

Estimating the sources of motor errors for adaptation and generalization

Max Berniker^{1,2} & Konrad Kording^{1,2}

Motor adaptation is usually defined as the process by which our nervous system produces accurate movements while the properties of our bodies and our environment continuously change. Many experimental and theoretical studies have characterized this process by assuming that the nervous system uses internal models to compensate for motor errors. Here we extend these approaches and construct a probabilistic model that not only compensates for motor errors but estimates the sources of these errors. These estimates dictate how the nervous system should generalize. For example, estimated changes of limb properties will affect movements across the workspace but not movements with the other limb. We provide evidence that many movement-generalization phenomena emerge from a strategy by which the nervous system estimates the sources of our motor errors.

Suppose I throw a stone with my right hand and it travels less than I had expected. This observation can have at least two interpretations: the stone may be heavier than I thought, or my arm may be weaker than I thought. If I believe the stone to be heavier, I will adjust accordingly the next time I throw the same stone, with my right or my left arm. Movements I make without the stone, clearly, shouldn't be influenced by my updated estimate of the stone's weight. If, on the other hand, I believe my right arm to be weaker, I will adjust future movements with this arm, with or without the stone, but not subsequent movements of my (unaffected) left arm. This example highlights what our intuition suggests and dynamics demands: the nervous system needs to estimate not only the necessary corrections (for example, throw with more effort) but also the sources of errors (for example, the stone is heavy or my arm is weak) when we adapt our movements.

Ideally, the nervous system should treat the body and the world as a coupled dynamical system whose properties constantly change over time. For instance, the strength of each of our muscles may vary over time (through fatigue, exercise, injury, aging and so forth). Similarly, the objects we interact with have variable properties and dynamics. To adapt and generalize, the nervous system needs to attribute observed movement errors to either the body or the world and then estimate the value of those properties.

Many studies have analyzed how human subjects adapt to perturbations and have found that they estimate these disturbances instead of simply memorizing a motor plan^{1,2}. This idea is captured by the notion of an internal model, a representation of how our motor system predicts the outcome of motor commands. Studies that examine internal models assume that motor errors, regardless of their source, are estimated with a general model of disturbances. Although some studies do recognize that multiple sources may be responsible for

motor errors^{3,4}, their internal models do not distinguish between them. This single model compensates for any changes in the properties of the world or the body. As a result, these models have no mechanism for representing that an estimated change in the limb should generalize across the workspace but not to the other limb, and that an estimated change in the world should generalize to the other hand.

Here we present a model that formalizes the estimation of the sources of motor errors by adapting its representation of both the body and the world. This source-estimation model constantly updates its parameter estimates and its uncertainty about those estimates through an application of Bayesian inference. To test the model, we compare its predictions to three classes of published motor adaptation studies: velocity-dependent force fields, rotating (Coriolis) rooms and inertial perturbations. Our model predicts many findings of these experiments. These results support our hypothesis that the nervous system adapts and generalizes by estimating the sources of motor errors.

RESULTS

Modeling sources of errors: body and world

As in the example of throwing a stone, a continuum of possible parameter values can explain many motor errors equally. In addition, owing to noisy perception and commands, the motor system operates under uncertain conditions. The nervous system therefore needs to assign credit to the likely parameters and estimate their properties. Bayesian inference is a systematic method for optimally solving such a credit-assignment problem. For example, if we knew the stone's weight with great certainty and errors in our estimate of its inertia could not account for the motor errors we observed, then other sources, such as limb properties, would be considered likely culprits. This would result in relatively small changes in our estimate of the stone's weight and large changes in our estimates of limb properties (for example, mass,

¹Rehabilitation Institute of Chicago, Department of Physical Medicine and Rehabilitation, Northwestern University, 345 E. Superior Street, ONT-931, Chicago, Illinois 60611, USA. ²These authors contributed equally to this work. Correspondence should be addressed to M.B. (mbernike@northwestern.edu).

Received 19 August; accepted 21 October; published online 16 November 2008; doi:10.1038/nn.2229

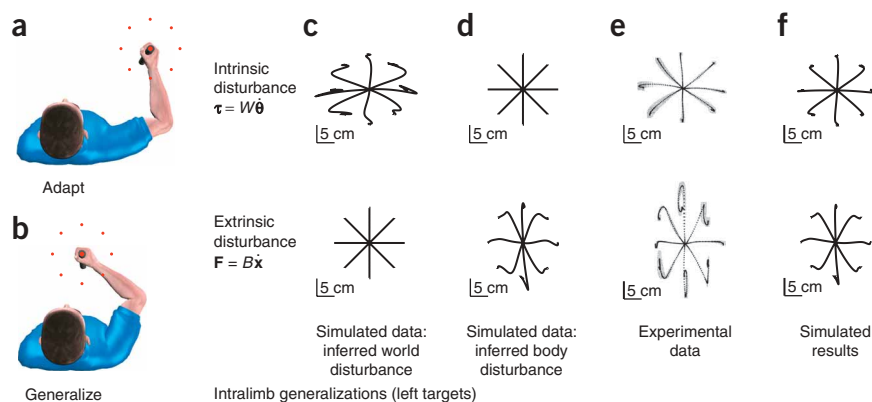


Figure 1 Simulated intralimb generalization. **(a)** The subject adapts to a velocity-dependent force field in the right region of the workspace. **(b)** The subject then makes reaches in the unexperienced left region of the workspace to test for generalization. **(c)** A disturbance is estimated as an entirely extrinsic velocity-dependent force field and the arm is exposed to an intrinsic field, where the torques, τ , are defined in terms of joint velocity, $\dot{\theta}$ (top); and an extrinsic field, where the forces, F , are defined in terms of the hand's velocity, \dot{x} (bottom). **(d)** A disturbance is adapted to as misestimates in limb properties and the arm is exposed to an intrinsic, velocity-dependent force field (top) and an extrinsic field (bottom). **(e)** Experimental results for intralimb generalization (reproduced from ref. 5) when subjects are exposed to an intrinsic and an extrinsic force field; top and bottom, respectively. **(f)** Simulated intralimb generalizations showing a mixture of inferred body and world estimates.

strength and so forth). Bayesian inference formalizes this update process by combining our belief in parameter estimates with a likelihood that these parameters can account for our observations. This results in updated parameter estimates that are maximally certain.

Our approach extends the traditional internal model to allow for the estimation of multiple sources of error. The source estimation model adapts by estimating the properties of the body and the world and generalizes by using these estimates for new movements. When adapting to perturbed movements, the source estimation model may produce the same compensatory corrections that a traditional internal model would. However, the two classes of model make different predictions for generalization to new movements. For example, when we adapt while throwing a stone and then throw it in a new direction, or throw it with our left arm, or throw a different stone, the source estimation model has the necessary information to generalize. Estimates about the world can transfer across limbs; they represent knowledge that is independent of the motor apparatus. Body estimates help us to generalize to new motor goals, but do not necessarily transfer from one limb to another; they represent knowledge of a specific motor apparatus.

Here we examine a nonlinear musculoskeletal model of the human limb (body) interacting with various imposed (world) disturbances. The source estimation model uses this simulated body–world system to predict adaptation and generalization. To estimate the body (limb) and world (experimental perturbations) parameters, we used an extended Kalman filter (see Methods). The source estimation model infers changes in the nominal values of a large number of parameters. We assume that the variations driving these changes are proportional to their nominal values to restrict the model to two free parameters, a proportionality constant for body parameters and another for world parameters. We further restrict the model to a single free parameter, α , the ratio of these values. For all the results shown below, α was set to 0.4; that is, the model assumes the body is more variable than the world. Thus, if all parameters are equally likely to account for an observed motor error, body parameters will receive more credit. We have thus obtained a model with a single free parameter that can predict motor

adaptation and generalization. Below we present the model's predictions for simulated experimental studies of human reaching movements (see Methods for details).

Velocity-dependent force field studies

We first compare the source estimation model with the results from a study on motor adaptation and generalization⁵ that asked questions highly relevant to the issues we consider here. In that study, researchers examined whether the nervous system represents disturbances in terms of a body-based (intrinsic) coordinate frame or a world-based (extrinsic) coordinate frame. A robot rendered velocity-dependent forces that perturbed the hand of subjects while they made reaching movements. After the subjects adapted, they were tested for generalization in a different part of the workspace (**Fig. 1a,b**). When the disturbing forces were defined in terms of the hand's velocity (an extrinsic coordinate frame), subjects generalized poorly. However, when the forces were defined in terms of the angular velocities of

the subject's limb (an intrinsic coordinate frame), subjects generalized well. Thus, the study concluded that the nervous system represents disturbances in an intrinsic (joint-based) coordinate system.

That study (and many subsequent ones) examined its hypothesis in terms of two coordinate frames, assumed to be mutually exclusive. Here we argue that the nervous system should not simply use one representation. Such an assumption is equivalent to the belief that the properties of either the body or the world are constant and known without error. Yet all of these properties are variable and need to be estimated.

We used the source estimation model to simulate the same adaptation and generalization experiment with equivalent force fields (see Methods). Once the model has adapted to the external field, it has updated estimates of all its parameters, both world and body properties. To what extent a parameter's estimate is updated depends on its ability to explain the motor errors and on the model parameter α . These inferred estimates are used to compensate for the actual external field. For any setting of this variable α , we will obtain a different set of predictions for movement generalization.

For illustrative purposes, we present how our model adapts with two extreme values of the parameter α . When the system attributes errors only to the world and the field is hand-based, reaches made in the left region of the workspace perfectly cancel the force field (**Fig. 1c**, bottom). However, if the field is joint-based, there are clear and distinct patterns of generalization that reveal the internal model's incorrect estimation of the velocity-dependent force field (**Fig. 1c**, top). This type of inference of sources of error occurs when the model is relatively certain of the body parameters ($\alpha \gg 1$). Similarly, if the system attributes errors only to the body ($\alpha \ll 1$) it will generalize well with joint-based fields and poorly with hand-based fields (**Fig. 1d**). These extreme settings of α correspond to the situation present in traditional internal model studies where only a single disturbance, or a single coordinate system, is estimated during adaptation.

The source estimation model predicts the experimental data well (compare **Fig. 1e** with **Fig. 1f**) under the assumption that changes in the body are more likely than changes in the world ($\alpha = 0.4$). To

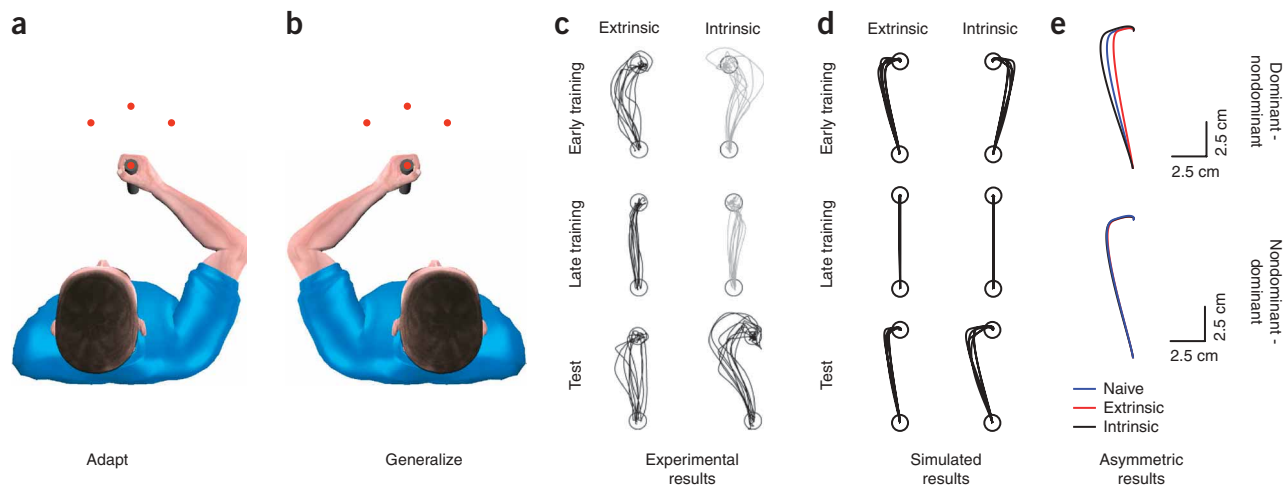


Figure 2 Simulated interlimb generalization. **(a)** Adaptation and **(b)** interlimb generalization with a curl field. **(c)** Experimental evidence (reproduced from ref. 6) showing early exposure, adaptation (late training) and interlimb generalization (test) in a clockwise (extrinsic) and counter-clockwise (intrinsic) field. **(d)** Simulated data with our model demonstrating similar findings as the result of a small estimated external field and large misestimates in limb parameters. **(e)** When we increase the body's uncertainty (decreasing α to 0.1), our model predicts asymmetric generalization. The top plots display generalization from the dominant to the nondominant limb; the bottom plots display generalization from the nondominant to the dominant limb with little transfer of the adapted skill.

quantify this similarity, we compared signed perpendicular errors at 25%, 50% and 75% of the movement displacements. Our model had an r.m.s. error of 4.81 cm, whereas the results based on a hypothesis of world-only or body-only inferences (**Figs. 1c** and **1d**, respectively) had r.m.s. errors of 11.0 cm and 7.70 cm, respectively.

A close inspection of the experimental generalization to an intrinsic field (**Fig. 1e**, top) shows that the trajectories are not perfectly straight but tend to be slightly curved into the direction predicted by generalization after estimating world errors (**Fig. 1c,f**, top). This indicates that the nervous system may be generalizing simultaneously in hand-based and joint-based coordinates, while placing a clear emphasis on the latter. The model we present here naturally shows this kind of mixed generalization.

Examining movements with only one limb cannot fully reveal how the nervous system adapts and generalizes. Consequently, generalization to the other limb (interlimb generalization) is also widely tested. In a notable study of interlimb generalization⁶, subjects once again adapted to a robot-rendered, velocity-dependent force field while making reaching movements (**Fig. 2a**). One group of subjects adapted to a hand-based clockwise curl field, and another group adapted to a counterclockwise curl field. All subjects were then tested using their other limb in a counterclockwise field to examine generalization (**Fig. 2b**).

The experiment aimed to distinguish between two hypotheses. If internal models use hand-based (extrinsic) coordinates, then they should transfer to the other hand when reaching in the same field and facilitate motor adaptation. But if internal models use joint-based coordinates (intrinsic), then making reaching movements with the other hand would generalize only if the field switched from a clockwise to a counterclockwise field upon switching hands. Subjects adapting to a counterclockwise curl field showed some ability to compensate for the same counterclockwise curl field when they switched to their other hand (**Fig. 2c**). However, when they adapted to a clockwise curl field and switched to their other hand, their paths were even more displaced than those of naive subjects (**Fig. 2c**). The investigators interpreted this as supporting evidence for internal models with an extrinsic reference frame during interlimb generalization.

When the source estimation model adapts to the same counterclockwise curl field, it will assign more credit to the limb properties but some to the world properties. However, only the estimated world parameters will transfer to the other limb. This estimated external disturbance is weak but still allows the other limb to partially compensate for the field during the initial exposure (**Fig. 2d**). By contrast, when adapting to the clockwise curl field, the estimated world parameters transfer to the other limb to estimate a clockwise disturbance. When a counterclockwise curl field is then presented to the other hand, the limb's trajectory is greatly displaced. The transferred estimates of the world disturbance are used to compute a compensatory motor command that only acts to strengthen the field's effects (**Fig. 2c,d**). In this respect, the source estimation model predicts the pattern of interlimb generalization well.

The source estimation model may help to explain another salient property of motor generalization: no transfer of adapted skills was found when subjects trained with their nondominant hand and switched to their dominant hand. As we have demonstrated, the nervous system seems to favor interpreting motor errors in terms of uncertainty and variability in limb properties. We expect this uncertainty to be larger for the less familiar, nondominant limb, and several studies have examined this notion^{7,8}. When α indicates an even higher uncertainty about the body for the nondominant hand, we also find that adaptation generalizes from the dominant to the nondominant hand but not vice versa (**Fig. 2e**, $\alpha = 0.1$). This is because motor errors with the nondominant hand are even more likely to be attributed to a misestimate of that limb's properties, information that does not facilitate subsequent movements with the dominant limb.

Rotating room studies

Experiments using robots to perturb human hand movements do not reveal all aspects of generalization behavior. In another set of experiments, researchers analyzed reaching movements made in a rotating room^{9–11}. The rotating room induces nonlinear Coriolis forces on the limb proportional to the limb's translational velocity (**Fig. 3a**). Though this study was distinct from the previous studies in a number of ways, the results were generally similar. Subjects initially produced motor

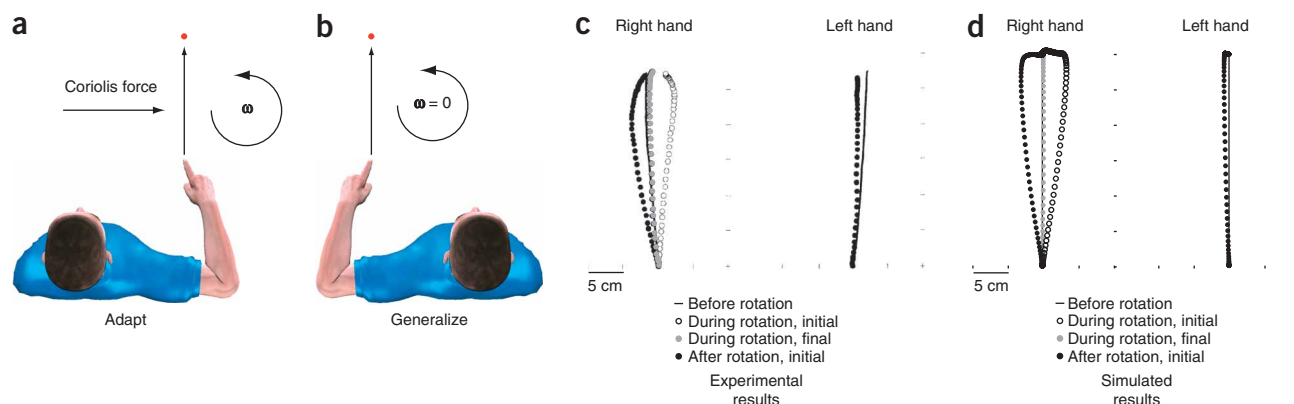


Figure 3 Simulated Coriolis room generalization. The room rotates with angular rate ω . (a) The subject adapts to a Coriolis force disturbance and (b) interlimb generalization is observed. (c) Experimental results of motor adaptation and interlimb generalization (reproduced from ref. 9). (d) Simulated experiment predicting similar results.

errors and then adapted and showed aftereffects consistent with a newly learned internal model. When subjects were then asked to make reaches with the alternate arm (Fig. 3b), their movements were relatively straight but rotated slightly, appearing as if they were made to a laterally displaced target (Fig. 3c). Instead of explaining the results in terms of coordinate frames, the investigators proposed that this shift in the reaching direction represents a transferred change in the motor system's kinematic plan and is thus evidence for distinct representations, one for movement dynamics and another for kinematics⁹. Finding evidence for two distinct representations, the experiment supports one of the central assumptions of our source estimation model.

Simulated reaches made with the source estimation model acting under a Coriolis force disturbance predict a similar pattern of adaptation and generalization (Fig. 3d). The model adapts to the disturbance and, when the room stops rotating, produces aftereffects consistent with the experimental data. As in the previous simulated results, the model adapts to the rotating room by inferring relatively large misestimates in the body's parameters and small estimates in the world parameters. When the left arm makes a subsequent reaching movement, it uses this relatively small, inferred world disturbance. However, this information causes the limb to compensate for a nonexistent external disturbance (the room is no longer rotating), and the movement direction is slightly altered, just as had been found experimentally. To quantify this similarity, we again compared perpendicular errors. Our model had an r.m.s. error of 3.81 mm, whereas the same results for a body-only or a world-only inference are 16.7 mm and 10.0 mm, respectively.

Similarly to the force field studies examined earlier, the rotating room studies found asymmetric generalization between the dominant and nondominant limbs. Again, this is consistent with the source estimation model. If we assume an increased uncertainty for the nondominant limb's parameters, then the model predicts generalization to the dominant limb that is nearly indistinguishable from a naive movement (see Supplementary Fig. 1 online). The nervous system's preference for relatively large uncertainty in the body (and even larger uncertainty in the nondominant limb) results in motor errors being largely attributed to

body parameters and only weakly to world parameters that would generalize to the other arm.

Inertial disturbance studies

Velocity-dependent disturbances were used in all the studies above. However, the nervous system is constantly faced with disturbances of other kinds as well. For example, movements with a mass introduce forces that also vary with the acceleration and orientation of the limb. We thus proceeded to compare the source estimation model with results from a motor adaptation paradigm that introduces inertial perturbations^{12–14}. In these studies, the subject wears an arm splint to which a mass is attached (Fig. 4a). This disturbance increases the subject's apparent limb inertia by an unknown amount, and reaching movements initially are distorted. Just as in other motor adaptation studies, the subject eventually adapts to the disturbance and produces relatively straight reaches. Similarly, when the disturbance is removed, the subject shows aftereffects indicative of the now absent inertia. In a study using this paradigm¹⁵, subjects adapted to reaching movements with the inertial disturbance and then the mass was attached to their opposite limb to observe evidence for the transfer of this motor skill (Fig. 4b). Initial movements made with the opposite limb after training performed better than under naive conditions (Fig. 4c). The researchers interpreted these results as indicating an intrinsic coordinate frame for interlimb generalization.

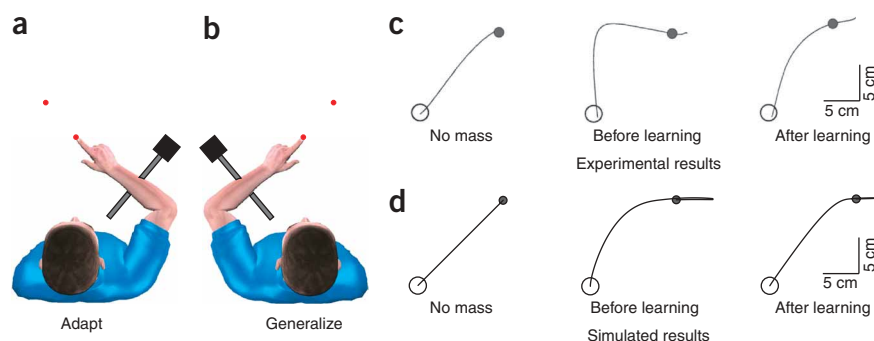


Figure 4 Simulated inertial perturbation generalization. (a) Subject adapts to an inertial disturbance with the dominant limb, and (b) interlimb generalization is observed with the nondominant limb. (c) Experimental results for initial-exposure reaches with nondominant left arm without a mass and with a mass before and after training with the dominant right arm (reproduced from ref. 15). (d) Simulated experiment demonstrating similar predictions for initial reaches with nondominant arm.




Figure 5 Comparison of aftereffects. Maximum perpendicular errors (PD) from a straight line are plotted versus trials. (a) Experimental data displaying reach errors (mean of six subjects) during adaptation (left portion) and aftereffects with and without holding a robot handle (black and gray lines, respectively; shaded regions indicate one standard error of the mean. Reproduced from ref. 4). (b) Simulated data for the same experiment (mean of eight simulated runs). When the simulated limb lets go of the handle, world estimates are neglected. (c) Simulated data with the two-module source estimation model.

When the source estimation model adapts and generalizes to a simulated inertial disturbance, the findings are similar to the experimental results (Fig. 4d). The model infers that the observed motor errors are due largely to an incorrect estimate of the limb's inertia and estimates a relatively small value for the external mass. When the nondominant limb is then tested, it makes reaches with the benefit of a small, inferred world disturbance (the added external mass). This results in the observed generalization. Again comparing perpendicular movement errors, our model has an r.m.s. error of 16.5 mm, whereas a purely world-only or body-only inference results in r.m.s. errors of 17.3 mm and 23.5 mm, respectively.

As is consistent with the previous studies, this transfer of adaptation was only observed when the initial training was performed with the dominant hand. Previously, this asymmetry was viewed as evidence that the nervous system is less proficient at developing an internal model for the nondominant limb. As in those studies, we find this asymmetry when we introduce more uncertainty into the representation of the nondominant hand (see Supplementary Fig. 1).

Generalization to natural hand movements

As our source estimation model adapts the properties of both the body and the world, a crucial test of its validity would be an experiment that independently tests for world and body adaptations. For instance, according to the source estimation model, if adapting to an experimental disturbance results in new estimates of limb parameters, then this information should influence movements made after the experiment has ended. A few recent studies have specifically examined these kind of aftereffects^{4,16}. In one study in particular⁴, this predicted phenomenon was found: after the subject released the robot handle, some aspect of the training persisted. In this study, subjects adapted to a velocity-dependent force field and were then tested with and without grasping the robot handle. The results provide strong evidence for the representation of a world model separate from the body model, just as we have hypothesized. Notably, the authors found large aftereffects even when subjects were not grasping the handle (Fig. 5a). Again, our model can offer a concise explanation for these results. When the field is turned off, the new estimates of both the limb and the body act to produce aftereffects with a magnitude approximately equivalent to the force field but an orientation opposite to it (Fig. 5b). However, when the model makes movements without holding the robot handle, the new world estimates are neglected, but the new body estimates persist. This results in aftereffects of a similar nature but of reduced magnitude (Fig. 5b).

The model results may seem to suggest that once the experiment has ended, subjects undergo a relatively slow period of 'de-adapting' to return to their nominal, preexperiment state (compare Fig. 5a to Fig. 5b). However, previous work has proposed that the nervous

system may not rely on a single representation of the body and may switch its model of the body, or of the world, in a context-dependent manner^{17,18}. These approaches can easily be combined with the source estimation model presented here. For example, a pair of body and world models could be grouped together within a specific module appropriate for a given context. Then several modules could be stored and combined at any time according to their abilities to account for the current context¹⁹.

To demonstrate this multiple-module system, we built a new source estimation model with two world-body modules. We assume that one module correctly represents unperturbed reaches (a correct limb model and no world disturbances), and a second module is used to account for new, perturbed motor contexts. When the force field is turned on, the large motor errors indicate a small likelihood that the unperturbed module represents this context, and the second module's likelihood is large (see Methods). As a result, the perturbed module adapts its world and body parameters. Once the force field is turned off, the motor errors (now in the opposite orientation) indicate that the perturbed module's likelihood of representing this context quickly decreases, as the unperturbed module's likelihood increases. The unperturbed module's high likelihood of representing the current context essentially allows it to 'win' over the perturbed module. This results in a quicker reduction in the observed aftereffects (Fig. 5c). Moreover, by quickly switching to the unperturbed module, the body and world models of both modules undergo minimal adaptation, allowing for their storage and later retrieval. These results, along with the results of our basic model (Fig. 5b), demonstrate that the source-estimation scheme comes close to capturing what could be the salient temporal and spatial features of adaptation and generalization under such conditions.

DISCUSSION

Many traditional studies of motor control posit a single internal model that adapts its parameters to reduce observed motor errors. Regardless of how these errors arise, the model parameters are adjusted to compensate for perceived disturbances and reduce future motor errors. However, since these disturbance models do not distinguish between errors arising from the properties of the body and the properties of the world, their ability to make rational predictions for generalization is limited. For example, these models have no mechanism for indicating that what the nervous system learns about the right arm should generalize across the workspace but not to the left arm. For this reason, these internal models have no option but to use the learned disturbance model to generalize under all circumstances.

By design, these internal models can only generalize along one coordinate frame. Capitalizing on the differences between extrinsic and intrinsic coordinate frames, experimental data has been used to examine which of these two coordinate frames best fits the observed data. Yet, if the nervous system is simultaneously generalizing with parameters that vary in both coordinate frames, then the evidence for these internal models should be mixed. Indeed, the experimental evidence is far from clearly supporting one coordinate frame over the other. Some intralimb generalization studies have argued for an intrinsic representation^{5,20–22}, whereas some evidence supports an extrinsic representation²³, and yet others have found evidence for what is argued to be a mixed representation³. Studies of interlimb generalization, in contrast, have found support for an extrinsic

representation^{6,24,25}. Further confounding the issue, generalization seems to be asymmetrical: newly acquired skills transfer from the dominant arm to the nondominant, but not vice versa^{7,8,12,26,27}. Jointly, these studies paint a complicated picture of inter- and intralimb generalization, dependent on the task as well as the limb used. As we have demonstrated, this may be due to an overly restrictive interpretation for describing internal models.

A recent line of work addresses how a wide range of properties of the nervous system can be understood using bayesian statistics. Neural coding^{28,29}, sensorimotor integration^{30,31}, sensory cue combination^{32–36} and cognitive decision making³⁷ have all been described in the same normative framework. In this study, we have extended these ideas to motor adaptation and generalization—where many variables need to be estimated simultaneously and the problem of ambiguity between the body and the world becomes important.

Many properties of the body and the world are variable and can contribute to motor errors. Moreover, they can contribute to errors in an ambiguous way. Here we have proposed that the nervous system therefore needs to estimate the sources of errors. Thus, we extended the traditional internal model representation to one of a coupled body–world representation. This model allows for the estimation of likely sources of motor errors and provides a commonsense method for generalizing on the basis of these estimations. Further, our model offers an explanation for the inconsistent findings on extrinsic and intrinsic representations for internal models and the longstanding observation of asymmetric motor generalization. In addition, a source estimation strategy may account for the apparently contradictory evidence found examining adaptation to kinematic and dynamic disturbances. For example, in one experiment¹⁴, kinematic and dynamic variables were relatively uncorrelated to each other and found to have little influence on their respective adaptations. In a follow-up experiment³⁸, the kinematic and dynamic disturbances were highly correlated and their respective influence on adaptation was substantially disrupted. In general, whenever two experimental perturbations can be explained by shared sources, adaptation is not independent when using a source estimation strategy.

The source estimation model we have presented here is a preliminary attempt to formalize the problems of motor adaptation and generalization. There are many modifications that future implementations could consider. For instance, in the current model we have assumed that the position and velocity of the limb's endpoint (hand) is the only sensory information available, but everything from grip force to muscle spindle fiber afferent signals could be incorporated into the model. Furthermore, at present we have almost entirely neglected the temporal aspects of adaptation and instead focused on the 'steady-state' conditions, wherein motor errors are sufficiently small as to justify the halt of adaptation. However, our model can be extended to include adaptation to parameters that have their own timescales³⁹. This may be useful in examining studies in which perturbations are gradually, as opposed to suddenly, presented to the subject. For instance, one study of interlimb generalization²⁴ found that subjects who adapted to a gradually applied perturbation showed little generalization, whereas subjects who adapted to a sudden-onset perturbation showed the familiar characteristics of extrinsic generalization. Assuming that the nervous system entertains abrupt changes in motor errors as more likely the result of sources in the world, our model should predict similar findings. Finally, we have at present grouped parameter properties according to body and world to minimize the number of model parameters to be tuned. Future versions of the model could include independent properties of all relevant parameters, whether they are body, world, extrinsic, intrinsic or other.

In this study, we have sidestepped the issue of motor learning by assuming that the goal of the motor system is to produce straight, minimum-jerk trajectories⁴⁰. This is a common assumption shared by many other studies and supported by a host of data, including the experimental studies modeled here. However, this clearly cannot be sufficient to explain all aspects of motor adaptation. Indeed, many studies of adaptation have demonstrated clear distinctions between motor behaviors before and after experimental disturbances^{41,42}. Rather than acting to maintain a desired trajectory, the nervous system seems to be optimizing some other criteria. Although the issues of optimal control and motor learning are the subjects of ongoing research^{43,44}, exactly what the nervous system is optimizing for is still unknown. Regardless, estimating the properties of the body and the world is a necessary first step toward reoptimizing a new motor plan.

METHODS

Generative model. The human upper limb is modeled as a nonlinear two-link, two-degree-of-freedom mechanism with joint angles \mathbf{q} as

$$I(\mathbf{q})\ddot{\mathbf{q}} + C(\mathbf{q}, \dot{\mathbf{q}})\dot{\mathbf{q}} + B_m\dot{\mathbf{q}} + K_m\mathbf{q} = \boldsymbol{\tau}_{\text{motor}} + \boldsymbol{\tau}_{\text{world}} + \boldsymbol{\tau}_{\text{noise}}$$

$I(\mathbf{q})$ and $C(\mathbf{q}, \dot{\mathbf{q}})$ are the inertial and velocity-dependent acceleration matrices, K_m and B_m represent the joint stiffness and damping that result from muscle forces as well as spinal and supraspinal feedback loops, $\boldsymbol{\tau}_{\text{motor}}$ is the commanded torque, $\boldsymbol{\tau}_{\text{world}}$ is an external disturbance from the world and $\boldsymbol{\tau}_{\text{noise}}$ is the vector of motor noise (with zero mean). As elsewhere⁵, we shall assume that the joint torque specified by the motor system is the sum of three components: (1) a feedback torque, (2) a feedforward controller to compensate for limb dynamics and (3) a compensatory term for estimated world disturbances.

$$\boldsymbol{\tau}_{\text{motor}} = [K_P(\mathbf{q}^*(t) - \mathbf{q}) + K_D(\dot{\mathbf{q}}^*(t) - \dot{\mathbf{q}})] \\ + [\hat{I}\ddot{\mathbf{q}}^*(t) + \hat{C}\dot{\mathbf{q}}^*(t) + \hat{B}_m\dot{\mathbf{q}}^*(t) + \hat{K}_m\mathbf{q}^*(t)] - [\hat{\boldsymbol{\tau}}_{\text{world}}]$$

\hat{I} , \hat{C} , \hat{B}_m and \hat{K}_m are estimates of the system variables. K_P and K_D are matrices that model the limb's (proportional and derivative) viscoelastic convergence to a commanded trajectory $\mathbf{q}^*(t)$.

We shall define the output of the system, $\mathbf{y}(t)$, as the position vector, \mathbf{x} and velocity vector, $d\mathbf{x}/dt$ of the limb's endpoint (or hand) in a cartesian reference frame, $\mathbf{y} = [\mathbf{x}(t), d\mathbf{x}(t)/dt]^T$. We shall further assume this output is corrupted by measurement noise, $\mathbf{n}(t)$, with zero mean and covariance R .

We follow the lead of previous studies and assume that the motor system can estimate a force field applied to the hand. This external perturbation is a velocity-dependent force field, \mathbf{F}_{ext} , modeled as $\hat{\mathbf{F}}_{\text{ext}} = \hat{B}_{\text{ext}}\frac{d\mathbf{x}}{dt}$. In addition, we will assume that an external mass, M , attached to the forearm at a known position can also be estimated as \hat{M} and that its influence on the limb's dynamics can be modeled appropriately. We only simulate estimations of the limb parameters I and B_m because changes in these two matrices are similar or equivalent to the disturbances that are imposed during the experimental studies we examine. Estimated variations from the nominal values for I and B_m , as well as the estimated values for B_{ext} and M (13 parameters in total) are used to compute $\boldsymbol{\tau}_{\text{motor}}$.

To compactly parameterize all these parameters, we introduce the vector \mathbf{p} and a vector of inferred parameter estimates, $\hat{\mathbf{p}}$. Therefore, we can express the estimated velocity-dependent field parameter, for example, as $\hat{B}_{\text{ext}} = [\hat{\mathbf{p}}_1, \hat{\mathbf{p}}_2, \hat{\mathbf{p}}_3, \hat{\mathbf{p}}_4]$. An estimate of the generative model can then be expressed as $\hat{\mathbf{y}}(\hat{\mathbf{p}}, t)$ (for brevity we shall refer to this estimate simply as $\hat{\mathbf{y}}(t)$).

To infer these parameters, \mathbf{p} , we will need to make one final assumption, that the parameters vary according a random walk with a small forgetting factor,

$$\mathbf{p}_i(t+\Delta) = a\mathbf{p}_i(t) + w_i(t)$$

where $a = 0.99$ and w_i is a zero-mean random variable drawn from a normal distribution with variance σ_i^2 . These variances represent 13 free parameters. Allowing each to vary independently would provide our model with freedom to capture the data well. However, to avoid the problem of overfitting, we restrict these values. We make the following assumptions: for each \mathbf{p}_i , we assume its

noise's magnitude is proportional to its nominal value, $\sigma_i = C p_i^0$, where we assume C takes one value for body parameters and another value for world parameters. As the magnitude of this parameter influences the rate of adaptation, a characteristic we are not investigating, we investigate only the ratio of its values and define this ratio as $\alpha = C_{\text{world}}/C_{\text{body}}$. This single free parameter was varied across all the simulations to find an approximate best fit to the data.

Optimal inferences. The task of the source estimation model is to infer changes, both in the body and the world, that may account for any perceived errors between what is observed and what is predicted. To develop our inference algorithm, we assume that

$$\hat{\mathbf{y}}(\mathbf{p}, t_n) - \hat{\mathbf{y}}(\hat{\mathbf{p}}, t_n) \approx \frac{\partial \hat{\mathbf{y}}}{\partial \mathbf{p}} (\mathbf{p} - \hat{\mathbf{p}}) \quad (1)$$

For small perturbations in \mathbf{p} , this is equivalent to a Taylor series approximation. However, as is demonstrated in our results, even when the quantity $(\mathbf{p} - \hat{\mathbf{p}})$ is large, the inference algorithm converges to correct values. Eqn (1) and our assumption of normal, uncorrelated noise and random walks in the parameters are the necessary assumptions to conclude that the Kalman filter is an optimal mechanism for performing inferences of \mathbf{p} . Therefore, to best estimate the parameter values based on the observed system output and our confidence in the current parameter values, we update the parameter values according to the rule

$$\hat{\mathbf{p}}_{n+1} = \hat{\mathbf{p}}_n + K_n (\mathbf{y} - \hat{\mathbf{y}})$$

where K_n is the Kalman gain matrix at the n^{th} update. This update is computed at discrete points in time $t_0, t_1, \dots, t_m, \dots, t_N$ during each movement. The standard update equations for the Kalman gain are

$$\begin{aligned} P_n^- &= A P_{n-1} A^T + Q \\ K_n &= P_n^- H_n^T [H_n P_n^- H_n^T + R]^{-1} \\ P_n &= [(P_n^-)^{-1} + H_n^T R^{-1} H_n]^{-1} \end{aligned}$$

where Q is the covariance of the parameter \mathbf{p} (defined above through α), R is the measurement error covariance and P_n is residual covariance of \mathbf{p} . A is a forgetting factor, aI . For our nonlinear model, we approximate H_n , the output matrix at the n^{th} time step, as

$$[H_n]_{ij} = (\hat{y}_i(\hat{\mathbf{p}}, t_n) - \hat{y}_i(\hat{\mathbf{p}} + \delta_j, t_n)) / \delta \quad (2)$$

where δ_j is a small perturbation in the j^{th} element of \mathbf{p} . This is a back-difference approximation to the gradient $\partial \hat{\mathbf{y}} / \partial \mathbf{p}$ evaluated at time t_n . While eqn (1) holds and eqn (2) is accurate, we have the guarantee that our update rule will yield a set of parameters $\hat{\mathbf{p}}$ that minimizes the variance in our estimates.

To simulate inferences made with multiple modules, we introduce a second set of parameters that are relevant for the module representing the force field, $\mathbf{p}_{\text{field}}$. The nominal module's parameters (no field present) will be denoted \mathbf{p} . We also introduce the parameter λ , a binary variable with a continuous probability distribution. If $\lambda = 1$, an external field exists; if $\lambda = 0$, there is no field. Using these separate parameters, the force field module and the nominal module each have their own estimated outputs, $\hat{\mathbf{y}}(\hat{\mathbf{p}}_{\text{field}})$ and $\hat{\mathbf{y}}(\hat{\mathbf{p}})$, respectively. The source estimation model's overall estimate is then based on a weighted combination of the two module's parameters, $\mathbf{p}_{\text{total}} = P(\lambda = 1) \mathbf{p}_{\text{field}} + P(\lambda = 0) \mathbf{p}$. The commands used to control the limb are based on the same weighted combination of parameters.

The probability of λ is computed by comparing output estimates based on the two modules,

$$P(\lambda = 1) = \frac{P(\hat{\mathbf{y}}(\mathbf{p}_{\text{field}}) | \lambda = 1) P(\lambda = 1)}{[P(\hat{\mathbf{y}}(\mathbf{p}_{\text{field}}) | \lambda = 1) P(\lambda = 1) + (\hat{\mathbf{y}}(\mathbf{p}) | \lambda = 0) P(\lambda = 0)]}$$

where the probability of the force field module's estimate is,

$$P(\hat{\mathbf{y}}(\mathbf{p}_{\text{field}}) | \lambda = 1) P(\lambda = 1) \propto \exp\left(-(\mathbf{y} - \hat{\mathbf{y}}(\mathbf{p}_{\text{field}}))^T R^{-1} (\mathbf{y} - \hat{\mathbf{y}}(\mathbf{p}_{\text{field}})) / 2\right)$$

and the probability of the nominal module's estimate (the no-field condition) is similarly defined. The probability of λ is computed at each update time t_n and

is then used to scale the inferences as

$$\hat{\mathbf{p}}_{\text{field}, n+1} = \hat{\mathbf{p}}_{\text{field}, n} + P(\lambda = 1) K_n (\mathbf{y} - \hat{\mathbf{y}})$$

Simulations. The limb parameter values were based on ref. 5. For all simulated force field experiments, the field was equivalent to that used in the study. Because of the size of the stiffness and damping parameters, the Coriolis forces were largely mitigated. To counteract this, in simulations the room rotated at five times the experimental rate. For similar reasons, in the simulated inertial disturbance study the perturbing mass was five times as large as the experimental one. For all simulated studies, we assumed \mathbf{q}^* was a minimum-jerk trajectory specified by the target locations and the movement times reported. Parameter estimates were updated six times per movement, and movement targets were randomly selected until the movement errors were sufficiently small. Software that computes this study's inferences and generates the figures in this manuscript is available as **Supplementary Software** online.

Note: Supplementary information is available on the Nature Neuroscience website.

ACKNOWLEDGMENTS

We would like to thank S. Mussa-Ivaldi, J. Patton, R. Shadmehr and C. Ghez for helpful discussions. This work has been supported by the US National Institutes of Health (grant RO1 NS057814-01) and by the Falk Trust.

Published online at <http://www.nature.com/natureneuroscience/>

Reprints and permissions information is available online at <http://npg.nature.com/reprintsandpermissions/>

1. Conditt, M.A., Gandolfo, F. & Mussa-Ivaldi, F.A. The motor system does not learn the dynamics of the arm by rote memorization of past experience. *J. Neurophysiol.* **78**, 554–560 (1997).
2. Conditt, M.A. & Mussa-Ivaldi, F.A. Central representation of time during motor learning. *Proc. Natl. Acad. Sci. USA* **96**, 11625–11630 (1999).
3. Ahmed, A.A., Wolpert, D.M. & Flanagan, J.R. Flexible representations of dynamics are used in object manipulation. *Curr. Biol.* **18**, 763–768 (2008).
4. Cothros, N., Wong, J.D. & Gribble, P.L. Are there distinct neural representations of object and limb dynamics? *Exp. Brain Res.* **173**, 689–697 (2006).
5. Shadmehr, R. & Mussa-Ivaldi, F.A. Adaptive representation of dynamics during learning of a motor task. *J. Neurosci.* **14**, 3208–3224 (1994).
6. Criscimagna-Hemminger, S.E. *et al.* Learned dynamics of reaching movements generalize from dominant to nondominant arm. *J. Neurophysiol.* **89**, 168–176 (2003).
7. Sainburg, R.L. & Kalakanis, D. Differences in control of limb dynamics during dominant and nondominant arm reaching. *J. Neurophysiol.* **83**, 2661–2675 (2000).
8. Sainburg, R.L. Evidence for a dynamic-dominance hypothesis of handedness. *Exp. Brain Res.* **142**, 241–258 (2002).
9. Dizio, P. & Lackner, J.R. Motor adaptation to Coriolis force perturbations of reaching movements: endpoint but not trajectory adaptation transfers to the nonexposed arm. *J. Neurophysiol.* **74**, 1787–1792 (1995).
10. Lackner, J.R. & Dizio, P. Rapid adaptation to Coriolis force perturbations of arm trajectory. *J. Neurophysiol.* **72**, 299–313 (1994).
11. Lackner, J.R. & Dizio, P. Gravitoinertial force background level affects adaptation to coriolis force perturbations of reaching movements. *J. Neurophysiol.* **80**, 546–553 (1998).
12. Wang, J. & Sainburg, R.L. Limitations in interlimb transfer of visuomotor rotations. *Exp. Brain Res.* **155**, 1–8 (2004).
13. Bock, O. Early stages of load compensation in human aimed arm movements. *Behav. Brain Res.* **55**, 61–68 (1993).
14. Krakauer, J.W., Ghilardi, M.F. & Ghez, C. Independent learning of internal models for kinematic and dynamic control of reaching. *Nat. Neurosci.* **2**, 1026–1031 (1999).
15. Wang, J. & Sainburg, R.L. Interlimb transfer of novel inertial dynamics is asymmetrical. *J. Neurophysiol.* **92**, 349–360 (2004).
16. Lackner, J.R. & DiZio, P. Motor control and learning in altered dynamic environments. *Curr. Opin. Neurobiol.* **15**, 653–659 (2005).
17. Haruno, M., Wolpert, D.M. & Kawato, M. Mosaic model for sensorimotor learning and control. *Neural Comput.* **13**, 2201–2220 (2001).
18. Schaal, S. & Schweighofer, N. Computational motor control in humans and robots. *Curr. Opin. Neurobiol.* **15**, 675–682 (2005).
19. Wolpert, D.M. & Kawato, M. Multiple paired forward and inverse models for motor control. *Neural Netw.* **11**, 1317–1329 (1998).
20. Shadmehr, R. & Mussa-Ivaldi, F.A. Spatial generalization from learning dynamics of reaching movements. *J. Neurosci.* **20**, 7807–7815 (2000).
21. Malfait, N., Shiller, D.M. & Ostry, D.J. Transfer of motor learning across arm configurations. *J. Neurosci.* **22**, 9656–9660 (2002).
22. Malfait, N., Gribble, P.L. & Ostry, D.J. Generalization of motor learning based on multiple field exposures and local adaptation. *J. Neurophysiol.* **93**, 3327–3338 (2005).
23. Krakauer, J.W. *et al.* Learning of visuomotor transformations for vectorial planning of reaching trajectories. *J. Neurosci.* **20**, 8916–8924 (2000).

24. Malfait, N. & Ostry, D.J. Is interlimb transfer of force-field adaptation a cognitive response to the sudden introduction of load? *J. Neurosci.* **24**, 8084–8089 (2004).
25. Burgess, J.K., Bareither, R. & Patton, J.L. Single limb performance following contralateral bimanual limb training. *IEEE Trans. Neural Syst. Rehabil. Eng.* **15**, 347–355 (2007).
26. Laszlo, J.I., Baguley, R.A. & Bairstow, P.J. Bilateral transfer in tapping skill in the absence of peripheral information. *J. Mot. Behav.* **2**, 261–271 (1970).
27. Bagesteiro, L.B. & Sainburg, R.L. Handedness: dominant arm advantages in control of limb dynamics. *J. Neurophysiol.* **88**, 2408–2421 (2002).
28. Gold, J.I. & Shadlen, M.N. The neural basis of decision making. *Annu. Rev. Neurosci.* **30**, 535–574 (2007).
29. Ma, W.J. *et al.* Bayesian inference with probabilistic population codes. *Nat. Neurosci.* **9**, 1432–1438 (2006).
30. Wolpert, D.M., Ghahramani, Z. & Jordan, M.I. An internal model for sensorimotor integration. *Science* **269**, 1880–1882 (1995).
31. Kording, K.P. & Wolpert, D.M. Bayesian integration in sensorimotor learning. *Nature* **427**, 244–247 (2004).
32. Ernst, M.O. & Banks, M.S. Humans integrate visual and haptic information in a statistically optimal fashion. *Nature* **415**, 429–433 (2002).
33. Ernst, M.O. & Bulthoff, H.H. Merging the senses into a robust percept. *Trends Cogn. Sci.* **8**, 162–169 (2004).
34. Yuille, A. & Kersten, D. Vision as Bayesian inference: analysis by synthesis? *Trends Cogn. Sci.* **10**, 301–308 (2006).
35. Kording, K.P. *et al.* Causal inference in multisensory perception. *PLoS ONE* **2**, e943 (2007).
36. Weiss, Y., Simoncelli, E.P. & Adelson, E.H. Motion illusions as optimal percepts. *Nat. Neurosci.* **5**, 598–604 (2002).
37. Tenenbaum, J.B., Griffiths, T.L. & Kemp, C. Theory-based Bayesian models of inductive learning and reasoning. *Trends Cogn. Sci.* **10**, 309–318 (2006).
38. Tong, C., Wolpert, D.M. & Flanagan, J.R. Kinematics and dynamics are not represented independently in motor working memory: evidence from an interference study. *J. Neurosci.* **22**, 1108–1113 (2002).
39. Kording, K.P., Tenenbaum, J.B. & Shadmehr, R. The dynamics of memory as a consequence of optimal adaptation to a changing body. *Nat. Neurosci.* **10**, 779–786 (2007).
40. Hogan, N. An organizing principle for a class of voluntary movements. *J. Neurosci.* **4**, 2745–2754 (1984).
41. Uno, Y., Kawato, M. & Suzuki, R. Formation and control of optimal trajectory in human multijoint arm movement. Minimum torque-change model. *Biol. Cybern.* **61**, 89–101 (1989).
42. Chib, V.S. *et al.* Haptic identification of surfaces as fields of force. *J. Neurophysiol.* **95**, 1068–1077 (2006).
43. Todorov, E. & Jordan, M.I. Optimal feedback control as a theory of motor coordination. *Nat. Neurosci.* **5**, 1226–1235 (2002).
44. Izawa, J. *et al.* Motor adaptation as a process of reoptimization. *J. Neurosci.* **28**, 2883–2891 (2008).