J Neurophysiol 128: 1278–1291, 2022.
First published October 12, 2022; doi:10.1152/jn.00273.2022

JNP JOURNAL OF NEUROPHYSIOLOGY

RESEARCH ARTICLE

Control of Movement

Control becomes habitual early on when learning a novel motor skill

Christopher S. Yang,[1,2] Noah J. Cowan,[3,4] and Adrian M. Haith[2]

[1]*Department of Neuroscience, Johns Hopkins University, Baltimore, Maryland;* [2]*Department of Neurology, Johns Hopkins University, Baltimore, Maryland;* [3]*Department of Mechanical Engineering, Johns Hopkins University, Baltimore, Maryland; and* [4]*Laboratory for Computational Sensing and Robotics, Johns Hopkins University, Baltimore, Maryland*

Abstract

When people perform the same task repeatedly, their behavior becomes habitual, or inflexible to changes in the goals or structure of a task. Although habits have been hypothesized to be a key aspect of motor skill acquisition, there has been little empirical work investigating the relationship between skills and habits. To better understand this relationship, we examined whether and when people's behavior would become habitual as they learned a challenging new motor skill: maneuvering an on-screen cursor with a nonintuitive bimanual mapping from hand to cursor position. After participants practiced using this mapping for up to 10 days, we altered the mapping between the hands and the cursor to assess whether participants could flexibly adjust their behavior or would habitually persist in performing the task the way they had originally learned. We found that participants' behavior became habitual within 2 days of practice, at which point they were still relatively unskilled. Further practice led to improved skill but did not alter the strength of habitual behavior. These data demonstrate that motor skills become habitual after relatively little training but can nevertheless further improve with practice. We suggest that building habits early in learning may be a crucial step in acquiring new motor skills.

NEW & NOTEWORTHY Habits and motor skills have often been thought to be deeply related, but very few studies have empirically examined the relationship between the two. We present evidence that habits emerge early in learning, long before a motor skill has been fully learned. Our results suggest that habits may play an integral role in the learning and performance of motor skills from even the early stages of acquiring a new skill.

motor skill; habit

INTRODUCTION

We have all experienced the frustration of having to overcome old habits when we need to alter the way we perform a task. In a recent striking example of this, YouTuber Destin Sandlin created a "backwards bicycle," a bicycle in which rotation of the handlebar in one direction causes the front tire to rotate in the opposite direction (i.e., opposite of a normal bicycle) (1). Although it is easy to understand how the handlebar moves the tire and it is trivial to rotate the handlebar, people find it difficult to ride the backwards bicycle, seemingly because they habitually try to balance themselves using the same movements they would perform on a normal bicycle.

Habits are generally defined as behaviors that, through extensive repetition, have become inflexible to changes in the goals or structure of a task (2–4). Specifically, the example above illustrates a particular type of habit known as a "slip-of-action" habit, a process whereby incorrect actions are obligatorily selected (see Ref. 5 for more details about different types of habits). Slip-of-action habits are particularly relevant to motor skill learning as they tend to occur when one action among many must be selected rapidly. Habits of this kind are thought to be beneficial for the performance of motor skills since they enable actions to be selected and generated more rapidly while also freeing up cognitive resources (6–10). As a result, neuroscientists and psychologists have long speculated about the potential relationship between habits and skills (5, 11–13). Although the repetitive practice required to master a skill might well lead to slip-of-action habits being formed, the relative timing and interdependence of these processes are not clear. One plausible scenario

Correspondence: C. S. Yang (cyang71@jhu.edu).
Submitted 26 June 2022 / Revised 27 September 2022 / Accepted 5 October 2022

is that habits might form early during learning, cementing a possibly rudimentary version of the skill that can be refined through further practice. Alternatively, habits might only occur as an entrenched version of a skill, forming only after a skill has been fully learned and overpracticed.

Relatively little is understood about the relationship between slip-of-action habits and skills largely because there has been a dearth of empirical studies examining this relationship (for brevity, we use the term "habit" to refer specifically to a slip-of-action habit from here onwards). In human behavioral studies, very few experimental protocols have successfully induced habits in motor behaviors (5, 14). Studies that have induced habits have employed simple tasks that require choices between a few discrete actions (14–19), such as deciding which button to press on a keypad, or whether or not to perform an action at all (14, 20). In such cases, habits are conceptualized as stimulus-response associations that have become obligatory through repetition (21–24).

It is unclear whether findings from discrete tasks can be generalized to more complex, real-world tasks that entail continuous state and action spaces, such as riding a bicycle. In the continuous domain, the analog of a stimulus-response association guiding behavior is a *controller*, a mapping from the instantaneous states of the environment and one's body to outgoing motor commands. Although it is conceptually straightforward to extend the concept of a habit from discrete tasks to continuous movement control, it is unknown whether habits form in the same way in both cases. A key tenet of the stimulus-response framework is that a particular stimulus and resulting response must be paired repeatedly for a habit to form, but in continuous control tasks there is a continuum of (i.e., infinitely many) possible states and actions and it is unclear whether one will ever repeat the same action in the same state often enough for a habit to form. To a limited extent, behavior that could be interpreted as habits has been studied in continuous control tasks such as reaching under mirror-reversed visual feedback (25, 26) or more real-world skills like javelin throwing (27), swimming (28), and weightlifting (29). However, such work has examined the process of replacing already highly skilled movements (baseline or well-practiced behavior) with new movements rather than how habits form when people initially learn a skill.

To better understand the relationship between habits and skills, we performed an experiment to investigate how quickly a learned motor skill would become habitual. Participants learned to maneuver an on-screen cursor toward a visual target using a nonintuitive mapping from the position of both hands to the location of a cursor: forward-backward movements of the left hand were mapped to left-right movements of the cursor and left-right movements of the right hand were mapped to forward-backward movements of the cursor. Previous work suggests that people learn this mapping by building a new controller de novo (30), in contrast to how people learn simpler perturbations like rotations of visual feedback by adapting an existing controller (31–34). Three separate groups of participants learned to control the bimanual mapping over 2, 5, or 10 days of practice. At the end of the final day of practice, we flipped the direction of the mapping between movement of the left hand and movement of the cursor (i.e., a mirror reversal) and assessed whether participants would habitually continue to control the cursor according to the originally practiced mapping or would be able to flexibly adjust their control to accommodate the new flipped mapping. If participants exhibited habitual behavior, we assessed whether their habit expression depended on how long they had practiced using the bimanual mapping to control the cursor, as well as how long this habit persisted while participants continued to practice the flipped mapping.

MATERIALS AND METHODS

Participants

A total of 32 right-handed participants were recruited for this study [age = 23.0 ± 4.3 yr (mean ± standard deviation); 13 male, 19 female], 13 for the 2-day group, 14 for the 5-day group, and 5 for the 10-day group (recruitment for the 10-day group was cut short because of the COVID-19 pandemic). All participants reported no prior history of neurological disorders. All methods were approved by the Johns Hopkins School of Medicine Institutional Review Board and were carried out in accordance with relevant guidelines and regulations. Written informed consent was obtained from all participants in the study.

Experimental Setup

Participants were seated in front of a table with both of their hands supported on the table by frictionless air sleds. The positions of participants' hands were monitored at 130 Hz with a Flock of Birds magnetic tracker (Ascension Technology, Shelburne, VT) placed near each hand's index finger. Participants viewed stimuli on a horizontal mirror that reflected an LCD monitor (60 Hz), and the mirror obscured vision of both hands.

Tasks

Participants learned to maneuver an on-screen cursor (circle of radius 2.5 mm) using one of two versions of a bimanual hand-to-cursor mapping. Half of the participants learned one version in which forward-backward movements of the left hand produced right-left movements of the cursor while right-left movements of the right hand produced forward-backward movements of the cursor (Fig. 1). The other half of participants learned a different version in which the mapping from hand to cursor movements was rotated by 180° relative to the other version (i.e., forward-backward movements of the left hand produced left-right movements of the cursor while left-right movements of the right hand produced forward-backward movements of the cursor). We used these two versions of the mapping to control for any possible asymmetries in behavior (e.g., due to biomechanics) that might have biased our results.

Three different groups practiced the bimanual mapping over 2, 5, or 10 days of training (Fig. 1). Training consisted of participants making point-to-point reaches toward randomly placed targets (circles with radius of 10 mm) within a 20×20-cm workspace. Participants were instructed to reach toward each target as quickly and accurately as possible,

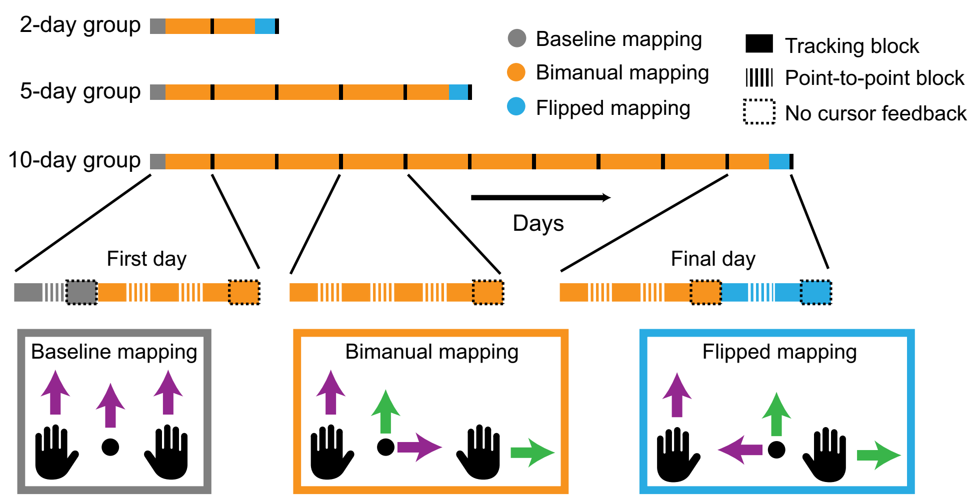

Figure 1. Tasks and experiment. Participants learned to control an on-screen cursor using a bimanual hand-to-cursor mapping (orange) over 2 (n = 13; n represents number of subjects), 5 (n = 14), or 10 (n = 5) days of practice. Half of the participants in each group practiced the depicted bimanual mapping, whereas the other half practiced an alternate version in which cursor movements were rotated 180° relative to the depicted mapping (this was done to control for potential asymmetries in behavior due to, for example, biomechanics). On each day, participants performed blocks of point-to-point reaching (hashed rectangle; 1 block = 100 trials) and continuous tracking (1 block = 5 min) both with (solid rectangle) and without (solid rectangle with dashed outline) visual feedback of the cursor. Learning was compared relative to a baseline mapping where the cursor was placed at the average position of the 2 hands (gray). At the end of each group's final training day, we flipped the left hand's mapping to cursor movements (blue) and assessed whether participants would habitually continue to control the cursor according to the bimanual mapping they originally learned.

with each trial consisting of one target location. Once the cursor was stationary (speed < 0.065 m/s) within the target for 1 s, the target appeared in a random direction 12 cm away. To encourage participants to move quickly, we provided feedback to the participants indicating whether their peak velocity exceeded 0.39 m/s on that trial. If this threshold was exceeded, the target turned yellow and a pleasant tone was played once the cursor reached the target, and if the threshold was not exceeded, the target did not turn yellow and no tone was played. The baseline point-to-point block consisted of 30 trials, and all other point-to-point blocks consisted of 100 trials.

In between blocks of point-to-point reaching, participants performed a continuous tracking task. In this task, a target moved continuously on the screen in a sum-of-sinusoids trajectory. The trajectory was composed of 12 sinusoids, 6 each in the x- and y-axes, parameterized by amplitude ($\vec{a}$), frequency ($\vec{\omega}$), and phase ($\vec{\phi}$) vectors. The target's position along a single axis, r, was computed as

$$r = \sum_{i=1}^{6} a_i \cos(2\pi t \omega_i + \phi_i). \tag{1}$$

For the x-axis, $\vec{a}$ = [2.31, 2.31, 2.31, 1.76, 1.30, 0.97] (cm) and $\vec{\omega}$ = [0.1, 0.25, 0.55, 0.85, 1.15, 1.55] (Hz). For the y-axis, $\vec{a}$ = [2.31, 2.31, 2.31, 1.58, 1.03, 0.81] (cm) and $\vec{\omega}$ = [0.15, 0.35, 0.65, 0.95, 1.45, 1.85] (Hz). The values of $\vec{\phi}$ were randomized between $[-\pi, \pi)$. Because movement frequencies were chosen to be prime multiples of 0.05 Hz, the target repeated its trajectory every 20 s. However, the trajectory was pseudorandom (unpredictable) because many sinusoids were combined to generate the trajectory, thereby making it challenging for participants to detect any rhythmicity in the task. Each tracking trial lasted 66 s, and during the first 5 s of each trial the target's amplitude ramped up linearly from 0 to its full value. Each block consisted of five trials. Periodically, participants performed a tracking block without visual feedback of their cursor.

To assess the extent to which participants' control of the bimanual mapping had become habitual, at the end of each group's final day of training we flipped the left hand's mapping to cursor movement (flip block): forward-backward movements of the left hand now resulted in left-right movements of the cursor instead of right-left (in the case of the 180°-rotated bimanual mapping, right-left movements of the cursor became left-right). We recruited three different groups of participants for this experiment rather than repeatedly assessing habit within the same individuals at different points during learning because we expected that, after a participant experienced the flipped mapping once, this might influence any future learning of the original bimanual mapping as well as future habit assessments using the flipped mapping. The order of all blocks during the experiment is depicted in Fig. 1.

Data Analysis

Software.

Data analyses were performed in MATLAB R2018b (The MathWorks, Natick, MA) and R version 4.0.2 (RStudio, Inc., Boston, MA; Ref. 35) using the Matrix (36), lme4 (37), lmerTest (38), and emmeans packages (39). Figures were generated with Adobe Illustrator (Adobe Inc., San Jose, CA).

Analysis of point-to-point data.

The cursor's position in each trial was smoothed with a third-order Savitzky–Golay filter with a window of seven samples (~50 ms). Path length was defined as the total distance that the cursor traveled in a single trial. Movement time was defined as the time between movement initiation (when the cursor left the start target) and termination (when the cursor was in the end target with speed < 0.065 m/s).

Reaction time was defined as the time between when the target appeared and the cursor left the starting target. Peak velocity was defined as the cursor's highest tangential velocity. We computed the tangential velocity by linearly resampling the cursor's position at the times recorded by the Flock of Birds and computing the distance traversed by the cursor between two consecutive samples divided by the time elapsed. Resampling was necessary because, occasionally, the recorded time at which a sample was collected by the Flock of Birds did not match the true time at which it was collected, causing the calculated velocity to be inaccurate. Velocity profiles were also smoothed with a third-order Savitzky–Golay filter.

Initial reach direction was defined as the direction of the instantaneous velocity vector 150 ms after movement initiation. Initial reach direction error was computed as the difference in angle between this instantaneous velocity vector and the vector pointing from the target on the previous trial to the target on the current trial. Probability density functions were estimated for reach direction errors with a kernel-smoothing function, implemented as the `ksdensity` function in MATLAB.

Variance estimation model. We measured the variability in participants' initial reach direction errors (i.e., how consistently straight participants reached toward the target) by fitting a mixture model to this data. In the model, we assumed that participants' reach direction errors, x, were generated by one of two causes: *1*) an error from a goal-directed reach toward the target (modeled as a von Mises distribution) or *2*) an error from a reach in a random direction (modeled as a uniform distribution). The probability density function of the mixture model, mix(·), was defined as

$$\mathrm{mix}(x|\mu,\kappa,\alpha) = \alpha \cdot \mathrm{vm}(x|\mu,\kappa) + (1-\alpha)\cdot \mathrm{unif}(x) \tag{2}$$

where α is a parameter valued between 0 and 1 weighting the probability density functions of the von Mises [vm(·)] and uniform [unif(·)] distributions. The probability density functions of the individual distributions were defined as

$$\mathrm{vm}(x|\mu,\kappa) = \frac{e^{\kappa\cos(x-\mu)}}{2\pi I_0(\kappa)}, \qquad \mathrm{unif}(x) = \frac{1}{2\pi}. \tag{3}$$

Here, μ and κ are the mean and concentration of the von Mises distribution and $I_n(\cdot)$ is the modified Bessel function of the first kind with order n, which in this case was 0.

The parameters μ, κ, and α were fit to the data from single participants in each block via maximum likelihood estimation. Specifically, we used the MATLAB function `fmincon` to determine the values of the parameters that would maximize the likelihood function below over the n trials within one block:

$$\hat{\mu},\hat{\kappa},\hat{\alpha} = \underset{\mu,\kappa,\alpha}{\mathrm{argmax}}\left\{\sum_{i=1}^{n}\log\left[\mathrm{mix}(x_i|\mu,\kappa,\alpha)\right]\right\}. \tag{4}$$

Then, using the fitted concentration parameter of the von Mises distribution, $\hat{\kappa}$, we computed the circular standard deviation, σ, as

$$\sigma = \sqrt{-2\ln(R)}, \qquad R = I_1(\hat{\kappa})/I_0(\hat{\kappa}). \tag{5}$$

We used σ as our measure of the variability of participants' reach direction errors.

Habit strength model. To assess whether participants exhibited habitual behavior during the flip block, we quantified each participant's tendency to reach toward the true target versus a virtual target flipped across the mirroring axis. More specifically, we assumed that for each trial participants' initial reach direction could be explained by at least one of three causes: *1*) a goal-directed reach toward the target, *2*) a habitual reach toward the mirrored target, and *3*) a reach aimed toward neither target (i.e., random movement). We modeled the first two causes as von Mises distributions with different means—ϕ_a and ϕ_m, set by the direction of the actual and mirrored targets, respectively—but the same concentration parameter, κ. We modeled the third cause, random movements, as a uniform distribution.

Assuming that each participant's behavior within one block could be modeled as a weighted mixture of these three distributions, mix′(·), we used the MATLAB function `fmincon` to determine the weights, α_a and α_m, that would maximize the following likelihood function over the n trials within one block:

$$\hat{\alpha}_a,\hat{\alpha}_m,\hat{\kappa} = \underset{\alpha_a,\alpha_m}{\mathrm{argmax}}\left\{\sum_{i=1}^{n}\log\left[\mathrm{mix}'(x|\phi_a,\phi_m,\kappa,\alpha_a,\alpha_m)\right]\right\} \tag{6}$$

where

$$\mathrm{mix}'(x|\phi_a,\phi_m,\kappa,\alpha_a,\alpha_m) = \alpha_a\cdot\mathrm{vm}(x|\phi_a,\kappa) + \alpha_m\cdot\mathrm{vm}(x|\phi_m,\kappa) + (1-\alpha_a-\alpha_m)\cdot\mathrm{unif}(x). \tag{7}$$

Here, x represents participants' reach directions and α_a and α_m correspond to the probabilities that a participant reached toward the actual and mirrored targets, respectively. Definitions for vm(·) and unif(·) can be found in *Eq. 3*. We used $\hat{\alpha}_m$ as our metric for the strength of habitual behavior. For Fig. 4*E*, instead of fitting this model to all trials in the flip block, we fit the model to either the first or second half of trials in this block.

We used the fitted weights from this approach to classify each trial as goal directed, habitual, or random. For each trial, we computed the probability that the reach direction was generated from each of the three mixture components under the fitted mixture model's probability density function [in essence, computing *P*(reach direction|goal directed), *P*(reach direction|habitual), and *P*(reach direction|random)]. Trials were classified as goal directed, habitual, or random based on which of these three probabilities was the highest. We used this classification to compute the reaction times of goal-directed versus habitual reaches in Fig. 4*D*.

As an alternative approach to assess whether participants exhibited habitual behavior without fitting a model to the data, we assessed whether their horizontal cursor movements were aimed away from the target. Using only the cursor's x-axis position, for each trial we determined whether the horizontal component of the cursor's instantaneous velocity vector was aimed toward or away from the target 150 ms after the horizontal position of the cursor deviated from the center of the starting target by 1 cm (i.e., the radius of the target). We classified cursor movements in each trial as moving away from the target if the velocity vector's direction was opposite of the direction of the target relative to the starting position (e.g., target located to the left but cursor moving toward the right). This method was unable to compute an initial horizontal reach direction on a small minority of trials (95 out of 3,200 trials) where the target on the current trial was either directly above or below the target from the previous trial. This was because either *1*) the cursor did not deviate 1 cm horizontally away from the center of the starting target (i.e., the radius of the target), making it impossible to detect the time of movement initiation, or *2*) the detected movement initiation time was <150 ms before the end of the trial, meaning that the trial ended before the time at which we assessed reach direction. These trials were excluded from the analysis.

Random reach model. We compared the model in *Eq. 7* with an alternative model that was designed to capture behavior where participants would move their right hand (which controlled vertical cursor movement) in the correct direction but would move their left hand (which controlled horizontal cursor movement) in a random direction. We modeled participants' reach directions, x, given the target's direction, ϕ_a, as a mixture of two weighted uniform distributions:

$$\mathrm{mix}^*(x|\phi_a,\alpha) = \begin{cases} \alpha \cdot \mathrm{unif}^*(x), & (\sin(x) \geq 0 \,\&\, \sin(\phi_a) \geq 0) \,|\, (\sin(x) < 0 \,\&\, \sin(\phi_a) < 0) \\ (1-\alpha) \cdot \mathrm{unif}^*(x), & (\sin(x) < 0 \,\&\, \sin(\phi_a) \geq 0) \,|\, (\sin(x) \geq 0 \,\&\, \sin(\phi_a) < 0) \end{cases} \tag{8}$$

where

$$\mathrm{unif}^*(x) = \frac{1}{\pi}. \tag{9}$$

Here, α is the probability that the cursor moved vertically in the correct direction. We used the MATLAB function `fmincon` to determine the α that maximized the following likelihood function:

$$\hat{\alpha} = \underset{\alpha}{\mathrm{argmax}} \left\{ \sum_{i=1}^{n} \log\left[\mathrm{mix}^*(x|\phi_a,\alpha)\right] \right\}. \tag{10}$$

The fits for the habit strength and random reach models were compared using the Bayesian information criterion (BIC). Model recovery analyses were performed by simulating data from both of these models, fitting both models to each simulated data set, comparing fits with BIC, and generating a confusion matrix. To generate data from *Eq. 7*, we used values for α_a and α_m that ranged between 0 and 1, and we fixed κ = 3. Lower κ's set higher variability for the von Mises distributions (i.e., harder to distinguish from the model in *Eq. 8*), so we fixed κ to be the lowest average κ that we observed in the late learning data from any group, as estimated in *Eq. 4*. To generate one confusion matrix, we simulated data from both the habit strength and random reach models 50 times where each simulation consisted of 100 simulated reach directions, matching the number of trials in the experimental data. Model recovery across different choices of parameters were compared by computing the accuracy of the confusion matrices.

Analysis of tracking data.

Data from two tracking trials (each from different subjects) were excluded from the analysis because of hardware failure. Unlike the point-to-point data, the tracking data were not smoothed with a filter as the noise introduced by our tracking hardware did not significantly impact our analyses. Tracking error was computed as the mean-squared error between the cursor's and target's positions. Time-domain trajectories of the cursor and target position (the first 60 s of each trial following the initial 5-s ramp period) were converted to phasors (complex numbers representing sinusoids) in the frequency domain via the discrete Fourier transform. An input-output transfer function was computed at every frequency by dividing the cursor's phasor by the target's phasor. This transfer function described the relationship between the cursor and target sinusoids in terms of gain (relative amplitude) and phase (difference in time).

Using these transfer functions, we sought to describe the direction that participants moved their cursor to track the target. In this task, participants' cursor movements would conventionally be described as phase lagged relative to the target with a positive gain (i.e., moving in the same direction as the target with a time delay). However, when a mirror reversal has been applied (such as in the flipped mapping), participants may habitually continue to use their original control policy, causing their movements to be flipped across the mirroring axis relative to before. Although the relationship between movements before and after the flip could be described as movements with positive gain but now in antiphase (i.e., moving in the same direction as the target but with more time delay), a better way to describe them would be to say that the movements have the same phase but a negative gain (i.e., moving with the same time delay but in the opposite direction of the target).

Given that conventional analysis methods always yield a positive gain to describe frequency-domain data, we used the method described in Ref. 40 to compute a signed gain, g, relating cursor and target movements. This was computed as the dot product between transfer functions:

$$g = a \cdot \hat{b} \tag{11}$$

where a is the transfer function for a given block of interest and $\hat{b}$ is a unit vector with the same phase as the transfer function at baseline. Computing the dot product implicitly fixes the phase of cursor movements to be the same as baseline across all blocks, allowing a signed gain to be computed. This assumption of fixed phase is valid for analyzing data in late learning as participants' phase lags under the bimanual mapping became more similar to baseline through practice. Different phases were fixed for each participant based on their individual baseline behavior.

We computed this signed gain between each axis of target and cursor movement, building a series of 2×2 matrices, $G(\omega)$, relating the transformation between the two trajectories where each matrix represented the transformation within a small bandwidth of frequencies, ω:

$$G(\omega) = \begin{bmatrix} g_{xx}(\omega) & g_{xy}(\omega) \\ g_{yx}(\omega) & g_{yy}(\omega) \end{bmatrix}. \quad (12)$$

Here, the first subscript of g indicates the axis of hand movement and the second subscript indicates the axis of target movement. These matrices were represented geometrically by plotting their column vectors in Fig. 3*C*, where the green (purple) arrows represent the first (second) column of the matrices. Figure 3*D* and Supplemental Fig. S4B (see ENDNOTE) were generated by plotting the element in the first row and first column of each matrix.

To quantify the strength of habitual behavior in Fig. 5*A*, we reanalyzed the gain between *x*-axis target and *x*-axis cursor movements from the flip blocks by fixing their phases to be the same as late learning. We did this because we expected any habitual behavior to maintain the same phase as at late learning, which might have differed slightly from the phase of behavior under the baseline mapping and thus provided a more accurate estimate of habitual behavior. We also calculated the normalized gain by dividing the amplitude of the tracking response by that from late learning. When analyzing these normalized gains at the group level, we excluded one outlier participant in the 10-day group who exhibited dramatically more negative gains than any other participants within any group (*bottom left* participant in Supplemental Fig. S3A). To compare the habitual behavior we observed between the point-to-point and tracking tasks, we correlated each participant's α_m from *Eq. 6* with their normalized gains (averaged over the highest 3 frequencies) from Fig. 5*B* via linear regression.

Statistics.

Because the 10-day group had only five subjects and was therefore underpowered, we only performed statistical analyses for the 2-day and 5-day groups. All results from the 10-day group are reported qualitatively. Most primary statistical analyses were performed by fitting linear mixed effects models to the data. For all analyses in Fig. 2, the models used group (2-day or 5-day) and block (2-day: *day 1* vs. *day 2*; 5-day: *day 2* vs. *day 5*) as fixed effects and subject as a random effect. For Fig. 4, *C* and *E*, models used the same group and subject effects but with a different set of blocks being compared [(late learning vs. flip block) and (first half of flip block

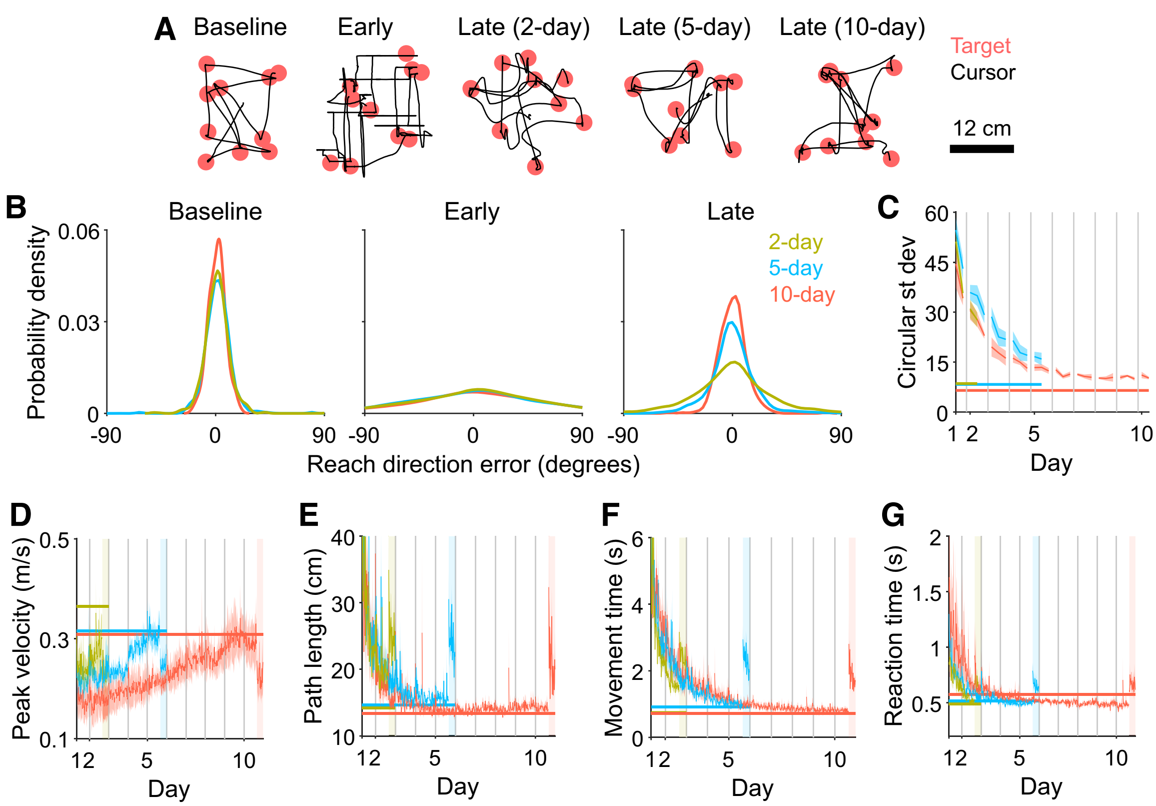

Figure 2. Performance in the point-to-point task under the bimanual mapping. *A*: examples of raw cursor trajectories (black line) from baseline, early learning, and late learning (last block before flip block). Targets are displayed as red circles. Data from 10 trials are shown for each block. *B*: kernel-smoothed probability density of reach-direction errors pooled over all subjects and trials for a given block. All blocks are the same as those shown in *A*. *C*: circular standard deviation of reach direction errors, computed by fitting a mixture model to the data in *B* (see *Analysis of point-to-point data* for more details). Each point corresponds to data from a single block, and error bars indicate SE across participants. Standard deviations under the baseline mapping for each group are shown as horizontal lines. Days are demarcated by gray vertical lines. Data from the flip block are not shown because a separate model was fit to this data, shown in Fig. 4. *D–G*: peak velocity (*D*), path length (*E*), movement time (*F*), and reaction time (*G*) of point-to-point movements throughout learning. Data were averaged across bins of 5 trials, and error bars indicate SE across participants. Values for each group under the baseline mapping are shown as horizontal lines. Different days are demarcated by gray vertical lines. Shaded areas indicate data from the flip block for each group.

vs. second half of flip block), respectively]. For Fig. 4*D*, models used the same group and subject effects but with reach type as an additional fixed effect (goal directed vs. habitual). Post hoc pairwise comparisons were performed with the Tukey test.

For data from the tracking task, mixed effects models were fit using the same effects as Fig. 2 but with an additional fixed effect of frequency. We also fit separate models to data from each frequency because behavior varied dramatically as a function of frequency. Post hoc pairwise comparisons were performed with the Tukey test. An additional Bonferroni correction factor of 6 was applied to the *P* values for pairwise comparisons to account for the separate models fits for each frequency. Additionally, to determine whether participants exhibited significantly negative gains in the tracking task, for each frequency we performed a series of one-sample *t* tests and corrected for multiple (6) comparisons with a Holm–Bonferroni correction at $\alpha = 0.05$.

RESULTS

Participants' Control of the Bimanual Mapping Improved with Practice in the Point-to-Point Task

Figure 2*A* shows representative raw cursor trajectories at baseline, early learning, and late learning for each group in the point-to-point task. As previously found (30), participants initially experienced great difficulty in coordinating their two hands together to move the cursor straight toward each target. But they gradually improved their performance with practice, eventually moving between targets in a straight line, similar to their performance when using an easy mapping in which the cursor appeared exactly halfway between the left and right hand ("baseline"; Fig. 1).

As our primary measure of task performance under the bimanual mapping, we quantified how precisely they aimed the cursor's initial movement toward the target (Fig. 2, *B* and *C*). Precision improved with practice, improving significantly between 2 and 5 days of practice [linear mixed effects model with post hoc Tukey test (see MATERIALS AND METHODS for details about statistical analyses): $t = -7.35$, $P < 0.0001$]. Although there were improvements in performance from *day 5* to *day 10* in the 10-day group, these improvements were relatively small. Other metrics of performance including peak velocity, path length, movement time, and reaction time also improved over multiple days of practice (Fig. 2, *D*–*G*). These data collectively suggest that participants became more skilled in performing point-to-point reaches under the bimanual mapping, with the bulk of this improvement occurring over the first 5 days of practice.

Participants' Control of the Bimanual Mapping Improved with Practice in the Tracking Task

Participants also performed a second task under the bimanual mapping in which they tracked a target moving in a pseudorandom trajectory (sum of sinusoids ranging in frequency from 0.1 to 1.85 Hz; Fig. 3*A*). Unlike the point-to-point task, where participants had an unlimited amount of time to plan their movements at the start of each trial, in the tracking task the target moved quickly and unpredictably, limiting the amount of time participants had to plan their movements; any movements planned at one moment would become outdated within tens of milliseconds as the target would move to a new, unpredictable location. We had participants perform this task because previous studies of habit formation have suggested that the expression of habits may be masked by deliberative, goal-directed processes that

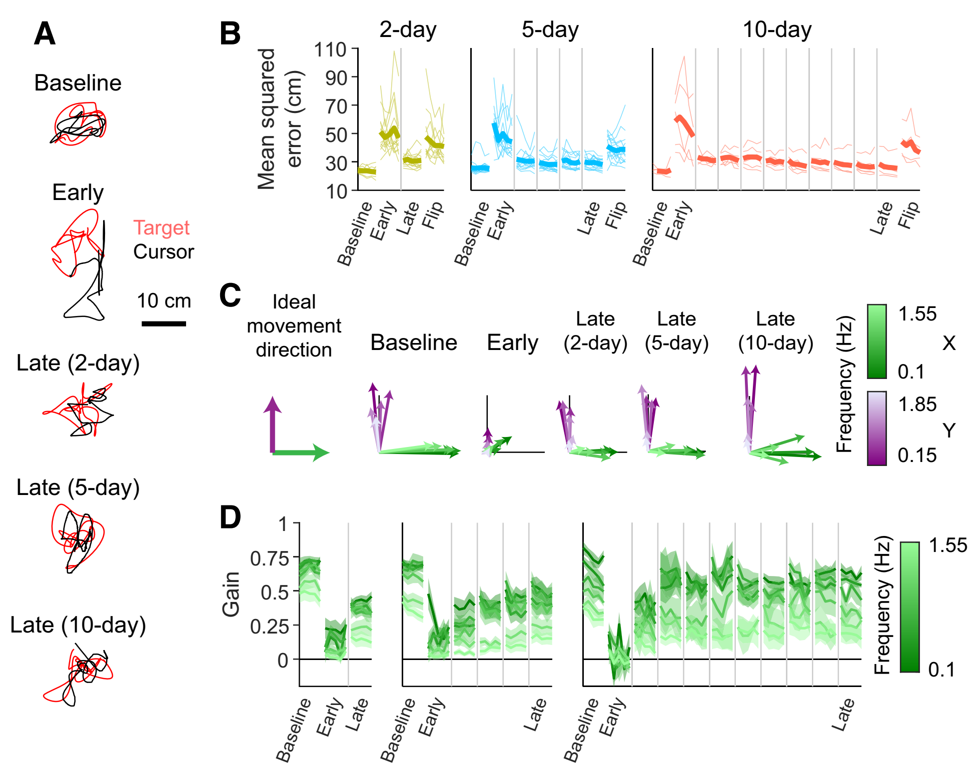

Figure 3. Performance in the tracking task under the bimanual mapping. *A*: example cursor (black) and target (red) trajectories from single trials. *B*: mean-squared error between cursor and target positions. Thin lines indicate individual participants, and thick lines indicate group averages. Data collected from different days are separated by gray lines. Data from only 1 or 2 blocks are shown for each day for ease of visualization. *C*: visualization of the cursor's movement direction and gain relative to the target under the bimanual mapping. Each arrow depicts the direction and gain averaged across participants at a single frequency of *x* (green)- or *y* (purple)-axis target movement. Lower and higher frequencies are depicted as darker and lighter colors, respectively. Black lines are scale bars indicating a movement gain of 0.5. Visualization of ideal movement direction is shown on *left*. Note that this visualization does not depict the ideal gain because the ideal gain changes as a function of phase lag (see Ref. 42). *D*: gain of horizontal cursor movements at frequencies of *x*-axis target movement (horizontal component of green arrows in *C*). Error bars indicate SE across participants. Data from flip block not shown because these gains were quantified differently, as shown in Fig. 5*B*.

might override habitual responses during the reaction time before movement, particularly if participants are allowed ample time to prepare their movements (41). In the tracking task, the amount of time participants had to prepare their movements depended on the frequency at which the target moved. At low frequencies the target oscillated slowly, providing people ample time to prepare their movements, but at high frequencies the target oscillated quickly, forcing people to respond quickly.

With practice, participants learned to reduce the positional error between the target and cursor (Fig. 3*B*). We examined participants' tracking performance at different frequencies of movement with a system identification approach (40, 42–45), which allowed us to separately examine the behavioral responses to target movements at different frequencies even though they occurred concurrently in the task. Specifically, we computed the gain and direction of cursor movement relative to target movement at each frequency, i.e., similar to reach direction in the point-to-point task (see MATERIALS AND METHODS for more details). Each arrow in Fig. 3*C* shows the gain and direction of cursor movements in response to target movement at a particular frequency. Ideally, participants would track the target by moving their cursor in the same direction as the target. Thus, cursor responses to positive *x*-axis target movement (green) should be pointed rightward whereas responses to positive *y*-axis target movement (purple) should be pointed upward, which indeed was the case at baseline. By late learning, all groups exhibited movement gains and directions that approached those of baseline performance.

To statistically compare each group's performance, we computed the gain of horizontal cursor movements at the frequencies of *x*-axis target movement (*x*-component of green arrows in Fig. 3*C*). Gains improved from *day 1* to *day 2* in the 2-day group (Fig. 3*D*; linear mixed effects model with post hoc Tukey test; $P < 0.05$ for 5 of 6 frequencies) and from *day 2* to *day 5* in the 5-day group ($P < 0.05$ for 3 of 6 frequencies). However, gains did not appear to improve past *day 5* in the 10-day group. These data demonstrate that, with practice, participants became able to successfully move their hands in the appropriate direction to track the target and, as in the point-to-point task, the bulk of this improvement occurred in the first 5 days.

Participants' Behavior in the Point-to-Point Task Was Habitual after Only Two Days of Practice

Having examined participants' improvement in task performance through practice, we next asked whether and when their behavior became habitual. Might participants' behavior become habitual around the same time that their task performance plateaued (i.e., by *day 5*), early in learning (i.e., by *day 2*), or only after their task performance had plateaued (i.e., by *day 10*)? Or, finally, might participants' behavior never have become habitual? To determine this, at the end of each group's final day of practice we had participants control the cursor under a new flipped mapping in which the mapping between the left hand and the cursor movement was reversed relative to what they had originally practiced ("flip" block), effectively amounting to a left-right mirror reversal applied on top of the originally practiced bimanual mapping. Participants were explicitly informed about the reversal of their left hand's mapping, and we tested whether participants would habitually continue to control the cursor according to the originally learned bimanual mapping or successfully alter their behavior according to the flipped mapping.

First, we assessed whether participants exhibited habitual behavior in the point-to-point task. On a given trial, if participants could successfully control the cursor under the flipped mapping, then we would expect their cursor's initial movement to be aimed toward the true target (i.e., goal directed). But if participants habitually controlled the cursor according to the original bimanual mapping, then we would expect their cursor's initial movement to be aimed toward a virtual target reflected directly across a vertical mirroring axis. We found that participants in all three groups exhibited both goal-directed and habitual behavior during the flip block (Fig. 4*A*) on different trials. We visualized how often participants' movements were aimed toward the virtual mirrored target as a heat map plotting the cursor's initial movement directions as a function of the target's direction (Fig. 4*B*). If participants reached toward the virtual mirrored target, their initial cursor directions would lie along the $y = -x$ line. Although none of the groups exhibited initial cursor directions along this line at late learning, all groups did exhibit such behavior during the flip block.

We estimated the proportion of trials in which participants initially reached toward the mirrored target by fitting a mixture model to the reach direction data (see *Analysis of point-to-point data* for more details; see Supplemental Fig. S1 for a model comparison and model recovery analysis), using this as a metric for how strongly participants exhibited habitual behavior. We found that the proportion of habitual movements was significantly higher in the flip block compared with late learning for the 2-day and 5-day groups (Fig. 4*C*; linear mixed effects model with post hoc Tukey test; 2-day: $t = 5.78$, $P < 0.0001$; 5-day: $t = 9.03$, $P < 0.0001$), and the 10-day group exhibited a similar trend. These data demonstrate that all groups exhibited habitual behavior in the point-to-point task.

Perhaps surprisingly, the proportion of reaches toward the mirrored target was not significantly different between the 2-day and 5-day groups (Fig. 4*C*; linear mixed effects model with post hoc Tukey test; 2-day vs. 5-day: $t = -2.56$, $P = 0.0635$), and the 10-day group did not appear to exhibit more habitual reaches than either of these groups. In other words, despite the fact that the 5-day group practiced using the original bimanual mapping for more than twice as long as the 2-day group, they did not exhibit more strongly habitual behavior in the point-to-point task. Moreover, the reaction times for goal-directed reaches were not significantly different from habitual reaches {Fig. 4*D*; linear mixed effects model; no main effect of group [$F(1,25) = 0.05$, $P = 0.8261$] or reach [$F(1,25) = 0.16$, $P = 0.6954$]}, suggesting that the lack of differences across groups in Fig. 4*C* could not be explained by differences in the amount of time participants had to plan their movements.

In the flip block, we noted that participants occasionally adopted a strategy of initially moving the cursor vertically (the axis along which the mapping had not changed)

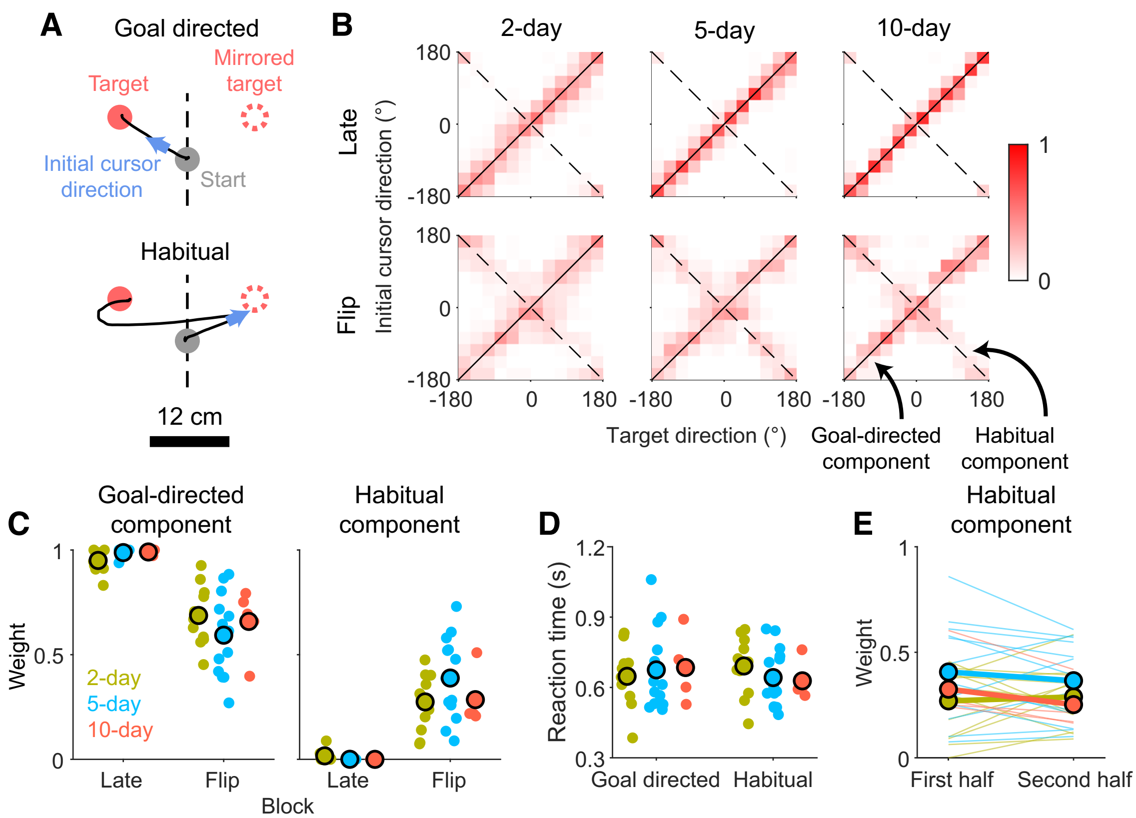

Figure 4. Analysis of habit in the point-to-point task. *A*: cursor trajectories (black line) from single trials in the flip block. The trajectories show trials in which the cursor's movement was initially aimed straight toward the target (*top*) or aimed toward a virtual target (*bottom*) mirrored across the vertical axis (dashed line). *B*: heat map of cursor's initial movement direction as a function of target directions. Data were pooled from all subjects and grouped into bins of 30° on both axes. We defined 0° as the positive *y*-axis (i.e., the mirroring axis). Within each target direction bin, we computed the fraction of trials that fell in a particular reach direction bin, plotting this fraction as color intensity in the heat map. To measure the proportion of trials in which participants exhibited habitual behavior, we fit a mixture model composed of 2 weighted von Mises distributions centered on either the $y = x$ (goal-directed behavior) or $y = -x$ (habitual behavior) line. *C*: fitted weights for the goal-directed (*left*) and habitual (*right*) components of the mixture model depicted in *B*. Fits for individual participants shown as small circles and group means shown as large circles. *D*: reaction time for reaches toward the actual target (goal directed) vs. a virtual mirrored target (habitual). *E*: weight of the habitual component of the mixture model when fit to data from either the first or second half of the flip block. Thin lines are individual participants, and thick lines are group means.

before initiating the horizontal component of their movement (Supplemental Fig. S2A). Therefore, as an alternative assay for habitual behavior, we computed the proportion of trials in which the horizontal component of the cursor's movement was initially directed away from the target. This analysis yielded qualitatively similar results, with behavior in the 2-day and 5-day groups being habitual (Supplemental Fig. S2B; linear mixed effects model with post hoc Tukey test; 2-day: $t = 8.54$, $P < 0.0001$; 5-day: $t = 11.00$, $P < 0.0001$) and no significant difference in the strength of habit between groups (2-day vs. 5-day: $t = -0.76$, $P = 0.8735$). The 10-day group also exhibited the same trend. Collectively, these results suggest that in the point-to-point task participants exhibited similarly strong habitual behavior regardless of whether they had practiced using the bimanual mapping for 2, 5, or 10 days.

Behavior in the Tracking Task Also Became Habitual after Only Two Days of Practice

We next examined participants' behavior in the tracking task to see whether they would exhibit similar habitual behavior under the flipped mapping, or whether habit effects might even be exacerbated given the imperative to generate movements rapidly while tracking the target. We compared the direction of participants' responses (i.e., cursor movement) to movements of the target between late learning and the flip block (Fig. 5*A*). During the flip block, if participants habitually behaved according to the original bimanual mapping, then their hand movement in response to horizontal target movement would be similar to late learning, and the horizontal movement of the cursor would therefore be directed opposite to the movement of the target. We expected that the extent of this effect might vary according to the frequency of target motion, with high frequencies being more likely to appear habitual owing to the need to respond more rapidly. We defined frequencies ≥0.85 Hz as "high frequency" as we have previously found that, above this frequency, participants habitually express baseline behavior when performing manual tracking under mirror reversal (40). We normalized the cursor's horizontal movement gain from the flip block by the gain from late learning such that normalized gains of −1 would indicate habitual behavior (Supplemental Fig. S3). [In subsequent analyses, we removed data from 1 outlier participant in the 10-day group

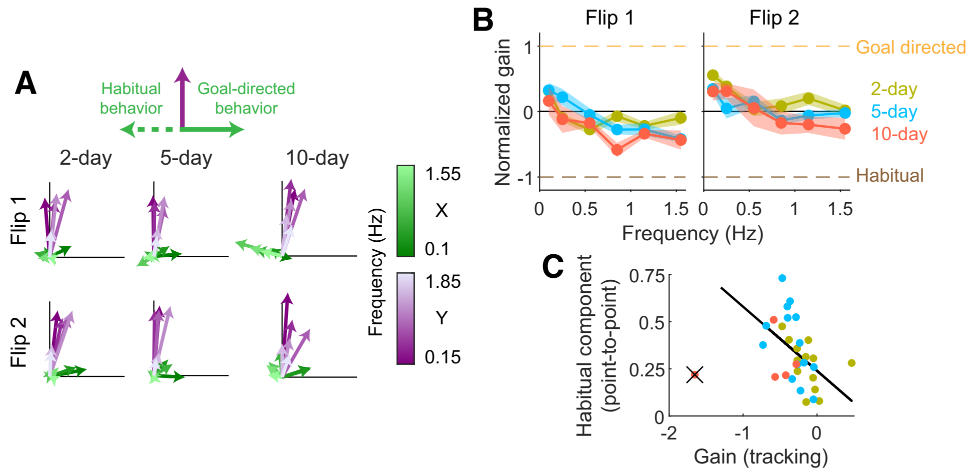

Figure 5. Analysis of habit in the tracking task. *A*: visualization of the cursor's movement direction and gain (relative to the target) while using the flipped mapping, similar to Fig. 3*C*. Each arrow depicts the average across participants at *x* (green)- and *y* (purple)-axis frequencies. Lower and higher frequencies are depicted as darker and lighter colors, respectively. Black lines are scale bars indicating a movement gain of 0.5. *Flip 1* and *Flip 2* are the first and second tracking blocks under the flipped mapping, respectively. At *top*, we depict the direction the green arrows should point if participants exhibit goal-directed (*right*) or habitual (*left*) behavior. *B*: gain of horizontal cursor movements under the flipped mapping normalized to the gain under the original bimanual mapping at late learning. If participants responded to target movement by moving their cursor in the same direction under both the original and flipped bimanual mappings, then the gain would be positive, approaching 1 if the movement gains were the same (goal directed; yellow dashed line). But if they moved their cursor in opposite directions under the 2 mappings, then the gain would be negative, approaching −1 (habitual; brown dashed line). Error bars indicate SE across participants. *C*: linear regression between average gains of the highest 3 frequencies from the first flipped tracking block and the proportion of habitual reaches from the flipped point-to-point block in Fig. 4*C*. Data from 1 outlier subject in the 10-day group (crossed out in black) was not used for fitting.

(*bottom left* participant in Supplemental Fig. S3), who exhibited dramatically more negative gains than we would expect if they exhibited habitual behavior.]

In the first flip block, the 2-day and 5-day groups exhibited negative gains (Fig. 5*B*; 1-sample *t* test with Holm–Bonferroni correction at $\alpha = 0.05$; 2-day: 2 of 6 frequencies, 5-day: 3 of 6 frequencies), particularly at higher frequencies, as we expected. The 10-day group also appeared to exhibit negative gains at similar frequencies. However, we did not find any evidence that the 5-day group exhibited significantly more negative gains than the 2-day group (linear mixed effects model with post hoc Tukey test: $P > 0.05$ for all comparisons of gains within frequencies). Although the 10-day group's movement at 0.85 Hz appeared to exhibit slightly more negative gains than the other groups, this effect was not consistent across frequencies.

Although the above analysis considers differences in behavior across groups, in previous work we have found that habitual behavior can vary greatly across individuals (41). We therefore also examined whether or not behavior was habitual at an individual participant level. We calculated the proportion of participants who exhibited significantly negative gains during the flip block. We found a mixture of habitual and nonhabitual participants in the 2-day and 5-day groups (Supplemental Fig. S3; 1-sample *t* test with Holm–Bonferroni correction at $\alpha = 0.05$; 2-day: 5 of 13 participants, 5-day: 7 of 14 participants). Although there was a slight increase in the proportion of participants who exhibited habitual behavior, it is difficult to conclude whether or not this trend was meaningful. Collectively, these data suggest that all three groups exhibited habitual behavior but groups with more practice were not more habitual than groups with less practice.

The Strength of Habitual Behavior Was Correlated between the Point-to-Point and Tracking Tasks

Might there be any relationship between the habitual behavior we observed in the point-to-point and tracking tasks? To examine this, we compared how strongly participants exhibited habitual behavior between the two tasks. First, we averaged the gains of each participant's tracking behavior at the highest three frequencies, given that we expected habitual behavior to be strongest at these frequencies (40). We then correlated each subject's average gain with the proportion of habitual reaches they made in the point-to-point task, as in Fig. 4*C*. Indeed, we found a correlation between tasks (Fig. 5*C*; slope = −0.34, Pearson's $r = 0.49$, $P = 0.0052$), suggesting that the tasks may have indeed assessed the same underlying habit (with the outlier included, slope = −0.12, Pearson's $r = 0.17$, $P = 0.2483$).

All Groups' Habits Were Similarly Resistant to Extinction

The preceding analysis quantifies the strength of habitual behavior either in terms of the probability that habitual behavior is expressed or in terms of the magnitude of the habitual response. However, an alternative way in which the "strength" of habitual behavior might vary with practice is by becoming more persistent, i.e., resistant to extinction. In other words, with increasing practice, the habits one forms may persist for longer. To assess whether habitual behavior would became more persistent with more training, we examined how participants' performance varied over the course of practicing the flipped mapping.

We first examined the persistence of participants' habits in the point-to-point task by fitting the mixture model from Fig. 4*C* to the first and last 50 reaches in the block instead of all 100 reaches. However, we found no evidence for any difference in the strength of habitual behavior between the first and last half of the block for all three groups (Fig. 4*E*; linear mixed effects model with post hoc Tukey test; 2-day: $t = -0.46$, $P = 0.9667$; 5-day: $t = 1.05$, $P = 0.7207$), suggesting that this aspect of habitual behavior did not extinguish over this period.

We then examined whether habitual behavior in the tracking task persisted when compared between the initial tracking block under the flipped mapping and a second tracking block performed after having practiced the flipped mapping in a point-to-point block (Fig. 1). At the group level, the 2-day and 5-day groups no longer exhibited significantly negative gains at any frequency in the second flip mapping (Fig. 5*B*; 1-sample t test with Holm–Bonferroni correction at $\alpha = 0.05$), and the 10-day group exhibited a similar trend, suggesting that the habits had been largely extinguished in all groups. However, at the level of individual participants, all the participants in the 10-day group who exhibited negative gains in the first tracking block still appeared to still do so in the second tracking block. Meanwhile, only one of five participants in the 2-day group and two of seven participants in the 5-day group still exhibited significantly negative gains. Although these data suggest that habitual behavior may have been more persistent in the 10-day group, they must be interpreted with caution given the small number of participants in this group.

DISCUSSION

In the present study, we examined the time course over which habitual behavior emerged as participants learned a new continuous motor skill, controlling a cursor under a novel bimanual hand-to-cursor mapping. Participants became more proficient in using this mapping by practicing with a combination of point-to-point reaches and continuous tracking, and their task performance plateaued after $\sim$5 days of practice. After 2, 5, or 10 days, we flipped the left hand's control of the cursor and tested whether participants would habitually continue to control the cursor according to the original mapping they had learned. We found that habitual behavior emerged after only 2 days of practice, which we observed in both the point-to-point and tracking tasks. We did not find compelling evidence, however, that habitual behavior became stronger with more practice.

An important caveat to our approach is that our analyses for the two different tasks (point to point and tracking) were not equally sensitive in detecting habitual behavior. In the point-to-point task, habitual behavior manifested as occasional trials in which participants acted habitually, and we were therefore able to detect the presence of even weak habitual behavior, which would have manifested as a very small proportion of trials being habitual. In the tracking task, however, we assessed habit based on the horizontal gain of the response to target movement. This gain likely reflected a mixture of goal-directed and habitual response components, either due to dual controllers operating in tandem or due to participants switching between goal-directed and habitual control at different moments, and would only be strictly negative if the habitual component of the response was strictly greater than the goal-directed component. A horizontal gain that was either zero or weakly positive could still be consistent with habitual control being present but dominated by goal-directed control. This may explain the apparent lack of habitual behavior at low frequencies (where goal-directed contributions to control were likely stronger) and in the second flip block. Importantly, however, only a strictly negative gain could be taken as unambiguous evidence of habitual behavior, since zero gain or weakly positive gain might have simply been due to control along that axis reverting to a more naive state.

Although we mainly focused on assessing how strong habitual behavior became during learning, to a limited extent we also assessed how persistent habits became, examining whether habits were extinguished while participants continued to practice the flipped mapping. Whereas a relatively short period of practice seemed to be sufficient to extinguish the habit in the tracking task, we did not observe extinction in the point-to-point task. It is possible that this discrepancy was attributable to the difference in sensitivity in detecting habitual behavior in these two tasks. It is also possible, however, that this difference was due to the order in which the tasks were performed; all of the point-to-point trials under the flipped mapping occurred between two blocks of tracking. Participants therefore had more experience with the new mapping by the time they performed the second tracking block than when they performed the second point-to-point block. Regardless of any possible differences in persistence across tasks, we emphasize that we did not find strong evidence to suggest that the persistence of habitual behavior within task depended on how long participants practiced the bimanual mapping.

One additional aspect of our experiment that we did not report in the results (since it was not directly relevant to our primary question of when control became habitual) was that at the end of each day participants performed an additional tracking block without visual feedback of the cursor's position. We used this block to examine the extent to which participants' learning could be attributed to improvements in feedforward control. However, we found that for all groups there was negligible improvement in mean-squared tracking error throughout learning and movement gains remained low (Supplemental Fig. S4), indicating that participants were not capable of expressing their learned behavior without visual feedback of the cursor.

It is often supposed that a newly learned behavior will become habitual only after it has been fully learned and then extensively repeated. Our findings clearly show that this is not the case. Instead, we found that behavior was habitual even relatively early in learning. Our findings parallel that of Hardwick et al. (41), who demonstrated that participants who learned a discrete arbitrary visuomotor association task exhibited improvements in their speed-accuracy trade-off (i.e., improved skill) over 20 days of practice, even though their behavior had become habitual after 4 days of practice. They further found that the habits could be explained as an all-or-none phenomenon (i.e., one either is or is not habitual), consistent with our observation that habitual behavior did not become stronger with more practice. The skill in Ref. 41 is quite rudimentary in that performance improvements amount only to speeding up action selection by tens of milliseconds, which might potentially have occurred via a specialized mechanism, perhaps associated with processing speed. In the present study, however, the improvements in skill after becoming habitual were more pronounced and unlikely to be explained by improvements in processing speed alone. The similarity of the results between the two studies suggests a potential common principle of habit formation in the process of learning new skill across both discrete and continuous domains.

Our findings have important implications for theoretical accounts of habits. A wide variety of theories have been proposed to explain the computational basis of habit formation, such as forming stimulus-response associations (21–24), caching expected future rewards (46–48), and caching computations/policies (9, 49). Central to these theories is the idea that habitual behavior is inflexible to change. Although behavioral inflexibility is (rightly) central to the definition of habits, our findings suggest that one may not need to break a habit to alter habitual behavior. To account for this, theoretical accounts of habits should allow scope for habitual behaviors to remain flexible to some degree. Learning rules that might accomplish this are incorporated into reinforcement learning-based frameworks, where habitual (model-free) behavior can be updated from experience (46–48, 50). Our results underscore that habitual behavior is not set in stone but can continue to evolve with experience, in this case over multiple days of practice. A similar computational principle is seen in studies of value-based decision-making, where model-based and model-free reinforcement learning occur in parallel. However, we emphasize that although model-free learning is often equated with habit (46–48), the type of learning in value-based decision is likely quite different from our task, since the learning here occurs gradually over multiple days whereas learning in value-based decision-making tasks can be observed in single trials (51). Furthermore, whereas prior theories have suggested that habitual control might only come to dominate goal-directed control after extensive repetition of the same task (46), we find that habitual control is prominent early in learning.

Conversely, our findings also have important implications for theories of motor control and motor learning. Most existing computational theories of motor learning apply only in very narrow settings and are devised to account for phenomena such as motor adaptation (52) or use-dependent learning (53, 54). These existing theories do not provide plausible models of de novo learning of the kind exhibited in our task. Although findings from visuomotor adaptation might suggest that cognitive strategies could play an important role in motor learning (55), our previous work examining learning in a mirror reversal task called into question the extent of explicit reaiming strategies in de novo learning (40). Our present results are broadly consistent with this point of view; one would expect cognitive strategies to be goal directed rather than habitual, but we, on the contrary, found that participants' behavior was habitual from early stages of learning. We therefore expect that new theories of motor learning will be required to account for de novo learning and the role that habits play in it.

One idea that could potentially be important in tying together skilled motor performance and habits is the concept of bounded rationality, which asserts that the brain has only limited resources at its disposal to solve problems, particularly cognitive resources. Rendering control habitual provides a means to reduce the cognitive load of a task. This is particularly important for complex skills that involve many component computations, in which case it would not be possible to perform the entire task deliberatively and instead most or all of these computations must be habitual (or, equivalently, automatized) (5). Recent theoretical work has characterized habits in the cognitive domain as a default behavioral policy that can be deviated from at the cost of cognitive effort (and, presumably, time) (50). We suggest that this framework could be fruitfully applied to motor behavior and is broadly consistent with our observations.

To conclude, a behavior becoming habitual is often viewed as the final step in learning: learned behavior must be repeated to render it habitual, at which point it becomes a persistent and dependable component of skilled performance. Although this idea may seem intuitively true, it has not previously been empirically tested. Our results challenge this view, suggesting that behaviors become habitual early in learning but maintain some flexibility to change with experience. We conclude that habits play an integral role in the learning and performance of motor skills from even the early stages of acquiring a new skill.

SUPPLEMENTAL MATERIAL

Data, code, and supplemental figures are available at https://doi.org/10.7281/T1/FWDYPW.

ACKNOWLEDGMENTS

We thank Yue Du, John Krakauer, Amy Bastian, and Christopher Fetsch for helpful discussions.

GRANTS

This material is based upon work supported by the National Science Foundation under Grant No. 1825489 (N.J.C.). This work was also supported by Facebook Technologies (A.M.H. and C.S.Y.). C.S.Y was supported by the Link Foundation Modeling, Simulation & Training Fellowship.

DISCLOSURES

No conflicts of interest, financial or otherwise, are declared by the authors.

AUTHOR CONTRIBUTIONS

C.S.Y. and A.M.H. conceived and designed research; C.S.Y. performed experiments; C.S.Y. analyzed data; C.S.Y., N.J.C., and A.M.H. interpreted results of experiments; C.S.Y. prepared figures; C.S.Y. drafted manuscript; C.S.Y., N.J.C., and A.M.H. edited and revised manuscript; C.S.Y., N.J.C., and A.M.H. approved final version of manuscript.

ENDNOTE

At the request of the authors, readers are herein alerted to the fact that additional materials related to this manuscript may be found at https://doi.org/10.7281/T1/FWDYPW. These materials are not a part of this manuscript and have not undergone peer review by the American Physiological Society (APS). APS and the journal editors take no responsibility for these materials, for the website address, or for any links to or from it.

REFERENCES

1. **Sandlin D.** The backwards brain bicycle—smarter every day 133. 2015. https://www.youtube.com/watch?v=MFzDaBzBlL0&t=2s.
2. **Wood W**, **Neal DT.** A new look at habits and the habit-goal interface. *Psychol Rev* 114: 843–863, 2007. doi:10.1037/0033-295X.114.4.843.
3. **Seger C**, **Spiering B.** A critical review of habit learning and the basal ganglia. *Front Syst Neurosci* 5: 66, 2011. doi:10.3389/fnsys.2011.00066.
4. **Hélie S**, **Cousineau D.** The cognitive neuroscience of automaticity: behavioral and brain signatures. In: *Advances in Cognitive and Behavioral Sciences*, edited by Sun MK. Hauppauge, NY: Nova Science Publishers, 2014, p. 141–159.
5. **Du Y**, **Krakauer JW**, **Haith AM.** The relationship between habits and motor skills in humans. *Trends Cogn Sci* 26: 371–387, 2022. doi:10.1016/j.tics.2022.02.002.
6. **Salmon DP**, **Butters N.** Neurobiology of skill and habit learning. *Curr Opin Neurobiol* 5: 184–190, 1995. doi:10.1016/0959-4388(95)80025-5.
7. **Bernacer J**, **Murillo JI.** The Aristotelian conception of habit and its contribution to human neuroscience. *Front Hum Neurosci* 8: 883, 2014. doi:10.3389/fnhum.2014.00883.
8. **Graybiel AM**, **Grafton ST.** The striatum: where skills and habits meet. *Cold Spring Harb Perspect Biol* 7: a021691, 2015. doi:10.1101/cshperspect.a021691.
9. **Haith AM**, **Krakauer JW.** The multiple effects of practice: skill, habit and reduced cognitive load. *Curr Opin Behav Sci* 20: 196–201, 2018. doi:10.1016/j.cobeha.2018.01.015.
10. **Marien H**, **Custers R**, **Aarts H.** Understanding the formation of human habits: an analysis of mechanisms of habitual behaviour. In: *The Psychology of Habit*, edited by Verplanken B. Cham, Switzerland: Springer International Publishing, 2018, p. 51–69. doi:10.1007/978-3-319-97529-0_4.
11. **Logan GD.** Skill and automaticity: elations, implications, and future directions. *Can J Psychol* 39: 367–386, 1985. doi:10.1037/h0080066.
12. **Bargh JA.** The ecology of automaticity: toward establishing the conditions needed to produce automatic processing effects. *Am J Psychol* 105: 181–199, 1992. doi:10.2307/1423027.
13. **Moors A**, **De Houwer J.** Automaticity: a theoretical and conceptual analysis. *Psychol Bull* 132: 297–326, 2006. doi:10.1037/0033-2909.132.2.297.
14. **de Wit S**, **Kindt M**, **Knot SL**, **Verhoeven AA**, **Robbins TW**, **Gasull-Camos J**, **Evans M**, **Mirza H**, **Gillan CM.** Shifting the balance between goals and habits: five failures in experimental habit induction. *J Exp Psychol Gen* 147: 1043–1065, 2018. doi:10.1037/xge0000402.
15. **Adams CD**, **Dickinson A.** Instrumental responding following reinforcer devaluation. *Q J Exp Psychol B* 33: 109–121, 1981. doi:10.1080/14640748108400816.
16. **Balleine BW**, **O'Doherty JP.** Human and rodent homologies in action control: corticostriatal determinants of goal-directed and habitual action. *Neuropsychopharmacology* 35: 48–69, 2010. doi:10.1038/npp.2009.131.
17. **Luque D**, **Molinero S**, **Watson P**, **López FJ**, **Le Pelley ME.** Measuring habit formation through goal-directed response switching. *J Exp Psychol Gen* 149: 1449–1459, 2020. doi:10.1037/xge0000722.
18. **Ceceli AO**, **Myers CE**, **Tricomi E.** Demonstrating and disrupting well-learned habits. *PLoS One* 15: e0234424, 2020. doi:10.1371/journal.pone.0234424.
19. **Popp NJ**, **Yokoi A**, **Gribble PL**, **Diedrichsen J.** The effect of instruction on motor skill learning. *J Neurophysiol* 124: 1449–1457, 2020. doi:10.1152/jn.00271.2020.
20. **Tricomi E**, **Balleine BW**, **O'Doherty JP.** A specific role for posterior dorsolateral striatum in human habit learning. *Eur J Neurosci* 29: 2225–2232, 2009. doi:10.1111/j.1460-9568.2009.06796.x.
21. **McDonald RJ**, **King AL**, **Hong NS.** Context-specific interference on reversal learning of a stimulus-response habit. *Behav Brain Res* 121: 149–165, 2001 [Erratum in *Behav Brain Res* 2002 128: 219–222, 2002]. doi:10.1016/S0166-4328(01)00160-7.
22. **Faure A**, **Haberland U**, **Condé F**, **Massioui NE.** Lesion to the nigrostriatal dopamine system disrupts stimulus-response habit formation. *J Neurosci* 25: 2771–2780, 2005. doi:10.1523/JNEUROSCI.3894-04.2005.
23. **Yin HH**, **Knowlton BJ.** The role of the basal ganglia in habit formation. *Nat Rev Neurosci* 7: 464–476, 2006. doi:10.1038/nrn1919.
24. **Robbins TW**, **Costa RM.** Habits. *Curr Biol* 27: R1200–R1206, 2017. doi:10.1016/j.cub.2017.09.060.
25. **Gritsenko V**, **Kalaska JF.** Rapid online correction is selectively suppressed during movement with a visuomotor transformation. *J Neurophysiol* 104: 3084–3104, 2010. doi:10.1152/jn.00909.2009.
26. **Telgen S**, **Parvin D**, **Diedrichsen J.** Mirror reversal and visual rotation are learned and consolidated via separate mechanisms: recalibrating or learning de novo? *J Neurosci* 34: 13768–13779, 2014. doi:10.1523/JNEUROSCI.5306-13.2014.
27. **Collins D**, **Morriss C**, **Trower J.** Getting it back: a case study of skill recovery in an elite athlete. *Sport Psychol* 13: 288–298, 1999. doi:10.1123/tsp.13.3.288.
28. **Hanin Y**, **Malvela M**, **Hanina M.** Rapid correction of start technique in an olympic-level swimmer: a case study using old way/new way. *J Swim Res* 16: 11–17, 2004.
29. **Carson HJ**, **Collins D**, **Jones B.** A case study of technical change and rehabilitation: intervention design and multidisciplinary team interaction. *Int J Sport Psychol* 45: 57–78, 2014.
30. **Haith AM**, **Yang CS**, **Pakpoor J**, **Kita K.** De novo motor learning of a bimanual control task over multiple days of practice. *J Neurophysiol* 128: 982–993, 2022. doi:10.1152/jn.00474.2021.
31. **Krakauer JW**, **Pine ZM**, **Ghilardi MF**, **Ghez C.** Learning of visuomotor transformations for vectorial planning of reaching trajectories. *J Neurosci*. 20: 8916–8924, 2000. doi:10.1523/JNEUROSCI.20-23-08916.2000.
32. **Fernandez-Ruiz J**, **Wong W**, **Armstrong IT**, **Flanagan JR.** Relation between reaction time and reach errors during visuomotor adaptation. *Behav Brain Res* 219: 8–14, 2011. doi:10.1016/j.bbr.2010.11.060.
33. **Bond KM**, **Taylor JA.** Flexible explicit but rigid implicit learning in a visuomotor adaptation task. *J. Neurophysiol* 113: 3836–3849, 2015. doi:10.1152/jn.00009.2015.
34. **Morehead JR**, **Qasim SE**, **Crossley MJ**, **Ivry R.** Savings upon re-aiming in visuomotor adaptation. *J Neurosci* 35: 14386–14396, 2015. doi:10.1523/JNEUROSCI.1046-15.2015.
35. **R Core Team.** *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing, 2020. https://www.R-project.org/.
36. **Bates D**, **Maechler M.** *Matrix: Sparse and Dense Matrix Classes and Methods*. 2019. https://CRAN.R-project.org/package=Matrix.
37. **Bates D**, **Mächler M**, **Bolker B**, **Walker S.** Fitting linear mixed-effects models using lme4. *J Stat Softw* 67: 1–48, 2015. doi:10.18637/jss.v067.i01.
38. **Kuznetsova A**, **Brockhoff PB**, **Christensen RH.** lmerTest package: tests in linear mixed effects models. *J Stat Softw* 82: 1–26, 2017. doi:10.18637/jss.v082.i13.
39. **Lenth R.** *emmeans: Estimated Marginal Means, aka Least-Squares Means*. 2020. https://CRAN.R-project.org/package=emmeans.
40. **Yang CS**, **Cowan NJ**, **Haith AM.** De novo learning versus adaptation of continuous control in a manual tracking task. *eLife* 10: e62578, 2021. doi:10.7554/eLife.62578.
41. **Hardwick RM**, **Forrence AD**, **Krakauer JW**, **Haith AM.** Time-dependent competition between goal-directed and habitual response preparation. *Nat Hum Behav* 3: 1252–1262, 2019. doi:10.1038/s41562-019-0725-0.
42. **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* 214: 1170–1180, 2011. doi:10.1242/jeb.048124.
43. **Madhav MS**, **Stamper SA**, **Fortune ES**, **Cowan NJ.** Closed-loop stabilization of the jamming avoidance response reveals its locally unstable and globally nonlinear dynamics. *J Exp Biol* 216: 4272–4284, 2013. doi:10.1242/jeb.088922.
44. **Sponberg S**, **Dyhr JP**, **Hall RW**, **Daniel TL.** Luminance-dependent visual processing enables moth flight in low light. *Science* 348: 1245–1248, 2015. doi:10.1126/science.aaa3042.

45. **Zimmet AM**, **Cao D**, **Bastian AJ**, **Cowan NJ.** Cerebellar patients have intact feedback control that can be leveraged to improve reaching. *eLife* 9: e53246, 2020. doi:10.7554/eLife.53246.
46. **Daw ND**, **Niv Y**, **Dayan P.** Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nat Neurosci* 8: 1704–1711, 2005. doi:10.1038/nn1560.
47. **Dolan RJ**, **Dayan P.** Goals and habits in the brain. *Neuron* 80: 312–325, 2013. doi:10.1016/j.neuron.2013.09.007.
48. **Hogarth L.** A critical review of habit theory of drug dependence. In: *The Psychology of Habit*, edited by Verplanken B. Cham, Switzerland: Springer International Publishing, 2018, p. 325–341. doi:10.1007/978-3-319-97529-0_18.
49. **Maisto D**, **Friston K**, **Pezzulo G.** Caching mechanisms for habit formation in active inference. *Neurocomputing* 359: 298–314, 2019. doi:10.1016/j.neucom.2019.05.083.
50. **Piray P**, **Daw ND.** Linear reinforcement learning in planning, grid fields, and cognitive control. *Nat Commun* 12: 4942, 2021. doi:10.1038/s41467-021-25123-3.
51. **Gläscher J**, **Daw N**, **Dayan P**, **O'Doherty JP.** States versus rewards: dissociable neural prediction error signals underlying model-based and model-free reinforcement learning. *Neuron* 66: 585–595, 2010. doi:10.1016/j.neuron.2010.04.016.
52. **Morehead JR**, **Orban de Xivry JJ.** A synthesis of the many errors and learning processes of visuomotor adaptation (Preprint). *bioRxiv* 2021.03.14.435278, 2021. doi:10.1101/2021.03.14.435278.
53. **Diedrichsen J**, **White O**, **Newman D**, **Lally N.** Use-dependent and error-based learning of motor behaviors. *J Neurosci* 30: 5159–5166, 2010. doi:10.1523/JNEUROSCI.5406-09.2010.
54. **Verstynen T**, **Sabes PN.** How each movement changes the next: an experimental and theoretical study of fast adaptive priors in reaching. *J Neurosci* 31: 10050–10059, 2011. doi:10.1523/JNEUROSCI.6525-10.2011.
55. **McDougle SD**, **Ivry RB**, **Taylor JA.** Taking aim at the cognitive side of learning in sensorimotor adaptation tasks. *Trends Cogn Sci* 20: 535–544, 2016. doi:10.1016/j.tics.2016.05.002.