

Exploration of joint redundancy but not task space variability facilitates supervised motor learning

Puneet Singh^a, Sumitash Jana^b, Ashitava Ghosal^{c,1}, and Aditya Murthy^{b,1}

^aCentre for Biosystems Science and Engineering, Indian Institute of Science, Bangalore 560012, India; ^bCentre for Neuroscience, Indian Institute of Science, Bangalore 560012, India; and ^cDepartment of Mechanical Engineering, Indian Institute of Science, Bangalore 560012, India

Edited by Michael E. Goldberg, Columbia University College of Physicians, New York, NY, and approved October 27, 2016 (received for review August 15, 2016)

The number of joints and muscles in a human arm is more than what is required for reaching to a desired point in 3D space. Although previous studies have emphasized how such redundancy and the associated flexibility may play an important role in path planning, control of noise, and optimization of motion, whether and how redundancy might promote motor learning has not been investigated. In this work, we quantify redundancy space and investigate its significance and effect on motor learning. We propose that a larger redundancy space leads to faster learning across subjects. We observed this pattern in subjects learning novel kinematics (visuomotor adaptation) and dynamics (force-field adaptation). Interestingly, we also observed differences in the redundancy space between the dominant hand and nondominant hand that explained differences in the learning of dynamics. Taken together, these results provide support for the hypothesis that redundancy aids in motor learning and that the redundant component of motor variability is not noise.

supervised learning | minimum-intervention principle | reaching | motor noise | motor control

Redundancy is a ubiquitous property that renders biological systems robust to disruptions. Goal-directed movements also display redundancy because a given movement, such as touching one's nose, can be made in many different ways with different combinations of joint angles. Although redundancy generates flexibility, it also poses a fundamental problem for the motor system because a large component of motor variability is attributed to muscle noise (1). Thus, if muscles operated independently, the total noise (reflected in the trajectory) would be a summation of the noise due to the component parts. However, if the covariation of muscles or joints occurs such that the effect of individual muscle or joint variability is mitigated, it is possible to maintain acceptable levels of variability while maintaining a reasonable degree of redundancy that allows flexibility in behavior. Consistent with this view, it has been observed in a wide range of tasks (2–7) that variability is not eliminated but optimized (8–11) to accumulate in a task-relevant dimensions using a minimum-intervention principle (12). Such variability, that is a consequence of redundancy, can be quantified as an uncontrolled manifold (13–15) in which task-independent variability is constrained to a redundant subspace (or “uncontrolled manifold”).

Although minimizing variability is expected to improve task-related performance, recent evidence suggests that motor variability paradoxically helps in motor learning (16–20). Such findings are supported by ideas in reinforcement learning (21, 22), which suggest that baseline variability increases exploration, which in turn facilitates learning. Interestingly, motor variability has also been shown to help learning during supervised error-based learning tasks (16), suggesting a more general role of variability in motor learning. In this study, we tested whether variability arising from joint redundancy plays a role in supervised motor learning.

Results

We used two experimental setups shown in Fig. S1A. The subject moves the end-effector of a robot manipulator or hand from an initial point to a task-space target point. As shown in Fig. 1, the experiment had three phases—a preadaptation baseline period, followed by a phase with either one of two kinds of perturbations: a visuomotor (kinematic perturbation) or an applied viscous curl force (dynamic perturbation) and finally a postadaptation phase when the perturbation was removed. We simultaneously measured the end point and joint angles while subjects reached to the target during the baseline period as shown in Fig. S1B. The map between the joint angles and the end point (x, y) point is many-to-one, i.e., there is redundancy. The joint variability in the baseline period was quantified into two components—the joint variability due to the redundancy termed as the null-space variability that did not affect end point and the joint variability that caused changes in end point termed as task-space variability (Fig. S1C). In this work, we studied the effect of these two types of variability on the learning of kinematic and dynamic perturbations. We have studied learning in (i) simple visuomotor adaptations (two directions), (ii) generalized visuomotor adaptation (eight directions), and (iii) generalized force-field adaptation when subjects used their dominant and nondominant hand.

Redundancy in Kinematic Learning. We trained 40 subjects to learn point-to-point reaching movements using their dominant hand, along two directions, in a visuomotor perturbation that was set using Eq. 1. In this experiment, the cursor was rotated by -45° from the hand trajectory (Fig. S2A). The trajectory of the hand

Significance

Theories of reinforcement learning claim that motor variability helps in motor learning and are supported by recent experimental work. In contrast, theories of motor control propose that variability is noise that needs to be suppressed. We attempt to reconcile these apparent contradictory positions. Using the formulation of the unconstrained manifold hypothesis, we show that motor variability has two components—a part arising out of the redundancy that does not affect task-space and another component related to task-space variability. We show that the motor variability component resulting from the redundancy determines both dynamic and kinematic learning ability across subjects without affecting task-space variability.

Author contributions: P.S., S.J., A.G., and A.M. designed research; P.S. and S.J. performed research; P.S., A.G., and A.M. contributed new reagents/analytic tools; P.S. analyzed data; and P.S., A.G., and A.M. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

¹To whom correspondence may be addressed. Email: asitava@mecheng.iisc.ernet.in or aditya@cns.iisc.ernet.in.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1613383113/-DCSupplemental.

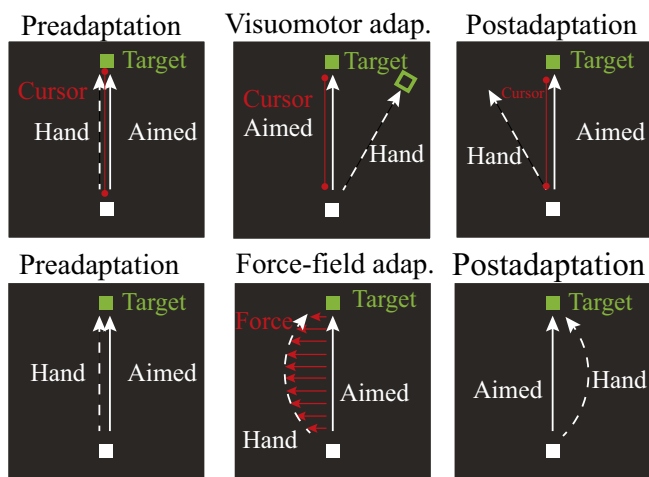


Fig. 1. Experiment setup and design. Experiments were divided into preadaptation (baseline), adaptation, and postadaptation (washout) epochs. Subjects adapted to a visuomotor perturbation (Top) or a force-field perturbation (Bottom) in separate experiments.

for a typical subject in the preadaptation (baseline) (Fig. S2B), visuomotor adaptation (Fig. S2C), and postadaptation epochs (Fig. S2D) are shown. Overall, the pattern of trajectories are consistent with previous work showing that whereas typical movements follow a nearly straight trajectory in the baseline condition, they show strong curved trajectories in the presence of a visuomotor perturbation. The curved trajectories gradually become straighter with practice over the course of about 60 trials (Fig. 24). In addition, as a consequence of motor learning, subjects showed a washout effect (postadaptation) where errors in trajectory invert direction when the learned visuomotor perturbation is turned off.

To quantify the error, we used the error at peak velocity along the trajectory. The reduction in peak velocity error (Eq. 9) was used as a metric to quantify the learning rate for each subject. To test whether the learning rate of a subject could be predicted on the basis of motor redundancy exhibited during the preadaptation (baseline period), we computed $N(J)$ (Eq. 8), the chosen measure of variability due to redundancy space called null-space variability. We divided the population into two groups based on their learning rate (above-mean and below-mean learning group), as quantified by the fitted exponential β value (Fig. 2B). In support of the hypothesis, the null-space variability was significantly different between the two learning groups (Fig. 2C). Interestingly, we also found a strong positive correlation between baseline null-space variability and learning rate in the visuomotor adaptation task (Fig. 2D; $r = 0.55$, $P = 0.0002$). However, we found no correlation between baseline task-space variability and learning rate ($r = 0.13$, $P = 0.43$).

Redundancy in Generalized Kinematic Learning. To test whether redundancy could aid in learning a generalized task, we next examined 10 subjects while they learned point-to-point reaching movements with the visuomotor perturbation (Eq. 1) along eight directions, where in each case, the cursor was rotated by 45° from the hand trajectory (Fig. S2E). The trajectory of the hand for a typical subject in the preadaptation (baseline) (Fig. S2F), visuomotor adaptation (Fig. S2G), and postadaptation epochs (Fig. S2H) are shown. Each individual subject's learning curve is shown in Fig. S3, and a representative subject is shown in Fig. 3A. The average behavior pooled across 10 subjects showed a similar learning pattern (Fig. 3B; goodness of fit, $r^2 = 0.90$). Again, the learning rate in the visuomotor perturbation and the

null-space and task-space variability in the preadaptation baseline period was computed. A significant correlation between the null-space variability with learning rate was observed (Fig. 3C; $r = 0.71$, $P = 0.021$). Interestingly, we found no correlation of the baseline task-space variability with the learning rate (Fig. 3D; $r = 0.42$, $P = 0.22$).

Redundancy and Dynamic Learning. To test whether redundancy could aid in learning in other types of perturbation, we trained 10 subjects to learn point-to-point reaching movements using their dominant hand, along 8 directions, in a force-field that was set using the force-field perturbation defined by Eq. 2. In this experiment, the perturbation was proportional to the velocity of the hand but perpendicular to the hand movement direction (Fig. S2I). The trajectory of the hand for a typical subject in the preadaptation (baseline) (Fig. S2J), force-field adaptation (Fig. S2K), and postadaptation epochs (Fig. S2L) are shown. Similar to the visuomotor perturbation, and consistent with literature, typical hand movements follow a nearly straight trajectory in the baseline condition, and they show strong curved trajectories in the presence of a viscous curl force-field. The curved trajectories gradually become straighter with practice over the course of about 200 trials (Fig. 4A). In addition, as a consequence of motor learning, subjects showed a washout effect, where errors in trajectory invert in direction when the learned force-field was turned off in the postadaptation period. This washout error converged to baseline levels within approximately 100 trials. The average behavior pooled across the 10 subjects showed a similar learning pattern (Fig. 4B; goodness of fit, $r^2 = 0.93$). The

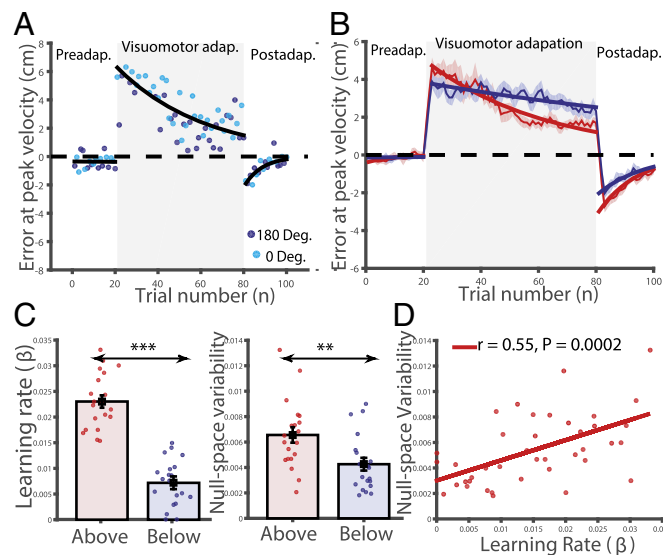


Fig. 2. Joint redundancy predicts learning rate in a visuomotor adaptation task. (A) Error at peak velocity in preadaptation, visuomotor adaptation, and postadaptation for a representative subject showing the progression of adaptation. Errors in each of the two directions are color-coded. Fitted exponential (black line) significantly accounts for most of the progression of errors across trials during adaptation. (B) Three-trial running mean \pm SEM across subjects (shading is SEM). Fitted exponential curves across subjects significantly account for most of the progression of errors during adaptation (red indicates above-mean learning group; blue indicates below-mean learning group). (C) Comparison of baseline null-space variability with learning rate between the above-mean learning group (red) and below-mean learning group (blue) reveal corresponding differences in null-space variability between groups. (D) Subject-by-subject comparison ($n = 40$) of baseline null-space variability with learning rate shows a significant positive relationship. Two subjects whose learning rates were negative have been clamped to 0. Asterisks indicate statistically significant differences ($*P < 0.05$, $**P < 0.005$, $***P < 0.0005$).

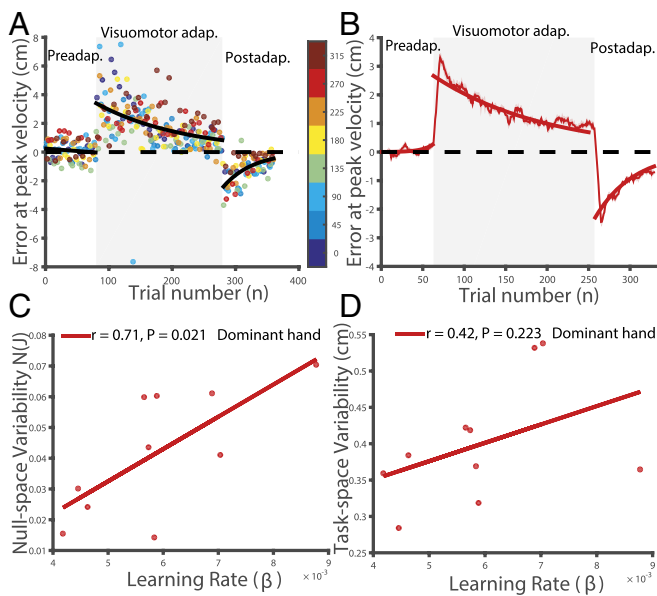


Fig. 3. Joint redundancy predicts the learning rate in a generalized visuomotor adaptation task. (A) Error at peak velocity in preadaptation, adaptation, and postadaptation showing the progression of adaptation for a subject. Errors in each of the eight directions are color-coded. (B) Eight-trial running mean \pm SEM across subjects (shading is SEM). Fitted exponential curves across subjects significantly account for most of the progression of errors in the adaptation. (C) The comparison of baseline null-space variability at peak velocity with learning rate shows a significant positive relationship between joint redundancy and motor learning. (D) The comparison of baseline task-space variability at peak velocity with learning rate shows no significant relationship between variability and motor learning.

learning curves of each individual subject with the dominant and nondominant hand are shown in Figs. S4 and S5, respectively. As before, the reduction in error at peak velocity (Eq. 9) was used as a metric to quantify the learning rate for each subject. In the force-field adaptation period, the mean learning rate was 0.006 ± 0.002 for the dominant hand and 0.004 ± 0.001 for the nondominant hand. To test whether the learning rate of a subject under a viscous perturbing force could be predicted on the basis of motor redundancy exhibited during the preadaptation (baseline period), we again computed null-space variability as in Eq. 8. We found a strong positive correlation between baseline null-space variability and learning rate in the force-field [Fig. 4C; $r = 0.69$, $P = 0.027$ (dominant hand), $r = 0.60$, $P = 0.067$ (nondominant hand)]. Consistent with that observed for the kinematic perturbation task, we found poor correlation between the learning rate and task-space variability in the baseline period [Fig. 4D; $r = 0.22$, $P = 0.53$ (dominant hand), $r = 0.13$, $P = 0.72$ (nondominant hand)].

Redundancy and Dominant and Nondominant Hand Learning. Differences between the learning of the dominant and nondominant arm across subjects is consistent with previous work (23) and was used as an additional approach to test the relation between null-space variability and motor learning. To test this notion, we compared the learning rate and null-space variability between the dominant and nondominant hand in 10 subjects. We observed that the mean learning rate for the nondominant hand (mean, 0.004 ± 0.001) was significantly less than the mean learning rate for the dominant hand (mean, 0.006 ± 0.002) (Fig. 5A; $P = 0.008$, $t(8) = 3.54$). Interestingly, null-space variability was also lesser in the nondominant hand (mean, 0.054 ± 0.036) compared with the dominant hand (mean, 0.12 ± 0.066) (Fig. 5B; $P = 0.035$, $t(8) = 2.54$). We observed no difference in

the mean task-space variability between the dominant and nondominant hand, suggesting that task-space variability did not influence learning rate. Furthermore, we observed a strong correlation (Fig. 5D; $r = 0.84$, $P = 0.002$) between the difference in the learning rate and the difference in the null-space variability of the dominant and nondominant hand, suggesting that the difference in the null-space variability could partly explain the difference in learning rate. We also observed an outlier subject whose learning rate was higher in the nondominant hand compared with the dominant hand (marked as a dotted line in Fig. 5A–C and marked in blue in Fig. 5D). Nevertheless, even for this subject, the null-space variability was greater in the nondominant hand compared with the dominant hand, in support of our hypothesis. Taken together, these findings indicate that the difference in learning rate between the dominant and nondominant hand may be a consequence of the greater redundancy in the dominant hand.

Discussion

In contrast to previous work that has studied joint redundancy and learning in isolation, this study tested the relation between these two variables under the assumption that the extra degrees of freedom conferred by the arm is used by the motor system to facilitate learning. We have shown that task-space variability in the baseline period while reaching a target in task-space has low correlation with learning during perturbations, whereas the variability in the null-space, resulting from redundancy, aids in learning. We interpret these results as indicating that exploration of redundancy aids in motor learning when a visuomotor perturbation or a force-field perturbation is present.

Joint Redundancy. The uncontrolled manifold hypothesis (UCM), that has its origin in the initial observations by Bernstein (24, 25)

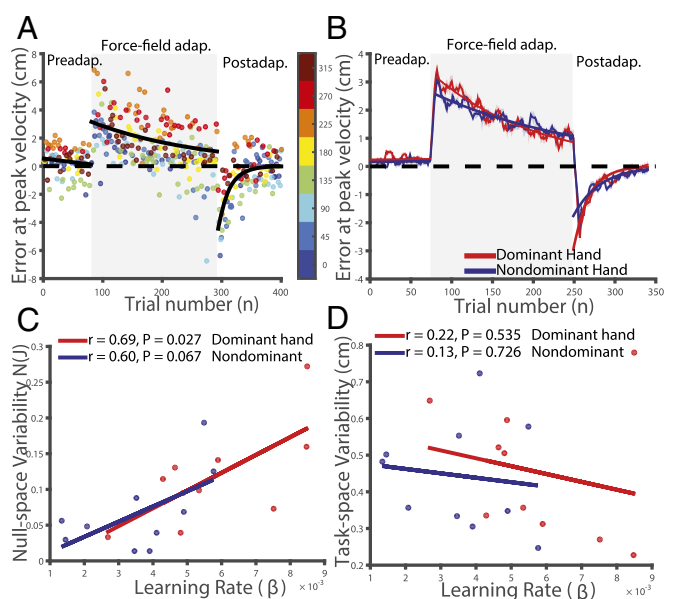


Fig. 4. Joint redundancy predicts rate of learning in a force-field adaptation task. (A) Error at peak velocity in preadaptation, adaptation, and postadaptation showing the progression of adaptation for a subject. Errors in each of the eight directions are color-coded. (B) Eight-trial running mean \pm SEM across subjects (shading is SEM). Fitted exponential curves across subjects significantly account for most of the progression of errors in the adaptation (red indicates dominant hand, blue indicates nondominant hand). (C) Comparison of baseline null-space variability at peak velocity with learning rate shows a relationship. (D) Comparison of baseline task-space variability at peak velocity with learning rate shows no significant relationship between task-space variability and motor learning.

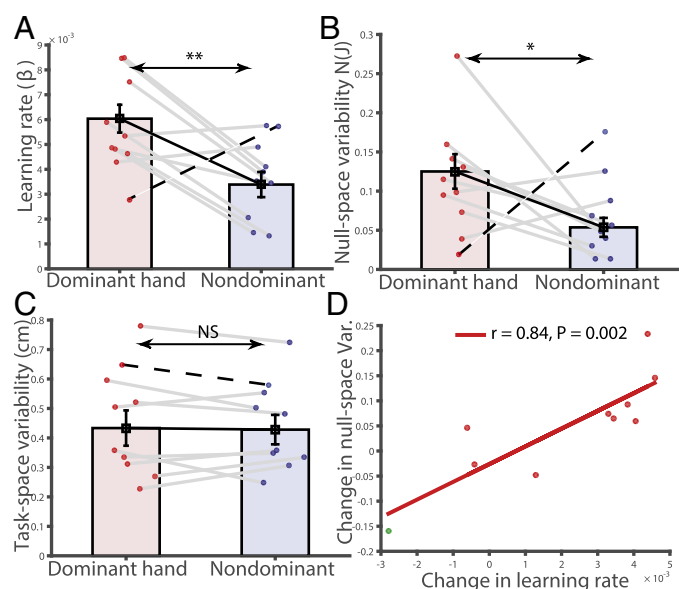


Fig. 5. Learning differences between the dominant hand and nondominant hand. (A) Learning differences in the dominant hand (red) and nondominant hand (blue) ($n = 10$) reveal faster learning in the dominant hand. (B) Baseline null-space variability in the dominant and nondominant hand reveals differences in null-space variability between hands. (C) However, baseline task-space variability in the dominant hand and nondominant hand reveals no differences in task-space variability between hands. (D) Correlation between the difference in null-space variability with the difference in learning rate between the dominant hand and nondominant hand shows a significant positive relationship. The outlier subject data are shown as a dotted line (A–C) and as a green dot (D).

is the dominant framework to understand and quantify joint redundancy (13–15). Such redundancy is now established as a ubiquitous feature of behavior observed across a variety of tasks (2–7). In the current study we followed the UCM framework to quantify joint redundancy. However, unlike previous work, we also quantified redundancy in individual subjects in addition to measuring the group response and observed large variability, suggesting that redundancy might be an idiosyncratic feature that is unique to each subject and/or session. In addition, unlike past work where redundancy was quantified at the maximum peak velocity or at the target location, we restricted our computation to the former in case of all perturbations. This choice was made because task-space variability is known to be highest at the peak velocity of the trajectory and smallest at the end point (target), particularly when the targets are small. Thus, quantifying redundancy at peak velocity is better suited to reveal the full scale of variability across subjects, which may be essential to understand its bearing on motor learning.

Previous work has suggested that the degree of redundancy can be task specific and can be optimized such that the motor system obeys the principle of minimum intervention in which the brain only controls task relevant variability but does not control task irrelevant or null-space variability. Although we did not explicitly study the control of null-space variability, we did observe interesting task specific differences indicating that null-space variability or task irrelevant variability also may be actively controlled and is not merely an epiphenomena of having more degrees of freedom than required for the task. For example, in our pool of subjects we observed that null-space variability was on average greater in the dynamics condition (mean, 0.11 ± 0.069) than that in the kinematic condition (mean, 0.04 ± 0.020) even in the preadaptation period when the reaching task was identical. This difference in null-space variability

is likely to be task-specific, the basis of which remains to be elucidated.

Inline with earlier work (23) that showed differences between the dominant and nondominant hand during force-field perturbation but not during visuomotor (kinematics) perturbation, we tested null-space variability in the dominant versus nondominant hand in the dynamic perturbation task. It is particularly interesting to note that despite similar force application and task-space variability (Fig. 5C), the null-space variability was greater in the dominant hand, and further, differences within subjects correlate with corresponding differences in the learning rates (Fig. 5D). The larger null-space variability seen in the dominant hand may thus provide a natural explanation of why learning might be more potent in the dominant hand and reinforces the belief that null-space variability not only reflects the biomechanical characteristics of the arm but may reflect active control.

Dynamic and Kinematic Learning. To study motor learning, we followed previous work that has tested the ability of subjects to implicitly adapt their motor behavior in the presence of dynamic and kinematic perturbations (26–29). In our learning paradigm, subjects learned the perturbations while making movements in two directions (Fig. 2) and also in all eight directions (Fig. 3), picked at random. Hence, unlike some learning paradigms that emphasize specific learning in one direction, our learning is expected to generalize across directions. Thus, we used a single exponential fit pooled across all directions to study the average rate of learning as a single variable even though our data suggest the presence of fast and slow learning phase that has been reported in the literature (30). In the future, we hope to study direction-specific motor learning to test whether joint redundancy better correlates with the fast versus slow learning phase. Nevertheless, our results revealed learning rates that are comparable to the literature (26–28, 31). The data also revealed that learning in the dominant hand was significantly faster than that in the nondominant hand (23). Second, trends indicate that some directions appear to be easier to learn, and like joint redundancy, there is a large subject-specific variability in the data, the implications of which are discussed in *Relating Joint Redundancy and Motor Learning*.

Relating Joint Redundancy and Motor Learning. The strong subject-wise correlations observed in both the dynamic and kinematic learning tasks support the hypothesis that joint redundancy supports motor learning. Although our data are fundamentally correlative in nature, we were able to exploit a feature in our experiment that involved the use of the dominant and nondominant hand that resulted in differences in learning rates and redundancy. Lending further credence to the hypothesis, we found that the smaller redundancy in the nondominant hand was associated with slower learning. Moreover, the differences in learning and redundancy were also correlated (Fig. 5D). In addition, we showed similar group-wise differences in the visuomotor adaptation task showing smaller redundancy for below-mean learners compared with above-mean learners. This result notwithstanding, we do not claim that redundancy is the sole source of differences in motor learning. Kinematic learning is faster than dynamic learning despite smaller redundancy in the kinematic task compared with the dynamic task, and is likely to reflect differences in the mechanisms involved in learning these two perturbations that involve the learning of different internal models (27), with joint redundancy being a common factor that confers greater flexibility to explore motor space. Additionally, we also observed differences in the degree of redundancy and learning across directions. Although these differences were not statistically significant (they were also poorly correlated), these results suggest that redundancy not only possesses an active component that correlates with learning

but also a passive component that reflects differences in the biomechanics.

Although we observed strong correlations between joint redundancy and motor learning, our results did not show a correlation between task-space variability and motor learning. In contrast, recent work by Wu et al. (16), using both reinforcement learning and error-based supervised learning, emphasizes the selective role of task relevant variability in motor learning. Although the basis of differences between these studies remain unclear, it is interesting to note that other studies suggest that task variability does not correlate with motor learning (see ref. 32 for a meta-analysis). Further research would be required to reconcile these points of view. In addition, although it is not still mechanically clear how joint redundancy facilitates motor learning, we suggest that such active exploration of task-irrelevant space may be essential to motor learning, while simultaneously ensuring optimal motor performance by minimizing task-space variability. This possibility notwithstanding, we believe that this relation between null-space variability and motor learning can be leveraged to enhance motor learning skills and motor rehabilitation.

Materials and Methods

Subjects. All subjects were paid for participation and gave informed consent in accordance with the guidelines of the Human Ethics Committee of the Indian Institute of Science that approved the protocol. Sixty naive subjects participated in the study (aged 22–70 y; 42 male and 18 female). In experiment 1, 40 subjects (30 male and 10 female, all right-handed) performed a visuomotor task and in experiment 2, 10 subjects (5 male and 5 female; all right-handed) performed a generalized visuomotor task. In experiment 3, 10 subjects (7 male and 3 female; 6 right-handed, 4 left-handed subjects) performed the force-field experiment first with their dominant and after a gap of ~5 d with their nondominant hands. Handedness was tested using the modified Edinburgh Handedness Index (33).

Experimental Setup. In experiment 1, subjects sat on a chair with their hand placed on the table (Fig. S1A). A monitor that displayed the targets and the cursor movement was placed in front of the subjects. The experiment was performed using the psychophysics toolbox (34) that displayed visual stimuli, sampled, and stored the data and other behavioral parameters. Hand positions and joint angles were recorded using an electromagnetic position and orientation tracking device (Liberty; Polhemus).

In all other experiments, subjects sat on a chair while their chins were supported by a chin rest, and their heads were locked with head bars on both sides of their temple as shown in Fig. S1A. Subjects looked down on a semitransparent mirror on which they saw the targets while they moved a robotic arm handle (KINARM End-Point Robot; BKIN) in a horizontal plane below the plane of the mirror. Targets were presented by an inverted monitor above the mirror setup, which gave the impression that the targets appeared in a virtual plane below the mirror aligned to the plane in which the robotic arm handle moved. All experiments were performed using TEMPO/VideoSYNC software (Reflective Computing) that displayed visual stimuli and sampled and stored the data and other behavioral parameters in real time at a resolution of 1.04 ms. Hand positions and joint angles were recorded (spatial resolution of 7.62 mm) using an electromagnetic position and orientation tracking device (Liberty; Polhemus) interfacing with TEMPO in real time at 240 Hz.

Experimental Paradigm. In experiment 1, trials were divided into three phases: baseline or preadaptation, adaptation, and postadaptation. All subjects performed ~10 practice trials. Subjects performed about 100 trials per session, with a typical session lasting about 15 min. Each trial started with the presentation of a square fixation box (1 cm) at the center of a screen where subjects had to fixate both their hand and eye. After successful fixation, a square target of length 1 cm was displayed randomly in any one of two locations that uniformly spanned a circle of 20-cm radius around the central fixation box. Trials were aborted if a premature movement was made. Auditory feedback was given when the subject performed a trial correctly.

In all other experiments, trials were divided into three phases: baseline or preadaptation, adaptation, and postadaptation. All subjects performed ~30 practice trials. Subjects performed about 400 trials per session, with a

typical session lasting between 1.5–2 h. Each trial started with the presentation of a square fixation box (0.4 cm) at the center of a screen where the subjects had to fixate both their eye and the robotic end-effector. After successful fixation, a square target (0.7 cm) was displayed randomly in any one of eight locations that uniformly spanned a circle of 15-cm radius around the central fixation box. Similar to experiment 1, trials were aborted if a premature movement was made. Auditory feedback was given when the subject performed a trial correctly.

To minimize the effect of transfer of learning (35), subjects were always exposed to washout trials after perturbation trials. In addition, typically and ~5-d gap was given between recording the dominant and nondominant hand.

Visuomotor Perturbation. During visuomotor perturbation (29), the cursor movement is rotated according to Eq. 1,

$$\begin{bmatrix} P_x \\ P_y \end{bmatrix} = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} \begin{bmatrix} p_x \\ p_y \end{bmatrix}, \quad [1]$$

where P_x, P_y correspond to the position of the cursor, p_x, p_y correspond to the actual position of the hand, and θ denotes the perturbation angle about the center of work space with θ equal to 45° . This perturbation also led to a trajectory error and to compensate the subjects altered their hand trajectory.

Force-Field Perturbation. During force-field perturbation (31), the robot produced viscous curl forces depending on the instantaneous hand velocity as in Eq. 2,

$$\begin{bmatrix} F_x \\ F_y \end{bmatrix} = \begin{bmatrix} 0 & -K \\ K & 0 \end{bmatrix} \begin{bmatrix} \dot{x} \\ \dot{y} \end{bmatrix}, \quad [2]$$

where F_x, F_y correspond to the forces exerted on the robotic arm, \dot{x}, \dot{y} correspond to the velocity components of hand, and K denotes the force perturbation coefficient along the orthogonal directions with K equal to 20 Ns/m. This force-field perturbation disturbed the hand trajectory.

Quantifying Redundancy. To calculate the joint-space variability in different directions, a 2D forward-kinematics model was created for the human arm. The model is given by

$$\begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} l_1 \cos(\theta_1) + l_2 \cos(\theta_2) + l_3 \cos(\theta_3) + l_4 \cos(\theta_4) \\ l_1 \sin(\theta_1) + l_2 \sin(\theta_2) + l_3 \sin(\theta_3) + l_4 \sin(\theta_4) \end{bmatrix}, \quad [3]$$

where the four joints rotations are clavicle protraction–retraction (θ_1), shoulder horizontal abduction–adduction (θ_2), elbow flexion–extension (θ_3), and wrist medial–lateral (θ_4) (Fig. S1B). In experiment 1, we also incorporated a fifth joint-index finger abduction–adduction, and accordingly extended Eqs. 3 and 5. The distribution of variability in joint-space was computed for baseline trials for each of the different directions at the maximum reach velocity. The mean joint configuration across trials, along each of the directions, was computed at the peak velocity v and is denoted by $\bar{\theta}^v$. The deviation of the joint configuration for a trial k , $\Delta\theta_k$, is obtained by subtracting the joint configuration at the peak velocity, θ_k^v , from the mean as below:

$$\Delta\theta_k = \bar{\theta}^v - \theta_k^v. \quad [4]$$

Based on the 2D forward-kinematics model, the Jacobian matrix at peak velocity was computed as

$$J(\bar{\theta}^v) = \begin{bmatrix} \frac{\partial x}{\partial \theta_1} & \frac{\partial x}{\partial \theta_2} & \frac{\partial x}{\partial \theta_3} & \frac{\partial x}{\partial \theta_4} \\ \frac{\partial y}{\partial \theta_1} & \frac{\partial y}{\partial \theta_2} & \frac{\partial y}{\partial \theta_3} & \frac{\partial y}{\partial \theta_4} \end{bmatrix}. \quad [5]$$

The joint configuration vector ξ_i lying in the null-space of the Jacobian matrix was computed from

$$J(\bar{\theta}^v) \xi_i = 0. \quad [6]$$

For each trial, the sum of the component of $\Delta\theta_k$ along the null-space directions is given by

$$\theta_R = \sum_{i=1}^m \langle \Delta\theta_k, \xi_i \rangle \xi_i, \quad m = 2 \text{ or } 3. \quad [7]$$

We quantify redundancy as the sum of the squares of θ_R across all of the trials divided by the number of trials n . Mathematically, this is written as

$$N(J) = \sum_{i=1}^n \frac{(\theta_R)^2}{n}. \quad [8]$$

In this work, the scalar $N(J)$ is used as a measure of redundancy space as θ_R , from which $N(J)$ is derived, lies in the null-space, and does not affect task-space motion (Fig. S1C).

Task-Space Variability. We quantified the task-space variability from the hand trajectory when it reached its peak velocity. The SD of the perpendicular distance of this point from the straight line joining the start and end of the trajectory was used as the metric of task-space variability.

Quantifying Learning. The error at the peak velocity was calculated as the perpendicular distance of the hand trajectory at peak velocity from the straight line joining the central fixation box to the target location. The error, denoted by $f(n)$, is related to the trial number by the following equation.

$$f(x) = ae^{-\beta x}. \quad [9]$$

The above equation represents a first-order learning process, and in this work, we concentrate on β , which represents the intrinsic learning for a subject. To compute the population learning in perturbation trials, errors were fitted with an exponential fit using robust least-squares

- Harris CM, Wolpert DM (1998) Signal-dependent noise determines motor planning. *Nature* 394(6695):780–784.
- Tseng Y, Scholz JP, Schoner G (2002) Goal-equivalent joint coordination in pointing: Affect of vision and arm dominance. *Motor Control* 6(2):183–207.
- Haggard P, Hutchinson K, Stein J (1995) Patterns of coordinated multi-joint movement. *Exp Brain Res* 107(2):254–266.
- Cole KJ, Abbs JH (1986) Coordination of three-joint digit movements for rapid finger-thumb grasp. *J Neurophysiol* 55(6):1407–1423.
- Black DP, Smith BA, Wu J, Ulrich BD (2007) Uncontrolled manifold analysis of segmental angle variability during walking: Preadolescents with and without down syndrome. *Exp Brain Res* 183(4):511–521.
- Yang JF, Scholz JP, Latash ML (2007) The role of kinematic redundancy in adaptation of reaching. *Exp Brain Res* 176(1):54–69.
- Park J, Wu YH, Lewis MM, Huang X, Latash ML (2012) Changes in multifinger interaction and coordination in Parkinson's disease. *J Neurophysiol* 108(3):915–924.
- Todorov E, Jordan MI (2002) Optimal feedback control as a theory of motor coordination. *Nat Neurosci* 5(11):1226–1235.
- Todorov E (2004) Optimality principles in sensorimotor control. *Nat Neurosci* 7(9):907–915.
- Liu D, Todorov E (2007) Evidence for the flexible sensorimotor strategies predicted by optimal feedback control. *J Neurosci* 27(35):9354–9368.
- Diedrichsen J, Shadmehr R, Ivry RB (2010) The coordination of movement: Optimal feedback control and beyond. *Trends Cogn Sci* 14(1):31–39.
- Todorov E, Jordan MI (2002) A minimal intervention principle for coordinated movement. *Advances in Neural Information Processing Systems 2002*, eds Becker S, Thrun S, Obermayer K (Neural Information Processing Systems Foundation, La Jolla, CA), pp 27–34.
- Scholz JP, Schöner G (1999) The uncontrolled manifold concept: Identifying control variables for a functional task. *Exp Brain Res* 126(3):289–306.
- Latash ML, Scholz JP, Schöner G (2002) Motor control strategies revealed in the structure of motor variability. *Exerc Sport Sci Rev* 30(1):26–31.
- Latash ML, Scholz JP, Schoner G (2007) Toward a new theory of motor synergies. *Motor Control* 11(3):276.
- Wu HG, Miyamoto YR, Castro LNG, Ölveczky BP, Smith MA (2014) Temporal structure of motor variability is dynamically regulated and predicts motor learning ability. *Nat Neurosci* 17(2):312–321.
- Herzfeld DJ, Shadmehr R (2014) Motor variability is not noise, but grist for the learning mill. *Nat Neurosci* 17(2):149–150.
- Tumer EC, Brainard MS (2007) Performance variability enables adaptive plasticity of 'crystallized' adult birdsong. *Nature* 450(7173):1240–1244.
- Cashaback JG, McGregor HR, Gribble PL (2015) The human motor system alters its reaching movement plan for task-irrelevant, positional forces. *J Neurophysiol* 113(7):2137–2149.
- Therrien AS, Wolpert DM, Bastian AJ (2016) Effective reinforcement learning following cerebellar damage requires a balance between exploration and motor noise. *Brain* 139(1):101–114.
- Sutton RS, Barto AG (1998) *Reinforcement Learning: An Introduction* (MIT Press, Cambridge, MA).
- Barto AG (1994) Reinforcement learning control. *Curr Opin Neurobiol* 4(6):888–893.
- Sainburg RL (2002) Evidence for a dynamic-dominance hypothesis of handedness. *Exp Brain Res* 142(2):241–258.
- Whiting HTA (1983) *Human Motor Actions: Bernstein Reassessed* (Elsevier, Amsterdam), Vol 17.
- Latash ML (2008) *Synergy* (Oxford Univ Press, New York).
- Shadmehr R, Smith MA, Krakauer JW (2010) Error correction, sensory prediction, and adaptation in motor control. *Annu Rev Neurosci* 33:89–108.
- Krakauer JW, Ghilardi MF, Ghez C (1999) Independent learning of internal models for kinematic and dynamic control of reaching. *Nat Neurosci* 2(11):1026–1031.
- Krakauer JW, Pine ZM, Ghilardi MF, Ghez C (2000) Learning of visuomotor transformations for vectorial planning of reaching trajectories. *J Neurosci* 20(23):8916–8924.
- Mazzoni P, Krakauer JW (2006) An implicit plan overrides an explicit strategy during visuomotor adaptation. *J Neurosci* 26(14):3642–3645.
- Smith MA, Ghazizadeh A, Shadmehr R (2006) Interacting adaptive processes with different timescales underlie short-term motor learning. *PLoS Biol* 4(6):e179.
- Shadmehr R, Mussa-Ivaldi FA (1994) Adaptive representation of dynamics during learning of a motor task. *J Neurosci* 14(5):3208–3224.
- He K, et al. (2016) The statistical determinants of the speed of motor learning. *PLoS Comput Biol* 12(9):e1005023.
- Salmaso D, Longoni AM (1985) Problems in the assessment of hand preference. *Cortex* 21(4):533–549.
- Brainard DH (1997) The psychophysics toolbox. *Spat Vis* 10:433–436.
- Criscimagna-Hemminger SE, Donchin O, Gazzaniga MS, Shadmehr R (2003) Learned dynamics of reaching movements generalize from dominant to nondominant arm. *J Neurophysiol* 89(1):168–176.

Supporting Information

Singh et al. 10.1073/pnas.1613383113

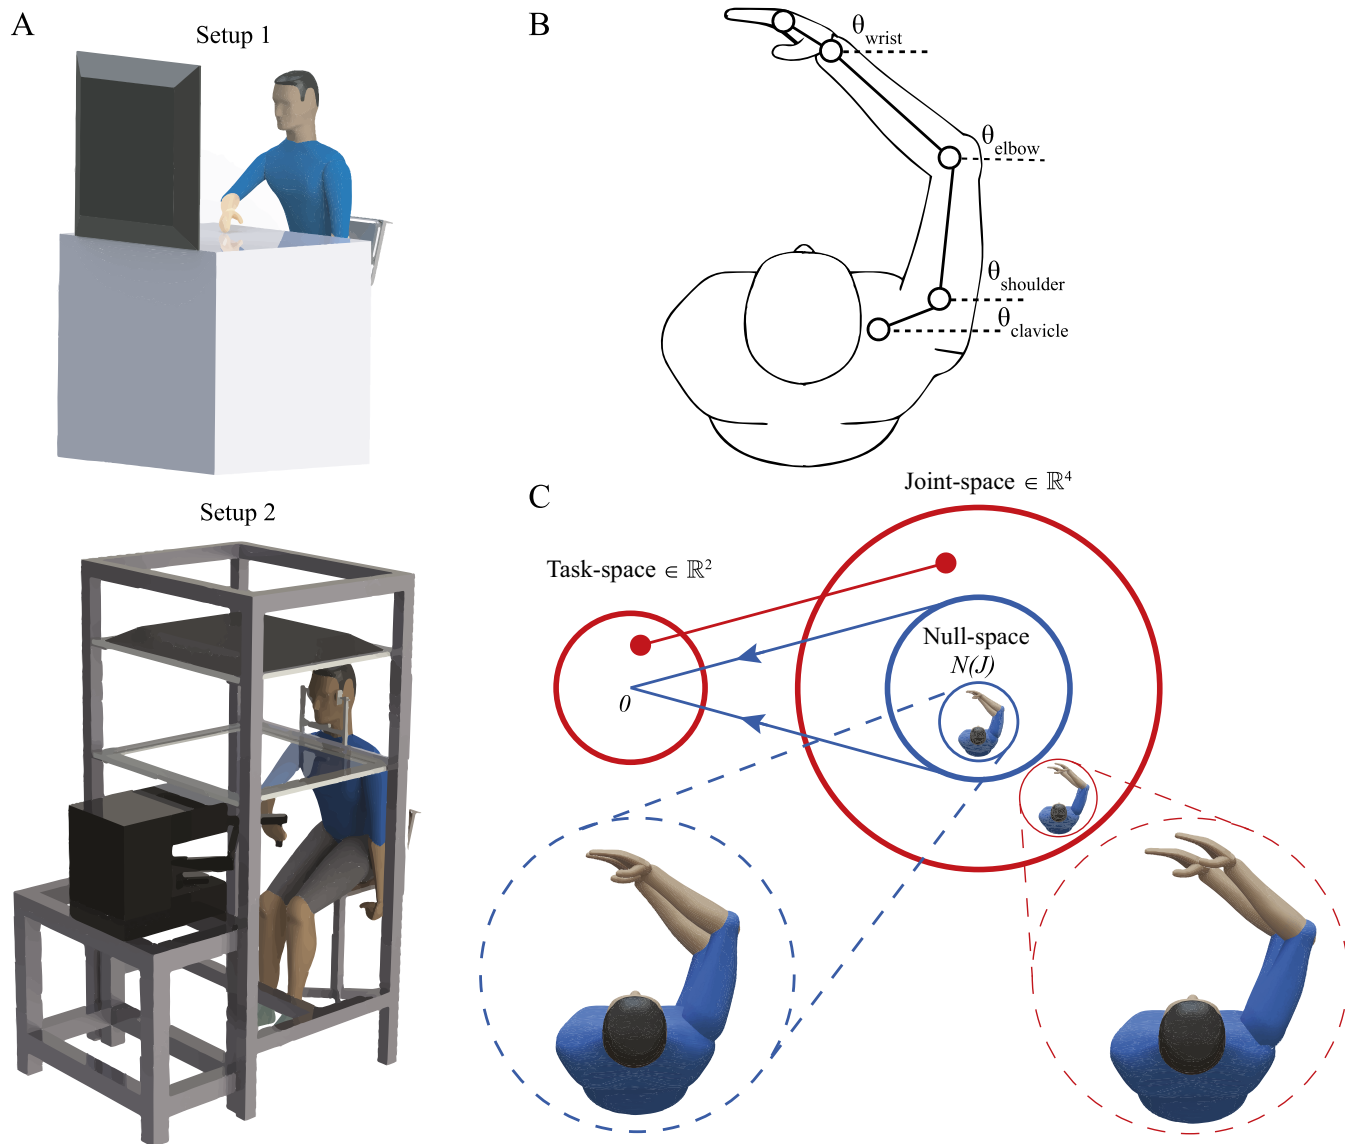


Fig. S1. Experiment setup and design. (A) Subjects made point-to-point reaching movements using two such experiment setups. (B) Trackers were used to measure joint rotation angles. (C) Illustration of null-space variability: the blue circle does not affect task-space variability; the red circle affects task-space variability.

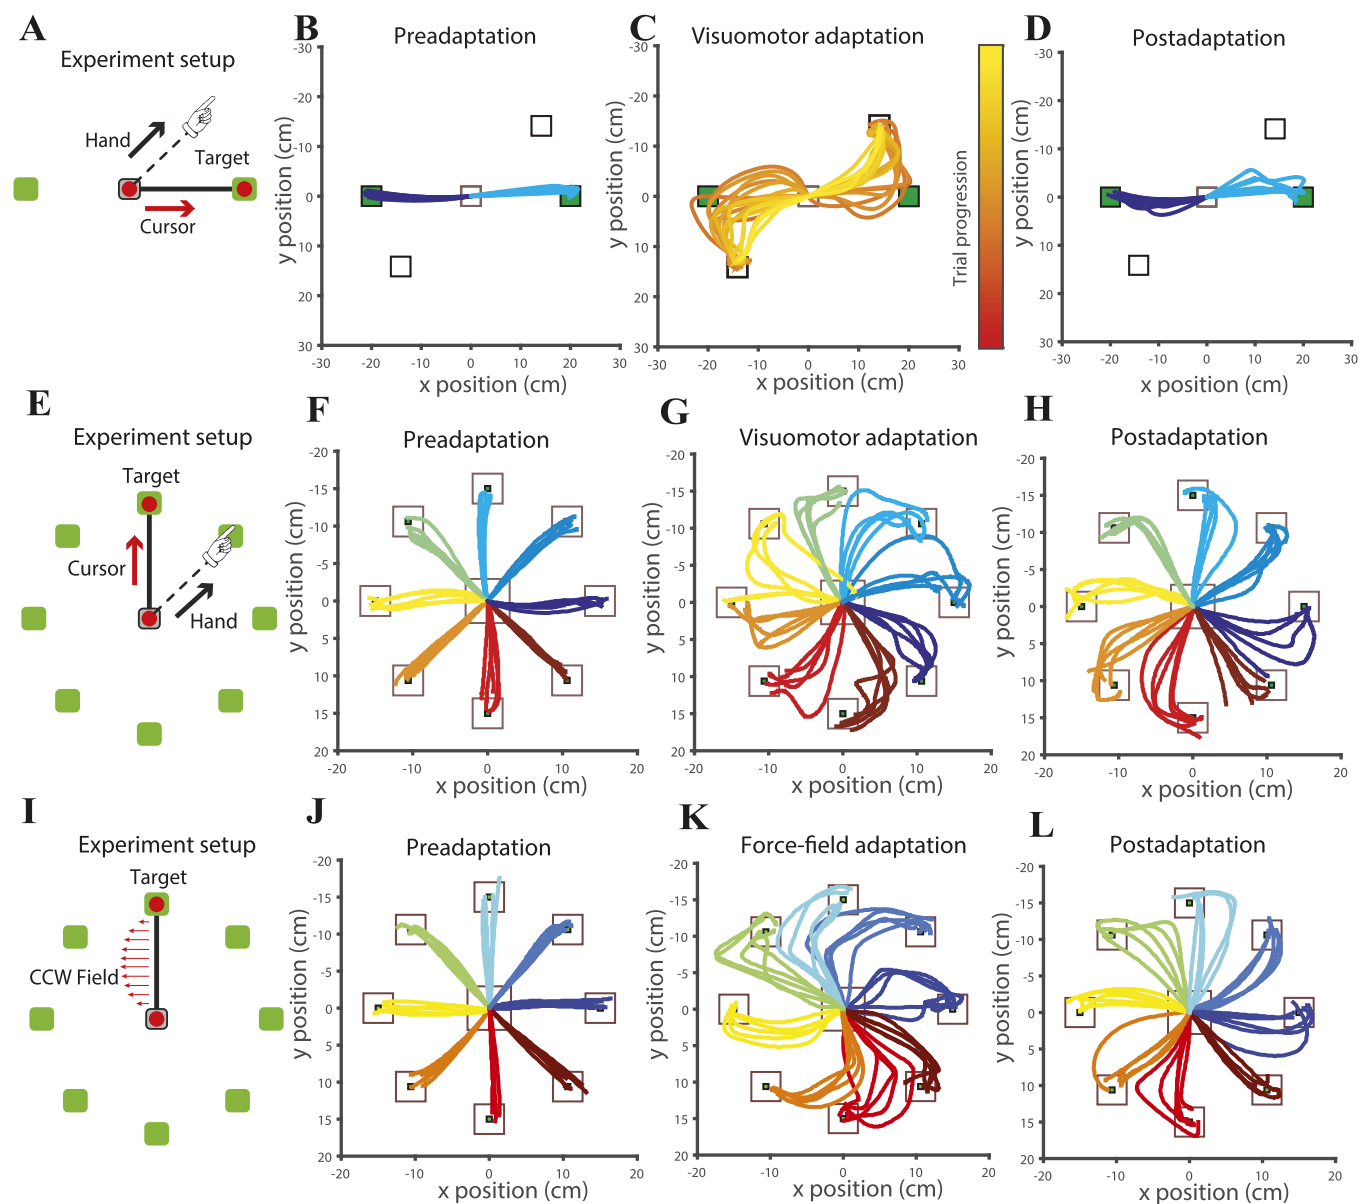


Fig. S2. Experiment setup and design. (A) Experimental apparatus and illustration of the visuomotor rotation task. (B) First five preadaptation trials from a subject showing baseline motor variability. (C) First 10 visuomotor adaptation trials (color-coded) from the same subject showing the disturbed hand trajectory. (D) First five postadaptation trials from the same subject showing the effect of adaptation. (E) Experimental apparatus and illustration of the visuomotor rotation task. (F) First five of preadaptation (baseline) trials from a subject showing motor variability in preadaptation. (G) First five of generalized visuomotor adaptation trials from the same subject showing the disturbed hand trajectory. (H) First five of postadaptation trials from the same subject showing the effect of adaptation. (I) Experimental apparatus and illustration of the viscous curl force-field. (J) First five of preadaptation (baseline) trials from a subject showing motor variability in preadaptation. (K) First five of force-field adaptation trials from the same subject showing the disturbed hand trajectory. (L) First five of postadaptation trials from the same subject showing the effect of adaptation.

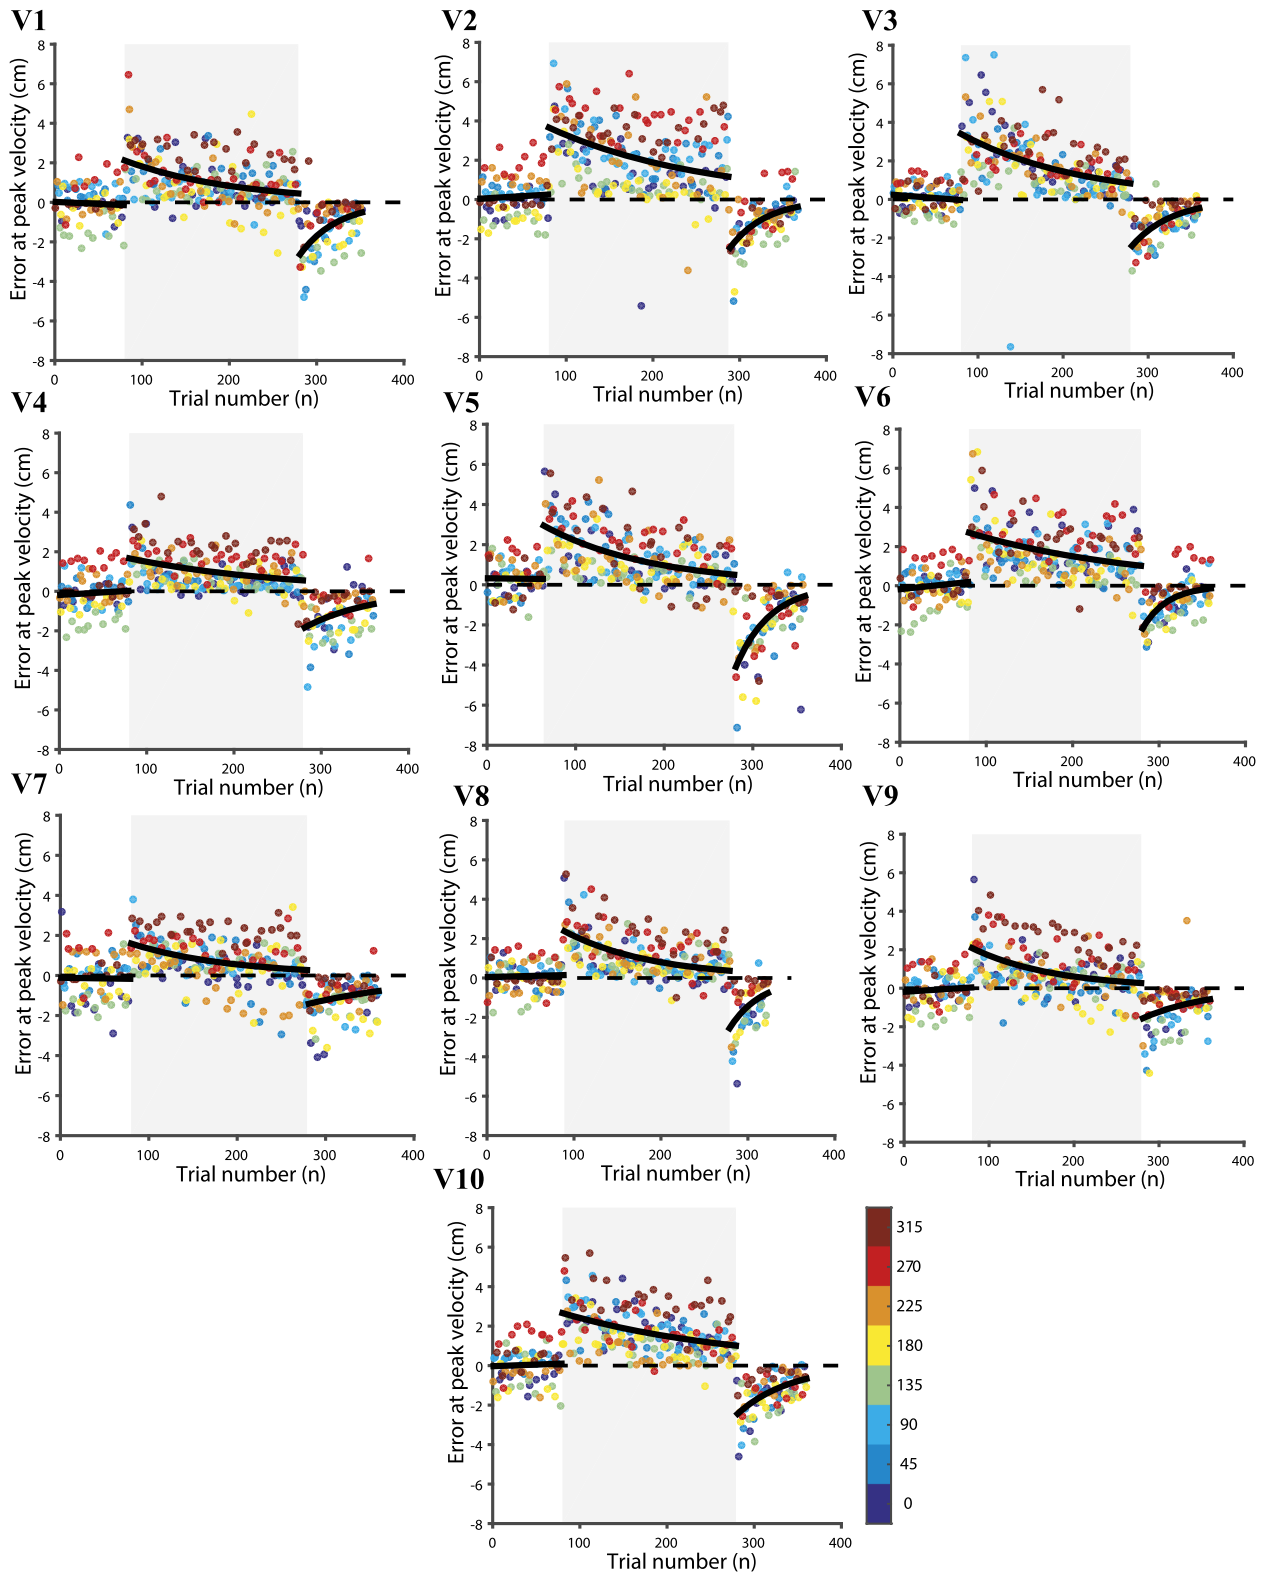


Fig. S3. Visuomotor adaptation error at peak velocity in preadaptation, adaptation, and postadaptation for all subjects (V1–V10) showing the progression of adaptation. Errors in each of the eight directions are color-coded.

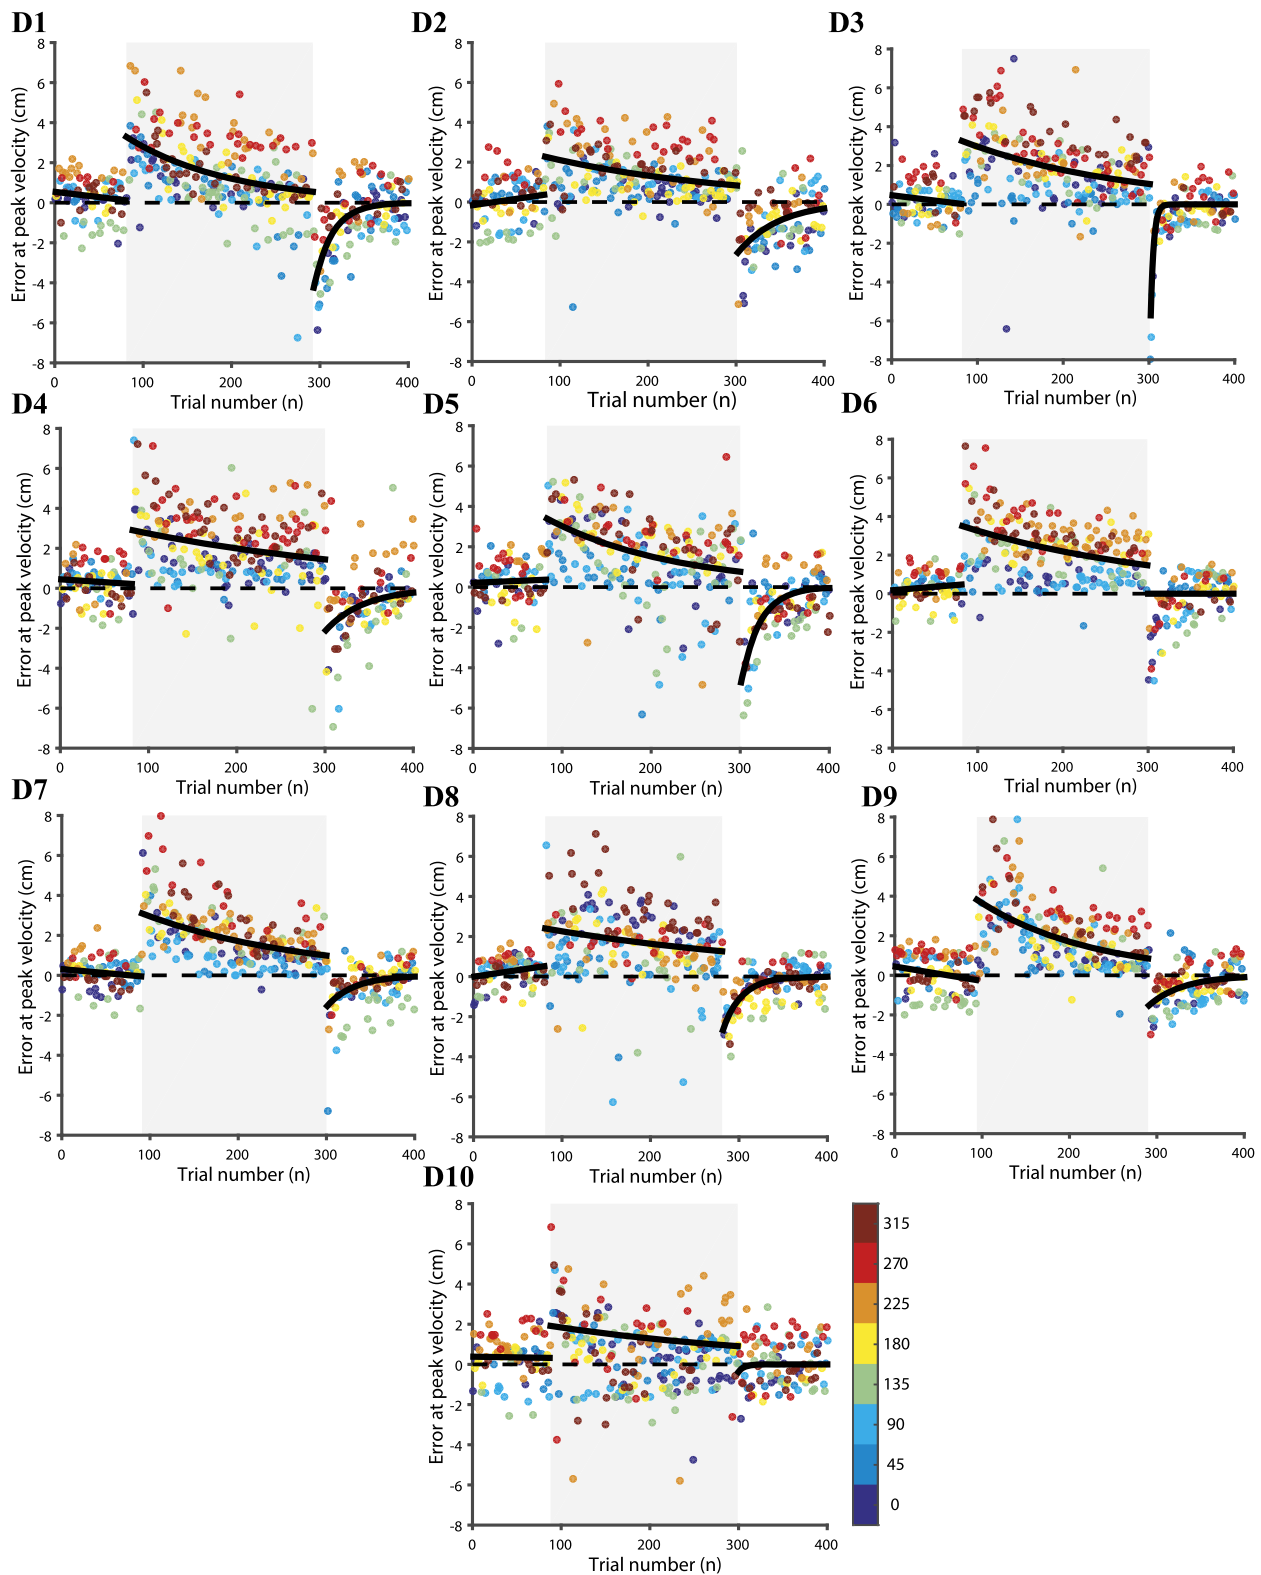


Fig. 54. Force-field adaptation error at peak velocity in preadaptation, adaptation, and postadaptation for all subjects with the dominant hand (D1–D10) showing the progression of adaptation. Errors in each of the eight directions are color-coded.

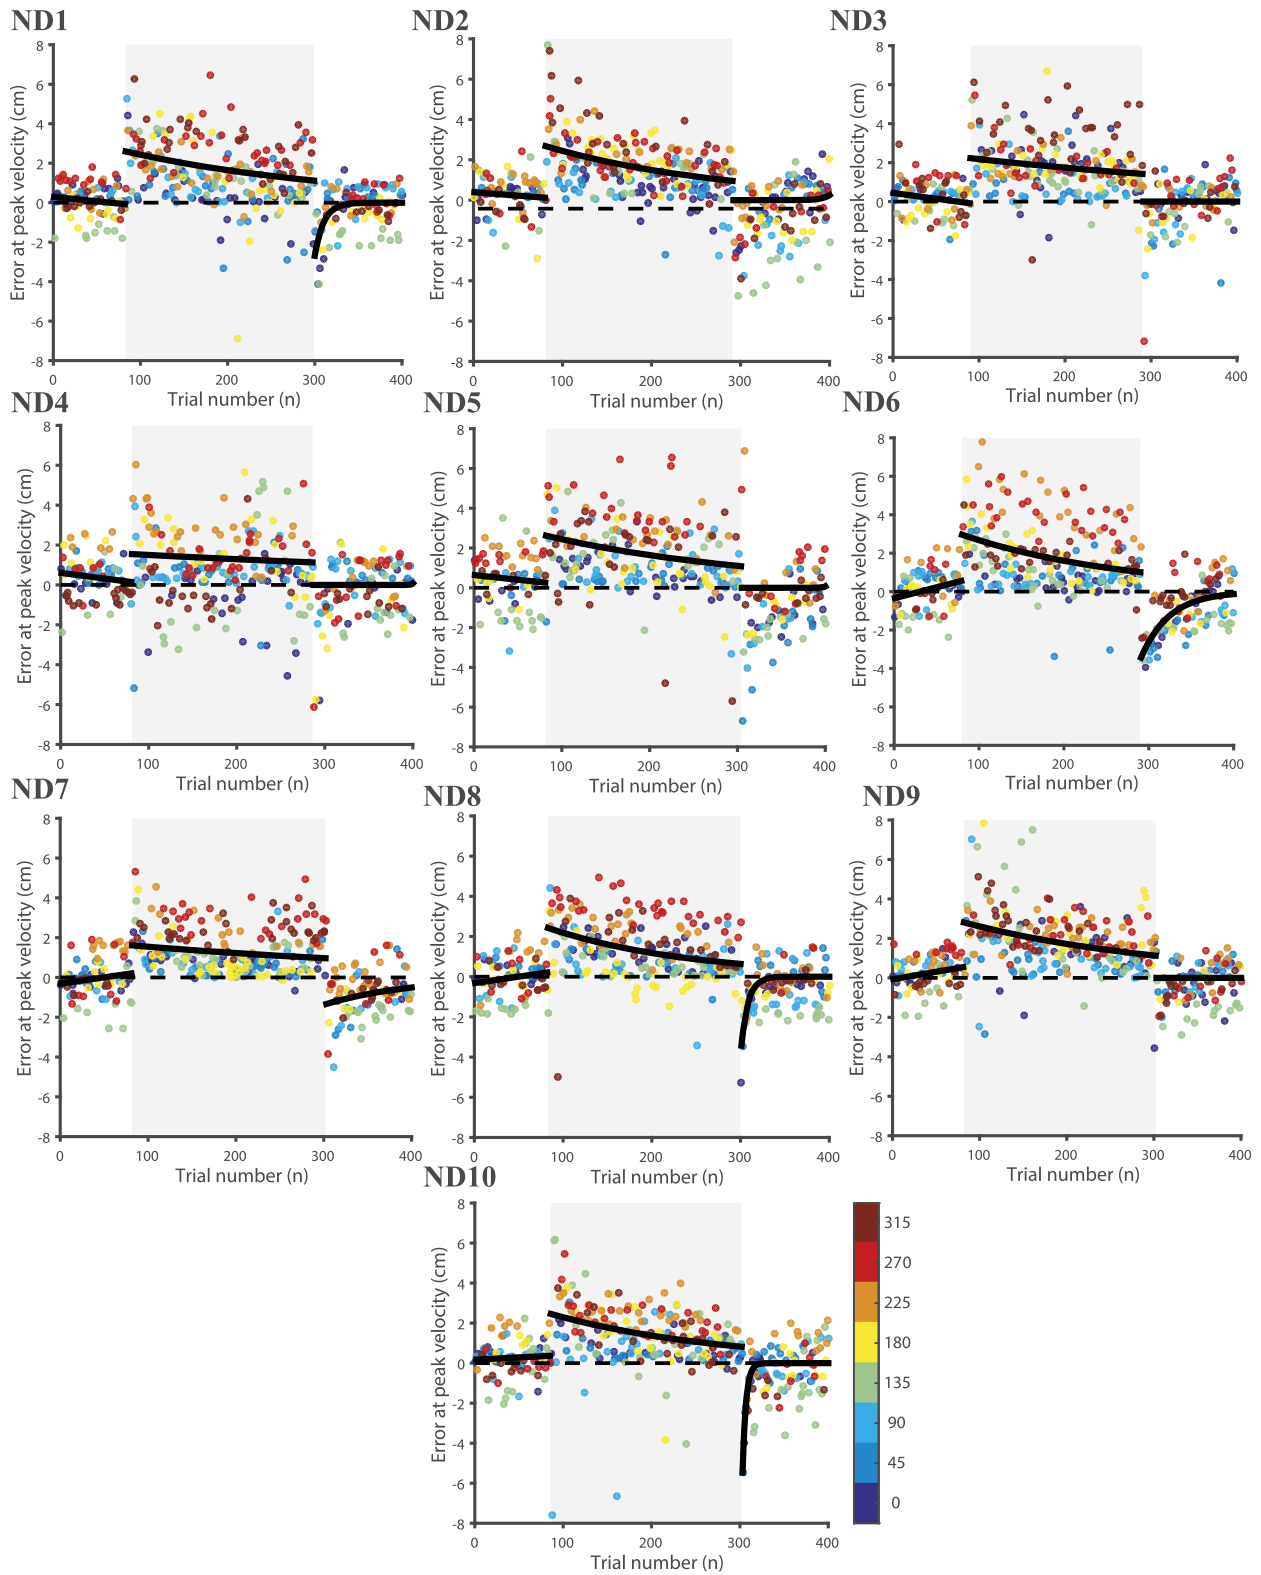


Fig. S5. Force-field adaptation error at peak velocity in preadaptation, adaptation, and postadaptation for same subjects with the nondominant hand (ND1–ND10) showing the progression of adaptation. Errors in each of the eight directions are color-coded.