Published in final edited form as:

Nature. 2000 October 12; 407(6805): 742–747. doi:10.1038/35037588.

Learning of Action Through Adaptive Combination of Motor Primitives

Kurt A. Thoroughman¹ and Reza Shadmehr

Department of Biomedical Engineering, Johns Hopkins University

Abstract

Understanding how the brain constructs movements remains a fundamental challenge in neuroscience. The brain may control complex movements through flexible combination of motor primitives¹, where each primitive is an element of computation in the sensorimotor map that transforms desired limb trajectories into motor commands. Theoretical studies have shown that a system's ability to learn actions depends on the shape of its primitives². Using a time-series analysis of error patterns, here we find evidence that humans learn dynamics of reaching movements through flexible combination of primitives that have Gaussian-like tuning functions encoding hand velocity. The wide tuning of the inferred primitives predicts limitations on the brain's ability to represent viscous dynamics. We find close agreement between the predicted limitations and subjects' adaptation to novel force fields. The mathematical properties of the derived primitives resemble the tuning curves of Purkinje cells in the cerebellum. Activity of these cells may encode primitives that underlie learning of dynamics.

Studies of reaching movements have demonstrated that humans construct motor commands based on a prediction of forces that will be experienced in the upcoming movement³. When novel forces are imposed on the arm, the prediction is in error and the arm does not follow the desired trajectory³⁻⁴. With practice the motor commands are modified⁵ and the trajectory approximates the desired path. The learning of dynamics, however, affects movements outside the region of training^{3,6-8}, suggesting that the brain builds a state-dependent approximation of external forces⁹, called an internal model (IM). Occasional movements with unexpectedly altered dynamics, termed catch trials, have been used to quantify how the IM generalizes^{3,4}. Catch trials, however, not only test the IM for a given movement but cause errors that in turn change the IM and affect future movements. We demonstrate that the effect of errors experienced in a given movement on subsequent movements can reveal characteristics of primitives with which motor commands are generated.

We consider the IM to be a sensorimotor map transforming desired arm trajectories into muscle forces¹⁰⁻¹² through a flexible combination of a set of primitives:

$$\hat{\mathbf{f}} = \mathbf{W}^T \mathbf{g}(\mathbf{x}^*, \dot{\mathbf{x}}^*, \ddot{\mathbf{x}}^*) \quad (1)$$

where T is the transpose operator, $\hat{\mathbf{f}}$ is a vector approximation of forces $\mathbf{f}$ to be produced by muscles to compensate for task dynamics, and $\mathbf{g}$ is a vector of scalar valued primitives $[g_1, \dots, g_i]^T$. Although $\mathbf{g}$ generally can depend on desired position, velocity, and acceleration $(\mathbf{x}^*, \dot{\mathbf{x}}^*, \ddot{\mathbf{x}}^*)$, here we investigated learning of viscous forces and therefore considered a simpler

Correspondence should be addressed to R.S. (e-mail: reza@bme.jhu.edu).

¹Present address: Volen Center of Complex Systems and Department of Biology, Brandeis University, Waltham, Massachusetts 02454, USA

subset of primitive functions that depended only on desired velocity. The IM is learned through experience dependent modification of the weight matrix W . Assuming a learning rule that minimizes $\tilde{\mathbf{f}}^2 \equiv \|\mathbf{f} - \hat{\mathbf{f}}\|^2$ is adjusted after a movement (indexed 1) according to:

$$\Delta W_1 = -\eta \mathbf{g}(\dot{\mathbf{x}}_1^*) \mathbf{f}_1^T \quad (2)$$

where η is a constant learning step. This adaptation changes the IM output in the subsequent movement (indexed 2):

$$\begin{aligned} \hat{\mathbf{f}}_2(W + \Delta W_1) - \hat{\mathbf{f}}_2(W) &= (W + \Delta W_1)^T \mathbf{g}(\dot{\mathbf{x}}_2^*) - W^T \mathbf{g}(\dot{\mathbf{x}}_2^*) \\ &= -\eta \mathbf{g}^T(\dot{\mathbf{x}}_1^*) \mathbf{g}(\dot{\mathbf{x}}_2^*) \mathbf{f}_1. \end{aligned} \quad (3)$$

The change in the IM output depends on experienced error and the mutual projection between evaluations of the primitives, but does not depend on the weight matrix. Since the primitives depend on desired velocity, when the two movements have the same desired trajectory (e.g., toward the same target), the change should be proportional to the error experienced. When the two movements are toward different targets, the change will also depend upon the breadth of the receptive fields of the primitives.

We first tested whether an error experienced in a given movement causes a proportional change in the IM for the next movement to the same target. We asked subjects to make reaching movements while holding a manipulandum¹³ which produced viscous forces $\mathbf{f} = B \dot{\mathbf{x}}$, where $B = \{0, 13; -13, 0\}$ N·s/m. Catch trials, movements during which $\mathbf{f} = 0$, were randomly interspersed among the targets. Our proxy for error was hand displacement perpendicular to target direction (perpendicular displacement; p.d.) measured 250 msec into the movement. The first movement in the field (1st in Figure 1A) had significant error (p.d.=2.38 cm), but with training (ct-1 in Figure 1A) became less disturbed (p.d.=0.45 cm). In the next movement toward this direction (90°), a catch trial (ct), there was a large error in the direction opposite the initial error, suggesting formation of an IM¹³. In the subsequent movement to 90° (ct+1), during which the force field was present, p.d. was substantially greater (p.d.=1.22 cm) than in ct-1, indicating partial unlearning of the IM as predicted by Equation 2. In agreement with Equation 3, there was a significant correlation between magnitude of movement errors in ct and the unlearning observed in ct+1 (Figure 1B, $r=0.65$). The physiological correlate of this unlearning was evident in the spatial tuning of movement-initiating muscle activations. The computational construct of an IM predicted that spatial tuning of EMG would undergo a specific rotation with training⁵. During training, the preferred direction of this tuning gradually rotated. However, between ct-1 and ct+1 the preferred direction rotated back toward the initial orientation (Figure 1C), indicating unlearning of the IM. This unlearning washed out by movement ct+3 (Figure 1D).

We next investigated the shape of primitives underlying IM formation by quantifying, independent of the model in Equation 3, the temporal dynamics of movement errors first within and then across directions. In a sequence of random target directions, the time series of movement errors for a given direction was fitted to the following system of equations:

$$\begin{aligned} z_{n+1} &= az_n + bu_n \\ y_n &= z_n + du_n. \end{aligned} \quad (4)$$

Here y represented error in the IM as quantified by p.d.; n was movement number, u indicated whether the force field was present ($u=-1$) or turned off ($u=1$). The hidden state of the system, z , represented the amount of movement error generated by the IM; actual error (y) also depended on whether the force field was applied. The implicit assumption in this initial model was that errors experienced in one target direction did not affect the IM for generating movements toward other targets. The best-fit model correlated to actual errors reasonably well

(Figure 2A, black line; across directions, mean $r=0.60$). The fit mimicked subjects' recovery from initial error, their large error in ct, and their jump in error from ct-1 to ct+1. Whereas this initial model smoothly decayed after jumps in error, subjects often generated a nonmonotonic change in error between catch trials. We hypothesized that this was because errors experienced in one target direction changed the IM for other directions.

To investigate whether locally experienced errors affected other directions of movement, we expanded both u and b in Equation 4 to eight-dimensional vectors. Each element of the input vector u flagged recently experienced dynamics in a particular target direction i : whether, since the last movement in the modeled target direction, a force field movement ($u(i)=-1$) or a catch trial ($u(i)=1$) had been most recently experienced, or if no movement had occurred in direction i ($u(i)=0$). Each element of b , $b(i)$, quantified the sensitivity to errors experienced in direction i . The expanded model now accounted for subtle changes in actual movement error (Figure 2A, red line; mean $r=0.81$). Confidence intervals on b suggested that there was a significant, non-zero influence of local errors on subsequent control in other directions. To calculate sensitivity across target directions, the elements of b were re-indexed by the angular distance between the direction in which errors were experienced and the modeled movement direction. This distance was represented by ϕ . Averaging $b(\phi)$ across movement directions (Figure 2B) demonstrated that errors experienced in a given movement maximally influenced the IM for that direction. This influence decayed in neighboring directions. Surprisingly, sensitivity became significantly negative when angular distances were larger than 90° . This indicated that when two force field movements were separated by angle ϕ , if ϕ was small, then errors experienced in the first movement improved the IM for the second movement. If ϕ was large, then errors in the first movement destructively interfered with the IM used to generate the second movement.

To explain this result, note that sensitivity of the IM to experienced errors, $b(\phi)$, was quantified in terms of p.d. Both the output and the error signal of the IM, however, are in terms of force (Equations 1 and 2). Because the force field is linear in velocity, the direction of force error corresponding to positive (clockwise) p.d. toward one target opposes the direction of force error corresponding to positive p.d. toward the opposite target. Interpreting the sensitivity of subjects' IM (Figure 2B) through the adaptation rule (Equation 3) suggests that both the positive values of b for $-45^\circ < \phi < 45^\circ$ and the negative values of b for large ϕ correspond to the same direction of force compensation. From this we deduced that the mutual projection $\mathbf{g}^T(\dot{\mathbf{x}}_i^*)\mathbf{g}^T(\dot{\mathbf{x}}_{i+1}^*)$ declines but always remains positive as the angular distance between two movements increases. This result rules out bases that encode velocity space linearly. Furthermore, because information experienced in each direction most strongly affects that direction and less so its neighbor, basis functions that have specific regions of preferred activity are more likely to underlie learning than global representations of dynamics.

We therefore investigated what conditions on $\mathbf{g}(\dot{\mathbf{x}}^*)$ were sufficient to generate the generalization function $b(\phi)$. A salient property of cells in the motor system is their directional tuning¹⁴ and modulation with hand speed¹⁵. In the cerebellum, a region which lesion¹⁶⁻¹⁹ and functional imaging studies^{20,21} have linked to learning and control of arm dynamics, many Purkinje cells simultaneously encode the direction and speed components of velocity²². These cells broadly encode hand velocity during planar reaching, firing maximally at preferred velocities (PVs) distributed in velocity space. This encoding precedes in time the actual movement, suggesting that these cells encode desired velocity. The behavior of each cell k could therefore be represented as a Gaussian with a center located at position $\mathbf{c}_k$ in desired velocity space. We simulated a controller attached to a biomechanical model of the arm that learned an IM with basis functions:

$$g_k(\dot{\mathbf{x}}^*) = \exp\left(-\|\dot{\mathbf{x}}^* - \mathbf{c}_k\|^2 / 2\sigma^2\right) \quad (5)$$

where σ is the standard deviation of the Gaussian. To accommodate the possibility that the exact shape of $\mathbf{b}(\phi)$ depended on the training paradigm, we trained subjects and the simulated controller with the identical set of targets and catch trials. A crucial component of the simulations was σ , the width of the primitives. When the Gaussians were narrow, the time series of errors generated by the simulation showed a spike after each catch trial and a smooth decay afterwards, similar to the scalar-input state-space model fit but unlike the performance of our subjects. Simulations driven by broad Gaussians, however, produced nonmonotonic changes that mimicked subjects' actual patterns of adaptation. Simulation results were fitted with Equation 4 to produce the generalization function $\mathbf{b}(\phi)$ (Figure 2C). The $\mathbf{b}(\phi)$ generated with narrow Gaussians rapidly dropped to zero as ϕ changed from zero. Learning with wide Gaussians, however, exhibited a generalization that was very similar to actual subject performance, including negative sensitivity for large ϕ . The correlation between $\mathbf{b}(\phi)$ in the simulations and the subject data was strongest for $\sigma=0.12$ m/s.

We next used the model to predict behavior beyond the data set with which the primitives were estimated. We noted that Gaussian width influences how force estimation generalizes across both directions and speeds. Simulations predicted that when learning relied upon wide Gaussians, reaching movements would not monotonically converge onto a straight line desired trajectory but would become S-shaped (Figure 3A). While the force field was linear in velocity, wide Gaussians produced an approximation that overestimated forces at low speeds and underestimated forces at high speeds (Figure 3B). Overestimation of the forces resulted in overcompensation of the field early in the movement; the magnitude of overcompensation depended on the Gaussian width (Figure 3C). With narrow Gaussians, the simulated IM did not overcompensate, but with wide Gaussians movements became S-shaped. To test this prediction, we trained 24 subjects in target sets without catch trials. Movements of subjects were S-shaped (Figure 3E-1), similar to movements made by simulations that learned with wide Gaussians (Figure 3A).

Simulations further predicted that the probability of catch trials influenced whether movements would become S-shaped (Figure 3D). If catch trials occurred with 17% probability, then even with Gaussians of $\sigma=0.12$ m/s there should be sufficient unlearning caused by each catch trial such that hand trajectories would converge toward a straight line, without overcompensation (Figure 3F). We tested this prediction by training subjects in target sets with 17% catch trial probability. As predicted, subjects did not show overcompensation (Figure 3E-2).

Because approximation of a high frequency signal with low frequency bases generally results in poor representation, we next explored the limitations of a system that learns with wide Gaussians. We simulated learning of nonlinear force fields:

$$\mathbf{f} = -13 \sqrt{\dot{x}^2 + \dot{y}^2} \begin{bmatrix} -\sin(m\theta) \\ \cos(m\theta) \end{bmatrix} \quad \theta = \arctan(\dot{y} / \dot{x}) \quad (6)$$

where $\dot{x}$ and $\dot{y}$ were components of hand velocity in Cartesian space. When $m=1$, the field was the curl pattern¹³ learned by subjects above. As m increased, the field's spatial frequency increased. For various σ , we simulated movements to 50 targets and correlated the IM learned by the simulation to the actual force field (Figure 4A). As the field's spatial frequency increased, the accuracy of the IM decreased for all basis function widths. However, the learning capability of wider bases collapsed at lower frequencies. This agrees with the recent finding that humans demonstrated a lesser ability to adapt in higher spatial frequency force fields²³.

To illustrate this deficiency, we trained an adaptive controller ($\sigma=0.12$ m/s) for 50 movements in a high spatial frequency field ($m=4$, Figure 4B). Because approximation is performed with wide bases, the IM learned by the controller (Figure 4C) cannot represent faithfully the rapidly changing forces. In particular, the simulation predicts that in movements toward 22.5° and 157.5° the IM expects resistive forces where the force field is assistive. We tested for this prediction in three subjects by training each with the same random pattern of targets presented to the simulation, then presenting two catch trials toward 22.5° and 157.5° . Subjects behaved as though they were expecting a resistive field in those directions, as illustrated by their hand velocities in the catch trials (Figure 4D).

Errors in learning of arm movements suggest that the brain composes motor commands with computational elements that are broadly tuned in arm velocity. When expressed in polar coordinates, the Gaussians exhibit a preferred direction of movement, much like the cosine tuning curves typically associated with cells in the motor system. However, our inferred bases have on average a half-width at half-height value of about 40° , significantly less than the 90° value required of cosines. Recent results²⁴ have demonstrated that tuning curves in monkey motor cortex have a median width of 56° , also much narrower than cosines. Motor cortical cells, however, have been reported to encode hand speed linearly¹⁵. Learning a linear force field with bases that have cosine directional tuning and linear speed tuning results in an IM that does not produce the S-shaped movements observed in our subjects. Wide Gaussians predict this curious behavior. They also explain why humans generalize sublinearly to fast movements after training in a linear force field at slow speeds⁸. The nonlinear encoding of speed inherent to Gaussians resemble tuning properties of Purkinje cells that encode arm movements in the cerebellum²². While several investigators have proposed a major role for the cerebellum in learning of IMs²⁵⁻²⁷, our results suggest a link between patterns of generalization and firing properties of cells in this area. Since the output of the cerebellum partially affects the motor cortex²⁹, the finding that the preferred direction of motor cortical cells rotates during learning of force fields²⁸ may be a consequence of changing input from the cerebellum³⁰.

Methods

Three groups of right-handed normal human subjects were trained to make horizontal movements while holding the handle of a lightweight robotic arm. All movements were toward a pseudorandomly chosen target, then back to a center target. Targets appeared at 0° , 45° , ... 315° , at 10 cm displacement. Desired movement duration was 500 ± 50 ms. Timing feedback was provided by changing the target color. All subjects initially practiced the task without any perturbing force (the null field). The first group of subjects ($n=40$) trained in target sets of 192 movements during which the force field $B=\{0, 13; -13, 0\}$ N·s/m was applied in 83% of the targets. The remaining 17% of targets (pseudorandomly selected) were catch trials during which the force field was turned off. In 13 of these subjects, EMG from anterior and posterior deltoid, biceps, and triceps were measured with surface electrodes, amplified, filtered, RMS rectified, averaged from -50 to 100 ms into the movement, multiplied by unit vectors pointing toward movement direction, and summed across movement directions to produce a preferred direction for each muscle⁵. The second group of subjects ($n=24$) trained in the same target sets as above, but with the force field always on (no catch trials). A final group of subjects ($n=3$) trained for 58 movements in a higher frequency force field (Equation 6; $m=4$), receiving catch trials only on targets 51 and 58 in directions 22.5° and 157.5° .

A simulated anthropomorphic controller¹³ made movements to the same sequence of targets experienced by subjects. The controller's IM learned to map desired hand velocities into forces. Desired hand trajectories were minimum jerk, 0.5 sec in duration, and 10 cm long. The IM approximated the imposed force field (Equation 1) with Gaussian bases (Equation 5; controllers

were differentiated by σ). Gaussian centers spanned a range of desired velocities (-0.5 to 0.5 m/s in the x- and y-direction) spaced one σ apart. Sensitivity analyses suggested there were no significant changes in the results when the density of the bases was increased by an order of magnitude. Weights were initially randomized, then updated using Equation 2 with $\eta=0.0025$. To compare the output of Gaussian basis functions with cosine tuning curves, unweighted output of each Gaussian was averaged during movement time. The collection of averaged outputs across movement directions formed a tuning curve for each Gaussian primitive. Tuning curves were aligned to the preferred direction of each Gaussian and averaged across all Gaussians, from which the half-width at half-height value was calculated.

Acknowledgements

We thank M. A. Smith, O. Donchin, and R. Nezafat. The work was supported by grants from the Office of Naval Research and the National Institutes of Health (to R.S.), and a research traineeship from the National Science Foundation (K.A.T.).

References

1. Mussa-Ivaldi FA, Giszter SF, Bizzi E. Linear combinations of primitives in vertebrate motor control. *Proc. Natl. Acad. Sci. U S A* 1994;91:7534. [PubMed: 8052615]
2. Schaal S, Atkeson CG. Constructive incremental learning from only local information. *Neural Computation* 1998;10:2047–1084. [PubMed: 9804671]
3. Shadmehr R, Mussa-Ivaldi FA. Adaptive representation of dynamics during learning of a motor task. *J. Neurosci* 1994;14:3208–3224. [PubMed: 8182467]
4. Ghez, C.; Krakauer, JW.; Sainburg, RL.; Ghilardi, MF. *The New Cognitive Neurosciences*. Gazzaniga, MS., editor. MIT Press; Cambridge, Massachusetts: 2000. p. 501-514.
5. Thoroughman KA, Shadmehr R. EMG correlates of learning internal models of reaching movements. *J. Neurosci* 1999;19:8573–8588. [PubMed: 10493757]
6. Gandolfo F, Mussa-Ivaldi FA, Bizzi E. Motor learning by field approximation. *Proc. Natl. Acad. Sci. U S A* 1996;93:3843–3846. [PubMed: 8632977]
7. Conditt MA, Gandolfo F, Mussa-Ivaldi FA. The motor system does not learn the dynamics of the arm by rote memorization of past experience. *J. Neurophysiol* 1997;78:554–560. [PubMed: 9242306]
8. Goodbody SJ, Wolpert DM. Temporal and amplitude generalization in motor learning. *J. Neurophysiol* 1998;79:1825–1838. [PubMed: 9535951]
9. Conditt MA, Mussa-Ivaldi FA. Central representation of time during motor learning. *Proc. Natl. Acad. Sci. U S A* 1996;96:11625–11630. [PubMed: 10500227]
10. Kawato M. Adaptation and learning in control of voluntary movement by the central nervous system. *Advanced Robotics* 1989;3:229–249.
11. Sanner RM, Kosha M. A mathematical model of the adaptive control of human arm motions. *Biol. Cybern* 1999;80:369–382. [PubMed: 10365428]
12. Bhushan N, Shadmehr R. Computational architecture of human adaptive control during learning of reaching movements in force fields. *Biol. Cybern* 1999;81:39–60. [PubMed: 10434390]
13. Shadmehr R, Brashers-Krug T. Functional stages in the formation of human long-term motor memory. *J. Neurosci* 1997;17:409–419. [PubMed: 8987766]
14. Georgopoulos AP, Schwartz AB, Kettner RE. Neural population coding of movement direction. *Science* 1986;233:1416–1419. [PubMed: 3749885]
15. Moran DW, Schwartz AB. Motor cortical representation of speed and direction during reaching. *J. Neurophys* 1999;82:2676–2692.
16. Conrad B, Matsunami K, Meyer-Lohmann J, Wiesendanger M, Brooks VB. Cortical load compensation during voluntary elbow movements. *Brain Res* 1974;71:507–514. [PubMed: 4219750]
17. Goodkin HP, Keating JG, Martin TA, Thach WT. Preserved simple and impaired compound movement after infarction in the territory of the superior cerebellar artery. *Can. J. Neurol. Sci* 1993;20:S93–S104. [PubMed: 8334599]

18. Bastian AJ, Martin TA, Keating JG, Thach WT. Cerebellar ataxia: abnormal control of interaction torques across multiple joints. *J. Neurophysiol* 1996;76:492–509. [PubMed: 8836239]
19. Lang CE, Bastian AJ. Cerebellar subjects show impaired adaptation of anticipatory EMG during catching. *J. Neurophysiol* 1999;82:2108–2119. [PubMed: 10561391]
20. Nezafat R, Shadmehr R, Holcomb HH. Neural correlates of motor memory retention over a 4 week period as quantified through PET. *Soc. Neurosci. Abs* 1999;25:2178.
21. Imamizu H, et al. Human cerebellar activity reflecting an acquired internal model of a new tool. *Nature* 2000;403:192–195. [PubMed: 10646603]
22. Coltz JD, Johnson MTV, Ebner TJ. Cerebellar Purkinje cell simple spike discharge encodes movement velocity in primates during visuomotor arm tracking. *J. Neurosci* 1999;19:1782–1803. [PubMed: 10024363]
23. Matsouka, Y. *Models of generalization in motor control*. PhD thesis. MIT; 1998.
24. Amirikian B, Georgopoulos AP. Directional tuning profiles of motor cortical cells. *Neurosci. Res* 2000;36:73–79. [PubMed: 10678534]
25. Houk JC, Wise SP. Distributed modular architectures linking basal ganglia, cerebellum, and cerebral cortex: their role in planning and controlling action. *Cereb. Cortex* 1995;5:95–110. [PubMed: 7620294]
26. Wolpert DM, Kawato M. Multiple paired forward and inverse models for motor control. *Neural Networks* 1998;11:1317–1329. [PubMed: 12662752]
27. Spaelstra J, Schweighofer N, Arbib MA. Cerebellar learning of accurate predictive control for faster-reaching movements. *Biol. Cybern* 2000;82:321. [PubMed: 10804064]
28. Gandolfo F, Li CSR, Benda BJ, Schioppa CP, Bizzi E. Cortical correlates of learning in monkeys adapting to a new dynamical environment. *Proc. Natl. Acad. Sci. (USA)* 2000;97:2259. [PubMed: 10681435]
29. Holdefer RN, Miller LE, Houk JC. Functional connectivity between cerebellum and primary motor cortex in the awake monkey. *J. Neurophysiol* 2000;84:585–590. [PubMed: 10899231]
30. Martin JH, Cooper SE, Hacking A, Ghez C. Differential effects of deep cerebellar nuclei inactivation on reaching and adaptive control. *J. Neurophysiol* 2000;83:1886–1899. [PubMed: 10758100]

Figure 1.

Catch trials induced short-term unlearning. **A:** Hand trajectories of single movements, averaged across all 40 subjects, during the first (labeled 1st), 38th (ct-1), 39th (ct), and 40th (ct+1) movements toward 90° (0° is at 3 o'clock). Intersection of error bars indicates parallel and perpendicular displacement (p.d.) 250 ms into the movement. **B:** Jumps in p.d. between before (ct-1) and after (ct+1) catch trials versus p.d. in catch trials (ct), averaged within target directions across subjects. **C:** Orientation of preferred direction of movement-initiating EMG in ct-1, ct, and ct+1. Asterisks indicate significant ($p < 0.05$) rotation of EMG preferred direction between ct-1 and ct+1. **D:** Errors in force field movements following ct. In all figures, error bars indicate 95% confidence intervals of the mean (CIM) across subjects.

Figure 2.

Sensitivity to movement error across target directions. **A:** Errors in subjects' reaches ($\cdot$) (p.d. at 250 msec) toward 90°, averaged across subjects ($n=40$). Catch trials are the large negative spikes (one labeled ct). Lines are best fits of scalar (black) and vector (red) state-space models to the data (y in Equation 4). **B:** Average sensitivity of the IM to errors experienced in previous movements (b in Equation 4) as a function of angular distance ϕ . Error bars (95% CIM) calculated through bootstrapping. **C:** Average sensitivity $b(\phi)$ for simulated adaptive controllers that constructed an IM with Gaussians of width σ . $b(\phi)$ in simulations and subjects were averaged across target directions.

Figure 3. Movement characteristics of systems that learn an IM with velocity encoding Gaussians. **A:** Simulated adaptive controllers trained in a target set without catch trials. The 11th (left) and 21st (right) movements toward 90° are shown for each controller. Learning with wide Gaussians produced S-shaped movements. **B:** Adaptive controllers' estimation of the force field after training for 100 movements in a target set without catch trials. Peak velocity of typical movements is 0.35 m/s. **C:** Time course of error (p.d. at 200 msec) for simulated controllers with various Gaussian widths, smoothed across a 13-movement window. Overcompensation (generation of S-shaped hand trajectories) occurs when p.d. becomes negative. **D:** Error (p.d.) averaged over the subset of movements 65-128 during which the force

field was on. With narrow Gaussians, overcompensation never occurred, regardless of catch trial probability. With wide Gaussians, movements became S-shaped below critical probabilities (near 17%). **E**: Parallel and perpendicular displacements averaged across movements and subjects. **E1**: Subjects (n=24) trained in target sets without catch trials, resulting in S-shaped movements. **E2**: Subjects (n=40) trained in target sets with 17% catch trial probability and did not produce S-shaped movements. **F**: A controller relying upon wide Gaussians ($\sigma=0.12$ m/s) trained in target sets with 17% catch trial probability did not produce S-shaped trajectories.

Figure 4.

Learning with wide Gaussians imposes limits on adaptation. **A:** Correlation between actual force field and IM acquired after 50 targets by adaptive controllers with various width Gaussians, versus the frequency of the nonlinear force field (m in Equation 6). **B:** A high spatial frequency field ($m=4$) in which subjects and a simulated controller trained. **C:** The IM learned by the controller with $\sigma=0.12$ m/s Gaussians after training for 50 movements. Note that for directions 22.5° and 157.5°, the controller predicts the field to be resistive, whereas actual forces are assistive. **D:** Hand velocities parallel to target direction of three subjects trained in the same high spatial frequency field in the same target set as the simulated controller. Plotted are velocities in baseline null field movements (mean $\pm$ standard deviation) and during catch trials in targets 51 and 58 toward 22.5° and 157.5°. In catch trials, the hand is moving significantly faster, consistent with an IM that expects a resistive field.