

The Structure and Acquisition of Sensorimotor Maps

Floris T. van Vugt¹ and David J. Ostry^{1,2}

1913

Abstract

■ One of the puzzles of learning to talk or play a musical instrument is how we learn which movement produces a particular sound: an audiomotor map. Existing research has used mappings that are already well learned such as controlling a cursor using a computer mouse. By contrast, the acquisition of novel sensorimotor maps was studied by having participants learn arm movements to auditory targets. These sounds did not come from different directions but, like speech, were only distinguished by their frequencies. It is shown that learning involves forming not one but two maps: a point map connecting sensory targets with motor commands and an error map linking sensory errors to motor corrections. Learning a point map is possible even when targets never repeat. Thus, although

participants make errors, there is no opportunity to correct them because the target is different on every trial, and therefore learning cannot be driven by error correction. Furthermore, when the opportunity for error correction is provided, it is seen that acquiring error correction is itself a learning process that changes over time and results in an error map. In principle, the error map could be derived from the point map, but instead, these two maps are independently acquired and jointly enable sensorimotor control and learning. A computational model shows that this dual encoding is optimal and simulations based on this architecture predict that learning the two maps results in performance improvements comparable with those observed empirically. ■

INTRODUCTION

When first learning to talk or to play a musical instrument, a fundamental challenge is to learn which movement to use to produce a particular sound, that is, a sensorimotor map. There is a substantial understanding of how existing sensorimotor maps are adjusted in situations in which these maps are already well learned at the outset, such as when participants respond to visual rotations during reaching or auditory perturbations during speech (Huberdeau, Krakauer, & Haith, 2015; Krakauer, 2009; Houde & Jordan, 1998). These perturbations require participants to make relatively minor adjustments to previously acquired mappings (Telgen, Parvin, & Diedrichsen, 2014), which is reflected in the fact that performance reaches asymptote rapidly. However, such perturbation responses may not provide insight into the process by which maps are acquired in the first place. Here, a paradigm is presented in which participants acquire an entirely novel audiomotor map. This is a challenging task in which performance improvements are seen over several days of training. Previous work has shown that learning is possible under these conditions (Liu, Mosier, Mussa-Ivaldi, Casadio, & Scheidt, 2011; Mussa-Ivaldi, Casadio, Danziger, Mosier, & Scheidt, 2011; Radhakrishnan, Baker, & Jackson, 2008; Mosier, Scheidt, Acosta, & Mussa-Ivaldi, 2005), but it remains unclear what is the structure of the acquired maps. Initially,

participants face two problems: Given a target sound, they do not know where to move to, and when they make an error, they have no basis on which to correct it. Monitoring participants during learning allows insight into how sensorimotor maps are structured to solve these two problems.

To learn a motor skill, it is necessary to learn the sensory effects of one's movements. It is also necessary to be able to correct errors when they occur, and this requires a different type of knowledge, namely, of a mapping between sensory errors and motor corrections. In the present paradigm, both of these have to be learned. The former map is essentially a function from sensory output s to motor commands m (i.e., $f(s) = m$), which here is referred to as a "point map," whereas the latter is a mapping between sensory errors Δs and motor corrections Δm (i.e., $f'(\Delta s) = \Delta m$), which is called an "error map." Knowledge of this latter type is what is required to compute a proportionate correction of an error as is done in error correction models (Herzfeld, Vaswani, Marko, & Shadmehr, 2014; Shadmehr, Smith, & Krakauer, 2010; Thoroughman & Shadmehr, 2000; Ghahramani, Wolpert, & Jordan, 1997). In principle, the error map could be computed from the point map, because it is its mathematical derivative. That is, there could be one single map on which both movement selection and error correction depend. However, a computational architecture in which limited informational units are allocated to either a point map or an error map (Figure 1) suggests that reaching would be most accurate when informational units

¹McGill University, ²Haskins Laboratories, New Haven, CT

are divided (equally) between the two maps, instead of one. This architecture motivates a set of predictions, in particular, that two separate maps should exist: a point map and an error map. Simulations based on this architecture, which assume that learning involves the addition of units to the maps (see Model simulations), predict that point map learning should be observable as reduced error when participants make reaching movements to random targets, even if these targets never repeat. Furthermore, the simulations predict that error map learning should result in improved convergence onto targets that are presented repeatedly; in particular, both the rate of convergence and asymptotic performance should improve.

This study tested the hypothesis that participants acquire two separate maps, a point map and an error map, by monitoring reaching movements to auditory targets. Auditory feedback was presented at the end of each movement. Conditions were created in which participants either could apply error correction or could not. In a random presentation condition, which assessed the acquisition of a point map, targets were different on each trial, thus precluding error correction. To probe learning of

the error map, the same target was presented on multiple subsequent trials, thus allowing engagement of error-corrective processes. It was hypothesized that this error-corrective process itself has to be learned: The formation of an error map is a learning process. Moreover, error map learning was hypothesized to be independent from the acquisition of a point map.

METHODS

Participants and Experimental Tasks

Eighty-eight participants were recruited. All participants were right-handed (as verified by the Edinburgh Handedness Inventory) and had no or only minimal musical training. Participants reported no neurological or hearing impairments. Each session lasted approximately 1.5 hr. Participants provided written consent, and all procedures were approved by the McGill University institutional review board.

Participants made reaching movements to auditory targets while holding a robot handle. A 2-degree-of-freedom

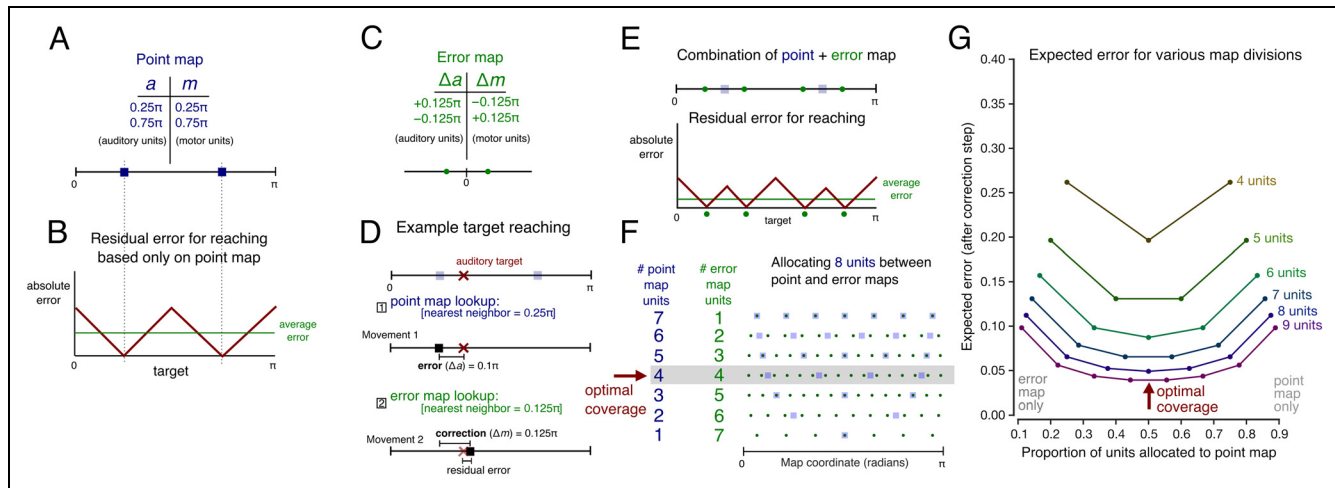
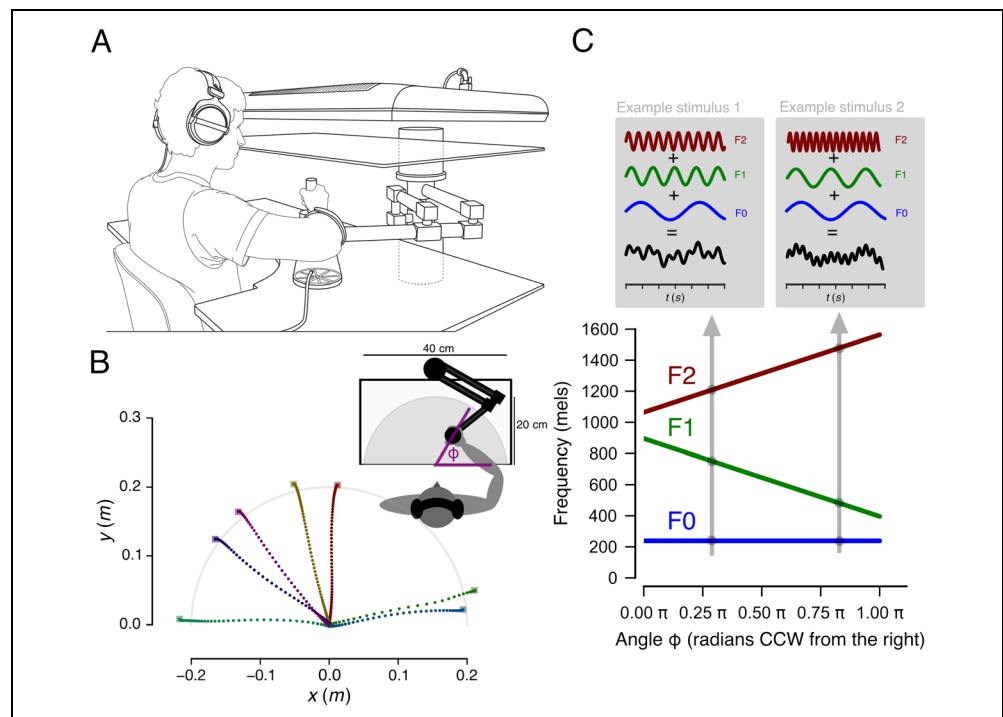


Figure 1. A computational architecture that demonstrates the optimal design of a sensory-motor controller includes both a point map and an error map instead of only one of the two. (A) Consider a controller that has access to a point map that is a lookup table of units connecting a particular movement with a sensory effect and an error map that connects sensory errors with motor corrections (C). Because we assume that the mapping between movements and sounds is one-on-one, we can, without loss of generality, represent these lookup tables as points on a motor/sensory space—here, for simplicity, a 1-D interval of movement angles $[0, \pi]$, which corresponds to those used in this empirical study. (B) When a target is presented, the point map is queried and returns the nearest neighbor entry. For the point map example above, this leads to a pattern of errors (red line) and, if the targets are equiprobable, an expected absolute error (green line). Note that the error varies with the distance between the target and the nearest neighbor. (D) On a given trial, first, an auditory target is presented. The controller selects the point map entry (out of two in this example, marked in blue) closest to the target and performs the corresponding movement (Movement 1). If the movement is not exactly on target, a nonzero error occurs, which is a difference vector that is then looked up in the error map (units not shown) resulting in a corrective movement (Movement 2). Although, by way of simplification, the controller aims to correct an error completely (gain = 1), because of the limited resolution of the error map, a residual error exists when the trial ends. (E) Example of a combined map that has two point map entries and a two-entry error map reproduced at each point map location. Note that there is a single error map that is not linked to particular targets but can be applied at any position in the workspace when error correction is possible. The red trace indicates the residual error for each target location given the above two-movement scheme (a point map movement and a corrective movement). (F) Adding more elements to either the point or error map increases performance as measured by a reduction in expected residual error, but if resources are limited, how can available units be best divided between these two maps? In the example case of 8 units, assuming the units are placed optimally on the maps, the best coverage of the space is achieved when the 8 units are divided equally between the point and error maps. (G) This is indeed true in general: The lowest cost (optimal division) is achieved for an equal split of units between a point map and an error map. This leads to the prediction that two maps encode information more efficiently than one. Although this is a simplified architecture, relaxing its assumptions to make it more realistic (such as the addition of noise in reaching or map lookup) will not alter the computational argument that dual encoding provides optimal spatial coverage.

Figure 2. Participants make movements to auditory targets. Before each trial, a target sound is presented. Participants then make a reaching movement to the presumed location of the sound. (A) Side view of the workspace. Participants made reaching movements from the center of the workspace to points on a semicircle (not visible to participants). When participants had stopped moving, auditory feedback was presented that depended on the angle between the starting point and the end of the movement. The amplitude of the movement was inconsequential to the sound. No acoustic location cues were present in the sounds because the same signal was presented to both ears. (B) Example movements for one participant. Square markers indicate the final hand position, whose angle relative to the starting point (ϕ) determines the feedback sound. For reference, the target circle is indicated in gray but not seen by the participant. (C) The sounds that were played to participants at the end of each movement consisted of three oscillators whose frequencies depended linearly on the angle at movement end ϕ (see gray inlays for a decomposition of two example stimuli that correspond to the movement angles indicated by the arrows). To ensure perceptual uniformity of the auditory space, we used the mel scale for sound frequency.



planar robotic arm (InMotion2; Interactive Motion Technologies, Watertown, MA) was used (Figure 2A) that sampled the position of the handle at 400 Hz. First, in a calibration phase, participants were asked to hold the robot handle in front of their body midline. The lateral coordinate of this position was captured and used as the movement start point throughout the experiment. A target circle was then defined as a half-circle around this midpoint, and during the experiment, participants made movements from the start point to points on this circle (Figure 2B). This circle was never shown visually, but participants were shown a schematic drawing of it before the experiment, and the robot demonstrated points on this circle (at 0° , 90° , and 180° counterclockwise from the right) by moving the participants' hand to them. Vision of the circle and the arm was blocked. All experiments followed the same schedule on a testing day and differed only in number of days and the way the auditory targets were chosen.

Auditory Stimuli

The sounds (target and feedback) consisted of three sine wave oscillators: one with fixed frequency (F0, 165 Hz) and the frequency of the other two signals (F1, F2) linearly decreased or increased, respectively, with the angle of the movement end point (Figure 2C). The frequency range of the F1 and F2 oscillators correspond to the first and second formant frequencies of vowel

sounds (Remez, Rubin, Pisoni, & Carrell, 1981). These particular signals were chosen because they provide a rich yet learnable stimulus that participants are already familiar with by virtue of its structural similarity to vowels. These more complex stimuli limit the benefit to performance that might arise because of prior expectations, which could readily have occurred had auditory stimuli been used that varied monotonically over the work space. To normalize the space-to-frequency mapping, angles were mapped linearly to frequencies in mel space, which is an interval scale encoding of frequency differences (Stevens, 1937). To correct for perceived loudness differences, the amplitude of each oscillator was adjusted using equal loudness curves (Robinson & Dadson, 1956) to 75 phons (so that each sound would be perceived to be as loud as a 1-kHz tone of 75 dB). In this way, equal displacements in the motor space cause a perceptually equal change in sound frequency and little or no change in sound intensity. Sounds were presented over Beyerdynamic (Heilbronn, Germany) DT770M headphones. This mapping between positions and sounds remained the same throughout all experiments. Because the mapping is one-to-one, without ambiguity, the sounds will be referred to by the angle they were mapped to on the interval of angles $[0, \pi]$ radians. The sounds themselves were not localized in space: The same sound was presented to both ears, hence there were no acoustic location cues, and only their frequency content contained information about the position they were mapped to.

Audiomotor Training

In audiomotor training trials, participants received a target sound and were instructed to move to the location that corresponded to the sound. At the start of each trial, the robot returned the handle to the starting point using a minimum-jerk trajectory and held it there for 500 msec. Then, the target sound was presented for 1000 msec, during which time the forces holding the participant's hand in place were gradually decreased, leaving the participant free to move when the sound ended. Movement onset was defined as the moment when handle was more than 5 cm away from the starting position, and movement end was when the velocity fell below 5% of peak velocity for 50 msec. At movement end, the robot held the participant hand at its current position using an attractor controller. The sound corresponding to the angle between the starting point and the end position was presented for 1000 msec. The amplitude of participants' movement was inconsequential to the sound. Other than the sound, no feedback was provided.

Movement Copy

Movement-copy trials served to familiarize the participant with the size of the target circle (because it was not shown visually) and to measure baseline motor accuracy in reaching toward particular directions. During each trial, the robot first brought the handle back to the starting position. Then, the hand was moved out to a target position on the target circle in a minimum-jerk trajectory of 900 msec, held there for 500 msec, and then moved back. A visual icon "MOVE" appeared that signaled participants to move to the indicated target location in a single, straight, swift movement. No forces were applied by the robot during the participant's movement. Movement onset and end were detected as in the audiomotor trials above. Target directions were equally spaced at 10% of the half-circle, including the end points, yielding 11 target locations presented in random order.

Audiomotor No-Feedback Trials

Audiomotor no-feedback trials were obtained before and after the audiomotor training sequence. These trials were identical to training trials except that no feedback sound was presented. Ten trials were administered before and after the training phase (Figure 2D).

Auditory Psychophysical Testing

Auditory psychophysical testing was completed away from the robot with the participant seated in front of a computer. On every trial, a train of four sounds of 200-msec duration each was presented with a 75-msec pause between sounds. Three of the sounds were identical, and one (either the second or third) was different. Parti-

cipants' task was to respond by pressing a button whether the mismatched sound was the second or third in the sequence. The three identical sounds were those that were mapped to the 0.5π angle in our auditory-motor mapping, and the mismatched sound corresponded to the angle $0.5\pi + \delta x$ where δx was 1 of 10 logarithmically spaced values between 0.0015π and 0.09π . Participants completed 200 trials (10 stimulus levels in each direction \times 10 repetitions) in blocks of 20 with a short break in between. No feedback was given about the accuracy of the participants' response.

To obtain auditory thresholds, psychophysical curves were fitted to the data offline using maximum likelihood. The psychophysical curves were sigmoid functions defined by the formula where p is the probability of giving the correct response, x is the stimulus level (angle that defines the sound), and erf is the mathematical error function. The fit parameters are m (curve midpoint), s (parameter controlling the slope), and λ (lapse rate). Psychometric curves were discarded when their fit was poor ($R^2 < .5$).

Experiments

Random Continuous Targets (Continuous-1d)

Auditory targets were chosen from a continuous uniform random distribution on the interval $[0, \pi]$ (in radians) with the constraint that the angular distance between the target on trial $n + 1$ and the movement end point of trial n was at least 0.3π radians. The rationale for this constraint originated from a pilot study observation that participants were more accurate on trials whose target was close to their previous movement end point, which allowed them to achieve better-than-chance performance without knowledge of the audiomotor mapping. The minimum distance requirements prevented participants from using such a strategy. Eighteen participants (10 women) were aged $21.4 (\pm 3.6)$ years. Participants were tested for one session.

Discrete Targets (Five Targets, 1 Day)

For each participant, a set of five unique targets was selected, and on each trial, one of these five was presented, with the constraints that (1) the same target could not be repeated on consecutive trials and (2) the angular distance between the target on trial $n + 1$ and the movement end point of trial n was at least 0.3π . The set of targets was determined randomly as follows: Five angles were placed equidistant on the interval $[0.05\pi, 0.95\pi]$ (including end points) and then jittered with zero-centered normal noise with an SD of 0.02π . When the absolute value of the jitter for any of the targets exceeded 0.1π in magnitude, then the noise vector was recomputed. Eighteen participants (12 women) were aged $22.4 (\pm 2.5)$ years. Participants were tested for one session.

Discrete Targets (Five Targets, 3 Days)

Target placement was identical to the “5-targets-1d” experiment except that the minimum angular distance between end points and subsequent targets was set to 0.05π to achieve greater uniformity of the target distribution. The set of targets was determined as during the 5-targets-1d experiment. Sixteen participants (10 women) were aged $22.1 (\pm 3.5)$ years. Participants were tested on 3 days with 2- to 4-day intervals in between.

Repeated Targets (1 Day)

Trials were divided into batches of 16 trials on which the same target sound was presented repeatedly. Every 16 trials, a new target was selected from a continuous uniform distribution on the interval $[0, \pi]$ with the constraint that the target had to be a minimum distance of 0.2π radians away from the last trial (previous batch) movement end point. Twenty-one participants (14 women) were aged $23.6 (\pm 4.2)$ years. Participants were tested on 1 day. For analysis, data from the first day of the 3-day experiment were included because the design was identical.

Repeated Targets (3 Days)

Target selection was identical to that of the 1-day experiment. Fifteen participants (10 women) were aged $23.9 (\pm 3.4)$ years. Participants were tested on 3 days with 2- to 4-day intervals in between.

Repeated-Targets Analysis

In the experiments involving repeated-target presentation, learning curves were fitted to individual participants. For a group of trials with the same target, we averaged the absolute angular error as a function of trial within batch and fitted a learning curve defined by the equation $e(t) = (a - b)\exp(-t/tc) + b$, where $e(t)$ represents the angular error on within-batch trial t and a is the intercept of the learning curve, b is the asymptote, and tc is the time constant. To enforce a and b to be non-negative, we transform them into pseudo-parameters using the exponential function.

Lag-1 autocorrelation (ACF1) was calculated as the Pearson correlation between the vector of movement end points and the vector of movements shifted by one in chronological order. We skipped the four initial movements to each target because these generally had larger errors than later trials and might have disproportionate effects on the correlation.

Statistics

Unless otherwise specified, we computed linear mixed models statistics with participant as a random factor and maximal random structure.

RESULTS

Participants made reaching movements to auditory targets that were presented over headphones (Figure 2). The same signal was presented to both ears, and therefore no acoustic location cues were given; however, the frequency content of the sounds depended on the angle of the movement. This situation is directly analogous to learning to speak, where the position of articulators (e.g., tongue) changes the frequencies of the produced sound, not the perceived physical location of the sound itself. As in learning to speak, the participants have to learn the appropriate movement direction to produce a particular combination of frequencies.

Model Simulations

Simulations were run based on the architecture introduced previously (Figure 1) to find predicted indices of the acquisition of point and error maps, respectively. Maps were modeled as lookup tables, which encode input–output pairs without making assumptions about the overall structure of the maps. The simulations assume that the learning process is composed of the addition of units to the point and/or error maps. Learning is assessed in terms of changes in predicted reaching accuracy. Alternative model frameworks (Pouget & Snyder, 2000; Jordan & Rumelhart, 1992) may be used to account for the data here, but here, a simple model was opted for that could account for the data with only minimal assumptions. Details about the model are included in the supplementary materials (<https://doi.org/10.6084/m9.figshare.5527963.v1>).

We simulated “participants” with various numbers of units in point and error maps, respectively, and investigated the effects of map density on performance. In a random target condition (Figure 3A) on each trial, a new target was selected at random and then its nearest neighbor was looked up in the model’s point map. Gaussian noise ($\sigma = 0.05\pi$) was added to both auditory input and motor output. As point map resolution increases (larger number of entries in a lookup table), the simulations predict a reduction in error (Figure 3A). The model presented here is not a model of learning, because it does not specify in what locations in the map units are added, but it can be used to extract predicted changes in performance as a result of increases in map density. For simplicity, the simulations assumed that units are added incrementally in random locations (Figure 3A, green trace). For reference, we also performed separate simulations for each point map density in isolation, where map units were put in optimal locations (Figure 3A, purple trace). Note that, in these latter simulations, units were not added to the map incrementally in any one simulation, but separate simulations were run for each map density.

To investigate the effect of error map learning, the point map was held constant (one unit), whereas the error map density was varied. Error correction was enabled in these

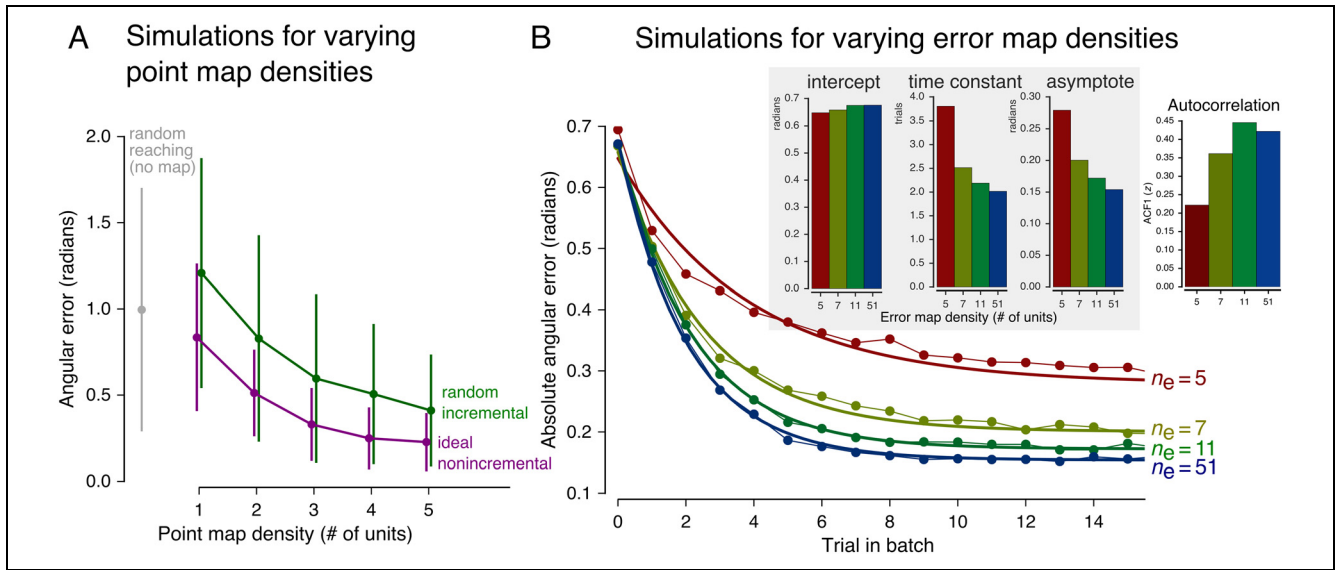


Figure 3. Model simulations show that adding units to the point map (n_p) results in a reduced error in reaching to random targets and adding units to the error map (n_e) results in improved convergence to repeating targets. (A) Random-target simulations reveal that increasing point map resolution (adding units to the point map) yields a decreased angular error in reaching to random targets. Simulations were performed by incrementally adding units to the point map in random locations (green trace). For reference, the purple trace indicates model performance when the same number of units is placed in optimal locations (which are different for different numbers of map units), and random reaching performance is shown in gray. Error bars indicate SD. (B) Increased error map resolution (adding units to the error map) results in improved convergence to the target, as measured by a smaller time constant and lower asymptotic performance, whereas the intercept of the learning curve did not decrease (see inlay barplots). In addition, as error map density increases, the trial-to-trial autocorrelation of movements to the same target increases. Dots represent the data points, and lines represent the fitted learning curves.

simulations by presenting the same target repeatedly for 16 trials. On the first trial of a batch, because the target is novel, the model performs a lookup in its point map. On a subsequent trial n in a batch ($n > 1$), the model aims to correct a proportion ($\eta = 0.3$) of the (signed) error experienced on the previous trial, by looking up the closest entry to $\eta \times e_{n-1}$ in its error map (lookup table) and applying that correction to the previous movement u_{n-1} . Gaussian noise ($\sigma = 0.05\pi$) is added to movement output and auditory inputs. To the average absolute error within batches, we fit the learning function $e(t) = (a - b)\exp(-t/tc) + b$, where $e(t)$ is the error on trial t within the batch, b is the asymptotic performance (at the end of the batch), a represents the intercept, and tc is the time constant. We varied the number of entries in the error map and observed that, with increasing error map resolution, convergence to the target is more rapid and asymptotic performance is improved (Figure 3B). Absolute errors are reported for the model as they are for the human participants so that both data sets are processed identically.

In summary, as units are added to the point map (point map learning), there should be a reduction in error when reaching to random targets (Figure 3A). As units are added to the error map (error map learning), there should be improved convergence when the same target is repeated for a number of consecutive trials, in particular, a faster rate of convergence and a lower asymptotic error (Figure 3B). Why would the time constant of learning curves change? The reason is that, although

on a given trial the model always aimed to correct a given proportion of the error, it would be less successful in generating exactly that correction when error map density is low and more successful when error map density is high. Similarly, the asymptote of the learning curve decreased with increasing error map density because a more dense error map would allow the encoding of smaller errors and enable the model to correct for them.

Learning a Point Map

Continuous Targets

To study point map acquisition empirically, a different target was presented on each trial. When targets are different on subsequent trials, the correction appropriate for an error on one trial cannot be applied directly to the next trial because it is a movement to a different target. Participants made reaching movements toward auditory targets, corresponding to angular positions chosen uniformly from the half-circle workspace (because the audiomotor mapping is one-to-one, we will refer to the movement angles and corresponding sound positions as points on the workspace $[0, \pi]$ without ambiguity). A decrease in absolute angular error was observed (angular distance between the auditory target location and the movement end point) over the course of audiomotor training (Figure 4A). Before and after learning, participants completed trials where no auditory feedback was presented. These trials probed the point map exclusively

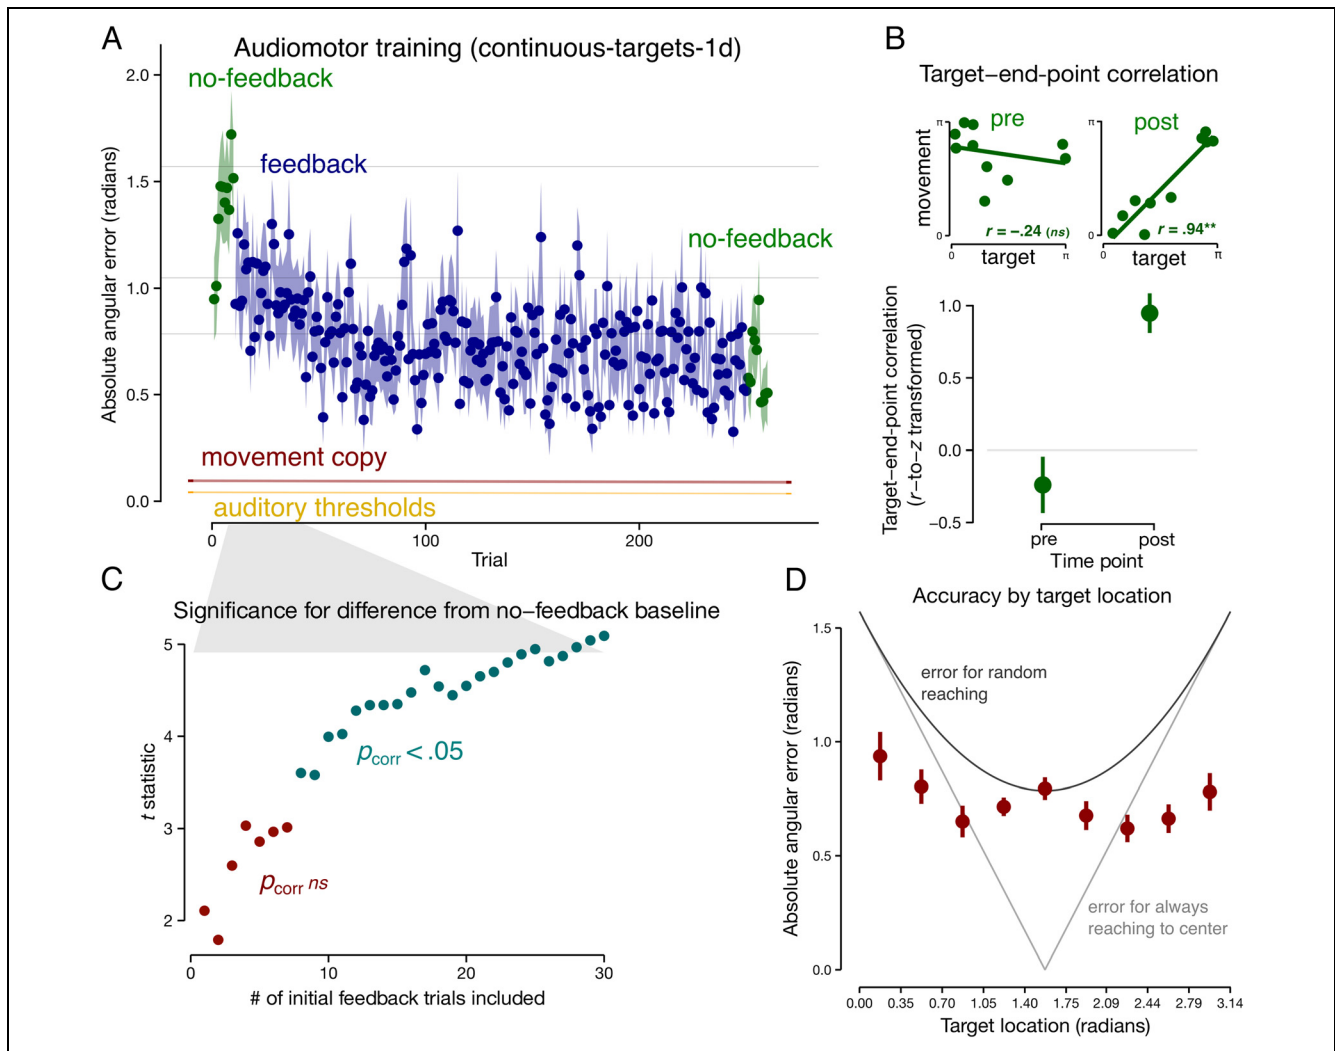


Figure 4. An audiomotor point map is acquired even when participants make movements to auditory targets that are never repeated. (A) Absolute angular error between the target and the movement end point is shown to improve as a function of trial (continuous-1d experiment). Dots indicate means across participants, and the shaded area indicates the SEM. Performance on auditory (yellow) and motor (red) tests remained stable and are shown for reference, translated into motor coordinates (in radians, mean and SEM; the variability is so small that it appears as a line). (B) Correlations between the target and the movement angle for a representative participant indicate that, at the outset, there was no prior information about the mapping (no correlation; top left), whereas after learning, movement directions vary systematically with auditory target location (top right). This was true for participants overall (main panel; data *r*-to-*z* Fisher-transformed; error bars indicate SEM). (C) To study at what point over the course of training the movements became different from chance reaching, we took the initial *n* trials with feedback, and computed a *t* test against the no-feedback pretraining trials. Color coding indicates that movement error reduction reaches significance after eight movements, indicating very rapid learning. (D) Absolute angular error for various bins of target location (in radians), indicating that participants' performance deviated from chance error levels that would be expected if participants would always reach to the center (gray; the theoretically optimal strategy if the mapping was unknown to them) or if participants would reach randomly (in black).

because error correction could not operate in the absence of feedback. Performance on no-feedback trials after training was significantly improved relative to before training ($F(1, 15) = 58.39, p < .00001$).

To assess whether this improvement was attributable to the mapping instead of perceptual or motor learning in isolation, measures of auditory perceptual and motor learning performance were added as covariates to the above analysis, specifically (1) the auditory psychophysical threshold (curve midpoint) before versus after learning and (2) the motor-copy error (trials in which we measured absolute angular error when participants reproduced movements

indicated by the robot; Figure 4A). Performance on these tests was not significant as predictors of the audiomotor error ($F(1, 57.9) = 1.00, p = .32$, for auditory thresholds, and $F(1, 218.19) = 2.63, p = .11$, for motor-copy trials; see also Supplementary Figure S1, <https://doi.org/10.6084/m9.figshare.5527963.v1>). Furthermore, motor-copy angular absolute error was not different after training relative to before ($F(1, 15) = 0.64, p = .43$) nor were auditory psychophysical thresholds ($F(1, 15) = 3.18, p = .09$). That is, evidence for learning was obtained in the audiomotor task that cannot be explained by motor learning or auditory perceptual learning in isolation and therefore can be attributed

to learning of the mapping between movements and sounds. The same was true for the other experiments reported below (see Supplementary materials A.2; <https://doi.org/10.6084/m9.figshare.5527963.v1>).

Did participants really learn the mapping from scratch? For individual participants, the (Spearman) correlation between target angle and movement angle was computed for the no-feedback trials before and after audiomotor training. The *r*-to-*z*-transformed correlation coefficients were not significantly different from zero before training ($t(15) = -1.23, p = .24$), indicating that participants' initial reaching was at chance level and they had no knowledge of the mapping. Correlation coefficients were significantly higher after the audiomotor training relative to before ($F(1, 15) = 43.03, p < .0001$; Figure 4B), confirming that participants learned the mapping. To gain insight into the earliest stages of learning of the audiomotor mapping, initial segments of audiomotor training trials were compared statistically with the before-training no-feedback trials using *t* tests, finding that, after as few as eight trials, the error on audiomotor trials was better than the no-feedback chance levels (Bonferroni-corrected $p < .05$; Figure 4C). This indicated that rudimentary knowledge of the mapping was acquired rapidly by par-

ticipants. Investigating how participants' errors varied across the workspace, the targets were divided into bins and the pattern of errors as a function of target bin was found to be different from those that would be expected if participants had no knowledge of the audiomotor mapping (Figure 4D).

In summary, although the same target was never repeated, participants remarkably showed the capacity to learn the mapping, and the learning was already evident within a handful of trials.

Five-Targets Random Presentation

It is unclear how much detail participants acquired about the audiomotor mapping. Here, the acquired point map was probed by having participants reach to a set of five targets, spaced equally across the workspace with a small jitter and presented in random order. The no-feedback trials showed a significant decrease in error after audiomotor training relative to before ($F(1, 322) = 64.33, p < .0001$; Figure 5A). The error for the initial audiomotor training trials was significantly lower than the preceding no-feedback trials after nine movements ($p < .05$, Bonferroni-corrected *t* tests) and onward. During the later trials

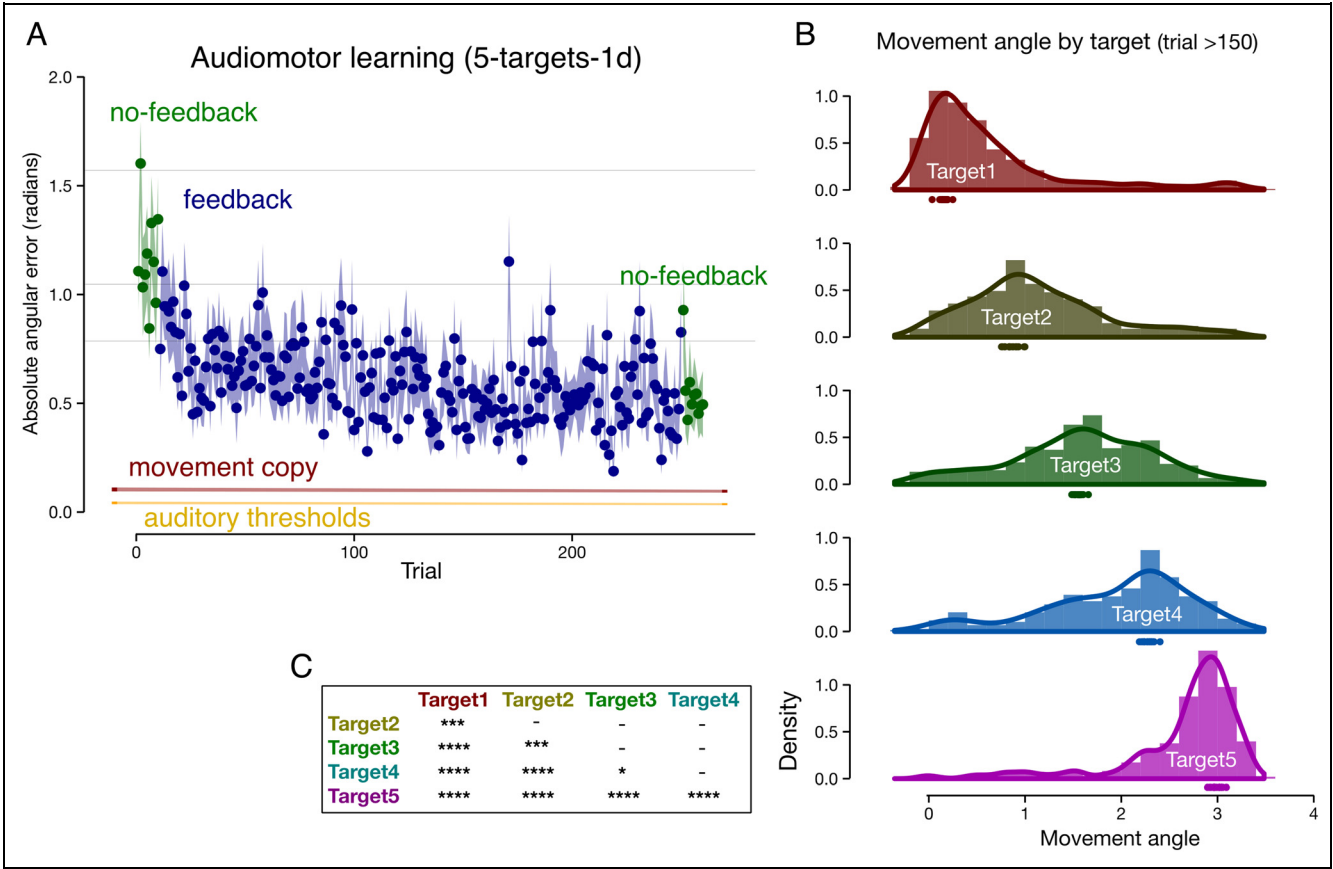


Figure 5. The acquired audiomotor point map encodes multiple targets (five or more). (A) Angular error over time (5-targets-1d). (B) Movement angles as a function of the target presented for all participants combined (targets are numbered in counterclockwise order starting from the right). Actual locations of the targets varied between participants and are indicated by dots below the histograms. Lines indicate kernel density estimates. (C) Pairwise comparisons of the reaching end points for the individual targets were all significant (*t* tests with Bonferroni-corrected *p* values: **p* < .05, ***p* < .01, ****p* < .001, *****p* < .0001).

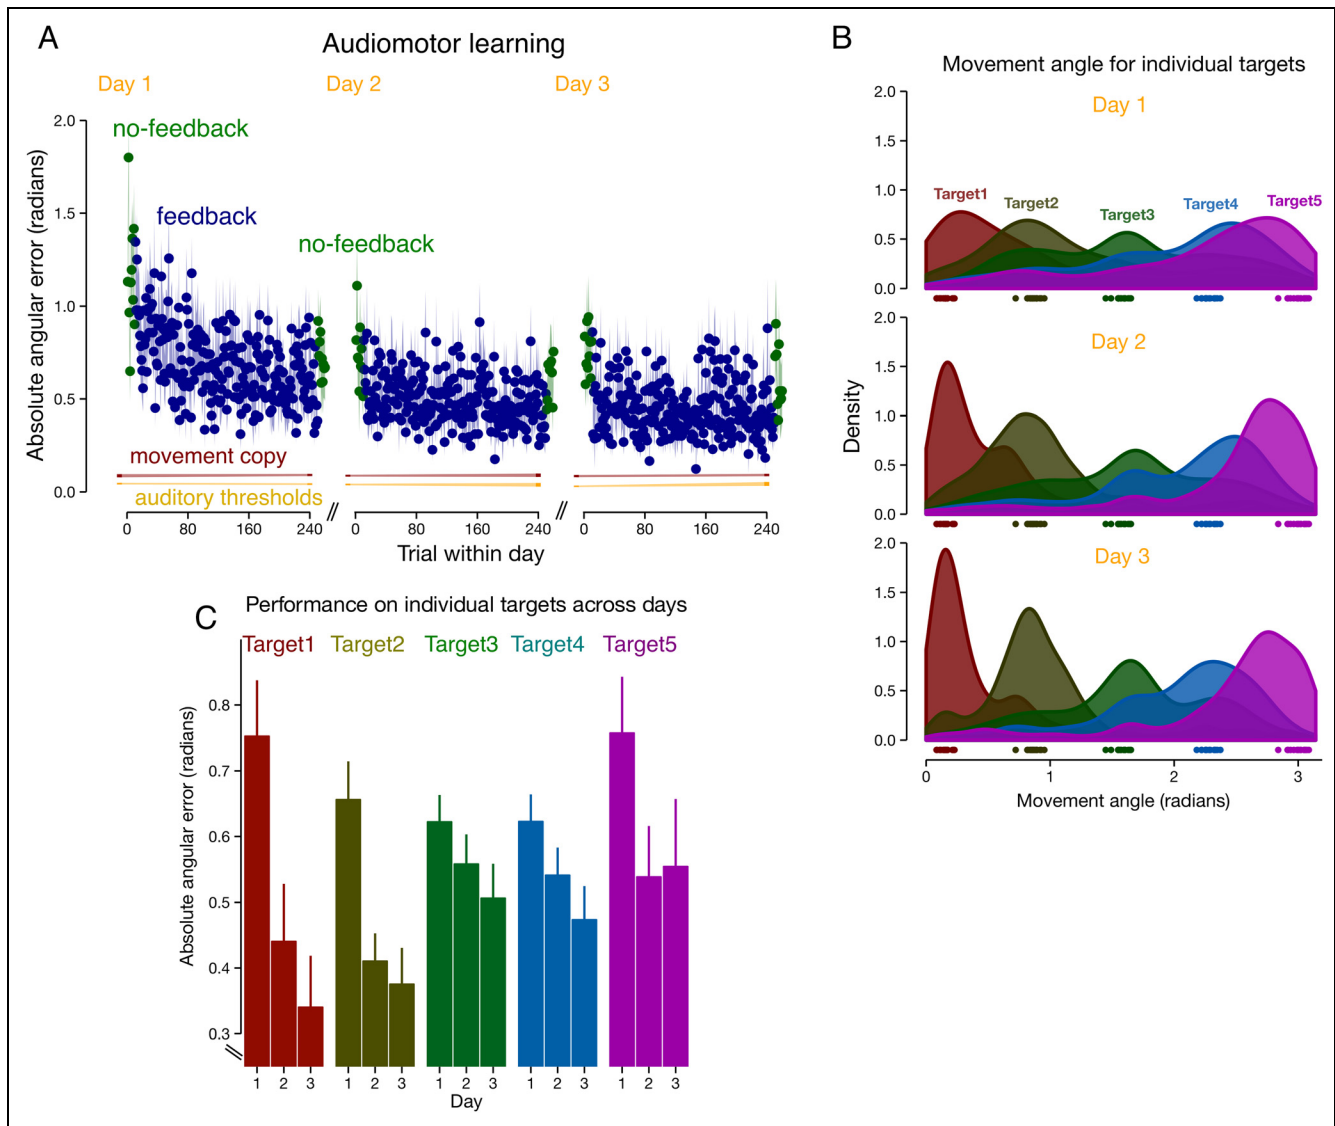


Figure 6. The audiomotor map is retained, and learning continues across days (5-targets-3d experiment). (A) Mean absolute error across trials on the 3 days. (B) Density estimation of movements to each of the five targets for all participants combined shows that movements become progressively more clustered around actual target locations (indicated with dots below the horizontal axes). (C) Average absolute angular error for the individual targets across the 3 days.

(Trial 150 and onward), the end points of movements toward the five individual targets were significantly segregated ($F(3.18, 50.82) = 105.23, p < .0001$, main effect of target; Figure 5B). Planned contrasts using Bonferroni-corrected t tests with pooled SD s revealed that, for all pairs of targets, the movement end points were significantly different from one another (Figure 5C). This shows that the point map acquired by participants in a single day contained sufficient information to encode multiple (at least five) distinct targets.

Five Targets, 3 Days

To study whether learning was retained and could continue to improve across days, participants were measured on 3 testing days with 2–4 intervening nontesting days.

Within-day reaching accuracy improvements on no-feedback trials differed over the 3 testing days ($F(2, 866.00) = 4.52, p = .01$). Planned contrasts revealed a reduction in error after training relative to before on the first day of training ($p < .001$) but not on Day 2 ($p = .27$) or Day 3 ($p = .24$; Figure 6A). There was also a reduction in before-training errors from Days 1 to 2 ($p < .001$) but not from Days 2 to 3 ($p = 1$). In line with the results of the 5-targets-1d experiment, reaches toward individual targets were significantly different for the first day and continued to be so for the subsequent days as evidenced by a tighter clustering of movements around actual target locations (Figure 6B). In support of this observation, the angular absolute error did not differ between targets ($F(4, 14.09) = 1.2, p = .35$) but decreased over days ($F(2, 18.10) = 37.85, p < .0001$; Figure 6C).

Planned contrasts showed a reduction in error between Days 1 and 2 ($p < .0001$) and between Days 2 and 3 ($p = .03$). In summary, learning is retained across days, and improvements continue to occur into the third day of training, revealing that the audiomotor map that is learned is persistent.

Learning an Error Map

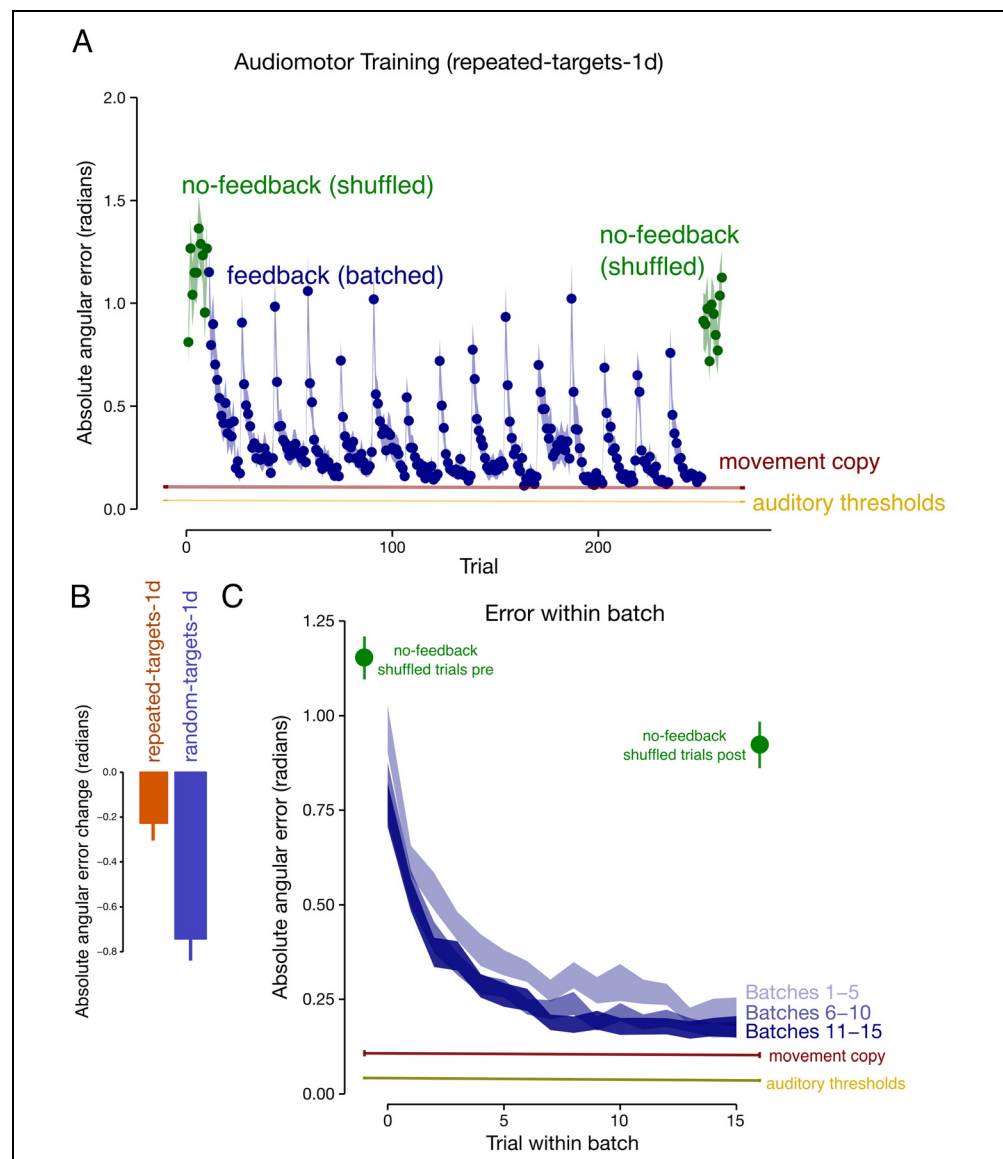
In the next series of experiments, the same target was repeated multiple times in a row (16 trials). This allowed participants to correct an error experienced on a given trial. Each new target sound was selected randomly and was different from all prior targets. In this process, it was hypothesized that participants would collect this error-corrective information and form an error map. We specifically test the hypothesis that error map formation is itself a learning process. Evidence for learning an error

map would not simply be convergence to the target but rather an improved capacity to correct errors across multiple targets. This improved convergence should be seen as a faster rate of learning and a better asymptotic performance.

Targets Presented Repeatedly (1 Day)

Reaching behavior was markedly different (Figure 7A) from the random presentation experiments (cf. Figures 4 and 5) with large errors in the beginning of a batch, which decreased rapidly until reaching an asymptote (Figure 7C). Error levels reached toward the end of a batch were much lower than those achieved during random presentation (cf. Figures 4A and 5A). On the basis of simulations, it was hypothesized that error map learning should be accompanied by a reduction in asymptotic performance, which is tested here below. However,

Figure 7. When error correction is made possible, convergence to targets is observed, and this convergence itself changes over time, which is an evidence for error map acquisition. (A) Audiomotor training trials indicate a sharp decline in error over the course of repeated presentations of a single target. (B) However, point map learning is less when targets are repeated than when they are random on every trial (cf. Figure 4), as shown by average reduction in angular absolute error during the no-feedback trials. (C) Average errors in groups of five batches. A sharp initial decline in error ends in asymptotic performance that is close to the movement-copy error (in red). Green dots indicate average error on no-feedback trials before and after audiomotor training.



point map learning may also have occurred in parallel, with the addition of each new target, and this point map learning could potentially have resulted in improvements in asymptotic performance. To verify that the improvements observed here are due to error map learning and not point map learning, we test for improvements in asymptotic performance while taking the first trial in each batch as a covariate. To quantify asymptotic performance, the absolute angular error was averaged for the last five repetitions of a target, and it was found that this error decreased in later batches ($F(1, 446.47) = 11.19, p = .0009$), as predicted by error map learning. This was true even when accounting for the reduction of error on the first trial to each new target as a covariate ($F(1, 470.96) = 1.07, p = .30$). Hence, the data indicate error map learning that could not be accounted for by improvements in performance because of point map learning.

Point map information was acquired as well during this condition, but to a lesser extent than seen in the random condition. The first trial for each new target was essentially a random trial comparable with the random experiments where participants only performed such trials. A statistical trend for a reduction in error of this first trial over time was observed ($F(1, 25.95) = 3.70, p = .07$). The no-feedback trials, another measure of point map learning, before and after training showed significant improvement ($F(1, 605.27) = 12.22, p = .0005$; Figure 7A), but this reduction in error in the no-feedback trials was significantly smaller than that observed in the “continuous-1d” experiment ($F(1, 930) = 25.28, p < .0001$). The difference in improvement in no-feedback trials in repeated and random presentations provides evidence against the possibility that participants learn a single map that gathers all sensorimotor information, because participants in both conditions made the same number of movements and received the same number of auditory feedback trials (Figure 7B).

In summary, performance improvements obtained during repeated presentations of the same target were not carried over to the point map as probed in no-feedback trials, which, to the contrary, revealed less learning than in the random experiments. Crucially, improvements in convergence were observed, that is, improvements in error-corrective behavior that could not be explained by an improvement in the point map, suggesting the existence of a separate error map that is independent of the point map.

Targets Presented Repeatedly (3 Days)

Another set of participants performed the repeated target experiment on each of the 3 days. In support of the idea that there is error map learning, it was observed that the error decreased more rapidly on later days, reaching lower asymptotes (Figure 8B). Learning curves were fitted to the error as a function of trial within each group of five batches (to obtain a sufficient signal-to-noise ratio).

The intercept of the learning curves showed a reduction over time ($F(1, 15.41) = 5.79, p = .029$), indicative of point map learning. Importantly, convergence onto the target was improved, as reflected in a lower asymptote ($F(2, 87.70) = 3.94, p = .02$) and decreased time constant of the learning curve (the slope; $F(2, 25.62) = 9.23, p = .001$); both of these were computed taking into account changes in curve intercept (Figure 8B). This improvement (from Days 1 to 3) could not be explained as a change in motor-copy error ($t = 1.09, p = .31$) or auditory threshold ($t = 1.49, p = .17$; Figure S1). The improvements in performance observed here empirically with repeated targets across days are consistent with those observed during simulations with increasing error map density.

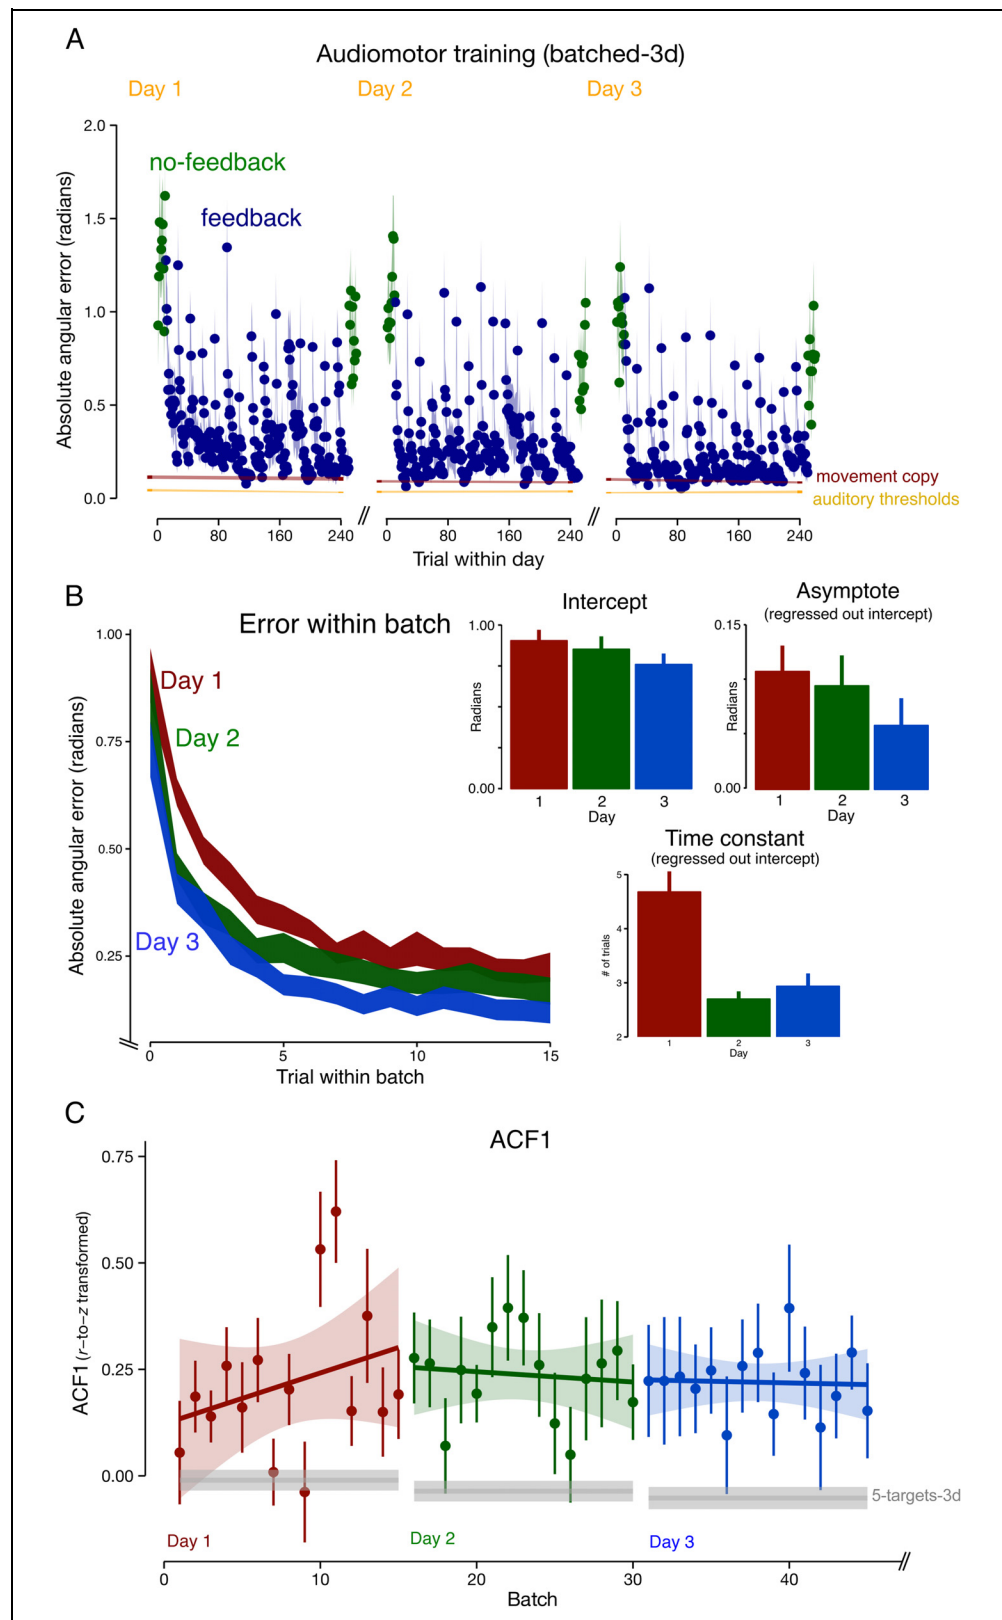
Again, point map information was also acquired as shown by performance on no-feedback trials. Performance in the no-feedback trials improved on all 3 days after audiomotor training (all $ps < .02$) relative to before. In addition, the error on the first trial for each new target was reduced over time ($F(1, 295.69) = 10.52, p = .001$).

The improvement in learning curve time constant (and asymptote) reflected an increased ability for convergence within the auditory–motor workspace. This was hypothesized to be due to the formation of an error map that enabled more accurate error correction. To test this idea, trial-by-trial correlations of the movements (ACF1) were investigated, which are typically observed in error-based learning models as well as in the simulations reported here (Figure 3B). As shown in Figure 8C, the autocorrelation increased during the first day and then remained at nonzero levels during Days 2 and 3. The absence of autocorrelation at the beginning of the first day suggests that there was little systematic use of error information to converge to the target. The autocorrelation then increased over the course of the first day ($t = 1.87, p = .06$), suggesting more gradual adjustment of movements based on errors. This nonzero autocorrelation was then maintained during the second and third days ($p < .001$ and $p = .025$, respectively). For comparison, we extracted, for each target separately, the chronological series of movements in the 5-targets-3d experiment (ignoring intervening reaches toward different targets). ACF1 was calculated in windows of 16 movements and then averaged to ensure compatibility with the repeated-target 3-day data by equating biases in short time series autocorrelation (Marriott & Pope, 1954). If the same learning process would operate in these experiments, one would expect to see similar autocorrelations. Contrary to this, a zero ACF1 was found in the 5-targets-3d experiment (Figure 8C).

Learning a Point Map and an Error Map Are Independent

The preceding section showed that error map acquisition cannot be explained by point map learning (Figure 8B). However, could error map learning contribute to point map acquisition? If error map learning contributed to

Figure 8. Error map learning is itself a learning process, observed as improved convergence onto the target across days. (A) Error as a function of trial across days. (B) Error as a function of repeated presentation of the same target across the 3 days. The barplot indicates that the time constant decreases (rate of convergence onto the target increases) and the asymptote decreases, and this cannot be explained by the intercept of learning curves (see barplots). (C) ACF1 calculated for each new target increases over the course of the first day (red) and then remains stable at a small but nonzero level for the following days (green and blue). For random trials, ACF1 remains at zero (in gray; calculated over chronologically ordered series of 16 movements to the same target in the 5-targets-3d experiment, omitting intermediate trials to different targets), suggesting that a different learning process operates in the repeated- and random-target studies.



point map learning, then repeated trials to the same target should contribute toward the acquisition of the point map. To test this idea, the error on the first trial (“random” trial) of every new target was tracked in the repeated-targets experiment (45 trials) and compared

with participants in the continuous target experiment who were tested with random trials only (Figure 9A). If only the first trial in each batch contributes to point map formation, then both groups should show similar learning. Indeed, it was found that both groups improve ($F(1,$

29.03) = 9.96, $p = .004$), but there was no difference between the groups or between the slopes ($F(1, 29.03) = 0.01$, $p = .94$), despite the fact that participants in the repeated-targets experiment performed 15 times more movements (and therefore received substantially more exposure to the auditory-motor workspace). This argues against the idea that participants learn a single mapping and instead shows that repeated movements to the same targets do not contribute to point map learning but rather that participants are engaged in a different type of learning contributing to an error map.

One alternative explanation could be that, if participants in the repeated-targets 3-day experiment would move to roughly the same direction for each presenta-

tion of the same target, then they would obtain little additional information about the mapping beyond the first movement to that target. To test whether this could explain the above finding, the spatial coverage of the set of movements (including the repeated-targets trials > 2) was computed, where more negative values indicate greater spatial coverage (for details about spatial coverage, see Supplementary Section S6, <https://doi.org/10.6084/m9.figshare.5527963.v1>). It is seen that the repeated-target participants' movements achieved greater spatial coverage than the random participants' movements, thus invalidating this potential confound (Figure 9B).

Separate Acquisition of Point and Error Maps

Point map acquisition was shown by improved reaching toward random targets. Error map formation was shown by improved convergence across days when targets were presented repeatedly. In neither version of the task could improved performance be reduced to motor or perceptual learning in isolation, because these remained stable, and therefore learning was attributable to the formation of a mapping between them. In principle, the error map could be computed as the spatial derivative of the point map. In other words, participants would learn only a single map. However, summarized below are the observations that suggest that this is not the case here.

First, repeated movements to the same targets did not contribute to map learning over and above the first movement to each new target. That is, performance on

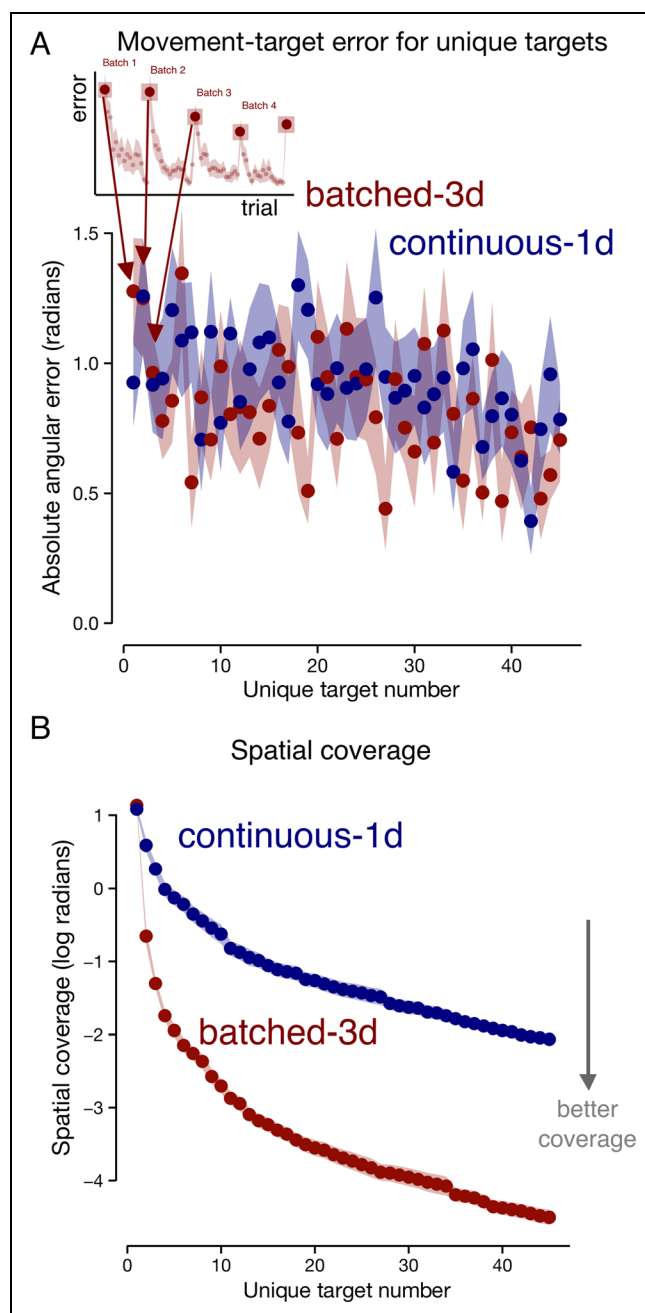


Figure 9. Point and error map learning are independent. (A) Repeated movements to the same target do not contribute to learning a point map. Independence of the mappings is indicated by plotting the error on the first trial of every new target (red; trial selection is illustrated in the inset figure where green markers indicate the first trial of every new target, which is shown in the main panel) against the errors that occurred when targets were selected randomly and never repeated (continuous-1d experiment; Figure 4). It is seen that, although participants who moved repeatedly to the same target showed a much better asymptotic performance at the end of each batch, the improvement from the start of one batch to the start of the next was no different than that observed for participants tested with random targets. This indicates that the learning process that operates with repetitions of the same target (error map learning) is independent of the process by which a point map is acquired. (B) Repeated-target participants received ample exposure to the audiomotor workspace. Spatial coverage computes how far any given point in the interval is to the closest movement; therefore, as the set of movements cover the space more densely, the distance to the closest point is less (see supplementary materials, <https://doi.org/10.6084/m9.figshare.5527963.v1>). For the repeated-targets experiment, we calculate the spatial coverage based on all intermediate movements but plot only the first trial to every new target to allow comparison with the continuous-1d experiment coverage. The repeated-target participants' coverage of the space is greater than that of the continuous-1d participants, indicating that their identical performance on the random trials (A) cannot be explained as a lack of spatial coverage. (C) A visual-auditory observation experiment reveals some learning indicating that a cognitive strategy, if used at all, could only lead to a limited improvement in performance.

the first movements to each new target in the repeated-targets experiment is indistinguishable from performance on the first 45 trials in the random experiment, despite the fact that repeated-target participants performed 15 times more trials than random-target participants. This shows that error corrections made in the course of repeated movements to the same target did not contribute to point map formation.

Second, learning error correction in the course of repeated movements to one target improves error-corrective behavior to novel targets; that is, learning the error-corrective process is not tied to particular locations in space. Specifically, error map acquisition is shown by improvements in time constant and asymptotic error when the same target was presented repeatedly, and these improvements cannot be explained by improvements in performance on the first trial of each new target (indicative of point map acquisition).

Third, in the repeated-targets experiments, trial-to-trial correlations of reaching movements were observed that would be expected in error-corrective learning and indeed were seen in simulations of error map learning, whereas these correlations were absent in the series of reaches toward the same target in the random-target presentation condition.

These observations, taken together, suggest that two distinct sensorimotor maps are formed. Random-target presentation favors the formation of a point map, and in that condition, error map formation is negligible because there is no direct opportunity for error corrections to be applied. Repeated presentation of the same target involves formation of both maps. A point map is formed, but it receives the equivalent of one new data point for every new target. The repeated movements to the same target do not feed information to the point map but instead contribute to the formation of an error map, which in turn allows, over the course of many new targets, improved convergence.

DISCUSSION

A paradigm is introduced here to study how participants initially acquire sensorimotor maps. Participants make arm movements to auditory targets. The sounds do not come from different physical locations but, like speech sounds, are distinguished only by their frequency content. Participants are in much the same situation as an infant learning to talk: They have to learn from scratch which movements to make to produce particular sounds and to learn how to correct errors when they occur.

One principal observation is that learning is possible even when targets never repeat. This learning is not directly driven by error correction, because, when targets were selected randomly, there was no opportunity to move to the same target again and no opportunity to correct an error. Instead, simulations (Figure 3) indicate that the learning observed here could be accounted for

by the acquisition of a mapping whose structure could be as simple as a lookup table in which information about movements and their sensory consequences is progressively accumulated.

A second observation is that, when participants are given the opportunity to correct errors by moving repeatedly to the same target, the error correction process itself has to be learned. This is because, when participants first come to this task, even if they experience an error on a given trial, they are unable to use error correction because, by design, the mapping from sensory errors to motor corrections is initially unknown. In the course of the experiment, across many batches, each to a novel target, participants learn to better correct their errors, indicating the acquisition of a mapping between sensory errors and motor corrections (an error map). It is seen that this learning of the error-corrective process is not linked to particular targets because, for each novel target, the learning rate and asymptotic performance improve.

A final observation is that the point and error maps are learned independently. Learning a point map does not contribute to learning an error map and vice versa.

Relation to Previous Work on Sensorimotor Adaptation

The point map is a function from sensory output s to motor commands m (i.e., $f(s) = m$), and the error map is a mapping between sensory errors Δs and motor corrections Δm (i.e., $f'(\Delta s) = \Delta m$). Although sensory information is needed for the formation of both maps, learning in the present context proceeds on a trial-to-trial basis, and both kinds of maps contribute to feedforward control, in one case, to produce movements to novel targets (point map) and, in the other, to produce movements that are appropriate adjustments to previous movements to correct for sensory errors (error map). The error map enables trial-to-trial corrections. It may also contribute to online control of movements, but in the present paradigm, such control was not investigated, and in other work, online control of movement was found to be partially independent of trial-to-trial control (Yousif & Diedrichsen, 2012).

In adaptation learning, trial-to-trial errors are thought to contribute to updating the feedforward controller (Nakanishi & Schaal, 2004; Wolpert & Ghahramani, 2000; Kawato, 1999). However, the error-corrective process itself is fully formed and is not updated during the learning process. These studies tap into an earlier stage of learning in which, by design, the needed mappings are initially unknown. This revealed properties of learning that differ from those involved in adaptation.

First, in this early stage of learning, the error-corrective process itself has to be learned, that is, that the mapping between sensory errors and motor corrections is initially unknown and then learned. The notion that error correction

is learned and not a static process is consistent with previous data. Braun, Aertsen, Wolpert, and Mehring (2009) train participants using random visuomotor rotations in batches of eight trials and find that the error-corrective process is not static because learning a subsequent, novel rotation is accelerated. Herzfeld and colleagues (2014) show that the error-corrective process that allows adaptation to alternating perturbations (such as visuomotor rotations or force fields) is influenced by the history of errors, which could give rise to faster learning similar to the present observation of accelerated convergence to novel targets.

A second property of this early stage of learning is that error correction is not needed to build a feedforward controller. Specifically, error correction does not produce improvements in performance when moving to novel random targets (Figure 9). Moreover, the point map can be learned in the absence of the opportunity for error correction, as is the case when the target changes on each trial. The present work is thus part of a growing literature documenting motor learning in the absence of error correction. In reinforcement-based learning, success or failure drives learning without the need to know either the direction or magnitude of error (Izawa & Shadmehr, 2011; Sutton & Barto, 1998). In use-dependent learning, performance improvements are observed in the context of repetition of movements alone (Diedrichsen, White, Newman, & Lally, 2010; Nudo, Milliken, Jenkins, & Merzenich, 1996). However, there is little opportunity for use-dependent learning when the target, and therefore the movement, is different on every trial (see below for a discussion of the potential role of reinforcement in the present learning).

Error-corrective trials do not feed into the point map but instead contribute to forming an error map, which is learned independently. Using two maps may allow a more efficient encoding of space than a single map, as shown by a computational architecture (Figure 1). In the context of that architecture, adding a third map (or more) would seem to yield a yet more efficient coding of space, but for the sensorimotor apparatus to use such a third map (a derivative error map), it would have to be able to compute a trial-to-trial change in error, a capability for which, to our knowledge, there is no empirical support. The closest operation the sensorimotor apparatus has been shown to be capable of is to compute the sign of the change of errors (Herzfeld et al., 2014). This study in principle does not rule out the existence of additional maps (in addition to the point and error maps) but is concerned with showing that there are at least two maps.

When targets never repeat, learning could in principle be driven by sensory prediction error: After each movement, the predicted sensory effect of that movement (the forward model) is updated based on the actual feedback. The updated sensory prediction is the previous prediction plus a fraction of the prediction error (dif-

ference between prediction and actual feedback; Synofzik, Lindner, & Thier, 2008; Tseng, Diedrichsen, Krakauer, Shadmehr, & Bastian, 2007; Synofzik, Thier, & Lindner, 2006). However, this account relies on the assumption that prediction errors update a forward model, and this assumption entails that all information is gathered in a single sensorimotor map. If only one map is learned that all movements are fed into, then repeated movements to the same target should yield better learning than making only one movement to each target, which is contradicted by the present data (Figure 9). Furthermore, learning dependent on sensory prediction error requires that participants generate sensory predictions, but the present data suggest that at least initially they do not, as demonstrated by the absence of a correlation between targets and movements.

Both point and error map learning display characteristics of generalization. Specifically, in point map learning, the target on every trial is different, yet participants' performance improves over the course of training. In error map learning, each set of movements involves a unique target, yet the convergence onto novel targets improves with training. Simulations (Figure 3) indicate that such generalization behavior would be expected, even if the map is as simple as a lookup table. This is because, as units are added with learning and map density is increased, the nearest neighbor to any given target will be closer and therefore reaching error will be smaller. Note that the lookup table account assumes the presence of a distance metric. The returned value is then used to probe the lookup table. Whether maps indeed represent instance-based learning (a lookup table) or instead encode the structure of the space remains a question for future research.

The question of whether reinforcement could drive the learning observed here merits consideration. In a restricted sense where participants repeat previously rewarded movements, reinforcement could not account for the learning here. On a given trial, when a participant produces a movement that is similar to the target sound, although they could experience it as rewarding, if they would in future trials attempt to repeat this movement, it will not necessarily be beneficial because the target may be different. However, in a more general sense, reinforcement paired with a map approximation function could drive learning. The question is what this map approximation function would be. One option proposed here is that it could be as simple as a lookup table.

Adapting Existing Sensorimotor Maps

The acquisition of sensorimotor mappings has been studied previously, but often in cases where participants already had either fully formed sensorimotor maps or prior expectations about the structure of the sensory-to-motor relationships. For example, in visuomotor adaptation, vision of the arm is experimentally displaced as participants make reaching movements (Krakauer, 2009).

The newly acquired sensorimotor map can be represented as the existing sensorimotor mapping plus a correction term that pertains to the particular experimental perturbation (Telgen et al., 2014). Indeed, this is typically how sensorimotor adaptation is computationally modeled (Cheng & Sabes, 2006; Ghahramani et al., 1997). The novelty of the work here is not simply that it pertains to the auditory (instead of visual) modality. Indeed, during speaking, people are known to compensate for altered auditory feedback (Houde & Jordan, 1998). The novelty of the paradigm is that there is no preexisting sensorimotor map that gets adjusted. Apart from work on babbling in early infancy, little is known about how these maps are initially formed.

Acquisition of sensorimotor mappings closer to that investigated here is reported in studies in which participants controlled a screen cursor using physical movements with the fingers or hands (Yamamoto, Hoffman, & Strick, 2006; Mosier et al., 2005; Sailer, Flanagan, & Johansson, 2005) but through a nontrivial transformation. The present work builds on these studies by elucidating the structure of the acquired maps.

Learning Musical Instruments

Learning a musical instrument (Herholz & Zatorre, 2012; Bangert & Altenmüller, 2003) also requires forming an audiomotor map, and therefore its neural underpinnings may be similar to the neural structures that enable sensorimotor map learning observed here (Zatorre, Chen, & Penhune, 2007). A large body of research documents differences between musicians and nonmusicians in brain morphology (Vaquero et al., 2016; Gaser & Schlaug, 2003) or brain networks documented using resting state scans (Palomar-García, Zatorre, Ventura-Campos, Bueichekú, & Ávila, 2017; Feinberg & Setsompop, 2013; Luo et al., 2012). These differences occur in a distributed network of areas including auditory cortices, primary motor cortex, premotor areas, superior temporal gyrus, somatosensory cortex, and the BG. However, it is unclear whether the observed differences are specifically due to musical training or to neuroanatomical factors that predispose one to become a musician.

A number of studies monitor participants directly as they learn to play a musical instrument, and these studies implicate a number of brain areas. Activity in the dorsal premotor cortex (dPMC) was found to be reduced late versus early in learning to play a musical melody on a keyboard (Chen, Rae, & Watkins, 2012); another study found increased activity in dPMC after nonmusicians had been trained to associate musical chords with keystrokes (Bermudez & Zatorre, 2005). Lega and colleagues applied TMS during a task in which participants learned associations between keystrokes and sounds and found that dPMC is causally involved in this learning (Lega, Stephan, Zatorre, & Penhune, 2016; for a related paradigm, see Säfström & Edin, 2006; Wise & Murray, 2000). Other areas

associated with musical training are the posterior superior temporal gyrus, which after drum training, showed increased connectivity with the rest of the brain (Amad et al., 2017).

There are differences between learning to play a musical instrument and the audiomotor task employed in the present studies that bear on the interpretation of reported patterns of neural activity. First, the acquisition documented in studies of learning to play musical instruments presumably relies on prior information, because even nonmusicians have structural expectations about the mappings between space and pitch (Rusconi, Kwan, Giordano, Umiltà, & Butterworth, 2006). Second, learning to play a musical instrument not only involves the formation of an audiomotor map but also entails changes to motor and perceptual function as well (Kraus & Chandrasekaran, 2010). Accordingly, it is not possible based on the current studies of musical instrument learning to dissociate areas whose activation specifically reflects map learning.

Summary

Participants acquired novel sensorimotor maps by making reaching movements to auditory targets and, in this process, form two independent mappings: a point map connecting sensory targets to motor commands and an error map that links motor corrections to sensory errors. This study identifies the structure of these maps, and the future challenge is to determine the learning rules by which their content is acquired.

Acknowledgments

We are indebted to Bilal Alchalabi and Eric di Tomasso for assistance in data collection and Paul Gribble, Mark Tiede, and members of the Motor Control Lab at McGill for valuable discussions. This work was supported by National Institute of Child Health and Human Development R01 HD075740, Les Fonds Québécois de la Recherche sur la Nature et les Technologies, Québec (FQRNT), and a Banting Fellowship BPF-NSERC-01098.

Reprint requests should be sent to David J. Ostry, Psychology Department, McGill University, Montreal, QC, H3A 1B1 Canada, or via e-mail: david.ostry@mcgill.ca.

REFERENCES

- Amad, A., Seidman, J., Draper, S. B., Bruchhage, M. M. K., Lowry, R. G., Wheeler, J., et al. (2017). Motor learning induces plasticity in the resting brain—Drumming up a connection. *Cerebral Cortex*, 27, 2010–2021.
- Bangert, M., & Altenmüller, E. O. (2003). Mapping perception to action in piano practice: A longitudinal DC-EEG study. *BMC Neuroscience*, 4, 26.
- Bermudez, P., & Zatorre, R. J. (2005). Conditional associative memory for musical stimuli in nonmusicians: Implications for absolute pitch. *Journal of Neuroscience*, 25, 7718–7723.
- Braun, D. A., Aertsen, A., Wolpert, D. M., & Mehring, C. (2009). Motor task variation induces structural learning. *Current Biology*, 19, 352–357.
- Chen, J. L., Rae, C., & Watkins, K. E. (2012). Learning to play a melody: An fMRI study examining the formation of auditory-motor associations. *Neuroimage*, 59, 1200–1208.

- Cheng, S., & Sabes, P. N. (2006). Modeling sensorimotor learning with linear dynamical systems. *Neural Computation*, 18, 760–793.
- Diedrichsen, J., White, O., Newman, D., & Lally, N. (2010). Use-dependent and error-based learning of motor behaviors. *Journal of Neuroscience*, 30, 5159–5166.
- Feinberg, D. A., & Setsompop, K. (2013). Ultra-fast MRI of the human brain with simultaneous multi-slice imaging. *Journal of Magnetic Resonance*, 229, 90–100.
- Gaser, C., & Schlaug, G. (2003). Brain structures differ between musicians and non-musicians. *Journal of Neuroscience*, 23, 9240–9245.
- Ghahramani, Z., Wolpert, D. M., & Jordan, M. I. (1997). Computational models of sensorimotor integration. *Advances in Psychology*, 119, 117–147.
- Herholz, S. C., & Zatorre, R. J. (2012). Musical training as a framework for brain plasticity: Behavior, function, and structure. *Neuron*, 76, 486–502.
- Herzfeld, D. J., Vaswani, P. A., Marko, M. K., & Shadmehr, R. (2014). A memory of errors in sensorimotor learning. *Science*, 345, 1349–1353.
- Houde, J. F., & Jordan, M. I. (1998). Sensorimotor adaptation in speech production. *Science*, 279, 1213–1216.
- Huberdeau, D. M., Krakauer, J. W., & Haith, A. M. (2015). Dual-process decomposition in human sensorimotor adaptation. *Current Opinion in Neurobiology*, 33, 71–77.
- Izawa, J., & Shadmehr, R. (2011). Learning from sensory and reward prediction errors during motor adaptation. *PLoS Computational Biology*, 7, e1002012.
- Jordan, M. I., & Rumelhart, D. E. (1992). Forward models: Supervised learning with a distal teacher. *Cognitive Science*, 16, 307–354.
- Kawato, M. (1999). Internal models for motor control and trajectory planning. *Current Opinion in Neurobiology*, 9, 718–727.
- Krakauer, J. W. (2009). Motor learning and consolidation: The case of visuomotor rotation. *Advances in Experimental Medicine and Biology*, 629, 405–421.
- Kraus, N., & Chandrasekaran, B. (2010). Music training for the development of auditory skills. *Nature Reviews Neuroscience*, 11, 599–605.
- Lega, C., Stephan, M. A., Zatorre, R. J., & Penhune, V. (2016). Testing the role of dorsal premotor cortex in auditory-motor association learning using transcranial magnetic stimulation (TMS). *PLoS One*, 11, e0163380.
- Liu, X., Mosier, K. M., Mussa-Ivaldi, F. A., Casadio, M., & Scheidt, R. A. (2011). Reorganization of finger coordination patterns during adaptation to rotation and scaling of a newly learned sensorimotor transformation. *Journal of Neurophysiology*, 105, 454–473.
- Luo, C., Guo, Z.-W., Lai, Y.-X., Liao, W., Liu, Q., Kendrick, K. M., et al. (2012). Musical training induces functional plasticity in perceptual and motor networks: Insights from resting-state fMRI. *PLoS One*, 7, e36568.
- Marriott, F. H. C., & Pope, J. A. (1954). Bias in the estimation of autocorrelations. *Biometrika*, 41, 390–402.
- 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, 4362–4372.
- Mussa-Ivaldi, F. A., Casadio, M., Danziger, Z. C., Mosier, K. M., & Scheidt, R. A. (2011). Chapter 3—Sensory motor remapping of space in human-machine interfaces. In A. M. Green, J. F. Kalaska, F. Lepore, & C. E. Chapman (Eds.), *Progress in brain research, Vol. 191, Enhancing performance for action and perception: Multisensory integration, neuroplasticity and neuroprosthetics, part I* (pp. 45–64). Amsterdam: Elsevier.
- Nakanishi, J., & Schaal, S. (2004). Feedback error learning and nonlinear adaptive control. *Neural Networks*, 17, 1453–1465.
- Nudo, R. J., Milliken, G. W., Jenkins, W. M., & Merzenich, M. M. (1996). Use-dependent alterations of movement representations in primary motor cortex of adult squirrel monkeys. *Journal of Neuroscience*, 16, 785–807.
- Palomar-García, M.-Á., Zatorre, R. J., Ventura-Campos, N., Bueichekú, E., & Ávila, C. (2017). Modulation of functional connectivity in auditory-motor networks in musicians compared with nonmusicians. *Cerebral Cortex*, 27, 2768–2778.
- Pouget, A., & Snyder, L. H. (2000). Computational approaches to sensorimotor transformations. *Nature Neuroscience*, 3(Suppl.), 1192–1198.
- Radhakrishnan, S. M., Baker, S. N., & Jackson, A. (2008). Learning a novel myoelectric-controlled interface task. *Journal of Neurophysiology*, 100, 2397–2408.
- Remez, R. E., Rubin, P. E., Pisoni, D. B., & Carrell, T. D. (1981). Speech perception without traditional speech cues. *Science*, 212, 947–949.
- Robinson, D. W., & Dadson, R. S. (1956). A re-determination of the equal-loudness relations for pure tones. *British Journal of Applied Physics*, 7, 166.
- Rusconi, E., Kwan, B., Giordano, B. L., Umiltà, C., & Butterworth, B. (2006). Spatial representation of pitch height: The SMARC effect. *Cognition*, 99, 113–129.
- Säfsström, D., & Edin, B. B. (2006). Acquiring and adapting a novel audiomotor map in human grasping. *Experimental Brain Research*, 173, 487–497.
- Sailer, U., Flanagan, J. R., & Johansson, R. S. (2005). Eye-hand coordination during learning of a novel visuomotor task. *Journal of Neuroscience*, 25, 8833–8842.
- Shadmehr, R., Smith, M. A., & Krakauer, J. W. (2010). Error correction, sensory prediction, and adaptation in motor control. *Annual Review of Neuroscience*, 33, 89–108.
- Stevens, S. S. (1937). A scale for the measurement of the psychological magnitude pitch. *Journal of the Acoustical Society of America*, 8, 185.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.
- Synofzik, M., Lindner, A., & Thier, P. (2008). The cerebellum updates predictions about the visual consequences of one's behavior. *Current Biology*, 18, 814–818.
- Synofzik, M., Thier, P., & Lindner, A. (2006). Internalizing agency of self-action: Perception of one's own hand movements depends on an adaptable prediction about the sensory action outcome. *Journal of Neurophysiology*, 96, 1592–1601.
- Telgen, S., Parvin, D., & Diedrichsen, J. (2014). Mirror reversal and visual rotation are learned and consolidated via separate mechanisms: Recalibrating or learning de novo? *Journal of Neuroscience*, 34, 13768–13779.
- Thoroughman, K. A., & Shadmehr, R. (2000). Learning of action through adaptive combination of motor primitives. *Nature*, 407, 742–747.
- Tseng, Y.-W., Diedrichsen, J., Krakauer, J. W., Shadmehr, R., & Bastian, A. J. (2007). Sensory prediction errors drive cerebellum-dependent adaptation of reaching. *Journal of Neurophysiology*, 98, 54–62.
- Vaquero, L., Hartmann, K., Ripollés, P., Rojo, N., Sierpowska, J., François, C., et al. (2016). Structural neuroplasticity in expert pianists depends on the age of musical training onset. *Neuroimage*, 126, 106–119.
- Wise, S. P., & Murray, E. A. (2000). Arbitrary associations between antecedents and actions. *Trends in Neurosciences*, 23, 271–276.
- Wolpert, D. M., & Ghahramani, Z. (2000). Computational principles of movement neuroscience. *Nature Neuroscience*, 3, 1212–1217.
- Yamamoto, K., Hoffman, D. S., & Strick, P. L. (2006). Rapid and long-lasting plasticity of input-output mapping. *Journal of Neurophysiology*, 96, 2797–2801.
- Yousif, N., & Diedrichsen, J. (2012). Structural learning in feedforward and feedback control. *Journal of Neurophysiology*, 108, 2373–2382.
- Zatorre, R. J., Chen, J. L., & Penhune, V. B. (2007). When the brain plays music: Auditory-motor interactions in music perception and production. *Nature Reviews Neuroscience*, 8, 547–558.