help model this system? Several groups have modeled the motor system with neural networks to investigate how, for example, task cues can be transformed into rich temporal sequential patterns that are necessary for creating behavior (81, 82). In particular, recurrent neural networks (RNNs), whose activity is dependent on their past activity (memory), have been used to study the motor system, as they can produce complex dynamics (79, 80).

In a highly influential study, Sussillo and colleagues showed that RNNs can learn to reproduce complex patterns of muscle activities recorded during a primate reaching task. By modifying the network's characteristics, such as forcing the dynamics to be smooth, the natural dynamics that emerged from the network closely resembled the one observed in the primates' primary motor cortex (M1) (83). In a similar way, RNNs can also be used to understand why a brain area shows a specific feature. For instance, M1 has low-tangled population dynamics when primates perform a cycling movement task using a hand-pedal, unlike muscle activity and sensory feedback (84). Building a network model which is trained on the same task, not only can replicate the same dynamics but it also unveils noise robustness as a possible underlying principle of the observed dynamics. Moreover, RNNs can be used to test various hypotheses about how the brain drives adaptive behavior by comparing neural or behavioral activity to the "neural" units. For instance, it has been investigated how prior beliefs are integrated into the neural dynamics of the network (85), how temporal flexibility is connected to a network's nonlinearities (86), and how robust trajectory shifting allows the translation between sequential categorical decisions (87).

Importantly for motor control, RNNs have also been used to study how sequential, independent movements, or tasks, are generated (88). Here, the authors also tackled the problem of cross task interference, which the RNN overcame by utilizing orthogonal subspaces (88). In related work, a novel learning rule that aimed to conserve network dynamics within subspaces (defined by activity of previously learned tasks) allowed for more robust learning of multiple tasks with a RNN (89).

RNNs are also capable of re-producing complex spatio-temporal behavior, such as speech (90, 91). This was achieved using a two-stage decoder with bidirectional, long short-term memory networks. First, the articulatory features were decoded by learning a mapping between sequences of neural activity (high-gamma amplitude envelope and low frequency component from EcoG recordings) and 33 articulatory kinematic features. Second, the kinematic representation is mapped to 32 acoustic features. Importantly, because subjects share the kinematic representation, the first stage of the decoder could be transferred to different participants, which requires less calibration data (91).

Modeling sensory systems.

Sensory inputs, which are delayed and live in a different coordinates to the motor system, need to be integrated for adaptive motor control (77, 78). Deep learning based models of sensory systems have several advantages: responses for arbitrary stimuli can be computed, their parts can be mapped to brain regions, and they have high predictability (93–95).

In general, there are two main approaches used for generating models of sensory systems: (i) data-driven models and (ii) task-driven models. Data-driven approaches use encoding models that are trained to predict neural activity from stimuli. As such, using non-deep-learning-based encoding models is a classical approach (96) where neural responses are approximated using tuning curves (97). Since stimulus features can be highly complex, DNNs can be especially useful in learning the stimulus-response mappings (94, 95, 98–

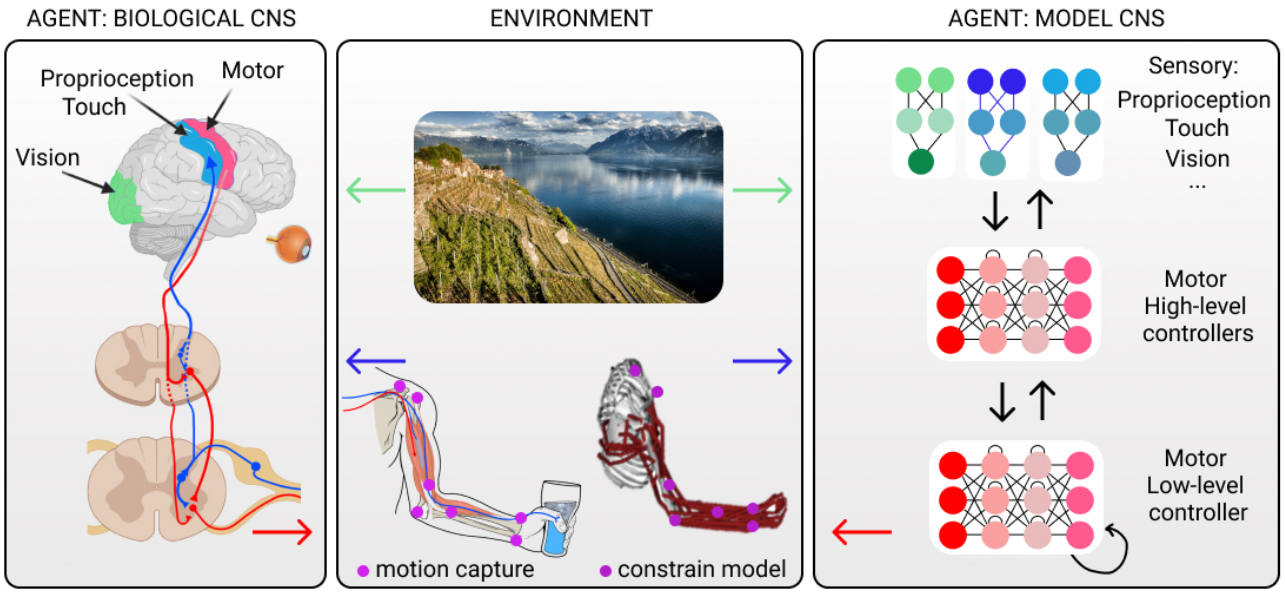

Figure. 3. Modeling the sensorimotor system. The biological agent's central nervous system (CNS, left panel) receives sensory input such as vision, proprioception, and touch from the environment (middle panel). The sensorimotor system can be modeled using hierarchical, artificial neural networks (right panel), such as feed-forward networks (blue and green), and recurrent networks (red) to represent the sensory and motor systems (respectively). The biological agent and environment provide crucial model constraints (derived from fine-grained behavioral measurements, biochemical and anatomical constraints, etc.). For instance, motion capture and modeling of the plant (body) in the environment (i.e., here the arm) can be used to constrain motor outputs. The model CNS can be used to make predictions about the biological system, and used to explain biological data. Anatomical drawings on the left created with BioRender.com; OpenSim arm adapted from (92).

101). Indeed, DNNs can outperform standard methods, such as the generalized linear model (GLM), in predicting the neural activity of the somatosensory and motor cortex (98). Yet, to fit complex deep learning models, one typically needs a lot of data, a limitation that can be overcome by the task-driven approach.

Task-driven modeling builds on the hypothesis that sensory systems have evolved to solve relevant ecological tasks (102, 103): if artificial systems are trained to solve the same complex tasks that animals face, they might converge towards representations observed in biological systems (94, 95, 104). Therefore, one can (potentially) learn large-scale models for limited neural datasets by taking advantage of transfer learning (e.g., train on ImageNet (105) to obtain better models of the visual (ventral) pathway). Second, the choice of tasks allows researchers to test diverse hypotheses such as the highly successful hypothesis that the ventral pathway in primates is optimized to solve object recognition (93, 94, 106). Specifically, hierarchical DNNs trained to solve an object recognition task can explain a large fraction of variance in neural data of different brain areas along the visual pathway (106) and auditory pathway (107). Crucially, for audition and vision, large scale datasets of relevant stimuli are readily available. How can this work be expanded to senses of importance to the motor system, such as touch and proprioception, where delivering relevant touch-only or proprioception-only information is more challenging (if not impossible)?

One possible way to overcome this issue consists of simulating touch and proprioceptive inputs using biophysical models. For instance, a physically-realistic model of a mouse's whisker array has been used to develop a synthetic dataset of whisker (touch) sweeps across different 3D objects (108). A muscle-spindle firing rate dataset has been generated from 3D movements based on a musculoskeletal model of a human arm (92). In this way, DNNs trained to perform object or character recognition suggested putative ways spatial and temporal information can be integrated across the hierarchy (108) and that network's architecture might play a main role in shaping its kinematic tuning properties (92). Both of these studies propose perception-based tasks, which provide baseline models, yet we envision the task space, complexity of the biophysical models, and the plant modelled will increase in future studies.

Modeling the sensorimotor system.

Although modeling the motor or sensory systems alone is of great importance, a crucial body of work stems from combining models of sensory, motor, and task-relevant signals. Deep reinforcement learning (DRL) is a powerful policy optimization method to train models of sensorimotor control (109).

For instance, DRL was used to solve a navigation task in zebrafish, and representations in units of the network resembled those observed in the brain (110). Not only did the optimized network have units that correlated with temperature and behavioral states, but it also predicted an additional functional class of neurons which was not previously identified with calcium imaging alone (110). DNNs trained to achieve chemotaxis behavior of *c. elegans* with DRL provided insights into neuropeptide modulation and circuits (111). In an artificial

agent trained to perform vector-based navigation, grid cells emerged as well as coding properties (112) that had been observed in rodents (113), and predicted from theory (114). Motor control has been studied using a virtual rodent trained with DRL to solve a few behavioral tasks. A closer look into the learned neural representations delineates a task-invariant and a task-specific class, which belong to the low- and high-level controller, respectively (115).

Moreover, DRL has been utilized to learn control policies for complex movements, such as locomotion in challenging environments (116–118) or complex dexterous hand manipulation (119). Eventually, the policy learned in simulation environments can be transferred to a real-world scenario achieving rather remarkable robustness in terrains not encountered during training (120) and evolving human-like grasping and object interaction behaviors without relying on any human demonstrations (121, 122). Lastly, real-world sensory feedback and measurements during natural tasks can also be used to constrain models and define goals. For example, reference motion data can be used to train musculoskeletal model using DRL to reproduce human locomotion behaviors (123, 124).

An alternative approach to modeling the sensorimotor system in neural networks is via engineering or theoretical principles. Neural network models have an increased capacity to model highly complex integrated systems, as evidenced by advances, like Spaun, a spiking, multi-million-neuron-network model that can perform many cognitive tasks (125–127). Additionally, optimal feedback control (OFC) theory is a powerful normative principle for deriving models, that predicts many behavioral phenomena (76, 77). OFC translated into DNNs accurately predicted neural coding properties in the motor cortex with biomechanical arm models and peripheral feedback (128). Recent experimental and modeling work has also begun to link different cortical regions to their functional roles in an OFC theory framework (129, 130). We believe that future work will combine neural recordings and perturbations to continue to test hypotheses generated by DNN-based OFC, and other DNN-based system-wide models, which are constrained by rich behavior (Figure 3).

Towards such hybrid models, Michaels et al. have developed an exceptional model which leverages visual feedback to produce grasping movements. This model combines a CNN to extract visual features and three RNNs (modules) to control a musculosketal human arm. Amongst different architectures, the one with sparse inter-module connectivity was the best in explaining the brain-related neural activity thereby revealing possible anatomical principles of cortical circuits. Moreover, behavioral deficits, observed in previous lesion studies of the corresponding cortical areas, could be predicted by silencing specific modules. Interestingly, slight deficits occurred in regularized networks (i.e., penalty on high firing rates) when inputs from the intermediary module to the output module were lesioned, whereas behavior was completely disrupted in non-regularized networks. This observation suggests that the minimization of the firing rate could be a potential organizational principle of cortical circuits and that M1 might autonomously generate movements (131). However, this contrasts recent evidence that continuous input from the thalamus (in mice) is necessary to perform movements (65).

DNNs can also be trained in a self-supervised way based on rich data sets (132), which is a highly attractive platform for studying sensorimotor learning. Sullivan et al. showed that DNNs can learn high-level visual representations when trained on the same naturalistic visual inputs that babies receive using head-mounted cameras (133). These learned representations are invariant to natural transformations and support generalization to unseen categories with few labelled examples (132). An important future direction for studying the motor system will be to not only focus on representational similarity, but also comparing the learning rules used by both biological systems and machine models.

Lastly, DNNs for sensorimotor control are also popular in robotics. Not only can robots serve as great testing grounds for control policies (134, 135), but they can also reveal limitations when going from "simulation to reality", such as challenges related to robustness (136, 137).

Outlook

In the past few years, neuroscience has tremendously benefited from advances in machine learning. Behavioral analysis got more accurate while also being much less time consuming. This has already revealed novel aspects of behavior, but we are just at the dawn of these developments. Furthermore, advances in deep learning shaped the way biological systems can be modeled. We believe that in the future these two aspects will become increasingly intertwined. Namely, the unreasonable effectiveness of data suggests that powerful models—which further approximate the neural code—can be trained with large-scale behavioral measurements that are now possible. Concretely, with new tools to measure behavior one can constrain biologically plausible agents (such as OpenSim (138) or mujoco models (139)) to generate the behavior in an artificial setting. This simulation can then be used to generate data that might not otherwise be available: i.e., muscle activity from the whole arm or body in parallel with a visual input. These data could be used to develop hierarchical DNN models of the system (as we illustrate in Figure 3). These artificial network models, constrained with complex tasks, will provide researchers with sophisticated tools for testing hypotheses and gaining insight about the mechanisms the brain might use to generate behavior.

Acknowledgments:

The authors declare no conflicts of interest. We thank members of the Mathis Lab, Mathis Group and Travis DeWolf for comments. MWM is the Bertarelli Foundation Chair of Integrative Neuroscience.

References

1. David Marr and Herbert Keith Nishihara. Representation and recognition of the spatial organization of three-

dimensional shapes. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, 200(1140):269–294, 1978.
2. Carlos E Vargas-Irwin, Gregory Shakhnarovich, Payman Yadollahpour, John MK Mislow, Michael J Black, and John P Donoghue. Decoding complete reach and grasp actions from local primary motor cortex populations. *Journal of neuroscience*, 30(29):9659–9669, 2010.
3. Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggang Wang, et al. Deep high-resolution representation learning for visual recognition. *IEEE transactions on pattern analysis and machine intelligence*, 2020.
4. Alexander Mathis, Pranav Mamidanna, Kevin M Cury, Taiga Abe, Venkatesh N Murthy, Mackenzie Weygandt Mathis, and Matthias Bethge. Deeplabcut: markerless pose estimation of user-defined body parts with deep learning. *Nature neuroscience*, 21:1281–1289, 2018.
5. Maria S. Esposito, Pablo Capelli, and Silvia Arber. Brainstem nucleus mdv mediates skilled forelimb motor tasks. *Nature*, 508:351–356, 2014.
6. Julien Bouvier, Vittorio Caggiano, Roberto Leiras, Vanessa Caldeira, Carmelo Bellardita, Kira Balueva, Andrea Fuchs, and Ole Kiehn. Descending command neurons in the brainstem that halt locomotion. *Cell*, 163(5):1191–1203, 2015.
7. Charles Sherrington. *The integrative action of the nervous system*. Cambridge University Press Archive, 1952.
8. Sandeep Robert Datta, David J Anderson, Kristin Branson, Pietro Perona, and Andrew Leifer. Computational neuroethology: a call to action. *Neuron*, 104(1):11–24, 2019.
9. Mackenzie W. Mathis and Alexander Mathis. Deep learning tools for the measurement of animal behavior in neuroscience. *Current Opinion in Neurobiology*, 60:1–11, 2020.
10. Talmo D Pereira, Joshua W Shaevitz, and Mala Murthy. Quantifying behavior to understand the brain. *Nature neuroscience*, pages 1–13, 2020.
11. Lukas von Ziegler, Oliver Sturman, and Johannes Bohacek. Big behavior: challenges and opportunities in a new era of deep behavior profiling. *Neuropsychopharmacology*, 46(1):33–44, 2021.
12. Alexander Mathis, Thomas Biasi, Steffen Schneider, Mert Yuksekgonul, Byron Rogers, Matthias Bethge, and Mackenzie W Mathis. Pretraining boosts out-of-domain robustness for pose estimation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1859–1868, 2021.
13. Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanas Phillips, Sara Beery, et al. Wilds: A benchmark of in-the-wild distribution shifts. *arXiv preprint arXiv:2012.07421*, 2020.
14. Tanmay Nath, Alexander Mathis, An Chi Chen, Amir Patel, Matthias Bethge, and Mackenzie W Mathis. Using deeplabcut for 3d markerless pose estimation across species and behaviors. *Nature protocols*, 14:2152–2176, 2019.
15. Talmo D Pereira, Diego E Aldarondo, Lindsay Willmore, Mikhail Kislin, Samuel S-H Wang, Mala Murthy, and Joshua W Shaevitz. Fast animal pose estimation using deep neural networks. *Nature methods*, 16(1):117, 2019.
16. Ahmet Arac, Pingping Zhao, Bruce H Dobkin, S Thomas Carmichael, and Peyman Golshani. Deepbehavior: A deep learning toolbox for automated analysis of animal and human behavior imaging data. *Frontiers in systems neuroscience*, 13:20, 2019.
17. Jacob M Graving, Daniel Chae, Hemal Naik, Liang Li, Benjamin Koger, Blair R Costelloe, and Iain D Couzin. Deepposekit, a software toolkit for fast and robust animal pose estimation using deep learning. *eLife*, 8:e47994, oct 2019. ISSN 2050-084X. doi: 10.7554/eLife.47994.
18. Semih Günel, Helge Rhodin, Daniel Morales, João Campagnolo, Pavan Ramdya, and Pascal Fua. Deepfly3d, a deep learning-based approach for 3d limb and appendage tracking in tethered, adult drosophila. *Elife*, 8:e48571, 2019.
19. Anqi Wu, Estefany Kelly Buchanan, Matthew Whiteway, Michael Schartner, Guido Meijer, Jean-Paul Noel, Erica Rodriguez, Claire Everett, Amy Norovich, Evan Schaffer, et al. Deep graph pose: a semi-supervised deep graphical model for improved animal pose tracking. *Advances in Neural Information Processing Systems*, 33, 2020.
20. Alexander Mathis, Steffen Schneider, Jessy Lauer, and Mackenzie Weygandt Mathis. A primer on motion capture with deep learning: Principles, pitfalls, and perspectives. *Neuron*, 108(1):44–65, 2020.
21. Gregoire Courtine, Bingbing Song, Roland R Roy, Hui Zhong, Julia E Herrmann, Yan Ao, Jingwei Qi, V Reggie Edgerton, and Michael V Sofroniew. Recovery of supraspinal control of stepping via indirect propriospinal relay connections after spinal cord injury. *Nature medicine*, 14(1):69–74, 2008.
22. Joachim Von Zitzewitz, Leonie Asboth, Nicolas Fumeaux, Alexander Hasse, Laetitia Baud, Heike Vallery, and Grégoire Courtine. A neurorobotic platform for locomotor prosthetic development in rats and mice. *Journal of neural engineering*, 13(2):026007, 2016.
23. Rubia Van den Brand, Janine Heutschi, Quentin Barraud, Jack DiGiovanna, Kay Bartholdi, Michèle Huerlimann, Lucia Friedli, Isabel Vollenweider, Eduardo Martin Moraud, Simone Duis, et al. Restoring voluntary control of locomotion after paralyzing spinal cord injury. *science*, 336(6085): 1182–1185, 2012.
24. Łukasz Kidziński, Bryan Yang, Jennifer L Hicks, Apoorva Rajagopal, Scott L Delp, and Michael H Schwartz. Deep neural networks enable quantitative movement analysis using single-camera videos. *Nature communications*, 11 (1):1–10, 2020.
25. McKenzie S White, Ross J Brancati, and Lindsey K Lepley. Relationship between altered knee kinematics and subchondral bone remodeling in a clinically translational model of acl injury. *Journal of Orthopaedic Research®*, 2020.
26. Stefan Williams, Zhibin Zhao, Awais Hafeez, David C Wong, Samuel D Relton, Hui Fang, and Jane E Alty. The discerning eye of computer vision: Can it measure parkinson's finger tap bradykinesia? *Journal of the Neurological Sciences*, 416:117003, 2020.
27. Jared M Cregg, Roberto Leiras, Alexia Montalant, Paulina Wanken, Ian R Wickersham, and Ole Kiehn. Brainstem neurons that command mammalian locomotor asymmetries. *Nature Neuroscience*, 23(6):730–740, 2020.
28. Eleanor Batty, Matthew Whiteway, Shreya Saxena, Dan Biderman, Taiga Abe, Simon Musall, Winthrop Gillis, Jeffrey Markowitz, Anne Churchland, John P Cunningham, et al. Behavenet: nonlinear embedding and bayesian neural decoding of behavioral videos. In *Advances in Neural Information Processing Systems*, pages 15680–15691, 2019.
29. Kevin Luxem, Falko Fuhrmann, Johannes Kürsch, Stefan Remy, and Pavol Bauer. Identifying behavioral structure from deep variational embeddings of animal motion. *bioRxiv*, 2020.
30. Robert Evan Johnson, Scott Linderman, Thomas Panier, Caroline Lei Wee, Erin Song, Kristian Joseph Herrera, Andrew Miller, and Florian Engert. Probabilistic models of larval zebrafish behavior reveal structure on many scales. *Current Biology*, 30(1):70–82, 2020.

31. Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9 (Nov):2579–2605, 2008.
32. Leland McInnes, John Healy, and James Melville. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv:1802.03426*, 2018.
33. Yuke Yan, James Michael Goodman, Dalton D Moore, Sara Solla, and Sliman J Bensmaia. Unexpected complexity of everyday manual behaviors. *Nat Commun*, 11: 3564, 2020.
34. Dmitry Kobak and Philipp Berens. The art of using t-sne for single-cell transcriptomics. *Nature communications*, 10 (1):1–14, 2019.
35. Dmitry Kobak and George C Linderman. Initialization is critical for preserving global data structure in both t-sne and umap. *Nature Biotechnology*, pages 1–2, 2021.
36. Gordon J. Berman, Daniel M. Choi, William Bialek, and Joshua W. Shaevitz. Mapping the stereotyped behaviour of freely moving fruit flies. *Journal of The Royal Society Interface*, 11(99), 2014. ISSN 1742-5689. doi: 10.1098/rsif.2014.0672.
37. Alexander I Hsu and Eric A Yttri. B-soid: An open source unsupervised algorithm for discovery of spontaneous behaviors. *bioRxiv*, page 770271, 2020.
38. Praneet C Bala, Benjamin R Eisenreich, Seng Bum Michael Yoo, Benjamin Y Hayden, Hyun Soo Park, and Jan Zimmermann. Openmonkeystudio: automated markerless pose estimation in freely moving macaques. *bioRxiv*, 2020.
39. Brian D DeAngelis, Jacob A Zavatone-Veth, and Damon A Clark. The manifold structure of limb coordination in walking drosophila. *Elife*, 8:e46409, 2019.
40. João C Marques and Michael B Orger. Clusterdv: a simple density-based clustering method that is robust, general and automatic. *Bioinformatics*, 35(12):2125–2132, 2019.
41. João C Marques, Simone Lackner, Rita Félix, and Michael B Orger. Structure of the zebrafish locomotor repertoire revealed with unsupervised behavioral clustering. *Current Biology*, 28(2):181–195, 2018.
42. Duncan S Mearns, Joseph C Donovan, António M Fernandes, Julia L Semmelhack, and Herwig Baier. Deconstructing hunting behavior reveals a tightly coupled stimulus-response loop. *Current Biology*, 30(1):54–69, 2020.
43. Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv:1312.6114*, 2013.
44. Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. http://www.deeplearningbook.org.
45. D. Blei, A. Kucukelbir, and J. McAuliffe. Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112:859 – 877, 2016.
46. Jacob M Graving and Iain D Couzin. Vae-sne: a deep generative model for simultaneous dimensionality reduction and clustering. *BioRxiv*, 2020.
47. Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in python. *Journal of machine learning research*, 12(Oct):2825–2830, 2011.
48. Simon RO Nilsson, Nastacia L. Goodwin, Jia Jie Choong, Sophia Hwang, Hayden R. Wright, Zane C. Norville, Xiaoyu Tong, Dayu Lin, B. Bentzley, N. Eshel, R. McLaughlin, and S. Golden. Simple behavioral analysis (simba) – an open source toolkit for computer classification of complex social behaviors in experimental animals. *bioRxiv*, 2020.
49. Oliver Sturman, Lukas von Ziegler, C. Schläppi, Furkan Akyol, M. Privitera, Daria Slominski, Christina Grimm, Laetitia Thieren, V. Zerbi, Benjamin F. Grewe, and J. Bohacek. Deep learning-based behavioral analysis reaches human accuracy and is capable of outperforming commercial solutions. *Neuropsychopharmacology*, 45:1942 – 1952, 2020.
50. Cristina Segalin, Jalani Williams, Tomomi Karigo, May Hui, Moriel Zelikowsky, Jennifer J Sun, Pietro Perona, David J Anderson, and Ann Kennedy. The mouse action recognition system (mars): a software pipeline for automated analysis of social behaviors in mice. *bioRxiv*, 2020.
51. Alexander B Wiltschko, Matthew J Johnson, Giuliano Iurilli, Ralph E Peterson, Jesse M Katon, Stan L Pashkovski, Victoria E Abraira, Ryan P Adams, and Sandeep Robert Datta. Mapping sub-second structure in mouse behavior. *Neuron*, 88(6):1121–1135, 2015.
52. Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017.
53. Yunpeng Chen, Yannis Kalantidis, Jianshu Li, Shuicheng Yan, and Jiashi Feng. Multi-fiber networks for video recognition. In *Proceedings of the european conference on computer vision (ECCV)*, pages 352–367, 2018.
54. Niek Andresen, Manuel Wöllhaf, Katharina Hohlbaum, Lars Lewejohann, Olaf Hellwich, Christa Thöne-Reineke, and Vitaly Belik. Towards a fully automated surveillance of well-being status in laboratory mice using deep learning: Starting with facial expression analysis. *Plos one*, 15(4): e0228059, 2020.
55. Nejc Dolensek, Daniel A Gehrlach, Alexandra S Klein, and Nadine Gogolla. Facial expressions of emotion states and their neuronal correlates in mice. *Science*, 368(6486):89–94, 2020.
56. Alexander B Wiltschko, Tatsuya Tsukahara, Ayman Zeine, Rockwell Anyoha, Winthrop F Gillis, Jeffrey E Markowitz, Ralph E Peterson, Jesse Katon, Matthew J Johnson, and Sandeep Robert Datta. Revealing the structure of pharmacobehavioral space through motion sequencing. *Nature Neuroscience*, 23(11):1433–1443, 2020.
57. Biagio Brattoli, Uta Buchler, Anna-Sophia Wahl, Martin E Schwab, and Bjorn Ommer. Lstm self-supervision for detailed behavior analysis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6466–6475, 2017.
58. James P Bohnslav, Nivanthika K Wimalasena, Kelsey J Clausing, David Yarmolinsky, Tomas Cruz, Eugenia Chiappe, Lauren L Orefice, Clifford J Woolf, and Christopher D Harvey. Deepethogram: a machine learning pipeline for supervised behavior classification from raw pixels. *bioRxiv*, 2020.
59. Christopher L Buckley and Taro Toyoizumi. A theory of how active behavior stabilises neural activity: Neural gain modulation by closed-loop environmental feedback. *PLoS computational biology*, 14(1):e1005926, 2018.
60. Ali Nourizonoz, Robert Zimmermann, Chun Lum Andy Ho, Sebastien Pellat, Yannick Ormen, Clément Prévost-Solié, Gilles Reymond, Fabien Pifferi, Fabienne Aujard, Anthony Herrel, et al. Etholoop: automated closed-loop neuroethology in naturalistic environments. *Nature Methods*, 17(10): 1052–1059, 2020.
61. Gary A Kane, Gonçalo Lopes, Jonny L Sanders, Alexander Mathis, and Mackenzie Mathis. Real-time, low-latency closed-loop feedback using markerless posture tracking. *Elife*, 9:e61909, 2020.
62. Brandon J Forys, Dongsheng Xiao, Pankaj Gupta, and Timothy H Murphy. Real-time selective markerless tracking of forepaws of head fixed mice using deep neural networks. *Eneuro*, 2020.
63. Jens Schweihoff, Matvey Loshakov, I. Pavlova, Laura Kück, L. A. Ewell, and M. Schwarz. Deeplabstream: Closing the loop using deep learning-based markerless, real-

time posture detection. *bioRxiv*, 2019.
64. Teppei Ebina, Keitaro Obara, Akiya Watakabe, Yoshito Masamizu, Shin-Ichiro Terada, Ryota Matoba, Masafumi Takaji, Nobuhiko Hatanaka, Atsushi Nambu, Hiroaki Mizukami, et al. Arm movements induced by noninvasive optogenetic stimulation of the motor cortex in the common marmoset. *Proceedings of the National Academy of Sciences*, 116(45):22844–22850, 2019.
65. Britton A Sauerbrei, Jian-Zhong Guo, Jeremy D Cohen, Matteo Mischiati, Wendy Guo, Mayank Kabra, Nakul Verma, Brett Mensh, Kristin Branson, and Adam W Hantman. Cortical pattern generation during dexterous movement is input-driven. *Nature*, 577(7790):386–391, 2020.
66. Cornelis Immanuel van der Zouwen, Joel Boutin, Maxime Fougere, Aurelie Flaive, Melanie Vivancos, Alessandro Santuz, Turgay Akay, Philippe Sarret, and Dimitri Ryczko. Freely behaving mice can brake and turn during optogenetic stimulation of the mesencephalic locomotor region. *bioRxiv*, 2020.
67. Carsen Stringer, Marius Pachitariu, Nicholas Steinmetz, Charu Bai Reddy, Matteo Carandini, and Kenneth D. Harris. Spontaneous behaviors drive multidimensional, brainwide activity. *Science*, 364(6437), 2019. ISSN 0036-8075. doi: 10.1126/science.aav7893.
68. Simon Musall, Matthew T. Kaufman, Ashley L. Juavinett, Steven Gluf, and Anne K. Churchland. Single-trial neural dynamics are dominated by richly varied movements. *Nat. Neurosci.*, 22, 2019. doi: https://doi.org/10.1038/s41593-019-0502-4.
69. Mackenzie W. Mathis. A new spin on fidgets. *Nature Neuroscience*, 22:1614–1616, 2019.
70. Bartul Mimica, Benjamin A Dunn, Tuce Tombaz, VPTNC Srikanth Bojja, and Jonathan R Whitlock. Efficient cortical coding of 3d posture in freely behaving rats. *Science*, 362(6414):584–589, 2018.
71. Cristopher M Niell and Michael P Stryker. Modulation of visual responses by behavioral state in mouse visual cortex. *Neuron*, 65:472–479, 2010.
72. Marcus Leinweber, D. R. Ward, Jan M. Sobczak, Alexander Attinger, and G. Keller. A sensorimotor circuit in mouse cortex for visual flow predictions. *Neuron*, 95:1420–1432.e5, 2017.
73. Alessandro Del Vecchio, F. Sylos-Labini, V. Mondì, P. Paolillo, Y. Ivanenko, F. Lacquaniti, and D. Farina. Spinal motoneurons of the human newborn are highly synchronized during leg movements. *Science Advances*, 6, 2020.
74. Alessandro Del Vecchio, Francesco Negro, Ales Holobar, Andrea Casolo, Jonathan P Folland, Francesco Felici, and Dario Farina. You are as fast as your motor neurons: speed of recruitment and maximal discharge of motor neurons determine the maximal rate of force development in humans. *The Journal of Physiology*, 597(9), 2019.
75. Coralie Hérent, Séverine Diem, Gilles Fortin, and Julien Bouvier. Independent respiratory and locomotor rhythms in running mice. *eLife*, 2020. doi: 10.7554/eLife.61919.
76. Stephen H Scott. Optimal feedback control and the neural basis of volitional motor control. *Nature Reviews Neuroscience*, 5:532–546, 2004.
77. Emanuel Todorov and Michael I Jordan. Optimal feedback control as a theory of motor coordination. *Nature neuroscience*, 5(11):1226–1235, 2002.
78. David W Franklin and Daniel M Wolpert. Computational mechanisms of sensorimotor control. *Neuron*, 72(3):425–442, 2011.
79. Tomaso Poggio, Andrzej Banburski, and Qianli Liao. Theoretical issues in deep networks. *Proceedings of the National Academy of Sciences*, 2020.
80. Guangyu Robert Yang and Xiao-Jing Wang. Artificial neural networks for neuroscientists: A primer. *arXiv:2006.01001*, 2020.
81. Mark M Churchland, John P Cunningham, Matthew T Kaufman, Justin D Foster, Paul Nuyujukian, Stephen I Ryu, and Krishna V Shenoy. Neural population dynamics during reaching. *Nature*, 487(7405):51–56, 2012.
82. Saurabh Vyas, Matthew D Golub, David Sussillo, and Krishna V Shenoy. Computation through neural population dynamics. *Annual Review of Neuroscience*, 43:249–275, 2020.
83. David Sussillo, Mark M Churchland, Matthew T Kaufman, and Krishna V Shenoy. A neural network that finds a naturalistic solution for the production of muscle activity. *Nature neuroscience*, 18(7):1025–1033, 2015.
84. Abigail A Russo, Sean R Bittner, Sean M Perkins, Jeffrey S Seely, Brian M London, Antonio H Lara, Andrew Miri, Najja J Marshall, Adam Kohn, Thomas M Jessell, et al. Motor cortex embeds muscle-like commands in an untangled population response. *Neuron*, 97(4):953–966, 2018.
85. Hansem Sohn, Devika Narain, Nicolas Meirhaeghe, and Mehrdad Jazayeri. Bayesian computation through cortical latent dynamics. *Neuron*, 103(5):934–947, 2019.
86. Jing Wang, Devika Narain, Eghbal A Hosseini, and Mehrdad Jazayeri. Flexible timing by temporal scaling of cortical responses. *Nature neuroscience*, 21(1):102–110, 2018.
87. Warasinee Chaisangmongkon, Sruthi K Swaminathan, David J Freedman, and Xiao-Jing Wang. Computing by robust transience: how the fronto-parietal network performs sequential, category-based decisions. *Neuron*, 93 (6):1504–1517, 2017.
88. Andrew J Zimnik and Mark M Churchland. Independent generation of sequence elements by motor cortex. *Nature neuroscience*, pages 1–13, 2021.
89. Lea Duncker, Laura Driscoll, Krishna V Shenoy, Maneesh Sahani, and David Sussillo. Organizing recurrent network dynamics by task-computation to enable continual learning. *Advances in Neural Information Processing Systems*, 33, 2020.
90. Gopala K Anumanchipalli, Josh Chartier, and Edward F Chang. Speech synthesis from neural decoding of spoken sentences. *Nature*, 568(7753):493–498, 2019.
91. Joseph G Makin, David A Moses, and Edward F Chang. Machine translation of cortical activity to text with an encoder–decoder framework. Technical report, Nature Publishing Group, 2020.
92. Kai J Sandbrink, Pranav Mamidanna, Claudio Michaelis, Mackenzie Weygandt Mathis, Matthias Bethge, and Alexander Mathis. Task-driven hierarchical deep neural network models of the proprioceptive pathway. *bioRxiv*, 2020.
93. Daniel LK Yamins, Ha Hong, Charles F Cadieu, Ethan A Solomon, Darren Seibert, and James J DiCarlo. Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proceedings of the National Academy of Sciences*, 111(23):8619–8624, 2014.
94. Daniel LK Yamins and James J DiCarlo. Using goal-driven deep learning models to understand sensory cortex. *Nature neuroscience*, 19(3):356–365, 2016.
95. Alexander JE Kell and Josh H McDermott. Deep neural network models of sensory systems: windows onto the role of task constraints. *Current opinion in neurobiology*, 55:121–132, 2019.
96. Marcel AJ van Gerven. A primer on encoding models in sensory neuroscience. *Journal of Mathematical Psychology*, 76:172–183, 2017.
97. MJ Prud'Homme and John F Kalaska. Proprioceptive activity in primate primary somatosensory cortex during active arm reaching movements. *Journal of neurophysiology*, 72(5):2280–2301, 1994.

98. Ari S Benjamin, Hugo L Fernandes, Tucker Tomlinson, Pavan Ramkumar, Chris VerSteeg, Raeed H Chowdhury, Lee E Miller, and Konrad P Kording. Modern machine learning as a benchmark for fitting neural responses. *Frontiers in computational neuroscience*, 12:56, 2018.
99. Santiago A Cadena, George H Denfield, Edgar Y Walker, Leon A Gatys, Andreas S Tolias, Matthias Bethge, and Alexander S Ecker. Deep convolutional models improve predictions of macaque v1 responses to natural images. *PLoS computational biology*, 15(4):e1006897, 2019.
100. Edgar Y. Walker, Fabian H Sinz, E. Cobos, Taliah Muhammad, Emmanouil Froudarakis, P. G. Fahey, Alexander S. Ecker, J. Reimer, Xaq Pitkow, and A. Tolias. Inception loops discover what excites neurons most using deep predictive models. *Nature Neuroscience*, pages 1–6, 2019.
101. Pouya Bashivan, Kohitij Kar, and J. DiCarlo. Neural population control via deep image synthesis. *Science*, 364, 2019.
102. Eero P Simoncelli and Bruno A Olshausen. Natural image statistics and neural representation. *Annual review of neuroscience*, 24(1):1193–1216, 2001.
103. Wilson S Geisler. Contributions of ideal observer theory to vision research. *Vision research*, 51(7):771–781, 2011.
104. Blake A Richards, Timothy P Lillicrap, Philippe Beaudoin, Yoshua Bengio, Rafal Bogacz, Amelia Christensen, Claudia Clopath, Rui Ponte Costa, Archy de Berker, Surya Ganguli, et al. A deep learning framework for neuroscience. *Nature neuroscience*, 22(11):1761–1770, 2019.
105. Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.
106. Martin Schrimpf, Jonas Kubilius, Michael J Lee, N Apurva Ratan Murty, Robert Ajemian, and James J DiCarlo. Integrative benchmarking to advance neurally mechanistic models of human intelligence. *Neuron*, 2020.
107. Alexander JE Kell, Daniel LK Yamins, Erica N Shook, Sam V Norman-Haignere, and Josh H McDermott. A task-optimized neural network replicates human auditory behavior, predicts brain responses, and reveals a cortical processing hierarchy. *Neuron*, 98(3):630–644, 2018.
108. Chengxu Zhuang, Jonas Kubilius, Mitra J Hartmann, and Daniel L Yamins. Toward goal-driven neural network models for the rodent whisker-trigeminal system. *Advances in Neural Information Processing Systems*, 30:2555–2565, 2017.
109. Matthew Botvinick, Jane X Wang, Will Dabney, Kevin J Miller, and Zeb Kurth-Nelson. Deep reinforcement learning and its neuroscientific implications. *Neuron*, 2020.
110. Martin Haesemeyer, Alexander F Schier, and Florian Engert. Convergent temperature representations in artificial and biological neural networks. *Neuron*, 103(6):1123–1134, 2019.
111. Jimin Kim and Eli Shlizerman. Deep reinforcement learning for neural control. *arXiv:2006.07352*, 2020.
112. Andrea Banino, Caswell Barry, Benigno Uria, Charles Blundell, Timothy Lillicrap, Piotr Mirowski, Alexander Pritzel, Martin J Chadwick, Thomas Degris, Joseph Modayil, et al. Vector-based navigation using grid-like representations in artificial agents. *Nature*, 557(7705):429–433, 2018.
113. Hanne Stensola, Tor Stensola, Trygve Solstad, Kristian Frøland, May-Britt Moser, and Edvard I Moser. The entorhinal grid map is discretized. *Nature*, 492(7427):72–78, 2012.
114. Martin Stemmler, Alexander Mathis, and Andreas VM Herz. Connecting multiple spatial scales to decode the population activity of grid cells. *Science Advances*, 1(11): e1500816, 2015.
115. Josh Merel, Diego Aldarondo, Jesse Marshall, Yuval Tassa, Greg Wayne, and Bence Olveczky. Deep neuroethology of a virtual rodent. *International Conference on Learning Representations*, 2019.
116. Nicolas Heess, Dhruva TB, Srinivasan Sriram, Jay Lemmon, Josh Merel, Greg Wayne, Yuval Tassa, Tom Erez, Ziyu Wang, SM Eslami, et al. Emergence of locomotion behaviours in rich environments. *arXiv:1707.02286*, 2017.
117. Xue Bin Peng, Glen Berseth, and Michiel Van de Panne. Terrain-adaptive locomotion skills using deep reinforcement learning. *ACM Transactions on Graphics (TOG)*, 35 (4):1–12, 2016.
118. Xue Bin Peng, Glen Berseth, KangKang Yin, and Michiel Van De Panne. Deeploco: Dynamic locomotion skills using hierarchical deep reinforcement learning. *ACM Transactions on Graphics (TOG)*, 36(4):1–13, 2017.
119. Divye Jain, Andrew Li, Shivam Singhal, Aravind Rajeswaran, Vikash Kumar, and Emanuel Todorov. Learning deep visuomotor policies for dexterous hand manipulation. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 3636–3643. IEEE, 2019.
120. Joonho Lee, Jemin Hwangbo, Lorenz Wellhausen, Vladlen Koltun, and Marco Hutter. Learning quadrupedal locomotion over challenging terrain. *Science robotics*, 5(47), 2020.
121. Ilge Akkaya, Marcin Andrychowicz, Maciek Chociej, Mateusz Litwin, Bob McGrew, Arthur Petron, Alex Paino, Matthias Plappert, Glenn Powell, Raphael Ribas, et al. Solving rubik's cube with a robot hand. *arXiv:1910.07113*, 2019.
122. OpenAI: Marcin Andrychowicz, Bowen Baker, Maciek Chociej, Rafal Jozefowicz, Bob McGrew, Jakub Pachocki, Arthur Petron, Matthias Plappert, Glenn Powell, Alex Ray, et al. Learning dexterous in-hand manipulation. *The International Journal of Robotics Research*, 39(1):3–20, 2020.
123. Bo Zhou, Hongsheng Zeng, Fan Wang, Yunxiang Li, and Hao Tian. Efficient and robust reinforcement learning with uncertainty-based value expansion. *arXiv:1912.05328*, 2019.
124. Seungmoon Song, Łukasz Kidziński, Xue Bin Peng, Carmichael Ong, Jennifer L Hicks, Serge Levine, Christopher Atkeson, and Scot Delp. Deep reinforcement learning for modeling human locomotion control in neuromechanical simulation. *bioRxiv*, 2020.
125. Chris Eliasmith, T. C. Stewart, Xuan Choo, Trevor Bekolay, T. DeWolf, Y. Tang, and Daniel Rasmussen. A large-scale model of the functioning brain. *Science*, 338:1202 – 1205, 2012.
126. Travis DeWolf, T. C. Stewart, J. Slotine, and C. Eliasmith. A spiking neural model of adaptive arm control. *Proceedings of the Royal Society B: Biological Sciences*, 283, 2016.
127. Feng-Xuan Choo. Spaun 2.0: Extending the world's largest functional brain model. 2018.
128. Timothy P Lillicrap and Stephen H Scott. Preference distributions of primary motor cortex neurons reflect control solutions optimized for limb biomechanics. *Neuron*, 77(1): 168–179, 2013.
129. Tomohiko Takei, Stephen G Lomber, Douglas J Cook, and Stephen H Scott. Causal roles of frontoparietal cortical areas in feedback control of the limb. *bioRxiv*, 2020.
130. Mackenzie Weygandt Mathis, Alexander Mathis, and Naoshige Uchida. Somatosensory cortex plays an essential role in forelimb motor adaptation in mice. *Neuron*, 93(6):1493 – 1503.e6, 2017. ISSN 0896-6273. doi: https://doi.org/10.1016/j.neuron.2017.02.049.
131. Jonathan A Michaels, Stefan Schaffelhofer, Andres Agudelo-Toro, and Hansjörg Scherberger. A goal-driven modular neural network predicts parietofrontal neural dy-

namics during grasping. *Proceedings of the National Academy of Sciences*, 2020.
132. Emin Orhan, Vaibhav Gupta, and Brenden M Lake. Self-supervised learning through the eyes of a child. *Advances in Neural Information Processing Systems*, 33, 2020.
133. Jess Sullivan, Michelle Mei, Amy Perfors, Erica H Wojcik, and Michael C Frank. Saycam: A large, longitudinal audiovisual dataset recorded from the infant's perspective. *PsyArXiv*, 2020. doi: 10.31234/osf.io/fy8zx.
134. Frederik Ebert, Chelsea Finn, Sudeep Dasari, Annie Xie, Alex X. Lee, and S. Levine. Visual foresight: Model-based deep reinforcement learning for vision-based robotic control. *ArXiv*, abs/1812.00568, 2018.
135. M. Pearson, B. Mitchinson, J. C. Sullivan, A. Pipe, and T. Prescott. Biomimetic vibrissal sensing for robots. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 366:3085 – 3096, 2011.
136. Douglas Heaven. Why deep-learning ais are so easy to fool. *Nature*, 574:163–166, 2019.
137. Dan Hendrycks, K. Zhao, Steven Basart, J. Steinhardt, and D. Song. Natural adversarial examples. *ArXiv*, abs/1907.07174, 2019.
138. Scott L Delp, Frank C Anderson, Allison S Arnold, Peter Loan, Ayman Habib, Chand T John, Eran Guendelman, and Darryl G Thelen. Opensim: open-source software to create and analyze dynamic simulations of movement. *IEEE transactions on biomedical engineering*, 54 (11):1940–1950, 2007.
139. Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033. IEEE, 2012.