Manuscript submitted to eLife

De novo learning and adaptation of continuous control in a manual tracking task

Christopher S Yang[1*], **Noah J Cowan**[2], **Adrian M Haith**[3]

***For correspondence:** christopher.yang@jhmi.edu (CSY)

[1]Department of Neuroscience, Johns Hopkins University, Baltimore, USA; [2]Department of Mechanical Engineering, Johns Hopkins University, Baltimore, USA; [3]Department of Neurology, Johns Hopkins University, Baltimore, USA

Abstract How do people learn to perform tasks that require continuous adjustments of motor output, like riding a bicycle? People rely heavily on cognitive strategies when learning discrete movement tasks, but such time-consuming strategies are infeasible in continuous control tasks that demand rapid responses to ongoing sensory feedback. To understand how people can learn to perform such tasks without the benefit of cognitive strategies, we imposed a rotation/mirror reversal of visual feedback while participants performed a continuous tracking task. We analyzed behavior using a system identification approach which revealed two qualitatively different components of learning: adaptation of a baseline controller and formation of a new task-specific continuous controller. These components exhibited different signatures in the frequency domain and were differentially engaged under the rotation/mirror reversal. Our results demonstrate that people can rapidly build a new continuous controller *de novo* and can flexibly integrate this process with adaptation of an existing controller.

Introduction

In many real-world motor tasks, skilled performance requires us to continuously control our actions in response to ongoing external events. For example, remaining stable on a bicycle depends on being able to rapidly respond to the tilt of the bicycle as well as obstacles in our path. The demand for continuous control in such tasks can pose additional challenges when it comes to learning them in the first place. In particular, new skills often require us to learn arbitrary relationships between our actions and their outcomes (like moving our arms to steer or flexing our fingers to brake) and it is thought that learning such mappings depends on the use of time-consuming cognitive strategies (***McDougle et al., 2016***). Continuous control tasks, however, require us to produce responses rapidly, leaving little time for deliberation about our actions. Therefore, it remains unclear how continuous motor skills are learned.

Studies of motor learning have revealed a variety of different processes that support motor learning in humans (***Krakauer et al., 2019***). One of the most well-characterized processes is adaptation, an error-driven learning mechanism by which task performance is improved by using sensory prediction errors to recalibrate motor output (***Mazzoni and Krakauer, 2006***; ***Tseng et al., 2007***; ***Shadmehr et al., 2010***). Adaptation is primarily characterized by the presence of aftereffects (***Redding and Wallace, 1993***; ***Shadmehr and Mussa-Ivaldi, 1994***; ***Kluzik et al., 2008***) and is known to support learning in a variety of laboratory settings including making simple movements under imposed visuomotor rotations (***Krakauer et al., 1999***; ***Fernández-Ruiz et al., 2011***; ***Morehead et al.,***

2015), prism goggles (*Martin et al., 1996*; *Fernández-Ruiz and Díaz, 1999*), split-belt treadmills (*Choi and Bastian, 2007*; *Finley et al., 2015*), force fields (*Lackner and Dizio, 1994*; *Shadmehr and Mussa-Ivaldi, 1994*), as well as in more complex settings such as path integration in gain-altered virtual reality (*Tcheang et al., 2011*; *Jayakumar et al., 2019*). However, it appears that adaptation can only adjust motor output to a limited extent; in the case of visuomotor rotations, implicit adaptation is only capable of 15–25° of compensation, even when much larger rotations are applied (*Taylor et al., 2010*; *Fernández-Ruiz et al., 2011*; *Taylor and Ivry, 2011*; *Bond and Taylor, 2015*). This suggests that other mechanisms are required when learning to compensate for perturbations that impose significant deviations from one's existing baseline motor repertoire.

In scenarios where adaptation is insufficient, people appear to adopt a re-aiming strategy to compensate for perturbed visual feedback. This strategy involves aiming one's movements towards a surrogate target rather than the true target of the movement. It has been shown that people use re-aiming strategies—or more generally, cognitive strategies—to compensate (at least in part) for visuomotor rotations (*Mazzoni and Krakauer, 2006*; *de Rugy et al., 2012*; *Taylor et al., 2014*; *Morehead et al., 2015*) and force fields (*Schween et al., 2019*). In principle, re-aiming enables people to compensate for arbitrary re-mappings of their environment, including large (90°) visuomotor rotations (*Bond and Taylor, 2015*) or mirror-reversed visual feedback (*Wilterson and Taylor, 2019*). However, implementing re-aiming is cognitively demanding and time-consuming process that significantly increases reaction times (*Haith et al., 2015*; *Leow et al., 2017*; *McDougle and Taylor, 2019*; *Fernández-Ruiz et al., 2011*). Indeed, although people can successfully compensate for a mirror reversal in point-to-point reaching tasks, this learning is not reflected in feedback response corrections to mid-movement perturbations (*Telgen et al., 2014*; *Kasuga et al., 2015*; *Gritsenko and Kalaska, 2010*), suggesting that the initial compensation was achieved through time-consuming re-aiming that could not be applied during a rapid online correction. By contrast, adaptation generalizes strongly to online corrective movements (*Ahmadi-Pajouh et al., 2012*; *Cluff and Scott, 2013*; *Telgen et al., 2014*). Thus, re-aiming strategies seem unlikely to account for how humans learn motor tasks that require continuous responses to perturbations and changing goals.

Instead, continuous tasks must likely be learned by building a new controller that implements the newly required mapping from sensory input to motor output—a process that has been termed *de novo* learning (*Figure 1A*) (*Costa, 2011*; *Telgen et al., 2014*; *Sternad, 2018*). This approach can be contrasted with adaptation, which parametrically adjusts the existing controller, and with re-aiming, which maintains the existing controller and provides it with artificial movement goals to obtain the desired motor output. It has been suggested that *de novo* learning is necessary when learning to compensate for a mirror-reversal of visual feedback. Theoretical and empirical findings have demonstrated that the learning rules which allow people to adapt to visuomotor rotations fail under mirror reversal (*Abdelghani et al., 2008*; *Hadjiosif et al., 2020*). Other studies have shown that mirror-reversal learning, unlike rotation learning, does not result in reach-direction aftereffects (*Gutierrez-Garralda et al., 2013*; *Lillicrap et al., 2013*) and likely has a distinct neural basis from rotation learning (*Schugens et al., 1998*; *Maschke et al., 2004*; *Morton and Bastian, 2006*; *Gutierrez-Garralda et al., 2013*). Although some studies have suggested the opposite, namely that mirror-reversals and rotations engage the same learning mechanisms (*Werner and Bock, 2010*; *Bock, 2013*), existing evidence overwhelmingly supports the view that these perturbations engage—at least in part—qualitatively distinct mechanisms.

While most studies indicate that mirror reversal is learned *de novo*, many of them used experimental paradigms involving discrete movements (e.g., point-to-point reaches, ball throwing) that can be learned via re-aiming; for a point-to-point reach, if one knows the mirroring axis, one can still use their existing baseline controller to generate movement and simply aim their hand across the axis to reach the target. To dissociate true *de novo* learning of a new controller from re-aiming it is necessary to consider tasks in which movement goals change more quickly than the time it takes for slow cognitive strategies to be applied. One such approach is to require participants to continuously track an unpredictable stimulus. Although several studies have tested continu-

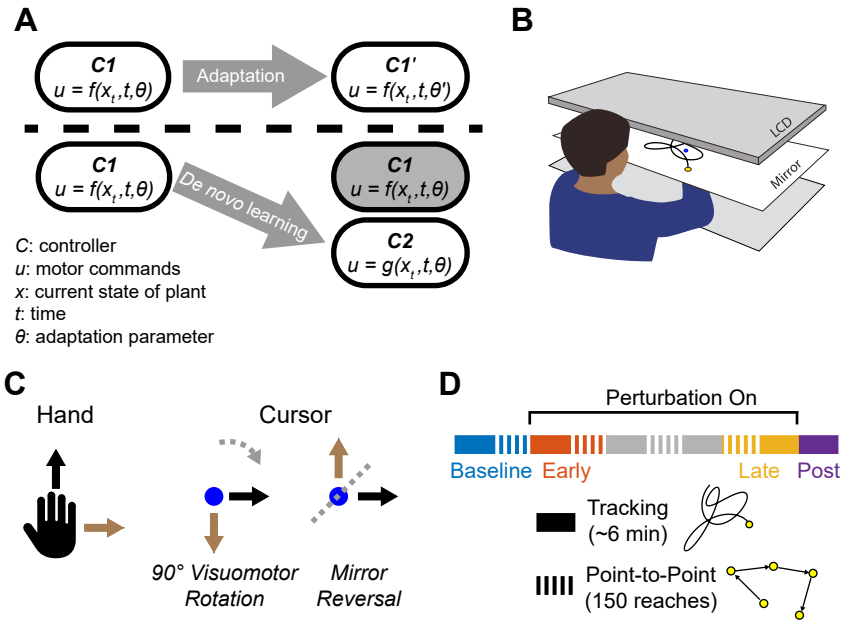

Figure 1. Conceptual overview and experimental design. **A.** We conceptualize adaptation as a parametric change to an existing controller (changing θ to θ') and *de novo* learning as building a new controller (g) to replace the baseline controller (f). **B.** Participants performed planar movements with their right hand while a target (yellow) and cursor (blue) were presented to them on an LCD display. Participants were asked to either move the cursor to a static target (point-to-point task) or track a moving target with the cursor (tracking task). **C.** Participants learned to control the cursor under one of two visuomotor perturbations: a 90° visuomotor rotation, or a mirror reversal. **D.** Participants alternated between point-to-point reaching (1 block = 150 reaches) and tracking (1 block = ~6 mins). We first measured baseline performance in both tasks under veridical visual feedback (blue), followed by interleaved tracking and point-to-point blocks with perturbed visual feedback from early learning (orange) to late learning (yellow). At the end of the experiment, we assessed aftereffects in the tracking task by removing the perturbation (purple).

ous tracking under mirror reversal (***Schugens et al., 1998***; ***Bock and Schneider, 2001***; ***Bock et al., 2001***), these studies used low-frequency stimuli (<0.35 Hz) which could potentially be tracked using intermittent "catch-up" movements that are strategically planned similar to explicit re-aiming of point-to-point movements (***Craik, 1947***; ***Miall et al., 1993a***; ***Russell and Sternad, 2001***). Therefore, to our knowledge, no study has yet demonstrated that mirror-reversal genuinely reflects *de novo* learning of a continuous controller.

Here, we sought to explicitly test whether a motor task could be learned by forming a *de novo* controller, rather than through adaptation or re-aiming. Participants learned to counter a mirror-reversal of visual feedback in both a point-to-point movement task and in a continuous tracking task in which a target moved in a pseudorandom sum-of-sinusoids trajectory (***Figure 1B-C***) (***Miall et al., 1993b***; ***Kiemel et al., 2006***; ***Roth et al., 2011***; ***Madhav et al., 2013***; ***Sponberg et al., 2015***; ***Yamagami et al., 2019***). The target in the tracking task moved continuously and unpredictably at frequencies ranging from 0.1–2.15 Hz—a high enough frequency that participants could not engage in time-consuming deliberate planning of the kind associated with re-aiming (***Fernández-Ruiz et al., 2011***; ***McDougle and Taylor, 2019***; ***Leow et al., 2017***; ***Haith et al., 2015***). Participants instead had to continuously generate movements to track the target. Critically, the sum-of-sines structure of the target motion allowed us to employ a frequency-based system identification approach to characterize changes in participants' motor controllers during mirror-reversal learning. We compared learning in this group to that of a second group of participants that learned to counter a visuomotor rotation, where presumably—unlike mirror reversal—adaptation would contribute to learning.

We hypothesized that participants learning to counter the mirror reversal would be able to learn a *de novo* controller that enables them to smoothly track the target. If, however, the mirror reversal can only be learned through a re-aiming strategy, participants would be incapable of performing

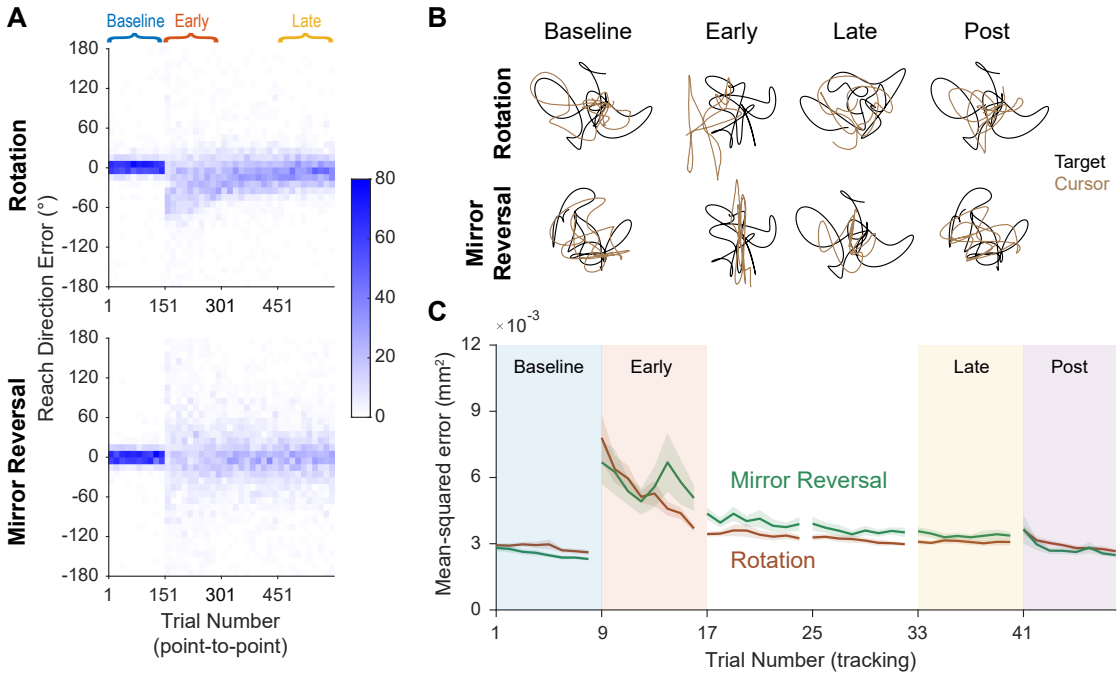

Figure 2. Task performance improves in the point-to-point and tracking tasks. **A.** Performance in the point-to-point task, as quantified by initial reach direction error, is plotted as heat maps for the rotation (top) and mirror-reversal groups (bottom). Each column shows the distribution of initial reach direction errors, pooled across all participants, over a (horizontal) bin of 15 trials. The intensity of color represents the number of trials in each 10° vertical bin (max possible value of 150 for any bin). **B.** Example tracking trajectories from a representative participant in each group. Target trajectories are shown in black while cursor trajectories are shown in brown. Each trajectory displays approximately 5 seconds of movement. **C.** Performance in the tracking task as quantified by average mean-squared positional error between the cursor and target during each trial. Error bars are SEM across participants.

Figure 2–video 1. Video of tracking behavior at different time points during learning.

the continuous feedback control required to track the target. Additionally, we hypothesized that participants *should* be able to counter the rotation in the tracking task due to the presence of sensory prediction errors—the only error signal needed to drive adaptation.

Results

Participants Learned to Compensate for the Rotation and Mirror Reversal but using Different Learning Mechanisms

Twenty participants used their right hand to manipulate an on-screen cursor under either a 90° visuomotor rotation ($n = 10$) or a mirror reversal ($n = 10$) about an oblique 45° axis (***Figure 1C***). These perturbations were designed such that, in both cases, motion of the hand in the x direction was mapped to cursor motion in the y direction and vice versa. Each group first practiced moving under their respective perturbation in a *point-to-point task*, reaching towards targets that appeared at random locations on the screen (***Figure 1D***), and we quantified their performance in each trial through the the error in their initial reach direction. For the rotation group, this error decreased as a function of training time and plateaued near 0°, demonstrating that participants successfully learned to compensate for the rotation (***Figure 2A***, upper panel). For the mirror-reversal group, the directional error did not show any clear learning (***Figure 2A***, lower panel), but performance was better than would be expected if participants had not attempted to compensate at all (which would manifest as reach errors uniformly distributed between ± 180°). Thus, both groups of participants compensated for the perturbations during point-to-point movements (at least partially), consistent with previous findings.

To test whether participants could compensate for these perturbations in a continuous control

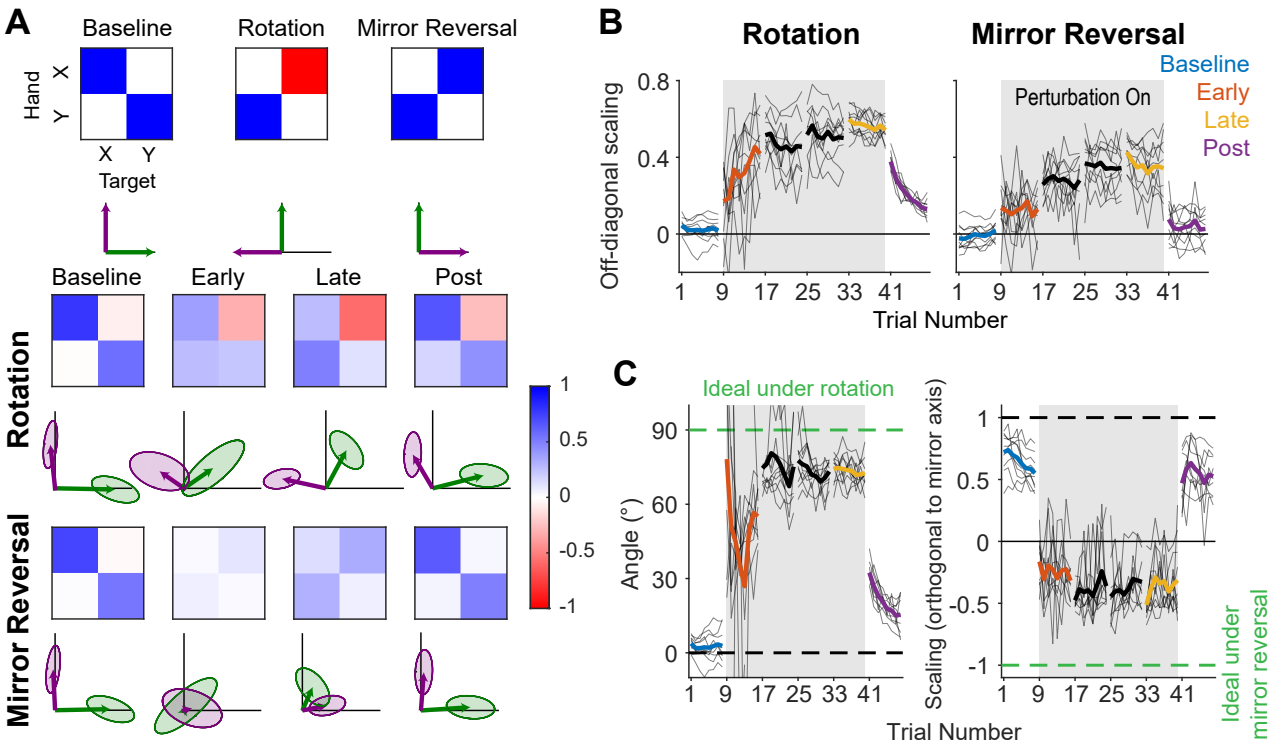

Figure 3. The rotation group exhibited reach-direction aftereffects while the mirror-reversal group did not. **A.** Alignment matrices relating target and hand movement established by trajectory alignment. The top row illustrates the ideal alignment matrices at baseline or to successfully compensate for each perturbation (blue represents positive values, red represents negative values). Alignment matrices (calculated from one trial averaged across participants) from the rotation (middle row) and mirror-reversal (bottom row) groups are depicted at different points during learning. Below each matrix, we visualized how the unit x and y vectors (black lines) would be transformed by the columns of the matrices (transformed x = green, transformed y = purple). Shaded areas are 95% confidence ellipses across participants. **B.** The average of the two off-diagonal elements of the estimated alignment matrices across all blocks of the experiment (baseline through post). Grey boxes indicate when the rotation or mirror reversal were applied. Thin black lines indicate individual participants and thick lines indicate the mean across participants. **C.** (Left: rotation group) Angular compensation for the rotation, estimated by approximating each alignment matrix with a pure rotation matrix. (Right: mirror-reversal group) Scaling factor orthogonal to the mirror axis. In each plot, dashed lines depict ideal performance when the perturbation is (green) or is not (black) applied. Thin black lines indicate individual participants and thick lines indicate the mean across participants.

task after having practiced them in the point-to-point task, we had them perform a manual tracking task. In this task, participants tracked a target that moved in a continuous sum-of-sinusoids trajectory at frequencies ranging between 0.1–2.15 Hz, with distinct frequencies used for x- and y-axis target movement. The resulting target motion was unpredictable and appeared random. Furthermore, the target's trajectory was altered every block by randomizing the phases of the component sinusoids, preventing participants from being able to learn a specific target trajectory. Example trajectories from single participants are presented in ***Figure 2B*** (see also ***Figure 2–video 1*** for a movie of tracking behavior).

As an initial assessment of how well participants learned to track the target, we measured the average mean-squared positional error (tracking error) between the target and cursor trajectories during every tracking trial. Tracking error improved with practice in both groups of participants, approaching similar levels of error by late learning (***Figure 2C***). Therefore, in both the point-to-point and tracking tasks, participants' performance improved with practice. However, this measure of tracking error is insensitive to changes in the direction of cursor movement, a feature of the behavior that is critical for distinguishing how participants learned to counter the rotation versus the mirror reversal.

Instead, we quantified participants' behavior in the tracking task using a different method,

estimating how participants translated target motion into hand motion at different points during learning. We aligned the hand and target tracking trajectories with a linear transformation matrix (alignment matrix) that, when applied to the target trajectory, minimized the discrepancy between the hand and target trajectories (see Methods for details). This approach was sensitive to detecting the directionality of hand movements, allowing us to more precisely quantify how participants altered their movements during learning.

Figure 3A shows the estimated matrices for both groups at different time points during the experiment, along with a visualization of how they affected the unit x and y vectors. At baseline, the estimated alignment matrices were close to the identity matrix, as would be expected if the hand trajectory is well aligned with the target trajectory. Under perturbed feedback, perfect tracking would be achieved when the alignment matrix is equal to the inverse of the matrix representing the applied perturbation. Indeed, by late learning, the alignment matrices resembled this inverse under both perturbations.

To test whether these changes were statistically significant, we focused on the off-diagonal elements of the matrices. These elements critically distinguish the different transformations from one another and from baseline. In late learning, both the rotation (linear mixed effects model [see "Statistics" in Methods for details about the structure of the model]: two-way interaction between group and block ($F(2, 36) = 7.56$, $p = 0.0018$; Tukey's range test: $p < 0.0001$) and mirror-reversal groups (Tukey's range test: $p < 0.0001$) exhibited off-diagonal values that were significantly different from their baseline values (***Figure 3B***), and in the appropriate direction to compensate for their respective perturbations.

From these matrices, we derived additional metrics associated with each perturbation to further characterize learning. For the rotation group, we estimated a compensation angle, θ, using a singular value decomposition approach (***Figure 3C***; see "Trajectory-alignment analysis" in Methods for details). At baseline, we found that $\theta = 3.8 \pm 1.0°$ (mean $\pm$ SEM), and this increased to $\theta = 72.5 \pm 1.9°$ by late learning. For the mirror-reversal group, to assess whether participants learned to flip the direction of their movements across the mirroring axis, we computed the scaling of the target trajectory along the direction orthogonal to the mirror axis (***Figure 3C***). This value was positive at baseline and negative by late learning, indicating that participants successfully inverted their hand trajectories relative to that of the target.

Lastly, we sought to confirm that the rotation and mirror reversal were learned using different mechanisms, as has been suggested by previous studies (***Gutierrez-Garralda et al., 2013***; ***Telgen et al., 2014***). We did so by assessing whether participants in each group expressed reach-direction aftereffects—the canonical hallmark of adaptation—at the end of the experiment, following removal of each perturbation in the tracking task (and with participants made explicitly aware of this). Again estimating alignment matrices, we found that the magnitude of the aftereffects was different for the two visuomotor perturbations (***Figure 3B***). The off-diagonal elements for the rotation group were significantly different from baseline (Tukey's range test: $p < 0.0001$), indicating clear aftereffects. These aftereffects corresponded to a compensation angle of $\theta = 32.4 \pm 1.4°$, similar to the magnitude of aftereffects reported for visuomotor rotation in point-to-point tasks (***Bond and Taylor, 2015***; ***Morehead et al., 2017***). For the mirror-reversal group, by contrast, the off-diagonal elements of the post-learning alignment matrix were not significantly different from baseline (Tukey's range test: $p = 0.2057$; baseline range: -0.11–0.11; post-learning range: -0.07–0.28), suggesting negligible aftereffects. (Note that aftereffects were not evident in the mean-squared error analysis in ***Figure 2C***, as that analysis was not designed to be sensitive to the small changes in movement direction associated with aftereffects.) The lack of aftereffects under mirror-reversal implies that participants did not counter this perturbation via adaptation and instead engaged *de novo* learning.

In summary, these data suggest that participants were able to compensate for both perturbations in the more challenging tracking task. Consistent with previous studies, the data support the idea that the rotation was learned via adaptation while the mirror reversal was learned via *de novo* learning.

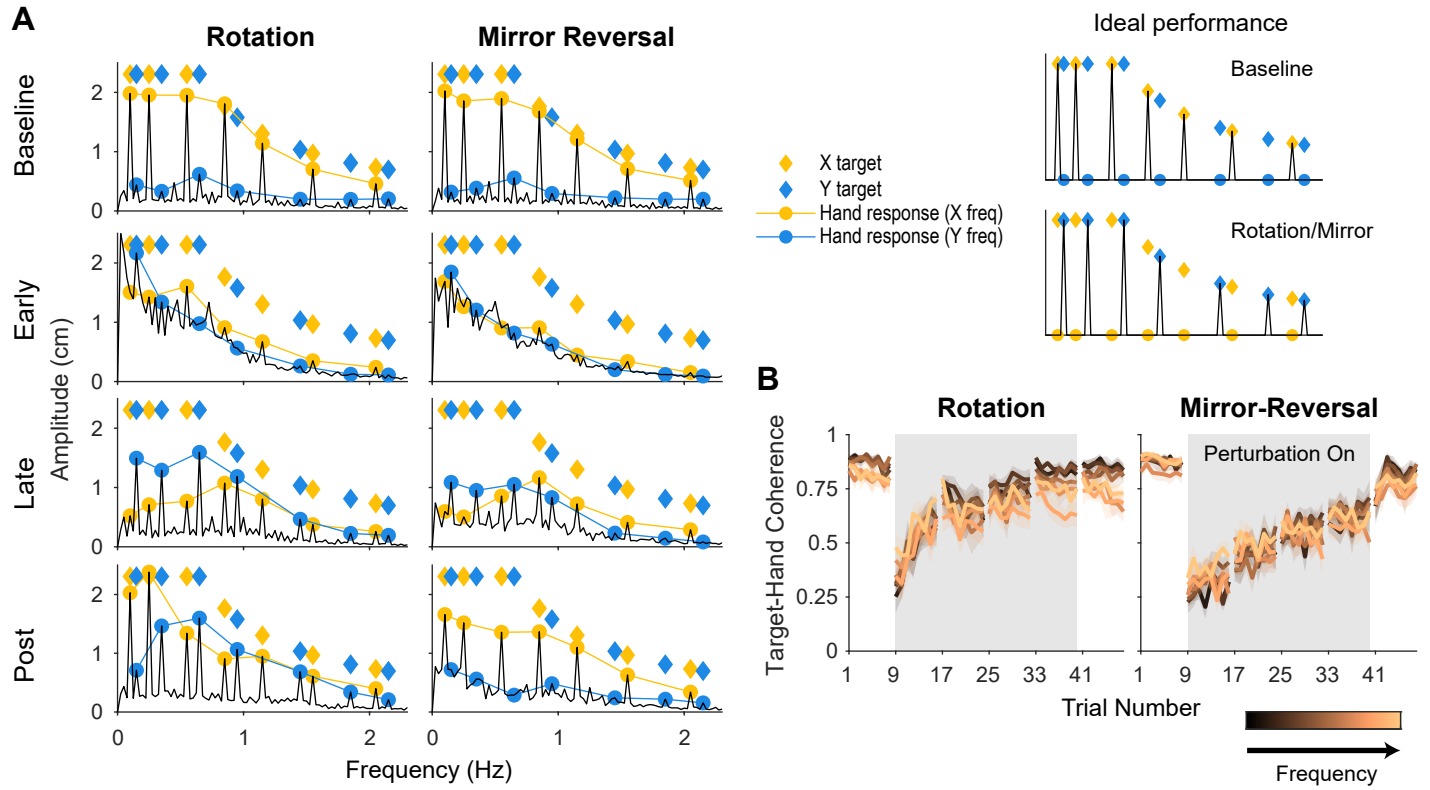

Figure 4. Tracking behavior was approximately linear, indicating that the hand tracked the target continuously. **A.** Amplitude spectra of x-hand trajectories (black line) averaged across participants from one trial in each listed block. In each plot, the amplitudes and frequencies of target motion are indicated by diamonds (yellow: x-target frequencies; blue: y-target frequencies). Hand responses at x- and y-target frequencies are highlighted as yellow and blue circles, respectively, and are connected by lines for ease of visualization. Ideal performance at baseline and under the rotation/mirror reversal are depicted on the right. **B.** Spectral coherence between x-target movement and both x- and y-hand movement (i.e., single-input multi-output coherence), which provides a measure of how linear the relationship is between target motion and hand motion. This coherence is proportional to the linear component of the hand's response to the target. Darker colors represent lower frequencies and lighter colors represent higher frequencies. Error bars are SEM across participants.

Figure 4–Figure supplement 1. Amplitude spectra and coherence plots for y-hand movements.

Figure 4–Figure supplement 2. Amplitude spectra of x-hand movements from single subjects.

Participants Performed Continuous Control to Track the Target

Although participants could learn to successfully perform the tracking task under both perturbations, it is not necessarily clear that they achieved this using continuous control; the target moved primarily at low frequencies (0.1–0.65 Hz) which by design had larger amplitude and lower velocity. This could potentially have allowed participants to track the target through intermittent "catch-up" movements that were strategically planned similar to explicit re-aiming of point-to-point movements (***Craik, 1947***; ***Miall et al., 1993a***; ***Russell and Sternad, 2001***). While it is difficult to rule this out based on the trajectory alignment analysis, the tracking task was designed to be amenable to a more more fine-grained analysis of participants' behavior using frequency-domain system identification that would allow us to examine the continuity of movements.

We examined the amplitude spectra of participants' hand movements after transforming the time-domain data to the frequency domain using the discrete Fourier transform and compared these amplitude spectra to that of the target (***Figure 4A***). This allowed us to check whether the dynamics of participants' tracking behavior was close to linear. If the relationship between hand and target movement was linear this would imply that the hand moved at the same frequencies as the target and, consequently, would suggest that participants were faithfully tracking the target using continuous movements. In contrast, if the relationship was nonlinear (e.g., the hand moved at

different frequencies than the target), this would suggest that participants were using an alternative tracking strategy. Note that coupling across axes (e.g., x-hand movements in both the x and y-axes at a given frequency) would *not* indicate nonlinear behavior; at a given frequency, the participant's goal during learning is to remap responses in one axis to the other. Thus, coupling would indicate an imperfect, though possibly linear, sensorimotor mapping.

At baseline, both groups of participants moved almost exclusively at the frequencies of the target (***Figure 4A***: x-hand data, ***Figure 4–Figure Supplement 1***: y-hand data, ***Figure 4–Figure Supplement 2***: single-subject data). More specifically, x-hand movements primarily occurred at x-target frequencies and y-hand movements primarily occurred at y-target frequencies, as would be expected for tracking behavior at baseline. The introduction of the perturbation led to a broadband increase in amplitude across all frequencies for both groups (***Figure 4A***, "Early"), indicating some nonlinear behavior as one might expect, particularly during early learning. However, the peaks at the target frequencies were still clearly identifiable. These nonlinearities abated with practice (***Figure 4A***, "Late") and remained modest after the perturbation was removed (***Figure 4A***, "Post").

We further quantified how linear participants' responses were by computing the spectral coherence between target and hand movement (***Roddey et al., 2000***). Although the coherence was low during early learning, it was close to the maximum value of 1 at baseline, late learning, and post-learning (***Figure 4B***: x-hand data, ***Figure 4–Figure Supplement 1***: y-hand data). Altogether, these data are consistent with a linear relationship between target motion and hand motion, suggesting that both the rotation and mirror-reversal groups' tracking movements were continuous and that participants did not use an intermittent strategy to perform the task.

Adaptation and *De Novo* Learning Exhibit Distinct Signatures in the Frequency Domain

Because tracking behavior was approximately linear, this provided validation for using linear systems analysis to more deeply explore how learning altered participants' control capabilities. While learning likely results in nonlinear changes to motor controllers, the movements we observed (i.e., the product of learning) were linear. This allowed us to treat each trial as a snapshot of participants' input–output relationship between target and hand movement. While recent studies have shown that learning can occur on very fast timescales (***Crevecoeur et al., 2020a***,b), we believe the magnitude of these learning effects are small enough that each 40 second trial approximately captures a single state of the input–output relationship. Furthermore, such within-trial learning effects would likely be restricted only to the earliest blocks of exposure to the perturbation.

To perfectly compensate for either the rotation or the mirror reversal, movement at x-target frequencies needed to be remapped from x-hand movements to y-hand movements, and vice versa at y-target frequencies. During early learning, participants in the rotation group did produce x-hand movements in response to y-target frequencies, but also inappropriately continued to produce x-hand movements at x-target frequencies (***Figure 4A***). By late learning, the amplitude of x-hand movement further increased at y-target frequencies and decreased at x-target frequencies. Behavior for the mirror-reversal group followed a similar pattern, albeit with less pronounced peaks in the amplitude spectrum during early learning.

After the perturbation was removed (Post-learning), the rotation group exhibited x-hand movements at both x- and y-target frequencies, unlike baseline where movements were largely restricted to x-target frequencies (***Figure 4A***), indicating aftereffects, consistent with our earlier, trajectory-alignment analysis. In contrast, the amplitude spectrum of the mirror-reversal group's x-hand movements was similar to baseline, confirming that any aftereffects were negligible. These features of the amplitude spectra, and the differences across groups, were qualitatively the same for y-hand movements (***Figure 4–Figure Supplement 1***) and were also evident in individual subjects (***Figure 4–Figure Supplement 2***).

Although the amplitude spectra illustrate important features of learning, they do not carry information about the directionality of movements and thus do not distinguish learning of the

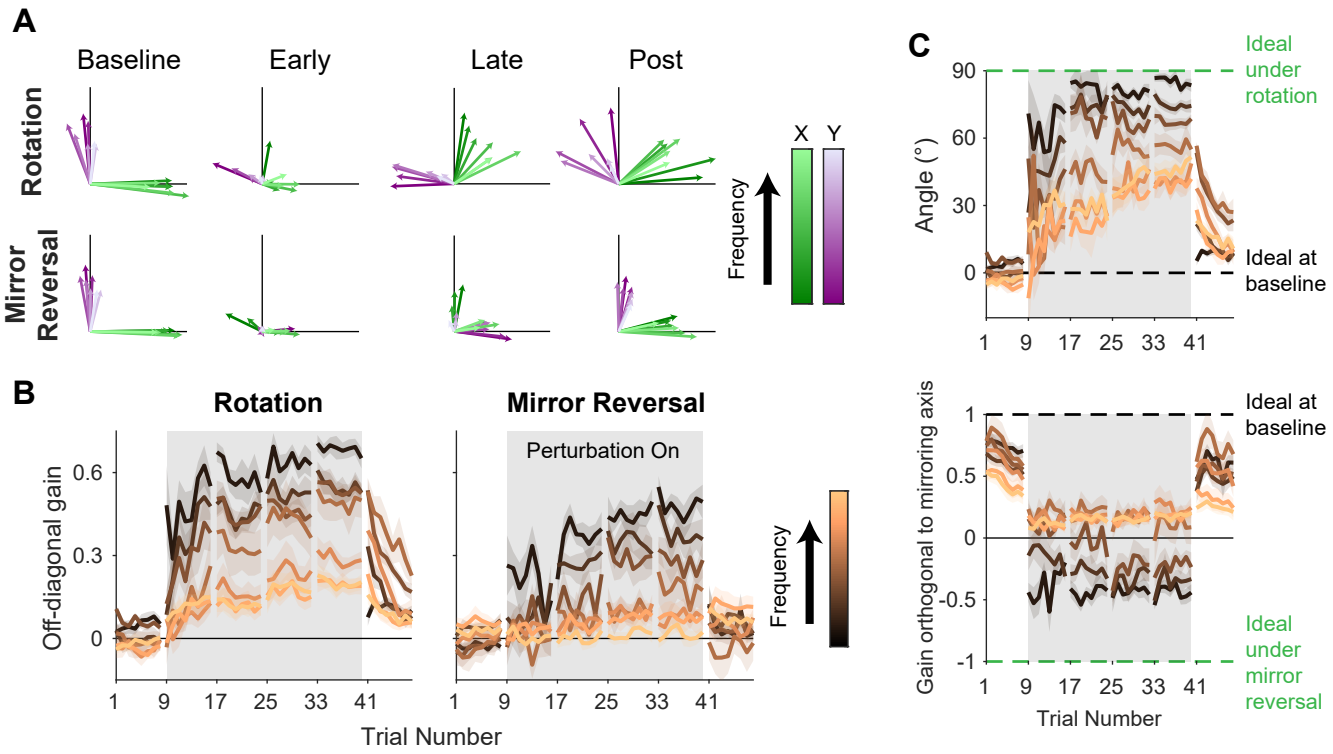

Figure 5. Adaptation and *de novo* learning exhibit distinct frequency-dependent signatures. We estimated how participants transformed target motion into hand movement across different frequencies (i.e., gain matrix analysis). For all panels, neighboring x- and y-target frequencies were paired together in numerical order. The resulting 7 frequency pairings were (x then y frequencies reported in each parentheses in Hz): (0.1, 0.15), (0.25, 0.35), (0.55, 0.65), (0.85, 0.95), (1.15, 1.45), (1.55, 1.85), (2.05, 2.15). See Methods for details on how gain matrices were fit. **A.** Visualizations of the estimated gain matrices relating target motion to hand motion across frequencies. These visualizations were generated from one trial of each listed block, averaged across participants. Green and purple arrows depict hand responses to x- and y-target frequencies, respectively. Darker colors represent lower frequencies and lighter colors represent higher frequencies. The gain matrices associated with these visualizations can be found in ***Figure 5–Figure Supplement 2***. **B.** Average of the two off-diagonal values of the gain matrices at different points during learning. Darker colors represent lower frequencies and lighter colors represent higher frequencies. Grey boxes indicate when the rotation or mirror reversal were applied. Error bars are SEM across participants. **C.** (Top) Estimated compensation angle for the rotation group as a function of frequency at different points during learning. (Bottom) Gain of movement orthogonal to the mirror axis for the mirror-reversal group. Green and black dashed lines show ideal compensation when the perturbation is or is not applied, respectively. Darker colors represent lower frequencies and lighter colors represent higher frequencies. Error bars are SEM across participants.

Figure 5–Figure supplement 1. Example gain matrices for each block and frequency

Figure 5–Figure supplement 2. Gain matrix analysis performed on single-subject data

two different perturbations; perfect compensation would lead to identical *amplitude* spectra for each perturbation. In order to distinguish these responses, we needed to determine not just the amplitude, but the direction of the response along each axis, i.e. whether it was positive or negative. We used *phase* information to disambiguate the direction of the response (the sign of the gain) by assuming that the phase of the response at each frequency would remain similar to baseline throughout learning. We then used this information to estimate signed gain matrices which describe the linear transformations relating target and hand motion (***Figure 5–Figure Supplement 1***). These matrices relay similar information as the alignment matrices in ***Figure 3*** except here, different transformations were estimated for different frequencies of movement. To construct these gain matrices, the hand responses from neighboring pairs of x- and y-target frequencies were grouped together. This grouping was performed because target movement at any given frequency was one dimensional, but target movement across two neighboring frequencies was two dimensional; examining hand/target movements in this way thus provided two-dimensional insight into how the rotation/mirroring of hand responses varied across the frequency spectrum

(see "Frequency-domain analysis" in Methods for details).

Similar to the trajectory-alignment analysis, these gain matrices should be close to the identity matrix at baseline but equal the inverse of the matrix describing the perturbation if participants are able to perfectly compensate for the perturbation. We again visualized these estimated frequency-dependent gain matrices through their effect on the unit x and y vectors (the columns of the gain matrices; ***Figure 5A***: average across subjects, ***Figure 5–Figure Supplement 2***: single subjects), only now we include a set of vectors for each pair of neighboring frequencies.

At baseline, participants in both groups responded to x- and y-target motion by moving their hands in the x- and y-axes, respectively, with similar performance across all target frequencies. Late in learning, for the rotation group, participants successfully compensated for the perturbation - apparent through the fact that all vectors rotated clockwise during learning. The extent of compensation, however, was not uniform across frequencies but was more complete at low frequencies (darker arrows) than at high frequencies (lighter arrows). For the mirror-reversal group, compensation during late learning occurred most successfully at low frequencies, apparent as the darker vectors flipping across the mirror axis relative to baseline. At high frequencies, however, responses failed to flip across the mirror axis and remained similar to baseline.

To quantify these observations statistically, we focused again on the off-diagonal elements of the estimated gain matrices. The rotation group's gain matrices were altered in the appropriate direction to counter the perturbation at all frequencies (***Figure 5B***; linear mixed effects model [see "Statistics" in methods for details about the structure of the model]: 3-way interaction between block and frequency, $F(12, 360) = 3.20$, $p = 0.0002$; data split by frequency for post hoc Tukey's range test: Bonferroni-adjusted $p < 0.0001$ for all frequencies). Although the mirror-reversal group's low-frequency gain matrices were also altered in the appropriate direction (Tukey's range test: Bonferroni-adjusted $p < 0.0003$ for lowest 3 frequencies), the high-frequency gain matrices were not significantly different from baseline (Tukey's range test: Bonferroni-adjusted $p > 0.2$ for highest 4 frequencies; baseline gain range: -0.18–0.18; late-learning gain range: -0.25–0.66).

Fitting a rotation matrix to the rotation group's gain matrix at each frequency revealed that participants' baseline compensation angle was close to 0° at all frequencies (***Figure 5C***). By late learning, compensation was nearly perfect at the lowest frequency but was only partial at higher frequencies. For the mirror-reversal group, the gains of participants' low-frequency movements orthogonal to the mirror axis were positive at baseline and became negative during learning, appropriate to counter the perturbation. At high frequencies, by contrast, the gain reduced slightly during learning but never became negative. Thus, both groups of participants were successful at compensating at low frequencies, but at high frequencies, the rotation group was only partially successful and the mirror-reversal group was largely unsuccessful.

Post-learning, the rotation group's off-diagonal gains were significantly different from baseline for all frequencies except the lowest frequency (***Figure 5B***; Tukey's range test: Bonferroni-adjusted $p < 0.005$ for highest 6 frequencies), indicating aftereffects. A similar trend was evident in participants' estimated compensation angles (***Figure 5C***). By contrast, the mirror-reversal group's matrices were not significantly different from baseline across all frequencies (***Figure 5B***; Tukey's range test: Bonferroni-adjusted $p > 0.9$ for all frequencies; baseline gain range: -0.18–0.18; post-learning gain range: -0.49 to 0.37). The gains orthogonal to the mirroring axis were also similar to baseline across all frequencies, indicating the absence of aftereffects (***Figure 5C***).

To summarize, compensation for the visuomotor rotation resulted in reach-direction aftereffects of similar magnitude to that reported in previous studies (***Figure 5***). Compensation was also expressed at both low and high frequencies of movement (***Figure 5B***). The fact that participants exhibited low-frequency compensation is, to some extent, not surprising because the low frequencies in our task required movements that were slower than that at high frequencies. However, previous studies have demonstrated that adaptation generalizes strongly to rapid online corrective movements (***Ahmadi-Pajouh et al., 2012***; ***Cluff and Scott, 2013***; ***Telgen et al., 2014***), which would be reflected in our experiment as an ability to perform tracking at high frequencies. Thus, these

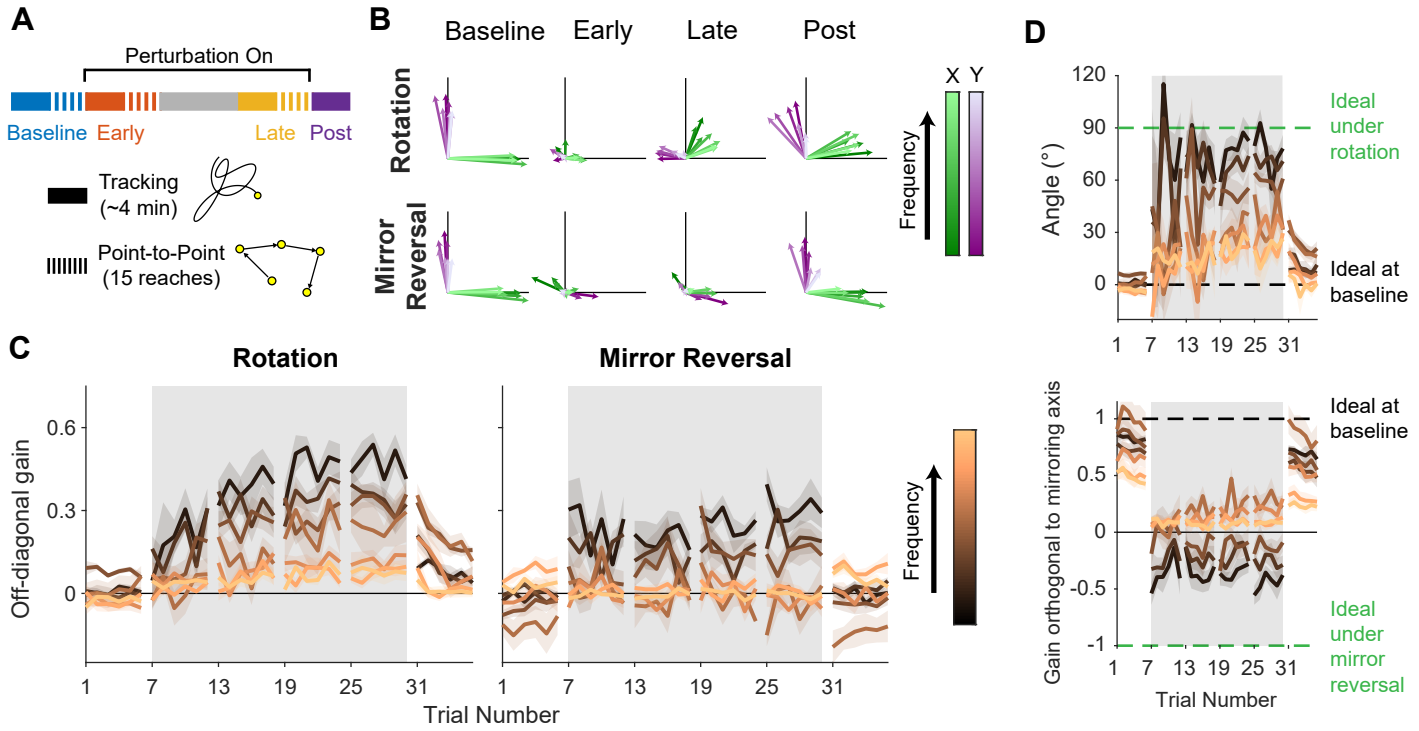

Figure 6. Making point-to-point reaches improves tracking performance, especially under mirror reversal. **A.** Participants learned to counter either a visuomotor rotation ($n = 10$) or mirror-reversal ($n = 10$). The experimental design was similar to the main experiment except point-to-point reaching practice was almost entirely eliminated; between the early- and late-learning tracking blocks, participants only performed 15 point-to-point reaches. The purpose of these reaches was not for training but simply to assess learning in the point-to-point task. **B–D.** Gain matrix analysis, identical to that in ***Figure 5***, performed on data from the follow-up experiment. **B.** Visualization of the gain matrix from one trial of each listed block, averaged across participants. **C.** Off-diagonal elements of the gain matrices, averaged across participants. **D.** Estimated rotation angle for the rotation group's gain matrices (upper) and gain orthogonal to mirroring axis for the mirror-reversal group (lower), averaged across participants. All error bars in this figure are SEM across participants.

Figure 6–Figure supplement 1. Gain matrix analysis performed on single-subject data for the follow-up experiment

data suggest that rotation learning engages adaptation, in agreement with previous studies. In contrast, the mirror-reversal group did not exhibit aftereffects and only expressed compensation at low frequencies. The lack of aftereffects suggests that participants did not learn through adaptation of an existing controller, but instead learned by building a new controller from scratch—i.e., through *de novo* learning (***Figure 1A***).

This *de novo* learning process also appeared to contribute to learning in the rotation group. The aftereffects in this group (~25°) only accounted for a fraction of the overall compensation achieved by this group (~70°), suggesting that an additional component of learning also contributed to learning to compensate for the rotation. Examining the time course of learning for both groups in ***Figure 5B***, while the rotation group's gains were overall higher than the mirror-reversal group's, there was a striking similarity in the frequency-dependent pattern of learning between the two groups. We therefore suggest that the residual learning not accounted for by adaptation was attributable to the same *de novo* learning process that drove learning under the mirror reversal.

Making Point-to-Point Reaches Improves Tracking Performance, Especially under Mirror Reversal

Although the data suggest participants did not heavily rely on an aiming strategy while tracking, participants likely did use such a strategy to learn to counter the rotation/mirror reversal while performing point-to-point reaches. How important might such cognitive strategies be for ultimately learning the tracking task? To test this, we performed a follow-up experiment with twenty additional

participants. This experiment was similar to the main experiment except for the fact that participants experienced the rotation or mirror reversal almost exclusively in the tracking task, receiving minimal practice in the point-to-point task (***Figure 6A***).

In comparison to the main experiment, the total amount of compensation expressed during learning was blunted in both groups (***Figure 6B-D***: average across subjects; ***Figure 6–Figure Supplement 1***: single subjects). Comparing the the off-diagonal gain at different frequencies of movement, the decrement in gain from the main experiment to the follow-up experiment was similar between groups (comparing gains for each group between ***Figures 5B*** and ***6C***). However, the overall trend in learning was different; while the rotation group's performance improved from early to late learning, particularly at low frequencies (linear mixed effects model [see "Statistics" in methods for details about the structure of the model]: 3-way interaction between block, group, and frequency, $F(18, 461) = 3.20$, $p < 0.0001$; Tukey's range test: Bonferroni-adjusted $p < 0.05$ for two out of seven frequencies), the mirror-reversal group's performance plateaued by early learning (Tukey's range test: Bonferroni-adjusted $p = 1$ for all frequencies; ***Figure 6C***). ***Figures 6B*** and ***D*** also demonstrate a similar difference in learning trends. These results are consistent with previous work comparing the generalization of rotation learning from pointing to tracking tasks and vice versa (***Abeele and Bock, 2003***). Ultimately, this suggests that training in the point-to-point task played a critical role in acquiring the ability to track the target under a mirror-reversal learning, but had a lesser impact for learning the rotation.

It is also worth noting that the follow-up experiment reproduced many of the same behavioral phenomena evident in the main experiment. The rotation group exhibited aftereffects, albeit only within the middle bandwidth of frequencies (Tukey's range test: Bonferroni-adjusted $p < 0.05$ for three out of seven frequencies), while the mirror-reversal group did not (Tukey's range test: Bonferroni-adjusted $p = 1$ for all seven frequencies). The rotation group exhibited compensation at high frequencies (Tukey's range test: Bonferroni-adjusted $p < 0.05$ for two out of three highest frequencies) whereas the mirror-reversal group did not (Tukey's range test: Bonferroni-adjusted $p = 1$ for highest three frequencies). In other words, the follow-up experiment provides evidence that the effects we observed in the main experiment are robust and replicable.

Discussion

In the present study, we tested whether participants could learn to successfully control a cursor to pursue a continuously moving target under either rotated or mirror-reversed visual feedback. Although previous work has established that participants can learn to compensate for these perturbations during point-to-point movements, this compensation often seems to be depend upon the use of re-aiming strategies—a solution that is time-consuming and therefore does not seem feasible in a task in which goals are constantly changing. Although it is possible that participants might have performed the tracking task through numerous, intermittent catch-up movements, we found that, under both perturbations, participants' tracking behavior was close to linear across all frequencies (as quantified by spectral coherence), implying instead that they tracked the target smoothly and continuously. As expected, we found that participants who learned to counter the visuomotor rotation exhibited strong aftereffects once the perturbation was removed, amounting to an approximately 25° rotation of hand motion relative to target motion—consistent with previous findings in point-to-point tasks (***Taylor et al., 2010***; ***Fernández-Ruiz et al., 2011***; ***Taylor and Ivry, 2011***; ***Bond and Taylor, 2015***). In contrast, participants who learned to compensate for the mirror-reversal showed no aftereffects, suggesting that they did not engage adaptation of their existing controller, but instead learned to compensate by establishing a *de novo* controller.

One potential caveat of our results is that if the mirror-reversal group expressed a quickly decaying aftereffect (e.g., decaying within a single trial), the temporal resolution of our analysis (tracking analysis performed in 40 second bins) may not have been fine enough to capture such a phenomenon. Indeed, prior work has demonstrated that motor learning can occur on very fast timescales, even when learning to counter random, unexpected perturbations (***Crevecoeur et al.,***

406 ***2020a***,b). However, the timescale of decay that would be necessary to give rise to such a result is
407 so starkly different from that in the rotation group (on the order of several minutes) that it seems
408 unlikely the same learning mechanism can account for behavior in both groups of participants. In
409 summary, our results corroborate previous findings which suggest that people learn to counter
410 rotations and mirror reversals of visual feedback in qualitatively different ways.

411 ## Frequency-Domain Signatures of Adaptation and *De Novo* Learning

412 The pattern of compensation under the rotation and mirror-reversal was frequency specific (***Figure***
413 ***5B***), with the nature of compensation at high frequencies revealing distinct signatures of adaptation
414 and *de novo* learning between the two groups. At low frequencies, both groups of participants
415 successfully compensated for their perturbations. But at high frequencies, only the rotation group
416 was able to compensate; behavior for the mirror-reversal group at high frequencies was similar
417 to baseline behavior. There were similarities, however, in the overall time course and frequency
418 dependence of learning under each perturbation (***Figure 5B***), with both groups exhibiting a steady
419 increase in compensation over time, particularly at lower frequencies. Additionally, both groups'
420 compensation exhibited a similar gradation as a function of frequency, decreasing as frequency
421 increased.

422 We believe these results show that distinct learning processes drove two separate components
423 of learning. One component, present only in the rotation group, was expressed uniformly at
424 all frequencies and likely reflects a parametric adjustment of an existing baseline controller, i.e.,
425 adaptation. This interpretation is consistent with studies demonstrating that adaptation of point-to-
426 point movements generalizes to corrections for unexpected perturbations (***Ahmadi-Pajouh et al.,***
427 ***2012***; ***Cluff and Scott, 2013***; ***Telgen et al., 2014***). A second component of learning contributed
428 to compensation in both groups of participants. This component was expressed primarily at
429 low frequencies, exhibited a gradation as a function of frequency, and was not associated with
430 aftereffects. We suggest this component corresponds to formation of a *de novo* controller for the
431 task. The mirror-reversal group's behavior in ***Figure 5B*** demonstrates the frequency-dependent
432 characteristics of this *de novo* learned controller and how those characteristics evolve with practice.

433 The failure to compensate at high frequencies under the mirror-reversal is consistent with the ob-
434 servation that people who have learned to make point-to-point movements under mirror-reversed
435 feedback are unable to generate appropriate rapid corrections to unexpected perturbations (***Tel-***
436 ***gen et al., 2014***; ***Kasuga et al., 2015***; ***Gritsenko and Kalaska, 2010***). However, importantly, in the
437 tracking task, participants were able to generate appropriate motor output continuously, rather
438 than in a discrete manner, as in the point-to-point task. Although compensation for the rotation
439 bore many hallmarks of adaptation, it also exhibited features of *de novo* learning seen in the
440 mirror-reversal group, suggesting that participants in the rotation group employed a combination
441 of the two learning processes. Our data thus support previous suggestions that residual learning
442 under a visuomotor rotation that cannot be attributed to implicit adaptation may rely on the same
443 mechanisms as those used for *de novo* learning (***Krakauer et al., 2019***).

444 ## Potential Control Architectures Supporting Multiple Components of Learning

445 The properties of adaptation and *de novo* learning we have identified here can potentially be
446 explained by the existence of two distinct control pathways, each capable of different forms of
447 plasticity but with differing sensorimotor delays. An inability to compensate at high frequencies
448 (when tracking an unpredictable stimulus; see ***Roth et al.*** (***2011***)) suggests higher phase lags,
449 potentially due to greater sensorimotor delays or slower system dynamics; as phase lags approach
450 the period of oscillation, it becomes impossible to exert precise control at that frequency. One
451 pathway may be fast but can only be recalibrated through adaptation while the other pathway is
452 slower but can be reconfigured to implement arbitrary new controllers.

453 These two control pathways might correspond to feedforward control (generating motor output
454 based purely on target motion) and feedback control (generating motor output based on the cursor

location and/or distance between cursor and target). Feedback control is slower than feedforward control due to the additional delays associated with observing the effects of one's earlier motor commands on the current cursor position. The observed pattern of behavior may thus be due to a fast but inflexible feedforward controller that responds rapidly to target motion, but always expresses baseline behavior (potentially recalibrated via implicit adaptation) interacting with a slow but reconfigurable feedback controller that responds to both target motion and the current cursor position. At low frequencies, the target may move slowly enough that any inappropriate feedforward control to track the target is masked by corrective feedback responses. But at high frequencies, the target may move too fast for feedback control to be exerted, leaving only inappropriate feedforward responses.

An alternative possibility is that there may be multiple feedforward controllers (and/or feedback controllers) incurring different delays. A fast but inflexible baseline controller (amenable to recalibration through adaptation) might interact with a slower but more flexible controller. This organization parallels dual-process theories of learning and action selection (***Hardwick et al., 2019***; ***Day and Lyon, 2000***; ***Huberdeau et al., 2015***) and raises the possibility that the *de novo* learning exhibited by our participants might be, in some sense, cognitive in nature. Cognitive processes are generally conceived of in terms of discrete associations that require time-consuming, deliberative processing to compute—consistent with the use of re-aiming strategies (***McDougle et al., 2016***; ***Huberdeau et al., 2015***). It is possible, however, for action selection to occur rapidly but still be considered cognitive. For instance, it has been proposed that stimulus-response associations can be "cached" in working memory, enabling a cognitive response to be deployed rapidly and without deliberation (***McDougle and Taylor, 2019***). Caching associations in this way appears to be limited to just 2–7 elements (***McDougle and Taylor, 2019***; ***Collins and Frank, 2012***), raising doubts as to whether such a mechanism could support a feedback controller that must generate output in continuous time and space. Nevertheless, recent theories have framed prefrontal cortex as a general-purpose network capable of learning to perform arbitrary computations on its inputs (***Wang et al., 2018***). From this perspective, it does not seem infeasible that such a network could learn to implement an arbitrary continuous feedback controller that could reasonably be considered cognitive.

Role of Re-Aiming Strategies in Acquisition versus Execution of a *De Novo* Controller

Although participants likely could not have used an aiming strategy to execute continuous tracking movements, they could have used such a strategy to acquire the controller necessary to perform these movements. In a follow-up experiment, we tested whether limited practice in the point-to-point task would impair how well participants could learn to counter the rotation/mirror reversal. Indeed, we found both groups' performance suffered from the lack of point-to-point practice, but whereas the rotation group exhibited moderate improvement from early to late learning, the mirror-reversal group did not improve past early learning. The mirror-reversal group's lack of learning suggests that re-aiming may be important for initially *building* a *de novo* controller. But with point-to-point practice, participants demonstrated tracking improvements under the mirror reversal, suggesting that re-aiming may not be a necessary component of the tracking controller after the initial phase of learning. Likewise, the rotation group may have exhibited a decrement in tracking performance without point-to-point training because this prevented the engagement of re-aiming strategies that may ordinarily facilitate learning of a *de novo* controller, even if the eventually learned controller does not rely on any form of re-aiming. Improvement in tracking performance in this rotation group could be due to adaptation, which is thought to be driven only by sensory prediction errors which would be present and, presumably, also driving learning even during the tracking task. The poor performance in both groups may also have been attributable to participants simply spending less total time practicing their respective perturbations. However, if time on task was the only variable driving learning, then the mirror-reversal group's performance should have improved from early learning to late learning, which it did not.

A re-aiming strategy may be important for building a *de novo* controller because this process may rely on the deliberative computations performed when planning upcoming movements. Alternatively, it may be easier for people to evaluate the quality of straight-line reaches (e.g., reach direction, movement time, task error) compared to random tracking movements, allowing them to update the parameters of a nascent controller using these quality metrics. In any case, our data suggest that cognitive strategies/processes may serve a critical role in facilitating *de novo* learning even if the eventually learned controller does not depend on re-aiming strategies.

System Identification as a Tool for Characterizing Motor Learning

Our characterization of learning made use of frequency-based system identification, a powerful tool that has been previously used to study biological motor control such as insect flight (***Sponberg et al., 2015***; ***Roth et al., 2016***), electric fish refuge tracking (***Cowan and Fortune, 2007***; ***Madhav et al., 2013***), human posture (***Oie et al., 2002***; ***Kiemel et al., 2006***), and human reaching (***Zimmet et al., 2019***). This approach has a number of practical advantages over other methods for studying motor control (e.g., point-to-point reaches)—including time-efficiency for data collection and the availability of a rich suite of tools in the frequency domain (***Schoukens et al., 2004***). However, to our knowledge, frequency-based system identification has not previously been applied to investigate motor learning. Here, we have demonstrated that this approach can not only recapitulate the results of previous studies but also extend these results to identify distinct components of control. Our approach is also general as it can be applied to assess learning of arbitrary linear visuomotor mappings (e.g., 15° rotation, body-machine interfaces (***Mussa-Ivaldi et al., 2011***)). Under previous approaches, characterizing the quality of movements under different types of learned mappings (rotation, mirror-reversal) has necessitated different *ad hoc* analyses that cannot be directly compared (***Telgen et al., 2014***). In contrast, our frequency-based approach provides a general method to characterize behavior under rotations, mirror-reversals, or any linear mapping from effectors to a cursor.

While the system identification approach used in the present study does capture learning, the results obtained using this approach do warrant careful interpretation. In particular, one must not interpret the empirical relationship that we measure between the target and hand as equivalent to the input–output relationship of the brain's motor controller. The former measures the response of the entire sensorimotor system to external input. The latter only measures how the controller sends motor commands to the body in response to input from the environment/internal feedback. Estimating the latter relationship requires a more nuanced approach that takes into account the closed-loop topology (***Roth et al., 2014***). Despite this, changes to the controller are still revealed using our approach; assuming that learning only drives changes in the input–output relationship of the controller—as opposed to, for example, the plant or the visual system—any changes in the overall target–hand relationship will reflect changes to the controller. Thus, our approach is a valid way to investigate learning.

Although the primary goal of our frequency-based analysis was to establish how participants mapped target motion into hand motion, system identification yields more detailed information than this; in principle, it provides complete knowledge of a linear system in that knowing how the system responds to sinusoidal input at different frequencies enables one to predict how the system will respond to arbitrary inputs. This data can be used to formally compare different possible control system architectures (***Zimmet et al., 2019***) supporting learning, and we plan to explore this more detailed analysis in future work.

De Novo Learning and Real-World Skill Learning

We have used the term "*de novo* learning" to refer to any mechanism, aside from implicit adaptation, that leads to the creation of a new controller. We propose that *de novo* learning proceeds initially through explicit processes before becoming cached or automatized into a more procedural form. There are, however, a number of alternative mechanisms that could be engaged to establish a new controller. One proposal is that *de novo* learning occurs by simultaneously updating forward and

inverse models by simple gradient descent (***Pierella et al., 2019***). Another possibility is that a new controller could be learned through reinforcement learning. In motor learning tasks, reinforcement has been demonstrated to engage a learning mechanism that is independent of implicit adaptation (***Izawa and Shadmehr, 2011***; ***Cashaback et al., 2017***; ***Holland et al., 2018***) potentially via basal-ganglia-dependent mechanisms (***Schultz et al., 1997***; ***Hikosaka et al., 2002***). Such reinforcement could provide a basis for forming a new controller. Although prior work on motor learning has focused on simply learning the required direction for a point-to-point movement, theoretical frameworks for reinforcement learning have been extended to continuous time and space to learn continuous controllers for robotics (***Doya, 2000***; ***Theodorou et al., 2010***; ***Smart and Kaelbling, 2000***; ***Todorov, 2009***), and such theories could be applicable to how people learned continuous control in our experiment.

Although we have described the mirror-reversal task as requiring *de novo* learning, we acknowledge that there are many types of learning which might be described as *de novo* learning that this task does not capture. For example, many skills, such as playing the cello, challenge one to learn how to *execute* new movement patterns that one has never executed before (***Costa, 2011***). This is not the case in the tracking task which only challenges one to *select* movements one already knows how to execute. Also, in many cases, one must learn to use information from new sensory modalities for control (***van Vugt and Ostry, 2018***; ***Bach-y-Rita and W. Kercel, 2003***), such as using auditory feedback to adjust one's finger positioning while playing the cello. Our task, by contrast, only uses very familiar visual cues. Nevertheless, we believe that learning a new controller that maps familiar sensory feedback to well-practiced actions in a novel way is a critical element of many real-world learning tasks (e.g., driving a car, playing video games) and should be considered a fundamental aspect of any *de novo* learning.

Ultimately, our goal is to understand real-world skill learning. We believe that studying learning in continuous tracking tasks is important to bring us closer to this goal since a critical component of many skills is the ability to continuously control an effector in response to ongoing external events, like in juggling or riding a bicycle. Studies of *well-practiced* human behavior in continuous control tasks has a long history, such as those examining the dynamics of pilot and vehicle interactions (***McRuer and Jex, 1967***). However, most existing paradigms for studying motor *learning* have examined only point-to-point movements. We believe the tracking task presented here offers a simple but powerful approach for characterizing how we learn a new continuous controller and, as such, provides an important new direction for advancing our understanding of how real-world skills are acquired.

Methods and Materials

Participants

40 right-handed, healthy participants over 18 years of age were recruited for this study (24.28 $\pm$ 5.06 years old; 19 male, 21 female), 20 for the main experiment (***Figures 2–5***) and 20 for the follow-up experiment (***Figure 6***). Participants all reported having no history of neurological disorders. All methods were approved by the Johns Hopkins School of Medicine Institutional Review Board.

Tasks

Participants made planar movements with their right arm, which was supported by a frictionless air sled on a table, to control a cursor on an LCD monitor (60 Hz). Participants viewed the cursor on a horizontal mirror which reflected the monitor (***Figure 1B***). Hand movement was monitored with a Flock of Birds magnetic tracker (Ascension Technology, VT, USA) at 130 Hz. The (positive) x axis was defined as rightward, and the y axis, forward. The cursor was controlled under three different hand-to-cursor mappings: 1) veridical, 2) 90° visuomotor rotation, and 3) mirror reversal about the 45° oblique axis in the $(x, y) = (1, 1)$ direction. Participants were divided evenly into two groups, one that experienced the visuomotor rotation ($n = 10$; 4 male, 6 female) and one that experienced the

mirror reversal ($n = 10$; 6 male, 4 female). Both groups were exposed to the perturbed cursors while performing two different tasks: 1) the point-to-point task, and 2) the tracking task.

Point-to-point task

To start a trial, participants were required to move their cursor (circle of radius 2.5 mm) into a target (grey circle of radius 10 mm) that appeared in the center of the screen. After 500 ms, the target appeared 12 cm away from the starting location in a random direction. Participants were instructed to move in a straight line, as quickly and accurately as possible to the new target. Once the cursor remained stationary (speed < 0.065 m/s) in the new target for 1 sec, the target appeared in a new location 12 cm away, but constrained to lie within a 20 cm $\times$ 20 cm workspace. Each block used different, random target locations from other blocks. Blocks in the main experiment consisted of 150 reaches while blocks in the follow-up experiment (***Figure 6***) consisted of 15 reaches. To encourage participants to move quickly to each target, we provided feedback at the end of each trial about the peak velocity they attained during their reaches, giving positive feedback (a pleasant tone and the target turning yellow) if the peak velocity exceeded roughly 0.39 m/s and negative feedback (no tone and the target turning blue) if the peak velocity was below that threshold.

Tracking task

At the start of each trial, a motionless target (grey circle of radius 8 mm) appeared in the center of the screen, and the trial was initiated when the participant's cursor (circle of radius 2.5 mm) was stationary (speed < 0.065 m/s) in the target. From then, the target began to move for 46 seconds in a continuous, pseudo-random trajectory. The first 5 seconds was a ramp period where the amplitude of the cursor increased linearly from 0 to its full value, and for the remaining 41 seconds, the target moved at full amplitude. The target moved in a two-dimensional, sum-of-sinusoids trajectory; fourteen sinusoids of different frequencies, amplitudes and phases were summed to determine target movement in the x-axis (frequencies (Hz): 0.1, 0.25, 0.55, 0.85, 1.15, 1.55, 2.05; amplitudes (cm): 2.31, 2.31, 2.31, 1.76, 1.30, 0.97, 0.73, respectively), and the y-axis (frequencies (Hz): 0.15, 0.35, 0.65, 0.95, 1.45, 1.85, 2.15; amplitudes (cm): 2.31, 2.31, 2.31, 1.58, 1.03, 0.81, 0.70, respectively). Different frequencies were used for the x- and y-axes so that hand movements at a given frequency could be attributed to either x- or y-axis target movements. All frequencies were prime multiples of 0.05 Hz to ensure that the harmonics of any target frequency would not overlap with any other target frequency. The amplitudes of the sinusoids for all but the lowest frequencies were proportional to the inverse of their frequency to ensure that each individual sinusoid had similar peak velocity. We set a ceiling amplitude for low frequencies in order to prevent target movements that were too large for participants to comfortably track. Note that due to the construction of the input signal, there were no low-order harmonic relations between any of the component sinusoids on the input, making it likely that nonlinearities in the tracking dynamics would manifest as easily discernible harmonic interactions (i.e. extraneous peaks in the output spectra). Moreover, care was taken to avoid "frequency leakage" by designing discrete Fourier transform windows that were integer multiples of the base period (20 s), improving our ability to detect such nonlinearities.

Participants were instructed to keep their cursor inside the target for as long as possible during the trial. The target's color changed to yellow anytime the cursor was inside the target to provide feedback for their success. One block of the tracking task consisted of eight, 46-second trials, and the same target trajectory was used for every trial within a block. For different blocks, we randomized the phases, but not the frequencies, of target sinusoids to produce different trajectories. We produced five different target trajectories for participants to track in the six tracking blocks. The trajectory used for baseline and post-learning were the same to allow a better comparison for aftereffects. All participants tracked the same five target trajectories, but the order in which they experienced these trajectories was randomized in order to minimize any phase-dependent learning effects.

Experiment

We first assessed the baseline control of both groups of participants by having them perform one block of the tracking task followed by one block of the point-to-point task under veridical cursor feedback. We then applied either the visuomotor rotation or mirror reversal to the cursor, and used the tracking task to measure their control capabilities during early learning. Afterwards, we alternated three times between blocks of point-to-point training and blocks of tracking. In total, each participant received 450 point-to-point reaches of training under perturbed cursor feedback. Finally, we measured aftereffects post-learning by returning to the veridical mapping and using the tracking task.

Data Analysis

Analyses were performed in MATLAB R2018b (The Mathworks, Natick, MA, USA) and R version 4.0.2 (RStudio, Inc., Boston, MA, USA) (***R Core Team, 2016***; ***Lenth, 2015***; ***Pinheiro et al., 2016***; ***Lenth, 2016***). Figures were created using Adobe Illustrator (Adobe Inc., San Jose, CA, USA).

Trajectory-alignment analysis

In the point-to-point task, we assessed performance by calculating the angular error between the cursor's initial movement direction and the target direction relative to the start position. To determine the cursor's initial movement direction, we computed the direction of the cursor's instantaneous velocity vector ~150 ms after the time of movement initiation. Movement initiation was defined as the time when the cursor left the start circle on a given trial.

In the tracking task, we assessed performance by measuring the average mean-squared error between the hand and target positions for every trial. For the alignment matrix analysis, we fit a matrix, $\hat{M} = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$, that minimized the mean-squared error between the hand and target trajectories for every trial. In the latter analysis, the mean-squared error was additionally minimized in time by delaying the target trajectory relative to the hand. (While the time-delay allowed for the fairest possible comparison between the hand and target trajectories in subsequent analysis, changing or eliminating the alignment *did not* qualitatively change our results.) We estimated $\hat{M}$ as

$$\hat{M} = \underset{M}{\operatorname{argmin}} \left\{ \begin{bmatrix} H_x \\ H_y \end{bmatrix} - M \begin{bmatrix} T_x \\ T_y \end{bmatrix} \right\} \tag{1}$$

where H and T represent hand and target trajectories. These estimated $\hat{M}$'s were averaged element-wise across participants to generate the alignment matrices shown in ***Figure 3A***. These matrices were visualized by plotting their column vectors, also shown in ***Figure 3A***.

The off-diagonal elements of each participant's alignment matrix were used to calculate the off-diagonal scaling, S, in ***Figure 3B***:

$$S_{\text{rotation}} = \frac{-b + c}{2}, \qquad S_{\text{mirror}} = \frac{b + c}{2}. \tag{2}$$

Compensation angles, θ, for the rotation group's alignment matrices were found using the singular value decomposition, $\text{SVD}(\cdot)$. This is a standard approach which, as described in ***Umeyama*** (***1991***), identifies a pure rotation, R, that best describes $\hat{M}$ irrespective of other transformations (e.g., dilation, shear) (***Figure 3C***, left). Briefly,

$$U \Sigma V^T = \text{SVD}(\hat{M}^T), \tag{3}$$

$$R = V U^T \tag{4}$$

where U and V contain the left and right singular vectors and Σ contains the singular values. Note that R is a rotation matrix only if $\det(\hat{M}^T) \geq 0$, but R is a reflection matrix when $\det(\hat{M}^T) < 0$. Although ***Umeyama*** (***1991***) have described a method whereby all R can be forced to be a rotation

matrix, we did not want to impose nonexistent structure onto R and, thus, did not analyze trials which yielded reflection matrices. However, this was not a major issue for the analysis as nearly all trials yielded rotation matrices. Subsequently, θ was calculated as

$$\theta = \mathrm{atan2}(R_{2,1}, R_{1,1}) \tag{5}$$

where $\mathrm{atan2}(\cdot)$ is the 2-argument arctangent.

Finally, for the mirror-reversal group, the scaling orthogonal to the mirror axis was found by computing how the matrix transformed the unit vector along the orthogonal axis (***Figure 3C***, right):

$$S_{\text{orthogonal}} = \frac{1}{2}\left(\begin{bmatrix}1 & -1\end{bmatrix}\begin{bmatrix}a & b\\ c & d\end{bmatrix}\begin{bmatrix}1\\ -1\end{bmatrix}\right) = \frac{1}{2}(a-b-c+d). \tag{6}$$

Frequency-domain analysis

To analyze trajectories in the frequency domain, we applied the discrete Fourier transform to the target and hand trajectories in every tracking trial. This produced a set of complex numbers representing the amplitude and phase of the signal at every frequency. We only analyzed the first 40 seconds of the trajectory that followed the 5-second ramp period so that our analysis period was equivalent to an integer multiple of the base period (20 s). This ensured that we would obtain clean estimates of the sinusoids at each target frequency. Amplitude spectra were generated by taking double the modulus of the Fourier-transformed hand trajectories at positive frequencies.

Spectral coherence was calculated between the target and hand trajectories (***Stoica and Moses, 2005***). To do so, we evaluated the single-input multi-output coherence at every frequency of target motion, determining how target motion in one axis elicited hand movement in both axes. This best captured the linearity of participants' behavior as using hand movement in only one axis for the analysis would only partially capture participants' responses to target movement at a given frequency. Calculations were performed using a 1040-sample Blackman-Harris window with 50% overlap between segments.

During each 40 s stimulus period, we assumed the relationship between target position and hand position behavior was well approximated by linear, time-invariant dynamics; this assumption was tested using the coherence analysis described above. Under this assumption, pure sinusoidal target motion at each frequency should be translated into pure sinusoidal hand motion at the same frequency but with different magnitude and phase. The relationship between hand and target can therefore be described in terms of a 2×2 matrix of transfer functions describing the behavior of the system at each possible frequency:

$$\begin{bmatrix}H_x(\omega)\\ H_y(\omega)\end{bmatrix} = P(\omega)\begin{bmatrix}T_x(\omega)\\ T_y(\omega)\end{bmatrix}, \quad P(\omega) = \begin{bmatrix}p_{xx}(\omega) & p_{xy}(\omega)\\ p_{yx}(\omega) & p_{yy}(\omega)\end{bmatrix}. \tag{7}$$

Here, $H(\omega)$ and $T(\omega)$ are the Fourier transforms of the time-domain hand and target trajectories, respectively, and ω is the frequency of movement. Each element of $P(\omega)$ represents a transfer function relating a particular axis of target motion to a particular axis of hand motion; the first and second subscripts represent the hand- and target-movement axes, respectively. Each such transfer function is a complex-valued function of frequency, which can further be decomposed into gain and phase components, e.g.:

$$p_{xy}(\omega) = g_{xy}(\omega)e^{\phi_{xy}(\omega)}, \tag{8}$$

where $g_{xy}(\omega)$ describes the gain (ratio of amplitudes) between y-target and x-hand motion as a function of frequency, and $\phi_{xy}(\omega)$ describes the corresponding difference in the phase of oscillation.

We estimated the elements of $P(\omega)$ for frequencies at which the target moved by first noting that, for x-axis frequencies ω, $T_y(\omega) = 0$. Consequently,

$$\begin{bmatrix}H_x(\omega)\\ H_y(\omega)\end{bmatrix} = \begin{bmatrix}p_{xx}(\omega)T_x(\omega)\\ p_{yx}(\omega)T_x(\omega)\end{bmatrix}, \tag{9}$$

and we can therefore estimate $p_{xx}(\omega)$ and $p_{yx}(\omega)$ as:

$$p_{xx}(\omega) = \frac{H_x(\omega)}{T_x(\omega)}, \qquad p_{yx}(\omega) = \frac{H_y(\omega)}{T_x(\omega)}. \tag{10}$$

These quantities are also known as phasors, complex numbers which describe the gain and phase relationship between the hand and target. We estimated $p_{yx}(\omega)$ and $p_{yy}(\omega)$ analogously at y-frequencies of target motion.

These estimates yielded two elements of the overall transformation matrix $P(\omega)$ at each frequency of target movement. In order to construct a full 2×2 matrix, we paired the gains from neighboring x- and y-frequencies, assuming that participants' behavior would be approximately the same at neighboring frequencies. For example, the pair of gains from the lowest x-target frequency (0.1 Hz) was grouped with the analogous gains from the lowest y-target frequency (0.15 Hz) to construct a 2×2 matrix.

The spatial transformation of target motion into hand motion at each frequency is described by the gain of each element of $P(\omega)$. However, gain and phase data can lead to certain ambiguities; for example, a positive gain with a phase of π radians is indistinguishable from a negative gain with a phase of 0. Conventionally, this is resolved by assuming that gain is positive. In our task, however, the sign of the gain was crucial to disambiguate the directionality of the hand responses (e.g., whether the hand moved left or right in response to upward target motion). We used phase information to disambiguate positive from negative gains. Specifically, we assumed that the phase lag of the hand response at a given frequency would be similar across both axes of hand movement and throughout the experiment, but the gain would vary:

$$p_{xx}(\omega) \approx g_{xx}(\omega)e^{\widetilde{\phi}(\omega)}, \qquad p_{yx}(\omega) \approx g_{yx}(\omega)e^{\widetilde{\phi}(\omega)}. \tag{11}$$

For a given movement frequency, $\widetilde{\phi}(\omega)$ was set to be the same as the mean phase lag during the baseline block, where the gain was unambiguously positive. This assumption enabled us to estimate a signed gain for each phasor using a least-squares approach. This method thus yielded estimated gains for each axis of hand motion, at each target frequency, and at each point during learning. As we did for the transfer-function matrix $P(\omega)$, we paired the estimated gains from neighboring frequencies to obtain a series of seven gain matrices describing, geometrically, how target motion was translated into hand motion from low to high frequencies. Similar to the alignment matrix analysis, visualizations of these gain matrices were constructed by plotting the column vectors of the matrices. Off-diagonal gain, rotation angle, and gain orthogonal to the mirroring axis—all shown in ***Fig 5***—were calculated in the same way as in ***Equations 2–6***.

Statistics

The primary statistical tests for the main and follow-up experiments were performed using linear mixed-effects models. Primary mixed-effects models were fit using data from three parts of the study: 1) alignment matrix analysis in the main experiment, 2) gain matrix analysis in the main experiment, and 3) gain matrix analysis in the follow-up experiment. The data used in these models were the off-diagonal values of the transformation and gain matrices. In all models, data from the first trial of baseline, the last trial of late learning, and the first trial of post-learning were analyzed. For the follow-up experiment, we also included data from the first trial of early learning. Additionally, we removed outliers (25 out of 560 data points) from the follow-up experiment data as the effects demonstrated by most subjects were either greatly magnified or attenuated by one or two subjects. Outliers were defined as data that was 1.5 interquartile ranges outside the data from a given trial and group. Outlier rejection was not performed on data from the main experiment. Using Wilkinson notation, the structure of the model for the alignment matrix analysis was [off-diagonal scaling] ~ [block of learning] * [perturbation group] while the structure for both gain matrix analyses was [off-diagonal gain] ~ [block of learning] * [perturbation group] * [frequency of movement]. Data were grouped within subjects.

We subsequently performed post-hoc statistical comparisons as needed for each of the linear mixed-effects models. For the alignment matrix analysis, we performed pairwise comparisons using Tukey's range test. For the gain matrix analysis in the main and follow-up experiments, there was a 3-way interaction between frequency and the other regressors, so we fit seven different mixed-effects models for each frequency of movement post-hoc. We performed pairwise comparisons on these frequency-specific models using Tukey's range test, Bonferroni correcting by the number of fitted models (i.e., 7).

Data and Code Availability

The data and code used to produce the results in this study can be found here: Johns Hopkins University Data Archive

Competing Interests

The authors declare no competing interests.

Acknowledgments

We thank Amanda Zimmet, John Krakauer, Amy Bastian, and Christopher Fetsch for immensely helpful discussions. This material is based upon work supported by the National Science Foundation under Grant No. 1825489. C.S.Y. was supported by NIH 5T32NS091018-17, 5T32NS091018-18, and the Link Foundation Modeling, Simulation & Training Fellowship.

References

Abdelghani MN, Lillicrap TP, Tweed DB. Sensitivity derivatives for flexible sensorimotor learning. Neural Comput. 2008 Aug; 20(8):2085–2111. doi: 10.1162/neco.2008.04-07-507.

Abeele S, Bock O. Transfer of Sensorimotor Adaptation between Different Movement Categories. Exp Brain Res. 2003 Jan; 148(1):128–132. doi: 10.1007/s00221-002-1317-0.

Ahmadi-Pajouh MA, Towhidkhah F, Shadmehr R. Preparing to Reach: Selecting an Adaptive Long-Latency Feedback Controller. J Neurosci. 2012 Jul; 32(28):9537–9545. doi: 10.1523/JNEUROSCI.4275-11.2012.

Bach-y-Rita P, W Kercel S. Sensory Substitution and the Human–Machine Interface. Trends in Cognitive Sciences. 2003 Dec; 7(12):541–546. doi: 10.1016/j.tics.2003.10.013.

Bock O. Basic principles of sensorimotor adaptation to different distortions with different effectors and movement types: a review and synthesis of behavioral findings. Front Hum Neurosci. 2013 Mar; 7:81. doi: 10.3389/fnhum.2013.00081.

Bock O, Schneider S. Acquisition of a sensorimotor skill in younger and older adults. Acta Physiol Pharmacol Bulg. 2001 Dec; 26(1-2):89–92.

Bock O, Schneider S, Bloomberg J. Conditions for interference versus facilitation during sequential sensorimotor adaptation. Exp Brain Res. 2001 Jan; 138(3):359–365. doi: 10.1007/s002210100704.

Bond KM, Taylor JA. Flexible Explicit but Rigid Implicit Learning in a Visuomotor Adaptation Task. J Neurophysiol. 2015 Jun; 113(10):3836–3849. doi: 10.1152/jn.00009.2015.

Cashaback JGA, McGregor HR, Mohatarem A, Gribble PL. Dissociating Error-Based and Reinforcement-Based Loss Functions during Sensorimotor Learning. PLoS Comput Biol. 2017 Jul; 13(7):e1005623. doi: 10.1371/journal.pcbi.1005623.

Choi JT, Bastian AJ. Adaptation Reveals Independent Control Networks for Human Walking. Nat Neurosci. 2007 Aug; 10(8):1055–1062. doi: 10.1038/nn1930.

Cluff T, Scott SH. Rapid Feedback Responses Correlate with Reach Adaptation and Properties of Novel Upper Limb Loads. J Neurosci. 2013 Oct; 33(40):15903–15914. doi: 10.1523/JNEUROSCI.0263-13.2013.

Collins AGE, Frank MJ. How Much of Reinforcement Learning Is Working Memory, Not Reinforcement Learning? A Behavioral, Computational, and Neurogenetic Analysis. Eur J Neurosci. 2012 Apr; 35(7):1024–1035. doi: 10.1111/j.1460-9568.2011.07980.x.

Costa RM. A Selectionist Account of de Novo Action Learning. Curr Opin Neurobiol. 2011 Aug; 21(4):579–586. doi: 10.1016/j.conb.2011.05.004.

Cowan NJ, Fortune ES. The Critical Role of Locomotion Mechanics in Decoding Sensory Systems. J Neurosci. 2007 Jan; 27(5):1123–1128. doi: 10.1523/JNEUROSCI.4198-06.2007.

Craik KJW. Theory of the human operator in control systems. I. The operator as an engineering system. Br J Psychol. 1947 Dec; 38:56–61.

Crevecoeur F, Mathew J, Bastin M, Lefèvre P. Feedback Adaptation to Unpredictable Force Fields in 250 ms. eNeuro. 2020; 7(2). doi: 10.1523/ENEURO.0400-19.2020.

Crevecoeur F, Thonnard JL, Lefèvre P. A Very Fast Time Scale of Human Motor Adaptation: Within Movement Adjustments of Internal Representations during Reaching. eNeuro. 2020; 7(1). doi: 10.1523/ENEURO.0149-19.2019.

Day BL, Lyon IN. Voluntary Modification of Automatic Arm Movements Evoked by Motion of a Visual Target. Exp Brain Res. 2000 Jan; 130(2):159–168. doi: 10.1007/s002219900218.

Doya K. Reinforcement Learning in Continuous Time and Space. Neural Comput. 2000 Jan; 12(1):219–245. doi: 10.1162/089976600300015961.

Fernández-Ruiz J, Díaz R. Prism Adaptation and Aftereffect: Specifying the Properties of a Procedural Memory System. Learn Mem. 1999 Jan; 6(1):47–53.

Fernández-Ruiz J, Wong W, Armstrong IT, Flanagan JR. Relation between Reaction Time and Reach Errors during Visuomotor Adaptation. Behav Brain Res. 2011 May; 219(1):8–14. doi: 10.1016/j.bbr.2010.11.060.

Finley JM, Long A, Bastian AJ, Torres-Oviedo G. Spatial and Temporal Control Contribute to Step Length Asymmetry during Split-Belt Adaptation and Hemiparetic Gait. Neurorehabil Neural Repair. 2015 Sep; 29(8):786–795. doi: 10.1177/1545968314567149.

Gritsenko V, Kalaska JF. Rapid Online Correction Is Selectively Suppressed during Movement with a Visuomotor Transformation. J Neurophysiol. 2010 Dec; 104(6):3084–3104. doi: 10.1152/jn.00909.2009.

Gutierrez-Garralda JM, Moreno-Briseño P, Boll MC, Morgado-Valle C, Campos-Romo A, Diaz R, Fernandez-Ruiz J. The Effect of Parkinson's Disease and Huntington's Disease on Human Visuomotor Learning. Eur J Neurosci. 2013 Sep; 38(6):2933–2940. doi: 10.1111/ejn.12288.

Hadjiosif AM, Krakauer JW, Haith AM. Did we get sensorimotor adaptation wrong? Implicit adaptation as direct policy updating rather than forward-model-based learning. bioRxiv. 2020 Jan; doi: 10.1101/2020.01.22.914473.

Haith AM, Huberdeau DM, Krakauer JW. The Influence of Movement Preparation Time on the Expression of Visuomotor Learning and Savings. J Neurosci. 2015 Apr; 35(13):5109–5117. doi: 10.1523/JNEUROSCI.3869-14.2015.

Hardwick RM, Forrence AD, Krakauer JW, Haith AM. Time-Dependent Competition between Goal-Directed and Habitual Response Preparation. Nat Hum Behav. 2019 Sep; 3(12):1252–1262. doi: 10.1038/s41562-019-0725-0.

Hikosaka O, Nakamura K, Sakai K, Nakahara H. Central Mechanisms of Motor Skill Learning. Curr Opin Neurobiol. 2002 Apr; 12(2):217–222. doi: 10.1016/s0959-4388(02)00307-0.

Holland P, Codol O, Galea JM. Contribution of Explicit Processes to Reinforcement-Based Motor Learning. J Neurophysiol. 2018 Jan; 119(6):2241–2255. doi: 10.1152/jn.00901.2017.

Huberdeau DM, Krakauer JW, Haith AM. Dual-Process Decomposition in Human Sensorimotor Adaptation. Curr Opin Neurobiol. 2015 Aug; 33:71–77. doi: 10.1016/j.conb.2015.03.003.

Izawa J, Shadmehr R. Learning from Sensory and Reward Prediction Errors during Motor Adaptation. PLoS Comput Biol. 2011 Mar; 7(3):e1002012. doi: 10.1371/journal.pcbi.1002012.

Jayakumar RP, Madhav MS, Savelli F, Blair HT, Cowan NJ, Knierim JJ. Recalibration of path integration in hippocampal place cells. Nature. 2019; 566(745):533–537. doi: 10.1038/s41586-019-0939-3.

Kasuga S, Telgen S, Ushiba J, Nozaki D, Diedrichsen J. Learning Feedback and Feedforward Control in a Mirror-Reversed Visual Environment. J Neurophysiol. 2015 Oct; 114(4):2187–2193. doi: 10.1152/jn.00096.2015.

875 **Kiemel T**, Oie KS, Jeka JJ. Slow Dynamics of Postural Sway Are in the Feedback Loop. J Neurophysiol. 2006 Mar;
876 95(3):1410–1418. doi: 10.1152/jn.01144.2004.

877 **Kluzik J**, Diedrichsen J, Shadmehr R, Bastian AJ. Reach Adaptation: What Determines Whether We Learn an
878 Internal Model of the Tool or Adapt the Model of Our Arm? J Neurophysiol. 2008 Sep; 100(3):1455–1464. doi:
879 10.1152/jn.90334.2008.

880 **Krakauer JW**, Ghilardi MF, Ghez C. Independent Learning of Internal Models for Kinematic and Dynamic Control
881 of Reaching. Nat Neurosci. 1999 Nov; 2(11):1026–1031. doi: 10.1038/14826.

882 **Krakauer JW**, Hadjiosif AM, Xu J, Wong AL, Haith AM. Motor Learning. Compr Physiol. 2019 Mar; 9(2):613–663.
883 doi: 10.1002/cphy.c170043.

884 **Lackner JR**, Dizio P. Rapid Adaptation to Coriolis Force Perturbations of Arm Trajectory. J Neurophysiol. 1994
885 Jul; 72(1):299–313. doi: 10.1152/jn.1994.72.1.299.

886 **Lenth RV**. estimability: Tools for Assessing Estimability of Linear Predictions; 2015, https://CRAN.R-project.org/
887 package=estimability, r package version 1.1-1.

888 **Lenth RV**. Least-Squares Means: The R Package lsmeans. J Stat Softw. 2016; 69(1):1–33. doi:
889 10.18637/jss.v069.i01.

890 **Leow LA**, Gunn R, Marinovic W, Carroll TJ. Estimating the Implicit Component of Visuomotor Rotation
891 Learning by Constraining Movement Preparation Time. J Neurophysiol. 2017 Jan; 118(2):666–676. doi:
892 10.1152/jn.00834.2016.

893 **Lillicrap TP**, Moreno-Briseño P, Diaz R, Tweed DB, Troje NF, Fernandez-Ruiz J. Adapting to Inversion of the Visual
894 Field: A New Twist on an Old Problem. Exp Brain Res. 2013 Jul; 228(3):327–339. doi: 10.1007/s00221-013-
895 3565-6.

896 **Madhav MS**, Stamper SA, Fortune ES, Cowan NJ. Closed-Loop Stabilization of the Jamming Avoidance Response
897 Reveals Its Locally Unstable and Globally Nonlinear Dynamics. J Exp Biol. 2013 Nov; 216(Pt 22):4272–4284.
898 doi: 10.1242/jeb.088922.

899 **Martin TA**, Keating JG, Goodkin HP, Bastian AJ, Thach WT. Throwing While Looking through Prisms. I. Focal Olivo-
900 cerebellar Lesions Impair Adaptation. Brain. 1996 Aug; 119 (Pt 4):1183–1198. doi: 10.1093/brain/119.4.1183.

901 **Maschke M**, Gomez CM, Ebner TJ, Konczak J. Hereditary Cerebellar Ataxia Progressively Impairs Force Adaptation
902 during Goal-Directed Arm Movements. J Neurophysiol. 2004 Jan; 91(1):230–238. doi: 10.1152/jn.00557.2003.

903 **Mazzoni P**, Krakauer JW. An Implicit Plan Overrides an Explicit Strategy during Visuomotor Adaptation. J
904 Neurosci. 2006 Apr; 26(14):3642–3645. doi: 10.1523/JNEUROSCI.5317-05.2006.

905 **McDougle SD**, Ivry RB, Taylor JA. Taking Aim at the Cognitive Side of Learning in Sensorimotor Adaptation Tasks.
906 Trends Cogn Sci (Regul Ed). 2016 Jul; 20(7):535–544. doi: 10.1016/j.tics.2016.05.002.

907 **McDougle SD**, Taylor JA. Dissociable Cognitive Strategies for Sensorimotor Learning. Nat Commun. 2019 Mar;
908 10(1):40. doi: 10.1038/s41467-018-07941-0.

909 **McRuer DT**, Jex HR. A Review of Quasi-Linear Pilot Models. IEEE Trans Hum Factors Electron. 1967 Sep;
910 HFE-8(3):231–249. doi: 10.1109/THFE.1967.234304.

911 **Miall RC**, Weir DJ, Stein JF. Intermittency in Human Manual Tracking Tasks. J of Mot Behav. 1993 04; 25:53–63.
912 doi: 10.1080/00222895.1993.9941639.

913 **Miall RC**, Weir DJ, Wolpert DM, Stein JF. Is the Cerebellum a Smith Predictor? J Mot Behav. 1993 Sep; 25(3):203–
914 216. doi: 10.1080/00222895.1993.9942050.

915 **Morehead JR**, Qasim SE, Crossley MJ, Ivry R. Savings upon Re-Aiming in Visuomotor Adaptation. J Neurosci.
916 2015 Oct; 35(42):14386–14396. doi: 10.1523/JNEUROSCI.1046-15.2015.

917 **Morehead JR**, Taylor JA, Parvin DE, Ivry RB. Characteristics of Implicit Sensorimotor Adaptation Revealed by
918 Task-Irrelevant Clamped Feedback. J Cogn Neurosci. 2017 Jun; 29(6):1061–1074. doi: 10.1162/jocn_a_01108.

919 **Morton SM**, Bastian AJ. Cerebellar Contributions to Locomotor Adaptations during Splitbelt Treadmill Walking.
920 J Neurosci. 2006 Sep; 26(36):9107–9116. doi: 10.1523/JNEUROSCI.2622-06.2006.

Mussa-Ivaldi FA, Casadio M, Danziger ZC, Mosier KM, Scheidt RA. Sensory motor remapping of space in human–machine interfaces. Prog Brain Res. 2011; 191:45–64. doi: 10.1016/B978-0-444-53752-2.00014-X.

Oie KS, Kiemel T, Jeka JJ. Multisensory Fusion: Simultaneous Re-Weighting of Vision and Touch for the Control of Human Posture. Brain Res Cogn Brain Res. 2002 Jun; 14(1):164–176.

Pierella C, Casadio M, Mussa-Ivaldi FA, Solla SA. The dynamics of motor learning through the formation of internal models. PLoS Comput Biol. 2019 Dec; 15(12):e1007118. doi: 10.1101/652727.

Pinheiro J, Bates D, DebRoy S, Sarkar D, R Core Team. nlme: Linear and Nonlinear Mixed Effects Models; 2016, http://CRAN.R-project.org/package=nlme, r package version 3.1-128.

R Core Team. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria; 2016, https://www.R-project.org/.

Redding GM, Wallace B. Adaptive Coordination and Alignment of Eye and Hand. J Mot Behav. 1993 Jun; 25(2):75–88. doi: 10.1080/00222895.1993.9941642.

Roddey JC, Girish B, Miller JP. Assessing the Performance of Neural Encoding Models in the Presence of Noise. J Comput Neurosci. 2000; 8(2):95–112.

Roth E, Hall RW, Daniel TL, Sponberg S. Integration of Parallel Mechanosensory and Visual Pathways Resolved through Sensory Conflict. Proc Natl Acad Sci USA. 2016 Nov; 113(45):12832–12837. doi: 10.1073/pnas.1522419113.

Roth E, Sponberg S, Cowan NJ. A Comparative Approach to Closed-Loop Computation. Current Opinion in Neurobiology. 2014; 25:54–62. http://dx.doi.org/10.1016/j.conb.2013.11.005, doi: 10.1016/j.conb.2013.11.005.

Roth E, Zhuang K, Stamper SA, Fortune ES, Cowan NJ. Stimulus Predictability Mediates a Switch in Locomotor Smooth Pursuit Performance for Eigenmannia Virescens. J Exp Biol. 2011 Apr; 214(Pt 7):1170–1180. doi: 10.1242/jeb.048124.

de Rugy A, Loeb GE, Carroll TJ. Muscle coordination is habitual rather than optimal. J Neurosci. 2012 May; 32(21):7384–91. doi: 10.1523/JNEUROSCI.5792-11.2012.

Russell DM, Sternad D. sinusoidal visuomotor tracking: intermittent servo-control or coupled oscillations? J Mot Behav. 2001 Dec; 33:329–49.

Schoukens J, Pintelon R, Rolain Y. Time Domain Identification, Frequency Domain Identification. Equivalencies! Differences? In: *Proceedings of the 2004 American Control Conference*, vol. 1; 2004. p. 661–666. doi: 10.23919/ACC.2004.1383679.

Schugens MM, Breitenstein C, Ackermann H, Daum I. Role of the Striatum and the Cerebellum in Motor Skill Acquisition. Behav Neurol. 1998; 11(3):149–157.

Schultz W, Dayan P, Montague PR. A Neural Substrate of Prediction and Reward. Science. 1997 Mar; 275(5306):1593–1599. doi: 10.1126/science.275.5306.1593.

Schween R, McDougle SD, Hegele M, Taylor JA. Explicit strategies in force field adaptation. bioRxiv. 2019 Jul; doi: 10.1101/694430.

Shadmehr R, Mussa-Ivaldi FA. Adaptive Representation of Dynamics during Learning of a Motor Task. J Neurosci. 1994 May; 14(5 Pt 2):3208–3224.

Shadmehr R, Smith MA, Krakauer JW. Error Correction, Sensory Prediction, and Adaptation in Motor Control. Annu Rev Neurosci. 2010; 33:89–108. doi: 10.1146/annurev-neuro-060909-153135.

Smart WD, Kaelbling LP. Practical Reinforcement Learning in Continuous Spaces. In: *Proceedings of the Seventeenth International Conference on Machine Learning* ICML '00, San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.; 2000. p. 903–910.

Sponberg S, Dyhr JP, Hall RW, Daniel TL. Luminance-Dependent Visual Processing Enables Moth Flight in Low Light. Science. 2015 Jun; 348(6240):1245–1248. doi: 10.1126/science.aaa3042.

Sternad D. It's Not (Only) the Mean That Matters: Variability, Noise and Exploration in Skill Learning. Curr Opin Behav Sci. 2018 Apr; 20:183–195. doi: 10.1016/j.cobeha.2018.01.004.

Stoica P, Moses RL. Spectral analysis of signals. Upper Saddle River, N.J.: Pearson/Prentice Hall; 2005.

968 **Taylor JA**, Ivry RB. Flexible Cognitive Strategies during Motor Learning. PLoS Comput Biol. 2011 Mar; 7(3). doi:
969 10.1371/journal.pcbi.1001096.

970 **Taylor JA**, Klemfuss NM, Ivry RB. An Explicit Strategy Prevails When the Cerebellum Fails to Compute Movement
971 Errors. Cerebellum. 2010 Dec; 9(4):580–586. doi: 10.1007/s12311-010-0201-x.

972 **Taylor JA**, Krakauer JW, Ivry RB. Explicit and implicit contributions to learning in a sensorimotor adaptation task.
973 J Neurosci. 2014 Feb; 34(8):3023–3032. doi: 10.1523/JNEUROSCI.3619-13.2014.

974 **Tcheang L**, Bülthoff HH, Burgess N. Visual influence on path integration in darkness indicates a multimodal
975 representation of large-scale space. Proceedings of the National Academy of Sciences. 2011; 108(3):1152–
976 1157.

977 **Telgen S**, Parvin D, Diedrichsen J. Mirror Reversal and Visual Rotation Are Learned and Consolidated via
978 Separate Mechanisms: Recalibrating or Learning de Novo? J Neurosci. 2014 Oct; 34(41):13768–13779. doi:
979 10.1523/JNEUROSCI.5306-13.2014.

980 **Theodorou E**, Buchli J, Schaal S. Reinforcement Learning of Motor Skills in High Dimensions: A Path Integral
981 Approach. In: *2010 IEEE International Conference on Robotics and Automation*; 2010. p. 2397–2403. doi:
982 10.1109/ROBOT.2010.5509336.

983 **Todorov E**. Efficient Computation of Optimal Actions. Proc Natl Acad Sci USA. 2009 Jul; 106(28):11478–11483.
984 doi: 10.1073/pnas.0710743106.

985 **Tseng YW**, Diedrichsen J, Krakauer JW, Shadmehr R, Bastian AJ. Sensory Prediction Errors Drive Cerebellum-
986 Dependent Adaptation of Reaching. J Neurophysiol. 2007 Jul; 98(1):54–62. doi: 10.1152/jn.00266.2007.

987 **Umeyama S**. Least-squares estimation of transformation parameters between two point patterns. IEEE
988 Transactions on Pattern Analysis and Machine Intelligence. 1991; 13(4):376–380.

989 **van Vugt FT**, Ostry DJ. The Structure and Acquisition of Sensorimotor Maps. J Cogn Neurosci. 2018 Mar;
990 30(3):290–306. doi: 10.1162/jocn_a_01204.

991 **Wang JX**, Kurth-Nelson Z, Kumaran D, Tirumala D, Soyer H, Leibo JZ, Hassabis D, Botvinick M. Prefrontal Cortex
992 as a Meta-Reinforcement Learning System. Nat Neurosci. 2018 Jun; 21(6):860–868. doi: 10.1038/s41593-018-
993 0147-8.

994 **Werner S**, Bock O. Mechanisms for visuomotor adaptation to left–right reversed vision. Hum Mov Sci. 2010 Apr;
995 29(2):172–178. doi: 10.1016/j.humov.2010.02.004.

996 **Wilterson SA**, Taylor JA. Implicit Visuomotor Adaptation Remains Limited after Several Days of Training. bioRxiv.
997 2019 Aug; doi: 10.1101/711598.

998 **Yamagami M**, Howell D, Roth E, Burden S. Contributions of feedforward and feedback control in a manual
999 trajectory-tracking task. IFAC-PapersOnLine. 2019 Jan; 51:61–66. doi: 10.1016/j.ifacol.2019.01.025.

1000 **Zimmet AM**, Bastian AJ, Cowan NJ. Cerebellar Patients Have Intact Feedback Control That Can Be Leveraged to
1001 Improve Reaching. bioRxiv. 2019 Nov; doi: 10.1101/827113.

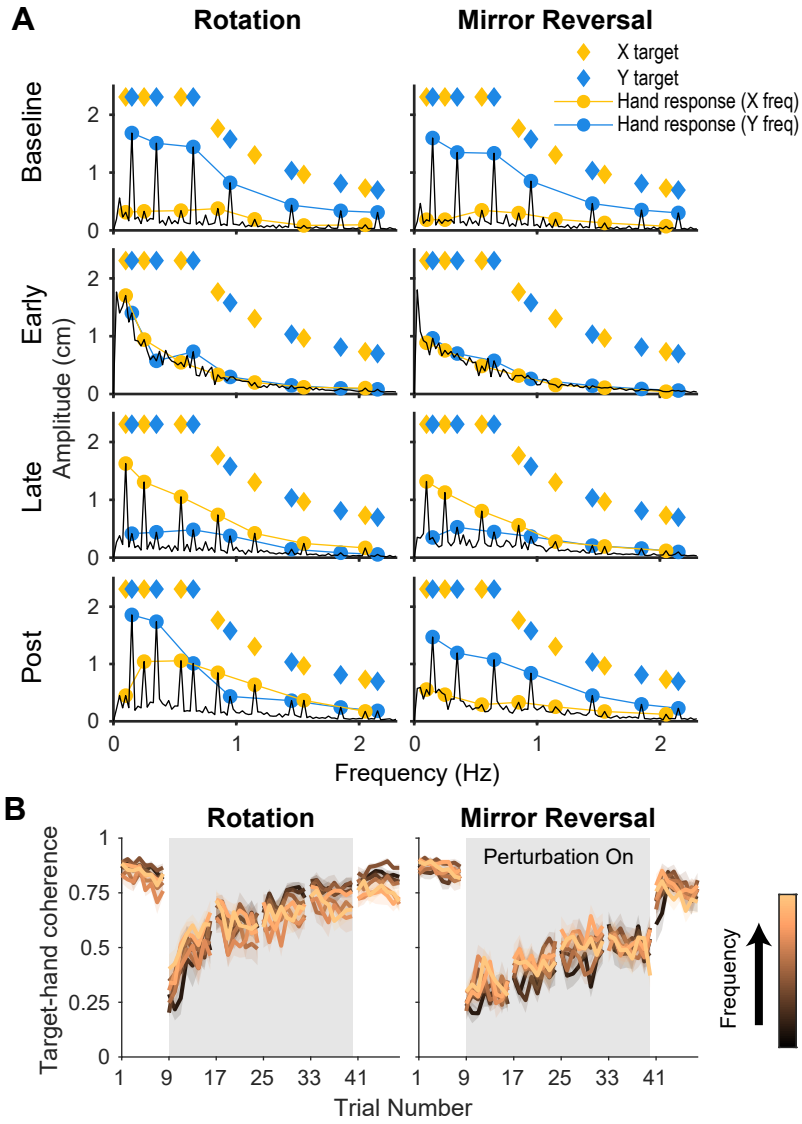

Figure 4–Figure supplement 1. Amplitude spectra and coherence plots for y-hand movements. **A.** Amplitude spectra of y-hand movements (black) averaged across participants from one trial in each listed block. The amplitudes and frequencies of target movement are indicated by diamonds (yellow: x-target; blue: y-target). Hand responses at the x- (yellow circles) and y-target (blue circles) frequencies are connected by lines, respectively, for ease of visualization. **B.** Spectral coherence between y-target movement and both x- and y-hand movement. Darker colors represent lower frequencies and lighter colors represent higher frequencies. Error bars are SEM across participants.

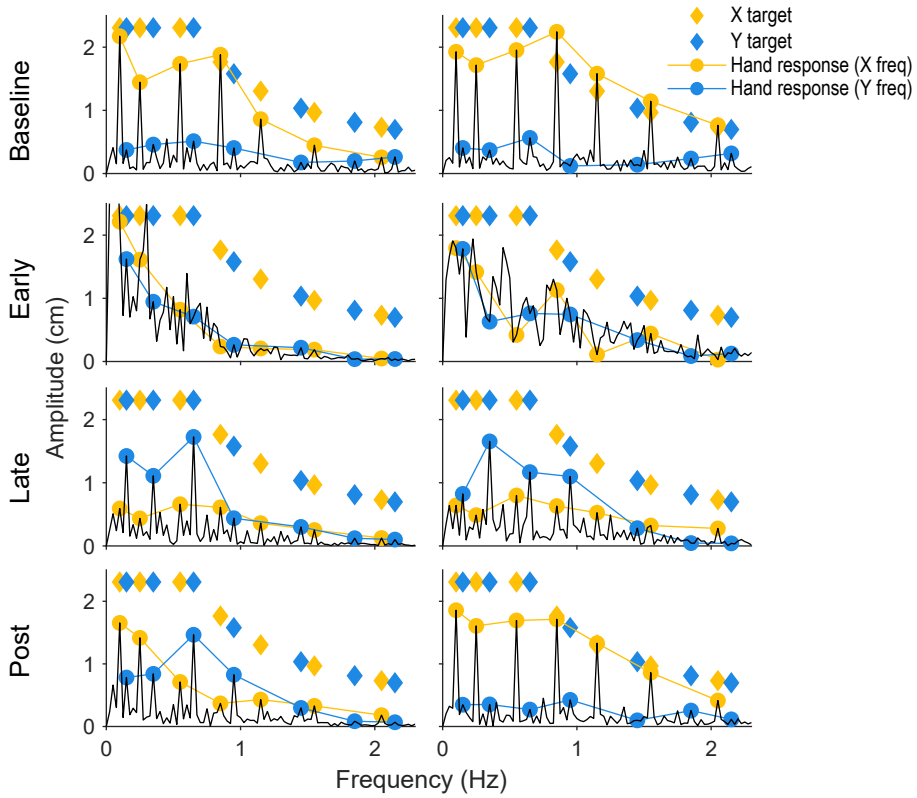

Figure 4–Figure supplement 2. Amplitude spectra of x-hand movements from single subjects. **A.** Amplitude spectra (black) averaged across participants from one trial in each listed block. The amplitudes and frequencies of target movement are indicated by diamonds (yellow: x-target; blue: y-target). Hand responses at the x- (yellow circles) and y-target (blue circles) frequencies are connected by lines, respectively, for ease of visualization.

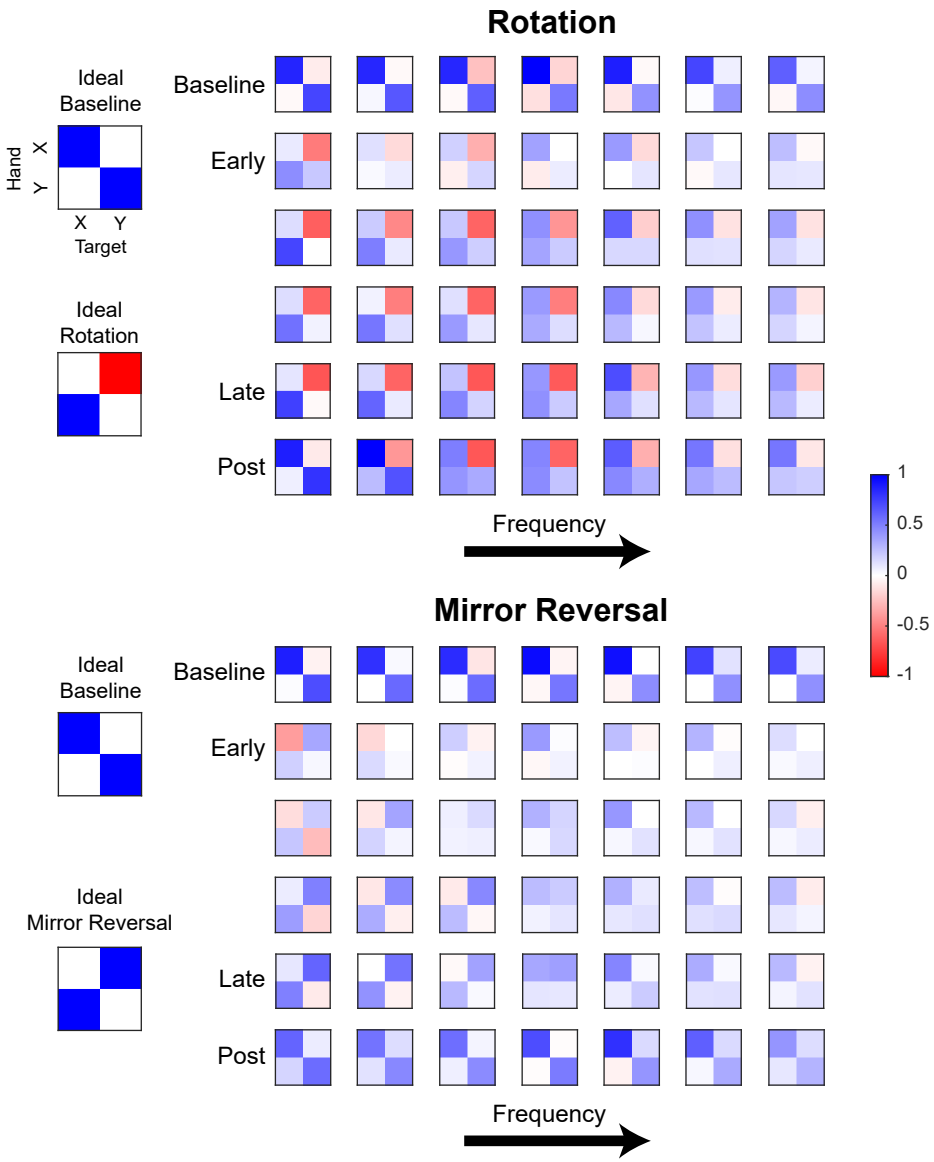

Figure 5–Figure supplement 1. Gain matrices fitted for different frequencies of hand movement. Each element within a 2×2 matrix is the gain estimated between hand and target movement at a particular frequency. Each row of matrices displays the data from one trial of a tracking block (averaged across participants) and each column is a frequency (frequencies increase from left to right). Although only one trial of each block is depicted, other trials within each block were qualitatively similar. ***Figure 5A*** was created by plotting the left column vector of each matrix as a green arrow and the right column vector as a purple arrow.

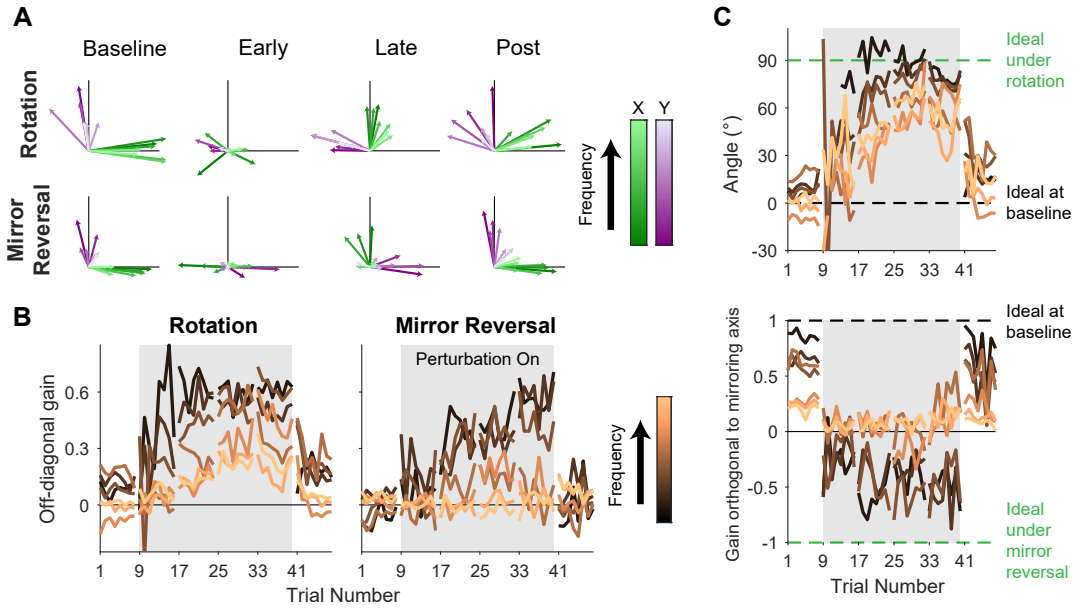

Figure 5–Figure supplement 2. Gain matrix analysis, identical to the one performed in ***Figure 5*** except performed on a single subject from each group. **A.** Visualizations of gain matrices from a single trial in each listed block. **B.** Average of the off-diagonal values of the gain matrix. **C.** Compensation angle for the rotation group and gain orthogonal to the mirror axis for the mirror-reversal group. Note compensation angles could not be estimated for every trial using our singular value decomposition approach, so several data points are missing in the figure (see "Trajectory-alignment analysis" for details on this approach).

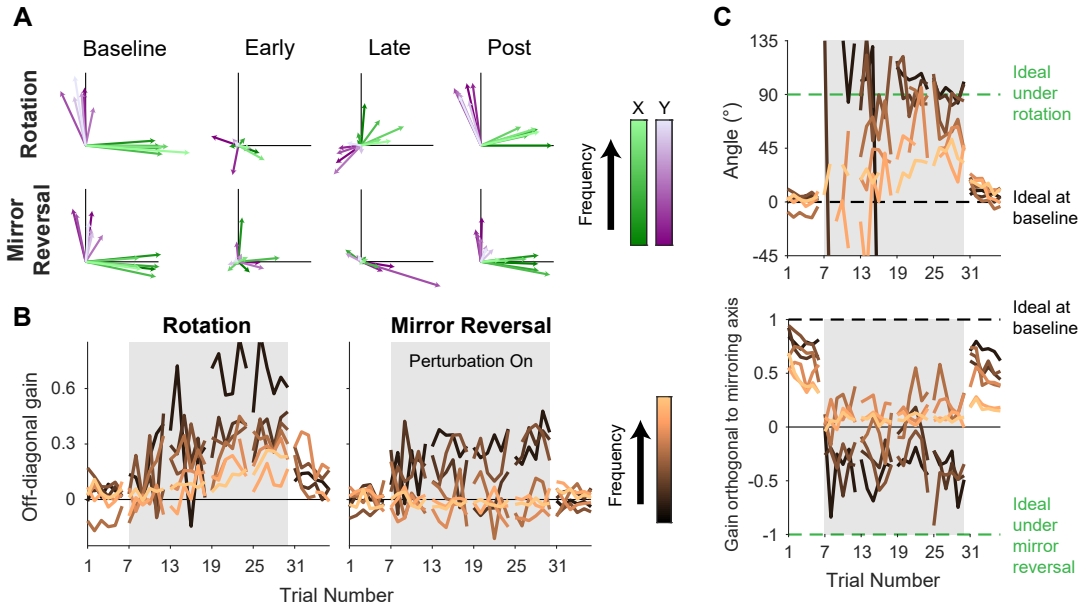

Figure 6–Figure supplement 1. Gain matrix analysis, identical to the one performed in ***Figure 6*** except performed on a single subject from each group. **A.** Visualizations of gain matrices from a single trial in each listed block. **B.** Average of the off-diagonal values of the gain matrix. **C.** Compensation angle for the rotation group and gain orthogonal to the mirror axis for the mirror-reversal group. Note compensation angles could not be estimated for every trial using our singular value decomposition approach (see "Trajectory-alignment analysis" for details on this approach), so several data points are missing in the figure.