

# Latent factors and dynamics in motor cortex and their application to brain-machine interfaces

In the fifty years since Evarts first recorded single neurons in motor cortex of behaving monkeys, great effort has been devoted to understanding their relation to movement. Yet these single neurons exist within a vast network, the nature of which has been largely inaccessible. With advances in recording technologies, algorithms, and computational power, the ability to study network-level phenomena is increasing exponentially. Recent experimental results suggest that the dynamical properties of these networks are critical to movement planning and execution. Here we discuss this dynamical systems perspective, and how it is reshaping our understanding of the motor cortices. Following an overview of key studies in motor cortex, we discuss techniques to uncover the “latent factors” underlying observed neural population activity. Finally, we discuss efforts to leverage these factors to improve the performance of brain-machine interfaces, promising to make these findings broadly relevant to neuroengineering as well as systems neuroscience.

**Title:** Latent factors and dynamics in motor cortex and their application to brain-machine interfaces

**Authors:** Chethan Pandarinath<sup>1,2\*</sup>, K. Cora Ames<sup>3,4</sup>, Abigail A. Russo<sup>3,5,6</sup>, Ali Farshchian<sup>7</sup>, Lee E. Miller<sup>7</sup>, Eva L. Dyer<sup>1,8</sup>, Jonathan C. Kao<sup>9,10</sup>

1. Wallace H. Coulter Department of Biomedical Engineering, Emory University and Georgia Institute of Technology, Atlanta, Georgia, USA
2. Department of Neurosurgery, Emory University, Atlanta, Georgia, USA
3. Department of Neuroscience, Columbia University, New York, New York, USA
4. Center for Theoretical Neuroscience, Columbia University, New York, New York, USA
5. Grossman Center for the Statistics of Mind, Columbia University, New York, New York, USA
6. Zuckerman Institute, Columbia University, New York, New York, USA
7. Department of Physiology, Northwestern University, Chicago, Illinois, USA
8. Department of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, Georgia, USA
9. Department of Electrical and Computer Engineering, University of California, Los Angeles, CA 90095
10. Neurosciences Program, University of California, Los Angeles, CA 90095

\*Correspondence: [chethan@gatech.edu](mailto:chethan@gatech.edu)

# Abstract:

In the fifty years since Evarts first recorded single neurons in motor cortex of behaving monkeys, great effort has been devoted to understanding their relation to movement. Yet these single neurons exist within a vast network, the nature of which has been largely inaccessible. With advances in recording technologies, algorithms, and computational power, the ability to study network-level phenomena is increasing exponentially. Recent experimental results suggest that the dynamical properties of these networks are critical to movement planning and execution. Here we discuss this dynamical systems perspective, and how it is reshaping our understanding of the motor cortices. Following an overview of key studies in motor cortex, we discuss techniques to uncover the “latent factors” underlying observed neural population activity. Finally, we discuss efforts to leverage these factors to improve the performance of brain-machine interfaces, promising to make these findings broadly relevant to neuroengineering as well as systems neuroscience.

## Introduction

Our knowledge of the motor cortices (MC) is rapidly evolving. Traditional models of motor cortical activity held that the firing rates of individual neurons “represent” externally-measurable movement covariates, such as hand or joint kinematics, forces, or muscle activity. Much effort in related studies was devoted to finding the “correct” coordinate system. However, the increased ability to record from many neurons simultaneously has revealed many features of population activity that are difficult to reconcile with a purely representational viewpoint. First, much of the observed, high-dimensional activity of neural populations in MC can be explained as a combination of a modest number of “latent factors” – abstract, time-varying patterns that cannot be observed directly, but represent the correlated activity of the neural population. Second, during movements, these factors appear to evolve in time by obeying consistent dynamic rules, much like the lawful dynamics that govern physical systems. Through this lens, the complex, often-puzzling responses of individual neurons are naturally explained as minor elements in a coordinated underlying dynamical system. These findings have provided a new framework for evaluating neural activity during many of the functions that are ascribed to MC, such as motor preparation and execution, motor learning, bimanual control, and the production of muscle activity.

Beyond their application to the motor cortices, the dynamical systems framework and related computational methods may have broad applicability throughout the brain. Over the past decade, the ability to record from large populations of neurons has increased exponentially (Stevenson and Kording, 2011; Sofroniew et al., 2016; Jun et al., 2017; Stringer et al., 2018). These data collection tools promise to further transform our understanding of the brain, but only if we can process and interpret the coming wave of massive datasets. Trying to interpret the “tuning” of 10,000 neurons is not only onerous but a missed opportunity – much of the brain’s computation is inaccessible from the activity of individual neurons, but instead instantiated via population-level dynamics. Fortunately, modeling neural populations as low-dimensional dynamical systems is providing new insights in many cortical areas, including areas that mediate cognitive processes such as decision-making (Mante et al., 2013; Raposo et al., 2014; Carnevale et al., 2015), interval timing (Remington et al., 2018), and navigation (Harvey et al., 2012; Morcos and Harvey, 2016). This has deep implications for systems neuroscience: moving forward, the central thrust in understanding how brain areas perform computations and mediate behaviors may be through uncovering their population structure and underlying dynamics. MC is a critical model for studying these phenomena, as its activity appears strongly governed by internal dynamics, yet is well-related to observable behavior. These characteristics make MC an excellent “proving ground” for tools that may be useful in a wide variety of brain areas.

Further, our increasing knowledge of latent factors and dynamics in MC creates new opportunities to harness cortical activity to build high-performance and robust brain-machine interfaces (BMIs) to restore mobility to people with paralysis. BMIs aim to restore function by directly interfacing with the brain and reading out neural activity related to a person’s movement intent. To date, the vast majority of BMIs that leverage MC activity have used a representational viewpoint, with the assumption that individual neurons represent external movement covariates. Incorporating knowledge of the latent structure and dynamics of MC population activity potentially offers the means to develop BMIs whose performance and long-term stability are greatly improved.

Our review is divided to cover three broad areas: (1) an overview of the dynamical systems view of MC, including key studies that have tested its applicability and demonstrated new insight into the structure of population activity in MC; (2) current techniques to uncover latent structure and dynamics from the activity

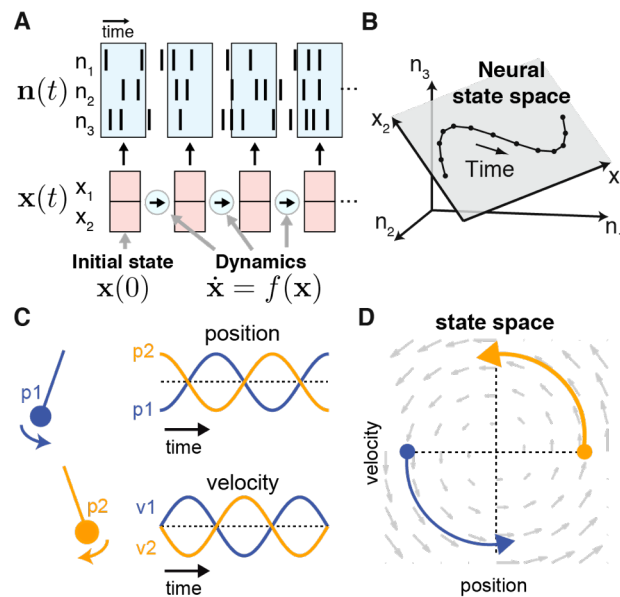
of neural populations; (3) recent efforts to leverage latent factors and dynamics to improve BMI performance.

### **The dynamical systems view and evidence in motor cortex**

Early work to understand the relationship between MC activity and movements drew inspiration from studies in sensory areas, such as the experiments of Hubel and Wiesel in visual cortex. In those experiments, the response of a neuron was modeled as a function of carefully-controlled features of the presented stimuli (Hubel and Wiesel, 1959). Similarly, studies in MC revealed that the responses of individual neurons (e.g., spike counts over hundreds of milliseconds) could be reasonably well-modeled as a function of kinetic or kinematic movement parameters (Evarts, 1968; Georgopoulos et al., 1982; Schwartz et al., 1988). A complication of the motor domain is that these movement covariates could only be studied by training animals to produce highly stereotypic movements, replete with many correlations across limb segments and measurement systems. Over the next decades, a long-simmering debate that had originated perhaps with Hughlings Jackson over which parameters of movement were represented (Jackson, 1873; Phillips, 1975) was given new fuel. Anatomical considerations argue for a strong, direct link between primary motor cortex and muscle activity (Landgren et al., 1962; Jankowska et al., 1975; Cheney and Fetz, 1985), supported by many studies which found that neural activity co-varies with muscle activation and kinetics (Evarts, 1968; Hepp-Reymond et al., 1999; Gribble and Scott, 2002; Holdefer and Miller, 2002). Yet correlates of higher-level parameters such as endpoint position (Riehle and Requin, 1989), velocity (Georgopoulos et al., 1982), speed (Churchland et al., 2006b), and curvature (Hocherman and Wise, 1991) could all be found as well. As this list became longer, some began to notice that these representations could also break down quite badly (Fu et al., 1995; Churchland and Shenoy, 2007a), and that such correlations could be spurious (Mussa-Ivaldi, 1988). This lead many to wonder whether viewing MC as a representational system is appropriate (Fetz, 1992; Scott, 2008; Churchland et al., 2010).

Rather than asking which parameters constitute the output of MC, one might instead view the system from a generative perspective: how does MC generate its output? From this perspective, MC is seen as a computational engine whose activity translates high-level movement intention into the complex patterns of muscle activity required to execute a movement (Todorov and Jordan, 2002; Scott, 2004; Shenoy et al., 2013). If so, how might this computation be performed?

For decades, theoreticians have posited that brain areas may perform computation through network-level phenomena in which information is distributed across the activity of many neurons, and processed via lawful dynamics that dictate how the activity of a neural population evolves over time (rev. in (Yuste, 2015)). We formalize this dynamical view in **Fig. 1A**. We assume that at a given time point  $t$ , the activity of an population of  $D$  neurons can be captured by a vector of spike counts  $\mathbf{n}(t) = [n_1(t), n_2(t), \dots, n_D(t)]$ . The neural population acts as a coordinated unit, with a  $K$ -dimensional internal ‘state’  $\mathbf{x}(t) = [x_1(t), x_2(t), \dots, x_K(t)]$ . In many brain areas,  $\mathbf{x}(t)$  has been observed to be much lower-dimensional than the total number of observed neurons (i.e.,  $K \ll D$ ; **Fig. 1B**; (Cunningham and Yu, 2014)). This dimensionality is likely somewhat constrained due to the recurrent connectivity of the network, which restricts the possible patterns of co-activation that may occur (rev. in (Gallego et al., 2017a)). A notable consideration, however, is that the observed dimensionality is often lower than might be expected due to network constraints alone – this particularly low dimensionality may be further induced by the simplicity of common behavioral paradigms (Gao and Ganguli, 2015; Gao et al., 2017).



**Figure 1.** Intuition for latent factors and dynamical systems. **(A)**  $\mathbf{n}(t)$  is a vector representing observed spiking activity. Each element of the vector captures the number of spikes a given neuron emits within a short time window around time  $t$ .  $\mathbf{n}(t)$  can typically be captured by the neural state variable  $\mathbf{x}(t)$ , an abstract, lower-dimensional representation that captures the state of the network. Dynamics are the rules that govern how the state updates in time. For a completely autonomous dynamical system without noise, if the dynamics  $f(\mathbf{x})$  are known, then the upcoming states are completely predictable based on an initial state  $\mathbf{x}(0)$ . **(B)** In a simple 3-neuron example, the ensemble's activity at each point in time traces out a trajectory in a 3-D state space, where each axis represents the activity of a given neuron. Not all possible patterns of activity are observed, rather, activity is confined to a 2-D plane within the 3-D space. The axes of this plane represent the neural state dimensions. Adapted from Cunningham & Yu, 2014. **(C)** Conceptual low-dimensional dynamical system: a 1-D pendulum. A pendulum released from point  $p1$  or  $p2$  traces out different positions and velocities over time, and the state of the system can be captured by two state variables (position and velocity). **(D)** The evolution of the system over time follows a fixed set of dynamic rules, i.e., the pendulum's equations of motion. Knowing the pendulum's initial state ( $\mathbf{x}(0)$ , filled circles) and the dynamical rules that govern its evolution ( $f(\mathbf{x})$ , grey vector flow-field) is sufficient to predict the system's state at all future time points.

The dynamical systems view posits an additional constraint: the evolution of the population's activity in time is largely determined by internal rules (dynamics). In the limit of an autonomous dynamical system (i.e., a system that operates independently of any external factors), and without noise, the system's evolution follows the equation  $\dot{\mathbf{x}} = f(\mathbf{x})$ , that is, its future state changes are completely dependent upon (and predicted by) the current state. A conceptual example of a low-dimensional system with simple rotational dynamics (a 1-D pendulum), and its related dynamical flow-field, is presented in **Fig. 1C,D**. We note that MC clearly cannot be autonomous. It must receive and process inputs, such as sensory information, to produce responsive behaviors. However, as discussed below, the model of an autonomous dynamical system is reasonable for MC activity during the execution of well-prepared movements. During behaviors that are unprepared, or where unpredictable events necessitate corrections (such as responding to task perturbations), MC activity may be well-modeled as an input-driven dynamical system, analogous to a pendulum started from particular initial conditions, and subject to external perturbations (Pandarinath et al., 2018).

The dynamical systems framework makes testable predictions about the nature of MC activity. First, it predicts that the initial conditions of the system, such as those observed during movement preparation,

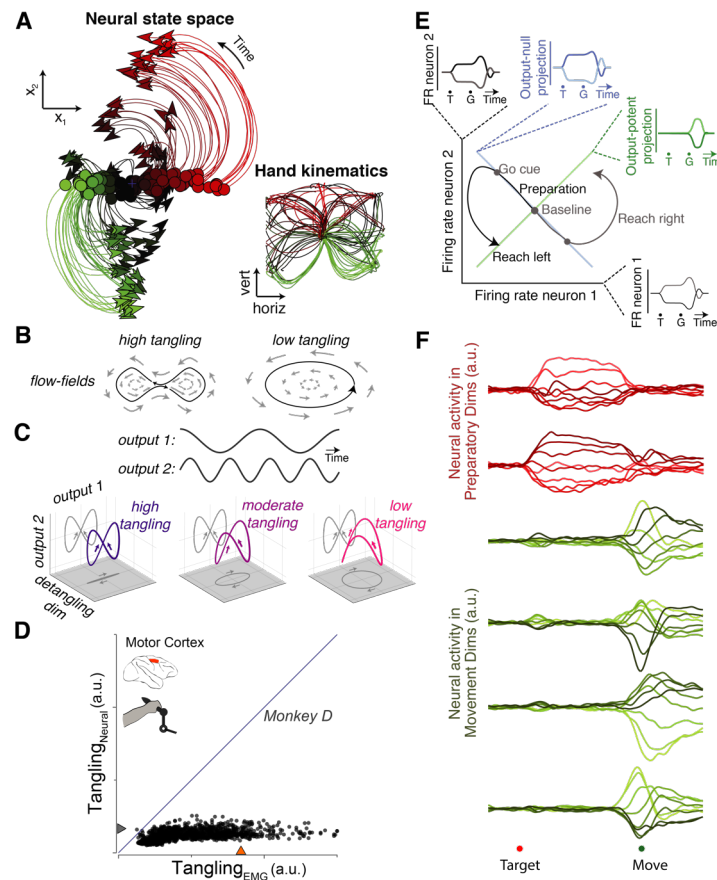
largely determine the subsequent evolution of activity. Second, the activity of neurons in MC should relate not only to the inputs and outputs of the system, but also to the computations being performed. Finally, distinct computations may be appropriated into different, non-overlapping, neural dimensions. Here, we explore experimental evidence related to each of these predictions.

Early studies exploring the dynamical systems hypothesis in MC examined whether preparatory activity served as an “initial condition” for the subsequent dynamics. In tasks with delay periods, where a subject has knowledge of the movement condition prior to execution, neural activity in MC approaches distinct “preparatory states” for distinct movements (Tanji and Evarts, 1976). In a dynamical system, initial conditions determine subsequent activity patterns, so the same dynamical “rules” can give rise to different activity patterns and behaviors if the initial condition is different. Similarly, in MC, an altered preparatory state relates to altered movement execution. If neural preparatory activity is not in the right state at the time of the go cue, either due to natural fluctuations (Churchland et al., 2006a; Afshar et al., 2011; Michaels et al., 2015, 2018), sub-threshold microstimulation during the delay period (Churchland and Shenoy, 2007b), or a change in the location of the target (Ames et al., 2014), the reaction time is delayed compared to well-prepared trials. This suggests that, if motor preparation is incorrect, subjects do not move until their preparation has been corrected. Furthermore, motor adaptation to a visuo-motor rotation (Vyas et al., 2018) or visuo-motor scaling (Stavisky et al., 2017a) has been shown to modify the motor preparatory state. These modifications correspond to altered execution trajectories, and the associated changes in preparation states and execution trajectories transfer from covert settings (BMI tasks without movement) to overt movements (normal reaching movements) (Vyas et al., 2018). Additional work has tested whether dynamical system models continue to predict observed neural activity during the transition from preparation to movement. A simple model in which preparatory activity seeds the initial condition for rotational dynamics during movement generation fits neural activity well in non-human primates (**Fig. 2A**; (Churchland et al., 2012; Elsayed et al., 2016; Michaels et al., 2016; Pandarinath et al., 2018)) and humans (Pandarinath et al., 2015; Pandarinath et al., 2018). Furthermore, recurrent neural networks trained to generate muscle activity after receiving preparatory input display dynamics similar to those recorded in MC (Hennequin et al., 2014; Sussillo et al., 2015; Kaufman et al., 2016), suggesting that a dynamical system that uses preparatory activity as the initial condition for subsequent movement dynamics may be a natural strategy for generating muscle activity during reaching.

Another important prediction of the dynamical systems model is that not all of the activity in MC must directly relate to task parameters or muscle activity, but may instead relate to internal processes that subserve the current computation. For example, the switch from movement preparation to movement generation is accompanied by a substantial change in dynamics (Kaufman et al., 2014; Elsayed et al., 2016). Recent work has posited that this change is accomplished by a large, condition-invariant translation in state-space which triggers the activation of movement generation dynamics. Indeed, this condition-invariant signal is not only present at the switch from preparation to generation, but is the largest aspect of the motor cortical response (Kaufman et al., 2016). Further, during movement generation itself, the dominant patterns of neural activity may also play a role in supporting neural dynamics, rather than directly encoding the output. One challenge for a dynamical system results when the flow-field is highly “tangled”: when there are points in the space in which very similar states lead to very different future behavior. If two nearby points lead to different paths, then small amounts of noise in the system can lead to dramatic differences in the evolution of the neural state (**Fig. 2B**). A robust dynamical system, therefore, must ensure that the tangling is low, potentially by adding in additional dimensions of activity whose job is to “pull apart” points of high tangling (**Fig. 2C**). In MC, while some components of neural activity resemble muscle-like signals during movement generation, the largest patterns of neural activity



during movement generation appear to function to reduce tangling. (**Fig. 2D**; (Russo et al., 2018)). Thus, within MC, evidence has been found that some signals primarily support and drive dynamics, rather than directly encoding input or output.



**Figure 2.** Overview of results supporting the dynamical systems view of motor cortex. (**A**) The neural state achieved during the delay period (green-red dots) predicts the subsequent trajectory of movement activity (green-red lines). Each dot/line is a single reach condition, recorded from a 108-reach condition task (inset). Adapted from Churchland et al., 2012. (**B**) In dynamical systems, places where neighboring points in state space have very different dynamics are indications of ‘tangling.’ Such regions would be highly-sensitive to noise – small perturbations yield very different trajectories. (**C**) Conceptual example illustrating tangling. Imagine a system that needs to produce two sine waves, one of which has double the frequency of the other. If it contains these sine waves with no additional dimensions, activity would trace out a figure eight, with a point of “high tangling” in the center. By adding in a third dimension, the system can move from a “high tangling” to a “low tangling” configuration, using the third dimension to separate the tangled points. Adapted from Russo et al., 2018. (**D**) While EMG often displays highly-tangled points (x-axis), MC’s neural activity maintains low tangling (y-axis). (**E**) Illustration of muscle-potent/muscle-null concept. Imagine a muscle which is driven with a strength equal to the sum of the firing rates of two units. If the units change in such a way that one unit’s firing rate decreases as the other increases, then the overall drive to the muscle will remain the same (muscle-null). If, on the other hand, the neurons increase or decrease together, then the drive to the muscle will change (muscle-potent). In this way, neural activity can change in the muscle-null space while avoiding causing a direct change in the command to the muscles. Adapted from Kaufman et al., 2014. (**F**) Neural activity in MC occupies a different set of dimensions during motor preparation than during movement. Red: neural activity across different reach conditions in “preparatory” dimensions. Green: neural activity across different reach conditions in “movement” dimensions. Adapted from Elsayed et al., 2016.

Finally, the dynamical systems framework predicts that to perform different computations, neural activity may leverage different dimensions (Mante et al., 2013). While this need not be a property of every possible dynamical system, leveraging different dimensions for different functions allows a system to better maintain independence between its different roles. It has long been observed that many neurons are active both during movement preparation and movement generation. How then does preparatory activity avoid causing movement? Traditional views held that preparatory activity lies below a movement-generation threshold (Tanji and Evarts, 1976; Erlhagen and Schöner, 2002) or is under the influence of a gate (Bullock and Grossberg, 1988; Cisek, 2006). However, sub-threshold activation fails to explain why preferred directions are minimally correlated between preparation and movement (Churchland et al., 2010; Kaufman et al., 2010), and there is little evidence for gating, as inhibitory neurons in MC are not preferentially activated during motor preparation (Kaufman et al., 2013). The dynamical systems model, by contrast, makes a different prediction: that unwanted movements can be avoided by avoiding specific neural dimensions. Some dimensions, termed “output-potent”, correspond to activation patterns that are output to the muscles, while others, termed “output-null”, do not (**Fig. 2D**; (Kaufman et al., 2014)). In MC, different dimensions are activated during movement preparation and generation (**Fig. 2E**; (Elsayed et al., 2016)). Furthermore, the dimensions that best correlate with muscle activity are preferentially active during movement generation, suggesting that output-potent dimensions are selectively avoided during preparation (Kaufman et al., 2014). Similarly, distinct dimensions may be explored during cortically-dependent movement versus non-cortically-dependent movement (Miri et al., 2017), and sensory feedback initially enters MC in different dimensions from those of the muscle-potent activity (Stavisky et al., 2017b). This division of different functions into different neural dimensions is not limited to only muscle-related activity. In BMIs, where output-potent dimensions can be specified explicitly, activity that is informative about experimentally-induced perturbations is initially orthogonal to the corresponding corrective responses (Stavisky et al., 2017b). During long-term BMI use, activity in output-potent dimensions is more stable than output-null dimensions (Flint et al., 2016). Neural activity also tends to occupy different dimensions for different movement categories: for example, pedaling in a forward vs. reverse direction (Russo et al., 2018), isometric force production versus limb movement (Gallego et al., 2017b), or moving with the contralateral versus ipsilateral arm (Ames and Churchland, 2018). Leveraging different neural dimensions for different functions may allow the MC flexibility to generate activity patterns that support a wide variety of functions without interfering with one another (Perich et al., 2017).

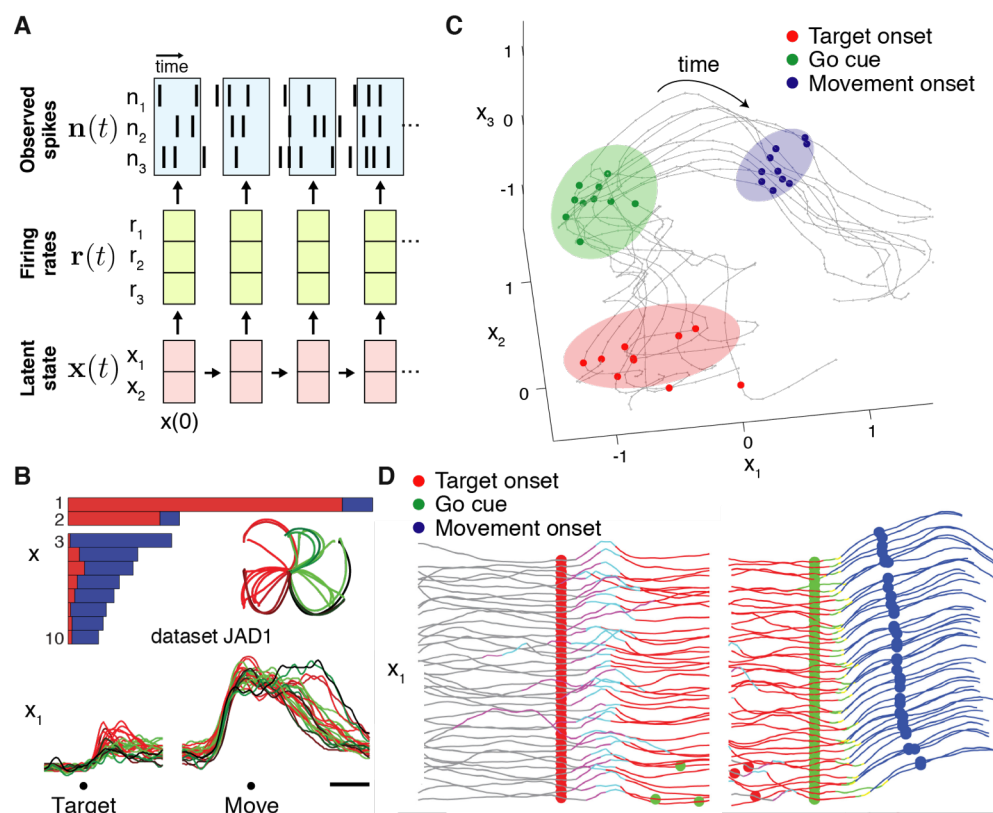
### **Methods for estimating and evaluating motor cortical dynamics**

As detailed above, much of the activity of MC neurons is naturally explained as a reflection of a low-dimensional dynamical system. Studying such dynamic processes requires techniques that can infer latent structure and its dynamics from observed, high-dimensional data. Related techniques have been applied to a wide variety of model systems and brain areas over the last two decades (rev. in (Cunningham and Yu, 2014)). However, MC holds particular value for testing these techniques, as its activity is closely tied to observable behavior (e.g., movement conditions, arm or hand kinematics, reaction times), which provides a key reference for validating the inferred state estimates. Here we review common techniques to estimate latent state and dynamics that have been applied in MC. We first present a general framework for discussion. Next, we review techniques that are applied to time points independently (i.e., do not explicitly model neural dynamics). Finally, we review techniques that do explicitly model neural dynamics, thereby resulting in better latent state estimates.



Techniques for latent state estimation typically view spiking activity as being “generated” by an underlying state  $\mathbf{x}(t)$  (**Fig. 3A**). A common assumption is that for any given trial, the observed high-dimensional spike counts  $\mathbf{n}(t)$  reflect a noisy sample from each neuron’s underlying firing rate distribution  $\mathbf{r}(t)$ , a distribution that is itself derived from the latent state  $\mathbf{x}(t)$ . For motor cortical data, the distinction between observations  $\mathbf{n}(t)$  and underlying rates  $\mathbf{r}(t)$  captures the empirical observation that the spiking of any given neuron across multiple repeats (trials) of the same movement is highly variable.

A standard approach to de-noising  $\mathbf{n}(t)$  and approximating  $\mathbf{r}(t)$  is trial-averaging. Trial-averaging assumes



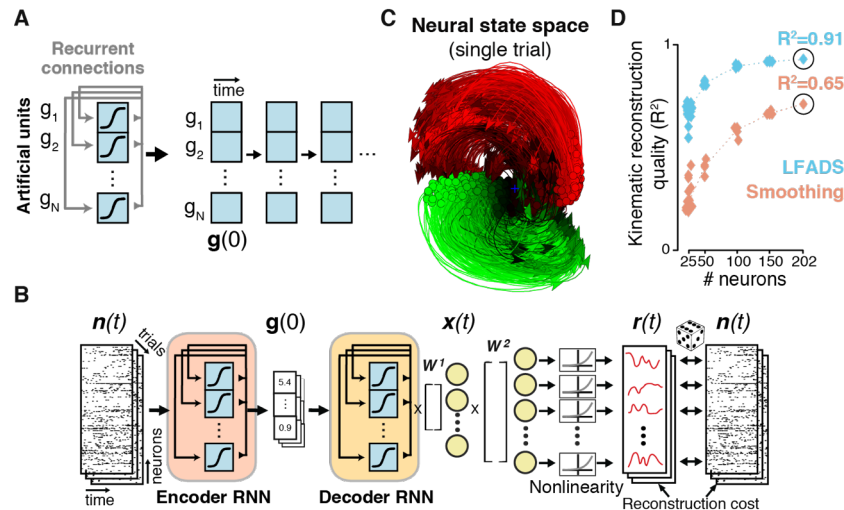
**Figure 3.** Applications of latent state and dynamics estimation methods to MC ensemble activity. **(A)** Generative model of observed neural activity. Population spiking activity is assumed to reflect an underlying latent state  $\mathbf{x}(t)$  whose temporal evolution follows consistent rules (dynamics). Firing rates for each neuron  $\mathbf{r}(t)$  are derived from  $\mathbf{x}(t)$ , and observed spikes  $\mathbf{n}(t)$  reflect a noisy sample from  $\mathbf{r}(t)$ . **(B)** dPCA applied to trial-averaged MC activity during a delayed reaching task separates condition-invariant and condition-variant dimensions. Each bar shows the total variance captured by each dimension, with red portions denoting condition-invariant fraction, and blue portions denoting condition-variant fraction. Traces show projection onto first dimension found by dPCA. Each trace corresponds to a single condition (inset: kinematic trajectories with corresponding colors). Adapted from Kaufman et al., 2016. **(C)** GPFA reveals single-trial state space trajectories during a delayed reaching task. Gray traces represent individual trials. Ellipses indicate across-trial variability of the neural state at reach target onset (red shading), go cue (green shading), and movement onset (blue shading). Adapted from Yu et al., 2009. **(D)** SLDS enables segmentation of individual trials by their dynamics. Each horizontal trace represents a single trial for the first state dimension found by the SLDS. Trace coloring represents time periods with distinct (discrete) dynamics for each trial, recognized in an unsupervised fashion. Switching between dynamic states reliably follow target onset and precede movement onset, with time lags that are correlated with reaction time. Adapted from Petreska et al., 2011.

all trials of a given movement condition are identical, and reduces single-trial noise in the estimate of  $\mathbf{r}(t)$  by averaging  $\mathbf{n}(t)$  across repeated trials.  $\mathbf{r}(t)$  is often further de-noised by convolving it with a smoothing kernel. A common approach to estimate the lower dimensional  $\mathbf{x}(t)$  is to perform Principal Component Analysis (PCA). Performing PCA on  $\mathbf{r}(t)$  rather than  $\mathbf{n}(t)$  is preferred – if performed on  $\mathbf{n}(t)$ , PCA often

results in poor latent factor estimation, because it simply maximizes the variance captured by the low-dimensional space, without separating variance that is shared among neurons from variance that is independent across neurons (rev. in (Yu et al., 2009)). When performed on  $\mathbf{r}(t)$ , PCA is typically accompanied by firing rate normalization, so that neurons with high rates (and thus higher variability) do not dominate the dimensionality reduction. Further, PCA can be extended by integrating some supervision into the dimensionality reduction step, e.g. by integrating information about task conditions in order to identify dimensions that capture neural variability related to particular task variables, using demixed PCA (dPCA, **Fig. 3B**; (Kobak et al., 2014; Kaufman et al., 2016; Gallego et al., 2017b)). The strategy of trial-averaging followed by PCA has led to several insights into latent structure and dynamics in MC (Churchland et al., 2012; Ames et al., 2014; Kaufman et al., 2014; Pandarinath et al., 2015; Elsayed et al., 2016; Kaufman et al., 2016; Gallego et al., 2017b; Russo et al., 2018).

However, circumventing the need to average over trials is critical for elucidating inherently single-trial phenomena, such as the trial-to-trial variability of real movements (and their corresponding error corrections), non-repeated behaviors such as natural behaviors, random target tasks, and tasks involving learning. Likewise, studying “internal” processes that vary substantially across trials and have limited behavioral correlates, such as decision-making, vacillation, and internal state estimates (Golub et al., 2015; Kaufman et al., 2015), is also impossible with trial-averaged data. Factor Analysis (FA; (Everitt, 1984)) is often favored for analyzing single-trial phenomena (Santhanam et al., 2009; Sadtler et al., 2014; Golub et al., 2015; Athalye et al., 2017; Golub et al., 2018). A key assumption of FA is that activity that is correlated across neurons represents “signal” (comprising the latent factors  $\mathbf{x}(t)$ ), and activity that is not correlated across neurons represents “noise.” This assumption matches the graphical model in **Fig. 3A**. A recent, complementary approach to capturing trial-dependent variability in  $\mathbf{n}(t)$  without corrupting the latent factors is to integrate information regarding trial ordering into the dimensionality reduction step, and introduce a set of “trial factors” that accommodate variability across trials, as in Tensor Components Analysis (TCA; (Williams et al., 2018)).

A key limitation of the above techniques (PCA, FA) is that they treat neighboring time points as though they are independent. However, as discussed, a core assumption of the dynamical systems framework is that time points are intimately related, and in particular, previous states are predictive of future states. Therefore, methods that simultaneously infer latent states and dynamics should provide more accurate state estimation, by leveraging the inter-dependencies of data points that are close in time. Two well-developed families of models are Gaussian Process-based approaches (**Fig. 3C**; (Yu et al., 2009; Lakshmanan et al., 2015; Zhao and Park, 2017; Duncker and Sahani, 2018)) and linear dynamical systems (LDS)-based approaches (Macke et al., 2011; Buesing et al., 2012; Aghagolzadeh and Truccolo, 2015; Gao et al., 2015; Kao et al., 2015; Gao et al., 2016; Kao et al., 2017). Gaussian Process approaches assume that the latent state  $\mathbf{x}(t)$  is composed of factors that vary smoothly and independently in time, with each factor having its own characteristic time constant. In comparison, LDS-based approaches assume that the latent state at a given time point is a linear function of the previous state (i.e.,  $\dot{\mathbf{x}} = \mathbf{A}\mathbf{x}$ ), which incorporates linear interactions between latent dimensions. One issue with the LDS approach is that the matrix  $\mathbf{A}$  is time-invariant, yet must capture the dynamics at all time points. In MC, this is potentially problematic, as activity during different behavioral phases (e.g., preparation and movement) is governed by very different dynamics (Kaufman et al., 2014; Elsayed et al., 2016; Kaufman et al., 2016). A promising approach to address this challenge is Switching LDS (SLDS), which assumes that at any given time point, the system’s evolution obeys one of a discrete set of possible dynamics, each of which must be learned (**Fig. 3D**; (Petreska et al., 2011; Linderman et al., 2017; Wei et al., 2018)).



**Figure 4.** LFADS uses recurrent neural networks to infer precise estimates of single-trial population dynamics. **(A)** A recurrent neural network (simplified) is a set of artificial neurons that implements a nonlinear dynamical system, with dynamics set by adjusting the weights of its recurrent connections. Conceptually, the RNN can be “unrolled” in time, where future states of the RNN are completely predicted based in an initial state  $\mathbf{g}(0)$  and its learned recurrent connectivity (compare to **Fig. 3A**). **(B)** The SAE framework consists of an encoding network and decoding network. The encoder (RNN) compresses single-trial observed activity  $\mathbf{n}(t)$  into a trial code  $\mathbf{g}(0)$ , which sets the initial state of the decoder RNN. The decoder attempts to re-create  $\mathbf{n}(t)$  based only on  $\mathbf{g}(0)$ . To do so, the decoder must model the ensemble’s dynamics using its recurrent connectivity. The output of the decoder is  $\mathbf{x}(t)$ , the latent factors, and  $\mathbf{r}(t)$ , the de-noised firing rates. **(C)** The de-noised single-trial estimates produced by LFADS uncover known dynamic features (such as rotations, **Fig. 2A**) on single trials. **(D)** Decoding the LFADS-de-noised rates using simple optimal linear estimation leads to vastly improved predictions of behavioral variables (hand velocities) over Gaussian smoothing, even with limited numbers of neurons. Adapted from Pandarinath et al., 2018.

An alternate approach to uncovering single-trial population dynamics uses recurrent neural networks (RNNs). Known as Latent Factor Analysis via Dynamical Systems (LFADS; (Sussillo et al., 2016a; Pandarinath et al., 2018)), the approach trains an RNN as a generative model of the observed spiking activity. RNNs are powerful nonlinear function approximators, capable of modeling complex, highly nonlinear dynamics through adjustment of their recurrent connectivity (**Fig. 4A**). LFADS uses a sequential autoencoder (SAE) framework (**Fig. 4B**), allowing the potentially nonlinear dynamics to be learned from noisy, single-trial neural population activity using stochastic gradient descent. This allows LFADS to accurately infer dynamics on a single-trial, moment-to-moment basis (**Fig. 4C**). A critical confirmation that these dynamics are accurate and meaningful is that they lead to dramatic improvements in the ability to predict behavior – as shown, the LFADS-inferred latent representations were considerably more informative about subjects’ reaching movements than was the population activity that was directly observed (**Fig. 4D**). These findings reinforce that population states, rather than the activity of individual neurons, may be a key factor in understanding how brain areas mediate behaviors, and further, that SAEs provide a powerful new avenue towards linking the activity of neural populations to the behaviors they mediate.

### Leveraging latent factors and dynamics for brain-machine interfaces

Brain machine interfaces (BMIs) aim to recover lost motor function by directly decoding movement intent from neuron spiking activity to control external devices or recover movement (Taylor et al., 2002; Carmena et al., 2003; Hochberg et al., 2006; Ethier et al., 2012; Hochberg et al., 2012; Collinger et al.,

2013; Sadtler et al., 2014; Gilja et al., 2015; Ajiboye et al., 2017; Pandarinath et al., 2017). BMIs have largely decoded neural activity into movement through a representational viewpoint: each neuron represents a reach direction and if the neuron fires, it votes for movement in that direction (Georgopoulos et al., 1982; Taylor et al., 2002; Gilja et al., 2012; Hochberg et al., 2012; Collinger et al., 2013). Efforts to decode EMG activity have been essentially similar, though in a higher-dimensional, more abstract space (Pohlmeyer et al., 2007; Ethier et al., 2012). However, as previously discussed, this representational model has important limitations in describing MC activity. Here we review recent studies that have asked whether representational assumptions also limit BMI performance, and if so, whether performance and robustness can be increased by incorporating MC latent factors and dynamics.

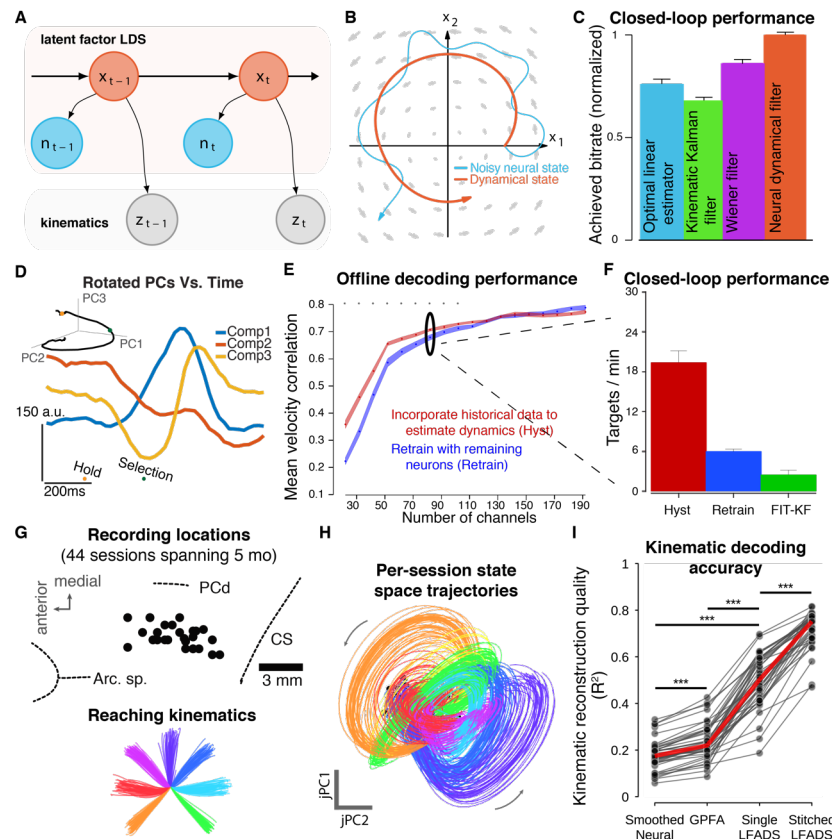
### *Using latent factors and dynamics to increase BMI performance*

The dynamical systems view holds that movement-related variables (such as kinematics or EMG activity) are among the many factors that influence the activity of MC neurons. However, BMI decoders based on the standard representational model do not take other factors into account when relating observed activity to movement intention. Recent work introduced a decoding architecture (graphically represented in **Fig. 5A**) that incorporates latent factors and their dynamics (modeled as a simple linear dynamical system; (Aghagolzadeh and Truccolo, 2015; Kao et al., 2015)). One advantage of this architecture is that modeling latent factors can account for the multiple, diverse influences on observed neural activity to better uncover movement-related variables. A second advantage is that latent factors may be more easily denoised than the observed high-dimensional activity, resulting in higher BMI performance. Briefly, the dynamical systems view holds the temporal evolution of MC states is largely predictable. If so, deviations from this prediction may largely correspond to noise. To denoise, MC dynamics can be used to adjust the latent factors so that they are more consistent with the dynamic predictions (**Fig. 5B**). In closed-loop BMI experiments, decoding the dynamically denoised latent factors significantly increased performance over prior approaches (**Fig 5C**), including previous representational decoders that (1) denoise activity by smoothing using an experimenter-chosen filter (optimal linear estimator, OLE; e.g., (Velliste et al., 2008; Collinger et al., 2013)), (2) denoise activity by incorporating prior knowledge about kinematic smoothness (kinematic Kalman filter; e.g., (Wu et al., 2003; Kim et al., 2008; Gilja et al., 2012; Hochberg et al., 2012; Gilja et al., 2015)), and (3) denoise activity by learning filtering parameters via least-squares regression (Wiener filter; e.g., (Carmena et al., 2003; Hochberg et al., 2006)).

BMI performance may also be increased through the use of non-movement signals that become apparent by examining latent factors. Recently, Even-Chen et al. (2017) exploited this idea to identify factors that reflect errors made during BMI control. The motivation for this work is that errors inevitably happen when controlling a BMI; however, instead of the user having to correct an error explicitly, it is possible to detect (or predict) its occurrence and automatically correct it (or prevent it). They applied PCA to identify an error-related signal in MC and found dimensions where projected neural data reflected task errors (example latent factors observed during errors are shown in **Fig. 5D**). In real-time experiments with



monkeys, these latent factors were decoded to both prevent and autocorrect errors, increasing BMI performance.



**Figure 5.** Improving BMI performance and longevity by leveraging neural dynamics. **(A)** Graphical model of decoder with dynamical smoothing. **(B)** Illustration of smoothing latent state estimates using neural dynamics. The instantaneous estimate of the latent state (blue) is augmented by a dynamical prior (grey flow-field) to produce a smoother, denoised estimate (orange). **(C)** Smoothing using neural dynamics results in better closed-loop BMI performance than other approaches. Performance is achieved information bitrate. Adapted from Kao et al., 2015. **(D)** Example of low-dimensional signals that can be used to augment intracortical BMIs. PCA applied to neural activity around the time of target selection identifies a putative “error signal”, allowing real-time detection and correction of user errors in a typing BMI. Adapted from Even-Chen et al., 2017. **(E)** Remembering dynamics from earlier recording conditions can extend performance as neurons are lost. Performance measure is (offline) mean velocity correlation. **(F)** Comparison of closed-loop performance when 110 channels are “lost” shows a >3x improvement achieved by remembering dynamics. FIT-KF is state of the art kinematic Kalman filter (Fan et al., 2014). Adapted from Kao et al., 2017. **(G)** Dynamic neural stitching with LFADS. A single model was trained on 44 recording sessions. Each session used a 24-ch recording probe. *Left:* recording locations in MC. *Right:* single-trial reaches from an example session. **(H)** Neural state space trajectories inferred by LFADS. Each trace of a given color is from a separate recording session (44 traces per condition). Inferred trajectories are consistent across five months. **(I)** Using LFADS to align five months of data (“Stitched”) significantly improves decoding vs. other tested methods.

### Using dynamics and latent factors to increase BMI longevity

Ideally, a BMI’s performance would be maintained indefinitely. However, neural recording conditions frequently change across days, and even within-day in pilot clinical trials, e.g., due to neuron death or electrode movement and failure (Barrese et al., 2013; Perge et al., 2013; Sussillo et al., 2016b; Downey et al., 2018), which can lead to decoding instability. Current approaches to this problem include decoding

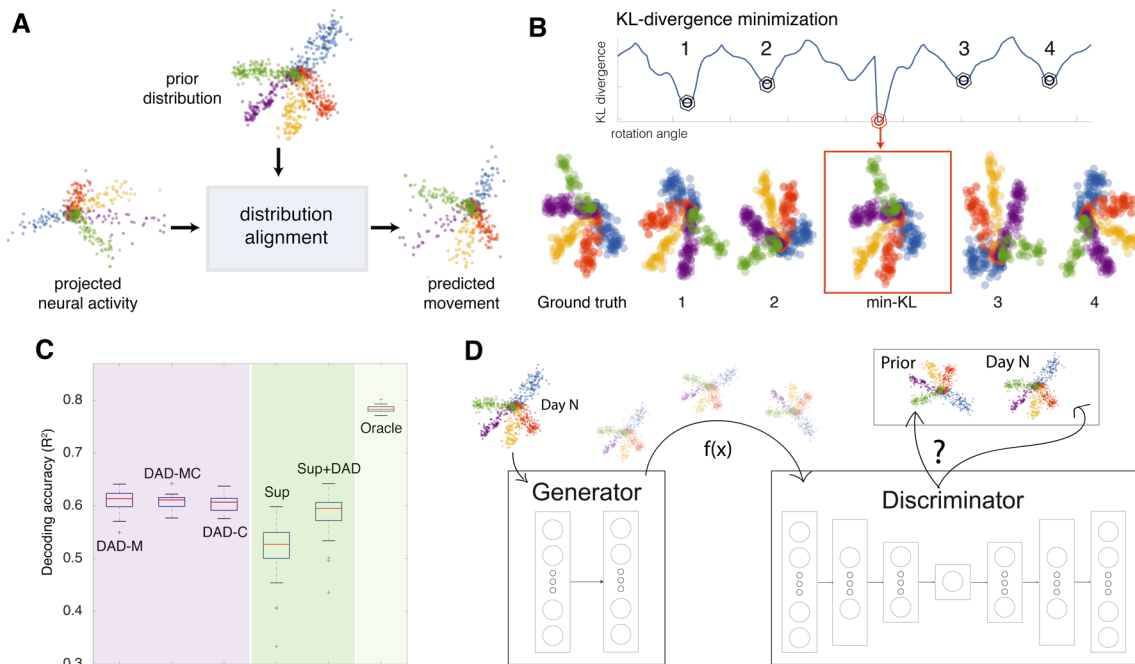


more stable neural signals (e.g., threshold crossings and local field potentials; (Flint et al., 2013; Nuyujukian et al., 2014; Gilja et al., 2015; Stavisky et al., 2015)), gradually updating decoder parameters using a weighted sliding average (Orsborn et al., 2012; Dangi et al., 2013), automated decoder recalibration by updating “tuning” estimates daily (Bishop et al., 2014) or continuous recalibration by retrospectively inferring the user’s intention among a set of fixed targets (Jarosiewicz et al., 2015), and training robust neural network decoders on a diversity of conditions using large data volumes (Sussillo et al., 2016b).

A separate class of approaches aims to exploit the underlying neural latent space, which, as a property of the neural population, should have a stable relationship with the user’s intention that is independent of the specific neurons observed at any moment (Gao and Ganguli, 2015; Dyer et al., 2017; Gallego et al., 2017b; Kao et al., 2017; Pandarinath et al., 2018). However, it is challenging to relate the observed neurons from a given recording condition to the underlying latent space. Recent studies using supervised alignment strategies have demonstrated the potential of latent dynamics to maintain BMI performance. Kao et al. (2017) exploited historical information about population dynamics (**Fig. 5E,F**), finding that even under severe neuron loss, aligning the remaining neurons to previously-learned dynamics could partially rescue closed-loop performance, effectively extending BMI lifetime. Alternatively, Pandarinath et al. (2018) learned a single LFADS model from 44 independently-recorded neural populations spanning many millimeters of MC and 5 months of recording sessions (**Fig. 5G,H**). They then used a single linear decoder to map these latent dynamics onto kinematics (**Fig. 5I**). This work demonstrated that, in the absence of learning, a single, consistent dynamical model describes neural population activity across long time periods and large cortical areas, and yields improved offline decoding performance for any given recording session than was otherwise possible.

Some settings lack data for supervised alignment, i.e., directly linking neural activity from new recording conditions to motor intent may be challenging (settings without structured behaviors, or where intent is less clear on a moment-by-moment basis). In these settings, unsupervised techniques may be useful for aligning data. Recently, Dyer et al. (2017) introduced a semi-supervised approach called Distribution Alignment Decoding (DAD; **Fig. 6A-C**). This approach aims to map neural data from a new recording condition (new data) onto a previously recorded low-dimensional movement distribution. To do this, DAD first reduces the dimensionality of the neural data (using PCA or a nonlinear manifold learning technique), and then searches for an affine transformation to match the low-dimensional neural data to movements, by minimizing the KL-divergence between the two datasets (**Fig. 6B**). Their results demonstrate that DAD can achieve similar performance to that of a supervised decoder that has access to information about neural state and movement, if the underlying data distribution contains asymmetries that facilitate alignment (**Fig. 6C**). While a powerful approach for unsupervised alignment, DAD solves a non-convex optimization problem with many local minima (**Fig. 6B**) by using a brute force search. To improve alignment and avoid having to first perform dimensionality reduction, neural network architectures such as generative adversarial networks (GANs; (Goodfellow et al., 2014; Molano-Mazon et al., 2018)) provide a potential method to learn nonlinear mappings from one distribution to another (**Fig. 6D**). By leveraging the fact that low-dimensional representations of neural activity are consistent across days and even

subjects, distribution alignment methods like DAD or GANs provide a strategy for decoding movements without labeled training data from new recording conditions.



**Figure 6.** Distribution alignment methods for stabilizing movement decoders across days and subjects. **(A)** Overview of distribution alignment approaches which first reduce data dimensionality, and then align low-dimensional projected neural activity onto a previously recorded movement distribution. **(B)** KL-divergence provides a robust metric for alignment, but also produces many local minima (points 1, 2, 3, 4) which make alignment difficult. The KL-divergence is displayed as a function of the angle used to rotate the data. **(C)** Comparison of the prediction accuracy of 2D kinematics for distribution alignment decoding and supervised methods. Left: results for DAD using movements from Subject M (DAD-M), from Subject C (DAD-C), and using movements from both Subjects M and C (DAD-MC). Right: results for a standard L2-regularized supervised decoder (Sup) and a combined decoder (Sup-DAD) which averages the results of the supervised and DAD decoders. All of these results are compared to an Oracle decoder which provides an upper bound for the best linear decoding performance for this task. **(D)** A schematic of a generative adversarial network strategy for distribution alignment across multiple days: Generator network (left) receives new data and learns a transformation of the data to match the prior (from a previous day).

## Conclusions

The increasing ability to monitor large numbers of neurons simultaneously will present new opportunities to study neural activity at the population level. Mounting evidence shows that this provides a qualitatively different window into the nervous system from that of single-neuron recordings, and that population-level dynamics likely underlie neural population activity across a wide range of systems. Here we reviewed recent evidence that such dynamics shape activity and drive behavior in MC, outlined key methods for inferring latent factors and dynamics that have been applied to MC activity, and showed how uncovering latent factors and dynamics can yield higher-performing and more robust BMIs. Continuing advances in recording technologies, algorithms, and computational power will enable studies of dynamics that were not previously possible, and further, may open new avenues for neural prostheses to address a wide variety of disorders of the nervous system.

## Acknowledgements

The authors would like to thank Steven Chase, Chandramouli Chandrasekaran, Juan Gallego, Matthew Kaufman, Daniel O'Shea, David Sussillo, Sergey Stavisky, Xulu Sun, Eric Trautmann, Jessica Verhein, Saurabh Vyas, Megan Wang, and Byron Yu for their feedback on the manuscript. This work was supported by a Burroughs Wellcome Fund Collaborative Research Travel Grant and NSF NCS 1835364 (CP), and NIH NINDS R01NS053603 and NSF NCS 1835345 (LEM).

## References

- Afshar A, Santhanam G, Yu BM, Ryu SI, Sahani M, Shenoy KV (2011) Single-trial neural correlates of arm movement preparation. *Neuron* 71:555-564.
- Aghagolzadeh M, Truccolo W (2015) Inference and Decoding of Motor Cortex Low-Dimensional Dynamics via Latent State-Space Models. *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 4320:1-1.
- Ajiboye AB, Willett FR, Young DR, Memberg W, Walters BC, Sweet JA, Hoyer HA, Keith MW, Peckham PH, Simeral JD, Donoghue JP, Miller JP, Hochberg LR, Kirsch RF (2017) Restoration of reaching and grasping in a person with tetraplegia through brain-controlled muscle stimulation: a proof-of-concept demonstration. *The Lancet* 6736:(In press).
- Ames KC, Churchland MM (2018) Population-level but not neuron-level similarity during movement of the contra- vs ipsi-lateral hand. In: *Computational and Systems Neuroscience*, pp 11-28. Denver, CO.
- Ames KC, Ryu SI, Shenoy KV (2014) Neural dynamics of reaching following incorrect or absent motor preparation. *Neuron* 81:438-451.
- Athalye VR, Ganguly K, Costa RM, Carmena JM (2017) Emergence of Coordinated Neural Dynamics Underlies Neuroprosthetic Learning and Skillful Control. *Neuron*:1-16.
- Barrese JC, Rao N, Paroo K, Triebwasser C, Vargas-Irwin C, Franquemont L, Donoghue JP (2013) Failure mode analysis of silicon-based intracortical microelectrode arrays in non-human primates. *Journal of Neural Engineering* 10:06014.
- Bishop W, Chestek CC, Gilja V, Nuyujukian P, Foster JD, Ryu SI, Shenoy KV, Byron MY (2014) Self-recalibrating classifiers for intracortical brain-computer interfaces. *Journal of neural engineering* 11:026001.
- Buesing L, Macke J, Sahani M (2012) Learning stable, regularised latent models of neural population dynamics. *Network*:1-24.
- Bullock D, Grossberg S (1988) Neural dynamics of planned arm movements: emergent invariants and speed-accuracy properties during trajectory formation. *Psychological review* 95:49.
- Carmena JM, Lebedev MA, Crist RE, O'Doherty JE, Santucci DM, Dimitrov DF, Patil PG, Henriquez CS, Nicolelis MaL (2003) Learning to control a brain-machine interface for reaching and grasping by primates. *PLoS Biology* 1:193-208.
- Carnevale F, de Lafuente V, Romo R, Barak O, Parga N (2015) Dynamic Control of Response Criterion in Premotor Cortex during Perceptual Detection under Temporal Uncertainty. *Neuron* 86:1067-1077.
- Cheney PD, Fetz EE (1985) Comparable patterns of muscle facilitation evoked by individual corticomotoneuronal (CM) cells and by single intracortical microstimuli in primates: evidence for functional groups of CM cells. *Journal of neurophysiology* 53:786-804.
- Churchland MM, Shenoy KV (2007a) Temporal complexity and heterogeneity of single-neuron activity in premotor and motor cortex. *Journal of neurophysiology* 97:4235-4257.
- Churchland MM, Shenoy KV (2007b) Delay of movement caused by disruption of cortical preparatory activity. *Journal of neurophysiology* 97:348-359.
- Churchland MM, Afshar A, Shenoy KV (2006a) A central source of movement variability. *Neuron* 52:1085-1096.
- Churchland MM, Santhanam G, Shenoy KV (2006b) Preparatory activity in premotor and motor cortex reflects the speed of the upcoming reach. *Journal of neurophysiology* 96:3130-3146.

- Churchland MM, Cunningham JP, Kaufman MT, Ryu SI, Shenoy KV (2010) Cortical Preparatory Activity: Representation of Movement or First Cog in a Dynamical Machine? In: *Neuron*, pp 387-400: Elsevier Inc.
- Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P, Ryu SI, Shenoy KV (2012) Neural population dynamics during reaching. *Nature*.
- Cisek P (2006) Integrated neural processes for defining potential actions and deciding between them: a computational model. *Journal of Neuroscience* 26:9761-9770.
- Collinger JL, Wodlinger B, Downey JE, Wang W, Tyler-Kabara EC, Weber DJ, McMorland AJC, Velliste M, Boninger ML, Schwartz AB (2013) High-performance neuroprosthetic control by an individual with tetraplegia. *The Lancet* 381:557-564.
- Cunningham JP, Yu BM (2014) Dimensionality reduction for large-scale neural recordings. *Nature Neuroscience* 17:1500-1509.
- Dangi S, Orsborn AL, Moorman HG, Carmena JM (2013) Design and analysis of closed-loop decoder adaptation algorithms for brain-machine interfaces. *Neural computation* 25:1693-1731.
- Downey JE, Schwed N, Chase SM, Schwartz AB, Collinger JL (2018) Intracortical recording stability in human brain-computer interface users. *Journal of neural engineering* 15:046016.
- Duncker L, Sahani M (2018) Temporal alignment and latent Gaussian process factor inference in population spike trains. *bioRxiv*:331751.
- Dyer EL, Azar MG, Perich MG, Fernandes HL, Naufel S, Miller LE, Körding KP (2017) A cryptography-based approach for movement decoding. *Nature Biomedical Engineering* 1:967.
- Elsayed GF, Lara AH, Kaufman MT, Churchland MM, Cunningham JP (2016) Reorganization between preparatory and movement population responses in motor cortex. *Nature Communications*:13239.
- Erlhagen W, Schöner G (2002) Dynamic field theory of movement preparation. *Psychological review* 109:545.
- Ethier C, Oby ER, Bauman MJ, Miller LE (2012) Restoration of grasp following paralysis through brain-controlled stimulation of muscles. *Nature* 485:368-371.
- Evarts EV (1968) Relation of pyramidal tract activity to force exerted during voluntary movement. *Journal of neurophysiology* 31:14-27.
- Even-Chen N, Stavisky SD, Kao JC, Ryu SI, Shenoy KV (2017) Augmenting intracortical brain-machine interface with neurally driven error detectors. *Journal of neural engineering* 14:066007.
- Everitt B (1984) An introduction to latent variable models. London: Chapman and Hall.
- Fan JM, Nuyujukian P, Cao JC, Chestek CA, Ryu SI, Shenoy KV (2014) Intention estimation in brain-machine interfaces. *Journal of neural engineering* 11:016004.
- Fetz EE (1992) Are movement parameters recognizably coded in the activity of single neurons? *Behavioral and Brain Sciences* 15:679-690.
- Flint RD, Wright ZA, Scheid MR, Slutzky MW (2013) Long term, stable brain machine interface performance using local field potentials and multiunit spikes. *Journal of neural engineering* 10:056005.
- Flint RD, Scheid MR, Wright ZA, Solla SA, Slutzky MW (2016) Long-term stability of motor cortical activity: implications for brain machine interfaces and optimal feedback control. *Journal of Neuroscience* 36:3623-3632.
- Fu Q, Flament D, Coltz J, Ebner T (1995) Temporal encoding of movement kinematics in the discharge of primate primary motor and premotor neurons. *Journal of Neurophysiology* 73:836-854.
- Gallego JA, Perich MG, Miller LE, Solla SA (2017a) Neural Manifolds for the Control of Movement. *Neuron* 94:978-984.
- Gallego JA, Perich MG, Naufel SN, Ethier C, Solla SA, Miller LE (2017b) Multiple tasks viewed from the neural manifold: Stable control of varied behavior. *bioRxiv*:176081.
- Gao P, Ganguli S (2015) On simplicity and complexity in the brave new world of large-scale neuroscience. *Current Opinion in Neurobiology* 32.
- Gao P, Trautmann E, Byron MY, Santhanam G, Ryu S, Shenoy K, Ganguli S (2017) A theory of multineuronal dimensionality, dynamics and measurement. *bioRxiv*:214262.



- Gao Y, Buesing L, Shenoy KV, Cunningham JP (2015) High-dimensional neural spike train analysis with generalized count linear dynamical systems. *Advances in Neural Information Processing System*:1-9.
- Gao Y, Archer E, Paninski L, Cunningham JP (2016) Linear dynamical neural population models through nonlinear embeddings.
- Georgopoulos AP, Kalaska JF, Caminiti R, Massey JT (1982) On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *The Journal of neuroscience : the official journal of the Society for Neuroscience* 2:1527-1537.
- Gilja V, Nuyujukian P, Chestek Ca, Cunningham JP, Yu BM, Fan JM, Churchland MM, Kaufman MT, Kao JC, Ryu SI, Shenoy KV (2012) A high-performance neural prosthesis enabled by control algorithm design. *Nature neuroscience* 15:1752-1757.
- Gilja V, Pandarinath C, Blabe CH, Nuyujukian P, Simeral JD, Sarma AA, Sorice BL, Perge JA, Jarosiewicz B, Hochberg LR, Shenoy KV, Henderson JM (2015) Clinical translation of a high-performance neural prosthesis. *Nature Medicine* 21.
- Golub MD, Byron MY, Chase SM (2015) Internal models for interpreting neural population activity during sensorimotor control. *Elife* 4:e10015.
- Golub MD, Sadtler PT, Oby ER, Quick KM, Ryu SI, Tyler-Kabara EC, Batista AP, Chase SM, Yu BM (2018) Learning by neural reassociation. *Nat Neurosci* 21.
- Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A, Bengio Y (2014) Generative adversarial nets. In: *Advances in neural information processing systems*, pp 2672-2680.
- Gribble PL, Scott SH (2002) Overlap of internal models in motor cortex for mechanical loads during reaching. *Nature* 417:938.
- Harvey CD, Coen P, Tank DW (2012) Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature* 484:62-68.
- Hennequin G, Vogels TP, Gerstner W (2014) Optimal control of transient dynamics in balanced networks supports generation of complex movements. *Neuron* 82:1394-1406.
- Hepp-Reymond M-C, Kirkpatrick-Tanner M, Gabernet L, Qi H-X, Weber B (1999) Context-dependent force coding in motor and premotor cortical areas. *Experimental brain research* 128:123-133.
- Hochberg LR, Serruya MD, Friehs GM, Mukand Ja, Saleh M, Caplan AH, Branner A, Chen D, Penn RD, Donoghue JP (2006) Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature* 442:164-171.
- Hochberg LR, Bacher D, Jarosiewicz B, Masse NY, Simeral JD, Vogel J, Haddadin S, Liu J, Cash SS, van der Smagt P, Donoghue JP (2012) Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. *Nature* 485:372-375.
- Hocherman S, Wise S (1991) Effects of hand movement path on motor cortical activity in awake, behaving rhesus monkeys. *Experimental brain research* 83:285-302.
- Holdefer R, Miller L (2002) Primary motor cortical neurons encode functional muscle synergies. *Experimental Brain Research* 146:233-243.
- Hubel DH, Wiesel TN (1959) Receptive fields of single neurones in the cat's striate cortex. *The Journal of physiology* 148:574-591.
- Jackson JH (1873) Observations on the localisation of movements in the cerebral hemispheres, as revealed by cases of convulsion, chorea, and "aphasia". *The West Riding Lunatic Asylum Medical Reports* 3:175-195.
- Jankowska E, Padel Y, Tanaka R (1975) Projections of pyramidal tract cells to alpha-motoneurons innervating hind-limb muscles in the monkey. *The Journal of Physiology* 249:637-667.
- Jarosiewicz B, Sarma AA, Bacher D, Masse NY, Simeral JD, Sorice B, Oakley EM, Blabe C, Pandarinath C, Gilja V, Cash SS, Eskandar EN, Friehs G, Henderson JM, Shenoy KV, Donoghue JP, Hochberg LR (2015) Virtual typing by people with tetraplegia using a self-calibrating intracortical brain-computer interface. *Science Translational Medicine* 7.
- Jun JJ, Steinmetz NA, Siegle JH, Denman DJ, Bauza M, Barbarits B, Lee AK, Anastassiou CA, Andrei A, Aydin Ç (2017) Fully integrated silicon probes for high-density recording of neural activity. *Nature* 551:232.



- Kao JC, Ryu SI, Shenoy KV (2017) Leveraging neural dynamics to extend functional lifetime of brain-machine interfaces. *Scientific reports* 7:7395.
- Kao JC, Nuyujukian P, Ryu SI, Churchland MM, Cunningham JP, Shenoy KV (2015) Single-trial dynamics of motor cortex and their applications to brain-machine interfaces. *Nature communications* 6.
- Kaufman MT, Churchland MM, Shenoy KV (2013) The roles of monkey M1 neuron classes in movement preparation and execution. *Journal of neurophysiology* 110:817-825.
- Kaufman MT, Churchland MM, Ryu SI, Shenoy KV (2014) Cortical activity in the null space: permitting preparation without movement. *Nature neuroscience* 17:440-448.
- Kaufman MT, Churchland MM, Ryu SI, Shenoy KV (2015) Vacillation, indecision and hesitation in moment-by-moment decoding of monkey motor cortex. *Elife* 4:e04677.
- Kaufman MT, Seely JS, Sussillo D, Ryu SI, Shenoy KV, Churchland MM (2016) The largest response component in motor cortex reflects movement timing but not movement type. *eNeuro* 3:ENEURO.0085-0016.2016.
- Kaufman MT, Churchland MM, Santhanam G, Yu BM, Afshar A, Ryu SI, Shenoy KV (2010) Roles of monkey premotor neuron classes in movement preparation and execution. *Journal of neurophysiology* 104:799-810.
- Kim S-P, Simeral JD, Hochberg LR, Donoghue JP, Black MJ (2008) Neural control of computer cursor velocity by decoding motor cortical spiking activity in humans with tetraplegia. *Journal of neural engineering* 5:455-476.
- Kobak D, Brendel W, Constantinidis C, Feierstein CE, Kepecs A, Mainen ZF, Rumo R, Qi X-L, Uchida N, Machens CK (2014) Demixed principal component analysis of population activity in higher cortical areas reveals independent representation of task parameters. *arXiv*.
- Lakshmanan KC, Sadtler PT, Tyler-Kabara EC, Batista AP, Yu BM (2015) Extracting low-dimensional latent structure from time series in the presence of delays. *Neural computation* 27:1825-1856.
- Landgren S, Phillips C, Porter R (1962) Minimal synaptic actions of pyrimidal impulses on some alpha motoneurons of the baboon's hand and forearm. *The Journal of physiology* 161:91-111.
- Linderman S, Johnson M, Miller A, Adams R, Blei D, Paninski L (2017) Bayesian learning and inference in recurrent switching linear dynamical systems. In: *Artificial Intelligence and Statistics*, pp 914-922.
- Macke JH, Buesing L, Cunningham JP, Yu BM, Shenoy KV, Sahani M (2011) Empirical models of spiking in neuronal populations. In *Advances in Neural Information Processing Systems (NIPS)* 24:1-9.
- Mante V, Sussillo D, Shenoy KV, Newsome WT (2013) Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* 503:78-84.
- Michaels JA, Dann B, Scherberger H (2016) Neural Population Dynamics during Reaching Are Better Explained by a Dynamical System than Representational Tuning. *PLOS Computational Biology* 12:e1005175.
- Michaels JA, Dann B, Intveld RW, Scherberger H (2015) Predicting reaction time from the neural state space of the premotor and parietal grasping network. *Journal of Neuroscience* 35:11415-11432.
- Michaels JA, Dann B, Intveld RW, Scherberger H (2018) Neural dynamics of variable grasp movement preparation in the macaque fronto-parietal network. *Journal of Neuroscience*:2557-2517.
- Miri A, Warriner CL, Seely JS, Elsayed GF, Cunningham JP, Churchland MM, Jessell TM (2017) Behaviorally Selective Engagement of Short-Latency Effector Pathways by Motor Cortex. *Neuron*:1-14.
- Molano-Mazon M, Onken A, Piasini E, Panzeri S (2018) Synthesizing realistic neural population activity patterns using Generative Adversarial Networks. *arXiv preprint arXiv:180300338*.
- Morcos AS, Harvey CD (2016) History-dependent variability in population dynamics during evidence accumulation in cortex. *Nature neuroscience* 19:1672-1680.
- Mussa-Ivaldi F (1988) Do neurons in the motor cortex encode movement direction? An alternative hypothesis. *Neurosci Lett* 91:106-111.
- Nuyujukian P, Kao JC, Fan JM, Stavisky SD, Ryu SI, Shenoy KV (2014) Performance sustaining intracortical neural prostheses. *Journal of Neural Engineering* 11:066003.
- Orsborn AL, Dangi S, Moorman HG, Carmena JM (2012) Closed-loop decoder adaptation on intermediate time-scales facilitates rapid BMI performance improvements independent of decoder

- initialization conditions. *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 20:468-477.
- Pandarinath C, Nuyujukian P, Blabe CH, Sorice BL, Saab J, Willett FR, Hochberg LR, Shenoy KV, Henderson JM (2017) High performance communication by people with paralysis using an intracortical brain-computer interface. *eLife* 6:1-27.
- Pandarinath C, Gilja V, Blabe CH, Nuyujukian P, Sarma AA, Sorice BL, Eskandar EN, Hochberg LR, Henderson JM, Shenoy KV (2015) Neural population dynamics in human motor cortex during movements in people with ALS. *eLife* 4.
- Pandarinath C, O'Shea DJ, Collins J, Jozefowicz R, Stavisky SD, Kao JC, Trautmann EM, Kaufman MT, Ryu SI, Hochberg LR, Henderson JM, Shenoy KV, Abbott LF, Sussillo D (2018) Inferring single-trial neural population dynamics using sequential auto-encoders. *Nature Methods* in press.
- Perge Ja, Homer ML, Malik WQ, Cash S, Eskandar E, Friebs G, Donoghue JP, Hochberg LR (2013) Intra-day signal instabilities affect decoding performance in an intracortical neural interface system. *Journal of neural engineering* 10:036004.
- Perich MG, Gallego JA, Miller LE (2017) A neural population mechanism for rapid learning. *bioRxiv*:138743.
- Petreska B, Yu B, Cunningham J, Santhanam G, Ryu S, Shenoy K, Sahani M (2011) Dynamical segmentation of single trials from population neural data. *Nips*:756-764.
- Phillips C (1975) Laying the ghost of 'muscles versus movements'. *Canadian Journal of Neurological Sciences* 2:209-218.
- Pohlmeyer EA, Solla SA, Perreault EJ, Miller LE (2007) Prediction of upper limb muscle activity from motor cortical discharge during reaching. *Journal of neural engineering* 4:369.
- Raposo D, Kaufman MT, Churchland AK (2014) A category-free neural population supports evolving demands during decision-making. *Nature neuroscience* 17:1784.
- Remington ED, Narain D, Hosseini EA, Jazayeri M (2018) Flexible sensorimotor computations through rapid reconfiguration of cortical dynamics. *Neuron* 98:1005-1019. e1005.
- Riehle A, Requin J (1989) Monkey primary motor and premotor cortex: single-cell activity related to prior information about direction and extent of an intended movement. *Journal of neurophysiology* 61:534-549.
- Russo AA, Bittner SR, Perkins SM, Abbott LF, Cunningham JP, Churchland MM, Russo AA, Bittner SR, Perkins SM, Seely JS, London BM, Lara AH (2018) Motor Cortex Embeds Muscle-like Commands in an Untangled Population Response. *Neuron*:1-14.
- Sadtler PT, Quick KM, Golub MD, Chase SM, Ryu SI, Tyler-kabara EC, Yu BM, Batista AP (2014) Neural constraints on learning. *Nature* in press:423-426.
- Santhanam G, Yu BM, Gilja V, Ryu SI, Afshar A, Sahani M, Shenoy KV (2009) Factor-analysis methods for higher-performance neural prostheses. *Journal of neurophysiology* 102:1315-1330.
- Schwartz AB, Kettner RE, Georgopoulos AP (1988) Primate motor cortex and free arm movements to visual targets in three-dimensional space. I. Relations between single cell discharge and direction of movement. *Journal of Neuroscience* 8:2913-2927.
- Scott SH (2004) Optimal feedback control and the neural basis of volitional motor control. *Nature Reviews Neuroscience* 5:532-546.
- Scott SH (2008) Inconvenient truths about neural processing in primary motor cortex. *The Journal of physiology* 586:1217-1224.
- Shenoy KV, Sahani M, Churchland MM (2013) Cortical control of arm movements: a dynamical systems perspective. *Annual review of neuroscience* 36:337-359.
- Sofroniew NJ, Flickinger D, King J, Svoboda K (2016) A large field of view two-photon mesoscope with subcellular resolution for in vivo imaging. *Elife* 5:e14472.
- Stavisky SD, Kao JC, Ryu SI, Shenoy KV (2017a) Trial-by-trial motor cortical correlates of a rapidly adapting visuomotor internal model. *Journal of Neuroscience*:1091-1016.
- Stavisky SD, Kao JC, Nuyujukian P, Ryu SI, Shenoy KV (2015) A high performing brain-machine interface driven by low-frequency local field potentials alone and together with spikes. *Journal of neural engineering* 12:036009.
- Stavisky SD, Kao JC, Ryu SI, Shenoy KV, Stavisky SD, Kao JC, Ryu SI, Shenoy KV (2017b) Motor Cortical Visuomotor Feedback Activity Is Initially Isolated from Downstream Targets in Output-

- Null Neural State Space Dimensions Article Motor Cortical Visuomotor Feedback Activity Is Initially Isolated from Downstream Targets in Output-Null Neural Neuron:1-14.
- Stevenson IH, Kording KP (2011) How advances in neural recording affect data analysis. *Nature neuroscience* 14:139.
- Stringer C, Pachitariu M, Steinmetz N, Reddy CB, Carandini M, Harris KD (2018) Spontaneous behaviors drive multidimensional, brain-wide population activity. *bioRxiv*:306019.
- Sussillo D, Churchland MM, Kaufman MT, Shenoy KV (2015) A neural network that finds a naturalistic solution for the production of muscle activity. *Nature neuroscience* 18:1025-1033.
- Sussillo D, Jozefowicz R, Abbott LF, Pandarinath C (2016a) LFADS - Latent Factor Analysis via Dynamical Systems. *arXiv*.
- Sussillo D, Stavisky SD, Kao JC, Ryu SI, Shenoy KV (2016b) Making brain-machine interfaces robust to future neural variability. *Nature communications* 7:13749.
- Tanji J, Evarts EV (1976) Anticipatory activity of motor cortex neurons in relation to direction of an intended movement. *Journal of Neurophysiology* 39:1062-1068.
- Taylor DM, Tillery SIH, Schwartz AB (2002) Direct cortical control of 3D neuroprosthetic devices. *Science* 296:1829-1832.
- Todorov E, Jordan MI (2002) Optimal feedback control as a theory of motor coordination. *Nature neuroscience* 5:1226-1235.
- Velliste M, Perel S, Spalding MC, Whitford aS, Schwartz aB (2008) Cortical control of a robotic arm for self-feeding. *Nature* 453:1098-1101.
- Vyas S, Even-Chen N, Stavisky SD, Ryu SI, Nuyujukian P, Shenoy KV (2018) Neural population dynamics underlying motor learning transfer. *Neuron* 97:1177-1186. e1173.
- Wei Z, Inagaki H, Li N, Svoboda K, Druckmann S (2018) An orderly single-trial organization of population dynamics in premotor cortex predicts behavioral variability. *bioRxiv*:376830.
- Williams AH, Kim TH, Wang F, Vyas S, Ryu SI, Shenoy KV, Schnitzer M, Kolda TG, Ganguli S (2018) Unsupervised Discovery of Demixed, Low-Dimensional Neural Dynamics across Multiple Timescales through Tensor Component Analysis. *Neuron*.
- Wu W, Black MJ, Gao Y, Serruya M, Shaikhouni A, Donoghue J, Bienenstock E (2003) Neural decoding of cursor motion using a Kalman filter. In: *Advances in neural information processing systems*, pp 133-140.
- Yu B, Cunningham J, Santhanam G, Ryu S, Shenoy K, Sahani M (2009) Gaussian-Process Factor Analysis for Low-Dimensional Single-Trial Analysis of Neural Population Activity. *Journal of neurophysiology* 102:614-635.
- Yuste R (2015) From the neuron doctrine to neural networks. *Nature Reviews Neuroscience* 16.
- Zhao Y, Park IM (2017) Variational latent Gaussian process for recovering single-trial dynamics from population spike trains. *Neural computation* 29:1293-1316.