

Robust Control in Human Reaching Movements: A Model-Free Strategy to Compensate for Unpredictable Disturbances

Frédéric Crevecoeur,^{1,2} Stephen H. Scott,^{3,4} and Tyler Cluff^{5,6}

¹Institute of Information Technologies, Electronics, and Applied Mathematics, ²Institute of Neuroscience, Université Catholique de Louvain, Louvain-la-Neuve 1348 Belgium, ³Centre for Neuroscience Studies, ⁴Department of Biomedical and Molecular Sciences, Queen's University, Kingston, Ontario K7L 3N6, Canada, ⁵Faculty of Kinesiology, and ⁶Hotchkiss Brain Institute, University of Calgary, Calgary, T2N 1N4 Alberta, Canada

Current models of motor learning suggest that multiple timescales support adaptation to changes in visual or mechanical properties of the environment. These models capture patterns of learning and memory across a broad range of tasks, yet do not consider the possibility that rapid changes in behavior may occur without adaptation. Such changes in behavior may be desirable when facing transient disturbances, or when unpredictable changes in visual or mechanical properties of the task make it difficult to form an accurate model of the perturbation. Whether humans can modulate control strategies without an accurate model of the perturbation remains unknown. Here we frame this question in the context of robust control (\mathcal{H}_∞ -control), a control strategy that specifically considers unpredictable disturbances by increasing initial movement speed and feedback gains. Correspondingly, we demonstrate in two human reaching experiments including males and females that the occurrence of a single unpredictable disturbance led to an increase in movement speed and in the gain of rapid feedback responses to mechanical disturbances on subsequent movements. This strategy reduced perturbation-related motion regardless of the direction of the perturbation. Furthermore, we found that changes in the control strategy were associated with co-contraction, which amplified the gain of muscle responses to both lengthening and shortening perturbations. These results have important implications for studies on motor adaptation because they highlight that trial-by-trial changes in limb motion also reflected changes in control strategies dissociable from error-based adaptation.

Key words: motor adaptation; optimal feedback control; reaching control; robust control

Significance Statement

Humans and animals use internal representations of movement dynamics to anticipate the impact of predictable disturbances. However, we are often confronted with transient or unpredictable disturbances, and it remains unknown whether and how the nervous system handles these disturbances over fast time scales. Here we hypothesized that humans can modulate their control strategy to make reaching movements less sensitive to perturbations. We tested this hypothesis in the framework of robust control, and found changes in movement speed and feedback gains consistent with the model predictions. These changes impacted participants' behavior on a trial-by-trial basis. We conclude that compensation for disturbances over fast time scales involves a robust control strategy, which potentially plays a key role in motor planning and execution.

Introduction

Humans and other animals use internal representations of movement dynamics, called internal models, to shape motor com-

mands in anticipation of their interactions with the environment. The update of these internal representations with practice contributes to the neural basis of motor learning and adaptation (Shadmehr et al., 2010; Wolpert et al., 2011). Indeed, several classic studies have shown that when exposed to novel movement dynamics, the resulting errors in internal models produce movement deviations. These movement deviations elicit systematic model updates through a reorganization of motor planning that is based on sensory feedback about the perturbed movement (Thoroughman and Shadmehr, 2000; Hwang et al., 2003). This powerful approach has captured many aspects of motor learning, including its dependency upon statistical properties of disturbances (Singh and Scott, 2003), and the dynamics of learning

Received April 5, 2019; revised Aug. 9, 2019; accepted Aug. 24, 2019.

Author contributions: F.C., S.H.S., and T.C. designed research; F.C. performed research; F.C. contributed unpublished reagents/analytic tools; F.C., S.H.S., and T.C. analyzed data; F.C., S.H.S., and T.C. wrote the paper.

This work was supported by F.R.S.-FNRS (Belgium, Grant 1.C.033.18F) to F.C. We thank Philippe Lefèvre for logistic and many other kinds of support.

S.H.S. is associated with BKIN Technologies that commercializes the robot that was used (KINARM). The remaining authors declare no competing financial interests.

Correspondence should be addressed to Frédéric Crevecoeur at frederic.crevecoeur@uclouvain.be.

<https://doi.org/10.1523/JNEUROSCI.0770-19.2019>

Copyright © 2019 the authors

across trials (Smith et al., 2006; Kording et al., 2007; Gonzalez Castro et al., 2014).

Less attention has been dedicated to understanding how humans control their reaching movements when their internal models are inaccurate or difficult to acquire, as arises due to partial or incomplete adaptation to the environmental dynamics, or due to transient and unpredictable disturbances. Here we question whether there is a neural control strategy that allows the nervous system to mitigate motor errors in situations where there is no model of the encountered disturbance available. The presence of such a strategy would impact our understanding of the neural bases and timescales of motor adaptation.

It has been suggested (Shadmehr and Mussa-Ivaldi, 1994; Burdet et al., 2001; Wagner and Smith, 2008) that the limb's intrinsic mechanical impedance counters deviations induced by inaccurate internal models. However, common techniques substantially overestimate the influence of limb impedance (Crevecoeur and Scott, 2014), because they incorporate the influence of not only the limb, but also of neural feedback (Burdet et al., 2000).

Stochastic optimal feedback control [linear quadratic Gaussian (LQG)] was introduced as a model of neural feedback control (Todorov and Jordan, 2002). Although this model handles sensorimotor noise efficiently, it does not consider possible *unmodeled* disturbances. An alternative strategy is *robust control* (\mathcal{H}_∞ -control; Basar and Bernhard, 1991), which can be defined as a model-free control design that is insensitive to model errors as possible. We refer to this strategy as a “model-free” strategy in the sense that it does not require knowledge of potential disturbances, and thus it is suitable for unmodeled perturbations. Model-free control has also been reported as a form of adaptation in a visuomotor rotation task (Huang et al., 2011), where actions were reinforced because they were successful. The parallel with robust control is that actions or strategies may thus be selected without a precise model of the environment.

Robust and LQG are two optimal solutions based on different assumptions about potential disturbances. In practice there is a continuum of control solutions, and the two can overlap with different sets of cost parameters. Our goal was not to demonstrate whether one control model fitted behavior better than the other. Instead, we examined how motor patterns were altered when a control strategy considered unmodeled disturbances, without any change in task demand. This is particularly important for motor adaptation studies, because a common assumption is that trial-to-trial changes in limb motion reflect adaptation of the internal models of movement dynamics. Here we investigated whether trial-to-trial changes in limb motion may include a change in control strategy.

We characterize a robust control design that predicts an increase in control gains for planar reaching movements. We then show that following a single perturbation trial, healthy humans increased the speed of their movements and the amplitude of their feedback responses to perturbations in a way that was consistent with the simulations of robust control. Interestingly, the changes in control strategy correlated with co-contraction, which produced a transient amplification of fast feedback responses to perturbation loads. Together our results highlighted a rapid and model-free compensation for unexpected disturbances that impacts learning and control.

Materials and Methods

Experimental procedures

A total of 23 healthy volunteers (11 females, between 19 and 45 years) participated in this study after providing written informed consent and following procedures approved by the Ethics Committee at the host institution (UCLouvain, Belgium).

Movements consisted of 15 cm forward reaches toward visual targets. In both experiments, the visual display, task instructions, and temporal constraints were as follows. Participants were instructed to grasp a robotic handle and move a hand-aligned cursor to the home target (radius 0.6 cm). The home target was initially red but turned green when the participant's cursor entered the target. The goal target was presented as an open red circle (radius: 1.2 cm) located 15 cm directly in front of the start position. After a random delay following stabilization in the home target (between 2 and 4 s, uniformly distributed), the goal target filled in and provided a “go signal” for participants to begin their movement. If participants reached the goal target in <0.6 s following the go cue, the target turned back to an open circle to indicate that they reached it too soon. When they reached the goal target after >0.8 s, the target remained red indicating that they took too long. The goal target turned green when participants reached it within the prescribed time window (0.6–0.8 s). The trial was considered successful if the hand-aligned cursor remained stable in this target for 1 s. There were no constraints on movement speed, only the arrival time, including the reaction time, was constrained. These instructions were used to maintain similar movement speeds but all trials were included in the dataset. The two experiments used the same protocol in the variants described below. Hereafter, baseline trials refer to trials performed in the null field, without any perturbation applied during movement.

Experiment 1. Participants ($n = 10$) first performed a practice series of 10–20 trials in the null field dependent on whether they felt comfortable with the task and instructions. They then performed 25 baseline trials without any perturbation (null field). Participants were explicitly told they would not encounter any disturbance during movement. These trials correspond to a “pre-exposure” phase before the trials with uncertainty about the task dynamics. Then, they performed a series of 6 blocks of 60 trials composed of 50 null field trials and 10 trials where curl force fields were randomly interleaved as catch trials [5 clockwise (CW), 5 counterclockwise (CCW)]. These six blocks corresponded to a “peri-exposure” phase where the dynamics were unpredictable and varied randomly from trial to trial. During this phase, participants were told that the robot could perturb their movements. Finally they performed another 25 null field trials (“post-exposure”), and were told that there would not be any perturbation applied during movements. Thus, the pre-exposure and post-exposure phases were similar in the sense that there was no uncertainty about the task dynamics.

The force field applied during peri-exposure was a standard curl force field where movement velocities were mapped onto a perturbation force as follows:

$$\begin{bmatrix} F_x \\ F_y \end{bmatrix} = \begin{bmatrix} 0 & L \\ -L & 0 \end{bmatrix} \begin{bmatrix} \dot{x} \\ \dot{y} \end{bmatrix}, \quad (1)$$

where the x and y coordinates correspond to lateral and forward axes relative to the reach path, $F_{x,y}$ are the force components along each axis and the dots indicate time derivatives. The value of L was 15 Ns/m for CW and -15 Ns/m for CCW perturbations.

Experiment 2. This experiment was designed to study feedback responses to step perturbations and gain insight into the neurophysiological basis of changes in control strategies (Scott, 2016). The step perturbations consisted of a rightward or leftward constant load (12 N, 10 ms linear buildup) applied when the hand path crossed a virtual line corresponding to one-third of the reach path. The step perturbations were always applied during null field trials. Participants first reached in a null field (baseline) with randomly interleaved step perturbations (2 blocks with 10 leftward, 10 rightward, and 40 null field trials per block), followed by four blocks with null field trials (30 per block), step perturbations (5 per direction), curl fields (Eq. 1), and orthogonal fields defined as follows (5 per direction for each FF, $L = \pm 13$ Ns/m):

$$\begin{bmatrix} F_x \\ F_y \end{bmatrix} = \begin{bmatrix} 0 & L \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \dot{x} \\ \dot{y} \end{bmatrix}. \quad (2)$$

The peri-exposure phase was then followed by two post-exposure blocks that consisted of null-field trials and step perturbations similar to the pre-exposure phase. We reasoned that the presence of orthogonal and curl fields would increase uncertainty about the dynamics during the peri-exposure phase compared with Experiment 1. However, as noted in the Results, participant behavior in Experiment 2 was distinct and the analyses were based on a classification of trials into one of two groups that depended on the number of preceding baseline trials. The median number of baseline trials that preceded each step perturbation or unperturbed trial was calculated for each participant, and used to separate trials into two groups that fell above or below the median. Note that this classification was independent of the trials, and only depended on the preceding sequence of trials.

Data collection and analysis

The two experiments were performed with an endpoint KINARM (BKIN Technologies). The two-dimensional coordinates of the hand-aligned cursor and components of the endpoint force were sampled at 1 kHz. The cursor velocity was obtained numerically with a fourth-order central-differences algorithm. In Experiment 2, we recorded the activity of mono-articular shoulder muscles (pectoralis major and posterior deltoid). The electrodes were attached to the skin above the muscle belly after light abrasion with alcohol. The signal was amplified (gain: 10^4), digitally bandpass filtered with a dual-pass, fourth-order Butterworth filter (10–500 Hz bandpass). EMG data were then normalized to the average activity across 1 s recorded when participants maintained postural control at the home target against a background load of 12 N applied three times in each direction. This calibration was performed at the end of the second and sixth blocks.

The variables extracted were the peak hand velocity calculated on a trial-by-trial basis, as well as the hand velocity at perturbation onset for step perturbation trials of Experiment 2. Kinematic variables were then averaged across trials for group data analysis. Similar to the analysis of movement kinematics, we classified both baseline and perturbation trials for Experiment 2 based on the number of baseline trials that preceded them, and split the data based on the median index. The two classes of trials are thus defined as the trials occurring “Early” or “Late” relative to the last perturbation trial. EMG activity was collapsed into the following epochs defined relative to perturbation onset: (Pre) pre-perturbation activity [−50, 0] ms; (1) short-latency [20, 50] ms; (2) long-latency [50, 100] ms; (3) early voluntary [120, 180] ms; and (4) voluntary [200, 250] ms.

Statistical design

Peak hand velocity from Experiment 1 was first analyzed within each participant. Changes in peak hand velocity were assessed with the Wilcoxon rank sum test to investigate whether the distribution of peak forward velocities in the peri-exposure phase was distinct from the distributions during the pre-exposure and post-exposure phases of the experiment. The p level chosen for this analysis was 0.05. In conjunction with the observation that a relatively large number of participants exhibited the same effect, this level of significance was sufficient to control for the rate of false-positive discoveries. In Experiment 1, we further analyzed group data using a repeated-measures ANOVA (rmANOVA) to assess the effect of changes in context (pre-, peri-, post-) on the average peak hand velocity averaged across trials. The trial-by-trial effect was also assessed based on the same test applied to the average peak velocity of trials in the pre-exposure and post-exposure phases, as well as trials that fell within 1–12 movements for Experiment 1, and up to 7 in Experiment 2 (only a minority of participants experienced several trials with larger indices). *Post hoc* analyses were reported based on comparisons between the data in each index group for peri-exposure trials and the peak hand velocity during pre-exposure. Sphericity was tested using Mauchly’s test and the results of the rmANOVA were further considered after Greenhouse–Geisser and Huynh–Feldt ϵ corrections when applicable. *Post hoc* tests were performed with one-sided paired t tests. One-sided compari-

sons were warranted by the model predictions, as we were interested in testing whether the peak forward velocities were higher during the peri-exposure phase. Possible after-effects in Experiment 1 were measured based on the maximum lateral displacement during unperturbed trials.

Group data of trials occurring early or late following perturbation trials in Experiment 2 were analyzed based on one-tail paired t tests. Again the one sided comparison was warranted because a directional effect was predicted in theory. Changes in the time profile of lateral hand velocities were assessed with a sliding t test. The moment when the contrast of lateral hand velocities dropped <0.01 was reported in the analysis after verifying they were followed by very small p values. Corrections for multiple comparisons do not apply in this case because consecutive samples are not independent, and because there is only one comparison performed at each time step. We also extracted the norm of the endpoint error averaged across trials for Early and Late trials, as well as the area of endpoint dispersion ellipses calculated at 800 ms following reach onset to characterize endpoint distributions. This analysis was performed on the data of Experiment 2 across Early and Late categories to calculate the variances on the same number of trials.

The analysis of EMG recordings was based on repeated-measures ANOVA with the classification (Early and Late) and response epochs (Pre, SLR, LLR, early Vol, and Vol) as factors. Finally, comparisons of the amplitude of after-effect and of the maximum hand displacement in force fields across experiments was based on Wilcoxon rank sum tests to assess differences in the distributions of individual trials for each participant. This nonparametric test on grouped data was used as the two experiments were performed with distinct groups of participants, and thus comparisons could not be paired. The rank sum tests were based on two-tailed comparisons because there was no reason to expect directional differences a priori. Significance was considered at the level $p < 0.05$, although our interpretations and the reported comparisons (with the exception of a few comparisons for an EMG data) were based on statistical differences at the level $p < 0.005$ (Benjamin et al., 2018).

Model

The system model describes the translation of a point mass ($m = 1$ kg) in the horizontal plane against dissipative viscous forces. The controlled force is modeled as a first-order low-pass response to the control vector as a linear approximation of muscle dynamics. We must acknowledge from the outset that the translation of a point mass is a simplified model of the complex and nonlinear nature of even simple biological systems like the multi-jointed arm. In practice, we have used linear models for simplicity and constrained experimental testing accordingly (Crevecoeur and Scott, 2013, 2014). In the context of the present paper, the linear model is in part imposed by the limitation of current robust control theories. We used a robust control design based on game-theoretic principles (\mathcal{H}_∞ -control; Basar and Bernhard, 1991), which has the advantage of being a natural extension of the more common stochastic optimal control framework (Bitmead et al., 1990). The downside of this approach is that it forces us to consider a linear state-space representation, as well as constraints imposed to the cost-function that guarantee a solution exists. An extension of this control design would be to linearize the system around a state or a trajectory and derive locally optimal controllers that are updated across time or space. This can also be done for stochastic optimal control (Li and Todorov, 2007). However, we believe that such an extension is beyond the scope of this study and suggest that the mathematical limitations do not hinder the use of this model to illustrate some underpinnings of biological control.

We first concentrate on force-field trials and note that for completeness, an additional variable capturing the external force is needed to model step-perturbations (Crevecoeur and Scott, 2013). The continuous-time differential equations are as follows:

$$\dot{x} = -k\dot{x} + \theta_x\dot{y} + F_x, \quad (3)$$

$$\dot{y} = -k\dot{y} + \theta_y\dot{x} + F_y, \quad (4)$$

$$\tau\dot{F}_x = u_x - F_x, \quad (5)$$

$$\tau\dot{F}_y = u_y - F_y. \quad (6)$$

where x and y are the coordinates of the workspace and the dot(s) represent the time derivative. The variables u_x and u_y are the control commands. The viscous constant was set to $k = 0.1$ Ns/m and the time constant of the linear muscle model was set to $\tau = 60$ ms (Brown and Loeb, 2000). The parameters θ_x and θ_y are unknown from the point of view of the controller. In the case of a baseline trial, these parameters are $\theta_x = \theta_y = 0$. For an orthogonal force field, we have $\theta_x = \pm L$ Nsm $^{-1}$, and $\theta_y = 0$. For a curl field we have $\theta_x = \pm L$ Nsm $^{-1}$ and $\theta_y = -\theta_x$.

In our experiment, all trials were randomly interleaved. Thus there was uncertainty about movement dynamics. It is important to stress that this uncertainty was random from trial to trial, but it was not random within a trial, as previously assumed (Izawa et al., 2008). Thus the uncertainty for a given trial is not zero on average and the problem of control with fixed model errors must be considered. We express this mathematically by rewriting the differential equation in an algebraic form and grouping the unknown terms (θ_x and θ_y) into a model disturbance that is unknown to the controller (ΔA). Defining the state vector as follows: $\mathbf{x} = [x, y, \dot{x}, \dot{y}, F_x, F_y]^T$ and the control vector as $\mathbf{u} = [u_x, u_y]^T$ we have (note that we used boldface \mathbf{x} for the state vector, and normal x for the coordinate):

$$\dot{\mathbf{x}} = (A + \Delta A)\mathbf{x} + B\mathbf{u}. \quad (7)$$

The model disturbance impacts the unforced dynamics (A) in agreement with the nature of the perturbation induced by the force field. In general there can also be disturbances to the matrix B . Taking noise into account, we rewrite this equation based on instantaneous differences instead of derivatives, and include stochastic disturbances as follows:

$$d\mathbf{x} = [(A + \Delta A)\mathbf{x} + B\mathbf{u}]dt + Cdw. \quad (8)$$

Equation 8 introduces the standard Brownian motion w of appropriate dimension and C is a scaling function that characterizes the noise properties. Observe that C is allowed to depend on control or state variable and capture signal-dependent noise. We now rearrange terms as follows:

$$d\mathbf{x} = [A\mathbf{x} + B\mathbf{u}]dt + \Delta A\mathbf{x}dt + Cdw, \quad (9)$$

or equivalently:

$$d\mathbf{x} = [A\mathbf{x} + B\mathbf{u}]dt + \varepsilon(\mathbf{x}, w), \quad (10)$$

where $\varepsilon(\mathbf{x}, w) = \Delta A\mathbf{x}dt + Cdw$ is the unknown disturbance that lumps together the impact of process noise and the fixed model error. Equation 10 expresses that the dynamics corresponds to the modeled dynamics, plus an unknown disturbance function.

The control problem consists in deriving a time-varying control law that minimizes a performance index. When the process disturbance is purely stochastic with zero mean and known covariance matrix, that is $\Delta A = 0$ and $\varepsilon(\mathbf{x}, w) \sim N(0, \Sigma_\varepsilon)$ the solution can be derived in the framework of stochastic optimal control (Todorov, 2005). When $\delta A \neq 0$, there is a fixed bias or error in the expected dynamics, and the formalism of stochastic optimal control does not apply. A dedicated control design consists in deriving a control solution in the worst-case scenario, including fixed model errors or possibly worse input disturbances. This is formalized in the framework of robust, or \mathcal{H}_∞ -control (Basar and Bernhard, 1991). It is important to stress that the two control designs produce a goal-directed, state-feedback control law that is consistent with behavioral results in a wide range of perturbation paradigms (Crevecoeur and Kurtzer, 2018).

Briefly, the controller design is expected to minimize a quadratic cost-function that captures the intended behavior (in this case a reaching task). We call $J(\mathbf{x}, \mathbf{u})$ the quadratic penalty at each time step. For robust control design, because the controller must minimize the cost-function, it is assumed that a second “player” is trying to maximize the same cost by manipulating the unknown disturbance function $\varepsilon(\mathbf{x}, w)$. Thus, the robust control problem is to find a control sequence \mathbf{u}^* , which minimizes the maximum of $J(\mathbf{x}, \mathbf{u})$ assuming a worst-case disturbance. This formulation requires considering the following augmented cost-function:

$$\tilde{J}(\mathbf{x}, \mathbf{u}, \varepsilon) = J(\mathbf{x}, \mathbf{u}) - \gamma^2 \varepsilon^T \varepsilon, \quad (11)$$

and the robust optimal controller in the sense of \mathcal{H}_∞ -control is defined as the control function that minimizes the time integral until the final time (t_f) of the running cost plus the final cost in the worst case scenario, or equivalently when this cost is maximized over ε :

$$\mathbf{u}^* = \operatorname{argmin}_{\mathbf{u}} (\max_{\varepsilon} \int_0^{t_f} \tilde{J}(\mathbf{x}, \mathbf{u}, \varepsilon) dt + \mathbf{x}_f^T Q \mathbf{x}_f). \quad (12)$$

The parameter γ must be jointly optimized to ensure that the solution exists and corresponds to \mathcal{H}_∞ -optimal control (Basar and Bernhard, 1991). This parameter was optimized through numerical search stopped when the relative improvement in its value before violating conditions of existence was $\leq 0.1\%$. When the process disturbance does not contain a fixed error, the expected value of the second term of the right-hand side of Equation 11 is equal to the variance of ε , which is constant, and the control problem reduces to the minimization of the expected value of $J(\mathbf{x}, \mathbf{u})$. This problem is handled in the usual framework of stochastic optimal control (Todorov, 2005).

Finally, the control problem was transformed into a discrete time system based on Euler integration with 10 ms time steps. The transformation into a discrete time system allowed us to consider sensorimotor delays through system augmentation. We use 50 ms (6 time steps) in agreement with the long-latency pathway through cortex (Scott, 2016). We use the index “ d ” to refer to the discrete-time representation and the subscript “ t ” indicates the time step. With these definitions, we write the discrete time system as follows:

$$\mathbf{x}_{t+1} = A_d \mathbf{x}_t + B_d \mathbf{u}_t + \varepsilon_t, \quad (13)$$

with ε_t containing the potential model bias and a noise term defined as $\xi_t \sim N(0, \alpha_1 \Sigma_\xi)$. The information available to the controller is as follows:

$$y_t = x_{t-h} + \sigma_t, \quad (14)$$

with $\sigma_t \sim N(0, \alpha_2 \Sigma_\sigma)$. The definition of the cost is scaled such that the penalty of the command has unity weight, and the target was mapped to the origin of the workspace. This cost-function corresponds to:

$$J(\mathbf{x}, \mathbf{u}) = \sum_{t=1}^N \alpha_3 (t/N)^6 (\mathbf{x}_t^T Q \mathbf{x}_t) + \sum_{t=1}^{N-1} \|\mathbf{u}_t\|^2. \quad (15)$$

Previous studies considered only a terminal penalty on the state, and no penalty during movement. This is not allowed for robust control, indeed a condition for the existence of solution in the sense of \mathcal{H}_∞ -control is that the matrix Q be positive definite. We use the following: $Q = \operatorname{diag}([10^6, 10^6, 10^5, 10^5, 1, 1])$ for the terminal cost, and scaled it by the factor $(t/N)^6$ to ensure low initial cost while positive definite, and a smooth build up until the terminal cost. The cost of forces was kept low compared with minimize their impact while keeping a positive definite Q . As the behavior of the controller depends on the ratio between state and command-related costs, we varied α_3 over an order of magnitude to investigate the impact of the cost-function on the model behavior. Similarly the model predictions were derived through changes in the intensity of the motor and sensory noise factors (α_2 and α_3).

In all, the only free parameters were α_i , $i = [1, 2, 3]$ and we report the impact of changing these parameters over a wide range to address the sensitivity of the corresponding model. In all cases, the same perturbations and noise were applied to simulations derived in the context of stochastic or robust optimal feedback control. In the case of stochastic disturbances the problem was solved in the context of extended linear-quadratic-Gaussian regulator (Todorov, 2005). Otherwise the control problem was solved in the framework of robust \mathcal{H}_∞ -control (Basar and Bernhard, 1991). In general, the derivation for \mathcal{H}_∞ -control is based on matrix recurrences that extend the framework of LQG to the cases where disturbances are not Gaussian. Details about these recurrences are provided in the references above, and an application to human reaching movements without experimental validation was proposed previously by Ueyama (2014). A detailed description of the derivation is beyond the scope of the present paper, but access to working code for replication can be found here: <http://modeldb.yale.edu/258846>. We should underlie that both control designs compensated for delays through system augmenta-

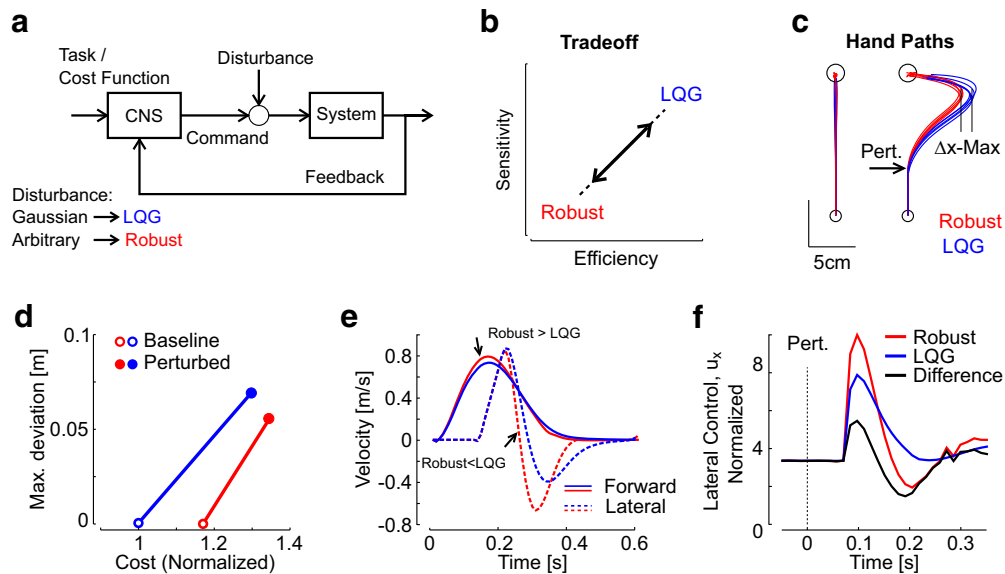


Figure 1. Model predictions. **a**, Schematic representation of a control system where the controller is the CNS, and includes a state estimator and a motor command generator. LQG minimizes the expected cost of movements when disturbances follow a zero mean Gaussian distribution, whereas robust control mitigates the impact of arbitrary (bounded) disturbances. **b**, Schematic illustration of the tradeoff between efficiency and sensitivity to model errors. LQG generates efficient movements under the assumption that there is no fixed model error, whereas robust control generates costly movement solutions that are insensitive to model errors. **c**, Simulated hand path to the target with (right) and without (left) step perturbation loads applied to the system. The movement amplitude, movement time, and perturbations were similar to the experimental design. The trajectories were generated with a robust controller (red) or a stochastic optimal controller (blue). **d**, Consequences of the tradeoff captured by the total cost computed averaged across 100 simulation runs in each condition (baseline, open; perturbed, filled). The logarithm of the cost was normalized to the cost of LQG control in unperturbed conditions for illustration. **e**, Forward and lateral velocity averaged across 100 simulation runs with identical color code as in **b**. The increase in control gains for the robust controller produces greater velocity toward the target (forward velocity: robust > LQG), and a rapid reduction of lateral velocity when step-force disturbances are applied to the system (lateral velocity: robust < LQG). **f**, Control response following the lateral perturbation applied to the system averaged across 100 simulation runs. Observe that the robust strategy yields a transient increase in control response.

tion and allowed the increase in control gains observed for the robust controller. In the absence of accurate delay compensation such an increase in control gains may be problematic (Michiels and Niculescu, 2007; Crevecoeur and Scott, 2014).

Results

Model

We first describe the properties of robust controllers (\mathcal{H}_∞ -control) by contrasting the predictions of this control design with those obtained from a stochastic optimal controller (LQG). Under the assumption of stochastic disturbances (which are zero on average; Fig. 1a), the control and estimation algorithms minimize the expected cost of movements (LQG; Todorov, 2005). When the model error includes nonrandom disturbances, such as for instance fixed biases in the model parameters, robust control minimizes the impact of these disturbances in the worst-case scenario (Fig. 1a; min-max or \mathcal{H}_∞ -control; Basar and Bernhard, 1991). Optimizing the performance index for the worst-case scenario effectively minimizes the sensitivity of the control design to any kind of disturbance (in particular to the worst case), and thus this design is advantageous where there is no model of the disturbance available.

The reason why robust control is not a good default choice is because, although this controller is less sensitive to unmodeled disturbances, it is also costlier because it typically produces conservative strategies. In our examples, robust control responds to noise as if it were disturbances, which is not desirable in the case of zero-mean stochastic disturbances. This is a well known property of control design: there is an inherent tradeoff between efficiency and robustness (Fig. 1b). Intuitively, a costly but safe control strategy is warranted in the presence of dynamic uncertainty (robust control), whereas a more predictable context may

promote an efficient strategy at the cost of sensitivity to unmodeled disturbances (LQG control).

The robust controller exhibited larger feedback gains, which for linear systems are the same control gains that steer the system to the target and respond to perturbations. This increase in control gains resulted in faster movement velocities toward the target and more vigorous responses to external perturbations applied to the virtual mass. Exemplar trajectories are presented in Figure 1c–e: the robust controller displayed higher forward velocity for unperturbed trials (Fig. 1e, forward velocity, solid), and limited the perturbation-evoked velocity following lateral disturbances (Fig. 1e, lateral velocity, dashed). The faster reduction of lateral velocity was due to a phasic increase in control response shown in Figure 1f. Indeed, because the actuator was linear, the increases in feedback gains amplified the perturbation response and generated a phasic modulation of the control response (Fig. 1d, dashed trace).

The tradeoff between sensitivity and robustness was apparent when we calculated the cost of the simulated movements. Compared with the LQG controller, we found the cost of movements generated by the robust controller was $\sim 15\%$ higher (normalized to the average cost of LQG control) even for unperturbed trials (Fig. 1d). However, the perturbation had a proportionally smaller impact on the robust controller. The increase in cost associated with step-load perturbations for this controller was $< 20\%$ compared with the LQG controller, which exhibited close to 30% increase because of larger lateral hand displacements.

We attempted to limit the dependency of our predictions on model parameters to the extent possible. First, the parameters of the biomechanical system were fixed a priori based on measured or standard values (see Materials and Methods). Second, the

noise covariance matrices and cost-function were set based on the fact that they generated trajectories that resembled human behavior when participants performed movements of the same amplitude under the same time constraints. Then, we scaled the cost-function and noise covariance matrices over an order of magnitude to assess how much they impacted the model predictions. Qualitatively, the following predictions were independent of the cost or noise parameters: (1) the robust controller always generated faster movement velocities toward the target; and (2) when disturbed by the same perturbation, the robust controller responded more vigorously, resulting in smaller lateral deviations and velocities than the stochastic optimal controller.

Experiment 1

We sought to investigate whether unexpected disturbances impacted human reaching control in a way that was consistent with the predictions of robust control presented above. We first contrasted the behavior of individual subjects across the pre-, peri-, and post-exposure phases to unpredictable disturbances (Fig. 2*a*). The prediction that uncertainty about dynamics modulated control gains was clearly borne out in our results. Figure 2*b* illustrates the modulation of forward hand velocities from one participant chosen to illustrate the main effect. The peak hand velocities during peri-exposure were larger than the same data in the pre-exposure and post-exposure conditions as emphasized by the arrows aligned with this participant's average velocity.

Group data from all participants are presented in Figure 2*c* based on 20 trials taken from the pre-, peri-, and post-exposure conditions. We selected 20 evenly spaced trials during the peri-exposure phase simply to represent the data with the same number of trials as the pre- and post-exposure phase. For statistical comparisons, all trials from the peri-exposure were included. The increase in hand speed during the peri-exposure blocks was significant for 7/10 participants when assessed on the distribution of individual trials (peri > pre; Fig. 2*d*, left). Moreover, the decrease back to near-baseline velocities (i.e., pre-exposure) in the post-exposure phase was significant for 8/10 participants (post < peri; Fig. 2*b*, right; Wilcoxon rank sum test, $p < 0.05$). Accordingly, group-level data revealed a highly significant effect of the condition (rmANOVA: $F_{(2,18)} = 14.3$, all corrected $p < 0.001$). *Post hoc* comparisons based on paired t tests confirmed the increase in forward hand velocity during the peri-exposure phase (one-sided comparisons: $Vel_{PRE} < Vel_{PERI}$: $t_{(9)} = 4.64$, $p < 0.001$; $Vel_{PERI} > Vel_{POST}$: $t_{(9)} = 6.11$, $p < 10^{-4}$; there was no statistical difference between Vel_{PRE} and Vel_{POST} , two-sided comparison: $t_{(9)} = 0.62$, $p = 0.55$). Average forward hand velocities in the peri-exposure and other phases of the task are shown in Figure 2*d*. These data only included the baseline trials that were preceded by >2 baseline trials to remove a possible influence of after-effects evoked by the perturbation. The analysis below characterizes the lateral deviations induced by after-effects in more detail. The filled and open dots indicate participants who did ($p < 0.05$) or did not display ($p > 0.05$) significant changes in forward velocity across the different phases of the experiment.

More interestingly, we uncovered a trial-by-trial change in control strategy, where a single perturbation was followed immediately by an increase in average forward hand velocity in the next trial. This result is illustrated in Figure 3*a* for a time series of 100 trials selected from an exemplar participant. Observe the participant's forward hand velocities often increased in the trial immediately after an unexpected force field. Moreover, the forward hand velocities decayed gradually during each sequence of unperturbed movements. This result was not expected a priori and was

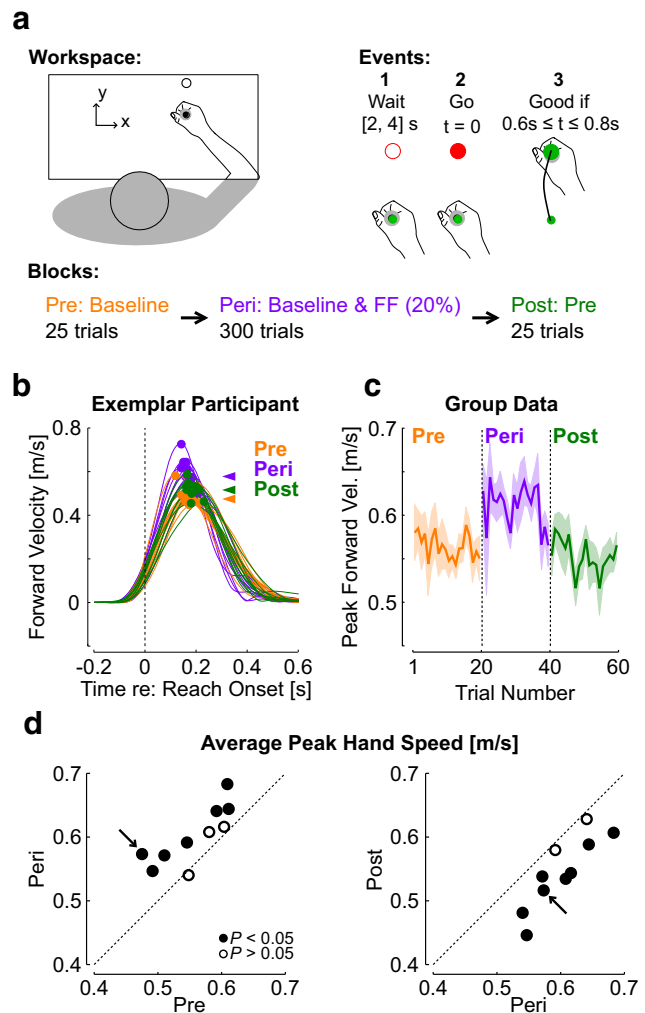


Figure 2. Experiment 1: behavior. **a**, Illustration of the experimental procedures: participants performed visually guided reaching movements in the forward direction while grasping the handle of a robotic manipulandum. They performed movements with or without force-field (curl field) trials randomly interleaved. See Materials and Methods for more details. **b**, Hand velocity of individual trials from one representative participant. Ten traces were randomly selected for illustration from the pre- (orange), peri- (purple), and post-exposure to uncertain context (green). The peak hand speed computed for each trial is highlighted with dots following the same color code and the triangles are aligned with the average peak velocity in each condition. **c**, Group data of forward hand velocity in the pre-, peri-, and post-exposure to randomly applied disturbances. The figure shows the mean \pm SEM across 20 trials in each phase (trials chosen from the peri-exposure were evenly spaced, $n = 10$ subjects). **d**, Left, Individual change in average peak velocity in the peri-exposure as a function of the average from the pre-exposure. Each dot represents one participant. The filled dots represent the participants who exhibited significant differences in the distribution of individual trials (Wilcoxon rank sum test, $p < 0.05$). Right, Same as the left for the data post-exposure as a function of the average in the peri-exposure. The arrow points to the participant chosen to illustrate the main effect in **b**. FF, Force Field.

observed during exploratory analyses as we investigated the time scale of changes in control strategies. We quantified this effect at the group level by sorting unperturbed trials as a function of their index following the last force-field perturbation. This analysis revealed a substantial increase in average peak forward velocity recorded in baseline trials following a force-field perturbation (Fig. 3*b*; rmANOVA, $F_{(13,117)} = 10.7$, all corrected $p < 10^{-6}$). Movement velocities decayed back to pre-/post-exposure values within ~ 10 undisturbed trials. *Post hoc* comparisons were used to contrast the average peak velocity in the post-exposure phase to

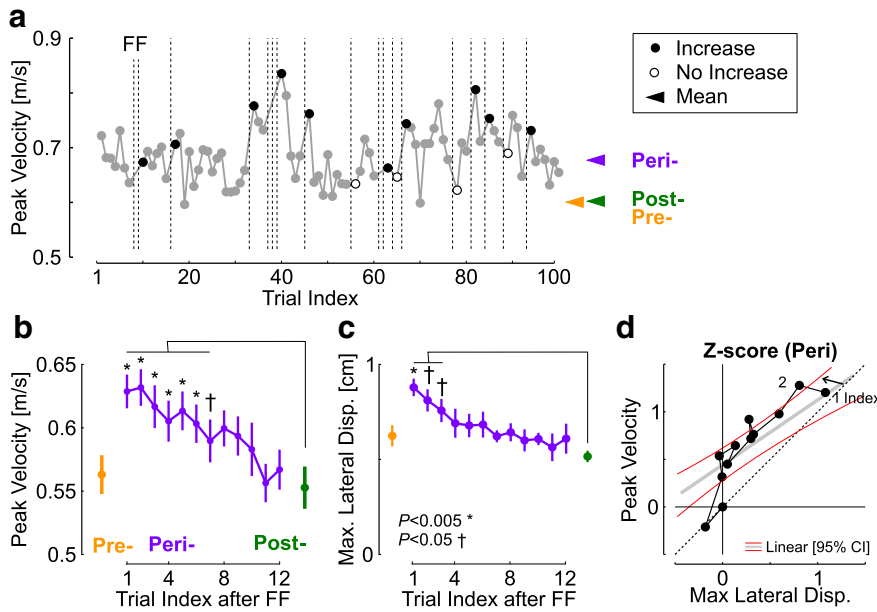


Figure 3. Trial-by-trial modulation. **a**, Illustration of trial-by-trial changes in control strategy for one selected series of trials. Dots represent the time series of peak hand velocity across baseline trials. Vertical lines indicate the occurrence of perturbation trials. Black dots highlight the baseline trials for which there was an increase in hand velocity following a force field. Open dots represent the baseline that followed a force field but did not display an increase in forward peak hand velocity after being exposed to a force field. Observe the relatively high proportion of baseline trials displaying an increase in hand velocity following force field, and the tendency for peak forward velocities to decay between two force-field trials. The arrows at the end of the series are aligned on the averaged forward velocities of baseline trials across pre-, peri-, and post-exposure phases for this participant. **b**, Trial-by-trial modulation of peak forward velocity in each condition. In the peri-exposure phase, trials were grouped as a function of their index after each force-field trials. Statistical comparisons are based on paired *t* tests between trials in each index and peak hand velocity in the post-condition ($*p < 0.005$, $\dagger p < 0.05$, after Bonferroni correction, $n = 10$). **c**, After-effect quantified as the maximum lateral deviation across the different phases of the experiment. Statistical comparisons were performed as for the peak forward hand velocity by comparing data of each trial index with the data from the post-exposure phase. Similar Bonferroni correction corrections were used as in **b**. **d**, Peak forward velocity as a function of maximum lateral displacement. Data were transformed into z-score for comparison relative to the mean and population of the data from trial Index 12. Dashed trace is the identity line; gray and red illustrate least-square regressions and 95% confidence interval. Only population averages are displayed but the regression and confidence intervals were calculated on individuals' averages. FF, Force Field.

groups of trials sorted by how closely they followed a perturbation trial (Indices 1–12, and post-exposure; paired, one-sided *t* tests). Significant differences at the level $p < 0.005$ after correcting for multiple comparisons are illustrated in Figure 3*b*. Because the data from each trial index was compared with the data from post-exposure, we performed 12 comparisons and consequently the actual *p* levels corresponded to 0.004 and 0.0004 for the reported levels 0.05 and 0.005, respectively. Thus, recent exposure to an unexpected disturbance produced a substantial and sustained increase in the forward hand velocity.

It was important to verify that this increase in forward hand velocity was not merely due to a possible after-effect evoked by the occurrence of force-field trials. To mitigate such an influence, we first excluded two baseline trials immediately following perturbations trials for the analysis presented in Figure 2*d*. Second, we verified that the analyses shown in Figure 3*b* were similar by taking only the forward hand velocity, based on the idea that an after-effect would primarily impact the lateral velocity. We found the same result as reported in Figure 3*b*. Finally, we quantified the after-effect by extracting the maximum lateral deviation, in trials that followed a force-field trial. Consistent with an after-effect, we found an increase in maximum lateral deviation following force-field trials. However, we found the after-effect decayed faster than changes in forward hand velocity because significant increases in lateral deviations were observed only for trial indices

up to 3 (Fig. 3*c*). We directly compared the decay in forward hand velocity with the reduction in maximum lateral displacement by transforming the data into z-score relative to the population from trial Index 12, and regressed the peak forward velocity as a function of the maximum lateral displacement (Fig. 3*d*). We found that the peak forward hand velocity decayed slower. Indeed, the linear regression had a slope that was significantly < 1 [value: 0.66, 95% CI: (0.46, 0.86)], and an intercept that was significantly > 0 [value: 0.5, 95% CI: (0.32, 0.7)]. Thus, as illustrated in Figure 3*d*, the lateral displacement following a force-field trial returned to baseline levels faster than changes in peak forward velocity. Collectively, these analyses suggest that alterations in forward hand velocity were not simply because of after-effects evoked by perturbation trials.

Experiment 2

Experiment 2 was designed to probe the model prediction that the increase in control gain should generate more vigorous responses to abrupt perturbations applied to the limb during movement (step-loads). We designed a similar series of blocks (pre-, peri-, and post-exposure) and varied the level of uncertainty by using different kinds of perturbations. However, the behavior in this experiment was distinct. In contrast to the results of Experiment 1, we did not observe any systematic increase in forward hand velocity when comparing the peri- and pre-exposure phases of the experiment (Fig. 4*b,c*).

Instead, the peak hand velocity decayed across trials during the pre-exposure phase, and then remained relatively constant during the peri-exposure phase (Fig. 4*b*). Congruent with Experiment 1, however, we did observe a significant decrease in forward hand velocity during the post-exposure phase (one-sided paired *t* test, $t_{(12)} = 2.14$, $p = 0.026$). This behavior can be explained in our framework if the presence of step-perturbations already promoted a more robust strategy during the pre-exposure phase, which reduced the contrast between pre- and peri-exposure phases.

Because there was no systematic modulation across phases of the experiment, we leveraged the trial-by-trial modulation of forward hand velocities observed in Experiment 1. We indexed both perturbation and baseline trials dependent on the preceding number of baseline trials. We then analyzed the kinematics and EMG for both unperturbed and perturbed trials based on this index. For the unperturbed trials, we found as in Experiment 1 a clear change in movement velocity toward the target that depended on the preceding number of unperturbed trials. Figure 5*a* reproduces the analysis of Figure 3*b* and reveals similar trial-by-trial modulation of forward velocity following unanticipated perturbation trials encountered by participants in Experiment 2. Because of the larger proportion of perturbation trials in this experiment, the series of unperturbed trials were shorter and the rmANOVA had to be performed on trial Indices 1–7 to ensure a

balanced test because perturbations were more frequent (rmANOVA, $F_{(6,72)} = 13.9$, all corrected $p < 10^{-4}$).

Interestingly, the activity of pectoralis major and posterior deltoid muscles followed a similar pattern, which is surprising because they were chosen for their antagonist actions. We extracted the average activity in a window of 50 ms before the crossing of the position threshold used for the step perturbation trials. We found a similar trial-by-trial decay in the activity of both muscles during movement, in particular in the 50 ms time window before the position threshold used to trigger the step loads (Fig. 5*b*; rmANOVA, $F_{(6,72)} = 6.4$, all corrected p values < 0.005). The concomitant variation of activity in both muscles indicated the presence of co-contraction. To further illustrate this change in activity, we plotted the average activity across trials with Indices 1, 3, or >4 chosen as they are representative of the range of variation of muscle activity (Fig. 5*c*). This figure shows that the overall activity throughout the movement depended on the preceding sequence of trials. The relationship between forward velocity and co-contraction was further assessed on a trial-by-trial basis as we found for both muscles a positive correlation between the mean EMG in the 50 ms window before threshold crossing and the peak velocity (Fig. 5*d,e*). The link between co-contraction and forward hand velocity can be deduced from the fact that both muscles exhibited positive correlation and concomitant trial-by-trial modulation. To complement this analysis, we calculated the linear regressions between forward hand velocities and the minimum activity across muscles, used as an index of co-contraction. We found significant regressions with positive slopes for all participants (means slope: 0.2, range: 0.07, 0.4).

We now turn back to the perturbation trials. We classified them based on a median split on their indices. Trials with an index below the median were classified as Early, and those with an index higher than the median were classified as Late (Fig. 6*a*). Our prediction was clear: if the increase in velocity following disturbances indeed reflected a more robust strategy, then perturbation trials classified as Early (presumably more robust) would also display more vigorous feedback responses when disturbed by a step load.

The distribution of trial indices for all participants is shown in Figure 6*b*. The median number of baseline trials between perturbation trials was equal to 1 for all participants. Thus, the classification separated perturbation trials that directly followed a force field or step load (index = 1) from perturbation trials that followed one or more baseline trials (index ≥ 2). First, the forward hand velocity at perturbation onset tended to be larger for trials classified as Early. This effect was present for 6/13 participants as assessed on the distribution of individual trials (Wilcoxon rank sum test, $p < 0.05$; Fig. 6*c*). Average forward hand velocities at perturbation onset were significantly higher for Early trials when examined at the group level (one-sided comparison: $t_{(12)} = 2.37$, $p = 0.018$). Similar results were observed when this analysis was performed based on the peak forward hand velocity instead of peak hand speed to reduce possible effects of

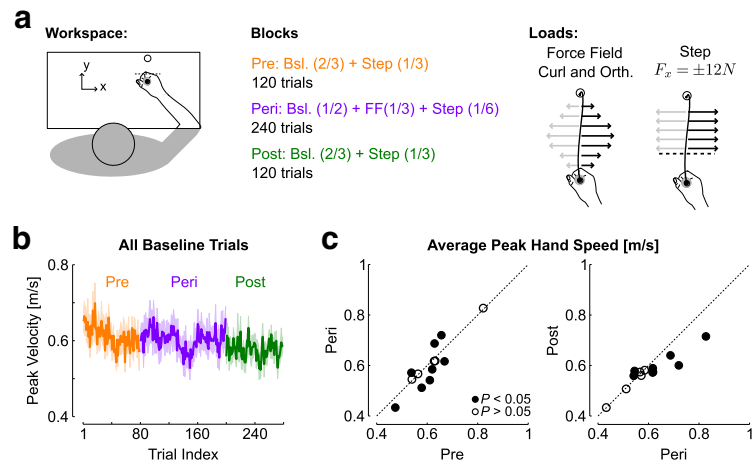


Figure 4. Behavior from Experiment 2. *a*, Illustration of the experimental procedure: the task and workspace were identical to Experiment 1. Pre- and post-exposure phases include step perturbations triggered when the cursor crossed a position threshold corresponding to one-third of the reach path, whereas the peri-condition includes step perturbations, orthogonal and curl force fields, all randomly interleaved. *b*, Trial by trial peak forward velocity during pre-, peri-, and post-exposure phases with the same color code as in Figure 3. *c*, Same as Figure 3*b* for the data from Experiment 2. There was no systematic difference in peak forward hand velocities between pre- and peri-exposure phases, whereas the reduction in hand velocity during post-exposure was weakly significant (paired t test, $p < 0.05$, $n = 13$).

lateral velocities induced by the after-effect: 6/13 participants exhibited increases based on trialwise distributions ($p < 0.05$), and the group effect was also significant (one-sided paired t test: $t_{(12)} = -3.46$, $p < 0.005$).

It is interesting to note that there were no systematic changes in endpoint distributions across Early and Late trials. For baseline trials, there was a small increase in endpoint variance for Early trials (t test on areas of endpoint dispersion ellipses: $t_{(12)} = 2.27$; $p = 0.04$), and no statistical difference in the norm of the endpoint error. For the perturbation trials, we found no effect of the categories on the norm of the endpoint error and on the dispersion ellipses (all $t_{(12)} < 1.4$, $p > 0.1$).

Strikingly the classification of trials separated not only the hand velocity toward the target (Fig. 6*c*), but also the lateral hand velocity and lateral deviation when a step-perturbation was applied to the hand. Indeed, the absolute lateral hand velocity following the perturbation quickly became smaller for trials classified as Early (Fig. 6*d*). We performed a sliding t test on the lateral hand velocity and found a strong reduction of the absolute lateral hand deviation for Early trials consistent with the model prediction (Fig. 1*e*). Vertical arrows correspond to the moment when $p < 0.01$, and p values dropped later to levels $< 10^{-4}$. As a consequence, the lateral displacement exhibited lesser excursion for these trials, and the difference in maximum lateral displacement was highly significant (Fig. 6*e–g*). Figure 6*g* shows the difference between maximum displacements in Late–Early trials, such that the peak lateral displacement in Early trials was always smaller in absolute value than the peak lateral hand displacement in Late trials (** $p < 0.0005$, one-tail paired t test). Thus, this analysis highlighted that trials with increased average forward hand velocity also exhibited larger, more forceful responses to the same lateral step-force disturbances. This modulation was consistent with the hypothesis that these trials were controlled with a more robust strategy as the perturbation loads impacted hand trajectories to a lesser extent.

We now turn back to the surface EMG data from step perturbation trials to address whether the change in behavior shown in Figure 6 resulted from a modulation of neural feedback gains

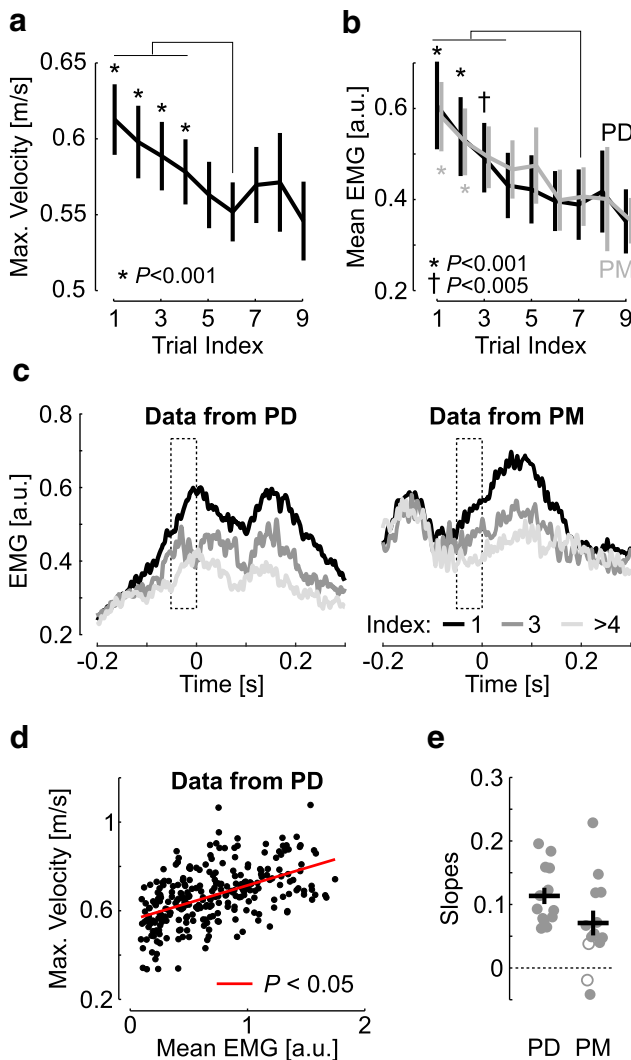


Figure 5. Behavior of Experiment 2: baseline trials. **a**, Peak forward velocity as a function of trial index similar to Figure 4d. Paired comparisons were based on paired *t* test with the data with minimum average velocity (Index 6). **b**, Same as **a** with the activity of each muscle. Paired comparisons were performed with the data from Index 7. **c**, Surface activity of PD (left) and PM (right) during baseline trials averaged across trials and subjects. Trials were separated dependent on the preceding number of baseline trials. The categories chosen to highlight the dependency on the trial sequence were trials with Indices 1 (black), 3 (gray), and >4 (light gray). The traces were aligned to the moment when the hand crossed the position threshold, and smoothed with a 5 ms-sliding window for illustration. The dashed rectangles shows the time window used in the subsequent analyses and corresponds to the Pre epoch (−50 to 0 ms before threshold crossing). **d**, Correlation between mean EMG in the same time window and peak velocity across trials. Data from one exemplar participant. **e**, Slopes of the correlations as computed in **d** for each subject. One dot is one participant, filled dots illustrate significant regressions, open dots illustrate nonsignificant regressions ($p > 0.05$). Observe that all but two fits were significant, and all but two display positive correlations of activity from antagonist muscles with forward hand velocity. Horizontal bars are population averages and vertical bars are 1 SEM ($n = 13$).

predicted in theory. Raw EMG activities for each muscle are shown in Figure 7 for stretch and shortening responses (after normalization; see Materials and Methods). As for unperturbed trials analyzed in Figure 5, we observed an increase in co-contraction for trials classified as Early compared with trials classified as Late in the Pre epoch (−50 to 0 ms re: perturbation onset), which corresponded to $25 \pm 4\%$ increase in pectoralis major (PM) and a $30 \pm 4\%$ increase in posterior deltoid activity

(PD; mean \pm SEM). We focus on the interaction between the classification criterion (Early or Late) and the different epochs (Nieuwenhuis et al., 2011). The rmANOVA performed on the stretch responses with classification labels (Early and Late) and epochs as factors revealed significant interactions between these factors for both muscles (PD: $F_{(4,48)} = 11.9, p < 0.001$; PM: $F_{(4,48)} = 3.8, p = 0.009$). Regarding the shortening response, we also observed a significant interaction between classes and epochs for both muscles (PM: $F_{(4,48)} = 5.4, p = 0.001$; PD: $F_{(4,48)} = 3.22, p = 0.02$).

Considering the dependency of the stretch reflex on the baseline activity for loading and unloading conditions (Pruszynski et al., 2009; Nashed et al., 2015), we compared the difference between the stretch responses and the baseline activity in the short- and long-latency epochs across classes. That is, we subtracted the raw activity collected during unperturbed trials classified as Early (high baseline) and Late (low baseline) trials from the perturbation responses corresponding to the same Early and Late classes. This difference reflected the response while taking changes in baseline activity into account. Stretch responses from PM and PD were averaged for this analysis and direct comparisons were performed based on paired *t* test. We found a significant increase the stretch response (one-sided comparison: $t_{(12)} = 2.34, p = 0.018$ for long-latency, $t_{(25)} = 2.74, p = 0.005$ for long-latency and short-latency pooled together, although short-latency alone was not significant: $t_{(12)} = 1.63, p = 0.063$). It is clear that the scaling of the muscle responses in these early epochs is small and only visible in Figure 7, *c* and *f*. The fact that this scaling was still observed based on rather low levels of coactivation indicated that it was a reliable effect. The phasic increase and decrease in the stretch response, along with the larger shortening response was consistent with the change in control gains predicted in the model (Fig. 1*f*).

Comparisons across Experiments 1 and 2

Collectively our experimental results demonstrated a change in control strategy impacting forward velocity and feedback responses to perturbations based on co-contraction and amplification of control gains. Thus, if two force-field trials occur in a row, we may expect that the lateral displacement in the second force field be reduced compared with the first one because of the use of a more robust strategy. Indeed, in Experiments 1 and 2 we found that a sequence of perturbations (in any direction) evoked a buildup in forward hand velocity, suggesting a gradual increase in robustness (Fig. 8*a–c*). We performed a similar analysis on the EMG data from Experiment 2 and also found a concomitant buildup of activity in both pectoralis and deltoid muscles, as expected (Pre epoch 50 ms before threshold, Indices 0–3: rmANOVA, $F_{(3,48)} > 17, p < 10^{-6}$). In parallel, when perturbations were separated based on direction, we also found in each experiment the presence of an after-effect, which increased following one or two perturbations experienced in a row (Fig. 8*d–f*). Thus, unexpected disturbance evoked both a gradual increase in control gains, as well as an increase in standard after-effect in the next baseline trial.

These observations are particularly important because they suggest that robust control and error-based adaptation occurred in parallel. Because these processes varied from trial to trial, their respective contributions to behavior are difficult to dissociate. We end this section by showing that robust control is in fact necessary to explain an apparent contradiction in the data. Recall that we used the same curl field randomly interspersed in each experiment. Thus, it was possible to compare participants’ be-

havior when facing the same disturbance across the two experiments, as well as the after-effect. First we observed that the hand paths from Experiment 2 were less deviated than in Experiment 1 (Fig. 9*a*). This result could be quantified by comparing the maximum absolute hand displacement from participants' average traces. We grouped data from CW and CCW perturbations to increase the power (because the comparisons could not be paired due to the fact our experiments involved distinct groups of participants), and found a significant trend toward smaller hand deviation for Experiment 2 (Fig. 9*b*; Wilcoxon rank sum test: $Z = 2.07$, $p = 0.038$). Taking lateral deviation as an index of error-based adaptation to the force field, one would conclude from this result alone that participants from Experiment 2 were more adapted. The contradiction is that the after-effects displayed from participants of Experiment 2 were significantly smaller than the after-effects displayed by participants of Experiment 1 (Fig. 9*a,c*; $Z = 3.11$, $p = 0.0019$).

How could participants from Experiment 2 appear more adapted but at the same time express smaller after-effects? Our framework may resolve this contradiction. Indeed, we believe that the increased frequency and variation of dynamical disturbances caused by curl and orthogonal force fields in Experiment 2 evoked an overall more robust strategy than in Experiment 1 where perturbation were less frequent. To verify this hypothesis, we observed that curl field trials in Experiment 2 occurred on average after <2 baseline trials, and this number never exceeded 10 (Fig. 9*d*). In contrast in Experiment 1, curl field trials occurred on average after ~ 4 baseline trials, and the distribution of indices was much broader (note that this could not be directly assessed based on velocity during curl field trials, because the increase in feedback gains and the impact of the force field have opposite effects). A statistical comparison on the distribution means also revealed highly significant differences across experiments (KS stat = 1, $p < 10^{-5}$). In light of these results, our explanation is that the increased frequency of perturbations in Experiment 2 produced a more robust strategy in these participants, which limited the lateral hand displacement without necessarily evoking internal model adaptation.

Discussion

We explored the hypothesis that compensation for unmodeled disturbances was supported by a robust neural control strategy. We studied the predictions of stochastic optimal control (LQG; Todorov, 2005) and a robust control design that can equivalently be described as a “min-max” or worst-case strategy (Basar and Bernhard, 1991) applied to linear models of planar reaching movements. The robust controller displayed an increase in control gains, resulting in faster movements toward the target and more vigorous responses to perturbations. Our experimental results supported these predictions: the occurrence of unexpected force-field disturbances evoked both faster movements and more vigorous

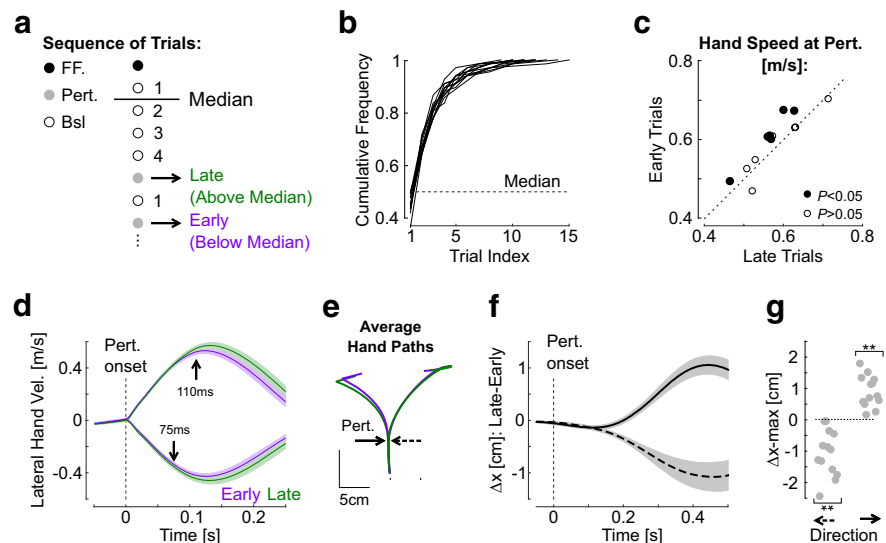


Figure 6. Behavior of Experiment 2: step perturbations. *a*, Illustration of the classification of perturbation trials dependent on the number of preceding baseline trials. The classification was a median split based on this index (below median: trials occurring Early; above median, trials occurring Late). *b*, Distribution of baseline trial indices. These indices for each step-perturbation trial correspond to the number of unperturbed trials that preceded them. *c*, Average forward hand velocity for perturbation trials measured at perturbation onset. The data are the average across trials classified as Early (preceded by a perturbation) as a function of the same variable for perturbation trials classified as Late (preceded by ≥ 1 baseline trials). *d*, Lateral hand velocity for the step perturbation trials (left or right) dependent on the classification shown in *a*. Vertical arrows depict the movement when the sliding t test on participant averages across trial types became dropped below $p < 0.01$. *e*, Average hand path for each category of trial and perturbation responses. Trials were aligned on perturbation onset. *f*, Grand average of differences in the x -coordinate. Displays are Late minus Early, indicating that trials classified as Early exhibited smaller lateral displacement in absolute value. *g*, Changes in average maximum lateral displacement from each individual. As in *g*, positive and negative data for positive and negative perturbations indicate that trials classified Early exhibited smaller lateral displacement. Asterisks indicate significant differences ($p < 0.005$, one-tail, paired comparisons, $n = 13$).

responses to perturbations. Thus, the neural controller was more robust in the sense that the feedback responses reduced the impact of the perturbations (step and force field). Thus the compensation for disturbances involved a model-free component.

Our results suggest a robust control strategy, which can be dissociated from automatic stiffening of the limb through impedance control (Hogan, 1984; Shadmehr and Mussa-Ivaldi, 1994; Burdet et al., 2001). Indeed, the increase in co-contraction could have only moderately altered the intrinsic properties of muscles ($\leq 30\%$ increase in baseline activity; Crevecoeur and Scott, 2014). Even when considering a change in mechanical impedance of the limb, the modulation of forward hand velocity as well as the phasic modulation of agonist and antagonist responses together supported a modulation of control gains.

We should underline that there is no disagreement between robust control and stochastic optimal control (Todorov and Jordan, 2002). In theory, LQG yields efficient control in the presence of Gaussian disturbances based on goal-directed, task-dependent state-feedback control law (Todorov, 2004; Franklin and Wolpert, 2011; Scott, 2012; Crevecoeur and Kurtzer, 2018). The control solution produced by robust control also consists of a goal-directed, state-feedback control law that can change flexibly should there be a change in the task or cost-function. Thus, previous evidence for flexible state-feedback control in humans supports robust control as well as LQG as models of sensorimotor coordination.

Clearly it is possible to modify the cost-function to fit the modulation of control gains in the context of LQG. However, we do not feel that such a model would provide much conceptual

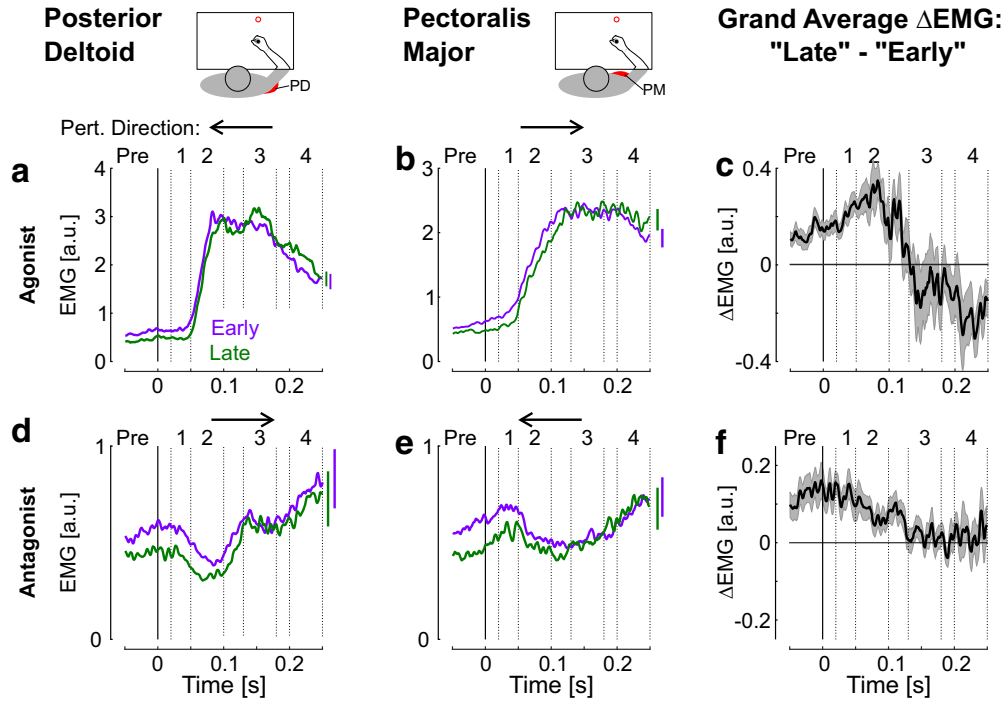


Figure 7. Muscle responses to perturbations. **a**, Response to the perturbation in PD. The solid vertical lines in the middle and bottom panels correspond to perturbation onset, and the dashed lines correspond to different epochs used in the binned analysis [(1) short-latency (20, 50 ms); (2) long-latency (50, 100 ms); (3) early voluntary (130, 180 ms); and (4) voluntary (200, 250 ms)]. Vertical lines at the right side of the plots show the SE across participants at 250 ms after perturbation onset. Arrows above each panel indicate the perturbation direction. **b**, Same as **a** for the PM (shoulder flexor). **c**, Grand average of the difference between muscle activity across Early and Late trials from **a** and **b**. The difference was averaged across muscles, and then across participants. **d–f**, Same as **a–c** for the antagonist response. All traces were smoothed with a 5 ms, centered moving average. Shaded areas represent 1 SEM across participants ($n = 13$). Muscle activities in **a**, **b**, **d**, and **e** correspond to the raw EMG data after normalization against the background load (see Materials and Methods). The traces shown in **d** and **f** are the difference across classes. Observe the pre-perturbation differences linked to the strategy and co-contraction (Fig. 5).

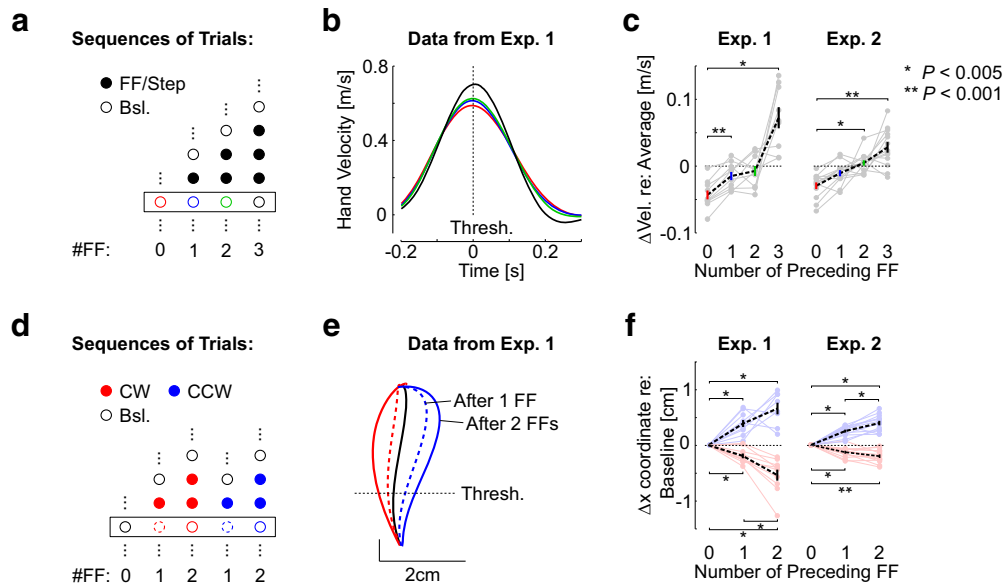


Figure 8. Change in strategy and learning. **a**, Baseline trials were assigned an index corresponding to the number of perturbation trials of any kind that preceded them. **b**, Average forward hand velocity across trials of the same index and across participants. Traces were aligned to the position threshold as above, which was on average close to peak velocity. Data from Experiment 1. **c**, Individual and average peak hand velocity across indices from Experiments 1 and 2. The data illustrates that the increase in control gains documented in each experiment builds up during a series of any kind of perturbation. Asterisks indicate significant comparisons at the indicated level after Bonferroni corrections for multiple comparisons ($n = 10$ for Experiment 1, $n = 13$ for Experiment 2). Vertical bars are the SE. **d**, Illustration of the labeling procedure. The trials were indexed based on the preceding number of perturbations, and on the direction of the perturbation. Step perturbations were included as CW or CCW disturbances in this analysis as they produced an after-effect similar to the force field (see Results). **e**, Illustration of hand paths following one (dashed) or two (solid) force-field trials, in CW (red) or CCW (blue) directions. **f**, Lateral coordinate at the position threshold for each index and direction. Individuals' grand average across all baseline trials was subtracted for illustration. Statistical comparisons followed the same procedure as for **c**. FF, Force Field.

advance, as there was no change in spatial, temporal, or accuracy requirements justifying a change in cost-function in the model. In contrast, considering robust control highlighted that changes in behavior could be explained by considering unmodeled dynamics, which was directly motivated by our experiments and did not require arbitrary fitting. Other behavioral signatures may provide further insight into the control strategy dependent on the context. Indeed, previous work highlighted incomplete correction for target jumps predicted in the framework of stochastic optimal control (Liu and Todorov, 2007). In our study, the simulations of LQG displayed incomplete corrections (Fig. 1*f*), whereas the robust controller was steering the system to the target regardless. This paralleled participants' behavior, as we did not see any systematic effect on endpoint distributions across conditions. We believe the details of error corrections dependent on the context is an interesting topic for future work.

An important finding was the trial-by-trial adjustments of control gains. An unexpected perturbation evoked an increase in control gains for the next trials, whereas a sequence of unperturbed trials tended to push these control gains back toward values corresponding to the predictable context (Figs. 3*b*, 5*b*). The theory provides a framework for interpreting these results. Solutions of LQG and robust controllers impose a tradeoff between efficiency and robustness of movement control (Fig. 1). The most efficient controller (LQG) is achieved at the cost of sensitivity to model errors, whereas the most robust controller generates costly control solutions, but is less sensitive to model errors. This result is well known in theory (Levine, 1996; Boulet and Duan, 2007). Thus, the adjustments of control strategies that we observed experimentally could be understood as the behavioral expression of competing mechanisms that promote either robustness through an increase in control gains or efficiency by decreasing the gains after a series of predictable trials. Reproducing the trial-by-trial change in the model could be achieved by altering the parameter γ , which determines the optimal level of disturbance attenuation, whereas suboptimal values of this parameter correspond to solutions that may resemble LQG. A simple algorithm could be to set γ to the optimal value following a disturbance (robust control), and relax this constraint in predictable contexts or after adaptation.

This trial-by-trial modulation correlated with co-contraction, which potentially presents two advantages. On the one hand, the elevated levels of muscle activity enable both stretch and shortening responses in antagonist muscles at short latencies, which may increase the range of the feedback response. On the other hand, an increase in baseline muscle activity, evoked by a background load or by co-contraction, is associated with "gain-scaling" (Stein et al., 1995; Pruszynski et al., 2009; Crevecoeur and Scott, 2014; Nashed et al., 2015), which may increase feedback gains by α -gamma coactivation (Vallbo, 1974). Indeed a modulation of spindles sensitivity is likely involved, because studies using galvanic stimulations (H-reflex) did not always observe gain-scaling with co-contraction (Nielsen and Kagamihara, 1993; Carroll et al., 2005). The tradeoff between efficiency and




Figure 9. Apparent adaptation. *a*, Left, Average hand displacements during curl field trials from Experiments 1 and 2 depicted with thin and thick traces, respectively. Blue and red traces correspond to CCW and CW perturbations. The maximum absolute lateral displacement was extracted from each participant's individual average and data from the two directions were pooled to increase power. Right, Hand traces in baseline trials following force-field trials to highlight the after-effect. This after-effect was extracted at the position threshold similar to Figure 8. The lateral component of the average baseline trajectory was subtracted for each participant. Ellipses represent 2D dispersion for the mean (i.e., the axes of the ellipses were scaled by the square root of sample size). *b*, Maximum absolute hand displacement across experiments. The absolute hand displacement was larger in Experiment 1 ($n = 10 \times 2$ directions) than in Experiment 2 ($n = 13 \times 2$ directions) at the level $p < 0.05$ (see also Results). Dots represent average traces for each participant for CCW (dark blue) and CW (dark red) trials. *c*, Same as *b* for the after-effect. The two distributions were statistically different at the level $p < 0.005$. *d*, Cumulative distribution of the number of baseline trials preceding each force-field trial per participant and in each experiment. Vertical lines represent participant averages. Statistical difference at the level $p < 0.005$ are highlighted.

robustness is also apparent in this scenario, as it is clear that there is a metabolic cost incurred with the increased activity in neural circuits and in muscles used to maintain higher levels of activity. Our results demonstrated that co-contraction was inhibited gradually but systematically in the absence of disturbances, likely to save energy.

Feedback modulation was observed during motor adaptation (Wagner and Smith, 2008; Franklin et al., 2012; Cluff and Scott, 2013). This modulation could also involve a change in robustness, but most importantly reflected knowledge of the environment following adaptation. Here we observed a specific role of co-contraction. Thus, this strategy differed from that observed after learning, but it may be a component of the early stages of adaptation as highlighted in previous studies (Milner and Franklin, 2005; Franklin et al., 2008). In the latter reference, Franklin and collaborators showed that unexpected muscle stretches during reaching evoked co-contraction, whereas the absence of unexpected disturbances was followed by a decrease in coactivation ("V" learning scheme). This update rule was shown to converge to novel patterns of activity suitable for adapting motor commands to altered dynamics. Here we reproduced similar adjustments across trials: unpredictable perturbations evoked co-contraction, and a sequence of similar (unperturbed) trials decreased it. Our contribution was to show that co-contraction may not only modulate the limb's intrinsic impedance (Hogan, 1984; Burdet et al., 2001), but also adjust the robustness of neural control through a change in feedback gains.

Behaviorally, robust control is consistent with previous reports exploring human motor control. For instance, Wei et al., 2010 studied adaptation to perturbations of different kinds and concluded that humans used a nonspecific strategy. This result was consistent with a model-free compensation as the perturbations were unpredictable, and such a default strategy must also be suitable for the worst-case disturbance. Robust control may also explain strategies observed by Hadjiosif and Smith (2015) in the context of object manipulation. In this experiment, participants encountered force fields where the coefficient mapping forward hand velocities onto lateral forces was normally distributed with differing levels of variance across conditions. Hadjiosif and Smith (2015) reported changes in grip control that scaled with the variance of encountered loads during movement. In general, a sto-

chastic optimal controller would not scale with the variance of the noise, as it optimizes the expected cost. In contrast, a robust controller would produce sufficient grip force in the worst-case scenario, and thus it would be sensitive to, and scale with, the variance of disturbances.

Our findings have important impact on our interpretations of the neural basis of motor learning. Indeed, although a reduction of perturbation-related motor errors is commonly viewed as evidence for adaptation, our experiments revealed that such changes also involved a robust (model-free) control strategy. The use of force channels to measure adaptation may partially overcome this difficulty, but because of spontaneous movement curvature, the robust strategy may also impact the force produced against the channel walls. In addition, we uncovered that the selection of a robust strategy potentially impacted learning (Fig. 9). We previously reported a tradeoff characterizing individual differences in motor learning (Cluff et al., 2019): some individuals exploited robust strategies to compensate for partial learning of a force field, whereas others who expressed better learning and reliance on their internal models also displayed less robust strategies when exposed to unpredictable loads. Whether the relationship between robustness and learning shown in Figure 9 was because of differences across the two groups or to interaction between these processes remains unknown but constitutes an important question for future work. Furthermore, it was recently reported that online control during unexpected force fields or visuomotor perturbations involved very rapid learning (Braun et al., 2009; Crevecoeur et al., 2018). Thus, a clear challenge is to disentangle robust control from rapid adaptation, because the function of both processes is to counter perturbation-related motor errors over fast timescales.

References

- Basar T, Bernhard P (1991) *Hoo-optimal control and related minimax design problems: a dynamic game approach*. Boston; Basel; Berlin: Birkhäuser.
- Benjamin DJ, Berger JO, Johannesson M, Nosek BA, Wagenmakers EJ, Berk R, Bollen KA, Brembs B, Brown L, Camerer C, Cesarini D, Chambers CD, Clyde M, Cook TD, De Boeck P, Dienes Z, Dreber A, Easwaran K, Efferson C, Fehr E, et al. (2018) Redefine statistical significance. *Nat Hum Behav* 2:6–10.
- Bitmead RR, Gevers M, Wertz V (1990) *Adaptive optimal control: the thinking man's GPC*. New York: Prentice Hall.
- Boulet B, Duan Y (2007) The fundamental tradeoff between performance and robustness: a new perspective on loop shaping. *IEEE Control Syst Mag* 27:30–44.
- Braun DA, Aertsen A, Wolpert DM, Mehring C (2009) Learning optimal adaptation strategies in unpredictable motor tasks. *J Neurosci* 29:6472–6478.
- Brown IE, Loeb GE (2000) Measured and modeled properties of mammalian skeletal muscle: IV. Dynamics of activation and deactivation. *J Muscle Res Cell Motil* 21:33–47.
- Burdet E, Osu R, Franklin DW, Yoshioka T, Milner TE, Kawato M (2000) A method for measuring endpoint stiffness during multi-joint arm movements. *J Biomech* 33:1705–1709.
- Burdet E, Osu R, Franklin DW, Milner TE, Kawato M (2001) The central nervous system stabilizes unstable dynamics by learning optimal impedance. *Nature* 414:446–449.
- Carroll TJ, Baldwin ER, Collins DF (2005) Task dependent gain regulation of spinal circuits projecting to the human flexor carpi radialis. *Exp Brain Res* 161:299–306.
- Cluff T, Scott SH (2013) Rapid feedback responses correlate with reach adaptation and properties of novel upper limb loads. *J Neurosci* 33:15903–15914.
- Cluff T, Crevecoeur F, Scott SH (2019) Tradeoffs in optimal control capture patterns of human sensorimotor control and adaptation. *bioRxiv*. Advance online publication. Retrieved August 09, 2019. doi: 10.1101/730713.
- Crevecoeur F, Kurtzer I (2018) Long-latency reflexes for inter-effector coordination reflect a continuous state-feedback controller. *J Neurophysiol* 120:2466–2483.
- Crevecoeur F, Scott SH (2013) Priors engaged in long-latency responses to mechanical perturbations suggest a rapid update in state estimation. *PLoS Comput Biol* 9:e1003177.
- Crevecoeur F, Scott SH (2014) Beyond muscles stiffness: importance of state estimation to account for very fast motor corrections. *PLoS Comput Biol* 10:e1003869.
- Crevecoeur F, Thonnard JL, Lefevre P (2018) A sub-movement timescale of human motor adaptation. *bioRxiv*. Advance online publication. Retrieved February 22, 2018. doi: 10.1101/269134.
- Franklin DW, Wolpert DM (2011) Computational mechanisms of sensorimotor control. *Neuron* 72:425–442.
- Franklin DW, Burdet E, Tee KP, Osu R, Chew CM, Milner TE, Kawato M (2008) CNS learns stable, accurate, and efficient movements using a simple algorithm. *J Neurosci* 28:11165–11173.
- Franklin S, Wolpert DM, Franklin DW (2012) Visuomotor feedback gains upregulate during the learning of novel dynamics. *J Neurophysiol* 108:467–478.
- Gonzalez Castro LN, Hadjiosif AM, Hemphill MA, Smith MA (2014) Environmental consistency determines the rate of motor adaptation. *Curr Biol* 24:1050–1061.
- Hadjiosif AM, Smith MA (2015) Flexible control of safety margins for action based on environmental variability. *J Neurosci* 35:9106–9121.
- Hogan N (1984) Adaptive-control of mechanical impedance by coactivation of antagonist muscles. *IEEE Trans Automat Control* 29:681–690.
- Huang VS, Haith A, Mazzoni P, Krakauer JW (2011) Rethinking motor learning and savings in adaptation paradigms: model-free memory for successful actions combines with internal models. *Neuron* 70:787–801.
- Hwang EJ, Donchin O, Smith MA, Shadmehr R (2003) A gain-field encoding of limb position and velocity in the internal model of arm dynamics. *PLoS Biol* 1:E25.
- Izawa J, Rane T, Donchin O, Shadmehr R (2008) Motor adaptation as a process of reoptimization. *J Neurosci* 28:2883–2891.
- Kording KP, Tenenbaum JB, Shadmehr R (2007) The dynamics of memory as a consequence of optimal adaptation to a changing body. *Nat Neurosci* 10:779–786.
- Levine WS (1996) *The control handbook*. Boca Raton, FL; New York: CRC; IEEE.
- Li W, Todorov E (2007) Iterative linearization methods for approximately optimal control and estimation of non-linear stochastic system. *Int J Control* 80:1439–1453.
- Liu D, Todorov E (2007) Evidence for the flexible sensorimotor strategies predicted by optimal feedback control. *J Neurosci* 27:9354–9368.
- Michiels W, Niculescu S-L (2007) *Stability and stabilization of time-delay systems*. Philadelphia: Society for Industrial and Applied Mathematics.
- Milner TE, Franklin DW (2005) Impedance control and internal model use during the initial stage of adaptation to novel dynamics in humans. *J Physiol* 567:651–664.
- Nashed JY, Kurtzer IL, Scott SH (2015) Context-dependent inhibition of unloaded muscles during the long-latency epoch. *J Neurophysiol* 113:192–202.
- Nielsen J, Kagamihara Y (1993) The regulation of presynaptic inhibition during co-contraction of antagonistic muscles in man. *J Physiol* 464:575–593.
- Nieuwenhuis S, Forstmann BU, Wagenmakers EJ (2011) Erroneous analyses of interactions in neuroscience: a problem of significance. *Nat Neurosci* 14:1105–1107.
- Pruszynski JA, Kurtzer I, Lillicrap TP, Scott SH (2009) Temporal evolution of “automatic gain-scaling”. *J Neurophysiol* 102:992–1003.
- Scott SH (2012) The computational and neural basis of voluntary motor control and planning. *Trends Cogn Sci* 16:541–549.
- Scott SH (2016) A functional taxonomy of bottom-up sensory feedback processing for motor actions. *Trends Neurosci* 39:512–526.
- Shadmehr R, Mussa-Ivaldi FA (1994) Adaptive representation of dynamics during learning of a motor task. *J Neurosci* 14:3208–3224.
- Shadmehr R, Smith MA, Krakauer JW (2010) Error correction, sensory prediction, and adaptation in motor control. *Annu Rev Neurosci* 33:89–108.
- Singh K, Scott SH (2003) A motor learning strategy reflects neural circuitry for limb control. *Nat Neurosci* 6:399–403.
- Smith MA, Ghazizadeh A, Shadmehr R (2006) Interacting adaptive pro-

- cesses with different timescales underlie short-term motor learning. *Plos Biol* 4:e179.
- Stein RB, Hunter IW, Lafontaine SR, Jones LA (1995) Analysis of short-latency reflexes in human elbow flexor muscles. *J Neurophysiol* 73:1900–1911.
- Thoroughman KA, Shadmehr R (2000) Learning of action through adaptive combination of motor primitives. *Nature* 407:742–747.
- Todorov E (2004) Optimality principles in sensorimotor control. *Nat Neurosci* 7:907–915.
- Todorov E (2005) Stochastic optimal control and estimation methods adapted to the noise characteristics of the sensorimotor system. *Neural Comput* 17:1084–1108.
- Todorov E, Jordan MI (2002) Optimal feedback control as a theory of motor coordination. *Nat Neurosci* 5:1226–1235.
- Ueyama Y (2014) Mini-max feedback control as a computational theory of sensorimotor control in the presence of structural uncertainty. *Front Comput Neurosci* 8:119.
- Vallbo AB (1974) Human muscle spindle discharge during isometric voluntary contractions: amplitude relations between spindle frequency and torque. *Acta Physiol Scand* 90:319–336.
- Wagner MJ, Smith MA (2008) Shared internal models for feedforward and feedback control. *J Neurosci* 28:10663–10673.
- Wei K, Wert D, Körding KP (2010) The nervous system uses nonspecific motor learning in response to random perturbations of varying nature. *J Neurophysiol* 104:3053–3063.
- Wolpert DM, Diedrichsen J, Flanagan JR (2011) Principles of sensorimotor learning. *Nat Rev Neurosci* 12:739–751.