Neural Population Mechanisms for Flexible Sensorimotor Control

Hari Kalidindi[1,2], Frédéric Crevecoeur[1,2,3]

[1] Institute of Communication Technologies, Engineering and Applied Mathematics (ICTEAM), UCLouvain, Avenue Georges Lemaitre 4-6, 1348 Louvain-la-Neuve, Belgium

[2] Institute of Neuroscience, UCLouvain, Avenue E. Mounier 53, 1200 Brussels, Belgium

[3] WEL Research Institute, Wavre, Belgium

Corresponding author: frederic.crevecoeur@uclouvain.be

ABSTRACT

Modern large-scale recordings have revealed that motor cortex activity during reaching follows low-dimensional dynamics, thought to reflect sensorimotor computations. However, the origin of these patterns, and how they flexibly reorganize across different tasks remain unclear. Here we demonstrate that the key features of neural activity naturally emerge in a linear model combining a random network with a biomechanical system. Remarkably, this model shows how a fixed network can flexibly control different behaviors by optimally mapping sensory and internal state onto task-specific network inputs. Finally, analytical decomposition of the controller reveals that low-dimensional network dynamics directly follows from propagation of low-dimensional feedback signals from the biomechanical plant through the network. These results provide a biologically plausible mechanism for flexible motor control in the nervous system which directly links neural population dynamics to behavior.

Introduction

A central challenge in neuroscience is to explain how neural activity generates behavior and adapts to changing environments[1]. To address this, recent advances have shifted focus from individual neurons to the collective dynamics of populations of neurons[2–4]. Empirical studies show that the population dynamics can be effectively captured by rather low-dimensional, time-varying signals, referred to as *neural factors*[5], or *latent dynamical states*[6,7]. These factors have been interpreted as signals that the brain must generate to orchestrate behavior[5,8]. Yet, fundamental questions remain: what are the biological or physical mechanisms that give rise to these factors, and how many factors are required for a given behavior?

Prevailing theories propose that neural factors emerge from recurrent connectivity among cortical and subcortical areas, forming attractors that shape population activity[7]. Artificial recurrent neural networks (RNN) capture this idea by developing stable patterns – such as tonic or multi-phasic activity – generating muscle commands when optimized on sensorimotor tasks[9–17]. However, these models rely on extensive training, with task-specific dynamics emerging only after exposure to relevant data[18]. Their behavior is often difficult to interpret due to non-linearities and dependency on both training data[18], and training algorithms[19]. In contrast, primates exhibit remarkable flexibility as they can rapidly adapt to novel contexts, or even update control within an ongoing movement[20–24]. To explain this flexibility, recent studies suggest that tonic, contextual inputs to the network can evoke specific low-dimensional patterns suited to a given task[9,25–33]. However, it remains unclear how these inputs are computed and adjusted across tasks (e.g., reaching versus shooting to a target, or lifting a weighted object).

Here we propose a unifying theoretical framework that explains both the emergence of low-dimensional factors in motor cortical dynamics, and how inputs to the network should be computed to accomplish the task at hand. We developed a linear model coupling a random network to a biomechanical limb and simulated various tasks using stochastic optimal control[32,34,35]. In this framework, the feedback control policy maps the feedback about the state of the whole system (network and limb) onto optimal network activations. We demonstrate that this approach directly links the parameters of the controller to the network activity across tasks without the need for example-driven (re)training. Remarkably, our model reproduced key features observed in motor cortex, including preparatory activity anticipating the upcoming movement[36,37], rotational dynamics

during movement[8,38], and orthogonal null and potent subspaces[39]. Furthermore, we show that this linear formulation offers a novel view on the origin of low-dimensional factors using the low-rank hypothesis of complex systems[40]: although the network connectivity is high-dimensional, its activity is mainly driven by inputs carrying sensory feedback and task parameters, yielding low-dimensional patterns of activation. Together, our results formalize the link between sensorimotor control and low-dimensional population activity for a fixed random network, offering a powerful basis to interpret the neural basis of flexible motor control.

Results

We constructed a super-system composed of a randomly connected recurrent network ($n = 100$ nodes) coupled to a limb (Fig. 1a). The network activity, $\boldsymbol{r}(t) \in \mathbb{R}^{100}$, was represented as a standard rate equation[9,12,28,32]. The network output controlled one of the two limb models: a point-mass moving in a cartesian plane or a two-link planar arm. Both models featured two-dimensional actuation, with cartesian forces applied along the plane for the point-mass, and joint torques applied to the shoulder and elbow joints for the planar arm. The limb state, $\boldsymbol{q}(t) \in \mathbb{R}^{\boldsymbol{6}}$, encompassed position, velocity and actuation forces/torques. Full details of the network and limb dynamics are provided in the Methods. For compact representation, the linear dynamics governing the network-body system were jointly described by:

$$\begin{bmatrix} \dot{\boldsymbol{r}} \\ \dot{\boldsymbol{q}} \end{bmatrix} = \begin{bmatrix} J & 0 \\ C & M \end{bmatrix} \begin{bmatrix} \boldsymbol{r} \\ \boldsymbol{q} \end{bmatrix} + \begin{bmatrix} \boldsymbol{\varepsilon} \\ 0 \end{bmatrix} + \begin{bmatrix} I \\ 0 \end{bmatrix} \boldsymbol{u}, \tag{1}$$

where $J \in \mathbb{R}^{100\times100}$ was the random state transition matrix, $\boldsymbol{\varepsilon} \in \mathbb{R}^{100}$ was a constant input vector inducing heterogenous spontaneous activity, C was the random linear readout that mapped the network activity to limb actuation, and M was the state transition matrix of limb dynamics (the explicit dependency on time was omitted from Eqn. 1 for clarity). The control input $\boldsymbol{u}(t) \in \mathbb{R}^{100}$ influenced only the network activity. The identity matrix I denotes that each control input component activated a single network unit.

The objective was to find inputs $\boldsymbol{u}(t)$ that steer the network activity to make the limb move optimally as defined by a cost-function. We adopted the formalism of receding-horizon to allow modelling changes in task instructions such as the unexpected occurrence of a "GO" cue presented during movement preparation (see Methods). This problem was solved as follows: we determined the control inputs that minimized the expected cost ($\mathcal{J}(\boldsymbol{q}, \boldsymbol{u}, t)$) over a time-horizon Δh defined by:

$$\mathcal{J}(\boldsymbol{q}, \boldsymbol{u}, t) = \int_t^{t+\Delta h} \boldsymbol{q}(s)^T Q(s) \boldsymbol{q}(s) + \boldsymbol{u}(s)^T R(s) \boldsymbol{u}(s) ds \tag{2}$$

We then applied the first time-step of the control law and repeated the process by moving the time-horizon forward. Here, $Q(.)$ is a positive-semidefinite matrix penalizing deviations from desired limb states, and $R(.)$ is a positive-definite matrix penalizing large network inputs.

We now simply consider the structure of the controller to interpret its effect on the network-body dynamics across task variations. For the given system dynamics (Eqn. 1) and quadratic cost

function (Eqn. 2), the optimal control inputs for each time-horizon can be computed using the classical Linear Quadratic Regulator (LQR) approach[35,41], which yields an optimal state-feedback control law that can be decomposed as follows:

$$\boldsymbol{u}(t) = -L_{net}(\boldsymbol{\Theta}, t)\boldsymbol{r}(t) - L_{limb}(\boldsymbol{\Theta}, t)\boldsymbol{q}(t). \tag{3}$$

The set $\boldsymbol{\Theta}$ encompasses all the system and cost parameters (i.e. $\boldsymbol{\Theta} = \{Q, R, J, C, M\}$), and L_{net} and L_{limb} are the control gains applied to the network and limb states, respectively. Substituting this feedback control law (Eqn. 3) into the system dynamics (Eqn. 1) gives the closed-loop system:

$$\begin{bmatrix} \dot{\boldsymbol{r}} \\ \dot{\boldsymbol{q}} \end{bmatrix} = \begin{bmatrix} J - L_{net}(\boldsymbol{\Theta}, t) & -L_{limb}(\boldsymbol{\Theta}, t) \\ C & M \end{bmatrix} \begin{bmatrix} \boldsymbol{r} \\ \boldsymbol{q} \end{bmatrix} + \begin{bmatrix} \boldsymbol{\varepsilon} \\ 0 \end{bmatrix}. \tag{4}$$

Rewriting the system in this closed-loop form offers an interpretation of the controller structure: this optimal control scheme enables multiple tasks (represented by $\boldsymbol{\Theta}$) by dynamically adjusting the coupling of network nodes (J became $J - L_{net}(\boldsymbol{\Theta}, t)$), and the feedback about the limb states through the matrix $-L_{limb}(\boldsymbol{\Theta}, t)$ (Fig. 1a). In essence, this approach leverages the linear structure of the system to compute the optimal network connectivity suitable for a given set of parameters without data-driven (re)training. Thus, we may expect that different tasks (and their corresponding sets of parameters) induce different network activations, resulting in various shapes and trajectories of the network states representing neural population activity (Fig. 1b). We show below that this model remarkably reproduces previously published properties (Fig. 1c), and makes quantitative predictions of neural activities associated with standard tasks (Fig. 1d).

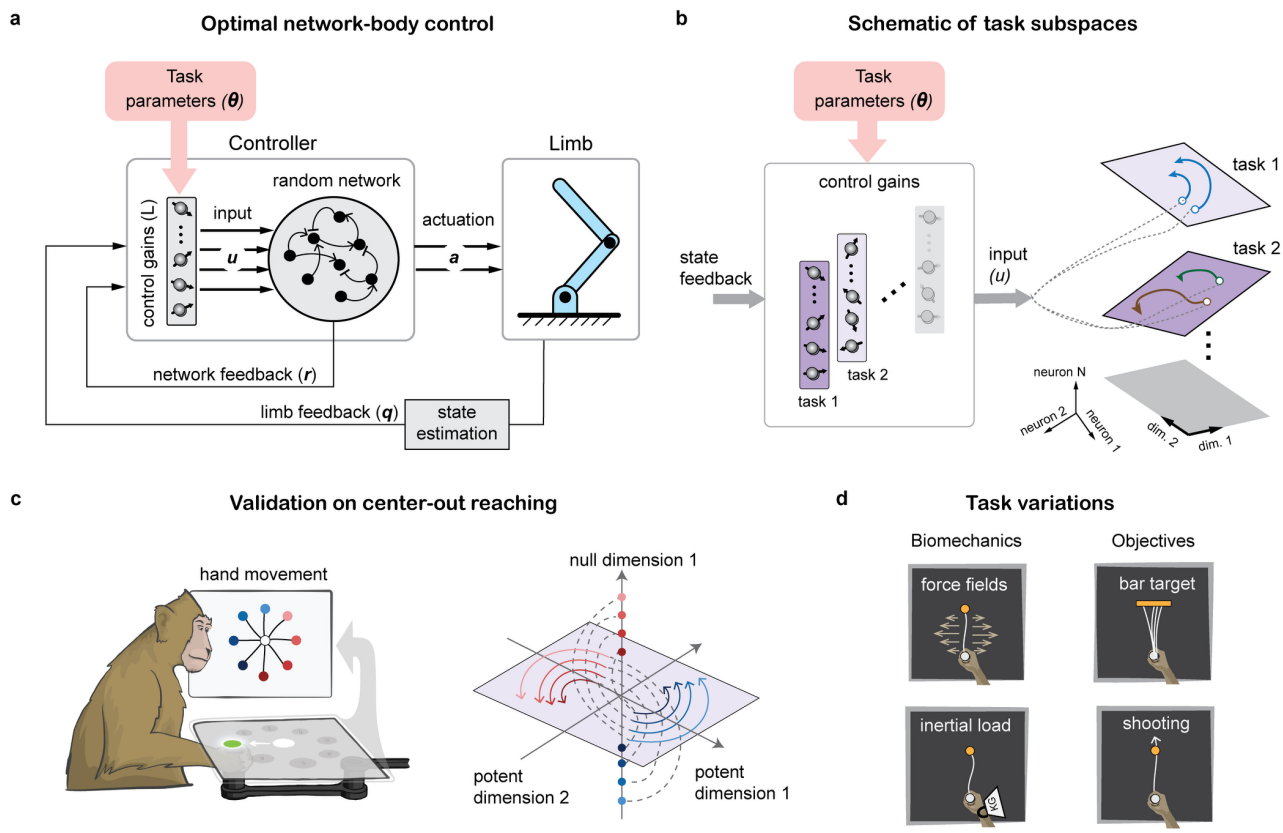

Figure 1 Schematic representation of the network-body system. a) Task-dependent control of network-body system through tunable control gains. b) illustration of how the tunable control gains may reshape the low dimensional neural trajectories based on task parameters. c) Illustration of a stereotypical prepare-to-reach task, showing the emergence of orthogonal subspaces for preparation and movement. d) Selected benchmark tasks taken from standard motor control literature used to test the malleability of sensorimotor control in primates dependent on changes in task objectives (flexibility) and biomechanics (adaptation).

Emergence of low-dimensional factors for preparation and execution

We first validate the model by showing that it reproduces hallmark features of monkey's primary motor cortex (M1) activity during a standard center-out reaching task[8,39] (Fig. 1c). Specifically, using a receding-horizon control scheme[42,43] (see Methods), the model generated realistic hand trajectories in a delayed prepare-to-reach paradigm, including bell-shape velocity profiles following the GO cue (Fig 2a). At the level of individual network nodes, we observed heterogenous activity patterns, with tonic levels of activity during the preparatory phase that quickly encoded the impending target, and a subsequent phasic activity linked to movement execution (Fig 2b). Both the amplitude and the phase of individual responses varied across nodes and movement directions.

To examine the dominant features of population activity and link them to previously reported neural factors, we performed principal component analysis (PCA) on the network activity during preparatory and movement epochs (see Methods). The dominant principal components (PCs) in the preparatory epoch encoded target information prior to movement onset (Fig. 2c). This finding demonstrates that directional encoding of the upcoming movement occurred within a subspace of $\mathbb{R}^n$ that did not produce any motor output. This observation is consistent with the documented existence of an *output null-space*, where neuronal activity unfolds along dimensions that do not produce any movement[5,39]. The same PCA applied to the population activity after the GO cue revealed rotational trajectories (Fig. 2d), closely resembling rotations reported in M1[8]. Notably, the control scheme showed distinct responses between preparatory versus movement epochs because the task objectives (Q) changed: during preparation deviation from start was penalized, whereas during movement deviation from the goal was penalized. As a result, two different controllers allowed the same state-feedback signals to produce different network trajectories depending on whether the system was preparing or executing a movement.

This task-dependent control scheme also captured another well-known feature of M1 activity: the orthogonality between preparatory and execution subspaces[39]. Applying PCA to network activity in each epoch separately, the preparatory-PCs explained little of the variance observed during the movement, and vice-versa. To quantify this, we defined an alignment index as proposed by Elsayed and colleagues[39] (see Methods), measuring how much variance in one epoch (preparatory or movement) was explained by the first ten PCs from the other epoch, normalized by the variance that these PCs explain in their own epochs. When the two subspaces are orthogonal, this alignment index is zero. In our model, the network activities exhibited a low alignment index (<0.1), consistent with recordings from M1[39]. Furthermore, we plotted functional connectivity to visualize correlations of time-series activity for each pair of nodes, separately during preparatory epoch (-150 to $-50\ ms$ before the GO cue) and movement epoch (50 to $150\ ms$ after the GO cue). Fig. 2f shows that the average pairwise-correlations across targets observed during preparation was not maintained during movement. Consequently, the pairwise-correlations among nodes during preparation were not linearly related to those observed during movement, resulting in a low Pearson correlation-coefficient (Fig 2g).

Thus, closed-loop control of the network-body system naturally gave rise to low-dimensional neural trajectories in orthogonal preparatory and movement subspaces, with rotational population

dynamics closely paralleling patterns of neural activities recorded in M1. Remarkably, these features emerged from a linear system combined with a quadratic cost. No additional constraint was required, such as enforcing smooth dynamics[9], or minimizing trajectory tangling[15,44] as expressed in previous models.

Variation in movement dynamics

Building on the validation of the model, we now illustrate its predictive power. Indeed, we can change any task parameter value (i.e. change a subset of $\mathbf{\Theta}$) and interrogate the model about the resulting neural activities. Using paradigmatic examples from motor control literature (Fig. 1d), we demonstrate how task parameters shape neural trajectories, enabling quantitative predictions about neural population dynamics across a range of tasks. Two subsets of task parameters are often studied: (1) the task objectives (the weighting matrices Q and R penalizing states or inputs in the cost function, expressed in Eqn. 2), and (2) the system dynamics (the matrix M). Based on the fact that the control gains, and hence the network inputs, are modulated as a function of the above-mentioned task parameters (Eqn. 3), we can explore how these parameters shaped the network dynamics.

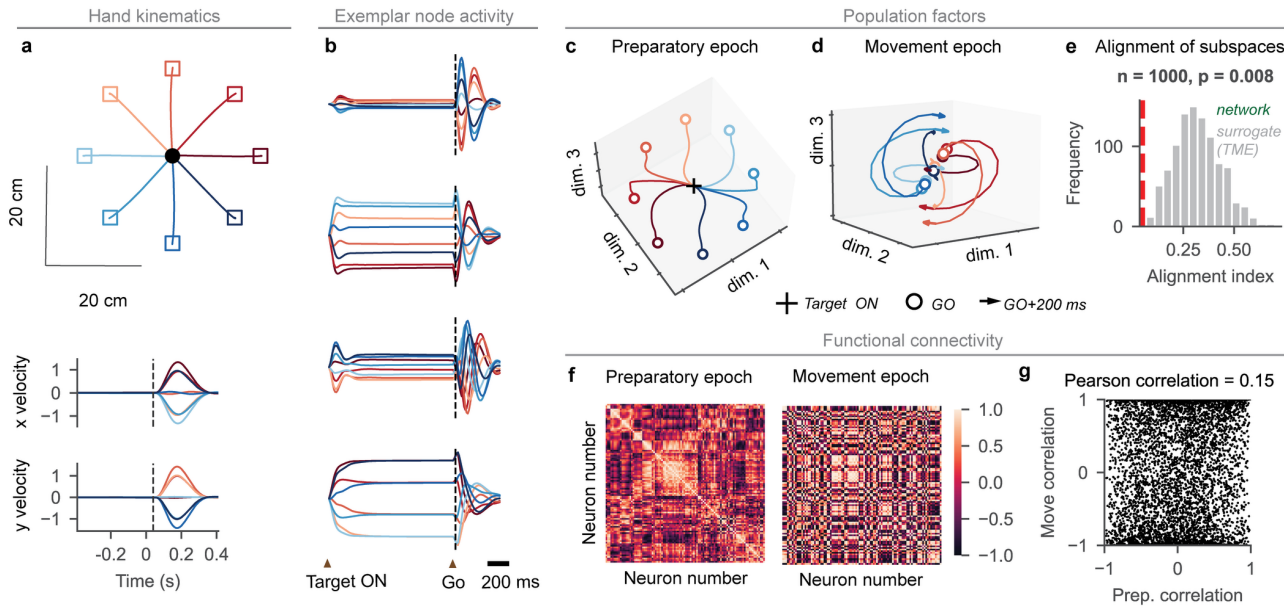

Figure 2 Feedback control produces meaningful behavior and reproduces M1-like preparation and movement-related "factors" a) Hand kinematics for the simulation of a point mass system b) exemplar neural firing rates from selected units, observe that the activity of some neurons depends on the target of the pending movement. c) neural activity along the principal modes in the preparatory period. d) same as c for the movement period. e) alignment between preparatory and movement period subspaces. The histogram shows distribution of alignment indices across surrogate data and the vertical line represent the alignment of subspaces from the principal modes of preparatory and execution epochs (low alignment corresponds to orthogonal). f) Pairwise correlations in preparatory and movement periods g) Linear relationship between pairwise-correlations from movement to preparatory period. The lack of pattern shows that activities from one epoch is statistically independent of that of the other epoch.

We first examine the influence of velocity-dependent force-fields, widely used to study motor adaptation[45], on the network dynamics. An orthogonal force-field of the form $f_x = \alpha\dot{y}$ modifies the limb dynamics by coupling forward speed and lateral force. Assuming that the controller has knowledge of this force field, the model was able to produce straight hand paths towards the target after the GO signal (Fig 3a), countering the applied lateral forces with scaled hand forces (Fig 3a, bottom-right). Note that we do not model the process of adaptation, but we assumed that the controller knows the force-field parameter. Although no movement occurred prior to the GO cue, a few individual nodes encoded the level of the force-field parameter (α) during the preparatory phase, without generating motor output (Fig. 3b). Hence, network activity in the output-null space already encoded the strength of force field parameter. This was followed by a scaling in the amplitude of rate responses in individual nodes during the movement epoch (Fig. 3b).

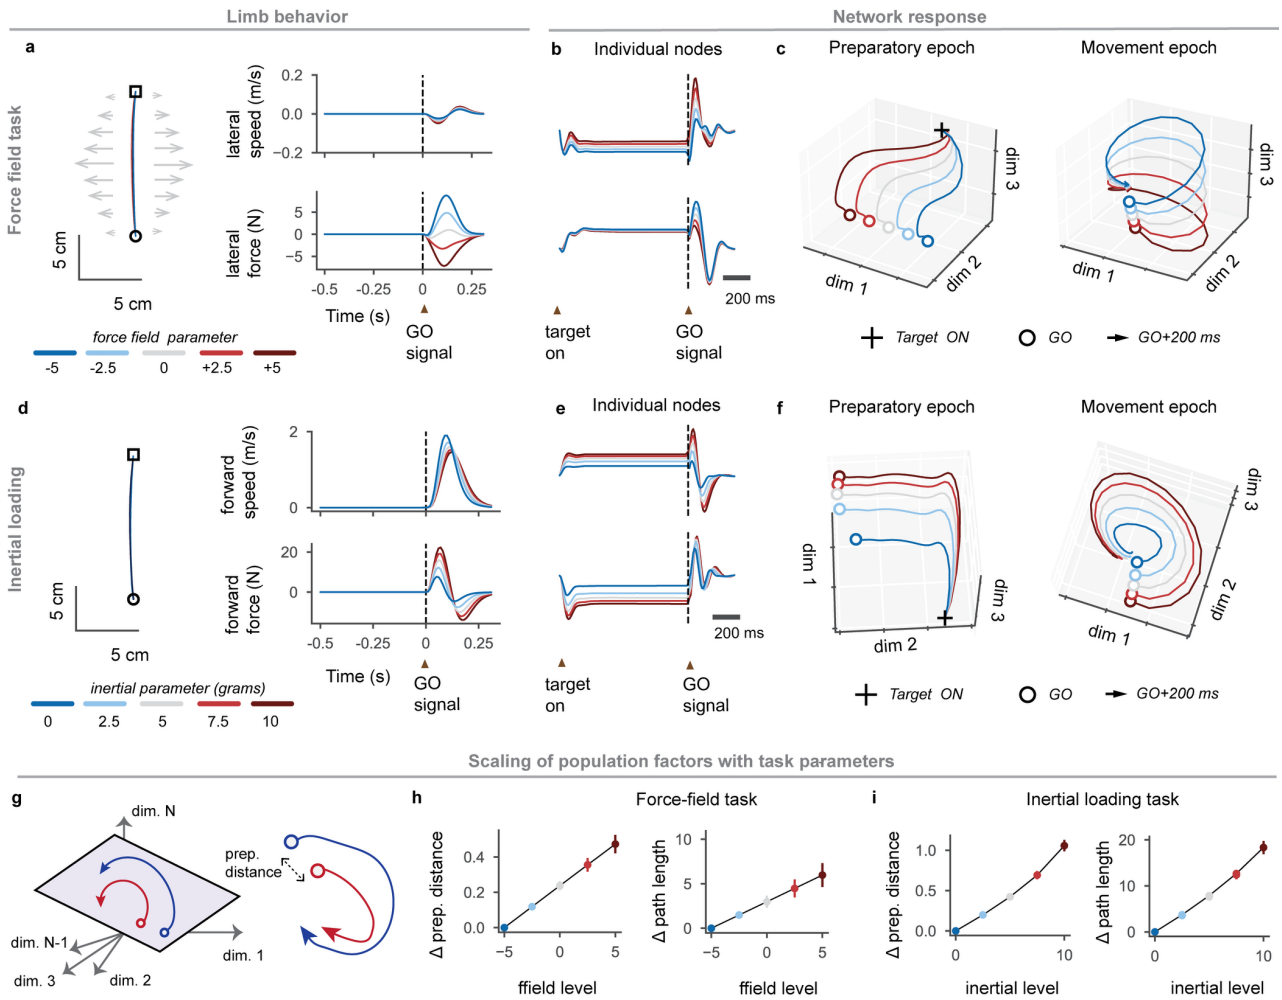

Figure 3 Flexible control to changes in biomechanical requirements. a) representation of the task with clockwise or counterclockwise force field of different strength, and hand kinematics with varying levels of lateral force applied to counter the force field. b) Exemplar neuron firing rates c) projection of the network trajectories in the principal mode of the preparation and movement epochs for five different levels of orthogonal force-field parameter. (d-e-f) same as above for five different levels of limb inertia. g) Illustration of the key measures of variation in neural factors: h) change in preparatory and movement related factors for force-field task. i) Same as h for inertial loading task. The circle markers depict mean across N=5 networks, and the error bars indicate the s.e.m.

At the population level, the dominant PCs encoded force field strength during the preparatory epoch, followed by rotational trajectories in the movement epoch (Fig. 3c). Notably, these rotations occupied different planes but maintained a consistent direction even when the force field was applied in opposite directions (-5 to $+5$). It is worthwhile noting that such invariant rotations were previously reported also in cycling tasks in monkeys, where the rotations in M1 maintained their direction while the EMG activity reversed with the cycling direction[44].

Next, we quantified the changes in neural population responses across force-fields by computing the path changes compared to the baseline parameter level ($\alpha = 0$) (Fig. 3g). The change in preparatory state was measured as the Euclidean distance between the preparatory state at the GO signal and at the baseline force-field parameter level. Movement-related changes in PCs were quantified as the sum of squared differences between points on the PC trajectory and the baseline trajectory. Across different networks, both preparatory and movement period PCs scale with the strength of the force-field parameter (α) (Fig 3h). Indeed, analyzing the corresponding control gains across varying force-field values revealed that these gains changed with α during both preparatory and execution epochs, thereby shifting network activity into different preparatory states that scaled with the force-field strength even when moving to the same spatial target.

We next asked whether this modulation in neural trajectories extended to other changes in limb dynamics. In a second set of simulations, we varied inertial loads by simulating different masses of the hand while reaching to the target (Fig. 3d). Similar to the force-field task, individual nodes encoded the inertial load level (Fig. 3e) during both preparatory and movement epochs. Applying PCA again showed that the dominant PCs in each epoch reflected the load magnitude (Fig. 3f). As before, both preparatory and movement PCs scaled with the level of loading (Fig. 3i), confirming that parameters altering limb dynamics produced feedback modulation that reshaped the network activity.

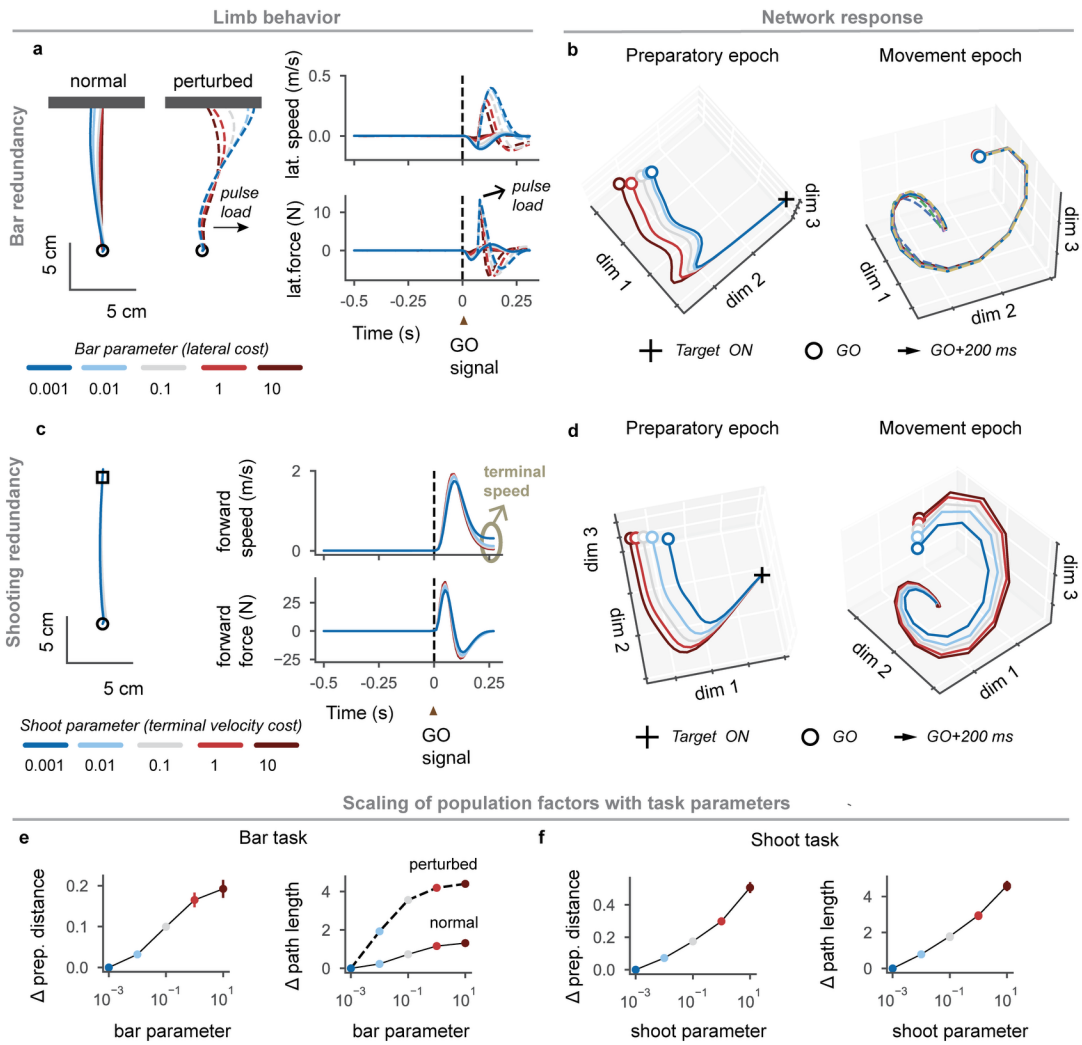

Figure 4 Flexible control to changes in task performance objectives a) representation of the task with varying levels of penalty on the lateral end-point coordinate allowing modulation of the feedback response b) projection of the network trajectories on the three principal modes from preparatory and movement epochs for five different levels of lateral terminal displacement cost parameter. (c-d) same as above for five different levels of shooting parameter captured by a cost on termina hand velocity. e) Change in preparatory and movement period factor properties for bar task. f) Same as e for the shooting task. The circle markers depict mean across N=5 networks, and the error bars indicate the s.e.m.

Variation in task objectives

We now use the same approach to show a modulation of population activity in response to variation in task objectives (varying Q matrix in Eqn. 3). To test this, we simulated two tasks typically used in human behavioral experiments: the bar task[46,47], and a shooting (or a hit) task[48].

In the bar task, the hand must reach a target with different levels of penalty on the lateral endpoint deviation, leading to variation in the lateral end position (Fig 4a). When a brief mechanical

load was applied to perturb the movement laterally, the model corrected errors more vigorously in directions where errors were more costly (high lateral cost), leading to smaller deviations and higher corrective forces (Fig. 4a, perturbed condition). In contrast when the cost was low, it corrected less (smaller force, larger hand deviation) as expected under the minimal intervention principle[34]. Despite the absence of motor output during preparation, some network nodes already encoded task objectives (i.e. lateral cost). At the population level, preparatory PCs reflected steady activity patterns corresponding to the lateral cost, followed by distinct trajectories during execution (Fig. 4b).

In the shooting task, the penalty on terminal velocity of reaching the target was systematically reduced, leading to shooting movements that hit the target without stopping on target (Fig 4c). Notably, the differences in kinematics were minimal (Fig. 4c), especially the differences in hand speed became apparent only near the end of the movement (Fig. 4c top-right panel). At the population level, preparatory PCs reflected changes in the cost on terminal velocity before the actual movement, followed by distinct movement related PCs (Fig. 4d). Across different network instantiations, both preparatory and movement-related PC trajectories scaled with cost parameters for both tasks, demonstrating that these parameters can be extracted from the low-dimensional features of the network activity (Fig. 4e, f). Movement PCs during perturbation showed larger path length modulation with the bar parameter (Fig. 4e), highlighting how state-feedback control integrates limb feedback to adjust network trajectories according to task objectives.

Low rank hypothesis applied to network-body system

Our developments and simulations suggest a critical role of state-feedback and task parameters in shaping population activity. First, the optimal input to the system clearly conveys feedback from the modelled limb into the network (Eqn. 3), which must influence the network trajectories. Second, our previous simulations show that both task parameters and perturbation responses together influence network trajectories (Fig. 4e). These results contrast with prevailing models that suggest neural population activity results from an autonomous dynamical system without limb feedback[9,27,28,30,49]. Here, we develop a more formal argument of why low-dimensional features emerge from limb feedback based on the spectral properties of the network-body system.

The dynamics of the full system (Eqn. 4) can be written compactly as $\dot{\boldsymbol{z}} = A_c \boldsymbol{z}$, where A_c represents the combined system dynamics of the network, limb and desired limb states under closed-

loop control. For simplicity we ignored the constant input ($\boldsymbol{\varepsilon}$) as it does not impact subsequent analyses. The matrix A_c has a block structure, consisting of a high-dimensional, partially random component ($A_c^{(1)}$) and a low-dimensional component ($A_c^{(2)}$) related to the limb feedback gains (Fig. 5a, see Methods).

$$\underbrace{\begin{bmatrix} J - BL_{net}(\boldsymbol{\Theta}) & -BL_{limb}(\boldsymbol{\Theta}) \\ C & M \end{bmatrix}}_{A_c} = \underbrace{\begin{bmatrix} J - BL_{net}(\boldsymbol{\Theta}) & 0 \\ C & M \end{bmatrix}}_{A_c^{(1)}} + \underbrace{\begin{bmatrix} 0 & -BL_{limb}(\boldsymbol{\Theta}) \\ 0 & 0 \end{bmatrix}}_{A_c^{(2)}} \tag{5}$$

Thibeault and colleagues [40] recently showed that many natural physical systems can be modelled as a combination of a random component and a low-dimensional component. They used Weyl's inequality to show such systems are effectively low-dimensional, provided the spectrum of the random component is small in comparison with the small rank component. By rewriting the Weyl's inequality from[40] while considering Eqn. 5, we get:

$$\left|\sigma_i(A_c) - \sigma_i\left(A_c^{(2)}\right)\right| \leq \left\|A_c^{(1)}\right\|_2 \tag{6}$$

for all $i \in \{1,2,\dots,n\}$, where σ_i is the ith singular value and $\|.\|_2$ is the spectral norm (largest singular value) of a matrix. The inequality indicates that, when the spectral norm of $A_c^{(1)}$ is sufficiently small, the singular values of the total connectivity matrix (A_c) remains close to the singular values of its low-dimensional component ($A_c^{(2)}$)[40]. In this scenario, the total matrix (A_c) also has attains a low effective rank, meaning that a small number of dominant features capture most of the connectivity structure.

We examined the dimensionality of network structure under the application of feedback control in the center-out reaching task. In line with the low rank property (Eqn. 6), the singular values of the matrix A_c decreased rapidly, and matched those of the limb-feedback connectivity term $A_c^{(2)}$ for the first four components (Fig. 5 a-b). Using a well-defined effective rank metric that is based on the uniformity of singular value distribution (see Methods), we found that the dimensionality of the network structure closely tracked that of the limb related input. Consequently, even though the system's initial connectivity was high-dimensional ($erank$: 69), the structure of the controller input forced low-dimensionality ($erank$: 3). Thus, it is expected that the network activity be also low-dimensional as their dimension is determined by the connectivity structure. Indeed, the effective rank of the covariance of network activity across time and targets was much lower ($erank$: 6.8) than the actual number of network nodes ($n = 100$) (Fig. 5c).

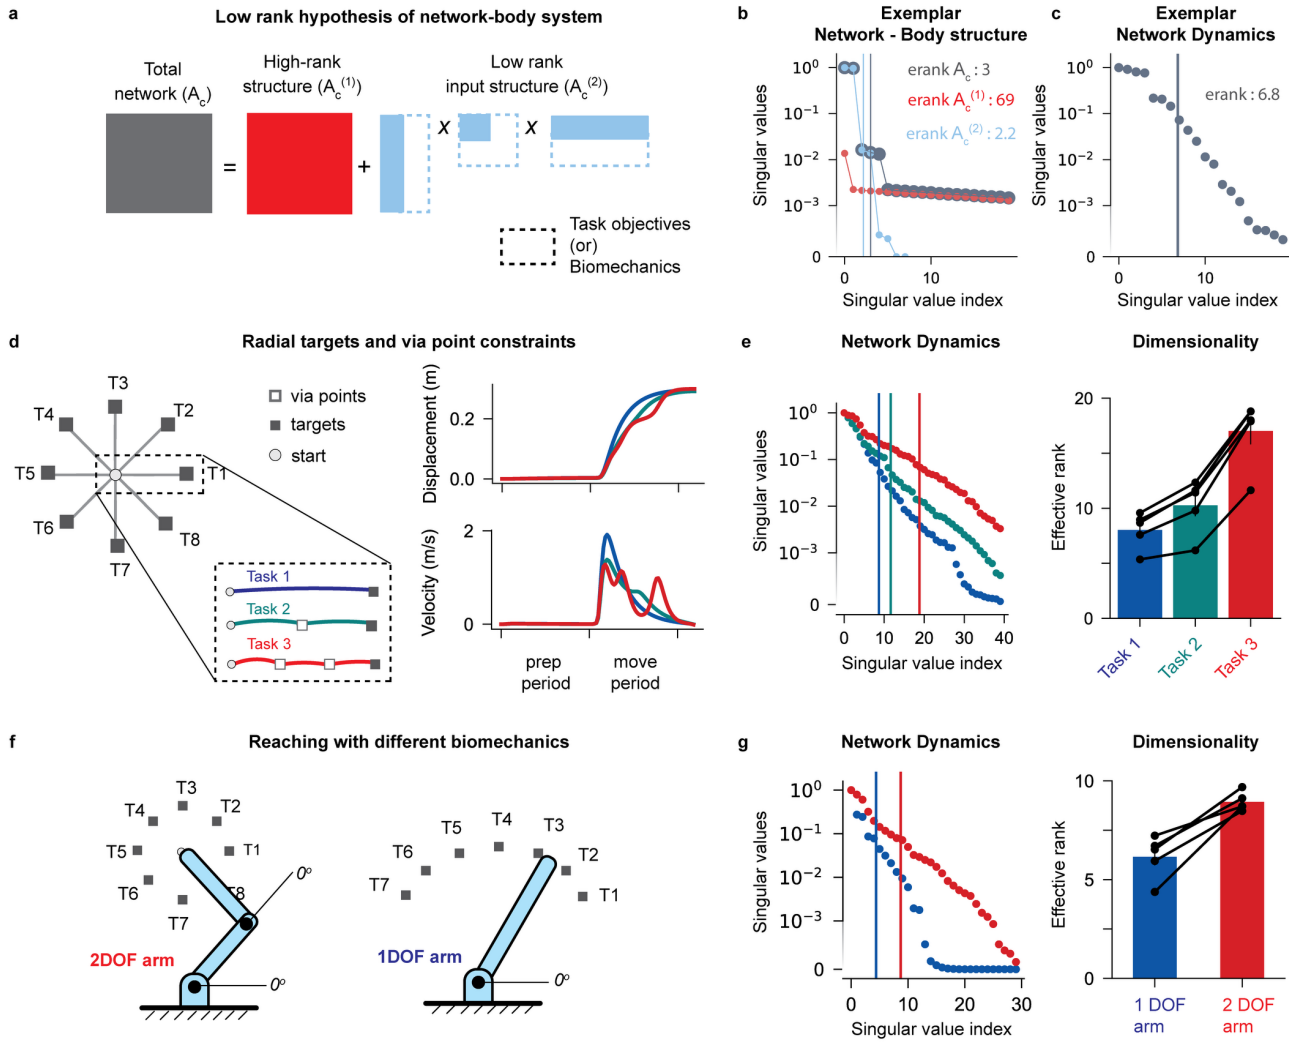

Figure 5 Testing low-rank hypothesis in the interacting network-body system a) decomposing network structure into the default connectivity and the limb input structure resulting from the structure of the LQR controller. We indicated the low-dimensional structure of $A_c^{(2)}$ by illustrating the singular value decomposition of this matrix. The dashed lines indicate that the rank of the low-dimensional components can change depending on the task definition and dimension of the effector. See main text (Eqn. 5) for the mathematical description of these block matrices. b) Singular values of the composite matrix (in gray), of the default connectivity (in red), and the limb input matrix (in cyan). c) Singular values of covariance matrix of network trajectories obtained after concatenation of their trajectories across the 8 different targets. Observe that the effective rank (<7) was low in comparison with the network size (n = 100). d) Illustration of the task with via-point used to modulate the complexity (left subpanel). Simulated hand displacement and velocity to a horizontal target in the three via-point tasks (right subpanels). e) Singular values of the network activity in the movement period (left), effective ranks of the covariance matrices of the network trajectories across five different instances of the model. The vertical lines are the effective ranks, the horizontal axis was truncated for clarity. f) Illustration of the eight-directional reaching task with a 2-degree of freedom (2-DOF) and a 1-degree of freedom (1-DOF) arm. g) Singular values of the covariance matrix of the network activities in the movement period (left), and effective ranks for five different instances of the model (right). The light gray lines on the bar plots correspond to individual network instances.

The foregoing spectral analysis connected the rapid decay of the singular values of the system matrix A_c (i.e., its structural properties) to the low-dimensionality of the network activity (i.e., its dynamics). A direct corollary is that increasing the dimensionality of the input signal (through $A_c^{(2)}$) containing task-relevant information should increase the effective rank of the network activity (dashed lines in Fig. 5a). We tested this in two ways: (1) first we increased task complexity by adding via-points along the movement path (2) then we modified the dimension of the plant by considering either a single joint system or a two-joints system.

We simulated the center-out reaching task with $0, 1$, or 2 intermediate via-points. Although hand paths were generally similar (Fig. 5d), the velocity profiles showed additional sub-movements as the number of via-points increased (Fig. 5e). Singular value analysis of the network activity showed a less rapid singular value decay as the number of via-points increased (Fig. 5f). Consequently, the effective rank, equivalent to the dimensionality of the network activities, increased with task-complexity (Fig. 5g), although the network structure, eight reach directions, and movement duration remained unchanged. This suggest that increasing task-demands selectively induced higher-dimensional feedback input signals, thereby increasing the dimensionality of the associated network dynamics.

We next varied limb biomechanics by simulating movements with either a one-dimensional (1-DOF) or a two-dimensional (2-DOF) limb to reach eight radial targets (Fig. 5f). The number of DOF directly modified the size of the state vector and of the feedback gain matrix $L_{limb}(\mathbf{\Theta})$. Again, singular value decomposition of the network activity showed that the 1-DOF simulations exhibited a steeper singular value decay, resulting in a lower effective rank than the 2-DOF simulations (Fig. 5g). Thus, increasing the biomechanical dimensions also increased the dimensionality of the corresponding network dynamics.

Together, these results support the view that low-rank motor cortical dynamics emerge naturally from optimal control of the network-body system, and provide a formal connection between the dimensionality of the input to the network, such as the number of degrees-of-freedom or the presence of intermediate goals, and the neural population activity. We made several additional observations that are reported in the Supplementary Material, which can be summarized as follows: our results appear robust against variations in the model structures for comparable network and inputs. The first supplementary figure shows that the results for the center-out reaching task were

similar when coupling the network to a two-link arm instead of a point-mass as shown in Fig. 2 (Supplementary Fig. 1). Then, we also observed consistent null and potent spaces, as well as rotations, in fully connected recurrent networks without any excitation-inhibition balanced structure, and randomly sampled Erdös-Rényi networks (see Methods and Supp. Fig. 2, 3), although the alignment index and functional connectivity also depended on the network structure. Finally, as a corollary of low rank hypothesis described above, we observed that for similar tasks and biomechanics, and therefore similar inputs to networks have comparable properties, the principal components were preserved across variations in the network connectivity (Supplementary Fig. 4). This latter observation is consistent with the observed similarity in neural population trajectories across individuals performing a task [47], assuming that different brains can be characterized by a similar network topology.

Discussion

Our results demonstrate that several features of neural population dynamics naturally emerge from optimal control of a simple, linear network-body system. By coupling a high-dimensional random network to a biomechanical plant and computing optimal control gains for each task, we replicated several hallmark features of motor cortical activity including rotational dynamics[8,38], preparatory states encoding the parameters of impending movements[5,36,37], and orthogonal null and potent subspaces[39]. Our model captured the fact that population activity reflects variations in task parameters, such as altered force fields[51,52], inertial load[53], or task demands before movement onset. Finally we demonstrated that this linear formulation establishes the critical role of the input to the network (here the state-feedback control signal) in the shaping of the network dynamics[54] based on the low-rank hypothesis of complex systems[40].

Clearly, the model was based on very simplifying assumptions: we assumed a fully observable state, direct readout for the actuation signals, and perfect knowledge of the network connectivity to derive the optimal controller. We believe that these assumptions were justified to formulate the basic properties of such a network-body system in the simplest case, and expect that future works unveil the properties of more biologically plausible models in detail. Clearly, considering a partially observable state (i.e., with delays and noise[55]), and indirect connections between the body and the network may enrich the predictions with insight into distributed processes across different brain regions[11,56,57] including cerebellum[58,59], somatosensory and parietal areas[56,60,61], and spinal cord[16,29,53]. Likewise, the question of robustness against imperfect knowledge of the network structure in the controller is an important challenge for prospective studies.

Importantly, these simplifying assumptions are not limitations since they can be addressed in more complex instances of the network-body model, by modifying the block-structure of the network, changing the body dynamics, or adding delays and noise following standard techniques[62]. Thus, the simplifying assumptions discussed above can be seen as an asset: our contribution has been to document comprehensively the remarkable explanatory power of the simplest instance of the model, thereby showing that the first basic, linear approximation already provides deep insight along with testable predictions about the shape and dimensions of neural factors that one may expect for any experiment that could be modeled in this framework.

Our approach differs from conventional neural network models which rely on data-driven training[9–13,15–17,19,63–67], with task-dependent network trajectories that can be recalled by contextual inputs[25] when similar situations arise. Here, because we formulated a linear model, the optimal network weights could be calculated analytically as optimal feedback control gains modifying both the way state feedback was fed into the network, as well as how different network nodes needed to be (de)coupled. This approach offers an original interpretation of how optimal controllers can be synthesized in the brain: variations in model parameters ($\boldsymbol{\Theta}$) yield distinct feedback control gains that modify input to the system and coupling of network nodes. Thus, task-selection (e.g. hold then move, slow versus fast, etc.,)[39,68], task-dependent control[15,57,69–72], and motor adaptation[51,53,73] can in principle be achieved in a fixed-network architecture by efficiently modulating the projections of somatosensory signals into the network according to the task parameters[74].

This view provides a mechanism by which humans express flexible, task-dependent control across fast timescales, i.e., between or even within trials[21,22,48,75,76] (<1s). While connectivity may be reshaped over longer timescales through synaptic plasticity[14,65,77,78], rapid adjustments may be plausibly explained by swiftly modifying feedback control inputs through a fixed network architecture. We do not rule out that real sensorimotor circuits may develop specialized connectivity for commonly encountered tasks over time[79], owing to plasticity in cortical and subcortical regions[80]. Rather, we propose that such experience-driven plasticity coexists with fast, task-dependent modulation of feedback driven input to the network, enabling primates to deploy flexible control policies across a wide range of contexts very rapidly.

Another key insight offered by the linear formulation was that low-dimensional neural factors emerge because limb and task parameters provide a dominant, low-rank input to the otherwise high-dimensional recurrent network. This aligns with the low-rank hypothesis of complex systems[40], which posits that a small subset of principal modes embedded within a large random network yields a low effective rank, making the system trajectories low-dimensional. The block structure of the network-body dynamics and the expression of optimal control gains (Eqn. 5) allowed decomposing the closed-loop system into components linked to the network or feedback input, and study the spectral properties of each component as well as their effect on the dimension of the network trajectories. A direct consequence of this relationship was that the dimensionality of neural trajectories scaled with task complexity or with the dimensionality of the effector (Fig. 5, via points and degrees of freedom,

respectively). Overall, this model portrays motor cortex as an input-driven dynamical system[32,54,56,58,74,81–84] in which task-dependent inputs shape neural activity.

In all this study shows that linear models as first approximation provide a powerful test bed to make interpretable predictions about the shape and dimensions of neural trajectories, providing a theoretical link between neural activity and sensorimotor control. This framework predicts that variations in task parameters lead to measurable changes in population activities, reflecting distinct control policies and biomechanical constraints. These normative insights may inform fundamental research and brain-machine development to overcome the challenges posed by increasingly large neural datasets.

Methods

We used the following convention throughout the paper: scalar quantities appear in italicized symbols, vectors in bold and italicized font, matrices in bold and capital symbols, and sets of parameters in capital symbols. The dependency on time is sometimes omitted for clarity.

Network dynamics.

We generated random networks comprising $n = 100$ units. The network activity was represented by the rate vector $\boldsymbol{r}(t)$ denoting the firing rate of all network units. The dynamics of this network was governed by the rate equation:

$$\tau_{net}\dot{\boldsymbol{r}} = -\boldsymbol{r}(t) + W\boldsymbol{r}(t) + \boldsymbol{\varepsilon} + \boldsymbol{u}(t) \tag{7}$$

Here $\tau_{net} = 20ms$ is the time constant of a single network unit, and W is the network adjacency matrix. For compactness, we describe the network dynamics by combining the first two terms into $J = \frac{1}{\tau_{net}}(W - I)$, where I is the identity matrix. Throughout this article, we refer to J as the network state transition matrix. The input to the network consists of two terms: $\boldsymbol{\varepsilon} \in \mathbb{R}^n$ is a constant vector drawn from a zero-mean normal distribution $\mathcal{N}(0,\ 5)$ generating an offset representing spontaneous activity. This constant offset was treated with system augmentation considering $\dot{\boldsymbol{\varepsilon}} = 0$. The vector $\boldsymbol{u}(t)$ is the control signal that must be calculated based on the cost function representing task demands. In our model, we assumed that all units receive direct control inputs. We observed that restricting direct control inputs to a subset of network units produced similar results. The network adjacency matrix W was initialized such that all its eigenvalues were negative. We studied three classes of random networks described below.

The first class of networks consisted of *Inhibition Stabilized Networks (ISN's),* comprising separate inhibitory and excitatory populations, with $n_E = 80$ excitatory units and $n_I = 20$ inhibitory units. The combined network activity was represented by the rate vector $\boldsymbol{r}(t) = \begin{bmatrix} \boldsymbol{r_E}(t) \\ \boldsymbol{r_I}(t) \end{bmatrix}$, where $\boldsymbol{r_E}(t)$ and $\boldsymbol{r_I}(t)$ denote the firing rates of the excitatory and inhibitory populations, respectively. To simulate inhibitory stabilized regime, the network adjacency matrix was constructed using the algorithm described by Hennequin and colleagues in[28,85].

The second class of networks consisted of *Fully Connected Recurrent Networks.* Here, the network adjacency matrix was drawn from a normal distribution $\mathcal{N}(0, g^2/\mathrm{n})$, where $\mathrm{g} = 0.8$. This configuration ensured stable recurrent dynamics, as described previously[9,11].

The third class of networks corresponded to *Erdös-Rényi* random networks, in which connections between unit pairs were determined probabilistically. Specifically, each connection was assigned with a given probability parameter, with $p = 0.4$ indicating that only 40% of total number of neuron pairs were connected (other parameter values yielded comparable results). Excitatory and inhibitory weights were assigned by randomly setting nonzero connections to $+0.5$ or -0.5 with equal probability. To ensure network stability, we first set the diagonal terms in the adjacency matrix to zeros. This was followed by incremental decrease in the diagonal terms until all eigen values became negative.

Overall, the given network drives limb movement by actuating muscles through an output vector $\boldsymbol{a}(t)$. The relationship between network activity and limb muscle actuation was defined as:

$$\boldsymbol{a}(t) = W_{out}\boldsymbol{r}(t) \tag{8}$$

Here $W_{\mathrm{out}} \in \mathbb{R}^{2\times n}$ is the readout matrix whose entries are drawn from a normal distribution, $\mathcal{N}(0, 0.5)$, allowing multiple network units to contribute at different amounts to limb actuation. Notably, each random network corresponded to different random samples of the adjacency matrix, while maintaining a fixed readout matrix.

Limb dynamics.

For the point mass approximation of limb dynamics, the limb state vector was defined as $\boldsymbol{q} = [x \quad \dot{x} \quad f_x \quad y \quad \dot{y} \quad f_y]^T$ where the coordinates x and y represent the displacement of point mass in the lateral and forward directions respectively, the dot represent their time derivatives, and f_x and f_y are the forces applied by the controlled actuator to the mass. These forces respond to the network's output $\boldsymbol{a}(t)$, with a_x and a_y representing the network's actuation in the x and y directions, respectively. Overall, the system of differential equations for the limb and the actuation dynamics was:

$$m\ddot{x} = -k\dot{x} + f_x \tag{9}$$

$$m\ddot{y} = -k\dot{y} + f_y \tag{10}$$

$$\tau_{act}\dot{f}_x = a_x - f_x \tag{11}$$

$$\tau_{act}\dot{f}_y = a_y - f_y \tag{12}$$

Here k is the dissipative constant (0.1 Nsm^{-1}), and Equations (11), (12) represent the first order muscle dynamics with the actuation time-constant $\tau_{act} = 0.1\ s$.

The limb state evolves according to the differential equation $\dot{\boldsymbol{q}} = M\boldsymbol{q}(t) + F\boldsymbol{a}(t)$, with:

$$M = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & -\frac{k}{m} & \frac{1}{m} & 0 & 0 & 0 \\ 0 & 0 & -\frac{1}{\tau_{act}} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & -\frac{k}{m} & \frac{1}{m} \\ 0 & 0 & 0 & 0 & 0 & -\frac{1}{\tau_{act}} \end{bmatrix} \tag{13}$$

$$F = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ \frac{1}{\tau_{act}} & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & \frac{1}{\tau_{act}} \end{bmatrix} \tag{14}$$

Additionally, we have simulated a linear approximation of a two-dimensional planar arm actuated at the shoulder and elbow joints. Details on the methods and simulations with this shoulder-elbow arm are provided in the Supplementary Material.

For simplicity, we can define a joint variable $\boldsymbol{z} = \begin{bmatrix} \boldsymbol{r} \\ \boldsymbol{q} \\ \boldsymbol{\varepsilon} \end{bmatrix}$ to jointly describe the limb and network states, including the constant offset term for spontaneous activity, expressed as $\dot{\boldsymbol{\varepsilon}} = 0$. The dynamics of this joint system can be expressed compactly as $\dot{\boldsymbol{z}} = A\boldsymbol{z} + Bu$, where:

$$\dot{\boldsymbol{z}} = \begin{bmatrix} \dot{\boldsymbol{r}} \\ \dot{\boldsymbol{q}} \\ \dot{\boldsymbol{\varepsilon}} \end{bmatrix}, A = \begin{bmatrix} J & 0 & 1 \\ C & \mathrm{M} & 0 \\ 0 & 0 & 0 \end{bmatrix} and\ B = \begin{bmatrix} I \\ 0 \\ 0 \end{bmatrix} \boldsymbol{u}. \tag{15}$$

Here 0 and 1 represents a block matrix filled with 0's and $1's$ respectively, and I represents an identity matrix with size equal to the number of network units. Here, $C = FW_{\mathrm{out}}$ represents the

transformation (FW_{out}) from network activity ($\boldsymbol{r}$) to the forces acting on the limb ($\boldsymbol{f}$). Note the $\tilde{A}$ was denoted with a tilde to indicate the default state-transition matrix of the network-body system. Applying closed-loop control, as follows, involves modifications in this state-transition matrix with feedback components.

Receding horizon control in a Linear Quadratic Regulator (LQR) form.

We solved the problem of finding optimal network control inputs by repeatedly minimizing a cost on limb states and inputs within a moving time horizon

$$\boldsymbol{u} \coloneqq \underset{\boldsymbol{u}}{\operatorname{argmin}}\, \mathcal{J}(\boldsymbol{z}, \boldsymbol{u}) \tag{16}$$

Given the quadratic form of the cost function (Eqn. 2), and the linear system dynamics (Eqn. 3), the optimal control reflects the well-known LQR problem [35,41], whose solutions takes the following state-feedback form

$$\boldsymbol{u}(t) = -L(A, B, Q, R, t)\boldsymbol{z}(t) \tag{17}$$

Here, L is a time-varying control gain matrix that linearly maps the available state measurements ($\boldsymbol{z}$) or their estimates onto control inputs. Importantly, the control gains are dependent on the system dynamics (A and B) and the task objectives (Q and R) [35,41]. Calling $\boldsymbol{\Theta}^{(i)} = \{A^{(i)}, B^{(i)}, Q^{(i)}, R^{(i)}\}$ the set of task parameters known to the controller for a given task (i), we can rewrite the control law as:

$$\boldsymbol{u}(t) = -L\left(\boldsymbol{\Theta}^{(i)}, t\right)\boldsymbol{z}(t) = \left[-L_{net}\left(\boldsymbol{\Theta}^{(i)}, t\right) \quad \middle| \quad -L_{limb}\left(\boldsymbol{\Theta}^{(i)}, t\right)\right]\begin{bmatrix}\boldsymbol{r}(t)\\ \boldsymbol{q}(t)\end{bmatrix} \tag{18}$$

This formulation expresses that the state-feedback control policy is specific to a given task parameter set $\boldsymbol{\Theta}$. Notice that in results section, we denoted the task parameter set with individual block components (J, C, M) of the system matrix (A), but they are essentially the same. Notably, the control law can be decomposed to separate the contribution of network or limb related feedback.

Overall, the Receding Horizon Control scheme [42] iteratively recalculates the optimal control gains at each step. This process involves:

1. Collecting the task parameters $\boldsymbol{\Theta}^{(i)}$ corresponding to the current task (i) at each time step (t)
2. Calculating the control gains over the prediction horizon $\left\{L\left(\boldsymbol{\Theta}^{(i)}, t\right), \dots, L\left(\boldsymbol{\Theta}^{(i)}, t + \Delta h\right)\right\}$, for the current time window (from t to $t + \Delta h$).
3. Applying the first control gain of the sequence, $\boldsymbol{u}(t) = L\left(\boldsymbol{\Theta}^{(i)}, t\right)\boldsymbol{z}(t)$

4. Updating the time to $(t+1)$ and repeating from step 1.

In this work, we assumed that the knowledge of task parameters is available and expect that adding a module to perform inference about these parameters when they are not known exactly may link our developments to more general contexts [86].

Prepare-to-reach task.

This task involved holding posture at a given starting location for an unspecified duration, then executing the movement towards the target upon receiving a "GO" signal. This paradigm effectively induces a transition between postural control (prepare) and reaching (move) control following a signal which occurs at a non-predictable time. To capture the fact that the occurrence of the “GO” signal was unpredictable, we employed a two-phases receding horizon cost-function with three terms as follows:

$$\mathcal{J}(\boldsymbol{q},\boldsymbol{u},t) = \underbrace{\int_{t}^{t+\Delta h_{prep}} \boldsymbol{q}(s)^T Q_{prep}(s)\boldsymbol{q}(s)ds}_{prep\ cost} + \underbrace{\int_{t+\Delta h_{prep}}^{t+\Delta h_{prep}+\Delta h_{move}} \boldsymbol{q}(s)^T Q_{move}(s)\boldsymbol{q}(s)ds}_{move\ cost} + \underbrace{\int_{t}^{t+\Delta h_{prep}+\Delta h_{move}} \boldsymbol{u}(s)^T R(s)\boldsymbol{u}(s)ds}_{input\ cost} \tag{19}$$

Here, $Q_{prep}(.)$ and $Q_{move}(.)$ penalize hand deviations from the start and goal locations respectively. Before the “GO” signal, the first term penalizes deviations from the starting position over a predicted preparatory horizon $\Delta h_{prep} = 50\ ms$. This horizon maybe much shorter than the actual preparatory period (typically $> 500\ ms$), to encourage rapid preparation for the uncertain GO signal. The second term penalizes deviations from the target over a fixed movement horizon $\Delta h_{move} = 500\ ms$, starting at $t + \Delta h_{prep}$. The third term penalizes control inputs over both horizons. This cost was recomputed at each time step during the preparatory period until the “GO” signal occurs.

Once the GO signal was delivered, the predicted preparatory horizon was replaced by the measured time of the GO signal (t_{GO}). The cost function for the movement period became:

$$\mathcal{J}(\boldsymbol{q},\boldsymbol{u},t) = \int_{t}^{t_{GO}+\Delta h_{move}} \boldsymbol{q}(s)^T Q_{move}(s)\boldsymbol{q}(s)ds + \int_{t}^{t_{GO}+\Delta h_{move}} \boldsymbol{u}(s)^T R(s)\boldsymbol{u}(s)ds \tag{20}$$

The first term penalizes the deviations from the target state, while the second term penalizes control inputs over the fixed movement horizon. By expressing the distinct objectives of preparatory and movement periods into separate cost terms, this optimization naturally yields two distinct sets of feedback control gains: one for maintaining posture while anticipating a movement during preparation, and another for executing the reach once the GO signal arrives.

Changes in Limb Mechanics

We investigated two specific scenarios: a velocity-dependent force field and an inertial loading task. In the Velocity-Dependent Force Field Task (FF), the limb was subjected to a force field where the force f_x was proportional to the effector's velocity $\dot{y}$, expressed as $f_x = \alpha\dot{y}$ as in standard motor adaptation paradigms. This introduced a coupling between forward velocity and lateral force, effectively modifying the limb dynamics matrix M to $M(\alpha)$.

In the Inertial Loading Task, we simulated the addition of inertial loads (m_{load}), which altered the limb dynamics. These additional loads influence the limb dynamics matrix M, resulting in a modified matrix $M(m_{load})$. This allowed us to examine how the control scheme adapted to changes in limb inertia. We simulated the prepare-to-reach task using the receding horizon control scheme across five different levels for each parameter: $\alpha \in \{-5, -2.5, 0, 2.5, 5\}\ NSm^{-1}$ and $m_{load} \in \{0.2, 0.4, 0.6, 0.8, 1.0\}\ KG$.

Changes in Task Objectives

To investigate the influence of task objectives on the network-body dynamics, we modulated matrix Q that assigns cost penalty on each state component.

In the spatial target redundancy task, we focused on the terminal lateral deviation at the final time step to simulate reaching to a target with different geometry (as the dot-bar task [46]). We systematically varied the parameter governing this deviation in the cost matrix Q across five levels: $\{0.001, 0.01, 0.1, 1, 10\}$. A low parameter value (e.g., 0.001) imposed minimal penalty on terminal lateral deviation, allowing for larger deviation in endpoint location. Conversely, higher parameter value (e.g., 10) imposed a strong penalty, necessitating corrective actions to achieve a low terminal displacement.

In the "shooting" task, we focused on the terminal hand velocity at the final time step. We varied the cost associated with this parameter across the same five levels: $\{0.001, 0.01, 0.1, 1, 10\}$. At low cost values (e.g., 0.001), stopping at the target was not required, resulting in non-zero terminal velocity—consistent with a "shooting" motion. At higher values (e.g., 10), terminal velocity was heavily penalized, necessitating deceleration or stopping of the hand as it reached the target.

Both biomechanical adaptation (modifying matrix M) and flexibility in task objectives (adjusting matrix Q) involved changes in the task parameter set $\mathbf{\Theta}^{(i)}$. As described in Equation (18), alterations in these parameters modified the control gains $L(t)$, leading to task-dependent network-body trajectories. By systematically varying the parameters associated with limb dynamics (k or m) and the entries in the cost function (Q), we assessed how the control system adapted its policy to maintain optimal performance across different biomechanical conditions and task objectives.

PCA in preparatory and movement epochs.

Following previous work by Elsayed and colleagues [39], we grouped the network responses into two matrices corresponding to the preparatory and movement epochs. Preparatory epoch responses were represented by a rectangular matrix $P \in \mathbb{R}^{n \times cd}$, where n is the number of units, c is the number of conditions (e.g., number of targets, force-field values, bar target penalties, via-point scenarios etc.,), and d is the total duration of the preparatory epoch. Similarly, movement execution period responses are represented by $E \in \mathbb{R}^{n \times cd}$. We performed PCA on the covariance matrix of preparatory response epochs P, treating each row (i.e., network unit) as a variable, to obtain the top ten components (prep-PCs) of the network trajectories in the n-dimensional space. Similarly, we obtained the top ten components from the movement epochs (move-PCs) by performing PCA on the covariance matrix of movement response E.

We compared how the pairwise correlations changed between the preparatory and movement epochs by computing the correlation in firing rates between each possible pair of units during each epoch. We then computed an average of these correlations across eight target conditions. This yielded two correlations for each pair of units, one for each of the preparation and movement epoch (Fig 2f). We tested if the correlation coefficient for each pair of neurons in one epoch linearly correlated with that of the other epoch. Pearson correlation coefficient was used to quantify the pairwise correlations between the preparatory and movement epochs across all neuron pairs (Fig 2g).

Alignment Index

We determined how aligned the top-ten preparatory and movement related principal components were by the alignment index, following the procedure described previously in [39] :

$$AI = \frac{Tr\left(PC_{move}^{T} Cov_{prep} PC_{move}\right)}{Tr\left(PC_{prep}^{T} Cov_{prep} PC_{prep}\right)} \tag{21}$$

Here PC_{prep} and $PC_{move} \in \mathbb{R}^{10\times cd}$ are matrices containing the top-ten preparatory and movement principal components, Cov_{prep} is the preparatory covariance matrix, and $Tr(.)$ is the trace operator. $Eqn. 21$ described the ratio of how much variance the top-ten movement principal components could explain of the preparatory activity with how much top-ten preparatory principal components could explain of the preparatory activity (the maximum preparatory variance captured by any ten linear projections). The alignment index ranges from 0, indicating the preparatory and movement principal components are orthogonal, to 1, indicating the preparatory and movement principal components are perfectly aligned.

Tensor Maximum Entropy (TME)

To determine whether the properties of the observed principal components (PCs) of network trajectories – specifically, low alignment indices and preserved PCs across networks – reflected population-level phenomena or merely, we employed the TME method [87]. TME method generates surrogate datasets that retain the second-order statistics of the original network data (covariances across neurons, conditions, and time) while randomizing the higher-order interaction terms. We generated surrogate datasets for each network across simulated tasks.

In the center-out reach simulated task, we computed the PCs of 1000 randomly sampled TME surrogate datasets, for each network activity dataset, during both the preparation and movement epochs, then calculated the alignment index. We compared these surrogate alignment indices with those obtained from the actual network data (Fig. 2, Supplementary Fig. 1, 2, 3). A t-test was used to test the statistical differences in the alignment index of the actual network activity and the surrogate data with significance level set to $p < 0.01$.

Furthermore, for each task, we computed the PCs of 1000 randomly sampled TME datasets per network and examined the correlations among these PCs across networks performing the same task (Supplementary Fig. 4).

Singular value decomposition analysis.

We performed singular value decomposition to analyze the dimensionality of the network connectivity matrices and population activities. Recall that the network-body dynamics can be jointly expressed as a closed-loop system of the form presented in Eqn. 4. For simplicity, we can ignore the spontaneous activity offset $\boldsymbol{\varepsilon}$, and the overall system can be represented as follows,

$$\begin{bmatrix} \dot{\boldsymbol{r}} \\ \dot{\boldsymbol{q}} \end{bmatrix} = \begin{bmatrix} J - BL_{net}(\boldsymbol{\Theta}) & -BL_{limb}(\boldsymbol{\Theta}) \\ C & M \end{bmatrix} \begin{bmatrix} \boldsymbol{r} \\ \boldsymbol{q} \end{bmatrix} \tag{24}$$

The above network-body dynamics jointly express closed-loop system of the form $\dot{\boldsymbol{z}} = A_c \boldsymbol{z}$, where A_c has a block structure corresponding to network and limb dynamics as well as their connection. Note that A_c represents the closed-loop state-transition matrix that includes control gains, and it is different from the open-loop state-transition matrix A described in Eqn. 15 before applying the feedback control law. This matrix was then decomposed into two terms ($A_c^{(1)}$ and $A_c^{(2)}$) as described in Eqn. 5.

Singular value spectra have been found for each of the decomposed terms and of the composite matrix (Figure 5 a, b). Additionally, singular value spectrum has also been computed for the network responses matrix in the movement epoch, $E \in \mathbb{R}^{n \times cd}$ (Figure 5c). Here $c = 8$ describes the number of target directions in the center-out reaching task for both the via-point (Fig. 5d) and biomechanical tasks (Fig. 5f).

Dimensionality

Given the singular value spectra, $\sigma_{i=1,2,\dots,n}$, of a matrix A, its dimensionality can be measured from the well-defined Roy and Vetterli's effective rank as described in [40]:

$$erank(A) = \exp([H(p_1, \dots, p_n)]) \tag{25}$$

Here $H(p_1, \dots, p_n) = -\sum_i p_i log p_i$ is the Shannon entropy of the singular value spectrum, measured as a function of the singular value mass function

$$p_i = \frac{\sigma_i}{\|\mathbf{A}\|_*} = \frac{\sigma_i}{\sum_{j=1}^{n} \sigma_j}, \forall i \in \{1, \dots, n\} \tag{26}$$

$erank$ measures the uniformity of the singular value distribution, resulting in a low effective rank when the singular values decrease rapidly.

Acknowledgements

HTK was supported by the Fonds de la Recherche Scientifique (FRS-FNRS) Chargé de recherche Grant CR 252 (FC 043127). FC was supported by the FRS-FNRS Grant 1.C.033.18 (FC 036239).

Code Availability

The code and data from simulations will be available in the public repository 'figshare' after publication

References

1. Krakauer, J. W., Ghazanfar, A. A., Gomez-Marin, A., MacIver, M. A. & Poeppel, D. Neuroscience Needs Behavior: Correcting a Reductionist Bias. *Neuron* **93**, 480–490 (2017).
2. Urai, A. E., Doiron, B., Leifer, A. M. & Churchland, A. K. Large-scale neural recordings call for new insights to link brain and behavior. *Nat Neurosci* **25**, 11–19 (2022).
3. Gao, P. & Ganguli, S. On simplicity and complexity in the brave new world of large-scale neuroscience. *Current Opinion in Neurobiology* **32**, 148–155 (2015).
4. Stringer, C. & Pachitariu, M. Analysis methods for large-scale neuronal recordings. *Science* **386**, eadp7429 (2024).
5. Churchland, M. M. & Shenoy, K. V. Preparatory activity and the expansive null-space. *Nat. Rev. Neurosci.* 1–24 (2024) doi:10.1038/s41583-024-00796-z.
6. Vyas, S., Golub, M. D., Sussillo, D. & Shenoy, K. V. Computation Through Neural Population Dynamics. *Annual Review of Neuroscience* **43**, 249–275 (2020).
7. Gallego, J. A., Perich, M. G., Miller, L. E. & Solla, S. A. Neural Manifolds for the Control of Movement. *Neuron* **94**, 978–984 (2017).
8. Churchland, M. M. *et al.* Neural population dynamics during reaching. *Nature* **487**, 51–56 (2012).
9. Sussillo, D., Churchland, M. M., Kaufman, M. T. & Shenoy, K. V. A neural network that finds a naturalistic solution for the production of muscle activity. *Nature Neuroscience* **18**, 1025–1033 (2015).
10. Feulner, B. *et al.* Small, correlated changes in synaptic connectivity may facilitate rapid motor learning. *Nat Commun* **13**, 5163 (2022).
11. Michaels, J. A., Schaffelhofer, S., Agudelo-Toro, A. & Scherberger, H. A goal-driven modular neural network predicts parietofrontal neural dynamics during grasping. *PNAS* **117**, 32124–32135 (2020).
12. Kalidindi, H. T. *et al.* Rotational dynamics in motor cortex are consistent with a feedback controller. *eLife* **10**, e67256 (2021).

13. Lillicrap, T. P. & Scott, S. H. Preference Distributions of Primary Motor Cortex Neurons Reflect Control Solutions Optimized for Limb Biomechanics. *Neuron* **77**, 168–179 (2013).
14. Gurnani, H., Liu, W. & Brunton, B. W. Feedback control of recurrent dynamics constrains learning timescales during motor adaptation. 2024.05.24.595772 Preprint at https://doi.org/10.1101/2024.05.24.595772 (2024).
15. Saxena, S., Russo, A. A., Cunningham, J. & Churchland, M. M. Motor cortex activity across movement speeds is predicted by network-level strategies for generating muscle activity. *eLife* **11**, e67620 (2022).
16. Guang, H. *et al.* Parallel feedback processing for voluntary control: knowledge of spinal feedback in motor cortex. 2024.05.12.593756 Preprint at https://doi.org/10.1101/2024.05.12.593756 (2024).
17. Aldarondo, D. *et al.* A virtual rodent predicts the structure of neural activity across behaviours. *Nature* **632**, 594–602 (2024).
18. Zador, A. *et al.* Catalyzing next-generation Artificial Intelligence through NeuroAI. *Nat Commun* **14**, 1597 (2023).
19. Codol, O., Krishna, N. H., Lajoie, G. & Perich, M. G. Brain-like neural dynamics for behavioral control develop through reinforcement learning. 2024.10.04.616712 Preprint at https://doi.org/10.1101/2024.10.04.616712 (2024).
20. Comite, A. D., Lefèvre, P. & Crevecoeur, F. Continuous evaluation of cost-to-go for flexible reaching control and online decisions. *PLOS Computational Biology* **19**, e1011493 (2023).
21. Crevecoeur, F., Mathew, J., Bastin, M. & Lefèvre, P. Feedback Adaptation to Unpredictable Force Fields in 250 ms. *eNeuro* **7**, (2020).
22. Braun, D. A., Aertsen, A., Wolpert, D. M. & Mehring, C. Learning Optimal Adaptation Strategies in Unpredictable Motor Tasks. *J. Neurosci.* **29**, 6472–6478 (2009).
23. Kalidindi, H. T. & Crevecoeur, F. Human reaching control in dynamic environments. *Current Opinion in Neurobiology* **83**, 102810 (2023).

24. Kalidindi, H. T. & Crevecoeur, F. Task dependent coarticulation of movement sequences. *eLife* **13**, (2024).
25. Mante, V., Sussillo, D., Shenoy, K. V. & Newsome, W. T. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* **503**, 78–84 (2013).
26. Beiran, M., Meirhaeghe, N., Sohn, H., Jazayeri, M. & Ostojic, S. Parametric control of flexible timing through low-dimensional neural manifolds. *Neuron* **111**, 739-753.e8 (2023).
27. Logiaco, L., Abbott, L. F. & Escola, S. Thalamic control of cortical dynamics in a model of flexible motor sequencing. *Cell Reports* **35**, 109090 (2021).
28. Kao, T.-C., Sadabadi, M. S. & Hennequin, G. Optimal anticipatory control as a theory of motor preparation: A thalamo-cortical circuit model. *Neuron* **109**, 1567-1581.e12 (2021).
29. Lindén, H., Petersen, P. C., Vestergaard, M. & Berg, R. W. Movement is governed by rotational neural dynamics in spinal motor networks. *Nature* **610**, 526–531 (2022).
30. Stroud, J. P., Porter, M. A., Hennequin, G. & Vogels, T. P. Motor primitives in space and time via targeted gain modulation in cortical networks. *Nat Neurosci* **21**, 1774–1783 (2018).
31. Sohn, H., Meirhaeghe, N., Rajalingham, R. & Jazayeri, M. A Network Perspective on Sensorimotor Learning. *Trends in Neurosciences* **44**, 170–181 (2021).
32. Schimel, M., Kao, T.-C. & Hennequin, G. When and why does motor preparation arise in recurrent neural network models of motor control? *eLife* **12**, (2023).
33. Zheng, C., Wang, Q. & Cui, H. Continuous sensorimotor transformation enhances robustness of neural dynamics to perturbation in macaque motor cortex. *Nat Commun* **16**, 3213 (2025).
34. Todorov, E. & Jordan, M. I. Optimal feedback control as a theory of motor coordination. *Nature Neuroscience* **5**, 1226–1235 (2002).
35. Åström, K. J. *Introduction to Stochastic Control Theory*. vol. 70 (Academic Press, 1970).
36. Kaufman, M. T., Churchland, M. M., Ryu, S. I. & Shenoy, K. V. Cortical activity in the null space: permitting preparation without movement. *Nature Neuroscience* **17**, 440–448 (2014).

37. Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Ryu, S. I. & Shenoy, K. V. Cortical Preparatory Activity: Representation of Movement or First Cog in a Dynamical Machine? *Neuron* **68**, 387–400 (2010).
38. Sabatini, D. A. & Kaufman, M. T. Reach-dependent reorientation of rotational dynamics in motor cortex. *Nat Commun* **15**, 7007 (2024).
39. Elsayed, G. F., Lara, A. H., Kaufman, M. T., Churchland, M. M. & Cunningham, J. P. Reorganization between preparatory and movement population responses in motor cortex. *Nature Communications* **7**, (2016).
40. Thibeault, V., Allard, A. & Desrosiers, P. The low-rank hypothesis of complex systems. *Nat. Phys.* 1–9 (2024) doi:10.1038/s41567-023-02303-0.
41. Todorov, E. Stochastic Optimal Control and Estimation Methods Adapted to the Noise Characteristics of the Sensorimotor System. *Neural Computation* **17**, 1084–1108 (2005).
42. Guigon, E. A computational theory for the production of limb movements. *Psychological Review* **130**, 23–51 (2023).
43. Maciejowski, J. M. *Predictive Control: With Constraints*. (Pearson Education, 2002).
44. Russo, A. A. *et al.* Motor Cortex Embeds Muscle-like Commands in an Untangled Population Response. *Neuron* **97**, 953-966.e8 (2018).
45. Thoroughman, K. A. & Shadmehr, R. Learning of action through adaptive combination of motor primitives. *Nature* **407**, 742–747 (2000).
46. Nashed, J. Y., Crevecoeur, F. & Scott, S. H. Influence of the behavioral goal and environmental obstacles on rapid feedback responses. *Journal of Neurophysiology* **108**, 999–1009 (2012).
47. Gallivan, J. P., Logan, L., Wolpert, D. M. & Flanagan, J. R. Parallel specification of competing sensorimotor control policies for alternative action options. *Nature Neuroscience* **19**, 320–326 (2016).
48. Česonis, J. & Franklin, D. W. Contextual cues are not unique for motor learning: Task-dependant switching of feedback controllers. *PLOS Computational Biology* **18**, e1010192 (2022).

49. Sussillo, D. & Abbott, L. F. Generating Coherent Patterns of Activity from Chaotic Neural Networks. *Neuron* **63**, 544–557 (2009).

50. Safaie, M. *et al.* Preserved neural dynamics across animals performing similar behaviour. *Nature* **623**, 765–771 (2023).

51. Sun, X. *et al.* Cortical preparatory activity indexes learned motor memories. *Nature* **602**, 274–279 (2022).

52. Ahmadi-Pajouh, M. A., Towhidkhah, F. & Shadmehr, R. Preparing to Reach: Selecting an Adaptive Long-Latency Feedback Controller. *J. Neurosci.* **32**, 9537–9545 (2012).

53. Kirk, E. A., Hope, K. T., Sober, S. J. & Sauerbrei, B. A. An output-null signature of inertial load in motor cortex. *Nat Commun* **15**, 7309 (2024).

54. Sauerbrei, B. A. *et al.* Cortical pattern generation during dexterous movement is input-driven. *Nature* **577**, 386–391 (2020).

55. Scott, S. H. A Functional Taxonomy of Bottom-Up Sensory Feedback Processing for Motor Actions. *Trends in Neurosciences* **39**, 512–526 (2016).

56. Omrani, M., Murnaghan, C. D., Pruszynski, J. A. & Scott, S. H. Distributed task-specific processing of somatosensory feedback for voluntary motor control. *eLife* **5**, e13141 (2016).

57. Michaels, J. A. *et al.* Sensory expectations shape neural population dynamics in motor circuits. 2024.12.22.629295 Preprint at https://doi.org/10.1101/2024.12.22.629295 (2025).

58. Israely, S. *et al.* Cerebellar output shapes cortical preparatory activity during motor adaptation. 2024.07.12.603354 Preprint at https://doi.org/10.1101/2024.07.12.603354 (2025).

59. Becker, M. I. & Person, A. L. Cerebellar Control of Reach Kinematics for Endpoint Precision. *Neuron* **103**, 335-348.e5 (2019).

60. Chowdhury, R. H., Glaser, J. I. & Miller, L. E. Area 2 of primary somatosensory cortex encodes kinematics of the whole arm. *eLife* **9**, e48198 (2020).

61. Guo, H., Kuang, S. & Gail, A. Sensorimotor environment but not task rule reconfigures population dynamics in rhesus monkey posterior parietal cortex. *Nat Commun* **16**, 1116 (2025).

62. Crevecoeur, F., Munoz, D. P. & Scott, S. H. Dynamic Multisensory Integration: Somatosensory Speed Trumps Visual Accuracy during Feedback Control. *The Journal of Neuroscience* **36**, 8598–8611 (2016).

63. Zimnik, A. J. & Churchland, M. M. Independent generation of sequence elements by motor cortex. *Nat Neurosci* **24**, 412–424 (2021).

64. Michaels, J. A., Dann, B. & Scherberger, H. Neural Population Dynamics during Reaching Are Better Explained by a Dynamical System than Representational Tuning. *PLoS Comput Biol* **12**, (2016).

65. Feulner, B., Perich, M. G., Miller, L. E., Clopath, C. & Gallego, J. A. A neural implementation model of feedback-based motor learning. *Nat Commun* **16**, 1805 (2025).

66. Merel, J., Botvinick, M. & Wayne, G. Hierarchical motor control in mammals and machines. *Nat Commun* **10**, 5489 (2019).

67. Chiappa, A. S. *et al.* Acquiring musculoskeletal skills with curriculum-based reinforcement learning. *Neuron* **112**, 3969-3983.e5 (2024).

68. Rodriguez, A. C., Perich, M. G., Miller, L. E. & Humphries, M. D. Motor Cortex Latent Dynamics Encode Spatial and Temporal Arm Movement Parameters Independently. *J. Neurosci.* **44**, (2024).

69. Miri, A. *et al.* Behaviorally Selective Engagement of Short-Latency Effector Pathways by Motor Cortex. *Neuron* **95**, 683-696.e11 (2017).

70. Heming, E. A., Cross, K. P., Takei, T., Cook, D. J. & Scott, S. H. Independent representations of ipsilateral and contralateral limbs in primary motor cortex. *eLife* **8**, e48190 (2019).

71. Donchin, O., Gribova, A., Steinberg, O., Bergman, H. & Vaadia, E. Primary motor cortex is involved in bimanual coordination. *Nature* **395**, 274–278 (1998).

72. Gallego, J. A. *et al.* Cortical population activity within a preserved neural manifold underlies multiple motor behaviors. *Nat Commun* **9**, 4233 (2018).

73. Perich, M. G., Gallego, J. A. & Miller, L. A Neural Population Mechanism For Rapid Learning. *Neuron* **100**, 964–976 (2018).

74. Azim, E. & Seki, K. Gain control in the sensorimotor system. *Current Opinion in Physiology* **8**, 177–187 (2019).

75. Kobak, D. & Mehring, C. Adaptation Paths to Novel Motor Tasks Are Shaped by Prior Structure Learning. *Journal of Neuroscience* **32**, 9898–9908 (2012).

76. Crevecoeur, F., Thonnard, J.-L. & Lefèvre, P. A Very Fast Time Scale of Human Motor Adaptation: Within Movement Adjustments of Internal Representations during Reaching. *eNeuro* **7**, ENEURO.0149-19.2019 (2020).

77. Sadtler, P. T. *et al.* Neural constraints on learning. *Nature* **512**, 423–426 (2014).

78. Ostry, D. J. & Gribble, P. L. Sensory Plasticity in Human Motor Learning. *Trends in Neurosciences* **39**, 114–123 (2016).

79. Singh, K. & Scott, S. H. A motor learning strategy reflects neural circuitry for limb control. *Nature Neuroscience* **6**, 399–403 (2003).

80. Pemberton, J., Chadderton, P. & Costa, R. P. Cerebellar-driven cortical dynamics can enable task acquisition, switching and consolidation. *Nat Commun* **15**, 10913 (2024).

81. O'Shea, D. J. *et al.* Direct neural perturbations reveal a dynamical mechanism for robust computation. 2022.12.16.520768 Preprint at https://doi.org/10.1101/2022.12.16.520768 (2022).

82. Vahidi, P., Sani, O. G. & Shanechi, M. M. Modeling and dissociation of intrinsic and input-driven neural population dynamics underlying behavior. *Proc Natl Acad Sci U S A* **121**, e2212887121 (2024).

83. Park, J. *et al.* Conjoint specification of action by neocortex and striatum. 2023.10.04.560957 Preprint at https://doi.org/10.1101/2023.10.04.560957 (2023).

84. Takei, T., Lomber, S. G., Cook, D. J. & Scott, S. H. Transient deactivation of dorsal premotor cortex or parietal area 5 impairs feedback control of the limb in macaques. *Current Biology* **31**, 1476-1487.e5 (2021).

85. Hennequin, G., Vogels, T. P. & Gerstner, W. Optimal Control of Transient Dynamics in Balanced Networks Supports Generation of Complex Movements. *Neuron* **82**, 1394–1406 (2014).

86. Heald, J. B., Lengyel, M. & Wolpert, D. M. Contextual inference underlies the learning of sensorimotor repertoires. *Nature* **600**, 489–493 (2021).

87. Elsayed, G. F. & Cunningham, J. P. Structure in neural population recordings: an expected byproduct of simpler phenomena? *Nature Neuroscience* **20**, 1310–1318 (2017).

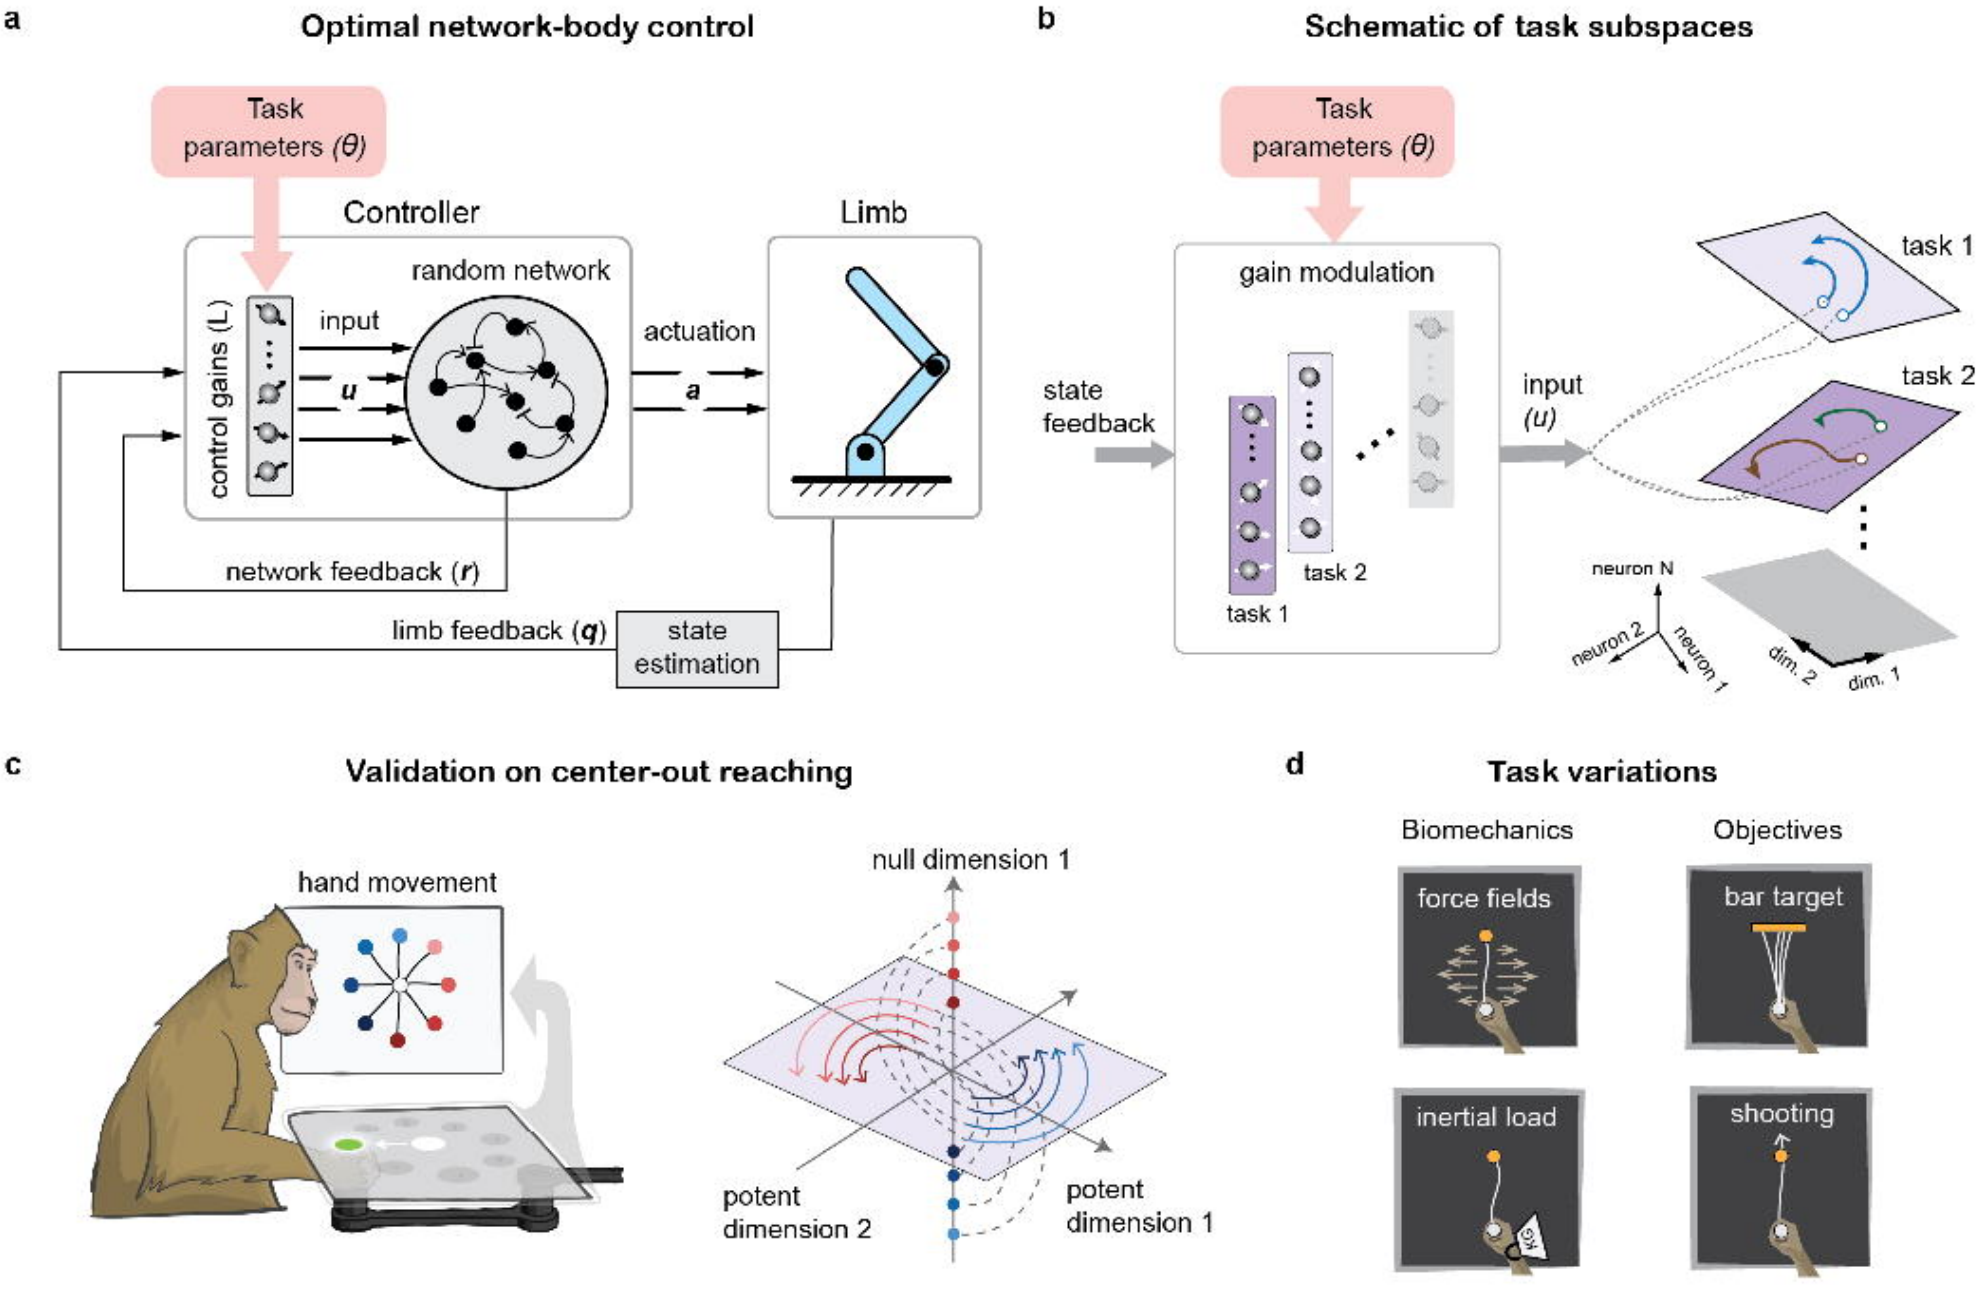
a
Optimal network-body control
Task parameters (θ)
Controller
random network
input
u
control gains (L)
actuation
a
Limb
network feedback (r)
limb feedback (q)
state estimation
b
Schematic of task subspaces
Task parameters (θ)
gain modulation
state feedback
task 1
task 2
input (u)
task 1
task 2
neuron N
neuron 2
neuron 1
dim. 2
dim. 1
c
Validation on center-out reaching
hand movement
null dimension 1
potent dimension 2
potent dimension 1
d
Task variations
Biomechanics
Objectives
force fields
bar target
inertial load
shooting

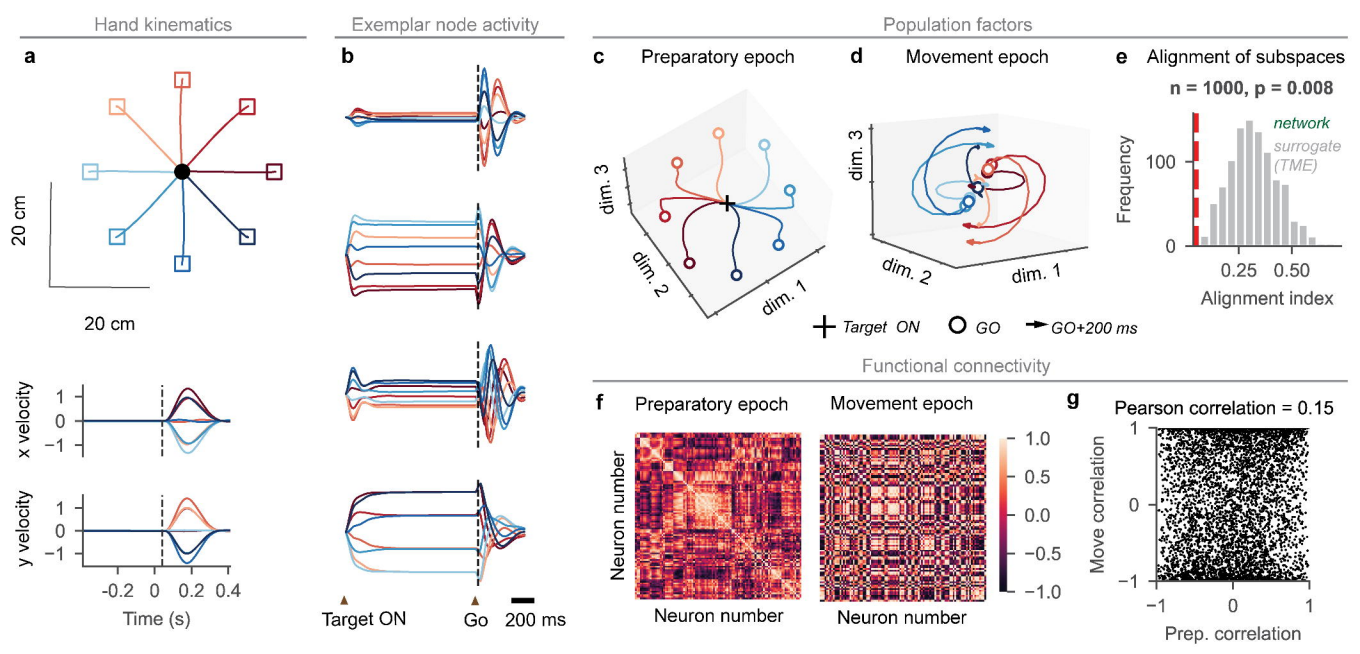
Hand kinematics
Exemplar node activity
Population factors
a
b
c
Preparatory epoch
d
Movement epoch
e
Alignment of subspaces
n = 1000, p = 0.008
network
surrogate (TME)
Frequency
100
0
0.25
0.50
Alignment index
20 cm
20 cm
x velocity
y velocity
1
0
−1
-0.2
0
0.2
0.4
Time (s)
Target ON
Go
200 ms
dim. 3
dim. 2
dim. 1
Target ON
GO
GO+200 ms
Functional connectivity
f
Preparatory epoch
Movement epoch
Neuron number
Neuron number
Neuron number
1.0
0.5
0.0
−0.5
−1.0
g
Pearson correlation = 0.15
Move correlation
Prep. correlation
1
0
−1

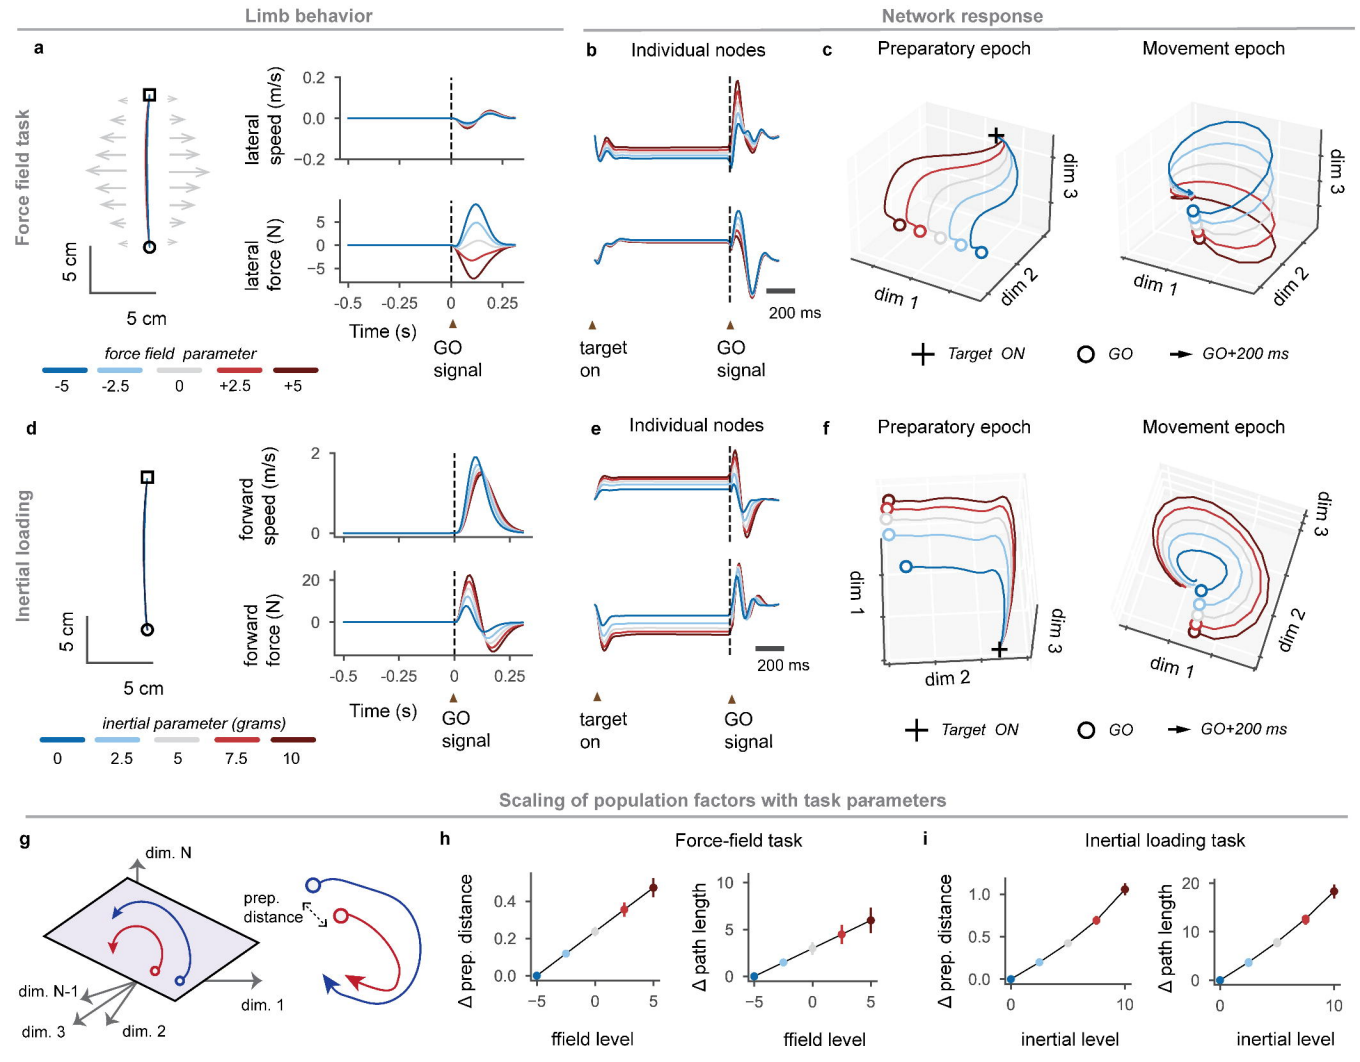
Limb behavior
Network response
a
Force field task
5 cm
5 cm
lateral speed (m/s)
0.2
0.0
−0.2
lateral force (N)
5
0
−5
-0.5
-0.25
0
0.25
Time (s)
GO signal
force field parameter
-5
-2.5
0
+2.5
+5
b
Individual nodes
200 ms
target on
GO signal
c
Preparatory epoch
Movement epoch
dim 1
dim 2
dim 3
Target ON
GO
GO+200 ms
d
Inertial loading
5 cm
5 cm
forward speed (m/s)
2
0
forward force (N)
20
0
-0.5
-0.25
0
0.25
Time (s)
GO signal
inertial parameter (grams)
0
2.5
5
7.5
10
e
Individual nodes
200 ms
target on
GO signal
f
Preparatory epoch
Movement epoch
dim 1
dim 2
dim 3
Target ON
GO
GO+200 ms
Scaling of population factors with task parameters
g
dim. N
prep. distance
dim. N-1
dim. 1
dim. 3
dim. 2
h
Force-field task
Δ prep. distance
0.4
0.2
0.0
−5
0
5
ffield level
Δ path length
10
5
0
−5
0
5
ffield level
i
Inertial loading task
Δ prep. distance
1.0
0.5
0.0
0
10
inertial level
Δ path length
20
10
0
0
10
inertial level

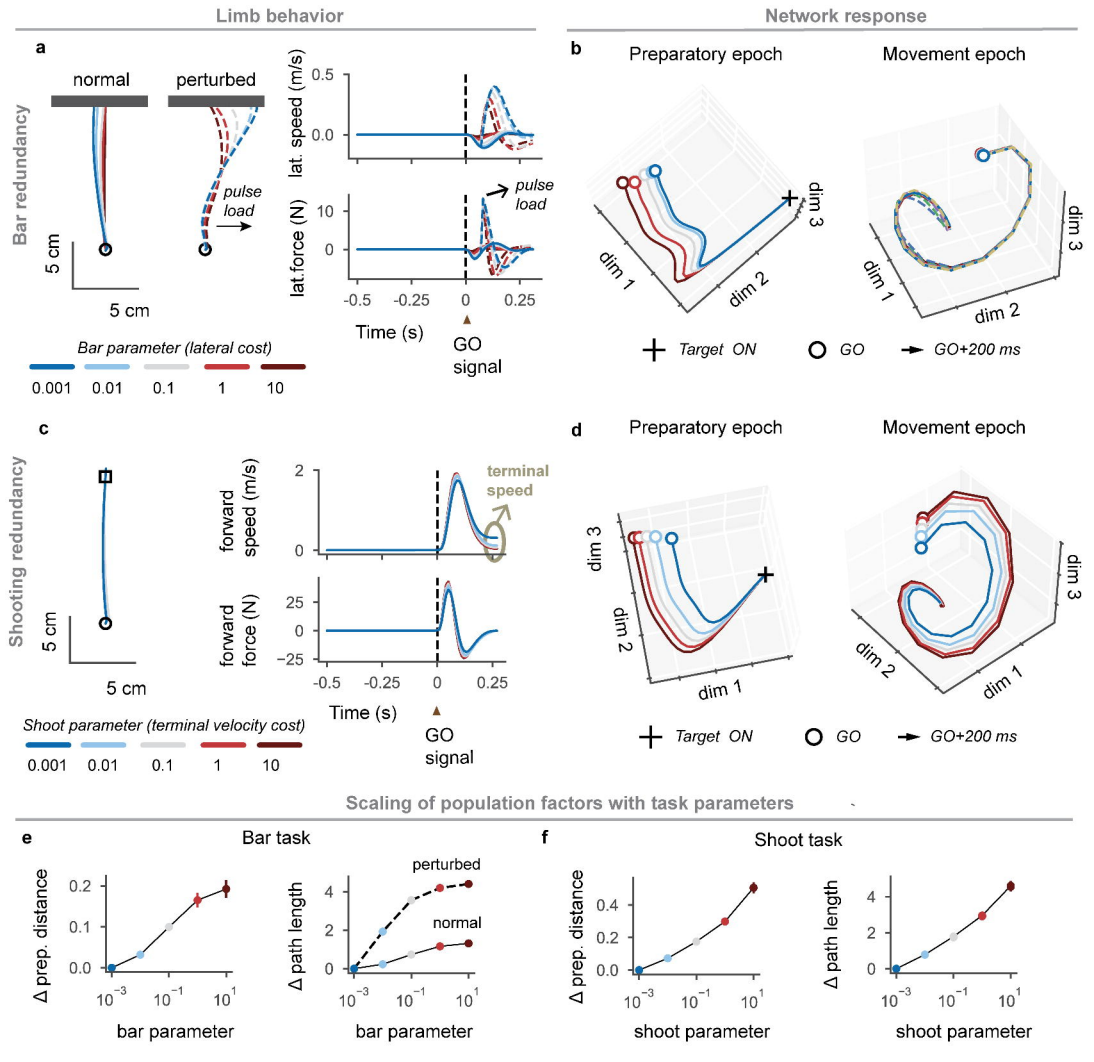
Limb behavior
Network response
a
normal
perturbed
Bar redundancy
pulse load
5 cm
5 cm
lat. speed (m/s)
lat.force (N)
pulse load
Time (s)
GO signal
Bar parameter (lateral cost)
0.001
0.01
0.1
1
10
b
Preparatory epoch
Movement epoch
dim 1
dim 2
dim 3
Target ON
GO
GO+200 ms
c
Shooting redundancy
5 cm
5 cm
forward speed (m/s)
forward force (N)
terminal speed
Time (s)
GO signal
Shoot parameter (terminal velocity cost)
0.001
0.01
0.1
1
10
d
Preparatory epoch
Movement epoch
dim 1
dim 2
dim 3
Target ON
GO
GO+200 ms
Scaling of population factors with task parameters
e
Bar task
Δ prep. distance
bar parameter
Δ path length
perturbed
normal
bar parameter
f
Shoot task
Δ prep. distance
shoot parameter
Δ path length
shoot parameter

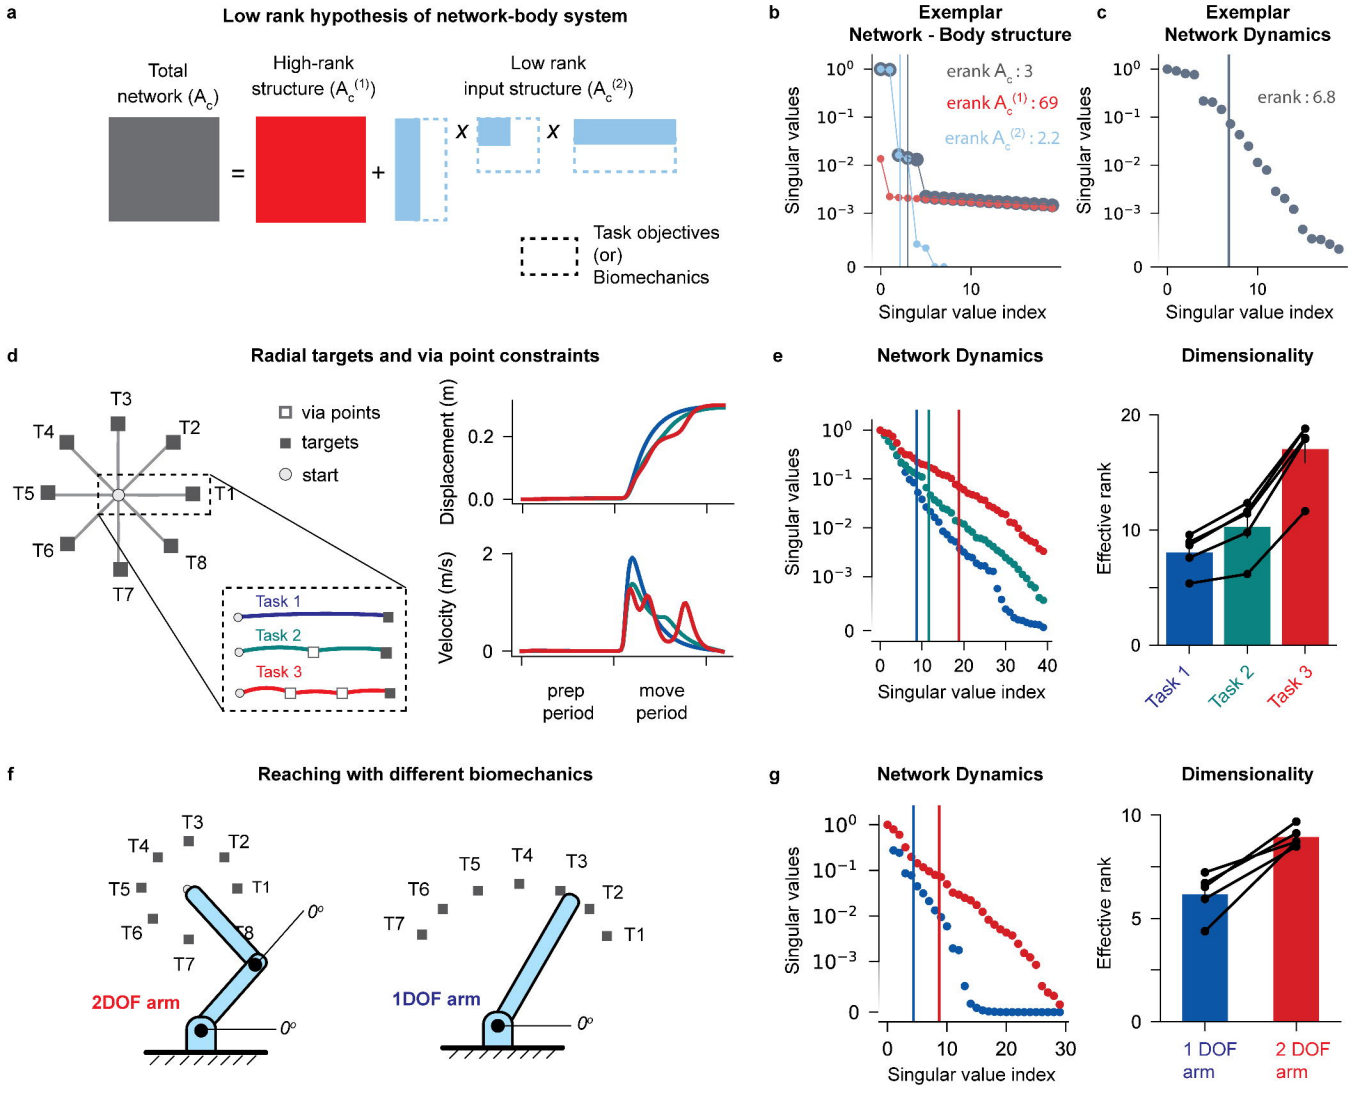
a
Low rank hypothesis of network-body system
Total network (A_c)
High-rank structure ($A_c^{(1)}$)
Low rank input structure ($A_c^{(2)}$)
Task objectives (or) Biomechanics
b
Exemplar Network - Body structure
erank A_c : 3
erank $A_c^{(1)}$: 69
erank $A_c^{(2)}$: 2.2
Singular values
Singular value index
c
Exemplar Network Dynamics
erank : 6.8
Singular values
Singular value index
d
Radial targets and via point constraints
T1
T2
T3
T4
T5
T6
T7
T8
via points
targets
start
Task 1
Task 2
Task 3
Displacement (m)
Velocity (m/s)
prep period
move period
e
Network Dynamics
Singular values
Singular value index
Dimensionality
Effective rank
Task 1
Task 2
Task 3
f
Reaching with different biomechanics
2DOF arm
1DOF arm
0°
g
Network Dynamics
Singular values
Singular value index
Dimensionality
Effective rank
1 DOF arm
2 DOF arm