RESEARCH ARTICLE | *Control of Movement*

Early stages of sensorimotor map acquisition: learning with free exploration, without active movement or global structure

F. T. van Vugt[1,2] and D. J. Ostry[1,2]

[1]*Psychology Department, McGill University, Montreal, Canada; and* [2]*Haskins Laboratories, New Haven, Connecticut*

van Vugt FT, Ostry DJ. Early stages of sensorimotor map acquisition: learning with free exploration, without active movement or global structure. *J Neurophysiol* 122: 1708–1720, 2019. First published August 21, 2019; doi:10.1152/jn.00429.2019.—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. The initial stages of map acquisition can be studied by having participants learn arm movements to auditory targets. The key question is what mechanism drives this early learning. Three learning processes from previous literature were tested: map learning may rely on active motor outflow (target), on error correction, and on the correspondence between sensory and motor distances (i.e., that similar movements map to similar sounds). Alternatively, we hypothesized that map learning can proceed without these. Participants made movements that were mapped to sounds in a number of different conditions that each precluded one of the potential learning processes. We tested whether map learning relies on assumptions about topological continuity by exposing participants to a permuted map that did not preserve distances in auditory and motor space. Further groups were tested who passively experienced the targets, kinematic trajectories produced by a robot arm, and auditory feedback as a yoked active participant (hence without active motor outflow). Another group made movements without receiving targets (thus without experiencing errors). In each case we observed substantial learning, therefore none of the three hypothesized processes is required for learning. Instead early map acquisition can occur with free exploration without target error correction, is based on sensory-to-sensory correspondences, and possible even for discontinuous maps. The findings are consistent with the idea that early sensorimotor map formation can involve instance-specific learning.

NEW & NOTEWORTHY This study tested learning of novel sensorimotor maps in a variety of unusual circumstances, including learning a mapping that was permuted in such as way that it fragmented the sensorimotor workspace into discontinuous parts, thus not preserving sensory and motor topology. Participants could learn this mapping, and they could learn without motor outflow or targets. These results point to a robust learning mechanism building on individual instances, inspired from machine learning literature.

error correction; motor learning; sensorimotor mapping; speech; topology

INTRODUCTION

When learning new motor skills, such as learning to speak or play a musical instrument, the learner is faced with the challenge of determining which movement to perform to make a particular sound: learning a sensorimotor map. Typically, studies of sensorimotor learning are based on cases where subjects come to the laboratory with a preexisting sensorimotor map, which is then experimentally perturbed, for example, by imposing a rotation of a visual cursor. By contrast, during the earliest stages of learning of a novel sensorimotor map subjects have no knowledge of the mapping between movements and sensory effects, and this phase can be studied by monitoring participants as they learn a novel association between movements and sounds (van Vugt and Ostry 2018a). Little is known about what drives this kind of learning. In other learning paradigms, a number of underlying mechanisms have been identified, and here we tested whether each of these is involved in the initial stage of map learning as well. Previous work on motor adaptation has shown that active motor outflow and a process of (target) error correction are key ingredients in many learning settings. The learner may also proceed by making certain assumptions about the sensorimotor map, such as that similar movements map to similar sensory effects, i.e., that the mapping preserves the distance relationships in the sensory and motor domains. Alternatively, we hypothesized that early map learning may not rely on any of these mechanisms and instead be based on accumulating a history of previous instances of movements and sensory effects, without computing proportionate error corrections or even requiring active motor outflow. These hypotheses are contrasted here in the context of learning novel sensorimotor maps from scratch, in which participants were studied as they first learn a mapping between arm movements and sounds. The results indicate that learning in this context is possible without active motor outflow or (target) error correction and that discontinuous mappings can be learned that do not preserve the topology of the sensory and motor spaces.

Many aspects of sensorimotor learning are thought to rely on active motor outflow, which is used here to mean a motor command that would be reflected in muscle activation. For example, in a classic study neonatal kittens were divided in pairs where one kitten moved actively while another yoked individual passively experienced the motion and the visual sensory experience of the first. The passive individual failed to develop similar sensory and motor capacities even though it

experienced the same sensory input (Held and Hein 1963), suggesting that motor outflow is required for learning. Similarly, some studies on humans show that although passive movement training may have some benefit (Beste and Dinse 2013), learning is only partial or limited (Beets et al. 2012; Lotze et al. 2003). However, other studies show that reinforcement motor learning is possible through passive somatosensory experience coupled with accuracy feedback (Bernardi et al. 2015), and learning new mechanical environments can occur through passive visual observation of another individual performing the task to be learned (Mattar and Gribble 2005). In the case of learning sensorimotor maps, it remains unclear whether active motor outflow is necessary. The present study tested passive groups who experienced the same kinematic movement trajectories as an active yoked participant learning the sensorimotor map and were also presented with the same sounds. If active motor outflow is critical, these groups should show impaired learning, but if mappings can be formed based only on sensory information (somatosensory and auditory), learning should occur.

When subjects learn a sensorimotor map, it may further be hypothesized that they make assumptions about the type of map they will encounter. For example, participants may anticipate that the novel sensorimotor mapping preserves distances in sensory and motor modalities (Ritter et al. 1989), which means in the present case that if two movements are similar (small angular distance), then the associated sounds will also be similar (in the frequency domain). The sensorimotor apparatus that learners are typically confronted with has this (homomorphic) property: for example, similar arm movements result in similar changes in position in visual coordinates, or similar motor commands sent to the vocal apparatus yield similar sounds (Guenther et al. 2006). Humans can learn complex artificial mappings (Mosier et al. 2005; Mussa-Ivaldi et al. 2011), but the mappings used in these studies still preserve distances. To our knowledge the question whether the sensorimotor map learning mechanism is rooted in the assumption that mappings preserve distances has not been addressed directly. The current setting where participants learned a novel sensorimotor map provided an opportunity to expose participants to a permuted mapping that violated these distance assumptions, by dividing the sensorimotor workspace into bins and then shuffling the way these bins mapped to sounds (Fig. 1D). In this permuted mapping, there can be movements that straddle the boundary between bins and therefore are close together in angular direction have very different associated sounds. The permutation applied here to the mapping left residual local structure (within bins), but the size of the perturbation bins was chosen such that this local structure would not be usable and this was verified empirically. If learning proceeds from the assumption that sensory and motor distances are preserved, then learning should be impaired in this condition, but if learning does not require the presence of global structure, then one would expect similar performance with the permuted map.

Another mechanism thought to drive motor learning is error correction, in which the difference between a target and actually produced movement is computed and then a proportionate correction is applied to update performance on the next trial (Seidler et al. 2013; Wolpert and Ghahramani 2000; Wolpert et al. 1995). The target can be an external sensory goal (task) or

internal (sensory prediction; Butcher and Taylor 2018; Mazzoni and Krakauer 2006), yielding two potential error signals, thought to drive explicit and implicit adaptation processes, respectively (Leow et al. 2018; Reichenthal et al. 2016; Taylor and Ivry 2011). When learning novel sensorimotor maps, learning is slow and shows large errors (van Vugt and Ostry 2018a), which leads to the hypothesis that sensorimotor map learning may be driven initially by task (target) error. Alternatively, one can hypothesize that sensorimotor map acquisition proceeds by learning associations between movements and sensory consequences, which at least in principle is possible using free exploratory movements without targets and hence without target error. Similar to this, some forms of sensorimotor learning have been reported previously that do not rely on error correction (Haith and Krakauer 2013) such as use-dependent learning (Diedrichsen et al. 2010). In this work, it will be tested whether sensorimotor map learning requires target error using a free exploration condition in which no targets are presented. If learning relies on target error, this condition should show impaired performance, but if error is not required, learning should be identical to a condition with targets.

The various conditions described above serve to evaluate whether active motor outflow, structure preservation, and/or target error correction play a role in learning sensorimotor maps. An alternative hypothesis is that learning is based on collecting a history of instances of sensory-to-sensory correspondences, yielding an instance-based sensory learning model such as described previously in machine learning (Aha et al. 1991; Atkeson et al. 1997a) and for sensorimotor map learning (van Vugt and Ostry 2018a). Such a learning model may be appropriate during the earliest stages of acquiring novel sensorimotor maps, because in such cases little to no prior information is available on which to base learning.

METHODS

A total of 165 participants was included (96 female, 67 male, 2 not reported) with average age of 23.3 (SD 4.2) yr, right-handed, normal hearing, and no prior neurological conditions. Participants were selected to not currently be taking music lessons. Various participants had received minimal music training in school. As for private lessons, participants were categorized into those that had received formal musical training in the past (>5 yr ago) or informal lessons; none was currently musically active. The distribution of subjects among these categories was not different across experimental conditions [$\chi^2(12) = 16.70$, $P = 0.16$]. All procedures were approved by the McGill University Faculty of Medicine Institutional Review Board. Participants provided written, informed consent.

A systematic mapping was defined between spatial locations and sounds, where the frequency content of sounds varied with the angle of the movement (see *Sounds*; of note, the sounds were not acoustically localized at different points in space but only differed by their frequency content). Participants were trained to learn this mapping by making movements and hearing the sound each movement was mapped to. The experimental conditions tested here were all variants of the following procedure. Participants were seated at a planar arm robot (Interactive Motion Technologies) and held the handle in their (dominant) right hand. On audiomotor trials, participants made a reaching movement to points on a hidden target circle of 20 cm, and at the end of the movement, a feedback sound was played that depended on the angle of the end point of the subject's movement (Fig. 1, *A* and *B*). Before each movement a target sound was presented, and the subjects' task was to move to the location that would
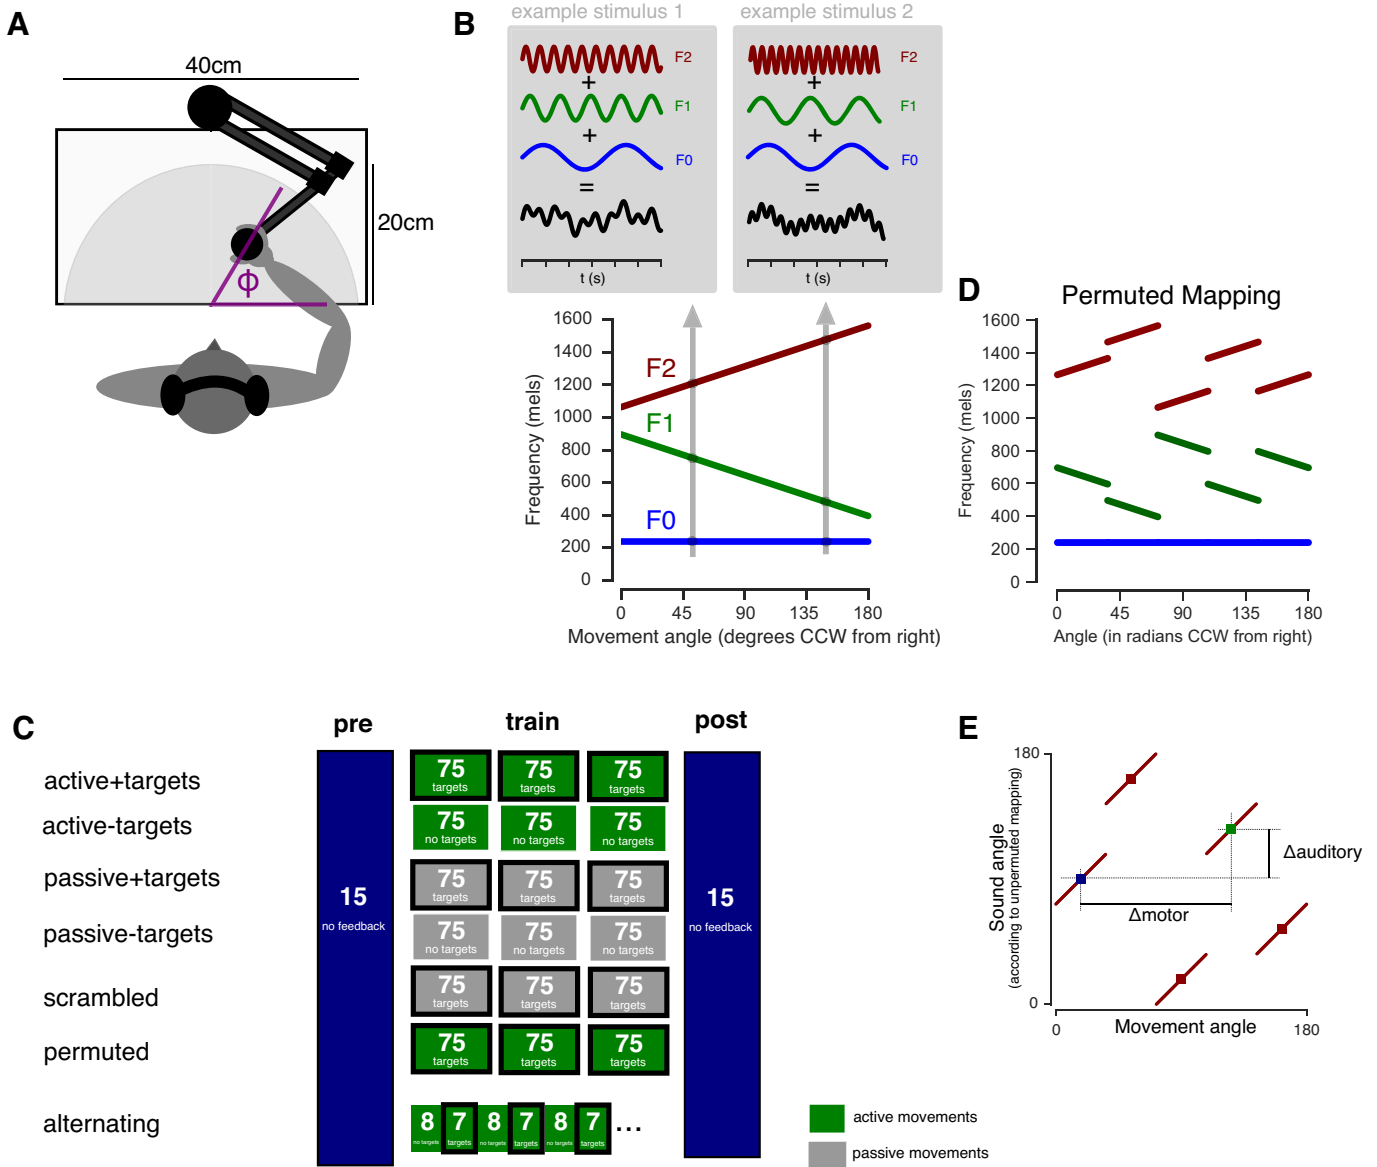
Fig. 1. Participants made reaching movements that were turned into sounds. *A*: participants were seated at a planar arm movement robot and made movements from a central position to points on a half-circle. Vision was blocked during the entire experiment. *B*: at movement end a feedback sound was delivered that consisted of 3 pure tone oscillators (F0, F1, and F2) presented simultaneously and whose frequencies depended on the angle of the movement. CCW, counterclockwise. *C*: to test whether target error is required for learning, we tested conditions were targets were absent altogether (active − targets) or where targets and feedback were present but not on the same trials (alternating). To test whether active motor outflow is necessary, participants in passive groups experienced the movements of yoked active participants with or without the associated target sounds. In a scrambled condition, participants also experienced movements, targets, and feedback but the time series of these were shuffled so that the target, movement, and feedback experienced at any moment originated from different trials. *D*: to test whether the subjects assumed a priori that the mapping they experienced would preserve the structure of auditory and motor spaces, a permuted mapping was designed in which the workspace was divided into 5 bins and the sound-to-movement mapping was shuffled across these bins. Note that the mapping is still one-to-one. *E*: this permuted mapping does not preserve distances between auditory and motor spaces, so that two points can be close together in auditory space but far apart in motor space, as indicated here, or vice versa.

yield that feedback sound. Target angles were sampled from a uniform distribution covering 0–180° and were different on every trial for every subject. There were no visual targets in any condition. All vision was blocked during the entire experiment: the experiment was conducted in a darkened room and participant's vision of their arm was blocked by a semitransparent mirror. Participants performed 3 blocks of 75 training trials yielding a total of 225 trials (Fig. 1*C*).

No-feedback trials: before and after these training trials learning was assessed by having participants perform movements to a set of 15 auditory targets spaced linearly (equidistant) covering the entire workspace (0–180°), presented once each in random order, but without receiving feedback. Thus all subjects received the same set of

targets in these blocks. By design, before training subjects had no information about the correspondence between movements and sounds because the mapping was arbitrary (but fixed during the experiment) and therefore they were expected to reach randomly. The baseline trials were important to ensure that reaching was indeed random and that there were no biases. The analyses that assessed learning were conducted without regard to this baseline block.

Learning conditions. Various conditions were tested that differed in the training phase, but the no-feedback trials before and after learning were always the same (Fig. 1*C*). In what follows, "targets" refers to auditory targets that were presented to subjects through headphones. There were no visual targets at any point. In the active + targets

condition ($n = 25$), participants received a target on each trial, made an active movement, and then received a feedback sound. In the active − targets condition ($n = 20$), participants did not receive a target sound but made an active movement (when a visual icon appeared, which contained no information about the target location) and received auditory feedback. In active trials without targets, subjects were told they could move freely to any location of their choosing but had to avoid moving across the workspace in a systematic way. (This restriction was based on a pilot test in which subjects explored the space in a patterned way, which was markedly different from the movement pattern in the active + targets condition thus precluding a fair comparison.) An alternating condition was tested ($n = 20$), in which participants first made seven movements without targets but with auditory feedback and then made eight movements with targets but without auditory feedback. In this way, both targets and feedback were present during learning, but never in the same trial, and therefore, we reasoned that target error (defined as the difference between the target and auditory feedback) could not be computed in this condition. At the transition between a set of target trials and the next set of feedback trials, it was possible in principle that subjects would use the target of the last trial and combine it with the feedback on the next trial. To prevent this, a masking sound was presented (4 bursts of pink noise, 4 s each, faded in and out during 1 s), which served to erode the potential auditory memory of the target sound. In this alternating targets/no-targets condition, unlike the other experimental conditions, subjects performed 6 blocks of 75 trials, half of which were target (no feedback) trials and half were no-target (feedback) trials. The rationale for this was that even though the alternating group performed more trials in total than the active + targets group, they were matched for the number of targets and feedback sounds.

In the passive conditions, each participant was yoked to an active participant previously tested, whose targets, movement trajectories, and feedback sounds had been recorded. These sounds were played back to the passive participant, and the kinematic trajectory was reproduced by the robotic arm. The passive participant thus experienced the same sounds and movement as the active participant but without themselves producing motor outflow. Passive movements were generated by the robot under position servo-control using the time series of positions measured from the trajectory of a yoked active subject. Two variants of this passive condition were tested: passive + targets (where the target sounds were presented, $n = 25$) and passive − targets (where the target sounds were omitted, $n = 25$). Note that in both passive groups, the yoked active participants were taken from the active + targets group, which was done so as to equate the two passive groups. In passive trials, subjects were told that another participant had made movements that resulted in specific sounds and that the robot would now reproduce their trajectories. The passive subjects' task was merely to pay attention to the movements and the sounds.

As a passive control group, a scrambled condition was tested ($n = 25$), which was identical to the passive + targets condition, except that within each block of 75 movements, the order of the target sounds, the movements, and the feedback sounds were each shuffled independently, so that on a given trial the target, movement, and feedback sound that were presented originated from different trials of the yoked active participant. In this way, the movement did not correspond to the feedback or the target. The hypothesis was that learning in this scrambled condition is not possible, because to learn the mapping, the subject needs to be presented with auditory feedback that maps to movements with some degree of consistency.

To test whether learning required that the mapping preserved the structure of the auditory and motor space, a permuted mapping was created. In this condition, the physical workspace was divided into five bins that were shuffled (Fig. 1D). To determine the auditory feedback sound of a given movement, the original end point angle was first mapped to the shuffled bin and then mapped to the sound in the usual way, yielding an irregular movement-to-sound mapping. In

particular, this mapping had the property that two movement-sound pairs could be close together in sound space but far apart in motor space, and vice versa (Fig. 1E). Participants in the permuted condition ($n = 25$) received targets and feedback sounds. In the permuted condition, the location of the targets during no-feedback testing was slightly displaced for the following reason. The workspace was divided into five bins of 36° each. That is, movement angles from 0 to 36° were mapped to *bin 1*, angles from 36 to 72° to *bin 2*, etc. Targets placed at the edges of the bins (e.g., at 36°) could be considered ambiguous, because they could arbitrarily be set to fall into either of the two bins. Therefore these targets were slightly displaced (~3° on average) to avoid falling on these edge locations.

Sounds. The sounds used in this study consisted of three pure tone oscillators presented simultaneously, whose frequencies depended on the direction of the associated movement (Fig. 1B). A mel-space mapping was used so that the same change in direction anywhere in the workspace caused a perceptually comparable difference in the frequency of the individual oscillators. Similarly, the sound amplitude was adjusted for loudness as a function of frequency using previously published equal loudness curves (Robinson and Dadson 1956) for 75 phon, so that the perceived intensity of the sounds associated with various directions was roughly equal. The same sound was presented to both ears, and therefore, there were no acoustic direction cues. The sounds associated with the different movement directions only differed by their frequency content.

Testing auditory and motor function in isolation. Before and after the auditory-motor trials described above, participants performed an auditory discrimination test and a motor copy test, which were designed to test whether there was any change in auditory or motor functioning, respectively. In the auditory test, participants' discrimination thresholds were tested with the sounds used in the auditory-motor experiment. On each trial, four sounds were presented (200-ms duration, 75-ms pause between sounds), which were identical except for one sound, which was always the second or third, and the subject's task was to respond by button press whether the second or third was the oddball. The three identical sounds were the sounds that corresponded to the middle of the workspace (90°, Fig. 1B) and the oddball was chosen to correspond to one of 10 sounds corresponding to logarithmically spaced angles of 90.27–106.20° (inclusive). Participants completed a total of 200 trials.

A motor copy test was performed to measure proprioceptive and motor function in isolation, without auditory stimulation. In this test (which was different from the audiomotor pre/posttests), on each trial, the robot moved the subjects' hand passively to a target location on the circle (900-ms smooth trajectory) and back, and then the subjects' task was to actively reproduce the movement they experienced. Vision of the arm was blocked during this procedure, and no sounds were presented. The target locations were 11 positions linearly spaced on the entire target circle (from 0 to 180°, inclusive). The targets were each presented once before training and once after training in random order.

Data analysis. The movement end point was determined online (defined as the moment when velocity remained below 5% of the peak velocity on a per-trial basis consecutively for 50 ms), and the angle from the starting point to the movement end point was computed to determine the feedback sound to be presented on each trial. The same movement angle was also used offline to calculate the reaching error: the absolute angular difference between the target angle and the angle at the end point of movement, relative to the starting point. Control analyses also investigated the deviation from the center of the workspace, computed as the absolute angular difference between the movement angle and the direction straight ahead, as well as the average or standard deviation of the movement angles, computed using the signed value of the movement angles.

Error was averaged for each subject and block (no-feedback pre- and posttraining). The main analyses focused on the posttraining no-feedback error scores only because the prescores were random,

which is a direct result of the design of this study where subjects learn a novel arbitrary map. Differences in learning between groups were assessed using an ANOVA with the posttraining no-feedback error as dependent variable, and a between-subjects factor group (7 levels for the groups: active + targets, active − targets, alternating, passive + targets, passive − targets, scrambled, and permuted). Pairwise *t* tests (Bonferroni-Holm corrected) were then computed to determine which pairs of groups showed different amounts of learning. Investigation of performance during the training phase was only meaningful for the active groups, and among these, reaching error could by design only be computed for the active + targets, permuted, and alternating groups and in the latter case only for the trials in which targets were presented. Levene's test was performed for homogeneity of variance, and adjustments are reported wherever required. Generalized effect size is reported for significant contrasts (η^2). In figures, error bars indicate means ± SE.

RESULTS

Participants made active movements to targets, receiving auditory feedback at the end of each movement (active +

targets condition). In this condition, a reduction in error was observed over the course of training (Fig. 2*A*). To assess learning, before and after training, participants performed reaching movements without receiving auditory feedback (no-feedback trials). Reaching error, defined as the absolute angular difference between the target angle and reaching end point angle, was lower after training [$t(24) = 6.82$, $P < 0.0001$]. To test hypotheses about the information that drives learning, various conditions were tested that systematically removed one or more of these sources of information (Figs. 1 and 2).

The first key finding is that all groups showed learning, as indicated by a reduction in error on no-feedback trials (all $t > 2.61$, $P < 0.015$; Fig. 2*D*) relative to the baseline (pre) no-feedback trials except for a scrambled control group that passively experienced movements and sounds of an active participant but in which the link between movements and sounds was removed by shuffling the trials [$t(24) = 0.04$, $P = 0.96$].
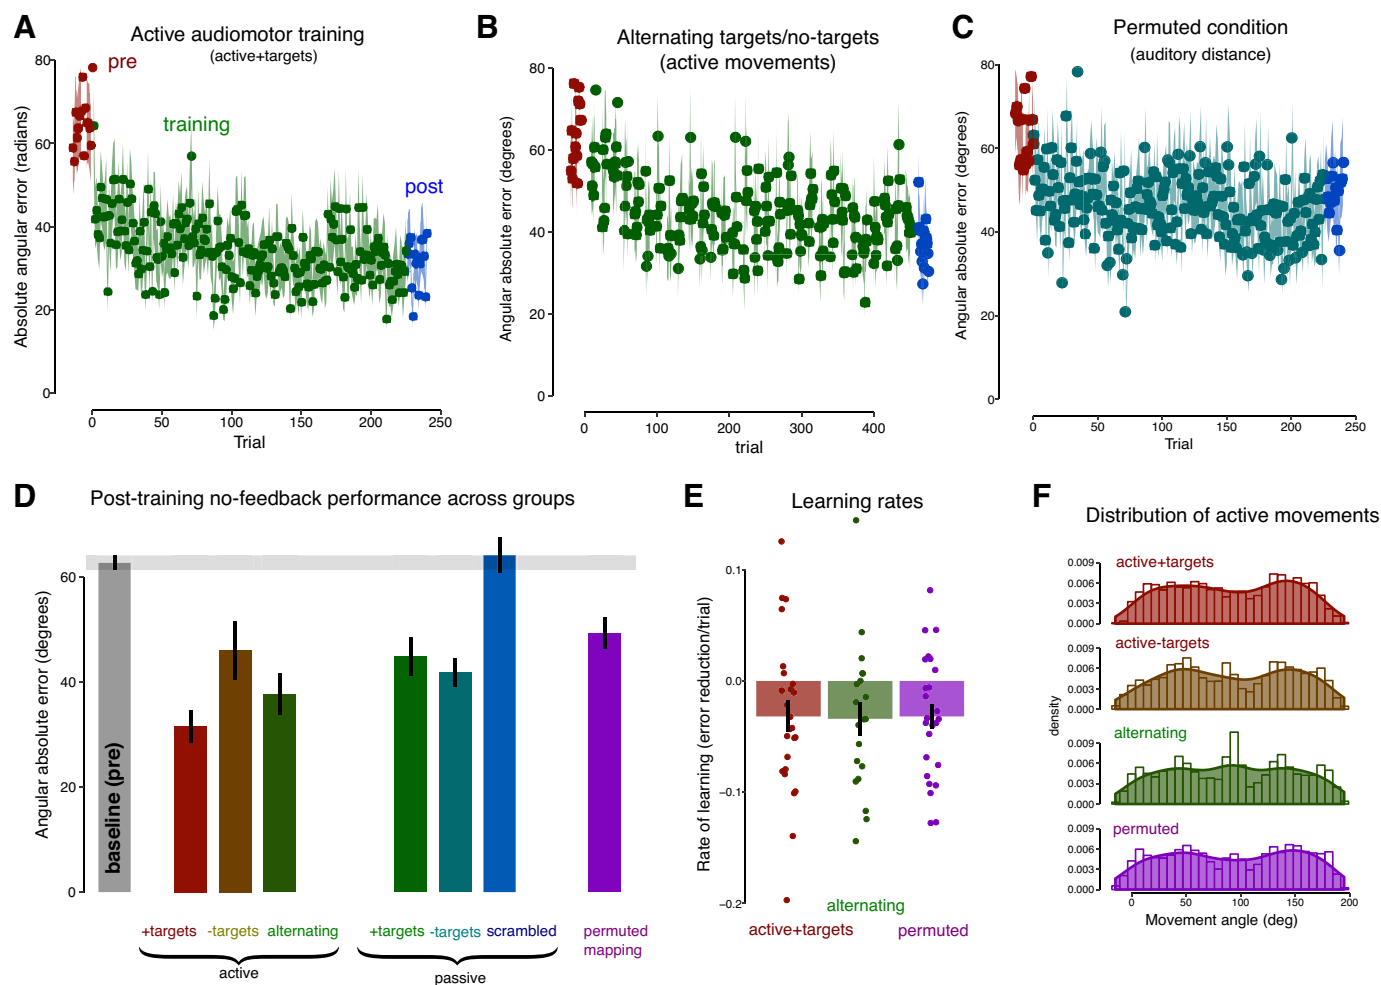

Fig. 2. *A*: participants learn to reach to auditory targets (active + targets group, $n = 25$). Before and after learning (green), all groups were tested using active no-feedback trials (red and blue). *B*: performance on target trials for the group where targets and feedback were presented in alternating fashion but never together for the same trial ($n = 20$). In between each set of target trials, these participants also performed no-target active reaching movements. *C*: participants ($n = 25$) were able to learn the permuted map. This learning was evident when the target-to-movement error was computed in auditory space but not when it was computed in motor space. This indicated that participants produced movements whose sounds were similar to the target sound instead of producing movements that were similar to the target movement. *D*: baseline performance was not different across groups and is shown combined in the gray bar. Colored bars indicate the performance of the various groups during the postactive no-feedback trials. All groups showed learning relative to baseline, except for the scrambled group. *E*: rate of learning as assessed during training. For each participant, a line was fit to the error during training and the slopes of these lines are shown to be the same across groups. *F*: as a control, the distribution of active movements was largely equated between the conditions. Error bars and shaded area indicate means ± SE.

In the primary analyses of learning, performance was assessed based on posttraining no-feedback trials, instead of pre- to postchanges, because the pretraining performance was random (by design) and analyses using pre- to posttrials were found to be less sensitive (see below).

An ANOVA was performed with the post-no-feedback error as a dependent variable and a between-subjects factor group. The difference between experimental conditions was significant [$F(6,158) = 8.50$, $P < 0.0001$, $\eta^2 = 0.24$]. Pairwise t tests (Bonferroni-corrected) revealed the following pattern of findings. The passive scrambled group showed greater error than all other groups (all $P < 0.04$), which confirms that learning is only possible when the link between movements and feedback is intact. The active and passive groups were not different (all pairwise comparisons $P > 0.1$ with or without targets), suggesting that the presence or absence of active motor outflow does not affect learning. The permuted group showed greater error than the active + targets group ($P = 0.008$), suggesting that the permuted mapping may be more difficult to learn, although learning was present in both conditions. There was a trend for the active group with targets to show lower error than the active group without targets ($P = 0.09$), which suggested that targets, and therefore potentially target error, could affect the amount of learning. The alternating condition served as a control for this, in which targets and feedback were both presented but never on the same trial, and therefore, error could not be computed. Learning in the alternating condition was not significantly different from either active condition ($P = 1.0$). This suggests that target error does not drive learning. Other than the comparisons mentioned above, no other differences between groups in this experiment were significant (all $P > 0.1$).

To investigate performance during training, learning rates were computed as follows. Errors were computed for each subject and each trial separately, and then a, line was fit using robust regression. Robust regression is a variant of ordinary least-squares regression, which uses an algorithm to estimate which data points are likely to be outliers and assigns these a lower weight so that they do not disproportionately affect the regression estimate (slope in our example) (Andersen 2008). For each subject, one slope value was thus obtained that reflected the amount of change in error over the course of training. These slopes were significantly different from zero across groups [$t(69) = 4.28$, $P < 0.001$] indicating that overall subjects reduced their reaching error over time and there were no significant differences between groups [$F(2,67) = 0.01$, $P = 0.99$]. Additionally, fitting exponential learning curves to the data indicated no difference in exponential learning rates between the groups ($P = 0.77$). This suggests that the three groups showed similar learning performance during training (Fig. 2E) and in particular that the permuted condition performed similarly to the nonpermuted groups. Note that performance during training can only be investigated for active groups for trials in which they receive targets, i.e., active + targets, alternating, and permuted groups.

The tests described above focused on a single day of training, in which relatively high levels of error were observed at the end of training. Three additional subjects were tested who each performed the learning task over the course of 21 training sessions, each on a separate day (with at most 3 intervening nontraining days and a maximum total duration of 5 wk, 225 trials per day, with targets and feedback). Improvements were seen throughout training (Fig. 3). This shows that learning an audiomotor map is a challenging task in which there are progressive improvements in performance over extended periods of training. Hence, the findings reported in the main experiments (single-day) reflect the earliest stages of sensorimotor map acquisition.

Four conditions involved active movements (active + targets, active − targets, alternating, and permuted). A potential confound was that differences between these groups, or absences thereof, could be driven by differences in the distribution of movements. However, there were no significant differences observed between the groups in average movement direction [$F(3,86) = 0.13$, $P = 0.94$], deviation from the workspace center [$F(3,86)) = 1.37$, $P = 0.26$], or standard deviation of the movement angles [$F(3,86) = 0.77$, $P = 0.52$; Fig. 2F]. This suggests that the different groups had comparable exposure to the sensorimotor workspace.

In the permuted condition, the map structure was divided into bins that were shuffled, thus removing the global structure but preserving local structure within each bin. It is possible in principle that subjects used this residual local structure for learning the mapping. However, two empirical findings suggest that this was not the case. In a first analysis, we reasoned that to use the local structure to drive learning, subsequent movements to same-bin targets would need to land in the same bin, because otherwise a different local structure would apply to the different movements. To probe how often this happened, targets were grouped according to the bin they fell in. Then, for each target bin, we looked at the series of each subject's movements to targets within the same bin, ignoring intervening trials to other bins. Our analysis was not restricted to cases where subsequent movements landed in the same bin. In fact, we identified pairs of trials where the targets were in the same bin, but in between these two trials there could be any number
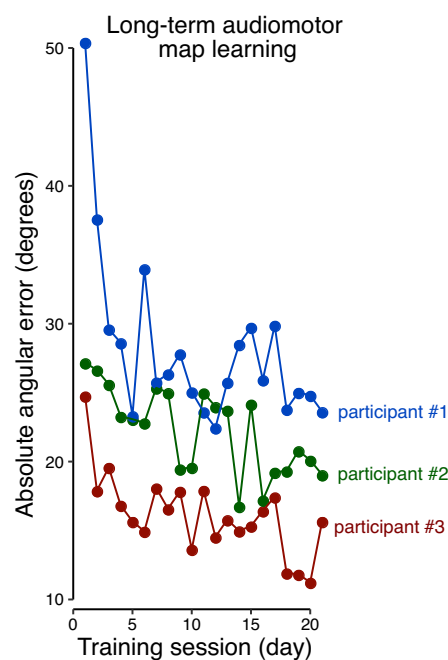

Fig. 3. Prolonged training measured in 3 participants leads to continual increases in performance.

of intervening trials to targets in different bins. For only 30.9% of all cases, two subsequent movements to the same target bin fell into the same bin (Fig. 4A). Note that chance level, if reaching is completely random, is 20%. This indicates that in the majority of trials where two targets were in the same bin, subsequent movements to these targets nevertheless land in different bins and therefore the local structure could not be used. In only a small portion of trials subsequent movements to similar targets land in the same bin, and therefore, the local structure of the map, within the bins, could actually be exploited. In a second, separate analysis, going back to all individual movements, we selected movements which landed in the correct bin, e.g., movements to a target in *bin 3* where the

subject correctly landed in *bin 3*. If subjects used the local structure of the map within this bin, we hypothesized that their movement angles should follow the local structure of this bin; that is, the movement directions should vary with the position of the target within the bin. However, there was no significant correlation between the target angle and the movement angle [for all subjects combined, $t(1438) = 1.11$, $P = 0.27$; see Fig. 4B]. Note that the average of the movement angles corresponds to the center of the bin, but the individual movements are distributed largely uniformly across the bin (Fig. 4C) and this pattern is stable over the course of training (Fig. 4D). Taken together, these two analyses suggest that the local structure that remained after the permutation was applied was not used in the learning.
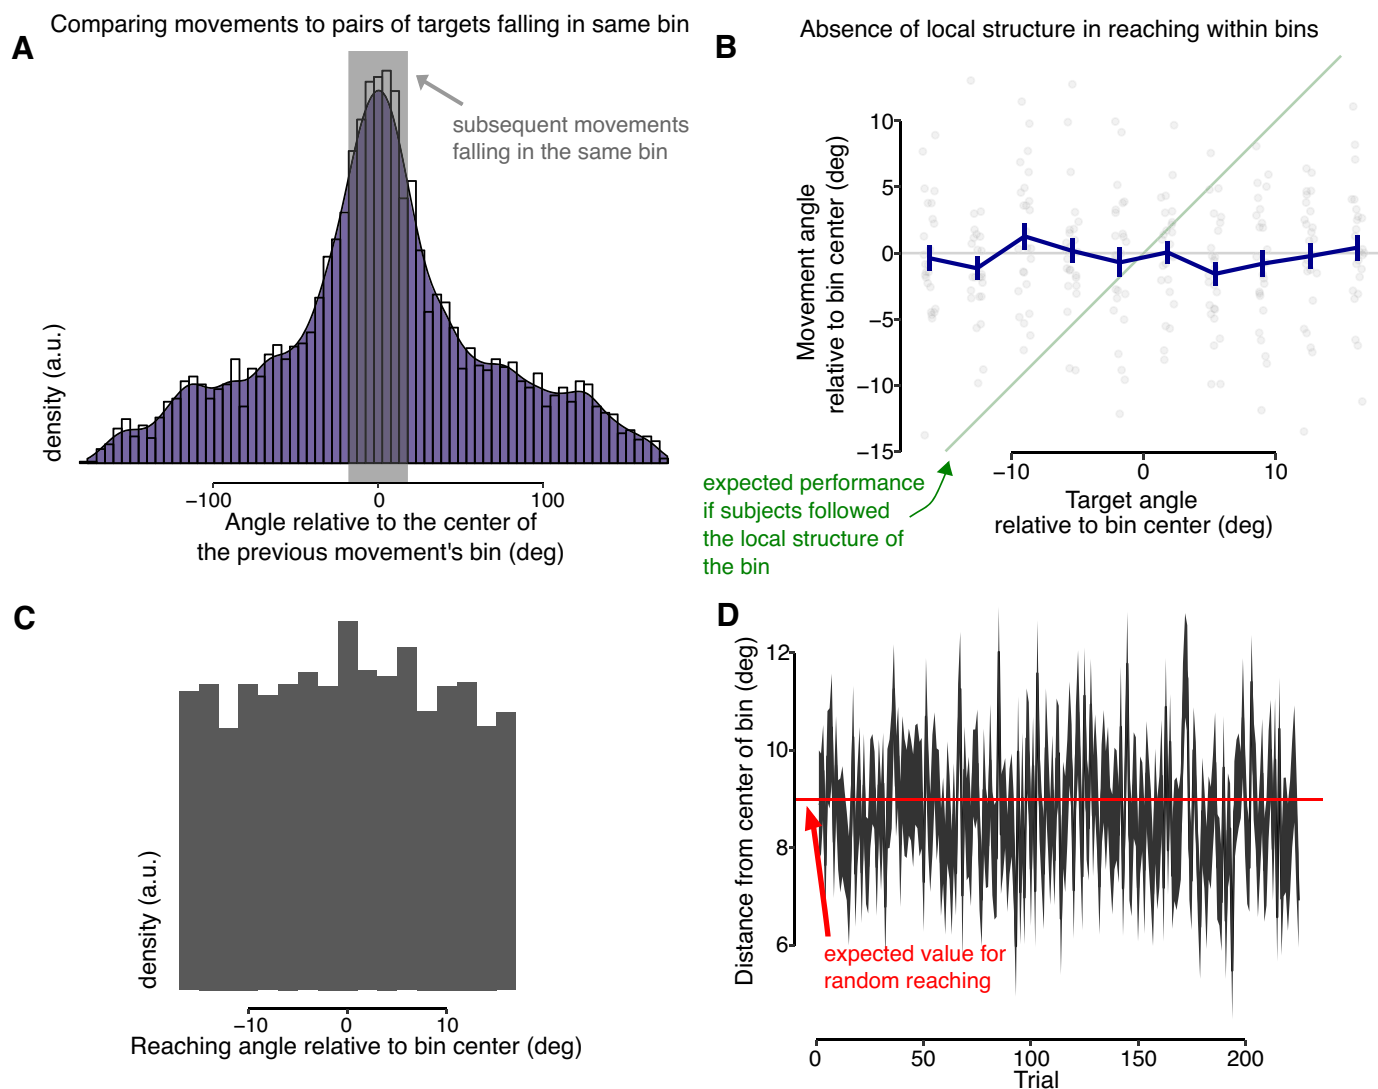

Fig. 4. The scale of the permutation corresponded roughly to the scale of the errors made by subjects in the task. *A*: distribution of movement angles relative to the center of the bin of the movement on the previous trial. The gray area indicates the width of the bin and represents subsequent movements that land in the same bin. This indicated that only in a small portion of trials did subjects subsequent movements to similar targets land in the same permutation bin; a.u., arbitrary units. *B*: if participants matched the local structure, then within each bin the reaching angle should increase with target angle. Here are indicated trials in which subjects landed in the correct bin. Each bin was then divided into 10 subbins (horizontal axis), and the subjects' average movements are indicated in gray dots (vertical axis) with averages (and SE) indicated in blue. Reaching angles did not increase with target angle, indicating that participant performance does not match the local structure within each bin. *C*: note that although subjects' average movements were in the center of the bin, the individual movements were distributed uniformly within each bin (here shown for all subjects combined; first 100 trials discarded for display). *D*: if subjects learn to reach to the center of the bins, this should reduce the distance between reach angle and the bin center, but that distance was stable over the course of training. Shown is average ± SE of the (absolute) distance between reach angle and the center of the bin. The red trace indicates the expected distance for uniform reaching (which is 9° for an 18° bin).

In the passive condition, the subjects were instructed to not make any movements themselves and let the robot guide them. However, in principle it is possible that subjects actively moved along with the robot. In that case, they potentially produced motor outflow and may have used this to drive learning. To test whether this was the case, in each passive trial, the force that was sent to the robot motors was logged, and at the same time, the forces acting on the handle were measured. If subjects produced active movements that precisely followed the movements of the handle, there should be little force recorded acting on the handle (similar to when subjects are not holding the handle at all, in which case forces should be near zero). If however subjects let themselves be moved by the robot, then a force produced by the robot motor should result in an observed force in the opposite direction,

caused by the inertia of the subject's arm (Fig. 5A). The empirical finding was that motor force was indeed strongly negatively correlated with the observed force on the handle (Fig. 5B). This correlation was robust in all subjects and did not differ between the groups [$F(2,70) = 1.49$, $P = 0.23$; Fig. 5C]. This is consistent with the idea that subjects do not actively move along with the robot. To further exclude that subjects are actively moving along with the robot, we performed the following analysis. The assumption was that if subjects were completely passive, the displacement we observe should be approximately equivalent to that of moving a point mass of ~3 kg (the arm is ~5% of body weight). If subjects are actively moving along, and based on the displacement, we estimated the mass that is being moved, it would seem less because of the assistive force. On each trial, the net force on the handle was
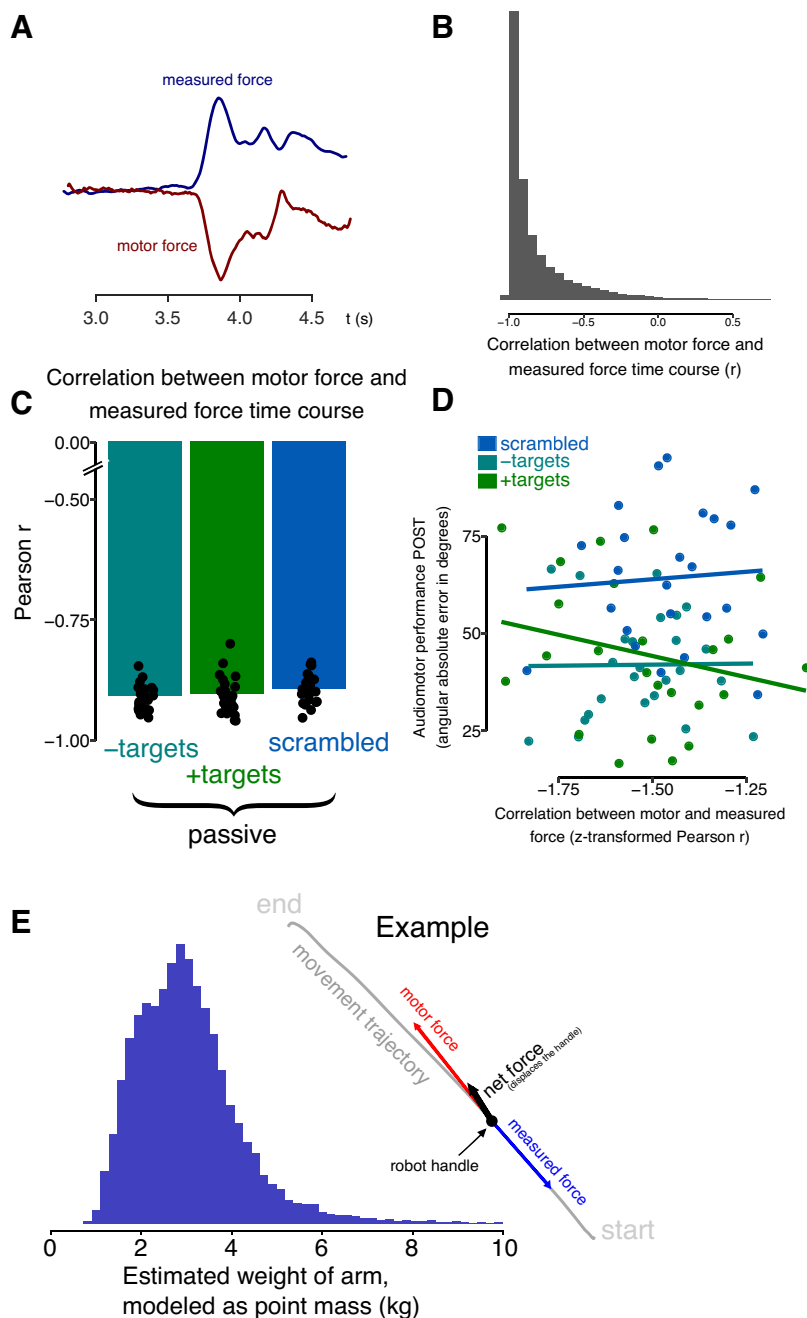

Fig. 5. Control analysis to confirm that passive subjects were not actively moving along with the robot handle. In each trial, forces on the handle were measured and compared with the force exerted by the motor to move the subject. *A*: time course of the motor force (red) and the measured force on the handle (blue) shows a strong negative correlation, as would be expected if the subject is not moving along with the handle. *B*: the distribution of correlations of all trials. *C*: a strong negative correlation was observed for all subjects and was not different between the passive groups. *D*: no relationship was observed between the motor versus observed force correlation and the amount of learning of the subject. *E*: to exclude that subjects actively moved along with the handle, we compared the programmed motor force (red arrow in inlay) and the measured force at the handle (blue arrow). The resulting net force (black arrow) was able to explain the observed displacement of a mass corresponding to roughly the weight of the subject arm, thus without needing to appeal to additional forces produced by the subject.

calculated as the sum of the commanded motor force and the force recorded by the force transducer on the robot handle (e.g., the sum of the traces in Fig. 5*A*). This net force is what causes the motion of the subjects' arm (Fig. 5*E*, *inset*). Modeling the arm as a variable point mass, one can then compute the predicted displacement, and subsequently fit for the mass that matches the actually observed displacement. The average estimated weight was 3.084 kg, which matches the weight of the arm (typically 5% of body weight). In sum, the observed forces are able to account for the observed handle displacement. An additional control analysis was performed based on the idea that if in spite of the evidence above, marginal active movements would occur in the passive groups, and these would contribute to learning, then the closer the motor-to-observed force correlation is to zero (indicative of more active contribution by the subject), the more subjects should learn. That is, the motor-to-observed force correlation during training should predict the amount of learning measured during the no-feedback trials. However, there was no significant correlation with learning (passive + targets: $r = -0.24$, $P = 0.25$; passive − targets: $r = 0.01$, $P = 0.96$; and scrambled: $r = 0.07$, $P = 0.75$; Fig. 5*D*).

Before and after the audiomotor testing, participants completed auditory discrimination testing. Response accuracy was on average 72.9% (SD 6.9) before audiomotor learning and 73.4% (SD 6.9) afterwards. This difference was not significant [$F(1,157) = 1.41$, $P = 0.24$], and there was no significant effect of group [$F(6,157) = 0.50$, $P = 0.81$] or interaction [$F(6,157) = 1.03$, $P = 0.41$]. To investigate whether auditory functioning could explain reaching performance, correlations were computed between audiomotor no-feedback trials after training and change in auditory test performance, which were not significant ($r = -0.06$, $P = 0.48$ for all subjects combined excluding the scrambled group, and carrying this out per group all correlation magnitudes $r < 0.52$, corrected $P > 0.12$). This showed that basic auditory functioning remained stable during this experiment and learning could not be explained by an improvement in auditory perceptual function alone (van Vugt and Ostry 2018a, 2018b).

Participants also performed a motor copy test to measure proprioceptive and motor function in isolation (without auditory stimulation), in which the robot moved their hand to a location on the target circle and then to the starting point and their task was to move back to the indicated location (all in the absence of vision). Reaching error in this test was 5.41° (SD 1.5) before training and 5.13° (SD 1.4) after, and there was a statistical trend for this improvement to be significant [$F(1,157) = 3.90$, $P = 0.05$, $\eta^2 = 0.009$], but there was no significant overall difference between groups [$F(6,157) = 0.84$, $P = 0.54$] and no difference in amount of change per group [interaction $F(6,157) = 0.41$, $P = 0.87$]. In particular, the active with targets group did not show significant improvement [$F(1,43) = 0.37$, $P = 0.55$]. This suggested that although there may be a small improvement in reaching to proprioceptive targets, amounting to less than half a degree, this improvement could not explain the differences in audiomotor learning between the groups observed here. A possible reason why there was no proprioceptive learning here was that the robot's inertia was compensated, and thus subjects did not need to learn this inertia.

The main analyses described above compared learning between the various conditions using the no-feedback performance after training (post) instead of a change score from pre to post. The reason was that the pretraining no-feedback trials were included only to ensure that subjects had no prior information about the mapping, by showing that reaching was essentially random and that there were no biases. Pre-no-feedback trials are random by design because subjects have no way of knowing the arbitrary location-to-sound mapping. To confirm that pretrials were random, first, reaching error was computed for 1,000 permutations of the target angles and the subject's movements, and the resulting errors were not different from those actually observed empirically [$t(197.35) = -1.37$, $P = 0.17$]. Second: reaching error on pre-no-feedback trials did not differ between the groups [$F(6,158) = 1.03$, $P = 0.41$, $\eta^2 = 0.04$]. Third, pre- versus postcomparisons are based on the idea that individual differences between subjects are subtracted out and only within-subject differences remain. However, there was no significant correlation between the pre- and postscores of the individual subjects ($r = 0.03$, $P = 0.72$), suggesting that pre- versus postcomparisons do not subtract individual differences but rather subtract a random value (the prescores) and thereby actually add noise. These three arguments show that the baseline no-feedback performance is essentially random, in line with previous studies (van Vugt and Ostry 2018a, 2018b), which confirmed that the learning task successfully studies sensorimotor map learning de novo. The main analyses in this paper therefore focused on the postmeasure in isolation as a measure of learning. The reason for that was that we expected that if the prescores are random, a pre- to poststatistical comparison analysis would be corrupted by additional random noise (the prescores), making it less powerful to detect subtle differences between the groups and thus unfairly bias the outcome toward our own hypotheses. To verify this, all the group comparisons were repeated also for the pre- to postcomparison, which indeed detected only a subset of the significant group differences of the post-only analysis reported above. The pre- to postimprovement scores were different across groups [$F(6,158) = 4.65$, $P = 0.0002$). Pairwise *t* tests between all groups revealed only two pairs of groups to be significantly different: scrambled versus active + targets ($P < .0001$) and scrambled versus alternating ($P = 0.004$). None of the other pairs of groups showed significant differences in learning. This shows that the post-only comparison is more sensitive to group differences than the pre- to postanalysis and therefore the most stringent test of our own hypotheses. Note that although we found no differences in initial performance between the groups, this does not rule out that due to sampling error subjects in some groups had greater learning ability at the outset.

Additional control analyses revealed that across groups, there were no significant differences in amount of learning between male and female participants [$F(1,126) = 0.08$, $P = 0.78$; excluding data from the scrambled condition where learning was precluded by design]. Similarly, we compared performance between participants who reported no musical training, informal musical experience, or having played before (>5 yr ago) and found no effect on learning [$F(2,137) = 0.08$, $P = 0.92$]. This suggests that gender or musical experience did not confound the results reported in the paper.

DISCUSSION

A first challenge facing a sensorimotor learner is to understand the link between movements and their sensory outcomes, for example, learning what are the visual effects of arm reaching movements or which sounds result from which vocal tract configuration. The present study investigated how sensorimotor maps are first acquired. Specifically, do the early stages of map learning require active motor outflow? Does learning assume that sensory and motor distances are preserved by the mapping? Does map learning rely on comparing movement feedback with an overt target (error correction)? Subjects were tested in various conditions that each removed one or several of these sources of information. The key result was that in all conditions except for a scrambled condition, learning was observed and was similar in magnitude. This showed that maps can be learned using free exploration without targets and they can be learned passively. In sum, as long as sounds are consistently linked to somatosensory states, the mapping can be learned, suggesting a robust learning mechanism that can operate with minimal requirements. Even a mapping that does not preserve distances across motor and sensory modalities can be learned, showing that the learner does not need to make assumptions about the structure of the sensorimotor map and instead learning can proceed through the acquisition of a collection of individual sensory-motor instances.

In the present data set, participants learned a permuted mapping that was discontinuous and therefore did not preserve distances in sensory and motor spaces. This finding shows that at least in the early stages of learning, the learners do not need to make assumptions about the structure of the mapping that they encounter. To our knowledge, previous work in motor learning has not investigated learning of discontinuous mappings but is mostly based on mappings that preserve topology (Ritter et al. 1989). What comes closest are studies where humans learned an arbitrary, artificial mapping to control a visual cursor on a screen using finger movements (Mosier et al. 2005), which constitutes a complex but continuous mapping. The fact that a permuted mapping can be learned and is learned at the same rate as a nonpermuted mapping (Fig. 2E) suggests that sensorimotor mappings are stored not as a singular entity but rather as a patchwork set of distinct informational units on which learning can operate more or less independently (reminiscent of previous proposals such as Wolpert and Kawato 1998). Sensory-to-motor information is tracked largely independently in these local portions of the sensorimotor workspace, because if not, the error experienced for a movement in one permutation bin would be used to (erroneously) update the mapping in another bin (refer to Fig. 1D) and learning would be erratic. This finding is in line with studies in sensorimotor adaptation showing that what is learned in one portion of the workspace has little effect (generalization) in distant parts (Brayanov et al. 2012; Gandolfo et al. 1996; Joiner et al. 2013; Mattar and Ostry 2007, 2010; Wu and Smith 2013) and that multiple transformations can be learned simultaneously (Rochet-Capellan and Ostry 2011). That is, the sensorimotor mapping can undergo changes locally that largely do not cascade into global changes across the workspace. A similar phenomenon was observed previously in the context of learning novel audiomotor maps, where instead of interpolating to unknown movements, subjects proportionally reproduced

known movements, which is reminiscent of instance-based learning (van Vugt and Ostry 2018b). How exactly in each local neighborhood incoming feedback information is used to update the mapping remains an open question; this could involve either a proportionate correction (sensory prediction error) or replacing the previous entry with the current feedback, a strategy that has previously been suggested for machine learning (Aha et al. 1991; Atkeson et al. 1997b; Schaal and Atkeson 1998). The permuted mapping result is also similar to findings that subjects can learn multiple, even opposing perturbations in parallel, but in those cases contextual cues are typically provided (Howard et al. 2010; Imamizu et al. 2007; Lee and Schweighofer 2009) and interference tends to be observed between the perturbations. The permuted mapping shows that in the initial stages of learning complex sensorimotor contingencies, the brain can track different portions of the sensorimotor workspace largely independently. This finding is surprising because it would seem adaptive for the motor system to restrict the search space for sensorimotor maps to those mappings that are physiologically plausible, i.e., preserve distances, and a similar idea has been proposed for neural interfaces (Sadtler et al. 2014). The present data show that although permuted mappings can be learned, learning may be less than in a nonpermuted condition (active with targets). However, the permuted group shows learning and indeed the rate at which they learn is not different from the active nonpermuted groups. The permutation that was applied to the mapping disrupted the global structure of the maps (across bins) but retained some local structure (within bins). In principle, the local structure could have been responsible for learning gains observed here. However, it was found that subsequent movements to similar targets in only a small portion of cases landed in the same bin; therefore, in the majority of cases, the local structure could not be used (Fig. 4). Moreover, it was found that the reaching pattern within each bin did not follow the local structure. These two findings suggest that the residual local structure did not drive the learning observed here.

Active motor outflow appears not to be necessary for sensorimotor map learning. Participants in the present study showed learning when they were passively presented with the movements and sounds experienced by a yoked active participant. Some previous work shows a benefit of somatosensory experience for learning (Beste and Dinse 2013). Previous work also demonstrated visuomotor adaptation learning can occur with passive arm movements alone although learning is less than in active conditions (Lei et al. 2016). Passive movements with rotated visual feedback are furthermore accompanied by sensory change of similar magnitude to that typically observed with active motor learning (Mostafa et al. 2019). Only few previous studies showed passive learning was possible to the same extent as active learning (Bernardi et al. 2015), and many studies instead show little or no learning (Beets et al. 2012; Held and Hein 1963; Lotze et al. 2003). Here, how could passive exposure improve subsequent active movements to sounds? The present data suggest that rather than exploiting the experienced links between active movements and auditory feedback, passive subjects based their learning on the link between auditory feedback and the somatosensory trace of the movement that was presented to them passively. Then when tested on active movements after training, these participants could use the somatosensory memory trace to reproduce the

same movement actively, a capacity widely observed previously (Fuentes and Bastian 2010). Active participants, during their learning trials, may have similarly used their somatosensory information instead of motor outflow during map formation so that sensorimotor maps are stored as sensory-to-sensory links instead of sensory-to-motor links. This aligns with recent work that suggests that sensorimotor adaptation applies to the actually performed movement rather than the planned movement (Gonzalez Castro et al. 2011); because in that case the only source of information about the actually experienced movement is the sensory inflow rather than the motor outflow. However, classical models of sensorimotor control typically start from the idea that what is learned is the relation between motor commands and sensory effects (Ghahramani et al. 1997; Wolpert et al. 2011). Similarly, models of sensorimotor learning in robotics have investigated map formation using self-generated exploratory movements such as exhibited by participants here (Saegusa et al. 2009), but to our knowledge, these are based on active motor outflow. A likely reason is that in these models, the robot's task is to learn the correlation between movements and somatosensory feedback. The apparent discrepancy between these existing ideas and the present data could be resolved by postulating that subjects learn sensory-to-somatosensory associations (in our case auditory-to-somatosensory) and when required to produce movements, these somatosensory traces are mapped to motor commands. For example, a subject in the active condition can keep track of the somatosensory experience of their own movements linked with the sounds they hear and thus form an auditory-to-somatosensory map. When they are subsequently asked to make active movements to targets (in the posttraining no-feedback trials), they can use the auditory-to-somatosensory map to generate a desired somatosensory trace, and then use their existing somatosensory-to-motor mapping to generate their movements. In sum, using those two mappings could yield auditory-to-motor links required for performance observed here. The fact that the somatosensory-to-motor map is preexisting and does not undergo changes is supported by the observation that subjects can perform the motor copy task before training and that their performance does not change substantially after training. A related question is that if learning in the active condition is based on linking somatosensory information to auditory feedback, and supposing that incoming sensory information is less reliable during active movement (Chapman et al. 1987; Collins et al. 1998), one might expect that there would be more learning in a passive condition. However, the passive groups did not show better learning during the no-feedback test phase in this experiment.

Sensorimotor map learning is also found to be possible without target error correction, because subjects were able to learn the mapping in a condition where no targets were presented during the learning phase. Previous studies approached the question whether error drives learning by creating conditions where errors were zero, for example, by giving subjects a perturbation coupled with an explicit strategy so that they produced no target error at the outset (Mazzoni and Krakauer 2006). Under such conditions, sensorimotor changes are still observed, which are thought to be driven by a discrepancy between predicted and actual sensory feedback, i.e., sensory prediction error (Shadmehr et al. 2010; Wolpert and Flanagan 2001). Sensory prediction error learning was possible in prin-

ciple in the present study as well, even without overt targets, because subjects could, before each movement, form a prediction of the auditory feedback, compare it with the feedback actually received, and then apply a proportionate correction. Learning was observed also in the permuted condition, where the sensorimotor mapping topology was discontinuous because it did not preserve distances in sensory and motor spaces (Fig. 1D). Sensory prediction error in this case could still be used, provided that it operates exclusively in the sensory domain, because there are no discontinuities in the auditory modality itself. In sum, the present data show that target error does not drive learning. The data do not rule out sensory-prediction error as a potential mechanism for learning. The data also do not force the conclusion that sensory-prediction error drives learning. An account of learning observed here based on sensory prediction error requires the assumption that subjects can correctly create a weighted combination of two sounds (the predicted sound and the actual feedback), which would be used as the new prediction for a given movement. The present result is in line with previous work in visuomotor adaptation that performance error does not drive learning (Lee et al. 2018).

Participants who received no targets at all showed a statistical trend toward impaired learning (relative to the group with targets), whereas this trend was not found in an alternating group who received intermittent trials with targets but no feedback. This suggests that the intermittent target trials had improved learning. The target trials themselves did not provide the subjects with information about the mapping because only targets were presented, not feedback, and the scrambled condition demonstrates that targets alone are not enough for learning but instead learning requires that feedback is linked in a systematic way to movements. Instead, the performance improvement of the alternating group (relative to the group that did not receive targets) could be due to the fact that they were required to query their sensorimotor map when targets were being presented. This idea is in line with a class of machine learning algorithms that defer processing of input data until a query is presented (Atkeson et al. 1997a), that is, until a targeted movement is required. Alternatively it could simply be that the presence of intermittent targets engaged attentional mechanisms that facilitated learning in general.

The present work tested participants learning a novel mapping between movements and sounds and therefore the conclusions pertain to the auditory-motor modalities. The learning process investigated here may be the same or different when learning mappings from movement to other sensory modalities (somatosensory, visual). Within the auditory domain, there is previous work examining the updating of auditory-motor representations (Kagerer and Contreras-Vidal 2009; Schmitz and Bock 2017), but these studies focus on the acoustic localization of sounds, which presumably relies on different processes than studied here because sounds used in the present study were not localized in space and differed only by their frequency contents.

Participants in the current study are in an usual situation because they have to build a sensorimotor map instead of having one at the outset. This is confirmed by their random performance at the outset, in line with previous studies using this setup (van Vugt and Ostry 2018a, 2018b) but contrasting with typical motor control paradigms where subjects, for example, reach to visual targets with rotated

visual feedback or experience force fields, where performance at the outset is biased (by the perturbation) and not random (Krakauer 2009). Participants in the current study were in much the same situation as an infant learning to speak by making predominantly random vocal movements thought to serve exploration of sensory-motor contingencies (Kuhl 2004; Locke 1989; von Hofsten 1982; Werker and Tees 1999). Some robotic learning architectures also use random movements to map out the sensory effects of motor commands (Schillaci and Hafner 2011). Due to the initial random performance, presumably subjects in the present study do not have an auditory-motor map at the outset and have to form one in the course of learning, whereas in previous studies an existing sensorimotor mapping is being updated. The key result from the present study is that this early phase of learning where sensorimotor maps are first formed does not require target error correction, active motor outflow, or continuous, topology-preserving mappings and instead shows that subjects may use a robust, instance-based learning mechanism.

ACKNOWLEDGMENTS

GRANTS

DISCLOSURES

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

AUTHOR CONTRIBUTIONS

F.T.v.V. and D.J.O. conceived and designed research; F.T.v.V. performed experiments; F.T.v.V. analyzed data; F.T.v.V. and D.J.O. interpreted results of experiments; F.T.v.V. prepared figures; F.T.v.V. drafted manuscript; F.T.v.V. and D.J.O. edited and revised manuscript; F.T.v.V. and D.J.O. approved final version of manuscript.

REFERENCES

Aha DW, Kibler D, Albert MK. Instance-based learning algorithms. *Mach Learn* 6: 37–66, 1991. doi:10.1007/BF00153759.

Andersen R. Modern methods for robust regression. In: *Quantitative Applications in the Social Sciences*. Los Angeles, CA: Sage, 2008, no. 07-152. OCLC: ocm85833319.

Atkeson CG, Moore AW, Schaal S. Locally weighted learning. *Artif Intell Rev* 11: 11–73, 1997a. doi:10.1023/A:1006559212014.

Atkeson CG, Moore AW, Schaal S. Locally weighted learning for control. In: *Lazy Learning*, edited by Aha DW. Dordrecht, The Netherlands: Springer, 1997b, p. 75–113.

Beets IA, Macé M, Meesen RLJ, Cuypers K, Levin O, Swinnen SP. Active versus passive training of a complex bimanual task: is prescriptive proprioceptive information sufficient for inducing motor learning? *PLoS One* 7: e37687, 2012. doi:10.1371/journal.pone.0037687.

Bernardi NF, Darainy M, Ostry DJ. Somatosensory contribution to the initial stages of human motor learning. *J Neurosci* 35: 14316–14326, 2015. doi:10.1523/JNEUROSCI.1344-15.2015.

Beste C, Dinse HR. Learning without training. *Curr Biol* 23: R489–R499, 2013. doi:10.1016/j.cub.2013.04.044.

Brayanov JB, Press DZ, Smith MA. Motor memory is encoded as a gain-field combination of intrinsic and extrinsic action representations. *J Neurosci* 32: 14951–14965, 2012. doi:10.1523/JNEUROSCI.1928-12.2012.

Butcher PA, Taylor JA. Decomposition of a sensory prediction error signal for visuomotor adaptation. *J Exp Psychol Hum Percept Perform* 44: 176–194, 2018. doi:10.1037/xhp0000440.

Chapman CE, Bushnell MC, Miron D, Duncan GH, Lund JP. Sensory perception during movement in man. *Exp Brain Res* 68: 516–524, 1987. doi:10.1007/BF00249795.

Collins DF, Cameron T, Gillard DM, Prochazka A. Muscular sense is attenuated when humans move. *J Physiol* 508: 635–643, 1998. doi:10.1111/j.1469-7793.1998.00635.x.

Diedrichsen J, White O, Newman D, Lally N. Use-dependent and error-based learning of motor behaviors. *J Neurosci* 30: 5159–5166, 2010. doi:10.1523/JNEUROSCI.5406-09.2010.

Fuentes CT, Bastian AJ. Where is your arm? Variations in proprioception across space and tasks. *J Neurophysiol* 103: 164–171, 2010. doi:10.1152/jn.00494.2009.

Gandolfo F, Mussa-Ivaldi FA, Bizzi E. Motor learning by field approximation. *Proc Natl Acad Sci USA* 93: 3843–3846, 1996. doi:10.1073/pnas.93.9.3843.

Ghahramani Z, Wolpert DM, Jordan MI. Computational models of sensorimotor integration. *Adv Psychol* 119: 117–147, 1997. doi:10.1016/S0166-4115(97)80006-4.

Gonzalez Castro LN, Monsen CB, Smith MA. The binding of learning to action in motor adaptation. *PLOS Comput Biol* 7: e1002052, 2011. doi:10.1371/journal.pcbi.1002052.

Guenther FH, Ghosh SS, Tourville JA. Neural modeling and imaging of the cortical interactions underlying syllable production. *Brain Lang* 96: 280–301, 2006. doi:10.1016/j.bandl.2005.06.001.

Haith AM, Krakauer JW. Model-based and model-free mechanisms of human motor learning. *Adv Exp Med Biol* 782: 1–21, 2013. doi:10.1007/978-1-4614-5465-6_1.

Held R, Hein A. Movement-produced stimulation in the development of visually guided behavior. *J Comp Physiol Psychol* 56: 872–876, 1963. doi:10.1037/h0040546.

Howard IS, Ingram JN, Wolpert DM. Context-dependent partitioning of motor learning in bimanual movements. *J Neurophysiol* 104: 2082–2091, 2010. doi:10.1152/jn.00299.2010.

Imamizu H, Sugimoto N, Osu R, Tsutsui K, Sugiyama K, Wada Y, Kawato M. Explicit contextual information selectively contributes to predictive switching of internal models. *Exp Brain Res* 181: 395–408, 2007. doi:10.1007/s00221-007-0940-1.

Joiner WM, Brayanov JB, Smith MA. The training schedule affects the stability, not the magnitude, of the interlimb transfer of learned dynamics. *J Neurophysiol* 110: 984–998, 2013. doi:10.1152/jn.01072.2012.

Kagerer FA, Contreras-Vidal JL. Adaptation of sound localization induced by rotated visual feedback in reaching movements. *Exp Brain Res* 193: 315–321, 2009. doi:10.1007/s00221-008-1630-3.

Krakauer JW. Motor learning and consolidation: the case of visuomotor rotation. *Adv Exp Med Biol* 629: 405–421, 2009. doi:10.1007/978-0-387-77064-2_21.

Kuhl PK. Early language acquisition: cracking the speech code. *Nat Rev Neurosci* 5: 831–843, 2004. doi:10.1038/nrn1533.

Lee JY, Schweighofer N. Dual adaptation supports a parallel architecture of motor memory. *J Neurosci* 29: 10396–10404, 2009. doi:10.1523/JNEUROSCI.1294-09.2009.

Lee K, Oh Y, Izawa J, Schweighofer N. Sensory prediction errors, not performance errors, update memories in visuomotor adaptation. *Sci Rep* 8: 16483, 2018. doi:10.1038/s41598-018-34598-y.

Lei Y, Bao S, Wang J. The combined effects of action observation and passive proprioceptive training on adaptive motor learning. *Neuroscience* 331: 91–98, 2016. doi:10.1016/j.neuroscience.2016.06.011.

Leow LA, Marinovic W, de Rugy A, Carroll TJ. Task errors contribute to implicit aftereffects in sensorimotor adaptation. *Eur J Neurosci* 48: 3397–3409, 2018. doi:10.1111/ejn.14213.

Locke JL. Babbling and early speech: continuity and individual differences. *First Lang* 9: 191–205, 1989. doi:10.1177/014272378900900606.

Lotze M, Braun C, Birbaumer N, Anders S, Cohen LG. Motor learning elicited by voluntary drive. *Brain* 126: 866–872, 2003. doi:10.1093/brain/awg079.

Mattar AA, Gribble PL. Motor learning by observing. *Neuron* 46: 153–160, 2005. doi:10.1016/j.neuron.2005.02.009.

Mattar AA, Ostry DJ. Modifiability of generalization in dynamics learning. *J Neurophysiol* 98: 3321–3329, 2007. doi:10.1152/jn.00576.2007.

Mattar AA, Ostry DJ. Generalization of dynamics learning across changes in movement amplitude. *J Neurophysiol* 104: 426–438, 2010. doi:10.1152/jn.00886.2009.

Mazzoni P, Krakauer JW. An implicit plan overrides an explicit strategy during visuomotor adaptation. *J Neurosci* 26: 3642–3645, 2006. doi:10.1523/JNEUROSCI.5317-05.2006.

Mosier KM, Scheidt RA, Acosta S, Mussa-Ivaldi FA. Remapping hand movements in a novel geometrical environment. *J Neurophysiol* 94: 4362–4372, 2005. doi:10.1152/jn.00380.2005.

Mostafa AA, Hart Bt, Henriques DY. Motor learning without moving: proprioceptive and predictive hand localization after passive visuoproprioceptive discrepancy training. *PLoS One* 14: e0221861, 2019. doi:10.1371/journal.pone.0221861.

Mussa-Ivaldi FA, Casadio M, Danziger ZC, Mosier KM, Scheidt RA. Sensory motor remapping of space in human-machine interfaces. In: *Progress in Brain Research. Enhancing Performance for Action and Perception Multisensory Integration, Neuroplasticity and Neuroprosthetics, Part 1*, edited by Green A, Chapman CE, Kalaska J, Leopre F. Amsterdam, The Netherlands: Elsevier, 2011, vol 191, chapt. 3, p. 45–64.

Reichenthal M, Avraham G, Karniel A, Shmuelof L. Target size matters: target errors contribute to the generalization of implicit visuomotor learning. *J Neurophysiol* 116: 411–424, 2016. doi:10.1152/jn.00830.2015.

Ritter HJ, Martinetz TM, Schulten KJ. Topology-conserving maps for learning visuo-motor-coordination. *Neural Netw* 2: 159–168, 1989. doi:10.1016/0893-6080(89)90001-4.

Robinson DW, Dadson RS. A re-determination of the equal-loudness relations for pure tones. *Br J Appl Phys* 7: 166–181, 1956. doi:10.1088/0508-3443/7/5/302.

Rochet-Capellan A, Ostry DJ. Simultaneous acquisition of multiple auditory-motor transformations in speech. *J Neurosci* 31: 2657–2662, 2011. doi:10.1523/JNEUROSCI.6020-10.2011.

Sadtler PT, Quick KM, Golub MD, Chase SM, Ryu SI, Tyler-Kabara EC, Yu BM, Batista AP. Neural constraints on learning. *Nature* 512: 423–426, 2014. doi:10.1038/nature13665.

Saegusa R, Metta G, Sandini G, Sakka S. 2009 Active motor babbling for sensorimotor learning. In: *2008 IEEE International Conference on Robotics and Biomimetics*. Bangkok, Thailand, 2009, p. 794–799.

Schaal S, Atkeson CG. Constructive incremental learning from only local information. *Neural Comput* 10: 2047–2084, 1998. doi:10.1162/089976698300016963.

Schillaci G, Hafner VV. Random movement strategies in self-exploration for a humanoid robot. In: *Proceedings of the 6th International Conference on Human-Robot Interaction*. New York, 2011, p. 245–246.

Schmitz G, Bock OL. Properties of intermodal transfer after dual visuo- and auditory-motor adaptation. *Hum Mov Sci* 55: 108–120, 2017. doi:10.1016/j.humov.2017.08.006.

Seidler RD, Kwak Y, Fling BW, Bernard JA. Neurocognitive mechanisms of error-based motor learning. *Adv Exp Med Biol* 782: 39–60, 2013. doi:10.1007/978-1-4614-5465-6_3.

Shadmehr R, Smith MA, Krakauer JW. Error correction, sensory prediction, and adaptation in motor control. *Annu Rev Neurosci* 33: 89–108, 2010. doi:10.1146/annurev-neuro-060909-153135.

Taylor JA, Ivry RB. Flexible cognitive strategies during motor learning. *PLOS Comput Biol* 7: e1001096, 2011. doi:10.1371/journal.pcbi.1001096.

van Vugt FT, Ostry DJ. The structure and acquisition of sensorimotor maps. *J Cogn Neurosci* 30: 290–306, 2018a. doi:10.1162/jocn_a_01204.

van Vugt FT, Ostry DJ. From known to unknown: moving to unvisited locations in a novel sensorimotor map. *Ann N Y Acad Sci* 1423: 368–377, 2018b. doi:10.1111/nyas.13608.

von Hofsten C. Eye-hand coordination in the newborn. *Dev Psychol* 18: 450–461, 1982. doi:10.1037/0012-1649.18.3.450.

Werker JF, Tees RC. Influences on infant speech processing: toward a new synthesis. *Annu Rev Psychol* 50: 509–535, 1999. doi:10.1146/annurev.psych.50.1.509.

Wolpert DM, Diedrichsen J, Flanagan JR. Principles of sensorimotor learning. *Nat Rev Neurosci* 12: 739–751, 2011. doi:10.1038/nrn3112.

Wolpert DM, Flanagan JR. Motor prediction. *Curr Biol* 11: R729–R732, 2001. doi:10.1016/S0960-9822(01)00432-8.

Wolpert DM, Ghahramani Z. Computational principles of movement neuroscience. *Nat Neurosci* 3, Suppl: 1212–1217, 2000. doi:10.1038/81497.

Wolpert DM, Ghahramani Z, Jordan MI. An internal model for sensorimotor integration. *Science* 269: 1880–1882, 1995. doi:10.1126/science.7569931.

Wolpert DM, Kawato M. Multiple paired forward and inverse models for motor control. *Neural Netw* 11: 1317–1329, 1998. doi:10.1016/S0893-6080(98)00066-5.

Wu HG, Smith MA. The generalization of visuomotor learning to untrained movements and movement sequences based on movement vector and goal location remapping. *J Neurosci* 33: 10772–10789, 2013. doi:10.1523/JNEUROSCI.3761-12.2013.