Primitives for Motor Adaptation Reflect Correlated Neural Tuning to Position and Velocity

Gary C. Sing[1], Wilsaan M. Joiner[2], Thrishantha Nanayakkara[1], Jordan B. Brayanov[1] & Maurice A. Smith[1,3*]

1 Neuromotor Control Lab, Harvard School of Engineering and Applied Sciences, Cambridge, MA 02138, USA, **2** Laboratory of Sensorimotor Research, National Eye Institute, National Institutes of Health, Bethesda, MD 20892, USA, **3** Center for Brain Science, Harvard University, Cambridge, MA 02138, USA

[*]To whom correspondence should be addressed. E-mail: mas@seas.harvard.edu

The motor commands required to control voluntary movements under various environmental conditions may be formed by adaptively combining a fixed set of motor primitives. Since this motor output must contend with state-dependent physical dynamics during movement, these primitives are thought to depend on the position and velocity of motion. Using a recently-developed "error-clamp" technique, we measured the fine temporal structure of changes in motor output during adaptation. Interestingly, these measurements reveal that motor primitives echo a key feature of the neural coding of limb motion – correlated tuning to position and velocity. We show that this correlated tuning explains why initial stages of motor learning are often rapid and stereotyped, whereas later stages are slower and stimulus-specific. With our new understanding of these primitives, we design novel dynamic environments that are intrinsically the easiest and most difficult to learn, suggesting a theoretical basis for the rational design of improved procedures for motor training and rehabilitation.

Introduction

Given the high dimensionality of sensory input to the nervous system, tractable interaction with our environment requires efficient representation of sensory information and response planning (Atick, 1992; Barlow, 2001; Smith and Lewicki, 2006). One mechanism for obtaining this efficiency is the simultaneous encoding of multiple features of the environment in neuronal firing patterns (Barlow and Foldiak, 1989). The encoding of information about limb motion in the nervous system is one such example of multi-featured coding. For example, most neurons in primary motor cortex respond to both the position and velocity of voluntary limb movements (Ashe and Georgopoulos, 1994; Paninski et al., 2004; Wang et al., 2007), and their preferred directions for tuning to position and velocity are positively-correlated with one another (Paninski et al., 2004; Wang et al., 2007). Likewise, muscle spindle afferents display positively-correlated responses to position and velocity of muscle stretch by transiently increasing their firing rates while being stretched and sustaining a partially-elevated static firing rate when that stretch is subsequently maintained (Edin and Vallbo, 1990; Matthews, 1933; Prochazka, 1999). The representation of position and velocity state variables throughout the nervous system in sensory receptors (Edin and Vallbo, 1990; Matthews, 1933; Prochazka, 1999), sensorimotor association areas (Ashe and Georgopoulos, 1994), the cerebellum (Shidara et al., 1993), and the motor cortex (Ashe and Georgopoulos, 1994; Paninski et al., 2004; Wang et al., 2007) (whether they be in muscle, joint, or extrinsic coordinates) may result from the fact that position and velocity are the state variables for the physics of both intrinsic limb motion and environmental interactions (Hollerbach and Flash, 1982; Hwang et al., 2006; Mussa-Ivaldi and Bizzi, 2000; Shadmehr and Mussa-Ivaldi, 1994), which provide a complete characterization of the motion state of the body.

These dynamics of body/environment interactions are thought to be stored in internal models (Bhushan and Shadmehr, 1999; Diedrichsen et al., 2007; Kawato, 1999; Lackner and Dizio, 1994; Shadmehr and Mussa-Ivaldi, 1994; Shidara et al., 1993; Wagner and Smith, 2008; Wolpert et al., 1998) which provide neural representations of these dynamics. Two types of internal models are believed to coexist in the motor system. An inverse-dynamics model directly maps desired motion to

the motor command required to achieve it (Shadmehr and Mussa-Ivaldi, 1994; Shidara et al., 1993), while a forward model predicts the sensory consequences of a motor command (Diedrichsen et al., 2007; Kawato, 1999; Miall et al., 1993; Wolpert and Miall, 1996). Internal models with accurate characterizations of these dynamics allow us to make accurate movements in our environment. In order to maintain accuracy, both types of these internal models (inverse-dynamics and forward) must be able to adapt when fatigue, injury, or novel loads change the dynamics of the motor system (Wolpert and Ghahramani, 2000). It has been hypothesized that such adaptation might be facilitated by an internal model representation consisting of a linear combination of state-dependent motor primitives (Donchin et al., 2003; Mussa-Ivaldi and Bizzi, 2000; Poggio and Bizzi, 2004; Thoroughman and Shadmehr, 2000), and that the properties of these primitives can be identified by studying how motor learning in one context is generalized to another (Conditt et al., 1997; Conditt and Mussa-Ivaldi, 1999; Donchin et al., 2003; Gandolfo et al., 1996; Goodbody and Wolpert, 1998; Krakauer et al., 1999; Krakauer et al., 2000; Poggio and Bizzi, 2004; Shadmehr and Mussa-Ivaldi, 1994; Thoroughman and Shadmehr, 2000). Since both types of internal models may approximate physical dynamics by mapping between motor commands and the resulting body motion, these models can be characterized by dependence on motion state, i.e. the effect of a particular motor activation is largely determined by the current position and velocity of motion, as dictated by the physical mechanics of the arm. It is therefore not surprising that the motor system learns to associate external force perturbations with the *underlying state* of the limb rather than with the times at which they occur (Conditt et al., 1997; Conditt and Mussa-Ivaldi, 1999). Even when discrete actions like button presses are learned during arm movements, they are consistently associated with motion state rather than time (Diedrichsen et al., 2007).

Given (1) the pervasive coupling between the encoding of position and velocity in neural representations of motion, and (2) that adaptive internal models critically depend on these state variables, surprisingly little attention has been given to how interactions between state variables might influence motor learning (Bays et al., 2005; Hwang et al., 2003). Here we studied how the temporal patterns of motor output evolve during motor adaptation, and in particular, how these temporal patterns depend on motion state. We hypothesized that the neural processing of position and velocity

signals might lead to specific patterns of state-dependent crosstalk in motor adaptation, determine which types of dynamics are easy and difficult to learn, as well as dictate the patterns of interference expressed between adaptations to different types of dynamics.

Results

Progression of Learning in a Velocity-Dependent Force-Field

We hypothesized that the evolution of motor output during adaptation might provide insight into the nature of the computational basis elements underlying this type of learning. In Experiment 1, we examined this evolution when novel environmental dynamics were imposed on voluntary reaching arm movements. Subjects were trained to make straight, 100mm, 500ms point-to-point movements in the horizontal plane while grasping a manipulandum (Figure 1A). As subjects made reaching movements, we introduced two novel velocity-dependent force-field (FF) environments in which the manipulandum produced perturbing forces perpendicular to the subjects' reach motion that were proportional in magnitude to the reach velocity (Figure 1B). These clockwise (CW) and counter-clockwise (CCW) FFs initially perturb movements off-course, but with practice, subjects gradually learn to make straighter movements by producing forces to counteract the FF environment (Shadmehr and Mussa-Ivaldi, 1994). Although it is well known that this adaptation proceeds over the course of 100-400 trials (Shadmehr and Mussa-Ivaldi, 1994; Smith et al., 2006; Smith and Shadmehr, 2005) and is dependent on the cerebellum (Smith and Shadmehr, 2005), it has been a challenge to measure how the motor output evolves over the course of adaptation. The main difficulty here is that these 500ms reaching movements are not entirely ballistic – after 150-250ms, subjects produce substantial corrective feedback responses to compensate for perturbation-induced deviations in their movements (Cordo, 1987; Cordo, 1990; Shadmehr and Brashers-Krug, 1997; Tseng et al., 2007; Wagner and Smith, 2008). These responses mask the learning-related feedforward changes in motor output because the net motor output, which can be directly measured, reflects both feedforward changes and feedback responses. Here we employed error-clamp probe trials (Bays and Wolpert, 2006; Joiner and Smith, 2008; Scheidt et al., 2000; Smith et al., 2006; Wagner and Smith, 2008) to unmask these

learning-related changes in motor output. During these 100mm long probe trials, 99% of lateral kinematic errors are clamped to 1.2mm or less so that the feedback responses to these errors can be effectively eliminated and adaptive changes in motor output can be directly measured (Figure 1C). We pseudorandomly interspersed these probe trials among FF training trials to obtain estimates of how subjects' motor outputs progressed.

Interestingly, over the course of training, we observed a stereotyped evolution in not only the magnitude but also the shape of the force patterns that subjects learned to produce in order to counteract the effects of the FFs they experienced (Figure 2A). Early in training (i.e. within the first 20 trials), the force pattern contains a transient peak in the middle of the movement, *appropriate* for the transient nature of a velocity-dependent FF. However, seemingly *inappropriate* for this perturbation, the force pattern also contains a static "tail" at movement termination (Figure 2B). As the training progresses, the perturbation-appropriate transient response gradually increases, while the inappropriate static response diminishes. After 220-240 trials, the force pattern becomes almost entirely transient (Figure 2C). Closer observation reveals that the shape of the early-learning force output strongly resembles a muscle spindle firing pattern that is responding positively to both the velocity of stretch (resulting in the transient peak) and amount of stretch (leading to the static "tail") (Edin and Vallbo, 1990; Matthews, 1933; Prochazka, 1999). With this joint dependence in mind, we find that linear regression of the average early and late force patterns shown in Figure 2BC onto the shapes of the corresponding average position and velocity traces produces surprisingly good fits (early: R^2=0.93; late: R^2=0.98).

These results suggest that the learned force pattern at any point during training can be efficiently represented using a simple linear combination of position and velocity signals, multiplied by stiffness (K) and viscosity (B) gains, respectively. The progression of force patterns can then be represented as a progression of points in the two-dimensional space of stiffness and viscosity gains (Figure 2D). In this gain-space, the "learning trajectory" evolves toward a goal of [0,1], corresponding to complete learning of a velocity-dependent (viscous) FF. However, instead of progressing directly toward this goal via specific increases in the viscosity gain, i.e. climbing the y-axis in Figure 2D, the learning trajectory instead deviates towards the center of the 1st quadrant

(representing positive contributions from both position and velocity). The trajectory then gradually reduces its inappropriate position dependence as the velocity dependence continues to grow.

Why do the force patterns produced early in training for a *velocity*-dependent perturbation exhibit substantial, inappropriate *position* dependence while force patterns produced late in training do not? In other words, why does end-movement force persist early in training? One possibility is that this end-movement 'force-tail' reflects adaptation to end-movement kinematic errors experienced on FF trials. If such errors persisted at movement termination, they might explain the persistent force-tail. However, as shown in the bottom graph of Figure 2B, these errors do not persist (see Figure S3 for analysis of persistent end-movement kinematic error on the very first FF trial). Instead, lateral errors induced by the FF perturbation go away just 600ms after movement onset, whereas the force tail persists well beyond 1200ms.

Viscoelastic Primitive Model

An alternative explanation for the inappropriate position dependence is that this behavior could emerge if the motor primitives responsible for motor adaptation exhibited a joint dependence on position and velocity, which is characteristic of neural representations of limb motion (Ashe and Georgopoulos, 1994; Edin and Vallbo, 1990; Matthews, 1933; Paninski et al., 2004; Prochazka, 1999; Shidara et al., 1993; Wang et al., 2007). Such a dependence on both position and velocity might generate crosstalk between position-dependent and velocity-dependent adaptation, providing an explanation for the inappropriate position dependence observed.

To explore this hypothesis, we constructed a simple model simulating trial-to-trial adaptation (i.e. the viscoelastic primitive model) where *n* motor primitives in a population $\mathbf{S} \in \mathbb{R}^{n\times 2}$ receive as inputs the time-varying position and velocity traces of the previous movement. The weighted outputs of these primitives (i.e. time-varying force patterns) are combined to obtain net motor output for the next movement. More specifically, each primitive element simultaneously responds to the position and velocity of a movement with a specific gain for each component, denoted by k_i and b_i for the i^{th} primitive and can thus be represented as a point in position/velocity gain-space ($\mathbf{S}_i = k_i, b_i$, Figure

3A). Note that the distribution of these points reflects the biased joint dependence of the motor primitives on position and velocity in our model; i.e., if the i^{th} primitive responds positively or negatively to position, it will also very likely respond to velocity in the same manner. The population of these primitives is modified by a weighting vector $\mathbf{w} \in \mathbb{R}^{n\times 1}$, leading to a net motor output of $\mathbf{y} = \sum_{i=1}^{n} \mathbf{S}_i \mathbf{w}_i = \mathbf{S}^T \mathbf{w}$. Thus, if the desired motor output goal is $\mathbf{y}^* \in \mathbb{R}^{2\times 1}$, motor (or force) error can be expressed as $\mathbf{y}^* - \mathbf{y}$, the difference between the goal and net motor output (or the difference between the desired and actual force outputs). A standard gradient descent learning rule updates the weights associated with each learning element to minimize the motor output error. This update rule can be expressed as $\Delta \mathbf{w}_i = \eta \mathbf{S}_i \ \mathbf{y}^* - \mathbf{y} \ = \eta \left|\mathbf{S}_i\right| \cdot \left|\mathbf{y}^* - \mathbf{y}\right| \cdot \cos\ \theta_{\mathbf{y}^* - \mathbf{y}} - \theta_{\mathbf{S}_i}$, where η is a learning constant much smaller than 1. This learning rule represents the projection of the error vector onto each primitive (see Supplemental Information for derivation). The maximal weight change for a given primitive will occur when the error vector and that primitive are aligned with each other (i.e. the angle of difference [AOD] between the two vectors is 0°), while no weight update will occur when the two vectors are orthogonal (i.e. AOD = 90°). This is schematically illustrated in the bottom right panels of Figure 3A. The motor primitives are depicted as having a zero-mean distribution, with a positive correlation between position and velocity gains. The location of each dot in these panels signifies the gain-space representation of a primitive, and the color signifies the cosine of the AOD for that primitive, which is proportional to its weight change. Early in learning, most of the primitive elements display large non-zero cosines of the AOD, shown as colored dots. Because the majority of these cosines are large, most primitives contribute substantially to initial learning. As a result, the initial learning step is composed of contributions from nearly the entire population and reflects the distribution of primitives more strongly than motor error (i.e. the learning step is taken largely in the direction where the distribution variance is maximal). In contrast, later in learning, a much smaller and more specific fraction of primitive elements makes substantial contributions to the learning process because there are now many fewer primitives aligned with the motor error (i.e. the learning step is now in the direction where distribution variance is minimal). This development leads to slower

learning that more specifically addresses motor errors. Overall, this model predicts that early adaptation should largely reflect the properties of the motor primitives, whereas late adaptation should reflect task-specific goals.

In our model, the statistical distribution of these primitives plays a crucial role in determining the dynamics of learning. If this distribution is arbitrarily narrow, i.e. a line distribution, then every primitive element is just a scaled version of every other primitive, meaning that the only difference in their responses to the position and velocity of an arm movement will be the magnitudes – the ratios remain the same. Therefore, during learning, primitives will display weight changes such that the resultant learning is confined to the direction of maximal variance. As training progresses, the rate of learning decreases as the AOD's approach 90°, as shown in Figure 3B. For a wider distribution, a large fraction of the primitives will initially maintain large non-zero weight changes, resulting in a bias of the initial learning direction towards the direction of maximal variance. However, as learning continues, the weight changes will become more specific, such that primitives more closely aligned with the goal will undergo weight increases, while primitives farther away will experience weight reductions, leading to a rotation of the learning direction towards the learning goal in line with the error vector. As the width of the distribution increases even more, this redirection occurs earlier and more gradually (Figure 3B). In the extreme case of an unbiased distribution, the learning trajectory will follow the error vector directly to the goal. Note that the learning trajectories produced by distributions of intermediate width capture the essential features of the learning behavior observed in Experiment 1. In particular, initial learning of a velocity-dependent perturbation manifests a position-dependent component, whereas late learning veers back towards the goal, reducing this position-dependent component.

Model Predictions Based on a Correlated Primitive Distribution

Several key predictions arise from considering how learning might depend on the distribution of motor primitives. First, if subjects are exposed to a *position*-dependent perturbation, this model predicts that while initial learning will exhibit an inappropriate *velocity*-dependent cross-adaptation,

late learning will substantially reflect purely position-dependent learning (Figure 3C, blue lines). Second, the model predicts that the initial learning of position-dependent and velocity-dependent force perturbations will be quite similar to one another, despite the fact that these two perturbations are essentially orthogonal (the correlation coefficient between average position and velocity traces - and thus the correlation between the force patterns associated with the position-dependent and velocity-dependent perturbations - over a 2.25s window, centered around the peak speed point is r=0.01, indicating near-orthogonality). We test these first two predictions in Experiment 2.

The viscoelastic primitive model also predicts that exposure to a FF for which the position and velocity dependence are *positively*-correlated (i.e. represented as goals in the first and third quadrants in position/velocity gain-space; Figure 3C, green lines) would result in an especially close correspondence between the error vector and the direction of maximal variance for the primitive distribution. This would yield small AOD's that would maximize the updates of the weights associated with each primitive and produce fast adaptation. This adaptation would be rapid and closely directed towards the FF learning goal, leading to learned force patterns that are both large in magnitude and very similar in shape to the perturbing FF. In contrast, learning a FF for which the position and velocity dependence are *negatively*-correlated would result in a substantially orthogonal relationship between the motor primitive distribution and error, resulting in slow adaptation (Figure 3C, purple lines). Nonetheless, the learning here would also be very closely directed towards the FF goal, again leading to learned force patterns very similar in shape to the perturbing FF (see Supplemental Information for the derivation of these FFs). We test these predictions in Experiments 3 and 4.

Since this viscoelastic primitive model is linear with respect to motion state, the positive-combination and negative-combination directions can be thought of as the eigendirections of the learning space. As such, any viscoelastic dynamics can be broken down into constituent positive-combination and negative-combination components, and the resulting adaptation can be predicted by summing these component-specific adaptations. Note that adaptation to the positive-combination component is predicted to be much faster than adaptation to the negative-combination component. Therefore, adaptation to any FF with a component in the positive-combination direction would

include a rapid component, and would be biased towards this positive-combination direction. Thus, initial adaptation to purely position-dependent and purely velocity-dependent FFs (which include components in the positive-combination direction) are predicted to be rapid and stereotyped, as indicated by the similarity between initial learning directions for these FFs (Figure 3C). In contrast, initial adaptation to a negative-combination FF (which lacks a positive-combination component) will be slower and in an orthogonal direction in position/velocity gain-space when compared to the other three types of FFs, as shown in Figure 3C.

The viscoelastic primitive model also predicts that interference from adaptation to position-dependent FFs onto learning velocity-dependent FFs (and vice versa) will not be monolithic, but instead will depend on the transition required between the previously-learned and current dynamics, as illustrated in Figure 7AB. For example, adaptation to a CW velocity-dependent FF (+V) following exposure to a CCW position-dependent FF (-P) requires both unlearning the -P FF and learning the +V, corresponding to a change in dynamics of Δ=[+P, +V]. On the other hand, adaptation to a -V FF after exposure to a -P FF requires a learning change of Δ=[+P, -V]. Since the Δ=[+P, +V] learning changes for the -P→+V transition are positively-correlated and aligned with the direction of maximal variance for the primitive distribution, while the Δ=[+P, -V] changes are negatively-correlated, the -P→+V learning should occur at a much more rapid rate than the -P→-V learning, according to the viscoelastic primitive model (Figure 7AB). These predictions are tested in Experiment 5.

The Effects of Different Learning Rules

A key feature of this model is that a variety of learning rules yields very similar results. In Figure 3, we used a simple gradient descent learning rule in which the weights associated with motor primitives opposing the error grow while weights associated with motor primitives aggravating the error are decreased. However, recent work suggests that motor activation can increase, albeit asymmetrically, in both agonist and antagonist directions in response to errors as a mechanism for regulating co-contraction (Franklin et al., 2008; Milner and Franklin, 2005). We find that incorporating this learning rule into the viscoelastic primitive model gives essentially identical results

to the gradient descent learning rule for the patterns of net motor output (compare Figure 4B with 4A). A Bayesian approach also produces very similar behavior (Figure 4C). In a Bayesian framework, the net motor output at trial *n+1* (i.e. the posterior distribution) is a combination of the net motor output at trial *n* (prior distribution) with the learning goal (state measurement). The probability distributions which characterize the certainty of the prior and measurement, determine how these two "inputs" are combined to give the posterior estimate. In our model, if the positively-correlated population of primitives serves as the prior distribution, as depicted by the distribution ellipses in Figure 4C, then we achieve essentially the same adaptive behavior as the gradient descent and co-contraction-based learning rules (see Supplemental Information for specifics of these learning rule implementations).

It is important to note, however, that higher-order learning rules which depend on higher-order derivatives of the weights with respect to motor errors could in theory generate patterns of motor output that effectively compensate for the learning element distribution. Since such higher-order learning schemes (Battiti, 1992) can produce output that is independent of this distribution, they would *not* reproduce the predictions shown in Figure 3, as seen in Figure 4D. Hybrid learning rules which combine first-order and higher-order learning rules (Battiti, 1992) will predict behavior that is partially comparable (Figure 4E), but do so specifically because of the partial dependence on the first-order learning rule, which is sensitive to the primitive distribution. High levels of noise are pervasive in the nervous system (Faisal et al., 2008) and such noise limits the computational power of neural networks (Maass and Orponen, 1998). Therefore, it may be difficult to implement learning schemes which depend on higher-order derivatives because the accurate estimation of these derivatives is computationally expensive (Battiti, 1992) and sensitive to noise. This underscores the importance of the distribution of learning elements as the key feature in our model and, in particular, the positive correlation between how these elements respond to position and velocity. Learning rules which are influenced by this correlation predict patterns of motor output consistent with the simulations shown in Figure 3 (Figure 4A-C).

Early Learning for Position- and Velocity-Dependent Force-Fields

In Experiment 2, we examined the force output patterns after just a single trial of exposure to a position-dependent or a velocity-dependent perturbation. After an error-clamp probe trial was used to measure a baseline force pattern, subjects were exposed to a single trial of either a CW or CCW position-dependent or velocity-dependent FF followed by a second probe trial. We attributed the difference between the force patterns measured in these two probe trials (Figure 5AB) to the motor adaptation resulting from the single-trial FF exposure. We found that the learning-induced changes in force production resulting from both single-state FFs show clear evidence of positively-correlated dependence on position and velocity. For the position-dependent perturbation (Figure 5A), the learning-related force pattern cannot be fully explained by position alone. Instead, this force pattern is significantly dependent on both position and velocity ($p<10^{-12}$ in both cases; $R^2=0.72$ for position alone (gray curve), partial R^2 for velocity=0.85, $R^2=0.96$ when both position and velocity are included in the regression (black curve)). Correspondingly, for the velocity-dependent perturbation (Figure 5B), we found this learning-related force pattern to be significantly dependent on both velocity and position ($p<10^{-12}$ in both cases; $R^2=0.60$ for velocity alone (gray curve), partial R^2 for position=0.56, $R^2=0.85$ when both velocity and position are included in the regression (black curve)). Note that like the data in Figure 2, this position-dependent component cannot be explained by persistent lateral kinematic errors (Figure S3).

Interestingly, the crosstalk between position and velocity-dependent adaptation is strong enough to make the learned force patterns that result from a single-trial position-dependent perturbation and from a single-trial velocity-dependent perturbation highly correlated with one another (r=0.81), even though the temporal profiles of the force perturbations inducing them are essentially orthogonal (r=0.01 between the two perturbations). This reflects the similar initial directions of position-dependent and velocity-dependent learning trajectories, as shown in Figures 3 and 4A-C, despite the learning goals being 90° apart in position/velocity gain-space.

Early Learning for Combination Force-Fields

In Experiment 3, we used the same probe/FF/probe paradigm as in Experiment 2 to measure the single-trial learning associated with combination FFs having either a positively-correlated dependence on position and velocity, like the representation of motion state in motor cortex and muscle spindle afferents, or a negatively-correlated dependence on those two states. The viscoelastic primitive model predicts the fastest and slowest adaptation, respectively, to these FFs. However, it also predicts that adaptation to both of these FFs should be essentially free of cross-talk – i.e., the direction of learning in position/velocity gain-space should be closely aligned with the FF perturbation directions. We found that single-trial exposure to both of these combination FFs produces learned force patterns that are not only composed of significant position and velocity contributions, but are also very similar in shape to the associated FF perturbations (Figure 5C; Positive-Combination FF: partial R^2 for position=0.79, partial R^2 for velocity=0.78, R^2=0.88 when both position and velocity are included in the regression as independent variables (black curve), R^2=0.87 when the learned force pattern is regressed onto the shape of the positive-combination FF (gray curve), $p<10^{-12}$ in all cases; Negative-Combination FF: partial R^2 for position=0.68, partial R^2 for velocity=0.25, R^2=0.71 when both position and velocity are included in the regression as independent variables (black curve), R^2=0.70 when the learned force pattern is regressed onto the shape of the negative-combination FF (gray curve), $p<10^{-12}$ in all cases). Figure 5D shows a representation of the similarity between the learned force outputs and the shapes of the perturbing FFs. Here the curved arrows, which represent the misalignments of learned force patterns with the learning goals, are much smaller for the combination FFs than for the single-state FFs.

Perhaps the most striking feature of the force patterns shown in Figure 5C is that the positive-combination (PC) FF induces much greater learning than the negative-combination (NC) FF. To quantify the magnitude of learning for both FFs, we projected the gain-space representation of the force patterns onto the corresponding FF directions (Figure 5D, see Methods). We found a significant difference between groups (one-way ANOVA, $F_{3,74}$=12.6, $p=9.9\times10^{-7}$). Single-trial exposure to the positive-combination FF produced three-fold greater adaptation than exposure to the

negative-combination FF. Positive-combination adaptation was significantly greater than for all other FFs (PC/Position: $p=0.030$, PC/Velocity: $p=0.022$, PC/NC: $p=2.0\times10^{-8}$, one-sided unpaired t-test), whereas the negative-combination adaptation was significantly less than for all other FFs (NC/Position: $p=5.2\times10^{-4}$, NC/Velocity: $p=2.5\times10^{-4}$, one-sided unpaired t-test). Double-trial learning data (obtained using a probe/FF/FF/probe paradigm) for all four FFs were also consistent with the single-trial learning data and the predictions made by the viscoelastic primitive model (Figure S6).

Extended Learning for Different Types of Dynamics

In Experiment 4, we exposed additional groups of subjects to extended periods (160 trials) of the CW and CCW position-dependent, positive-combination, and negative-combination FFs to compare the associated motor adaptation with that of the velocity-dependent FF. Figure 6C depicts the learning curves for each of these FFs (adaptation index is calculated in the same way as in Experiment 3), and the inset shows the adaptation averaged over the latter half of training (trials 81-160). As in the single-trial learning experiments, we found a significant difference between groups (one-way ANOVA, $F_{3,133}=22.0$, $p=1.3\times10^{-11}$); participants learned the positive-combination FF significantly better than the other three FFs (PC/Position: $p=0.028$, PC/Velocity: $p=0.022$, PC/NC: $p=6.6\times10^{-10}$, one-sided unpaired t-test), and they learned the negative-combination FF significantly worse (NC/Position: $p=3.4\times10^{-6}$, NC/Velocity: $p=4.1\times10^{-6}$, one-sided unpaired t-test). Moreover, the gain-space trajectories shown in Figure 6D closely resemble those predicted by the viscoelastic primitive model (Figures 3C, 6B), reflecting the model's ability to characterize the interplay between the motor primitive distribution and task-specific goal. The learning trajectories associated with the position-dependent and velocity-dependent FFs are curved, reflecting the partial alignment of the task-specific goal with the direction of maximal variance in the primitive distribution, while the positive-combination and negative-combination trajectories are straighter, reflecting full or no alignment, respectively, between the direction of maximal variance, and the learning goal (Figure 6D).

Patterns of Interference Between Different Types of Dynamics

In order to test the prediction that a positively-correlated change in dynamics (Δ=[+P,+V] or Δ=[-P,-V]) should produce a faster pattern of interference than a negatively-correlated change in dynamics (Δ=[+P,-V] or Δ=[-P,+V]), we recruited 20 additional subjects (divided into two subgroups) for an additional set of experiments (Experiment 5). The experimental design is illustrated in Figure 7. These subjects were alternately exposed to position- and velocity-dependent FFs. After a baseline session of 160 null-field trials in each of two directions (90° and 270°), subgroup 1 learned: +P, +V, -P, -V, and +P FFs for 120 trials each (in that order) in the 270° direction, and +P, -V, -P, +V, and +P FFs in the 90° direction (Figure 7). As in Experiments 1 and 4, error-clamp probe trials were interspersed throughout FF exposure with a frequency of about 1 in 6 trials to measure the adaptation to each FF environment. For subgroup 1, the 2-D gain-space representations of the FF sequence in the 270° direction proceeds in a CCW fashion, while the 90° sequence proceeds in a CW fashion (Figure 7). For subgroup 2, the FF sequences for 90° and 270° movements were swapped for a balanced experimental design. This experimental design allowed us to examine each of the four possible changes in dynamics between position- and velocity-dependent FFs (Δ=[+P,+V], Δ=[-P,-V], Δ=[+P,-V], Δ=[-P,+V]) and the eight possible transitions between dynamics (+P→+V, +P→-V, -P→+V, -P→-V, +V→+P, +V→-P, -V→+P, -V→-P).

Experimental results are displayed in Figures 7CD and 8. The gain-space trajectories displayed in Figure 7CD show that transitions requiring positively-correlated changes in dynamics (such as +P→-V) have larger initial learning rates (displayed as the longer, brighter arrows with positive slopes) than those transitions requiring negatively-correlated changes (such as +P→+V, displayed as the shorter, darker arrows with negative slopes). Each individual red, brown, or black arrow depicts a gain-space representation of the learning achieved during a block of ten trials. The purple and green arrows correspond to the adaptation rate during early learning - the first 3 trials of exposure to the FF – for positively- and negatively-correlated changes in dynamics, respectively. Note that the purple arrows are consistently longer than the green arrows, indicating that positively-correlated transitions display higher learning rates than negatively-correlated transitions. Note also

that in this interference paradigm, although the initial learning rates are faster for positively-correlated changes in dynamics, the *initial* amount of interference is greater. Take, for example, the transition -P→+V, which is displayed in the upper-left of Figure 7A – it requires a change in dynamics of Δ=[+P,+V], and is thus in the positive-combination direction. Since this transition begins with a slightly negative value for the velocity-gain in the 3rd quadrant, the velocity-gain is initially biased against that which is to be learned (+V). In general, transitions with positively-correlated changes in dynamics begin at points that are initially biased against the current FF goal but display higher learning rates, whereas transitions requiring negatively-correlated changes are initially biased towards the current FF goal, but display slower learning rates.

Both of these effects are readily apparent in the learning curves displayed in Figure 8EHK. In these panels, the learning curves associated with positively-correlated transitions (plotted as thick solid lines) generally begin lower but end higher than the curves associated with negatively-correlated transitions (plotted as dotted lines). These curves, along with Figure 8B, are calculated and binned as in Experiment 4. We calculated normalized adaptation curves (third column of Figure 8, see Methods) in order to make a fair comparison between the learning rates displayed by the various transitions in this interference paradigm. These curves (Figure 8FIL) show that the transitions requiring positively-correlated changes in dynamics (thick solid lines) display faster normalized adaptation rates than their negatively-correlated counterparts (dotted lines), corresponding to lower overall interference. The size of this effect is as great or greater than the overall interference between velocity- and position-dependent dynamics, as shown in Figure 8A-C.

For interference of velocity-dependent dynamics onto position-dependent dynamics, we found that the negatively-correlated transitions displayed slower initial learning rates (Figure 8D; one-sided unpaired t-test, $p=0.0059$) and slower overall learning (Figure 8F; one-sided unpaired t-test, $p=0.0073$). For interference of position-dependent dynamics onto velocity-dependent dynamics, we also found that the negatively-correlated transitions displayed slower initial learning rates (Figure 8G; one-sided unpaired t-test, $p=0.002$) and slower overall learning (Figure 8F; one-sided unpaired t-test, $p=8.1\times10^{-9}$). When all the data is combined (Figure 8J-L), the differences between the patterns of interference for positively-correlated vs. negatively-correlated transitions are even clearer – the initial

learning rates for the positively-correlated transitions are 73% higher on average (Figure 8J; one-sided unpaired t-test, $p=9.4\times10^{-5}$), and the overall learning rate is 25% higher on average (Figure 8L; one-sided unpaired t-test, $p=1.3\times10^{-7}$).

Discussion

Here we have shown that neural representations of body motion are tightly coupled with key features of motor adaptation. In particular, the changes in motor output during the course of adaptation can be explained by a population of motor primitives which exhibit positively-correlated responses to position and velocity, as do neurons in primary motor cortex and muscle spindle afferents. This positive correlation predicts a specific pattern of crosstalk during adaptation to different classes of dynamics that is borne out in our data. In particular, purely position-dependent and purely velocity-dependent dynamics induce adaptation with both appropriate and inappropriate components, whereas combination FFs induce purely appropriate adaptation. This model also explains a general principle of learning: that early adaptation tends to be rapid and stereotyped, whereas late adaptation is slower and more task-specific. Interestingly, the viscoelastic primitive model predicts the rate at which different classes of dynamics can be learned; specifically, that positive-combination dynamics are intrinsically easiest to learn, while negative-combination dynamics are hardest to learn. In keeping with this prediction, we find a greater than three-fold difference in initial adaptation rates for these classes of dynamics. Furthermore, the model predicts that the patterns of interference between adaptations to position-dependent and velocity-dependent FFs (and vice versa) are not monolithic, but instead depend on whether the transition between consecutively-learned FFs requires a positively-correlated vs. negatively-correlated change in dynamics in the position/velocity gain-space.

Primitives Depend on a Combination of Position and Velocity

Previous work attempting to characterize the primitives underlying motor adaptation has shown that these primitives respond to motion state rather than time (Conditt et al., 1997; Conditt and

Mussa-Ivaldi, 1999; Diedrichsen et al., 2007). However, prior investigations on the primitives for motor adaptation have focused on single state variables rather than on combinations of them. For example, studies have found local tuning of motor primitives in velocity-space, although the precise width of this tuning has remained somewhat controversial (Donchin et al., 2003; Gandolfo et al., 1996; Krakauer et al., 2000; Thoroughman and Shadmehr, 2000). Another line of work has focused on comparing the ability of the motor system to learn different single-state dynamics. This work has revealed that motor primitives are more sensitive to position and velocity than to acceleration because novel acceleration-dependent dynamics are more poorly learned and generalized than other novel dynamics (Hwang and Shadmehr, 2005; Hwang et al., 2006). However, studies focusing on single-state representations of dynamics are unable to account for how the temporal structure of motor output might evolve during adaptation. Indeed, it is only by considering the motor primitives' joint dependence on multiple states (i.e. position and velocity) that we are able to gain insight into this evolution.

Interference Between Adaptations to Different State-Dependent Dynamics

We were able to uncover a complex pattern of interference between different state-dependent adaptations by examining all eight possible transitions between position-dependent and velocity-dependent adaptations. Some transitions displayed high interference, and some displayed low interference. The viscoelastic primitive model accounts for this pattern of interference as well as previous work that examined interference between tasks based on the same kinematic variable (i.e. position-dependent visuomotor rotation and position-dependent FF) (Tong et al., 2002), or that only studied the low-interference, +P→-V transition without examining the high-interference +P→+V transition (Bays et al., 2005) (for further discussion on Bays et al, 2005, see Supplemental Information). A key observation about the patterns of interference is that transitions with positively-correlated changes in dynamics have greater initial force error than the transitions with negatively-correlated changes, as discussed in the Results section. This finding that larger initial force errors are not synonymous with slower learning rates suggests that the practice of defining interference as higher initial errors at task transition should perhaps be reconsidered (Brashers-Krug et al., 1996;

Donchin et al., 2002; Lee and Schweighofer, 2009; Miall et al., 2004; Shadmehr and Brashers-Krug, 1997).

Muscle Synergies

Several studies investigating postural control, force production, and movement generation have found that focal stimulation of the spinal cord elicits activation of specific muscle combinations, termed synergies (Kargo and Giszter, 2000; Mussa-Ivaldi et al., 1994; Tresch and Bizzi, 1999). These synergies are also observed across different natural behaviors in animals (Bizzi et al., 2008; d'Avella et al., 2003; Ting and Macpherson, 2005). It has been hypothesized that these synergies constitute the functional building blocks for these behaviors, in the sense that motor activation patterns during a variety of behaviors are formed by different combinations of a small number of shared synergies (Bizzi et al., 2008; d'Avella et al., 2003; Mussa-Ivaldi and Bizzi, 2000; Mussa-Ivaldi et al., 1994). However, it is not known to what extent changes in motor output resulting from motor adaptation can be attributed to changes in the organization of the synergies themselves versus changes in the relative activations of different, existent synergies (Ting and McKay, 2007).

Interestingly, focal stimulation of the spinal cord, which is thought to activate a specific muscle synergy, produces motor output which is known to depend strongly on the position of the limb (Kargo and Giszter, 2000; Mussa-Ivaldi et al., 1994; Tresch and Bizzi, 1999). However, it is unclear how this motor output depends on other components of motion state, particularly velocity. Although our current work indicates that motor learning, as well as the tuning of motor cortical cells (Ashe and Georgopoulos, 1994; Paninski et al., 2004; Wang et al., 2007) and muscle spindle afferents (Edin and Vallbo, 1990; Matthews, 1933; Prochazka, 1999), is driven by a joint dependence on position and velocity, such dependence has not been observed for muscle synergies because the velocity-dependence of the motor output of these muscle synergies has not been characterized. Nonetheless, it is likely that these spinal-level muscle synergies are based on adaptive, supraspinal motor primitives. If so, one would expect that the motor output patterns associated with these synergies would often depend on position and velocity in a correlated fashion given that the majority of the motor primitives do so. If this were true, then perhaps adaptation requiring positively-correlated changes in dynamics

of position- and velocity-dependence could be achieved by varying the activation of different synergies, whereas adaptation requiring negatively-correlated changes in dynamics could require changes in the synergies themselves. Thus, the rapid adaptation we see in the positive-combination direction may correspond to re-weighting the synergies contributing to motor output, while the slower adaptation in the negative-combination direction may correspond to a reorganization of synergy composition, or even the formation of a new synergy.

Just as muscle synergies may undergo reorganization, it is possible that the state-dependence of the motor primitive population may also be reshaped. Our current data suggest that if such reshaping occurs, it operates on a timescale that is slower than 240 trials because it appears that a single fixed distribution explains both the single-trial and extended learning data presented in this study.

Early Learning Tends to Be Non-Specific

A key feature of motor adaptation is that early learning of different dynamics tends to be non-specific. This explains both the present finding that single-trial exposure to position-dependent and velocity-dependent FFs produces remarkably similar force patterns, and previous work showing little difference in the single-trial adaptive responses to force pulses administered over a range of different positions in reaching arm movements (Fine and Thoroughman, 2006). Interestingly, the viscoelastic primitive model predicts not only the stereotypic shape of the adaptive responses observed in this study, but their magnitudes as well (Figure S12E). This underscores the idea that the motor system adapts much more strongly to state-dependent dynamics than non-state-dependent dynamics because although the shape of the narrow, 70ms force pulses used in this study *cannot* be well-characterized by position-dependent and velocity-dependent components, the adaptation resulting from exposure to these force pulses can. Details are given in the Supplemental Information.

The Ability to Learn Different Dynamic Environments Can Be Predicted

Our viscoelastic primitive model not only gives us insight into the evolution of motor adaptation, but it also allows us to predict how quickly and to what extent any dynamics composed of

elastic and viscous components can be learned, and how different viscoelastic dynamics will interfere with one another. In the current paper, we verified these predictions by testing the most extreme combinations of the elastic and viscous components. Prediction of the "learnability" of a dynamic environment can now be based on an understanding of the intrinsic relationship between the dynamics of the environment and the motor primitive distribution, rather than solely on interference from other learning experiences, as in previous work (Bays et al., 2005; Hwang et al., 2003).

The ability to create the easiest possible motor learning task may be of practical value for motor rehabilitation. For example, the difficulty that stroke victims display in recovering motor output appropriate for common tasks results from reduced rates of motor learning for the affected limb. This difficulty is unfortunately compounded by learned nonuse of the affected limb that results from the motor disability and corresponding learning deficit (Han et al., 2008; Taub et al., 2002). Training with the easiest possible learning task may help reduce this learned nonuse and speed recovery. In general, a better understanding of motor learning will allow for the rational design of improved procedures to facilitate motor training and rehabilitation. The tight coupling that we observe between the neural representations of body motion and key features of motor learning suggests that greater understanding of neural representations will improve our understanding of learning and vice-versa.

Materials and Methods

All experiment participants gave informed consent, and experimental protocols were approved by the Harvard University Institutional Review Board. Note that a more detailed version of these methods is available in the Supplementary Information.

In all experiments, subjects were instructed to grasp the handle of a robotic manipulandum while making 500ms, 10cm straight, reaching arm movements in the 90˚ and 270˚ directions (Figure 1A); however, all 90° movements were error clamp trials, and only 270° movements were analyzed. As the subjects moved the manipulandum, we applied perturbative FF environments (Figure 1B) of the form:

$$\begin{bmatrix} \mathbf{F}_x \\ \mathbf{F}_y \end{bmatrix} = \mathbf{K} \cdot \mathbf{p} + \mathbf{B} \cdot \mathbf{v} = \begin{bmatrix} 0 & -K \\ K & 0 \end{bmatrix} \cdot \begin{bmatrix} x \\ y \end{bmatrix} + \begin{bmatrix} 0 & -B \\ B & 0 \end{bmatrix} \cdot \begin{bmatrix} \dot{x} \\ \dot{y} \end{bmatrix} \quad (1)$$

where the eight different FFs (CW or CCW versions of the four FF types) were: Position-Dependent: K=±45N/m, B=0Ns/m; Velocity-Dependent: K=0N/m, B=±15Ns/m; Positive-Combination (PC): K=±21.2N/m, B=±13.2Ns/m; Negative-Combination (NC): K=$\mp$35N/m, B=±9.4Ns/m. To move in a straight line, subjects needed to produce compensatory forces that were equal and opposite to the robot-produced forces. Learning was quantified with a metric found by projecting the gain-space representation of each learned force pattern, measured using error-clamp probe trials, onto the associated gain-space representation of the perturbing FF.

When plotting the gain-space representation of learning, we normalized the axes such that the velocity-dependent FF of $15\frac{N \cdot s}{m}$ and a position-dependent FF of $45\frac{N}{m}$ were represented by [0,1] and [1,0], respectively. These values were chosen to equate peak perturbing force between these two types of FFs (see Supplemental Information).

Experiment 1

We instructed 93 subjects to make 160 null-field baseline trials (i.e. trials where no FF was applied) to familiarize themselves with the task. We then applied velocity-dependent viscous curl fields (Figure 1B) represented by Equation 1 (K=0N/m, B=±15Ns/m), during which we gave error-clamp trials 20% of the time in order to measure the force output. As the training progressed, subjects consistently improved their performance on this dynamic task (Figure 2A).

Viscoelastic Primitive Model

For Figures 3 and 6, we implemented the zero-mean, positively-correlated primitive distribution as a population of 5000 primitives, where this population had identical variances for its sensitivity to position and velocity signals, with a correlation of 0.8 between the two variances. The first-order gradient descent learning rule is given by:

$$\Delta \mathbf{w}_i = \eta \mathbf{S}_i \ \ \mathbf{error} \ = \eta \left| \mathbf{S}_i \right| \cdot \left| \mathbf{error} \right| \cdot \cos \ \theta_{\mathbf{error}} - \theta_{\mathbf{S}_i} \tag{2}$$

where $\Delta \mathbf{w}_i$ is the change in the weight for the i^{th} primitive, η is the learning rate, $\mathbf{S}_i$ represents the position and velocity sensitivity of the i^{th} primitive, and $\theta_{\mathbf{error}}$ is the angle for the error vector.

For Figure 4, we implemented all of the learning rules with a zero-mean, positively-correlated primitive distribution of 350 primitives with identical population characteristics to that described above. The first-order gradient descent learning rule applied was identical, as well. We implemented the co-contraction learning rule as:

$$\Delta \mathbf{w}_i = \begin{cases} \eta_1 \mathbf{S}_i \cdot \mathbf{error} - c & \mathbf{S}_i \cdot \mathbf{error} > 0 \\ \eta_2 \mathbf{S}_i \cdot \mathbf{error} - c & \mathbf{S}_i \cdot \mathbf{error} \le 0 \end{cases} \tag{3}$$

where $\eta_1 > 0$, $\eta_2 < 0$, $\left|\eta_1\right| > \left|\eta_2\right|$, and $c > 0$. We implemented the Bayesian learning rule as:

$$\begin{bmatrix} \mu_{Post,x} \\ \mu_{Post,y} \end{bmatrix} = \ \Sigma_{Prior}^{-1} + \Sigma_M^{-1} \ ^{-1} \Sigma_M^{-1} \begin{bmatrix} \mu_{Mx} \\ \mu_{My} \end{bmatrix} + \ \Sigma_{Prior}^{-1} + \Sigma_M^{-1} \ ^{-1} \Sigma_{Prior}^{-1} \begin{bmatrix} \mu_{Prior} \\ \mu_{Prior} \end{bmatrix} \tag{4}$$

$$\Sigma_{Post} = \ \Sigma_{Prior}^{-1} + \Sigma_M^{-1} \ ^{-1} \tag{5}$$

where $\begin{bmatrix} \mu_{Post,x} \\ \mu_{Post,y} \end{bmatrix}$ and Σ_{Post} are the mean and covariance of the posterior, $\begin{bmatrix} \mu_{Mx} \\ \mu_{My} \end{bmatrix}$ and Σ_M are the mean and covariance of the measurement distribution, and $\begin{bmatrix} \mu_{Prior} \\ \mu_{Prior} \end{bmatrix}$ and Σ_{Prior} are the mean and covariance matrix of the prior.

We implemented the "pure" second-order learning rule as:

$$\Delta \mathbf{w} = -\eta \ \ \nabla^2 \mathbf{error} \ ^{-1} \ \ \nabla \mathbf{error} \tag{6}$$

where η is the learning rate, $\nabla \mathbf{error} = \mathbf{S} \ \ \mathbf{S}^T \mathbf{w} - \mathbf{y}^*$ is the gradient, and $\nabla^2 \mathbf{error} = \mathbf{S}\mathbf{S}^T$ is the Hessian matrix. We implemented the second-order gradient descent learning rule as:

$$\Delta \mathbf{w} = -\eta k \hat{\mathbf{g}} = -\eta \left(\frac{\hat{\mathbf{g}}^T \ \ \nabla \mathbf{error}}{\hat{\mathbf{g}}^T \ \ \nabla^2 \mathbf{error} \ \ \hat{\mathbf{g}}} \right) \hat{\mathbf{g}} \tag{7}$$

where k is the step size determined by the Hessian, and $\hat{\mathbf{g}} = \nabla\mathbf{error} / \|\nabla\mathbf{error}\|$ is the unit vector pointing in the gradient direction. See Supplemental Information for a detailed explanation of these learning rules.

Experiment 2

To observe the learning after a single trial of exposure to either a position-dependent or velocity-dependent FF, we had 39 subjects (Position-Dependent: 23 subjects; Velocity-Dependent: 16 subjects) perform sets of 90° and 270° movements; however, all 90° were error-clamps, and only the 270° direction movements were analyzed. After subjects made 200 null-trial movements, they repeated 3-trial movement groups with 2-5 null-field washout trials between groups in order to estimate the amount of single-trial learning. In particular, the first and last movements in the groups were error-clamp probe trials, used to obtain the subject's force output before and after the sandwiched movements, which were either CW or CCW position-dependent or velocity-dependent FF trials. Representative force patterns were found for each subject by averaging together 24 and 9 repetitions of these movement groups for the single-trial learning, respectively. We characterized the position and velocity contributions of these subject-specific force patterns, and took the mean across subjects to find the average gain-space representation of the single-trial (Figure 5). The force patterns were also combined to find the average lateral force patterns produced after single-trial for either a position-dependent or velocity-dependent FF.

Experiment 3

See Supplemental Information for how the exact values of the PC and NC FFs were determined. To observe the learning after a single trial of either a PC or NC FF, we had 39 subjects (PC: 23 subjects; NC: 16 subjects) repeat 3-trial movement groups in the 270° direction, mimicking the paradigm followed in Experiment 2.

Experiment 4

After a 200-trial baseline period, we instructed subjects to learn three FFs (CW and CCW position-dependent, PC, and NC) for 160 trials in the 270˚ direction (Position-Dependent: 34 subjects; PC: 16 subjects; NC: 16 subjects). Error-clamp probe trials were interspersed during exposure to the FF with a frequency of 20%. The force outputs observed during these probe trials were averaged together into 20-trial bins for each subject. We characterized the position and velocity contributions to these subject-specific force patterns, and took the mean of these contributions across subjects to find the gain-space representation for each bin (Figure 6D). To compare the learning of each FF, we projected each binned point in gain-space onto the vector representing the FF target (Figure 6C).

Experiment 5

20 subjects participated in the interference experiments. The different FFs used in this experiment were of the same form as Equation 1, with K=±45N/m, B=0Ns/m for the position-dependent FFs, and K=0N/m, B=±15 Ns/m for the velocity-dependent FFs. The raw adaptation curves were normalized by the following equation:

$$LC_{normal} = \frac{LC_{raw} - baseline}{1 - baseline} \tag{8}$$

This normalization allowed us to compare learning curves that had different baselines, such as the curves associated with -P→+V and -P→-V (Figure 8H).

The gain-space trajectories in Figure 7CD are obtained by averaging data across movement directions and subjects. Half of the subjects learned the CW sequence in the 270° movement direction, along with the CCW sequence in the 90° direction, while the other half of subjects did the reverse.

Acknowledgments

We wish to thank Mark Wagner and Emery Brown for helpful discussions. This work was supported by grants from the McKnight Endowment for Neuroscience, the Alfred P. Sloan Foundation, and the Wallace H. Coulter Foundation to MAS.

Competing Interests

The authors have declared that no competing interests exist.

References

1. Ashe, J., and Georgopoulos, A.P. (1994). Movement Parameters and Neural Activity in Motor Cortex and Area 5. Cereb. Cortex *4*, 590-600.
2. Atick, J.J. (1992). Could information theory provide an ecological theory of sensory processing? Network-Comp. Neural. *3*, 213-251.
3. Barlow, H. (2001). The exploitation of regularities in the environment by the brain. Behav. Brain Sci. *24*, 602-607.
4. Barlow, H., and Foldiak, P. (1989). Adaptation and decorrelation in the cortex (Workingham, England: Addison-Wesley).
5. Battiti, R. (1992). First- and Second-Order Methods for Learning: Between Steepest Descent and Newton's Method. Neural Comput *4*, 141-166.
6. Bays, P.M., Flanagan, J.R., and Wolpert, D.M. (2005). Interference between velocity-dependent and position-dependent force-fields indicates that tasks depending on different kinematic parameters compete for motor working memory. Exp. Brain Res. *163*, 400-405.
7. Bays, P.M., and Wolpert, D.M. (2006). Actions and consequences in bimanual interaction are represented in different coordinate systems. J Neurosci *26*, 7121-7126.
8. Bhushan, N., and Shadmehr, R. (1999). Computational nature of human adaptive control during learning of reaching movements in force fields. Biol Cybern *81*, 39-60.
9. Bizzi, E., Cheung, V.C.K., d'Avella, A., Saltiel, P., and Tresch, M.C. (2008). Combining modules for movement. Brain Res Rev *57*, 125-133.
10. Brashers-Krug, T., Shadmehr, R., and Bizzi, E. (1996). Consolidation in Human Motor Memory. Nature *382*, 252-255.
11. Conditt, M.A., Gandolfo, F., and Mussa-Ivaldi, F.A. (1997). The Motor System Does Not Learn the Dynamics of the Arm by Rote Memorization of Past Experience. J. Neurophysiol. *78*, 554-560.
12. Conditt, M.A., and Mussa-Ivaldi, F.A. (1999). Central representation of time during motor learning. Proc. Natl. Acad. Sci. USA *96*, 11625-11630.
13. Cordo, P.J. (1987). Mechanisms controlling accurate changes in elbow torque in humans. J Neurosci *7*, 432-442.
14. Cordo, P.J. (1990). Kinesthetic Control of a Multijoint Movement Sequence. J. Neurophysiol. *63*, 161-172.
15. d'Avella, A., Saltiel, P., and Bizzi, E. (2003). Combinations of muscle synergies in the construction of a natural motor behavior. Nat. Neurosci. *6*, 300-308.
16. Diedrichsen, J., Criscimagna-Hemminger, S.E., and Shadmehr, R. (2007). Dissociating Timing and Coordination as Functions of the Cerebellum. J. Neurosci. *27*, 6291-6301.
17. Donchin, O., Francis, J.T., and Shadmehr, R. (2003). Quantifying Generalization from Trial-by-Trial Behavior of Adaptive Systems that Learn with Basis Functions: Theory and Experiments in Human Motor Control. J. Neurosci. *23*, 9032-9045.
18. Donchin, O., Sawaki, L., Madupu, G., Cohen, L.G., and Shadmehr, R. (2002). Mechanisms influencing acquisition and recall of motor memories. J Neurophysiol *88*, 2114-2123.
19. Edin, B.B., and Vallbo, Å.B. (1990). Dynamic Response of Human Muscle Spindle Afferents to Stretch. J. Neurophysiol. *63*, 1297-1306.
20. Faisal, A.A., Selen, L.P.J., and Wolpert, D.M. (2008). Noise in the nervous system. Nat. Rev. Neurosci. *9*, 292-303.
21. Fine, M.S., and Thoroughman, K.A. (2006). Motor Adaptation to Single Force Pulses: Sensitive to Direction but Insensitive to Within-Movement Pulse Placement and Magnitude. J. Neurophysiol. *96*, 710-720.
22. Franklin, D.W., Burdet, E., Tee, K.P., Osu, R., Chew, C.-M., Milner, T.E., and Kawato, M. (2008). CNS learns stable, accurate, and efficient movements using a simple algorithm. J. Neurosci. *28*, 11165-11173.

23. Gandolfo, F., Mussa-Ivaldi, F.A., and Bizzi, E. (1996). Motor learning by field approximation. Proc. Natl. Acad. Sci. USA *93*, 3843-3846.
24. Goodbody, S.J., and Wolpert, D.M. (1998). Temporal and Amplitude Generalization in Motor Learning. J. Neurophysiol. *79*, 1825-1838.
25. Han, C.E., Arbib, M.A., and Schweighofer, N. (2008). Stroke Rehabilitation Reaches a Threshold. PLoS Comput. Biol. *4*, e1000133.
26. Hollerbach, M.J., and Flash, T. (1982). Dynamic interactions between limb segments during planar arm movement. Biol Cybern *44*, 67-77.
27. Hwang, E.J., Donchin, O., Smith, M.A., and Shadmehr, R. (2003). A Gain-Field Encoding of Limb Position and Velocity in the Internal Model of Arm Dynamics. PLoS Biol. *1*, 209-220.
28. Hwang, E.J., and Shadmehr, R. (2005). Internal models of limb dynamics and the encoding of limb state. J. Neural. Eng. *2*, S266-S278.
29. Hwang, E.J., Smith, M.A., and Shadmehr, R. (2006). Adaptation and generalization in acceleration-dependent force fields. Exp. Brain Res. *169*, 496-506.
30. Joiner, W.M., and Smith, M.A. (2008). Long-Term Retention Explained by a Model of Short-Term Learning in the Adaptive Control of Reaching. J. Neurophysiol. *100*, 2948-2955.
31. Kargo, W.J., and Giszter, S.F. (2000). Rapid correction of aimed movements by summation of force-field primitives. J. Neurosci. *20*, 409-426.
32. Kawato, M. (1999). Internal models for motor control and trajectory planning. Curr Opin Neurobiol *9*, 718-727.
33. Krakauer, J.W., Ghilardi, M.-F., and Ghez, C. (1999). Independent learning of internal models for kinematic and dynamic control of reaching. Nat. Neurosci. *2*, 1026-1031.
34. Krakauer, J.W., Pine, Z.M., Ghilardi, M.-F., and Ghez, C. (2000). Learning of Visuomotor Transformations for Vectorial Planning of Reaching Trajectories. J. Neurosci. *20*, 8916-8924.
35. Lackner, J.R., and Dizio, P. (1994). Rapid adaptation to Coriolis force perturbations of arm trajectory. J Neurophysiol *72*, 299-313.
36. Lee, J.-Y., and Schweighofer, N. (2009). Dual Adaptation Supports a Parallel Architecture of Motor Memory. Journal of Neuroscience *29*, 10396-10404.
37. Maass, W., and Orponen, P. (1998). On the effect of analog noise in discrete-time analog computations. Neural Comput *10*, 1071-1095.
38. Matthews, B.H.C. (1933). Nerve endings in mammalian muscle. J. Physiol. *78*, 1-53.
39. Miall, R.C., Jenkinson, N., and Kulkarni, K. (2004). Adaptation to rotated visual feedback: a re-examination of motor interference. Exp. Brain Res. *154*, 201-210.
40. Miall, R.C., Weir, D.J., Wolpert, D.M., and Stein, J.F. (1993). Is the Cerebellum a Smith Predictor? J Motor Beh *25*, 203-216.
41. Milner, T.E., and Franklin, D.W. (2005). Impedance control and internal model use during the initial stage of adaptation to novel dynamics in humans. J. Physiol. *567*, 651-664.
42. Mussa-Ivaldi, F.A., and Bizzi, E. (2000). Motor learning through the combination of primitives. Philos T Roy Soc B *355*, 1755-1769.
43. Mussa-Ivaldi, F.A., Giszter, S.F., and Bizzi, E. (1994). Linear combinations of primitives in vertebrate motor control. Proc. Natl. Acad. Sci. USA *91*, 7534-7538.
44. Paninski, L., Fellows, M.R., Hatsopoulos, N.G., and Donoghue, J.P. (2004). Spatiotemporal Tuning of Motor Cortical Neurons for Hand Position and Velocity. J. Neurophysiol. *91*, 515-532.
45. Poggio, T., and Bizzi, E. (2004). Generalization in vision and motor control. Nature *431*, 768-774.
46. Prochazka, A. (1999). Quantifying proprioception. Prog. Brain Res. *123*, 133-142.
47. Scheidt, R.A., Reinkensmeyer, D.J., Conditt, M.A., Rymer, W.Z., and Mussa-Ivaldi, F.A. (2000). Persistence of Motor Adaptation During Constrained, Multi-Joint Arm Movements. J. Neurophysiol. *84*, 853.
48. Shadmehr, R., and Brashers-Krug, T. (1997). Functional Stages in the Formation of Human Long-Term Motor Memory. J. Neurosci. *17*, 409-419.
49. Shadmehr, R., and Mussa-Ivaldi, F.A. (1994). Adaptive Representation of Dynamics during Learning of a Motor Task. J. Neurosci. *14*, 3208-3224.
50. Shidara, M., Kawano, K., Gomi, H., and Kawato, M. (1993). Inverse-dynamics model eye movement control by Purkinje cells in the cerebellum. Nature *365*, 50-52.
51. Smith, E.C., and Lewicki, M.S. (2006). Efficient Auditory Coding. Nature *439*, 978-982.
52. Smith, M.A., Ghazizadeh, A., and Shadmehr, R. (2006). Interacting Adaptive Processes with Different Timescales Underlie Short-Term Motor Learning. PLoS Biol. *4*, e179.

53. Smith, M.A., and Shadmehr, R. (2005). Intact ability to learn internal models of arm dynamics in Huntington's disease but not cerebellar degeneration. J. Neurophysiol. *93*, 2809-2821.
54. Taub, E., Uswatte, G., and Elbert, T. (2002). New treatments in neurorehabilitation founded on basic research. Nat. Rev. Neurosci. *3*, 228-236.
55. Thoroughman, K.A., and Shadmehr, R. (2000). Learning of action through adaptive combination of motor primitives. Nature *407*, 742-747.
56. Ting, L.H., and Macpherson, J.M. (2005). A limited set of muscle synergies for force control during a postural task. J. Neurophysiol. *93*, 609-613.
57. Ting, L.H., and McKay, J.L. (2007). Neuromechanics of muscle synergies for posture and movement. Curr Opin Neurobiol *17*, 622-628.
58. Tong, C., Wolpert, D.M., and Flanagan, J.R. (2002). Kinematics and dynamics are not represented independently in motor working memory: evidence from an interference study. J Neurosci *22*, 1108-1113.
59. Tresch, M.C., and Bizzi, E. (1999). Responses to spinal microstimulation in the chronically spinalized rat and their relationship to spinal systems activated by low threshold cutaneous stimulation. Exp. Brain Res. *129*, 401-146.
60. Tseng, Y.W., Diedrichsen, J., Krakauer, J.W., Shadmehr, R., and Bastian, A.J. (2007). Sensory prediction errors drive cerebellum-dependent adaptation of reaching. J Neurophysiol *98*, 54-62.
61. Wagner, M.J., and Smith, M.A. (2008). Shared internal models for feedforward and feedback control. J. Neurosci. *28*, 10663-10673.
62. Wang, W., Chan, S.S., Heldman, D.A., and Moran, D.W. (2007). Motor Cortical Representation of Position and Velocity During Reaching. J. Neurophysiol. *97*, 4258-4270.
63. Wolpert, D.M., and Ghahramani, Z. (2000). Computational principles of movement neuroscience. Nat. Neurosci. *3 Suppl*, 1212-1217.
64. Wolpert, D.M., and Miall, R.C. (1996). Forward Models for Physiological Motor Control. Neural Netw *9*, 1265-1279.
65. Wolpert, D.M., Miall, R.C., and Kawato, M. (1998). Internal models in the cerebellum. Trends Cogn Sci *2*, 338-347.

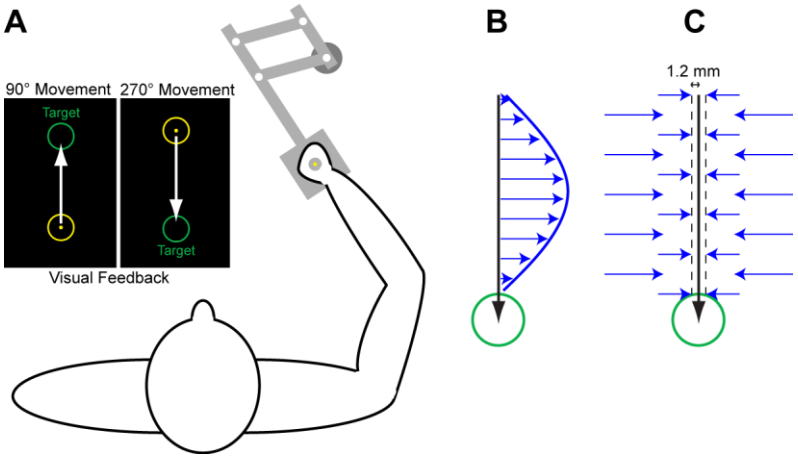

Figure 1. Illustration of Experimental Paradigm.

(**A**) Task diagram. Subjects were instructed to grasp the handle a robotic arm and make arm-reaching movements in the 90° and 270° directions. The position of their hand was displayed by a cursor on the screen. They were asked to move the cursor into targets that appeared within 450-550ms.
(**B**) Illustration of a FF trial, where lateral force applied by the robot is proportional to the movement speed. Note that this is the viscous curl-field represented by Equation 1 and (K=0N/m, B=±15Ns/m). Movements in the negative y-direction will generate perturbative forces in the positive x-direction.
(**C**) Illustration of an error-clamp probe trial. As subjects make a movement toward the target (green circle), the robotic arm applies a damped spring force (blue arrows) to the subject's hand when any lateral deviation from the midline (solid black line) occurs. This effectively restricts the subjects' movements to the confines of a virtual "channel" (bounded by black dotted lines; 99% of lateral displacements are 1.2mm or less).

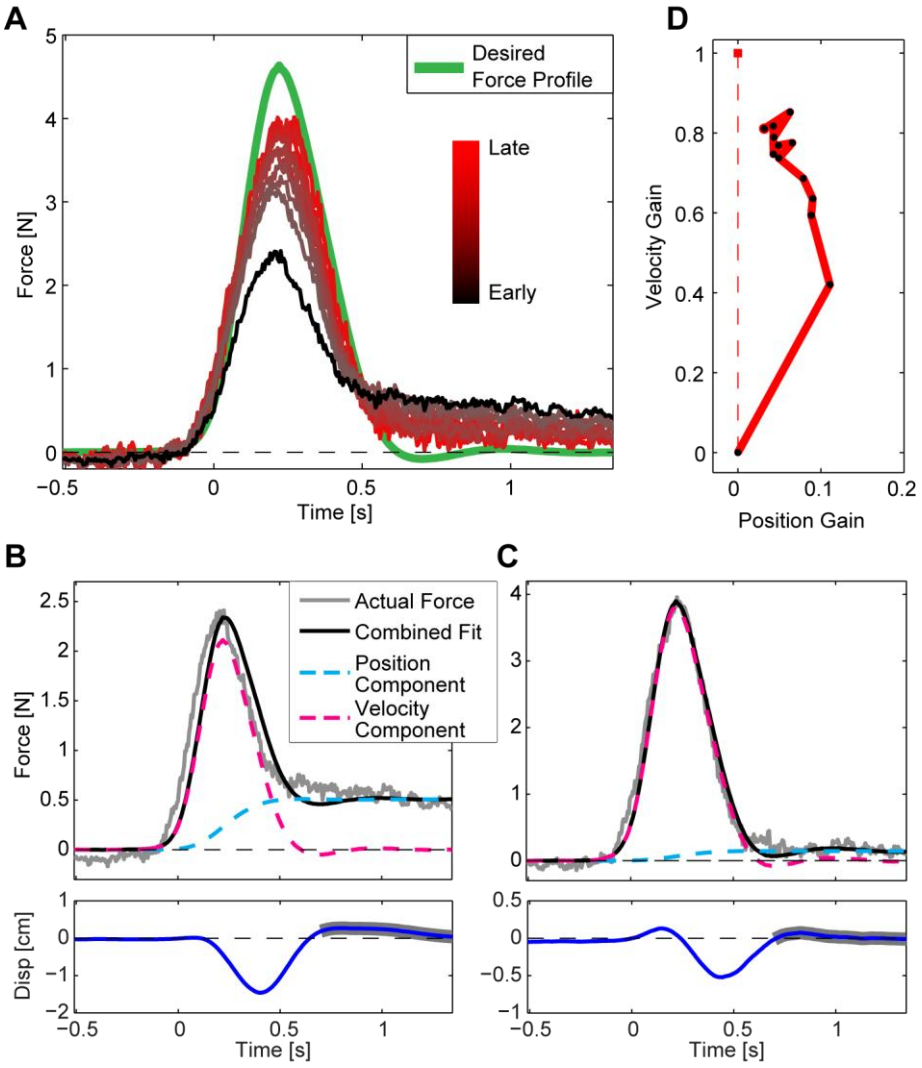

Figure 2. Motor Adaptation to a Velocity-Dependent Force–Field Exhibits Position and Velocity Dependence.

(**A**) Progression of lateral force output during the learning of a velocity-dependent FF. Subjects need to produce the green, desired force pattern to move in a straight line. Each force pattern shown is the average of the learned force patterns measured for each subject in a 20-trial bin.

(**B-C**) The average force patterns (gray) measured during the first (B) or last (C) bin of learning can be well-approximated with a curve (black) that contains significant contributions from position (cyan) and velocity (magenta). The lateral displacements of the FF trials within each bin were at or below zero just 700ms after movement onset (highlighted regions in bottom graphs), but the end-movement force persisted.

(**D**) The gains associated with linear regressions of the force patterns shown in A onto the corresponding position and velocity traces are plotted in a position/velocity gain-space (black dots). Note the deviation of this gain-space trajectory into the first quadrant. The red square on the y-axis represents the velocity-dependent FF goal.

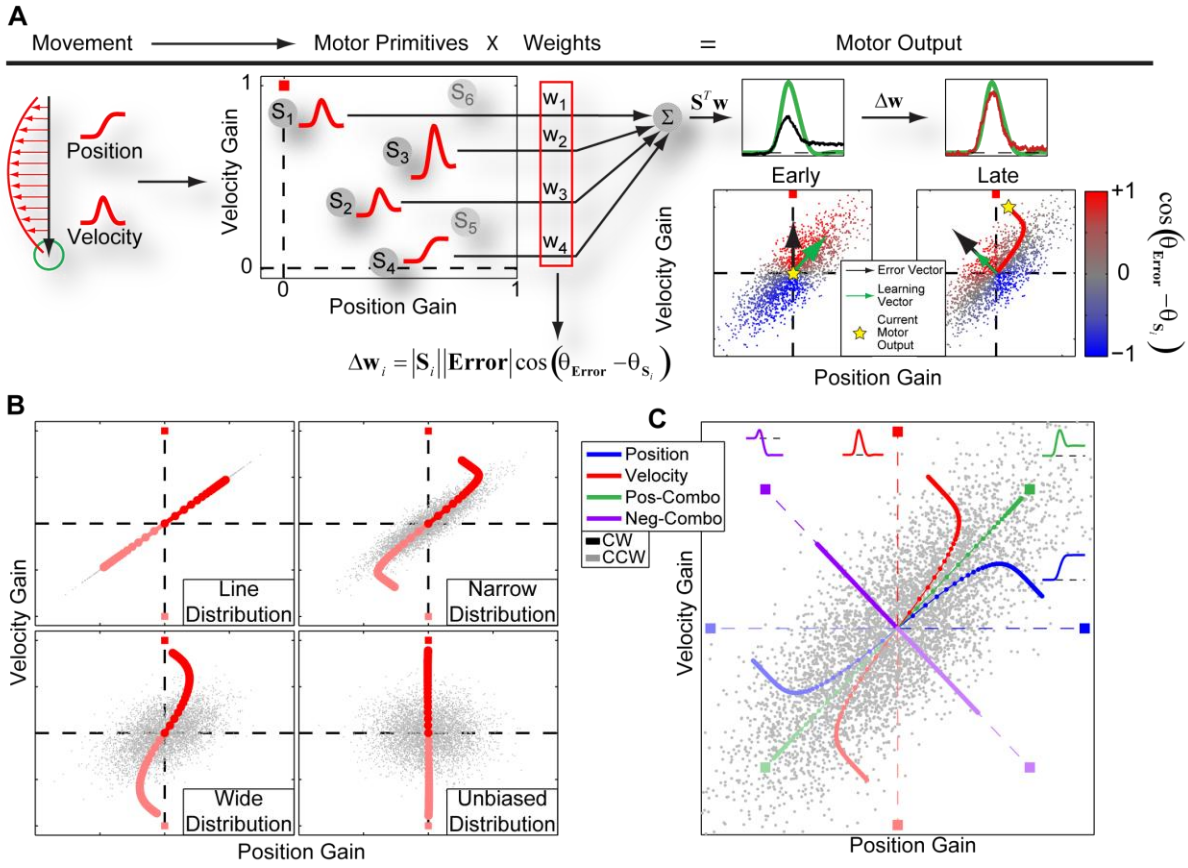

Figure 3. Viscoelastic Primitive Model.

(**A**) Schematic of the viscoelastic primitive model. Subjects make an arm-reaching movement in a perturbative FF environment (red arrows) towards a green target. The motor learning primitives (S_i, gray dots) process the position and velocity of that movement with different gains, as indicated by their locations in position/velocity gain-space. Each primitive's contribution is weighted by some value w_i, and is then summed with the contributions of the other primitives to find the overall motor output (force patterns from Figure 2A). Each primitive's weight is modified by a simple learning rule proportional to the cosine of the angle between the error and primitive vectors. Early in learning, the majority of the primitives make substantial positive contributions to the motor output (red-colored dots, "Early" bottom panel). As a result, the learning step (green arrow) is biased more towards the mean of the primitive distribution than the motor error (black arrow). However, later in learning, the number of primitives that can help to reduce the error is much smaller (many more grayish primitives, "Late" bottom panel), leading to a learning that is slower but more strongly aligned with the motor error.

(**B**) Effect of distribution width on learning behavior. Faded trajectories are for learning of opposite FFs.

(**C**) Simulated learning trajectories for different FFs in position/velocity gain-space, mediated by the motor primitives (gray dots).

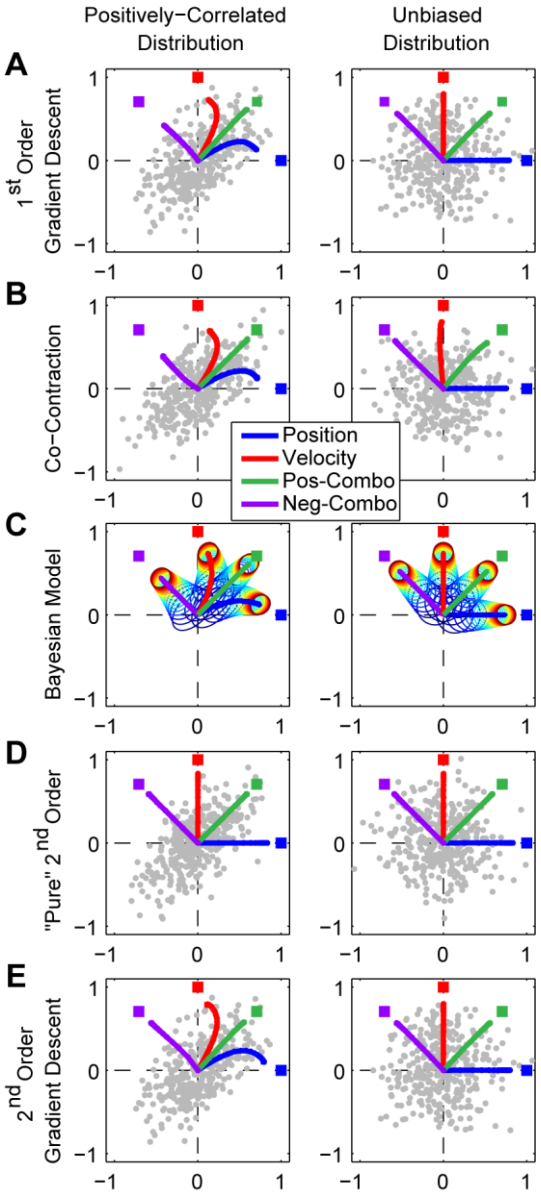

Figure 4. Alternative Learning Algorithms.

Gray dots represent sample primitives, each row utilizes the same learning rule, each column assumes the same general primitive distribution (i.e. positively-correlated or unbiased), and squares are FF goals. A first-order gradient descent learning rule (**A**), a co-contraction learning rule where weight changes can only be positive (**B**), and a Bayesian learning framework (**C**), all yield comparable behavioral predictions. A "pure" second-order learning rule does not yield comparable predictions (**D**), whereas a hybrid rule which determines learning steps based on both first-order and second-order derivatives (e.g. second-order gradient descent) produces predictions that are partially comparable (**E**). When an unbiased distribution is assumed, every learning rule will yield behavioral predictions not seen in the experimental evidence. The ellipses in the Bayesian learner framework (B) indicate the prior distribution at each time step.

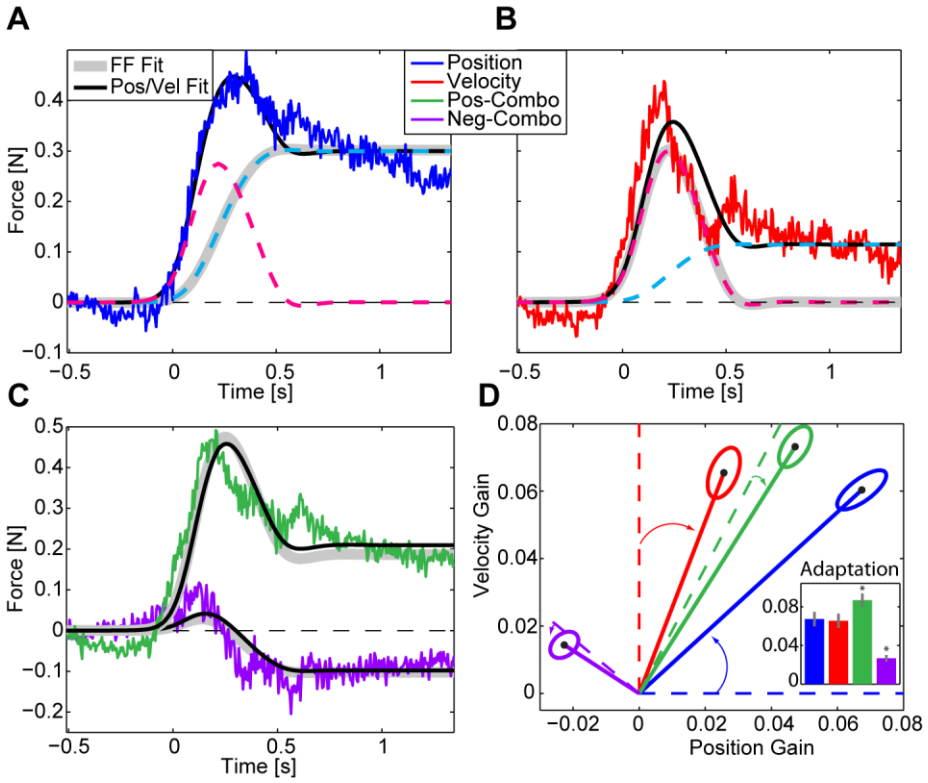

Figure 5. Single-Trial Learning of Various Force-Fields.

(**A**) Average learned force pattern (blue) after a single trial of exposure to a position-dependent FF. The force pattern cannot be fully explained by a fit to just position (gray curve); instead, fitting to *both* the position (cyan) and velocity (magenta) traces of the arm movements independently leads to a significantly closer fit (black curve).

(**B**) Learned force pattern after a single trial of learning a velocity-dependent. Again, the force pattern cannot be fully explained by a fit to just velocity (gray curve); instead, fitting to *both* the velocity (magenta) and position (cyan) traces of the arm movements independently leads to a closer fit (black curve).

(**C**) Learned force patterns after a single trial of exposure to either a positive-combination (green) or a negative-combination FF (purple). For these combination FFs, the FF-specific fit regresses the force patterns onto the shape of the respective perturbing FF (gray).

(**D**) Position/velocity dependence of single-trial learning. The curved arrows represent the amount of misalignment between the learned force patterns (black dots, solid lines) and the FF directions (dotted lines). The inset displays the amount of learning after a single trial of exposure to the various FF environments. Error ellipses represent standard error.

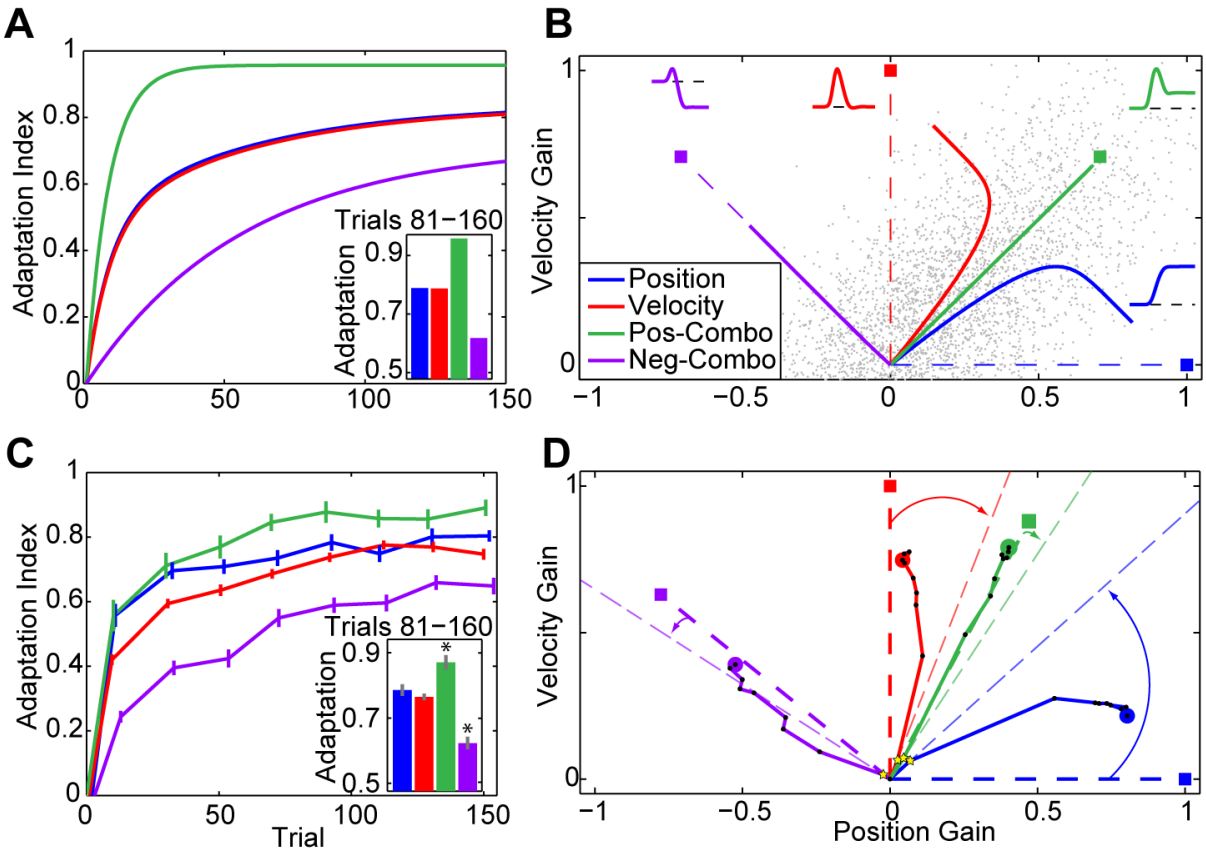

Figure 6. Extended Learning of Various Force-Fields.

(**A**) Model-predicted learning curves for different FFs.

(**B**) Model-predicted learning trajectories (colored curves) in position/ velocity gain-space as mediated by the motor primitives (gray dots). Learning is directed towards the FF goals (squares, force patterns).

(**C**) Experimental learning curves for different FFs. Each data point represents the average learning across subjects within 20-trial bins. The x-coordinates of the data points are the average locations of error-clamp probe trials within each bin (maximum standard error of these average locations is 1.3). The velocity-dependent, position-dependent, and negative-combination curves have one, two and three trial offsets, respectively, to facilitate viewing of the error bars, which represent standard errors. The inset displays the averaged learning during the second half of FF exposure (trials 81-160).

(**D**) Experimental learning trajectories (solid lines) in position/velocity gain-space. Note the similarity to the model-predicted pattern show in B. Curved arrows show the misalignment between the single-trial learning direction (dash-dot lines) and the FF direction (dotted lines). Each black dot represents the force output during a 20-trial bin, while the yellow stars denote the single-trial learning shown in Figure 3D.

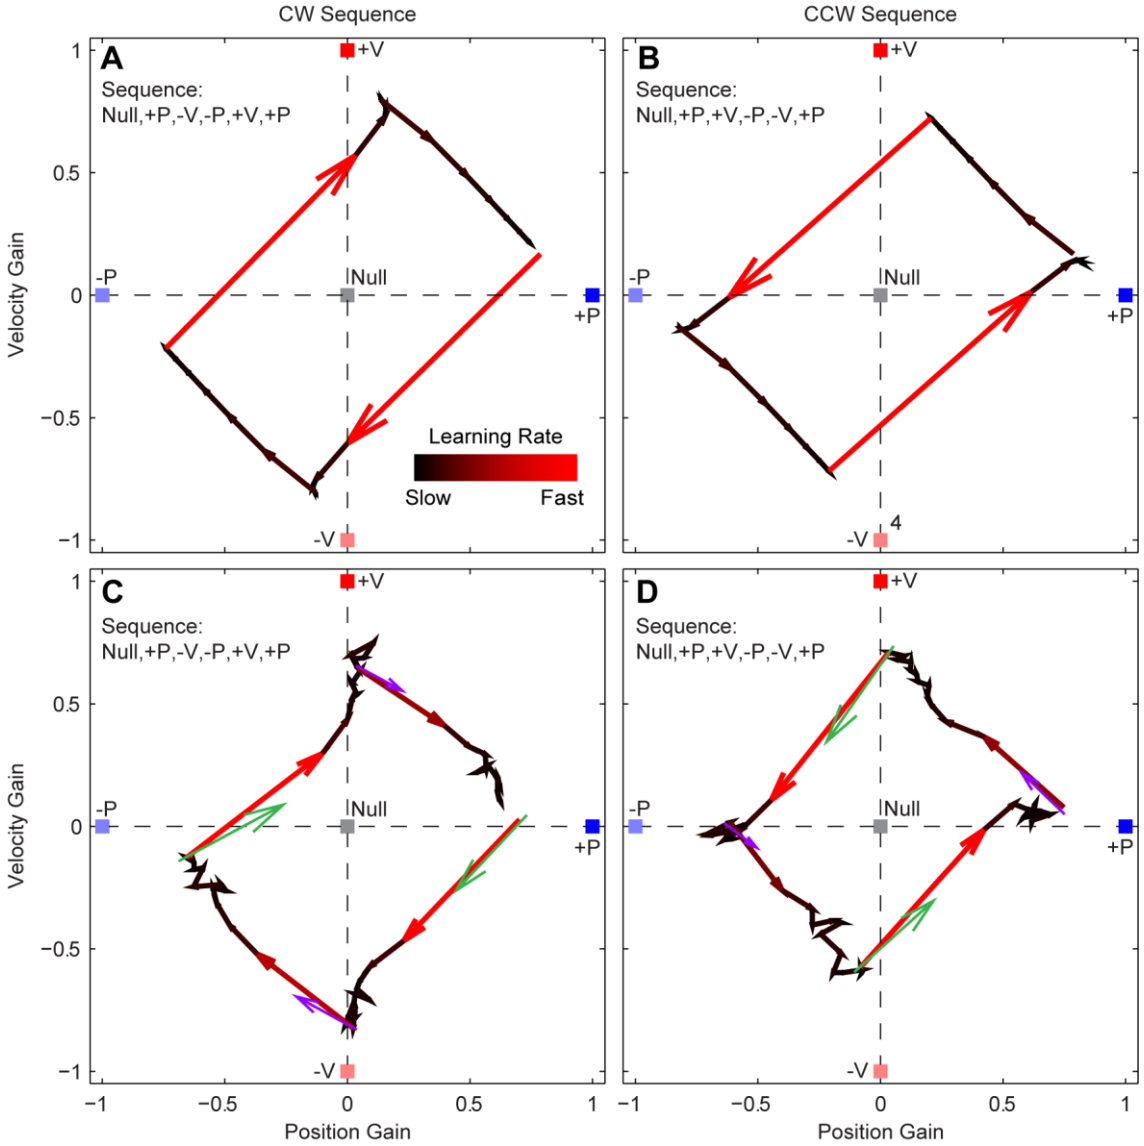

Figure 7. Interference Experiment Between Position-Dependent and Velocity-Dependent Force-Fields: Gain-Space Trajectories

(**A**) Simulated gain-space trajectories for CW sequence of FFs. The sequence of FFs experienced is indicated in the top-left corner of each panel. FF goals are indicated by squares (+P: blue, -P, light blue; +V: red, -V: light red). Longer and brighter arrows indicate greater learning changes in ten-trial intervals. (**B**) Simulated gain-space trajectories for CCW sequence of FFs. (**C**) Experimentally-observed gain-space trajectories for CW sequence of FFs. Red, black, and brown arrows indicate the learning changes every ten trials – if an error-clamp was not available at a particular trial, the point in gain-space was obtained by interpolation of neighboring data. Green and purple arrows represent the learning-related changes in the positive- and negative-combination directions, respectively, for adaptation after the first three trials of exposure to each FF. Subject averages are displayed. (**D**) Experimental gain-space trajectories for CCW sequence of FFs.

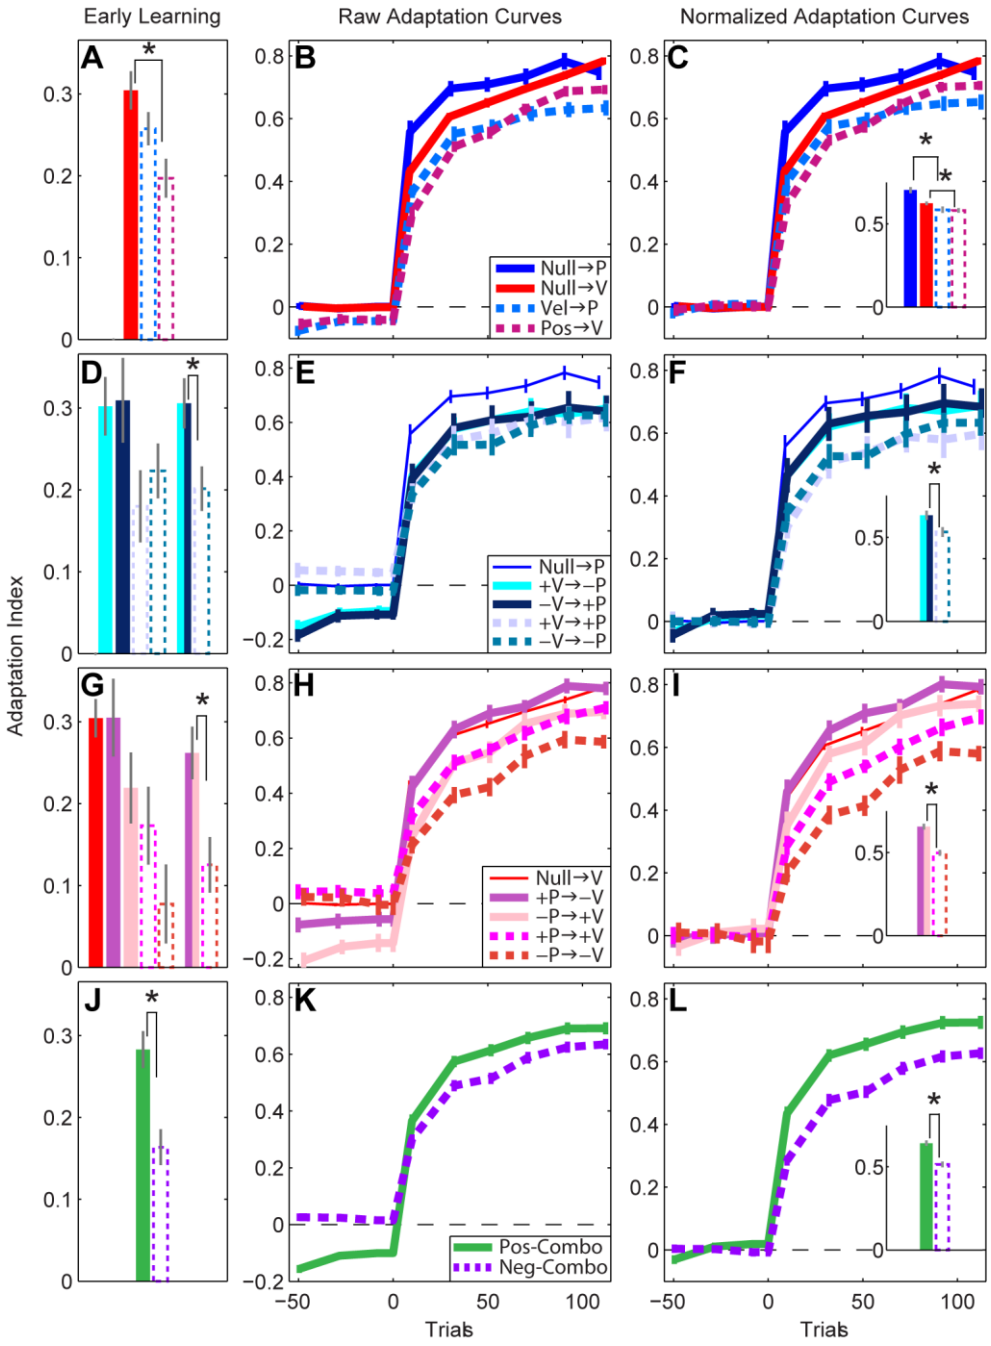

Figure 8. Interference Experiment Between Position-Dependent and Velocity-Dependent Force-Fields: Learning Curves

(**A-C**) Learning for Null→P, Null→V, V→P, and P→V transitions. First column displays learning during the first three trials of exposure to the new FF. Three-trial learning data was not available for the Null→P transition. Second column displays raw adaptation curves. Third column displays normalized adaptation curves (see Methods) – inset displays average value of normalized learning curves starting from Trial 0. *$p<0.0035$. (**D-F**) Learning for all transitions to P FFs. Solid curves indicate either Null→P or positive-combination transitions. Dotted curves indicate negative-combination transitions. *$p<0.0059$ (**G-I**) Learning for all transitions to V FFs. Solid curves indicate either Null→V or positive-combination transitions. Dotted curves indicate negative-combination transitions. *$p<0.0020$ (**J-L**) Learning curves grouped by gain-space direction. *$p<9.4\times10^{-5}$