

Neural Manifolds for the Control of Movement

Juan A. Gallego,^{1,2} Matthew G. Perich,³ Lee E. Miller,^{1,3,4} and Sara A. Solla^{1,5,*}

¹Department of Physiology, Northwestern University, Chicago, IL 60611, USA

²Neural and Cognitive Engineering Group, Centre for Robotics and Automation CSIC-UPM, Arganda del Rey 28500, Spain

³Department of Biomedical Engineering, Northwestern University, Evanston, IL 60208, USA

⁴Department of Physical Medicine and Rehabilitation, Northwestern University, Chicago, IL 60611, USA

⁵Department of Physics and Astronomy, Northwestern University, Evanston, IL 60208, USA

*Correspondence: solla@northwestern.edu

<http://dx.doi.org/10.1016/j.neuron.2017.05.025>

The analysis of neural dynamics in several brain cortices has consistently uncovered low-dimensional manifolds that capture a significant fraction of neural variability. These neural manifolds are spanned by specific patterns of correlated neural activity, the “neural modes.” We discuss a model for neural control of movement in which the time-dependent activation of these neural modes is the generator of motor behavior. This manifold-based view of motor cortex may lead to a better understanding of how the brain controls movement.

Since the work of Herbert Jasper (Jasper et al., 1958) and Ed Evarts (Evarts, 1968), cortical function has been studied by recording single-neuron activity while animals perform a variety of behaviors, including decision making (Newsome et al., 1989), sensation (Wurtz, 1969), and movement (Georgopoulos et al., 1982; Humphrey et al., 1970).

In the motor system, the main focus of this article, single neuron studies typically involved recordings during repeated, stereotypical movements. Many of these experiments sought explicit representations that relate single-neuron activity to specific movement covariates, including but not limited to target position, endpoint and joint kinematics, endpoint forces, and muscle activity (Evarts, 1968; Georgopoulos et al., 1982; Humphrey et al., 1970; Morrow et al., 2007; Thach, 1978). Although some of these efforts involved the decoding of population activity (Georgopoulos et al., 1982), they were restricted to models of non-interacting neurons whose individual activity was associated with specific movement covariates.

However, some of these studies also identified single neurons whose activity did not represent movement parameters (Churchland and Shenoy, 2007; Fetz, 1992; Scott, 2008). If neurons in primary motor cortex (M1) were to represent movement parameters, those representations ought to be most evident in corticomotoneuronal (CM) cells, which make direct connections onto spinal motoneurons (Fetz, 1992). Yet, many of these CM cells do not represent any specific movement covariate (Fetz et al., 1989).

The ultimate role of M1 is to generate movement, not to represent it (Churchland et al., 2012; Cisek, 2006; Scott, 2004); thus, it is not surprising that many M1 neurons do not relate to any single movement covariate. The search for representations at the single-neuron level might actually divert us from understanding the neural control of movement. Early neural network simulations indicated that individual neurons need not explicitly encode movement covariates when the goal of M1 population activity is to generate realistic muscle activation patterns (Fetz, 1992).

The role of neurons that do not explicitly represent any movement covariate can be explained by recent work based on optimal feedback control theory, which postulates that the goal

of motor cortex is to produce a desired movement and force, taking into account the state of the muscles. This hypothesis avoids the need for explicit representation of movement covariates by single neurons, though some neurons may still represent movement covariates or high-level task characteristics as a byproduct of the necessary control signals (Scott, 2008; Todorov, 2000).

Recent and accelerating technical developments provide the experimental tools for monitoring the activity of large numbers of neurons, as well as the statistical and modeling tools for analyzing how these neural populations perform the computations necessary to plan and execute movement (Gao and Ganguli, 2015). The challenge of understanding the neural control of movement by analyzing neural population activity is formidable, as population activity in any specific area not only reflects its intrinsic dynamics, but must also respond to its inputs and generate output projections based on the computations being performed (Sussillo et al., 2015). A simplification arises from the fact that neural computations are based on the joint activity of interconnected neurons (Fetz, 1992; Hatsopoulos et al., 1998; Shenoy et al., 2013); the resulting population activity is thus likely constrained by the connectivity of the underlying network.

Here we argue that the underlying network connectivity constrains these possible patterns of population activity (Okun et al., 2015; Sadtler et al., 2014; Tsodyks et al., 1999) and that the possible patterns are confined to a low-dimensional manifold (Stopfer et al., 2003; Yu et al., 2009) spanned by a few independent patterns that we call “neural modes.” These neural modes capture a significant fraction of population covariance. It is the activation of these neural modes, rather than the activity of single neurons, that provides the basic building blocks of neural dynamics and function (Luczak et al., 2015; Sadtler et al., 2014; Shenoy et al., 2013).

We thus propose a generative model of the activity of individual neurons based on the activation of neural modes, and explain how the parameters of the model can be identified using dimensionality reduction methods. We then review work showing that these neural modes span task-specific neural manifolds in




Figure 1. The Neural Manifold Hypothesis

(A) Latent variables as a generative model for population activity. The relative area of the blue/green regions in each neuron represents the relative magnitude of the contribution of each latent variable to the neuron's activity. (B) Spikes from three recorded neurons during task execution as a linear combination of two latent variables. (C) Trajectory of time-varying population activity in the neural space of the three recorded neurons (black). The trajectory is mostly confined to the neural manifold, a plane shown in gray and spanned by the neural modes u_1 and u_2 . (D) A curved, nonlinear neural manifold, shown in blue. Linear methods would capture a flat, local approximation to a small task-specific region of the manifold. (E) Linear manifolds for two different tasks shown as gray and purple planes. Are these two planes local linear approximations to different regions within a large, continuous manifold (transparent surface with blue contour), or are they distinct task-specific manifolds that may or not share neural modes?

premotor and motor cortices. We propose that neural manifolds spanned by a surprisingly small number of neural modes are likely to simplify the neural control of movement and speculate on the potential learning mechanisms underlying the emergence of this low-dimensional organization.

From Single Neurons to Neural Manifolds

Current multi-electrode arrays (MEAs) allow for the simultaneous recording of about a hundred neurons. This is much more than the small numbers recorded with single electrodes but still a tiny fraction of the total number of neurons involved in movement generation. Despite this limitation, brain-machine interfaces (BMIs) based on these MEAs are able to predict reasonably well many behavioral variables (Carmena et al., 2003; Ethier et al., 2012; Serruya et al., 2002).

What is the underlying reason for this success? Intuitively, it is the high degree of correlation and redundancy across individual neural activity. This intuition has been recently made precise in elegant arguments on the low dimensionality of the stereotypical motor behaviors used in most motor control studies (Gao and Ganguli, 2015). The relatively small number of independent signals needed to control behavior during the execution of such tasks only requires a small number of independent

neural signals. These neural signals are the “latent variables” (Cunningham and Yu, 2014) that describe the dynamics of the “neural modes.”

The participation of individual neurons in neural modes is illustrated in Figure 1A. Note that each neural mode includes a large fraction of the neurons in the population and that a given neuron can participate in several neural modes. In this view, the time-dependent activity of individual neurons is simply a reflection of the latent variables (Figure 1B) (Kaufman et al., 2016; Kobak et al., 2016; Macke et al., 2011). Consider the “neural space” in Figure 1C; each axis represents the activity of one of the N recorded neurons (here, $N = 3$). Assuming that network connectivity constrains the possible patterns of population activity (Okun et al., 2015; Sadler et al., 2014; Tsodyks et al., 1999), the population dynamics will not explore the full high-dimensional neural space but will instead remain confined to a low-dimensional surface within the full space, the “neural manifold.” In the simplest linear case, the neural manifold is flat, as is the hyperplane in Figure 1C, spanned by the two neural modes, u_1 and u_2 .

This geometrical picture illustrates a possible generative model for the dynamics of individual neurons: the activity $n_i(t)$ of the i th neuron, $1 \leq i \leq N$, results from a linear combination of latent variables $L_j(t)$ plus additive noise ε_i :

$$n_i(t) = \sum_j u_{ij} L_j(t) + \varepsilon_i. \quad \text{Equation 1}$$

Here, $L_j(t)$ is the j th latent variable, the time-dependent activation of the j th neural mode. Each latent variable results from projecting the neural population activity onto the corresponding neural mode. The coefficient u_{ij} in the linear combination quantifies the contribution of the j th latent variable to the activity of the i th neuron. These “participation weights” relate to the internal connectivity of the network (Okun et al., 2015). The noise term ε_i represents intrinsic neural noise, and potentially other processes not accounted for in the model. By construction, neural population activity remains within the neural manifold except for small fluctuations (see how close the actual black trajectory is to the gray trajectory projected into the manifold in Figure 1C).

Dimensionality reduction techniques allow us to study neural population dynamics by finding a set of neural modes that span the neural manifold and identify relevant population features (Cunningham and Yu, 2014). Common linear techniques for dimensionality reduction, such as principal component analysis (PCA) and factor analysis (FA), identify neural modes as dominant patterns of covariation across neurons and yield the parameters of the generative model (Equation 1; Figure 1B).

As an illustration, we show that neural data recorded during an isometric wrist task (Figures 2A and 2B) is largely accounted for by the latent variables in Figure 2C. The low dimensionality of the neural manifold follows from the rapid increase of the explained variance with the number of neural modes (Figure 2D).

Neural Manifolds: A Framework to Study Neural Control of Movement

The concept of the neural manifold and its associated latent variables has been used in a series of recent studies that abandon the search for movement representation by single neurons to consider instead movement planning and execution based

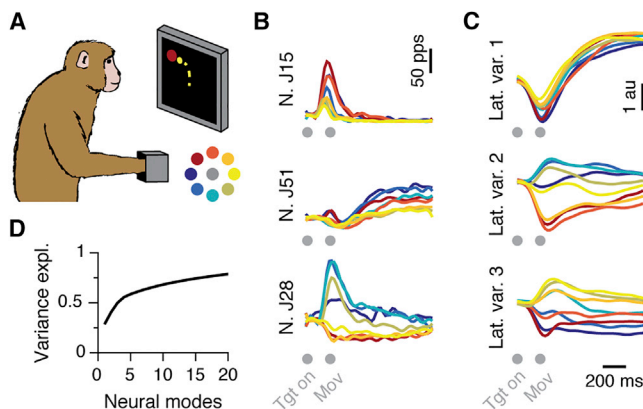


Figure 2. Latent Variables for an Isometric Center-Out Task

(A) Monkey performing an isometric eight-target center-out wrist task. The targets are color coded (inset). The same target color code is used for the curves shown in (B) and (C).

(B) Peristimulus time histogram for three out of the 68 neurons recorded in a single session.

(C) Latent variables track the activation of the three leading neural modes, identified with PCA. In (B) and (C), target presentation ("Tgt on") is followed by movement onset ("Mov") after about 140 ms.

(D) Variance explained as function of the number of PCA neural modes. Ten modes sufficed to account for about 70% of the variance.

on the activation of a few neural modes (Ahrens et al., 2012; Bruno et al., 2015; Churchland et al., 2010a, 2010b, 2012; Churchland and Shenoy, 2007; Elsayed et al., 2016; Kaufman et al., 2014, 2016; Michaels et al., 2016; Overduin et al., 2015; Sadtler et al., 2014; Santhanam et al., 2009; Sussillo et al., 2015).

One of the earliest findings of a neural manifold for movement control comes from Shenoy and colleagues (Santhanam et al., 2009), who analyzed population activity recorded with an MEA implanted in the arm area of dorsal premotor cortex (PMd) during a delayed center-out reach task. Single-neuron activity in PMd correlates with the direction toward the end point of an upcoming reach movement (Riehle and Requin, 1989; Shen and Alexander, 1997). Shenoy and colleagues used FA to obtain neural modes that accounted for the observed shared variance of individual neurons. They found that a three-dimensional manifold sufficed to identify target-specific clusters of latent activity during the delay period (Figure 3A).

A subsequent study (Churchland et al., 2010b) showed a systematic decrease in the trial-to-trial variability in the neural dynamics of both PMd and primary visual cortex (V1) following stimulus onset, as demonstrated in two-dimensional visualizations of the latent variables (Figure 3B). The low-dimensional manifold was characterized using Gaussian process factor analysis (GPFA), a method that combines FA with temporal smoothing through a Gaussian kernel, to extract the low-dimensional trajectories defined by the latent variables during individual trials. The method was proposed and compared to static methods like PCA and FA in an earlier paper (Yu et al., 2009) that identified variability reduction following target presentation in PMd data.

The notion of a neural manifold and its associated latent variables was subsequently used by Churchland, Shenoy, and colleagues (Churchland et al., 2010a, 2012) to explain how neural activity in both PMd and M1 during movement planning (Riehle

and Requin, 1989) does not generate movement during the delay preceding a go signal (Cisek and Kalaska, 2005). To explain how M1 could prepare movement without causing it, the same group (Kaufman et al., 2014) identified a six-dimensional neural manifold using PCA, then built a linear model that related these latent variables to three "muscle synergies" (d'Avella et al., 2003; Tresch and Jarc, 2009), also identified by PCA. Based on this linear model, they divided the neural manifold into a "potent" space, whose activity controls muscle activity, and a "null" space, whose activity does not affect muscle activity (Kaufman et al., 2014) (Figure 3C). They showed that preparatory activity lies in the null space; this condition-dependent activity provides an initialization from which the population dynamics evolve to generate the desired movement (Churchland et al., 2010a, 2012; Kaufman et al., 2014; Shenoy et al., 2013). In a recent follow-up study, the same group expanded this analysis to show that preparatory and movement activity lie in orthogonal spaces within the manifold and that population dynamics evolve from one to the other (Elsayed et al., 2016) (Figure 3D).

The separation between potent and null spaces was also used by Slutzky and colleagues to investigate the long-term stability of BMIs (Flint et al., 2016). They found that the stability of all recorded neurons was not uniformly necessary to achieve stable BMI control and showed that neural activity in the potent space was significantly more stable than neural activity in the null space. This finding provided evidence supporting the postulate that optimal feedback control allows the brain to control activity in the potent space, while activity in the null space can vary from trial to trial (Todorov and Jordan, 2002).

Together, these studies strongly support the existence of low-dimensional manifolds in motor cortices. The notion that latent variables may constitute the building blocks of population activity allows us to consider the activity of individual neurons as one-dimensional samples of the population dynamics. These studies also suggest that the constraints embodied by the neural manifolds simplify movement generation by providing a small number of signals that are independently controlled to achieve a desired behavior (Giszter, 2015; Mussa-Ivaldi and Solla, 2004; Overduin et al., 2015; Shenoy et al., 2013; Thoroughman and Shadmehr, 2000).

Emergence of Neural Manifolds through Learning

Neural manifolds embody patterns of correlated activity; we hypothesize that these correlations reflect the connectivity of the underlying network of neurons. Since long-term learning can alter cortical connectivity (Fu et al., 2012; Rioult-Pedotti et al., 1998), we address potential connections between learning and the emergence of neural manifolds.

A revealing connection between neural manifolds and learning was discovered by Batista, Yu, and colleagues (Sadtler et al., 2014), who used a BMI paradigm in monkeys to address the question of why some motor skills are easier to learn than others. They used FA to identify a ten-dimensional neural manifold and built a linear decoder from these latent variables into a two-dimensional center-out task. Once the monkeys had learned the task, the decoder was modified in one of two ways; see Figure 4A. When the modification required the use of existing neural modes (a within-manifold perturbation,

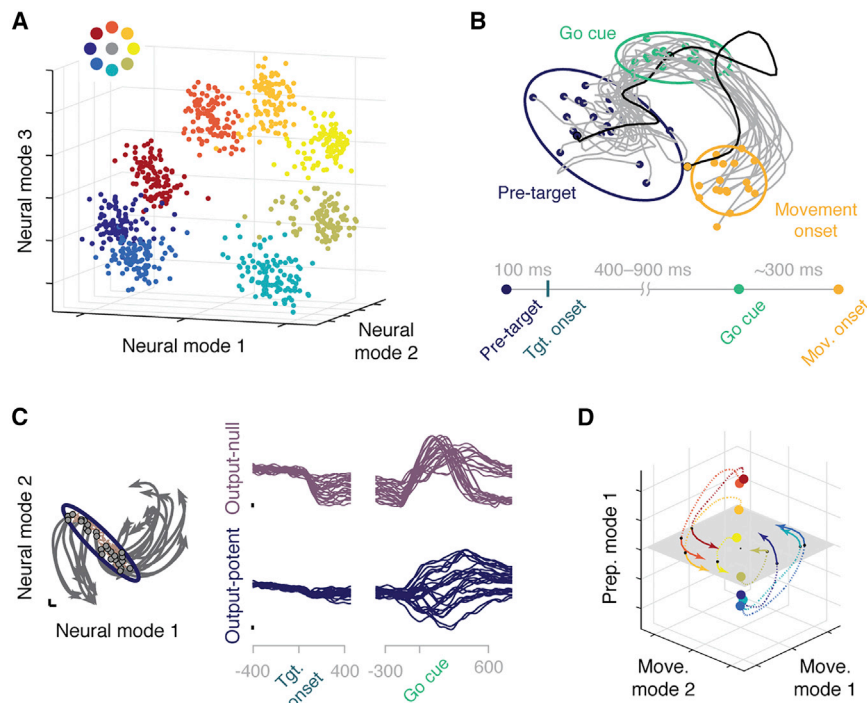


Figure 3. Neural Modes in Motor Cortices
(A) Preparatory activity in PMd for an eight-target reach task. A neural manifold spanned by three neural modes reveals target-specific clusters. The targets are color coded (inset). Each point represents activity measured during the delay period for one trial. Adapted from [Santhanam et al. \(2009\)](#).
(B) Population variability in PMd is reduced by stimulus presentation. For a delayed reach task (timeline at bottom of panel), population activity trajectories for individual trials are shown in gray in a two-dimensional manifold. The “Pre-target” confidence ellipsoid shrinks during “Go cue” as trajectories get closer together. Adapted from [Churchland et al. \(2010b\)](#).
(C) The null space allows for movement preparation without execution. Population activity trajectories for each reach condition are shown in gray in a two-dimensional manifold. Preparatory activity sets the corresponding initial conditions (gray circles within purple ellipsoid). Null (purple, top) and potent (dark blue, bottom) latent variables, defined with respect to EMG activity, are shown on the right. Adapted from [Kaufman et al. \(2014\)](#).
(D) Neural modes associated with movement preparation and execution span different spaces. Two movement modes span a plane (in gray), while the orthogonal preparation mode spans a line that contains the initial conditions (colored circles). Projections of the full trajectories onto the movement manifold resemble the traces in (C). Trajectories are color coded for each target, as in (A). Adapted from [Elsayed et al. \(2016\)](#).

Figure 4B), the monkeys easily adapted in a single session. In contrast, when the modification required the acquisition of new neural modes (an outside-manifold perturbation, Figure 4C), the task proved significantly harder to learn. The same group subsequently showed that monkeys could learn new neural modes (E.R. Oby et al., 2015, Soc. Neurosci., abstract) if they were guided to generate these new patterns progressively, over many days. These results suggest that short-term adaptation may be based on the generation of new combinations of preexisting neural modes, while long-term learning may require generating new neural modes. If neural modes arise from neural circuitry that constrains spatiotemporal patterns of activity, it is not surprising that modifying a neural manifold is more difficult than simply learning to use an existing neural manifold in novel ways.

A connection between the emergence of neural manifolds and learning also arises from a simulation ([Sussillo et al., 2015](#)) in which a recurrent neural network (RNN) was trained to output the correct activity pattern for seven muscles recorded during a reach task ([Churchland et al., 2012](#)) (Figure 4D). In the RNN model, training refers to learning the connections between neurons to obtain a network able to perform the desired function, in analogy with the modification of synaptic connectivity in biological learning. Latent variables were then identified for the data generated by the trained RNN and compared to those for the experimental data. The experimental data were best reproduced by the simplest RNN connectivity that allowed the network to output the measured muscle activity. A similar result was subsequently achieved with an RNN trained to output the horizontal and vertical components of hand velocity for a delayed center-out reach task ([Michaels et al., 2016](#)).

In a complementary simulation study ([Hennequin et al., 2014](#)), Gerstner and colleagues proposed a neural network model with random excitatory recurrent connections and inhibitory connections that were learned to achieve stable population dynamics. This balanced network generated population activity patterns whose associated latent variables also exhibited the damped oscillations experimentally observed in the latent variables during reaching ([Churchland et al., 2012](#)). It is quite interesting that a network model not trained to produce a specific output—kinematics or muscle activity—but to stabilize its internal dynamics also exhibits oscillatory latent variables.

Neural Manifolds in Non-motor Brain Cortices

Although we have focused on neural manifolds in motor cortices, it is important to remark that neural manifolds seem to be widely present across the brain. A number of studies have shown that the largely heterogeneous activity patterns of individual neurons in monkey ([Kobak et al., 2016](#); [Machens et al., 2010](#); [Mante et al., 2013](#); [Markowitz et al., 2015](#)) and rat ([Durstewitz et al., 2010](#)) prefrontal cortex, monkey ([Churchland et al., 2010b](#); [Cowley et al., 2016](#)) and rat ([Forsberg et al., 2016](#)) V1, rat olfactory cortex ([Kobak et al., 2016](#)), rat thalamus ([Chapin and Nicolelis, 1999](#)), rat parietal cortex ([Raposo et al., 2014](#)), locust olfactory system ([Stopfer et al., 2003](#)), *Aplysia* pedal ganglion ([Bruno et al., 2015](#)), and perhaps the entire zebrafish brain ([Ahrens et al., 2012](#)) can be explained as generated by a small set of latent variables associated with neural modes. In all these studies, neural modes and their time-varying activation helped describe previously unexplained mechanisms of neural function.

Studies of sensory cortices have provided clear evidence of network connectivity constraints on the activity of individual




Figure 4. Neural Manifolds and Learning

(A) Existing neural modes facilitate the adaptation to variations of a learned task. In a BMI paradigm, ten latent variables representing the recorded activity of M1 neurons were mapped onto x and y components of cursor velocity. (B and C) Once monkeys were proficient with the BMI, a (B) within-manifold or (C) outside-manifold perturbation was imposed on the decoder inputs. Adaptation to (B) required a change in the relative activation of existing neural modes, while (C) required the acquisition of new neural modes through changes in neural comodulation patterns. Adapted from Sadtler et al. (2014). (D) A recurrent neural network was trained to generate condition-specific EMG patterns selected through inputs representing preparatory activity. The latent variables that described the population activity of the trained recurrent network closely resembled those associated with the experimental data. Adapted from Sussillo et al. (2015).

neurons. In the cat V1, the instantaneous activity of strongly tuned neurons is tightly linked to the population activity measured with optical imaging, both in response to stimulus presentation and during spontaneous activity (Tsodyks et al., 1999). Experiments in auditory and somatosensory cortices of awake and anesthetized rats further examined the relation between spontaneous and stimulus-evoked activity (Luczak et al., 2009) and found a surprising degree of conservation across these distinct regimes. The authors identified neural modes for spontaneous population activity and found that neural modes in response to stimuli lay within the same neural manifold. Interestingly, the evoked responses sampled a smaller portion of the manifold than the spontaneous activity, an organization of population activity also found in monkey V1 (Cowley et al., 2016).

Recent experiments in both mouse and monkey V1, and in rat auditory cortex demonstrated that the correlation between single-neuron activity and population activity is the same during both spontaneous and stimulus-evoked activity (Okun et al., 2015). These authors demonstrated that correlations between individual neurons and the population significantly predicted all pairwise correlations among the neurons. The degree of correla-

tion between individual neurons and the population displayed three interesting properties: (1) it predicted the responses of individual neurons during optogenetic stimulation of the surrounding population, suggesting that the relationship between population activity and that of a single neuron is causal; (2) it predicted the estimated number of synapses that a neuron receives from its neighbors, relating activity correlations to the underlying network structure; (3) it correlated strongly with the participation weight of the first FA factor on the activity of each neuron. Together, these results provide convincing if not direct evidence that network connectivity underlies the interactions among neurons captured by dimensionality reduction methods and the resulting neural modes.

Open Questions

The prevalence of neural manifolds across brain areas highlights their importance as a tool to understand brain function and reinforces the idea that neural processing may be built upon latent variables rather than on the activity of single neurons. The concept of neural manifolds is not restricted to flat surfaces; the manifold might well be a nonlinear surface within the neural space (Figure 1D). However, for stereotypical laboratory tasks such as those discussed here, linear methods such as PCA or FA capture a large fraction of the neural covariance. If the task-specific neural dynamics explore only a limited region within the manifold, a local linear approximation to a nonlinear neural manifold would work quite well (see the gray and purple planes in Figure 1E). For complex behaviors whose dynamics explore a larger region of neural space, linear methods may provide poor estimates of the neural manifold. Nonlinear methods for dimensionality reduction such as locally linear embedding (LLE) (Roweis and Saul, 2000), isomap (Tenenbaum et al., 2000), or autoencoder neural networks (Hinton and Salakhutdinov, 2006) might be needed.

These nonlinear methods have already been used in non-motor brain areas. Analysis of population data from the locust antennal lobe during an odor identification task using LLE revealed odor-specific neural manifolds that contain trajectories corresponding to different concentrations of the same odor (Stopfer et al., 2003). The analysis of the population activity of retinal ganglion cells using isomap (Ganmor et al., 2015) identified activity clusters corresponding to similar visual stimuli. It remains an open question whether nonlinear methods might reveal mechanisms for the neural control of movement that linear methods have not revealed.

The studies discussed here focus on neural manifolds associated with specific tasks. Since organisms are able to execute a rich repertoire of motor tasks, how are the corresponding neural manifolds organized with respect to each other within the neural space? The neural space may contain distinct neural manifolds, each associated with a specific task. What then is the relationship among these distinct manifolds? Is each manifold spanned by its own unique neural modes, or do tasks in a class share some neural modes that represent common features? These are virtually unexplored questions; note however a recent suggestion that preparation and execution of a movement correspond to orthogonal but related manifolds (Elsayed et al., 2016). Whether a similar finding applies to manifolds corresponding to the execution of different tasks is yet unknown.

An alternative is that all motor behaviors might lie within a single, possibly nonlinear, universal neural manifold, with each task sampling a different region (Figure 1E). Because the tasks commonly studied are simple and stereotypical, the manifolds extracted from the recorded neural data may capture only a small region of the universal manifold, a region spanned by task-specific neural modes. This picture agrees with theoretical arguments that limiting the complexity of behavior also constrains the dimensionality of the neural manifolds (Gao and Ganguli, 2015). In this view, recording larger populations of neurons may not yield new insights; only by recording population activity during naturalistic, complex behaviors will we uncover the true underlying neural manifold.

We have argued that motor behaviors are generated by the time-varying activation of a small set of neural modes, population-wide activity patterns arising from network connectivity that define a low-dimensional manifold in neural space. We posit that the transition from a neuron-centric to a manifold-centric view of neural activity fosters an important advance in our understanding of brain function. Experiments involving longer and more complex motor tasks will require the development of increasingly sophisticated techniques for recording and data analysis; all these will be critical to advance our understanding of the relation between network connectivity, the resulting neural manifolds, and motor behaviors.

ACKNOWLEDGMENTS

This work was supported in part by Grant NS053603 from the National Institute of Neurological Disorder and Stroke (L.E.M.), by Grant NS094748 from the National Institute of Neurological Disorder and Stroke (S.A.S.), by Grant FP7-PEOPLE-2013-IOF-627384 from the Commission of the European Union (J.A.G.), and by Grant F31-NS092356 from the National Institute of Neurological Disorder and Stroke and Grant T32-HD07418 from the National Center for Medical Rehabilitation Research (M.G.P.). We thank Carolina Masumoto for the monkey illustration.

REFERENCES

- Ahrens, M.B., Li, J.M., Orger, M.B., Robson, D.N., Schier, A.F., Engert, F., and Portugues, R. (2012). Brain-wide neuronal dynamics during motor adaptation in zebrafish. *Nature* 485, 471–477.
- Bruno, A.M., Frost, W.N., and Humphries, M.D. (2015). Modular deconstruction reveals the dynamical and physical building blocks of a locomotion motor program. *Neuron* 86, 304–318.
- Carmena, J.M., Lebedev, M.A., Crist, R.E., O'Doherty, J.E., Santucci, D.M., Dimitrov, D.F., Patil, P.G., Henriquez, C.S., and Nicolelis, M.A.L. (2003). Learning to control a brain-machine interface for reaching and grasping by primates. *PLoS Biol.* 1, E42.
- Chapin, J.K., and Nicolelis, M.A.L. (1999). Principal component analysis of neuronal ensemble activity reveals multidimensional somatosensory representations. *J. Neurosci. Methods* 94, 121–140.
- Churchland, M.M., and Shenoy, K.V. (2007). Temporal complexity and heterogeneity of single-neuron activity in premotor and motor cortex. *J. Neurophysiol.* 97, 4235–4257.
- Churchland, M.M., Cunningham, J.P., Kaufman, M.T., Ryu, S.I., and Shenoy, K.V. (2010a). Cortical preparatory activity: representation of movement or first cog in a dynamical machine? *Neuron* 68, 387–400.
- Churchland, M.M., Yu, B.M., Cunningham, J.P., Sugrue, L.P., Cohen, M.R., Corrado, G.S., Newsome, W.T., Clark, A.M., Hosseini, P., Scott, B.B., et al. (2010b). Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nat. Neurosci.* 13, 369–378.
- Churchland, M.M., Cunningham, J.P., Kaufman, M.T., Foster, J.D., Nuyujukian, P., Ryu, S.I., and Shenoy, K.V. (2012). Neural population dynamics during reaching. *Nature* 487, 51–56.
- Cisek, P. (2006). Preparing for speed. Focus on “Preparatory activity in premotor and motor cortex reflects the speed of the upcoming reach”. *J. Neurophysiol.* 96, 2842–2843.
- Cisek, P., and Kalaska, J.F. (2005). Neural correlates of reaching decisions in dorsal premotor cortex: specification of multiple direction choices and final selection of action. *Neuron* 45, 801–814.
- Cowley, B.R., Smith, M.A., Kohn, A., and Yu, B.M. (2016). Stimulus-driven population activity patterns in macaque primary visual cortex. *PLoS Comput. Biol.* 12, e1005185.
- Cunningham, J.P., and Yu, B.M. (2014). Dimensionality reduction for large-scale neural recordings. *Nat. Neurosci.* 17, 1500–1509.
- d'Avella, A., Saltiel, P., and Bizzi, E. (2003). Combinations of muscle synergies in the construction of a natural motor behavior. *Nat. Neurosci.* 6, 300–308.
- Durstewitz, D., Vittoz, N.M., Floresco, S.B., and Seamans, J.K. (2010). Abrupt transitions between prefrontal neural ensemble states accompany behavioral transitions during rule learning. *Neuron* 66, 438–448.
- Elsayed, G.F., Lara, A.H., Kaufman, M.T., Churchland, M.M., and Cunningham, J.P. (2016). Reorganization between preparatory and movement population responses in motor cortex. *Nat. Commun.* 7, 13239.
- Ethier, C., Oby, E.R., Bauman, M.J., and Miller, L.E. (2012). Restoration of grasp following paralysis through brain-controlled stimulation of muscles. *Nature* 485, 368–371.
- Evarts, E.V. (1968). Relation of pyramidal tract activity to force exerted during voluntary movement. *J. Neurophysiol.* 31, 14–27.
- Fetz, E.E. (1992). Are movement parameters recognizably coded in the activity of single neurons? *Behav. Brain Sci.* 15, 679–690.
- Fetz, E.E., Cheney, P.D., Mewes, K., and Palmer, S. (1989). Control of forelimb muscle activity by populations of corticomotoneuronal and rubromotoneuronal cells. *Prog. Brain Res.* 80, 437–449, discussion 427–430.
- Flint, R.D., Scheid, M.R., Wright, Z.A., Solla, S.A., and Slutsky, M.W. (2016). Long-term stability of motor cortical activity: implications for brain machine interfaces and optimal feedback control. *J. Neurosci.* 36, 3623–3632.
- Forsberg, L.E., Bonde, L.H., Harvey, M.A., and Roland, P.E. (2016). The second spiking threshold: dynamics of laminar network spiking in the visual cortex. *Front. Syst. Neurosci.* 10, 65.
- Fu, M., Yu, X., Lu, J., and Zuo, Y. (2012). Repetitive motor learning induces coordinated formation of clustered dendritic spines in vivo. *Nature* 483, 92–95.
- Ganmor, E., Segev, R., and Schneidman, E. (2015). A thesaurus for a neural population code. *eLife* 4, 1–19.
- Gao, P., and Ganguli, S. (2015). On simplicity and complexity in the brave new world of large-scale neuroscience. *Curr. Opin. Neurobiol.* 32, 148–155.
- Georgopoulos, A.P., Kalaska, J.F., Caminiti, R., and Massey, J.T. (1982). On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *J. Neurosci.* 2, 1527–1537.
- Giszter, S.F. (2015). Motor primitives—new data and future questions. *Curr. Opin. Neurobiol.* 33, 156–165.
- Hatsopoulos, N.G., Ojakangas, C.L., Paninski, L., and Donoghue, J.P. (1998). Information about movement direction obtained from synchronous activity of motor cortical neurons. *Proc. Natl. Acad. Sci. USA* 95, 15706–15711.
- Hennequin, G., Vogels, T.P., and Gerstner, W. (2014). Optimal control of transient dynamics in balanced networks supports generation of complex movements. *Neuron* 82, 1394–1406.
- Hinton, G.E., and Salakhutdinov, R.R. (2006). Reducing the dimensionality of data with neural networks. *Science* 313, 504–507.
- Humphrey, D.R., Schmidt, E.M., and Thompson, W.D. (1970). Predicting measures of motor performance from multiple cortical spike trains. *Science* 170, 758–762.

- Jasper, H., Ricci, G.F., and Doane, B. (1958). Patterns of cortical neuronal discharge during conditioned responses in monkeys. In *Ciba Foundation Symposium - Neurological Basis of Behaviour*, G. Wolstenholme and C. O'Connor, eds. (John Wiley & Sons), pp. 277–294.
- Kaufman, M.T., Churchland, M.M., Ryu, S.I., and Shenoy, K.V. (2014). Cortical activity in the null space: permitting preparation without movement. *Nat. Neurosci.* 17, 440–448.
- Kaufman, M.T., Seely, J.S., Sussillo, D., Ryu, S.I., Shenoy, K.V., and Churchland, M.M. (2016). The largest response component in the motor cortex reflects movement timing but not movement type. *eNeuro* 3, ENEURO.0085-16.2016.
- Kobak, D., Brendel, W., Constantinidis, C., Feierstein, C.E., Kepecs, A., Mainen, Z.F., Qi, X.L., Romo, R., Uchida, N., and Machens, C.K. (2016). Demixed principal component analysis of neural population data. *eLife* 5, 1–36.
- Luczak, A., Barthó, P., and Harris, K.D. (2009). Spontaneous events outline the realm of possible sensory responses in neocortical populations. *Neuron* 62, 413–425.
- Luczak, A., McNaughton, B.L., and Harris, K.D. (2015). Packet-based communication in the cortex. *Nat. Rev. Neurosci.* 16, 745–755.
- Machens, C.K., Romo, R., and Brody, C.D. (2010). Functional, but not anatomical, separation of “what” and “when” in prefrontal cortex. *J. Neurosci.* 30, 350–360.
- Macke, J.H., Buesing, L., Cunningham, J.P., Yu, B.M., Shenoy, K.V., and Sahani, M. (2011). Empirical models of spiking in neuronal populations. *Adv. Neural Inf. Process. Syst.* 24, 1–9.
- Mante, V., Sussillo, D., Shenoy, K.V., and Newsome, W.T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* 503, 78–84.
- Markowitz, D.A., Curtis, C.E., and Pesaran, B. (2015). Multiple component networks support working memory in prefrontal cortex. *Proc. Natl. Acad. Sci. USA* 112, 11084–11089.
- Michaels, J.A., Dann, B., and Scherberger, H. (2016). Neural population dynamics during reaching are better explained by a dynamical system than representational tuning. *PLoS Comput. Biol.* 12, e1005175.
- Morrow, M.M., Jordan, L.R., and Miller, L.E. (2007). Direct comparison of the task-dependent discharge of M1 in hand space and muscle space. *J. Neurophysiol.* 97, 1786–1798.
- Mussa-Ivaldi, F.A., and Solla, S.A. (2004). Neural primitives for motion control. *IEEE J. Ocean. Eng.* 29, 640–650.
- Newsome, W.T., Britten, K.H., and Movshon, J.A. (1989). Neuronal correlates of a perceptual decision. *Nature* 341, 52–54.
- Okun, M., Steinmetz, N.A., Cossell, L., Iacuruso, M.F., Ko, H., Barthó, P., Moore, T., Hofer, S.B., Mscis-Flogel, T.D., Carandini, M., and Harris, K.D. (2015). Diverse coupling of neurons to populations in sensory cortex. *Nature* 521, 511–515.
- Overduin, S.A., d'Avella, A., Roh, J., Carmena, J.M., and Bizzi, E. (2015). Representation of muscle synergies in the primate brain. *J. Neurosci.* 35, 12615–12624.
- Raposo, D., Kaufman, M.T., and Churchland, A.K. (2014). A category-free neural population supports evolving demands during decision-making. *Nat. Neurosci.* 17, 1784–1792.
- Riehle, A., and Requin, J. (1989). Monkey primary motor and premotor cortex: single-cell activity related to prior information about direction and extent of an intended movement. *J. Neurophysiol.* 61, 534–549.
- Riout-Pedotti, M.S., Friedman, D., Hess, G., and Donoghue, J.P. (1998). Strengthening of horizontal cortical connections following skill learning. *Nat. Neurosci.* 1, 230–234.
- Roweis, S.T., and Saul, L.K. (2000). Nonlinear dimensionality reduction by locally linear embedding. *Science* 290, 2323–2326.
- Sadtler, P.T., Quick, K.M., Golub, M.D., Chase, S.M., Ryu, S.I., Tyler-Kabara, E.C., Yu, B.M., and Batista, A.P. (2014). Neural constraints on learning. *Nature* 512, 423–426.
- Santhanam, G., Yu, B.M., Gilja, V., Ryu, S.I., Afshar, A., Sahani, M., and Shenoy, K.V. (2009). Factor-analysis methods for higher-performance neural prostheses. *J. Neurophysiol.* 102, 1315–1330.
- Scott, S.H. (2004). Optimal feedback control and the neural basis of volitional motor control. *Nat. Rev. Neurosci.* 5, 532–546.
- Scott, S.H. (2008). Inconvenient truths about neural processing in primary motor cortex. *J. Physiol.* 586, 1217–1224.
- Serruya, M.D., Hatsopoulos, N.G., Paninski, L., Fellows, M.R., and Donoghue, J.P. (2002). Instant neural control of a movement signal. *Nature* 416, 141–142.
- Shen, L., and Alexander, G.E. (1997). Preferential representation of instructed target location versus limb trajectory in dorsal premotor area. *J. Neurophysiol.* 77, 1195–1212.
- Shenoy, K.V., Sahani, M., and Churchland, M.M. (2013). Cortical control of arm movements: a dynamical systems perspective. *Annu. Rev. Neurosci.* 36, 337–359.
- Stopfer, M., Jayaraman, V., and Laurent, G. (2003). Intensity versus identity coding in an olfactory system. *Neuron* 39, 991–1004.
- Sussillo, D., Churchland, M.M., Kaufman, M.T., and Shenoy, K.V. (2015). A neural network that finds a naturalistic solution for the production of muscle activity. *Nat. Neurosci.* 18, 1025–1033.
- Tenenbaum, J.B., de Silva, V., and Langford, J.C. (2000). A global geometric framework for nonlinear dimensionality reduction. *Science* 290, 2319–2323.
- Thach, W.T. (1978). Correlation of neural discharge with pattern and force of muscular activity, joint position, and direction of intended next movement in motor cortex and cerebellum. *J. Neurophysiol.* 41, 654–676.
- Thoroughman, K.A., and Shadmehr, R. (2000). Learning of action through adaptive combination of motor primitives. *Nature* 407, 742–747.
- Todorov, E. (2000). Direct cortical control of muscle activation in voluntary arm movements: a model. *Nat. Neurosci.* 3, 391–398.
- Todorov, E., and Jordan, M.I. (2002). Optimal feedback control as a theory of motor coordination. *Nat. Neurosci.* 5, 1226–1235.
- Tresch, M.C., and Jarc, A. (2009). The case for and against muscle synergies. *Curr. Opin. Neurobiol.* 19, 601–607.
- Tsodyks, M., Kenet, T., Grinvald, A., and Arieli, A. (1999). Linking spontaneous activity of single cortical neurons and the underlying functional architecture. *Science* 286, 1943–1946.
- Wurtz, R.H. (1969). Visual receptive fields of striate cortex neurons in awake monkeys. *J. Neurophysiol.* 32, 727–742.
- Yu, B.M., Cunningham, J.P., Santhanam, G., Ryu, S.I., Shenoy, K.V., and Sahani, M. (2009). Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *J. Neurophysiol.* 102, 614–635.