

Dynamical Feedback Control: Motor cortex as an optimal feedback controller based on neural dynamics

Christopher Versteeg^{1*}, Lee E. Miller^{1,2,3,4}

¹ Department of Biomedical Engineering, McCormick School of Engineering, Northwestern University, Evanston, IL, USA

² Department of Physiology, Feinberg School of Medicine, Northwestern University, Chicago, IL, USA

³ Department of Physical Medicine and Rehabilitation, Feinberg School of Medicine, Northwestern University, Chicago, IL, USA

⁴ Shirley Ryan AbilityLab, Chicago, IL, USA

* Corresponding Author Contact: chrissversteeg@gmail.com

Keywords: Motor Control, Dynamics, Optimal Feedback Control, Review

Abstract:

Primates have an unparalleled ability to produce a wide range of dexterous movements, each sensitive to context and robust to perturbation. Recent progress in understanding the neural basis of movement generation comes from two largely independent ideas, “Optimal Feedback Control” and “Neural Dynamical Systems”. The optimal control framework was largely inspired by research programs from the ‘80s showing that the brain doesn’t seem to plan a simple desired movement trajectory, but instead produces movements by transforming sensory information into motor output that satisfies an optimality criterion. The more recent idea that the motor cortex acts as a dynamical system came about only as it became possible to analyze large numbers of simultaneously recorded single neurons. These two framings of the motor system have been largely incommensurate, neither able to contribute much to the understanding of the other. In this review, we reconcile these two views into a single model we call “Dynamical Feedback Control”. We propose that the dynamics in the motor cortex emerge from a sensorimotor transformation that couples the motor cortex to sensory input from the periphery, and to contextual inputs from other cortical and subcortical areas. Dynamics in motor cortex can be thought to approximate gains of a feedback controller, and by moving the neural state to different regions of state space, the motor system can rapidly alternate between different controllers. The DFC framework presents a new lens to interpret neural dynamics, and to understand how ensembles of neurons generate flexible and responsive patterns of muscle activity.

Introduction:

Many of the questions that we ask about the brain are determined by the perspective from which we view it. Questions that are entirely sensible from one vantage point may be meaningless in another; for example, from the perspective of neuroscientists that study single neuron activity, the question “Are these neural trajectories tangled?” (Russo et al., 2018) is meaningless. What is

considered a reasonable question can change depending on your viewpoint (Kaiser, 2012).

Two major conceptual frameworks have guided recent study of the motor system: Optimal feedback control (OFC) and neural dynamical systems (NDS). The former posits that the brain generates movement through a control system built from a feedback controller, fundamentally dependent on input from a forward model and state estimator (Scott, 2004; Shadmehr & Krakauer, 2008; Todorov & Jordan, 2002). The brain can adjust feedback gains according to planning and context, thereby altering the nature of the transformation of state estimates into motor output. Most research in this field has been at the level of movement psychophysics (Mazzoni & Krakauer, 2006; Sha et al., 2006; Taylor et al., 2014), only infrequently examining the firing rates of single neurons (Cross et al., 2021; Kalidindi et al., 2021; Pruszynski, 2014; Pruszynski et al., 2011). OFC presents what might be considered an “algorithm-level” description of how the motor system controls movements but is largely agnostic to how this algorithm is implemented by neurons.

In contrast, the dynamical systems framework posits that motor cortical areas are pattern-generating circuits whose firing patterns emerge due to intrinsic dynamics of the network and its extrinsic inputs (Shenoy et al., 2013). Under the dynamical systems hypothesis, the neural state evolves according to the structure of a dynamical landscape, much like a ball rolling predictably along a curved, multi-dimensional track. Time-varying descending commands to muscles can be read out by corticospinal neurons from the movements of this neural state, a form of pattern generation. In NDS, planning and contextual inputs encode different movements simply by establishing different starting points; different tracks for the neural state space ball to roll down. In contrast with OFC, NDS has its roots in analysis of multi-electrode neural recordings, not psychophysical experiments or control theory. In summary, NDS incorporates both an algorithm-level (pattern-generation) and an implementation-level (dynamics of neural circuits) description of movement generation.

While these two frameworks both focus on the sensorimotor system, questions that OFC models pose, such as “what are the feedback gains that produce this movement?” seem uninterpretable under NDS. Conversely, common questions in NDS, such as “how do rotational dynamics contribute to the generation of motor output?” have rarely been posed in the framework of OFC (although see Kalidindi et al., 2021). To bridge this gap, we propose a hybrid model called Dynamical Feedback Control (DFC), that provides a unifying framework for algorithmic and implementation-level descriptions of the sensorimotor system; we propose that the algorithm of optimal feedback control is implemented by the dynamics of motor circuitry. In building this argument, we review OFC and NDS, with the goal of convincing the reader that 1) the brain embodies the major components of an optimal feedback control system and 2) that dynamical systems analysis provides a powerful set of tools to describe the computations performed by neural circuits. Ultimately, we propose an experiment to test the key predictions of DFC and provide a roadmap for future study to further unify our understanding of dynamics and feedback control.

OFC Model:

Minimum intervention principle:

There are many ways that a movement can be “optimal”, perhaps by achieving minimal distance travelled, minimal endpoint jerk, or minimal energy expended (Huang et al., 2012; Sha et al., 2006). A major feature of OFC is a rule known as the minimum intervention principle, which states that a good control system should correct only task-relevant parameters, while allowing irrelevant parameters to vary (Todorov & Jordan, 2002). This property emerges from two considerations: 1) correcting for irrelevant perturbations increases effort-dependent noise and 2) to do so increases total energetic expenditure, which runs counter to expectation and observation (Todorov & Jordan, 2003).

Tasks designed to exploit an uncontrolled manifold exemplify the minimum intervention principle. In one such task, a subject is instructed to exert force on two buttons using different fingers (F1 and F2). The sum of the forces may be controlled using any combination of F1 and F2, leading to “task-relevant” (F1+F2) and “task-null” (F1-F2) dimensions of control. The minimum intervention principle predicts that an optimal controller should allow variance in the task-null dimension in order to achieve tighter control in the task-relevant dimension.

Human behavior generally follows the minimum intervention principle. In very early studies of motor psychophysics, Bernstein (Bernstein, 1967) described remarkably precise hammer strikes arising despite highly variable trajectories, suggesting that the hammer’s trajectory was irrelevant as long as it produced an accurate strike (Biryukova & Sirotkina, 2020). Across a wide variety of tasks, this seems to be a principle of human movement: Task-null dimensions have large variance while task-relevant dimensions are well controlled. Any implementation of OFC needs to account for this feature.

Components of the OFC model:

Optimal feedback control theories have been around for decades (Todorov & Jordan, 2002), but in recent years a common picture has emerged. Figure 1A depicts one prevailing model in which the planning module sets feedback controller gains that are specific to the desired movement and context (Shadmehr & Krakauer, 2008). The feedback controller transmits a motor command to the muscles, a copy of which is sent to the forward model. The forward model transforms this “efference copy” signal into the predicted sensory consequences given the current state of the limb. The ensuing movement generates actual sensory signals (“reafference”) that travel from the periphery to the brain. The state estimator combines the predicted and reafferent signals, weighted by their relative confidences, to produce an estimate of the state of the limb. This state estimate passes through the feedback controller to generate a new motor output and the cycle repeats itself.

However, different movements require different feedback gains; even a seemingly trivial change such as reaching to different targets requires a different set of gains. Clearly, this does not limit our ability to make rapidly varying reaching movements. More generally, we are able to switch rapidly between previously learned behaviors, although learning a new behavior may require some time. Inputs from the planning module must be able to modify the controller quickly to allow the expansive and context-dependent repertoire of movements that the brain can produce.

Modern OFC proposes that different regions of the brain embody the four major components necessary to generate movement under this framework: a planning module, a feedback controller, a forward model, and a state estimator (Shadmehr & Krakauer, 2008). The behavior of the cerebellum has many of the hallmarks of a forward model, the motor cortex appears able to enact feedback control policy, perhaps with dual proprioceptive and visual controllers, and the state estimator may be constructed by circuits in anterior (area 3a) and posterior parietal (areas 5/7) cortex (Figure 1B). As the planning module must account for a wide variety of potential cost and reward signals, it is unlikely that it is localized to any single area, but instead may be distributed throughout the premotor and prefrontal cortices, as well as the basal ganglia.

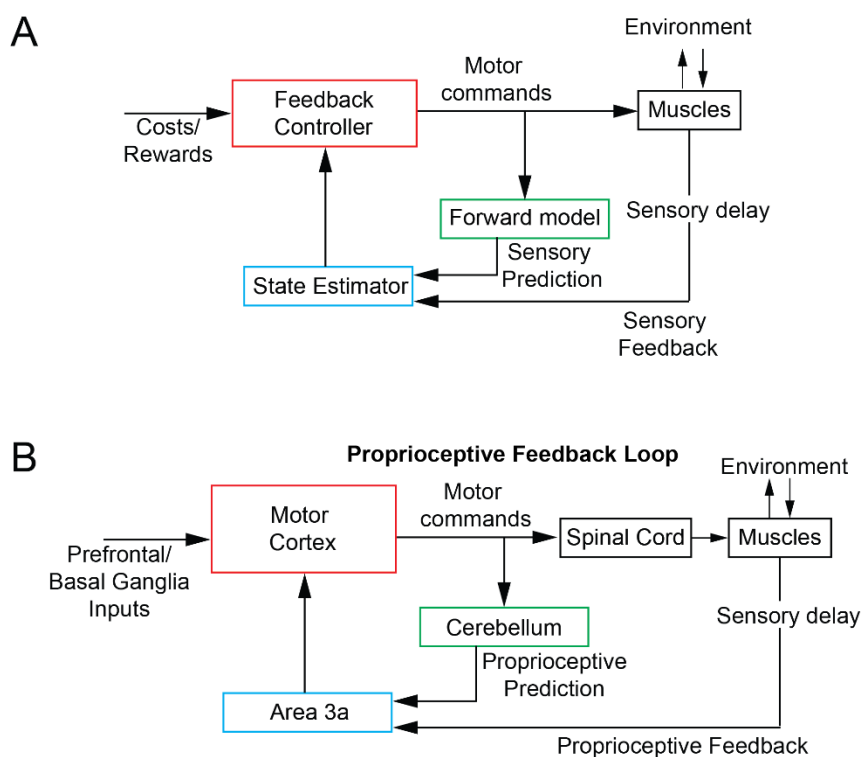


Figure 1: Features and neural correlates of Optimal Feedback Control for reaching. A: Model of optimal feedback control system, consisting of feedback controller (red), forward model (green), state estimator (blue), and musculoskeletal system. Colors denote the role of brain areas in the subsequent portions of the figure. B: Proposed proprioceptive feedback control loop. Motor cortex acts as a feedback controller, cerebellum a forward model, and area 3a a state estimator.

Cerebellum as a forward model:

Forward models, which predict sensory consequences of movement, are necessary for control systems that include substantial feedback delays. Running the feedback controller on delayed proprioceptive information (~40-50 ms for signals from the distal arm) can produce oscillations about the setpoint (Kawato, 1999; Wolpert et al., 1998). Relying on predicted feedback while

waiting for actual sensory feedback can mitigate the instability caused by sensory delays.

Among its other functions, many studies point to the cerebellum as the neural implementation of a forward model (Figure 2A) (Kawato, 1999). The intention tremor typical of patients with damage to the anterior lobes of the cerebellum is reminiscent of a feedback controller with long sensing delays. These patients also lack anticipatory increases in grip force during a predictable ball-drop task (Nowak et al., 2007; Serrien & Wiesendanger, 1999). Another clinical sign of cerebellar damage is difficulty in coordinating movements that involve multiple joints (Izawa et al. 2012; Nowak et al. 2007), which may reflect difficulties in predicting the interaction torques arising between mechanically-coupled limb segments (Bastian et al., 1996). In such a model, the sensory and efference copy inputs take the form of cerebellar mossy fibers, and the deep cerebellar nuclei provide the predicted sensory signals. Sensory prediction errors trigger updates to the forward model (through inferior olive activity and Purkinje cell complex spikes), which alter the synaptic weights of the parallel fibers onto Purkinje cells. For further review, see (Shadmehr, 2020).

Under OFC, disruptions of the forward model should eliminate sensory prediction and cause motor output to behave as though it has only lagged sensory information. Transcranial magnetic stimulation (TMS) can be used to disrupt processes in circumscribed regions of the brain and test how the resulting loss of function affects behavior. Human subjects have undergone cerebellar TMS while reaching (Miall et al., 2007). In these experiments, subjects initiated a slow reach to the right in response to an audio cue. At a random time into this reach, another audio cue signaled a reach to a distant visual target (Figure 2B). The timing of this cue caused the direction to the target relative to the hand to vary across trials. Although movements were made without vision of the arm, subjects could easily compensate for the varying position at the start of the second reach.

However, on TMS trials, the reach was typically not directed accurately at the target. Instead, its trajectory was appropriate to reach the target beginning from where the hand had been ~100 ms in the past (Figure 2B, red line). Disrupting the cerebellum did not eliminate the reach but caused it to proceed as though it had been planned with lagged somatosensory information. This experiment supports the existence of *both* a forward model in the cerebellum *and* a state estimator elsewhere in the brain that combines predicted and actual sensory information to compute a single state estimate for use by the controller.

Motor cortex as a feedback controller:

At a high level, a feedback controller maps states to actions that drive the system to a desired location in state space. A reach feedback controller might map a state estimate derived from vision, proprioception, and predicted state to the pattern of muscle activation that produces a desired movement.

To perform the feedback control presented in Figure 1B, the candidate feedback controller must 1) receive sensory inputs 2) project to the muscles and 3) send an efference copy signal to the forward model. Motor cortical neurons receive substantial somatosensory input (Cross et al., 2021; Pavlides et al., 1993), and old-world primates (and humans) have direct projections from

motor cortex to motor neurons (Lemon, 1997; Rathelot & Strick, 2009). The cortico-pontine mossy fibers provide the cerebellum with efference copy inputs (Ramnani, 2006). Thus, motor cortex presents a promising location for the feedback controller.

Long-latency reflexes may provide a useful model of OFC, as these automatic sensory-motor loops can mediate complex behavior yet have obvious analogies with classical feedback control models without the added complexity of voluntary influences. They are capable even of complex obstacle avoidance and rapid (<100ms) target switching in response to proprioceptive perturbations (Figure 2C, D (Nashed et al., 2014)). Anatomical, psychophysical, and electrophysiological evidence suggests that long-latency reflexes take a transcortical route through M1 (Asanuma, 1975; Zarzecki & Asanuma, 1979). When these reflexes are altered, such as by changing the interactions between joints, there are corresponding changes in firing rates of M1 neurons (Evarts & Tanji, 1974; Pruszynski et al., 2011). Furthermore, TMS over M1 of humans can potentiate these reflexes, suggesting a causal role for M1, not simply a correlation with reflexes generated elsewhere (Pruszynski et al., 2011).

One recent experiment examined long-latency reflexes evoked in humans during reaching (Figure 2C). Subjects were asked to reach to one of two targets while avoiding a visible obstacle. On some trials, the experimenters bumped the subject's hand to the left with a fixed magnitude perturbation, forcing the subject to correct their reach. On trials when the hand happened to be pushed farther to the left, subjects tended to correct to the left target, while on trials when the hand was pushed a smaller distance, the subjects moved to the right. The elbow extensor muscle activation differed between these two conditions in less than 100 ms after the bump (Figure 2D), a latency much lower than that of a voluntary choice (Scott, 2016). Perhaps the mapping from the arm's sensory state to the appropriate motor output has been pre-computed, to include responses to any "likely" perturbation; small bumps result in reflexive forces that guide the hand to the right target, while large bumps result in forces to the left (following the orange arrows, Figure 2C). Importantly, the orange arrows represent the approximate direction that combined muscle forces would accelerate the hand while correcting for a perturbation in these *particular* task conditions. Different targets or obstacles would require a different landscape. The motor system seems to be able to change the mapping from state-estimate to force output rapidly, based on the particulars of the reaching task.

The short latency of transcortical reflexes limits the potential influence of inputs from premotor areas, inputs that voluntary movements depend on. However, recent results suggest a tighter link between the feedback control model of long-latency reflexes and voluntary movements than had previously been appreciated (Maeda et al., 2018; Scott et al., 2015). This prospect seems contrary to our subjective sense of "agency", the feeling that we are the ones consciously controlling how our movements play out. In Box 1, we discuss this relationship between voluntary and reflexive control, suggesting that these two seemingly disparate processes may indeed be two sides of the same coin.

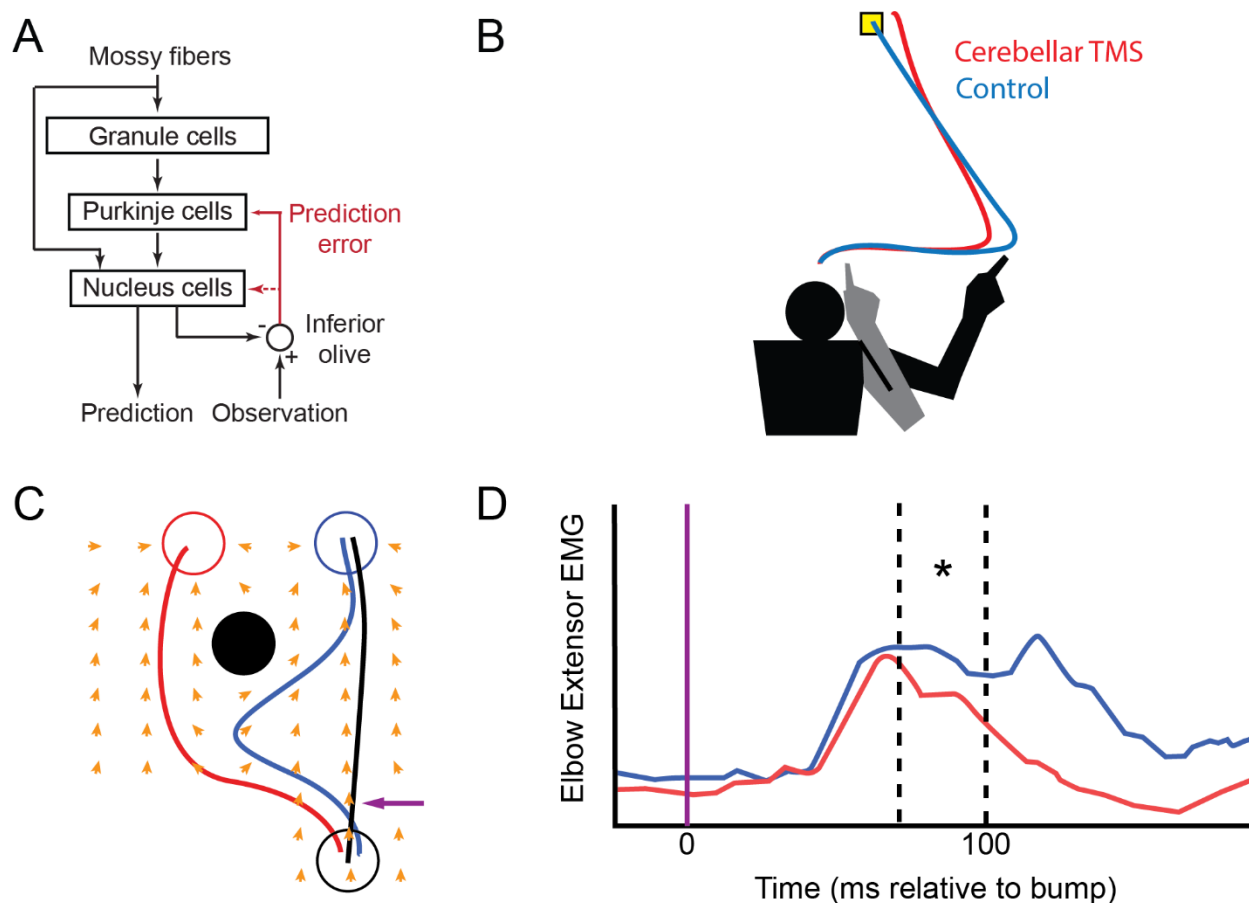


Figure 2: Evidence for cerebellum and motor cortex as linked forward model and feedback controller A: Model of cerebellum that generates a prediction of sensory input from efference copy (adapted from (Shadmehr, 2020), figure 1B). B: While redirected reaches made without vision (blue trace) normally acquired the target (yellow square) accurately, cerebellar transcranial magnetic stimulation caused trajectories (red trace) that appeared to be based on lagged somatosensory information (adapted from Miall et al. 2007, figure 1A, D). C: Adapted from Nashed et al., 2014, figures 2B and 6B. Human subjects were asked to reach from black circle to either of two outer targets (red and blue circles) while avoiding an obstacle (filled black circle). The unperturbed trajectory is shown in black. Leftward perturbation (purple arrow) displaced the hand, producing altered trajectories and subsequent responses (colored lines). The mapping from sensory to motor states is represented stylistically by orange arrows that approximate the motion of the hand as a function of its location and the task parameters. D: Mean elbow extensor muscle activity for trials that ended at the right (blue) and left targets (red). Statistically significant differences (region between dashed lines denoted by *) occurred less than 100 ms after bump was applied at time 0 (purple line).

Area 3a as a proprioceptive state estimator:

State estimators combine different sources of information based on the reliability of each input stream. For instance, a hunter shooting at a bird on a cloudy day must combine his noisy visual input with a prediction of where he expects the bird to be given its previous position and velocity. Similarly, a proprioceptive state estimator needs to combine afferent proprioceptive information with predictions based on efference copy and a forward model. This state estimator

should satisfy three criteria: 1) It should project to M1. 2) It should receive lagged proprioceptive signals from the periphery. 3) It should receive proprioceptive predictions from the cerebellum. Current evidence suggests that area 3a may act as a proprioceptive state estimator.

Area 3a projects strongly to M1 (Huffman & Krubitzer, 2001). Area 3a receives substantial proprioceptive input from the limbs (Jones & Porter, 1980; Phillips et al., 1971; Yumiya et al., 1974). While there is no direct evidence in primates that it also receives predicted proprioceptive signals, recordings in mice show that VL neurons encode sensory predictions as well as actual somatosensory signals (Dooley et al., 2021). Contrary to expectation, the primary thalamic inputs to area 3a are not from somatosensory thalamus (VPL and VPS), but instead from cerebellar thalamus (VL) (Padberg et al., 2009). More definitive evidence that it acts as a state estimator could be obtained by recording from area 3a while varying the relative reliability of predicted and actual proprioceptive information (e.g., by applying a predictable bump to the hand during a reach). We predict that as a monkey learns to expect this perturbation, neurons in 3a will begin to predict it, even on catch trials in which it does not occur. Importantly, this learned sensory prediction is distinct from alterations of the motor plan (i.e., a change in the feedback controller) that also occur during motor adaptation.

Summary of evidence for OFC in the brain

The large body of evidence reviewed above, including studies ranging from human psychophysics and TMS, to single neuron recording and modeling lend considerable support to the OFC model of movement control. A recent study attacked the question directly by using selective cooling to turn off particular regions of the brain thought to constitute OFC building blocks (Takei et al., 2021). Cooling area 5 (a putative visual state estimator) resembled disruption of a state estimator, while cooling motor cortex caused what looked like reduced gains in a feedback controller.

The motor system appears to act as a feedback system with at least one feedback controller in the motor cortices, with the cerebellum acting as a forward model, and a proprioceptive state estimator, probably in area 3a. However, knowing that motor cortex is a feedback controller does not tell us how that feedback controller is implemented by circuits of neurons. For insight into this question, we turn next to the more recent theory of neural dynamical systems.

Neural Dynamical Systems:

Discerning how ensembles of neurons perform computation is one of the critical challenges facing the fields of both neuroscience and artificial intelligence. Firing rates of individual neurons, in the brain and in artificial neural networks, have complex temporal patterns that are often difficult to interpret or to match to any observable task-related variable (Churchland et al., 2006; Fetz, 1992; Russo et al., 2018; Sussillo et al., 2015; Sussillo & Barak, 2013). This difficulty in interpreting single neuron activity has led to the adoption of population-level analyses that attempt to understand the network through discovery of lower-dimensional “latent spaces”. Because many neurons have highly correlated firing rates, a large percentage of the firing rate variance of a population of neurons can be explained by many, *many* fewer latent dimensions than the total number of neurons in the circuit (P. Gao & Ganguli, 2015). These patterns of correlations across neurons can be captured by “covariance matrices”, and there are a

number of related linear (principal component analysis and Gaussian-process factor analysis (Cunningham & Ghahramani, 2014; Cunningham & Yu, 2014; Yu et al., 2009)) and nonlinear (autoencoders, t-SNE (Kramer, 1991; Pandarinath et al., 2018; van der Maaten & Hinton, 2008)) methods used to compute and study these latent signals.

Preparatory activity sets the initial conditions of a neural state space dynamical system:

A prominent idea that emerged concurrently with this latent-space view is that cortical networks act as “dynamical systems” (Shenoy et al., 2013). The fundamental property of a dynamical system is that the state at some time in the future is determined by its current state, as well as any inputs to the system (Hirsch et al., 2013; Sussillo, 2014). In this sense, the current state of an ensemble of neurons recorded from the motor cortex appears to determine its future state with remarkable accuracy (Shenoy et al., 2013; Sussillo et al., 2016). The movement of the neural state within these latent spaces is often quite similar across trials (Kaufman et al., 2014; Pandarinath et al., 2018; Russo et al., 2018) as though it were playing out a fixed pattern, analogous to the dynamical trajectory of the three dimensional Lorenz attractor (Figure 3A). The particular path taken by the neural state seems to depend on the state of the system determined by planning-related activity prior to the go cue (Kaufman et al., 2014).

A recent experiment using optogenetic stimulation provided causal evidence of the critical role of preparatory activity in setting these initial conditions. Mice were trained to lick one of two water ports designated by the frequency of an auditory tone. Shortly after the tone there was a go cue, allowing the mouse to lick. Neural activity preceding each lick clustered in two discrete locations in motor cortical state space that were predictive of the eventual lick direction (Inagaki et al., 2019). Optogenetic stimulation of this region of the brain during the planning period perturbed the activity, which typically returned to its unperturbed state. However, occasionally, the preparatory activity “jumped” from a location encoding one choice to the opposite one. On these trials, the mouse played out the behavior that corresponded to the post-stimulation activity, rather than that of the original cued behavior. This experiment supports the theory that the dynamics of this region build “attractors” that are the starting locations for two different movements.

Further evidence of the coexistence of quite different dynamics within the cortical landscape comes from researchers who found a rapid switch between fundamentally different covariance patterns in mouse forelimb motor cortex in the transition between innate treadmill walking and a learned reaching task (Miri et al., 2017). These patterns, reflecting the functional interactions between cortical neurons, are thought to underlie differential control of muscles by distinct short-latency corticospinal pathways. Movement planning may amount to choosing an initial position within the landscape that produces specific trajectories of neural activity.

The evidence from these two experiments in mice suggests that activity in the motor cortex simply moves between existing regions of neural state space that are associated with the dynamics appropriate for different behaviors. The alternative, if one assumes that these dynamics are an emergent property of the synaptic weights in a network, is that massive numbers of synapses are altered each time a different movement is planned. Unlike these examples of rapid changes in neural covariance patterns, monkeys forced to adopt novel covariance patterns in a Brain-Computer Interface experiment required weeks of careful

coaching to learn the behavior (Sadler et al., 2014). We discuss the differences between these observations and their implications at greater length in Box 2.




Figure 3: Dynamical systems and motor potent dimensions help explain how the brain produces patterned muscle activity. A: Lorenz attractor, an autonomous dynamical system whose state develops as a function of a set of differential equations. B: One-dimensional projection from the Lorenz system, which generates a pattern over time. C: Low-dimensional trajectories of neural activity. Preparatory activity (blue) falls on a line determined by specifics of the motor plan. The trajectories of the neural state play out based on those initial conditions (green). Adapted from (Kaufman et al. 2014), Figure 3B.

Motor-potent spaces translate neural dynamics into muscle activity:

To produce a movement, dynamical trajectories of the neural state must be transformed into time-varying muscle commands. In NDS, a lower-dimensional control signal is “read out” of the cortical population activity by downstream neural systems, much like the time-varying projection of the Lorenz state onto the Y dimension (Figure 3B, green axis). By choosing the correct dynamics (the correct shape in Fig 3A) and the correct decoding axis (Fig 3B), arbitrary muscle patterns can be generated.

A remaining puzzle is why the substantial modulation of neural activity in both primary and premotor cortices after target presentation but prior to the go-cue does not cause undesired muscle activity during motor planning. Two major hypotheses have been proposed. The first is that there is a well-timed motor gating signal that prevents early, aberrant movements during motor planning (Benjamin et al., 2010; Duque et al., 2017; Duque & Ivry, 2009; Evarts & Tanji, 1974). Dimensionality analysis offers an alternate hypothesis. Because there are many more neurons than muscles, any linear readout between the two includes a “null space”, the directions along which the M1 neural state can move without changing muscle activity (Figure 3, red and blue axes). In contrast, the “output-potent space” is the set of directions that do change the muscle activity (Figure 3, green axis). Because there are only about 50 muscles in the arm and millions of M1 neurons, the vast majority of directions in the M1 neural state space are “output-

null”. Preparatory activity restricted to the null-space would cause no muscle activity.

By fitting linear models relating the M1 activity to EMG signals, Kaufman et al., demonstrated that most of the preparatory activity is indeed, output-null (Figure 3C) (Kaufman et al., 2014). This implies that the subspace of preparatory activity, and therefore the initial conditions for every dynamical trajectory that M1 is able to produce, is embedded within dimensions that do not project to muscles. Movement planning amounts to choosing an initial position within this landscape, resulting in dynamics that generate a pattern in the output-potent space producing appropriate muscle activity.

Neural dynamical systems and sensory feedback

Despite its appeal, the dynamical view of the motor system is difficult to reconcile with the substantial OFC literature. In particular, the lack of a prominent role for sensory inputs means that NDS can't address questions involving feedback control. More generally, the presence of sensory feedback complicates the interpretation of apparent “dynamics” in motor cortex. For example, a recent study demonstrated that neural population activity in the somatosensory cortex (Kalidindi et al., 2021) possesses many features attributed to dynamics in M1, in particular, rotations in neural state space (Russo et al., 2018). For these reasons, it is difficult to say what proportion of the dynamical behavior that we see in M1 is due to its own dynamics, versus dynamics “inherited” from other areas, including sensory input.

A recent study from the Hantman lab highlights the importance that inputs may play in the dynamics observed in M1, in particular, inputs from the motor thalamus (Sauerbrei et al., 2019). Mice were trained to retrieve a food pellet and bring it to their mouth, a behavior typically accompanied by consistent patterns of motor cortical activity across trials. Optogenetically inactivating the thalamic inputs to motor cortex eliminated the activity normally seen in motor cortex and caused the mice's reaching movements to fail. At least in the mouse, motor cortical dynamics are apparently contingent on thalamic inputs.

While these data are difficult to understand from the NDS perspective, through the lens of optimal feedback control they have a clear interpretation. Without a state estimate from VL, the sensory state will be outside the set of sensory inputs that the motor system has learned to handle. Not able to interpret the state of the limb, the motor system cannot produce normal motor output. However, while OFC highlights the importance of these sensory signals, its explanation is agnostic to how neurons might perform this computation. To describe how the algorithm of OFC might be implemented in the brain, we need the tools of NDS. This experiment underscores the need for a single model through which the combined body of literature of NDS and OFC can be understood.

Dynamical Feedback Control:

The differing priority of sensory input represents a disconnect between the OFC and NDS frameworks; OFC models see sensory feedback as fundamental while most models of neural dynamics are agnostic to it. To connect these two theories, we propose that signals from the state estimator project to a subspace of M1 neural activity we call the “sensory-potent space”. This space is complementary and orthogonal to a “non-sensory space”, the set of M1 dimensions that receive no sensory inputs. Together the sensory-potent and non-sensory spaces

comprise the complete M1 neural space. By moving signals from specific sensory-potent dimensions into output-potent dimensions, dynamics in M1 might approximate a specific set of feedback gains. In this scheme, the neural dynamics seen during behavior are “contingent” dynamics, contingent on the sensory inputs entering motor cortex and the contextual inputs that comprise the preparatory activity.

But how does motor cortex find dynamics that generate one particular movement from our behavioral repertoire? The possibility that the motor cortex might *change* the dynamics of a specific region of neural state space in the time it takes to prepare a movement seems unlikely, especially if those changes require alteration of many synaptic weights (see Box 2). Inspired by preparatory activity described by NDS, we propose instead that inputs encoding the associated costs, payoffs, and task requirements (from the planning module) *move* the neural state to regions where the dynamics already embody different mappings from limb state to motor output. In DFC, choosing the initial conditions of a dynamical system in NDS and choosing a feedback controller in OFC are in fact two different descriptions of the same process.

A simple dynamical feedback controller:

To demonstrate this idea, we will break down one of the simplest feedback loops in the body, the monosynaptic stretch reflex (Figure 4). In this example, the Ia afferent signals the lengthening velocity of the quadriceps muscle, and the α -motor neuron firing rate determines the activation of the quadriceps muscle. The synaptic weight between the Ia afferent and the α -motor neuron acts as a feedback gain.

This circuit can also be viewed as a simple dynamical system. At time t , the Ia afferent fires an action potential, followed by the α -motor neuron at time $t+1$. The neural state develops according to the dynamics dictated by the simple monosynaptic circuit (Figure 4B, black lines furthest into the page). Inputs to this dynamical system come from the environment via the muscle spindles, causing a predictable dynamical transformation that moves the 2D neural state from the sensory to the motor dimension.

In practice, the stretch reflex must be more complex than this; there must be some way of inactivating it, lest attempts to move voluntarily recruit the reflex and brake the movement. Evolution has devised a way to add context to this reflex, allowing it to treat self-generated movements differently from those imposed externally. To our simple 2D system we add a third dimension, an input causing presynaptic inhibition of the Ia terminal (Figure 4A, green (Meunier and Pierrot-Deseilligny 1989)). The circuit is still a feedback loop, now with gain that depends on where the neural state sits along the Inhibition dimension (Figure 4B). When Inhibition is zero, the stretch reflex occurs unimpeded (Figure 4B, black trace). When Inhibition is large, the stretch reflex does not occur, as its gain is zero (Figure 4B, light grey trace). At intermediate levels of inhibition, the reflex occurs with a gain that is functionally appropriate for the context.

The circuit is also now a slightly more complex dynamical system. At time t , the projection of the neural state along the Ia afferent dimension moves into the dimension of the α -motor neuron with dynamics determined by the projection of the neural state onto the Inhibition dimension (i.e., the context). Finding the dynamics within the Ia- α plane at different values of Inhibition is

equivalent to finding the feedback gain for that transformation.

This simple example illustrates the key components of the DFC model. Its instantiation in the brain includes higher-dimensional versions of the motor-potent, sensory-potent, and context dimensions (Figure 4C). There exists an output-potent subspace in M1 which transmits signals to the muscles analogous to the α dimension, and a somatosensory subspace into which the state estimator projects proprioceptive information (both actual and predicted), analogous to the 1a dimension. Inputs from (at least) the basal ganglia (via thalamus) and premotor areas provide inputs analogous to the Inhibition dimension. We might designate these inputs context and planning subspaces, depending on the information they encode. The extremely high number of non-sensory and output-null dimensions provide many “scratch” dimensions on which dynamics can be sculpted to produce appropriate sets of feedback gains. By driving the state to a location within the context/planning subspace, these inputs set the initial conditions of a dynamical system with specific transformations from sensory to motor dimensions (Figure 4C). Equivalently, this preparatory activity sets the feedback gains of a complex sensorimotor transformation, thereby implementing the feedback controller predicted by OFC.

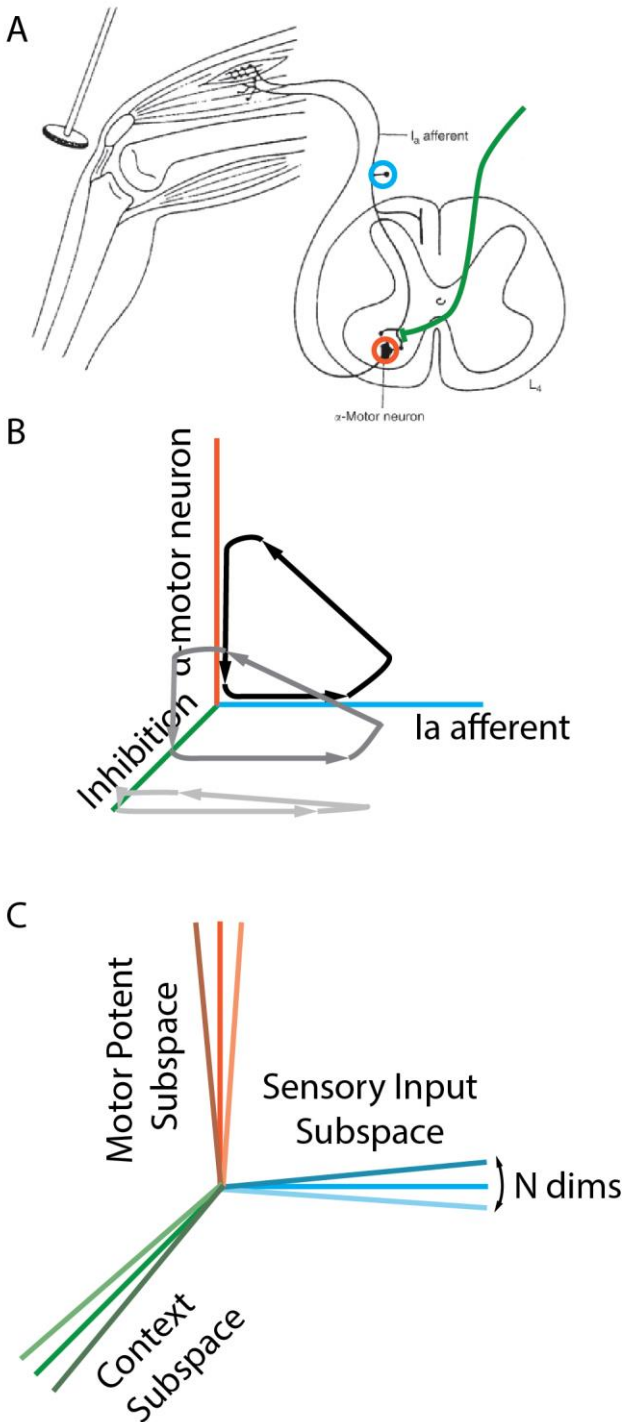


Figure 4: Stretch reflex as a simple dynamical feedback controller. A: Schematic diagram of the stretch reflex. Blue circle denotes the cell body of the Ia afferent muscle spindle. Red circle denotes the alpha-motor neuron projecting to the quadriceps muscle. Green axon represents presynaptic inhibitory axon. B: Dynamical landscape of this simple circuit. Red axis denotes the firing rate of the α -motor neuron, blue axis the Ia afferent firing rate, and green axis the firing rate of the presynaptic inhibitory neuron. Black and grey lines show the movement of neural activity in state space during a tendon tap at different levels of inhibition. C: Generalization of this model to motor cortex, with α -motor neuron firing rate replaced by an M-dimensional output-potent subspace and Ia afferent dimension replaced by an N-dimensional sensory-potent subspace. Inhibition dimension is replaced by a C-dimensional context and planning subspace.

A prototypical reach viewed through Dynamical Feedback Control

Of course, voluntary reaching is more complicated than even a complex reflex. How does DFC address the situation of a monkey presented with a target in one of two directions? In this case, the monkey may plan a movement but must delay its execution until receiving an auditory go cue. Before the go cue, while the monkey's hand is stationary and sensory inputs are unchanging, inputs from premotor areas and basal ganglia push motor cortex into a preparatory state that depends on the motor plan and costs/rewards, respectively (Z. Gao et al., 2018; Kaufman et al., 2014; Li et al., 2016). Movements along these preparatory dimensions place the neural state in a region whose local dynamics approximate a feedback controller that drives the limb to a desired location. The neural state remains in this preparatory location until the auditory cue nudges it, initiating its movement into the output-potent dimensions.

Following movement preparation, the neural state moves into output-potent dimensions, sending motor commands to the muscles and efference copy signals to the cerebellum. The cerebellum uses these signals to predict the sensory consequences of the motor commands. These predicted proprioceptive signals travel to area 3a; where they are combined with lagged signals to provide a combined state estimate, which projects into the sensory-potent dimensions of M1. Motor cortex transforms this combined state estimate into motor commands with gains determined by the dynamics of the M1 cortical circuit. These dynamics are likely not wholly intrinsic to M1, but also emergent from the recurrent connections with thalamus and other cortical areas. Importantly, by operating on state estimates derived from both prediction and lagged sensory signals, the same feedback controller can operate in predictive and reflexive modes (see Box 1). Thus, the movement unfurls through a recurrent loop connecting M1 to the periphery: motor outputs generate sensory inputs generate motor outputs. The dynamics that we observe in M1 are therefore the dynamics of the motor cortex coupled to the mechanics of the arm.

In this reflex-centric view of voluntary movement, reaches to the left and right would require different mappings from state to action (i.e., different feedback gains) in order for the sensory state at movement onset to produce the appropriate muscle activity. In DFC, the preparatory neural state moves along planning dimensions into different dynamical regions of the landscape for these two reach directions. The dynamics in these regions of state space build different sensory-motor mappings that generate different movements. By mapping task-relevant, but not task-null, sensory-potent dimensions onto corrective output-potent motor dimensions, the emergent dynamics of the motor system could correct only those errors that will hurt task performance, i.e., the minimum intervention principle hypothesized by OFC. How the brain modifies the dynamical landscape is still unclear, but this model of the motor cortex closely resembles a reinforcement learning policy; work connecting motor control and deep reinforcement learning may be essential to understanding the motor cortex.

Implications of Dynamical Feedback Control

What feedback transformation occurs during reaching?

This combined perspective allows researchers to ask Optimal Feedback Control questions in the language of Neural Dynamical Systems, but models are only useful if they can make novel, falsifiable predictions. NDS makes two key predictions. First, motor cortical dynamics should mirror corresponding sensory-motor transformations. In other words, movements of the M1

neural state from sensory-potent dimensions into output-potent dimensions should predict the actual motor activity. This finding would suggest that dynamics help to produce descending commands to muscles. Second, the transformation from sensory-potent to output-potent dimensions (i.e., the feedback gain) should change with the task requirements, not *in-place*, but through a translation of the neural state along context dimensions to a *different* dynamical landscape. Confirmation of these two predictions would demonstrate the utility of DFC as a multi-level description of the feedback control implemented by motor cortex.

We propose to use a 2D, visually guided reaching task while recording from M1. The monkey begins a trial by holding a robotic manipulandum in a target near to the body aligned left/right with the center of the screen (Figure 5A). We show the monkey a distant, midline target that is either narrow or wide (spanning the entire upper screen), chosen randomly across trials. After a random delay, we provide an auditory go cue. After acquiring the target, the monkey receives a liquid reward. On some trials, we apply a left or right perturbation to the monkey's hand during the reach and record the perturbation evoked neural activity and reflexively generated corrective force. Bumps during a reach to a narrow target are task-relevant, while bumps during a reach to a wide target are task-null. Therefore, we expect that the corrective motor response will differ between target widths. Our question is: can the dynamics of the sensory-motor transformation in M1 predict the target-dependent corrective response?

To determine the sensory-motor mapping in M1, we need to know which of its dimensions are sensory-potent and which are output-potent. By relating increases in sensory-potent dimensions to increases in output-potent dimensions, we can empirically estimate a feedback gain, the strength with which a given sensory input is transformed into motor output. We can map the dimensions of the sensory-potent subspace by recording M1 activity during perturbations of the monkey's hand at rest (Figure 5B).

Determining the output-potent space is a bit more complicated. If we were to simply measure M1 activity during voluntary movement and ignore the highly correlated reafferent inputs to M1 (as has been done previously) our estimates would be inaccurate. Instead, we exclude the previously identified sensory subspace activity (Figure 5B) from the neural space, then fit a model that relates the remaining M1 activity to handle forces. This will give us a motor subspace that maps neural activity to right and left force generation and is orthogonal to the somatosensory subspace (Figure 5C).

We can use these low-D sensory and motor subspaces to examine dynamics during the task. Specifically, we want to project the neural activity for a given target width onto the plane defined by single sensory and motor dimensions. The pair of dimensions should be related; for example, we expect that an error in a task-relevant sensory dimension (bump left) should be transformed into a projection of the neural state onto the output-potent dimension that corrects that error (move right; Figure 5D). The dynamics projected onto this plane will show how the sensory dimension moves into the motor dimension, or seen through the OFC lens, the feedback gain coupling the sensory input to the motor output. We would measure gain for a given movement by computing the ratio between the (motor) projection of the neural state onto the force dimension and its (sensory) projection onto the sensory-potent dimension. For a given sensory projection, trials with a small motor projection have a small feedback gain, while those

with a large motor projection will have a large feedback gain.

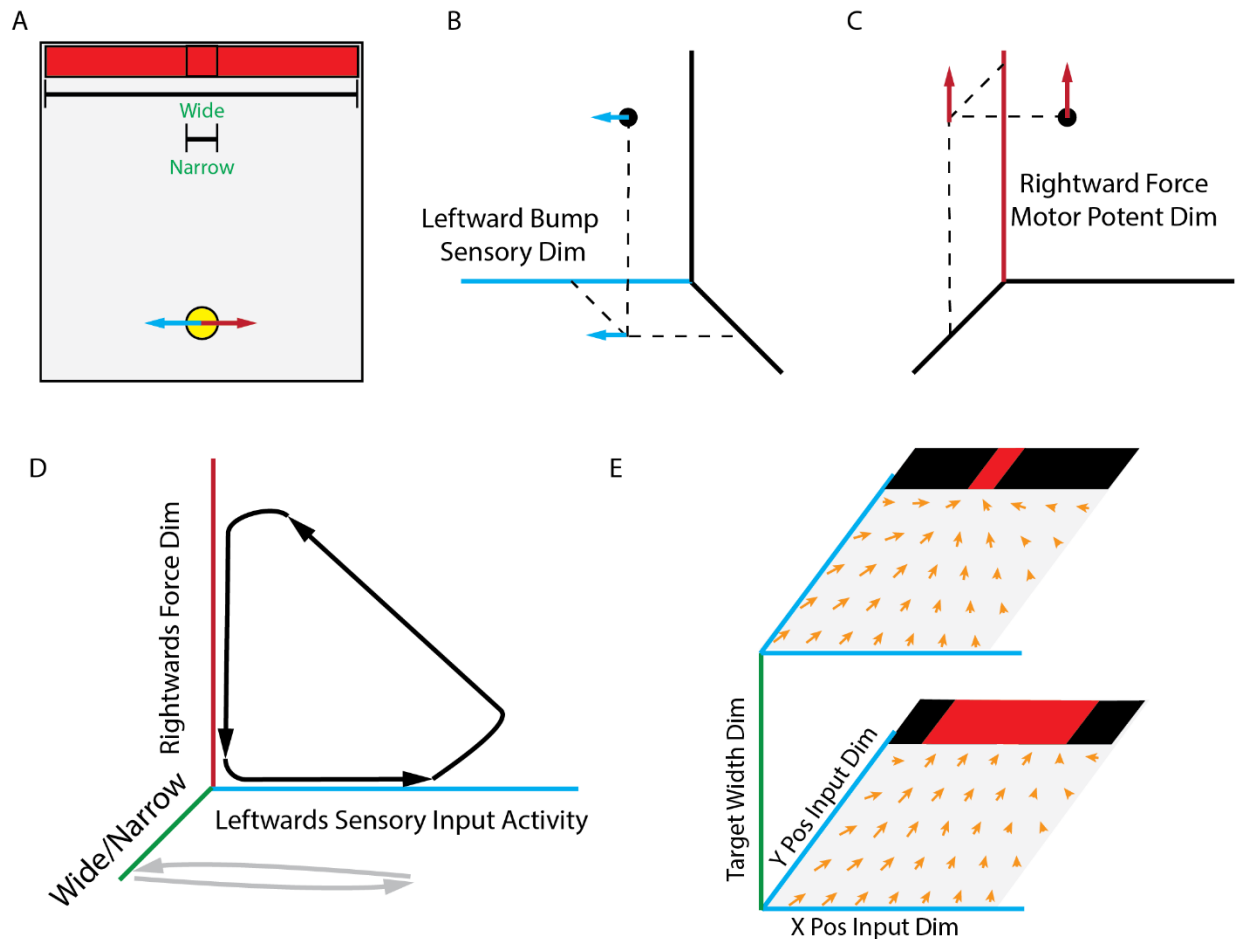


Figure 5: Example experiment under Dynamical Feedback Control. A: Task diagram. Narrow and wide targets represented by red square and rectangle, respectively. Yellow circle represents the cursor, controlled by hand position. Blue arrow represents a force perturbation that moves the monkey's hand to the left. Red arrow is a force generated by the monkey to the right. B: Schematic of responses in M1 to leftward bump perturbations. Black axes represent the high-dimensional neural space in M1. Blue axis represents the leftward bump sensory-potent dimension in M1. Black circle represents the neural state prior to the bump, which moves along the blue axis in response to the bump (blue arrows here and in panel A). C: Black axes represent the non-sensory subspace of M1 activity. Red axis represents the rightward force output-potent dimension in this reduced M1 space. Red arrows indicate movement of the neural state (black circle) corresponding to rightward force generation. D: Diagram of expected results (analogous to Figure 4B). Blue axis represents the leftward bump sensory-potent dimension in M1 (of Figure 5B). Red axis represents the rightward force output-potent dimension (Figure 5C). Green axis represents the context dimension, along which the target width is encoded during the preparatory period. Black line represents the expected dynamical behavior of M1 as a response to a leftward bump during a reach to a narrow target. Grey line corresponds to behavior for a wide target. E) Diagram of hand motion that arises from projections onto M1 motor-potent

dimensions that are a function of two sensory-potent dimensions in M1 that encode hand position (blue axes) and a context dimension that encodes target width (green axis). Orange arrows indicate the mapping from sensory to motor states as in Fig 2A.

OFC theory assumes a perpendicular bump will cause a task-relevant error primarily during reaches to the narrow target. Therefore, we predict that a given neural projection onto the bump-related sensory-potent dimension will generate a larger projection onto the corrective motor-potent dimension for narrow-target trials than for the wide-target trials (Figure 5D). This would indicate that the dynamics of the circuit (equivalently, the feedback transformations) are tuned to correct only for task-relevant perturbations.

Changes in the location of the neural state along the green axis during the preparatory period should encode the target type (Figure 5D, green axis); these movements will be accompanied by changes in the feedback gain of the sensory-motor transformation. Furthermore, we predict that the magnitude of the projection of the preparatory activity onto the context dimension (green) will correlate with the gain of the sensory-motor transformation for single trials. A trial having preparatory activity with a shorter projection along the green context axis will have a larger feedback gain than one with a longer projection, even for trials with the same target widths.

This experiment would allow us to test 1) whether dynamics transform M1 sensory-potent dimensions into output-potent dimensions in a way that predicts the corrective forces generated by the monkey, and 2) whether different locations along the preparatory dimensions of M1 constitute different regions of the dynamical landscape, each tuned to produce appropriate sensory-motor transformations to generate the movement and correct for task-relevant errors while ignoring task-null errors.

Interestingly, there is nothing constraining this model to correct only mechanical perturbations; it is equally well equipped to deal with more abstract perturbations such as a target that changes location mid-reach. If the landscape itself cannot be altered quickly enough, an altered target location could be dealt with by moving the neural state to a region of the context subspace that encodes the new target (Figure 5E). In this view, DFC presents motor cortex (and the other brain areas to which it is coupled) as a dynamical system embodying plans for all likely movements, each existing in a range of possible contexts. We describe this “multiverse” concept in more detail in Box 3.

Summary and Extensions of DFC:

Under DFC, inputs from planning modules designate the location within the dynamical landscape used to generate a movement. To understand how these locations are chosen, we need to understand how basal ganglia and premotor inputs affect the M1 neural state, i.e., the BG-M1 and Premotor-M1 input dimensions. These dimensions are analogous to the sensory-input dimensions to M1 described above, except they represent the dimensions in M1 that receive inputs from other brain areas. The analytical tools that have been used to find communication dimensions between other brain areas (Perich et al., 2018; Semedo et al., 2019) can be used to compute these M1 input dimensions from basal ganglia and premotor areas as well. Given what we know about the roles of BG and PFC in the motor system, we predict that the variables

encoded in the BG-input dimensions to M1 should relate to costs of movement, while the variables encoded along PFC-input dimensions should relate to the motor plan itself. Using DFC as a guide, we can map the functional consequences of inputs from other brain areas on M1 feedback control.

Importantly, we do not mean to suggest that the sensory-motor dynamics that we expect to find in M1 are “intrinsic” to M1 itself. Rather, these dynamics probably arise from recurrent coupling of M1 to a variety of other brain regions. Our model highlights one specific example of inter-region coupling that is experimentally tractable, namely, the coupling of M1 and the somatosensory periphery. Future work examining the connections of M1 to other areas (e.g., thalamocortical coupling) will be critical to untangling intrinsic M1 dynamics from those that arise from recurrent connections with other brain areas.

The DFC model makes apparent that tools to estimate motor potent spaces and neural covariance matrices can be biased by the effect of somatosensory inputs, as the structure of the neural population activity arises from both the dynamics of the feedback controller *and* the sensory information that flows through it. Including these inputs into the DFC model allows more accurate estimates of these characteristics of motor cortex.

There are some important limitations to this theory. Many components presented here, though based on existing evidence, remain speculative. Further work to characterize area 3a is needed to confirm that it receives both predicted and actual somatosensory signals, and that it combines these signals as a state estimator. In addition, any discussion of motor-potent projections makes the implicit assumption that these properties change only slowly, perhaps on the order of the time course of motor learning. There is evidence, however, that spinal interneurons process descending corticospinal signals even during motor planning (Prut & Fetz, 1999) and these signals are further modulated by the nonlinear properties of motoneurons (Dum & Strick, 1996; Heckman et al., 2005; Naufel et al., 2019; Shalit et al., 2012). Until we have a better understanding of how processing in the spinal cord affects descending cortical signals, we must interpret motor cortical signals with caution.

Groups that study the neural control of movement from the perspectives of OFC and NDS are often not in close communication with one another. Dynamical Feedback Control may help to bridge this gap between the high-level motor control theory presented by OFC and the empirically derived dynamical landscape of NDS. We propose it as a unifying theory that can explain our current understanding of the motor system at multiple conceptual levels and guide future inquiry.

Acknowledgements

We thank S. Bensmaia, R. Chowdhury, A. Miri, E. Oby, H. Riecke, S. Solla, M. Tresch and Y. Yang for their feedback and discussions on earlier drafts of this manuscript. This work was supported by grants from NINDS to Miller (R01 NS095162, R01 NS095251) and Versteeg (F31 NS092356).

Box 1: Voluntary movement as a complex transcortical reflex

A recent hypothesis suggests that voluntary and reflexive control of the arm share a sensory-motor mapping (Scott et al., 2015). Supporting this hypothesis is the observation that adaptation of long-latency reflexes is accompanied by a corresponding adaptation of voluntary reaches (Maeda et al., 2018; Pruszynski, 2014; Scott et al., 2015). Conversely, learning to reach in an altered force environment causes adaptation of reflexes (Maeda et al., 2018). This bidirectional transfer of motor learning suggests that updates are not simply applied concurrently to two separate models for voluntary and reflex control, but instead to a single, shared neural circuit.

Extending this logic, shared circuitry for reflexive and voluntary control of movement might allow a precomputed sensory-motor reflex mapping (as in Figure 2C) to produce controlled voluntary movements on trials without perturbations. The hand would move to the target by simply following the contour of the dynamical sensory-motor mapping, the direction and magnitude of endpoint motion determined in part by predicted sensory inputs based on efference copy, as well as direct efferent control.

However, when we make voluntary movements, we feel a sense of agency that is not present during reflexive movements (Haggard, 2017). This feeling is difficult to reconcile with a motor system acting as a simple feedback controller. Interestingly, a major component of this sense of agency depends on how well reafferent sensory information matches prediction (Blakemore et al., 1998, 2000); this corresponds to the function of the state estimator in the OFC model. When predicted proprioceptive information is accurate, as in a well-learned voluntary reach with no perturbations, the state estimator allows the feedback controller to circumvent sensory conduction delays by providing predicted proprioceptive signals (black trajectory in Figure 2C), and perhaps a sense of agency. In contrast, when predicted somatosensory information is inaccurate, such as when a subject begins to learn a predictable force field, the discrepancy between predicted and actual sensory information may produce a feeling that “something” else moved your limb. We hypothesize that this feeling of agency should return as learning progresses and prediction accuracy improves; to our knowledge this has never been tested.

Box 2: Learning: Fast and Slow

The location of preparatory activity within the neural state space seems to set the appropriate dynamics to allow movements in particular directions (see Fig 3C and (Churchland et al., 2006; Kaufman et al., 2014). Furthermore, there is evidence that this neural state space consists of multiple, fixed attractors, not a continuously remodeled dynamical landscape (Inagaki et al. 2019). However, when we interact with the world, we sometimes find that our actions don't produce the intended outcomes. Often this mismatch has to do with changes in the environment; in many cases we can compensate in a few minutes, but occasionally we are presented with something that takes much longer to learn. What differentiates changes we can learn quickly from those requiring days or weeks?

There are many mechanisms that can change the strength of existing synapses or form new ones, with time courses ranging from tens of milliseconds to hours or days. Activity dependent long-term potentiation and depression are thought to underlie the persistent changes associated with

memory, including motor learning. Learning to make movements against a novel but predictable force field is a paradigm used since the early '90's to study these processes (Scheidt et al., 2000; Shadmehr & Mussa-Ivaldi, 1994). Adaptation to new movement dynamics can happen in tens of minutes, but the learning remains fragile until the protein synthesis required for consolidation and stabilization of the underlying synaptic weight occurs (Citri & Malenka, 2008; Ranganathan et al., 2014; Shadmehr et al., 1995). A recent experiment examined the changes in the neural state space when monkeys learned to reach in an altered force environment. With roughly the time course of the behavioral adaptation, preparatory activity in motor cortex began to move in dimensions orthogonal to those of the control condition, pushing the neural state into a different region of the dynamical landscape (Sun et al., 2020). With roughly the time course of the behavioral adaptation, preparatory activity in motor cortex began to move in dimensions orthogonal to those of the control condition, pushing the neural state into a different region of the dynamical landscape (Sun et al., 2020). As in the mice experiments reviewed in the main text, we propose that in these experiments, dynamics appropriate for the curl field already existed, learned during the extensive behavioral training prior to the neural recordings.

Another set of monkey experiments studied the neural state changes during two types of motor adaptation that had very different time courses. In those experiments, monkeys were initially trained to control a cursor using a BCI decoder with an “intuitive mapping”, one based on the natural covariance structure of the neural activity during hand movements. Subsequently, this decoder was altered in one of two ways. Successful adaptation either did or did not require the neural state to move “outside the manifold” and thereby break the natural covariance patterns among neurons (Sadler et al., 2014). Monkeys could learn a “within-manifold” perturbation within a single experimental session. However, outside-manifold perturbations required weeks of coaching through progressive steps (Oby et al., 2019). Even for the within-manifold perturbations, monkeys “reassociated” their fixed repertoire of neural activity patterns to the new movements, a suboptimal strategy that preserves the covariance of the population activity on the manifold at the expense of task performance (Golub et al., 2018). Thus, for both within and outside manifold perturbations, the monkey finds it difficult to break the structure of activity across neurons, perhaps because doing so requires significant, coordinated synaptic weight changes. Remarkably, in one experiment after having learned the out of manifold perturbation, the monkey could then switch readily between it and the intuitive decoder (Oby, personal communication).

We suggest that some movement planning and some motor learning can happen relatively quickly by finding a region of the existing dynamical landscape with appropriate dynamics (some more easily found than others). When the appropriate dynamics don't exist, the landscape itself needs to be altered. In this case, the brain may adopt a process analogous to training a policy in a deep reinforcement learning (RL) network. Deep RL models can learn expressive control policies that function across highly variable task conditions, such as altered force environments (OpenAI et al., 2019; Tobin et al., 2017). These policies are analogous to optimal feedback controllers; both transform sensory inputs into control signals, subject to loss and optimality criteria. In the brain, learning at the edge of the known landscape likely requires synaptic changes, changes that, as in RL, are reinforced by trial success and failure.

Box 3: The dynamical systems multiverse:

We have attempted to paint a picture of a sensorimotor brain containing a multi-dimensional landscape in which planning a reaching movement involves selecting a track with contours that move the neural activity (and hand) to the intended target. The walls of this track defend against small perturbations; larger perturbations may bump the brain state from the “left” track onto the neighboring “right” one (Figure 2A). While perturbations within a given track are handled by a single feedback controller, large perturbations (or target shifts or even abstract rule changes) might push the neural state to a track that embodies a different set of sensorimotor transformations, the “track” for an entirely different feedback controller (Figure 5E and the mouse optogenetic experiment (Inagaki et al., 2019)). This view of the motor cortex implies the simultaneous presence of tracks that embody a variety of control strategies for many possible movements and contexts, accessible by a simple translation in neural state space.

Of course, the world is considerably more complex than left and right movements. In this view, the neural landscape must somehow include all likely movements, perturbations, and contexts: many narrow, parallel tracks separated by substantial walls and perhaps several shallow tracks all within one broad track. Perturbations push the neural state from one track to another more easily for those tracks with lower walls (Finkelstein et al., 2021). We propose that this is possible because of the ultra high-dimensional nature of motor cortex, with theoretically as many dimensions as there are neurons. Its virtually limitless state space allows for the storage, retrieval, and execution of a vast number of previously learned motor behaviors.

This is not to say that rapid changes to the dynamical landscape cannot occur. To execute a particular movement with greater urgency or care, one might desire a narrow track with high walls as opposed to a wide sloping plane. Perhaps there is an alternate means to modulate the general form of the tracks within the landscape beyond the adjustment of many specific synaptic weights. Diffuse neuromodulation may play this role.

A large body of literature has demonstrated the effects of neuromodulators, which can up or down regulate neuronal excitability and alter synaptic efficacy (Jankowska et al., 2000; Vitrac & Benoit-Marand, 2017). Since dynamics emerge from large populations of interconnected neurons, a critical undertaking will be to understand how neuromodulatory effects, presently understood best at the level of single synapses, influence the circuit dynamics themselves. Research using artificial neural networks has begun to model the consequences at the network level of these cellular and synaptic level changes (Shine et al., 2021; Tsuda et al., 2021). An appealing function of these neuromodulators might be to apply “filters” that sharpen or smooth the geometry of the landscape. For instance, increases in vigor (related to the levels of striatal dopamine (da Silva et al., 2018; Panigrahi et al., 2015) might modify the dynamical landscape to produce larger projections into the muscle-potent space, thereby increasing movement speeds at the cost of accuracy. In addition to the motor system’s ability to rapidly explore the known landscape and to extend its edges through prolonged training and practice, neuromodulation might allow the brain to tune the dynamical landscape globally, tailoring it to current environmental or psychological conditions.

- Asanuma, H. (1975). Recent developments in the study of the columnar arrangement of neurons within the motor cortex. *Https://Doi.Org/10.1152/Physrev.1975.55.2.143*, 55(2), 143–156. <https://doi.org/10.1152/PHYSREV.1975.55.2.143>
- Bastian, A. J., Martin, T. A., Keating, J. G., & Thach, W. T. (1996). Cerebellar Ataxia: Abnormal Control of Interaction Torques Across Multiple Joints. *JOURNAL OF NEUROPHYSIOLOGY*, 76(1).
- Benjamin, P. R., Staras, K., & Kemenes, G. (2010). What Roles Do Tonic Inhibition and Disinhibition Play in the Control of Motor Programs? *Frontiers in Behavioral Neuroscience*, 4(JUN). <https://doi.org/10.3389/FNBEH.2010.00030>
- Bernstein, N. (1967). *The co-ordination and regulation of movements*. Pergamon Press.
- Biryukova, E., & Sirotkina, I. (2020). Forward to Bernstein: Movement Complexity as a New Frontier. *Frontiers in Neuroscience*, 0, 553. <https://doi.org/10.3389/FNINS.2020.00553>
- Blakemore, S.-J., Wolpert, D., & Frith, C. (2000). Why can't you tickle yourself? *NeuroReport*, 11(11), R11.
- Blakemore, S.-J., Wolpert, D. M., & Frith, C. D. (1998). Central cancellation of self-produced tickle sensation. *Nature Neuroscience*, 1(7), 635–640. <https://doi.org/10.1038/2870>
- Churchland, M. M., Yu, B. M., Ryu, S. I., Santhanam, G., & Shenoy, K. V. (2006). Neural Variability in Premotor Cortex Provides a Signature of Motor Preparation. *Journal of Neuroscience*, 26(14), 3697–3712. <https://doi.org/10.1523/JNEUROSCI.3762-05.2006>
- Citri, A., & Malenka, R. C. (2008). Synaptic Plasticity: Multiple Forms, Functions, and Mechanisms. *Neuropsychopharmacology*, 33(1), 18–41. <https://doi.org/10.1038/sj.npp.1301559>

- Cross, K. P., Cook, D. J., & Scott, S. H. (2021). Convergence of proprioceptive and visual feedback on neurons in primary motor cortex. *BioRxiv*, 2021.05.01.442274.
<https://doi.org/10.1101/2021.05.01.442274>
- Cunningham, J. P., & Ghahramani, Z. (2014). *Linear Dimensionality Reduction: Survey, Insights, and Generalizations*. <https://doi.org/10.1002/cssc.201100629>
- Cunningham, J. P., & Yu, B. M. (2014). *Dimensionality reduction for large-scale neural recordings*. *17*(11). <https://doi.org/10.1038/nn.3776>
- da Silva, J. A., Tecuapetla, F., Paixão, V., & Costa, R. M. (2018). Dopamine neuron activity before action initiation gates and invigorates future movements. *Nature*, *554*(7691), 244–248. <https://doi.org/10.1038/nature25457>
- Dooley, J. C., Sokoloff, G., Blumberg, M. S., & Dooley, J. (2021). *Developmental onset of a cerebellar-dependent forward model of movement in motor thalamus 2 3*.
<https://doi.org/10.1101/2021.06.25.449956>
- Dum, R. P., & Strick, P. L. (1996). Spinal cord terminations of the medial wall motor areas in macaque monkeys. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *16*(20), 6513–6525. <https://doi.org/10.1523/jneurosci.0768-06.2006>
- Duque, J., Greenhouse, I., Labruna, L., & Ivry, R. B. (2017). Physiological markers of motor inhibition during human behavior. *Trends in Neurosciences*, *40*(4), 219.
<https://doi.org/10.1016/J.TINS.2017.02.006>
- Duque, J., & Ivry, R. B. (2009). Role of corticospinal suppression during motor preparation. *Cerebral Cortex*, *19*(9), 2013–2024. <https://doi.org/10.1093/cercor/bhn230>
- Evarts, E. V., & Tanji, J. (1974). Gating of motor cortex reflexes by prior instruction. *Brain Research*, *71*(2–3), 479–494. [https://doi.org/10.1016/0006-8993\(74\)90992-5](https://doi.org/10.1016/0006-8993(74)90992-5)

- Fetz, E. E. (1992). Are movement parameters recognizably coded in the activity of single neurons? *Behavioral and Brain Sciences*, *15*(4), 679–690.
<https://doi.org/10.1017/S0140525X00072599>
- Finkelstein, A., Fontolan, L., Economo, M. N., Li, N., Romani, S., & Svoboda, K. (2021). Attractor dynamics gate cortical information flow during decision-making. *Nature Neuroscience*, *24*(6), 843–850. <https://doi.org/10.1038/s41593-021-00840-6>
- Gao, P., & Ganguli, S. (2015). On simplicity and complexity in the brave new world of large-scale neuroscience. *Current Opinion in Neurobiology*, *32*, 148–155.
<https://doi.org/10.1016/j.conb.2015.04.003>
- Gao, Z., Davis, C., Thomas, A. M., Economo, M. N., Abrego, A. M., Svoboda, K., De Zeeuw, C. I., & Li, N. (2018). A cortico-cerebellar loop for motor planning. *Nature*, *563*(7729), 113–116. <https://doi.org/10.1038/s41586-018-0633-x>
- Golub, M. D., Sadtler, P. T., Oby, E. R., Quick, K. M., Ryu, S. I., Tyler-Kabara, E. C., Batista, A. P., Chase, S. M., & Yu, B. M. (2018). Learning by neural reassociation. *Nature Neuroscience*, *21*(4), 607–616. <https://doi.org/10.1038/s41593-018-0095-3>
- Haggard, P. (2017). Sense of agency in the human brain. *Nature Reviews Neuroscience*, *18*(4), 196–207. <https://doi.org/10.1038/nrn.2017.14>
- Heckman, C. J., Gorassini, M. A., & Bennett, D. J. (2005). Persistent inward currents in motoneuron dendrites: Implications for motor output. *Muscle & Nerve*, *31*(2), 135–156.
<https://doi.org/10.1002/MUS.20261>
- Hirsch, M. W., Smale, S., & Devaney, R. L. (2013). Differential Equations, Dynamical Systems, and an Introduction to Chaos. *Differential Equations, Dynamical Systems, and an Introduction to Chaos*. <https://doi.org/10.1016/C2009-0-61160-0>

- Huang, H. J., Kram, R., & Ahmed, A. A. (2012). Reduction of metabolic cost during motor learning of arm reaching dynamics. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *32*(6), 2182–2190.
<https://doi.org/10.1523/JNEUROSCI.4003-11.2012>
- Huffman, K. J., & Krubitzer, L. (2001). Thalamo-cortical connections of areas 3a and M1 in marmoset monkeys. *Journal of Comparative Neurology*, *435*(3), 291–310.
<https://doi.org/10.1002/cne.1031>
- Inagaki, H. K., Fontolan, L., Romani, S., & Svoboda, K. (2019). Discrete attractor dynamics underlies persistent activity in the frontal cortex. *Nature*, *566*(7743), 212–217.
<https://doi.org/10.1038/S41586-019-0919-7>
- Jankowska, E., Hammar, I., Chojnicka, B., & Hedén, C. H. (2000). Effects of monoamines on interneurons in four spinal reflex pathways from group I and/or group II muscle afferents. *European Journal of Neuroscience*, *12*(2), 701–714.
<https://doi.org/10.1046/j.1460-9568.2000.00955.x>
- Jones, E. G., & Porter, R. (1980). What is area 3a? *Brain Research Reviews*, *2*(1–3), 1–43.
[https://doi.org/10.1016/0165-0173\(80\)90002-8](https://doi.org/10.1016/0165-0173(80)90002-8)
- Kaiser, D. (2012). In retrospect: The Structure of Scientific Revolutions. *Nature* *2012* *484*:7393, *484*(7393), 164–165. <https://doi.org/10.1038/484164a>
- Kalidindi, H. T., Cross, K. P., Lillicrap, T. P., Omrani, M., Falotico, E., Sabes, P. N., & Scott, S. H. (2021). Rotational dynamics in motor cortex are consistent with a feedback controller. *ELife*, *10*. <https://doi.org/10.7554/ELIFE.67256>

- Kaufman, M. T., Churchland, M. M., Ryu, S. I., & Shenoy, K. V. (2014). Cortical activity in the null space: Permitting preparation without movement. *Nature Neuroscience*, *17*(3), 440–448. <https://doi.org/10.1038/nn.3643>
- Kawato, M. (1999). Internal models for motor control and trajectory planning. *Current Opinion in Neurobiology*, *9*(6), 718–727. [https://doi.org/10.1016/S0959-4388\(99\)00028-8](https://doi.org/10.1016/S0959-4388(99)00028-8)
- Kramer, M. A. (1991). Nonlinear principal component analysis using autoassociative neural networks. *AIChE Journal*, *37*(2), 233–243. <https://doi.org/10.1002/aic.690370209>
- Lemon, R. N. (1997). Mechanisms of Cortical Control of Hand Function. *The Neuroscientist*, *3*(6), 389–398. <https://doi.org/10.1177/107385849700300612>
- Li, N., Daie, K., Svoboda, K., & Druckmann, S. (2016). Robust neuronal dynamics in premotor cortex during motor planning. *Nature*, *532*(7600). <https://doi.org/10.1038/nature17643>
- Maeda, R. S., Cluff, T., Gribble, P. L., & Pruszynski, J. A. (2018). Feedforward and feedback control share an internal model of the arm's dynamics. *Journal of Neuroscience*, *38*(49), 10505–10514. <https://doi.org/10.1523/JNEUROSCI.1709-18.2018>
- Mazzoni, P., & Krakauer, J. W. (2006). An Implicit Plan Overrides an Explicit Strategy during Visuomotor Adaptation. *Journal of Neuroscience*, *26*(14), 3642–3645. <https://doi.org/10.1523/JNEUROSCI.5317-05.2006>
- Miall, R. C., Christensen, L. O. D., Cain, O., & Stanley, J. (2007). Disruption of State Estimation in the Human Lateral Cerebellum. *PLOS Biology*, *5*(11), e316. <https://doi.org/10.1371/JOURNAL.PBIO.0050316>
- Miri, A., Warriner, C. L., Seely, J. S., Elsayed, G. F., Cunningham, J. P., Churchland, M. M., & Jessell, T. M. (2017). Behaviorally Selective Engagement of Short-Latency Effector Pathways by Motor Cortex. *Neuron*, *95*(3), 683-696.e11.

<https://doi.org/10.1016/J.NEURON.2017.06.042/ATTACHMENT/DE856AFD-ACEA-4F45-8E45-8C7E55317B51/MMC3.MP4>

Nashed, J. Y., Crevecoeur, F., & Scott, S. H. (2014). Rapid online selection between multiple motor plans. *Journal of Neuroscience*, *34*(5), 1769–1780.

<https://doi.org/10.1523/JNEUROSCI.3063-13.2014>

Naufel, S., Glaser, J. I., Kording, K. P., Perreault, E. J., & Miller, L. E. (2019). A muscle-activity-dependent gain between motor cortex and emg. *Journal of Neurophysiology*, *121*(1), 61–73.

<https://doi.org/10.1152/JN.00329.2018/ASSET/IMAGES/LARGE/Z9K0011948910009.JPEG>

Nowak, D. A., Timmann, D., & Hermsdörfer, J. (2007). Dexterity in cerebellar agenesis.

Neuropsychologia, *45*(4), 696–703.

<https://doi.org/10.1016/j.neuropsychologia.2006.08.011>

Oby, E. R., Golub, M. D., Hennig, J. A., Degenhart, A. D., Tyler-Kabara, E. C., Yu, B. M., Chase, S. M., & Batista, A. P. (2019). New neural activity patterns emerge with long-term learning. *Proceedings of the National Academy of Sciences of the United States of America*, *116*(30), 15210–15215. <https://doi.org/10.1073/PNAS.1820296116/-/DCSUPPLEMENTAL>

OpenAI, Akkaya, I., Andrychowicz, M., Chociej, M., Litwin, M., McGrew, B., Petron, A., Paino, A., Plappert, M., Powell, G., Ribas, R., Schneider, J., Tezak, N., Tworek, J., Welinder, P., Weng, L., Yuan, Q., Zaremba, W., & Zhang, L. (2019). Solving Rubik's Cube with a Robot Hand. *ArXiv:1910.07113 [Cs, Stat]*. <http://arxiv.org/abs/1910.07113>

- Padberg, J., Cerkevich, C., Engle, J., Rajan, A. T., Recanzone, G., Kaas, J., & Krubitzer, L. (2009). Thalamocortical connections of parietal somatosensory cortical fields in macaque monkeys are highly divergent and convergent. *Cerebral Cortex*, *19*(9), 2038–2064. <https://doi.org/10.1093/cercor/bhn229>
- Pandarinath, C., O’Shea, D. J., Collins, J., Jozefowicz, R., Stavisky, S. D., Kao, J. C., Trautmann, E. M., Kaufman, M. T., Ryu, S. I., Hochberg, L. R., Henderson, J. M., Shenoy, K. V., Abbott, L. F., & Sussillo, D. (2018). Inferring single-trial neural population dynamics using sequential auto-encoders. *Nature Methods* *2018 15:10*, *15*(10), 805–815. <https://doi.org/10.1038/s41592-018-0109-9>
- Panigrahi, B., Martin, K. A., Li, Y., Graves, A. R., Vollmer, A., Olson, L., Mensh, B. D., Karpova, A. Y., & Dudman, J. T. (2015). Dopamine Is Required for the Neural Representation and Control of Movement Vigor. *Cell*, *162*(6), 1418–1430. <https://doi.org/10.1016/j.cell.2015.08.014>
- Pavlidis, C., Miyashita, E., & Asanuma, H. (1993). Projection from the sensory to the motor cortex is important in learning motor skills in the monkey. *Journal of Neurophysiology*, *70*(2), 733–741. <https://doi.org/10.1152/jn.1993.70.2.733>
- Perich, M. G., Gallego, J. A., & Miller, L. E. (2018). A Neural Population Mechanism for Rapid Learning. *Neuron*, *100*(4), 964-976.e7. <https://doi.org/10.1016/j.neuron.2018.09.030>
- Phillips, C. G., Powell, T. P., & Wiesendanger, M. (1971). Projection from low-threshold muscle afferents of hand and forearm to area 3a of baboon’s cortex. *The Journal of Physiology*, *217*(2), 419–446.
- Pruszynski, J. A. (2014). Primary motor cortex and fast feedback responses to mechanical perturbations: A primer on what we know now and some suggestions on what we should

- find out next. *Frontiers in Integrative Neuroscience*, 8, 1–7.
<https://doi.org/10.3389/FNINT.2014.00072>
- Pruszynski, J. A., Kurtzer, I., Nashed, J. Y., Omrani, M., Brouwer, B., & Scott, S. H. (2011). Primary motor cortex underlies multi-joint integration for fast feedback control. *Nature*, 478(7369), 387–390. <https://doi.org/10.1038/nature10436>
- Prut, Y., & Fetz, E. E. (1999). Primate spinal interneurons show pre-movement instructed delay activity. *Nature*, 401(6753), 590–594. <https://doi.org/10.1038/44145>
- Ramnani, N. (2006). The primate cortico-cerebellar system: Anatomy and function. *Nature Reviews Neuroscience*, 7(7), 511–522. <https://doi.org/10.1038/nrn1953>
- Ranganathan, R., Wieser, J., Mosier, K. M., Mussa-Ivaldi, F. A., & Scheidt, R. A. (2014). Learning Redundant Motor Tasks with and without Overlapping Dimensions: Facilitation and Interference Effects. *Journal of Neuroscience*, 34(24), 8289–8299. <https://doi.org/10.1523/JNEUROSCI.4455-13.2014>
- Rathelot, J.-A., & Strick, P. L. (2009). *Subdivisions of primary motor cortex based on cortico-motoneuronal cells*. www.pnas.org/cgi/content/full/
- Russo, A. A., Bittner, S. R., Perkins, S. M., Seely, J. S., London, B. M., Lara, A. H., Miri, A., Marshall, N. J., Kohn, A., Jessell, T. M., Abbott, L. F., Cunningham, J. P., & Churchland, M. M. (2018). Motor Cortex Embeds Muscle-like Commands in an Untangled Population Response. *Neuron*, 97(4), 953-966.e8. <https://doi.org/10.1016/j.neuron.2018.01.004>
- Sadtler, P. T., Quick, K. M., Golub, M. D., Chase, S. M., Ryu, S. I., Tyler-Kabara, E. C., Yu, B. M., & Batista, A. P. (2014). Neural constraints on learning. *Nature*, 512(7515), 423–426. <https://doi.org/10.1038/nature13665>

- Sauerbrei, B. A., Guo, J.-Z., Cohen, J. D., Mischiati, M., Guo, W., Kabra, M., Verma, N., Mensh, B., Branson, K., & Hantman, A. W. (2019). Cortical pattern generation during dexterous movement is input-driven. *Nature* 2019 577:7790, 577(7790), 386–391.
<https://doi.org/10.1038/s41586-019-1869-9>
- Scheidt, R. A., Reinkensmeyer, D. J., Conditt, M. A., Rymer, W. Z., & Mussa-Ivaldi, F. A. (2000). Persistence of Motor Adaptation During Constrained, Multi-Joint, Arm Movements. *Journal of Neurophysiology*, 84(2), 853–862.
<https://doi.org/10.1152/jn.2000.84.2.853>
- Scott, S. H. (2004). Optimal feedback control and the neural basis of volitional motor control. *Nature Reviews Neuroscience*, 5(7), 532–544. <https://doi.org/10.1038/nrn1427>
- Scott, S. H. (2016). A Functional Taxonomy of Bottom-Up Sensory Feedback Processing for Motor Actions. *Trends in Neurosciences*, 39(8), 512–526.
<https://doi.org/10.1016/j.tins.2016.06.001>
- Scott, S. H., Cluff, T., Lowrey, C. R., & Takei, T. (2015). Feedback control during voluntary motor actions. *Current Opinion in Neurobiology*, 33, 85–94.
<https://doi.org/10.1016/j.conb.2015.03.006>
- Semedo, J. D., Zandvakili, A., Machens, C. K., Yu, B. M., & Kohn, A. (2019). Cortical Areas Interact through a Communication Subspace. *Neuron*, 102(1), 249-259.e4.
<https://doi.org/10.1016/j.neuron.2019.01.026>
- Serrien, D. J., & Wiesendanger, M. (1999). Role of the cerebellum in tuning anticipatory and reactive grip force responses. *Journal of Cognitive Neuroscience*, 11(6), 672–681.
<https://doi.org/10.1162/089892999563634>

- Sha, D., Patton, J., & Mussa-Ivaldi, F. A. (2006). MINIMUM JERK REACHING MOVEMENTS OF HUMAN ARM WITH MECHANICAL CONSTRAINTS AT ENDPOINT. *International Journal of Computers, Systems, and Signals*, 7(1).
- Shadmehr, R. (2020). Population coding in the cerebellum and its implications for learning from error. *BioRxiv*, 2020.05.18.102376. <https://doi.org/10.1101/2020.05.18.102376>
- Shadmehr, R., Brashers-krug, T., & Y, O. M. (1995). Interference in learning internal models of inverse dynamics in humans. *Eds.*, *Advances in Neural Information Processing Systems* 7, 111–117.
- Shadmehr, R., & Krakauer, J. W. (2008). A computational neuroanatomy for motor control. *Experimental Brain Research*, 185(3), 359–381. <https://doi.org/10.1007/s00221-008-1280-5>
- Shadmehr, R., & Mussa-Ivaldi, F. A. (1994). Adaptive representation of dynamics during learning of a motor task. *Journal of Neuroscience*, 14(5), 3208–3224. <https://doi.org/10.1523/JNEUROSCI.14-05-03208.1994>
- Shalit, U., Zinger, N., Joshua, M., & Prut, Y. (2012). Descending Systems Translate Transient Cortical Commands into a Sustained Muscle Activation Signal. *Cerebral Cortex*, 22(8), 1904–1914. <https://doi.org/10.1093/CERCOR/BHR267>
- Shenoy, K. V., Sahani, M., & Churchland, M. M. (2013). Cortical Control of Arm Movements: A Dynamical Systems Perspective. *Annual Review of Neuroscience*, 36(1), 337–359. <https://doi.org/10.1146/annurev-neuro-062111-150509>
- Shine, J. M., Müller, E. J., Munn, B., Cabral, J., Moran, R. J., & Breakspear, M. (2021). Computational models link cellular mechanisms of neuromodulation to large-scale

- neural dynamics. *Nature Neuroscience*, 24(6), 765–776. <https://doi.org/10.1038/s41593-021-00824-6>
- Sun, X., O’Shea, D. J., Golub, M. D., Trautmann, E. M., Vyas, S., Ryu, S. I., & Shenoy, K. V. (2020). *Skill-specific changes in cortical preparatory activity during motor learning* (p. 2020.01.30.919894). <https://doi.org/10.1101/2020.01.30.919894>
- Sussillo, D. (2014). Neural circuits as computational dynamical systems. *Current Opinion in Neurobiology*, 25, 156–163. <https://doi.org/10.1016/j.conb.2014.01.008>
- Sussillo, D., & Barak, O. (2013). Opening the Black Box: Low-Dimensional Dynamics in High-Dimensional Recurrent Neural Networks. *Neural Computation*, 25(3), 626–649. https://doi.org/10.1162/NECO_a_00409
- Sussillo, D., Churchland, M. M., Kaufman, M. T., & Shenoy, K. V. (2015). A neural network that finds a naturalistic solution for the production of muscle activity. *Nature Neuroscience*, 18(7), 1025–1033. <https://doi.org/10.1038/nn.4042>
- Sussillo, D., Jozefowicz, R., Abbott, L. F., & Pandarinath, C. (2016). *LFADS - Latent Factor Analysis via Dynamical Systems*. <http://arxiv.org/abs/1608.06315>
- Takei, T., Lomber, S. G., Cook, D. J., & Scott, S. H. (2021). Transient deactivation of dorsal premotor cortex or parietal area 5 impairs feedback control of the limb in macaques. *Current Biology*, 31(7), 1476-1487.e5. <https://doi.org/10.1016/J.CUB.2021.01.049>
- Taylor, J. A., Krakauer, J. W., & Ivry, R. B. (2014). Explicit and implicit contributions to learning in a sensorimotor adaptation task. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, 34(8), 3023–3032. <https://doi.org/10.1523/JNEUROSCI.3619-13.2014>

- Tobin, J., Fong, R., Ray, A., Schneider, J., Zaremba, W., & Abbeel, P. (2017). Domain Randomization for Transferring Deep Neural Networks from Simulation to the Real World. *ArXiv:1703.06907 [Cs]*. <http://arxiv.org/abs/1703.06907>
- Todorov, E., & Jordan, M. I. (2002). Optimal feedback control as a theory of motor coordination. *Nature Neuroscience*, *5*(11), 1226–1235. <https://doi.org/10.1038/nn963>
- Todorov, E., & Jordan, M. I. (2003). A Minimal Intervention Principle for Coordinated Movement. *Advances in Neural Information Processing Systems*, *15*, 27–34.
- Tsuda, B., Pate, S. C., Tye, K. M., Siegelmann, H. T., & Sejnowski, T. J. (2021). *Neuromodulators enable overlapping synaptic memory regimes and nonlinear transition dynamics in recurrent neural networks* (p. 2021.05.31.446462). <https://doi.org/10.1101/2021.05.31.446462>
- van der Maaten, L., & Hinton, G. (2008). Visualizing data using t-SNE. *Journal of Machine Learning Research*, *9*, 2579–2605.
- Vitrac, C., & Benoit-Marand, M. (2017). Monoaminergic Modulation of Motor Cortex Function. *Frontiers in Neural Circuits*, *11*, 72. <https://doi.org/10.3389/fncir.2017.00072>
- Wolpert, D. M., Miall, R. C., & Kawato, M. (1998). Internal models in the cerebellum. *Trends in Cognitive Sciences*, *2*(9), 338–347. [https://doi.org/10.1016/S1364-6613\(98\)01221-2](https://doi.org/10.1016/S1364-6613(98)01221-2)
- Yu, B. M., Cunningham, J. P., Santhanam, G., Ryu, S. I., Shenoy, K. V., & Sahani, M. (2009). Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity (Journal of Neurophysiology (2009) 102, (614-635) DOI: 10.1152/jn.90941.2008). *Journal of Neurophysiology*, *102*(3), 2008. <https://doi.org/10.1152/jn.z9k-9684-corr.2009>

Yumiya, H., Kubota, K., & Asanuma, H. (1974). Activities of neurons in area 3a of the cerebral cortex during voluntary movements in the monkey. *Brain Research*, 78(2), 169–177.

[https://doi.org/10.1016/0006-8993\(74\)90544-7](https://doi.org/10.1016/0006-8993(74)90544-7)

Zarzecki, P., & Asanuma, H. (1979). Proprioceptive Influences on Somatosensory and Motor Cortex. *Progress in Brain Research*, 50(C), 113–119. [https://doi.org/10.1016/S0079-6123\(08\)60812-2](https://doi.org/10.1016/S0079-6123(08)60812-2)