

When and why does motor preparation arise in recurrent neural network models of motor control?

Reviewed Preprint

v3 • August 13, 2024

Revised by authors

Reviewed Preprint

v2 • May 20, 2024

Reviewed Preprint

v1 • August 10, 2023

Marine Schimel , Ta-Chu Kao, Guillaume Hennequin

Computational and Biological Learning Lab, Department of Engineering, University of Cambridge, Cambridge, U.K. • Meta Reality Labs

 https://en.wikipedia.org/wiki/Open_access
 Copyright information

Abstract

Summary

During delayed ballistic reaches, motor areas consistently display movement-specific activity patterns prior to movement onset. It is unclear why these patterns arise: while they have been proposed to seed an initial neural state from which the movement unfolds, recent experiments have uncovered the presence and necessity of ongoing inputs during movement, which may lessen the need for careful initialization. Here, we modelled the motor cortex as an input-driven dynamical system, and we asked what the optimal way to control this system to perform fast delayed reaches is. We find that delay-period inputs consistently arise in an optimally controlled model of M1. By studying a variety of network architectures, we could dissect and predict the situations in which it is beneficial for a network to prepare. Finally, we show that optimal input-driven control of neural dynamics gives rise to multiple phases of preparation during reach sequences, providing a novel explanation for experimentally observed features of monkey M1 activity in double reaching.

eLife assessment

This **important** study provides a new perspective on why preparatory activity occurs before the onset of movement. The authors report that when there is a cost on the inputs, the optimal inputs should start before the desired network output for a wide variety of recurrent networks. The authors present **compelling** evidence by combining mathematically tractable analyses in linear networks and numerical simulation in nonlinear networks.

<https://doi.org/10.7554/eLife.89131.3.sa4>

During the production of ballistic movements, the motor cortex is thought to operate as a dynamical system whose state trajectories trace out the appropriate motor commands for downstream effectors (Shenoy et al., 2013 ; Miri et al., 2017 ; Russo et al., 2018 ). The extent to which these cortical dynamics are controlled by exogenous inputs before and/or during movement is the subject of ongoing study.

On the one hand, several experimental and modelling studies point to a potential role for exogenous inputs in motor preparation. First, cortical state trajectories are empirically well described by a low-dimensional dynamical system evolving near-autonomously during movement (Churchland et al., 2012 [↗](#); Pandarinath et al., 2018 [↗](#); Schimel et al., 2022 [↗](#)), such that there is a priori no reason to suspect that inputs are required for motor production. Rather, inputs would be required during preparation to bring the state of the cortical network into a suitable initial condition. This input-driven seeding process is corroborated by observations of movement-specific primary motor cortex (M1) activity arising well before movement initiation (Lara et al., 2018 [↗](#); Kaufman et al., 2014 [↗](#); Churchland et al., 2012 [↗](#); Meirhaeghe et al., 2023 [↗](#); **Figure 1A** [↗](#)), and associated models demonstrate the critical role of preparatory inputs therein (Sussillo et al., 2015 [↗](#); Hennequin et al., 2014 [↗](#); Kao et al., 2021 [↗](#)).

On the other hand, recent studies in mice have shown that the motor cortex receives critical pattern-generating input from the thalamus during movement production (Sauerbrei et al., 2020 [↗](#)), and recurrent neural network (RNN)-based modelling of the motor feed-back loop involved in reaching movements suggests that sensory feedback may also contribute significantly to the observed dynamics of M1 (Kalidindi et al., 2021 [↗](#)). Moreover, most published network models of delayed reaches are able to perform the task just as well with-out preparatory inputs, i.e. with external inputs forcefully confined to the movement epoch – an illustrative example is shown in **Figure 1B** [↗](#). Thus, the relative contributions of preparatory vs. movement-epoch inputs to the dynamics implemented by M1 (potentially as part of a broader set of areas) remain unclear.

In addition to the specific form that inputs to cortical dynamics might take, one may ask more broadly about the computational role of motor preparation. Motor preparation is known to benefit behaviour (e.g. by short-ening reaction times and enabling more accurate execution Riehle and Requin, 1989 [↗](#); Churchland and Shenoy, 2007 [↗](#); Michaels et al., 2015 [↗](#)) and may facilitate motor learning (Sheahan et al., 2016 [↗](#); Sun et al., 2022 [↗](#)). However, from the perspective of cortical dynamics, preparation also introduces additional constraints. Specifically, the high density of M1 neurons projecting directly to the spinal cord (Dum and Strick, 1991 [↗](#)) suggests that motor cortical outputs control lower-level effectors with little intermediate processing. For preparatory processes to avoid triggering premature movement, any pre-movement activity in the motor and dorsal premotor (PMd) cortices must therefore engage the pyramidal tract neurons in a way that ensures their activity patterns will not lead to any movement.

While this can be achieved by constraining neural activity to evolve in a nullspace of the motor output (Kaufman et al., 2014 [↗](#)), the question nevertheless arises: what advantage is there to having neural dynamics begin earlier in a constrained manner, rather than unfold freely just in time for movement production?

Here we sought a normative explanation for motor preparation at the level of motor cortex dynamics: we asked whether preparation arises in recurrent neural networks (RNNs) performing delayed reaching tasks, and what factors lead to more or less preparation. Such an explanation could not be obtained from previous network models of delayed reaches, as they typically assume from the get-go that the cortical network receives preparatory inputs during a fixed time window preceding the go cue (Sussillo et al., 2015 [↗](#); Kao et al., 2021 [↗](#)). In this case, pre-movement activity is by design a critical determinant of the subsequent behaviour (Sussillo et al., 2015 [↗](#); Kao et al., 2021 [↗](#); Zimnik and Churchland, 2021 [↗](#)). In this work, we removed this modelling assumption and studied models in which the correct behaviour could in principle be obtained without explicit motor preparation.

To study the role of motor preparation, and that of exogenous inputs in this process, we followed an optimal control approach (Harris and Wolpert, 1998 [↗](#); Todorov and Jordan, 2002 [↗](#); Yeo et al., 2016 [↗](#)). We considered the dynamics of an RNN model of M1 coupled to a model arm (Todorov and Li, 2003 [↗](#)), and used a standard control cost functional to quantify and optimize performance




Figure 1

Control is possible under different strategies.

(A) Trial-averaged firing rate of two representative monkey M1 neurons, across 8 different movements, separately aligned to target onset (left) and movement onset (right). Neural activity starts separating across movements well before the animal starts moving. **(B)** Top: a RNN model of M1 dynamics receives external inputs $u(t)$ from a higher-level controller, and outputs control signals for a biophysical two-jointed arm model. Inputs are optimized for the correct production of eight center-out reaches to targets regularly positioned around a circle. Bottom: firing rate of a representative neuron in the RNN model for each reach, under two extreme control strategies. In the first strategy (left, solid lines), the external inputs $u(t)$ are optimized whilst being temporally confined to the preparatory period. In the second strategy (right, dashed lines), they are optimized whilst confined to the movement period. Although slight differences in hand kinematics can be seen (compare corresponding solid and dashed hand trajectories), both control policies lead to successful reaches. These introductory simulations are shown for illustration purposes; the particular choice of network connectivity and the way the control inputs were found are described in the Results section.

in a delayed-reaching task. We used the iLQR algorithm (Li and Todorov, 2004) to find the spatiotemporal patterns of network inputs that minimize this cost functional, for any given network connectivity. Critically, these inputs could arise both before and during movement; thus, our framework allowed for principled selection amongst a continuum of motor strategies, going from purely autonomous motor generation following preparation, to purely input-driven unprepared dynamics.

We considered an inhibition-stabilized network – which was shown previously to capture prominent aspects of monkey M1 activity (Hennequin et al., 2014; Kao et al., 2021) – and found that optimal control of the model requires preparation, with optimal inputs arising well before movement begins. To understand what features of network connectivity lead to optimal preparatory control strategies, we first turned to low-dimensional models, which could be more easily dissected. We then generalized insights from those systems back to high-dimensional networks using tools from control theory, and found that preparation can be largely explained by two quantities summarizing the dynamical response properties of the network.

Finally, we studied the optimal control of movement *sequences*. Consistent with recent experimental findings (Zimnik and Churchland, 2021), we observed that optimal control of compound reaches leads to input-driven preparatory activity in a dedicated activity subspace prior to each movement.

Overall, our results show that preparatory neural activity patterns arise from optimal control of reaching movements at the level of motor cortical circuits, thus providing a possible explanation for a number of observed experimental findings.

Model

A model of cortical dynamics for reaching movements

We considered a simple reaching task, in which the hand must move from a resting location to one of eight radially located targets in a 2D plane as fast as possible (Figure 1). The target had to be reached within 600 ms of a go cue that follows a delay period of varying (but known) duration. We modelled the trajectory of the hand via a two-jointed model arm (Li and Todorov, 2004; Kao et al., 2021), driven into motion by a pair of torques $\mathbf{m}(t)$ (Methods). We further assumed that these torques arise as a linear readout of the momentary firing rates $\mathbf{r}(t)$ of a population of M1 neurons,

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

where $\mathbf{C}$ was a randomly generated readout matrix, projecting the neural activity into the output space. We modelled the dynamics of $N = 200$ M1 neurons using a standard rate equation,

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

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

where the momentary population firing rate vector $\mathbf{r}(t)$ was obtained by passing a vector of internal neuronal activations $\mathbf{x}(t)$ through a rectified linear function $\phi[\cdot]$, element-wise. In Equation 2, $\mathbf{h}$ is a constant input that establishes a baseline firing rate of 5 Hz on average, with a standard deviation of 5 Hz across neurons, $\mathbf{u}(t)$ is a task-dependent control input (see below), and $\mathbf{W}$ denotes the matrix of recurrent connection weights. Throughout most of this work, we considered inhibition-stabilized M1 dynamics (Hennequin et al., 2014; Methods), which have previously been shown to produce activity resembling that of M1 during reaching (Kao et al.,

2021). Thus, our model can be viewed as a two-level controller, with the arm being controlled by M1, and M1 being controlled by external inputs. Note that each instantiation of our model corresponds to a set of $\mathbf{W}$, $\mathbf{C}$, and $\mathbf{h}$, none of which are specifically optimized for the task.

To prepare or not to prepare?

Previous experimental (Churchland et al., 2012; Shenoy et al., 2013) and modelling (Hennequin et al., 2014; Sussillo et al., 2015; Pandarinath et al., 2018) work suggests that fast ballistic movements rely on strong dynamics, which are observed in M1 and well modelled as near-autonomous (although the underlying dynamical system may not be anatomically confined to M1, as we discuss later). Network-level models of ballistic control thus rely critically on a preparation phase during which they are driven into a movement-specific state that seeds their subsequent autonomous dynamics (Kao et al., 2021; Sussillo et al., 2015). However, somewhat paradoxically, the same recurrent neural network models can also solve the task in a completely different regime, in which task-related inputs arise during movement only, with no preparatory inputs whatsoever. We illustrate this dichotomy in **Figure 1**. The same center-out reach can be produced with control inputs to M1 that arise either prior to movement only (full lines), or during movement only (dashed lines). In the latter case, no reach-specific preparatory activity is observed, making the model inconsistent with experimental findings. But what rationale is there in preparing for upcoming movements, then?

To address this question, we formulated delayed reaching as an optimal control problem, and asked what external inputs are required, and at what time, to drive the hand into the desired position with minimum control effort. Specifically, we sought inputs that were as weak as possible yet accurately drove the hand to the target within an allotted time window. We also penalized inputs that caused premature movement before the go cue.

Thus, we solved for spatio-temporal input trajectories that minimized a cost functional capturing the various task requirements. Our cost was composed of three terms: $\mathcal{J}_{\text{target}}$ penalizes deviations away from the target, with an “urgency” weight that increases quadratically with time, thus capturing the implicit incentive for animals to perform fast reaches in such experiments (which are normally conducted in sessions of fixed duration). $\mathcal{J}_{\text{null}}$ penalizes premature movement during preparation, as measured by any deviation in position, speed and acceleration of the hand. Finally, $\mathcal{J}_{\text{effort}}$ penalizes control effort in the form of input magnitude throughout the whole trial, thus promoting energy-efficient control solutions amongst a typically infinite set of possibilities (Kao et al., 2021; Sterling and Laughlin, 2015). Note that $\mathcal{J}_{\text{effort}}$ can be viewed as a standard regularization term, and must be included to ensure the control problem is well defined. The total objective thus had the following form :

$$\begin{aligned} \mathcal{J}[\mathbf{u}(t)] = & \underbrace{\int_0^T \|\boldsymbol{\theta}(t) - \boldsymbol{\theta}^*\|^2 \frac{t^2}{T^2} \frac{dt}{T}}_{\mathcal{J}_{\text{target}}} \\ & + \underbrace{\alpha_{\text{null}} \int_{-\Delta_{\text{prep}}}^0 \left(\|\boldsymbol{\theta}(t) - \boldsymbol{\theta}_0\|^2 + \|\dot{\boldsymbol{\theta}}(t)\|^2 + \|\mathbf{m}(t)\|^2 \right) \frac{dt}{T}}_{\mathcal{J}_{\text{null}}} \\ & + \underbrace{\alpha_{\text{effort}} \int_{-\Delta_{\text{prep}}}^T \|\mathbf{u}(t)\|^2 \frac{dt}{NT}}_{\mathcal{J}_{\text{effort}}}, \end{aligned} \quad (4)$$

where $\boldsymbol{\theta}$ and $\dot{\boldsymbol{\theta}}$ denote the position and velocity of the hand in angular space, Δ_{prep} was the duration of the delay period and T that of the movement period. As $\mathcal{J}_{\text{target}}$ and $\mathcal{J}_{\text{null}}$ depend on $\mathbf{u}(t)$ implicitly through Equations 1 and 2, $\mathcal{J}$ is a function of $\mathbf{u}$ only. Importantly, we allowed for

inputs within a time window beginning Δ_{prep} ms before, and ending T ms after the go cue (set at $t = 0$). Therefore, both preparation-only and movement-only input strategies (cf. **Figure 1**) could potentially arise, as well as anything in between.

Note that this control objective assumes that the delay duration (Δ_{prep}) is known ahead of time, an assumption that does not hold for many delayed-reaching tasks in monkeys where the delay is uncertain. We make this assumption for computational tractability and later discuss extensions to the uncertain case (Discussion).

Here, we solved for the optimal control inputs using the iterative linear quadratic regulator algorithm (iLQR; Li and Todorov, 2004), an efficient trajectory optimization algorithm that is well-suited for handling the nonlinear nature of both the arm's and the network's dynamics. As our primary goal was to assess the role of preparation in a normative way, we did not study the putative circuit dynamics upstream of M1 that might lead to the computation of these optimal inputs.

We balanced the various components of our cost functional by choosing α_{null} and α_{effort} to qualitatively match the behavioural requirements of a typical reach- and-hold task. Specifically, we tuned them jointly so as to ensure (i) stillness during preparation, and (ii) reach duration of approximately ~ 400 ms, with the hand staying within 0.5cm of the target for ~ 200 ms after the end of the reach. We ensured that the main qualitative features of the solution, i.e. the results presented below, were robust to the choice of hyperparameter values within the fairly large range in which the above soft-constraints are satisfied (Supplementary Material S1).

Results

Preparation arises as an optimal control strategy

Using the above control framework, we assessed whether the optimal way of performing a delayed reach involves preparation. More concretely, does the optimal control strategy of the model described in Equation 2 involve any preparatory inputs during the delay period? For any single optimally performed reach, we found that network activity began changing well before the go cue (**Figure 2B**, bottom), and that this was driven by inputs that arose early (**Figure 2B**, middle). Thus, although preparatory network activity cancels in the read-out (such that the hand remains still; **Figure 2B**, top) and therefore does not contribute directly to movement, it still forms an integral part of the optimal reach strategy. To quantify how much the optimal control strategy relied on inputs prior to movement, we defined the *preparation index* as the ratio of input magnitude during the delay period to that during the remainder of the trial:

$$\text{prep. index} = \sqrt{\frac{\int_{-\Delta_{\text{prep}}}^0 \|\mathbf{u}(t)\|^2 dt}{\int_0^T \|\mathbf{u}(t)\|^2 dt}}. \quad (5)$$

We found that the preparation index rose sharply as we increased the delay period, and eventually plateaued at ~ 1.3 for delay periods longer than 300 ms (**Figure 2C**). Similarly, the total cost of the task was highest in the absence of preparation, and decreased until it also reached a plateau at $\Delta_{\text{prep}} \sim 300$ ms (**Figure 2C**, black). This appears somewhat counterintuitive, as having a larger Δ_{prep} means that both $\mathcal{J}_{\text{effort}}$ and $\mathcal{J}_{\text{null}}$ are accumulated over a longer period. To resolve this paradox, we examined each component of the cost function. We found that the overall decrease in cost with increasing preparation time was driven by a concurrent decrease in both $\mathcal{J}_{\text{tgt}}$ and $\mathcal{J}_{\text{effort}}$. The former effect was due to the model producing faster reaches (**Figure 2C** inset; hand position for a reach with (red) and without (blue) preparation) while the latter arose from

smaller control inputs being necessary when preparation was allowed. Together, these results suggest that the presence of a delay period changes the optimal control strategy for reaching, and increases performance in the task.

The results above show that delaying the reach beyond ~ 300 ms brings little benefit; in particular, all components of the cost stabilize past that point (**Figure 2C**). We thus wondered what features the optimally controlled dynamics would display as Δ_{prep} increased beyond 300 ms. Would the network defer preparation to a last minute surge, or prepare more gently over the entire preparatory window? Would the network produce the same neural activity patterns? We found that the optimal controller made very little use of any preparation time available up to 300 ms before the go cue: with longer preparation times, external input continued to arise just a couple of hundred milliseconds before movement initiation, and single neuron firing rates remained remarkably similar (**Figure 3A**). This was also seen in PCA projections of the firing rates, which traced out similar trajectories irrespective of the delay period (**Figure 3B**). We hypothesized that this behaviour is due to the network dynamics having a certain maximum characteristic timescale, such that inputs that arrive too early end up being “forgotten” – they increase $\mathcal{J}_{\text{effort}}$ and possibly $\mathcal{J}_{\text{null}}$ without having a chance to influence $\mathcal{J}_{\text{tgt}}$. We confirmed this by varying the characteristic time constant (τ in **Equation 2**). For a fixed Δ_{prep} , we found that for larger (resp. lower) values of τ , the optimal control inputs started rising earlier (resp. later) and thus occupied more (resp. less) of the allotted preparatory period (**Figure S3**).

Understanding optimal control in simplified models

Having established that the ISN model of M1 relies on preparatory inputs to solve the delayed reaching task, we next tried to understand *why* it does so. To further unravel the interplay between the structure of the network and the optimal control strategy, i.e. what aspects of the dynamics of the network warrant preparation, we turned to simpler, two-dimensional models of cortical dynamics. These 2D models are small enough to enable detailed analysis (Supplementary Material S3), yet rich enough to capture the two dominant dynamical phenomena that arise in ISN dynamics: nonnormal amplification (Murphy and Miller, 2009; Goldman, 2009; Hennequin et al., 2012) and oscillations (Brunel, 2000; Dayan and Abbott, 2001). Specifically, networks of E and I neurons have been shown to embed two main motifs of effective connectivity which are revealed by appropriate orthogonal changes of basis: (i) feedforward (“nonnormal”) connectivity whereby a “source mode” of E-I imbalance feeds into a “sink mode” in which balance is restored, and (ii) anti-symmetric connectivity that causes the two populations to oscillate.

To study the impact of each of these prototypical connectivity motifs on movement preparation, we implemented them separately, i.e. as two small networks of two units each, with an overall connectivity scale parameter w which we varied (**Figure 4A** and **D**; Methods). As both nonnormal and oscillatory dynamics arise from linear algebraic properties of the connectivity matrix, we considered linear network dynamics for this analysis ($\phi(x) = x$ in **Equation 3**). Moreover, to preserve the existence of an output nullspace in which preparation could in principle occur without causing premature movement, we reduced the dimensionality of the motor readout from 2D (where there would be no room left for a nullspace) to 1D (leaving a 1D nullspace), and adapted the motor task so that the network now had to move the hand position along a single dimension (**Figure 4B** and **E**, top). Analogous to the previous arm model, we assumed that the hand’s acceleration along this axis was directly given by the 1D network readout.

We found that optimal control of both dynamical motifs generally led to preparatory dynamics, with inputs arising before the go cue (**Figure 4B** and **E**, bottom). In the feedforward motif, the amount of preparatory inputs appeared to depend critically on the orientation of the readout. When the readout was aligned with the sink (brown) mode (**Figure 4B**, left), the controller prepared the network by moving its activity along the source (orange) mode (**Figure 4C**, left). This placed the network in a position from which it had a natural propensity to generate large

Figure 2

Optimal control of the ISN network.

(A) Illustration of the different terms in the control cost function, designed to capture the different requirements of the task. “Tgt” marks the time of target onset, “Go” that of the go cue (known in advance) and “End” the end of the trial. (B) Time course of the hand velocity (top), optimal control inputs (middle; 10 example neurons), and firing rates (bottom, same neurons) during a delayed reach to one of the 8 targets shown in **Figure 1A**. Here, the delay period was set to $\Delta_{\text{prep}} = 300$ ms. Note that inputs arise well before the go cue, even though they have no direct effect on behaviour at that stage. (C) Dependence of the different terms of the cost function on preparation time. All costs are normalized by the total cost at $\Delta_{\text{prep}} = 0$ ms. The inset shows the time course of the hand’s average distance to the relevant target when no preparation is allowed (blue) and when preparation is allowed (red). Although the target is eventually reached for all values of Δ_{prep} , the hand gets there faster with longer preparation times, causing a decrease in $\mathcal{J}_{\text{tgt}}$ – and therefore also in $\mathcal{J}_{\text{tot}}$. Another part of the decrease in $\mathcal{J}_{\text{tot}}$ is due to a progressively lower input energy cost $\mathcal{J}_{\text{effort}}$. On the other hand, the cost of staying still before the Go cue increases slightly with Δ_{prep} . (D) We define the preparation index as the ratio of the norms of the external inputs during preparation and during movement (see text). The preparation index measures how much the optimal strategy relies on the preparatory period. As more preparation time is allowed, this is used by the optimal controller and more inputs are given during preparation. For longer preparation times, this ratio increases sub-linearly, and eventually settles.





Figure 3

Conservation of the optimal control strategy across delays

(A) Optimal control inputs to ten randomly chosen neurons in the model RNN (left) and their corresponding firing rates (right) for different preparation times Δ_{prep} (ranging from 0 to 800 ms; c.f. labels). (B) Projection of the movement-epoch population activity for each of the 8 reaches (panels) and each value of Δ_{prep} shown in A (darker to lighter colors). These population trajectories are broadly conserved across delay times, and become more similar for larger delays.






Figure 4

Analysis of the interplay between the optimal control strategy and two canonical motifs of E-I network dynamics:

nonnormal transients driven by feedforward connectivity (A–C), and oscillations driven by anti-symmetric connectivity (D–F). **(A)** Activity flow field (10 example trajectories) of the nonnormal network, in which a “source” unit (orange) drives a “sink” unit (brown). We consider two opposite readout configurations, where it is either the sink (left) or the source (right) that drives the acceleration of the hand. **(B)** Temporal evolution of the hand position (top; the dashed horizontal line indicates the reach target), hand acceleration (middle) and optimal control inputs to the two units (bottom; colours matching panel A), under optimal control given each of the two readout configurations shown in A (left vs. right). The dashed vertical line marks the go cue, and the gray bar indicates the delay period. While the task can be solved successfully in both cases, preparatory inputs are only useful when the sink is read out. **(C)** Network activity trajectories under optimal control. Each trajectory begins at the origin, and the end of the delay period is shown with a black cross. **(D–F)** Same as (A–C), for the oscillatory network. **(G–H)** Preparation index (top) and total amount of preparatory inputs (bottom) as a function of the scale parameter w of the network connectivity, for various readout configurations (colour-coded as shown in the top inset). The nonnormal network (top) prepares more when the readout is aligned to the most controllable mode, while the amount of preparation in the oscillatory network (bottom) is independent of the readout direction. The optimal strategy must balance the benefits from preparatory inputs which allow to exploit the intrinsic network dynamics, with the constraint to remain still. This is more difficult when the network dynamics are strong and pushing activity out of the readout-null subspace, explaining the decrease in preparation index for large values of w in the oscillatory network.

activity transients along the readout dimension (c.f. flow field in [Figure 4A](#)); here, these transients were exploited to drive the fast upstroke in hand acceleration and throw the hand towards the target location. Note that this strategy reduces the amount of input the controller needs to deliver during the movement, because the network itself does most of the work. Nevertheless, in this case the network's own impulse response was not rich enough to accommodate the phase reversal required to subsequently slow the hand down and terminate the movement. Optimal control therefore also involved inputs during the movement epoch, leading to a preparatory index of ~ 0.54 ([Figure 4G](#), dark blue).

When it was instead the source mode that was read out ([Figure 4B](#), right), the only dimension along which the system could prepare without moving was the sink mode. Preparing this way is of no benefit, because the flow field along the sink mode has no component along the source (readout) mode. Thus, here the optimal strategy was to defer control to the movement epoch, during which the transient growth of network activity along the readout rested entirely on adequate control inputs. This led to a preparation index of ~ 0 ([Figure 4G](#), pale green). Although the network did react with large activity excursions along the sink mode ([Figure 4C](#), right), these were inconsequential for the movement. Importantly, of the two extreme readout configurations discussed above, the first one yielded a smaller overall optimal control cost (by a factor of ~ 1.5). Thus, at a meta-control level, ideal downstream effectors would read out the sink mode, not the source mode. Note that while increasing the connectivity strength initially led to more preparation ([Figure 4H](#)), a plateau was eventually reached for $w \geq 4$. Indeed, while stronger dynamics initially make preparation more beneficial, they also make it more difficult for preparatory activity to remain in the readout nullspace.

We obtained similar insights for oscillatory network dynamics ([Figure 4D-F](#)). A key difference however was that the flow field was rotationally symmetric such that no distinction could be made between “source” and “sink” units – indeed the optimal control strategy yielded the same results (up to a rotation of the state space) irrespective of which of the two units was driving the hand's acceleration (compare left and right panels in [Figure 4D-F](#)). Nevertheless, the optimal controller consistently moved the network's activity along the output-null axis during preparation, in such a way as to engage the network's own rotational flow immediately after the go cue ([Figure 4F](#)). This rotational flow drove a fast rise and decay of activity in the readout unit, thus providing the initial segment of the required hand acceleration. The hand was subsequently slowed down by modest movement-epoch control inputs which eventually receded, leading to a preparation index of ~ 0.58 . Interestingly, the preparation index showed a decrease for very large w ([Figure 4G](#)), which did not reflect smaller preparatory inputs ([Figure 4H](#)) but rather reflected the larger inputs that were required during movement to cancel the fast oscillations naturally generated by the network.

The above results highlight how the optimal control strategy is shaped by the dynamical motifs present in the network. Crucially, we found that the optimal way to control the movement depends not only on the strength and flow of the internal network dynamics, but also on their interactions with the readout.

Control-theoretic properties predict the amount of preparation

Our investigation of preparation in a low-dimensional system allowed us to isolate the impact of core dynamical motifs, and highlighted how preparation depends on the geometry of the flow field, and its alignment to the readout. However, these intuitions remain somewhat qualitative, making them difficult to generalize to our high-dimensional ISN model.

To quantify the key criteria that appear important for preparation, we turned to tools from control theory. We reasoned that, for a network to be able to benefit from preparation and thus exhibit a large preparation index, there must be some advantage to using early inputs that do not immediately cause movement, relative to using later inputs that do. We hypothesized that this

advantage could be broken down into two criteria. First, there must exist activity patterns that are momentarily output-null (i.e. do not immediately cause movement) yet seed output-potent dynamics that subsequently move the arm. The necessity of this criterion was obvious in the 2D nonnormal network, which did not display any preparation when its nullspace was aligned with its “sink” mode. In the language of control theory, this criterion implies that the nullspace of the readout must be sufficiently “observable” – we captured this in a scalar quantity α (Methods; Kao and Hennequin, 2019 [↗](#); Skogestad and Postlethwaite, 2007 [↗](#)). Second, there must be a sizeable cost to performing the movement in an entirely input-driven manner without relying on preparation. In other words, the network should be hard to steer along the readout direction, i.e. the readout must be of limited “controllability” – we captured this in another scalar quantity β (Methods).

We illustrate the meaning of these two metrics in **Figure 5A** [↗](#) and **B** [↗](#) for a 2-dimensional example network that combines nonnormality and oscillations. We show two extreme choices of readout direction (**Figure 5A** [↗](#), dashed black): the one that maximizes α (top) and the one that minimizes it (bottom). In the first case, the readout nullspace (dashed orange) is very observable, i.e. trajectories that begin in the nullspace evolve to produce large transients along the readout (solid orange & inset). In the second case, the opposite is true. For each case, we also assessed the controllability of the readout (β). The controllability of a direction corresponds to how much variance activity trajectories exhibit along that direction, when they are randomly and isotropically initialized (**Figure 5B** [↗](#)). In other words, a very controllable direction is one along which network trajectories have a natural tendency to evolve.

We then assessed how well α and β could predict the preparation index of individual networks. In 2D networks, we found that a simple function that grows with α and decreases with β could accurately predict preparation across thousands of networks Supplementary Material S3.2. To assess whether these insights carried over to high-dimensional networks, we then generated a range of large ISNs with parameterically varied connectivity strengths and decay timescales (Methods). We then regressed the preparation index against the values of α and β computed for each of these networks (as controllability and observability are only defined for linear networks, we set $\phi(x) = x$ for this investigation). We found that a simple linear mapping, $\text{prep. index} = k_0 + k_\alpha \alpha + k_\beta \beta$, with parameters fitted to one half of the ISN networks, accurately predicted the preparation indices of the other half (**Figure 5C** [↗](#); $R^2 = 0.93$, 5-fold cross-validated). Interestingly, we observed that although α and β (which are both functions of the network connectivity) were highly correlated across different networks, discarding either variable in our linear regression led to a significant drop in R^2 (**Figure 5C** [↗](#), inset). Importantly, it was their difference that best predicted the preparation index ($k_\alpha > 0$ and $k_\beta < 0$), consistent with our hypothesis that the preparation index is a relative quantity which increases as the nullspace becomes more observable, but decreases as readout dimensions become more controllable.

We were able to confirm the generality of this predictive model by generating networks with other types of connectivity (oscillatory networks, and networks with unstructured random weights), which displayed dynamics very different from the ISNs (see Supplementary Material S4). Interestingly, despite the different distribution of α and β parameters in those networks, we could still capture a large fraction of the variance in their preparation index ($R^2 = 0.69$) using the linear fit obtained from the ISNs alone. This confirms that α and β can capture information about the networks’ dynamics in a universal manner.

Note that we do not make any claims about the specific functional form of the relationship between α , β and the preparation index. Rather, we claim that there should be a broad trend for the preparation index to increase with α and decrease with β , and we acknowledge that this relationship could in general be nonlinear. Indeed, in 2D networks, we found that the preparation index was in fact better predicted by the ratio of α over β than by their difference (Supplementary Material S3.2).




Figure 5

Predicting the preparation index from the observability of the output nullspace (α) and the controllability of the readout (β , see details in text).

(A) Illustration of the observability of the output nullspace in a synthetic 2-dimensional system. The observability of a direction is characterized by how much activity (integrated squared norm) is generated along the readout by a unit-norm initial condition aligned with that direction. The top and bottom panels show the choices of readout directions (dotted black) for which the corresponding nullspace (dotted orange) is most (maximum α) and least (minimum α) observable, respectively. Trajectories initialized along the null direction are shown in solid orange, and their projections onto the readout are shown in the inset. **(B)** Illustration of the controllability of the readout in the same 2D system as in (A). To compute controllability, the distribution of activity patterns collected along randomly initialized trajectories is estimated (heatmap); the controllability of a given direction then corresponds to how much variance it captures in this distribution. Here, the network has a natural propensity to generate activity patterns aligned with the dashed white line ('most controllable' direction). The readout directions are repeated from panel A (dotted black). The largest (resp. smallest) value of β would by definition be obtained when the readout is most (resp. least) controllable. Note the tradeoff in this example: the choice of readout that maximizes α (top) does not lead to the smallest β . **(C)** The values of α and β accurately predict the preparation index ($R^2 = 0.93$) for a range of high-dimensional ISNs (maroon dots) with different connectivity strengths and characteristic timescales (Methods). The best fit (after z-scoring) is given by $f(\alpha, \beta) = (16.94 \pm 0.02)\alpha - (15.97 \pm 0.02)\beta$ (mean $\pm$ sem were evaluated by bootstrapping). This confirms our hypothesis that optimal control relies more on preparation when α is large and β is small. Note that α and β alone only account for 34.8% and 30.4% of the variance in the preparation index, respectively (inset). Thus, α and β provide largely complementary information about the networks' ability to use inputs, and can be combined into a very good predictor of the preparation index. Importantly, even though this fit was obtained *using ISNs only*, it still captures 69% of preparation index variance across networks from other families (blue dots; Methods).

Finally, as the above results highlight that the amount of preparation depends on the alignment between internal dynamics and readout, one may wonder how much our conclusions depend on our use of a random unstructured readout matrix. First, we note that the effect of the alignment on preparation index is greatly amplified in the low-dimensional networks (**Figure 4G** [↗](#)). In high-dimensional networks, any random readout matrix C will pick out activity dimensions in the RNN that are sufficiently aligned with the most controllable directions of the dynamics to encourage preparation. Second, we performed additional simulations where we meta-optimized the readout so as to minimize the average optimal cost per movement. The resulting system is more observable overall (as it allows the network to solve the task at a lower cost) but relies just as much on preparation (**Figure S7** [↗](#)).

Modelling movement sequences

Having gained a better understanding of what features lead a network to prepare, we next set out to assess whether optimal control could also explain the neural preparatory processes underlying the generation of movement *sequences*. We revisited the experimental studies of [Zimnik and Churchland \(2021\)](#) [↗](#), where mon-keys were trained to perform two consecutive reaches. Each trial started with the display of both targets, followed by an explicitly enforced delay period before the onset of the first reach. A distinction was made between “double” reaches in which a pause was enforced between reaches, and “compound” reaches in which no pause was required. This study concluded that, rather than the whole movement sequence unrolling from a single preparatory period, each reach was instead successively seeded by its own preparatory activity.

Here, we asked whether such an independent, successive preparation strategy would arise as an optimal control solution, in the same way that single-reach preparation did. Importantly, we could not answer this question by directly examining network inputs as we did for single reaches. Indeed, any network input observed before the second reach could be contributing either to the end of the first movement, or to the preparation of the next. In fact, the issue of teasing apart preparatory vs. movement-related activity patterns also arose in the analysis of the monkey data. To address this, [Zimnik and Churchland \(2021\)](#) [↗](#) exploited the fact that monkey M1 activity just before and during single reaches is segregated into two distinct subspaces. Thus, momentary activity patterns (during either single or double reaches) can be unambiguously labelled as preparatory or movement-related depending on which of the two subspaces they occupied. We performed a similar analysis (Methods) and verified that preparatory and movement activity patterns in the model were also well segregated in their respective subspaces in the single-reach task (**Figure 6A-B** [↗](#)). We then assessed the occupancy of the preparatory subspace during double reaching in the model, and took this measure as a signature of preparation.

To model optimal control of a double reach, we modified our cost functional to account for the presence of two consecutive targets (see Methods). We considered the same set of eight targets as in our single-reach task, and modelled all possible combinations of two targets (one example shown in **Figure 6** [↗](#)). We set the hyperparameters of the cost function such that both targets could be reached by the resulting optimal controller, in a way that matched important qualitative aspects of the monkeys’ behaviour (in particular, such that both reaches were performed at similar velocities, with the second reach lasting slightly longer on average; **Figure 6B-C** [↗](#); top).

We projected the network activity onto preparatory and movement subspaces identified using single and double reaches activity (Methods). For double reaches with a long (600ms) pause, the preparatory subspace was transiently occupied twice, with the two peaks occurring just before the onset of each movement in the sequence (**Figure 6B** [↗](#); bottom).

Notably, the occupancy during the “compound” reach (without pause; **Figure 6C** [↗](#)) also started rising prior to the first movement before decaying very slightly and peaking again before the second reach, indicating two independent preparatory events. This is somewhat surprising, given that a movement sequence can also be viewed as a single “compound” movement, for which we




Figure 6

The model executes a sequence of two reaches using a holistic strategy.

(A) Hand velocity during one of the reaches, with the corresponding hand trajectory shown in the inset. (B–C) We identified two 6-dimensional orthogonal sub-spaces, capturing 79% and 85% of total activity variance during single-reach preparation and movement respectively. (B) First principal component of the model activity for the 8 different reaches projected into the subspaces identified using preparatory (top) and movement-epoch (bottom) activity. (C) Occupancy (total variance captured across movements) of the orthogonalized preparatory and movement subspaces, in the model (top) and in monkey motor cortical activity (bottom; reproduced from Lara et al., 2018 for monkey Ax). We report mean $\pm$ s.e.m., where the error is computed by bootstrapping from the neural population as in Lara et al. (2018). We normalize each curve separately to have a maximum mean value of 1. To align the model and monkey temporally, we re-defined the model's 'movement onset' time to be 120 ms after the model's hand velocity crossed a threshold – this accounts for cortico-spinal delays and muscle inertia in the monkey. Consistent with Lara et al. (2018)'s monkey M1 recordings, preparatory subspace occupancy in the model peaks shortly before movement onset, rapidly dropping thereafter to give way to pronounced occupancy of the movement subspace. Conversely, there is little movement subspace occupancy during preparation. (D) Behavioural (top) and neural (middle) correlates of the delayed reach for one example of a double reach with an enforced pause of 0.6 s. The optimal strategy relies on preparatory inputs preceding each movement. (E) Same as (C), for double reaches. The onsets of the monkey's two reaches are separately aligned to the model's using the same convention as in (C). The preparatory subspace displays two clear peaks of occupancy. This double occupancy peak is also observed in monkey neural activity (bottom; reproduced from Zimnik and Churchland, 2021, with the first occupancy peak aligned to that of the model). (F) Same as (D), for compound reaches with no enforced pause in between. Even though the sequence could be viewed as a single long movement, the control strategy relies on two periods of preparation. Here, inputs before the second reach are used to reinject energy into the system after slowing down at the end of the first reach. (G) Even though no explicit delay period is enforced in-between reaches during the compound movement, the preparatory occupancy rises twice, before the first reach and once again before the second reach. This is similar to observations in neural data (bottom; reproduced from Zimnik and Churchland, 2021).

have shown previously a unique preparatory phase is sufficient (**Figure 2**). In our model, this behaviour can be understood to arise from the requirement that the hand stop briefly at the first target. To produce the second reach, the hand needs to accelerate again, which requires transient re-growth of activity in the network. Given that the network's dynamical repertoire exhibits limited timescales, this is most easily achieved by reinjecting inputs into the system.

In summary, our results suggest that the “independent” preparation strategy observed in monkeys is consistent with the optimal control of a two-reach sequence. While [Zimnik and Churchland \(2021\)](#) showed that RNNs trained on this task used this “independent” strategy, this was by design as the network was only cued for the second reach after the first one had started. In addition to replicating this proof of concept that it is possible to prepare whilst moving, our model also shows how and why independent preparation might arise as an optimal control solution.

Discussion

In this work, we proposed a model for the dynamics of motor cortex during a delayed reaching task in non-human primates. Unlike previous work, we treated M1 as an input-driven nonlinear dynamical system, with generic connectivity not specifically optimized for the task, but with external inputs assumed to be optimal for each reach. Motivated by a large body of evidence suggesting that preparation is useful before delayed reaches ([Churchland et al., 2010](#); [Lara et al., 2018](#); [Afshar et al., 2011](#); [Shenoy et al., 2013](#)), but also evidence for thalamic inputs being necessary for accurate movement execution ([Sauerbrei et al., 2020](#)), we used this model to investigate whether and why neural circuits might rely on motor preparation during delayed reaching tasks. Interestingly, preparation arose as an optimal control strategy in our model, with the optimal solution to the task relying strongly on inputs prior to movement onset. Moreover, the benefits of preparation were dependent on the network connectivity, with preparation being more prevalent in networks whose rich internal dynamics can be advantageously seeded by early external inputs. We were able to quantify this intuition with a predictive model relating the dynamical response properties of a network to the amount of preparation it exhibits when controlled optimally. Finally, we found that prominent features of the monkeys' neural activity during sequential reaches arose naturally from optimal control assumptions. Specifically, optimally controlled networks relied on two phases of preparation when executing sequences of two reaches, corroborating recent experimental observations in monkey M1 ([Zimnik and Churchland, 2021](#)). Together, our results provide a normative explanation for the emergence of preparatory activity in both single and sequential reaching movements.

In recent years, task-optimized recurrent neural networks have become a very popular tool to model neural circuit dynamics. Classically, those models incorporate only those inputs that directly reflect task-related stimuli (e.g. motor target, go cue, etc). This requires assumptions about the form of the inputs, such as modelling them as simple step functions active during specific task epochs. However, as local neural circuits are part of a wider network of brain areas, a large portion of their inputs come from other brain areas at intermediate stages of the computation and may therefore not be directly tied to task stimuli. Thus, it is not always obvious what assumptions can reasonably be made about the inputs that drive the circuit's dynamics. Our optimization framework, which does not require us to make any specific assumptions about when and how inputs enter the network (although it does allow to incorporate prior information in the form of constraints), allows to bypass this problem and to implicitly model unobserved inputs from other areas. Importantly, our framework allows to ask questions – such as “why prepare” – that are difficult to phrase in standard input-driven RNN models. We note, however, that in the investigation we have presented here, the lack of imposed structure for the inputs also implied that the model could not make use of mechanisms known to contribute certain aspects of

preparatory neural activity. For example, our model did not exhibit the usual visually-driven response to the target input, nor did it have to use the delay epoch to keep such a transient sensory input in memory (Guo et al., 2014 [DOI](#); Li et al., 2015 [DOI](#)).

The main premise of our approach is that one can some-how delineate the dynamical system which M1 implements, and attribute any activity patterns that it cannot autonomously generate to external inputs. Just where the anatomical boundary of “M1 dynamics” lie – and therefore where “external inputs” originate – is unclear, and our results must be interpreted with this limitation in mind. Operationally, previous works in reaching monkeys have shown that M1 data can be mathematically well described by a dynamical system that appears largely autonomous during movement. These works have emphasized that those abstract dynamics, while inferred from M1 data alone, may not be anatomically confined to M1 itself. Instead, they may involve interactions between multiple brain areas, and even possibly parts of the body through delayed sensory feedback. Here, we too tend to think of our M1 models in this way, and therefore attribute external input to brain areas that are one-step removed from this potentially broad motor-generating network. Nevertheless, a more detailed multi-area model of the motor-generating circuitry including e.g. spinal networks (Prut and Fetz, 1999 [DOI](#)) could enable more detailed comparisons to multi-region neural data. In a similar vein, our model makes no distinction between external inputs that drive movement-specific planning computations, and other types of movement-unspecific inputs that might drive the transition from planning to execution (e.g. “trigger” inputs, Kaufman et al., 2016 [DOI](#)). Incorporating such distinctions (e.g. by temporally modulating the cost in individual input channels depending on specific task events, or by having separate channels for movement-unspecific inputs) might allow to ask more targeted questions about the role and provenance of external inputs.

A major limitation of our study is the specific choice of a quadratic penalty on the external input in our control objective. Whilst there are possible justifications for such a cost (e.g. regularization of the dynamics to promote robustness of the control solution), its use here is mainly motivated by mathematical tractability. Other costs might be conceivably more relevant and might affect our results. For example, studies of motor cortex have long thought of its dynamics as converting relatively simple inputs reflecting high-level, temporally stable plans, into detailed, temporally-varying motor commands. Thus, a potentially relevant form of a penalty for external inputs would be their temporal complexity. Such a penalty would have the advantage of encouraging a clearer separation between the inputs and the RNN activations; indeed, in our current model, we find that the optimal controls themselves have a temporal structure, part of which could be generated by a dynamical system and thus potentially absorbed into our “M1 dynamics”. To address this, we note that our optimization framework can be adjusted to penalize the magnitude of the temporal *derivative* of the external input $\|\dot{\mathbf{u}}\|^2$, instead of $\|\mathbf{u}\|^2$. We experimented with this extension and found qualitatively different optimal inputs and M1 firing rates, which evolved more slowly and plateaued for sufficiently long preparation (**Figure S8A-D** [DOI](#)) – this is in fact more consistent with monkey M1 data (e.g. Elsayed et al., 2016 [DOI](#)). Despite these qualitative difference in the specific form of preparation, our main conclusion stands that input-driven preparation continues to arise as an optimal solution (**Figure S8E-F** [DOI](#)).

Another important assumption we have made is that the optimal controller is aware of the duration of the delay period. While this made solving for the optimal control inputs easier, it made our task more akin to a self-initiated reach (Lara et al., 2018 [DOI](#)) than to a typical delayed reach with unpredictable, stochastic delay durations. Future work could revisit this assumption. As a first step towards this, we now briefly outline pilot experiments in this direction. We used an exponential distribution of delays (with mean 300 ms) and devised two modified versions of our model that dealt with the resulting uncertainty in two different ways. In the first strategy, at any time during preparation, the model would estimate the most probable time-to-go-cue given that it hadn’t arrived yet (in this case, this is always 300 ms in the future) and would plan an optimal sequence of inputs accordingly. In the second strategy, the network would prudently assume the

earliest possible go cue (i.e. the next time step) and plan accordingly. In both cases, only the first input in the optimal input sequence would be used at each step, and complete replanning would follow in the next step, as the model re-assesses the situation given new information (i.e. whether the actual go cue arrived or not; this is a form of “model predictive control”, [Rawlings et al., 2017](#)). Preparatory inputs arose in both settings, but we found that only the latter strategy led to activity patterns that plateaued early during preparation (see **Figure S9**).

Throughout the main text, we have referred to Δ_{prep} as the task-enforced delay period. However, a more accurate description may be that it corresponds to a delay period determined by an internally set go signal, which can lag behind the external go cue. While we would not expect a large difference between those two signals, the way in which we define Δ_{prep} becomes important as it approaches 0 ms (limit of a quasi-automatic reach; [Lara et al., 2018](#)). Indeed, in this limit, our model exhibits almost no activity in the preparatory subspace (as defined in **Figure 6** – see further analyses in **Figure S10**). In contrast, monkey M1 activity was found to transiently occupy the preparatory subspace even in this case ([Lara et al., 2018](#)). Evidence for a delay between the earliest possible response to sensory cues and the trigger of movement was also observed in [Kaufman et al. \(2016\)](#), as well as in human behavioral studies ([Haith et al., 2016](#)). Thus, one may wish to explicitly incorporate this additional delay in the model in order to make it more realistic. Note however that [Haith et al. \(2016\)](#) showed that this internal delay could be shortened without affecting movement accuracy, suggesting that part of the processing that empirically occurs inbetween the internal and external go cues may not be necessary, but rather reflect a decoupling between the end of preparation and the trigger of movement. This may be important to consider when attempting to compare the model to e.g. reaction times from behavioral experiments.

Dynamical systems have a longstanding history as models of neural populations ([Dayan and Abbott, 2001](#)). However, understanding how neural circuits can perform various computations remains a challenging question. Recently, there has been increased interest in trying to understand the role of inputs in shaping cortical dynamics. This question has been approached both from a data-driven perspective ([Malonis et al., 2021](#); [Soldado-Magraner et al., 2023](#)), and in modelling work with e.g. [Driscoll et al. \(2022\)](#) showing how a single network can perform different tasks by reorganizing its dynamics under the effect of external inputs and [Dubreuil et al. \(2021\)](#) relating network structure to the ability to process contextual inputs. To better understand how our motor system can generate flexible behaviours ([Logiaco et al., 2021](#); [Stroud et al., 2018](#)), and to characterize learning on short timescales ([Sohn et al., 2020](#); [Heald et al., 2023](#)), it is important to study how network dynamics can be modulated by external signals that allow rapid adaptation to new contexts without requiring extensive modifications of the network's connectivity. The optimal control approach we proposed here offers a way to systematically perform such evaluations, in a variety of tasks and under different assumptions regarding how inputs are allowed to impact the dynamics of the local circuit of interest. While our model's predictions will depend on e.g. the choice of connectivity or the design of the cost function, an exciting direction for future work will be to obtain those parameters in a data-driven manner, for instance using recently developed methods to infer dynamics from data ([Pandarinath et al., 2018](#); [Schimel et al., 2022](#)), and advances in inverse reinforcement learning and differentiable control ([Amos et al., 2018](#)) to infer the cost function that behaviour optimizes. These could additionally be combined with more biomechanically realistic effectors, such as the differentiable arm models from [Codol et al. \(2023\)](#).

Methods

Experimental model and subject details

In [Figure 1](#), we showed data from two primate datasets that were made available to us by Mark Churchland, Matthew Kaufman and Krishna Shenoy. Details of animal care, surgery, electrophysiological recordings, and behavioral task have been reported previously in [Churchland et al. \(2012\)](#); [Kaufman et al. \(2014\)](#) (see in particular the details associated with the J and N “array” datasets). The subjects of this study, J and N, were two adult male macaque monkeys (*Macaca mulatta*). The animal protocols were approved by the Stanford University Institutional Animal Care and Use Committee. Both monkeys were trained to perform a delayed reaching task on a fronto-parallel screen. At the beginning of each trial, they fixated on the center of the screen for some time, after which a target appeared on the screen. After a variable delay period (0–1000 ms), a go cue appeared instructing the monkeys to reach toward the target. Recordings were made in the dorsal premotor cortex and in the primary motor cortex using a pair of implanted 96-electrode arrays. In [Figure 6](#), we also reproduced data from [Lara et al. \(2018\)](#) and [Zimnik and Churchland \(2021\)](#). Details of animal care, surgery, electrophysiological recordings, and behavioral task for those data can be found in the Methods section of the respective papers.

Arm model

To simulate reaching movements, we used the planar two-link arm model described in [Li and Todorov \(2004\)](#). The two links have lengths L_1 and L_2 , masses M_1 and M_2 , and moments of inertia I_1 and I_2 respectively. The lower arm’s center of mass is located a distance D_2 from the elbow. By considering the geometry of the upper and lower limb, the position of the hand and elbow can be written as vectors $\mathbf{y}_h(t)$ and $\mathbf{y}_e$ given by

$$\begin{aligned}\mathbf{y}_h &= \begin{pmatrix} L_1 \cos \theta_1 + L_2 \cos(\theta_1 + \theta_2) \\ L_1 \sin \theta_1 + L_2 \sin(\theta_1 + \theta_2) \end{pmatrix} \text{ and} \\ \mathbf{y}_e &= \begin{pmatrix} L_1 \cos \theta_1 \\ L_1 \sin \theta_1 \end{pmatrix}.\end{aligned}\quad (6)$$

The joint angles $\boldsymbol{\theta} = (\theta_1; \theta_2)^T$ evolve dynamically according to the differential equation

$$\mathbf{m}(t) = \mathcal{M}(\boldsymbol{\theta})\ddot{\boldsymbol{\theta}} + \mathcal{X}(\boldsymbol{\theta}, \dot{\boldsymbol{\theta}}) + \mathcal{B}\dot{\boldsymbol{\theta}}, \quad (7)$$

where $\mathbf{m}(t)$ is the momentary torque vector, $\mathcal{M}$ is the matrix of inertia, $\mathcal{X}$ accounts for the centripetal and Coriolis forces, and $\mathcal{B}$ is a damping matrix representing joint friction. These parameters are given by

$$\mathcal{M}(\boldsymbol{\theta}) = \begin{pmatrix} a_1 + 2a_2 \cos \theta_2 & a_3 + a_2 \cos \theta_2 \\ a_3 + a_2 \cos \theta_2 & a_3 \end{pmatrix} \quad (8)$$

$$\mathcal{X}(\boldsymbol{\theta}, \dot{\boldsymbol{\theta}}) = a_2 \sin \theta_2 \begin{pmatrix} -\dot{\theta}_2(2\dot{\theta}_1 + \dot{\theta}_2) \\ \dot{\theta}_1^2 \end{pmatrix} \quad (9)$$

$$\mathcal{B} = \begin{pmatrix} 0.05 & 0.025 \\ 0.025 & 0.05 \end{pmatrix} \quad (10)$$

with $a_1 = I_1 + I_2 + M_2 L_1^2$, $a_2 = M_2 L_1 D_2$, and $a_3 = I_2$.

iLQR algorithm

Throughout this work, we used the iLQR algorithm ([Li and Todorov, 2004](#)) to find the locally optimal inputs that minimize our cost function. iLQR is a trajectory optimization algorithm that can handle nonlinear dynamics and non-quadratic costs. iLQR works in an iterative manner, by

linearizing the dynamics and performing a quadratic approximation of the cost at each iteration, thus turning the control problem into a local linear-quadratic problem whose unique solution is found using LQR (Kalman et al., 1960). The LQR solver uses a highly efficient dynamic programming approach that exploits the sequential structure of the problem. Our implementation of iLQR followed from Li and Todorov (2004) [with](#), with the difference that we performed regularization of the local curvature matrix as recommended by Tassa (2011) [with](#).

Generation of the high-dimensional readouts and networks

Generation of inhibitory-stabilized networks

Simulations in **Figures 1** [with](#), **3** [with](#), **5** [with](#) and **6** [with](#) were conducted using inhibition-stabilized networks (ISN). Those were generated according to the procedure described in Hennequin et al. (2014) [with](#) minor adjustments. In brief, we initialized strongly connected chaotic networks with sparse and log-normally distributed excitatory weights, and stabilized them through progressive $\mathcal{H}_2$ -optimal adjustments of the inhibitory weights until the spectral abscissa of the connectivity matrix fell below 0.8. This yielded strongly connected but stable networks with a strong degree of non-normality. When considering a larger range of ISNs (**Figure 5** [with](#)), we independently varied both the variance of the distribution of initial excitatory weights and the spectral abscissa below which we stopped optimizing the inhibitory weights.

Generation of additional networks in Figure 5 [with](#)

To assess the generality of our findings in **Figure 5** [with](#), we additionally generated randomly connected networks by sampling each weight from a Gaussian distribution with $\sigma = R/\sqrt{N}$, where the spectral radius R was varied between 0 and 0.99. We also sampled skew-symmetric networks by drawing a random network $\mathbf{S}$ as above, and setting $\mathbf{W} = (\mathbf{S} - \mathbf{S}^T)/2$. We varied the radius R of the $\mathbf{S}$ matrices between 0 and 5. Moreover, we considered diagonally shifted skew-symmetric networks $\mathbf{W} = (\mathbf{S} - \mathbf{S}^T)/2 + \lambda \mathbf{I}$ where λ denotes the real part of all the eigenvalues and was varied between 0 and 0.8.

The elements of the readout matrix $\mathbf{C}$ mapping neural activity onto torques were drawn from a normal distribution with zero mean and standard deviation $\sigma_C = 0.05/\sqrt{N}$. This was chosen to ensure that firing rates of standard deviation on the order of 30Hz would be decoded into torques of standard deviation ~ 2 N/m, which is the natural variation required for the production of the reaches we considered.

Details of Figure 4 [with](#)

To more easily dissect the phenomena leading to the presence or absence of preparation, we turned to 2D linear networks in **Figure 4** [with](#). We modelled nonnormal networks with a connectivity $\mathbf{W} = \begin{bmatrix} 0 & 0 \\ w & 0 \end{bmatrix}$ and oscillatory networks with connectivity $\mathbf{W} = \begin{bmatrix} 0 & -w \\ w & 0 \end{bmatrix}$. The activity of the two units evolved as

$$\tau \dot{\mathbf{x}}(t) = -\mathbf{x}(t) + \mathbf{W}\mathbf{x}(t) + \mathbf{u}(t) \quad (11)$$

and directly influenced the acceleration of a one-dimensional output $y(t)$ according to

$$\ddot{y}(t) = \mathbf{C}_i \mathbf{x}(t) \quad (12)$$

where $\mathbf{C}_i = [\cos \theta_C \sin \theta_C]$ was a row matrix reading the activity of the network along an angle θ_C from the horizontal (first unit). Our setup aimed to mirror the reaching task studied in this work. We thus optimized inputs to minimize the following cost function :

$$\begin{aligned} \mathcal{J}[\mathbf{u}] =: & \underbrace{\int_0^T \|y(t) - y^*\|^2 \frac{t^2}{T^2} \frac{dt}{T}}_{\mathcal{J}_{\text{target}}} \\ & + \underbrace{\alpha_{\text{null}} \int_{-\Delta_{\text{prep}}}^0 (\|y(t)\|^2 + \|\dot{y}(t)\|^2 + \|\ddot{y}(t)\|^2) \frac{dt}{T}}_{\mathcal{J}_{\text{null}}} \\ & + \underbrace{\alpha_{\text{effort}} \int_{-\Delta_{\text{prep}}}^T \|\mathbf{u}(t)\|^2 \frac{dt}{2T}}_{\mathcal{J}_{\text{effort}}}. \end{aligned} \quad (13)$$

where $y^* = 20$ was the target position.

Computing networks' controllability and observability to predict preparation in Figure 5 [↗](#)

As part of our attempt to predict how much a network will prepare given its intrinsic properties only, we computed the prospective potency of the nullspace α , and the controllability of the readout β . For a stable linear dynamical system described by

$$\frac{d\mathbf{x}}{dt} = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{u}(t) \quad (14)$$

$$\mathbf{y}(t) = \mathbf{C}\mathbf{x}(t), \quad (15)$$

the system's observability Gramian $\mathbf{Q}$ can be computed as the unique positive-definite solution of the Lyapunov equation

$$\mathbf{A}^T \mathbf{Q} + \mathbf{Q} \mathbf{A} + \mathbf{C}^T \mathbf{C} = 0. \quad (16)$$

The prospective potency of the nullspace $\mathbf{C}^\perp$ is then defined as

$$\alpha \triangleq \frac{\text{Tr}(\mathbf{C}^\perp \mathbf{Q} \mathbf{C}^{\perp T})}{N - 2} \quad (17)$$

Note that this measure α is invariant to the specific choice of basis for the nullspace $\mathbf{C}^\perp$. Similarly, to assess the controllability of the readout, we first computed the controllability Gramian of the system $\mathbf{P}$, which is the solution of

$$\mathbf{A} \mathbf{P} + \mathbf{P} \mathbf{A}^T + \mathbf{B} \mathbf{B}^T = 0, \quad (18)$$

with $\mathbf{B} = \mathbf{I}$ in our system. We then defined the controllability of the readout as

$$\beta \triangleq \frac{\text{Tr}(\mathbf{C} \mathbf{P} \mathbf{C}^T)}{2}. \quad (19)$$

Details of Figure 6 [↗](#)

Cost function

We modelled sequences of reaches by modifying our cost functional to account for the presence of two targets, as

$$\begin{aligned}
 \mathcal{J}[\mathbf{u}] = & \underbrace{\int_0^{\Delta_{\text{move}}^{(1)} + \tau} \|\boldsymbol{\theta}(t) - \boldsymbol{\theta}_1^*\|^2 \frac{t^2}{T^2} dt}_{\mathcal{J}_{\text{target}}^{(1)}} \\
 & + \alpha_{\text{pause}} \underbrace{\int_{\Delta_{\text{move}}^{(1)}}^{\Delta_{\text{move}}^{(1)} + \tau} \|\dot{\boldsymbol{\theta}}(t)\|^2 dt}_{\mathcal{J}_{\text{pause}}} \\
 & + \underbrace{\int_{\Delta_{\text{move}}^{(1)} + \tau}^T \|\boldsymbol{\theta}(t) - \boldsymbol{\theta}_2^*\|^2 \frac{(t - \Delta_{\text{move}}^{(1)} - \tau)^2}{T^2} dt}_{\mathcal{J}_{\text{target}}^{(2)}} \\
 & + \underbrace{\alpha_{\text{null}} \int_{-\Delta_{\text{prep}}}^0 \|\boldsymbol{\theta}(t) - \boldsymbol{\theta}_0\|^2 + \|\dot{\boldsymbol{\theta}}(t)\|^2 + \|\mathbf{m}(t)\|^2 dt}_{\mathcal{J}_{\text{null}}} \\
 & + \underbrace{\alpha_{\text{effort}} \int_{-\Delta_{\text{prep}}}^T \|\mathbf{u}(t)\|^2 dt}_{\mathcal{J}_{\text{effort}}}
 \end{aligned} \tag{20}$$

$$\tag{21}$$

where τ describes how long the monkey's hands had to stay on the intermediate target before performing its second reach. We used $\tau = 600$ ms and $\alpha_{\text{pause}} = 100$ for the double reaches in which a pause was explicitly enforced during the experiment. For compound reaches, the experiment did not require monkeys to stop for any specific duration. However, to ensure that the hand stopped on the target in the model (as it does in experiments when monkeys touch the screen) rather than fly through it, we set $\tau = 6$ ms and $\alpha_{\text{pause}} = 100$ when modelling compound reaches.

Preparatory subspace analysis

Lara et al. (2018) [proposed](#) an analysis to identify preparatory and movement-related subspaces. This analysis allows to assess when the neural activity enters those subspaces, independently of whether it is delay-period or post-go-cue activity.

The method identifies a set of preparatory dimensions and a set of movement dimensions, constrained to be orthogonal to one another, as in [Elsayed et al. \(2016\)](#). These are found in the following manner: the trial-averaged neural activity is split between preparatory and movement-related epochs, yielding two matrices of size $N \times MT$ where N is the number of neurons, T is the number of time bins and M is the number of reaches. One then optimizes the $W_{\text{prep}} \in \mathbb{R}^{N \times d_{\text{prep}}}$ and $W_{\text{mov}} \in \mathbb{R}^{N \times d_{\text{mov}}}$ (where d_{prep} and d_{mov} are the predefined dimensions of the two subspaces) such that the subspaces respectively capture most variance in the preparatory and movement activities, while being orthogonal to one another. This is achieved by maximizing the following objective :

$$\mathcal{C}(W_{\text{prep}}, W_{\text{mov}}) = \frac{1}{2} \left(\frac{\text{Tr}(W_{\text{prep}}^T C_{\text{prep}} W_{\text{prep}})}{Z_{\text{prep}}(d_{\text{prep}})} + \frac{\text{Tr}(W_{\text{mov}}^T C_{\text{mov}} W_{\text{mov}})}{Z_{\text{mov}}(d_{\text{mov}})} \right) \tag{22}$$

where $C_{\text{prep/mov}}$ are the covariance matrices of the neural activity during the preparatory and movement epochs, respectively. The normalizing constant $Z_{\text{prep}}(d_{\text{prep}})$ denotes the maximum amount of variance in preparatory activity that can be captured by any subspace of dimension d_{prep} (this is found via SVD), and similarly for $Z_{\text{mov}}(d_{\text{mov}})$. The objective is maximized under the constraints $W_{\text{prep}}^T W_{\text{mov}} = 0$, $W_{\text{prep}}^T W_{\text{prep}} = I$ and $W_{\text{mov}}^T W_{\text{mov}} = I$. We set subspace dimensions $d_{\text{prep}} = d_{\text{mov}} = 6$, although our results were robust to this choice.

The occupancy of the preparatory subspace was defined as

$$\text{occupancy}^{\text{prep}}(t) = \sum_{k=1}^{d_{\text{prep}}} \text{var}_{\theta}(x_k^{\text{prep}}(t, \theta))$$

and that of the movement subspace was defined as

$$\text{occupancy}^{\text{mov}}(t) = \sum_{k=1}^{d_{\text{mov}}} \text{var}_{\theta}(x_k^{\text{mov}}(t, \theta)).$$

For single reaches, we defined preparatory epoch responses as the activity in a 300ms window before the end of the delay period, and movement epoch responses as the activity in a 300ms window starting 50ms after the go cue. We normalized all neural activity traces using the same procedure as Churchland et al. (2012) [\[1\]](#); Elsayed et al. (2016) [\[2\]](#). For double reaches, we followed Zimnik and Churchland (2021) [\[3\]](#) and used neural activity traces from both single reaches and the first reach of double-reach sequences. Note that we did not include any activity from the second reaches in the sequence, or from compound reaches, when defining the subspaces.

Parameter table

Acknowledgements

We are grateful to Matthew T. Kaufman and Mark M. Churchland for sharing data for the monkey experiments. We thank Kristopher Jensen, David Liu, Javier Antorán, and Rory Byrne for helpful comments on the manuscript. M.S. was funded by an EPSRC DTP studentship, and part of this work was performed using resources provided by the Cambridge Tier-2 system operated by the University of Cambridge Research Computing Service (<http://www.hpc.cam.ac.uk> [\[4\]](#)) funded by EPSRC Tier-2 capital grant EP/P020259/1. For the purpose of open access, the authors have applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising from this submission.

Competing interests

T-C.K. is currently a research scientist at Meta Reality Labs, but only contributed to this work while studying at the University of Cambridge.

Supplementary material

symbol	fig. 1	fig. 2	fig. 3	fig. 5	fig. 4	fig. 6	unit	description
L_1	30				-	30	cm	length of the upper arm in model
L_2	30				-	30	cm	length of the forearm in model
I_1	0.025				-	0.025	kg.m ⁻²	inertia of upper arm
I_2	0.045				-	0.045	kg.m ⁻²	inertia of forearm
M_1	1.4				-	1.4	kg	mass of upper arm
M_2	1.0				-	1.0	kg	mass of forearm
D_2	16				-	16	cm	elbow to lower arm center of mass distance
r	12				20	12	cm	radius of the target reach
$\mu_{\mathbf{h}}$	20				0	-	mV	mean baseline firing rate
$\sigma_{\mathbf{h}}$	5				0	-	mV	s.t.d of the baseline firing rate
α_{effort}	5E-7				1E-5	5E-7	-	coeff. of input cost
α_{null}	1				1	10	-	coeff. of cost of moving during the delay
α_{pause}	-					100	-	coeff. of cost of moving between reaches
τ				150			ms	single-neuron time constant
$\Delta_{\text{move}}^{(1)}$	-					300	ms	duration of the first reach
Δ_{prep}	500	300	-		300	500	ms	delay period time
T	1100	900	-		900	2000 — 1406	ms	total movement duration
N	200				-	200	-	number of neurons
p_{con}	0.2				-	0.2	-	connection probability (E neurons)
p_{E}	80				-	80	-	percentage of E neurons
p_{I}	20				-	20	-	percentage of I neurons

Table 1

Parameters used for the various simulations.

S1 Choice of the hyperparameters of the model

Our cost function for the delayed single-reaching task was composed of 3 components. The relative weightings of the different terms in our cost, which are hyperparameters of the model, affect the way in which the task is solved. To ensure robustness of our results to hyperparameter changes, we scanned the space of α_{null} and α_{effort} (as the solution is invariant to scaling of the cost, only those relative weightings matter), and evaluated the solutions found across this hyperparameter space for a delayed reach of 300 ms.

Our evaluation was based on multiple criteria. We considered the target to have been successfully reached if the mean distance to the target in the last 200 ms of the movement was lower than 5 mm (for a reach radius of 12 cm). We considered that the requirement to stay still during the delay period was satisfied if the mean torques during preparation were smaller than 0.02 N/m. We computed the preparation index and total cost as described in Equation (5) and Equation (4). We moreover computed the total input energy per neuron as $\frac{1}{N} \int_{-\Delta_{\text{prep}}}^T \|\mathbf{u}\|^2 dt$, and the maximum velocity as $\max_t \sqrt{\dot{x}(t)^2 + \dot{y}(t)^2}$. These various quantities are shown for a range of hyperparameters in [Figure S1](#), with the choice of hyperparameters used throughout our simulations marked with a red star. This shows that the behaviour of the model is consistent across a range of hyperparameter settings around the one we used.

In [Figure S2](#), we illustrate the output of the model for several hyperparameter settings. One can notice that for very small values of α_{effort} the reach is successful, but executed with larger torques and velocity than is necessary – e.g the red and yellow reaches are equally successful but the red one is much faster – which comes at the cost of larger inputs. We chose the set of hyperparameters for our simulations such as to lie in an intermediate regime in which the task is solved successfully, but without requiring more inputs than necessary.

Figure S1

Correlates of the behaviour and control strategy across a wide range of hyperparameters. The “reach success” and “holding success” are set to 1 if the success criterion (see text) is satisfied and 0 otherwise. The task is executed successfully over a wide range of hyperparameters. The red star denotes the set of hyperparameters used in the main text simulations. This configuration was chosen to lie in a region in which the task can be successfully solved, with the performance being robust to small changes in the hyperparameters.





Figure S2

Illustration of the behaviour for several hyperparameter settings. (Left) Hand position along the horizontal axis, with the dotted line denoting the position of the target. (Middle) Temporal profile of the hand velocity. (Right) Temporal profile of the torques driving the hand.



S2 Investigation of the effect of the network decay timescale

Figure 3 highlighted that preparatory inputs tend to consistently arise late during the delay period. We hypothesized that this may be a reflection of the intrinsic tendency of the network dynamics to decay, such that inputs given too early may be “lost”. To test this, we changed the characteristic timescale of the dynamics *during preparation only*, leading to the following dynamics:

$$\tau_t \frac{d\mathbf{x}(t)}{dt} = -\mathbf{x}(t) + \mathbf{W}\phi[\mathbf{x}(t)] + \mathbf{h} + \frac{\tau_t}{\tau_{\text{mov}}} \mathbf{u}(t) \text{ where } \begin{cases} \tau_t = \tau_{\text{prep}} & \text{if } t \leq 0 \\ \tau_t = \tau_{\text{mov}} & \text{if } t \geq 0 \end{cases} \quad (\text{S1})$$

with $\tau_{\text{mov}} = 150\text{ms}$. This allowed us to evaluate whether having dynamics decaying more slowly during preparation led to inputs starting earlier. Note that we also rescaled the inputs during preparation by $\frac{\tau_{\text{prep}}}{\tau_{\text{mov}}}$, to ensure that the effective cost of the inputs was not affected by the timescale change.

As shown in **Figure S3**, inputs started rising earlier when the network’s decay timescale was longer. This was consistent with the hypothesis that the length of the window of preparation that the optimal controller uses depends on the network’s intrinsic timescale.




Figure S3

Illustration of the effect of the characteristic neuronal timescale on the temporal distribution of the inputs.

We modified the characteristic neuronal timescale of the ISN during preparation only and assessed how that changed the temporal distribution of inputs for 3 different timescales ($\tau_{\text{prep}} = 50\text{ms}$, $\tau_{\text{prep}} = 150\text{ms}$, $\tau_{\text{prep}} = 300\text{ms}$, top to bottom). As hypothesized, inputs start earlier during the preparation window when the decay timescale of the network was longer.

S3 Additional results in the 2D system

Our visualization of the behaviour of 2D networks in **Figure 4** [↗](#) allowed us to identify features of the dynamics that were well-suited to predicting preparation. Below, we compute α and β numerically and analytically in 2D oscillatory and nonnormal networks, to gain insights into how these quantities vary with the networks' dynamics. We then show how preparation can be predicted highly accurately across a large number of 2D systems, using only those quantities to summarize the network dynamics.

S3.1 Controllability and observability computations

In **Figure S4** [↗](#), we computed α and β numerically, as a function of the connectivity strength and the choice of readout, for the nonnormal and the oscillatory motifs shown in **Figure 4** [↗](#). This highlights the very different behaviours of the two networks, which are to some extent also reflected in higher-dimensional models. In particular, we find a strong effect of the alignment between the readout and the network dynamics in nonnormal networks, while α and β are independent of θ_C in oscillatory networks. Interestingly, we see that β is constant across all oscillatory networks, while α increases with w .

As the reduced 2D model is more amenable to mathematical analysis than its high-dimensional counterpart, we can gain further insights into the origin of these differences by computing $\alpha(w, \theta_C)$ and $\beta(w, \theta_C)$ analytically.




Figure S4

Illustration of α and β as a function of θ_C and w in the 2D networks.

Recall that the observability Gramian Q of a linear input-driven dynamical system satisfies

$$A^T Q + Q A + C^T C = 0 \quad (S2)$$

and the controllability Gramian satisfies

$$A P + P A^T + B B^T = 0, \quad (S3)$$

and that we defined $\alpha = \text{Tr}(C^\perp Q C^{\perp T})$ and $\beta = \text{Tr}(C P C^T)$ where $C^\perp$ denotes the nullspace of the readout matrix. Below, we compute these quantities for the 2D oscillatory and nonnormal networks, with $B = I$ and C a unit-norm vector whose direction we parametrize via a quantity θ_C . Note that we ignore the effect of dt and τ in the mathematical analysis, as those quantities can straightforwardly be included in the final result via a rescaling of w and B .

Oscillatory network

In the case of $A = -I + S$ where S is a skew-symmetric network (i.e. $S^T = -S$), Equation (S3) is solved by $P = I/2$ independently of the value of S . This explains why $\beta = \text{Tr}(C P C^T) = \frac{1}{2} \|C\|^2$ is independent of both the connectivity strength w and the orientation of the readout θ_C for skew-symmetric networks (see [Figure S4](#); bottom right). Practically, this means that skew-symmetric networks are equally controllable in all directions: when driven by random inputs, these networks display isotropic activity of equal variance along all directions. Moreover, as w controls the oscillation frequency of the network, but does not change the decay timescale of the eigenmodes, the amount of variance generated by a random stimulation is independent of w . Interestingly, we can see in [Figure S4](#) (top right) that α displays a different behaviour, and increases with w . As highlighted above, skew-symmetric systems are rotationally symmetric. Without loss of generality, we can thus define our 1D vector to read out the first unit, i.e. $C = [1 \ 0]$.

The observability Gramian must satisfy

$$A^T Q + Q A + C^T C = 0 \implies \begin{bmatrix} -1 & w \\ -w & -1 \end{bmatrix} Q + Q \begin{bmatrix} -1 & -w \\ w & -1 \end{bmatrix} = \begin{bmatrix} -1 & 0 \\ 0 & 0 \end{bmatrix}. \quad (S4)$$

This can be found in closed-form by solving the 2D system of equations, yielding

$$Q = \begin{bmatrix} \frac{1}{4} + \frac{1}{4(1+w^2)} & -\frac{w}{4(1+w^2)} \\ -\frac{w}{4(1+w^2)} & \frac{w^2}{4(1+w^2)} \end{bmatrix}. \quad (S5)$$

From there, we obtain $\alpha = \text{Tr}(C^\perp Q C^{\perp T}) = \frac{w^2}{4(1+w^2)}$. As found empirically, this quantity will initially increase before plateauing towards $1/4$ as w becomes large.

One might wonder why observability displays such a dependence on the oscillatory frequency of the network, even though the network is rotationally symmetric, and w does not affect the decay timescale. As highlighted in Equation (S2), controllability and observability Gramian would be identical for a skew-symmetric system if $C = I$. However, a feature of the systems we consider is the existence of a nullspace, i.e. the fact that the readout C only targets a subset of dimensions across the whole space (implying that $C^T C$ is a low-rank matrix). Intuitively, the reason why α increases with w while β is constant in skew-symmetric networks can be understood as follows: α is computing how much *readout activity* a vector initialized in the nullspace of C will generate, while β is computing the amount of energy that will be generated *across all directions* by a vector initialized in the readout space. Thus, assuming once again $C = [1 \ 0]$ and $C^\perp = [0 \ 1]$, the activity of vectors initialized along C and $C^\perp$ respectively and evolving autonomously from there is given by $v_C(t) = [e^{-t} \cos(wt) \ e^{-t} \sin(wt)]$ and $v_{C^\perp}(t) = [-e^{-t} \sin(wt) \ e^{-t} \cos(wt)]$.

From there, we can compute $\beta = \int_0^\infty \|\mathbf{v}_C(t)\|^2 dt = \int_0^\infty e^{-2t} dt = \frac{1}{2}$. Thus, as found above, only the decay timescale of the envelope (fixed to 1 here) affects the value of β .

Importantly, α will instead have a dependence on w arising from the fact that it depends on the size of the activity *projected into the readout*, as

$$\alpha = \int_0^\infty \|\mathbf{v}_{C^\perp}(t)^T \mathbf{C}\|^2 dt = \int_0^\infty e^{-2t} \sin^2(wt) dt \quad (\text{S6})$$

$$= \frac{1}{2} \int_0^\infty e^{-2t} (1 - \cos(2wt)) dt \quad (\text{S7})$$

$$= \frac{1}{2} \int_0^\infty e^{-2t} - \Re e^{-2(1-iw)t} dt \quad (\text{S8})$$

$$= \frac{1}{4} - \frac{1}{4(1+w^2)} \quad (\text{S9})$$

$$= \frac{w^2}{4(1+w^2)}. \quad (\text{S10})$$

The dependence of this quantity on w can be understood by the fact that activity patterns initialized in the readout nullspace benefit from the existence of rotational dynamics, which allows them to be read-out before the activity decays completely.

Nonnormal network

In the nonnormal network, we have $\mathbf{A} = -\mathbf{I} + \mathbf{W} = \begin{bmatrix} -1 & 0 \\ w & -1 \end{bmatrix}$. The nonnormal 2D system, unlike its oscillatory counterpart, does not have rotational symmetry. Thus, to remain general, we will consider $\mathbf{C}(\theta_C) = \cos \theta_C \sin \theta_C$, and $\mathbf{C}^\perp(\theta_C) = [-\sin \theta_C \cos \theta_C]$. Solving Equation (S3) for $\mathbf{B} = \mathbf{I}$ leads to an expression for the controllability Gramian of the nonnormal system as

$$\mathbf{P}(w) = \begin{bmatrix} \frac{1}{2} & \frac{w}{4} \\ \frac{w}{4} & \frac{1}{2} + \frac{w^2}{4} \end{bmatrix}. \quad (\text{S11})$$

Similarly, computation of the observability Gramian leads to

$$\mathbf{Q}(w, \theta_C) = \begin{bmatrix} \frac{w^2 \sin^2 \theta_C}{4} + \frac{\cos \theta_C \sin \theta_C w}{2} + \frac{\cos^2 \theta_C}{2} & \frac{w \sin^2 \theta_C}{4} + \frac{\cos \theta_C \sin \theta_C}{2} \\ \frac{w \sin^2 \theta_C}{4} + \frac{\cos \theta_C \sin \theta_C}{2} & \frac{\sin^2 \theta_C}{2} \end{bmatrix}. \quad (\text{S12})$$

We can then compute

$$\alpha(\theta_C, w) = \mathbf{C}^{\perp T} \mathbf{Q} \mathbf{C}^\perp = \frac{w^2}{4} \sin^4 \theta_C \quad (\text{S13})$$

and

$$\beta(\theta_C, w) = \mathbf{C}^T \mathbf{P} \mathbf{C} = \frac{(w \sin \theta_C + \cos \theta_C)^2 - \cos^2 \theta_C + 2}{4}. \quad (\text{S14})$$

This highlights the dependence of α and β on θ_C , which can also be seen in [Figure S4](#) (left). Interestingly, these expressions also make evident the supralinear scaling of α and β with w in nonnormal networks. Note however that we never investigate preparation in the very large w

regime, as the simulation of such networks with discretized dynamics is prone to numerical issues.

S3.2 Predicting preparation in 2D networks

To assess how well preparation could be predicted from the control-theoretic properties α and β (c.f. main text) of 2D networks, we generated 20000 networks with weight matrix

$$\mathbf{W}(a, \omega, w_{\text{ff}}) = \begin{bmatrix} a & \frac{1}{2}(w_{\text{ff}} + \sqrt{w_{\text{ff}}^2 + 4\omega^2}) \\ \frac{1}{2}(w_{\text{ff}} - \sqrt{w_{\text{ff}}^2 + 4\omega^2}) & a \end{bmatrix} \quad (\text{S15})$$

where $a \sim \mathcal{U}(0, 0.8)$, $\omega \sim \mathcal{U}(0, 4)$, and $w_{\text{ff}} \sim \mathcal{U}(0, 4)$. Equation (S15) implies that $\mathbf{W}$ has a pair of complex-conjugate eigenvalues $a \pm i\omega$, and also embeds a feedforward coupling of strength w_{ff} from the second to the first dimension. For each network configuration, we computed the corresponding values of α and β . To confirm our intuition that the preparation index should increase with α and decrease with β , we first attempted to fit $\text{prep. index} = c_0 + c_1 \frac{\alpha}{\beta}$. Interestingly, we found that while this quantity was positively correlated with the preparation index across networks, a substantial fraction of variance remained unexplained (test $R^2 = 0.16$). Labelling the preparation index by the rotational frequency of the network highlighted that a substantial fraction of the variance across networks came from this timescale of oscillations (**Figure S5** [↗](#), left). Indeed, a regression model of the form $\text{prep. index} = c_0 + c_1 \omega \frac{\alpha}{\beta}$ captured 80% of the variance in preparation index, yielding an accurate fit across networks with only two free parameters (**Figure S5** [↗](#), right).

We stress that the predictive power of these simple fits is remarkable given that the preparation index comes out of a complex process of optimization over control inputs. Thus, the control-theoretic quantities α and β appear to appropriately summarize the benefits of preparation for individual networks.

The fact that the preparation index also grows with ω can be understood by considering the alignment between the activity trajectories which the network can autonomously generate and those that are required for solving the motor task. Indeed, a network that is intrinsically unable to generate outputs with the right oscillatory timescale would have to rely on movement-related inputs, i.e. would have a low preparation index. As observed here, the network's characteristic frequency has a big impact in 2D networks, consistent with ω determining the *only* oscillatory pattern that the network can generate on its own. For high-dimensional networks, however, we did not have to incorporate such a measure of compatibility between task requirements and network dynamics (c.f. **Figure 5** [↗](#)). We speculate that this is due to averaging effects. Indeed, larger networks possess a wide range of intrinsic oscillatory timescales, and the readout matrix – which here was not aligned to the network's dynamics in any specific way – is expected to read out a little bit of all frequencies, including task-appropriate ones.

Figure S5

Predicting the preparation index from characteristic network quantities.

We evaluated how well the preparation index could be predicted as a linear function of $\frac{\alpha}{\beta}$ (left). A substantial amount of residual variance appeared to arise from variability in the oscillation frequency ω (color). Accounting for this frequency by regressing the preparation index against $\omega \frac{\alpha}{\beta}$ gave a better fit (right).





Figure S6

Preparation arises across a range of network architectures :

neural correlates of the reach are shown for 5 different networks (A), alongside the loss and prep. index as a function of Δ_{prep} . (A) Eigenvalue spectrum (top), internal network activations (middle) and inputs (bottom) for different network types. The unconnected network does not rely on preparatory inputs at all. The random network with weights draw from $\mathcal{N}(0, 0.95/\sqrt{N})$ uses very little delay period inputs while the skew-symmetric network with $\mathcal{N}(0, 4/\sqrt{N})$ shows a substantial amount of inputs during the delay period. The inhibition-stabilized network can be seen to rely most on preparation, more so than the similarity transformed ISN. (B) Loss (top) and preparation index (bottom) as a function of delay period length for the different networks. The unconnected and random networks can be seen to benefit very little from longer preparation times. Indeed, even as Δ_{prep} increases, their amount of preparatory inputs remains very close to 0. On the other hand, the skew-symmetric network and the ISN use preparatory inputs (bottom), which allow them to have a lower loss for larger values of Δ_{prep} . Interestingly, the surrogate ISN prepares considerably less than the full ISN.



S4 Comparison across networks

Our main investigation was largely focused on on behaviour of inhibition-stabilized networks, which are believed to constitute good models of M1. We however found that the expression we derived to obtain a network's preparation index from its control-theoretic properties generalized across to other types of networks. Below, we detail the other network families we considered, and show how their dynamics qualitatively differ from the ISN, although their preparation can be predicted using the same quantities.

We modelled three additional classes of networks: randomly connected networks with either (i) unstructured or (ii) skew-symmetric connectivities, (iii) a surrogate network obtained by applying a similarity transformation to the ISN that preserved its eigenvalue spectrum but eliminated any “nonnormality” (i.e. we found $\mathcal{T}$ such that $\tilde{\mathbf{A}} = \mathcal{T}^{-1}\mathbf{A}\mathcal{T}$ where $\tilde{\mathbf{A}}$ was a diagonal matrix with the same eigenvalues as $\mathbf{A}$). Note that we did not apply the transformation to the readout or input matrices, such that the transfer function of the system was changed by our transformation. This was voluntary, as we were interested in the effect that transforming the dynamics would have on the input-output response. These networks were chosen for the diversity of dynamical motifs they exhibit: combinations of rapidly and slowly decaying modes, oscillations, and transient dynamics. Moreover, each of these network families could be sampled from in a straightforward manner, allowing to compute results across many instantiations of each network type. We again used random readout matrices not specifically adjusted to the dynamics of the network nor to the motor task. To get an intuition for how different networks solve the task, we generated one network from each family and qualitatively compared their inputs and internal activations when performing the same delayed reach ([Figure S6A](#)). We first considered an unconnected network, i.e. a network whose recurrent weights were all 0. Unsurprisingly, this network had no use for a preparation phase. Indeed, there is no benefit to giving early inputs as the network is unable to amplify them. More surprisingly, a random network with a much stronger connectivity – as can be seen in its eigenvalue spectrum forming a small ball of radius close to 1 ([Figure S6A](#) (top)) – also displayed very little preparation. The strong, visually apparent similarity between the inputs to the random and unconnected networks suggests that the optimal way of controlling the random network relies largely on ignoring its internal dynamics and solving the task almost entirely in an input-driven regime. The example skew-symmetric network, which had imaginary eigenvalues only (ranging between -5.5 and 5.5), displayed considerably more preparation, but still relied on strong inputs during the movement phase that resembled those of the unconnected and random networks. Finally, the ISN relied much more on preparation; the small inputs it receives are strongly amplified into large activity patterns owing to its strong, nonnormal recurrent connectivity. Interestingly however, the similarity transformed ISN lost much of that ability to amplify inputs, instead displaying dynamics resembling that of the skew-symmetric network. This highlights the effect of the ISN's nonnormal dynamics in shaping the network's activity and optimal inputs.

Next, we assessed more directly how beneficial preparation was for the different networks. We evaluated how the total loss and preparation index evolved as a function of the delay period length ([Figure S6B](#)). As expected, the control of networks that relied on preparation (skew-symmetric and ISN) benefitted more from longer delays. The ISN has markedly lower control cost and higher preparation index than other networks, reflecting the fact that even weak (thus energetically cheap) inputs were sufficient to produce internal activity and thus output torques of the desired magnitude (c.f. [Figure S6A](#), right).

The above results give a sense of the range of possible dynamics that different types of networks display. Interestingly, despite these differences, we showed in [Figure 5](#) that the preparation index could be predicted with a simple formula across all networks.




Figure S7

Illustration of the effect of optimizing the readout matrix such as to minimize the cost of the reaches, across all movements.

To evaluate the effect that our choice of random readout directions has on our conclusions, we additionally compare to a model with the same dynamics, but where the readout was optimized such as to minimize the cost across movements (i.e. $\mathcal{L}(\mathbf{C}) = \sum_{j \in \text{targets}} \mathcal{J}^{(j)}(\mathbf{C})$), under the constraint that its norm was fixed. In **(A)**, we see that this leads to an increase in the observability of the system (compare the observability of the modes of the optimized system in black with those of the random readout in red). In **(B)** and **(C)**, we see that this leads to an output of similar amplitude **(B)**, but that is generated using smaller inputs **(C)**. Importantly, we see that the system still relies on preparatory inputs. Thus, the exact choice of readout does not alter the network strategy, but can help the system perform the same movements in a more efficient manner.




Figure S8

Comparison of the effect of penalizing temporal input smoothness versus input norm.

We compare the effect of using a cost over inputs that penalizes input norm, versus using a cost that penalizes the “temporal complexity” of the inputs – defined here as the temporal derivative of the inputs (i.e. $\|u(t+1) - u(t)\|^2$ in discrete time). This is achieved by augmenting the dynamical system to include an input integration stage, which then feeds into the original dynamical system; this way, the input to the augmented system – of which we continue to penalize the squared norm to enable the iLQR framework – is the derivative of the input to the original system. We perform this comparison in linear RNNs, across a range of different preparation times. We show the activation of an example neuron in **(A)** and **(B)**, and activity in an example input channel in **(C)** and **(D)**. Each color denotes a different reach. We see that the rates vary more slowly when penalizing the temporal complexity **(A)** versus the input norm **(B)**, exhibiting a plateau for longer preparation times that is more similar to neural recordings. This is a reflection of the fact that the inputs themselves vary more slowly when the temporal complexity is penalized (compare **C** and **D**). As we do not penalize the input norm within our definition of temporal complexity, the optimal strategy is for the network to rely on steady inputs, which is different from the strategy used when the norm is penalized (compare **C** and **D**). We note that, under this different choice of input penalty, preparation nevertheless remains optimal, with the normalized loss (shown in **E**) decreasing as Δ_{prep} increases, and the preparation index (shown in **F**) increasing as Δ_{prep} increases.




Figure S9

Control strategy adopted by the model when planning with uncertain delays.

We investigate an extension of the model, that includes uncertainty about the arrival of the Go cue, and involves replanning at regular intervals to update the model with new information (e.g whether the Go cue has arrived). In **(A)**, we model this by assuming that the network is adopting a strategy whereby it plans to be ready to move *as early as possible* following the target onset. We optimize the inputs using this assumption, and replan every 20ms to update the model with the available information (which here corresponds to the actual Go cue only arriving 500ms after Target onset). We plot the activity of two example neurons (left and right panels, respectively), for each of the reaches (each color denotes a different reach). We can see that the neuronal activations start ramping up at the beginning of the task, and plateau before the actual target onset. In **(B)**, we use a similar optimization strategy, but use a different “mental model” for the network, whereby we assume that, until it sees the actual Go cue, the model is always assuming that the delay period will be equal to the most likely a posteriori preparation time. Under the assumption of exponentially distributed delays with a mean of 150ms, this corresponds to always replanning assuming a delay of 150ms. We see that the network then adopts a different strategy, which does not include ramping/plateauing of the neural activity.




Figure S10

Comparison of the occupancy of the preparatory and movement subspaces across different delay periods.

Occupancy (normalized by the maximum value across preparatory and movement occupancies) of the preparatory and movement subspaces identified using a delay period of 500ms, for the activity generated using $\Delta_{\text{prep}} = 0\text{ms}$ (left), $\Delta_{\text{prep}} = 100\text{ms}$ (center), and $\Delta_{\text{prep}} = 500\text{ms}$ (right). We see that the network does not rely on preparatory activity when $\Delta_{\text{prep}} = 0\text{ms}$.

References

- Afshar A., Santhanam G., Byron M. Y., Ryu S. I., Sahani M., Shenoy K. V. (2011) **Single-trial neural correlates of arm movement preparation** *Neuron* **71**:555–564
- Amos B., Jimenez I., Sacks J., Boots B., Kolter J. Z. (2018) **Differentiable mpc for end-to-end planning and control** *Advances in Neural Information Processing Systems* :8289–8300
- Brunel N. (2000) **Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons** *J. Comput. Neurosci* **8**:183–208
- Churchland M. M., Cunningham J. P., Kaufman M. T., Foster J. D., Nuyujukian P., Ryu S. I., Shenoy K. V. (2012) **Neural population dynamics during reaching** *Nature* **487**:51–56
- Churchland M. M., Cunningham J. P., Kaufman M. T., Ryu S. I., Shenoy K. V. (2010) **Cortical preparatory activity: representation of movement or first cog in a dynamical machine?** *Neuron* **68**:387–400
- Churchland M. M., Shenoy K. V. (2007) **Delay of movement caused by disruption of cortical preparatory activity** *Journal of neurophysiology* **97**:348–359
- Codol O., Michaels J. A., Kashefi M., Pruszynski J. A., Gribble P. L. (2023) **Codol, O., Michaels, J. A., Kashefi, M., Pruszynski, J. A., and Gribble, P. L. (2023). Motornet: a python toolbox for controlling differentiable biomechanical effectors with artificial neural networks.**
- Dayan P., Abbott L. F. (2001) **Theoretical neuroscience**
- Driscoll L., Shenoy K., Sussillo D. (2022) **Flexible multitask computation in recurrent networks utilizes shared dynamical motifs** *bioRxiv*
- Dubreuil A., Valente A., Beiran M., Mastrogiuseppe F., Ostojic S. (2021) **The role of population structure in computations through neural dynamics** *bioRxiv* :2020–7
- Dum R. P., Strick P. L. (1991) **The origin of corti-cospinal projections from the premotor areas in the frontal lobe** *Journal of Neuroscience* **11**:667–689
- Elsayed G. F., Lara A. H., Kaufman M. T., Churchland M. M., Cunningham J. P. (2016) **Reorganization between preparatory and movement population responses in motor cortex** *Nature communications* **7**:1–15
- Goldman M. S. (2009) **Memory without feedback in a neural network** *Neuron* **61**:621–634
- Guo Z. V., Li N., Huber D., Ophir E., Gutnisky D., Ting J. T., Feng G., Svoboda K. (2014) **Flow of cortical activity underlying a tactile decision in mice** *Neuron* **81**:179–194
- Haith A. M., Pakpoor J., Krakauer J. W. (2016) **Independence of movement preparation and movement initiation** *Journal of Neuroscience* **36**:3007–3015
- Harris C. M., Wolpert D. M. (1998) **Signal-dependent noise determines motor planning** *Nature* **394**:780–784

- Heald J., Wolpert D., Lengyel M. (2023) **The computational and neural bases of context-dependent learning** *Ann. Rev. Neurosci.* :1–27
- Hennequin G., Vogels T. P., Gerstner W. (2012) **Nonnormal amplification in random balanced neuronal networks** *Physical Review E* **86**
- Hennequin G., Vogels T. P., Gerstner W. (2014) **Optimal control of transient dynamics in balanced networks supports generation of complex movements** *Neuron* **82**:1394–1406
- 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**
- Kalman R. E., et al. (1960) **Contributions to the theory of optimal control** *Bol. soc. mat. mexicana* **5**:102–119
- Kao T.-C., Hennequin G. (2019) **Neuroscience out of control: control-theoretic perspectives on neural circuit dynamics** *Cur. Op. Neurobiol* **58**:122–129
- Kao T.-C., Sadabadi M. S., Hennequin G. (2021) **Optimal anticipatory control as a theory of motor preparation: a thalamo-cortical circuit model** *Neuron* **109**
- 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**:440–448
- Kaufman M. T., Seely J. S., Sussillo D., Ryu S. I., Shenoy K. V., Churchland M. M. (2016) **The largest response component in the motor cortex reflects movement timing but not movement type** *Eneuro* **3**
- Lara A. H., Elsayed G. F., Zimnik A. J., Cunningham J. P., Churchland M. M. (2018) **Conservation of preparatory neural events in monkey motor cortex regardless of how movement is initiated** *Elife* **7**
- Li N., Chen T.-W., Guo Z. V., Gerfen C. R., Svoboda K. (2015) **A motor cortex circuit for motor planning and movement** *Nature* **519**:51–56
- Li W., Todorov E. (2004) **Iterative linear quadratic regulator design for nonlinear biological movement systems** *ICINCO (1)* :222–229
- Logiaco L., Abbott L., Escola S. (2021) **Thalamic control of cortical dynamics in a model of flexible motor sequencing** *Cell Reports* **35**
- Malonis P. J., Hatsopoulos N. G., MacLean J. N., Kaufman M. T. (2021) **M1 dynamics share similar inputs for initiating and correcting movement** *bioRxiv*
- Meirhaeghe N., Riehle A., Brochier T. (2023) **Parallel movement planning is achieved via an optimal preparatory state in motor cortex** *Cell Reports* **42**
- Michaels J. A., Dann B., Intveld R. W., Scherberger H. (2015) **Predicting reaction time from the neural state space of the premotor and parietal grasping network** *Journal of Neuroscience* **35**:11415–11432
- 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**:683–696

- Murphy B. K., Miller K. D. (2009) **Balanced amplification: a new mechanism of selective amplification of neural activity patterns** *Neuron* **61**:635–648
- Pandarinath C. *et al.* (2018) **Inferring single-trial neural population dynamics using sequential auto-encoders** *Nature methods* **15**:805–815
- Prut Y., Fetz E. E. (1999) **Primate spinal interneurons show pre-movement instructed delay activity** *Nature* **401**:590–594
- Rawlings J. B., Mayne D. Q., Diehl M. (2017) **Model predictive control: theory, computation, and design**
- Riehle A., Requin J. (1989) **Monkey primary motor and premotor cortex: single-cell activity related to prior information about direction and extent of an intended movement** *Journal of neurophysiology* **61**:534–549
- Russo A. A. *et al.* (2018) **Motor cortex embeds muscle-like commands in an untangled population response** *Neuron* **97**:953–966
- Sauerbrei B. A., Guo J.-Z., Cohen J. D., Mischiati M., Guo W., Kabra M., Verma N., Mensh B., Branson K., Hantman A. W. (2020) **Cortical pattern generation during dexterous movement is input-driven** *Nature* **577**:386–391
- Schimmel M., Kao T.-C., Jensen K. T., Hennequin G. (2022) **iLQR-VAE : control-based learning of input-driven dynamics with applications to neural data** *In International Conference on Learning Representations*
- Sheahan H. R., Franklin D. W., Wolpert D. M. (2016) **Motor planning, not execution, separates motor memories** *Neuron* **92**:773–779
- Shenoy K. V., Sahani M., Churchland M. M. (2013) **Cortical control of arm movements: a dynamical systems perspective** *Annual review of neuroscience* **36**
- Skogestad S., Postlethwaite I. (2007) **Multivariable feedback control: analysis and design**
- Sohn H., Meirhaeghe N., Rajalingham R., Jazayeri M. (2020) **A network perspective on sensorimotor learning** *Trends in Neurosciences*
- Soldado-Magraner J., Mante V., Sahani M. (2023) **Inferring context-dependent computations through linear approximations of prefrontal cortex dynamics** *bioRxiv*
- Sterling P., Laughlin S. (2015) **Principles of neural design**
- Stroud J. P., Porter M. A., Hennequin G., Vogels T. P. (2018) **Motor primitives in space and time via targeted gain modulation in cortical networks** *Nature neuroscience* **21**:1774–1783
- Sun X., O’Shea D. J., Golub M. D., Trautmann E. M., Vyas S., Ryu S. I., Shenoy K. V. (2022) **Cortical preparatory activity indexes learned motor memories** *Nature* **602**:1–6
- 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**:1025–1033
- Tassa Y. (2011) **Theory and Implementation of Biomimetic Motor Controllers**

Todorov E., Jordan M. I. (2002) **Optimal feedback control as a theory of motor coordination** *Nature neuro-science* **5**:1226–1235

Todorov E., Li W. (2003) **Optimal control methods suitable for biomechanical systems** *In Proceedings of the 25th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (IEEE Cat. No. 03CH37439)* :1758–1761

Yeo S.-H., Franklin D. W., Wolpert D. M. (2016) **When optimal feedback control is not enough: Feedforward strategies are required for optimal control with active sensing** *PLoS computational biology* **12**

Zimnik A. J., Churchland M. M. (2021) **Independent generation of sequence elements by motor cortex** *Nature neuroscience* **24**:412–424

Editors

Reviewing Editor

Jörn Diedrichsen

Western University, London, Canada

Senior Editor

Floris de Lange

Donders Institute for Brain, Cognition and Behaviour, Nijmegen, Netherlands

Reviewer #1 (Public Review):

In this work, the authors investigate an important question - under what circumstances should a recurrent neural network optimised to produce motor control signals receive preparatory input before the initiation of a movement, even though it is possible to use inputs to drive activity just-in-time for movement?

This question is important because many studies across animal models have shown that preparatory activity is widespread in neural populations close to motor output (e.g. motor cortex / M1), but it isn't clear under what circumstances this preparation is advantageous for performance, especially since preparation could cause unwanted motor output during a delay.

They show that networks optimised under reasonable constraints (speed, accuracy, lack of pre-movement) will use the input to seed the state of the network before movement and that these inputs reduce the need for ongoing input during the movement. By examining many different parameters in simplified models they identify a strong connection between the structure of the network and the amount of preparation that is optimal for control - namely, that preparation has the most value when nullspaces are highly observable relative to the readout dimension and when the controllability of readout dimensions is low. They conclude by showing that their model predictions are consistent with the observation in monkey motor cortex that even when a sequence of two movements is known in advance, preparatory activity only arises shortly before movement initiation.

Overall, this study provides valuable theoretical insight into the role of preparation in neural populations that generate motor output, and by treating input to motor cortex as a signal that is optimised directly this work is able to sidestep many of the problematic questions relating to estimating the potential inputs to motor cortex.

Reviewer #2 (Public Review):

This work clarifies neural mechanisms that can lead to a phenomenology consistent with motor preparation in its broader sense. In this context, motor preparation refers to activity that occurs before the corresponding movement. Another property often associated with preparatory activity is a correlation with global movement characteristics such as reach speed (Churchland et al., *Neuron* 2006), reach angle (Sun et al., *Nature* 2022), or grasp type (Meirhaeghe et al., *Cell Reports* 2023). Such activity has notably been observed in premotor and primary motor cortices, and it has been hypothesized to serve as an input to a motor execution circuit. The timing and mechanisms by which such 'preparatory' inputs are made available to motor execution circuits remain however unclear in general, especially in light of the presence of a 'trigger-like' signal that appears to relate to the transition from preparatory dynamics to execution activity (Kaufman et al. *eNeuron* 2016, Iganaki et al., *Cell* 2022, Zimnik and Churchland, *Nature Neuroscience* 2021).

The preparatory inputs have been hypothesized to fulfill one or several (non-mutually-exclusive) possible objectives. Two notable hypotheses are that these inputs could be shaped to maximize output accuracy under regularization of the input magnitude; or that they may help the flexible re-use of the neural machinery involved in the control of movements in different contexts.

Here, the authors investigate in detail how the former hypothesis may be compatible with the presence of early inputs in recurrent network models driving arm movements, and compare models to data.

Strengths:

The authors are able to deploy an in-depth evaluation of inputs that are optimized for producing an accurate output at a pre-defined time while using a regularization term on the input magnitude, in the case of movements that are thought to be controlled in a quasi-open loop fashion such as reaches.

First, the authors have identified that optimal control theory is a great framework to study this question as it provides methods to find and analyze exact solutions to this cost function in the case of models with linear dynamics. The authors not only use this framework to get an exact assessment of how much pre-movement input arises in large recurrent networks, but also give insight into the mechanisms by which it happens by dissecting in detail low-dimensional networks. The authors find that two key network properties - observability of the readout's nullspace and limited controllability - give rise to optimal inputs that are large before the start of the movement (while the corresponding network activity lies in the nullspace of the readout). Further, the authors numerically investigate the timing of optimized inputs in models with nonlinear dynamics, and find that pre-movement inputs can also arise in these more general networks. The authors also explore how some variations on their model's constraints - such as penalizing the input roughness or changing task contingencies about the go cue timing - affect their results. Finally, the authors point out some coarse-grained similarities between the pre-movement activity driven by the optimized inputs in some of the models they studied, and the phenomenology of preparation observed in the brain during single reaches and reach sequences. Overall, the authors deploy an impressive arsenal of tools and a very in-depth analysis of their models.

Outstanding questions that could lead to interesting follow-up work:

Like all great pieces of research, this article makes it clear where current limitations lie and therefore opens up opportunities for future work.

(1) Though the optimal control theory framework is ideal for determining inputs that minimize output error while regularizing the input norm or other simple input features, it cannot easily account for some other varied types of objectives - especially those that may lead to a complex optimization landscape. For instance, the reusability of parts of the circuit, sparse use of additional neurons when learning many movements, and ease of planning (especially under uncertainty about when to start the movement), may be alternative or additional reasons that could help explain the preparatory activity observed in the brain. It is interesting to note that inputs that optimize the objective chosen by the authors arguably lead to a trade-off in terms of other desirable objectives. Specifically, the inputs the authors derive are time-dependent, so a recurrent network would be needed to produce them and it may not be easy to interpolate between them to drive new movement variants. In addition, these inputs depend on the desired time of output and therefore make it difficult to plan, e.g. in circumstances when timing should be decided depending on sensory signals. Finally, these inputs are specific to the full movement chain that will unfold, so they do not permit reuse of the inputs e.g. in movement sequences of different orders. Of note, the authors have pointed out in the discussion how their framework may be extended in future work to account for some additional objectives, such as inputs' temporal smoothness or some strategies for dealing with go cue timing uncertainty.

(2) Relatedly, if the motor circuits were to balance different types of objectives, the activity and inputs occurring before each movement may be broken down into different categories that may each specialize into their own objective. For instance, previous work (Kaufman et al. eNeuron 2016, Iganaki et al., Cell 2022, Zimnik and Churchland, Nature Neuroscience 2021) has suggested that inputs occurring before the movement could be broken down into preparatory inputs 'stricto sensu' - relating to the planned characteristics of the movement - and a trigger signal, relating to the transition from planning to execution - irrespective of whether the movement is internally timed or triggered by an external event. The current work does not address which type(s) of early input may be labeled as 'preparatory' or may be thought of as a part of 'planning' computations, or whether these inputs may come from several different source circuits. Future research could investigate these questions using a different approach, for instance, by including structural constraints from brain architecture into a neural network model.

(3) While the authors rightly point out some similarities between the inputs that they derive and observed preparatory activity in the brain, notably during motor sequences, there are also some differences. For instance, while both the derived inputs and the data show two peaks during sequences, the data reproduced from Zimnik and Churchland show preparatory inputs that have a very asymmetric shape that really plummets before the start of the next movement, whereas the derived inputs have larger amplitude during the movement period - especially for the second movement of the sequence. In addition, the data show trigger-like signals before each of the two reaches. Finally, while the data show a very high correlation between the pattern of preparatory activity of the second reach in the double reach and compound reach conditions, the derived inputs appear to be more different between the two conditions. Note that the data would be consistent with separate planning of the two reaches even in the compound reach condition, as well as the re-use of the preparatory input between the compound and double reach conditions. Therefore, different motor sequence datasets - notably, those that would show even more coarticulation between submovements - may be more promising for finding a tight match between the data and the author's inputs. In the future, further analyses in these datasets could help determine whether the coarticulation could be due to simple filtering by the circuits and muscles downstream of M1, planning of movements with adjusted curvature to mitigate the work performed by the muscles while

permitting some amount of re-use across different sequences, or - as suggested by the authors - inputs fully tailored to one specific movement sequence that maximize accuracy and minimize the M1 input magnitude.

(4) Though iLQR is a powerful optimization method to find inputs optimizing the author's cost function, it also has some limitations. First, given that it relies on a linearization of the dynamics at each timestep, it has a limited ability to leverage potential advantages of nonlinearities in the dynamics. Second, the iLQR algorithm is not a biologically plausible learning rule and does not account for biological constraints affecting the circuits that produce and process these inputs. Therefore, it might be difficult for the brain to learn to produce the inputs that it finds. Consequently, when observing differences between model and data, this can confound the question of whether it comes from a difference of assumed objective or a difference of optimization procedure or circuit implementation. It remains unclear whether using alternative algorithms with different limitations - for instance, using variants of BPTT to train a separate RNN to produce the inputs in question - could impact some of the results.

(5) Under the objective considered by the authors, the amount of input occurring before the movement might be impacted by the presence of online sensory signals for closed-loop control. Even if the inputs include some sensory activity and/or the RNN activity could represent all general variables (e.g. sensory) whose states can be decoded from M1, the model does not currently include mechanisms that process imperfect (delayed, noisy) sensory feedback to adapt the output in a trial-specific manner. The information related to such sensory feedback cannot be anticipated, and therefore the related input would have to reach the motor cortex after preparation. Thus, it is an open question whether the objective and network characteristics suggested by the authors could also explain the presence of large preparatory activity before e.g. grasping movements that are thought to be more sensory-feedback-driven (Meirhaeghe et al., Cell Reports 2023).

(6) More broadly, with the type of objectives that the authors assume the inputs fulfill, some M1 properties that lead to strong preparation - notably, limited readout controllability - may not be favorable for control in general, so it would be interesting if other objectives and assumptions could robustly lead to strong preparation under more general M1 properties.'

<https://doi.org/10.7554/eLife.89131.3.sa2>

Reviewer #3 (Public Review):

I remain enthusiastic about this study. The manuscript is well-written, logical, and conceptually clear. To my knowledge, no prior modeling study has tackled the question of 'why prepare before executing, why not just execute?' Prior studies have simply assumed, to emulate empirical findings, that preparatory inputs precede execution. They never asked why. The authors show that, when there are constraints on inputs, preparation becomes a natural strategy. In contrast, with no constraint on inputs, there is no need for preparation as one could get anything one liked just via the inputs during movement. For the sake of tractability, the authors use a simple magnitude constraint: the cost function punishes the integral of the squared inputs. Thus, if small inputs before movement can reduce the size of the inputs needed during movement, preparation is a good strategy. This occurs if (and only if) the network has strong dynamics (otherwise feeding it preparatory activity would not produce anything interesting). All of this is sensible and clarifying.

As discussed in the prior round of reviews, the central constraint that the authors use is a mathematically tractable stand-in for a range of plausible (but often trickier to define and evaluate) constraints, such as simplicity of inputs (or inputs being things that other areas could provide). The manuscript now embraces this fact more explicitly and also gives some

results showing that other constraints (such as on the derivative of activity, which is one component of complexity) can have the same effect. The manuscript also now discusses and addresses a modest weakness of the previous manuscript: the preparatory activity in their simulations is often overly complex temporally, lacking the (rough) plateau typically seen for data. Depending on your point of view, this is simply 'window dressing', but from my perspective it was important to know that their approach could yield more realistic-looking preparatory activity.

The most recent version of the manuscript also has a useful section in the Discussion on the topic of preparation when there is no external delay, which I found helpful given prior behavioral and physiological studies arguing that preparation can 1) be very brief, but 2) is always present. These findings mesh nicely with the authors' central result that preparation is a good network strategy, and that it would thus be normative for there to be at least a brief interval of preparation even when not imposed externally.

<https://doi.org/10.7554/eLife.89131.3.sa1>

Author response:

The following is the authors' response to the previous reviews.

Reviewer 1:

Thank you for your review and pointing out multiple things to be discussed and clarified! Below, we go through the various limitations you pointed out and refer to the places where we have tried to address them.

(1) It's important to keep in mind that this work involves simplified models of the motor system, and often the terminology for 'motor cortex' and 'models of motor cortex' are used interchangeably, which may mislead some readers. Similarly, the introduction fails in many cases to state what model system is being discussed (e.g. line 14, line 29, line 31), even though these span humans, monkeys, mice, and simulations, which all differ in crucial ways that cannot always be lumped together.

That is a good point. We have clarified this in the text (Introduction and Discussion), to highlight the fact that our model isn't necessarily meant to just capture M1. We have also updated the introduction to make it more clear which species the experiments which motivate our investigation were performed in.

(2) At multiple points in the manuscript thalamic inputs during movement (in mice) is used as a motivation for examining the role of preparation. However, there are other more salient motivations, such as delayed sensory feedback from the limb and vision arriving in the motor cortex, as well as ongoing control signals from other areas such as the premotor cortex.

Yes – the motivation for thalamic inputs came from the fact that those have specifically been shown to be necessary for accurate movement generation in mice. However, it is true that the inputs in our model are meant to capture any signals external to the dynamical system modeled, and as such are likely to represent a mixture of sensory signals, and feedback from other areas. We have clarified this in the Discussion, and have added this additional motivation in the Introduction.

(3) Describing the main task in this work as a delayed reaching task is not justified without caveats (by the authors' own admission: line 687), since each network is

optimized with a fixed delay period length. Although this is mentioned to the reader, it's not clear enough that the dynamics observed during the delay period will not resemble those in the motor cortex for typical delayed reaching tasks.

Yes, we completely agree that the terminology might be confusing. While the task we are modeling is a delayed reaching task, it does differ from the usual setting since the network has knowledge of the delay period, and that is indeed a caveat of the model. We have added a brief paragraph just after the description of the optimal control objective to highlight this limitation.

We have also performed additional simulations using two different variants of a model-predictive control approach that allow us to relax the assumption that the go-cue time is known in advance. We show that these modifications of the optimal controller yield results that remain consistent with our main conclusions, and can in fact in some settings lead to preparatory activity plateaus during the preparation epoch as often found in monkey M1 (e.g. in Elsayed et al. 2016). We have modified the Discussion to explain these results and their limitations, which are summarized in a new Supplementary Figure (S9).

(4) A number of simplifications in the model may have crucial consequences for interpretation.

a) Even following the toy examples in Figure 4, all the models in Figure 5 are linear, which may limit the generalisability of the findings.

While we agree that linear models may be too simplistic, much prior analyses of M1 data suggest that it is often good enough to capture key aspects of M1 dynamics; for example, the generative model underlying jPCA is linear, and Sussillo et al. (2015) showed that the internal activity of nonlinear RNN models trained to reproduce EMG data aligned best with M1 activity when heavily regularized; in this regime, the RNN dynamics were close to linear. Nevertheless, this linearity assumption is indeed convenient from a modeling viewpoint: the optimal control problem is more easily solved for linear network dynamics and the optimal trajectories are more consistent across networks. Indeed, we had originally attempted to perform the analyses of Figure 5 in the nonlinear setting, but found that while the results were overall similar to what we report in the linear regime, iLQR was occasionally trapped into local minima, resulting in more variable results especially for inhibition-stabilized network in the strongly connected end of the spectrum. Finally, Figure 5 is primarily meant to explore to what extent motor preparation can be predicted from basic linear control-theoretic properties of the Jacobian of the dynamics; in this regard, it made sense to work with linear RNNs (for which the Jacobian is constant).

b) Crucially, there is no delayed sensory feedback in the model from the plant. Although this simplification is in some ways a strength, this decision allows networks to avoid having to deal with delayed feedback, which is a known component of closed-loop motor control and of motor cortex inputs and will have a large impact on the control policy.

This comment resonates well with Reviewer 3's remark regarding the autonomous nature (or not) of M1 during movement. Rather than thinking of our RNN models as anatomically confined models of M1 alone, we think of them as models of the dynamics which M1 implements possibly as part of a broader network involving "inter-area loops and (at some latency) sensory feedback", and whose state appears to be near-fully decodable from M1 activity alone. We have added a paragraph of Discussion on this important point.

(5) A key feature determining the usefulness of preparation is the direction of the readout dimension. However, all readouts had a similar structure (random Gaussian initialization). Therefore, it would be useful to have more discussion regarding how the

structure of the output connectivity would affect preparation, since the motor cortex certainly does not follow this output scheme.

We agree with this limitation of our model — indeed one key message of Figure 4 is that the degree of reliance on preparatory inputs depends strongly on how the dynamics align with the readout. However, this strong dependence is somewhat specific to low-dimensional models; in higher-dimensional models (most of our paper), one expects that any random readout matrix C will pick out activity dimensions in the RNN that are sufficiently aligned with the most controllable directions of the dynamics to encourage preparation.

We did consider optimizing C away (which required differentiating through the iLQR optimizer, which is possible but very costly), but the question inevitably arises what exactly should C be optimized for, and under what constraints (e.g. fixed norm or not). One possibility is to optimize C with respect to the same control objective that the control inputs are optimized for, and constrain its norm (otherwise, inputs to the M1 model, and its internal activity, could become arbitrarily small as C can grow to compensate). We performed this experiment (new Supplementary Figure S7) and obtained a similar preparation index; there was one notable difference, namely that the optimized readout modes led to greater observability compared to a random readout; thus, the same amount of “muscle energy” required for a given movement could now be produced by a smaller initial condition. In turn, this led to smaller control inputs, consistent with a lower control cost overall.

Whilst we could have systematically optimized C away, we reasoned that (i) it is computationally expensive, and (ii) the way M1 affects downstream effectors is presumably “optimized” for much richer motor tasks than simple 2D reaching, such that optimizing C for a fixed set of simple reaches could lead to misleading conclusions. We therefore decided to stick with random readouts.

Additional comments:

(1) The choice of cost function seems very important. Is it? For example, penalising the square of $u(t)$ may produce very different results than penalising the absolute value.

Yes, the choice of cost function does affect the results, at least qualitatively. The absolute value of the inputs is a challenging cost to use, as iLQR relies on a local quadratic approximation of the cost function. However, we have included additional experiments in which we penalized the squared derivative of the inputs (Supplementary Figure S8; see also our response to Reviewer 3's suggestion on this topic), and we do see differences in the qualitative behavior of the model (though the main takeaway, i.e. the reliance on preparation, continues to hold). This is now referred to and discussed in the Discussion section.

(2) In future work it would be useful to consider the role of spinal networks, which are known to contribute to preparation in some cases (e.g. Prut and Fetz, 1999).

(3) The control signal magnitude is penalised, but not the output torque magnitude, which highlights the fact that control in the model is quite different from muscle control, where co-contraction would be a possibility and therefore a penalty of muscle activation would be necessary. Future work should consider the role of these differences in control policy.

Thank you for pointing us to this reference! Regarding both of these concerns, we agree that the model could be greatly improved and made more realistic in future work (another avenue for this would be to consider a more realistic biophysical model, e.g. using the MotorNet library). We hope that the current Discussion, which highlights the various

limitations of our modeling choices, makes it clear that a lot of these choices could easily be modified depending on the specific assumptions/investigation being performed.

Reviewer 2:

Thank you for your positive review! We very much agree with the limitations you pointed out, some of which overlapped with the comments of the other reviewers. We have done our best to address them through additional discussion and new supplementary figures. We briefly highlight below where those changes can be found.

(1) Though the optimal control theory framework is ideal to determine inputs that minimize output error while regularizing the input norm, it however cannot easily account for some other varied types of objectives especially those that may lead to a complex optimization landscape. For instance, the reusability of parts of the circuit, sparse use of additional neurons when learning many movements, and ease of planning (especially under uncertainty about when to start the movement), may be alternative or additional reasons that could help explain the preparatory activity observed in the brain. It is interesting to note that inputs that optimize the objective chosen by the authors arguably lead to a trade-off in terms of other desirable objectives. Specifically, the inputs the authors derive are time-dependent, so a recurrent network would be needed to produce them and it may not be easy to interpolate between them to drive new movement variants. In addition, these inputs depend on the desired time of output and therefore make it difficult to plan, e.g. in circumstances when timing should be decided depending on sensory signals. Finally, these inputs are specific to the full movement chain that will unfold, so they do not permit reuse of the inputs e.g. in movement sequences of different orders.

Yes, that is a good point! We have incorporated further Discussion related to this point. We have additionally included a new example in which we regularize the temporal complexity of the inputs (see also our response to Reviewer 3's suggestion on this topic), which leads to more slowly varying inputs, and may indeed represent a more realistic constraint and lead to simpler inputs that can more easily be interpolated between. We also agree that uncertainty about the upcoming go cue may play an important role in the strategy adopted by the animals. While we have not performed an extensive investigation of the topic, we have included a Supplementary Figure (S9) in which we used Model Predictive Control to investigate the effect of planning under uncertainty about the go cue arrival time. We hope that this will give the reader a better sense of what sort of model extensions are possible within our framework.

(2) Relatedly, if the motor circuits were to balance different types of objectives, the activity and inputs occurring before each movement may be broken down into different categories that may each specialize into one objective. For instance, previous work (Kaufman et al. eNeuron 2016, Iganaki et al., Cell 2022, Zimnik and Churchland, Nature Neuroscience 2021) has suggested that inputs occurring before the movement could be broken down into preparatory inputs 'stricto sensu' - relating to the planned characteristics of the movement - and a trigger signal, relating to the transition from planning to execution - irrespective of whether the movement is internally timed or triggered by an external event. The current work does not address which type(s) of early input may be labeled as 'preparatory' or may be thought of as a part of 'planning' computations.

Yes, our model does indeed treat inputs in a very general way, and does not distinguish between the different types of processes they may be composed of. This is partly because we do not explicitly model where the inputs come from, such that our inputs likely englobe multiple processes. We have added discussion related to this point.

(3) While the authors rightly point out some similarities between the inputs that they derive and observed preparatory activity in the brain, notably during motor sequences, there are also some differences. For instance, while both the derived inputs and the data show two peaks during sequences, the data reproduced from Zimnik and Churchland show preparatory inputs that have a very asymmetric shape that really plummets before the start of the next movement, whereas the derived inputs have larger amplitude during the movement period - especially for the second movement of the sequence. In addition, the data show trigger-like signals before each of the two reaches. Finally, while the data show a very high correlation between the pattern of preparatory activity of the second reach in the double reach and compound reach conditions, the derived inputs appear to be more different between the two conditions. Note that the data would be consistent with separate planning of the two reaches even in the compound reach condition, as well as the re-use of the preparatory input between the compound and double reach conditions. Therefore, different motor sequence datasets - notably, those that would show even more coarticulation between submovements - may be more promising to find a tight match between the data and the author's inputs. Further analyses in these datasets could help determine whether the coarticulation could be due to simple filtering by the circuits and muscles downstream of M1, planning of movements with adjusted curvature to mitigate the work performed by the muscles while permitting some amount of re-use across different sequences, or - as suggested by the authors - inputs fully tailored to one specific movement sequence that maximize accuracy and minimize the M1 input magnitude.

Regarding the exact shape of the occupancy plots, it is important to note that some of the more qualitative aspects (e.g the relative height of the two peaks) will change if we change the parameters of the cost function. Right now, we have chosen the parameters to ensure that both reaches would be performed at roughly the same speed (as a way to very loosely constrain the parameters based on the observed behavior). However, small changes to the hyperparameters can lead to changes in the model output (e.g one of the two consecutive reaches being performed using greater acceleration than the other), and since our biophysical model is fairly simple, changes in the behavior are directly reflected in the network activity. Essentially, what this means is that while the double occupancy is a consistent feature of the model, the exact shape of the peaks is more sensitive to hyperparameters, and we do not wish to draw any strong conclusions from them, given the simplicity of the biophysical model. However, we do agree that our model exhibits some differences with the data. As discussed above, we have included additional discussion regarding the potential existence of separate inputs for planning vs triggering the movement in the context of single reaches.

Overall, we are excited about the suggestions made by the Reviewer here about using our approach to analyze other motor sequence datasets, but we think that in order to do this properly, one would need to adopt a more realistic musculo-skeletal model (such as one provided by MotorNet).

(4) Though iLQR is a powerful optimization method to find inputs optimizing the author's cost function, it also has some limitations. First, given that it relies on a linearization of the dynamics at each timestep, it has a limited ability to leverage potential advantages of nonlinearities in the dynamics. Second, the iLQR algorithm is not a biologically plausible learning rule and therefore it might be difficult for the brain to learn to produce the inputs that it finds. It remains unclear whether using alternative algorithms with different limitations - for instance, using variants of BPTT to train a separate RNN to produce the inputs in question - could impact some of the results.

We agree that our choice of iLQR has limitations: while it offers the advantage of convergence guarantees, it does indeed restrict the choice of cost function and dynamics that we can use. We have now included extensive discussion of how the modeling choices affect our results.

We do not view the lack of biological plausibility of iLQR as an issue, as the results are agnostic to the algorithm used for optimization. However, we agree that any structure imposed on the inputs (e.g by enforcing them to be the output of a self-contained dynamical system) would likely alter the results. A potentially interesting extension of our model would be to do just what the reviewer suggested, and try to learn a network that can generate the optimal inputs. However, this is outside the scope of our investigation, as it would then lead to new questions (e.g what brain region would that other RNN represent?).

(5) Under the objective considered by the authors, the amount of input occurring before the movement might be impacted by the presence of online sensory signals for closed-loop control. It is therefore an open question whether the objective and network characteristics suggested by the authors could also explain the presence of preparatory activity before e.g. grasping movements that are thought to be more sensory-driven (Meirhaeghe et al., Cell Reports 2023).

It is true that we aren't currently modeling sensory signals explicitly. However, some of the optimal inputs we infer may be capturing upstream information which could englobe some sensory information. This is currently unclear, and would likely depend on how exactly the model is specified. We have added new discussion to emphasize that our dynamics should not be understood as just representing M1, but more general circuits whose state can be decoded from M1.

Reviewer #2 (Recommendations For The Authors):

Additionally, thank you for pointing out various typos in the manuscript, we have fixed those!

Reviewer 3:

Thank you very much for your review, which makes a lot of very insightful points, and raises several interesting questions. In summary, we very much agree with the limitations you pointed out. In particular, the choice of input cost is something we had previously discussed, but we had found it challenging to decide on what a reasonable cost for “complexity” could be. Following your comment, we have however added a first attempt at penalizing “temporal complexity”, which shows promising behavior. We have only included those additional analyses as supplementary figures, and we have included new discussion, which hopefully highlights what we meant by the different model components, and how the model behavior may change as we vary some of our choices. We hope this can be informative for future models that may use a similar approach. Below, we highlight the changes that we have made to address your comments.

The main limitation of the study is that it focuses exclusively on one specific constraint - magnitude - that could limit motor-cortex inputs. This isn't unreasonable, but other constraints are at least as likely, if less mathematically tractable. The basic results of this study will probably be robust with regard such issues - generally speaking, any constraint on what can be delivered during execution will favor the strategy of preparing - but this robustness cuts both ways. It isn't clear that the constraint used in the present study - minimizing upstream energy costs - is the one that really matters. Upstream areas are likely to be limited in a variety of ways, including the complexity of inputs they can deliver. Indeed, one generally assumes that there are things that motor cortex can do that upstream areas can't do, which is where the real limitations should come from.

Yet in the interest of a tractable cost function, the authors have built a system where motor cortex actually doesn't do anything that couldn't be done equally well by its inputs. The system might actually be better off if motor cortex were removed. About the only thing that motor cortex appears to contribute is some amplification, which is 'good' from the standpoint of the cost function (inputs can be smaller) but hardly satisfying from a scientific standpoint.

The use of a term that punishes the squared magnitude of control signals has a long history, both because it creates mathematical tractability and because it (somewhat) maps onto the idea that one should minimize the energy expended by muscles and the possibility of damaging them with large inputs. One could make a case that those things apply to neural activity as well, and while that isn't unreasonable, it is far from clear whether this is actually true (and if it were, why punish the square if you are concerned about ATP expenditure?). Even if neural activity magnitude an important cost, any costs should pertain not just to inputs but to motor cortex activity itself. I don't think the authors really wish to propose that squared input magnitude is the key thing to be regularized. Instead, this is simply an easily imposed constraint that is tractable and acts as a stand-in for other forms of regularization / other types of constraints. Put differently, if one could write down the 'true' cost function, it might contain a term related to squared magnitude, but other regularizing terms would be very likely to dominate. Using only squared magnitude is a reasonable way to get started, but there are also ways in which it appears to be limiting the results (see below).

I would suggest that the study explore this topic a bit. Is it possible to use other forms of regularization? One appealing option is to constrain the complexity of inputs; a long-standing idea is that the role of motor cortex is to take relatively simple inputs and convert them to complex time-evolving inputs suitable for driving outputs. I realize that exploring this idea is not necessarily trivial. The right cost-function term is not clear (should it relate to low-dimensionality across conditions, or to smoothness across time?) and even if it were, it might not produce a convex cost function. Yet while exploring this possibility might be difficult, I think it is important for two reasons.

First, this study is an elegant exploration of how preparation emerges due to constraints on inputs, but at present that exploration focuses exclusively on one constraint. Second, at present there are a variety of aspects of the model responses that appear somewhat unrealistic. I suspect most of these flow from the fact that while the magnitude of inputs is constrained, their complexity is not (they can control every motor cortex neuron at both low and high frequencies). Because inputs are not complexity-constrained, preparatory activity appears overly complex and never 'settles' into the plateaus that one often sees in data. To be fair, even in data these plateaus are often imperfect, but they are still a very noticeable feature in the response of many neurons. Furthermore, the top PCs usually contain a nice plateau. Yet we never get to see this in the present study. In part this is because the authors never simulate the situation of an unpredictable delay (more on this below) but it also seems to be because preparatory inputs are themselves strongly time-varying. More realistic forms of regularization would likely remedy this.

That is a very good point, and it mirrors several concerns that we had in the past. While we did focus on the input norm for the sake of simplicity, and because it represents a very natural way to regularize our control solutions, we agree that a “complexity cost” may be better suited to models of brain circuits. We have addressed this in a supplementary investigation. We chose to focus on a cost that penalizes the temporal complexity of the inputs, as $\|u(t+1) - u(t)\|^2$. Note that this required augmenting the state of the model, making the computations quite a bit slower; while it is doable if we only penalize the first temporal derivative, it would not scale well to higher orders.

Interestingly, we did find that the activity in that setting was somewhat more realistic (see new Supplementary Figure S8), with more sustained inputs and plateauing activity. While we have kept the original model for most of the investigations, the somewhat more realistic nature of the results under that setting suggests that further exploration of penalties of that sort could represent a promising avenue to improve the model.

We also found the idea of a cost that would ensure low-dimensionality of the inputs across conditions very interesting. However, it is challenging to investigate with iLQR as we perform the optimization separately for each condition; nevertheless, it could be investigated using a different optimizer.

At present, it is also not clear whether preparation always occurs even with no delay. Given only magnitude-based regularization, it wouldn't necessarily have to be. The authors should perform a subspace-based analysis like that in Figure 6, but for different delay durations. I think it is critical to explore whether the model, like monkeys, uses preparation even for zero-delay trials. At present it might or might not. If not, it may be because of the lack of more realistic constraints on inputs. One might then either need to include more realistic constraints to induce zero-delay preparation, or propose that the brain basically never uses a zero delay (it always delays the internal go cue after the preparatory inputs) and that this is a mechanism separate from that being modeled.

I agree with the authors that the present version of the model, where optimization knows the exact time of movement onset, produces a reasonably realistic timecourse of preparation when compared to data from self-paced movements. At the same time, most readers will want to see that the model can produce realistic looking preparatory activity when presented with an unpredictable delay. I realize this may be an optimization nightmare, but there are probably ways to trick the model into optimizing to move soon, but then forcing it to wait (which is actually what monkeys are probably doing). Doing so would allow the model to produce preparation under the circumstances where most studies have examined it. In some ways this is just window-dressing (showing people something in a format they are used to and can digest) but it is actually more than that, because it would show that the model can produce a reasonable plateau of sustained preparation. At present it isn't clear it can do this, for the reasons noted above. If it can't, regularizing complexity might help (and even if this can't be shown, it could be discussed).

In summary, I found this to be a very strong study overall, with a conceptually timely message that was well-explained and nicely documented by thorough simulations. I think it is critical to perform the test, noted above, of examining preparatory subspace activity across a range of delay durations (including zero) to see whether preparation endures as it does empirically. I think the issue of a more realistic cost function is also important, both in terms of the conceptual message and in terms of inducing the model to produce more realistic activity. Conceptually it matters because I don't think the central message should be 'preparation reduces upstream ATP usage by allowing motor cortex to be an amplifier'. I think the central message the authors wish to convey is that constraints on inputs make preparation a good strategy. Many of those constraints likely relate to the fact that upstream areas can't do things that motor cortex can do (else you wouldn't need a motor cortex) and it would be good if regularization reflected that assumption. Furthermore, additional forms of regularization would likely improve the realism of model responses, in ways that matter both aesthetically and conceptually. Yet while I think this is an important issue, it is also a deep and tricky one, and I think the authors need considerable leeway in how they address it. Many of the cost-function terms one might want to use may be intractable. The authors may have to do what makes sense given technical limitations. If some things can't be done technically, they

may need to be addressed in words or via some other sort of non-optimization-based simulation.

Specific comments

As noted above, it would be good to show that preparatory subspace activity occurs similarly across delay durations. It actually might not, at present. For a zero ms delay, the simple magnitude-based regularization may be insufficient to induce preparation. If so, then the authors would either have to argue that a zero delay is actually never used internally (which is a reasonable argument) or show that other forms of regularization can induce zero-delay preparation.

Yes, that is a very interesting analysis to perform, which we had not considered before! When investigating this, we found that the zero-delay strategy does not rely on preparation in the same way as is seen in the monkeys. This seems to be a reflection of the fact that our “Go cue” corresponds to an “internal” go cue which would likely come after the true, “external go cue” – such that we would indeed never actually be in the zero delay setting. This is not something we had addressed (or really considered) before, although we had tried to ensure we referred to “delta prep” as the duration of the preparatory period but not necessarily the delay period. We have now included more discussion on this topic, as well as a new Supplementary Figure S10.

I agree with the authors that prior modeling work was limited by assuming the inputs to M1, which meant that prior work couldn't address the deep issue (tackled here) of why there should be any preparatory inputs at all. At the same time, the ability to hand-select inputs did provide some advantages. A strong assumption of prior work is that the inputs are 'simple', such that motor cortex must perform meaningful computations to convert them to outputs. This matters because if inputs can be anything, then they can just be the final outputs themselves, and motor cortex would have no job to do. Thus, prior work tried to assume the simplest inputs possible to motor cortex that could still explain the data. Most likely this went too far in the 'simple' direction, yet aspects of the simplicity were important for endowing responses with realistic properties. One such property is a large condition-invariant response just before movement onset. This is a very robust aspect of the data, and is explained by the assumption of a simple trigger signal that conveys information about when to move but is otherwise invariant to condition. Note that this is an implicit form of regularization, and one very different from that used in the present study: the input is allowed to be large, but constrained to be simple. Preparatory inputs are similarly constrained to be simple in the sense that they carry only information about which condition should be executed, but otherwise have little temporal structure. Arguably this produces slightly too simple preparatory-period responses, but the present study appears to go too far in the opposite direction. I would suggest that the authors do what they can to address these issue via simulations and/or discussion. I think it is fine if the conclusion is that there exist many constraints that tend to favor preparation, and that regularizing magnitude is just one easy way of demonstrating that. Ideally, other constraints would be explored. But even if they can't be, there should be some discussion of what is missing - preparatory plateaus, a realistic condition-invariant signal tied to movement onset - under the present modeling assumptions.

As described above, we have now included two additional figures. In the first one (S8, already discussed above), we used a temporal smoothness prior, and we indeed get slightly more realistic activity plateaus. In a second supplementary figure (S9), we have also considered using model predictive control (MPC) to optimize the inputs under an uncertain go cue arrival time. There, we found that removing the assumption that the delay period is known came with new challenges: in particular, it requires the specification of a “mental model” of

when the Go cue will arrive. While it is reasonable to expect that monkeys will have a prior over the go time arrival cue that will be shaped by the design of the experiment, some assumptions must be made about the utility functions that should be used to weigh this prior. For instance, if we imagine that monkeys carry a model of the possible arrival time of the go cue that is updated online, they could nonetheless act differently based on this information, for instance by either preparing so as to be ready for the earliest go cue possible or alternatively to be ready for the average go cue. This will likely depend on the exact task design and reward/penalty structure. Here, we added simulations with those two cases (making simplifying assumptions to make the problem tractable/solvable using model predictive control), and found that the “earliest preparation” strategy gives rise to more realistic plateauing activity, while the model where planning is done for the “most likely go time” does not. We suspect that more realistic activity patterns could be obtained by e.g. combining this framework with the temporal smoothness cost. However, the main point we wished to make with this new supplementary figure is that it is possible to model the task in a slightly more realistic way (although here it comes at the cost of additional model assumptions). We have now added more discussion related to those points. Note that we have kept our analyses on these new models to a minimum, as the main takeaway we wish to convey from them is that most components of the model could be modified/made more realistic. This would impact the qualitative behavior of the system and match to data but – in the examples we have so far considered – does not appear to modify the general strategy of networks relying on preparation.

On line 161, and in a few other places, the authors cite prior work as arguing for "autonomous internal dynamics in M1". I think it is worth being careful here because most of that work specifically stated that the dynamics are likely not internal to M1, and presumably involve inter-area loops and (at some latency) sensory feedback. The real claim of such work is that one can observe most of the key state variables in M1, such that there are periods of time where the dynamics are reasonably approximated as autonomous from a mathematical standpoint. This means that you can estimate the state from M1, and then there is some function that predicts the future state. This formal definition of autonomous shouldn't be conflated with an anatomical definition.

Yes, that is a good point, thank you for making it so clearly! Indeed, as previous work, we do not think of our “M1 dynamics” as being internal to M1, but they may instead include sensory feedback / inter-area loops, which we summarize into the connectivity, that we chose to have dynamics that qualitatively resemble data. We have now incorporated more discussion regarding what exactly the dynamics in our model represent.

Round 2 of reviews

Reviewer 3:

My remaining comments largely pertain to some subtle (but to me important) nuances at a few locations in the text. These should be easy for the authors to address, in whatever way they see fit.

Specific comments:

(1) The authors state the following on line 56: "For preparatory processes to avoid triggering premature movement, any pre-movement activity in the motor and dorsal pre-motor (PMd) cortices must carefully exclude those pyramidal tract neurons."

This constraint is overly restrictive. PT neurons absolutely can change their activity during preparation in principle (and appear to do so in practice). The key constraint is looser: those changes should have no net effect on the muscles. E.g., if d is the vector of changes in PT neuron firing rates, and b is the vector of weights, then the constraint is

that $b'd = 0$. $d = 0$ is one good way of doing this, but only one. Half the d 's could go up and half could go down. Or they all go up, but half the b 's are negative. Put differently, there is no reason the null space has to be upstream of the PT neurons. It could be partly, or entirely, downstream. In the end, this doesn't change the point the authors are making. It is still the case that d has to be structured to avoid causing muscle activity, which raises exactly the point the authors care about: why risk this unless preparation brings benefits? However, this point can be made with a more accurate motivation. This matters, because people often think that a null-space is a tricky thing to engineer, when really it is quite natural. With enough neurons, preparing in the null space is quite simple.

That is a good point – we have now reformulated this sentence to instead say “to avoid triggering premature movement, any pre-movement activity in the motor and dorsal premotor (PMd) cortices must engage the pyramidal tract neurons in a way that ensures their activity patterns will not lead to any movement”.

(2) Line 167: 'near-autonomous internal dynamics in M1'.

It would be good if such statements, early in the paper, could be modified to reflect the fact that the dynamics observed in M1 may depend on recurrence that is NOT purely internal to M1. A better phrase might be 'near-autonomous dynamics that can be observed in M1'. A similar point applies on line 13. This issue is handled very thoughtfully in the Discussion, starting on line 713. Obviously it is not sensible to also add multiple sentences making the same point early on. However, it is still worth phrasing things carefully, otherwise the reader may have the wrong impression up until the Discussion (i.e. they may think that both the authors, and prior studies, believe that all the relevant dynamics are internal to M1). If possible, it might also be worth adding one sentence, somewhere early, to keep readers from falling into this hole (and then being stuck there till the Discussion digs them out).

That is a good point: we have now edited the text after line 170 to make it clear that the underlying dynamics may not be confined to M1, and have referenced the later discussion there.

(3) The authors make the point, starting on line 815, that transient (but strong) preparatory activity empirically occurs without a delay. They note that their model will do this but only if 'no delay' means 'no external delay'. For their model to prepare, there still needs to be an internal delay between when the first inputs arrive and when movement generating inputs arrive.

This is not only a reasonable assumption, but is something that does indeed occur empirically. This can be seen in Figure 8c of Lara et al. Similarly, Kaufman et al. 2016 noted that "the sudden change in the CIS [the movement triggering event] occurred well after (~150 ms) the visual go cue... (~60 ms latency)" Behavioral experiments have also argued that internal movement-triggering events tend to be quite sluggish relative to the earliest they could be, causing RTs to be longer than they should be (Haith et al. Independence of Movement Preparation and Movement Initiation). Given this empirical support, the authors might wish to add a sentence indicating that the data tend to justify their assumption that the internal delay (separating the earliest response to sensory events from the events that actually cause movement to begin) never shrinks to zero.

While on this topic, the Haith and Krakauer paper mentioned above good to cite because it does ponder the question of whether preparation is really necessary. By showing that they could get RTs to shrink considerably before behavior became inaccurate, they showed that people normally (when not pressured) use more preparation time than they

really need. Given Lara et al, we know that preparation does always occur, but Haith and Krakauer were quite right that it can be very brief. This helped -- along with neural results -- change our view of preparation from something more cognitive that had to occur, so something more mechanical that was simply a good network strategy, which is indeed the authors current point. Working a discussion of this into the current paper may or may not make sense, but if there is a place where it is easy to cite, it would be appropriate.

This is a nice suggestion, and we thank the reviewer for pointing us to the Haith and Krakauer paper. We have now added this reference and extended the paragraph following line 815 to briefly discuss the possible decoupling between preparation and movement initiation that is shown in the Haith paper, emphasizing how this may affect the interpretation of the internal delay and comparisons with behavioral experiments.

<https://doi.org/10.7554/eLife.89131.3.sa0>