Optimal anticipatory control of movement as a theory of motor preparation: a thalamo-cortical circuit model

Ta-Chu Kao[1], Mahdieh S. Sadabadi[1,2], and Guillaume Hennequin[@1]

[1]Computational and Biological Learning Lab, Department of Engineering, University of Cambridge, Cambridge, U.K.
[2]Department of Automatic Control and Systems Engineering, University of Sheffield, Sheffield, U.K.

[@] Corresponding author (g.hennequin@eng.cam.ac.uk)

Summary

Across a range of motor and cognitive tasks, cortical activity can be accurately described by low-dimensional dynamics unfolding from specific initial conditions on every trial. These "preparatory states" largely determine the subsequent evolution of both neural activity and behaviour, and their importance raises questions regarding how they are — or ought to be — set. Here, we formulate motor preparation as optimal prospective control of future movements. The solution is a form of internal control of cortical circuit dynamics, which can be implemented as a thalamo-cortical loop gated by the basal ganglia. Critically, optimal control predicts selective quenching of variability in components of preparatory population activity that have future motor consequences, but not in others. This is consistent with recent perturbation experiments performed in mice, and with our novel analysis of monkey motor cortex activity during reaching. Together, these results suggest optimal anticipatory control of movement.

Fast ballistic movements (e.g. throwing) require spatially and temporally precise commands to the musculature. Many of these signals are thought to arise from internal dynamics in the primary motor cortex (M1; Figure 1A; Evarts, 1968; Todorov, 2000; Scott, 2012; Shenoy et al., 2013; Omrani et al., 2017). In turn, consistent with state trajectories produced by a dynamical system, M1 activity during movement depends strongly on the "initial condition" reached just before movement onset, and variability in initial condition predicts behavioural variability (Churchland et al., 2006a; Afshar et al., 2011; Pandarinath et al., 2018). An immediate consequence of this dynamical systems view is the so-called "optimal subspace hypothesis" (Churchland et al., 2010b; Shenoy et al., 2013): the network dynamics that generate movement must be seeded with an appropriate initial condition prior to each movement. In other words, accurate movement production likely requires fine adjustment of M1 activity during a phase of movement preparation (Figure 1B, green).

The optimal subspace hypothesis helps to make sense of neural activity during the preparation epoch, yet several unknowns remain. What should the structure of the optimal preparatory subspace be? How does this structure depend on the dynamics of the cortical network during the movement epoch, and on downstream motor processes? Must preparatory activity converge to a single movement-specific state and be held there until movement initiation, or is some slack allowed? What are the dynamical processes and associated circuit mechanisms responsible for motor preparation? These questions can be (and have been partially) addressed empirically, e.g. through analyses of neural population recordings in reaching monkeys (Churchland et al., 2010b; Ames et al., 2014; Elsayed et al., 2016) or optogenetic dissection of circuits involved in motor preparation (Li et al., 2016; Guo et al., 2017; Gao et al., 2018; Sauerbrei et al., 2019). Yet, for lack of an appropriate theoretical scaffold, it has been difficult to interpret these experimental results within the broader computational context of motor control.

Here, we bridge this gap by considering motor preparation as an integral part of motor control. We show that optimal control theory, which has successfully explained behaviour (Todorov and Jordan, 2002; Scott et al., 2015) and neural activity (Todorov, 2000; Lillicrap and Scott, 2013) during the movement epoch, can also be brought to bear on motor preparation. Specifically, we argue that there is a prospective component of motor control that can be performed in anticipation of the movement (i.e. during preparation). This leads to a concrete normative formulation of the optimal subspace hypothesis. Our theory specifies the inputs that must be given to the movement-generating network during preparation to ensure that (i) any subsequent motor errors are kept minimal and (ii) movements can be triggered rapidly.

We provide a full circuit implementation of the optimal anticipatory control strategy, for a specific model of M1 that we have proposed previously (Hennequin et al., 2014 — though the framework is general). In particular, we propose that cortex is actively controlled via a tha-

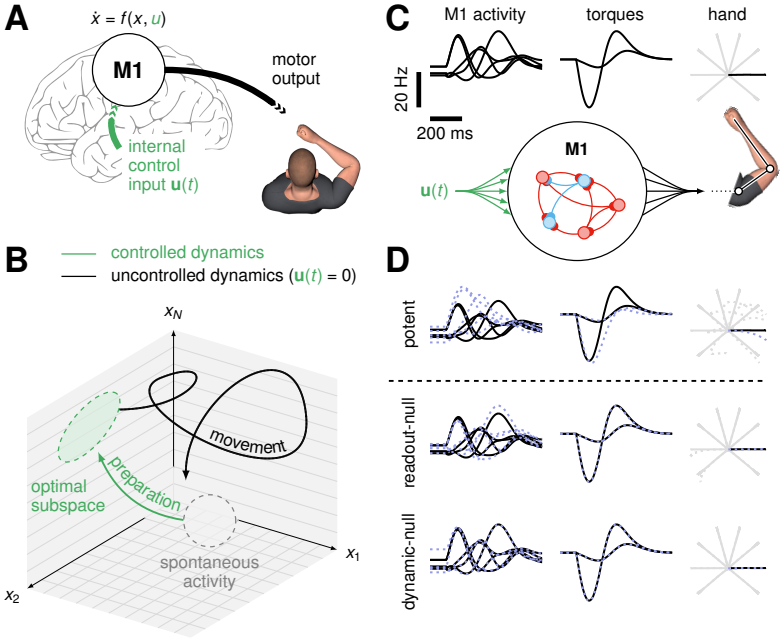

Figure 1: **Preparation & execution of ballistic movements.** **(A)** Under a dynamical systems view of motor control (Shenoy et al., 2013), movement is generated by internal dynamics in M1. Prior to movement, the population activity state $\mathbf{x}(t)$ must be controlled into an optimal, movement-specific subspace in a phase of movement preparation; this requires internally generated control inputs $\mathbf{u}(t)$. **(B)** Schematic state space trajectory during movement preparation and execution. **(C)** Schematics of our M1 model of motor pattern generation. The dynamics of an excitation-inhibition network (Hennequin et al., 2014) unfold from movement-specific initial conditions, resulting in firing rate trajectories (left; 5 neurons shown) which are linearly read out into joint torques (middle), thereby producing hand movements (right). The model is calibrated for the production of eight straight center-out reaches; firing rates and torques are shown only for the movement colored black. To help visualize initial conditions, firing rates are artificially clamped for the first 100 ms. **(D)** Effect of three qualitatively different types of small perturbations of the initial condition on the three processing stages leading to movement, as already shown in (C). Unperturbed traces are shown as solid lines, perturbed ones as dashed lines. Perturbations of all types on the initial condition have the exact same size, but different consequences. "Potent" perturbations (top) result in errors at every stage. "Readout-null" perturbations (middle) cause sizeable changes in internal network activity but not in the readout. "Dynamic-null" perturbations are inconsequential at every stage.

lamocortical loop during motor preparation, with thalamic afferents providing the desired optimal control inputs. This is consistent with the causal role of thalamus in the preparation of directed licking in mice (Guo et al., 2017). Moreover, we posit that the basal ganglia operate an ON/OFF switch on the thalamocortical loop (Jin and Costa, 2010; Cui et al., 2013; Halassa and Acsády, 2016; Logiaco et al., 2019), thereby flexibly controlling the timing of both movement planning and initiation.

We further analyze the model, and formulate predictions which we have successfully tested in data. At the most abstract level, our core prediction is that the "optimal subspace" is likely high dimensional, with many different initial conditions giving rise to the same correct movement. This has an important consequence for preparatory control: at the population level, only a few components of preparatory activity impact future motor outputs, and it is these components only that need active controlling. In contrast, one expects substantial pre-movement variability in other, inconsequential components. Concretely, we predict that following a perturbation, preparatory activity should recover only in state space directions that matter for subsequent movement, but not (necessarily) in others. We find that this prediction agrees with the effects of optogenetic perturbations reported by Svoboda and colleagues, in a directed licking task in mice (Li et al., 2016). Furthermore, the existence of a preparatory nullspace predicts selective variability quenching at preparation onset: trial-by-trial variability should drop predominantly in components that have motor consequences. We perform novel analyses of monkey M1 and dorsal premotor cortex (PMd) activity recorded during reaching, and find that the structure of variability quenching supports our main prediction. Finally, our model also predicts that population activity should evolve in orthogonal subspaces during preparation and movement, which is one of the most prominent features of perimovement activity in reaching monkeys (Kaufman et al., 2014; Elsayed et al., 2016).

Beyond motor control, there is a broader set of cortical computations that are also thought to rest on low dimensional circuit dynamics, with initial conditions largely determining behaviour (Pandarinath et al., 2018; Sohn et al., 2019). These computations, too, may hinge on careful preparation of the state of cortex in appropriate subspaces. Our framework, and control theory more

generally, may provide a useful language for reasoning about putative algorithms and neural mechanisms (Kao and Hennequin, 2019).

Results

A model of movement generation

We begin with a network model of motor cortex in which a detailed balance of excitation and inhibition enables the production of rich, naturalistic activity transients (Hennequin et al., 2014, Figure 1C). This network serves as a pattern generator for the production of movement. Specifically, the network directly controls the two joint torques of a two-link arm (Section S2.2), via a linear readout of the momentary network firing rates:

$$\mathbf{m}(t) = \mathbf{C}\mathbf{r}(t). \tag{1}$$

Here, $\mathbf{m}(t)$ is a vector containing the momentary torques, and $\mathbf{r}(t)$ is the population firing rate vector (described below). We assume that the output torques are artificially silenced during movement preparation. The network has $N = 200$ neurons, whose momentary internal activations $\mathbf{x}(t) = (x_1, x_2, \ldots, x_N)^T$ evolve according to (Dayan and Abbott, 2001; Section S2.1):

$$\tau \frac{d\mathbf{x}}{dt} = -\mathbf{x}(t) + \mathbf{W}\,\mathbf{r}(t) + \overline{\mathbf{h}} + \mathbf{h}(t) \tag{2}$$

$$\mathbf{r}(t) = \phi[\mathbf{x}(t)]\,. \tag{3}$$

Here, τ is the single-neuron time constant, $\mathbf{W}$ is the synaptic connectivity matrix, and $\phi[x]$ (applied to $\mathbf{x}$ element-wise) is a rectified-linear activation function converting internal activations into momentary firing rates. The network is driven by two different inputs shared across all movements: a constant input $\overline{\mathbf{h}}$ responsible for maintaining spontaneous activity, and a transient input $\mathbf{h}(t)$ arising at movement onset and decaying through movement. The latter input models the dominant, condition-independent timing-related component of monkey M1 activity during movement (Kaufman et al., 2016). We note that, while the network model is generally nonlinear, it can be well approximated by a linear model ($\mathbf{r} = \mathbf{x}$) as only a small fraction of neurons are silent at any given time (see below); our formal analyses here rely on linear approximations, but all simulations are based on Equations 2 and 3 with nonlinear ϕ.

We calibrated the model for the production of eight rapid straight reaches with bell-shaped velocity profiles (Figure 1C; details in Section S2.3, see also Figure S1). This would later enable comparison with data recorded in monkeys performing similar movements. To perform this calibration, we noted that—in line with the dynamical systems view of movement generation (Shenoy et al., 2013)— movements produced by our model depend strongly on the "initial condition", i.e. the cortical state $\mathbf{x}$ just before movement onset (Churchland et al., 2010b; Afshar et al., 2011). We thus "inverted" the model numerically, by finding eight different initial conditions and a common readout matrix $\mathbf{C}$ such that the dynamics of the nonlinear model (Equations 2 and 3), seeded with each initial condition, would produce the desired movement. Importantly, we constrained $\mathbf{C}$ so that its nullspace contained the network's spontaneous activity state, as well as all eight initial conditions. This constraint ensures that movement does not occur spontaneously, and is a minimum requirement (though not a guarantee) for movement not to occur prematurely during preparation.

We re-analyzed population recordings of monkey M1/PMd during reaching (data courtesy of Mark Churchland, Matt Kaufman and Krishna Shenoy; Section S5), and found that our model captures several essential aspects of movement-related neural dynamics (Figure S5). First, kinematically similar reaches are produced from similar preparatory end-states in both model and monkey. Second, neurons exhibit heterogeneous, multiphasic oscillatory activity that often grow transiently from the preparatory end-state before shrinking back to spontaneous levels (Hennequin et al., 2014). Third, these transients can be summarized as state-space rotations at the population level, as revealed by jPCA (Churchland et al., 2012). Finally, canonical correlations analysis (Sussillo et al., 2015) indicates substantial overlap between monkey and model population activity across time and conditions (Section S5).

Control of fast movements: preparatory control of cortical activity

Having calibrated our network model of movement generation, we now turn to preparatory dynamics. As stated previously, Shenoy et al.'s dynamical systems perspective suggests that accurate movement execution likely requires careful seeding of the generator's dynamics with an appropriate, reach-specific initial condition (Afshar et al., 2011). In our model, this means that the activity state $\mathbf{x}(t)$ of the cortical network must be steered towards the initial condition corresponding to the intended movement (Figure 1B, green). This process, which we call "preparatory control", forms the core of this study.

An important first step towards formalizing preparatory control and unravelling putative circuit mechanisms is to understand how deviations from "the right initial condition" impact the subsequent movement. Mathematical analysis of our model reveals that depending on the direction in state space along which the deviation occurs, there may be strong motor consequences or none at all (Figure 1D; Section S3). Some preparatory perturbations propagate through the dynamics of the generator network during the movement epoch, modifying its activity trajectories, eventually leading to er-

rors in torques and hand motion ("potent perturbations"; Figure 1D, top). Other preparatory perturbations cause subsequent deviations in cortical state trajectories, too, but these deviations are correlated across neurons in such a way that they cancel in the readout, leaving the movement unaltered ("readout-null"; Figure 1D, middle). Yet other preparatory perturbations are outright rejected by the recurrent dynamics of the network. These perturbations have little impact on neuronal activity, let alone on torques and hand motion ("dynamic-null"; Figure 1D, bottom).

The existence of readout-null and dynamic-null directions imply that, for each movement, many initial conditions give rise to the correct hand trajectory. This substantially lightens the computational burden of preparatory ballistic control: there are only a few potent directions in state space along which cortical activity needs active controlling prior to movement. For our model with only two readout torques, the effective dimensionality of this potent subspace is approx. 5 (Section S3, see also Figure S3 and Figure 4). Thus, taking into account the energetic cost of neural control, preparatory dynamics should aim at preferentially eliminating errors in preparatory states along those few directions that matter for movement.

We now formalize these insights in a normative model of preparatory motor control. We assume that, prior to movement, the initial condition for cortical dynamics is progressively reached during a preliminary phase of movement preparation. In this phase, the cortical network receives additional movement-specific control inputs $\mathbf{u}(t)$ (Figure 1A and B, green) which are rapidly switched off to initiate movement:

$$\tau \frac{d\mathbf{x}}{dt} = -\mathbf{x}(t) + \mathbf{W}\,\mathbf{r}(t) + \overline{\mathbf{h}} + \mathbf{h}(t) + \mathbf{u}(t) \qquad (4)$$

How should these preparatory inputs $\mathbf{u}(t)$ be chosen? At any time t during preparation, we can assign a "prospective motor error" $\mathcal{C}(\mathbf{x})$ to the current cortical state $\mathbf{x}(t)$, equal to the total error in movement that *would* result if movement was initiated at this time (i.e. if control inputs were suddenly switched off; Section S3). An ideal controller would supply the cortical network with such control inputs $\mathbf{u}(t)$ as necessary to lower the prospective motor error as fast as possible. This would enable accurate movement production in short order. We therefore propose the following cost functional:

$$\mathcal{J}[\mathbf{u}(t)] = \int_0^\infty \mathcal{C}(\mathbf{x}(t)) + \lambda \mathcal{R}(\mathbf{u}(t))\ dt \qquad (5)$$

where $\mathcal{R}(\mathbf{u})$ is an energetic cost which penalizes large control signals, and λ sets its relative importance in the overall cost. Note that $\mathbf{x}(t)$ depends on $\mathbf{u}(t)$ via Equation 4.

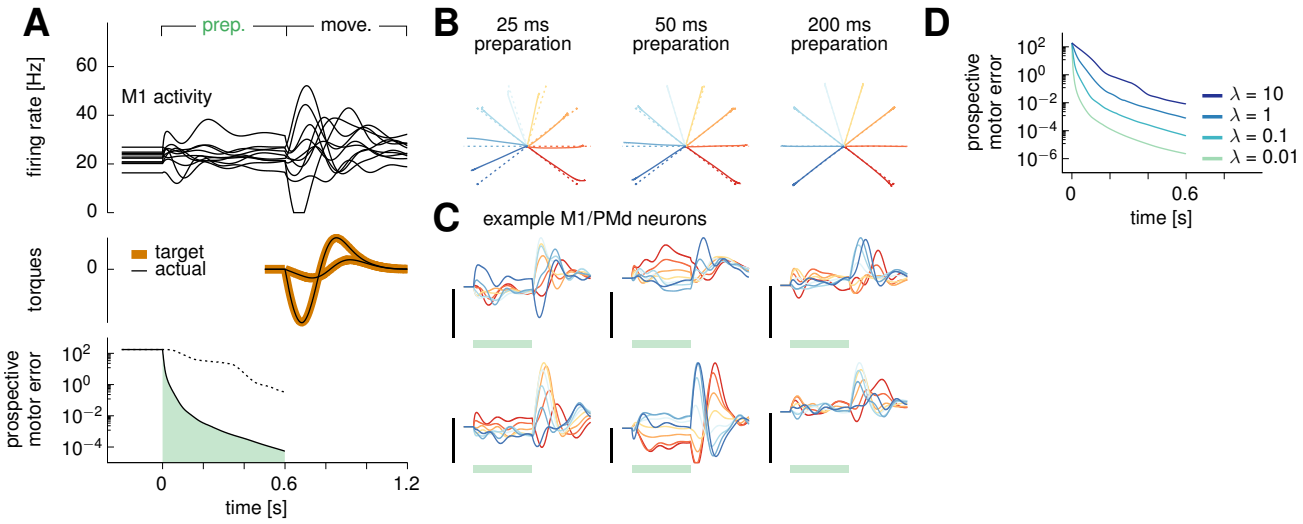

Figure 2: **Optimal preparatory control. (A)** Dynamics of the model during optimal preparation and execution of a straight reach at 0-degree angle. Optimal control inputs are fed to the cortical network during preparation, and subsequently withdrawn to elicit movement. Top: firing rates of a selection of ten model neurons. Middle: generated torques (black), compared to targets (brown). Bottom: the prospective motor error $\mathcal{C}[\mathbf{x}(t)]$ quantifies the accuracy of the movement if it were initiated at time t during the preparatory phase. Under the action of optimal control inputs, $\mathcal{C}[\mathbf{x}(t)]$ decreases very fast, until it becomes small enough that an accurate movement can be triggered. The control input is calculated so as to minimize the green area under the curve over an infinite preparation horizon (only 600 ms of which are shown here), plus an energy cost that prevents control inputs from growing unrealistically large (Equation 5). The dashed line shows the evolution of the prospective cost for the naive static strategy (see text). **(B)** Hand trajectories for each of the eight reaches (solid), following optimal preparation over a window of 25 ms (left), 50 ms (center) and 200 ms (right). Dashed lines show target movements. **(C)** Firing rates of six example neurons, for each movement (color-coded as in B). Green bars mark the 600 ms preparation window, black scale bars indicate 20 Hz. **(D)** Prospective motor error (averaged over the eight reaches) during preparation for different values of the energy penalty parameter λ. In (A-C), we used $\lambda = 0.1$.

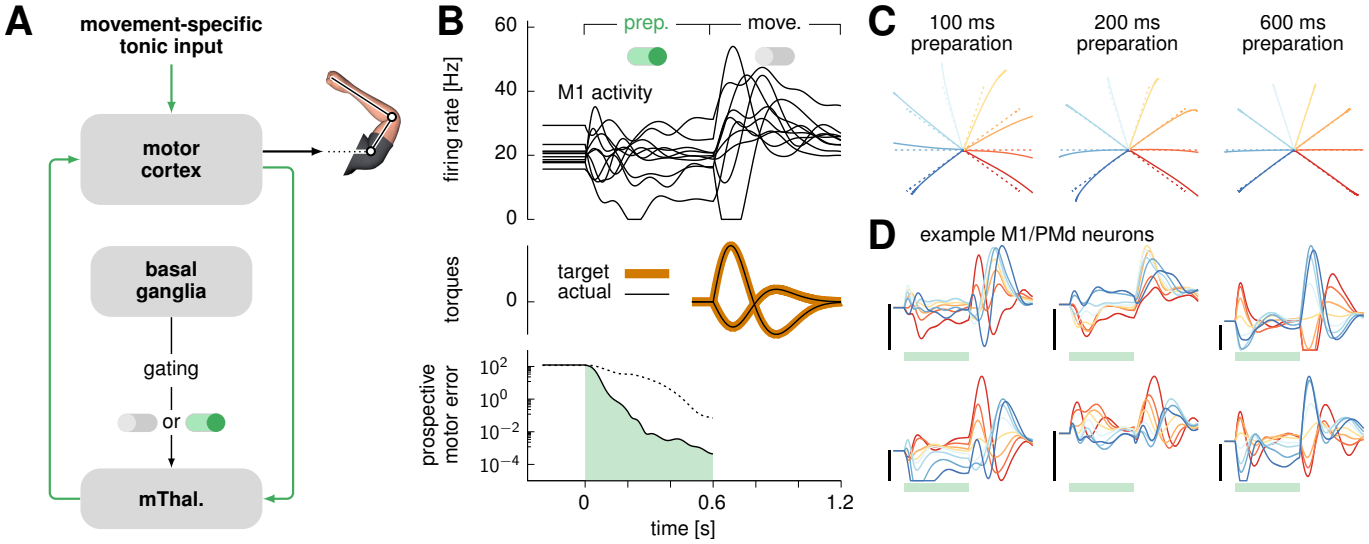

Figure 3: **Optimal movement preparation via a gated thalamo-cortical loop. (A)** Proposed circuit architecture for the optimal movement preparation (cf. text). **(B)** Cortical activity (top; 10 example neurons), generated torques (middle), and prospective motor error (bottom) during the course of movement preparation and execution in the circuit architecture shown in (A). Prospective motor error for the naive strategy is shown with a dotted line as in Figure 2A. **(C)** Hand trajectories (solid) compared to target trajectories (dashed) for the eight reaches, triggered after 100 ms (left), 200 ms (middle) and 600 ms (right) of motor preparation. **(D)** Firing rates of six example cortical neurons in the model. Green bars mark the 600 ms preparation window, during which the thalamus is disinhibited. Vertical scale bars denote 20 Hz.

Optimal preparatory control

When (i) the prospective motor error $\mathcal{C}$ is quadratic in the output torques $\mathbf{m}$, (ii) the energy cost $\mathcal{R}$ is quadratic in $\mathbf{u}$, and (iii) the network dynamics are linear, then minimizing Equation 5 corresponds to the well-known linear quadratic regulator problem in control theory (Skogestad and Postlethwaite, 2007). The optimal solution is a combination of a constant input and instantaneous (linear) state feedback,

$$\mathbf{u}_{\text{opt}}(t) = \mathbf{u}^\star + \mathbf{K}\,\delta\mathbf{x}(t), \tag{6}$$

where $\delta\mathbf{x}(t)$ is the momentary deviation of $\mathbf{x}$ from the desired initial condition. In Equation 6, the constant input $\mathbf{u}^\star$ is movement specific, but the optimal gain matrix $\mathbf{K}$ is generic; both can be derived in algebraic form (Section S4.2). Thus, even though the actual movement occurs in "open loop" (without corrective sensory feedback), optimal movement preparation occurs in closed loop, with the state of the pattern generator being controlled via internal feedback in anticipation of the movement.

When applied to our model system, the optimal control inputs lead to naturalistic transient dynamics in the cortical network during motor preparation (Figure 2A, top, and Figure 2C). The prospective motor error decreases very quickly to negligible values (Figure 2A, bottom; note the small green area under the curve) as $\mathbf{x}(t)$ is driven into the appropriate subspace. After the preparatory feedback loop is switched off and movement begins, the system accurately produces the desired torques and hand trajectories (Figure 2A, middle, and Figure 2B, right). Indeed, movements are ready to be performed after as little as 50 ms of preparation time (Figure 2B). We note, though, that it is possible to achieve an arbitrarily small motor cost, and therefore arbitrarily fast preparation, by decreasing the energy penalty factor λ in Equation 5 (Figure 2D). However, this is at the price of unrealistically large control inputs, i.e. large energetic costs $\mathcal{R}(\mathbf{u})$ (Section S4.3).

Importantly, feedback control vastly outperforms a naive preparation strategy which uses a simpler, constant feedforward input $\mathbf{u}(t) = \mathbf{u}^\star$ and ignores the error feedback term ($\mathbf{K} = 0$ in Equation 6). Under this naive strategy, network activity successfully settles in the desired initial condition *eventually*, but undergoes large initial transients in directions of high cost at preparation onset. These transients dramatically delay the eventual decay of the prospective motor error (Figure 2A, bottom, dashed line).

Circuit model for preparatory control: a gated thalamocortical loop

So far we have not discussed the source of optimal preparatory inputs $\mathbf{u}(t)$, other than saying that they close a feedback loop from the cortex onto itself (Equation 6). While such a loop could in principle be absorbed in local modifications of recurrent cortical connectivity (Sussillo and Abbott, 2009), this would preclude the flexible, near-instant ON/OFF switching of the control loop required at onset of preparation (ON) and movement (OFF). If, instead, the preparatory loop were

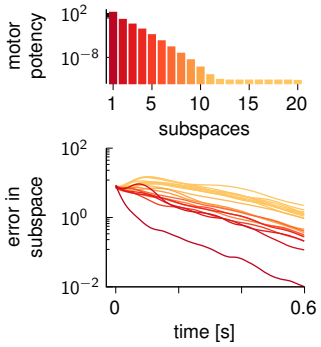

Figure 4: **Selective elimination of preparatory errors along potent directions.** **Top**: 200 orthogonal state-space directions were identified, going from the most to the least "motor potent" (c.f. text). These directions were ordered and grouped into twenty 10-dimensional subspaces. The average motor potency in each subspace is shown here, as measured by the prospective motor error $\mathcal{C}$. **Bottom**: the state of the cortical network in the thalamo-cortical model of Figure 3 was artifically set to deviate randomly from the target movement-specific initial state at time $t = 0$, prior to movement preparation. The squared Euclidean deviation from target (averaged over trials and movements) is decomposed into contributions from the twenty subspaces, and shown using a consistent color code.

to pass through another brain area, fast modulation of excitability in that relay area would provide a rapid and flexible switch. We therefore propose the circuit model shown in Figure 3A, where the motor thalamus acts as a relay station for cortical feedback (Guo et al., 2017; Nakajima and Halassa, 2017). The loop is gated ON/OFF at preparation onset/offset by the (dis)-inhibitory action of basal ganglia outputs (Jin and Costa, 2010; Cui et al., 2013; Halassa and Acsády, 2016; Logiaco et al., 2019). Specifically, cortical excitatory neurons project to 160 thalamic neurons, which make excitatory backprojections to a pool of 100 excitatory (E) and 100 inhibitory (I) neurons in cortex layer 4. In turn, these layer 4 neurons provide both excitation and inhibition to the main cortical network, thereby closing the control loop. Here, inhibition is necessary to capture the negative nature of optimal feedback. In addition to thalamic input, the cortical network also receives a movement-specific constant drive during preparation (analogous to $\mathbf{u}^\star$ in Equation 6 for the standard LQR algorithm).

The detailed patterns of synaptic efficacies in the thalamo-cortical loop are obtained by solving the same control problem as above, based on the minimization of the cost functional in Equation 5 (mathematical details in Section S4). Importantly, the solution must now take into account some key biological constraints: (i) feedback must be based on the activity of the cortical E neurons only, (ii) thalamic and layer-4 neurons have intrinsic dynamics that introduce lag, and (iii) the sign of each connection is constrained by the nature of the presynaptic neuron (E or I).

The circuit model we have obtained enables flexible, near-optimal anticipatory control of the reaching movements (Figure 3B). During spontaneous activity, thalamic neurons are silenced due to strong inhibition from basal ganglia outputs (not explicitly modelled), keeping the thalamocortical loop open (inactive) by default. At the onset of movement preparation, rapid and sustained disinhibition of thalamic neurons restores responsiveness to cortical inputs, thereby closing the control loop (ON/OFF switch in Figure 3B, top). This loop drives the cortical network into the appropriate preparatory subspace, rapidly reducing prospective motor errors (Figure 3B). To trigger movement, the movement-specific tonic input to cortex is shut off, and the basal ganglia resume sustained inhibition of the thalamus. Thus, the loop re-opens, which sets off the uncontrolled dynamics of the cortical network from the right initial condition to produce the desired movement (Figure 3C).

Selective elimination of preparatory errors

The neural trajectories under optimal preparatory control display a striking property, also observed in monkey M1 and PMd recordings (Ames et al., 2014; Lara et al., 2018; Discussion): by the time movement is ready to be triggered (approx. 100 ms), the firing rates of most neurons have not yet converged to the values they would attain after a long preparation time (Figure 3B and D) — that is, $\|\delta\mathbf{x}(t)\| \gg 0$. We found we could readily explain this effect by systematically searching for potent and null directions (thus generalizing Figure 1D) and examining how much each direction contributes to preparatory deviations $\delta\mathbf{x}(t)$. Specifically, we calculated a full set of orthogonal directions ranking from most to least motor-potent, with "potency" measured by the prospective motor error $\mathcal{C}$ (an analytical solution exists for linear systems; Section S3). For easier visualization, we grouped these successive directions into twenty 10-dimensional subspaces. We found that motor potency decreases steeply from one subspace to the next (Figure 4, top), indicating that preparatory deviations are inconsequential along most state-space directions. If this could explain early motor readiness despite unsettled firing rates, one would expect $\delta\mathbf{x}(t)$ to shrink very quickly along directions with motor consequences, while persisting (or even growing) along motor-null directions. To verify this, we artificially set $\mathbf{x}(t)$ at preparation onset to deviate from the target initial state in a random direction in each trial. Examining the dynamics of $\delta\mathbf{x}(t)$ in the various subspaces (averaging squared contributions across many trials), we confirmed that errors are selectively eliminated along directions with motor consequences, while they linger or even grow in other, inconsequential directions (Figure 4, bottom).

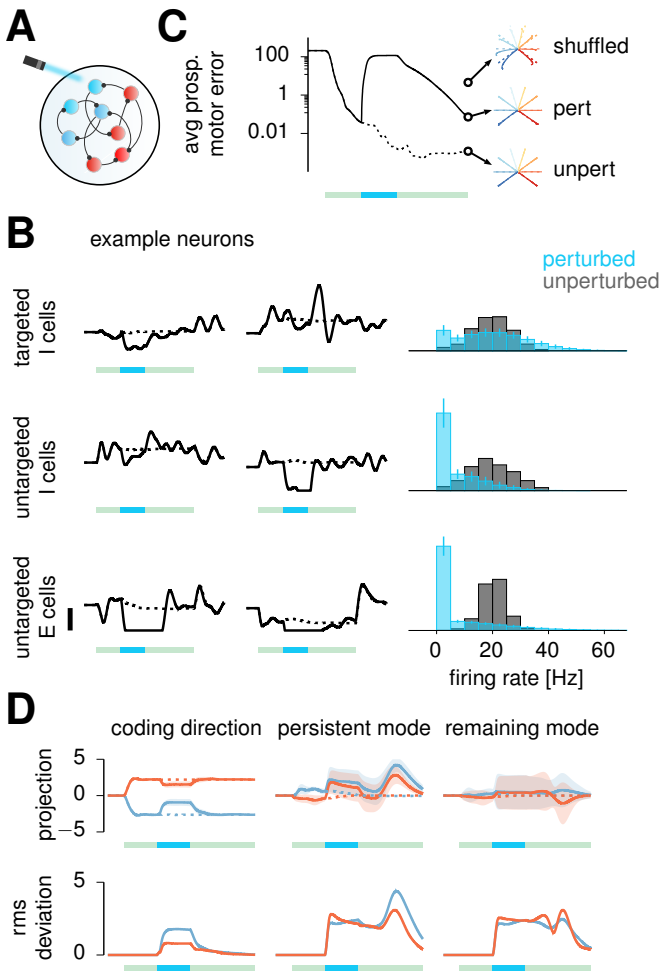

Figure 5: **Effect of perturbations.** **(A)** Illustration of perturbation via "photoinhibition": a subset (60%) of I neurons in the model are driven by strong positive input. **(B)** Left: firing rates (solid: perturbed; dashed: unperturbed) for a pair of targeted I cells (top), untargeted I cells (middle) and E cells (bottom). Green bars (1.6 s) mark the movement preparation epoch, and embedded turquoise bars (400 ms) denote the perturbation period. Right: population histograms of firing rates observed at the end of the perturbation (turquoise), and firing rates observed at the same time in unperturbed trials (gray). Error bars show one standard deviation across 300 experiments, each with a different random set of targeted I cells. **(C)** Prospective motor error (averaged across movements and perturbation experiments) in perturbed (solid) vs. unperturbed (dashed) conditions. Subsequent hand trajectories are shown for one experiment of each condition (middle and bottom insets). These are compared with the reaches obtained by randomly shuffling the deviation of the final perturbed preparatory state from target initial condition across neurons, and simulating the cortical dynamics thereafter (top inset; see also text). Target reaches are shown as dashed lines. **(D)** Analysis of perturbation and recovery for the case of two movements (left and right reaches; same color code as in (C)). Top: cortical activity projected along the coding direction (left), the persistent mode (middle), and the remaining mode (right; see text). Traces are shown as mean (solid) ± std. (shaded) across perturbation experiments, and compared with the unperturbed condition (dashed). Bottom: corresponding root-mean-square deviation between perturbed and unperturbed projections across experiments. Green and turquoise bars as in (B).

Selective recovery from photoinhibition

The selective elimination of preparatory errors also makes predictions for the manner in which the circuit should recover from perturbations. In recent years, neurophysiologists have begun to dissect the causal circuit mechanisms responsible for movement planning, by systematically perturbing activity in various brain structures (Li et al., 2016; Guo et al., 2017; Gao et al., 2018; Sauerbrei et al., 2019), and analyzing population recordings during recovery. We performed similar perturbations and analyses in our circuit model of Figure 3 to shed new light on preparatory motor control (Figure 5).

As our E/I cortical circuit model operates in the inhibition-stabilized regime (Hennequin et al., 2014; Tsodyks et al., 1997; Ozeki et al., 2009; Sanzeni et al., 2019), we were able to use the same photoinhibition strategy as used in these experimental studies to silence the cortical network (Figure 5A). We provided strong excitatory input to a random subset (60%) of inhibitory neurons, for a duration of 400 ms starting 400 ms after preparation onset. We found that "photoinhibition" has mixed effects on the targeted neurons: some are caused to fire at higher rates, but many are paradoxically suppressed (Figure 5B, top). For E cells and untargeted I cells, though, the effect is uniformly suppressive, as shown in Figure 5B (middle and bottom).

The perturbation transiently resets the prospective motor error to pre-preparation level, thus nullifying the benefits of the first 400 ms of preparation (Figure 5C). Following the end of the perturbation, the prospective motor error decreases again, but does not fully recover to its final value in unperturbed trials (Figure 5C, solid vs. dashed). Nevertheless, it recovers to sufficiently low values as to enable accurate movement production (compare middle and bottom hand trajectories in Figure 5C). This is due to the selective elimination of preparatory errors discussed earlier (Figure 4): indeed, shuffling the deviation of $\mathbf{x}(t)$ from the target initial state across neurons (i.e. uniformizing the distribution of errors in different state space directions) leads to a much higher prospective motor error, and eventually to impaired hand trajectories (Figure 5C, top right).

We next performed an analysis qualitatively similar to that conducted by Li et al. in the context of a task where mice had to report the location of a tactile stimulus through directed licking after a delay period. For comparison with the lick-left and lick-right movement conditions, we trimmed our model to only two movements, the left and the right reaches. We identified a "coding direction" (CD) that maximally separates firing

rates in left and right reaches towards the end of movement preparation in unperturbed trials (Section S4.9). Similarly, we identified a "persistent" mode (PM) that maximally separates firing rates in perturbed and unperturbed conditions towards the end of preparation, regardless of the reach direction. Finally, as in Li et al. (2016), we chose a third mode orthogonal to the PM and the CD which captures most of the remaining variance across the two reaches and perturbation conditions.

As Li et al. observed in the anterior lateral motor cortex (ALM) of mice, the CD and the PM modes are nearly orthogonal (89-degree angle) even though not constrained to be so. Moreover, the perturbation causes cortical activity to transiently deviate from unperturbed trajectories nearly equally along each of the three modes (Figure 5D). Remarkably, however, activity recovers promptly along the CD, but not along the other two modes — as in Li et al. (2016). In fact, the perturbation even grows transiently along the PM during early recovery. Such selective recovery can again be understood from optimal preparation eliminating errors along directions with motor consequences, but (owing to energy constraints) not in other inconsequential modes. Indeed, the CD is by definition a motor-potent direction: its contribution to the preparatory state is what determines whether the movement will be a left reach or a right reach. In contrast, the PM and the third mode are approximately motor-null (respectively 4339 times and 2286 times less motor-potent than the CD, by our measure $\mathcal{C}$ of motor potency). Thus, the dynamics of the closed-loop circuit have no incentive to quench perturbation-induced deviations along these modes. In sum, the dynamics of mouse ALM during this task are consistent with our model of optimal preparatory control.

Selective quenching of variability during motor preparation

The selective elimination of motor-potent preparatory errors is a central feature of the optimal control algorithm. Yet, experimental tests of this prediction in reaching monkeys would involve circuit perturbations analogous to those performed in mice (c.f. above), at a spatial resolution finer than currently achievable (O'Shea et al., 2018). Nevertheless, we reasoned that other signatures of this prediction might be present in the fine structure of trial-by-trial variability in population activity, which can be readily estimated in existing datasets.

To flesh out the specific predictions of our optimal control hypothesis for variability, we introduced stochasticity in the input to every neuron, in the form of independent Gaussian noise processes (Section S5.6). These noisy inputs propagate through the recurrent dynamics, and cause variability in the firing rate of each neuron across time and trials (Hennequin et al., 2018). We found that firing rate variability drops at preparation onset (Figure 6A, lightest purple), consistent with the drop in Fano factor previously reported in monkey M1/PMd (Churchland et al., 2006b, 2010b) — and reproduced here in our own analysis of the monkey data (Figure 6B). Although variability suppression in the model was relatively mild on average, we found that the effect grew as we added "phantom muscles" to the model. Specifically, we noted that the solution to the optimal preparation problem (Equation 6) does not depend on the details of the desired "muscle" (i.e. torques) activities, but only on the way they are read out from the cortical population (the readout matrix $\mathbf{C}$ in Equation 1). We thus artificially increased the dimensionality of the network readout, and observed increasingly strong variability suppression (Figure 6A, shades of purple).

A formal mathematical analysis of variability in the model explains this quenching effect (Section S5.6), and can be summarized as follows. For a linearized model under white noise input (a useful limit), the momentary deviation from trial-average response (residual) at any time in any trial is the superposition of the network's responses to a series of past consecutive noisy impulses ("impulse responses"). At preparation onset, the *effective* recurrent connectivity of the cortical network is altered by the sudden addition of the thalamocortical loop pathway (Nakajima and Halassa, 2017). This modifies the impulse response of the system, such that activity is pulled towards the desired preparatory end-state more strongly than during spontaneous dynamics. Thus, in each of the superimposing responses contributing to the momentary activity fluctuations, residuals decay faster overall than they would normally outside the preparation epoch. This results in overall variability suppression. Notably, with a higher-dimensional readout (more muscles), a greater number of state space directions are constrained by the optimal control strategy, i.e. are caused to decay faster than usual by the thalamocortical loop. This explains why variability suppression grows with the number of muscles (Figure 6A, darker purples).

Importantly, because prospective errors decay faster along motor-potent preparatory directions (Figure 4), the model should exhibit selective quenching of trial-by-trial variability in these directions. To verify this, we decomposed firing rate variability in the model into contributions from potent and non-potent state space directions, using the same orthogonal basis as used previously in Figure 4. We confirmed that during the preparation epoch, variability is more strongly quenched in potent directions than in null directions (Figure 6C).

We then sought to test this model prediction by analyzing the structure of variability in the monkey data. We reasoned that, although it is difficult to determine the motor potency of a given state space direction in the recorded population, one subspace is almost certain potent: the subspace spanned by all pairwise differ-

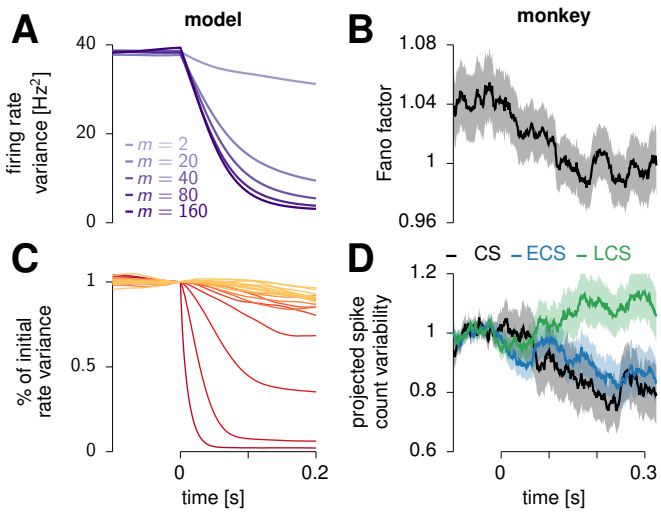

Figure 6: **Selective quenching of across-trial spike count variability in model and monkey.** **(A)** Across-trial firing rate variance in the model, averaged across cells and reach conditions. Different lines denote different numbers of "muscles" involved in the movement (see text). **(B)** Fano factor averaged over all neurons and reach conditions in the monkey data (shaded area: ± s.e.m.). **(C)** Firing rate variance decomposed in the twenty 10-dimensional subspaces of Figure 4 (same color code), in our base model with two "muscles" ($m = 2$). Traces are normalized to the variance obtained at preparation onset ($t = 0$). **(D)** Spike count co-variability in the monkey data, projected onto the coding subspace (CS, black), the early-change subspace (ECS, blue), and the late-change subspace (LCS, green). Values are normalized by the average projected variance in the 100 ms window preceding preparation onset ($t = 0$). See text for details. Shaded areas denote ± s.e.m. (boostrap). In (B) and (D), 123 neurons (single and multi-units) were analyzed across 8 straight reaches. Spikes were aligned onto target onset, and only trials with a delay period longer than 400 ms were analyzed (an average of ≈ 26 per condition).

ences between the eight reach-specific preparatory endpoints. Indeed, the contribution of this subspace to late preparatory activity is what determines the upcoming reach. Therefore we call this subspace the "coding subspace", in analogy with Li et al.'s "coding direction" (c.f. above). Conversely, since accurate reaches can be made as early as 100 ms after movement instruction, fluctuations in trial-averaged firing rates that occur late in the delay period are likely inconsequential. We thus extracted the state space directions corresponding to these late changes, to obtain a subspace we call the "late-change subspace" (LCS; Section S5.6). Finally, we also reasoned that another independent estimation of the potent subspace could be obtained by considering *early* changes in mean preparatory activity. Indeed, under optimal control, these changes should occur predominantly along potent dimensions (Figure 4). We thus collected the within-condition activity differences between early and mid-preparation, to obtain an "early-change subspace" (ECS), presumably potent. To avoid double-counting the CS (likely the most reliable estimate of potent directions), we further constrained the ECS and the LCS to be orthogonal to the CS.

Having identified putative potent and null directions using trial-averaged responses only, we next partioned across-trial variability in these different subspaces. Specifically, for each time t, neuron i, reach condition m and trial k, we counted the number of spikes $c_{imk}(t)$ that fell in a 150 ms-long window centered on t. We then constructed normalized residuals,

$$\tilde{c}_{imk}(t) = \frac{c_{imk}(t)}{\sqrt{\mu_{im}(t)}} - \sqrt{\mu_{im}(t)} \qquad (7)$$

where μ_{im} is the average of c_{imk} across trials (note that in particular, the mean of $\tilde{c}_{imk}$ over trials is zero). This construction recovers the standard Fano factor as the across-trial mean of $\tilde{c}^2_{imk}$, but also lets us generalize the Fano factor to measure spike count co-variability more specifically along any state space direction $\mathbf{d}$, as

$$\mathcal{V}(\mathbf{d}, t) = \left\langle (\mathbf{d}^T \tilde{\mathbf{c}}_{\bullet mk}(t))^2 \right\rangle_{mk}. \qquad (8)$$

Here, $\tilde{\mathbf{c}}_{\bullet mk} = (\tilde{c}_{1mk}, \tilde{c}_{2mk}, \ldots, \tilde{c}_{Nmk})^T$ denotes the population vector of normalized spike count residuals in condition m and trial k, and $\langle\cdot\rangle_{mk}$ denotes an average over all conditions and trials. To compute variability in each of the three subspaces, we averaged $\mathcal{V}(\mathbf{d}, t)$ over a set of orthogonal directions $\mathbf{d}$ defining the subpace, weighted by their contributions to the subspace (Section S5.6). We found that, as predicted by our model, the suppression of shared variability in M1/PMd activity is selective. Spiking variability is most strongly reduced in the two putative potent subspaces (Figure 6D, "CS" and "ECS"), but not in the presumed null subspace (Figure 6D, "LCS"). Together, these results support our optimal preparatory control hypothesis.

Reorganization between preparatory and movement responses

Finally, our model also accounts for a prominent feature of monkey M1/PMd responses during motor preparation and execution: across time and reach conditions, activity spans orthogonal subspaces during the two epochs. To show this, we followed Elsayed et al. and performed principal components analysis (PCA) on model and monkey trial-averaged activity during the two epochs separately (Section S5.4; Figure 7A). We then examined the fraction of variance explained by both sets of principal components (prep-PCs and move-PCs) during each epoch. Consistent with the monkey data, prep-PCs accounted for most of the activity vari-

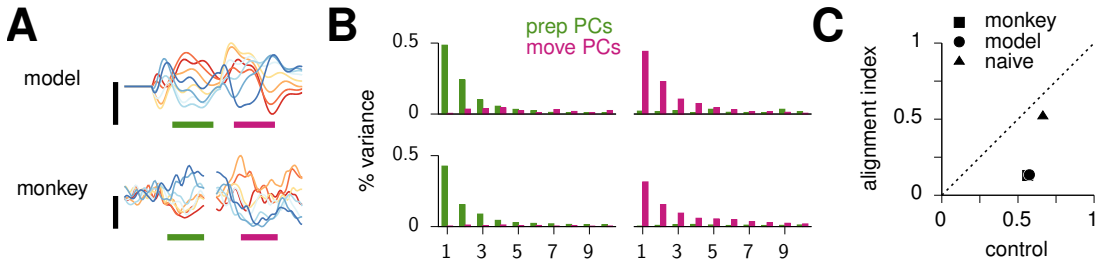

Figure 7: **Reorganization between preparatory and movement activity in model and monkey. (A)** Example single-neuron PSTHs in model (top) and monkey M1/PMd (bottom), for each of the eight movement conditions. For the model, these were already shown in Figure 2B. The monkey performed similar, though not identical, straight reaches (Figure S5). **(B)** Fraction of variance explained during movement preparation (left) and execution (right) by principal components calculated from preparatory (green) and movement-related (magenta) trial-averaged activity (Section S5.4). Only the first 10 components are shown for each. Variance is across reach conditions and time in 300 ms prep. and move. windows indicated by green and magenta bars in (A). **(C)** Alignment index (calculated as in Elsayed et al., 2016) for the monkey data (square), our optimal circuit model of Figure 3 (circle, overlapping with the square) and the naive control strategy based on static control inputs (triangle). Control values refer to the average alignment index between random subspaces drawn as in Elsayed et al., 2016 (see text and Section S5.4).

ance during preparation (by construction; Figure 7B, left), but accounted for little variance during movement (Figure 7B, right). Similarly, move-PCs captured little of the preparatory-epoch activity variance. We further quantified this (lack of) overlap using Elsayed et al.'s "alignment index", defined as the amount of preparatory-epoch activity variance captured by the top K move-PCs. This is further normalized by the maximum amount of variance that any K-dimensional subspace can capture. Here, K is chosen such that the top K prep-PCs capture 85% of activity variance during the preparatory-epoch (model, $K = 4$; monkey, $K = 12$). Both model and monkey data had a low alignment index (≈ 0.1), much below that expected if the prep. and move. subspaces had been chosen randomly within a relevant constrained subspace ("random" control in Elsayed et al., 2016; Figure 7C).

In the model, orthogonality between prep. and move. subspaces arises primarily due to the "nonnormality" of the connectivity matrix $\mathbf{W}$ (Hennequin et al., 2014; Trefethen and Embree, 2005). As shown previously in Hennequin et al. (2014), the dynamics of our generator network amplify a select set of initial conditions, from which population activity quickly grows, rotates away, oscillates in new orthogonal dimensions, and eventually decays. Since multiphasic torque patterns must be assembled from the network's activity during movement, our model calibration procedure (cf. above) tends to discover movement-specific initial conditions that belong precisely to this special set. This explains why movement-related activity is orthogonal to *late* preparatory activity in our model. Under the optimal control law, activity converges fast to the relevant subspace, hence the low alignment index overall. Under the naive strategy, convergence is much slower (Figure 2A, dashed line), indeed yielding a comparatively large alignment index (Figure 7C, triangle).

Discussion

Neural population activity in cortex can be accurately described as arising from low-dimensional dynamics (Churchland et al., 2012; Mante et al., 2013; Carnevale et al., 2015; Seely et al., 2016; Barak, 2017; Cunningham and Byron, 2014; Michaels et al., 2016). These dynamics unfold from a specific initial condition on each trial, and indeed these "preparatory states" predict the subsequent evolution of both neural activity and behaviour in single trials of the task (Churchland et al., 2010b; Pandarinath et al., 2018; Remington et al., 2018; Sohn et al., 2019). In addition, motor learning may rely on these preparatory states partitioning the space of subsequent movements (Sheahan et al., 2016).

How are appropriate initial conditions reached in the first place? Here, we have formalized movement preparation as an optimal control problem, showing how to translate anticipated motor costs phrased in terms of muscle kinematics into costs on neural activity in M1. Optimal preparation minimizes these costs, and the solution is feedback control: the cortical network must provide corrective feedback to itself, based on the prospective motor error associated with its current state. In other words, optimal preparation may rely on an implicit forward model (Wolpert et al., 1995; Desmurget and Grafton, 2000; Scott, 2012), whereby the future motor consequences of *preparatory activity* (not motor commands, as in classical theories) are predicted and fed back for online correction of the cortical trajectory.

Thalamic control of cortical dynamics

The mathematical structure of the optimal control solution suggested a circuit model based on cortico-cortical feedback. We have proposed that optimal feed-

back can be implemented as a cortico-thalamo-cortical loop, switched ON during movement preparation and OFF again at movement onset. The ON-switch occurs through fast disinhibition of those thalamic neurons that are part of the loop. Our model thus predicts a large degree of specificity in the synaptic interactions between cortex and thalamus (Halassa and Sherman, 2019), as well as a causal involvement of the thalamus in movement preparation (Guo et al., 2017; Sauerbrei et al., 2019). Furthermore, the dynamical entrainment of thalamus with cortex predicts tuning of thalamic neurons to task variables, consistent with a growing body of work showing specificity in thalamic responses (Nakajima and Halassa, 2017; Guo et al., 2017; Rikhye et al., 2018). For example, we predict that neurons in the motor thalamus should be tuned to movement properties, for much the same reasons that cortical neurons are (Todorov, 2000; Lillicrap and Scott, 2013; Omrani et al., 2017).

Thalamic control of cortical dynamics offers a particularly attractive way of performing nonlinear computations (Sussillo and Abbott, 2009; Logiaco et al., 2019). Although both preparatory and movement-related dynamics are approximately linear in our model, the transition from one to the other (orchestrated by the basal ganglia) is highly nonlinear. Indeed, our model can be thought of as a switching linear dynamical system (Linderman et al., 2017). Moreover, gated thalamocortical loops are a special example of achieving nonlinear effects through gain modulation. Here, it is the thalamic population only that is subjected to abrupt and binary gain modulation, but changes in gain could also affect cortical neurons. This was proposed recently as a way of expanding the dynamical repertoire of a cortical network (Stroud et al., 2018).

Switch-like nonlinearities may have relevance beyond movement preparation, e.g. for movement execution. In our model, different movement patterns are produced by different initial conditions seeding the same generator dynamics. However, we could equally well have generated each reach using a different movement-epoch thalamocortical loop. This would account for the recent demonstration that thalamus drives cortex during the production of skilled movements in mice (Sauerbrei et al., 2019). Logiaco et al. have recently explored this possibility, showing that gated thalamocortical loops provide an ideal substrate for flexible sequencing of multiple movements. In their model, each movement is achieved by its own loop (involving a shared cortical network), and the basal ganglia orchestrate a chain of thalamic disinhibitory events, each spatially targetted to activate those neurons that are responsible for the next loop in the sequence (Logiaco et al., 2019). Interestingly, their cortical network must still be properly initialized prior to each movement chunk, as it must in our model. For this, they proposed a generic preparatory loop similar to the one we have studied here. However, theirs does not take into account the degeneracies in preparatory states induced by prospective motor costs, which ours exploits. In sum, our two models address complementary facets of motor control (preparation and sequenching), and could be combined into a single model.

Sloppy preparation for accurate movements

Two elements might mitigate the need for exquisite control of cortical preparatory states. First, ongoing movements can be corrected rapidly based on sensory feedback, presumably enabling compensation for a "bad start" (Scott et al., 2015). Second, as we found, the mapping from initial condition to movement may be many-to-one, and optimal control dictates that only those components of the initial condition that matter for the subsequent movement ought to be controlled during preparation. Indeed, this feature of our model readily explains two distinctive features of preparatory activity in reaching monkeys: (i) that pre-movement activity on zero-delay trials needs not reach the state achieved for long movement delays (Ames et al., 2014), and (ii) that nevertheless movement is systematically preceded by activity in the same preparatory subspace irrespective of whether the reach is self-initiated, artifically delayed, or reactive and fast (Lara et al., 2018). In our model, preparatory activity converges rapidly in the subspace that matters, such that irrespective of the delay (above 50 ms), preparatory activity is always found to have some component in this common subspace as in Lara et al. (2018). Moreover, exactly which of the many acceptable initial conditions is reached by the end of the delay depends on the delay duration. Thus, our model predicts that different late-preparation states will be achieved for short and long delays, consistent with the results of Ames et al. (2014). Moreover, the state our model achieves at the end of a preparation epoch also depends on activity prior to preparation onset; therefore, it also depends on whether preparation started from scratch, or was initiated by a change in target that interrupted a previous preparatory process half way, as observed by Ames et al..

Selective suppression of cortical variability during preparation, and putative mechanisms

Degeneracies are ubiquitous in biological control (Edelman and Gally, 2001; Todorov and Jordan, 2002; O'Leary et al., 2014; Hennig et al., 2018). In motor control, movement trajectories are highly degenerate w.r.t. the goal; e.g. there are very many ways to reach for a cup of coffee. Todorov and Jordan viewed motor variability through the lens of stochastic optimal control, arguing that only those motor fluctuations that interfere with task goals should be corrected, while other aspects of the movement can vary freely. Here, we have shown that principles of optimal control can also explain the structure of *neural* variability in M1/PMd during prepa-

ration, due to an analogous degeneracy in preparatory activity. In particular, optimal elimination of prospective motor errors during preparation predicts suppression of trial-by-trial variability, which occurs in monkey M1/PMd (Nawrot et al., 2001; Churchland et al., 2006b; Rickert et al., 2009; Churchland et al., 2010a) and mouse ALM (Inagaki et al., 2019). More importantly, variability suppression should be coordinated across the population in such a way that fluctuations are suppressed faster along directions that matter more ("potent directions"). Here, we have successfully verified this novel prediction in monkey data.

In our model, selective quenching of variability arises from much the same mechanism as we have recently proposed for primary visual cortex (V1; Hennequin et al., 2018). There, we argued that external stimuli change the operating point of the network dynamics, and since these dynamics are nonlinear (Ahmadian et al., 2013; Rubin et al., 2015), the result is a modification of effective connectivity. We went on to show that this resulted in greater inhibitory dominance (as in Stringer et al., 2016), and therefore to quenching of fluctuations. Here, variability suppression is also due to increased effective negative feedback (Section S4.10), but this occurs due to the sudden addition of a thalamocortical pathway at preparation onset — as opposed to an exogenous stimulus.

Outlook

To explore the ramifications of optimal control for motor preparation, we had to commit to a concrete model of movement-generating dynamics in M1 (Hennequin et al., 2014). While our model captures several salient features of movement-related activity (Section S5), as well as key qualitative aspects of preparatory dynamics, more quantitative predictions would require detailed, data-driven modelling of M1 (Pandarinath et al., 2018). Such future extensions of our control-theoretic framework could help elucidate the role of the numerous brain areas that collectively control movement (Svoboda and Li, 2018), and make sense of their hierarchical organization in nested loops.

Acknowledgments

This work was supported by a Seed Award from the Wellcome Trust (202111/Z/16/Z), a Trinity-Henry Barlow scholarship (T-C.K.), and a scholarship from the Ministry of Education, ROC Taiwan (T-C.K.). We are grateful to Mark Churchland, Matthew Kaufman and Krishna Shenoy for sharing their data and for discussions; to Omri Barak for suggesting an estimation of potent and null subspaces from data; to Karel Svoboda and his lab for hospitality and discussions.

References

Afshar, A., Santhanam, G., Byron, M. Y., Ryu, S. I., Sahani, M., and Shenoy, K. V. (2011). Single-trial neural correlates of arm movement preparation. *Neuron*, 71:555–564.

Ahmadian, Y., Rubin, D. B., and Miller, K. D. (2013). Analysis of the stabilized supralinear network. *Neural Comput*, 25:1994–2037.

Ames, K. C., Ryu, S. I., and Shenoy, K. V. (2014). Neural dynamics of reaching following incorrect or absent motor preparation. *Neuron*, 81:438–451.

Barak, O. (2017). Recurrent neural networks as versatile tools of neuroscience research. *Curr Opin Neurobiol*, 46:1–6.

Carnevale, F., de Lafuente, V., Romo, R., Barak, O., and Parga, N. (2015). Dynamic control of response criterion in premotor cortex during perceptual detection under temporal uncertainty. *Neuron*, 86:1067–1077.

Churchland, M. M., Afshar, A., and Shenoy, K. V. (2006a). A central source of movement variability. *Neuron*, 52:1085–1096.

Churchland, M. M., Byron, M. Y., 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. (2010a). Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nat Neurosci*, 13:369–378.

Churchland, M. M., Byron, M. Y., Ryu, S. I., Santhanam, G., and Shenoy, K. V. (2006b). Neural variability in premotor cortex provides a signature of motor preparation. *J Neurosci*, 26:3697–3712.

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(7405):51.

Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Ryu, S. I., and Shenoy, K. V. (2010b). Cortical preparatory activity: representation of movement or first cog in a dynamical machine? *Neuron*, 68:387–400.

Cui, G., Jun, S. B., Jin, X., Pham, M. D., Vogel, S. S., Lovinger, D. M., and Costa, R. M. (2013). Concurrent activation of striatal direct and indirect pathways during action initiation. *Nature*, 494:238–242.

Cunningham, J. P. and Byron, M. Y. (2014). Dimensionality reduction for large-scale neural recordings. *Nat Neurosci*, 17:1500.

Dayan, P. and Abbott, L. F. (2001). *Theoretical neuroscience*. Cambridge, MA: MIT Press.

Desmurget, M. and Grafton, S. (2000). Forward modeling allows feedback control for fast reaching movements. *Trends Cogn Sci*, 4(11):423–431.

Edelman, G. M. and Gally, J. A. (2001). Degeneracy and complexity in biological systems. *Proc Natl Ac Sci USA*, 98(24):13763–13768.

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.

Evarts, E. V. (1968). Relation of pyramidal tract activity to force exerted during voluntary movement. *J Neurophysiol*, 31:14–27.

Gao, Z., Davis, C., Thomas, A. M., Economo, M. N., Abrego, A. M., Svoboda, K., De Zeeuw, C. I., and Li, N. (2018). A cortico-cerebellar loop for motor planning. *Nature*, 563:113–116.

Guo, Z. V., Inagaki, H. K., Daie, K., Druckmann, S., Gerfen, C. R., and Svoboda, K. (2017). Maintenance of persistent activity in a frontal thalamocortical loop. *Nature*, 545:181–186.

Halassa, M. M. and Acsády, L. (2016). Thalamic inhibition: diverse sources, diverse scales. *Trends Neurosci*, 39:680–693.

Halassa, M. M. and Sherman, S. M. (2019). Thalamocortical circuit motifs: a general framework. *Neuron*, 103:762–770.

Hennequin, G., Ahmadian, Y., Rubin, D. B., Lengyel, M., and Miller, K. D. (2018). The dynamical regime of sensory cortex: stable dynamics around a single stimulus-tuned attractor account for patterns of noise variability. *Neuron*, 98:846–860.

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.

Hennig, J. A., Golub, M. D., Lund, P. J., Sadtler, P. T., Oby, E. R., Quick, K. M., Ryu, S. I., Tyler-Kabara, E. C., Batista, A. P., Byron, M. Y., et al. (2018). Constraints on neural redundancy. *eLife*, 7:e36774.

Inagaki, H. K., Fontolan, L., Romani, S., and Svoboda, K. (2019). Discrete attractor dynamics underlies persistent activity in the frontal cortex. *Nature*, 566:212–217.

Jin, X. and Costa, R. M. (2010). Start/stop signals emerge in nigrostriatal circuits during sequence learning. *Nature*, 466:457–462.

Kao, T.-C. and Hennequin, G. (2019). Neuroscience out of control: control-theoretic perspectives on neural circuit dynamics. *Curr Opin Neurobiol*, 58:122–129.

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(4):0085–16.2016.

Lara, A. H., Elsayed, G. F., Zimnik, A. J., Cunningham, J. P., and Churchland, M. M. (2018). Conservation of preparatory neural events in monkey motor cortex regardless of how movement is initiated. *eLife*, 7:e31826.

Li, N., Daie, K., Svoboda, K., and Druckmann, S. (2016). Robust neuronal dynamics in premotor cortex during motor planning. *Nature*, 532(7600):459–464.

Lillicrap, T. P. and Scott, S. H. (2013). Preference distributions of primary motor cortex neurons reflect control solutions optimized for limb biomechanics. *Neuron*, 77(1):168–179.

Linderman, S., Johnson, M., Miller, A., Adams, R., Blei, D., and Paninski, L. (2017). Bayesian learning and inference in recurrent switching linear dynamical systems. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54, pages 914–922.

Logiaco, L., Abbott, L. F., and Escola, S. (2019). A model of flexible motor sequencing through thalamic control of cortical dynamics. *bioRxiv*.

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.

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.

Nakajima, M. and Halassa, M. M. (2017). Thalamic control of functional cortical connectivity. *Curr Opin Neurobiol*, 44:127–131.

Nawrot, M., Rodriguez, V., Heck, D., Riehle, A., Aertsen, A., and Rotter, S. (2001). Trial-by-trial variability of spike trains in vivo and in vitro. In *Soc Neurosci Abstr*, volume 27.

O'Leary, T., Williams, A. H., Franci, A., and Marder, E. (2014). Cell types, network homeostasis, and pathological compensation from a biologically plausible ion channel expression model. *Neuron*, 82:809–821.

Omrani, M., Kaufman, M. T., Hatsopoulos, N. G., and Cheney, P. D. (2017). Perspectives on classical controversies about the motor cortex. *J Neurophysiol*, 118:1828–1848.

Ozeki, H., Finn, I. M., Schaffer, E. S., Miller, K. D., and Ferster, D. (2009). Inhibitory stabilization of the cortical network underlies visual surround suppression. *Neuron*, 62(4):578–592.

O'Shea, D. J., Kalanithi, P., Ferenczi, E. A., Hsueh, B., Chandrasekaran, C., Goo, W., Diester, I., Ramakrishnan, C., Kaufman, M. T., Ryu, S. I., et al. (2018). Development of an optogenetic toolkit for neural circuit dissection in squirrel monkeys. *Sci Rep*, 8:1–20.

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., et al. (2018). Inferring single-trial neural population dynamics using sequential auto-encoders. *Nat Methods*, 15:805–815.

Remington, E. D., Egger, S. W., Narain, D., Wang, J., and Jazayeri, M. (2018). A dynamical systems perspective on flexible motor timing. *Trends Cogn Sci*, 22:938–952.

Rickert, J., Riehle, A., Aertsen, A., Rotter, S., and Nawrot, M. P. (2009). Dynamic encoding of movement direction in motor cortical neurons. *J Neurosci*, 29:13870–13882.

Rikhye, R. V., Gilra, A., and Halassa, M. M. (2018). Thalamic regulation of switching between cortical representations enables cognitive flexibility. *Nat Neurosci*, 21:1753–1763.

Rubin, D. B., Van Hooser, S. D., and Miller, K. D. (2015). The stabilized supralinear network: a unifying circuit motif underlying multi-input integration in sensory cortex. *Neuron*, 85:402–417.

Sanzeni, A., Akitake, B., Goldbach, H. C., Leedy, C. E., Brunel, N., and Histed, M. H. (2019). Inhibition stabilization is a widespread property of cortical networks. *bioRxiv*, page 656710.

Sauerbrei, B. A., Guo, J.-Z., Cohen, J. D., Mischiati, M., Guo, W., Kabra, M., Verma, N., Mensh, B., Branson, K., and Hantman, A. W. (2019). Cortical pattern generation during dexterous movement is input-driven. *Nature*, pages 1–6.

Scott, S. H. (2012). The computational and neural basis of voluntary motor control and planning. *Trends Cogn Sci*, 16:541–549.

Scott, S. H., Cluff, T., Lowrey, C. R., and Takei, T. (2015). Feedback control during voluntary motor actions. *Curr Opin Neurobiol*, 33:85–94.

Seely, J. S., Kaufman, M. T., Ryu, S. I., Shenoy, K. V., Cunningham, J. P., and Churchland, M. M. (2016). Tensor analysis reveals distinct population structure that parallels the different computational roles of areas m1 and v1. *PLoS Comput Biol*, 12.

Sheahan, H. R., Franklin, D. W., and Wolpert, D. M. (2016). Motor planning, not execution, separates motor memories. *Neuron*, 92:773–779.

Shenoy, K. V., Sahani, M., and Churchland, M. M. (2013). Cortical control of arm movements: a dynamical systems perspective. *Ann. Rev. Neurosci.*, 36:337–359.

Skogestad, S. and Postlethwaite, I. (2007). *Multivariable feedback control: analysis and design*, volume 2. Wiley New York.

Sohn, H., Narain, D., Meirhaeghe, N., and Jazayeri, M. (2019). Bayesian computation through cortical latent dynamics. *Neuron*, 103:934–947.

Stringer, C., Pachitariu, M., Steinmetz, N. A., Okun, M., Bartho, P., Harris, K. D., Sahani, M., and Lesica, N. A. (2016). Inhibitory control of correlated intrinsic variability in cortical networks. *eLife*, 5:e19695.

Stroud, J. P., Porter, M. A., Hennequin, G., and Vogels, T. P. (2018). Motor primitives in space and time via targeted gain modulation in cortical networks. *Nat Neurosci*, 21(12):1774–1783.

Sussillo, D. and Abbott, L. F. (2009). Generating coherent patterns of activity from chaotic neural networks. *Neuron*, 63:544–557.

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.

Svoboda, K. and Li, N. (2018). Neural mechanisms of movement planning: motor cortex and beyond. *Curr Opin Neurobiol*, 49:33–41.

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(11):1226–1235.

Trefethen, L. N. and Embree, M. (2005). *Spectra and pseudospectra: the behavior of nonnormal matrices and operators*. Princeton University Press.

Tsodyks, M. V., Skaggs, W. E., Sejnowski, T. J., and McNaughton, B. L. (1997). Paradoxical effects of external modulation of inhibitory interneurons. *J Neurosci*, 17(11):4382–4388.

Wolpert, D. M., Ghahramani, Z., and Jordan, M. I. (1995). An internal model for sensorimotor integration. *Science*, 269:1880–1882.