Contents lists available at ScienceDirect

Neural Networks

journal homepage: www.elsevier.com/locate/neunet

Theoretical limits on the speed of learning inverse models explain the rate of adaptation in arm reaching tasks

Victor R. Barradas [a,*], Yasuharu Koike [a], Nicolas Schweighofer [b]

[a] Institute of Innovative Research, Tokyo Institute of Technology, 4259 R2-16 Nagatsuta-cho, Midori-ku, Yokohama, Kanagawa 226-8503, Japan
[b] Biokinesiology and Physical Therapy, University of Southern California, 1540 Alcazar Street, CHP 155, Los Angeles, CA 90089-9006, USA

ARTICLE INFO

Keywords:
Motor learning
Inverse models
Gradient descent
Isometric tasks

ABSTRACT

An essential aspect of human motor learning is the formation of inverse models, which map desired actions to motor commands. Inverse models can be learned by adjusting parameters in neural circuits to minimize errors in the performance of motor tasks through gradient descent. However, the theory of gradient descent establishes limits on the learning speed. Specifically, the eigenvalues of the Hessian of the error surface around a minimum determine the maximum speed of learning in a task. Here, we use this theoretical framework to analyze the speed of learning in different inverse model learning architectures in a set of isometric arm-reaching tasks. We show theoretically that, in these tasks, the error surface and, thus the speed of learning, are determined by the shapes of the force manipulability ellipsoid of the arm and the distribution of targets in the task. In particular, rounder manipulability ellipsoids generate a rounder error surface, allowing for faster learning of the inverse model. Rounder target distributions have a similar effect. We tested these predictions experimentally in a quasi-isometric reaching task with a visuomotor transformation. The experimental results were consistent with our theoretical predictions. Furthermore, our analysis accounts for the speed of learning in previous experiments with incompatible and compatible virtual surgery tasks, and with visuomotor rotation tasks with different numbers of targets. By identifying aspects of a task that influence the speed of learning, our results provide theoretical principles for the design of motor tasks that allow for faster learning.

1. Introduction

Humans are capable of producing fast and accurate movements to skillfully execute a variety of motor tasks. To accomplish such tasks, the central nervous system (CNS) uses feedforward mechanisms, as it cannot rely solely on delayed sensory feedback to guide the execution of movements (Kawato et al., 1987; Schmidt et al., 1979). Growing evidence indicates that the CNS generates feedforward motor commands via inverse models, which compute motor commands to achieve a desired sensory consequence given the current state of the body and environment (Golub et al., 2015; Kawato, 1999; Pierella et al., 2019; Schweighofer et al., 1998; Shidara et al., 1993; Wolpert et al., 1998).

Inverse models must be able to adapt to changes in the environment or the body to maintain successful task execution. This adaptation faces two main challenges: (1) errors in task space do not directly inform how motor commands should be adjusted to eliminate the errors, and (2) the models are generally not unique because of the redundancy of the motor system (more motor commands than task variables). Accordingly, several computational models such as direct inverse modeling (Kuperstein, 1988; Sanger, 2004), distal learning (Jordan & Rumelhart, 1992), and feedback error learning (Kawato, 1990) have been proposed as mechanisms to solve these challenges. These models differ in their architectures and mechanisms to adequately relate task errors to changes in motor commands. However, despite differences in their theoretical underpinnings, these models are all learned based on attempts to minimize an error quantity, such as performance error (distal learning) or motor error (direct inverse modeling and feedback error learning). Thus, from a computational perspective, the inverse learning problem can be formulated as a function fitting problem (mapping from desired motor outcomes to motor commands), where the strengths of neural connections in the inverse model circuit are the parameters that are tuned to fit the function.

It is now well accepted that the CNS encodes movement errors in the cerebellum, suggesting that the errors are used to drive motor adaptation in a supervised learning fashion (Doya, 1999; Gilbert & Thach, 1977; Ito, 1984; Kawato et al., 2011; Kitazawa et al., 1998; Schlerf et al.,

[*] Corresponding author.
E-mail address: barradas.v.aa@m.titech.ac.jp (V.R. Barradas).

https://doi.org/10.1016/j.neunet.2023.10.049
Received 31 October 2022; Received in revised form 8 September 2023; Accepted 30 October 2023
Available online 31 October 2023
0893-6080/© 2023 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY-NC license (http://creativecommons.org/licenses/by-nc/4.0/).

2012; Schweighofer et al., 1998, 2004). This implies that the CNS uses some form of gradient descent on the errors to tune the parameters in the inverse model. In gradient descent, the parameters are adjusted in small steps in the direction which minimizes the error locally, given by first order derivatives of the error function with respect to the parameters. Thus, learning can be visualized as a path of successive steps in the multi-dimensional space of inverse model parameters. This path lies on a surface composed by the values of the error at each point in the inverse model parameter space.

Theoretical results in statistical learning establish that the shape of the error surface limits the speed of learning (Bishop, 1995). Round or isotropic surfaces allow faster learning because the gradient points in the general direction of the minimum of the error function, resulting in a relatively straight path to the minimum (see also Schweighofer et al. (2001), Spanne & Jörntell (2015)). In contrast, less round or anisotropic surfaces bring about slower learning because in large regions of the parameter space the gradient is misaligned with the direction of the minimum. The shape of the error surface can be described locally by the eigenvalues of a Hessian matrix, which contains the second order derivatives of the error with respect to the inverse model parameters (Bishop, 1995).

Here, we provide a theoretical analysis of the speed of learning an inverse model using gradient-following optimization by analyzing the relationship between the Hessian matrix of the learning system, the goals of the task, and the forward physics of the musculoskeletal system. In particular, we analyze learning both in isometric arm tasks with visuomotor transformations, which change the relationship between produced forces and visual feedback, and in musculoskeletal transformations, which change the role of muscles in force production. In these simple motor learning tasks, the CNS can be conceptualized as an inverse model represented as a feedforward neural network. In addition, the visuomotor and musculoskeletal transformations and the musculoskeletal system can be mathematically integrated in the forward physics of the arm as a linear mapping (Berger et al., 2013; Valero-Cuevas, 2009). As a result, the system composed of the inverse model and the forward physics of the arm (the musculoskeletal system and a transformation) can be viewed mathematically as a single feedforward network. We demonstrate that the Hessian matrix for this composite network can be readily computed, allowing us to make predictions about the speed of learning an inverse model for these tasks.

Analyzing the Hessian of the composite network allows us to identify how the goals and forward physics of the task influence the speed of learning under different transformations. In particular, we show that both 1) changes in the force manipulability ellipse of the arm (the ability to generate end-effector forces given a joint configuration (Yoshikawa, 1985)) brought about by different transformations, and 2) the shape of the distribution of target forces can limit the speed of learning.

In addition, we show that our theoretical results are valid for a range of assumptions about learning inverse models. Specifically, our results generalize to different learning architectures (distal learning and direct inverse learning), to different inverse model architectures (radial basis function networks and multi-layer perceptron networks), and to different gradient-following algorithms (backpropagation and node perturbation (Werfel et al., 2005; Williams, 1992)). Importantly, we also show that our results are compatible with the theory of optimal control in motor learning, in which the objective of learning is the minimization of both errors and effort (Emken et al., 2007; Todorov & Jordan, 2002).

We then confirm our theoretical predictions via two experiments that tested the role of the force manipulability ellipse and the target distribution on the speed of learning a quasi-isometric task in a virtual environment. In the first experiment, we compare the speed of learning a visuomotor rotation with isotropic or anisotropic manipulability ellipses. In the second experiment, we compare the speed of learning a visuomotor rotation with targets in a circular or elliptical configuration.

Finally, we use our theoretical framework to account for differences in the speed of learning in two previous studies. First, we study a virtual surgery task, a type of musculoskeletal transformation in which the pulling forces of muscles in a mapping between muscle activations and end-point forces are changed, simulating tendon transfer surgeries (Berger et al., 2013). Our results suggest that the shapes of the manipulability ellipses of the system after applying the virtual surgeries account, at least in part, for the differences in learning speed. Second, we address the speed of learning in visuomotor rotation experiments with different number of targets (Krakauer et al., 2000). We show that our framework accounts for these differences in learning speed because the number of targets has a direct influence on the shape of the distribution of targets.

2. Methods

2.1. Computational modeling framework

2.1.1. Overview

We propose a computational framework to study the speed of motor learning during isometric arm reaching tasks, including visuomotor transformations and compatible and incompatible virtual surgeries (Berger et al., 2013). We focus exclusively on feedforward control, which enables us to model the tasks as a static mapping between target and output forces, and to ignore transient forces. In the main text, we derive theoretical results based on the distal learning architecture (Jordan & Rumelhart, 1992) (Fig. 1a), in which a forward model of the arm is used to estimate the relationship between task errors and motor commands to learn the inverse model. We then show in Appendix A that our framework generalizes to direct inverse modeling (Kuperstein, 1988; Sanger, 2004). Note that feedback error learning is not directly applicable to this feedforward control isometric task.

For simplicity, we use a model of the upper extremity that simulates production of isometric planar forces to reach different targets in a

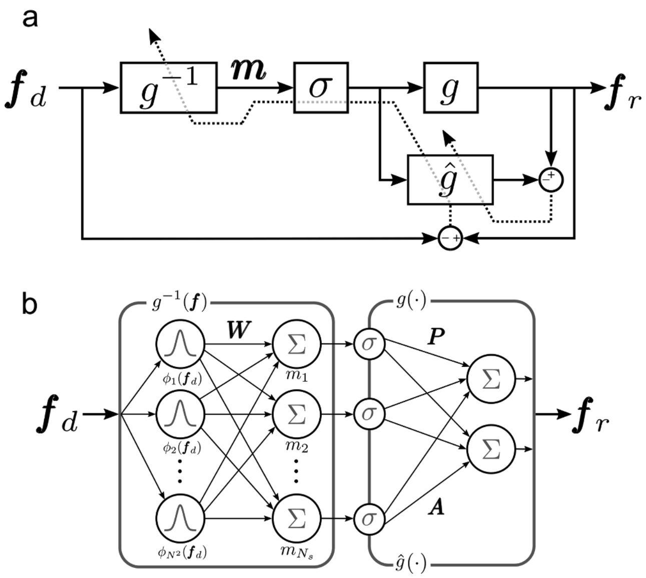

Fig. 1. Structure of computational model for the distal learning of isometric arm tasks in which the inverse model is implemented with an RBF network a. Model diagram. The CNS learns an inverse model g^{-1} of the motor system g by using a forward model $\hat{g}$ to update g^{-1} based on the minimization of the error $\boldsymbol{f}_d$ - $\boldsymbol{f}_r$. b. Structure of the inverse model g^{-1} and the motor system g. The input to g^{-1} is an RBF representation of a 2D force. The inverse model is a single-layer network and is fully defined by $\boldsymbol{W}$, the weights of the connections between the input and the output units. The output of g^{-1} is the motor command $\boldsymbol{m}$, which is the input to g. The motor system g and the forward model $\hat{g}$ are linear transformations $\boldsymbol{P}$ and $\boldsymbol{A}$, respectively, of $\boldsymbol{m}$ after a non-linear transformation σ. Note that the CNS and motor system as a whole can therefore be seen as a two-layer network with linear outputs.

virtual environment. Because the tasks are isometric, the biomechanics of the arm can be approximated adequately as a linear mapping between motor commands and output forces (Barradas et al., 2020; Berger et al., 2013; Valero-Cuevas, 2009; Zhou & Rymer, 2004).

2.1.2. Model structure and implementation

The model learns an inverse model g^{-1} of g, the motor system after applying a transformation T, which represents a visuomotor transformation or a virtual surgery. The inverse model g^{-1} receives a target 2-D force $\boldsymbol{f}_\mathrm{d}$ to generate a motor command $\boldsymbol{m}$, which can represent joint torques, muscle activations or muscle synergy activations (Barradas et al., 2020). Then, the motor command $\boldsymbol{m}$ propagates through the forward physics of the motor system g, producing a realized force $\boldsymbol{f}_\mathrm{r}$ (Fig. 1a). Note that, despite its biological implausibility, we also consider the case in which $\boldsymbol{m}$ represents joint torques because it allows a totally analytical derivation of the properties of the speed of learning in the system, which provides insight into the more complex cases where $\boldsymbol{m}$ represents muscle or synergy activations. Specifically, the motor system g is represented as:

$$g(\boldsymbol{m}) = \boldsymbol{P}\sigma(\boldsymbol{m}) = \boldsymbol{f}_\mathrm{r}, \tag{1}$$

where $\boldsymbol{P}$ represents the forward physics of the isometric task, and $\boldsymbol{m}$ is an N_m-dimensional motor command vector produced by the inverse model (joint torques, muscle activations or muscle synergy activations). Because the task is isometric, forces can be modeled as a linear function of the motor command $\boldsymbol{m}$, or the transformed motor command $\sigma(\boldsymbol{m})$ (Barradas et al., 2020; Berger et al., 2013; Valero-Cuevas, 2009; Zhou & Rymer, 2004). Thus, if the motor system is unperturbed, then $\boldsymbol{P} = \boldsymbol{M}$, where $\boldsymbol{M}$ is a matrix containing a linear mapping between motor commands and planar forces ($2 \times N_\mathrm{m}$) (Barradas et al., 2020; Berger et al., 2013). A transformation T, such as a visuomotor transformation or virtual surgery, can alter the forward physics of the task such that $\boldsymbol{P} = T(\boldsymbol{M})$ (see Sections 2.2.2 *Motor command to force mapping*, 2.2.3 *Anisotropic scaling task*, and 2.2.5 *Virtual surgery task* for details). If $\boldsymbol{m}$ represents joint torques, the activation function $\sigma(\boldsymbol{m})$ is a linear function. Otherwise, it is a non-linear function that maps each element of $\boldsymbol{m}$ into a non-negative muscle or synergy activation. The non-linear function can be a sigmoid or a rectified linear unit (RELU) function.

We first consider the case in which the inverse model g^{-1} is learned via a radial basis function (RBF) network, as these kind of networks are biologically plausible (Thoroughman & Shadmehr, 2000), and simplify theoretical derivations. With this inverse model architecture, and for isometric tasks, the structures and connections of g^{-1} and g allow us to view g^{-1} and g as a single neural network with one hidden layer and linear outputs (Fig. 1b). We note, however, that our approach is not limited to the RBF architecture, and show that it extends to multi-layer perceptron networks, for instance (see Section 2.1.5 *Generalizability of the results to muscular models, realistic error function, and neural network structure*). In the RBF network, the inverse model is given by:

$$g^{-1}(\boldsymbol{f}) = \boldsymbol{W}\boldsymbol{\phi}(\boldsymbol{f}) = \boldsymbol{m} \tag{2}$$

$$\phi_i(\boldsymbol{f}) = \exp\left(- \| \boldsymbol{\Sigma}(\boldsymbol{f}_{ci} - \boldsymbol{f}) \|^2 \right), \tag{3}$$

where $\boldsymbol{W}$ is a matrix of weights ($N_\mathrm{m} \times N^2$), $\boldsymbol{\phi}(\boldsymbol{f})$ is a vector of N^2 two-dimensional RBFs $\phi_i(\boldsymbol{f})$ evaluated at $\boldsymbol{f}$, $\boldsymbol{f}_{ci}$ are the centers of the RBFs placed on a $N \times N$ square grid centered on the origin of the force space, and $\boldsymbol{\Sigma}$ determines the shape of the RBFs.

The distal learning model includes an internal forward model $\hat{g}$, which is used to estimate the relationship between task errors and changes in the motor commands produced by g^{-1} (Jordan & Rumelhart, 1992). The forward model $\hat{g}$ uses an efferent copy of $\boldsymbol{m}$ to produce $\widehat{\boldsymbol{f}}_\mathrm{r}$, an estimate of $\boldsymbol{f}_\mathrm{r}$, to compute a sensory prediction error $\boldsymbol{f}_\mathrm{r} - \widehat{\boldsymbol{f}}_\mathrm{r}$. Since the forward physics of the arm are linear, we define $\hat{g}$ as a linear operator on $\sigma(\boldsymbol{m})$:

$$\widehat{g}(\boldsymbol{m}) = \boldsymbol{A}\sigma(\boldsymbol{m}) = \widehat{\boldsymbol{f}}_\mathrm{r}, \tag{4}$$

where $\boldsymbol{A}$ has the same dimensions as $\boldsymbol{P}$, so that $\hat{g}$ has the same mathematical formulation as g (Eq. (1)).

2.1.3. Learning of the inverse model

Learning in the isometric tasks aims to minimize the task performance error given by the cost function

$$J = \frac{1}{2}\| \boldsymbol{f}_\mathrm{d} - \boldsymbol{f}_\mathrm{r} \|^2. \tag{5}$$

In distal learning, the learner uses the forward model $\hat{g}$ to compute an approximate gradient of J because the exact gradient of J is not available, as $\boldsymbol{f}_\mathrm{r}$ is a function of g, which is generally unknown to the learner. Assuming that $\hat{g}$ is a good enough representation of g, learning converges to the minimum of J (Jordan & Rumelhart, 1992). As stated above, we can view g^{-1} and $\hat{g}$ as a composite neural network (Fig. 1b). Therefore, the approximate gradient for the network can be computed using methods such as backpropagation, stochastic methods like weight or node perturbation (Mazzoni et al., 1991; Werfel et al., 2005), or other reinforcement learning-based methods like REINFORCE (Williams, 1992). Here, we used backpropagation to illustrate the properties of the speed of learning under gradient-following algorithms.

The learning procedure starts with the forward propagation of a given $\boldsymbol{f}_\mathrm{d}$ through g^{-1} and the actual motor system g, resulting in $\boldsymbol{f}_\mathrm{r}$ and a task performance error $\boldsymbol{f}_\mathrm{d}$ - $\boldsymbol{f}_\mathrm{r}$. Then, the task performance error is backpropagated through $\hat{g}$ and g^{-1}, which provides $\widehat{\nabla}_{\boldsymbol{W}}J$, an approximation of $\nabla_{\boldsymbol{W}}J$, the gradient of J with respect to the weights $\boldsymbol{W}$ of g^{-1} (Eq. (2)). Lastly, $\widehat{\nabla}_{\boldsymbol{W}}J$ is used to update $\boldsymbol{W}$:

$$\boldsymbol{W}_{n+1} = \boldsymbol{W}_n - \eta_\mathrm{I}\widehat{\nabla}_{\boldsymbol{W}}J, \tag{6}$$

where $\boldsymbol{W}_n$ is the weight matrix at the current trial, $\boldsymbol{W}_{n+1}$ is the updated weight matrix, and η_I is the learning rate parameter. Notice that the forward model $\hat{g}$ is not updated in this backward pass; $\hat{g}$ is only used to compute $\widehat{\nabla}_{\boldsymbol{W}}J$, the approximation of $\nabla_{\boldsymbol{W}}J$.

We assume that the forward model is learned at a much faster rate than the inverse model, and that learning of the forward model is perfect. That is, we set $\boldsymbol{A} = \boldsymbol{P} = T(\boldsymbol{M})$ according to the transformation T used in the task. Therefore, the approximation of the gradient becomes equal to the actual gradient, $\widehat{\nabla}_{\boldsymbol{W}}J = \nabla_{\boldsymbol{W}}J$. This is justified because learning forward models is much simpler than learning inverse models, and inverse models can be learned via distal learning even when using imperfect forward models (Jordan & Rumelhart, 1992). This assumption allows us to apply the methods described in Section 2.1.4 to estimate the limits on the speed of learning the inverse model g^{-1} when the forward physics are subject to T.

2.1.4. Estimation of limit on learning speed

In this section we demonstrate that the speed of learning in the isometric task is limited by the shapes of the force manipulability ellipse and the distribution of targets in the task.

The cost function in Eq. (5) can be approximated locally around its minimum as a second order Taylor expansion (Bishop, 1995):

$$J(\boldsymbol{w}) = J(\boldsymbol{w}^*) + \frac{1}{2}(\boldsymbol{w} - \boldsymbol{w}^*)^\mathrm{T}\boldsymbol{H}(\boldsymbol{w} - \boldsymbol{w}^*) \tag{7}$$

$$\boldsymbol{H}_{ij} = \frac{\partial J}{\partial w_i \partial w_j}, \tag{8}$$

where $\boldsymbol{w}$ is a vector of weights, $\boldsymbol{w}^*$ is the vector of weights at the minimum of cost function J, and $\boldsymbol{H}$ is the Hessian of J, a matrix of second order derivatives of J with respect to $\boldsymbol{w}$. In this second order approximation, the gradient of J is $\nabla_{\boldsymbol{W}}J = \boldsymbol{H}(\boldsymbol{w} - \boldsymbol{w}^*)$. Therefore, the update equation of $\boldsymbol{w}$ using gradient descent can be approximated by the following learning rule:

$$\boldsymbol{w}_{n+1} = \boldsymbol{w}_n - \eta \boldsymbol{H}(\boldsymbol{w}_n - \boldsymbol{w}^*), \tag{9}$$

where η is the learning rate parameter. Previous work has demonstrated that in the vicinity of $\boldsymbol{w}$*, there is an upper limit on the value of η, beyond which updates to $\boldsymbol{w}$ become unstable. This value of η also determines the slowest time constant of learning, which limits the speed of convergence to $\boldsymbol{w}$* and is a function of $\lambda_{\boldsymbol{H}\mathrm{min}}/\lambda_{\boldsymbol{H}\mathrm{max}}$, the ratio of the smallest to largest eigenvalue of $\boldsymbol{H}$ (Bishop, 1995) (see Appendix C for a derivation of this result). From a geometric point of view, $\boldsymbol{H}$ describes the shape of the cost function J around $\boldsymbol{w}$*. Namely, it can be shown that the contours of constant value of J are ellipsoids, and that the sizes of the axes of these ellipsoids are proportional to the square roots of the non-zero eigenvalues of $\boldsymbol{H}$. Therefore, if both eigenvalues $\lambda_{\boldsymbol{H}\mathrm{min}}$ and $\lambda_{\boldsymbol{H}\mathrm{max}}$ have similar magnitudes, the shape of the cost function J is relatively round, allowing faster convergence to the minimum. On the other hand, a large disparity in the magnitudes of $\lambda_{\boldsymbol{H}\mathrm{min}}$ and $\lambda_{\boldsymbol{H}\mathrm{max}}$ corresponds to a highly anisotropic J, producing slower convergence to the minimum (see Appendix C for an illustration of this effect). Importantly, because the ratio $\lambda_{\boldsymbol{H}\mathrm{min}}/\lambda_{\boldsymbol{H}\mathrm{max}}$ is a descriptor of the shape of J, it characterizes the slowest rate of convergence independently of the value of η.

Based on these theoretical results relating the speed of learning and the Hessian $\boldsymbol{H}$ of the learning system, we aimed to derive $\boldsymbol{H}$ and its eigenvalues in the distal learning system for isometric tasks. In particular, we investigate how visuomotor and musculoskeletal transformations affect the structure of $\boldsymbol{H}$ and its eigenvalues, allowing us to estimate the limits on the speed of learning given by the slowest time constant of learning for each task. For mathematical simplicity, we first consider the case in which the activation function $\sigma(\boldsymbol{m})$ is linear, that is $\sigma(\boldsymbol{m}) = \boldsymbol{m}$. This can be interpreted as $\boldsymbol{m}$ representing a joint torque vector, and $\boldsymbol{P}$ representing a mapping between joint torques and end-point forces of the limb (i.e. the inverse Jacobian of the limb). As indicated in Section 2.1.3 *Learning of the inverse model*, we assume that $\hat{g}$ learns a perfect model of the perturbation as soon as the perturbation is introduced, namely, $\hat{g}$ becomes equal to g. This allows us to compute the gradient $\nabla_W J$ and the Hessian $\boldsymbol{H}$ analytically, and to estimate the limits on the speed of learning in terms of the goals and forward physics of the tasks.

Using Eqs. (2) and (5), the cost function is:

$$J = \sum_i \frac{1}{2} \| \boldsymbol{f}_{\mathrm{d}i} - \boldsymbol{P}\boldsymbol{W}\boldsymbol{\phi}(\boldsymbol{f}_{\mathrm{d}i}) \|^2, \tag{10}$$

where $\boldsymbol{f}_{\mathrm{d}i}$ corresponds to each target in the task, with $i = 1, \ldots., N_{\mathrm{T}}$, and N_{T} is the number of targets. Using standard matrix algebra (see Appendix D), the gradient $\nabla_W J$ is

$$\frac{\partial J}{\partial \boldsymbol{W}} = \nabla_W J = \sum_i \left[\left(\boldsymbol{P}^{\mathrm{T}}\boldsymbol{P}\right) \boldsymbol{W} \left(\boldsymbol{\phi}(\boldsymbol{f}_{\mathrm{d}i}) \boldsymbol{\phi}(\boldsymbol{f}_{\mathrm{d}i})^{\mathrm{T}} \right) - \boldsymbol{P}^{\mathrm{T}} \boldsymbol{f}_{\mathrm{d}i} \boldsymbol{\phi}(\boldsymbol{f}_{\mathrm{d}i})^{\mathrm{T}} \right], \tag{11}$$

and the Hessian $\boldsymbol{H}_{\boldsymbol{W}}$ is

$$\boldsymbol{H}_W = \frac{\partial^2 J}{\partial \boldsymbol{W}^2} = \sum_i \left[\left(\boldsymbol{P}^{\mathrm{T}}\boldsymbol{P}\right) \otimes \left(\boldsymbol{\phi}(\boldsymbol{f}_{\mathrm{d}i}) \boldsymbol{\phi}(\boldsymbol{f}_{\mathrm{d}i})^{\mathrm{T}} \right) \right], \tag{12}$$

where the symbol ⊗ denotes the Kronecker product. The derivation of Eqs. (11) and (12) is provided in Appendix D. We now use the associative property of the Kronecker product:

$$\boldsymbol{H}_W = \left(\boldsymbol{P}^{\mathrm{T}}\boldsymbol{P}\right) \otimes \sum_i \left[\boldsymbol{\phi}(\boldsymbol{f}_{\mathrm{d}i}) \boldsymbol{\phi}(\boldsymbol{f}_{\mathrm{d}i})^{\mathrm{T}} \right]. \tag{13}$$

The term on the right side of the Kronecker product is a scaled sample covariance matrix of the RBF representation of targets $\boldsymbol{\phi}(\boldsymbol{f}_{\mathrm{d}i})$:

$$\boldsymbol{K}_{\boldsymbol{\phi\phi}} = \frac{1}{N_{\mathrm{T}}} \sum_i \left[\boldsymbol{\phi}(\boldsymbol{f}_{\mathrm{d}i}) \boldsymbol{\phi}(\boldsymbol{f}_{\mathrm{d}i})^{\mathrm{T}} \right], \tag{14}$$

where N_{T} is the number of targets in the task. Therefore:

$$\boldsymbol{H}_{\mathrm{W}} = \left(\boldsymbol{P}^{\mathrm{T}}\boldsymbol{P}\right) \otimes \left(N_{\mathrm{T}}\boldsymbol{K}_{\boldsymbol{\phi\phi}}\right) = N_{\mathrm{T}}\left(\boldsymbol{P}^{\mathrm{T}}\boldsymbol{P} \otimes \boldsymbol{K}_{\boldsymbol{\phi\phi}}\right) \ . \tag{15}$$

The eigenvalues of the Kronecker product of two matrices are the set of products between each eigenvalue of the first matrix and each eigenvalue of the second matrix (Horn & Johnson, 1994). That is, $\lambda_{\boldsymbol{H}} = \{N_{\mathrm{T}}\lambda_{\boldsymbol{PP}}\lambda_{\boldsymbol{K}} \mid \lambda_{\boldsymbol{PP}} \in \Lambda_{\boldsymbol{PP}}, \lambda_{\boldsymbol{K}} \in \Lambda_{\boldsymbol{K}}\}$ where $\Lambda_{\boldsymbol{PP}}$ and $\Lambda_{\boldsymbol{K}}$ are the sets of eigenvalues of $\boldsymbol{P}^{\mathrm{T}}\boldsymbol{P}$ and $\boldsymbol{K}_{\boldsymbol{\phi\phi}}$, respectively. Because both $\boldsymbol{P}^{\mathrm{T}}\boldsymbol{P}$ and $\boldsymbol{K}_{\boldsymbol{\phi\phi}}$ are positive semi-definite (all eigenvalues are non-negative) the largest and smallest eigenvalues of $\boldsymbol{H}$ are $\lambda_{\boldsymbol{H}\mathrm{max}} = N_{\mathrm{T}}\lambda_{\boldsymbol{PP}\mathrm{max}}\lambda_{\boldsymbol{K}\mathrm{max}}$ and $\lambda_{\boldsymbol{H}\mathrm{min}} = N_{\mathrm{T}}\lambda_{\boldsymbol{PP}\mathrm{min}}\lambda_{\boldsymbol{K}\mathrm{min}}$, respectively. Thus,

$$\frac{\lambda_{\boldsymbol{H}\mathrm{min}}}{\lambda_{\boldsymbol{H}\mathrm{max}}} = \frac{\lambda_{\boldsymbol{PP}\mathrm{min}}\lambda_{\boldsymbol{K}\mathrm{min}}}{\lambda_{\boldsymbol{PP}\mathrm{max}}\lambda_{\boldsymbol{K}\mathrm{max}}}. \tag{16}$$

Thus, the ratio $\lambda_{\mathrm{min}}/\lambda_{\mathrm{max}}$ of $\boldsymbol{H}_{\boldsymbol{W}}$ is a function of the $\lambda_{\mathrm{min}}/\lambda_{\mathrm{max}}$ ratios of $\boldsymbol{P}^{\mathrm{T}}\boldsymbol{P}$ and $\boldsymbol{K}_{\boldsymbol{\phi\phi}}$. The matrix $\boldsymbol{P}^{\mathrm{T}}\boldsymbol{P}$ is related to the properties of the motor system g, and how it interacts with the perturbation T. The matrix $\boldsymbol{K}_{\boldsymbol{\phi\phi}}$ is related to the distribution of target forces $\boldsymbol{f}_{\mathrm{d}}$.

Next, we show how the ratio $\lambda_{\boldsymbol{PP}\mathrm{min}}/\lambda_{\boldsymbol{PP}\mathrm{max}}$ is directly related to the force manipulability ellipse of the arm. The force manipulability ellipse is a mapping from a unit ball of joint torques to endpoint forces. It characterizes the force generation capabilities of the arm, showing that the arm can more readily produce forces along the major axis of the ellipse than in other directions (Yoshikawa, 1985). If $\boldsymbol{P}$ corresponds to a mapping between joint torques and end-point forces, the mapping can be described by the force manipulability ellipse. The axes of the force manipulability ellipse have sizes that correspond to the singular values $s_{\boldsymbol{P}}$ of $\boldsymbol{P}$ (Yoshikawa, 1985). It is also known that $s_{\boldsymbol{P}i} = \sqrt{\lambda_{\boldsymbol{PP}i}}$. Thus,

$$\frac{\lambda_{\boldsymbol{H}\mathrm{min}}}{\lambda_{\boldsymbol{H}\mathrm{max}}} = \left(\frac{s_{\boldsymbol{P}\mathrm{min}}}{s_{\boldsymbol{P}\mathrm{max}}}\right)^2 \frac{\lambda_{\boldsymbol{K}\mathrm{min}}}{\lambda_{\boldsymbol{K}\mathrm{max}}}. \tag{17}$$

The force manipulability ellipse concept can be generalized to the case where $\boldsymbol{P}$ corresponds to a mapping between muscle or synergy activations and forces. In this case the ellipse is the image of a unit ball of muscle or synergy activations in the space of end-point forces.

We have therefore demonstrated analytically that $\lambda_{\boldsymbol{H}\mathrm{min}}/\lambda_{\boldsymbol{H}\mathrm{max}}$ can be controlled by independently adjusting the force manipulability ellipse of $\boldsymbol{P}$ (varying the ratio of the magnitudes of the ellipse axes), or by adjusting the distribution of targets $\boldsymbol{f}_{\mathrm{d}}$ or other RBF parameters in the $\boldsymbol{\phi}$ function. Notice that the transformation T in $\boldsymbol{P}$ can alter the shape of the force manipulability ellipse. This happens if T is a scaling transformation or a virtual surgery, as described in Sections 2.2 *Simulation methods* and 2.3 *Experimental methods.*

2.1.5. Generalizability of results to realistic error function, learning algorithms, muscular models, and neural network structure

The theoretical results derived above were obtained for a learning system that minimizes performance errors. However, there is ample evidence that suggests that the CNS actually minimizes a trade-off between performance errors and effort during learning (Emken et al., 2007; Ganesh et al., 2010; Todorov & Jordan, 2002). In Appendix E, we show that our theoretical results generalize to learning systems that optimize a trade-off between the performance error and the effort term, which penalizes large motor commands.

Additionally, the above results are derived from the idealized case of a linear system with joint torques as the motor command. In Appendix F we demonstrate that introducing the nonlinearity $\sigma(\boldsymbol{m})$ to the motor command produces qualitatively similar results to the linear case. That is, we show that our theoretical results are also useful to model the speed of learning in systems with pulling muscles. Including the error-effort trade-off and the musculoskeletal model in the calculation of the Hessian $\boldsymbol{H}$ of the distal learning system results in

$$\boldsymbol{H} = \left[\left(\boldsymbol{P}^{\mathrm{T}}\boldsymbol{P}\right) \circ \left(\sigma'(\boldsymbol{m}^*) \sigma'(\boldsymbol{m}^*)^{\mathrm{T}} \right) + \boldsymbol{\Gamma} \right] \otimes \boldsymbol{\phi}(\boldsymbol{f}_{\mathrm{d}i}) \boldsymbol{\phi}(\boldsymbol{f}_{\mathrm{d}i})^{\mathrm{T}}, \tag{18}$$

where σ’ is the first derivative of the non-linearity σ, which ensures that muscle activations are non-negative, and $\boldsymbol{\Gamma}$ is a matrix containing terms related to the minimization of effort. The derivation of Eq. (18) is shown in Appendix F.

Furthermore, the above theoretical results are based on a learning algorithm that can directly compute the gradient $\nabla_{\boldsymbol{W}} J$, such as backpropagation. However, the way in which the CNS actually computes the gradient or approximations to the gradient is still a subject of research. In Appendix B, we show that the node perturbation algorithm, a biologically plausible learning algorithm that produces parameter updates that on average follow the gradient, produces qualitatively similar results to the backpropagation algorithm.

Finally, although we defined above the inverse model g^{-1} as an RBF network, the above results can generalize to other inverse model architectures. For instance, the RBF network in the inverse model can be replaced by a typical multi-layer feedforward network or any other function approximator. To see why this is the case, we can visualize the new inverse model to include a function h on the network input with tunable parameters $\boldsymbol{w}_{\text{pre}}$ such that

$$g^{-1}(\boldsymbol{f}) = \boldsymbol{W}h(\boldsymbol{f}, \boldsymbol{w}_{\text{pre}}) = \boldsymbol{m}. \tag{19}$$

This functional form also encompasses the inverse model in Eq. (2), as $h(\boldsymbol{f},\boldsymbol{w}_{\text{pre}})$ can describe the output of the RBF units. Therefore, if we assume that a set of valid parameters $\boldsymbol{w}_{\text{pre}}$ is known prior to learning and learning only occurs on $\boldsymbol{W}$, then all the above arguments about the speed of learning are also valid for the new architecture. This idea is similar to our treatment of the RBF network, where the RBF unit parameters (center and shape) are not considered for the calculation of the gradient. Of course, the parameters $\boldsymbol{w}_{\text{pre}}$ are generally not known a priori, so they must be learned using gradient descent or other methods. However, it follows that the dynamics of learning $\boldsymbol{W}$ must be influenced by the shapes of the force manipulability ellipse and the target distribution in the same way that we have described above.

2.2. Simulation methods

2.2.1. Overview

Based on the theoretical results of our computational framework, we simulated learning in four isometric arm tasks: (1) anisotropic scaling task, 2) anisotropic target distribution task, 3) virtual surgery task, and 4) target number task. In the anisotropic scaling and virtual surgery tasks we control the shape of the manipulability ellipse of the arm in different conditions. In the anisotropic target distribution and target number tasks, we control the target distributions in different conditions. The virtual surgery and target number tasks were included to explain experimental findings reported in previous studies in the context of our theoretical results. We simulated all tasks using the distal learning and direct inverse modeling paradigms (Appendix A). In Section 2.1 *Computational modeling framework* we described a network in which the motor commands are joint torques for illustration purposes. However, our main interest is using a musculoskeletal system for a more accurate representation of the CNS. Therefore, we used sigmoidal activation functions in the outputs of the inverse model so that the motor command corresponds to muscle activations (Appendix F). We also considered network architectures in which muscle synergies are explicitly included in the model. This is important, because previous results in the virtual surgery task attribute differences in the speed of learning between conditions to the presence of muscle synergies in the CNS (Berger et al., 2013). Finally, we also considered a cost function that minimizes a trade-off between error and effort in the task in all simulations (Appendix E, F).

We ran all four simulated tasks with 15 simulated participants. Each simulated subject consists of a mapping between muscle activations and planar forces ($\boldsymbol{M}$) and a set of muscle synergies ($\boldsymbol{S}$). We used experimental data from 15 human participants during an isometric reaching task to obtain $\boldsymbol{M}$ and $\boldsymbol{S}$ for each simulated participant, as described in Barradas et al. (2020). This allowed us to create simulated participants that produce realistic muscle activations (Appendix G).

We also computed $\lambda_{\boldsymbol{H}\min}/\lambda_{\boldsymbol{H}\max}$, the ratio of the smallest and largest eigenvalue of $\boldsymbol{H}$, for each simulated subject in each task according to the model structure described above using Eq. (18) and as described in Appendix F.

2.2.2. Motor command to force mapping

Each simulated participant was defined by a mapping between motor commands $\boldsymbol{m}$ and forces $\boldsymbol{f}_{\text{r}}$. Each mapping was obtained from a real human participant in a previous study (Barradas et al., 2020), as described in Appendix G. We focus on two cases: $\boldsymbol{m}$ represents activations of muscles that cross the shoulder and elbow, or $\boldsymbol{m}$ represents muscle synergy activations of the same muscles. In both cases, the mapping is defined by $\boldsymbol{P}$ in Eq. (1).

Muscle activations as motor commands. In the case of a muscle activation to force mapping $\boldsymbol{P} = \boldsymbol{M}$, where each column of $\mathbf{M}$ indicates the force produced by each muscle, which is scaled by the elements of $\sigma(\boldsymbol{m})$. Notice that the nonlinearity σ ensures that $\boldsymbol{m}$ will produce only pulling forces as defined by $\boldsymbol{M}$. The $\boldsymbol{M}$ mappings taken from the real human participants had 10 muscles ($N_{\text{m}} = 10$).

Synergy activations as motor commands. In the case of a muscle synergy activation to force mapping, $\boldsymbol{m}$ is a N_{s}-dimensional vector of muscle synergy activations, and $\boldsymbol{P} = \boldsymbol{MS}$, where $\boldsymbol{S}$ is an $N_{\text{m}} \times N_{\text{s}}$ matrix representing a set of muscle synergies, and N_{s} is the number of muscle synergies. Notice that muscle synergies can be represented as an additional layer in the network in Fig. 1b. However, because muscle synergies are often modeled as a linear transformation of synergy activations, it is possible to incorporate them into $\boldsymbol{P}$. This collapses the additional network layer, resulting in the same model structure illustrated in Fig. 1b. This simplification is convenient for our analyses, as it allows us to treat the case in which $\boldsymbol{m}$ are muscle synergy activations using the same computational framework. Each simulated participant was associated to a set of synergies $\boldsymbol{S}$ obtained from the same real participant from which the mapping $\boldsymbol{M}$ was taken (Barradas et al., 2020). The number of synergies N_{s} ranged from 3 to 5. Notice that $\boldsymbol{S}$ is in general different from $\boldsymbol{S}_{\text{ext}}$, described in Section 2.2.5 *Virtual surgery task*.

2.2.3. Anisotropic scaling task

Based on the theoretical results in Eq. (17), we designed a task with visuomotor transformations that control the shape of the force manipulability ellipse of $\boldsymbol{P}$ to show its effects on the speed of learning (Fig. 2d, e). In our models, visuomotor transformations modify the motor command-force mapping $\boldsymbol{P}$ by applying a linear transformation in force space:

$$\boldsymbol{P}' = \boldsymbol{TP} \tag{20}$$

where $\boldsymbol{T}$ is a 2 × 2 matrix that constitutes the visuomotor transformation. We designed $\boldsymbol{T}$ as the combination of a rotation and a scaling transformation. We use a rotation in addition to the scaling so that we can observe learning even in the absence of scaling:

$$\boldsymbol{T} = \boldsymbol{T}_{\text{rot}}\boldsymbol{T}_{\text{sc}}, \tag{21}$$

where $\boldsymbol{T}_{\text{rot}}$ is a 30° counterclockwise rotation, and $\boldsymbol{T}_{\text{sc}}$ is a linear transformation that alters the shape of the manipulability ellipse, given by the ratio of the singular values $s_{\min}/s_{\max}$ of $\boldsymbol{P}$'. The $\boldsymbol{T}_{\text{sc}}$ transformation is a composition of a rotation and a scaling that allows to directly control $s_{\min}/s_{\max}$. Because the axes of the manipulability ellipse of $\boldsymbol{P}$ are not necessarily aligned with the horizontal and vertical directions, $\boldsymbol{T}_{\text{sc}}$ must first align $\boldsymbol{P}$ to the axes of the ellipse:

$$\boldsymbol{T}_{\text{sc}} = \boldsymbol{T}_{\text{align}}^{-1}\boldsymbol{T}_{\text{scale}}\boldsymbol{T}_{\text{align}} \tag{22}$$

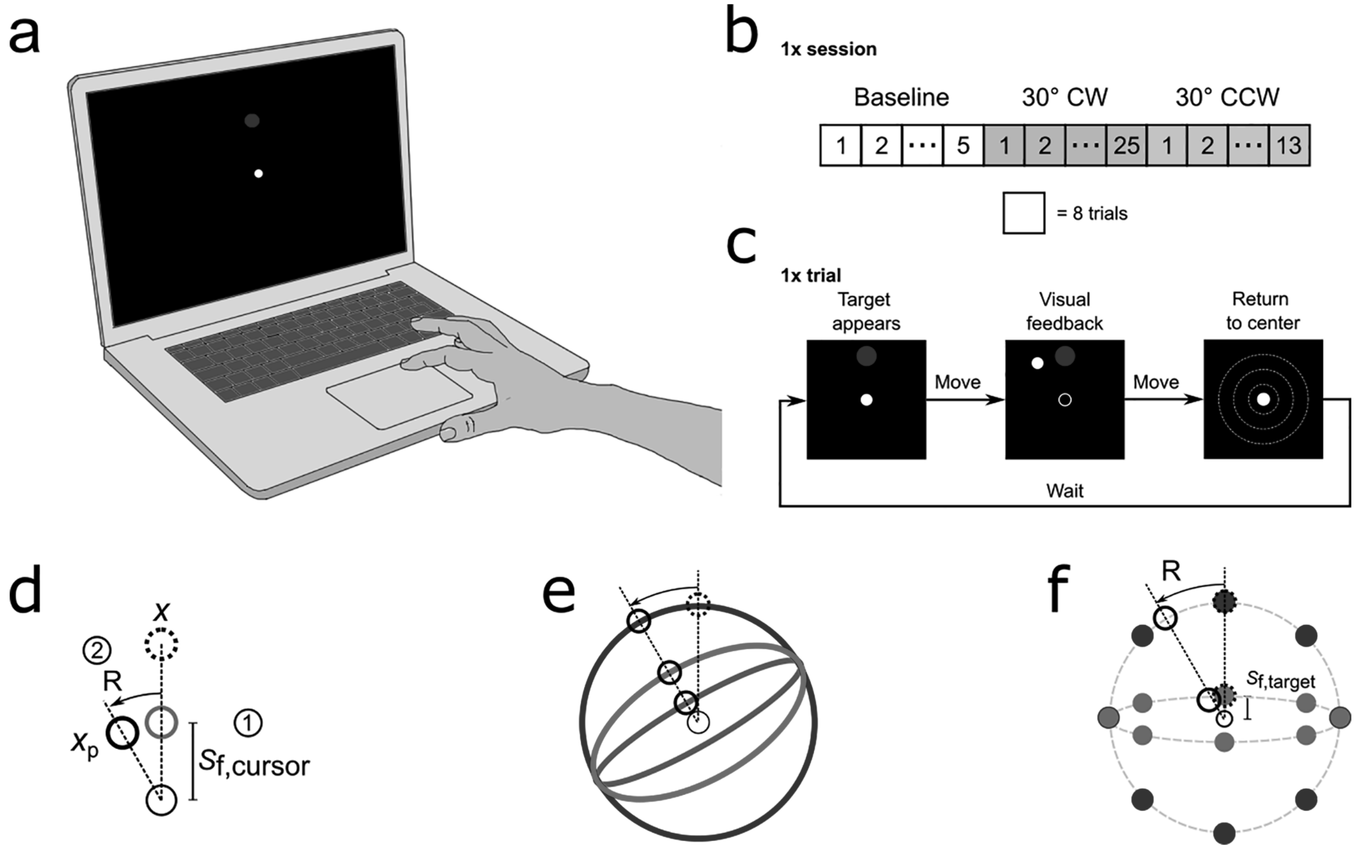

Fig. 2. Experimental setup and protocols. a. Experimental setup. Participants performed a reaching task using the trackpad of a laptop computer. b. Session structure. Each experimental session was conducted on a different day and was composed of 13 blocks of 8 trials. The first 5 blocks were baseline blocks, where cursor feedback was veridical. The next 25 blocks were perturbation blocks, where the position of the cursor was transformed by transformation $\boldsymbol{T}$ (experiment 1: scaling and rotation; experiment 2: rotation). The final 13 blocks were counter-perturbation blocks, where the rotation component of $\boldsymbol{T}$ was switched in direction. Both experiment 1 and experiment 2 had identical session structure. c. Trial structure. Each trial began with the cursor at the center of the virtual environment. After a target appeared, participants quickly moved the cursor to the target. Participants only received visual feedback of the final position of the cursor. At the end of the trial, the cursor and the target disappeared, and participants returned the cursor to the center of the virtual environment. Only a circle indicating the distance of the cursor to the center was visible during this stage. A new trial started after keeping the cursor at the center for 500 ms. d. Experiment 1: anisotropic scaling. During the perturbation and counter-perturbation phases of experiment 1, visual feedback of the cursor was transformed by scaling the position of the cursor in the vertical axis of the screen by a factor $s_{f,cursor}$ and rotating the resulting position 30° around the center of the screen in the counterclockwise or clockwise direction. e. Experimental conditions in experiment 1. Each participant performed the reaching task in three separate sessions. In each session the scaling factor $s_{f,cursor}$ was different (1.0, blue; 0.5, red; 0.2, green) but the rotation angle was the same. f. Experiment 2: target distributions. Each participant performed the reaching task with a visuomotor rotation in two separate sessions. In each session the distribution of target positions was different. In the first condition targets were uniformly and radially distributed around the center of the screen ($s_{f,target} = 1$). Thus, this condition is equivalent to the condition in experiment 1 in which $s_{f,cursor} = 1$ (blue). In the second condition the targets were arranged into an elliptical shape by scaling the vertical position of the targets in the first condition by $s_{f,target} = 0.2$ (red).

$$\boldsymbol{T}_{\text{scale}} = \begin{bmatrix} 1 & 0 \\ 0 & s_{\text{f,cursor}} \end{bmatrix} . \tag{23}$$

Therefore, the $\boldsymbol{T}_{\text{sc}}$ transformation first rotates $\boldsymbol{P}$ according to a rotation $\boldsymbol{T}_{\text{align}}$ such that the left singular vector of $\boldsymbol{P}$ corresponding to $s_{\max}$ becomes aligned with the x axis, then scales the aligned $\boldsymbol{P}$ according to a scaling $\boldsymbol{T}_{\text{scale}}$ such that $s_{\min}/s_{\max}$ changes, and finally, rotates $\boldsymbol{P}$ again according to the inverse rotation $\boldsymbol{T}_{\text{align}}^{-1}$ such that the left singular vector of $\boldsymbol{P}$ returns to its original orientation. The transformation $\boldsymbol{T}_{\text{scale}}$ is defined by $s_{\text{f,cursor}}$, a scaling factor. We created three conditions in which $s_{\text{f,cursor}} = [0.2\ 0.5\ 1]$.

2.2.4. Anisotropic target distribution task

Based on the theoretical results in Eq. (17), we designed a task with targets taken from distributions of different shapes to show their effects on the speed of learning (Fig. 2f). The task was a 30° counterclockwise visuomotor rotation. We defined two different target distributions. In the first one, targets were placed uniformly on the outline of a circle on the horizontal plane, constituting an isotropic distribution. In the second one, targets were placed on the outline of an ellipse, constituting an anisotropic distribution. The targets in the anisotropic distribution were defined by applying a scaling transformation on the targets of the isotropic distribution:

$$\boldsymbol{f}_{\text{d}}^{'} = \boldsymbol{T}_{\text{scale}}\boldsymbol{f}_{\text{d}} , \tag{24}$$

where $\boldsymbol{f}_{\text{d}}$ are the target forces in the isotropic distribution, $\boldsymbol{f}_{\text{d}}$' are the target forces in the anisotropic distribution, and $\boldsymbol{T}_{\text{scale}}$ is a scaling transformation as defined in Eq. (23). The value of the scaling factor $s_{\text{f, target}}$ was set to 0.2.

2.2.5. Virtual surgery task

Following the methods described in Berger et al. (2013), we designed virtual surgeries that were either incompatible or compatible with the muscle synergies $\boldsymbol{S}_{\text{ext}}$ extracted from the models. A virtual surgery modifies the forward physics of the arm ($\boldsymbol{P}$) by applying a linear transformation in muscle space:

$$\boldsymbol{P}^{'} = \boldsymbol{MT} \quad \text{or} \quad \boldsymbol{P}^{'} = \boldsymbol{MTS} , \tag{25}$$

where $\boldsymbol{T}$ is an orthonormal $N_\mathrm{m} \times N_\mathrm{m}$ matrix that constitutes the transformation or virtual surgery, and $\boldsymbol{S}$ is a $N_\mathrm{m} \times N_\mathrm{s}$ matrix representing the synergies embedded in the model. N_s is the number of synergies in the model. The first equation corresponds to the case in which $\boldsymbol{m}$ are muscle activations, and the second equation to the case in which $\boldsymbol{m}$ are muscle synergy activations.

Incompatible virtual surgeries are designed such that muscle activations produced by $\boldsymbol{S}$ are restricted to generate forces along only one dimension of the force space, while the resulting muscle activation-force mapping $\boldsymbol{MT}$ spans the whole force space. Therefore, theoretically, any force can be produced by a new combination of muscle activations, but in practice, forces are biased towards one dimension of the plane. In contrast, compatible virtual surgeries are designed such that muscle activations produced by $\boldsymbol{S}$ span the whole force space. We show that incompatible surgeries produce highly anisotropic manipulability ellipses of $\boldsymbol{P}$, whereas compatible surgeries do not affect the shape of the manipulability ellipse. Therefore, both kinds of surgeries have different effects on the speed of learning in the task.

For each simulated subject, we extracted muscle synergies $\boldsymbol{S}_\mathrm{ext}$ using non-negative matrix factorization (Lee & Seung, 1999) on the set of muscle activations that the simulated subject produced in response to a set of desired forces with magnitudes between 5 and 20 N in all directions of the horizontal plane. The number of synergies was selected as in previous studies (Barradas et al., 2020; Berger et al., 2013). We then designed an incompatible and a compatible surgery based on the simulated subject's $\boldsymbol{M}$ and extracted $\boldsymbol{S}_\mathrm{ext}$. We designed both surgeries according to the methods described in Berger et al. (2013). The incompatible surgeries were built by finding $\boldsymbol{T}_\mathrm{I}$ that maps muscle activations in the column space of $\boldsymbol{S}_\mathrm{ext}$ into the null space of $\boldsymbol{M}$. In contrast, the compatible virtual surgeries were built by finding $\boldsymbol{T}_\mathrm{C}$ that transforms the muscle activations into different muscle activations outside of the null space of $\boldsymbol{M}$. Details can be found in Berger et al. (2013). We randomly generated incompatible and compatible surgeries and applied them to the forward physics of the system to calculate the initial average error in force direction under each surgery. We used eight radially and uniformly distributed force targets. We selected surgeries that produced an initial average error in force direction close to 60°. This was decided based on previously reported results for the mean initial error during incompatible and compatible virtual surgeries (Berger et al., 2013). We followed the experiment schedule defined in Berger et al. (2013).

2.2.6. Target number task

We simulated a visuomotor rotation reaching task in which target sets with different numbers of targets are presented. This simulated task is based on a previous experiment (Krakauer et al., 2000). We show that different numbers of targets produce differences in the eigenvalues of the covariance matrix of targets $\boldsymbol{K}_{\boldsymbol{\phi\phi}}$, which influences the speed of learning. We used 3 target sets with 1, 4 and 8 targets each. The targets in the 4- and 8-target sets were radially and uniformly distributed. The visuomotor rotation was a counterclockwise 30° rotation. We followed the experiment schedule defined in Krakauer et al. (2000).

2.2.7. Simulation procedure

We ran simulations of each experiment for 15 simulated participants based on mappings between muscle activations and planar forces $\boldsymbol{M}$ and muscle synergy matrices $\boldsymbol{S}$ computed in a previous study (Barradas et al., 2020). We trained each simulated subject to adequately perform the baseline task without perturbations while producing plausible muscle activations (see Appendix G for details). Each simulation trial was defined as one cycle of target presentation, force production, and a learning step. The simulations followed the protocol defined in the corresponding experiment (number of baseline and perturbation trials defined in Section 2.3 *Experimental methods*, in Berger et al. (2013), and in Krakauer et al. (2000), except for counter-perturbation and washout phases, which we did not simulate.

The model takes the learning rate of the inverse model η_I as a parameter (Eq. (6)). We defined a set of learning rate values ($\eta_\mathrm{I} = [1, 2, 5] \times 10^{-a}$ where $a = [6, 5, 4, 3, 2]$, that is, 15 different values spanning different orders of magnitude) and ran simulations for each value. For each value of η_I and simulated subject, we ran 10 different simulations with a different target order. The target order in each simulation was pseudo-randomly generated within cycles of the defined targets. We averaged the learning curves corresponding to the error in force direction obtained from each set of these 10 simulations. We then further averaged the error in direction for every sequence of eight targets to obtain smooth learning curves (except for the multi-target experiment, in which there are only 18 trials per condition). For each simulated subject, we selected the value of η_I that best fit the mean experimental results in the corresponding task, as measured by the root mean square error (RMSE) between the subject's and the mean experimental learning curves. Note that for a given simulated participant, the value of η_I was the same for all conditions in a given simulated task. Finally, we averaged the resulting learning curves across all 15 simulated participants.

2.3. Experimental methods

Based on the insights provided by our computational framework, we hypothesized that anisotropy in the force manipulability ellipse and in the distribution of targets would slow down motor learning with respect to more isotropic force manipulability ellipses and target distributions. To test these hypotheses, we performed experiments based on the anisotropic scaling and anisotropic target distribution tasks described in Sections 2.2.3 and 2.2.4. Experiments for the virtual surgery task and the number of targets task have been performed in previous studies (Berger et al., 2013; Krakauer et al., 2000).

2.3.1. Participants

Twenty-three right-handed participants [mean age, 27.9 yr (SD 5.6); 18 men] were included in the study after providing written informed consent. All 23 participants participated in Experiment 1, whereas a subset of 21 participants participated in Experiment 2. Additionally, we used data from 15 participants reported in a previous study [mean age, 27.9 yr (SD 8.8); 13 men] (Barradas et al., 2020) to model learning in the isometric tasks. All procedures were approved by the Ethical Review Board of the Tokyo Institute of Technology.

2.3.2. Experimental setup

Because of the COVID-19 pandemic, which restricted in-person experiments, the experiments were implemented on an online platform (OnPoint) that participants accessed remotely using a laptop computer (Tsay et al., 2021). This online platform has been shown to produce similar results to in-person experiments in studies of adaptation to visuomotor transformations (Avraham et al., 2021; Tsay et al., 2021, 2022). Participants were instructed to place the laptop on a table or desk and to sit comfortably in front of it with the laptop centered at the body midline (Fig. 2a). A virtual environment was displayed on the laptop screen. The virtual environment consisted of a circular white cursor, a ring-shaped white target at the center of the screen, and several circular blue targets displayed on a black background.

2.3.3. Experimental protocols

Participants used their right index finger and the built-in trackpad of a laptop computer to control the position of a cursor on the visual display. The experimental task required the displacement of the cursor from a center position to one of eight targets uniformly distributed around the center of the screen. Displacements of the cursor in the virtual environment were amplified by a factor of 3 with respect to the default settings of each participant's laptop computer to ensure that the targets in all tasks were reachable. Participants were instructed to make quick center-out movements with the cursor towards the target, and to attempt to stop inside the target. Participants were also informed that for some trials visual feedback of the cursor would be perturbed and not match their expected feedback. Participants were instructed to continue aiming toward the displayed target and to ignore the perturbation, which allows the observation of the implicit component of motor adaptation (Morehead et al., 2017).

Each trial started by displaying the target at the central position, which roughly corresponded to the center of the trackpad. After the cursor was placed inside the central target for 500 ms, one of the targets appeared. Visual feedback of the cursor position was suppressed during the reaching movement and was only provided at the end of the trial as a static image of the cursor at the position where the cursor stopped. Visual feedback of the cursor at the end of the trial was held on the screen for 1 s, after which it was suppressed. Next, the central target and a white ring appeared. The radius of the white ring indicated the distance between the cursor and the central target and was used to eliminate direction-specific feedback of the cursor. Participants were asked to return to the central target by reducing the diameter of the ring. The central target turned into a white circle when the cursor was placed inside of it. After this, the next trial started. Fig. 2c illustrates the structure of a trial.

2.3.4. *Experiment 1: anisotropic scaling*

Participants completed three experimental sessions on different days, with 1 to 4 days between sessions, depending on the availability of each participant. In each session, participants performed a total of 43 target cycles (344 trials), divided into baseline, perturbation and counter-perturbation phases, which consisted of 5, 25 and 13 target cycles, respectively (Fig. 2b). Each target cycle contained a sequence of eight targets distributed uniformly around the center of the screen in a pseudorandom order (Fig. 2f, blue targets). Targets were located at a distance equivalent to 1/4 of the height of the screen of each participant's laptop. In the baseline phase, cursor feedback was veridical. In the perturbation phase, a composition of a 30° counterclockwise rotation and a scaling transformation was applied to the position of the cursor (Fig. 2d). In the counter-perturbation phase, the rotation was switched to the clockwise direction, and the scaling transformation was the same as in the perturbation phase. The counter-perturbation was included to induce retrograde interference (Krakauer et al., 2005), minimizing possible savings in learning the rotation aspect of the perturbation in subsequent sessions. In a similar way to the simulated anisotropic scaling task, the perturbations were defined as:

$$\boldsymbol{x}_{\mathrm{P}} = \boldsymbol{T}\boldsymbol{x} = \boldsymbol{T}_{\mathrm{rot}}\boldsymbol{T}_{\mathrm{scale}}\boldsymbol{x}, \tag{26}$$

where $\boldsymbol{x}$ and $\boldsymbol{x}_{\mathrm{P}}$ are two-dimensional vectors containing the position of the cursor on the screen before and after applying the perturbation, $\boldsymbol{T}$ is a visuomotor transformation consisting of a scaling and rotation, $\boldsymbol{T}_{\mathrm{rot}}$ is a rotation matrix, and $\boldsymbol{T}_{\mathrm{scale}}$ is a scaling matrix as defined in Eq. (23), The value of the scaling factor $s_{\mathrm{f,cursor}}$ in $\boldsymbol{T}_{\mathrm{scale}}$ was different in all three sessions (0.2, 0.5 or 1, Fig. 2e), and the order of the sessions was counterbalanced across participants. Notice that scaling transformations produce errors both in the direction and the extent of a reach.

2.3.5. *Experiment 2: target distribution*

A subset of the participants in Experiment 1 completed an additional experimental session after 2 - 4 weeks from the last session in Experiment 1. The trial structure in Experiment 2 was identical to that of Experiment 1, except for the position of the targets and the nature of the perturbation and the counter-perturbation. The positions of the targets were defined by scaling the vertical position of the targets defined in Experiment 1 according to the anisotropic target distribution described in Section 2.2.4 (scaling factor $s_{\mathrm{f,target}} = 0.2$, Fig. 2f, purple targets). The perturbation was a 30° counterclockwise rotation of the cursor position, and the counter-perturbation was a 30° clockwise rotation, both with no scaling component.

The results of Experiment 2 were compared to the results of the condition in Experiment 1 in which $s_{\mathrm{f,cursor}} = 1$. This is because this condition is equivalent to using a pure rotation and a round target distribution.

2.3.6. *Outcome measures*

We measured the error in reaching direction between the target and the final cursor position as the outcome variable. Similar to the simulations, we averaged the error for all targets in a target cycle. For each subject and condition, we fitted the error curve with a first order exponential to obtain a learning rate parameter. We used the *fit* nonlinear least squares fitting function in Matlab R2019b. We compared the computed learning rate parameter to $\lambda_{\boldsymbol{H}\mathrm{min}}/\lambda_{\boldsymbol{H}\mathrm{max}}$ based on the Hessian of the network of the corresponding condition.

2.3.7. *Statistical analysis*

The main outcomes of the experiments were the degree of adaptation to the visuomotor transformations (initial error – final error in reaching direction) and the computed decay rate parameter of the exponential fit to the time course of the error in direction, estimated as indicated in Section 2.3.6 *Outcome measures*. For Experiment 1, we tested the null hypothesis that the means of the outcome variables in the three anisotropic scaling conditions were equal by means of a repeated measures ANOVA. In cases where the outcome variables did not satisfy the normality assumptions of the ANOVA test, we performed a longitudinal non-parametric test using the nparLD package in R (Noguchi et al., 2012).We then performed post-hoc multiple comparisons using Bonferroni corrections. In Experiment 2, we tested the null hypothesis that the means of the outcome variables in the two target distribution scaling conditions were equal by means of a paired *t* test. In cases where the outcome variables did not satisfy the normality assumptions of the *t* test, we performed a Wilcoxon signed rank test. The significance threshold was set to $p = 0.05$. All analyses were performed in R 4.2.1.

3. Results

3.1. *High anisotropy in manipulability ellipsoids is associated with lower degrees of adaptation to visuomotor transformations*

Fig. 3 shows the performance of a representative participant before, at the onset and at the end of the transformation block for all conditions in experiments 1 and 2. Notice that errors for each target at the onset of transformations with higher degrees of scaling were widely different due to the anisotropy of the transformation (Fig. 3b, c). In experiment 1 (anisotropic scaling task), we found that scaling the manipulability ellipse had a statistically significant effect on the degree of adaptation to the visuomotor transformation (initial error – final error) [$s_{\mathrm{f,cursor}} = 1.0$, initial error: 27.4° (SD 3.4), final error: 18.9° (SD 6.1); $s_{\mathrm{f,cursor}} = 0.5$, initial error: 27.3° (SD 2.3), final error: 22.5° (SD 4.9); $s_{\mathrm{f,cursor}} = 0.2$, initial error: 30.2° (SD 3.3), final error: 28.3° (SD 4.1); $F(2,44) = 12.98$, $p < 1\times10^{-4}$, repeated measures ANOVA]. Post hoc tests indicated that the degree of adaptation was significantly larger when no scaling factor was used ($s_{\mathrm{f,cursor}} = 1.0$) [$s_{\mathrm{f,cursor}} = 1.0$ vs $s_{\mathrm{f,cursor}} = 0.5$, $p = 0.027$; $s_{\mathrm{f,cursor}} = 1.0$ vs $s_{\mathrm{f,cursor}} = 0.2$, $p < 1\times10^{-3}$]. However, there was no statistically significant difference in the degree of adaptation between scaling factors $s_{\mathrm{f,cursor}} = 0.5$ and $s_{\mathrm{f,cursor}} = 0.2$ ($p = 0.12$).

Fig. 4b shows the time course of learning as measured by the mean absolute error in the direction of the reach in all conditions of experiment 1 across all participants. We found that scaling the manipulability ellipse had a statistically significant effect on the learning rate parameter of a single exponential fit to each participant's learning curve (Fig. 4a) [$s_{\mathrm{f,cursor}} = 1.0$, learning rate: 0.0186 (SD 0.0196); $s_{\mathrm{f,cursor}} = 0.5$, learning rate: 0.0077 (SD 0.0081); $s_{\mathrm{f,cursor}} = 0.2$, learning rate: 0.0027 (SD 0.0036); $F = 21.9$, df = 1.77, $p < 1\times10^{-5}$, nparLD (data in condition $s_{\mathrm{f,cursor}} = 1.0$ not normal according to Shapiro-Wilk test)]. Post hoc tests indicated that the learning rate parameter was significantly larger when no scaling factor was used ($s_{\mathrm{f,cursor}} = 1.0$) [$s_{\mathrm{f,cursor}} = 1.0$ vs $s_{\mathrm{f,cursor}} = 0.5$,

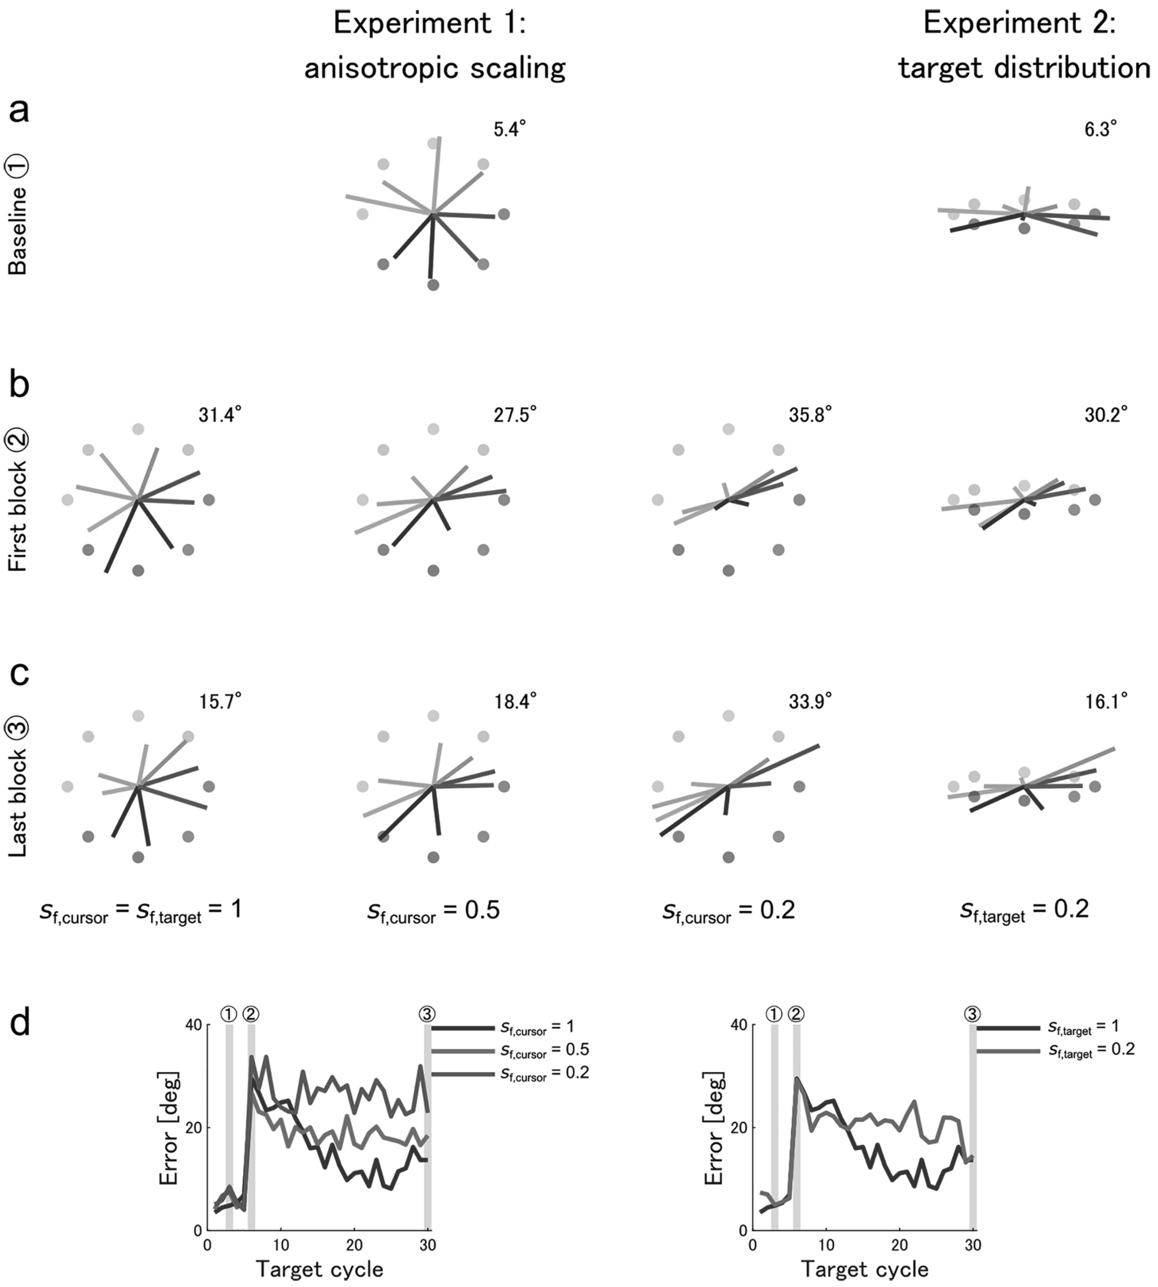

Fig. 3. Cursor positions at the end of a target cycle before, at the onset, and at the end of training in all conditions of experiments 1 and 2 for a representative participant. a. Baseline block. b. First block of training. c. Last block of training. The mean absolute value of the error in final direction is indicated for each condition and target cycle. d. Performance error measured as the mean absolute value of the error in initial direction for one block, comprising 8 targets. Shaded areas indicate the blocks corresponding to the cursor positions in a, b and c.

$p < 10^{-3}$; $s_{f,cursor} = 1.0$ vs $s_{f,cursor} = 0.2$, $p < 10^{-5}$]. Additionally, the learning rate parameter in scaling condition $s_{f,cursor} = 0.5$ was significantly larger than in condition $s_{f,cursor} = 0.2$ ($p = 0.039$). Because the anisotropic scaling task induces errors in both direction and extent of reach, we also inspected the time courses of the magnitude of the endpoint error. The time course of the magnitude of the error in all task conditions resembled the time courses of the error in direction (results not shown). The differences in the extent of adaptation and learning rates in all the experimental conditions suggest faster learning rates for tasks with the least anisotropic virtual manipulability ellipsoids.

3.2. High anisotropy in target distributions is associated with lower degrees of adaptation to visuomotor transformations

In experiment 2 (anisotropic target distribution task), we found that the round target layout was associated with a statistically significantly larger degree of adaptation to the visuomotor rotation than the elliptical target layout [$s_{f,target} = 1.0$, initial error: 27.9° (SD 3.0), final error: 18.9° (SD 6.0); $s_{f,target} = 0.2$, initial error: 27.0° (SD 4.4), final error: 23.3° (SD 6.3); $p = 0.005$, paired t test]. Fig. 4d shows the time course of learning in all conditions of experiment 2. We found that the shape of the target distribution had a statistically significant effect on the learning rate parameter of each participant's learning curve. The mean of the learning rate parameter was greater for the circular target layout than for the elliptical target layout (Fig. 4c) [$s_{f,target} = 1.0$, learning rate: 0.0186 (SD

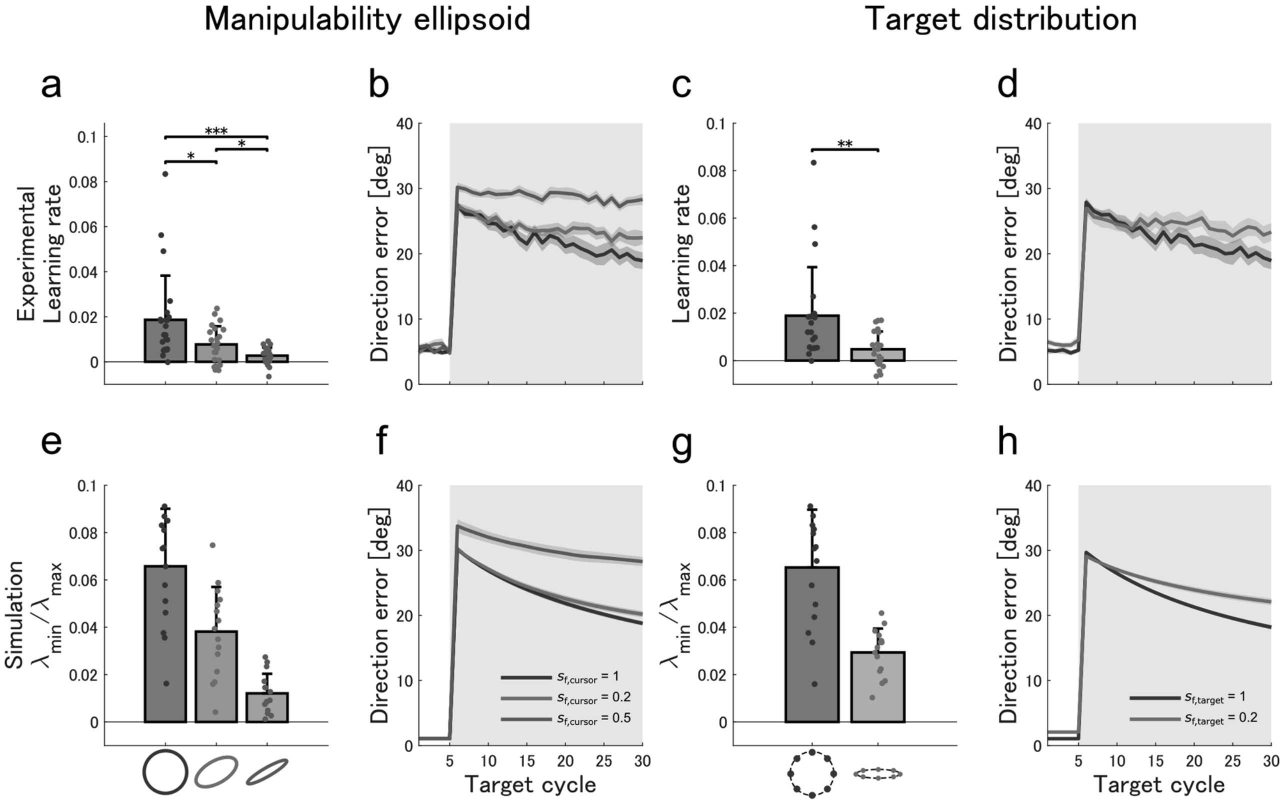

Fig. 4. Experimental and simulation results in the anisotropic scaling and anisotropic target distribution tasks. a. Experimental learning rates in Experiment 1, and c. Experiment 2. We estimated the experimental learning rates by fitting the learning curves of each participant to simple exponentials. Bars indicate the mean, and error bars the standard deviation. Asterisks indicate significant differences between conditions: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$. b, d. Experimental mean error in initial direction during learning in Experiments 1 and 2. e. Estimates of the slowest time constant of learning given by $\lambda_{\boldsymbol{H}\min}/\lambda_{\boldsymbol{H}\max}$ for the three scaling conditions in Experiment 1, and g. the two scaling conditions in Experiment 2. We generated the estimates of the slowest time constant of learning using the Hessian matrix of networks with sigmoidal hidden units for the 15 simulated participants (Eq. (18)). Bars indicate the mean, and error bars the standard deviation. f. Simulated mean error in initial direction during learning in the three scaling conditions in Experiment 1 (mean of the 15 simulated participants), and h. in Experiment 2. The simulation results were fitted to the mean experimental curves by adjusting the learning rate η_I in Eq. (6) for each simulated participant. Error bars indicate the standard error.

0.0196); $s_{f,target} = 0.2$, learning rate: 0.0048 (SD 0.0075); $p < 10^{-3}$, Wilcoxon signed rank test (data in condition $s_{f,target} = 1.0$ not normal according to Shapiro-Wilk test)]. The differences in the extent of adaptation and learning rates in all the experimental conditions suggest faster learning rates for tasks with rounder target distributions.

3.3. Hessian analysis predicts the effects of the manipulability ellipse and target distribution on the speed of learning a visuomotor transformation

We computed the Hessian of the network defined by the computational models of 15 simulated participants to estimate the slowest time constant of learning for each condition in Experiments 1 and 2 using the distal learning and direct inverse modeling architectures. We focused on the case where the motor command $\boldsymbol{m}$ corresponds to muscle synergy activations. For distal learning, we used Eq. (18), considering sigmoidal hidden activation functions (to use a musculoskeletal model) and a cost function that minimizes a trade-off between error and effort to compute the Hessian and its eigenvalues (Appendix F). Results using the direct inverse modeling architecture are provided in Appendix A. Consistent with our hypothesis, derived from the assumption of a network with linear hidden activation functions, the estimated slowest time constant of learning is larger for conditions with rounder manipulability ellipsoids or target distributions in the non-linear network (Fig. 4e, g) [$\lambda_{\boldsymbol{H}\min}/\lambda_{\boldsymbol{H}\max}$ in anisotropic scaling: $s_{f,cursor} = 1.0$: 0.066 (SD 0.024), $s_{f,cursor} = 0.5$: 0.038 (SD 0.019), $s_{f,cursor} = 0.2$: 0.012 (SD 0.008); $\lambda_{\boldsymbol{H}\min}/\lambda_{\boldsymbol{H}\max}$ in anisotropic target distribution: $s_{f,target} = 1.0$: 0.066 (SD 0.024), $s_{f,target} = 0.2$: 0.029 (SD 0.010)]. The measured learning rates in the experiments show a similar pattern to the estimated slowest time constant of learning ($\lambda_{\boldsymbol{H}\min}/\lambda_{\boldsymbol{H}\max}$) in both the distal learning and direct inverse modeling architectures, with a monotonic relationship between the scaling factor of the manipulability ellipsoid and target distributions, and the observed speed of learning (Fig. 4a, c).

Furthermore, for distal learning, the simulated learning curves in both the anisotropic scaling and target distribution tasks (Fig. 4f, h) closely match the learning curves observed experimentally after selecting an appropriate learning rate η_I (Eq. (6)) for each simulated subject (anisotropic scaling task: $R^2 = 0.90$, anisotropic target distribution task: $R^2 = 0.90$). We also found that the mean simulated learning curves for individual targets in each task condition qualitatively resemble the asymmetric pattern of initial errors observed experimentally (results not shown). Note, however, that the simulated learning curves using direct inverse modeling (Fig. A.2c) have a lower quality fit to the experimental observations (anisotropic scaling: $R^2 = 0.78$, anisotropic target distribution: $R^2 = 0.83$).

3.4. Hessian analysis predicts differences in learning rates in virtual surgery tasks

We applied the computational framework to a virtual surgery task (Berger et al., 2013). We found that the distal learning framework can only predict differences in the learning rates of an incompatible and a compatible virtual surgery when the model includes explicitly defined muscle synergies, (i.e., the motor command $\boldsymbol{m}$ corresponds to muscle synergy activations and $\boldsymbol{P} = \boldsymbol{MS}$) [distal learning, non-linear network;

$\lambda_{H\min}/\lambda_{H\max}$ with synergies, compatible: 0.029 (SD 0.016); incompatible: 0.017 (SD 0.014); $\lambda_{H\min}/\lambda_{H\max}$ of compatible larger for 13 of the simulated participants] (Fig. 5a). If muscle synergies are not included in the model (i.e., the motor command $\boldsymbol{m}$ corresponds to muscle activations and $\boldsymbol{P} = \boldsymbol{M}$), the framework predicts that the learning speed in the incompatible task is slightly larger than in the compatible task, which does not align with experimental observations [distal learning, non-linear network; $\lambda_{H\min}/\lambda_{H\max}$ without synergies, compatible: 0.094 (SD 0.036); incompatible: 0.112 (SD 0.048); $\lambda_{H\min}/\lambda_{H\max}$ of incompatible larger for 12 of the simulated participants]. In fact, for linear networks without muscle synergies, there is no difference in the value of $\lambda_{H\min}/\lambda_{H\max}$between the compatible and incompatible tasks. This is because in the case where synergies are not included in the model, the eigenvalues of $\boldsymbol{P}^{\mathrm{T}}\boldsymbol{P}$ are not a function of $\boldsymbol{T}$, given that for non-zero eigenvalues $\lambda(\boldsymbol{P}^{\mathrm{T}}\boldsymbol{P}) = \lambda(\boldsymbol{P}\boldsymbol{P}^{\mathrm{T}})$, and $\boldsymbol{P}\boldsymbol{P}^{\mathrm{T}} = \boldsymbol{M}\boldsymbol{T}\boldsymbol{T}^{\mathrm{T}}\boldsymbol{M}^{\mathrm{T}} = \boldsymbol{M}\boldsymbol{M}^{\mathrm{T}}$, as by definition, $\boldsymbol{T}$ is an orthonormal matrix ($\boldsymbol{T}\boldsymbol{T}^{\mathrm{T}} = \mathrm{I}$). However, the direct inverse modeling framework predicts differences in the learning rates of both tasks whether muscle synergies are explicitly defined in the model or not [direct inverse modeling, $\lambda_{H\min}/\lambda_{H\max}$ without synergies, compatible: 0.05 (SD 0.039); incompatible: 0.007 (SD 0.006); $\lambda_{H\min}/\lambda_{H\max}$ of compatible larger for 14 of the simulated participants] (Fig. A.2a). In both simulated frameworks, the simulated learning curves in the compatible and incompatible surgeries closely match the learning curves observed experimentally after selecting an appropriate learning rate η_{I} for each simulated subject (distal learning: $R^2 = 0.88$; direct inverse modeling: $R^2 = 0.86$) (Figs. 5b and A1.2c).

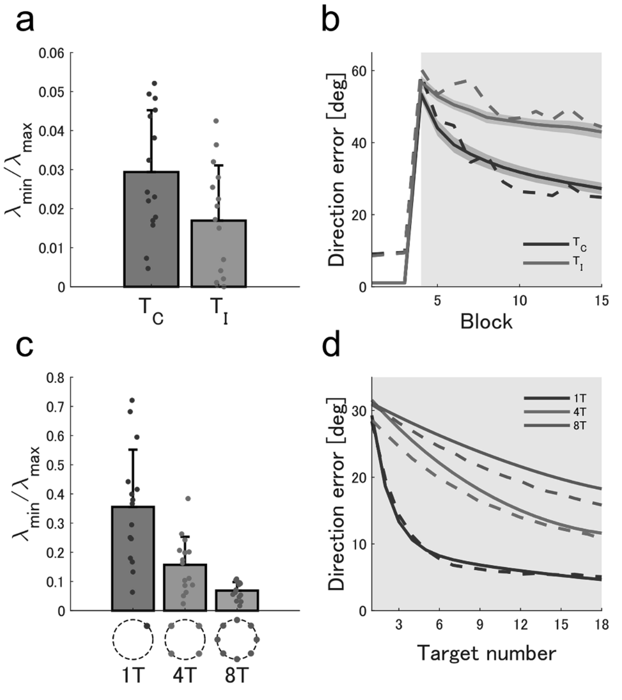

Fig. 5. Simulation results in the virtual surgery and multi-target tasks. a. Estimates of the slowest time constant of learning given by $\lambda_{H\min}/\lambda_{H\max}$ for the compatible ($\boldsymbol{T}_{\mathrm{C}}$) and incompatible ($\boldsymbol{T}_{\mathrm{I}}$) virtual surgeries. We generated the estimates using networks with sigmoidal hidden units and muscle synergies in the distal learning framework for the 15 simulated participants (Eq. (18)). Bars indicate the mean, and error bars the standard deviation. b. Simulated mean error in initial direction during learning in T_{C} and T_{I}. Simulated learning curves were fitted to the mean experimental curves by adjusting the learning rate η_{I} in Eq. (6). Solid lines correspond to the mean error in the 15 simulated participants. Dashed lines correspond to the mean error in experimental observations (Berger et al., 2013). Error bars indicate the standard error. c. Estimates of the slowest time constant of learning given by $\lambda_{H\min}/\lambda_{H\max}$ for the multi target task with 1, 4 and 8 targets (1T, 4T and 8T). We generated the estimates as indicated in a. d. Simulated mean error in initial direction during learning in 1T, 4T and 8T. Learning curves are double exponentials fitted to the mean simulated curves, which in turn were fitted to the mean experimental curves by adjusting the learning rate η_{I} in Eq. (6). Dashed lines correspond to double exponential fits to the experimental observations (Krakauer et al., 2000).

3.5. Hessian analysis predicts differences in learning rates in multi-target task

Finally, we applied the computational framework to a visuomotor rotation task in which experimental conditions are defined by the number of targets in the task (2, 4 or 8 targets) (Krakauer et al., 2000). We found that the framework predicts that the speed of learning decreases as the number of targets increases, consistent with experimental observations [distal learning, non-linear network; $\lambda_{H\min}/\lambda_{H\max}$ with synergies, 1 target: 0.356 (SD 0.196); 4 targets: 0.157 (SD 0.096); 8 targets: 0.069 (SD 0.029); $\lambda_{H\min}/\lambda_{H\max\ 1\mathrm{T}} > \lambda_{H\min}/\lambda_{H\max\ 4\mathrm{T}}$ and $\lambda_{H\min}/\lambda_{H\max\ 4\mathrm{T}} > \lambda_{H\min}/\lambda_{H\max\ 8\mathrm{T}}$ for all 15 simulated participants] (Fig. 5c). The direct inverse modeling framework also predicts a monotonic relationship between the number of targets and the speed of learning. However, the monotonic pattern of this relationship is appreciably different from the experimental observations (Fig. A.2a). The simulated learning curves in the number of targets task with the distal learning architecture closely match the learning curves observed experimentally after selecting an appropriate learning rate η_{I} for each simulated subject ($R^2 = 0.92$) (Fig. 5d). However, the simulated curves using direct inverse learning have a lower quality fit to the experimental observations ($R^2 = 0.77$) (Fig. A.2c).

4. Discussion

In this study, we showed that analyzing models of human motor learning through the lens of optimization theory allows us to predict how fast the CNS can learn new motor tasks. Specifically, we analyzed learning in a set of isometric force tasks for two different inverse learning paradigms, distal learning and direct inverse modeling. The simplicity of the forward physics of the isometric tasks allowed us to analytically derive the Hessian of a linear approximation of the learning system. We found that the eigenvalues of the Hessian matrix, and thus the speed of learning, are a function of the shapes of the force manipulability ellipsoid of the arm and the distribution of target forces. We extended these theoretical results to a non-linear representation of the motor system that includes muscles. Furthermore, our modeling framework establishes that these results are generalizable to other task conditions, and even different isometric tasks such as virtual surgeries. Although in the main text we have used backpropagation to illustrate the properties of the speed of learning under gradient-following algorithms, we also showed that our results do not depend on this choice and are also valid for the more biologically plausible node perturbation algorithm (Appendix B).

4.1. Testable predictions

Based on our theoretical results, we hypothesized that altering the manipulability ellipsoid of the arm or the force target distributions would directly impact the learning speed during a visuomotor rotation (VMR) in the isometric reaching task. We tested our hypothesis by having participants perform the VMR task under four different conditions. In the first three conditions, we altered the roundness of the manipulability ellipsoid in a virtual environment. In the fourth condition we used a target distribution in the shape of an ellipse. We found that rounder manipulability ellipsoids and target distributions allow for faster learning in the VMR task, confirming our hypothesis. Previous research supports our theoretical and experimental conclusions by showing in simulation that mechanical biases in the pulling directions of muscles influence the speed of learning (Hagio & Kouzaki, 2018). These biases can be interpreted as a factor that reshapes the generalized manipulability ellipsoid in the space of muscle activations (see Section 2.1.4 *Estimation of limit on learning speed*).

Our results suggest that the differences in the speed of learning in the task brought about by factors such as manipulability ellipsoids and target distributions may be an inherent computational property of learning in the isometric task. Both of these factors predict differences in the speed of learning for both the distal learning and direct inverse modeling paradigms, indicating that the differences may arise regardless of a particular architecture or optimized cost function (performance error vs motor error). However, although both inverse learning models are affected in a qualitatively similar way by these factors, the predicted relative speed of learning in different tasks is quantitatively different (Figs. 4e, g, 5a, c, A1. 1a), which is reflected in the simulated learning curves. Other factors seem to be particular to a given architecture. For example, the activation levels of muscles are involved in shaping the Hessian and thus the speed of learning in distal learning (Eq. (F.4)), but not in direct inverse modeling (Eqs. (A.5), (6)). Similarly, another architecture-specific factor that could influence the speed of learning is the nature of the RBFs encoding the network inputs, as mentioned above. Therefore, our framework can be used as a tool to generate systematic hypotheses to infer the type of algorithms that the CNS uses for motor learning in different tasks. For instance, our results show that the distal learning paradigm is able to describe the relative speeds of learning in all of the VMR tasks more adequately with respect to the data than the direct inverse learning paradigm. However, both paradigms provide similarly adequate descriptions of the speed of learning in the virtual surgery task.

Some of the factors that influence the speed of learning identified by our framework may be obvious based on previous theoretical and experimental studies. For example, the number of targets in a VMR task has been shown to affect the speed of learning because learning of rotations is only generalized locally. Learning around one target does not contribute to learning around faraway targets, increasing the overall number of trials needed to learn to counteract the perturbation for all targets (Krakauer et al., 2000). Since the models we used encode the positions of targets with local RBFs, it is not surprising that our simulations can reproduce the effects of the number of targets on the speed of learning. However, it is notable that this effect is predictable by examining the Hessian of the learning models. Increasing the number of targets tends to increase the rank of the covariance matrix of the RBF inputs, resulting in a decrease in the ratio of the smallest and largest eigenvalues of the matrix, which is associated with slower learning (Eq. (14), fewer targets than RBF units produce a rank deficient covariance matrix, and thus, as the number of targets increases, the rank increases). This suggests that other unexplored factors that affect the speed of learning, such as the properties of neuron populations in the inverse model (number of neurons, distribution of centers of neurons, or shape and size of receptive fields of neurons) could be identified by analyzing the Hessian of the learning models, even if their effect cannot be easily interpreted through intuitive mechanisms.

Additionally, our framework could potentially allow us to generate hypotheses about the role of different structures in the CNS on the speed of learning. Here, we found that a distal learning approach requires muscle synergies to be actual neural structures in the CNS to reproduce the results of virtual surgery experiments. These neural circuits in conjunction with the virtual surgeries reshape the manipulability ellipsoid of the arm in the virtual environment, and thus, compatible and incompatible surgeries elicit a difference in learning speeds. Otherwise, without these neural structures, the virtual surgeries have an unspecified effect on the manipulability ellipsoid, and the differences between compatible and incompatible surgeries are not captured by the model (see in-line equations in Section 3.4 *Hessian analysis predicts differences in learning rates in virtual surgery tasks*). In contrast, the direct inverse modeling approach can reproduce the experimental results with or without neural structures for muscle synergies in the model. This is because the virtual surgeries themselves differentially change the distribution of realized forces, which influences the speed of learning (Appendix A). However, other mechanisms may also explain the differences in learning speeds between the compatible and incompatible surgeries, with or without explicitly defined muscle synergies. For example, a recent study has shown that these differences can be explained by differences in the efficiency of learning a forward model of the surgeries due to task-relevant and irrelevant motor noise (Dal'Bello & Izawa, 2022).

Note, however, that care must be taken when comparing the results on the speed of learning in different architectures. Although both inverse learning paradigms, distal learning and direct inverse modeling, rely on gradient descent to learn an inverse, the way in which the inverse model is updated in both paradigms is different. Distal learning attempts to minimize predicted errors on performance via a forward model of the physics of the body, whereas direct inverse modeling attempts to minimize the error between the realized motor command and the motor command it would use to produce the realized action. That is, the quantities optimized in both paradigms are different. Therefore, our analysis pertains only to the learning rates of the quantity optimized in each paradigm. Our results are directly applicable in the case of distal learning, as the performance error is one of the main variables of interest in motor learning. In contrast, in the case of direct inverse modeling, our analysis corresponds to the learning rate in the motor error, which in general is not possible to measure experimentally, and does not directly relate to task performance. However, previous studies have shown that in linear systems, minimization of the motor error also leads to minimization of the performance error given appropriate initial conditions and/or exploration noise (Rolf & Steil, 2014). Furthermore, our direct inverse modeling simulations show that learning rates in the motor and performance domains correlate across tasks (Fig. A. 2), so our analysis is still useful provided that the minimization of performance and motor errors is coupled. This condition would be violated in the case of failure to form an adequate inverse model (Jordan & Rumelhart, 1992) or the occurrence of other failure modes (Sanger, 2004).

4.2. Future perspectives

One open question is the relation between the timescales of learning given by the eigenvalues of the Hessian and the different timescales of learning processes identified in the CNS. Motor adaptation is often described as the superposition of slow and fast processes that allow fast error corrections in the face of a changing environment and retention of the learned task in a stable environment (Smith et al., 2006). Models employing one fast and one slow process have been quite successful in describing motor adaptation in many tasks (Coltman et al., 2019; Lee & Schweighofer, 2009), although processes with multiple timescales have been identified (Kim et al., 2015). The neural basis of these processes is still not well understood, but may originate from different circuits in the CNS (Kim et al., 2015; Sarwary et al., 2018). However, in our framework, the multiple timescales defined by the Hessian describe the inherent dynamics of learning as an optimization process, without separate neural circuits operating at different timescales, as discussed below. Therefore, further research that considers the inherent dynamics of learning to study learning processes with different timescales seems necessary.

The dynamics of learning an inverse model are characterized by the eigenvalues of the Hessian and are reflected in the resulting learning curves as different timescales. The number of non-zero eigenvalues of the Hessian (its rank) then defines the number of timescales involved in learning. In the distal learning model, the eigenvalues of the Hessian are defined by the eigenvalues of the $\boldsymbol{P}^{\mathrm{T}}\boldsymbol{P}$ and $\boldsymbol{\phi}\boldsymbol{\phi}^{\mathrm{T}}$ matrices, representing the manipulability ellipsoid and target distribution, respectively (Eq. (16)). $\boldsymbol{P}^{\mathrm{T}}\boldsymbol{P}$ is in general only rank 2, whereas the rank of $\boldsymbol{\phi}\boldsymbol{\phi}^{\mathrm{T}}$ can be as large as the number of RBF input units, and depends on the number of targets. In our simulations, the rank of the Hessian is in general 16, as most experiments modeled use 8 targets ($\mathrm{rank}(\boldsymbol{H}) = \mathrm{rank}(\boldsymbol{P}^{\mathrm{T}}\boldsymbol{P})\cdot\mathrm{rank}(\boldsymbol{\phi}\boldsymbol{\phi}^{\mathrm{T}})$). We found that considering only the ratio of the slowest and fastest time constants predicts the relative speed of learning of several tasks quite well. However, it is possible that an analysis of other timescales (i.e., other eigenvalues) could provide a more specific description of the learning curves in the tasks. This could be especially important in tasks in which the time frame of the experiment does not allow convergence to zero error, as is the case in some of the experiments we conducted here. Additionally, considering other timescales could provide further differentiation between candidate models to describe learning in a task.

4.3. Limitations

One limitation of our framework is that it assumes that updates to the inverse model are obtained based on the batch gradient, which attempts to minimize the error for all target forces simultaneously. That is, errors for the whole set of targets are known and used for each learning update. However, in humans, learning takes place mainly on a trial-by-trial basis, although a memory of previous errors can contribute to learning (Herzfeld et al., 2014; Sugiyama et al., 2023). This shorter learning timeframes are more akin to online or stochastic gradient descent. Therefore, because the batch and the single-trial gradient are different in general, the actual limit of the speed of learning in the model may differ from the theoretical estimates. In practice, however, we found that using this assumption still adequately predicts the relative learning speed in different tasks both in simulations that use online learning, and experiments.

A further limitation in the analytical results of our computational framework is that we have made a simplifying assumption regarding the role of the target distribution. The eigenvalues of the Hessian matrix of the distal learning system are a function of the distribution of RBF inputs to the network rather than simply the two-dimensional distribution of target forces (Eq. (17)). Therefore, variables such as the number, centers and shapes of the RBF units also impact the eigenvalues of the Hessian. Unfortunately, these properties of the RBFs are hard or impossible to control in an experimental setup. However, because the RBF units encode the desired forces, the distribution of RBF inputs must also encode information about the distribution of target forces. This makes the distribution of target forces an ideal proxy for the distribution of RBF inputs, as it is an easily controllable variable.

Another limitation comes from our choices in the experimental design. Due to restrictions arising from the COVID-19 pandemic, we conducted the experiments using an online platform (Tsay et al., 2021). Although the results obtained from the online platform are similar to in-person results for VMR tasks (Avraham et al., 2021; Tsay et al., 2021, 2022), the arm posture in our experiment should be controlled, as it is related to the shape of the manipulability ellipse of the arm because the posture determines the arm Jacobian (Yoshikawa, 1985). We instructed the participants to adopt a similar posture across all experimental conditions, but did not control for the posture. However, considering that changes in arm posture during the experiment are unlikely to have as dramatic an effect on the shape of the manipulability ellipse as the visuomotor transformations that we used, any effects of changes in arm posture are likely not important.

An additional limitation stemming from our choices on experimental design is that our computational framework considers an isometric task in which the forward physics of the arm are linear. We referred to the experimental task as "quasi-isometric", as the cursor movements are controlled by small finger and/or arm movements (ideally around 1 cm in end-point movement extent, without appreciable changes in arm joint angles). Therefore, even though the forward physics of the arm are not linear because of these movements, the movement extent is small enough that a local linear approximation may be adequate for its description. Nevertheless, previous research in motor control and learning in arm reaching tasks shows that modeling reaching tasks as a linear system can have good descriptive and predictive power (Smith et al., 2006).

4.4. Conclusion

In conclusion, we have developed a framework to theoretically identify factors that influence the speed of learning. This could potentially be used in more complex tasks to systematically identify body and environment parameters that facilitate learning, as well as to evaluate the difficulty of the task based on the elements of the task. Additionally, it could be used to systematically find candidate hypotheses about model architecture and components in the CNS to describe learning in a task. Our results show the relevance of the theory of artificial neural networks to understand the mechanisms underlying learning in the CNS, as advocated in recent perspectives (Lillicrap et al., 2020; Marblestone et al., 2016; Richards et al., 2019).

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Funding

This work was supported by the Japan Society for the Promotion of Science (JSPS) [grant numbers KAKENHI 21K17789 to VRB and KAKENHI 19H05728 to YK] and the National Science Foundation [BCS grant number 2216344 to NS].

Supplementary materials

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.neunet.2023.10.049.

References

Avraham, G., Morehead, J. R., Kim, H. E., & Ivry, R. B. (2021). Reexposure to a sensorimotor perturbation produces opposite effects on explicit and implicit learning processes. *PLoS Biology, 19*(3), 1–26. https://doi.org/10.1371/journal.pbio.3001147

Barradas, V. R., Kutch, J. J., Kawase, T., Koike, Y., & Schweighofer, N. (2020). When 90% of the variance is not enough: Residual EMG from muscle synergy extraction influences task performance. *Journal of Neurophysiology*, 2180–2190. https://doi.org/10.1152/jn.00472.2019

Berger, D. J., Gentner, R., Edmunds, T., Pai, D. K., & D'Avella, A. (2013). Differences in adaptation rates after virtual surgeries provide direct evidence for modularity. *Journal of Neuroscience, 33*(30), 12384–12394. https://doi.org/10.1523/JNEUROSCI.0122-13.2013

Bishop, C. M. (1995). *Neural networks for pattern recognition*. Oxford University Press.

Coltman, S. K., Cashaback, J. G. A., & Gribble, P. L. (2019). Both fast and slow learning processes contribute to savings following sensorimotor adaptation. *Journal of Neurophysiology, 121*(4), 1575–1583. https://doi.org/10.1152/jn.00794.2018

Dal'Bello, L. R., & Izawa, J. (2022). Computational role of exploration noise in error-based de novo motor learning. *Neural Networks, 153*, 349–372. https://doi.org/10.1016/j.neunet.2022.06.011

Doya, K. (1999). What are the computations of the cerebellum, the basal ganglia and the cerebral cortex? *Neural Networks, 12*(7–8), 961–974. https://doi.org/10.1016/S0893-6080(99)00046-5

Emken, J. L., Benitez, R., Sideris, A., Bobrow, J. E., & Reinkensmeyer, D. J. (2007). Motor adaptation as a greedy optimization of error and effort. *Journal of Neurophysiology, 97*(6), 3997–4006. https://doi.org/10.1152/jn.01095.2006

Ganesh, G., Haruno, M., Kawato, M., & Burdet, E. (2010). Motor memory and local minimization of error and effort, not global optimization, determine motor behavior. *Journal of Neurophysiology, 104*(1), 382–390. https://doi.org/10.1152/jn.01058.2009

Gilbert, P. F. C., & Thach, W. T. (1977). Purkinje cell activity during motor learning. *Brain Research, 128*(2), 309–328. https://doi.org/10.1016/0006-8993(77)90997-0

Golub, M. D., Yu, B. M., & Chase, S. M. (2015). Internal models for interpreting neural population activity during sensorimotor control. *eLife, 4*, 1–28. https://doi.org/10.7554/eLife.10015. DECEMBER2015.

Hagio, S., & Kouzaki, M. (2018). Modularity speeds up motor learning by overcoming mechanical bias in musculoskeletal geometry. *Journal of the Royal Society Interface, 15*(147). https://doi.org/10.1098/rsif.2018.0249

Herzfeld, D. J., Vaswani, P. A., Marko, M. K., & Shadmehr, R. (2014). A memory of errors in sensorimotor learning. *Science, 345*(6202), 1349–1354.

Horn, R. A., & Johnson, C. R. (1994). *Topics in matrix analysis*. Cambridge University Press.

Ito, M. (1984). *The cerebellum and neural control*. Raven Press.

Jordan, M. I., & Rumelhart, D. E. (1992). Forward models: Supervised learning with a distal teacher. *Cognitive Science, 16*(3), 307–354. https://doi.org/10.1207/s15516709cog1603_1

Kawato, M. (1990). Feedback-error-learning neural network for supervised motor learning. *Advanced Neural Computers*, 365–372. https://doi.org/10.1016/b978-0-444-88400-8.50047-9

Kawato, M. (1999). Internal models for motor control and trajectory planning. *Current Opinion in Neurobiology, 9*(6), 718–727. https://doi.org/10.1016/S0959-4388(99)00028-8

Kawato, M., Furukawa, K., & Suzuki, R. (1987). A hierarchical neural-network model for control and learning of voluntary movement. *Biological Cybernetics, 57*(3), 169–185. https://doi.org/10.1007/BF00364149

Kawato, M., Kuroda, S., & Schweighofer, N. (2011). Cerebellar supervised learning revisited: Biophysical modeling and degrees-of-freedom control. *Current Opinion in Neurobiology, 21*(5), 791–800. https://doi.org/10.1016/j.conb.2011.05.014

Kim, S., Ogawa, K., Lv, J., Schweighofer, N., & Imamizu, H. (2015). Neural substrates related to motor memory with multiple timescales in sensorimotor adaptation. *PLoS Biology, 13*(12). https://doi.org/10.1371/journal.pbio.1002312

Kitazawa, S., Kimura, T., & Yin, P. B. (1998). Cerebellar complex spikes encode both destinationsand errors in arm movements. *Nature, 392*, 2–5. February 1998.

Krakauer, J. W., Ghez, C., & Ghilardi, M. F. (2005). Adaptation to visuomotor transformations: Consolidation, interference, and forgetting. *Journal of Neuroscience, 25*(2), 473–478. https://doi.org/10.1523/JNEUROSCI.4218-04.2005

Krakauer, J. W., Pine, Z. M., Ghilardi, M. F., & Ghez, C. (2000). Learning of visuomotor transformations for vectorial planning of reaching trajectories. *Journal of Neuroscience, 20*(23), 8916–8924. https://doi.org/10.1523/jneurosci.20-23-08916.2000

Kuperstein, M. (1988). Neural model of adaptive hand-eye coordination for single postures. *Science, 239*(4845), 1308–1311.

Lee, D. D., & Seung, H. S. (1999). Learning the parts of objects by non-negative matrix factorization. *Nature, 401*(6755), 788–791. https://doi.org/10.1038/44565

Lee, J. Y., & Schweighofer, N. (2009). Dual adaptation supports a parallel architecture of motor memory. *Journal of Neuroscience, 29*(33), 10396–10404. https://doi.org/10.1523/JNEUROSCI.1294-09.2009

Lillicrap, T. P., Santoro, A., Marris, L., Akerman, C. J., & Hinton, G. (2020). Backpropagation and the brain. *Nature Reviews Neuroscience, 21*(6), 335–346. https://doi.org/10.1038/s41583-020-0277-3

Marblestone, A. H., Wayne, G., & Kording, K. P. (2016). Toward an integration of deep learning and neuroscience. *Frontiers in Computational Neuroscience, 10*, 1–41. https://doi.org/10.3389/fncom.2016.00094. SEP.

Mazzoni, P., Andersen, R. A., & Jordan, M. I. (1991). A more biologically plausible learning rule for neural networks. *Proceedings of the National Academy of Sciences, 88*, 4433–4437. May.

Morehead, J. R., Taylor, J. A., Parvin, D. E., & Ivry, R. B. (2017). Characteristics of implicit sensorimotor adaptation revealed by task-irrelevant clamped feedback. *Journal of Cognitive Neuroscience, 139*. https://doi.org/10.1162/jocn

Noguchi, K., Gel, Y. R., Brunner, E., & Konietschke, F. (2012). NparLD : An R software package for the nonparametric analysis of longitudinal data in factorial experiments. *Journal of Statistical Software, 50*(12). https://doi.org/10.18637/jss.v050.i12

Pierella, C., Casadio, M., Mussa-Ivaldi, F. A., & Solla, S. A. (2019). The dynamics of motor learning through the formation of internal models. *PLoS Computational Biology, 15*(12), 1–25. https://doi.org/10.1371/journal.pcbi.1007118

Richards, B. A., Lillicrap, T. P., Beaudoin, P., Bengio, Y., Bogacz, R., Christensen, A., Clopath, C., Costa, R. P., de Berker, A., Ganguli, S., Gillon, C. J., Hafner, D., Kepecs, A., Kriegeskorte, N., Latham, P., Lindsay, G. W., Miller, K. D., Naud, R., Pack, C. C., & Kording, K. P. (2019). A deep learning framework for neuroscience. *Nature Neuroscience, 22*(11), 1761–1770. https://doi.org/10.1038/s41593-019-0520-2

Rolf, M., & Steil, J. J. (2014). Explorative learning of inverse models: A theoretical perspective. *Neurocomputing, 131*, 2–14. https://doi.org/10.1016/j.neucom.2013.04.050

Sanger, T. D. (2004). Failure of motor learning for large initial errors. *Neural Computation, 16*(9), 1873–1886. https://doi.org/10.1162/0899766041336431

Sarwary, A. M. E., Wischnewski, M., Schutter, D. J. L. G., Selen, L. P. J., & Medendorp, W. P. (2018). Corticospinal correlates of fast and slow adaptive processes in motor learning. *Journal of Neurophysiology, 120*(4), 2011–2019. https://doi.org/10.1152/jn.00488.2018

Schlerf, J., Ivry, R. B., & Diedrichsen, J. (2012). Encoding of sensory prediction errors in the human cerebellum. *Journal of Neuroscience, 32*(14), 4913–4922. https://doi.org/10.1523/JNEUROSCI.4504-11.2012

Schmidt, R. A., Zelaznik, B. H., Frank, J. S., & Quinn, J. T. (1979). Motor-output variability: A theory for the accuracy of rapid motor acts. *Psychological Review, 86*(5), 415–451. https://doi.org/10.1037/0033-295X.86.5.415

Schweighofer, N., Arbib, M. A., & Kawato, M. (1998). Role of the cerebellum in reaching movements in humans. I. Distributed inverse dynamics control. *European Journal of Neuroscience, 10*(1), 86–94. https://doi.org/10.1046/j.1460-9568.1998.00007.x

Schweighofer, N., Doya, K., Fukai, H., Chiron, J. V., Furukawa, T., & Kawato, M. (2004). Chaos may enhance information transmission in the inferior olive. *Proceedings of the National Academy of Sciences of the United States of America, 101*(13), 4655–4660. https://doi.org/10.1073/pnas.0305966101

Schweighofer, N., Doya, K., & Lay, F. (2001). Unsupervised learning of granule cell sparse codes enhances cerebellar adaptive control. *Neuroscience, 103*(1), 35–50. https://doi.org/10.1016/S0306-4522(00)00548-0

Shidara, M., Kawano, K., Gomi, H., & Kawato, M. (1993). Inverse-dynamics model eye movement control by Purkinje cells in the cerebellum. *Nature, 365*(6441), 50–52.

Smith, M. A., Ghazizadeh, A., & Shadmehr, R. (2006). Interacting adaptive processes with different timescales underlie short-term motor learning. *PLoS Biology, 4*(6), e179. https://doi.org/10.1371/journal.pbio.0040179

Spanne, A., & Jörntell, H. (2015). Questioning the role of sparse coding in the brain. *Trends in Neurosciences, 38*(7), 417–427. https://doi.org/10.1016/j.tins.2015.05.005

Sugiyama, T., Schweighofer, N., & Izawa, J. (2023). Reinforcement learning establishes a minimal metacognitive process to monitor and control motor learning performance. *Nature Communications, 14*(1). https://doi.org/10.1038/s41467-023-39536-9. Article 1.

Thoroughman, K. A., & Shadmehr, R. (2000). Learning of action through adaptive combination of motor primitives. *Nature, 407*(6805), 742–747. https://doi.org/10.1038/35037588

Todorov, E., & Jordan, M. I. (2002). Optimal feedback control as a theory of motor coordination. *Nature Neuroscience, 5*(11), 1226–1235. https://doi.org/10.1038/nn963

Tsay, J. S., Haith, A. M., Ivry, R. B., & Kim, H. E. (2022). Interactions between sensory prediction error and task error during implicit motor learning. *PLoS Computational Biology, 18*(3), 1–24. https://doi.org/10.1371/journal.pcbi.1010005

Tsay, J. S., Lee, A. S., Ivry, R. B., & Avraham, G. (2021). Moving outside the lab: The viability of conducting sensorimotor learning studies online. BioRxiv, 1–24.

Valero-Cuevas, F. J. (2009). A mathematical approach to the mechanical capabilities of limbs and fingers. In *Progress in Motor Control, 629* pp. 405–421). https://doi.org/10.1007/978-0-387-77064-2

Werfel, J., Xie, X., & Seung, H. S. (2005). Learning curves for stochastic gradient descent in linear feedforward networks. *Neural Computation, 17*(12), 2699–2718. https://doi.org/10.1162/089976605774320539

Williams, R. J. (1992). Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning, 8*(3), 229–256. https://doi.org/10.1023/A:1022672621406

Wolpert, D. M., Miall, R. C., & Kawato, M. (1998). Internal models in the cerebellum. *Trends in Cognitive Sciences, 2*(9), 338–347. https://doi.org/10.1016/S1364-6613(98)01221-2

Yoshikawa, T. (1985). Manipulability of robotic mechanisms. *The International Journal of Robotics Research, 4*(2), 3–9.

Zhou, P., & Rymer, W. Z. (2004). Factors governing the form of the relation between muscle force and the EMG: A simulation study. *Journal of Neurophysiology, 92*(5), 2878–2886. https://doi.org/10.1152/jn.00367.2004