Neuron, volume *64*

Supplemental Data

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

Gary C. Sing, Wilsaan M. Joiner, Thrishantha Nanayakkara, Jordan B. Brayanov, and Maurice A. Smith

Supplemental Information

Supplemental Methods

All experiment participants gave informed consent, and experimental protocols were approved by the Harvard University Institutional Review Board.

In all experiments, subjects were instructed to grasp the handle of a robotic manipulandum while making 500 ms, 10 cm 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 force-field 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 force-fields (CW or CCW versions of the four force-field types) were: Position-Dependent Field: K=±45 N/m, B=0 Ns/m; Velocity-Dependent Field: K=0 N/m, B=±15 Ns/m; Positive-Combination Field: K=±21.2 N/m, B=±13.2 Ns/m; Negative-Combination Field: K=∓ 35 N/m, B=±9.4 Ns/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 force-field.

When plotting the gain-space representation of learning, we adjusted the axes such that the velocity-dependent force-field of $15\frac{N \cdot s}{m}$ and a position-dependent force-field of $45\frac{N}{m}$ were represented by [0,1] and [1,0], respectively. We chose a magnitude of $45\frac{N}{m}$ for the position-dependent field so that it had similar strength to the velocity-dependent field (the peak perturbing forces were about 4.5 N in both cases). This normalization, which divides the position axis by $45\frac{N}{m}$ and the velocity axis by $15\frac{N \cdot s}{m}$ (Figure S1), roughly corresponds to rescaling these axes by the peak force produced and was made in all of the gain-space plots of data (Figures 2D, 5D, 6D, 7CD, S2A, S7AD, S8D, S12C). The learning metrics displayed in Figures 5D, 6C, 1st and 3rd columns of Figure 8, S2B, S5B, S7BC, S8D were calculated in this adjusted gain-space.

Any forces produced during learning are results of both feedforward adaptation and within-trial error correction. Isolating feedforward adaptation is necessary for an accurate representation of trial-to-trial adaptation. Using error-clamp probe trials, kinematic error within-trial can largely be removed by restricting 99% of lateral errors to no more than 1.2 mm (Figure 1C). These probe trials not only eliminate the fuel driving within-trial kinematic error correction, but also allow for direct observation of the feedforward adaptation. By applying a damped spring force (K=6000 N/m, B=250 Ns/m) to counteract any lateral movement deviation generated by subjects as they move in the error-clamp trials, we can then approximate the subject's lateral force output as the opposite of the robot clamping-force. All force patterns analyzed were appropriately baseline-subtracted, and those force patterns associated with "converse" force-fields (e.g. K=-45 N/m) were flipped so that they could be compared with the complementary force-fields (e.g. K=+45 N/m).

Note that two types of motor errors (both time-varying functions) are discussed in the text: errors in motor output (force) and errors in motion trajectory (kinematic). The focus of the viscoelastic primitive model is on force errors, which can be represented as an error vector of position and velocity gains (as shown in Figure 2D).

Experiment 1

We instructed 93 subjects (39 male, 4 left-handed, median age 23, age range 18-64) to make 160 null-field baseline trials (i.e. trials where no force-field was applied) to familiarize themselves with the task. We then applied velocity-dependent viscous curl fields (Figure 1B) represented by Equation 1 (K=0 N/m, B=±15 Ns/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:

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

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

$$\begin{bmatrix} \mu_{Post,x} \\ \mu_{Post,y} \end{bmatrix} = \left(\Sigma_{Prior}^{-1} + \Sigma_M^{-1} \right)^{-1} \Sigma_M^{-1} \begin{bmatrix} \mu_{M,x} \\ \mu_{M,y} \end{bmatrix} + \left(\Sigma_{Prior}^{-1} + \Sigma_M^{-1} \right)^{-1} \Sigma_{Prior}^{-1} \begin{bmatrix} \mu_{Prior,x} \\ \mu_{Prior,y} \end{bmatrix} \tag{4}$$

$$\Sigma_{Post} = \left(\Sigma_{Prior}^{-1} + \Sigma_M^{-1} \right)^{-1} \tag{5}$$

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

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

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

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

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

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 (Battiti, 1992). For a detailed explanation of these learning rules, see the "Learning Algorithm Alternatives" section below.

Experiment 2

To observe the learning after a single trial of exposure to either a position-dependent or velocity-dependent field, we had 39 subjects (Position-Dependent Field: 23 subjects, 7 male, all right-handed, median age 20, age-range 18-34; Velocity-Dependent Field: 16 subjects, 5 male, all right-handed, median age 20, age range 18-26) 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 two different tasks aimed at estimating the amount of single-trial and double-trial learning (see the "Double-Trial Learning" section below for double-trial learning discussion). In particular, subjects repeated 3-trial and 4-trial movement groups with 2-5 null-field washout trials in between each group. 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 field trials. Representative force patterns were found for each subject by averaging together 24 and 9 repetitions of these movement groups for the single-trial and double-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) and double-trial learning (Figure S8). The force patterns were also combined to find the average lateral force patterns produced after single-trial or double-trial learning for either a position-dependent or velocity-dependent field.

Experiment 3

See the "Calculating Positive-Combination and Negative-Combination Force-Fields" section for how the positive-combination and negative-combination fields were determined. To observe the learning after a single trial of either a positive-combination or negative-combination force-field, we had 39 subjects (Positive-Combination: 23 subjects, 10 male, 2 left-handed, median age 21, age range 18-24; Negative-Combination: 16 subjects, 5 male, 1 left-handed, median age 20, age range 18-34) repeat 3-trial and 4-trial movement groups in the 270° direction, mimicking the paradigm followed in Experiment 2 (see the "Double-Trial Learning" section below for double-trial learning discussion, Figure S8).

Experiment 4

After subjects familiarized themselves with the task through 200 null-field baseline trials, we instructed subjects to learn three force-fields (CW and CCW position-dependent, positive-combination, and negative-combination) for 160 trials in the 270° direction (Position-Dependent Field: 34 subjects, 24 male, 2 left-handed, median age 26, age range 18-60; Positive-Combination Field: 16 subjects, 7 male, all right-handed, median age 19, age range 18-30; Negative-Combination Field: 16 subjects, 7 male, all right-handed, median age 19, age range 18-30). Error-clamp probe trials were interspersed during exposure to the force-field 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 of these subject-specific force patterns, and took the mean of these contributions across subjects to find the gain-space representation of each bin (Figure 6D). Furthermore, to compare the learning of each force-field, we projected each binned point in gain-space onto the vector representing the force-field target (Figure 6C).

Experiment 5

Twenty subjects participated in the interference experiments (9 male, 2 left-handed, median age 19, age range 18-31). The different force-fields used in this experiment were of the same form as

Equation 1, with K=±45 N/m, B=0 Ns/m for the position-dependent force-fields, and K=0 N/m, B=±15 Ns/m for the velocity-dependent force-fields. 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.

Further Elaboration on the Viscoelastic Primitive Model

The weight update rule for this model is derived from minimization of the squared error in the system with respect to the weights.

$$\frac{\partial \frac{1}{2}\|\mathbf{error}\|^2}{\partial \mathbf{w}} = \frac{\partial \frac{1}{2}\|\mathbf{y}^* - \mathbf{S}^T\mathbf{w}\|^2}{\partial \mathbf{w}} = \frac{\partial}{\partial \mathbf{w}}\left[\frac{1}{2}\left(\mathbf{y}^* - \mathbf{S}^T\mathbf{w}\right)^T\left(\mathbf{y}^* - \mathbf{S}^T\mathbf{w}\right)\right] = -\mathbf{S}\left(\mathbf{y}^* - \mathbf{S}^T\mathbf{w}\right) \tag{9}$$

To find the weight changes that minimize the error, we take steps in the negative gradient direction. This expression is essentially a projection of the error vector $\left(\mathbf{y}^* - \mathbf{S}^T\mathbf{w}\right)$ onto each individual primitive. The weight update rule for an individual primitive then becomes the dot product of the error vector and the primitive vector:

$$\Delta\mathbf{w}_i = \eta\mathbf{S}_i\left(\mathbf{error}\right) = \eta\left|\mathbf{S}_i\right| \cdot \left|\mathbf{error}\right| \cdot \cos\left(\theta_{\mathbf{error}} - \theta_{\mathbf{S}_i}\right) \tag{10}$$

The updated weights scale the output of each primitive, which receive as inputs the reach position, p_y, and reach velocity, v_y.

$$\mathbf{y} = \sum_{i=1}^{n} \mathbf{S}_i{}^T\mathbf{w}_i = \sum_{i=1}^{n}\begin{bmatrix}\mathbf{k}_i p_y & \mathbf{b}_i v_y\end{bmatrix}^T \mathbf{w}_i \tag{11}$$

We write p_y and v_y here to denote the position and velocity along the y-axis because all of the movements studied in this task were point-to-point movements targeted along the y-axis. Note that an apparent implication of this representation is that force output, **y**, is only dependent on the position and velocity along the y-axis, suggesting a planar generalization function in velocity-space that is independent of x-velocity (i.e. with x-axis isoclines that monotonically increase up the positive y-axis, and monotonically decrease down the negative y-axis). However, previous work has shown that generalization functions in velocity-space are not planar, but are instead local and can be constructed with motor primitives displaying Gaussian tuning (Donchin et al., 2003; Thoroughman and Shadmehr, 2000) in velocity-space. In this particular study, we maintain a simple linear representation across position and velocity space; nevertheless, extension of the viscoelastic primitive model to local tuning is straightforward.

$$\mathbf{y} = \sum_{i=1}^{n} \mathbf{S}_i{}^T\mathbf{w}_i = \sum_{i=1}^{n}\begin{bmatrix}\mathbf{k}_i\mathbf{g}_i p_y & \mathbf{b}_i\mathbf{h}_i v_y\end{bmatrix}^T \mathbf{w}_i \tag{12}$$

Here, $\mathbf{k}_i\mathbf{g}_i$ and $\mathbf{b}_i\mathbf{h}_i$ are the gain-encoding basis elements in position- and velocity-space, respectively, where $\mathbf{g}_i$ and $\mathbf{h}_i$ represent the Gaussian shapes of the basis elements, and $\mathbf{k}_i$ and $\mathbf{b}_i$ represent the heights of these elements. As discussed above, these basis elements would be expected to display local tuning that may be well-approximated by Gaussian tuning functions in position- and velocity-space.

Calculating Positive-Combination and Negative-Combination Force-Fields

We have made the fundamental claim that early learning reflects the properties of the primitive distribution. Using simple behavioral experiments such as the single-trial learning experiments, it is possible to gain insight into these properties. The viscoelastic primitive model output after a single trial of learning is represented as:

$$Force(1) = \eta \mathbf{S}^T \mathbf{S} \mathbf{y}^* \tag{13}$$

where $\eta \mathbf{S} \mathbf{y}^*$ is the weight change. The gain on $\mathbf{y}^*$ in Equation 13, $\eta \mathbf{S}^T \mathbf{S}$, is a scalar multiple of $\mathbf{S}^T \mathbf{S}$, which can in turn be viewed as a transformation matrix. If we let $\mathbf{S} = \begin{bmatrix} \mathbf{k} & \mathbf{b} \end{bmatrix}$ with $\mathbf{k} = \begin{bmatrix} k_1 & k_2 & k_3 & ... \end{bmatrix}^T$ and $\mathbf{b} = \begin{bmatrix} b_1 & b_2 & b_3 & ... \end{bmatrix}^T$, then:

$$\mathbf{S}^T \mathbf{S} = \begin{bmatrix} k_1 & k_2 & k_3 & ... \\ b_1 & b_2 & b_3 & ... \end{bmatrix} \begin{bmatrix} k_1 & b_1 \\ k_2 & b_2 \\ k_3 & b_3 \\ \vdots & \vdots \end{bmatrix} = \begin{bmatrix} \mathbf{k}^T \\ \mathbf{b}^T \end{bmatrix} \begin{bmatrix} \mathbf{k} & \mathbf{b} \end{bmatrix} = \begin{bmatrix} \mathbf{k}^T \mathbf{k} & \mathbf{k}^T \mathbf{b} \\ \mathbf{b}^T \mathbf{k} & \mathbf{b}^T \mathbf{b} \end{bmatrix} \tag{14}$$

This transformation matrix receives as an input the target, $\mathbf{y}^*$, and outputs the appropriate force profile, $Force(1)$, after learning one trial. Although the various inner products between $\mathbf{k}$ and $\mathbf{b}$ are unknown, we are interested mainly in their population characteristics; accordingly, we take the expectations of these inner products in Equation 14 to find their means.

$$\begin{aligned} E\left[\mathbf{k}^T \mathbf{k}\right] &= E\left[\left(\left(\mathbf{k} - \mu_{\mathbf{k}}\right) + \mu_{\mathbf{k}}\right)^T \left(\left(\mathbf{k} - \mu_{\mathbf{k}}\right) + \mu_{\mathbf{k}}\right)\right] \\ &= E\left[\left(\mathbf{k} - \mu_{\mathbf{k}}\right)^T \left(\mathbf{k} - \mu_{\mathbf{k}}\right)\right] + \mu_{\mathbf{k}} \left(E\left[\mathbf{k}^T\right] - \mu_{\mathbf{k}}\right) + \mu_{\mathbf{k}} \left(E\left[\mathbf{k}\right] - \mu_{\mathbf{k}}\right) + E\left[\mu_{\mathbf{k}}^2\right] \\ &= \mu_{\mathbf{k}}^2 + Var\left(\mathbf{k}\right) \end{aligned} \tag{15}$$

$$\begin{aligned} E\left[\mathbf{k}^T \mathbf{b}\right] &= E\left[\left(\left(\mathbf{k} - \mu_{\mathbf{k}}\right) + \mu_{\mathbf{k}}\right)^T \left(\left(\mathbf{b} - \mu_{\mathbf{b}}\right) + \mu_{\mathbf{b}}\right)\right] \\ &= E\left[\left(\mathbf{k} - \mu_{\mathbf{k}}\right)^T \left(\mathbf{b} - \mu_{\mathbf{b}}\right)\right] + \mu_{\mathbf{b}} \left(E\left[\mathbf{k}^T\right] - \mu_{\mathbf{k}}\right) + \mu_{\mathbf{k}} \left(E\left[\mathbf{b}\right] - \mu_{\mathbf{b}}\right) + E\left[\mu_{\mathbf{k}} \mu_{\mathbf{b}}\right] \\ &= \mu_{\mathbf{k}} \mu_{\mathbf{b}} + Cov\left(\mathbf{k}, \mathbf{b}\right) \end{aligned} \tag{16}$$

Given that our distribution is centered at 0, we can therefore write the expectation of Equation 14 as:

$$E\left[\mathbf{S}^T \mathbf{S}\right] = \begin{bmatrix} Var\left(\mathbf{k}\right) & Cov\left(\mathbf{k}, \mathbf{b}\right) \\ Cov\left(\mathbf{k}, \mathbf{b}\right) & Var\left(\mathbf{b}\right) \end{bmatrix} \tag{17}$$

This symmetric matrix is precisely the covariance matrix.

Transformation matrices have certain vector inputs (i.e. eigenvectors) that elicit scaled versions of the inputs as responses. Furthermore, in a two-dimensional system such as our 2-D gain-space, a symmetric transformation matrix will have at most two of these eigenvectors, of which one that will cause a maximal response and one that will cause a minimal response. The directions of these eigenvectors are called the *eigendirections* of the matrix, where the major eigendirection leads to the maximal response, and the minor eigendirection leads to the minimal response. If these directions (or at least close approximations) can be found for our transformation matrix (Equation 17), then not only can one design force-fields that produce maximal or minimal force outputs, corresponding to reduced or

increased difficulty in learning, but one can also show that the learning in those eigendirections will be scaled version of the inputs (i.e. the outputs will be composed of position and velocity contributions in the same ratio as the force-fields).

The exact eigenvectors of this transformation matrix in Equation 17 depend on the correlation between $Var(\mathbf{k})$ and $Var(\mathbf{b})$, but in general, the major eigenvector is equal to:

$$\mathbf{v}_{major} = \begin{bmatrix} \dfrac{Var(\mathbf{k}) - Var(\mathbf{b}) + \sqrt{\left(Var(\mathbf{k}) - Var(\mathbf{b})\right)^2 + 4Cov(\mathbf{k},\mathbf{b})^2}}{2Cov(\mathbf{k},\mathbf{b})} \\ 1 \end{bmatrix} \tag{18}$$

In the extreme case where $Var(\mathbf{k})$ and $Var(\mathbf{b})$ are perfectly correlated with each other (i.e. $\rho = 1$), then the major eigenvector reduces to $\mathbf{v}_{major} = \begin{bmatrix} \dfrac{\sigma_k}{\sigma_b} \\ 1 \end{bmatrix} = \begin{bmatrix} \sigma_k \\ \sigma_b \end{bmatrix}$, indicating that the major eigendirection is $\dfrac{\sigma_b}{\sigma_k}$.

If a position-dependent or velocity-dependent field corresponds to a target of $\mathbf{y}^* = [1;0]$ or $\mathbf{y}^* = [0;1]$ in gain-space, respectively, the single-trial learning of both force-fields can then be represented as:

$$\begin{bmatrix} \mathbf{F}_{pos}(1) \\ \mathbf{F}_{pos}(2) \end{bmatrix} = \begin{bmatrix} Var(\mathbf{k}) \\ Cov(\mathbf{k},\mathbf{b}) \end{bmatrix}, \quad \begin{bmatrix} \mathbf{F}_{vel}(1) \\ \mathbf{F}_{vel}(2) \end{bmatrix} = \begin{bmatrix} Cov(\mathbf{k},\mathbf{b}) \\ Var(\mathbf{b}) \end{bmatrix} \tag{19), (20}$$

As Figure 5D shows, both of these single-trial force patterns are not aligned with their respective force-field goals, but instead are deviated towards the center of the first quadrant. An important property of transformation matrices is that any input not perfectly aligned with any of the eigendirections will be "rotated" towards the *major* eigendirection. For instance, take a matrix $\mathbf{A} \in \mathbb{R}^{2\times 2}$ with eigenvectors $\mathbf{v}_{major}, \mathbf{v}_{minor}$ corresponding to eigenvalues $\lambda_{major}, \lambda_{minor}$, with $\lambda_{major} > \lambda_{minor}$. Arbitrary input $\mathbf{v}_0 \in \mathbb{R}^{2\times 1}$ can be rewritten as $\mathbf{v}_0 = \alpha\mathbf{v}_{major} + \beta\mathbf{v}_{minor}$ because the two eigenvectors constitute a complete basis set in 2-D space (i.e. they are orthogonal to each other). Applying the transformation matrix A to the input v_0 gives $\mathbf{A}v_0 = \mathbf{A}\left(\alpha\mathbf{v}_{major} + \beta\mathbf{v}_{minor}\right) = \alpha\mathbf{A}\mathbf{v}_{major} + \beta\mathbf{A}\mathbf{v}_{minor} = \alpha\lambda_{major}\mathbf{v}_{major} + \beta\lambda_{minor}\mathbf{v}_{minor}$.

Given that $\lambda_{major} > \lambda_{minor}$, $\therefore \left(\dfrac{\alpha\lambda_{major}}{\beta\lambda_{minor}}\right) > \left(\dfrac{\alpha}{\beta}\right)$, meaning that the output of the transformation matrix contains a greater amount of the major eigenvector than does the input (i.e. the output has been rotated towards the major eigendirection). In the case of our single-trial learning paradigm, the major eigendirection *must* lie between the single-trial force patterns because of this induced learning rotation. To estimate this eigendirection, we take the geometric mean of the slopes representing the single-trial position and velocity force patterns (Equations 19 and 20):

$$\sqrt{\frac{\mathbf{F}_{pos}(2)}{\mathbf{F}_{pos}(1)} \cdot \frac{\mathbf{F}_{vel}(2)}{\mathbf{F}_{vel}(1)}} = \sqrt{\frac{Cov(\mathbf{k},\mathbf{b})}{Var(\mathbf{k})} \cdot \frac{Var(\mathbf{b})}{Cov(\mathbf{k},\mathbf{b})}} = \frac{\sigma_{\mathbf{b}}}{\sigma_{\mathbf{k}}} = 0.62s \tag{21}$$

If we define the relative amount of cross-adaptation, *C*, as the ratio between the inappropriate adaptation term and the appropriate adaptation term, then *C* is determined both by the correlation (ρ) in the distribution between position and velocity and the ratio of the standard deviations.

$$C_{pos} = \frac{\mathbf{F}_{pos}(2)}{\mathbf{F}_{pos}(1)} = \frac{Cov(\mathbf{k},\mathbf{b})}{Var(\mathbf{k})} = \rho\left(\frac{\sigma_{\mathbf{b}}}{\sigma_{\mathbf{k}}}\right) \quad (22)$$

$$C_{vel} = \frac{\mathbf{F}_{vel}(1)}{\mathbf{F}_{vel}(2)} = \frac{Cov(\mathbf{k},\mathbf{b})}{Var(\mathbf{b})} = \rho\left(\frac{\sigma_{\mathbf{k}}}{\sigma_{\mathbf{b}}}\right) \quad (23)$$

Note that the cross-adaptation for position-dependent learning increases both when the distribution becomes more correlated, and as the distribution's variance in velocity increases, compared to the variance in position.

Figure S1. Force-Field Ellipses.

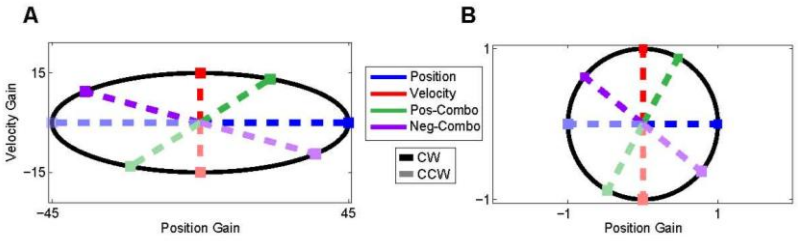

(**A**) Unadjusted "force-field ellipse" used to determine the force-field magnitudes for the positive-combination and negative-combination force-fields. Once the directions of the combination force-fields were found, we scaled them such that they were of the same magnitude as the position-dependent and velocity-dependent fields (i.e. so that they satisfied Equations 24 and 25). Squares are the final force-field goals. Faded dotted lines and squares represent CCW force-fields.
(**B**) Adjusted force-field ellipse. The axes are adjusted so that the position-dependent and velocity-dependent force-fields can be represented as [1,0] and [0,1], respectively. Faded dotted lines and squares represent CCW force-fields.

We now have an estimate of the major eigendirection for the maximal-learning (positive-combination) field (Equation 21), which is composed of positive position and velocity components and happens to be the major eigendirection if the variances were perfectly correlated; however, we have not yet determined the *magnitude* of the force-field. To do so, we require that the gain-space representation of the force-field lies on the circumference of the "force-field ellipse," which intersects the unadjusted gain-space axes (i.e. the axes are not divided by 45 and 15 as Figures 2D, 5D, 6D, S3A, S5AD, S6D are, see Methods) at the locations of the position-dependent and velocity-dependent force-field targets (Figure S1). This is equivalent to solving the following set of equations:

$$\frac{B}{K} = Force\ Field\ Direction, \quad \frac{K^2}{45^2} + \frac{B^2}{15^2} = 1 \quad (24), (25)$$

These two equations ensure both that the new force-field maintains the goal direction, and that the force-field magnitude in gain-space is the "same" as the single-state force-fields. For a force-field direction of 0.62s, solving these equations gives us a positive-combination field of $K = \pm 21.2\frac{N}{m}$, $B = \pm 13.2\frac{Ns}{m}$.

Our transformation matrix (Equation 17) is a symmetric matrix, meaning that the eigendirections are orthogonal to each other. Therefore, our estimate of the positive-combination direction should also be orthogonal to the minor eigendirection (negative-combination direction), or the minimal-output direction. However, orthogonality in position/velocity gain-space is a problematic issue because there is no reason to expect that *n* units of position gain should be treated equivalently to *n* units of velocity gain when these units are different from each other (N/m vs. Ns/m). To address this question, we assumed that a position gain of 45 N/m was comparable to a velocity gain of 15 Ns/m, which effectively weights velocity three times as heavily as position (for rational, see Experiment 1 Methods). We then looked at the single-trial learning of four different negative-combination force-fields, of which one was orthogonal in the adjusted gain-space to the positive-combination field (NC Field 1 (orthogonal): K=∓ 39.6 N/m, B=±7.05 Ns/m; NC Field 2: K=∓ -35 N/m, B=±9.4 Ns/m; NC Field 3: K=∓ 29.0 N/m, B=±11.46 Ns/m;

NC Field 4: K=∓ 21.2 N/m, B=±13.2 Ns/m; Figure S2). Despite the wide range of force-field directions, learning for *all* of the negative-combination fields was slower than the learning of the position-dependent, velocity-dependent, and positive-combination fields (Figure S2B). The learning of NC Fields 1 and 2 was rotated counter-clockwise from the fields, while the learning of NC Fields 3 and 4 was rotated clockwise, indicating that all four force-fields consistently rotated *towards* the major eigendirection, or equivalently, *away* from a putative minor eigendirection in between NC Fields 2 and 3. Although all fields were roughly orthogonal to the positive-combination direction, only NC Field 2 produced learning that was close to a scaled version of the input; therefore, we chose to investigate this field further by taking more data and analyzing this field in the main text.

Figure S2. Negative-Combination Force-Fields.

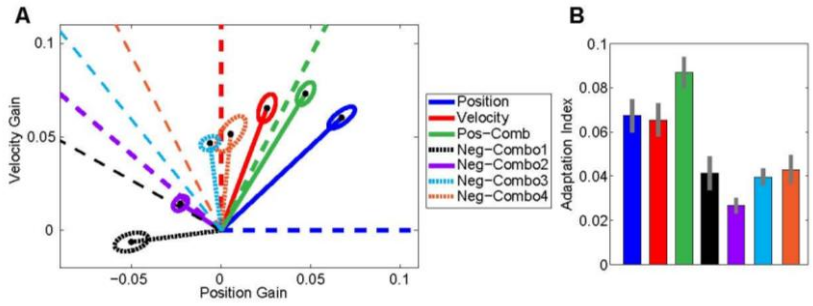

(**A**) Position/velocity dependence of single-trial learning for position-dependent, velocity-dependent, positive-combination, and four negative-combination force-fields. We compared the outputs of these four negative-combination force-fields to their respective inputs to see which one was the closest to a scaled version of the input. Negative-Combination Fields 1, 3, and 4 (dotted ellipses) were not mentioned in the main text. Error ellipses represent standard error.

(**B**) The amount of learning after a single trial of exposure to the various force-field environments. Errorbars represent standard error.

Negative-Combination Experiments

To determine which negative-combination field produced learning that was a scaled version of the input, we had 28 subjects learn four different negative-combination force-fields, with 6-8 subjects for each field (NC Field 1: 6 subjects, 1 male, all right-handed, median age 20.5, age range 18-26; NC Field 2: 8 subjects, 2 male, all right-handed, median age 19.5, age range 18-22; NC Field 3: 8 subjects, 2 male, 1 left-handed, median age 21, age range 18-26; NC Field 4: 6 subjects, 3 male, all right-handed, median age 19, age range 18-22). We had these subjects repeat 3-trial and 4-trial movement groups in the 270° direction, mimicking the paradigm followed in Experiment 2. After choosing to investigate NC Field 2 further, we collected data from an additional 8 subjects (8 subjects, 3 male, 1 left-handed, median age 20, age range 18-23).

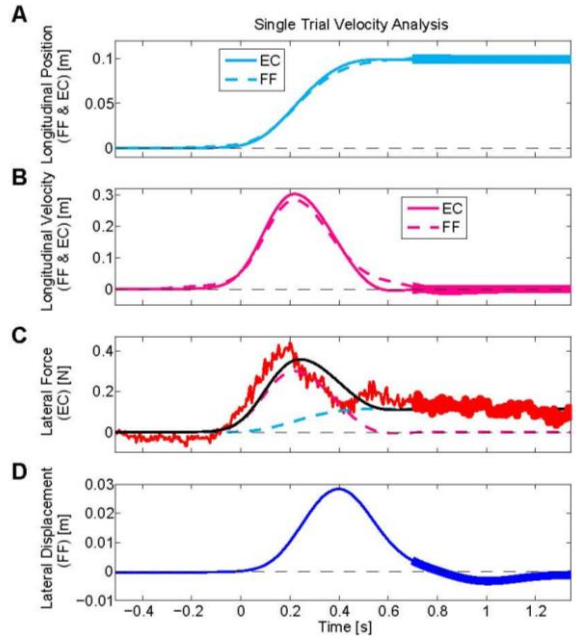

Figure S3. Force and Error Persistence After Movement Onset for Single-Trial Velocity Learning.

(**A**) Longitudinal position traces for the error-clamp and force-field trials used to measure single-trial velocity force-field learning.

(**B**) Longitudinal velocity traces for the error-clamp and force-field trials previously described.

(**C**) Lateral force production measured on error-clamp trials reveals persistence of force, even after movement termination (bold region starting at 0.7 sec).

(**D**) Lateral displacement measured on force-field trials is nearly gone at movement termination (bold region) and even crosses the zero-line.

Force and Error Persistence At Movement End for Single-Trial Velocity Learning

A possible alternative explanation for the persistence of force output measured during error-clamp trials early in velocity force-field learning is that the motor system is merely attempting to compensate for persistent lateral kinematic error in the previous (non-error clamp) trial. However, Figure 2B,C reveals that this error goes away as early as 600 ms after movement onset. Furthermore, analysis of the lateral kinematic error for single-trial velocity learning reveals that 700 ms after movement onset, the lateral error is substantially gone, and actually crosses the zero-line, while the persistence of positive force is maintained even beyond to 1200 ms after movement onset (Figure S3). The force data (Figure S3C) is taken from the error-clamp measurements during the single-trial velocity learning, the lateral error (Figure S3D) is taken from the force-field trials preceding the second error-clamp trial (see Methods for Experiment 2), and the longitudinal position (Figure S3A) and velocity (Figure S3B) are taken from both the force-field and error-clamp trials.

Static Force-Trail Induces Lateral Error in Force-Field Trials

Our analysis in Figure 2 shows that the static force-tails do not *result from* static kinematic errors, which are small and oppositely-directed from what would be required to cause persistent end-movement forces that resemble the force-tails observed. However, one would expect these force-tails to *result in* small kinematic errors at movement end. Note that these end-movement kinematic errors should be oppositely-directed to the errors caused by the force-field during movement because upon movement termination, the static force-tails would over-compensate the force-field perturbation. These kinematic errors would be small because the magnitudes of the force-tails observed are small (less than 0.6 N) given the stiffness of the arm. Despite the small size of these errors, we would expect a systematic relationship between the magnitude of the force-tails measured in error-clamp trials and end-movement kinematic error measured in force-field trials. Analysis of this relationship is presented below in Figure S4. We find (1) that these end-movement kinematic errors are indeed oppositely-directed to the force-field-induced errors, and (2) that as the force-tail decreases over the training period, lateral errors also decrease in a corresponding fashion ($r=0.90$, $p=6.0\times10^{-5}$correlation between force and lateral error measured between 750 and 850 ms) , suggesting that these end-movement errors (which tend to be maximal at about 800 ms after movement onset) may be driven by the static force-tails, shown in Figure 2A and Figure S4.

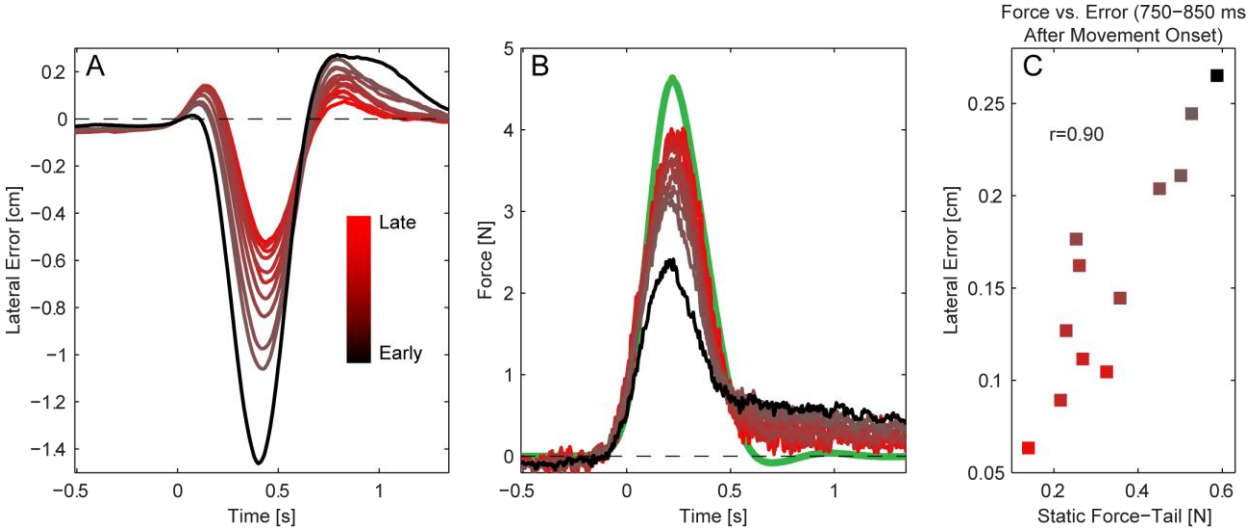

Figure S4. Static Force-Tail Induces Late Lateral Errors in Force-Field Trials

(**A**) Progression of lateral error in force-field trials during the learning of a velocity-dependent force-field. Each error trace is the average of the error traces for each subject averaged over a 20-trial bin. (**B**) Progression of lateral force output in error-clamp trials during the learning of a velocity-dependent force-field. Duplicate of Figure 2A. (**C**) Average lateral force output versus average lateral error from 750-850ms for each 20-trial bin.

Pattern of Errors Across Different Force-Fields Cannot Explain Pattern of Learning Rates

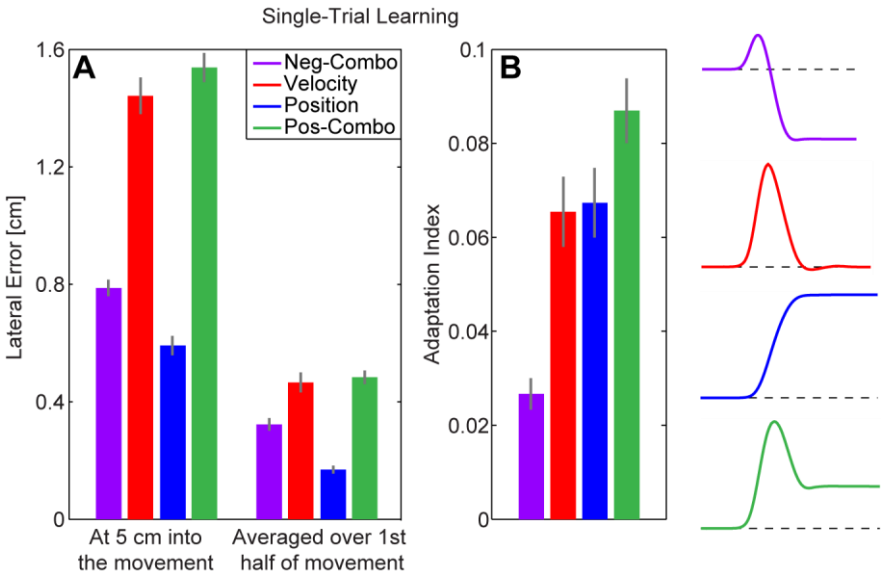

Figure S5. Patterns of Single-Trial Errors and Learning Rates

(**A**) Magnitude of lateral kinematic errors induced during single-trial exposure to a force-field at 5 cm into the movement, and over the first 5 cm of movement. Errorbars represent standard errors.

(**B**) Amount of learning after a single trial of exposure to the various force-field environments. Duplicate of inset in Figure 5D, except the order of presentation is rearranged from lowest to highest learning rates.

A seemingly reasonable alternative explanation to our findings that positive- and negative-combination force-fields are learned more quickly than position- and velocity-dependent dynamics (other than the viscoelastic primitive model), is that the pattern of learnability could be explained by the patterns of errors experienced during the early stages of movement. To determine if trial-to-trial adjustments driven by errors associated with the first half of the movement can account for our results, we calculated two metrics of early-movement error for the single-trial learning data set: (1) the lateral error at 5cm into the movement (the halfway point), and (2) the average lateral error during the first 5cm of the movement. Analysis of the data reveals a different pattern of errors across the different force-fields (Figure S5). The reason this occurs is that velocity-dependent dynamics have a much stronger effect on errors during the first half of the movement than position-dependent dynamics because maximal force-field strength for position-dependent dynamics is not reached until the end of the movement, whereas the force-fields with large velocity-dependent contributions have already reached maximal strength by the movement midpoint. Accordingly, we found that the position-dependent field displayed significantly smaller mid-movement errors than both the velocity-dependent and negative-combination fields despite showing slightly higher learning rates than the velocity-dependent field and significantly higher learning rates than the negative-combination field. These data, which are displayed in Figure R5, show that the pattern of errors across the different force-fields does not explain the learnability of these dynamics.

Experiment 1 Notes

The first bin of trials while learning the velocity-dependent force-field contained significant contributions from both position and velocity (Figure 2B, position partial R^2=0.64, velocity partial R^2=0.92, $\theta_{B/K}$=75.3°, where the gain-space angle ($\theta_{B/K}$) is the angle of the vector connecting the origin and the gain-space representation of the force output in the first bin). The last bin of learning the velocity-dependent field contained much more velocity than position contributions (Figure 2C, position partial R^2=0.19, velocity partial R^2=0.98, $\theta_{B/K}$=87.8°).

Experiments 2 and 3 – Notes

The single-trial learned force patterns were fit significantly by position and velocity (Position-Dependent Field: R^2=0.96, position partial R^2=0.94, velocity partial R^2=0.85; Velocity-Dependent Field: R^2=0.82, position partial R^2=0.56, velocity partial R^2=0.77; Positive-Combination Field: R^2=0.88, position partial R^2=0.79, velocity partial R^2=0.78; Negative-Combination Field: R^2=0.71, position partial R^2=0.68, velocity partial R^2=0.25; $p<1\times10^{-12}$ for all fields and both states). Furthermore, the initial

learning directions as expressed in the gain-space angle (Position-Dependent Field: $\theta_{B/K}$=42.8°; Velocity-Dependent Field: $\theta_{B/K}$=70.5°; Positive-Combination Field: $\theta_{B/K}$=57.7°; Negative-Combination Field: $\theta_{B/K}$=140.3°) were significantly different from each other for the different groups (PC/Position: $p=6.6\times10^{-5}$; PC/Velocity: p=0.0036; PC/NC: $p=1.6\times10^{-12}$; NC/Position: $p=1.1\times10^{-14}$; NC/Velocity: $p=4.0\times10^{-8}$; Position/Velocity: $p=1.1\times10^{-7}$, one-sided unpaired t-test).

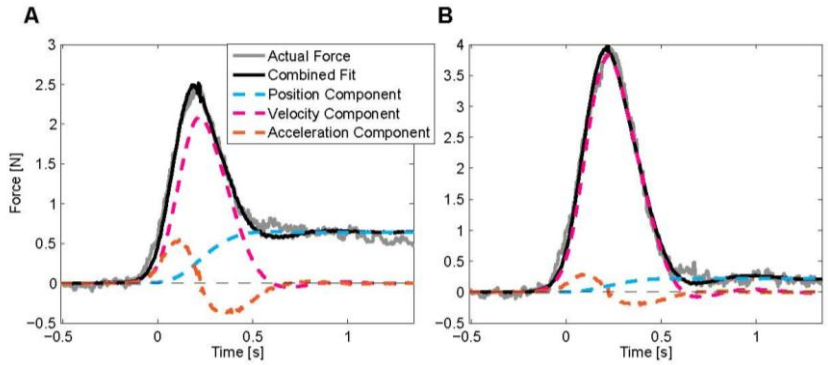

Figure S6. Contributions of Position, Velocity, and Acceleration Signals to Motor Output in Velocity-Dependent Field Learning.
Note that this figure is similar to panels B and C in Figure 2 in the main paper, except that here lateral forces are regressed onto position, velocity, and acceleration, rather than position and velocity alone.
(**A**) The average force pattern (gray line) measured during the first bin of learning a velocity-dependent force-field displays clear position (cyan line), velocity (magenta line), and acceleration (orange line) components. The force pattern can be approximated (black line) with a positive, linear combination of these three variables. (**B**) A characterization of the position, velocity, and acceleration contributions of the averaged force pattern during the last bin of learning.

Acceleration Dependence

The characterization of position and velocity contributions in the force outputs in Figure 2B,C produces excellent fits (early: R^2=0.93; late: R^2=0.98). Adding acceleration traces to these fits results in highly significant but relatively small improvements in the representation of these force profiles (see Figure S4). Early in training, although the acceleration signal's contribution is significant ($p<\times10^{-12}$) and increases the overall variance accounted for from 93% to 99% (position partial R^2=0.93, velocity partial R^2=0.98, acceleration partial R^2=0.81), it only accounts for 6% of the force profile variance, whereas position accounts for 19% and velocity accounts for 74%. The acceleration portion is 10-fold smaller than the velocity contribution and 3-fold smaller than that of position. Late in training, the velocity contribution dominates the other two (position partial R^2=0.35, velocity partial R^2=0.99, acceleration partial R^2=0.28). Accordingly, we focus only on the contributions of the position and velocity state variables.

CW and CCW Learning

We instructed subjects to learn CW and CCW versions of the four different force-fields (see Equation 1). After flipping the force outputs produced after exposure to the CCW force-fields, we averaged together the force outputs to obtain the data used to generate our figures. However, unlike model predictions, the learning of a CW force-field is not merely a flipped version of the learning of a CCW force-field – this can be readily seen in Figure S5 – the learning extents and trajectories are different. Nevertheless, the positively-correlated primitive distribution is reflected in the gray dotted ellipses (Figure S5AD) connecting learning data at similar training progression. These ellipses are not circular, which would reflect an unbiased primitive distribution, but instead have major and minor axes that resemble the major and minor eigenvector directions of the transformation matrix (Equation 17), respectively.

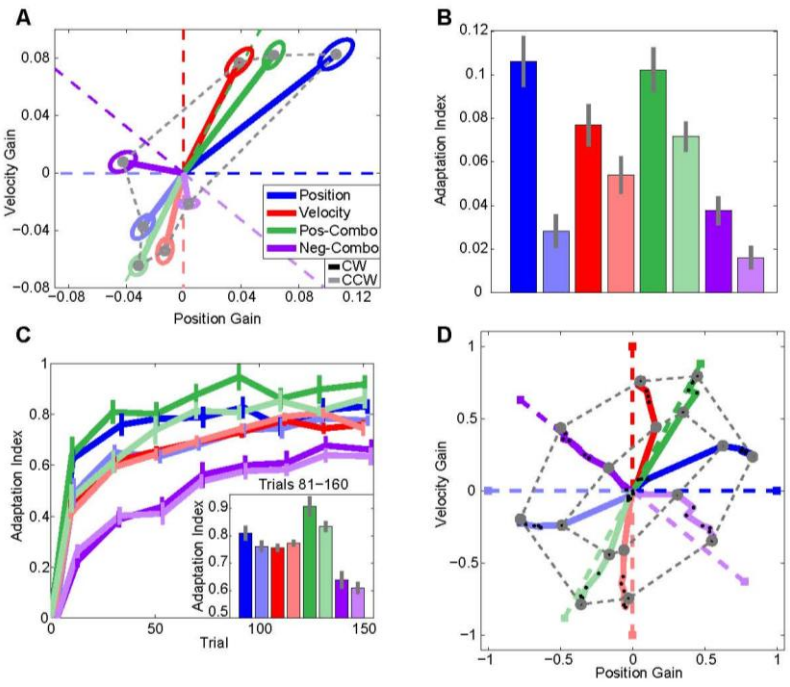

Figure S7. Learning of Clockwise and Counter-Clockwise Force-Fields. (**A**) Position/velocity dependence of single-trial learning for the CW and CCW position-dependent, velocity-dependent, positive-combination, and negative-combination force-fields. Error ellipses represent standard error. Faded trajectories, ellipses, and dotted lines correspond to CCW force-fields. Gray dotted lines form an ellipse that reflects the positive correlation of the primitive distribution.

(**B**) The amount of learning after a single trial of exposure to the various force-field environments. Errorbars represent standard error.

(**C**) Experimental learning curves for different force-fields. Each data point represents the average learning across subjects within 20-trial bins. The x-placements of the data points are the average locations of error-clamp probe trials within a particular bin. The velocity-dependent, position-dependent, and negative-combination curves have one, two and three trial offsets, respectively, to enable better viewing of errorbars, which represent standard errors. The inset displays the averaged learning during the second half of force-field exposure (trials 81-160).

(**D**) Experimental learning trajectories (solid lines) in position/velocity gain-space. Each black dot represents the force output during a 20-trial bin, while the colored squares and dotted lines represent the force-field goals and directions, respectively. The black squares denote the single-trial learning shown in **A**. Faded trajectories, dotted lines, and colored squares correspond to CCW force-fields. Gray dotted lines (connecting the first and last 20-trial bins) form ellipses that reflect the positive correlation of the primitive distribution.

Double-Trial Learning

During Experiments 2 and 3, we looked closely at both the single-trial and double-trial learning for the four different force-field environments. However, we withheld the double-trial learning data from the main text in order to simplify the presentation of the viscoelastic primitive model. Here, we present the withheld data and the associated analysis (Figure S6). Like the single-trial learning, all four double-trial learned force patterns were represented well by combining position and velocity contributions (Position-Dependent Field: R^2=0.97, position partial R^2=0.96, velocity partial R^2=0.89; Velocity-Dependent Field: R^2=0.85, position partial R^2=0.59, velocity partial R^2=0.81; Positive-Combination Field: R^2=0.90, position partial R^2=0.80, velocity partial R^2=0.85; Negative-Combination Field: R^2=0.71, position partial R^2=0.68, velocity partial R^2=0.27; $p<1\times10^{-12}$ for all fields and both states). The data suggest that the positive-combination learning is better than the position-dependent and velocity-dependent learning (PC/Position: p=0.22; PC/Velocity: p=0.20, one-sided unpaired t-test), and show that the negative-combination learning is significantly worse than the other three fields (NC/Position: $p=2.8\times10^{-7}$; NC/Velocity: $p=4.8\times10^{-5}$; NC/PC: $p=1.8\times10^{-6}$, one-sided unpaired t-test).

The gain-space angles for the double-trial learning match up very well to those for the single-trial learning (Position-Dependent Field: $\theta_{B/K}$=44.3°; Velocity-Dependent Field: $\theta_{B/K}$=69.6°; Positive-Combination Field: $\theta_{B/K}$=61.2°; Negative-Combination Field: $\theta_{B/K}$=138.2°; single-trial learning gain-space angles above) and were also significantly different from each other (PC/Position: $p=2.9\times10^{-6}$; PC/Velocity: p=0.017; PC/NC: $p=3.5\times10^{-11}$; NC/Position: $p=4.4\times10^{-14}$; NC/Velocity: $p=4.0\times10^{-8}$; Position/Velocity: $p=6.6\times10^{-1}$, one-sided unpaired t-test).

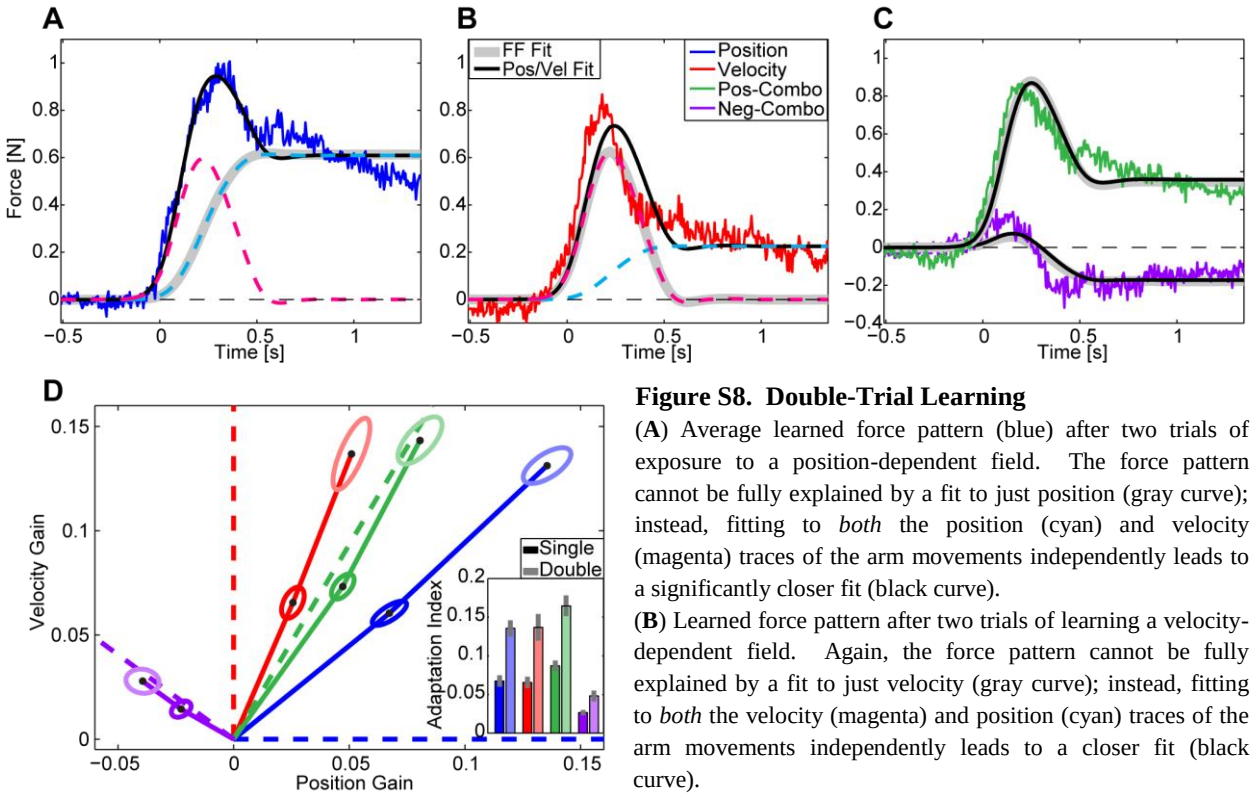

Figure S8. Double-Trial Learning

(**A**) Average learned force pattern (blue) after two trials of exposure to a position-dependent field. 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 two trials of learning a velocity-dependent field. 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 two trials of exposure to either a positive-combination (green) or a negative-combination field (purple). For these combination fields, the force-field-specific fit regresses the force patterns onto the shape of the respective perturbing force-field (gray).

(**D**) Position/velocity dependence of single-trial learning. The black dots and solid lines represent the learned force patterns, while the dotted lines represent the force-field direction. The inset displays the amount of learning after one and two trials of exposure to the various force-field environments. Error ellipses represent standard error.

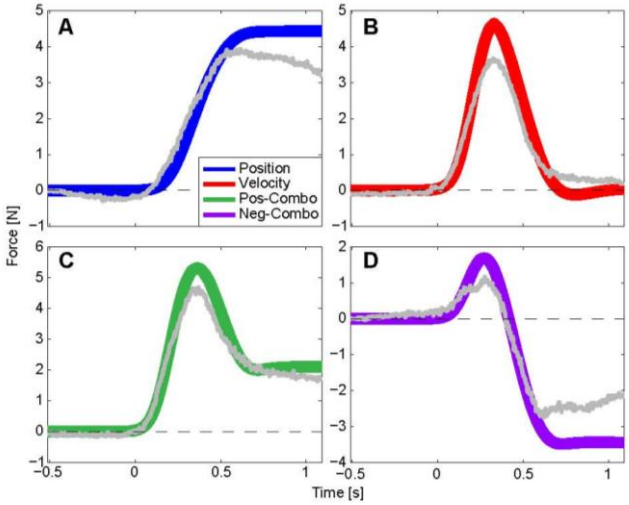

Late Learning

In Experiments 1 and 4, we exposed subjects to various force-fields for extended periods of time. Their learning trajectories are shown in Figure 6D. An alternative representation of their motor output in the last 20-trial bin of learning besides the gain-space representation is the average force patterns during that bin (Figure S9).

Figure S9. Late Learning for Different Force-Fields.

(**A-D**) The gray traces represent the average force patterns for the last half of extended learning (trials 81-160) in Experiment 4. The colored traces represent the desired force-field goals.

Learning Algorithms and Primitive Distribution Alternatives

We designed our viscoelastic primitive model with a motor primitive distribution centered at the origin in position/velocity gain-space and a positive correlation between the learning elements' responses to position and velocity (Figure 3). As noted in the main text, several predictions arise from this model. In Figure 3, a first-order gradient descent learning rule was applied to the viscoelastic primitive model, and the predictions we tested corresponded to this implementation. However, as we mentioned in the

main text, several other reasonable alternatives for learning rules exist. We applied several of these learning rules to the viscoelastic primitive model, and generally found that they made very similar predictions to when the first-order gradient descent learning rule was applied.

For example, recent work has shown that motor activation can increase, albeit asymmetrically, in both agonist and antagonist directions in response to large motor errors (Franklin et al., 2008). Specifically, when a large motor error is experienced, muscles that agonize correction of the error are activated more strongly in subsequent trials. However, muscles that antagonize correction of error also receive increased activation, but this increase is smaller than for the agonist muscles. This causes a "V-shaped" relationship between adaptation and error for each muscle where the slope of the "V" that would agonize the correction of the error is greater in magnitude than the other slope. The net effects are an increase in co-contraction coupled with an increase in the net motor output that would oppose the error. This co-contraction-based learning rule would give essentially the same net motor output as a simple gradient descent learning rule with a learning rate that corresponds to the sum of the slopes of the "V." The difference between these two rules is that the amount of co-contraction roughly corresponding to the mean activity of the learning elements in our model would differ. An implementation of this co-contraction-based learning rule is shown below:

$$\Delta \mathbf{w}_i = \begin{cases} \eta_1 \mathbf{S}_i \cdot \mathbf{error} - c & \mathbf{S}_i \cdot \mathbf{error} > 0 \\ \eta_2 \mathbf{S}_i \cdot \mathbf{error} - c & \mathbf{S}_i \cdot \mathbf{error} \le 0 \end{cases} \tag{26}$$

where $\eta_1 > 0$, $\eta_2 < 0$, $|\eta_1| > |\eta_2|$, and $c > 0$ in keeping with the above description. This learning rule would cause negative motor errors to drive positive weight changes (leading to an increase in error), and positive errors to drive even larger positive weight changes (leading to a net decrease in error). We found that this yields essentially identical predictions about the pattern of crosstalk and learning rate changes that occur in pure and hybrid force-fields (compare Figure 4B to Figure 4A).

A Bayesian framework can also yield comparable predictions. Grossly speaking, the Bayesian learning rule combines uncertain prior expectations with noisy sensory information in an optimal way. In the Gaussian case, the best posterior estimate combines the means of these distributions via a simple weighted average, where the weights correspond to the relative uncertainties of these distributions. For our implementation of the Bayesian model, we define one of the inputs, termed the prior distribution, as the primitive distribution at the beginning of each learning trial. We define the other input, termed the state measurement, as the learning goal for that particular trial. The explicit Bayesian learning rule for the combined output of these two inputs, termed the posterior distribution, is defined by its mean and variance:

$$\begin{bmatrix} \mu_{Post,x} \\ \mu_{Post,y} \end{bmatrix} = \left(\Sigma_{Prior}^{-1} + \Sigma_M^{-1}\right)^{-1} \Sigma_M^{-1} \begin{bmatrix} \mu_{M,x} \\ \mu_{M,y} \end{bmatrix} + \left(\Sigma_{Prior}^{-1} + \Sigma_M^{-1}\right)^{-1} \Sigma_{Prior}^{-1} \begin{bmatrix} \mu_{Prior,x} \\ \mu_{Prior,y} \end{bmatrix} \tag{27}$$

$$\Sigma_{Post} = \left(\Sigma_{Prior}^{-1} + \Sigma_M^{-1}\right)^{-1} \tag{28}$$

where $\begin{bmatrix} \mu_{Post,x} \\ \mu_{Post,y} \end{bmatrix}$ and Σ_{Post} are the mean and covariance of the posterior, $\begin{bmatrix} \mu_{M,x} \\ \mu_{M,y} \end{bmatrix}$ and Σ_M are the mean and covariance of the measurement distribution, and $\begin{bmatrix} \mu_{Prior,x} \\ \mu_{Prior,y} \end{bmatrix}$ and Σ_{Prior} are the mean and covariance matrix of the prior. If the prior distribution is positively-correlated between position and velocity (i.e. its off-diagonal terms are positive), and the measurement distribution is isotropic (i.e. the covariance matrix is a scaled identity matrix), then the posterior distribution will also be positively-correlated, and the posterior mean will be biased towards the direction of maximal variance of the prior distribution. Figure

4C illustrates that the gain-space trajectories for this Bayesian learning rule for a positively-correlated primitive distribution and a uniform primitive distribution are extremely similar to the trajectories obtained with a first-order gradient descent learning rule (Figure 4A) and the co-contraction learning rule (Figure 4B) for these two distributions.

It is important to note that higher-order learning rules which depend on higher-order derivatives of the weights with respect to motor errors could in theory generate motor output that effectively compensates for the primitive distribution. A "pure" second-order learning rule (Battiti 1992) in which a step is made toward the minimum of the error space (approximated with a second-order model) provides learning which is entirely independent of the primitive distribution in our model (Figure 4D). Here the weight update is based on a combination of the gradient and second-derivative, also known as the Hessian matrix. The update rule is given by:

$$\Delta\mathbf{w} = -\eta\left(\nabla^2\mathbf{error}\right)^{-1}\left(\nabla\mathbf{error}\right) \tag{29}$$

where η is the learning rate, $\nabla\mathbf{error} = \mathbf{S}\left(\mathbf{S}^T\mathbf{w} - \mathbf{y}^*\right)$ is the gradient, and $\nabla^2\mathbf{error} = \mathbf{S}\mathbf{S}^T$ is the Hessian matrix.

Hybrid learning rules which take steps determined by a combination of first-order and higher-order learning rules (Battiti 1992) might predict somewhat comparable behavior to first-order learning rules, but are able to do so precisely because of the dependence on the first-order learning rule. For example, the second-order gradient descent learning rule determines the step direction according to the first derivative (i.e. gradient), but determines the step size using the second derivative (i.e. Hessian) instead of using a fixed step size:

$$\Delta\mathbf{w} = -\eta k\hat{\mathbf{g}} = -\eta\left(\frac{\hat{\mathbf{g}}^T\left(\nabla\mathbf{error}\right)}{\hat{\mathbf{g}}^T\left(\nabla^2\mathbf{error}\right)\hat{\mathbf{g}}}\right)\hat{\mathbf{g}} \tag{30}$$

where k is the step size determined by the Hessian, and $\hat{\mathbf{g}} = \nabla\mathbf{error} / \left\|\nabla\mathbf{error}\right\|$ is the unit vector pointing in the gradient direction. This particular learning rule predicts curved trajectories when learning position-dependent and velocity-dependent force-fields, but does *not* predict slower negative-combination learning (Figure 4E).

Other choices of primitive distributions in the viscoelastic primitive model will yield similar results to the distribution we used as long as positive-correlation between position and velocity is maintained. For example, a primitive distribution with two groups of primitives, one centered in the 1st quadrant and one centered in the 3rd quadrants, as shown in Figure S8, yields essentially identical predictions to the zero-mean, elliptical distribution used in the main text, although the shape of this bimodal distribution is quite different.

Figure S10. Primitive Distribution Alternative.

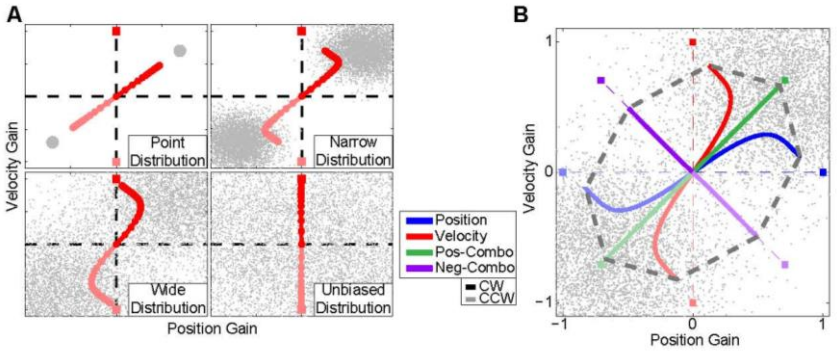

(**A**) Effect of distribution correlation/width on learning behavior. (**B**) Simulated learning trajectories (colored curves) in position/velocity gain-space as mediated by the motor primitives (gray dots). Learning is directed towards the force-field goals (squares, colored dotted lines). Faded trajectories and colored squares correspond to CCW force-fields. Gray dotted lines form an ellipse that reflects the positive correlation of the primitive distribution.

Other Potential Cause of Interference in Experiment 5

In Experiment 5, we found that interference from position-dependent dynamics onto velocity-dependent dynamics (and vice versa) was not monolithic, but instead depended on whether the transition between the two consecutively-learned force-fields required a positively- or negatively-correlated change in dynamics in position/velocity gain-space. However, other factors could potentially influence the patterns of observed interference. For instance, subjects in subgroup 1 learned a force-field sequence of +P, +V, -P, -V, +P in the 270° direction, and a force-field sequence of +P, -V, -P, +V, +P in the 90° direction. A consequence of this experimental paradigm is that the *same* position-dependent force-fields were always learned across both directions (e.g. +P was the first and fifth force-field learned in both directions, and -P was the third force-field learned in both directions), while *opposite* velocity-dependent force-fields were always learned across both directions (e.g. +V and -V were the second force-fields learned for the 270° and 90° directions, respectively). The consistent position-dependent fields across the two movement directions should not present a problem, but it might be possible the opposite velocity-dependent field in the opposite movement direction could produce interference on top of the interference generated by the previously learned dynamics in the same movement direction. Does learning opposite velocity-dependent force-fields in directions separated by 180° adversely affect learning?

We investigated this question by recruiting an additional 32 subjects, divided into two subgroups. Subgroup 1 learned the same velocity-dependent force-field in the 270° and 90° directions for 120 trials, while subgroup 2 learned the opposite velocity-dependent force-field in the two opposing directions for 120 trials. Error-clamp trials were interspersed with a frequency of 1 in 5 trials to estimate the feedforward adaptation to the velocity-dependent force-fields. The difference was not significant (two-sided unpaired t-test, $p=0.36$, data not shown), and the nominal difference in adaptation coefficients, 0.03, for the two groups was small. Furthermore, we found no significant difference at any point in learning between the two learning curves (two-sided unpaired t-test, minimum $p>0.2$ across all points). These findings strongly suggest that learning opposite viscous force-fields in the 270° and 90° directions did not interfere with learning.

Discussion of Other Data Sets – Bays et al., 2005

Previous work has shown that after adaptation to a clockwise position-dependent force-field (+P), learning a counterclockwise velocity-dependent force-field (-V) can retroactively interfere with the memory of the +P force-field (Bays et al., 2005). However, this interference is only partial – while the interference is greater than if there were no interfering force-field, it is less than the interference caused by learning the opposite position-dependent force-field (i.e. -P) following adaptation to the +P force-field.

Here we show that the viscoelastic primitive model predicts this finding. We simulated the three different experimental paradigms described above, and plotted the gain-space trajectories of the results in Figure S11A-C. In the no-interference group (Group 1, control), subjects learned a +P force-field on Day 1, and then relearned the +P force-field on Day 2. In the -P interference group (Group 2), subjects successively learned +P and -P force-fields on Day 1, and then relearned the +P force-field on Day 2. In the -V interference group (Group 3), subjects successively learned +P and -V force-fields on Day 1, and then relearned the +P force-field on Day 2. We assumed 45% retention of learning from Day 1 to Day 2, based on the results of a previous study measuring 24 hr retention after adaptation to a velocity-dependent force-field (Joiner and Smith, 2008). When comparing the +P force-field adaptation curves for Day 2 (Figure S11E), we find that (as in Bays et al., 2005) the -V interference group (magenta curve) displays worse performance than the no interference group (red curve) but better performance than the -P interference group (green curve).

Figure S11. Simulation of Bays et al. 2005 Experiment

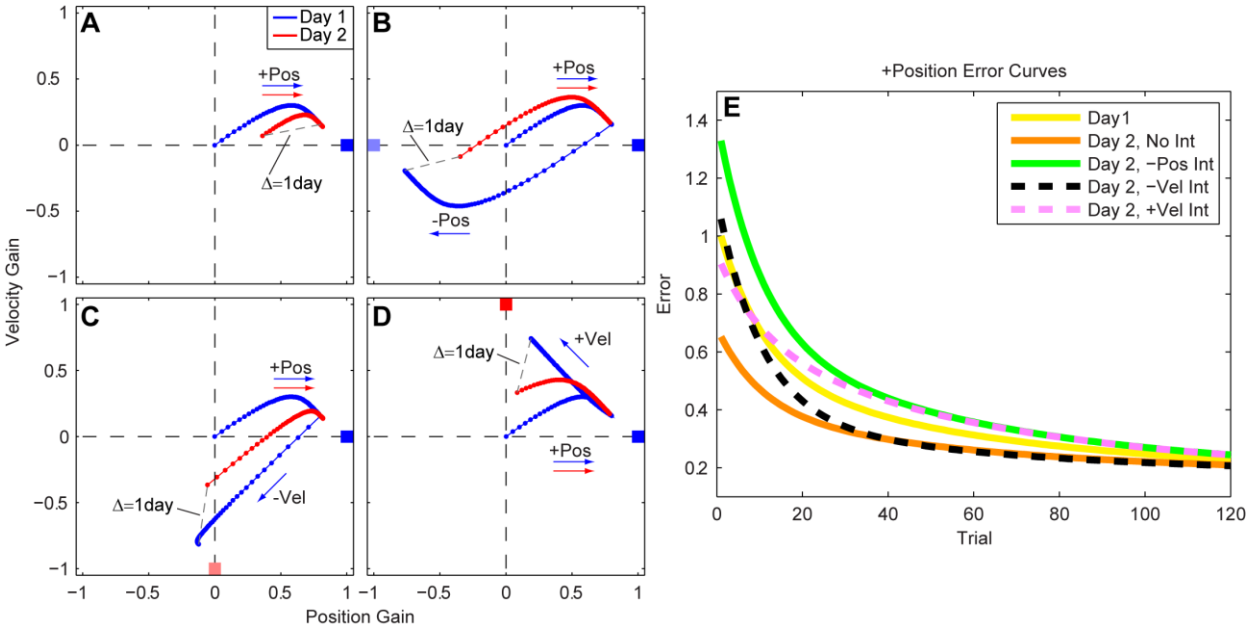

(**A**) Group 1 subjects learned a +P force-field (blue square) on Day 1, and then relearned the +P force-field on Day 2. Each dot represents the gain-space representation of motor output on each trial. The blue dots represent the learning on Day 1, while the red dots represent the learning on Day 2. The gray dotted line connecting the motor output at the end of Day 1 to the motor output at the beginning of Day 2 represents the amount of learning that decayed away. This 55% percent decay was estimated from Joiner and Smith, 2008. (**B**) Group 2 subjects successively learned +P and -P force-fields (light blue square) on Day 1, and then relearned the +P force-field on Day 2. (**C**) Group 3 subjects successively learned +P and-V force-fields (light red square) on Day 1, and then relearned the +P force-field on Day 2. (**D**) Bays et al. 2005 did not study this potential experiment. The subjects in this group would successively learn +P and +V force-fields (red square) on Day 1, and then relearn the +P force-field on Day 2. (**E**) Error curves while exposed to the +P force-field.

The Bays et al. study concluded that the level of interference between adaptations to different dynamics depends on the degree to which their representations conflict in working memory. However, our modeling results go beyond this previous work and predict that interference between position- and velocity-dependent force-fields will not be monolithic. Instead, the amount and pattern of interference 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 force-field (+V) following exposure to a CCW position-dependent force-field (-P) requires both unlearning the -P field and learning the +V, corresponding to a change in dynamics of Δ=[+P,+V]. On the other hand, adaptation to a -V force-field after exposure to a -P force-field requires a change in dynamics of Δ=[+P,-V]. Since the Δ=[+P,+V] change for the -P→+V transition is positively-correlated and aligned with the direction of maximal variance for the primitive distribution, while the Δ=[+P, -V] change is 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). This prediction suggests that if Bays et al. had used a +V (rather than -V) force-field to interfere with the +P force-field, then they would have observed a substantially greater amount of interference (Figure S11DE, compare pink and black dotted curves).

Discussion of Other Data Sets – Goodbody and Wolpert, 1998

A previous study instructed subjects to learn a velocity-dependent force-field at a particular (slow) velocity until asymptote, and then tested the generalization of this adaptation to a faster (but same length) movement. They proposed, among other hypotheses, the idea that the motor controller might learn the velocity-dependent force-field *not* as a function of velocity, but instead as a nonlinear function of position (given that there is a consistent, nonlinear relationship between position and velocity). If this were true, then after reaching asymptotic learning, subjects should produce the same compensatory motor

output for both slow and fast 15 cm movements because both of these movements traverse the same set of positions. However, they found that fast movements displayed larger forces, indicating that the learned dynamics were represented as a linear function of velocity rather than a nonlinear function of position.

At first glance, these findings appear to be at odds with our current work, which shows that substantial cross-adaptation occurs between position- and velocity-dependent dynamics (i.e. velocity-dependent dynamics are learned both as a function of position and velocity). However, note that the Goodbody and Wolpert study looked at the representation of learned dynamics only after learning had reached asymptote. Our current findings show that the inappropriate cross-adaptation, which is present during early learning, substantially disappears by the time asymptotic learning is achieved so that a velocity-dependent force-field has an almost entirely velocity-dependent representation at asymptote, which is consistent with their findings. Note that a non-linear position-dependent representation does not occur in our model because the learning elements display linear responses to both position and velocity signals.

Goodbody and Wolpert's finding that a velocity-dependent force-field is not learned as a nonlinear function of position could result from (a) an inability to represent dynamics as a function of position or (b) a relative inability to form non-linear representations. In Experiments 3-5, we show that subjects can indeed learn linear, position-dependent dynamics, suggesting the second explanation - that the motor system strongly prefers linear representations of dynamics – is correct. This provides some justification for the linear nature of the learning elements in the viscoelastic primitive model.

Discussion of Other Data Sets – Fine and Thoroughman, 2006

Previous work reported that the single-trial adaptive responses following unexpected force pulses administered over a range of different positions during reaching arm movements (Figure S12A) were remarkably similar to each other (Fine & Thoroughman, 2006), even though the kinematic errors experienced on the perturbed trials were very different (Figures 1B & 1D from this paper illustrate these points). We applied the viscoelastic primitive model to this data set in order to determine whether it could account for these findings.

Figure S12. Viscoelastic Primitive Model Accounts for Stereotypic Responses to a Range of Force Pulses.

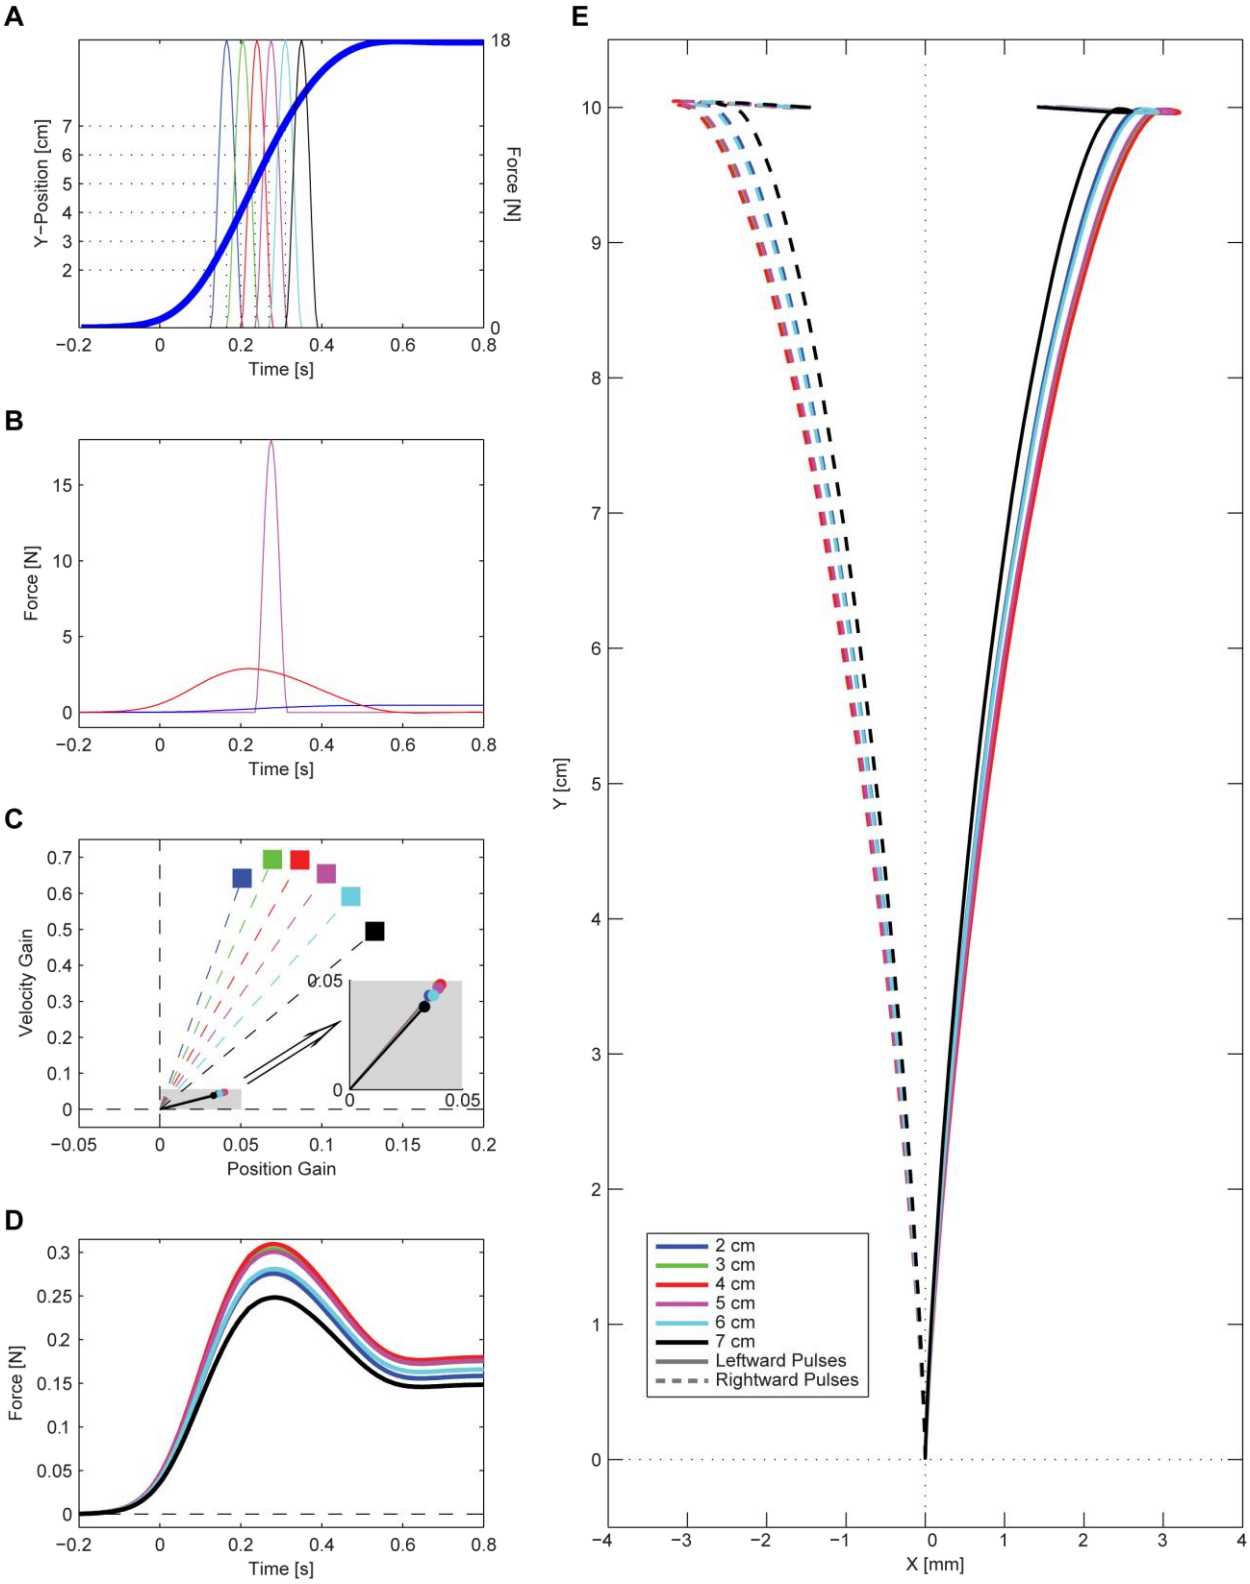

These simulations reproduced data from Fine and Thoroughman, 2006 in which force pulses in the direction perpendicular to motion were applied at 2, 3, 4, 5, 6, and 7 cm into the movement, as illustrated in panel (**A**). These force pulses were 70 ms in duration and 18 N in magnitude. The thick blue line represents the average position trace. The left y-axis indicates the y-position, the horizontal dotted lines indicate the y-position at which the pulses were applied, and the vertical dotted lines indicate the corresponding time points. The right y-axis indicates the magnitude of these force pulses. (**B**) Projection of one of these force pulses (the 5 cm pulse) into position- and velocity-dependent components shows that these components are much smaller in magnitude than the original force pulse. (**C**) The position and velocity components of the "equivalent" viscoelastic force-fields for all six force pulses are transformed into the normalized position/velocity gain-space used in the main paper, and the motor output after a single trial of exposure to these force-fields are simulated. The inset shows a magnified view of the single-

trial learning. (**D**) The force patterns associated with the simulated single-trial learning for all the "equivalent" force-fields. (**E**) The aftereffect trajectories that would result from the force profiles illustrated in panel D. We used a standard model for the mechanics of the human arm and robot arm, as well as the arm stiffness and viscosity that humans display (Smith and Shadmehr, 2005), to simulate these trajectories.

We hypothesized that the viscoelastic primitive model could account for these results despite the fact that these force pulses cannot be well-characterized by their position- and velocity-dependent components. This would be true if the adaptive responses to position-dependent and velocity-dependent dynamics are much stronger than adaptive responses to the other dynamics present in these force pulses. As expected, we found that the position- and velocity-dependent components of these force pulses were quite small compared to the overall size of the pulses, as shown in Figure S12B. We determined the magnitudes of these components by projecting the force pulse onto the position- and velocity-axes in gain-space. This can be accomplished by calculating the dot products of the force pulse patterns with unit vectors in the direction of the position and velocity traces, respectively. Not surprisingly, force pulses administered later in the movement expressed more position-dependence, and force pulses administered near the middle of the movement expressed more velocity-dependence, as illustrated by the colored squares in Figure S12C. Nevertheless, we found that the position- and velocity-dependent components are both positive for all force pulses (Figure S12C). We then simulated how the viscoelastic primitive model would adapt to each of these "equivalent" force-fields for a single trial (inset of Figure S12C), in order to estimate the force patterns associated with single-trial adaptation (Figure S12D). Note that the predicted *shapes* of adapted motor output patterns are nearly identical (r=0.9998, minimum correlation coefficient between these force patterns). The positive-combination force-fields inducing this adaptation are already close in position/velocity gain-space, and the resulting adaptations are even closer (Figure S12C). Since Fine & Thoroughman, 2006 looked at the aftereffects to study this single-trial adaptation, we used a standard model for the mechanics of the human arm and robot arm, as well as the arm stiffness and viscosity that humans display (Smith and Shadmehr, 2005), to simulate the trajectories (Figure S12E) that would result from the adaptive changes to the baseline force patterns illustrated in Figure S12D. The aftereffects predicted by the viscoelastic primitive models (Figure S12E) display remarkably close correspondence to the experimental data from Fine & Thoroughman, 2006 (see Figure 1D in their paper). Note that the format of this figure is the same as the format of our Figure S12E – each color represents the aftereffect associated with an 18 N, 70 ms force pulse applied at a particular distance into the movement, and the dashed lines represent aftereffects associated with rightward as opposed to leftward force pulses. The shapes of the aftereffect trajectories are all very similar to one another (as shown in F&T Figure 1D) despite the pulsed-trial trajectories being quite different (F&T Figure 1B). Our model accounts for not only the stereotypic shape of these aftereffect trajectories across different perturbation positions, but their magnitudes as well. We find it somewhat remarkable that the viscoelastic primitive model accurately predicts the adaptive responses to perturbations for which the force patterns are only poorly characterized by their position- and velocity-dependent components. Despite the fact that these components do not account for the majority of the force pulse applied, they appear to account for the adaptive response to the force pulse quite well. This underscores the idea that the motor system adapts much more strongly to state-dependent dynamics than non-state-dependent dynamics.