

From known to unknown: moving to unvisited locations in a novel sensorimotor map

1921

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

¹Department of Psychology, McGill University, Montreal, Quebec, Canada. ²Haskins Laboratories, New Haven, Connecticut

Address for correspondence: Floris T. van Vugt, Department of Psychology, McGill University, 2001 McGill College, Montreal, QC H3A 1G1, Canada. floris.vanvugt@mail.mcgill.ca

Sensorimotor learning requires knowledge of the relationship between movements and sensory effects: a sensorimotor map. Generally, these mappings are not innate but have to be learned. During learning, the challenge is to build a continuous map from a set of discrete observations, that is, predict locations of novel targets. One hypothesis is that the learner linearly interpolates among discrete observations that are already in the map. Here, this hypothesis is tested by exposing human subjects to a novel mapping between arm movements and sounds. Participants were passively moved to the edges of the workspace receiving the corresponding sounds and then were presented intermediate sounds and asked to make movements to locations they thought corresponded to those sounds. It is observed that average movements roughly match linear interpolation of the space. However, the actual distribution of participants' movements is best described by a bimodal reaching strategy in which they move to one of two locations near the workspace edge where they had prior exposure to the sound–movement pairing. These results suggest that interpolation happens to a limited extent only and that the acquisition of sensorimotor maps may not be driven by interpolation but instead relies on a flexible recombination of instance-based learning.

Keywords: motor learning; sensorimotor mapping; generalization; music; reaching movements; speech

Introduction

When first learning to talk or play a musical instrument, an important challenge is to learn which movement produces which sound, that is, a sensorimotor map. There is substantial knowledge about how maps, once learned, are adjusted, such as when adapting to a visuomotor rotation or novel force dynamics. Typically, in these settings, subjects are studied that already have fully formed sensorimotor maps. Experimental paradigms involve perturbations to these maps, which result in adaptation. However, less is known about the earliest phases of learning in which these maps still need to be learned.^{1–3} Learning novel sensorimotor maps can be studied by letting human subjects make arm movements to locations in space that each corresponds to a particular sound.⁴ This constitutes a challenging but feasible learning problem.

The nature of the map acquisition process remains elusive. The learner is exposed to only a

discrete set of observations about the sensorimotor workspace, namely previous movements and corresponding sensory results, but then often needs to move to targets that they have not heard before. That is, the learner needs to make predictions about unknown parts of the sensorimotor workspace. One hypothesis is that the learner stores a history of movements and sounds, and, when novel sounds are presented, they interpolate between these known instances to reach to intermediate locations. Thus, in principle, a mapping could be learned reasonably well if only parts of the sensorimotor workspace were known.

Instances of an interpolation process have been studied previously in the context of sensorimotor mappings that are largely known at the outset. Sensorimotor learning appears to generalize to novel movements or locations within the workspace,^{5,6} but this generalization falls off steeply with distance from the training direction.^{7–10} Generalization

appears to be encoded in intrinsic coordinates.¹⁰ Some studies suggest that generalization is tied to a particular context.^{11,12} Generalization also depends on the nature of the sensorimotor transformation being learned.⁶ In these cases, generalization was typically studied in situations where sensorimotor maps were already learned and underwent small modifications or perturbations. However, it remains unclear how generalization occurs when one learns sensorimotor maps from scratch.

The hypothesis that generalization is based on interpolation is tested here by limiting the exposure participants receive to specific instances of the relationship between movements and sounds and then testing them across the whole of the sensory space. Specifically, participants received what will be referred to as *anchor trials*, in which their arm was

moved to positions at the edges of the workspace, and then they were presented with the corresponding sounds (Fig. 1). On *probe trials*, one of a set of intermediate sounds was presented, and participants were invited to move to the location they believed corresponded to the sound. On these trials, they did not receive feedback and therefore the only information they received about the pairing between movements and sounds was from anchor trials at the edges. In principle, participants had no way of knowing which location should correspond to these probe sounds, because there are infinite mappings that would be consistent with the sound–movement mapping at the anchor locations. The central question addressed here is how participants interpolate sounds that correspond to unknown locations in the workspace.

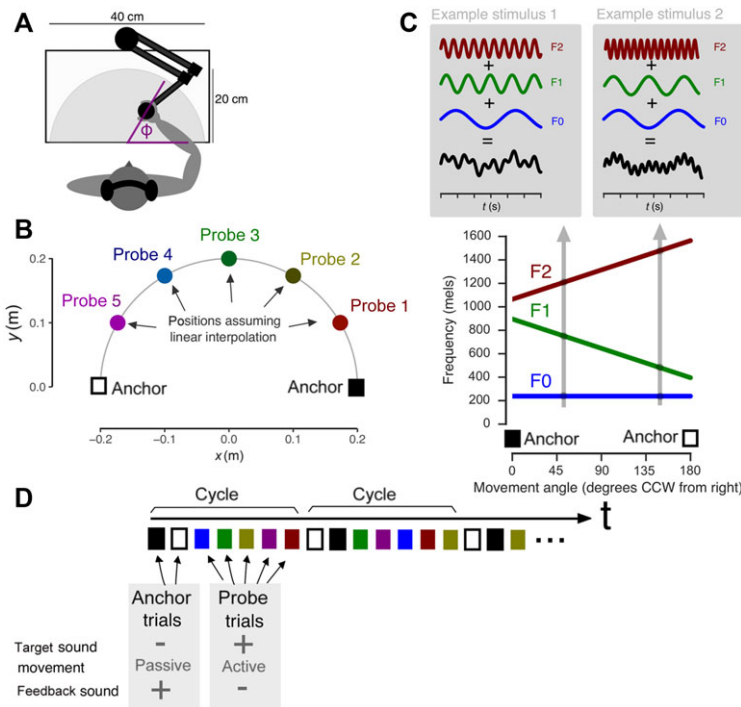


Figure 1. Participants make movements to auditory targets. Before each trial, a target sound is presented. Participants then make a reaching movement to the presumed location of the sound. (A) Top view of the workspace. Participants made reaching movements from the center of the workspace to points on a semicircle (not visible to participants). (B) Overview of the anchor locations and the locations that would correspond to the probe sounds assuming a linear mapping. (C) Sounds were not localized in space but, like speech sounds, were distinguished only by their frequency content. The sounds 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). (D) On every set of trials (cycle), first, the two anchors were presented, in random order, where the subject was moved passively with auditory feedback. Then, the five probes were presented in random order, and subjects were invited to make active movements but they did not receive auditory feedback.

Experimental procedures

Subjects and experimental tasks

A total of 27 participants were recruited. All participants were right-handed (handedness quotients 80–100% according to Edinburgh Handedness Inventory), except for two who were 88% and 20% left-handed. Participants reported no neurological or hearing impairments. Each session lasted approximately 1.5 hours. Participants provided written consent, and all procedures were approved by the McGill University Institutional Review Board, according to ethical principles stated in the Helsinki Declaration.

Participants filled out a questionnaire to list their musical training. Twelve participants listed no musical training of any kind. Three participants listed only vocal training (choir or singing lessons) for 1–2 years and between 8 and 20 years ago. The remaining participants listed mostly piano-only training (five participants) followed by violin (two participants), and others indicated various instruments. The instrumental training varied between 1 and 12 years, but in all cases the last training was more than 5 years ago, except for one participant who was currently still practicing.

Participants made reaching movements while holding a robot handle. A 2-degree-of-freedom planar robotic arm (InMotion2; Interactive Motion Technologies) was used (Fig. 1A) that sampled the position of the handle at 400 Hz. A *target circle* was defined as a half-circle around a subject-calibrated starting point, and during the experiment participants made movements from the center point to points on this circle (Fig. 1A). Vision of the circle and the arm was blocked for the duration of the experiment.

Auditory stimuli

The sounds consisted of three sine wave oscillators: one had a fixed frequency (F_0 , 165 Hz); the frequencies of the other two (F_1 , F_2) decreased or increased linearly, in relation to the angle of the movement endpoint (Fig. 1C). In order to normalize the space-to-frequency mapping, angles were mapped linearly to frequencies in mel-space.¹³ In order to correct for perceived loudness differences, the amplitude of each oscillator was adjusted using equal loudness curves¹⁴ to 75 phon so that each sound would be perceived to be as loud as a 1-kHz

tone of 75 dB. In this way, equally spaced locations about the target circle were associated with perceptually equal changes in sound frequency and little or no change in sound intensity. Sounds were presented over Beyerdynamic DT770M headphones. Since the mapping is one-to-one, the sounds will be referred to by the angle they were mapped to on the interval $[0, 180]$ degrees. 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 had information about the position they were mapped to.

Audiomotor testing

On anchor trials, indicated by a red color cue, the robot moved the participant's arm passively to one of the two anchor locations ($+90$ or -90 degrees from straight ahead; Fig. 1B) along a smooth minimum-jerk trajectory of 20 cm length and 900 ms duration. The robot then held their arm at the end position and presented the corresponding sound for 1000 milliseconds.

On probe trials, one of five target sounds was presented that, under linear interpolation, would correspond to the locations -60 , -30 , 0 , $+30$, and $+60$ degrees from straight ahead (Fig. 1B). Note that, although for convenience we refer to the sounds using the mapping location, this location had no physical meaning, since these sounds were never presented as feedback for movements. A green color cue indicated that subjects were expected to move. The target sound was presented for 1000 ms, after which color cue disappeared, and participants could initiate their movement. Movement end was when the velocity fell below 5% of peak velocity for 50 milliseconds. The robot then returned the subjects hand to the starting point. Subjects received no auditory feedback or any other kind of feedback for their probe movements. At no point were the subjects able to see their arms.

Anchor and probe trials were presented in cycles of seven trials. First, the two anchors were presented in random order. Then, the five probe trials were presented in quasirandom order (with the only restriction that the first probe could not be one of the two edge sounds corresponding to $+60$ and -60 degrees from straight ahead). The participants performed three blocks with 12 cycles each, for a total of 252 trials.

A subset of participants ($n = 15$) also performed probe-only blocks before and after the main testing described above. In these blocks, four cycles of the five probes were presented in random order without feedback for a total of 20 trials. This was done to assess participants' baseline performance as well as performance for the same target locations following the audiomotor testing.

Movement and auditory testing

Movement-copy and auditory psychophysical testing was performed before and after audiomotor testing to measure baseline motor accuracy in reaching toward particular directions and to assess whether auditory functioning was stable.⁴

Movement-copy. The participants' hand was moved out from the starting position to a target position on the target circle in a smooth movement of 900 ms, held there for 500 ms, and then moved back. A visual icon signaled participants to reproduce the movement they just experienced. Eleven target directions were presented that were equally spaced on the half circle, including the endpoints, and presented in random order. Participants were not able to see their arm, and no visual information of any kind was available.

Auditory testing. On every trial, a train of four sounds of 200 ms duration each was presented with a 75-ms pause between sounds. Three of the sounds were identical, and one (either the second or third) was different. Participants' task was to respond by pressing a button to indicate whether the mismatched sound was the second or third in the sequence. The three identical sounds were those that were mapped to the 90-degree angle in our auditory-motor mapping, and the mismatched sound corresponded in the mapping to the position angles of 90.27–106.20 degrees in 10 logarithmically spaced steps. Participants completed a total of 200 trials. No feedback was given about the accuracy of the participants' responses.

Data analysis

The main variable of interest is the movement angle on a given trial (since this provides an estimate of the participant's prediction of the position of the target sound). Movement angles associated with different auditory target locations are then compared with the linearly interpolated location of the target sound, and the absolute deviation between them is

computed, in order to assess whether participants follow linear interpolation. We computed ANOVA unless otherwise specified, and checked for sphericity using Mauchly's test, and, if significant, we reported the Greenhouse–Geisser corrected P -value of the ANOVA. In figures, the error bars and shaded lines indicate the standard error of the mean.

Model fitting

Participants' average movement direction may obscure patterns in the distribution of individual movements. In order to test hypotheses about what movements participants would make to novel targets (interpolation or lack thereof), a number of models of the distribution of movements to different target sounds were formulated that are laid out in the results section. Models were fitted for maximum likelihood. Models were compared according to the Bayesian information criterion (BIC) and Akaike information criterion (AIC),¹⁵ which quantify how good the model fits the data (likelihood) while taking into account the number of parameters used to get this fit in order to penalize overfitting. BIC is defined as $\ln(n)k - 2\ln(L)$, where n is the number of observations, k is the number of free parameters, and L is the maximum likelihood fit. AIC is defined as $2k - 2\ln(L)$. While the absolute values of AIC and BIC are meaningless, models can be compared by their difference in AIC or BIC (ΔAIC , ΔBIC). It is considered that, the lower the AIC or BIC, the better the model accounts for the data.

Results

Participants made reaching movements to auditory targets based on being presented with the movement-to-sound correspondence only for the workspace edges (on anchor trials) (Fig. 1). Reaching during probe trials revealed that on average participants moved to different locations for the five probes ($F(4,104) = 35.93$, $P < 0.0001$), and these locations on average followed the order of the probe sounds (Fig. 2A). There was some indication that these locations became more distinct over time (Fig. 2B), but overall there was little evidence for a change in performance over the course of testing. When comparing the deviation of subjects' reaching patterns from a linear interpolation, this deviation showed a tendency to decrease: comparing the first 50 with the last 50 trials revealed a trend for a decrease of deviation ($t(26) = 1.76$, $P = 0.08$), but

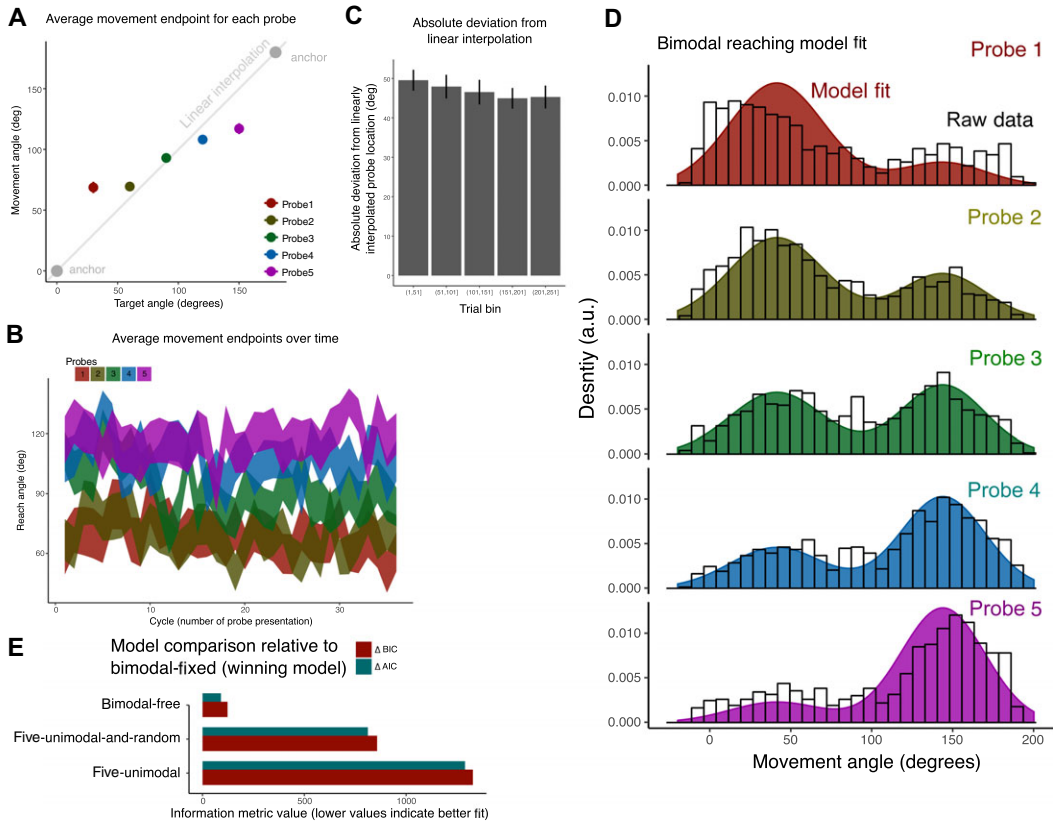


Figure 2. Probe trial reaching averages across subjects suggests, on first inspection, that participants interpolate to obtain probe locations. (A) Average reaching direction for each of the five probes. Error bars indicate standard error of the mean, but in some cases are covered by the dots. The gray line indicates the reaching position corresponding to linear interpolation between the anchor locations (gray dots). Subjects move on average to different locations for the five probe sounds, and the spatial ordering of these locations aligns with the auditory continuum of the probe sounds. (B) The reaching pattern over the course of the experiment indicates a tendency for the average movement endpoints to become more distinct for the five probes. (C) The deviation from a linear interpolation decreases, indicating a small trend that, reaching over time, more resembles a linear interpolation. (D) The distribution of individual reaches suggests that subjects use a bimodal reaching strategy in which they reach to one of only two locations, and the relative probability depends on the target sound. Each panel depicts the distribution of reaching endpoints for each of the probes; the raw data are represented as density histograms in black. The winning model, bimodal-fixed (see text), is shown overlaid on the figure. (E) Comparison of candidate models (see text for details). The model that best accounts for the data is the bimodal-fixed.

the magnitude of this change was small (Fig. 2C and see Discussion).

Modeling the movement distribution

On first inspection, it may be tempting to conclude that subjects interpolated locations for the probe sounds based on their exposure in the anchor trials. However, individual reaching movements followed a bimodal pattern, as confirmed by Hartigan’s dip test,¹⁶ which revealed that the reaching distributions were significantly nonunimodal for probes 1 and 3 ($P < 0.01$) and marginally significant for probe 2

($P = 0.076$). This indicates that the average reaching direction for a given probe is not representative of the reaching strategy. A modeling approach was therefore adopted in which a number of models were formulated to characterize the distribution of all reaching movements. Each model presents a different hypothetical interpolation strategy (or lack thereof).

The *five-unimodal* model embodied the idea that participants would produce unimodal reaching to a different location for each of the five probes. The

model contains 10 parameters: a mean and standard deviation for each probe. In the *bimodal-free* model, participants reached to one of only two locations, each with its own mean and standard deviation, and then the probability of reaching to either one or the other was parametrized by a probability that differed for each probe. That is, a total of nine parameters were estimated. However, the *bimodal-free* model left the weighting of the two response modes completely free, ignoring the fact that the probe sounds were related to the anchor sounds along an acoustic continuum. In order to accommodate this, a bimodal-fixed model was formulated which was like the bimodal-free model except that the weighting of the two response modes was assumed to reflect the acoustic similarity of the given probe to each of the two anchors. This acoustic similarity was encoded in a function “sim” that represented a distance metric in sound space that subjects were hypothesized to have at their disposal. Specifically, each mode was yoked to an anchor, a_0 or a_1 , so that given a probe r_i the probability of reaching to the mode associated with anchor a_j was $P = \text{sim}(r_i, a_j) / (\text{sim}(r_i, a_0) + \text{sim}(r_i, a_1))$ where “sim” refers to the acoustic similarity between the sounds. Since the sounds varied in pitch and were controlled for loudness (see Methods), the similarity was defined as the summed pitch differences in mel space between the corresponding oscillators of two sounds. The bimodal-fixed model had four free parameters (mean and standard deviation of the two modes). As a control, a unimodal reaching model, *five-unimodal-and-random*, was tested which was similar to the five-unimodal model in assuming that subjects would reach unimodally but on some proportion of trials would reach to a random location. This model used a total of 11 parameters (10 parameters as the five-unimodal model and one additional parameter that estimated the proportion of random reaching trials).

The model that best accounted for the data was the bimodal-fixed model, both according to AIC and BIC. The runner up was bimodal-free, which yielded worse model performance ($\Delta\text{BIC} = 123$, $\Delta\text{AIC} = 90$) (Fig. 2E), followed by the unimodal reaching models that did substantially worse ($\Delta\text{BIC}, \Delta\text{AIC} > 800$). We reran this analysis with the subset of subjects who had no prior musical training whatsoever, and, in this case, the bimodal models outperform the unimodal models as they do here; however,

bimodal-free outperformed the bimodal-fixed by a small margin.

No-feedback performance before and after exposure

A subgroup of participants performed probe trials before being exposed to anchors. On these trials, there was no difference in average reaching patterns between the probes ($F(1,14) = 0.08$, $P = 0.79$), indicating that participants had no knowledge of the audiomotor map before learning (Fig. 3A). Bayes factor analysis was performed on the no-feedback trials before exposure with the reach angle as dependent variable and the probe (factor with five levels) as independent variable. The Bayes factor for the probe factor was 0.018 ($\pm 1\%$), a value that is considered very strong evidence in favor of the null hypothesis, which in this case is that reaching patterns are not different for the probes. After exposure, the Bayes factor in the same analysis was 3327.02 ($\pm 0.69\%$), indicating very strong evidence in favor of the alternative hypothesis that reaching patterns are different across the probes. In no-feedback trials after exposure, movements differed by probe ($F(1,14) = 12.93$, $P = 0.003$), indicating that subjects had acquired information about the mapping. Comparing pairs of movements to either the same probes or to different probes revealed that before exposure, although across subjects movements to different probes could not be differentiated, within subjects there was a tendency for such differentiation that was strengthened over time (Fig. 3B). Immediately after presentation of the first two anchors, subjects' reaching patterns changed so that movements to different probes became more distinct, but then there was no further change for the duration of the experiment (Fig. 3C) in spite of repeated exposure to anchor trials.

Auditory-only and movement-only testing

Before and after exposure, participants were tested for auditory discrimination, and their performance on this test did not improve from before to after exposure, but instead showed a small decline ($t(26) = 2.49$, $P = 0.02$) (Fig. 4A). Participants also performed a movement-copy task where they had to reproduce passively experienced movements. Their replication of movements was reliable (Fig. 4B) and did not change after exposure ($t(26) = -0.54$, $P = 0.6$). This suggested that any changes observed in audiomotor trials, learning or other, were not a

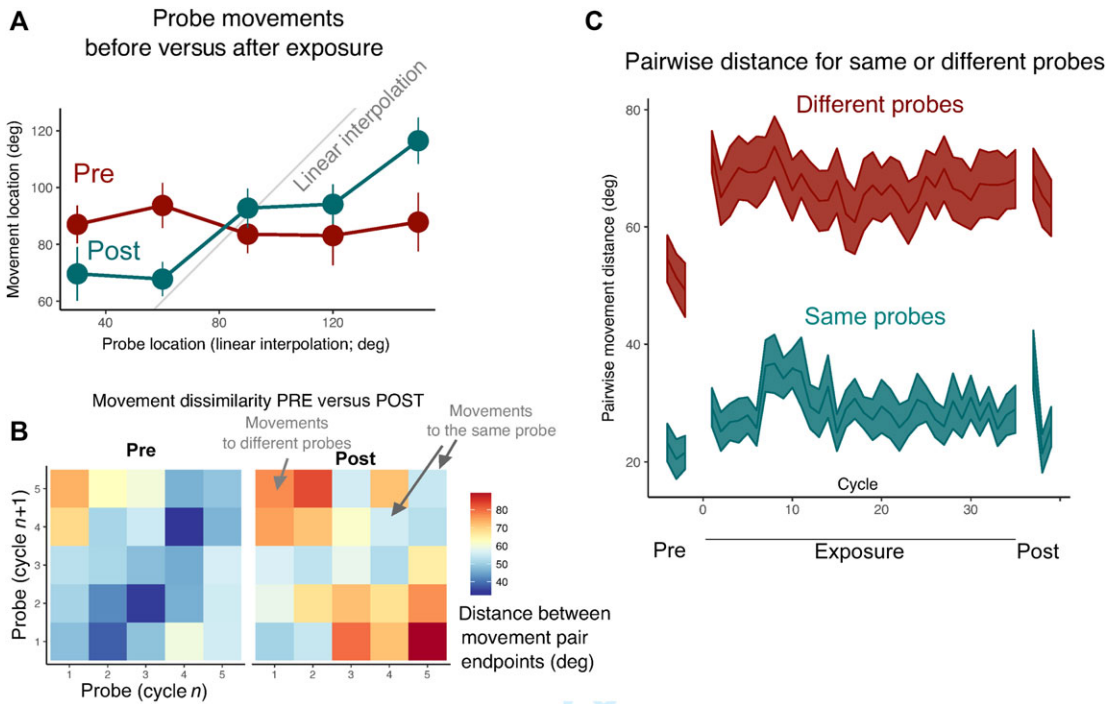


Figure 3. No-feedback trials before and after exposure show that subjects on average do not differentiate between targets, but nevertheless they show some consistency in their reaching pattern. This consistency is strengthened upon presentation of the first anchor trial and then remains constant. A subset of individuals performed probe-only trials (without feedback) before and after audiomotor testing. (A) Movement endpoints for the various probes. The horizontal axis indicates what the probe position would be given linear interpolation, and the gray line indicates the identity line where subjects would reach if they linearly interpolated the probes. Before learning, on average subjects move to the same location for all probes, but after audiomotor testing they distinguish the probes in their movements. (B) Cycle-by-cycle distance between pairs of movements. Each panel represents the comparison of movements on two cycles. The horizontal axis represents the probes of a given cycle n , and the vertical axis represents the probes of cycle $n + 1$. The color in each cell codes the absolute distance between the pair of reaching movements to those probes. Diagonal elements indicate movements to the same probe, off-diagonal elements indicate movements to different probes. It is evident that, even before audiomotor testing (pre), there is some structure to the participant’s reaching, reflecting an internal consistency in the way sounds are allocated in space, but this is strengthened after testing (post). (C) Tracking over time the distances between pairs of movements to same or different probes shows that, immediately upon exposure to the first anchors, the distance between pairs of movements to different probes increases, and then remains constant for the rest of the experiment in spite of repeated exposure to the anchors.

reflection of change in motor or auditory functioning alone.

Discussion

A prerequisite for sensorimotor learning is knowledge of the relation between movements and sensory effects: a sensorimotor map. There is increasing evidence that these maps are not innate but have to be learned, as for example, in the case of visually guided movements^{17–19} or learning to talk.^{20–22} Studying the acquisition of such maps is challenging, because in typical (adaptation) motor learning settings adults are tested that already have pre-existing

sensorimotor maps. Several studies approached this challenge by investigating the acquisition of novel sensorimotor maps by mapping wrist rotations to cursor movement on a screen,^{23,24} but, to our knowledge, these studies have not directly investigated how newly learned mappings generalize to enable movements to untrained locations. Other studies investigated participants learning to play the piano,^{25,26} which presumably requires acquisition of a sensorimotor map, but, since keys correspond to discrete pitches, there is little need for interpolation. Here, participants were exposed to a novel mapping between movements and sounds, which is

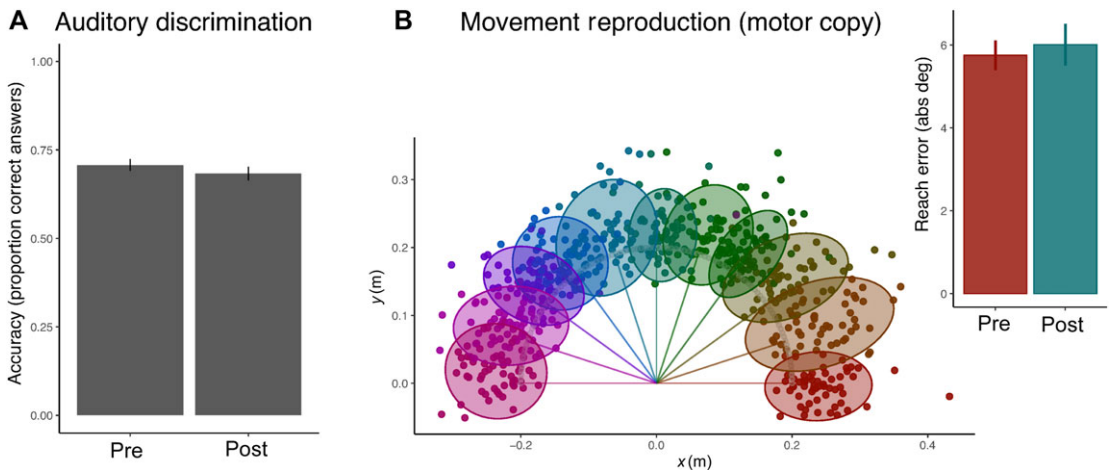


Figure 4. Auditory-only and movement-only tests reveal no improvement after audiomotor testing relative to before. (A) Auditory discrimination thresholds obtained using the sounds of the auditory–motor task reveal no improvements. (B) Motor copy tests reveal that participants can accurately reproduce movements. The dots indicate the endpoint locations for all subjects and are color-coded according to the target. The targets are indicated here as straight lines from the center using the same coloring scheme. The shaded ellipses represent the 2D error variance for each target. The motor copy performance does not change before versus after exposure (inlay).

challenging but feasible and allows studying how mappings are learned from scratch.⁴ In order to investigate the process by which sensorimotor maps generalize, we explored how subjects fill in regions of the sensorimotor workspace where they had no prior experience.

The principal observation in the present data set is that, although participants appear to reach to intermediate (interpolation) targets on average, the distribution of individual movements indicates that they achieve such averages by moving to one of only two locations with differing probabilities. That is, instead of generating unimodal movement distributions for each probe target, instead the participants appear to move to a weighted combination of two unimodal movement distributions in the vicinity of the anchor locations where they have prior exposure to the movement-to-sound mapping. Note that participants are not simply reproducing the anchors. The two locations of the bimodal reaching do not correspond to the anchor locations exactly, and the variability in movements to each of the two locations is larger than the movement-copy error.

One basic conclusion from this observation is that participants are at the very least able to judge the similarity between pairs of sounds, that is, compute a distance metric in auditory space. If instead participants perceived each sound categorically (as

is often observed with speech sounds²⁷), they could presumably not have produced sets of movements that align with the continuum of the probe sounds as was observed here.

In light of the apparent availability of auditory distance information, the finding that participants reach to a set of only two locations seems surprising. If participants know where a probe falls in the continuum between two anchor sounds, one would expect that they could produce the corresponding weighted average of the two anchor movements. Indeed, the motor system has been shown to have the capacity to maintain simultaneously in memory multiple movements to multiple targets and recombine them when task demands or the context changes.^{28,29} Averaging such movements would seem a relatively straightforward step for the motor system. It would follow therefore that the bimodal reaching strategy adopted by the subjects is not a memory limitation but originates elsewhere.

One possibility is that the motor system is reluctant to move to locations for which it has no prior experience, preferring to instead produce movements for which it can reliably predict the sensory result: in the present case, the movements close to the edges of the workspace. Indeed, there is in principle an infinite number of possible sensorimotor mappings that are consistent with the

sound–movement coupling the subject experienced at the workspace edges.^{30,31} It could therefore be risky to move far away from these edges. Given this constraint, the weighted bimodal reaching that subjects exhibit here might actually be the optimal strategy in a trade-off between two constraints: (1) producing movements that most closely match the target, that is, reducing error, and (2) producing movements that are closest to the known location in the sensorimotor workspace, that is, reducing risk. This account is in line with evidence suggesting that the motor system continually weighs the risk of novel and potentially more rewarding movements against the payoff for known movements,^{32,33} reminiscent of an exploration–exploitation trade-off.^{34,35} Within the context of the present study, it is possible that exploration has been discouraged because participants were aware that, for probe movements, no auditory feedback would be presented. Moving to novel locations, therefore, would not yield novel information about the sensorimotor workspace in the way it would in more natural settings.

Another possibility for the bimodal reaching pattern is that the sounds at the edges of the workspace are more distinct and therefore easier to memorize. Within the auditory continuum, the sounds at the edges are most different from all of the other sounds. However, this explanation is contradicted by the observation that movements to the sounds at the extremities are biased toward the workspace center instead of being biased outward toward the more stereotypical (Fig. 2A), as might be expected if reaching were based on memorization accuracy.

The pattern of findings reported in this study can lead to hypotheses about the neural mechanisms underlying learning of sensorimotor maps. The process of filling in the unvisited locations of the sensorimotor workspace likely requires retrieval of auditory and motor information associated with previous visits to nearby regions of the sensorimotor workspace. This information is then recombined in auditory space to compute a weighted average of the nearby auditory information, presumably involving auditory working memory areas of the brain. However, the present data suggest that information is not combined in the motor space, as participants in this case only select one of two movements and thus do not interpolate in motor space. Therefore, one could predict that motor or somatosensory working

memory areas are not involved in moving to novel targets.

Participants in the present study exhibited little learning in spite of being exposed to more and more anchor trials in the course of the experiment. This stands in contrast to continuing learning observed previously when subjects made reaching movements with feedback throughout the workspace.⁴ This difference in learning performance would be predicted if learning sensorimotor maps consists of forming a look-up table connecting movements with sounds. In such a case, once the anchor trials have been entered into the look-up table, repeated exposure to these anchors yields little additional information and thus little to no improvement in performance. Note that a small (and not statistically significant) amount of learning was seen here, which may challenge this account.

Learning audiomotor maps is also studied in songbirds, and recent work has investigated generalization. When songbirds are exposed to altered auditory feedback, they adapt their vocal output and generalize this behavior to new sequential contexts,³⁶ and this generalization can furthermore be modulated by the instructions given in order to flexibly respond to differential task demands in different contexts.³⁷ Although these studies deal with generalization, they focus on a different phenomenon than the one studied here, because the birds do not make novel movements in unexplored parts of the sensorimotor workspace but instead make known movements in novel sequential contexts.

In summary, the present data suggest that, when learning sensorimotor maps from scratch, there is only limited interpolation, and instead movements remain close to the parts of the workspace where the learner has exposure. This is consistent with some previous work on movement generalization showing that there is little change in movements beyond those visited during training.^{7,8} Instead of generating novel movements, the brain appears to prefer to use familiar movements, but selects them in such a fashion as to minimize target error.

Acknowledgments

This work was supported by National Institute of Child Health and Human Development (R01 HD075740), Fonds de recherche du Québec—

Nature et technologies (FRQNT), and a Banting Fellowship BPF-NSERC-01098. We are indebted to Bilal Alchalabi for assistance in data collection and the members of the Motor Control Lab at McGill for valuable discussions.

Competing interests

The authors declare no competing interests.

References

- Huberdeau, D.M., J.W. Krakauer & A.M. Haith. 2015. Dual-process decomposition in human sensorimotor adaptation. *Curr. Opin. Neurobiol.* **33**: 71–77.
- Krakauer, J.W. 2009. Motor learning and consolidation: the case of visuomotor rotation. *Adv. Exp. Med. Biol.* **629**: 405–421.
- Houde, J.F. & M.I. Jordan. 1998. Sensorimotor adaptation in speech production. *Science* **279**: 1213–1216.
- van Vugt, F.T. & D.J. Ostry. 2018. Structure and acquisition of novel sensorimotor maps. *J. Cogn. Neurosci.* **30**: 290–306.
- Goodbody, S.J. & D.M. Wolpert. 1998. Temporal and amplitude generalization in motor learning. *J. Neurophysiol.* **79**: 1825–1838.
- Krakauer, J.W., Z.M. Pine, M.-F. Ghilardi & C. Ghez. 2000. Learning of visuomotor transformations for vectorial planning of reaching trajectories. *J. Neurosci.* **20**: 8916–8924.
- Mattar, A.A.G. & D.J. Ostry. 2007. Modifiability of generalization in dynamics learning. *J. Neurophysiol.* **98**: 3321–3329.
- Mattar, A.A.G. & D.J. Ostry. 2010. Generalization of dynamics learning across changes in movement amplitude. *J. Neurophysiol.* **104**: 426–438.
- Gandolfo, F., F.A. Mussa-Ivaldi & E. Bizzi. 1996. Motor learning by field approximation. *Proc. Natl. Acad. Sci. USA* **93**: 3843–3846.
- Malfait, N., D.M. Shiller & D.J. Ostry. 2002. Transfer of motor learning across arm configurations. *J. Neurosci.* **22**: 9656–9660.
- Krakauer, J.W., P. Mazzoni, A. Ghazizadeh, *et al.* 2006. Generalization of motor learning depends on the history of prior action. *PLoS Biol.* **4**: e316.
- Cothros, N., J. Wong & P.L. Gribble. 2009. Visual cues signaling object grasp reduce interference in motor learning. *J. Neurophysiol.* **102**: 2112–2120.
- Stevens, S.S. 1937. A scale for the measurement of the psychological magnitude pitch. *J. Acoust. Soc. Am.* **8**: 185.
- Robinson, D.W. & R.S. Dadson. 1956. A re-determination of the equal-loudness relations for pure tones. *Br. J. Appl. Phys.* **7**: 166.
- Burnham, K.P. & D.R. Anderson. 2004. Multimodel inference: understanding AIC and BIC in model selection. *Sociol. Methods Res.* **33**: 261–304.
- Hartigan, J.A. & P.M. Hartigan. 1985. The dip test of unimodality. *Ann. Stat.* **13**: 70–84.
- Corbetta, D., S.L. Thurman, R. Wiener, *et al.* 2014. Mapping the feel of the arm with the sight of the object: on the embodied origins of infant reaching. *Cogn. Sci.* **5**: 576.
- Bushnell, E.W. 1985. The decline of visually guided reaching during infancy. *Infant Behav. Dev.* **8**: 139–155.
- Clifton, R.K., D.W. Muir, D.H. Ashmead & M.G. Clarkson. 1993. Is visually guided reaching in early infancy a myth? *Child Dev.* **64**: 1099–1110.
- Kuhl, P.K. 1994. Learning and representation in speech and language. *Curr. Opin. Neurobiol.* **4**: 812–822.
- Oller, D.K. 2000. *The Emergence of the Speech Capacity*. Taylor & Francis.
- Moulin-Frier, C., S.M. Nguyen & P.-Y. Oudeyer. 2014. Self-organization of early vocal development in infants and machines: the role of intrinsic motivation. *Cogn. Sci.* **4**: 1006.
- Mosier, K.M., R.A. Scheidt, S. Acosta & F.A. Mussa-Ivaldi. 2005. Remapping hand movements in a novel geometrical environment. *J. Neurophysiol.* **94**: 4362–4372.
- Liu, X., K.M. Mosier, F.A. Mussa-Ivaldi, *et al.* 2011. Reorganization of finger coordination patterns during adaptation to rotation and scaling of a newly learned sensorimotor transformation. *J. Neurophysiol.* **105**: 454–473.
- Bangert, M. & E.O. Altenmüller. 2003. Mapping perception to action in piano practice: a longitudinal DC-EEG study. *BMC Neurosci.* **4**: 26.
- Lahav, A., A. Boulanger, G. Schlaug & E. Saltzman. 2005. The power of listening: auditory–motor interactions in musical training. *Ann. N.Y. Acad. Sci.* **1060**: 189–194.
- Niziolek, C.A. & F.H. Guenther. 2013. Vowel category boundaries enhance cortical and behavioral responses to speech feedback alterations. *J. Neurosci.* **33**: 12090–12098.
- Gallivan, J.P., N.A.R. Bowman, C.S. Chapman, *et al.* 2016. The sequential encoding of competing action goals involves dynamic restructuring of motor plans in working memory. *J. Neurophysiol.* **115**: 3113–3122.
- Gallivan, J.P., L. Logan, D.M. Wolpert & J.R. Flanagan. 2016. Parallel specification of competing sensorimotor control policies for alternative action options. *Nat. Neurosci.* **19**: 320–326.
- Jordan, M.I. & D.E. Rumelhart. 1992. Forward models: supervised learning with a distal teacher. *Cogn. Sci.* **16**: 307–354.
- Jordan, M.I. & D.M. Wolpert. 1999. *Computational Motor Control*. MIT Press.
- Dhawale, A.K., M.A. Smith & B.P. Ölveczky. 2017. The role of variability in motor learning. *Annu. Rev. Neurosci.* **40**: 479–498.
- Pekny, S.E., J. Izawa & R. Shadmehr. 2015. Reward-dependent modulation of movement variability. *J. Neurosci.* **35**: 4015–4024.
- Sutton, R.S. & A.G. Barto. 1998. *Reinforcement Learning: An Introduction*. Adaptive computation and machine learning. Cambridge, MA: MIT Press.
- Peters, J. & S. Schaal. 2008. Reinforcement learning of motor skills with policy gradients. *Neural Netw.* **21**: 682–697.
- Hoffmann, L.A. & S.J. Sober. 2014. Vocal generalization depends on gesture identity and sequence. *J. Neurosci.* **34**: 5564–5574.
- Tian, L.Y. & M.S. Brainard. 2017. Discrete circuits support generalized versus context-specific vocal learning in the songbird. *Neuron* **96**: 1168–1177.e5.