De Novo Sensorimotor Skill Learning Through Reuse of Movement Components

G. A. Gabriel [1*], F. Mushtaq [1,2*], J. R. Morehead [1,2]

[1] School of Psychology, University of Leeds, LS29JT;

[2] Centre for Immersive Technologies, University of Leeds

*Correspondence to: George Gabriel (g.a.gabriel@leeds.ac.uk) or Faisal Mushtaq (f.mushtaq@leeds.ac.uk)

Keywords: Motor Learning; Skill; Generalisation; Electromyography

Abbreviated Title: *De Novo Learning Through Reuse*

Figures: *12*

Pages: *45*

Wordcount - Abstract: *249*

Wordcount - Introduction: *649*

Wordcount - Discussion: *1398*

Abstract

Recent theories of sensorimotor control propose that complex movements are constructed by combining simpler output-generating neural processes. Composing together reusable elements allows a rich variety of movements to be generated using a limited set of neural representations, but it may also impose limits on how new sensorimotor skills can be learned. If skill learning requires modification of existing compositional units, then other skills using the same components may be harmed when the new one is learned. Alternatively, if new skills simply reuse existing units, the repertoire of available units may impose limits on what skills can be learned. In this study, we used a myoelectric interface to explore whether learning a novel coordinated control task requires modification of the existing repertoire of output-generating components. Over five sessions on five consecutive days, participants learned to trace a series of trajectories using a computer cursor controlled by coordinated activation of two muscles. The timing of the generated cursor trajectory and its shape relative to the target improved for conditions trained with post-trial visual feedback. Improvements in timing transferred to all untrained conditions, but improvements in shape transferred only to untrained conditions requiring the trained order of muscle activation. Notably, all muscle outputs in the final session could already be generated during the first session. These results suggest that participants learned the new task by independently improving the selection and timing of existing output-generating components. This tendency toward reuse may help to protect existing skills from interference during learning of new skills.

Significance Statement

Complex movements may be generated by selecting and combining elements from a repertoire of movement-generating neural processes. It remains unclear if these processes are reused or modified when learning entirely new motor skills. In this study, we used a novel human-computer interface task to test this. Over five sessions, participants learned to trace a series of cursor trajectories by coordinating the activation of two muscles. We found that participants tended to reuse the same muscle contractions for different variants of the task, and that performance improvements were attributable to improvements in the choice and timing of muscle contractions. Our results demonstrate that learning of new complex movements does not necessarily require changes to the neural motor repertoire.

Introduction

Skilled sensorimotor behaviours often involve coordinated control of multiple muscles (Diedrichsen et al., 2010; Turvey, 1990). When tying our shoelaces, for example, the muscles controlling our fingers, hands, and arms must activate in precise spatial and temporal patterns to successfully shape the laces into a bow. An extensive literature on sensorimotor adaptation explains how such well-practised coordinated muscle activation can be adjusted to counteract a perturbing influence (Krakauer et al., 2019; Morehead & Xivry, 2021). However, our understanding of how these movements are learned to begin with remains relatively limited.

Recently, it has been argued that new coordinated control skills cannot be learned by adapting existing controllers (Haith et al., 2022; Yang et al., 2021, 2022; Gastrock et al., 2023; Gonda et al., 2019; Shmuelof et al., 2012), and that new controllers must instead be developed in a process termed de novo learning (literally, “learning anew”) (Krakauer et al., 2019). An important though easily overlooked question for this class of learning is how it balances reuse of existing learning with acquisition of new learning. If new skills are learned in isolation from existing ones, prior learning cannot be applied to speed up the acquisition of a new skill; when learning to tie our shoelaces and to write by hand, control of the fingers would have to be learned twice. Alternatively, if there is too much overlap between the new and existing skills, learning one skill could lead to changes in performance on the other (McCloskey & Cohen, 1989; Mermillod et al., 2013); mastering cursive could help or hinder our ability to tie a bow. Understanding the mechanisms of de novo learning therefore requires an explanation of how new learning proceeds without harming prior learning.

Recent explanations of skilled sensorimotor control propose that skilled movements are produced by combining elements from a repertoire of output-generating neural processes (Gallego et al., 2017, 2018; Golub et al., 2018; Hwang et al., 2013; Poggio & Bizzi, 2004). Elements of this repertoire could be implemented in various ways, including through spatiotemporal muscle synergies (Giszter, 2015; Overduin et al., 2015; Tresch & Jarc, 2009) or composable dynamical modes in motor-cortical neural populations (Chang et al., 2023; Ebitz & Hayden, 2021; Vyas et al., 2020; Gallego et al., 2017). While the neural implementation of the elements will determine how new sensorimotor skills are learned,

the same general mechanisms of de novo learning could apply under multiple implementations. For the purposes of this study, we consider two such possible mechanisms, independently of the implementation of the compositional elements.

Firstly, de novo learning could involve discovering a set of task-appropriate compositional elements and then developing the ability to reliably select and combine them (Diedrichsen & Kornysheva, 2015; Dhawale et al., 2017; Kim et al., 2021). In this case, interference between new and existing learning would be minimised, as the compositional units underlying multiple motor skills would not be modified. Alternatively, de novo skill learning may proceed by developing new or substantially modified compositional units (Oby et al., 2019). In this case, performance on pre-existing skills which used the modified elements would be particularly susceptible to disruption during de novo learning.

In the present study, we sought to determine whether changes to a shared repertoire of compositional units were necessary for learning of a novel feed-forward coordinated control task. We developed a myoelectric interface task in which participants controlled the horizontal velocity of a computer cursor using muscle activity recorded from the left shin and right hand. Over five sessions on consecutive days, participants were trained to follow two cursor trajectories, and we tested generalisation of learning to four untrained trajectories. Improvements in the shape of the untrained trajectories showed a distinctive pattern of generalisation indicative of reuse of a consistent set of muscle outputs. These results suggest that early de novo learning need not entail modifications to compositional units, and that re-use of existing units may be preferred.

Materials and Methods

Participants

A total of 20 participants (age 19-35, median 23 years; 12 male) each completed one session per day of an electromyographic control task for five consecutive days. Participants completed a pre-session questionnaire describing their prior experience with playing computer games, playing musical instruments, participating in sports, and driving. This information was not used to select participants for inclusion or exclusion from the study. All participants had no known neurological disorder and provided written consent through the online study sign-up process.

Participants were assigned to one of two conditions, labelled “congruent” (10 participants, 6 male) or “incongruent” (10 participants, 6 male). The two groups completed identical sessions but each used a differently configured myoelectric interface. During post-hoc analysis, one male participant was excluded from each group due to having extremely large mean cursor amplitude (more than double the maximum required amplitude) or mean cursor peak time after the end of each trial. The presented analyses use the remaining 18 participants unless otherwise stated.

Participants completed the Edinburgh handedness inventory, and we allowed participation regardless of handedness. Two participants (one male, one female) from the incongruent group and one female participant from the congruent group were identified as strongly left-handed.

Experiment Setup

Participants sat on a chair with both feet placed on blocks and arms resting on armrests (Figure 1). Abduction of the fingers of the right hand and dorsiflexion of the left foot were limited using velcro straps. The straps were placed around the fingers of the right hand and over the bridge of the left foot. Participants self-adjusted each strap to the maximum comfortable tension at the start of each session. Participants viewed visual feedback for the experiment tasks on a computer monitor (Dell AW2521HFLA, 24.5-inch, 1080 × 1920 pixels, 244Hz) at a distance of approximately 1.1 metres. The framerate of the task feedback was consistently above 110Hz.

Two bipolar EMG channels were recorded in real-time at 2048Hz using a biosignal amplifier (OTBioelettronica Quattrocento) and custom interface code written in Python. One channel was recorded from the left shin (tibialis anterior) and one from the right hand (abductor digiti minimi) of each participant. The shin electrodes were centred at points 7cm apart along a vertical line approximately 6cm below the tibial tuberosity and 2cm to the lateral side of the anterior margin of the tibia. The hand electrodes were centred at points 3cm apart approximately equidistant from the pisiform bone and the base of the fifth metacarpal bone. A reference electrode was placed on the ulnar styloid process of the right wrist. Locations of the electrodes were marked on the skin in ink and re-marked each session to allow consistent placement of the electrodes. All electrodes were self-adhesive solid gel type (Skintact F-261, 26mm diameter), and were further secured using micropore tape (hand electrodes) or kinesiology tape (shin electrodes).

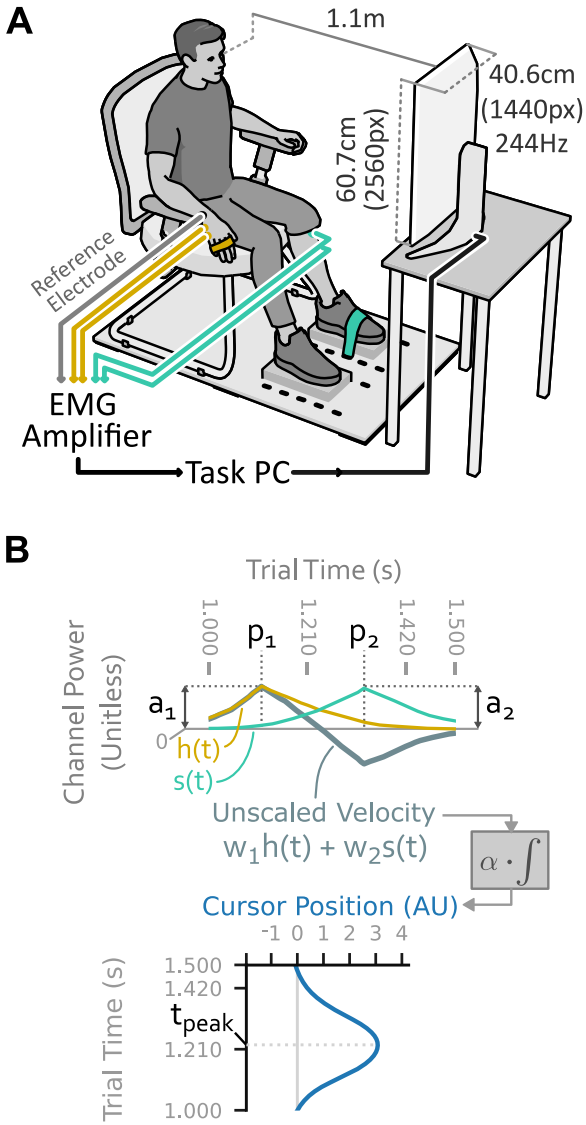

Figure 1 – Experiment setup and EMG interface *(A) Participants sat in a chair approximately 1.1m from the computer screen on which task feedback was displayed. A bipolar EMG channel was recorded in real time from the right abductor digiti minimi (yellow) and the left tibialis anterior (green), and used to control the horizontal position of an on-screen cursor. (B) The smoothed and scaled EMG signals generated by the hand, h(t), and shin, s(t), were combined in a weighted sum to produce the unscaled cursor velocity. The output cursor trajectory relative to the centre of the screen was given by the integral of this velocity signal multiplied by a constant velocity scale factor,* $\alpha = 2500\ pixels \cdot s^{-1}$*. The times (*$p_1$ *and* p_2*) and amplitudes (*a_1 *and* a_2*) of the per-channel peaks, and the time of the cursor trajectory peak (*t_{peak}*) are each used to quantify performance in subsequent analyses.*

129

Experiment Tasks

Each session comprised a series of calibration tasks and experiment tasks. Detailed instructions were presented to the participants through simultaneous on-screen text and audio narration. Instruction transcripts are available in the experiment data repository. All sessions were identical in structure (illustrated in Figure 2B), except for the addition of a single practice block and more detailed instructions in the first session.

Calibration

To set the power range of the two muscles, participants completed a maximum and minimum contraction task at the start of each session. The values recorded during the first session were used to set the gain of the electromyographic interface for all sessions. Maximum and minimum power level data from other sessions was used to track cross-session changes in signal-dependent noise for post-hoc analysis, but the gain of the interface was not changed after the first session.

Before calibration, participants were shown videos demonstrating how to activate the target muscles through abduction of the right little finger or dorsiflexion of the left foot. Calibration was completed separately for each of the two muscles. During calibration of a muscle, participants were shown a streaming lineplot on screen representing the instantaneous smoothed EMG power from that muscle. On maximum contraction trials, participants were instructed to contract the target muscle as strongly as possible using the demonstrated action, and to hold the contraction until told to release. The contraction stage lasted ten seconds from commencement of instructions to completion, and the mean of the smoothed EMG signal over the final four seconds of the contraction was set as the maximum power level of the muscle. The maximum contraction task could be repeated at the discretion of the experimenter if the participant appeared to have performed a sub-maximal or inconsistent contraction upon inspection of the EMG trace.

In the minimum power calibration task, participants were instructed to relax the muscle as fully as possible. This stage lasted fifteen seconds, and the mean of the smoothed EMG signal over the final ten seconds of data was used to define a noise threshold for the muscle. This task could also be

repeated at the discretion of the experimenter if movement or incomplete relaxation of the muscle was suspected.

Checkpoints were used after every four blocks of trials to re-assess the baseline noise level using the above-described minimum power calibration method. If the noise level during a checkpoint was found to be greater than the original baseline, the experimenter checked electrode adhesion and participant seating position and repeated the checkpoint.

Power Cycle

To allow participants to calibrate the strength of their muscle contractions to the range required in the main tasks, a “power cycle” task was completed after calibration in each session. For this task, one of the two EMG channels was selected, and the scaled power in that channel was used to control the vertical position of a small black circle. A larger grey circle moved with sinusoidal velocity up and down a line of 640 pixels height for three cycles over 50 seconds. The participant was instructed to move the black circle to keep it inside the grey circle by flexing the target muscle. Online visual feedback of the position of both circles was provided throughout. The same task was completed for each muscle in each session.

Free Movement and Path Following Tasks

The main task in each session involved controlling the velocity (not position, as was the case in the power cycle task) of an on-screen cursor. The cursor was constrained to move horizontally in the “cursor zone” vertically positioned at 1/3 screen height. To allow participants to familiarise themselves with the interface before starting the main path following task, a 30 second “free movement” stage was included after calibration in each session. During this stage, participants were shown the cursor and allowed to practice moving it using the velocity control interface. Online visual feedback of the cursor position was provided throughout.

In the main task (Figure 2A), participants were shown a grey box descending towards the cursor zone at a constant speed. The box descended for one second before reaching the cursor zone, after which it continued descending through the cursor zone for 0.5s. At either 0.5s or 0.8s before arrival at the

cursor zone, the box was replaced with a curving target path of equal height. This period is referred to as “preview time”. The path was represented by a series of 25 circular targets of six-pixel radius, uniformly vertically distanced along the shape of the path. Participants were instructed to move the cursor so that its tip (the highest point of the triangular cursor) touched as many of the descending circles as possible. Participants were specifically instructed to try and hit all targets rather than ignoring some of them. On all trials, the cursor became invisible 0.25s before the path arrived at the cursor zone. Online feedback of cursor position was therefore not available during the path-following segment of each trial, and participants had to hit the targets without seeing the current cursor position.

The horizontal cursor position was reset to the centre of the screen at the start of each trial. If the cursor moved more than 10 pixels from the centre of the screen while visible, the trial was abandoned, a warning buzzer sounded, and a written instruction was displayed informing the participant that they moved the cursor too early. Abandoned trials were repeated at the end of the block.

Each trial was followed by either a feedback stage or a no-feedback rest stage of three seconds duration. Performance feedback, when given, consisted of a trace of the actual cursor trajectory that the participant followed, aligned with the target circles. Hit targets were indicated as green filled circles, while missed targets were indicated as black unfilled circles.

Six different path variants were used, all of which were scaled versions of the same shape (Figure 2C). Three magnitudes of paths corresponding to 90, 180, and 270 pixels amplitude (1, 2, and 3 arbitrary units) and two directions (leftward peak and rightward peak) were used.

Two types of trial blocks were used: training blocks, and test blocks (Figure 2B). In training blocks, only the small-rightward and large-rightward paths were used, all trials had 0.8s preview time, and feedback was given after each trial. Training blocks included 30 trials of each of the two conditions, in pseudorandom order. In test blocks, all six path variants were used, both 0.5s and 0.8s preview times were used, and no feedback was given after each trial. Test blocks featured five trials of each of the 12 conditions, pseudorandomised such that the same condition was not repeated with fewer than 3 trials of other conditions between the repetitions. The first session included an additional

practice block before the first training block, in which participants practiced one trial of each of the twelve conditions with post-trial feedback.

A single round of the study consisted of two alternating train-test block pairs, interspersed with 20s rest periods, followed by a noise “checkpoint” as described above. Participants were also shown a session leaderboard for 10 seconds at the end of each round, featuring their and other anonymised participants’ cumulative numbers of targets hit up to that stage of the corresponding session. Three rounds were completed in each session. At the end of each session, the participants were shown a leaderboard featuring the total number of targets they had hit across all sessions, together with other participants’ totals.

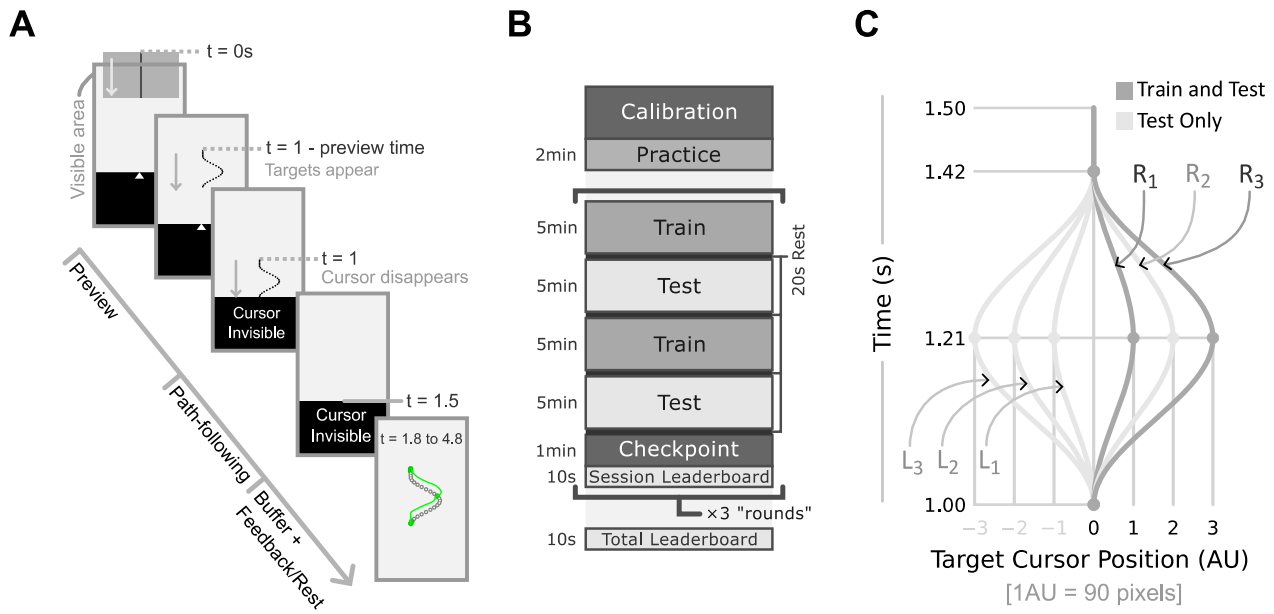

Figure 2 – Session structure and task conditions *(A) Each trial of the main task proceeded through several stages. A grey box with a central vertical line moved down the screen at a constant speed; the box was replaced with a series of 25 circular targets vertically distributed along the path trajectory; the cursor became invisible 0.7s after trial start; the targets continued to move steadily down the screen, reaching the cursor zone at 1s after trial start; and participants had to move the cursor to hit the targets without online visual feedback of cursor position. In training trials, feedback of targets hit (green circles) and produced cursor trajectory (green curve) were given for 3s. On probe trials, no feedback was given, and participants rested for 3s. (B) Participants followed a consistent schedule of tasks in each of the five experiment sessions, with additional instructions provided at the start of the first session (see main text). (C) Six different path shapes were used in the main task blocks. During training blocks, 30 alternating trials of each of the small and large amplitude rightward paths (R_1 and R_3,with magnitudes of 1 and 3 arbitrary units respectively) were used, all with 0.8s preview time and post-trial feedback. During the test blocks, all six path types were used with both the 0.5s and 0.8s preview times, with five repetitions of each of the 12 conditions per block and no post-trial feedback.*

EMG Interface

To create a low-latency control signal using the bipolar EMG signals, on each frame, a weighted average of the latest samples of rectified EMG data was computed for each channel. Two variants of the interface were used for different tasks. For the free movement and path following tasks, a weighted average of the rectified bipolar EMG was taken using a 256-sample triangular smoothing kernel which assigned greatest weight to the most recent EMG sample. For the power cycle task, a longer uniformly weighted kernel of 4096 samples was used. In both cases, the noise threshold for each channel (as identified during calibration) was subtracted from its smoothed EMG signal, and resulting values less than zero were set to zero. The thresholded and smoothed EMG channels were then scaled such that 35% of the participant's maximum power level produced an output signal of 1. We refer to the resulting time-varying signals as the channel profiles, $h(t)$ and $s(t)$, from the hand and shin muscle respectively (Figure 1B).

Two variants of the thresholding method were used for different tasks. For the main free movement and path-following tasks, a log-normal distribution was fitted to the baseline EMG data recorded during calibration, and the noise threshold was set at the 99.99th percentile of this distribution plus 1% of the EMG power at maximal contraction. This thresholding method provided robustness to noise without excessively reducing the dynamic range of the control signals. For the power cycle task, only the 1% of maximal contraction threshold was used. This was chosen to prevent the introduction of a noticeable "dead zone" in the controller, given that position control rather than velocity control was used in the power-cycle task.

For the path-following task, each control signal was linearly mapped to a one-dimensional velocity value, such that a control signal of 1 resulted in a cursor velocity of 2500 pixels per second. Unscaled control signals with magnitude greater than 1 were not capped. The entire screen width could therefore be traversed in 0.432 seconds without exceeding 35% of maximal power. The two channels each controlled velocity in opposite directions: one for positive (rightward) velocity, one for negative (leftward) velocity. A weighted sum of the two channels' control signals then determined the unscaled

243 velocity of the cursor. The integral of this unscaled velocity signal (multiplied by 2500 pixels per
244 second) gives the cursor trajectory used to complete the path-following task.

245 For participants from the congruent group, the left shin was mapped to leftward cursor velocity (i.e.
246 negative velocity values), while the right hand was mapped to positive velocity ($w_1 = 1$ and $w_2 = -1$
247 in Figure 1B). For participants in the incongruent group, the signs of the velocity for each channel
248 were flipped ($w_1 = -1$ and $w_2 = 1$ in Figure 1B). The same laterality of electrode placement (left shin
249 and right hand) was used for both groups.

250 Data Analysis

251 *Performance Measures*

252 To quantify improvements in the cursor trajectory shape independently of its timing relative to the
253 target trajectory, we computed a peak-aligned version of the output trajectories (illustrated in Figure
254 3A). For each input channel, we defined the peak amplitude of the channel profile as its maximum
255 value occurring between 0.7s and 1.8s after trial start and during any period of more than 16 samples
256 of consecutively non-zero activity (if such a period exists for the given trial). We also defined the
257 channel initiation as the time at which that consecutive interval of samples started. We then defined
258 the peak amplitude of the cursor trajectory as its maximal amplitude occurring between the identified
259 peak times of the two input channels. We next generated an interpolated version of the channel,
260 shifting it in time such that the identified peak occurred at 1.21s after trial start. This interpolation also
261 reduced the sampling rate of the cursor trajectory from 2048Hz to 1000Hz to reduce the computational
262 load for subsequent analyses.

263 Two basic performance measures were computed using the peak-aligned cursor trajectory. Firstly,
264 the peak-aligned target hit percentage is the percentage of the target points along the cursor trajectory
265 that the peak-aligned trajectory intersected. Secondly, the root-mean-squared peak-aligned cursor
266 trajectory error (denoted ϵ) is the root-mean-squared error between the observed peak-aligned cursor
267 trajectory and the target trajectory for that trial. The time of the cursor trajectory peak (computed
268 before peak alignment) is also used as a basic performance measure.

RMS Error Model

Two Bayesian regression models are used repeatedly throughout the analyses to produce estimates for the mean of the performance measures and channel features in each session. In all cases, we sampled the posterior distributions for the models using a NUTS Markov chain Monte Carlo sampler, implemented in Python.

For the RMS error, we used a hierarchical model with a common component shared across participants in the same congruence condition. In all applications of the model, we only used data from no-feedback trials where the hand and shin channel peaks were in the correct order (as determined by the target trajectory direction and the participant's congruence condition). We also centred, log-transformed, and re-scaled the values to have a sample standard deviation of 1 across all participants combined. The model is then as follows:

$$\mu_c \sim Normal(0,0.75)$$

$$\mu_{p,s,d} \sim Normal(\mu_c[p], 0.75)$$

$$\sigma_{p,s,d} \sim HalfNormal(1)$$

$$y_i \sim Normal\left(\mu_{p[i],s[i],d[i]}, \sigma_{p[i],s[i],d[i]}\right)$$

Where μ_c is a congruence group specific parameter, indexed by participant p; $\mu_{p,s,d}$ are the specific contributions to mean RMS for each participant, scale condition, and day (session); $\sigma_{p,s,d}$ is the specific standard deviation for each participant, scale condition, and day; and y_i is the (appropriately transformed) RMS of one observed trial.

Trajectory peak time also uses the same model, with y_i representing the (appropriately transformed) trajectory peak time of one observed trial.

To determine if there are differences in RMS error for different preview time conditions in session 1, we use a similar model where d is replaced by q, representing either 0.5s or 0.8s preview time.

Channel Peak Feature Model

For the peak amplitudes and times of each channel, we used a simpler model, transforming the data in the same way as for the RMS error model. This model was applied separately to data for each channel, again rejecting trials in which the order of channel peak activation was incorrect.

$$\mu_{\{p,s,d\}} \sim Normal(0,1)$$

$$\sigma_{\{p,s,d\}} \sim HalfNormal(1)$$

$$y_i \sim Normal\left(\mu_{\{p[i],s[i],d[i]\}}, \sigma_{\{p[i],s[i],d[i]\}}\right)$$

Where $\mu_{p,s,d}$ are the specific contributions to mean RMS for each participant, scale condition, and day (session); $\sigma_{p,s,d}$ is the specific standard deviation for each participant, scale condition, and day; and y_i is the (appropriately transformed) channel peak amplitude (or time) of one observed trial.

Channel Peak Time Correlation Model

To estimate the correlation between peak times in the first and second-activating channels, we used a multivariate normal model. Prior to model fitting, we subtracted the first channel activity start time from the peak times of both channels in each trial, and re-centred the resulting dataset to have sample mean of zero.

$$\mu_1, \mu_2 \sim Normal(0, 0.1)$$

$$\sigma_1, \sigma_2 \sim Exponential(0.5)$$

$$R \sim LKJcorr(1)$$

$$S = \begin{pmatrix} \sigma_1 & 0 \\ 0 & \sigma_2 \end{pmatrix} R \begin{pmatrix} \sigma_1 & 0 \\ 0 & \sigma_2 \end{pmatrix}$$

$$\begin{pmatrix} p_1 \\ p_2 \end{pmatrix} \sim MvNormal\left(\begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix}, S\right)$$

Where μ_1 and μ_2 are the prior means for the two-dimensional multivariate normal; σ_1 and σ_2 are the prior standard deviations for the covariance matrix; R is an LKJ prior for the unit standard deviation covariance matrix; S is the prior over covariance matrices, with scaled standard deviation; and $p_1$and p_2 are the channel peak times for the first and second-activating channels, with initiation time subtracted and re-centred to have sample mean of 0.

Bayes Factors

The reported Bayes factors are computed from the posterior distributions of the parameters of interest. Unless otherwise stated in the figure caption, the Bayes factors in favour of a reduction in a mean feature x value from session a to session b are computed using the formula:

$$\frac{\Pr(x_b < x_a \mid X)\Pr(x_b \geq x_a)}{\Pr(x_b \geq x_a \mid X)\Pr(x_b < x_a)}$$

Where X is the observed data. All prior and posterior probabilities are estimated by sampling from the respective distributions, and computing the proportion of samples satisfying the relevant inequality.

Results

Performance on trained conditions improved gradually over multiple sessions

Participants controlled the horizontal velocity of a computer cursor using bipolar EMG signals recorded from a muscle of the right hand (abductor digiti minimi) and a muscle of the left shin (tibialis anterior). Each participant was randomly assigned to one of two groups: congruent or incongruent. For the congruent group, each channel affected cursor velocity in the direction matching the laterality of the source muscle on the body (i.e., left shin to leftward velocity, and right hand to rightward velocity). For the incongruent groups, the muscles' laterality was unchanged, but the direction of their velocity contributions was reversed.

In the main task, participants were instructed to move the cursor to hit a series of circular targets that descended at a constant speed down the screen. Online visual feedback of cursor position was not provided in this task, but post-trial feedback of target hits and output cursor trajectory was provided during training blocks.

To determine whether the shape of participants' cursor trajectories improved independent of their timing relative to the target trajectory, we aligned the amplitude peaks of the output and target trajectories (Figure 3A) and computed performance measures based on these peak-aligned trajectories. Over five sessions, participants' mean peak-aligned target hit percentages for the trained trajectories (R_1 and R_3) showed a steady but statistically unreliable increase (Figure 3C). This is likely due to qualitative improvements in cursor trajectory shape which do not consistently result in more targets being hit, as demonstrated by the shape of session 5 median cursor trajectories compared to those of session 1 (Figure 3B, dark curves).

To provide a more sensitive measure of trajectory shape quality, we computed root-mean-squared error between the peak-aligned output and target trajectories. Marginalising across participants, we observed statistically robust reductions in the posterior mean and standard deviation of the peak-aligned RMS error between sessions 1 and 5 (Figure 3D, inset Bayes factors). These reductions are less robust for the incongruent group, likely because this group had lower variability in session 1 than

did the congruent group (Figure 3D, Bayes factors in caption). These results are consistent with improvements in timing-independent task performance arising from both changes in mean trajectory shape and changes in the variability of the generated trajectory shapes.

A separate feature of performance in the path-following task is the temporal alignment between the generated cursor trajectory and the target trajectory. To quantify how well participants aligned their outputs with the target, we computed the times at which each trajectory reached its largest amplitude (i.e. its peak) in comparison to the ideal peak time (Figure 4A). The per-participant and cross-participant marginal mean peak times show a statistically robust improvement, as measured by Bayes factors in favour of a reduction from session 1 to session 5 (Figure 4B, upper plots). Variability in peak timing may also show a reduction between sessions 1 and 5, but this is less statistically reliable for the incongruent group than the congruent group (Figure 4B, lower plots).

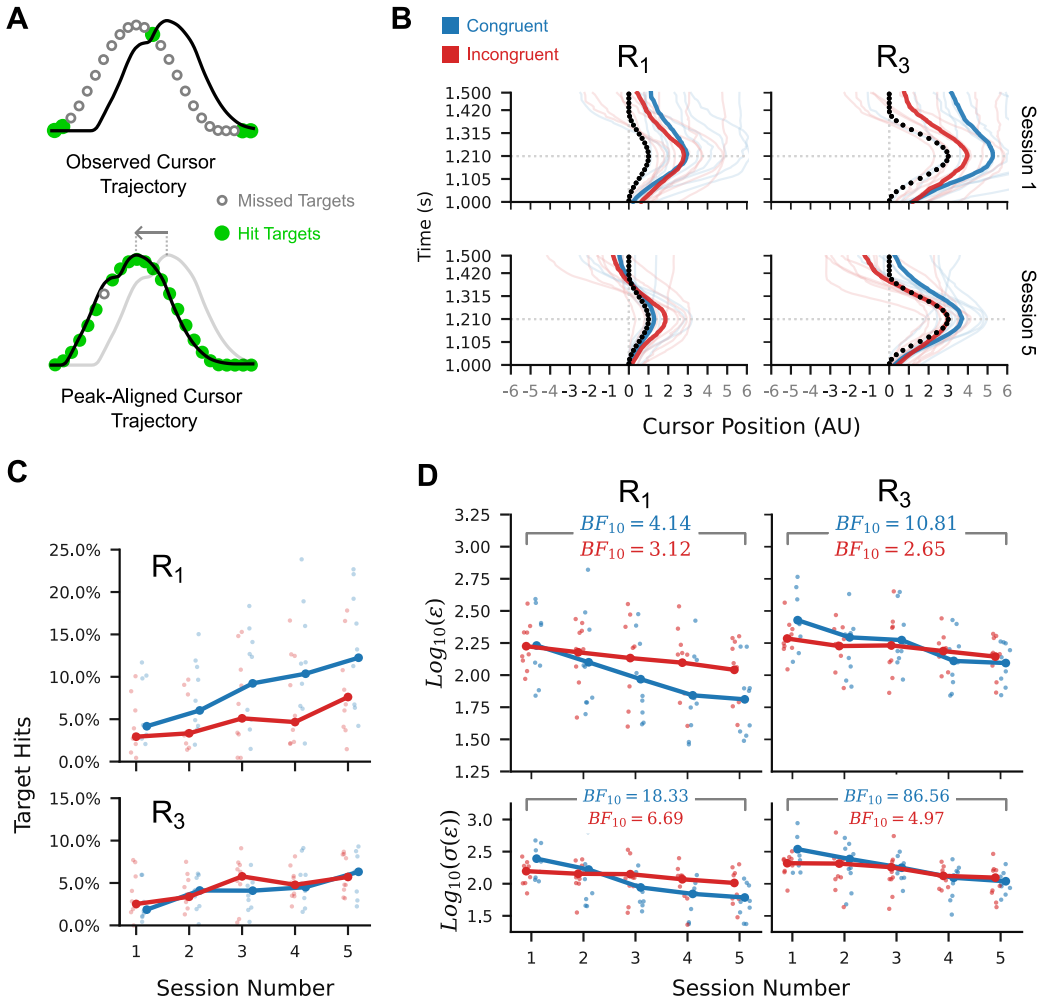

Figure 3 – Error and variability in trajectory shapes reduced across sessions for the trained conditions. *(A) Alignment of the peak of the observed cursor trajectory with the peak of the target trajectory. (B) Session 1 and 5 median cursor trajectories per participant (faint lines) and across participants (dark lines) in the conditions trained with post-trial feedback (R_1 and R_3). Inset black circles show the target points (to scale) for each trial condition. (C) Peak-aligned target hit percentages for the for the trained trial conditions. Lines show cross-participant medians within each condition group, and points show per-session means for individuals. Session 1 to session 5 increase in median hit percentage was 6.8% (SD=7.6%) for congruent R_1; 4.4% (SD=5.0%) for incongruent $_{R1}$; 2.5% (SD=3.9%) for congruent R_3; 2.6% (SD=3.3%) for incongruent R3. (D) Per-participant posterior means (top, faint points) and standard deviations (bottom, faint points) of peak-aligned RMS error in the generated cursor trajectories. Marginal means across participants (dark points and lines) show a steady reduction across sessions. Inset numbers are Bayes factors in favour of a reduction in marginal mean RMS from session 1 to session 5 for each of the two participant groups individually. Bayes factors in favour of the incongruent group having lower σ(ϵ) on session 1 than the congruent group are 11.23 for R_1, and 5.58 for R_3. Bayes factors in favour of the incongruent group having lower ϵ on session 1 than the congruent group are 1.09 for R_1, and 1.47 for R_3. All data is from "no-feedback" trials.*

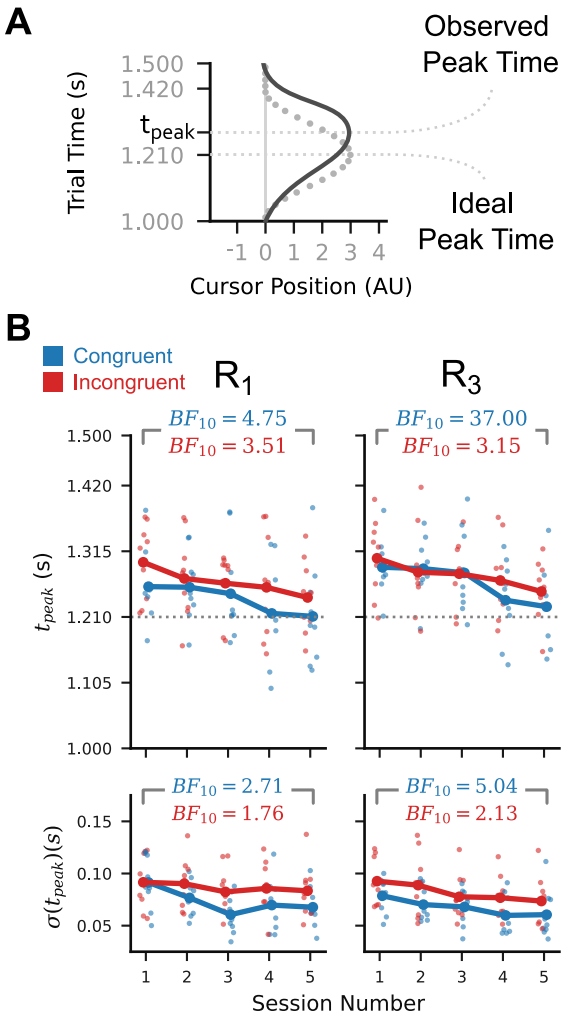

Figure 4 – Mean cursor trajectory peak time improved for trained conditions across sessions*. (A) Demonstration of the difference between target and observed trajectory peak times. The ideal trajectory peak time for all path scales was 1.21s after trial start. Note that this is slightly before the half-way point of the target trajectory, due to the flat "tail" section at the end of the target trajectory. (B) Posterior mean peak times (top row) and standard deviations in peak times (bottom row) for individual participants (faint points) and marginalised over participants within each congruence condition (dark lines). For both congruence conditions there is a statistically robust reduction in peak time from session 1 to session 5 on trajectories R_1 and R_2, as shown by the inset Bayes factors. All data is from "no-feedback" trials.*

363

364 **Performance improvements transferred to the untrained medium amplitude rightward**
365 **condition**

366 The above analyses demonstrate that, despite the unfamiliarity of the task and interface, participants
367 were able to execute the task and, on average, improved their performance on trained conditions
368 over the five sessions. We next sought to determine whether this learning could have arisen due to
369 the formation of habitual responses to the trained trajectories, rather than emergence of a new
370 controller as is purported to occur in de novo sensorimotor skill learning. To achieve this, we
371 assessed whether the performance improvements observed for the trained conditions arose
372 concurrently in the untrained task conditions. We reasoned that, while de novo skill learning could
373 support condition-general improvements in performance, habit formation should not (Du et al.,
374 2022).

375 The medium-magnitude rightward path (R_2) can be executed by generating EMG outputs with an
376 intermediate magnitude between those used for the small and large rightward paths. Unlike the other
377 three untrained conditions, R_2 also requires the same order of channel activation as the trained paths,
378 determined by the participants' congruence conditions. All performance results for R_2 were
379 qualitatively similar to those for R_1 and R_3: there was a statistically unreliable increase in peak-aligned
380 target hit percentage across the sessions (Figure 5A); the median shape of the output cursor
381 trajectories on "no-feedback" trials improved from session 1 to session 5 (Figure 5B); and we found
382 statistically robust Bayes factors in favour of reductions in the mean and standard deviation of peak-
383 aligned RMS (Figure 5C) and trajectory peak time (Figure 5D) between sessions 1 and 5.

384 Although performance on R_2 improved according to the selected measures, these trends do not show
385 that the outputs being generated for the R_2 trials were specific to that condition. Performance
386 improvements for this condition could also arise if participants used the same outputs for R_2 as for R_3
387 or R_1, as outputs suitable for these conditions will approximate the trajectory shape required by the
388 R_2 condition. It is unclear from inspection of the median cursor trajectories for the R_2 condition (Figure
389 5B) whether the generated outputs are distinct from those of the other two conditions. To clarify this,

we consider the per-participant distributions of input channel peak amplitudes (not to be confused with cursor trajectory peak amplitudes). These distributions show idiosyncratic differences dependent upon path condition. For some participants there is a clear difference in the amplitude distributions for each of the three conditions in session 5 (Figure 6A), while for other participants the medium rightward trial distribution in session 5 almost perfectly coincides with those of the large or small rightward trial conditions (Figure 6B). This suggests that, while some participants generated condition specific outputs, others simply re-used one or both outputs from the trained conditions to execute the R_2 condition. Grouping participants by whether their peak channel amplitudes had distinguishable or indistinguishable distributions for the three rightward trial conditions, we observed in both cases statistically robust reductions in RMS error (Figure 6C), trajectory peak time (Figure 6D), and the standard deviation thereof. This demonstrates that transfer of performance improvements to the untrained rightward condition was achieved either by production of untrained intermediate outputs, consistent with non-habitual learning, or by simple re-use of outputs suited to other conditions.

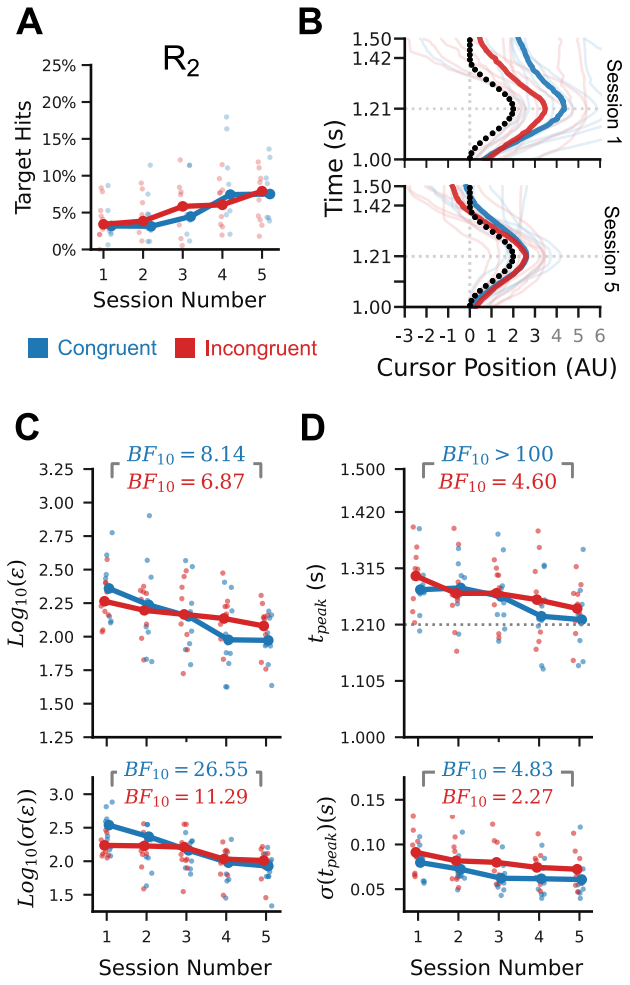

Figure 5 – Performance improvements transferred to the untrained R_2 condition. *(A) Peak-aligned target hit percentage for R_2. Session 1 to 5 increase is 3.8% (SD=4.5%) for congruent; 3.3% (SD=3.5%) for incongruent (B) Changes in median R_2 trial trajectory shape for individual participants (faint lines) and across participants (dark lines). (C) Changes in the logarithm of the posterior mean (top) and standard deviation (bottom) of peak-aligned RMS error in cursor trajectory for individual participants (faint points) and marginalising across participants (dark lines). Inset Bayes factors are in favour of a reduction in the marginal values from session 1 to session 5. (D) Similar to C, but for trajectory peak time. All data is from "no-feedback" trials.*

403

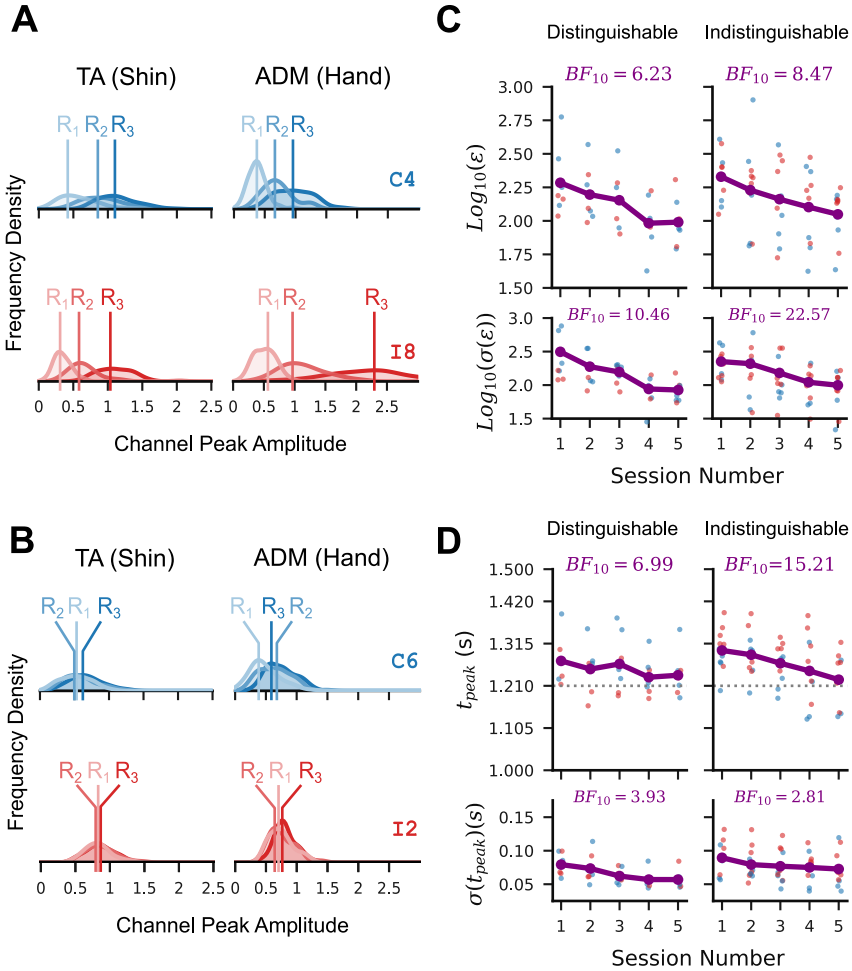

Figure 6 – Participants applied untrained condition-specific outputs or re-used existing non-specific outputs to achieve transfer of performance gains to the untrained R_2 condition. *A) KDE approximated distributions of observed per-channel peak amplitudes in session 5 for two example participants (C4, congruent; I8 incongruent), separated by trial condition. Vertical lines indicate the modes of each distribution. The distributions for these participants are clearly distinguishable. B) Distributions as in A, but for two other participants (C6 congruent, I2 incongruent) showing extensive overlap for the three trial conditions. C) All participants were assigned to one of two groups based on whether the peak amplitude distributions in at least one of the two channels were distinguishable. Participants were assigned to the "distinguishable" group (N = 7; 3 congruent) if the modal amplitudes for the three conditions were all more than 0.2 apart and assigned to the "indistinguishable" group otherwise (N = 11; 6 congruent). Plots show posterior Log-RMS marginalised across all participants in each distinguishable/indistinguishable group (purple points and lines). Red points are marginal posterior means for participants in the incongruent group, blue for the congruent group. Inset Bayes factors are in favour of a reduction from session 1 to session 5. D) Similar to C, but for trajectory peak time. Participants I3 and C8 are excluded from all analyses in this figure.*

404

Improvements in peak timing but not trajectory shape transferred to the untrained leftward conditions

To further clarify which learning processes resulted in the observed performance improvements on the rightward paths, we assessed how these improvements transferred to the leftward conditions (L_1, L_2, and L_3). These conditions were all untrained and required reversed order of input channel activation compared to the rightward conditions. As such, if the learning for the rightward conditions was specific to the trained order of muscle activation, we would expect little transfer to the leftward conditions.

Despite robust reductions in mean peak-aligned RMS for R_1 and R_2 between sessions 1 and 5, we found only weak evidence in favour of a reduction for L_1 and L_3 in the congruent group and weak evidence in favour of no change or an increase in the incongruent group (Figure 7A). The standard deviation of RMS for the leftward path showed similarly weak evidence of a reduction for both congruence conditions, with Bayes factors around a quarter the magnitude of those observed for R_1 and R_3 (Figure 3D).

Although the leftward path conditions require a different order of input channel activation than the trained rightward conditions, there is robust evidence in favour of an improvement in mean peak timing and a reduction in peak time standard deviation in the leftward conditions for the congruent participants (Figure 7B). The Bayes factors for the incongruent group also favour an improvement in peak time and a reduction in standard deviation of peak time, with evidence approximately as strong as in the corresponding trained rightward conditions.

Participants had different condition-specific biases in channel peak amplitude and timing

One possible reason for the limited transfer of trajectory shape improvements to the leftward conditions is that these trajectories may be intrinsically more difficult to generate than the corresponding magnitude rightward trajectories. If so, performance on these conditions during the first session should be worse than that of the rightward conditions. We observed no such bias in the

difference in log-RMS for rightward and leftward path trials of equal magnitude (Figure 8A), with individual participants instead showing idiosyncratic biases distributed approximately evenly around zero. This suggests that the limited transfer of trajectory shape quality to the leftward conditions is not a consequence of intrinsic differences in difficulty between the two directions.

A condition-dependent bias was observed for trajectory peak time in session 1. Participants from the congruent group tended to have a later peak time for leftward than for rightward trials, while the pattern was reversed for participants from the incongruent group (Figure 8B). This bias is likely a result of differences in the strength of inputs generated by each input channel. If participants tended to activate the shin channel more vigorously than the hand channel, this would result in an earlier trajectory peak when the shin channel was activated first and later when it was activated second. Consequently, the different congruence conditions will show opposite biases in peak time for the leftward and rightward conditions, due to their opposite mappings from muscle laterality to cursor velocity. Several participants showed such a difference in channel peak amplitudes which persisted in session 5 (Figure 10).

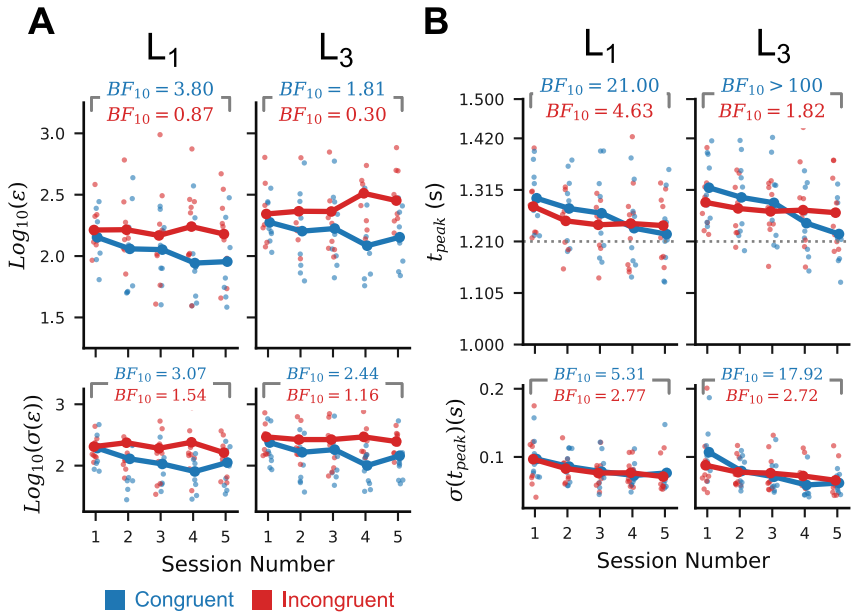

Figure 7 – Trajectory peak timing but not RMS error improved for the untrained leftward path conditions. *(A) Trends in log-transformed mean peak-aligned cursor trajectory RMS (top) and standard deviation thereof (bottom) for the L_1 and L_3 conditions. Qualitatively similar results were found for L_2 (not shown). B) Trends in trajectory peak time (top) and standard deviation thereof (bottom) for the L_1 and L_3 conditions. Inset Bayes factors are in favour of a reduction from session 1 to session 5.*

445

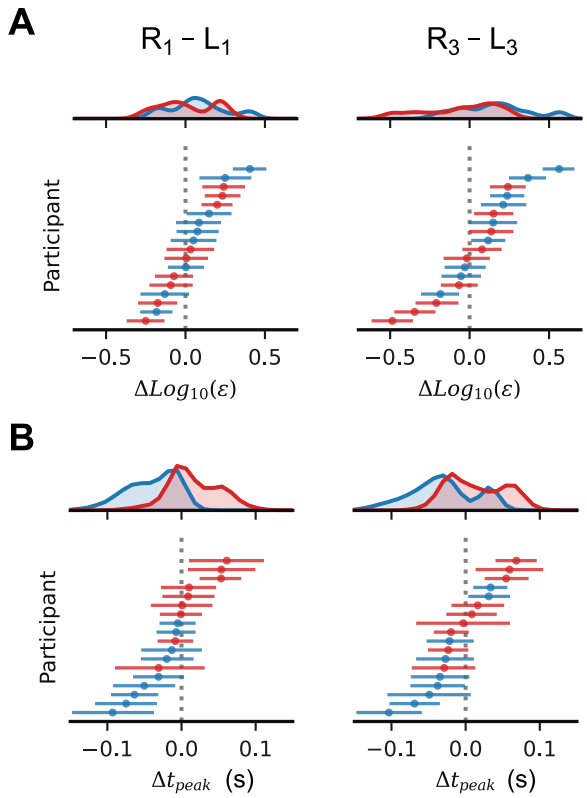

Figure 8 – Participants had idiosyncratic biases in error across the two trajectory directions, but neither the rightward nor leftward paths were intrinsically more difficult. *(A) Marginal posterior distributions for the difference in log-RMS between the rightward and leftward "no-feedback" trials. Individual horizontal lines are per-participant 95% posterior credible intervals. Shaded curves represent posterior density of the difference across all participants. Red features represent participants from the incongruent group, blue features represent participants from the congruent group. Columns correspond to different trajectory magnitudes. (B) Marginal posterior distributions for the difference in trajectory peak time between the rightward and leftward "no-feedback" trials. Features are as in A.*

446

Per-channel features were consistently similar for rightward and leftward paths of equal magnitude

The patterns of transfer described in the preceding section suggest that the learning processes responsible for determining trajectory shape and peak timing are partly independent. To clarify how these processes give rise to performance improvements, we now assess learning-related changes in the properties of the per-channel control signals.

To generate an appropriately timed cursor trajectory, the input channel profiles must themselves be appropriately timed. This could be achieved by triggering the channel inputs relative to some fixed movement initiation time, or by timing the second-activating channel relative to the first. To check for evidence of the latter case, we estimated the correlation between channel peak times after subtraction of movement initiation time. For the trained conditions, the resulting correlations were reliably positive for all participants in both the first and last sessions, with no consistent change in the posterior mean correlation coefficient between these sessions across participants (Figure 9A). Similar positive correlations are seen for the corresponding leftward paths, suggesting that the strategy of relative timing was consistently applied regardless of task condition, and was not strongly affected by training.

Although the second-activating channel is timed relative to the first-activating channel, the interval between activation of the two channels may vary across sessions without affecting the observed correlations. Improvements in trajectory peak time could therefore have been influenced by changes in the interval between channel activations. To check for such a change, we computed the posterior mean difference in channel peak times for each participant in each session and trial condition. The corresponding Bayes factors broadly support no change in cross-participant marginal mean inter-peak interval between sessions 1 and 5 (Figure 9B).

As the relative timing of the two channel outputs was consistent in the trained and untrained directions, a factor other than channel timing must be responsible for both the observed improvements in performance on the trained conditions and the lack of transfer to the leftward conditions. Direct comparison of the per-participant mean channel profiles for leftward and rightward trial conditions

suggests that the choice of profile shapes rather than their timing played an important role in both effects. In particular, the mean profiles for a given channel tend to be very similar regardless of whether that channel is activated in the context of a leftward or rightward path trial (Figure 10; light versus dark traces). This re-use of output profiles across directions has different consequences for the timing and the shape of the resulting cursor trajectory.

As the hand and shin channel profiles tend to differ in shape (including with respect to their amplitude), the cursor velocity resulting from taking a weighted sum of the two will differ depending on the order in which the profiles are generated. Consequently, the generated cursor trajectory will differ in shape when the same channel profiles are used in reversed order (Figure 11). Notably, the timing of the cursor trajectory peak is not as strongly affected by reversing the order of the channels. If each channel profile is approximately symmetrical about its peak and has approximately equal duration, activating the two channels at the same times but in reversed order will result in a cursor trajectory that reaches its peak at approximately the same time. These observations explain why we observed transfer of improvements in peak timing to the leftward path conditions but did not observe transfer of improvements in cursor trajectory shape to the leftward condition.

Participants did not learn to generate entirely new outputs

The performance improvements observed in the preceding sections are consistent with improved trajectory shape and reduced variability in the generated shape. We now consider whether this was achieved by developing the ability to generate entirely new outputs that could not be generated at the beginning of training.

Comparing the per-channel amplitude and peak timing in session 5 and session 1, we observe that outputs very similar to those used in session 5 could already be generated during the first session (Figure 12). This suggests that the observed improvements in task performance were not due to the participants learning to generate new per-channel outputs.

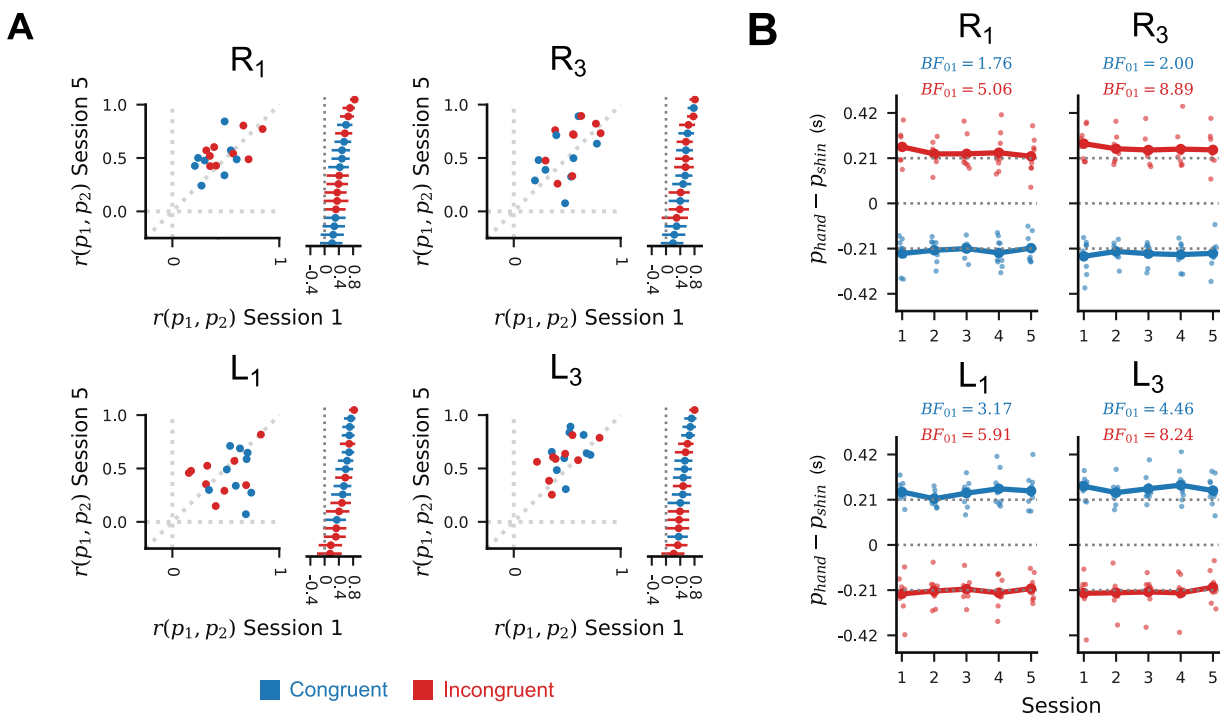

Figure 9 – Channel peaks are timed relative to each other, and inter-peak interval is consistent across sessions. *(A) Scatter plots show the per-participant marginal posterior correlation coefficients between first and second channel peak times in session 1 (horizontal axis) versus those in session 5 (vertical axis). Stacked lines show 99% posterior credible intervals around the session 1 correlation coefficient for each participant. Blue markers are for participants in the congruent group, while red markers are for participants in the incongruent group. (B) Points show posterior mean differences in per-channel peak times (i.e. inter-peak time intervals) for each path scale condition across sessions, in the test blocks. Dotted lines at* $\pm 0.21s$ *are ideal inter-peak times. Inset Bayes factors represent evidence in favour of no change in the inter-peak intervals for sessions 1 and 5.*

498

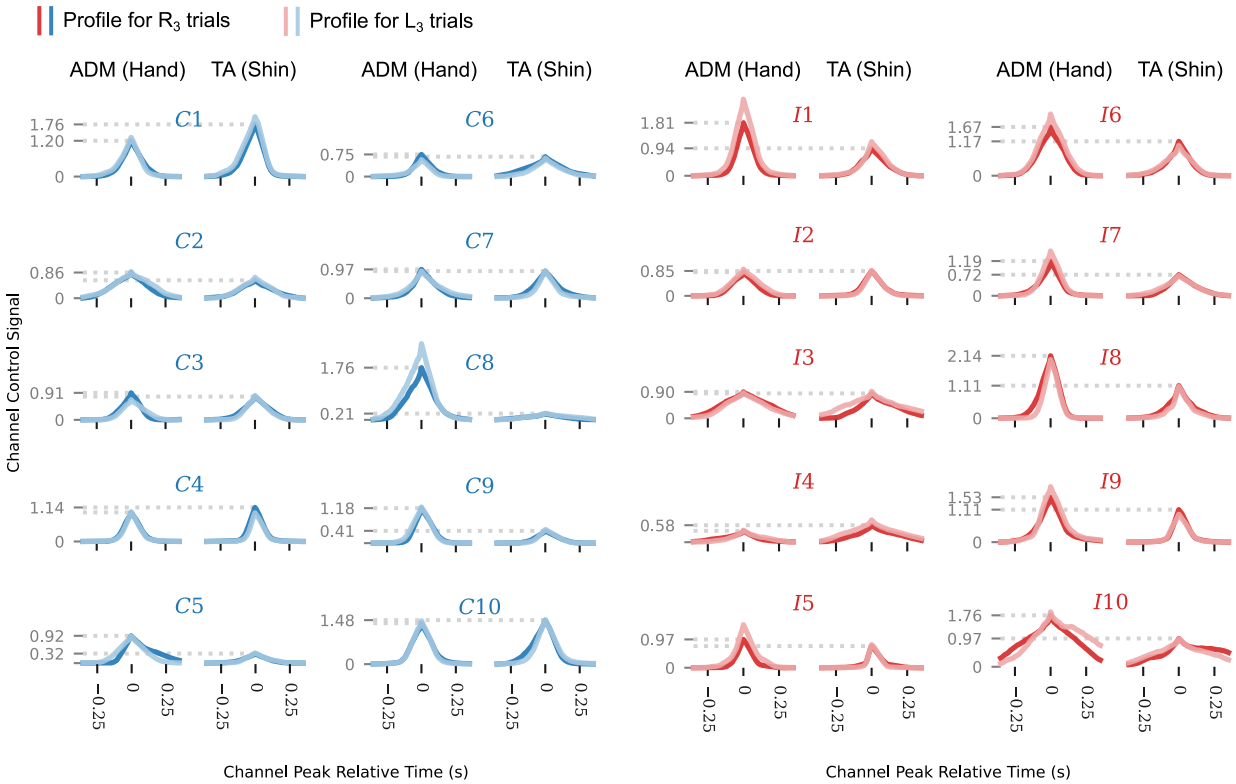

Figure 10 – Per-participant mean channel profiles were re-used across trajectory directions. *Plots show the mean peak-aligned per-participant channel profiles in R_3 trials (dark lines) and L_3 test trials in session 5. Dotted lines indicate the mean peak amplitudes for the R_3 trials. Red traces correspond to incongruent group participants, and blue traces correspond to congruent group participants. Results for other sessions and path magnitudes are qualitatively similar.*

499

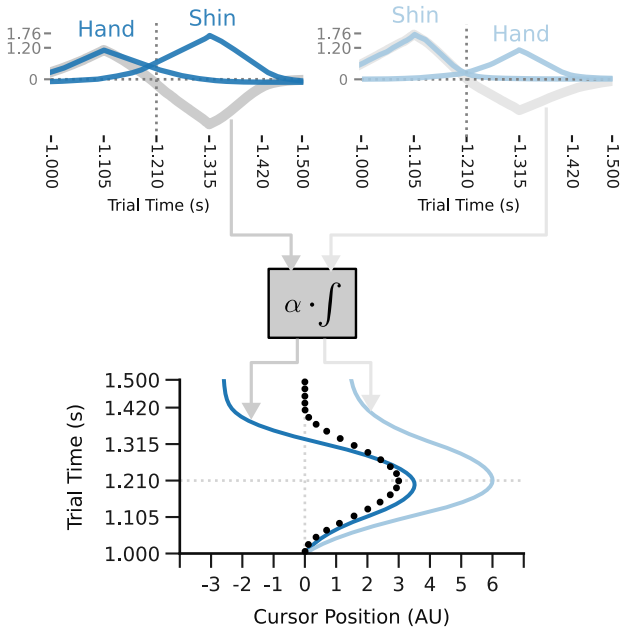

Figure 11 – Reuse of channel profiles suitable for rightward paths can result in transfer of peak timing and non-transfer of trajectory error to leftward paths*. Two pairs of mean channel profiles for R_3 trials (top left) and L_3 trials (top right) are very similar within channel but different across channels. For the rightward path condition, the hand channel is activated first, and the shin channel second, while the order is reversed for the leftward path condition. This leads to two different velocity profiles (top, grey lines), even when the peaks of the channels in each condition occur at the same two times. The cursor velocities are integrated and scaled to give the output cursor trajectory. This results in a different cursor trajectory in each condition, even though they used near identical per-channel outputs and relative timing.*

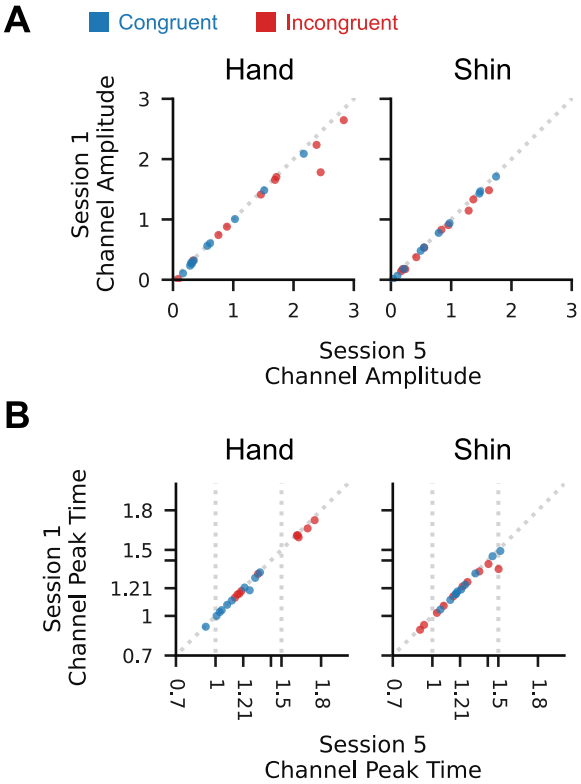

Figure 12 – Participants could produce per-channel outputs in session 1 which closely resembled those used in session 5. *(A) Points show the channel amplitude of a selected trial from session 5 plotted against that of another trial from session 1. The trials were selected by computing the minimum difference between the channel amplitude of a trial in session 5 and the channel amplitudes of all trials in session 1, then selecting the pair whose magnitude difference was the 99th percentile value. Plotted points therefore show differences in amplitude greater than or equal in magnitude to 99% of other trials. The consistently small difference between the session 5 and session 1 values demonstrates that participants could produce outputs matching the amplitude of those used in session 5 during the first session. (B) As in A, but for per channel peak time. Again, the peak times in session 5 could be generated during the first session.*

Discussion

We found that participants gradually improved both the shape and timing of cursor trajectories over five consecutive days of practice. This learning took place without substantial changes to the learner's repertoire of basic output-generating processes. This is notable, as this task had never been practised by the participants before the first session, and we selected muscles which are rarely coordinated together in natural movements. Distinctive patterns of transfer were observed for conditions practised without post-trial feedback. In particular, while improvements in the timing of the cursor trajectory peak were equally strong in all path conditions, improvements in cursor trajectory shape were clear only for the rightward path conditions. Moreover, although the channel outputs observed in the final session resulted in better task performance, they were consistently within the range of those generated by each participant in the first session. Based on these observations, we concluded that improvements in performance on the path following task did not require the emergence of new compositional units. We now briefly discuss how these improvements may have been achieved, and the implications of these possible mechanisms for de novo learning in general.

Transfer of trajectory shape improvements

In both direction conditions, the mean output profile for a given channel was very similar. This reuse of channel profiles across directions suggests that participants learned an association between the amplitude of the target trajectory and the pair of channel profiles to generate for that trajectory. For around half the participants in each condition group, this association generalised to the untrained medium-rightward condition, as shown by the distinct distributions of trajectory peak amplitudes in Figure 6. However, many participants reverted to reuse of small or large amplitude outputs in the medium rightward condition. Both cases are consistent with the existence of a relatively stable set of output-generating processes which can be selected by some other process or processes. The observed improvements in trajectory shape may therefore have arisen in part from improvements in these selection processes, regardless of whether selection appropriately distinguished between scale conditions.

528 **Transfer of trajectory timing improvements**

529 The transfer of peak timing across direction may be attributable to a separate mechanism of control.
530 Previous studies of sequence learning have suggested that, when the elements of the sequence
531 overlap in time, the later elements are likely to be timed relative to the state of the preceding element,
532 rather than relative to a common movement initiation time (Kornysheva, 2016). In the context of the
533 path-following task, we observed positive correlations in channel peak times in all task conditions,
534 consistent with relative timing. Improvements in trajectory peak timing in the trained rightward
535 conditions transferred to the leftward conditions despite the latter having reversed channel activation
536 order. This suggests that the physical state of the first-activating muscle does not determine the timing
537 of the second-activating muscle.

538 As the two channels usually have different profiles (see Figure 10), partly due to the different
539 physiological characteristics of the two muscles, using the precise state of one muscle to trigger
540 activation of the other would generally not produce the same timing when the order of the muscles
541 was reversed. Although we concluded that the second-activating channel is timed relative to the first-
542 activating channel during all five sessions, this does not imply that the second-activating channel is
543 timed relative to the intensity of activity in the first-activating channel. Rather, we suggest that the
544 second-activating channel may have been timed relative to some qualitative feature of the first-
545 activating channel profile, such as its peak or the time at which power in that channel started to reduce
546 after the peak. It should be noted that this feature-relative timing could be achieved during movement
547 planning or based on feedback received during execution. Further studies will be required to test
548 these possibilities.

549 The observed independence of channel peak timing and amplitude is consistent with results from
550 neuroimaging studies of discrete sequence learning. For discrete finger presses, the production of
551 ordered output sequences is attributed to a hierarchical representation in which the complete
552 sequence is built up from successively smaller sub-sequences (Rosenbaum et al., 1983; Sakai et al.,
553 2003). Separate representations of the order and timing of finger press sequences are found in
554 bilateral pre-motor areas during movement preparation, with integrated representations of both order

and timing arising in primary motor areas contralateral to the active hand during sequence execution (Kornysheva & Diedrichsen, 2014; Yewbrey et al., 2023; Yokoi & Diedrichsen, 2019). We would anticipate a similar independence in the neural representation of order and timing for the two muscles used in the present study. This suggests a plausible neural basis for independent improvements in the timing and the selection of amplitudes for the coordinated motor outputs in the path following task.

As the second-activating channel is timed relative to the first-activating channel, the peak time of the generated cursor trajectory is influenced by both the relative lag between the two channels and the onset time of the first channel. We found little evidence for a change in the average time interval between channel peaks in any condition (Figure 9B). This suggests that the main channel timing-related behavioural feature affecting trajectory peak timing was the timing of the onset of the whole movement. Although the rightward and leftward direction paths require activation of a different muscle at the start of the movement, we observed that improvements in the timing of the cursor trajectory achieved by practising the R_1 and R_3 conditions were preserved in the untrained leftward conditions. This suggests that the mechanisms responsible for onset timing are independent of the muscle being activated. This may partly explain why improvements in trajectory peak timing transferred to the untrained leftward conditions despite their requiring reversed order of muscle activation relative to the rightward conditions.

Is this learning “de novo”?

Together, the above-described results suggest that learning of a new continuous control task can be achieved by improving the selection and timing of outputs that are already in the repertoire of the learner. It remains unclear, however, if this is true of de novo learning tasks in general, or if it is a consequence of specific features of the task used in this study.

Previous work has attributed the performance improvements arising from de novo learning to the acquisition of a new controller or the development of a new policy determining which actions to take in which circumstances (Haith et al., 2022; Yang et al., 2021). Our analyses suggest that some participants formed general associations between task conditions and different output magnitudes and channel activation orders, consistent with the emergence of a context-dependent policy. However,

this does not imply that the same processes were responsible for learning in the present study as in all previous de novo learning studies.

Unlike in many previous de novo learning paradigms (Haith et al., 2022; Mosier et al., 2005; Ranganathan et al., 2014; Shmuelof et al., 2012; Yang et al., 2021), the EMG interface used in the present study controlled the velocity of the output cursor rather than its position. Using the velocity control interface, the path-following task could be completed by generating a pair of appropriately timed pulses of EMG activity. The mechanisms involved in learning to generate a well-timed sequence of discrete motor outputs are likely to differ from those involved in learning more continuous control tasks (Krakauer et al., 2019). As such, although the learning observed in this study appears to fit the definition of de novo learning, the mechanisms supporting learning in the path-following task may not be identical to that observed in other de novo learning studies. It is already well understood that sensorimotor learning is supported by multiple interacting learning processes (Hennig et al., 2021), and we suggest that de novo learning should be similarly understood as arising from a range of learning mechanisms, differently recruited by different tasks.

Our study clarifies some of the possible mechanisms underlying learning of novel sensorimotor skills. We have demonstrated that continuous control of one degree of freedom can be improved by adjusting the selection and timing of constituent motor outputs. Notably, for some participants the acquired ability retained some of the flexibility which distinguishes skilled from habitual behaviour. These results suggest that modification of basic output-generating components may, in some circumstances, be unnecessary for the development of new sensorimotor skills. This flexible re-use of elements in the output-generating repertoire may enable new learning while also preserving existing skills.

Data Availability

All data and analysis scripts will be made publicly available via an Open Science Framework repository upon acceptance of the manuscript.

Author Contributions

G. A. Gabriel: Conceptualisation, Formal Analysis, Investigation, Methodology, Software, Visualisation, Writing – original draft

J. R. Morehead: Methodology, Resources, Supervision, Writing – review and editing

F. Mushtaq: Methodology, Resources, Supervision, Writing – review and editing

References

Chang, J. C., Perich, M. G., Miller, L. E., Gallego, J. A., & Clopath, C. (2023). *De novo motor learning creates structure in neural activity space that shapes adaptation* (p. 2023.05.23.541925). bioRxiv. https://doi.org/10.1101/2023.05.23.541925

Dhawale, A. K., Smith, M. A., & Ölveczky, B. P. (2017). The Role of Variability in Motor Learning. *Annual Review of Neuroscience*, *40*(1), 479–498. https://doi.org/10.1146/annurev-neuro-072116-031548

Diedrichsen, J., & Kornysheva, K. (2015). Motor skill learning between selection and execution. *Trends in Cognitive Sciences*, *19*(4), 227–233. https://doi.org/10.1016/j.tics.2015.02.003

Diedrichsen, J., Shadmehr, R., & Ivry, R. B. (2010). The coordination of movement: Optimal feedback control and beyond. *Trends in Cognitive Sciences*, *14*(1), 31–39. https://doi.org/10.1016/j.tics.2009.11.004

Du, Y., Krakauer, J. W., & Haith, A. M. (2022). The relationship between habits and motor skills in humans. *Trends in Cognitive Sciences*, *26*(5), 371–387. https://doi.org/10.1016/j.tics.2022.02.002

Ebitz, R. B., & Hayden, B. Y. (2021). The population doctrine in cognitive neuroscience. *Neuron*, *109*(19), 3055–3068. https://doi.org/10.1016/j.neuron.2021.07.011

Gallego, J. A., Perich, M. G., Miller, L. E., & Solla, S. A. (2017). Neural Manifolds for the Control of Movement. *Neuron*, *94*(5), 978–984. https://doi.org/10.1016/j.neuron.2017.05.025

Gallego, J. A., Perich, M. G., Naufel, S. N., Ethier, C., Solla, S. A., & Miller, L. E. (2018). Cortical population activity within a preserved neural manifold underlies multiple motor behaviors. *Nature Communications*, *9*(1), Article 1. https://doi.org/10.1038/s41467-018-06560-z

Gastrock, R. Q., Hart, B. M. ’t, & Henriques, D. Y. P. (2023). *Distinct learning, retention, and generalization in de novo learning* (p. 2023.10.02.560506). bioRxiv. https://doi.org/10.1101/2023.10.02.560506

Giszter, S. F. (2015). Motor primitives—New data and future questions. *Current Opinion in Neurobiology*, *33*, 156–165. https://doi.org/10.1016/j.conb.2015.04.004

Golub, M. D., Sadtler, P. T., Oby, E. R., Quick, K. M., Ryu, S. I., Tyler-Kabara, E. C., Batista, A. P., Chase, S. M., & Yu, B. M. (2018). Learning by neural reassociation. *Nature Neuroscience*, *21*(4), 607–616. https://doi.org/10.1038/s41593-018-0095-3

Gonda, S., Shkedy Rabani, A., Horesh, N., & Shmuelof, L. (2019). Fast and specific: Insights into the acquisition and generalization of motor acuity. *Journal of Neurophysiology*, *122*(6), 2354–2363. https://doi.org/10.1152/jn.00558.2018

Haith, A. M., Yang, C. S., Pakpoor, J., & Kita, K. (2022). De novo motor learning of a bimanual control task over multiple days of practice. *Journal of Neurophysiology*, *128*(4), 982–993. https://doi.org/10.1152/jn.00474.2021

Hennig, J. A., Oby, E. R., Losey, D. M., Batista, A. P., Yu, B. M., & Chase, S. M. (2021). How learning unfolds in the brain: Toward an optimization view. *Neuron*, *109*(23), 3720–3735. https://doi.org/10.1016/j.neuron.2021.09.005

Hwang, E. J., Bailey, P. M., & Andersen, R. A. (2013). Volitional Control of Neural Activity Relies on the Natural Motor Repertoire. *Current Biology*, *23*(5), 353–361. https://doi.org/10.1016/j.cub.2013.01.027

Kim, H. E., Avraham, G., & Ivry, R. B. (2021). The Psychology of Reaching: Action Selection, Movement Implementation, and Sensorimotor Learning. *Annual Review of Psychology*, *72*(1), 61–95. https://doi.org/10.1146/annurev-psych-010419-051053

Kornysheva, K. (2016). Encoding Temporal Features of Skilled Movements—What, Whether and How? In J. Laczko & M. L. Latash (Eds.), *Progress in Motor Control: Theories and Translations* (pp. 35–54). Springer International Publishing. https://doi.org/10.1007/978-3-319-47313-0_3

Kornysheva, K., & Diedrichsen, J. (2014). Human premotor areas parse sequences into their spatial and temporal features. *ELife*, *3*, e03043. https://doi.org/10.7554/eLife.03043

Krakauer, J. W., Hadjiosif, A. M., Xu, J., Wong, A. L., & Haith, A. M. (2019). Motor Learning. In *Comprehensive Physiology* (pp. 613–663). John Wiley & Sons, Ltd. https://doi.org/10.1002/cphy.c170043

McCloskey, M., & Cohen, N. J. (1989). Catastrophic Interference in Connectionist Networks: The Sequential Learning Problem. In G. H. Bower (Ed.), *Psychology of Learning and Motivation* (Vol. 24, pp. 109–165). Academic Press. https://doi.org/10.1016/S0079-7421(08)60536-8

Mermillod, M., Bugaiska, A., & Bonin, P. (2013). The stability-plasticity dilemma: Investigating the continuum from catastrophic forgetting to age-limited learning effects. *Frontiers in Psychology*, *4*.

Morehead, J. R., & Xivry, J.-J. O. de. (2021). *A Synthesis of the Many Errors and Learning Processes of Visuomotor Adaptation* (p. 2021.03.14.435278). bioRxiv. https://doi.org/10.1101/2021.03.14.435278

Mosier, K. M., Scheidt, R. A., Acosta, S., & Mussa-Ivaldi, F. A. (2005). Remapping Hand Movements in a Novel Geometrical Environment. *Journal of Neurophysiology*, *94*(6), 4362–4372. https://doi.org/10.1152/jn.00380.2005

Oby, E. R., Golub, M. D., Hennig, J. A., Degenhart, A. D., Tyler-Kabara, E. C., Yu, B. M., Chase, S. M., & Batista, A. P. (2019). New neural activity patterns emerge with long-term learning. *Proceedings of the National Academy of Sciences*, *116*(30), 15210–15215. https://doi.org/10.1073/pnas.1820296116

Overduin, S. A., d'Avella, A., Roh, J., Carmena, J. M., & Bizzi, E. (2015). Representation of Muscle Synergies in the Primate Brain. *Journal of Neuroscience*, *35*(37), 12615–12624. https://doi.org/10.1523/JNEUROSCI.4302-14.2015

Poggio, T., & Bizzi, E. (2004). Generalization in vision and motor control. *Nature*, *431*(7010), Article 7010. https://doi.org/10.1038/nature03014

Ranganathan, R., Wieser, J., Mosier, K. M., Mussa-Ivaldi, F. A., & Scheidt, R. A. (2014). Learning Redundant Motor Tasks with and without Overlapping Dimensions: Facilitation and Interference Effects. *Journal of Neuroscience*, *34*(24), 8289–8299. https://doi.org/10.1523/JNEUROSCI.4455-13.2014

Rosenbaum, D. A., Kenny, S. B., & Derr, M. A. (1983). Hierarchical control of rapid movement sequences. *Journal of Experimental Psychology: Human Perception and Performance*, *9*(1), 86–102. https://doi.org/10.1037/0096-1523.9.1.86

Sakai, K., Kitaguchi, K., & Hikosaka, O. (2003). Chunking during human visuomotor sequence learning. *Experimental Brain Research*, *152*(2), 229–242. https://doi.org/10.1007/s00221-003-1548-8

Shmuelof, L., Krakauer, J. W., & Mazzoni, P. (2012). How is a motor skill learned? Change and invariance at the levels of task success and trajectory control. *Journal of Neurophysiology*, *108*(2), 578–594. https://doi.org/10.1152/jn.00856.2011

Tresch, M. C., & Jarc, A. (2009). The case for and against muscle synergies. *Current Opinion in Neurobiology*, *19*(6), 601–607. https://doi.org/10.1016/j.conb.2009.09.002

Turvey, M. T. (1990). Coordination. *American Psychologist*, *45*(8), 938–953. https://doi.org/10.1037/0003-066X.45.8.938

Vyas, S., Golub, M. D., Sussillo, D., & Shenoy, K. V. (2020). Computation Through Neural Population Dynamics. *Annual Review of Neuroscience*, *43*(1), 249–275. https://doi.org/10.1146/annurev-neuro-092619-094115

Yang, C. S., Cowan, N. J., & Haith, A. M. (2021). De novo learning versus adaptation of continuous control in a manual tracking task. *ELife*, *10*, e62578. https://doi.org/10.7554/eLife.62578

Yang, C. S., Cowan, N. J., & Haith, A. M. (2022). Control becomes habitual early on when learning a novel motor skill. *Journal of Neurophysiology*, *128*(5), 1278–1291. https://doi.org/10.1152/jn.00273.2022

Yewbrey, R., Mantziara, M., & Kornysheva, K. (2023). Cortical Patterns Shift from Sequence Feature Separation during Planning to Integration during Motor Execution. *Journal of Neuroscience*, *43*(10), 1742–1756. https://doi.org/10.1523/JNEUROSCI.1628-22.2023

Yokoi, A., & Diedrichsen, J. (2019). Neural Organization of Hierarchical Motor Sequence Representations in the Human Neocortex. *Neuron*, *103*(6), 1178-1190.e7. https://doi.org/10.1016/j.neuron.2019.06.017