npj Science of Learning

https://doi.org/10.1038/s41539-026-00428-4

Article in Press

Hypothesis testing governs strategic motor learning

Received: 17 December 2025

Accepted: 24 April 2026

Cite this article as: Ding, W., Niyogi, A., Taylor, J.A. *et al.* Hypothesis testing governs strategic motor learning. *npj Sci. Learn.* (2026). https://doi.org/10.1038/s41539-026-00428-4

Wei Ding, Anjuli Niyogi, Jordan A. Taylor & Jonathan S. Tsay

Hypothesis Testing Governs Strategic Motor Learning

Wei Ding[1, 2,*], Anjuli Niyogi[2], Jordan A. Taylor[3,4] and Jonathan S. Tsay[2, 5 *]

[1] *Department of Psychological and Cognitive Sciences, Tsinghua University, Beijing, China*
[2] *Department of Psychology, Carnegie Mellon University, Pittsburgh, PA, USA*
[3] *Department of Psychology, Princeton University, Princeton, NJ, USA*
[4] *Princeton Neuroscience Institute, Princeton University, Princeton, NJ, USA*
[5] *Neuroscience Institute, Carnegie Mellon University, Pittsburgh, PA, USA*

Correspondence(*): teresading999@gmail.com & xiaotsay2015@gmail.com

Abstract

How do people discover an effective movement strategy when the environment abruptly changes—such as when using the trackpad on an unfamiliar laptop? Strategic adaptation is often described as a reinforcement learning process characterized by two key features: random exploration followed by gradual error reduction. We propose a different view in which strategic adaptation operates through hypothesis testing: learners generate specific action–outcome hypotheses about the environmental change, discount those that conflict with feedback, and continue testing alternatives until they discover the correct rule. To adjudicate between these accounts, we conducted two large-scale experiments using a visuomotor rotation task designed to isolate strategic adaptation under different target arrangements (N = 560). Individual learning trajectories showed pronounced exploration but were far from random, exhibiting structured, multimodal error distributions. Moreover, participants did not converge on the solution gradually; instead, they discovered it abruptly. Critically, strategic adaptation depended on target arrangement: some configurations steered participants toward the correct rotational hypothesis, whereas others led them to alternate between rotational and translational hypotheses. Together, these findings position hypothesis testing as a core mechanism governing strategic motor learning.

Introduction

Successful goal-directed action depends on multiple learning processes [1]. Among these, strategic motor adaptation is especially critical for efficiently adjusting our movements in response to changes in the body and in the environment [2–4]. For example, when switching to an unfamiliar ping-pong paddle that sends shots long, players may *deliberately* aim shorter to compensate, adopting an *explicit strategy* to maintain accuracy.

Strategic adaptation is often described as a reinforcement learning process characterized by two key behavioral features when the environment changes: random exploration to find a crude solution and gradual error reduction that restores successful performance [3–8]. Support for this view comes in part from the classic visuomotor rotation task, which perturbs the mapping between visual feedback and movement direction. When learning performance is averaged across participants, compensation for the perturbation follows a gradual, power-law trajectory [3,9–11]. Variability around this gradual trajectory is typically attributed to large, randomly distributed exploration and minimal sensorimotor noise [12].

But individual learning curves tell a very different story. Learners often engage in rich, exploratory behavior that is anything but gradual, with variability that is neither minimal nor random [13,14]. Instead, movement errors cluster into multiple modes, often far from the solution [15,16]. Such structured exploration is difficult to reconcile with reinforcement learning, which predicts errors largely confined between baseline and solution states, with only modest random variability. Consequently, averaging across participants may be masking the very behavioral signatures that define strategic adaptation [17–20].

Inspired by individual learning data, we have proposed a novel framework for strategic adaptation: hypothesis testing [14]. When confronted with a new visuomotor mapping—such as the altered bounce from an unfamiliar ping-pong paddle—learners generate candidate hypotheses for the environmental change (e.g., that the paddle induces an added rotation or translation), test them against sensory feedback, discard those that fail, and iteratively refine both their actions and hypotheses through practice until a successful strategy is discovered, often abruptly [15,21]. In this framework, strategic adaptation unfolds through structured and deliberate exploration: The multimodal errors observed in individual learning trajectories arise as learners actively generate, evaluate, and revise candidate action–outcome solutions—patterns that may appear random without careful consideration.

Beyond individual trajectories, the hypothesis testing framework predicts systematic group-level differences in learning as a function of the training environment, which can either steer learners toward a single hypothesis (constrained) or leave multiple hypotheses viable (unconstrained) [22]. Specifically, in unconstrained environments, learning should proceed rapidly because participants can adopt whichever hypothesis appears most plausible; when probed, performance should cluster around multiple candidate rules (e.g., rotation versus translation). In contrast, in constrained environments, learning should proceed more slowly because progress is limited until the correct hypothesis is discovered; when probed, performance should converge on the same rule. Thus, environmental constraints shape both the rate of learning and the diversity of rules adopted.

To adjudicate between reinforcement learning and hypothesis testing accounts, we conducted two large-scale experiments (N = 560) using a visuomotor rotation task designed to bias learning towards strategic motor adaptation [23]. Participants were instructed to explore the workspace to align a 60° rotated cursor to the target without being told the nature of the perturbation (Figure 1a-b).

Critically, we manipulated the degree of environmental constraint by varying the spatial arrangement of the training targets (Figure 1c–d). With widely spaced *outer targets*, the action–outcome cues constrained the solution space to a single (rotational) hypothesis (Figure 1e, left). In contrast, closely spaced *inner targets*

created a more unconstrained environment, allowing multiple hypotheses (rotation and translation) to generate successful performance during training (Figure 1e, right). To infer the rule(s) participants adopted, we interleaved a no-feedback probe target throughout the experiment. This probe allowed us to determine whether learners adopted the experimentally-imposed rotational rule (aiming 60° away clockwise from the target) or an alternative rule—such as an approximate translation (manifesting as ~60° counterclockwise, which we term a "sign flip"), a mirror reversal (aiming 180° away from the target), or a default baseline response (aiming directly to the target; 0°) (Figure 1f). More broadly, the probe target provided a window into the space of hypotheses entertained during learning.

In summary, these experiments test whether strategic adaptation unfolds through reinforcement learning—marked by brief random exploration followed by prolonged gradual error reduction (Figure 1g)—or through deliberate exploration among discrete action–outcome hypotheses until a solution is abruptly discovered (Figure 1h). Moreover, under a hypothesis testing account, learning should proceed through rapid adoption of one among several expedient hypotheses in inner-target (unconstrained) environments but show slower convergence on a single hypothesis in outer-target (constrained) environments.

Results

Experiment 1. To evaluate whether hypothesis testing governs strategic adaptation, we employed a visuomotor rotation task specifically designed to isolate this deliberate learning process (N = 280) (Figure 1a-b) [23]: After a baseline block with veridical feedback, participants reached toward two interleaved training targets while receiving cursor feedback rotated 60° relative to their movement direction. Endpoint cursor feedback was delayed by 1000 ms—a manipulation known to minimize implicit adaptation—ensuring that learning performance is biased toward deliberate, strategic (explicit) processes. Participants were not informed about the nature of the perturbation; instead, they were instructed to explore the workspace to learn the visuomotor mapping and align the perturbed cursor with the target.

To vary the degree of environmental constraint, we manipulated the spatial separation between the two training targets (Figure 1c): In the outer-target group (constrained environment), the wide target separation provided stronger geometric cues that the perturbation reflected the experimentally imposed rotation, narrowing the set of viable hypotheses (Figure 1e, left). In contrast, the inner-target group (unconstrained environment) offered weaker, more ambiguous spatial cues, where multiple hypotheses remained viable (Figure 1e, right).

While the reinforcement learning framework offers no clear prediction about how environmental constraints should shape learning and generalization (Figure 1g), the hypothesis testing framework predicts a dissociation (Figure 1h): the outer-target (constrained) group should learn more slowly at the training targets as learners reject simpler hypotheses, yet ultimately converge on the imposed rotational rule at the probe target; conversely, the inner-target group should learn more rapidly by adopting one of several expedient hypotheses, but express divergent rules at the probe target (e.g., the imposed rotation or approximate translations such as sign flips).

Following a baseline phase to familiarize participants with the visuomotor rotation task, participants were exposed to the rotated visual feedback. Both group average learning trajectories showed a canonical, gradual learning curve at the two training targets: movements starting at the baseline target and progressing toward the 60° solution, a pattern taken to support the reinforcement learning framework.

However, a striking dissociation emerged between groups: The inner-target (unconstrained) group learned more rapidly at the training targets (Figure 2a, left) yet showed poor generalization to the probe target (Figure 2a, right). In contrast, the outer-target (constrained) group learned more slowly at the training

targets but exhibited robust generalization at the probe target. This pattern is difficult to reconcile with a reinforcement learning account and instead uniquely predicted by the hypothesis-testing framework.

This dissociation was verified statistically. Across the whole perturbation block, the inner-target group consistently outperformed the outer-target group at the training locations (cluster-based permutation test: cycles 9 – 40, $\Delta\mu = 6.4°$, $d_{circ} = 0.1$, $p_{perm} < .001$) (Figure 2a, left): During early adaptation, the inner-target group adapted nearly twice that of the outer-target group (Inner: 43.1°± 2.4°; Outer: 28.5°±3.2°; p<.001), demonstrating faster learning under the unconstrained environment. By late adaptation, both groups had converged on the 60° solution, nullifying the visuomotor rotation (Inner: 56.9° ± 1.1; Outer: 54.4° ± 1.9; $p = .035$). A one-sample V-test confirmed that movement angles during the aftereffect phase were clustered around 0° in both groups ($p < .001$), confirming that learning in both groups reflects strategic, rather than implicit, learning processes.

The pattern of group performances reversed in the generalization probes (Figure 2a, right) (interaction between Target Geometry x Target Type: $F = 41.9$, $p < .001$, $\eta^2 = 0.13$): Participants trained with inner-targets showed poor (inconsistent) generalization compared to those trained in the outer-target group (cluster-based permutation test: cycles 9 – 40, $\Delta\mu = -35.1°$, $d_{circ} = -0.6$, $p_{perm}<.001$). Although neither group achieved full generalization (60°), the outer-target group exhibited a persistent advantage (late adaptation: Inner: 10.0° ± 4.2; Outer: 43.4° ± 3.6; p<.001), demonstrating that a more constrained environment yielded greater (consistent) generalization despite slower learning.

In stark contrast to the smooth group-averaged curves, individual learning trajectories revealed pronounced exploratory behavior (Figure 2b; additional participants shown in Figure 4e). Motor variability during early adaptation was 46.8° ± 2.1° in the inner-target group and 64.2° ± 2.3° in the outer-target group (Figure 2c). These values far exceeded baseline sensorimotor noise (main effect of Phase: $F = 856.6$, $p < .001$, $\eta^2 = 0.76$). Variability then declined over training (late adaptation SD; Inner: 12.8° ± 1.1, Outer: 20.3° ± 2.0). This variability exhibited rich, evolving multimodal structure at both training and probe targets (Figure 2d).

To determine whether this exploratory behavior reflected random (unimodal) exploration predicted by reinforcement learning or directed (multimodal) exploration predicted by hypothesis testing, we fit von Mises distributions to movement-angle data across the perturbation phase (see early and late phases shown separately in Figure S3). This analysis revealed that movements to the training targets were strikingly multimodal (Figure 2e, top; Best BIC, inner = −5369.6 with 5 modes; outer = −8563.3 with 6 modes). Beyond reaches to the optimal solution (60°) and a uniform component reflecting diffuse exploration, three prominent peaks emerged: sign flips (–60°), mirror reversals (180°), and default reaches directly to the target (0°). The same multimodal structure appeared at the probe target (Figure 2e, bottom; Best BIC: inner = −5851.3 with 7 modes; outer = −5432.0 with 6 modes; full mode breakdown in Table S2). Such error distributions are difficult to reconcile with reinforcement learning accounts, which predict errors largely confined between baseline (0°) and the goal (60°), with only diffuse random variability. Instead, they are consistent with learners actively testing specific action–outcome hypotheses.

Why, then, did the inner- and outer-target groups differ in both learning and generalization? We addressed this in two ways. First, motor variability was substantially higher in the outer-target group during training (main effect of Group: $F = 26.9$, $p < .001$, $\eta^2 = 0.09$; early adaptation post-hoc t-test: $t = 5.6$, $p < .001$), suggesting that the constrained geometry prompted greater exploration, and, consequently, slower convergence on the solution.

Second, the groups differed strikingly in their movement-angle distributions during training (Figure 2e-f; JSD = 0.06, $p = .003$): the outer-target group explored a wider hypothesis space—showing more sign flips (–60°), more mirror reversals (180°), a stronger uniform component, and fewer reaches to the correct solution (60°). At the probe target, this pattern reversed (JSD = 0.13, $p < .001$): the outer-target group

produced fewer sign flips and more movements near the correct 60° rule. The inner-target group showed that learners did not settle on a single solution but instead oscillated between rotation and sign-flips (approximate translation).

Together, these results provide a process-level account of the dissociation between learning and generalization. In the outer-target condition, the task geometry constrains the solution space, requiring learners to explore more broadly before identifying the rotational rule—resulting in slower learning but consistent convergence on a single solution. In contrast, the inner-target condition is diagnostically ambiguous, allowing multiple hypotheses (e.g., rotation or translation) to yield successful performance at the training targets. As a result, learners explore more narrowly, learn more quickly, but often settle on different, locally effective strategies that do not generalize in the same way. This dissociation at the group level, coupled with multimodal performance at the individual level, is incompatible with a reinforcement learning account and instead support hypothesis testing as a core process governing strategic motor learning.

Experiment 2. One key finding of Experiment 1 is that target geometry modulates both learning and generalization—a pattern consistent with the hypothesis testing account. However, one methodological issue tempers this interpretation: the inner- and outer-target groups differed not only in target geometry but also in their spatial distance from the probe target. Because generalization declines with distance [24–27], the outer-target group's superior generalization could simply reflect that their probe target was closer to the trained locations (110° in the outer vs. 170° in the inner condition).

To address this, we designed Experiment 2 (N = 280), again manipulating the level of environmental constraint via target geometry while equating the spatial distance between training and probe targets (110° for both groups; Figure 1d). Participants in both groups adapted to *eight* training targets, allowing us to vary geometric structure while holding generalization distance constant. Again, in the outer-target (constrained) group, the target geometry provided stronger cues that prompted the imposed rotation hypothesis. In contrast, the inner-target (unconstrained) group—marked by fewer outer targets and more inner targets—rendered multiple hypotheses plausible.

We reproduced group-level effects of target geometry under these stringent conditions. Inner group learned faster during early adaptation (Inner: 33.8° ± 3.1; Outer: 27.4° ± 4.0; $p = .012$). Both groups successfully compensated for the perturbation at the training targets (late adaptation: inner 52.2° ± 1.9.; outer 52.1° ± 2.0; $p = .65$), and neither exhibited implicit aftereffects ($p < .001$), confirming that performance reflected primarily strategic learning processes. Although the inner-target advantage was subtle, significant group differences still emerged (Figure 3a, left; cycles 13–16: $\Delta\mu = -3.0°$, $d_{circ} = 0.0$, $p_{perm} = .006$; cycles 27–32: $\Delta\mu = -2.1°$, $d_{circ} = -0.1$, $p_{perm} = .010$).

Crucially, this pattern reversed at the probe target (Figure 3a, right; interaction between Target Geometry x Target Type: F=6.3, p=0.012, η^2=0.02): the inner-target group generalized more inconsistently than the outer-target group. Significant group differences were observed throughout cycles 13–30 ($\Delta\mu = 14.9°$, $d_{circ} = 0.2$, $p_{perm} < .001$) and cycles 35-40 ($\Delta\mu = 13.8°$, $d_{circ} = 0.2$, $p_{perm} = .001$). Together, the dissociation between learning and generalization persisted even when generalization distance was equated between groups.

Individual learning curves showed pronounced early motor variability (Figure 3b: Inner: 58.7° ± 2.5, Outer: 62.9° ± 2.5), which declined systematically over time (Late SD: Inner: 20.9° ± 1.5, Outer: 22.0° ± 1.8). These values were significantly larger than baseline sensorimotor noise (main effect of Phase: $F = 829.8$, $p < .001$, $\eta^2 = 0.75$). The same evolving multimodal structure was evident in trial-by-trial heatmaps of movement-angle distributions (Figure 3d). Critically, this early exploratory behavior was multimodal (Figure 3e, top; Best BIC: inner = −8729.6 with 6 modes; outer = −9382.1 with 7 modes; von Mises parameters in Table S3; early and late phases shown separately in Figure S4), with several modes far from the 60° solution. This same multimodal structure appeared at the probe target (Figure 3e, bottom; Best BIC:

inner = −5762.4 with 5 modes; outer = −5844.4 with 6 modes), consistent with learners evaluating discrete hypotheses rather than randomly exploring or gradually reducing error, as predicted by reinforcement learning.

We then asked whether the two groups differed in how they explored—that is, in the hypotheses they sampled. Unlike Experiment 1, distributional differences in movement angles between groups were small (Figure 3e; training target: JSD = 0.03, p = .694; probe target: JSD = 0.07, p = .088), and motor variability did not differ significantly between groups (Figure 3c; F = 2.33, p = .12, η^2 = 0.008; post-hoc t-test: t = 1.2, p = .23)—an issue we revisit with a more sensitive cross-experiment analysis. That said, here we replicated the pattern from Experiment 1 (Figure 3f): the outer-target group explored a wider hypothesis space—showing more sign flips (–60°), more mirror reversals (180°), a stronger uniform component, and fewer reaches to the correct solution (60°). At the probe target, this pattern reversed: the outer-target group produced fewer sign flips and more movements near the correct 60° rule. Moreover, learners in the inner-target group did not settle on a single solution but instead adopted one of several divergent rules, including rotation and sign-flip (approximate translation) solutions.

Taken together, the constraint-dependent effects on strategic adaptation at the group level, combined with multimodal performance at the individual level, provide convergent evidence that strategic motor learning is governed by hypothesis testing rather than reinforcement learning.

Combined Analysis Across Experiments. While the dissociable effects of target geometry on learning and generalization emerged in both experiments, they were more pronounced in Experiment 1. This difference likely reflects the strength of the geometric contrast between conditions: In Experiment 1, the two-target geometries differed completely in target arrangement, whereas in Experiment 2, the eight-target geometries substantially overlapped and differed only in the arrangement of the inner targets.

Across experiments, an intriguing trend emerged: the ratio of inner (unconstrained) to outer (constrained) targets appeared to form a continuum in learning and generalization (2 outer → 8 outer → 8 inner → 2 inner). Pure inner-target environments (2-inner) provided the most unconstrained setting—supporting rapid learning but yielding the least consistent generalization—whereas pure outer-target environments (2-outer) imposed strong constraints, producing slower learning but the most consistent generalization. The mixed eight-target conditions fall between these extremes. Thus, these data suggest that hypothesis adoption varies continuously with target geometry: arrangements weighted toward outer targets steered learners toward the imposed rotational rule, whereas arrangements resembling inner-target geometry rendered the perturbation more ambiguous.

This continuum was verified statistically in multiple ways (N = 560). First, at the training targets, early adaptation increased continuously from 2-outer to 8-outer to 8-inner to 2-inner—precisely mirroring the proportion of outer (constrained) versus inner (unconstrained) targets (Figure 4a; mixed-effects regression: β = 6.0, p = .005). At the probe target, this pattern reversed: early adaptation decreased in the same graded sequence, reflecting increasingly inconsistent rule adoption (mixed-effects regression: β = –5.7, p < .001). Second, exploratory behavior showed the same continuum: summing the von Mises weights for all non-solution (non–60°) modes at the training target revealed a graded increase from 2-inner to 2-outer, mirroring the proportion of constrained outer targets; this pattern reversed at the probe target (Figure 4b).

Third, we examined individual learning curves using an unsupervised machine-learning approach, clustering participants with soft dynamic time warping (see Methods). We restricted this analysis to participants who successfully learned during training (N = 476 out of 560; 85% learner rate), a conservative approach to ensure that clusters reflected genuine strategic differences rather than inattention or task disengagement (similar cluster structure also emerged when all participants were included; see Figure S1).

Indeed, this unsupervised approach—operating on fine-grained individual data—was enabled by our large online sample and would be difficult to implement in standard in-person studies.

We observed seven distinct profiles (Figure 4c & 4e, C1–C7). Although all clusters achieved successful performance at the training target, they diverged markedly at the probe target: some converged on the correct solution through distinct exploratory trajectories (C1: 62.2% of learners; C2: 5.3%; C3: 6.5%), whereas others never reached the correct strategy—settling instead on approximate translation (sign-flip) solutions (C4: 11.3%; C5: 6.1%; C6: 2.5%) or defaulting to straight-to-target movements (C7: 6.1%). These qualitative differences were echoed in the distinct movement-angle distributions observed during early and late adaptation (Figure 4f–g), revealing the rich and diverse hypothesis testing behavior that defines strategic adaptation.

Critically, the dissociable impact of target geometry on learning and generalization emerged from a continuous shift in the mixture of strategies (Figure 4d; Table S4; $\chi^2(18) = 75.8$, $p < .001$): Groups with different target geometries comprised mixtures of exploratory strategies whose proportions shifted continuously with the ratio of outer to inner targets (2 outer → 8 outer → 8 inner → 2 inner). Outer-target geometries produced a higher proportion of learners who rapidly converged on the rotational rule (C1). Conversely, inner-target geometries yielded more participants in clusters that either failed to reach the correct solution (C4–C7) or required extended exploration before finding it (C2–C3).

Together, these diverse individual strategies, coupled with the continuous effects of target geometry constraining learning and generalization, underscore hypothesis testing as a core process governing strategic motor adaptation.

Discussion

Summary of Results and Implications. Strategic adaptation is often described as a form of reinforcement learning, characterized by random, diffuse exploration followed by gradual error reduction. Here we propose an alternative view: strategic adaptation operates through a more deliberate, structured process of hypothesis testing. In this framework, learners generate candidate action–outcome hypotheses about the environmental change, evaluate them against sensory feedback, and iteratively refine both their actions and beliefs until a solution is abruptly discovered.

Our group-level results support this hypothesis testing framework. Using a visuomotor rotation task that biases learning toward strategic adaptation, we found that outer-target geometry (a constrained environment) slowed learning at trained targets yet, when probed, promoted consistent adoption of the imposed rotational rule. In contrast, inner-target geometry (an unconstrained environment) accelerated learning but yielded more divergent rule adoption, with many participants expressing mirror reversals, approximate translations (sign flips), or default reaches at the probe target. This group-level pattern cannot be fully explained by a reinforcement learning account, which offers no principled prediction about how environmental constraints should differentially shape learning and generalization (see Supplement for preliminary evidence that similar hypothesis-testing processes arise even under smaller perturbations; Figure S2).

Evidence for hypothesis testing also emerges at the level of individual behavior. First, individual learning curves reveal substantial exploratory variability—far exceeding the levels of sensorimotor noise. Second, movement-angle distributions were strikingly multimodal, exhibiting 6–7 peaks, whereas reinforcement learning predicts a unimodal distribution largely confined between baseline and solution states. Third, learners showed remarkable diversity in the temporal dynamics of strategy discovery: some discovered the

60° (clockwise or counterclockwise) rotation immediately, others exhibited a discrete moment of insight, and still others explored extensively yet never discovered the imposed rule. Future work can leverage richer probe designs and subjective reports to more precisely identify the rules learners adopt on each trial (see Supplemental Materials for free-response vignettes illustrating the diverse hypotheses reported by participants).

This interpretation fundamentally reframes how large errors in strategic adaptation should be understood. Prior work has attributed errors such as sign flips or mirror reversals to attentional lapses in executing an already learned (cached) strategy [26], or to rapid use-dependent processes such as the reuse of a default movement under time pressure. Crucially, both accounts predict that such errors should occur uniformly across learning or even increase later as fatigue or repeated movements accumulate. Instead, we observe the opposite pattern: sign flips—and related errors such as mirror reversals—are most prevalent early in learning and reliably diminish over time. We therefore interpret these behaviors *not* as execution failures, but as signatures of active hypothesis testing, reflecting learners' exploration of alternative action–outcome mappings while discovering an effective re-aiming strategy.

Our results also clarify the role of cognition in motor learning. Although cognition is widely acknowledged to play a central role early in motor learning [28], it has largely been characterized at a phenomenological level—as learning that is more effortful, deliberate, or explicit. Here, we move beyond description to specify the nature of this cognitive involvement. We show that a core component of strategic motor learning is *reasoning* about the spatial transformation imposed by the environment [29]. An important direction for future work is to determine how policies derived from reasoning are refined with practice—becoming faster, more accurate, and less effortful—and how these policies are subsequently retrieved automatically across contexts [14].

More broadly, these findings place strategic adaptation within the broader class of hypothesis-driven cognition. Across many domains, people learn by generating and evaluating structured hypotheses: in problem solving, they test and revise predictions about an opponent's next move to guide their own [30,31]; in intuitive physics, humans evaluate alternative action–outcome predictions to determine how objects will fall or collide [32]; and in concept learning, they test different categorical boundaries to uncover latent rules that group stimuli into meaningful categories [33,34]. Our results suggest that strategic adaptation operates according to the same computational logic: Learners actively generate, evaluate, and discard candidate action–outcome mappings, using feedback to infer the structure of the perturbation. Viewed through this lens, strategic adaptation is best understood as a form of 'embodied problem solving', in which hypothesis testing is implemented through action rather than abstract deliberation.

A central contribution of this work is highlighting how group averages can obscure the very learning processes they aim to characterize. We argue that these processes are more clearly revealed by examining individual learning trajectories, an approach long established in other areas of learning research but rarely adopted in motor learning [17,35–39]. To support this shift in focus, we applied a suite of analytical tools (e.g., von Mises fitting, dynamic time warping) that allow us to characterize the diversity and temporal structure of individual behavior and to identify hypothesis testing as a central mechanism in sensorimotor learning.

The presence of stable individual differences in learning strategies raises fundamental questions about their origins. To our knowledge, the cognitive and dispositional traits associated with different sensorimotor learning profiles have not been examined, likely in part because motor learning studies have traditionally relied on modest sample sizes. Our online experimental approach [40], together with the individual-level analytical framework introduced here, makes large-scale investigation feasible. By linking learning phenotypes to broader cognitive characteristics, such as working memory capacity, exploration–

exploitation preferences, or tolerance for uncertainty, future research can move beyond describing how individuals differ to explaining why these differences arise [41,42].

This work raises important questions about how hypothesis testing during strategic adaptation varies across the lifespan. During development, substantial work has characterized the maturation of implicit motor adaptation [43]; by contrast, the developmental time course of strategic adaptation remains largely unknown. A key open question is whether the emergence of strategic adaptation tracks the maturation of spatial reasoning abilities. In aging, we have found that older adults show selective impairments in strategic discovery in unconstrained environments (e.g., two-target conditions), while performance is preserved in more constrained environments (e.g., eight-target conditions) [44]. We speculate that this pattern reflects an age-related reduction in the ability to generate or evaluate more flexible or inventive hypotheses for solving visuomotor mappings—a possibility that needs to be directly tested.

Our findings also contribute to a growing literature highlighting the central role of the goal in learning [22,45–48]. Prior work shows that target geometry shapes implicit adaptation. We extend this to strategic adaptation [49]: the arrangement of targets can bias learners toward particular action–outcome hypotheses. Specifically, task geometry determines not only how readily competing hypotheses can be ruled out, but also how quickly learners adopt a "good-enough" hypothesis to reduce error.

Beyond its theoretical implications, manipulating goal features may provide a principled means of guiding learners toward more effective sensorimotor strategies in clinical and coaching contexts [50–53]. In some cases, it may be desirable to promote adoption of a specific strategy by ruling out competing hypotheses through a more constrained environment. In other cases, encouraging learners to consider multiple candidate strategies may be beneficial, motivating the use of more unconstrained contexts that support broader exploration. An important direction for future work is to translate these findings to ecological settings and evaluate how environmental constraints shape hypothesis testing and the adoption of good-enough solutions in real-world behavior [54–56].

Formalizing the Hypothesis Testing Model. An important direction for future research is to formalize the hypothesis testing account. We envision this framework as comprising two core components: a hypothesis space and an inference procedure. The hypothesis space consists of a discrete set of spatial transformations—including rotations, mirror reflections, translations, and gain changes—each associated with a prior belief. On each trial, an action is sampled in proportion to the learner's beliefs over candidate hypotheses, which may be guided by how diagnostic an action may be relative to competing alternative hypotheses [57,58]. The selected action is then executed with sensorimotor noise and evaluated through sensory feedback corrupted by perceptual noise. Inference then proceeds by updating these beliefs in light of sensory feedback, favoring hypotheses that make more specific predictions about the observed outcomes (i.e., the size principle) [58]. The resulting posterior guides subsequent action selection, enabling learners to rapidly converge on a strategy that counteracts the perturbation.

With the model now specified more formally, we can consider mechanistically how target geometry shapes hypothesis testing. Learners maintain priors over candidate spatial transformations, with simpler or more familiar rules (e.g., translations) often receiving greater weight than less familiar ones (e.g., rotations), though these priors vary across individuals. Because inner-target geometries are consistent with multiple transformations and are therefore less diagnostic, several hypotheses remain viable. Learners can then adopt whichever rule is most expedient—often the one with the highest prior for that individual—leading to faster improvements but greater diversity in the rules expressed across participants. In contrast, outer-target geometries provide stronger cues that differentiate competing hypotheses, reducing the set of rules that produce acceptable performance. Learners must therefore test and reject plausible alternatives—often those with high prior weight—before converging on a rule that generalizes across the workspace, slowing early learning but yielding more consistent adoption of the imposed rotation rule.

Moreover, hypothesis testing provides a principled account of individual differences in learning. Distinct learning phenotypes may arise from differences in prior beliefs and inference dynamics. Learners who exhibit sudden insight—shifting abruptly from baseline to successful compensation—may begin with a strong prior favoring no perturbation and update beliefs only after sufficient evidence accumulates, producing a rapid transition to the correct strategy with minimal exploration. In contrast, learners who explore more broadly may hold more diffuse priors over spatial hypotheses, leading to greater early exploration across multiple candidate transformations. Learners who converge on sign-flip (approximate translation) rules may be individuals who assign higher prior weight to simpler transformations—potentially reflecting greater everyday familiarity with translations than rotations.

Importantly, formalizing the hypothesis testing model also yields testable predictions about how learning transfers across contexts—often referred to as structural learning, meta-learning, or functional learning [25,59–63]. Hypothesis testing predicts that once a rotational mapping is inferred, belief in this hypothesis is strengthened, thereby facilitating the acquisition of future perturbations with similar rotational structures [64]. Conversely, such consolidation should impede learning of geometrically distinct mappings, such as reflections or translations. More broadly, this represents just one of many promising directions for future modeling work, including evaluating hypothesis testing across different training environments—such as compositional perturbation structures [64], conditions that limit cognitive resources (e.g., dual-tasking [10,65] or constrained preparation time), or interactions with other learning processes [66,67].

Consideration of Alternative Process Models. Could our data arise from mechanisms other than hypothesis testing? One recent proposal argues that strategic adaptation occurs through a discrete moment of insight, in which behavior jumps abruptly from baseline to the correct solution with no intervening exploration [68,69]. However, our data do not seem to support this account. Only a small subset of individuals exhibited such abrupt ‘aha’ moments, whereas most participants showed systematic exploratory errors. More importantly, although insight models can *describe* the form of behavioral change, they cannot *explain* why such insights arise. In contrast, hypothesis testing specifies a psychological process that gives rise to these “aha” moments: A subset of learners may entertain a small set of candidate hypotheses— “no perturbation” and the 60° rotation—and begin with a strong prior favoring no perturbation. As discrepant sensory evidence accumulates, posterior belief gradually shifts from the initially favored baseline hypothesis toward the imposed rotation hypothesis. Once its support exceeds a decision threshold, learners infer a rotated environment and abruptly re-aim to restore performance, producing the observed “aha” moments [70,71].

A second class of alternatives posits that strategic adaptation unfolds via model-free mechanisms, such as win–stay, lose–shift heuristics [72–75] or Q-value–based reinforcement learning [76,77]. While these processes could, in principle, generate broad exploration early in learning, they do not provide a compelling account of the structured, multimodal error distributions we observed—such as systematic sign flips or mirror reversals—where errors cluster around distinct solutions rather than being randomly distributed. Moreover, these model-free accounts typically rely on extended trial-and-error sampling, making them too slow to reconcile with the rapid learning observed in our data.

Finally, strategic adaptation has also been framed in terms of model-based reinforcement learning. Although the term is often used broadly [78], the core idea is that learners maintain an internal model linking actions to outcomes [79]. From this definition, we view our framework as a specific instantiation of model-based reinforcement learning. Indeed, many models described as Bayesian hypothesis testing can be reformulated as instances of model-based reinforcement learning, differing primarily in their algorithmic interpretation rather than their underlying computational structure [80]. An important direction for future work is to evaluate whether strategic learning is supported not only by hypothesis testing but also by additional mechanisms—such as meta-learning, choice stickiness, or structured priors—that may operate in parallel.

It is important to note that we are not claiming that hypothesis testing provides a superior account across all regimes of motor learning. For example, the gradual error reduction framework—typically formalized as a state-space or Kalman filter model—offers an excellent description of the implicit processes that support gradual sensorimotor recalibration [11,48,81]. Reinforcement learning effectively captures skill acquisition, as learners refine speed, precision, and timing through repetition [82,83]. Our aim is not to supplant these frameworks but to broaden the vocabulary of motor learning to include hypothesis testing—a process that likely complements [84], and at times competes with [5], these other processes (for a review, see ref. 85).

Methods

Participants and Apparatus. We recruited 560 participants through Prolific, an online participant recruitment website (312 females, aged 25.1 ± 3.4 years, all participants self-reported being right-handed). This sample size is substantially larger than that of typical in-person motor learning studies [40]. Before beginning the study, all participants provided informed consent in accordance with the Declaration of Helsinki and with policies approved by Carnegie Mellon University's Institutional Review Board (FWA00004206). Participation was in exchange for monetary compensation.

The experiment was created using the OnPoint platform, a package for running customized online motor learning experiments with JavaScript [40]. Participants completed the web-based experiment via an internet browser with their laptop's trackpad (Figure 1a). Our past online studies have shown that neither the type of browser nor the type of pointing device impacts motor performance [89]. The size and position of the visual stimuli were scaled to the individual's monitor size. Specifically, the workspace was centered on the screen midpoint, and target eccentricity was fixed at 25% of the screen height. All other stimuli were scaled according to the same proportion of screen height. For ease of interpretation, all stimulus parameters detailed below were based on an average 13-inch computer monitor.

General Procedure. We employed the Delayed Feedback Task to bias learning towards strategic adaptation [23,86,87] (Figure 1b). On each trial, participants positioned a white cursor (0.3 cm diameter) inside a white starting ring (0.5 cm diameter) at the center of the screen. Once the cursor entered the ring, it filled in. After holding the cursor there for 500 ms, the cursor disappeared, and a target (0.5 cm diameter) appeared on an invisible ring with a 4.5 cm radius. The target cued participants to make a rapid, straight movement to slice through it. When the hidden cursor crossed the invisible 4.5 cm ring, an auditory beep signaled movement termination. Visual cursor feedback reappeared 1000 ms after movement termination—a delay shown to minimize implicit processes underlying motor adaptation [88]—and remained visible for 500 ms. Participants then returned to the start position for the next trial.

Experiment 1. Participants (N = 280; 155 females; mean age = 25.9 ± 3.8 years) were randomly assigned to one of two groups. In the inner-target (unconstrained) group (N = 140), two blue training targets were positioned close together (20° apart) along with a single probe target positioned opposite the midpoint between the training targets (Figure 1c). In the outer-target (constrained) group (N = 140), the training targets were positioned farther apart (140° apart) with the same probe target. To ensure the effects were not specific to particular target locations, we counterbalanced target placement across participants (Table S1).

Trials were organized into triplets (3 trials/cycle): Two reaches to blue training targets with visual feedback followed by one reach to an orange probe target without feedback. The task comprised 132 trials (44 reaches/target), divided into three blocks: baseline (24 trials; 8 reaches/target), perturbation (96 trials; 32 reaches/target), and aftereffect blocks (12 trials; 4 reaches/target). During the baseline block, cursor feedback was veridical—aligned with the participant's actual movement direction—and was the only block in which visual feedback was provided for the probe target. In the perturbation block, a 60° visuomotor

rotation was applied to cursor feedback relative to the participant's movement direction, with the direction of rotation (clockwise vs counterclockwise) counterbalanced across participants. Rotated feedback was only provided at the training target; no feedback was given at the probe target. In the aftereffect block, no visual feedback was provided at either training or probe targets.

Before beginning the experiment, participants were given a general directive explaining the task structure and overall goal. They were instructed: "Move the white dot to the center of your screen. It will be hidden at first, but it will appear once you get close enough to the center. When your blue or orange target appears, quickly slice through it with the white dot. You'll hear a beep when the movement is completed. If there's feedback showing the final cursor position, it will appear shortly after the beep. After several attempts, the white dot will no longer follow the direction of your movement. When this happens, try your best to make the white dot hit the target. Throughout the entire experiment, your goal is to hit the target with your cursor."

Prior to the familiarization phase, participants were told, "Get ready to move quickly. The blue or orange target will then appear. Aim to hit the blue or orange target as quickly and accurately as possible. You'll hear a beep when the movement is completed." Before the perturbation phase, they were informed, "Your white cursor will appear somewhere different from where you move your hand. Your main objective is to identify the movement that gets your white cursor on the blue or orange target."

Before the first generalization trial in the perturbation phase, participants were informed: "Your white cursor will be hidden, but STILL different from where you move. Identify the movement needed to get your hidden cursor onto the ORANGE target. Use the same strategy you applied for the blue target." Before the aftereffect phase, participants were informed: "Your white cursor will now be hidden but will follow the direction of your movement. Aim DIRECTLY toward the blue or orange target to hit it as quickly and accurately as possible."

To ensure participants clearly understood the task, we embedded three instruction checks throughout the perturbation phase. After the first failed check, participants received a 10-second time penalty. Those who failed either of the subsequent two checks or who remained idle for more than one minute at any point were removed from the task and informed they could not continue. Restarting the experiment was not permitted.

Experiment 2. All procedures in Experiment 2 were identical to those in Experiment 1, except that Experiment 1 included *two* training targets, whereas Experiment 2 included *eight*. Participants (N = 280; 156 females; mean age = 24.9 ± 3.4 years) were randomly assigned to one of two groups. In the inner-target group (N = 140), the eight training targets were spaced 20° apart (Figure 1d). In the outer-target group (N = 140), the eight training targets were spaced 5° apart. In both groups, the nearest target lay 110° from the probe target, thereby equating this factor across inner- and outer-target conditions. To ensure the effects were not specific to particular target locations, we counterbalanced target placement across participants (Table S1). Trials were organized into triplets (3 trials/cycle). The task comprised 132 trials (11 trials/training target; 44 trials/probe target), consisting of baseline (24 trials; 2 trials/training target, 8 trials/probe target), perturbation (96 trials; 8 trials/training target, 32 trials/probe target), and aftereffect blocks (12 trials; 1 trial/training target, 4 trials/probe target).

Data analysis. All data processing and statistical analyses were conducted in Python 3.9.7. Analyses focused on hand position in screen coordinates (i.e., the position of the computer pointer) at the point where movement amplitude reached the target radius (4.5 cm). These data were used to calculate our main dependent variable, movement angle, defined as the angular distance between the movement and target positions. Movement angles across perturbation directions (clockwise vs. counterclockwise) were standardized so that positive values always indicated changes that counteracted the perturbation (i.e., movement angles for the counterclockwise rotation groups were multiplied by –1).

We summarized movement angles across three phases: early adaptation, late adaptation, and aftereffect. Early adaptation was defined as the average change in movement angle over the first 4 cycles of the perturbation block (cycles 9–12). Late adaptation was defined as the average movement angle over the last 4 cycles of the perturbation block (cycles 37-40). Aftereffect was defined as the average movement angle over all 4 cycles of the no-feedback aftereffect block (cycles 41-44). Because movement angle was measured in polar (angular) coordinates, mean angles were computed using circular statistics [96,97], and variability was quantified using circular standard error of the mean (SEM). All key results were also replicated under linear, non-circular statistics.

To evaluate group differences between inner (unconstrained) and outer-target (constrained) groups, we conducted a 2 (Target Type: training vs. probe) × 2 (Target Geometry: inner vs. outer) mixed-design ANOVAs on early adaptation and late adaptation phases. Target Type was a within-subjects factor and Target Geometry was a between-subjects factor. We report the F statistic, corresponding p-value, and effect sizes as partial eta squared (η^2). Post hoc tests were conducted using Watson two-sample tests for unpaired-sample circular comparisons and paired-sample within-sample circular comparisons.

Identifying Clustered Differences in Learning Functions. To identify differences between inner- and outer-target geometries without imposing arbitrary phase boundaries (early vs late), we applied a cluster-based permutation test to movement angles [91–93]. Training and probe trials were analyzed separately. Statistical comparisons were performed using the Watson two-sample test (U^2-values), which yields a p-value assessing whether two circular distributions differ significantly. Movement data were analyzed in two-cycle units. Units with a $p < .05$ were marked as candidates for clustering, and clusters were defined as two or more adjacent significant units. The cluster-level test statistic was defined as the sum of U^2-values across constituent units. Cluster significance was assessed against a null distribution generated from 5,000 permutations in which group labels were shuffled. For each permutation, the clustering procedure was repeated and the largest cluster sum recorded. Clusters were deemed reliable if their summed statistic exceeded the 95th percentile of the null distribution ($\alpha = .05$), thereby controlling the false positive rate. Effect sizes for each significant cluster were reported as (1) the absolute circular mean difference ($\Delta\mu$, in degrees) between groups, (2) circular Cohen's d and (3) the permutation-based cluster p-value [94,95].

Characterizing Multimodal Motor Error Distributions. To characterize the multimodal structure of participants' exploratory behavior, we fit a mixture of von Mises distributions to the movement angle data using an Expectation–Maximization (EM) algorithm [98,99]. EM iterations were capped at 1,000 steps with a convergence tolerance of 1×10^{-8}. The optimal number of mixture components (k=2–10) was determined via the Bayesian Information Criterion (BIC) [100]. For each k, the model was fit with 5,000 random initializations. For each component, we estimated three parameters: weight (proportion of movement angles assigned to that component), mean (degrees), and concentration parameter κ (the circular analogue of precision). We imposed four constraints to prevent degenerate model fits: (1) a minimum weight of 0.1% per component, (2) a minimum κ of 4.0, (3) a uniform distribution component to capture random responses with its weight constrained to not exceed 20%, and (4) a minimum angular separation of 25° between peaks. Additionally, to assess whether motor errors were distributed differently between the inner-target and outer-target groups, we computed Jensen–Shannon Divergence (JSD) [101]. A significant JSD ($p < 0.05$) would indicate that the groups produced distinct underlying distributions.

Characterizing the Effect of Target Geometry as a Continuum. To assess whether learning varied systematically across the four target geometries in a continuous manner, we conducted mixed-effects ANOVA in which Target Geometry was encoded along a hypothesized continuum (2 outer → 8 outer → 8 inner → 2 inner). We ran these analyses separately for training and probe trials. Evidence for a continuum would appear as a significant beta coefficient across geometries.

Identifying Subgroups of Learners. We sought to identify distinct subgroups of learners using a two-step procedure. First, because studies of strategic adaptation often include participants who show minimal learning or erratic exploration without convergence [44,87,90], we removed 'non-learners' prior to clustering. Specifically, we applied a one-sample V-test at the subject level to late-adaptation angles (μ_0 = 60°, $p < .05$) and a Watson two-sample test between the late-adaptation and aftereffect angle distributions ($p < .05$), yielding a sample of 247 out of 280 participants in Experiment 1 (88% learners) and 229 out of 280 participants in Experiment 2 (82% learners). Second, to identify distinct types of learners within this filtered learner sample, we applied Soft Dynamic Time Warping (Soft-DTW; $\gamma = 0.1$)[102] to each participant's learning function. Ward's hierarchical clustering was then performed on the resulting distance matrix. The optimal number of clusters was selected using both Silhouette Scores [103] and Within-Cluster Sum of Squares (WCSS) [104]. Finally, a chi-square test evaluated whether cluster frequencies differed across target-geometry conditions.

Data availability statement. All data is available on the Open Science Framework (OSF) via the OpenMotor repository (https://osf.io/aknqj/).

Code availability statement. All code is available on the Open Science Framework (OSF) via the OpenMotor repository (https://osf.io/aknqj/).

Acknowledgements. Not applicable.

Author contributions. W.D. and J.S.T. designed the study. W.D. conducted the experiments and analyzed the data. W.D., A.N., J.A.T. and J.S.T. interpreted the results. W.D. drafted the manuscript. A.N., J.A.T. and J.S.T. revised the manuscript critically for important intellectual content. All authors read and approved the final manuscript.

Competing interests. The authors declare no competing interests.

References

1. Shadmehr, R., Smith, M. A. & Krakauer, J. Error correction, sensory prediction, and adaptation in motor control. *Annu. Rev. Neurosci.* 33, 89 – 108 (2010).

2. McDougle, S. D., Ivry, R. B. & Taylor, J. A. Taking aim at the cognitive side of learning in sensorimotor adaptation tasks. *Trends Cogn. Sci.* 20, 535 – 544 (2016).

3. Taylor, J. A. & Ivry, R. B. Flexible cognitive strategies during motor learning. *PLoS Comput. Biol.* 7, e1001096 (2011).

4. Taylor, J. A. & Ivry, R. B. The role of strategies in motor learning. *Ann. N. Y. Acad. Sci.* 1251, 1 – 12 (2012).

5. Albert, S. T. et al. Competition between parallel sensorimotor learning systems. *eLife* 11, e65361 (2022).

6. Izawa, J. & Shadmehr, R. Learning from sensory and reward prediction errors during motor adaptation. *PLoS Comput. Biol.* 7, e1002012 (2011).

7. Uehara, S., Mawase, F., Therrien, A. S., Cherry-Allen, K. M. & Celnik, P. Interactions between motor exploration and reinforcement learning. *J. Neurophysiol.* 122, 797 – 808 (2019).

8. Cashaback, J. G. et al. The gradient of the reinforcement landscape influences sensorimotor learning. *PLoS Comput. Biol.* 15, e1006839 (2019).

9. Coltman, S. K., van Beers, R. J., Medendorp, W. P. & Gribble, P. L. Sensitivity to error during visuomotor adaptation is similarly modulated by abrupt, gradual, and random perturbation schedules. *J. Neurophysiol.* 126, 934 – 945 (2021).

10. Haith, A. M., Huberdeau, D. M. & Krakauer, J. W. The influence of movement preparation time on the expression of visuomotor learning and savings. *J. Neurosci.* 35, 5109 – 5117 (2015).

11. Smith, M. A., Ghazizadeh, A. & Shadmehr, R. Interacting adaptive processes with different timescales underlie short-term motor learning. *PLoS Biol.* 4, e179 (2006).

12. Albert, S. T. & Shadmehr, R. Estimating properties of the fast and slow adaptive processes during sensorimotor adaptation. *J. Neurophysiol.* 119, 1367 – 1393 (2018).

13. Townsend, M., Mon-Williams, M., Mushtaq, F. & Morehead, R. Explicit aiming solutions are gained through insight. *Neural Control of Movement*, Victoria, Canada (2023).

14. Tsay, J. S. et al. Fundamental processes in sensorimotor learning: Reasoning, refinement, and retrieval. *eLife* 13, e91839 (2024).

15. Chen, Y., Abram, S., Ivry, R. B. & Tsay, J. S. Indirect feedback hinders explicit sensorimotor adaptation. *Proc. R. Soc. B* 292, 20251407 (2025).

16. McDougle, S. D. et al. Neural signatures of prediction errors in a decision-making task are modulated by action execution failures. *Curr. Biol.* 29, 1606 – 1613 (2019).

17. Gallistel, C. R., Fairhurst, S. & Balsam, P. The learning curve: Implications of a quantitative analysis. *Proc. Natl Acad. Sci.* USA 101, 13124 – 13131 (2004).

18. Abe, M. O. & Sternad, D. Directionality in distribution and temporal structure of variability in skill acquisition. *Front. Hum. Neurosci.* 7, 225 (2013).

19. Sternad, D. It's not (only) the mean that matters: variability, noise and exploration in skill learning. *Curr. Opin. Behav. Sci.* 20, 183 – 195 (2018).

20. Sternad, D., Abe, M. O., Hu, X. & Müller, H. Neuromotor noise, error tolerance and velocity-dependent costs in skilled performance. *PLoS Comput. Biol.* 7, e1002159 (2011).

21. Niyogi, A., Cisneros, E., Ivry, R. & Tsay, J. S. Computational mechanisms underlying strategy discovery. *Society for Neuroscience* (2024).

22. Taylor, J. A. & Ivry, R. B. Context-dependent generalization. *Front. Hum. Neurosci.* 7, 171 (2013).

23. Brudner, S. N., Kethidi, N., Graeupner, D., Ivry, R. B. & Taylor, J. A. Delayed feedback during sensorimotor learning selectively disrupts adaptation but not strategy use. *J. Neurophysiol.* 115, 1499 – 1511 (2016).

24. Heald, J. B., Lengyel, M. & Wolpert, D. M. Contextual inference underlies the learning of sensorimotor repertoires. *Nature* 600, 489 – 493 (2021).

25. Heald, J. B., Lengyel, M. & Wolpert, D. M. Contextual inference in learning and memory. *Trends Cogn. Sci.* 27, 43 – 64 (2023).

26. McDougle, S. D. & Taylor, J. A. Dissociable cognitive strategies for sensorimotor learning. *Nat. Commun.* 10, 40 (2019).

27. Poh, E. & Taylor, J. A. Generalization via superposition: combined effects of mixed reference frame representations for explicit and implicit learning in a visuomotor adaptation task. *J. Neurophysiol.* 121, 1953 – 1966 (2019).

28. Fitts, P. M. & Posner, M. I. *Human Performance* (Brooks/Cole, 1967).

29. Anguera, J. A., Reuter-Lorenz, P. A., Willingham, D. T. & Seidler, R. D. Contributions of spatial working memory to visuomotor learning. *J. Cogn. Neurosci.* 22, 1917 – 1930 (2010).

30. Cowley, M. & Byrne, R. M. Chess masters' hypothesis testing. *In Proc. 26th Annual Meeting of the Cognitive Science Society* (2004).

31. Newell, A. & Simon, H. A. *Human Problem Solving* (Prentice-Hall, 1972).

32. Allen, K. R., Smith, K. A. & Tenenbaum, J. B. Rapid trial-and-error learning with simulation supports flexible tool use and physical reasoning. *Proc. Natl Acad. Sci.* USA 117, 29302 – 29310 (2020).

33. Anderson, J. R. Acquisition of cognitive skill. *Psychol. Rev.* 89, 369 – 406 (1982).

34. Steyvers, M., Tenenbaum, J. B., Wagenmakers, E. J. & Blum, B. Inferring causal networks from observations and interventions. *Cogn. Sci.* 27, 453 – 489 (2003).

35. Glautier, S. Revisiting the learning curve (once again). *Front. Psychol.* 4, 982 (2013).

36. Mercado, E., III. Neural and cognitive plasticity: From maps to minds. *Psychol. Bull.* 134, 109 – 137 (2008).

37. Braver, T. S., Cole, M. W. & Yarkoni, T. Vive les differences! Individual variation in neural mechanisms of executive control. *Curr. Opin. Neurobiol.* 20, 242 – 250 (2010).

38. Donner, Y. & Hardy, J. L. Piecewise power laws in individual learning curves. *Psychon. Bull. Rev.* 22, 1308 – 1319 (2015).

39. Maggi, S. et al. Tracking subjects' strategies in behavioural choice experiments at trial resolution. *eLife* 13, e86491 (2024).

40. Tsay, J. S., Ivry, R. B., Lee, A. & Avraham, G. Moving outside the lab: The viability of conducting sensorimotor learning studies online. *Neurons Behav. Data Anal. Theory* 5, 1 – 22 (2021).

41. Seidler, R. D., Mulavara, A. P., Bloomberg, J. J. & Peters, B. T. Individual predictors of sensorimotor adaptability. *Front. Syst. Neurosci.* 9, 100 (2015).

42. Ranganathan, R., Cone, S. & Fox, B. Predicting individual differences in motor learning: A critical review. *Neurosci. Biobehav. Rev.* 141, 104852 (2022).

43. Malone, L. A. et al. The control of movement gradually transitions from feedback control to feedforward adaptation throughout childhood. *npj Sci. Learn.* 10, 13 (2025).

44. Cisneros, E., Karny, S., Ivry, R. B. & Tsay, J. S. Differential Aging Effects on Implicit and Explicit Sensorimotor Learning. *bioRxiv* (2024).

45. Leow, L. A., Marinovic, W., de Rugy, A. & Carroll, T. J. Task errors contribute to implicit aftereffects in sensorimotor adaptation. *Eur. J. Neurosci.* 48, 3397 – 3409 (2018).

46. Molinaro, G. & Collins, A. G. A goal-centric outlook on learning. *Trends Cogn. Sci.* 27, 1150 – 1164 (2023).

47. Tsay, J. S., Haith, A. M., Ivry, R. B. & Kim, H. E. Interactions between sensory prediction error and task error during implicit motor learning. *PLoS Comput. Biol.* 18, e1010005 (2022).

48. Padmanabhan, S., Shadmehr, R., Klatzky, R. L. & Tsay, J. S. Goal uncertainty attenuates sensorimotor adaptation. *J. Neurophysiol.* 135, 1 – 10 (2026).

49. Villavicencio, P., Tsay, J. S. & de la Malla, C. Target configuration determines how and what we learn during sensorimotor adaptation. *npj Sci. Learn.* 10, 89 (2025).

50. Leech, K. A., Roemmich, R. T., Gordon, J., Reisman, D. S. & Cherry-Allen, K. M. Updates in motor learning: implications for physical therapist practice and education. *Phys. Ther.* 102, pzab250 (2022).

51. Roemmich, R. T. & Bastian, A. J. Closing the loop: from motor neuroscience to neurorehabilitation. *Annu. Rev. Neurosci.* 41, 415 – 429 (2018).

52. Tsay, J. S. & Winstein, C. J. Five Features to Look for in Early-Phase Clinical Intervention Studies. *Neurorehabil. Neural Repair* 35, 3 – 9 (2021).

53. Abram, S. J., Tsay, J. S., Yosef, H., Reisman, D. S. & Kim, H. E. The Detrimental Effect of Stroke on Motor Adaptation. *Neurorehabil. Neural Repair* 39, 213 – 225 (2025).

54. Cisek, P. & Green, A. M. Toward a neuroscience of natural behavior. *Curr. Opin. Neurobiol.* 86, 102859 (2024).

55. Tsay, J. et al. Bridging the gap between experimental control and ecological validity in human sensorimotor science. *J. Physiol.* 602, 3857 – 3876 (2024).

56. Fooken, J. et al. Perceptual-cognitive integration for goal-directed action in naturalistic environments. *J. Neurosci.* 43, 7511 – 7522 (2023).

57. Klayman, J. & Ha, Y. W. Hypothesis testing in rule discovery: Strategy, structure, and content. J. Exp. Psychol. *Learn. Mem. Cogn.* 15, 596 – 604 (1989).

58. Tenenbaum, J. B. & Griffiths, T. L. Generalization, similarity, and Bayesian inference. *Behav. Brain Sci.* 24, 629 – 640 (2001).

59. Bond, K. M. & Taylor, J. A. Structural learning in a visuomotor adaptation task is explicitly accessible. *eNeuro* 4, ENEURO.0122-17.2017 (2017).

60. Braun, D. A., Mehring, C. & Wolpert, D. M. Structure learning in action. *Behav. Brain Res.* 206, 157 – 165 (2010).

61. Wolpert, D. M., Diedrichsen, J. & Flanagan, J. R. Principles of sensorimotor learning. *Nat. Rev. Neurosci.* 12, 739 – 751 (2011).

62. Wolpert, D. M. & Flanagan, J. R. Motor learning. *Curr. Biol.* 20, R467 – R472 (2010).

63. Sugiyama, T., Uehara, S. & Izawa, J. Meta-learning of human motor adaptation via the dorsal premotor cortex. *Proc. Natl Acad. Sci.* USA 121, e2417543121 (2024).

64. Tian, L. Y. et al. Neural representation of action symbols in primate frontal cortex. *bioRxiv* (2025).

65. Taylor, J. A. & Thoroughman, K. A. Divided attention impairs human motor adaptation but not feedback control. *J. Neurophysiol.* 98, 317 – 326 (2007).

66. Chen, Y. & Taylor, J. A. Revisiting the explicit-implicit additivity assumption in visuomotor adaptation. *bioRxiv* (2025).

67. Marius 't Hart, B. et al. Measures of implicit and explicit adaptation do not linearly add. *eNeuro* 11, ENEURO.0190-24.2024 (2024).

68. Townsend, M. et al. An "Aha!" moment precedes the strategic response to a visuomotor rotation. *bioRxiv* (2025).

69. Telgen, S. J. *Motor Learning During Reaching Movements: Model Acquisition and Recalibration.* PhD thesis, UCL (2015).

70. Kaplan, C. A. & Simon, H. A. In search of insight. *Cogn. Psychol.* 22, 374 – 419 (1990).

71. Becker, M., Sommer, T. & Cabeza, R. Insight predicts subsequent memory via cortical representational change and hippocampal activity. *Nat. Commun.* 16, 4341 (2025).

72. Cashaback, J. G., McGregor, H. R., Mohatarem, A. & Gribble, P. L. Dissociating error-based and reinforcement-based loss functions during sensorimotor learning. *PLoS Comput. Biol.* 13, e1005623 (2017).

73. Therrien, A. S., Wolpert, D. M. & Bastian, A. J. Effective reinforcement learning following cerebellar damage requires a balance between exploration and motor noise. *Brain* 139, 101 – 114 (2016).

74. van Mastrigt, N. M., van der Kooij, K. & Smeets, J. B. Pitfalls in quantifying exploration in reward-based motor learning and how to avoid them. *Biol. Cybern.* 115, 365 – 382 (2021).

75. van Mastrigt, N. M. et al. Implicit reward-based motor learning. *Exp. Brain Res.* 241, 2287 – 2298 (2023).

76. Collins, A. G. & Frank, M. J. How much of reinforcement learning is working memory, not reinforcement learning? A behavioral, computational, and neurogenetic analysis. *Eur. J. Neurosci.* 35, 1024 – 1035 (2012).

77. Collins, A. G. The tortoise and the hare: Interactions between reinforcement learning and working memory. *J. Cogn. Neurosci.* 30, 1422 – 1432 (2018).

78. Collins, A. G. E. & Cockburn, J. Beyond dichotomies in reinforcement learning. *Nat. Rev. Neurosci.* 21, 576 – 586 (2020).

79. Daw, N. D., Niv, Y. & Dayan, P. Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nat. Neurosci.* 8, 1704 – 1711 (2005).

80. Eckstein, M. K. et al. The interpretation of computational model parameters depends on the context. *eLife* 11, e75474 (2022).

81. Burge, J., Ernst, M. O. & Banks, M. S. The statistical determinants of adaptation rate in human reaching. *J. Vis.* 8, 20 (2008).

82. Haith, A. M. Policy-Gradient Reinforcement Learning as a General Theory of Practice-Based Motor Skill Learning. *bioRxiv* (2025).

83. Seethapathi, N., Clark, B. C. & Srinivasan, M. Exploration-based learning of a stabilizing controller predicts locomotor adaptation. *Nat. Commun.* 15, 9498 (2024).

84. Miyamoto, Y. R., Wang, S. & Smith, M. A. Implicit adaptation compensates for erratic explicit strategy in human motor learning. *Nat. Neurosci.* 23, 443 – 455 (2020).

85. Therrien, A. S. & Wong, A. L. Mechanisms of human motor learning do not function independently. *Front. Hum. Neurosci.* 15, 785992 (2022).

86. Schween, R. & Hegele, M. Feedback delay attenuates implicit but facilitates explicit adjustments to a visuomotor rotation. *Neurobiol. Learn. Mem.* 140, 124 – 133 (2017).

87. Tsay, J. S., Schuck, L. & Ivry, R. B. Cerebellar degeneration impairs strategy discovery but not strategy recall. *Cerebellum* 22, 1223 – 1233 (2023).

88. Kitazawa, S., Kohno, T. & Uka, T. Effects of delayed visual information on the rate and amount of prism adaptation in the human. *J. Neurosci.* 15, 7644 – 7652 (1995).

89. Tsay, J. S. et al. Large-scale citizen science reveals predictors of sensorimotor adaptation. *Nat. Hum. Behav.* 8, 510 – 525 (2024).

90. Jang, J., Shadmehr, R. & Albert, S. T. A software tool for at-home measurement of sensorimotor adaptation. *bioRxiv* (2023).

91. Breska, A. & Ivry, R. B. Context-specific control over the neural dynamics of temporal attention by the human cerebellum. *Sci. Adv.* 6, eabb1141 (2020).

92. Tsay, J. S., Parvin, D. E. & Ivry, R. B. Continuous reports of sensed hand position during sensorimotor adaptation. *J. Neurophysiol.* 124, 1122 – 1130 (2020).

93. Wang, T., Avraham, G., Tsay, J. S., Thummala, T. & Ivry, R. B. Advanced feedback enhances sensorimotor adaptation. *Curr. Biol.* 34, 1076 – 1085 (2024).

94. Mardia, K. V. & Jupp, P. E. *Directional Statistics* (Wiley, 2000).

95. Pewsey, A., Neuhäuser, M. & Ruxton, G. D. *Circular Statistics in R* (Oxford Univ. Press, 2013).

96. Batschelet, E. Circular Statistics in Biology (Academic Press, 1981).

97. Fisher, N. I. *Statistical Analysis of Circular Data* (Cambridge Univ. Press, 1995).

98. Dempster, A. P., Laird, N. M. & Rubin, D. B. Maximum likelihood from incomplete data via the EM algorithm. *J. R. Stat. Soc. B* 39, 1 – 22 (1977).

99. Banerjee, A., Dhillon, I. S., Ghosh, J. & Sra, S. Clustering on the Unit Hypersphere using von Mises-Fisher Distributions. *J. Mach. Learn. Res.* 6, 1345 – 1382 (2005).

100. Schwarz, G. Estimating the dimension of a model. *Ann. Stat.* 6, 461 – 464 (1978).

101. Lin, J. Divergence measures based on the Shannon entropy. *IEEE Trans. Inf. Theory* 37, 145 – 151 (1991).

102. Cuturi, M. & Blondel, M. Soft-DTW: a differentiable loss function for time-series. *In Proc. 34th International Conference on Machine Learning* 894 – 903 (PMLR, 2017).

103. Rousseeuw, P. J. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *J. Comput. Appl. Math.* 20, 53 – 65 (1987).

104. McQueen, J. B. Some methods of classification and analysis of multivariate observations. *In Proc. 5th Berkeley Symposium on Mathematical Statistics and Probability* 281 – 297 (1967).

Figure 1. Experimental Overview. (a) Web-Based Experimental Setup**.** Participants completed a visuomotor adaptation task using their computer trackpad. **(b)** The Delayed Feedback Task isolates strategic motor adaptation. Endpoint cursor feedback, rotated by 60° (hollow black circle), was displayed 1000 ms after the hand (in screen coordinates) reached the target distance—a manipulation known to minimize implicit adaptation. Participants were instructed to explore different movements to align the rotated cursor with the target. Importantly, they were not told about the nature of the perturbation. **(c)** Experiment 1: Inner vs. outer target geometries. Participants reached to two training targets (blue circles) spaced either far apart (outer group, 140°) or close together (inner group, 20°). A single probe target (orange circle) was positioned opposite the midpoint between the training targets. **(d)** Experiment 2: Inner vs. outer target geometries. Eight training targets were arranged either along the outer edge of the circle (outer group, 5° spacing) or more uniformly across the workspace (inner group, 20° spacing). Unlike Experiment 1, the distance between training and probe targets was equated across groups. The gray circle was not shown to participants and is included here only for illustrative purposes. **(e)** Example movements (cursor pointer) and cursor outcome (hollow black circles) at the two training targets in Experiment 1 for the Outer- and Inner-target groups. **(f)** Possible movements at the no-feedback probe target, each reflecting a learned action–outcome hypothesis: the experimentally-imposed rotation (aiming 60° away clockwise from the target; * denotes the imposed rule) or one of several alternative rules, including an approximate translation (60° counterclockwise; also known as "sign flips"), a mirror reversal (aiming 180° away from the target), or a default reach-to-target response (aiming directly to the target; 0°). **(g)** Reinforcement learning process. Movement errors are expected to be unimodal and confined between baseline and the 60° solution, shifting smoothly from early to late learning at both the training (blue) and probe (orange) targets, with variability reflecting diffuse exploration and sensorimotor noise. This framework does not predict qualitative differences in the rules learned across target geometries. **(h)** Hypothesis testing process. At the training targets (blue, left), learners sample multiple discrete hypotheses early in learning, producing multimodal error distributions, and converge on the experimentally-imposed rotation by late adaptation. At the probe target (orange, right), the predictions diverge by condition: in the outer (constrained) environment, learners converge on the imposed 60° rotation, whereas in the inner (unconstrained) environment, learners may settle on either the imposed rotation or a simpler, alternative rule (e.g., an approximate translation, or "sign flip"), yielding multiple peaks in late adaptation. R1–R4 denote Rules 1–4 shown in panel (f).

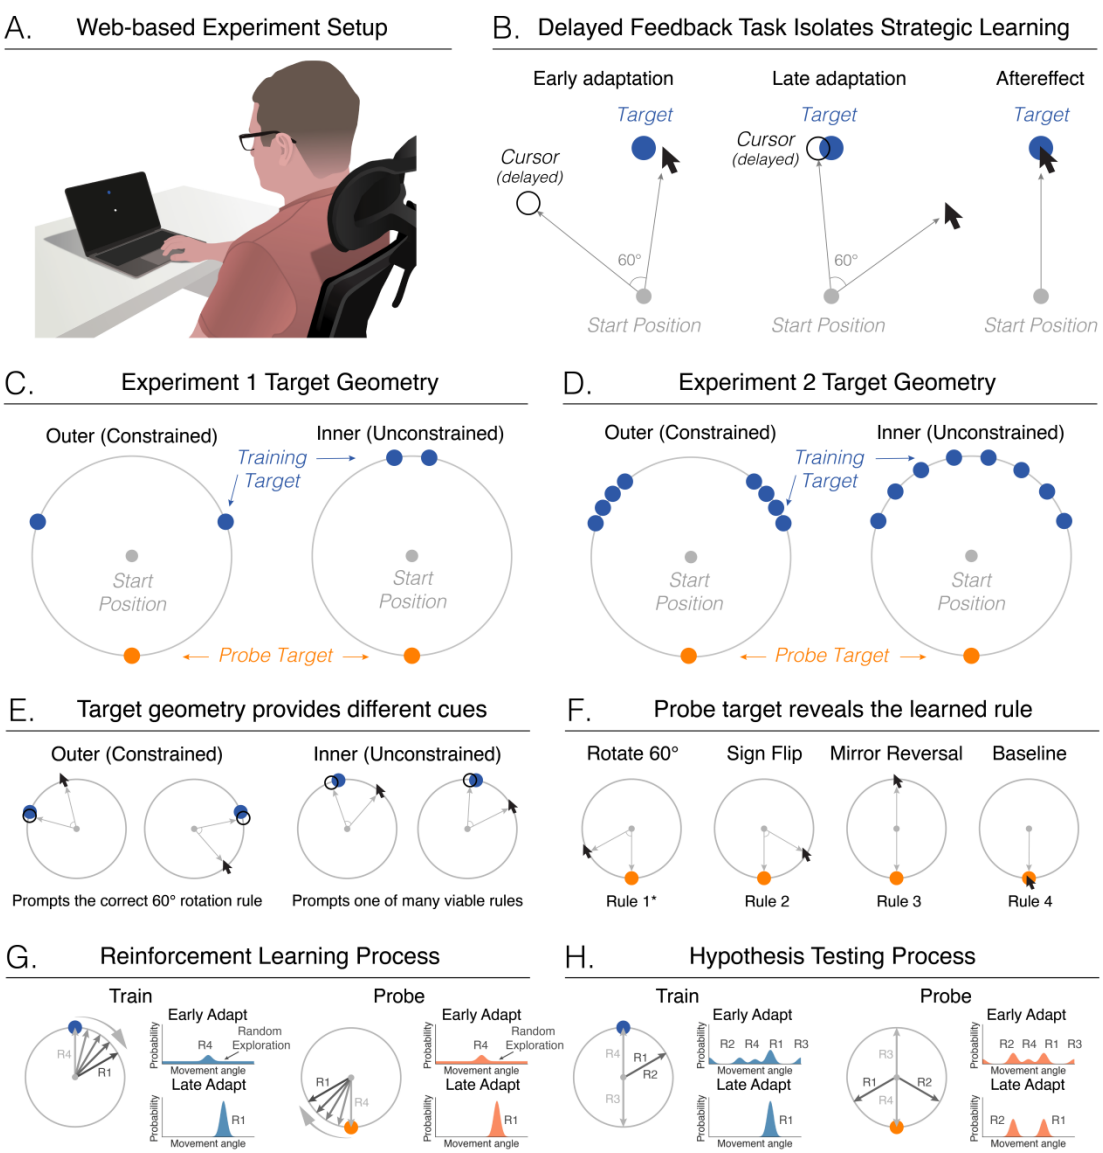

Figure 2. Hypothesis Testing governs Strategic Motor Adaptation. (a) Learning curves for training **(left)** and probe targets **(right)** in Experiment 1. Mean movement angle (°) is plotted for the outer-target (dark green) and inner-target (light green) groups; shaded regions denote SEM. Vertical dashed lines indicate perturbation onset and offset. Horizontal gray bars mark significant group differences identified via cluster-based permutation tests ($p < .05$). **(b)** Representative learner from the outer-target group. Filled green dots denote movements to training targets; open grey dots denote movements to probe targets. **(c)** Within-participant standard deviation of movement angles (motor variability) at the training target during baseline, early adaptation, and late adaptation. Error bars reflect SEM. **(d)** Heatmaps show the density of movement angles across participants on each trial for the outer-target and inner-target groups, separately for training (left) and probe targets (right). Color intensity reflects the proportion of participants at each angle (darker = higher density). Black lines denote the group average, as in panel A. **(e)** Movement angle distributions for training (top) and probe (bottom) targets (density). Solid black lines show von Mises mixture fits; dashed red lines mark individual mixture components; histograms are rescaled to match the density range. Vertical dotted lines denote the targets (0°, Rule 4 in Figure 1f), the required rotation solution (60°, Rule 1) and the mirror reversal solution (180°, Rule 3). Sign flips (–60°, Rule 2)—correct-magnitude reaches in the opposite direction—are highlighted in red for illustration. **(f)** Distributional differences (density) between inner- and outer-target groups for the training and probe targets **(left)**. Negative values indicate greater density in the inner-target group; positive values indicate greater density in the outer-target group. Group differences are corroborated by von Mises component weights (binned in 60° intervals) and the uniform (U) component **(right)**.

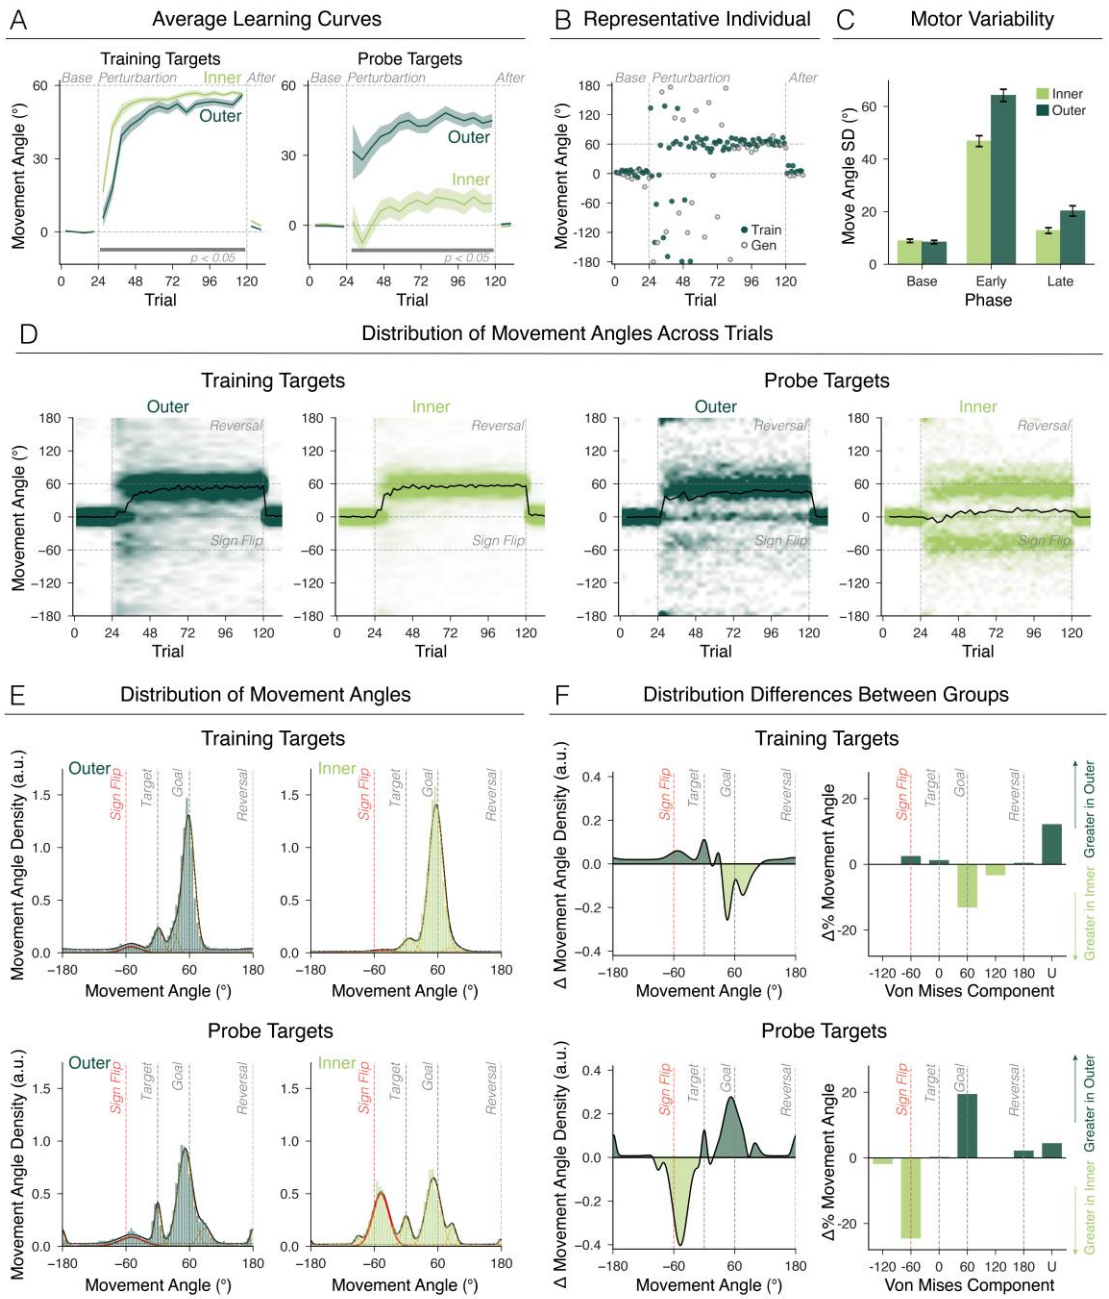
A
Average Learning Curves
Training Targets
Probe Targets
Base
Perturbation
Inner
Outer
After
Movement Angle (°)
Trial
p < 0.05
B
Representative Individual
Train
Gen
C
Motor Variability
Move Angle SD (°)
Base
Early
Late
Phase
D
Distribution of Movement Angles Across Trials
Reversal
Sign Flip
E
Distribution of Movement Angles
Movement Angle Density (a.u.)
Movement Angle (°)
Sign Flip
Target
Goal
Reversal
F
Distribution Differences Between Groups
Δ Movement Angle Density (a.u.)
Δ% Movement Angle
Von Mises Component
Greater in Outer
Greater in Inner

Figure 3. Replication of Hypothesis Testing in Strategic Motor Learning. (a) Learning curves for training **(left)** and probe targets **(right)** in Experiment 2. Mean movement angle (°) is plotted for the outer-target (purple) and inner-target (pink) groups; shaded regions denote SEM. Vertical dashed lines indicate perturbation onset and offset. Horizontal gray bars mark significant group differences identified via cluster-based permutation tests ($p < .05$). **(b)** Representative learner from the outer group. Filled purple dots correspond to movements to the training targets; open gray dots correspond to movements to the probe targets. **(c)** SD of within-subject training-target movement angles across phases. "Base" = baseline, "Early" = early adaptation, "Late" = late adaptation. Error bars indicate SEM. **(d)** Heatmaps show the density of movement angles across participants on each trial for the outer-target and inner-target groups, separately for training (left) and probe targets (right). Color intensity reflects the proportion of participants at each angle (darker = higher density). Black lines denote the group average, as in panel A.**(e)** Movement angle distributions for training (top) and probe (bottom) targets. Solid black lines show von Mises mixture fits; dashed red lines mark individual mixture components; histograms are rescaled to match the density range. Vertical dotted lines denote the targets (0°), the required rotation solution (60°) and the mirror reversal solution (180°). Sign flips (–60°)—correct-magnitude reaches in the opposite/incorrect direction—are highlighted in red for illustration. **(f)** Distributional differences of von Mises fit (density) between inner- and outer-target groups for training and probe targets **(left)**. Negative values indicate greater density in the inner-target group; positive values indicate greater density in the outer-target group. These differences are further illustrated by comparing the von Mises component weights (binned in 60° intervals) and the uniform component **(right)**. Critically, outer-target learners exhibit greater exploration at the training targets (i.e., greater sampling of non-solution regions), a pattern that reverses at the probe targets, consistent with the hypothesis-testing account.

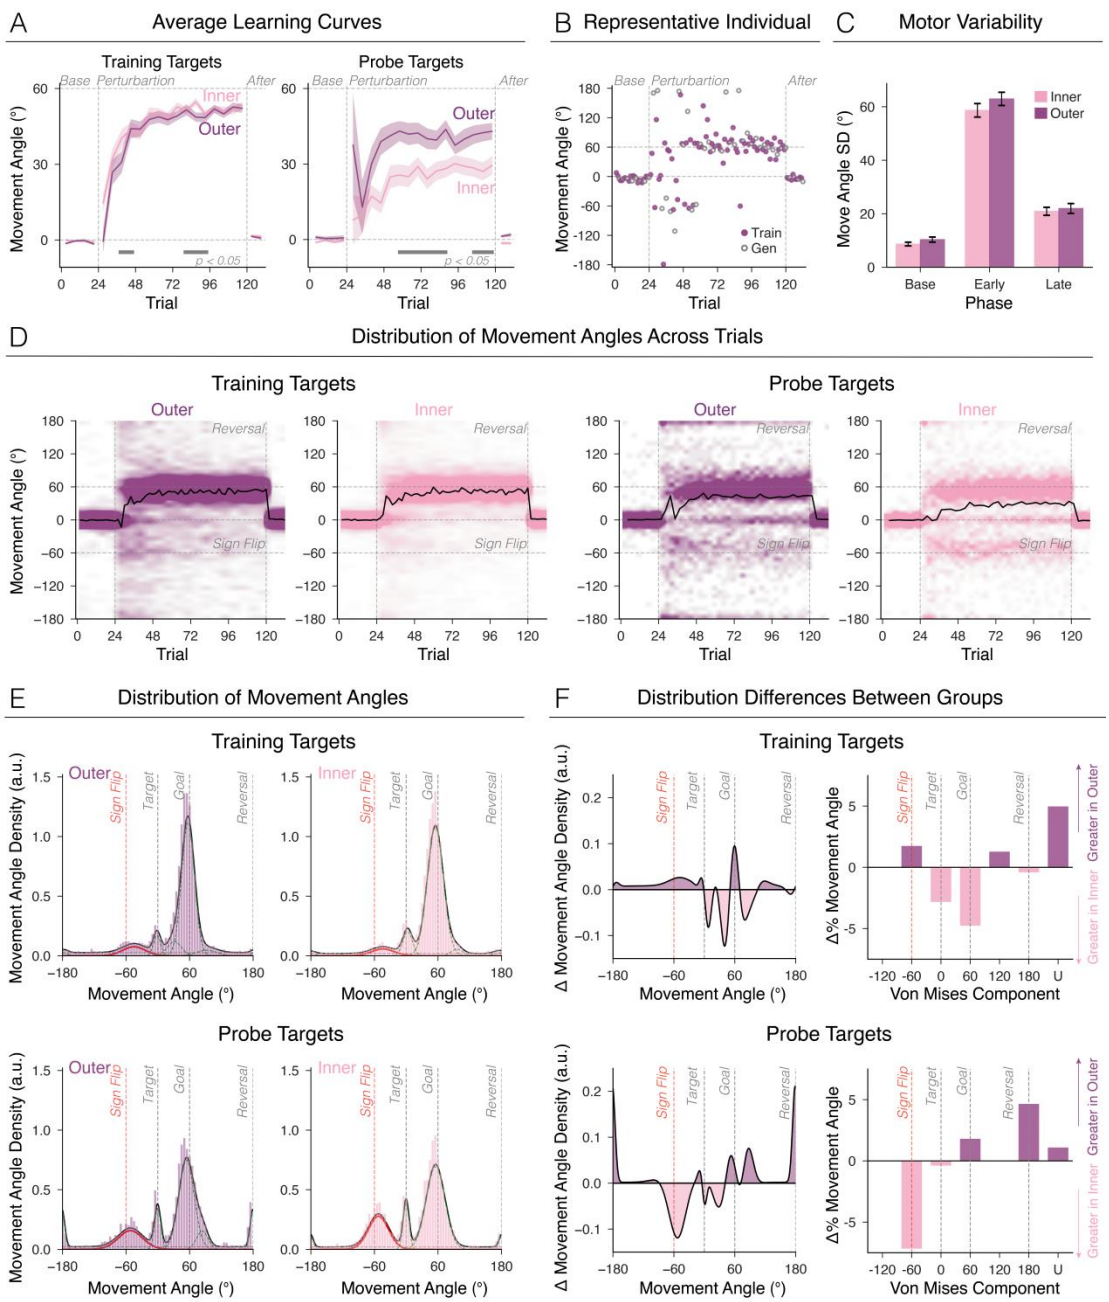

Figure 4. Environmental Constraint Modulates Individual Learning Strategies. (a) Mean movement angle during early adaptation across the four experimental conditions: Experiment 1 (2-inner, 2-outer) and Experiment 2 (8-inner, 8-outer). Error bars indicate ± SEM across participants. (b) Total von Mises mixture weight at the non-solution (60°) locations across the four experimental conditions, indexing the degree of motor exploration across groups. (c) Dendrogram and t-SNE embedding showing learner clusters identified via pairwise distances computed using Soft Dynamic Time Warping. Colors denote cluster membership (C1–C7). (d) Cluster distribution as a function of target geometry. Stacked bar plots show the proportion of participants in each cluster across the four target-geometry conditions. Percentages < 5% are omitted for clarity. (e) Representative learning trajectories for each cluster, marked by a star in panel C. Gray dots indicate training trials; colored dots denote probe performance. (f, g) Movement-angle distributions during early (f) and late (g) adaptation across the seven learner clusters. Solid lines with shade and dashed lines represent the distribution of movement angles at training and probe target locations, respectively.

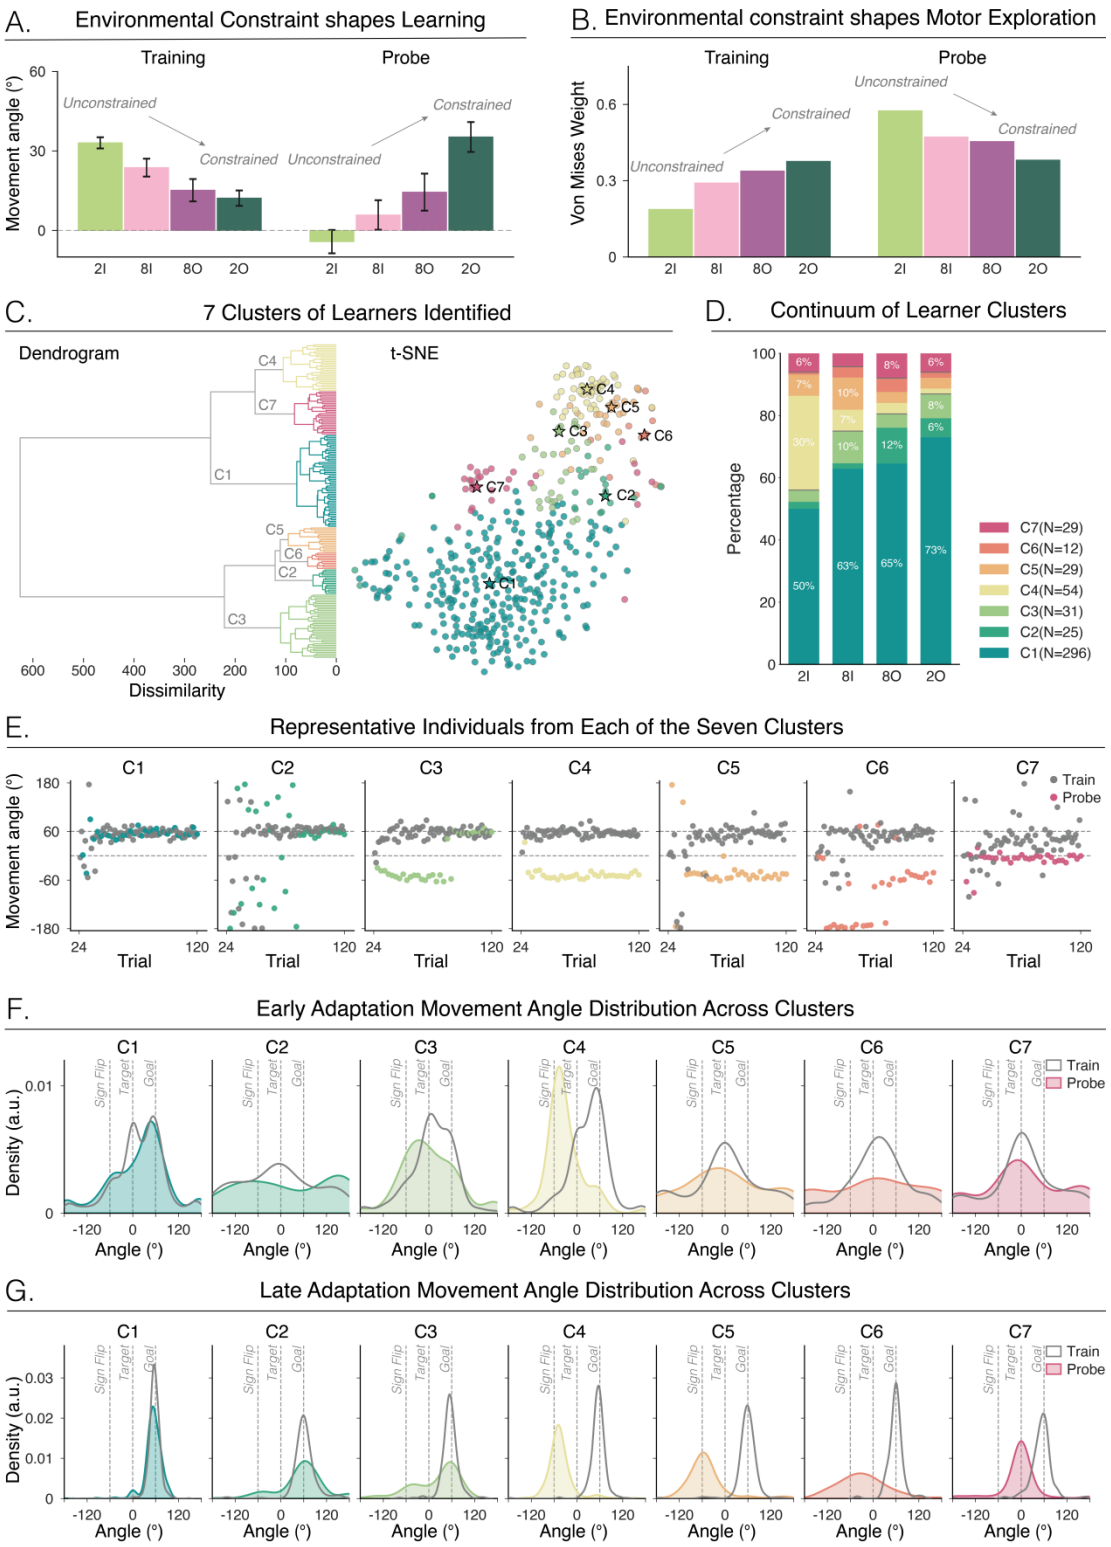
A.
Environmental Constraint shapes Learning
Training
Probe
Movement angle (°)
Unconstrained
Constrained
2I
8I
8O
2O
B.
Environmental constraint shapes Motor Exploration
Von Mises Weight
C.
7 Clusters of Learners Identified
Dendrogram
t-SNE
Dissimilarity
C1
C2
C3
C4
C5
C6
C7
D.
Continuum of Learner Clusters
Percentage
C7(N=29)
C6(N=12)
C5(N=29)
C4(N=54)
C3(N=31)
C2(N=25)
C1(N=296)
E.
Representative Individuals from Each of the Seven Clusters
Trial
Train
Probe
F.
Early Adaptation Movement Angle Distribution Across Clusters
Density (a.u.)
Sign Flip
Target
Goal
Angle (°)
G.
Late Adaptation Movement Angle Distribution Across Clusters