Visuomotor memory is not bound to visual motion

Matthew Warburton[1]*, Carlo Campagnoli[1], Mark Mon-Williams[1,2,3], Faisal Mushtaq[1,4], J. Ryan Morehead[1,4]

[1]School of Psychology, University of Leeds, [2]Bradford Institute for Health Research, Bradford, West Yorkshire, UK, [3]National Centre for Optics, Vision and Eye Care, University of South-Eastern Norway, Kongsberg, Hasbergs vei 36, 3616, Norway, [4]Centre for Immersive Technologies, University of Leeds

* Corresponding author: pscmwa@leeds.ac.uk

Abstract

Sensory feedback plays a critical role in motor control and motor learning, with both processes adjusting outgoing motor commands relative to the error between actual sensory feedback and the predictions of a forward model. However, models of motor control rarely specify the exact nature of these predictions. We hypothesized that large differences in low-level perceptual feedback would delineate contextual boundaries for motor memory, as sufficiently different feedback would require a distinct mapping between motor commands and sensory states. We tested this hypothesis by measuring transfer of visuomotor adaptation across contexts where hand movements caused visual motion in opposite directions (180° away). Instead of observing that visual feedback is bound to distinct internal models, we found nearly complete transfer of learning across the contexts. Finally, we found evidence that the motor memory was bound to the planned displacement of the hand, rather than visual features of the task space.

Keywords: Visuomotor adaptation, generalization, video games, learning

Introduction

Forward models are thought to play a crucial role in guiding sensorimotor behavior through the prediction of the sensory consequences of movement [1–3]. These predictions have been shown to be useful for attenuating self-generated sensory signals [4,5] and also for robust control of ongoing behavior [6,7]. In the latter case, feedback control theories propose that behavior is guided by a sensory state estimate that combines observed and predicted sensory feedback [6,8,9], with provisions for multiple internal models when the relationship between a motor command and the predicted sensory feedback is context-dependent [10,11]. However, most models lack clarity on the level at which these sensory predictions operate – for example whether a prediction of movement-contingent visual feedback relies on early visual processing carried out in the primary visual cortex [12], complex motion analysis specific to higher visual areas like MT [13], or mixed reference frames that combine gaze and body positions [14].

This lack of clarity on the characterization of internal model sensory predictions is felt when considering behaviors that differ only in terms of their visual feedback. We recently compared the same center-out reaching task in a mouse *Pointing* context, where a hand movement translates a cursor across a static background, to a first-person shooter (FPS) game inspired *Looking* visual context, where the same movement pans and tilts the view of the environment while the cursor remains central to the screen [15]. Movement kinematics and performance in these contexts were highly correlated within individuals,

despite profound visual differences, suggesting that both contexts may be making use of the same internal model to plan and correct movements on the fly. If true, this ability for the same internal model to operate across visual feedback contexts cannot be accounted for by prominent models of motor control, which pair forward models of sensory feedback with a movement control policy.

Here we directly probe whether an internal model is shared between mouse Look and Point feedback by measuring the transfer of visuomotor adaptation between these contexts. Transfer of adaptation is a well-established method to infer properties of internal models [16–18]. If learning under a visuomotor perturbation, thought to rely on sensory prediction errors [19,20], transfers across Point and Look contexts, it will require a reconsideration of current models of sensorimotor control and learning. Specifically, it is unclear how much learning in a forward model [8,21] and/or control policy [22,23] contributes to sensorimotor adaptation. Critically, adaptation that transfers across mouse Point and Look conditions cannot be explained by learning of a forward model that makes sensory predictions based on low-level processing of visual features or on more complex visual motion signals. Such learning would require that internal model adaptation involve a higher, more abstract level of representation such as visuospatial displacement vectors.

To investigate whether internal models are sensitive to the specificity of visual feedback provided for a common movement, we introduced a 30° visuomotor rotation (VMR) between a reaching movement, made with a computer mouse or trackpad, and the visual feedback of the reach, requiring that participants adapt their movements to reinstate good performance [24,25] (see supplemental videos). We first compared the adaptive response to the perturbation and the exhibited after-effects, finding nearly identical learning profiles between contexts. Next, we examined transfer of learning between contexts, observing near-complete transfer of learning upon a context switch both when the perturbation was either left on or turned off, confirming the use of a shared internal model across visual contexts. Finally, we found that the learning generalized around the planned movement of the hand, rather than the visual movement of the cursor or the location of the target in the workspace.

Results

To assess participants' adaptation to perturbations in an FPS-style experiment, we utilized a framework allowing visual feedback to be compared between and within participants [15]. Participants used their mouse or trackpad (learning did not differ between input devices across all experiments, so is not considered hereafter) to control a cursor on the screen, appearing like an FPS-style aiming reticle. Upon clicking a start-point, a target immediately appeared, and participants attempted to move to and click on the target (Figure 1a). Computer mouse or trackpad input could control the appearance of the game in two contexts (Figure 1b). In the Point context, movements translated the in-game cursor across a static background. This is consistent with how people normally interact with their computer's desktop environment, and standard practice for the study of visuomotor adaptation [24,26–28]. In the Look context, input movements panned and tilted the in-game camera's view of the virtual environment while the cursor remained central to the screen, consistent with how FPS games are designed. For example, moving the mouse *forward* would cause the in-game character's view to pan *up*. Critically, the two contexts were equated such that the same input movement in either context would reach a given target, allowing unbiased comparison of behavior between the contexts (see [15] for details).

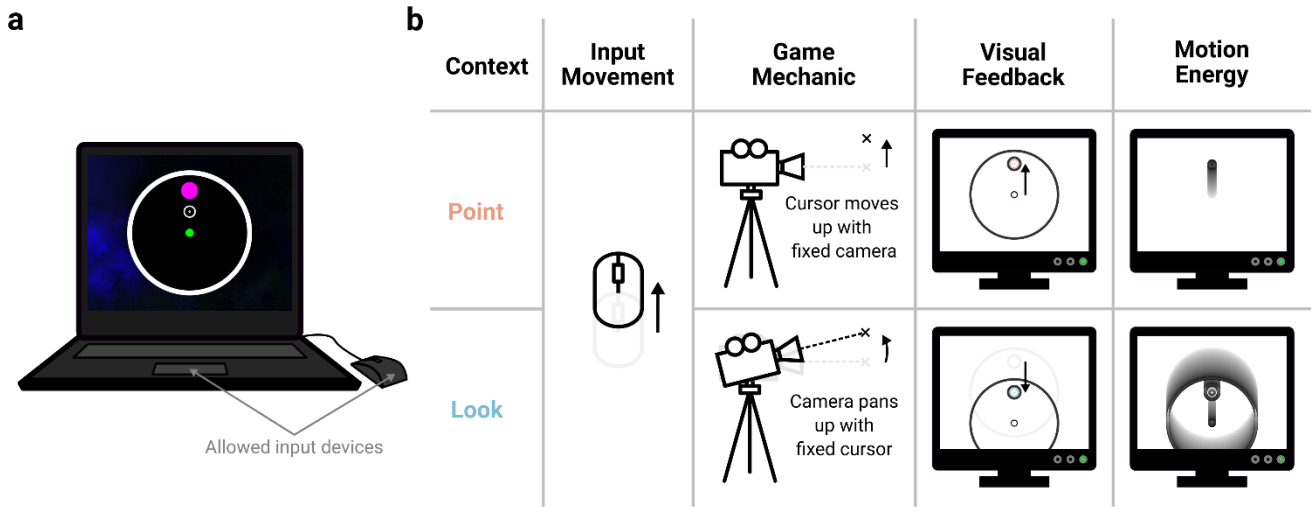

Figure 1. Task setup. **(a)** A visual representation of the task as seen by a participant in the Point mode. Participants used their personal computers and could interact with the task using either a trackpad or computer mouse. Participants left-clicked a start-point and were shown one of four targets to move to. Participants attempted to move to and left-click the target within a continuously staircased time-limit. **(b)** Differences between the contexts, illustrated for a common 'up' movement of the input device. In the Point mode, the in-game camera is stationary while a cursor moves across the background, giving visual feedback of only the cursor moving. In the Look mode, the in-game camera pans while the cursor is fixed central to the camera's view, giving visual feedback where everything except the cursor moves. This gives rise to very different experienced motion energy as the movement evolves over time.

Participants adapt similarly in the Point and Look contexts.

In Experiment 1, we began by assessing whether there were differences in how participants adapted to a perturbation in either context (Figure 2). The perturbation causes different visual errors depending on the context, which the participant must use for both online feedback corrections and trial-to-trial adaptation (Figure 2a). During a trial, upon clicking the start-point, a target immediately appeared in one of four locations, located on an imaginary circle in 90° increments. Participants had to move to and click on the target within a time limit to 'shoot' it, otherwise the target would disappear. The time limit was continually staircased throughout the experiment to maintain a ~50% success rate, inducing the type of time pressure players of FPS games might experience.

Participants first completed a baseline block with veridical movements, showing average movements that were straight to the target (Figure 2b). A visuomotor perturbation was then introduced, where the visual feedback of a movement was rotated around the start position by 30° from the input movement, requiring participants to compensate for this perturbation to be successful. Early movements (first ten cycles) during this phase showed a large corrective movement near to the target but had straightened by the end of the block (last ten cycles). Once the perturbation ceased, both groups showed similar "hooked" movements in the opposite direction to overcome after-effects, which were still present by the end of the washout period.

This learning is quantified in the average hand angles (measured at peak speed) per group over the time course of the experiment (Figure 2c). During the late rotation period, where learning appeared to have become asymptotic, participants in both groups compensated by 21.3° [Point: 95% CI = 19.4° – 23.2°, Look: 19.9° – 22.8°]. To assess differences over the learning block, hand angles over pre-defined regions were assessed using a 2 (context, Point vs Look) x 2 (period, early vs late rotation) mixed ANOVA (Figure 2d).

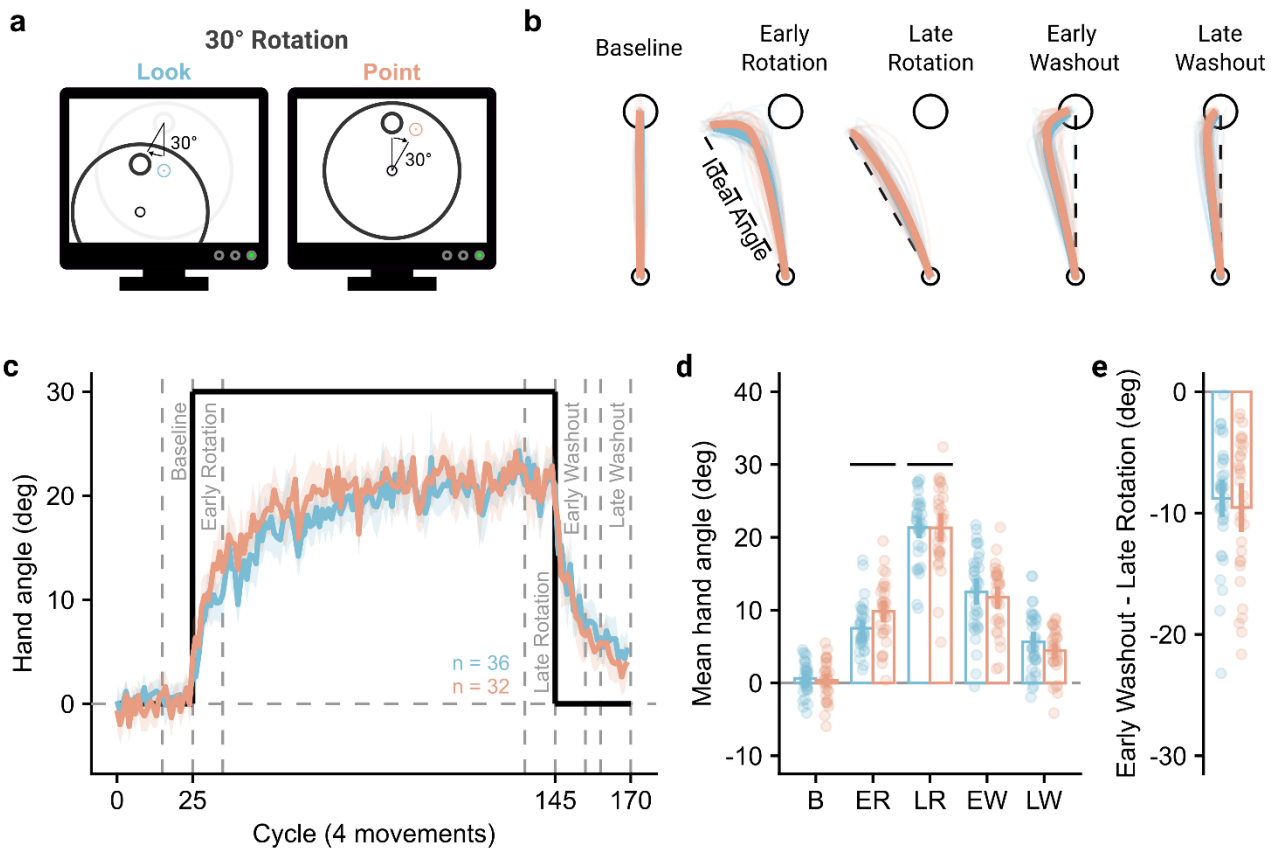

Figure 2. Participants adapted similarly across contexts. **(a)** During the experiment, participants (n = 68) experienced a visuomotor rotation that meant in-game cursor movements were rotated 30° from the veridical movement direction. **(b)** Participant's hand movements during the perturbation phase reflected compensation for the applied perturbation, with hand paths straightening over the rotation block, and also reflected after-effects, with deviations from baseline reaches at the end of a washout period. Faint lines show individual participants' average hand path during the period, and thick lines show the group average over participants. **(c)** Time course of the mean hand angle averaged over cycles (4 movements). The line shows the average hand angle from the target direction over participants, with the shaded area showing 95% confidence intervals. The labels show periods of cycles over which measures of learning were operationalised. **(d)** Mean hand angles over periods of the experiment. The period labels are abbreviated from those defined in panel (e). **(e)** Within-participant difference between early washout and asymptotic learning. For panels (f) and (g), points show individual participants, bars show the group mean, and vertical bars show the 95% confidence intervals.

The ANOVA showed a main effect of period ($F(1, 66) = 560.36$, $p < .001$, $BF_{10} > 100$, $\eta_p^2 = .90$), indicating participants had greater hand angles during late rotation, but no main effect of context ($F(1, 66) = 1.47$, $p = .230$, $BF_{10} = 0.51$, $\eta_p^2 = .02$). There was, however, a significant interaction between context and period ($F(1, 66) = 4.77$, $p = .032$, $BF_{10} = 1.77$, $\eta_p^2 = .07$). Participants in the Point group reached a higher compensatory hand angle than those in the Look group during early learning (estimated marginal mean difference [95% confidence interval] = 2.30° [0.46°, 4.15°], $t(66) = 2.49$, $p = .015$, $BF_{10} = 3.33$), but were not significantly different during late learning (-0.03° [-2.46°, 2.40°], $t(66) = -0.03$, $p = .980$ $BF_{10} = 0.25$). We also characterized the after-effect as the difference between early washout and late learning, to account for any potential differences in asymptotic learning that may have arisen (Figure 2e). The decrease in hand angle from late learning to early washout was not significantly different between the Point and Look groups (-0.74° [-3.29°, 1.81°], $t(66) = 0.58$, $p = .565$, $BF_{10} = 0.29$).

Learning transfers from a trained context to the other untrained context.

Experiment 1 found some evidence that the rate of adaptation was faster for the Point context early in the rotation phase, but this did not persist to asymptotic learning. Overall, behavior was strikingly similar between contexts, reaching similar asymptotic levels and decaying in a similar fashion upon removal of the perturbation. The similarity of learning, as well as the similarity across multiple movement metrics in a previous study [15], is consistent with (but insufficient to prove) the use of a shared internal model. To test this more directly, we trained participants to counteract a perturbation in one context before switching to the other, untrained, context. If these contexts do not share an internal model, we would expect poor transfer of learning to the untrained context.

In Experiment 2a, participants were assigned a context to learn to counteract the perturbation in. An experimental trial mimicked Experiment 1 but followed a different block schedule. Participants performed a pair of unperturbed baseline blocks in both the untrained and trained context before learning to counteract a 30° visuomotor rotation in the trained context. To observe the expression of after-effects following a context switch, participants then performed an unperturbed block in the untrained context, before finally performing an unperturbed block in the trained context.

There appeared to be a greater difference between contexts than in Experiment 1, with greater hand angles for the group training in the Point context (Figure 3a). A 2 (trained context, Point vs Look) x 2 (period, early vs late rotation) mixed ANOVA comparing pre-defined periods of the learning block (Figure 3b) showed there was a main effect of period ($F(1, 60) = 604.56$, $p < .001$, $BF_{10} > 100$, $\eta_p^2 = .91$), indicating hand angles were greater during late rotation, and a main effect of trained context ($F(1, 60) = 13.35$, $p < .001$, $BF_{10} = 51.24$, $\eta_p^2 = .18$), indicating hand angles were greater for participants trained in the Point context, with no significant interaction between period and context ($F(1, 60) = 0.32$, $p = .575$, $BF_{10} = 0.29$, $\eta_p^2 = .01$).

We next checked for differences in the expression of after-effects. If the learning was bound specifically to the trained context, we would expect a sharp drop in hand angle upon switching to the untrained context. The reduction in hand angle from late learning to early transfer (Figure 3c) was not significantly different between trained contexts (0.41° [-1.45°, 2.27°], $t(60) = 0.44$, $p = .660$, $BF_{10} = 0.28$). Critically, there was no significant difference in the reduction of hand angle when participants switched context, compared to when participants did not switch context in Experiment 1, for either the group trained in the Point context (-1.02° [-3.31°, 1.28°], $t(62) = 0.88$, $p = .380$, $BF_{10} = 0.36$) or Look context (-2.16° [-4.40°, 0.07°], $t(64) = 1.93$, $p = .058$, $BF_{10} = 1.21$). This indicates there was no evidence for a reduction in after-effects when participants switched to an untrained context.

Similar analyses were performed for Experiment 2b, where the perturbation was left on as participants were switched to the untrained context. Hand angles again appeared to be greater for those trained in the Point context (Figure 3d). The mixed ANOVA comparing rotation periods (Figure 3e) showed there was a main effect of period ($F(1, 64) = 557.21$, $p < .001$, $BF_{10} > 100$, $\eta_p^2 = .90$), with greater hand angles during late rotation, a main effect of trained context ($F(1, 64) = 10.02$, $p = .002$, $BF_{10} = 14.01$, $\eta_p^2 = .14$), with greater hand angles for participants trained in the Point context, and no significant interaction between period and context ($F(1, 64) = 0.18$, $p = .674$, $BF_{10} = 0.27$, $\eta_p^2 < .01$).

Despite the perturbation being left on in this task, we would still expect a substantial drop in hand angle immediately after transfer if the contexts did not share an internal model. The change in hand angle from

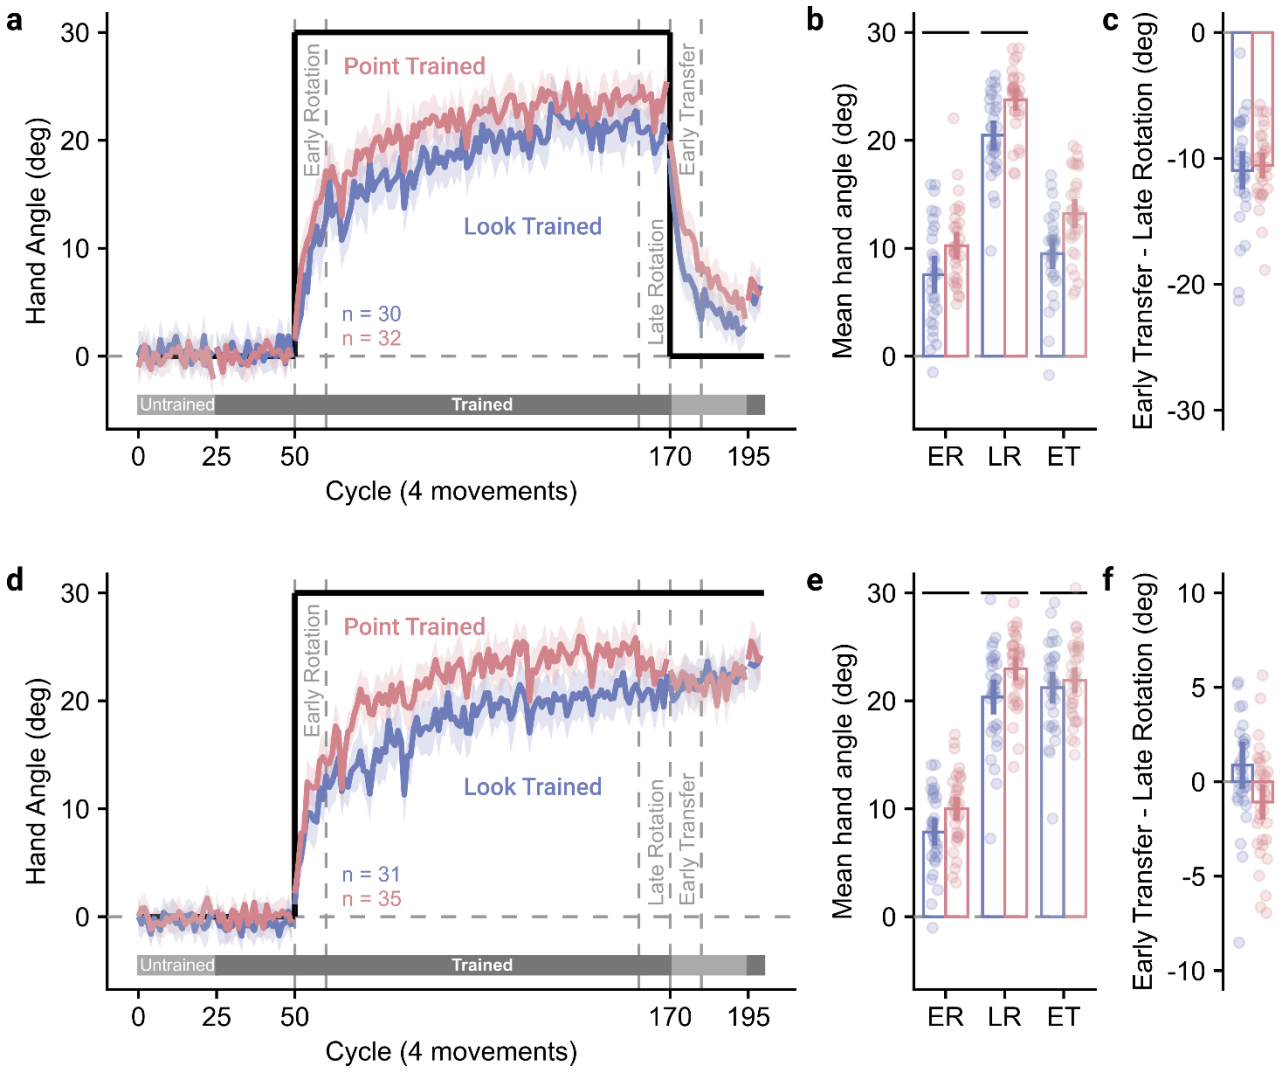

Figure 3. Learning transfers between context. **(a)** Time course of the mean hand angle averaged over cycles (4 movements). The line shows the average hand angle from the target direction over participants, with the shaded area showing 95% confidence intervals. The labels show periods of cycles over which measures of learning were operationalised. Shaded horizontal bars underneath show whether participants were currently completing trials in the trained or untrained context. **(b)** Mean hand angles over periods of the experiment. The period labels are abbreviated from those defined in panel (a). **(c)** Within-participant difference between early transfer and late learning. For panels (b) and (c), points show individual participants, bars show the group mean, and vertical bars show the 95% confidence intervals. **(d-f)** As (a-c) with the perturbation left on when switched to the untrained context.

late learning to early transfer (Figure 3f) was significantly lower for participants who trained in the Point context compared to the Look context (-1.95 [-3.53°, -0.37°], t(64) = -2.47, p = .016, BF_{10} = 3.18). Participants who trained in the Point context showed a small but significant reduction in hand angle from late learning to early transfer (-1.07 [-2.05°, -0.09°], t(34) = -2.23, p = .032, BF_{10} = 1.56), whereas participants who trained in the Look context showed no significant difference in hand angle between late learning and early transfer (0.88 [-0.43°, 2.19°], t(30) = 1.37, p = .180, BF_{10} = 0.45).

To explore this further, we performed linear mixed-effect regressions per trained context between hand angle (corrected for late learning) and cycle over the full untrained context transfer block (cycles 170-195). For participants trained in the Point mode, the intercept was significantly lower than asymptotic learning

upon switching to the Look mode (β = -1.28 [-2.33°, -0.24°], t = -2.41, p = .019), with no significant change over cycles (β = 0.01 [-0.03°, 0.05°], t = 0.53, p = .599). For participants trained in the Look mode, the intercept was not significantly different than asymptotic learning upon switching to the Point mode (β = 0.11 [-0.82°, 1.05°], t = 0.23, p = .816), but hand angle significantly increased over cycles (β = 0.07 [0.03°, 0.12°], t = 3.15, p = .002). Both regressions suggest asymptotic learning is slightly lower in the Look context – with those transferred to it experiencing a rapid decrease in hand angle and those transferred from it able to increase their compensatory angle slightly. Nevertheless, both experiments clearly show that the bulk of learning transfers to the untrained context.

Learning generalizes around the planned movement of the hand

The experiments so far are consistent in that learning proceeds similarly across context, with slightly greater learning in the Point context, and learning transfers between trained and untrained contexts, despite the difference in visual feedback. We suggest this shows the contexts do share a common internal model, and that learning generalizes around a feature common to both tasks. Visuomotor rotations induce learning that generalizes locally around the *planned* movement vector [29,30], decaying according to a Gaussian function as the difference in movement angle increases. Whether the plan is anchored in hand-space or cursor-space is unclear – the two are typically confounded. To investigate this, we used a manipulation taken from the FPS gaming domain, where a subset of games and gamers 'invert' their mapping, such that they move their mouse or joystick *forward* to look *down* in the game (Figure 4a). The advantage of this Inverted Look mapping is that it allows the, usually confounded, planned vectors of the hand and cursor to be dissociated (Figure 4b). The Inverted Look mapping should therefore allow us to dissociate whether learning generalizes locally around the planned vector of the hand, where it would generalize roughly around the trained movement direction; or the planned vector of the cursor, where it would generalize around the direction opposite the trained movement.

In Experiment 3, participants were split into two groups where the spatial generalization of learning was assessed with either Look or Inverted Look probes (Figure 4c, shown for Inverted Look probes). Trials during the first baseline block proceeded as in the previous experiments, with two exceptions – there was no time limit, and the upcoming context was signaled by a color cue when participants returned to the start-point. Participants were then introduced to a no-feedback return procedure for the Point context – upon outcome feedback being presented the cursor disappeared, the participant moved their hand back to a comfortable position, and their cursor was reset to the center of the screen. From this trial onwards, Point trials never had return feedback, whereas the probed context (either Look or Inverted Look) always had return feedback. Participants then completed the baseline generalization block, where 15 generalization cycles were performed. Each cycle had three trials in the probed context, first to either the left or right target and then to both the top and bottom targets, followed by 15 Point trials to the top target. No perturbation was applied during this block, and it allowed the later generalization block to be baseline-corrected, removing systematic biases in reaches [31].

Participants next experienced 120 Point trials to the top target (Figure 4c, rotation block), where they learned to counteract a 30° visuomotor rotation. Critically, because participants did not have return feedback during these trials, learning should only generalize locally around this trained movement direction, which was common between groups. Both groups learned to counteract the perturbation similarly (Figure 4d). A 2 (probed context, Look vs Inverted Look) x 2 (period, early vs late) mixed ANOVA comparing pre-defined periods of the learning block reiterated this, showing there was a main effect of

period ($F(1, 74) = 251.69$, $p < .001$, $BF_{10} > 100$, $\eta_p^2 = .77$), indicating hand angles were greater during late rotation, but no main effect of probed context ($F(1, 74) = 0.96$, $p = .331$, $BF_{10} = 0.35$, $\eta_p^2 = .01$) or interaction between period and probed context ($F(1, 74) = 0.04$, $p = .849$, $BF_{10} = 0.24$, $\eta_p^2 < .01$).

Participants then completed the generalization block, which was identical to the baseline generalization block except that the 15 Point trials per cycle had the 30° visuomotor rotation applied. Because unlearning will occur during the unperturbed probe trials, enough relearning trials (established in a pilot study) were allowed to 'top up' learning for the top target. During the generalization cycles (Figure 4e), hand angles during the probe trials were greater for the Look group, indicating a greater magnitude of generalization overall, but lower over the top-up trials, though the latter effect is driven by two participants in the Inverted Look group who overcompensated for the rotation (average hand angles over the last 5 rotation top-up trials of 43.17° and 54.61°).

Taking just the probe trials in the same and opposite direction to the trained movement (Figure 4f), learning generalized around the trained movement direction for both groups, indicating learning is bound to planned hand vectors. A 2 (probed context, Look vs Inverted Look) x 2 (movement direction from trained, 0° vs 180°) mixed ANOVA showed a significant main effect of movement direction ($F(1, 74) = 197.04$, $p < .001$, $BF_{10} > 100$, $\eta_p^2 = .73$), indicating hand angle was greater for the trained movement direction; probed context ($F(1, 74) = 9.64$, $p = .003$, $BF_{10} = 9.66$, $\eta_p^2 = .12$), indicating hand angle was greater for the Look probes; and a significant interaction ($F(1, 74) = 13.88$, $p < .001$, $BF_{10} = 95.26$, $\eta_p^2 = .16$). This interaction was followed up by assessing the differences in hand angle between trained and untrained movement direction. Hand angle was greater for the trained movement direction compared to the untrained movement direction in both Look (11.61° [9.81°, 13.40°], $t(74) = 12.90$, $p < .001$, $BF_{10} > 100$) and Inverted Look groups (6.74° [4.85°, 8.63°], $t(74) = 7.11$, $p < .001$, $BF_{10} > 100$; Figure 4g). Therefore, both groups show a robust transfer of learning from the trained Point context to the probed Look or Inverted Look context, supporting the previous experiment's results, and are consistent in that learning generalizes around the trained movement direction of the hand.

This difference was, however, significantly greater for the Look group compared to the Inverted Look group (4.87° [2.26°, 7.47°], $t(74) = 3.73$, $p < .001$, $BF_{10} = 69.94$). A mixed-effect linear regression, assessing how hand angle for the trained input direction changed over generalization cycle, showed there was no significant difference between probed context on the first generalization cycle ($\beta = -2.17$ [-5.02°, 0.69°], $t = -1.48$, $p = .140$). The Look group showed no significant change in hand angle over cycles ($\beta = 0.04$ [-0.12°, 0.20°], $t = 0.45$, $p = .651$), but the Inverted Look had a significant decrease in hand angle over cycles ($\beta = -0.29$ [-0.53°, -0.05°], $t = -2.40$, $p = .017$). Therefore, while hand angle is not significantly lower for the Inverted Look group at the start of the generalization block, the groups diverged over the cycles. This could indicate that, for the Inverted Look group, transferred learning itself reduced over cycles, which could happen if the number of top-up trials was not enough for the group to re-establish a consistent amount of implicit learning, or that participants formed an explicit strategy to overcome the implicit learning shown in the trained input direction, which was possible given participants had full vision during the probe trials.

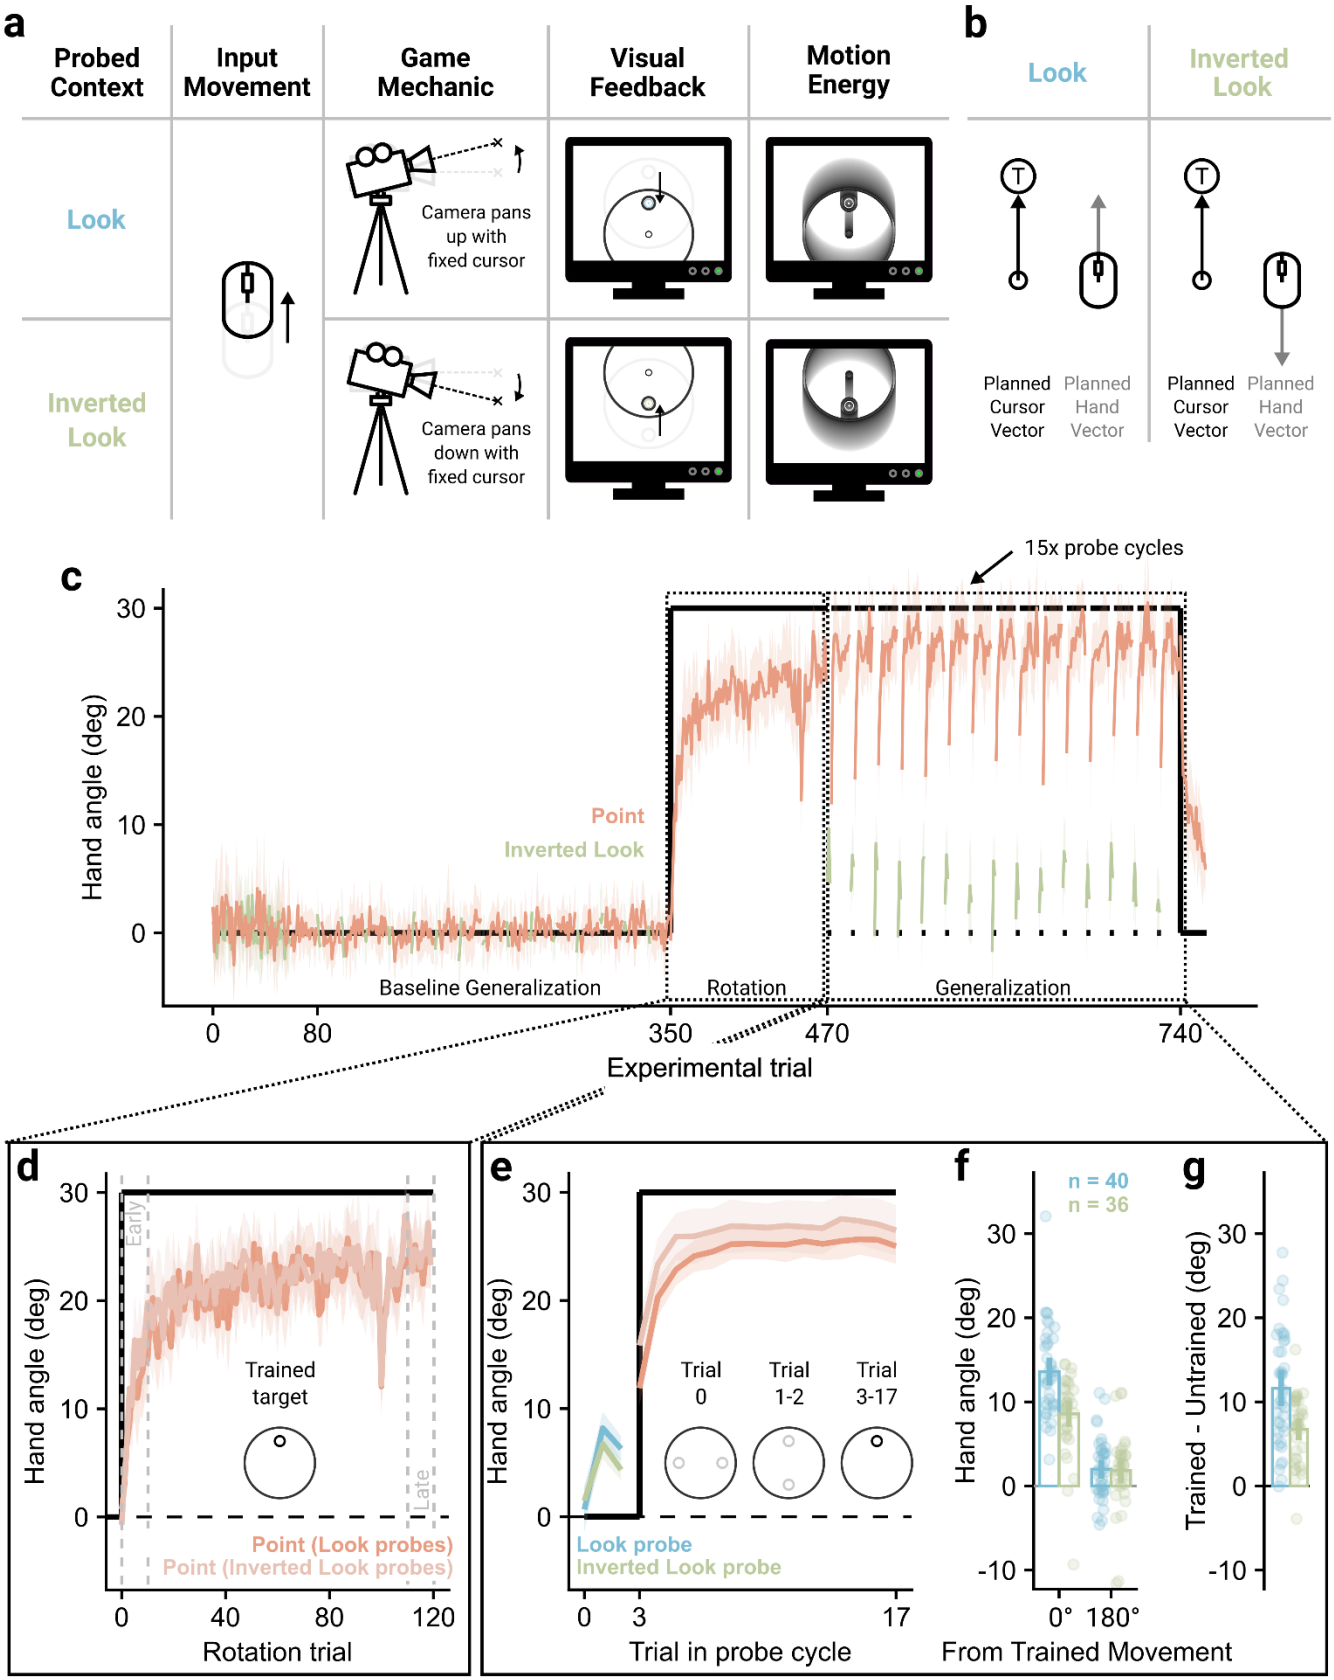
a
Probed Context
Input Movement
Game Mechanic
Visual Feedback
Motion Energy
Look
Inverted Look
Camera pans up with fixed cursor
Camera pans down with fixed cursor
b
Look
Inverted Look
Planned Cursor Vector
Planned Hand Vector
Planned Cursor Vector
Planned Hand Vector
c
15x probe cycles
Hand angle (deg)
Point
Inverted Look
Baseline Generalization
Rotation
Generalization
Experimental trial
d
Hand angle (deg)
Early
Late
Trained target
Point (Look probes)
Point (Inverted Look probes)
Rotation trial
e
Hand angle (deg)
Trial 0
Trial 1-2
Trial 3-17
Look probe
Inverted Look probe
Trial in probe cycle
f
n = 40
n = 36
Hand angle (deg)
From Trained Movement
g
Trained - Untrained (deg)

Figure 4. Learning generalizes around the planned hand vector. **(a)** A new context, Inverted Look, was added. When participants move their mouse forward, their in-game view pans down. **(b)** This allows the planned movement of the hand and cursor to be dissociated, where in the Look context they are confounded. **(c)** Hand angle across the experiment, for the group with Inverted Look probes. Participants performed a baseline block, and a baseline generalization with unperturbed feedback. Participants then learned to counteract an imposed 30° visuomotor rotation to a single trained target in the Point context, and completed a generalization block (with the same structure as the baseline generalization block). Participants finally completed an unperturbed washout block to the trained target with Point feedback. **(d)** Participants in the Look and Inverted Look probe groups both learned to counteract the perturbation similarly during the rotation block. **(e)** During a generalization cycle, the first three trials had unperturbed feedback in the probed context, followed by 15 Point trials with rotated feedback. The first probe trial was to a left or right target, the second and third to the top and bottom targets, and the 15 top-up trials to the trained target. Participants probed with Look feedback showed a larger magnitude of generalization during probe trials and lower hand angles during re-training trials compared to the Inverted Look probe group. **(f)** When considering trials 1-2 from (e), participants in both groups showed larger amounts of generalization to the trained movement of the hand compared to the untrained movement. **(g)** Within-participant comparisons between trained and untrained movement direction showed larger hand angles for the trained movement. For panels (c-e), lines show mean hand angle and shaded regions show 95% confidence intervals. For panels (f-g), bars show mean hand angles and vertical lines show 95% confidence intervals.

Discussion

Sensory predictions play an important role in motor control and motor learning [4,6,19]. Here we examined transfer of visuomotor adaptation across two task contexts with movement-contingent visual feedback that was categorically distinct, but with matched motoric elements. We demonstrated that adaptation proceeds similarly in these two contexts, and that this learning readily transfers between the contexts. Finally, we showed that learning generalizes around the planned movement of the hand, rather than the planned movement of a cursor or the visual workspace direction of a target. We believe these results have meaningful implications for the theory of internal models, specifically regarding the character of sensory predictions and where in the visuomotor hierarchy they are made. Moreover, this first demonstration of the similarity and transfer of visuomotor adaptation between and across mouse Look and Point contexts (along with our work showing similar kinematics [15]) should give other researchers confidence to perform experiments within the very interesting and popular context of FPS-style mouse Look games.

Sensory predictions

Internal models are ubiquitous in modern theories of motor control and learning [1–3]. The sensory predictions generated by forward models can address a number of issues [6], for example distinguishing self and externally-generated sensory input [32], stabilizing state estimates from noisy sensory input [33], overcoming the long time-delays associated with sensory processing for real-time feedback corrections [34], and guiding learning [35]. In some cases the details around the sensory predictions are exceptionally well-defined, for example in work assessing the cancellation of self-generated sensory feedback in the weakly electric fish [5]. In general, however, the output of the forward model is typically stated to be a prediction of the sensory consequences of a movement, without specifying what form that may take.

Based on our current and previous results [15], it seems unlikely that sensory predictions are at the level of low-level visual features. Under such a scenario, both visual contexts would require distinct forward models as the sensory predictions would be necessarily different. While models based on a mixture-of-experts [10,11] could accommodate this, by distributing learning from one context across others, it would require that participants already had the requisite internal models for the Look context, yet only 19 of 68 participants in Experiment 1 reported playing FPS games, and 14 report playing no games at all. While participants had chance to practice the Look movements before the experimental task, tasks thought to elicit *de novo* acquisition of a new internal model typically show extended periods of poor performance

[36], so it seems unlikely participants with little to no experience with FPS games could acquire such internal models and perform the task well so rapidly.

Instead, we suggest sensory predictions are made at a level abstracted from the exact visual feedback expected. While work has emphasized the role of the forward model is to predict the *state* of the body and world [37], this is unsatisfying as state variables for the cursor like its position and velocity never change in *extrinsic* space for the Look context, as it is clamped to the screen center, so only vary relative to other features in the scene. Further, unlike traditional tablet-based setups, there is no one-to-one mapping between hand and cursor position in the Point context when moving an optical mouse or trackpad – only the movement vector between initial and final hand and cursor positions correlate.

Previous work has found that movement vectors play an important role in motor control and planning. Invariant features of goal-directed reaches, like their straightness in extrinsic space and single-peaked tangential velocity, suggest movements are planned in terms of hand trajectories [38–41]. The direction and extent, identifying a movement vector in polar coordinates, appear to be relevant parameters that are specified separately during planning [42,43], with the origin centered on the hand's starting position [44,45]. During adaptation to a visuomotor rotation, the movement vector is the predominant feature remapped [25,46,47], with learning centered on the *planned* movement vector [29,30,48]. By dissociating hand and cursor movements in Experiment 3, we showed that the locus of learning is specifically the planned movement of the hand. We believe the transfer of learning observed in Experiments 2 and 3 suggest that sensory predictions and their comparison with sensory feedback are made at the level of movement vectors between cursor and target in visuospatial coordinates. Under this scheme, the sensory predictions would be consistent between the visual contexts, allowing learning to transfer through shared hand movement vectors.

Contextual cues

Recent work has found effective contextual cues that allow multiple mappings between movements and sensory feedback to be learned concurrently. Participants can learn to counteract opposite signed perturbations, which would normally interfere [49,50], when paired with contextual cues relevant to the movement, for example distinct physical workspace locations [51,52], planned movements [53], or (planned) lead-in or follow-through movements [50,54–56]. Cues related to visual aspects of the movement, like the visual workspace location that feedback is displayed in [51] or where participant gaze is directed [57], also appear effective. The visual contexts studied here may, therefore, serve as effective contextual cues, given visual feedback of the movement progresses differently over time and maintaining fixation on the target would require different gaze patterns between the contexts. In Experiments 2 and 3, we did not train participants to associate different visual contexts with different perturbations concurrently, instead finding that, without a reason to differentiate learning, blocked learning in one context showed near-complete transfer to the other. Additional analyses showed that this was true regardless of whether the participant played FPS games.

The difficulty with using the visual contexts as presented here as effective contextual cues is that, prior to movement onset, the two appear identical. This means additional arbitrary cues would have to be used, like the ring color used for Experiment 3, but such cues are generally ineffective at allowing opposing perturbations to be learned [49,51,58]. However, we do know that a subset of gamers express an Inverted Look mapping in the context of an FPS game, moving their mouse forward to reach a target directly down,

yet switch to typical Point movements when moving their cursor down on the in-game menu or in other non-FPS games. Future research could look at the time scale and cues necessary to allow opposing mappings to be retained long-term and contextually switched in this manner.

Studying learning in FPS games

Video games offer an exciting opportunity to study situated skilled behavior by simply monitoring the inputs, typically from the mouse and keyboard, that are used to interact with the game, allowing relatively passive large-scale data collection. For example, data collected from a commercial FPS aim-trainer game, which uses tasks isolating specific behaviors like aiming, has been used to investigate the long-term effects of practice [59] and the kinematic correlates of skill [60], and actual gameplay data has been analyzed to understand behavioral differences between novice and professional players [61]. Such approaches can allow the development a player's game skill to be tracked [62,63], offering insights about long-term motor learning that are hard or impossible to make with restricted laboratory studies. While previous studies have observed long-term motor learning [62,64], these typically focus on broad measures of improvement, for example the time required to roll cigars [65]. In contrast, the kinematic data collected can be used to decompose broad measures of improvement into a number of spatial and temporal metrics that concurrently improve and correlate well will more traditional Pointing movements [15]. We believe this, combined with our current findings, shows that both visual contexts predominantly tap into the same underlying processes, suggesting that findings from FPS games would generalize to the wider field. Further, movement kinematics inform us about more than just the motor system – they have been used to study the dynamics of a wide range of cognitive processes [66–69]. Given that FPS games typically engage a range of perceptual, cognitive, and motor abilities [70], the kinematic data recorded in actual gameplay could allow for novel insights into the development of skill across a number of domains.

Methods

Participants

In total, 276 participants were recruited to take part in the experiments. After exclusions (n = 2 in Experiment 2: 1 participant recorded no successful movements after the first block, 1 participant's data did not upload; n = 2 in Experiment 3: participant data did not upload), the final sample was 272 participants. Participants were recruited through Prolific.co, an online recruitment platform, and were restricted to those living in the UK or USA, who had English as a first language, and had a Prolific approval rating of 95% or above. Participants were paid between £6.25 - £8 upon completion for expected completion times between 50 – 60 minutes depending on experiment. Experiments were approved by the School of Psychology Ethics Committee at the University of Leeds, and participants gave informed consent via a web form prior to starting the study. Given the experiment was completed online, which has been associated with noisier responses within and between participants on similar paradigms [71], we recruited more participants per group (between 30 – 39) than comparable laboratory studies [50,54,57]. Specifically, we aimed to recruit a typical sample from previous experiments per input device, so continued recruiting until we had at least n=15 per input device per experimental group.

Apparatus and Experimental Procedure

Participants completed experiments using their own personal computer (restricted by Prolific's screening criteria to those using a laptop or desktop) and completed reaching movements in the game using a mouse

or trackpad. In the main text we present data from both mouse and trackpad, as we did not observe any meaningful differences between these input devices. Experiments were created using a recently developed framework allowing within and between-participant comparison of Pointing and Looking movements [15], which uses the Unity game engine (2019.4.15f) and the Unity Experiment Framework [72]. The experiments were delivered via a WebGL build hosted on a web page. No calibration procedure was performed to equate stimuli size across disparate participant monitor dimensions, so the physical size of the task was not constant across participants. The experiments were designed to be visible on a 4:3 monitor aspect ratio, with the height always taking up 4 arbitrary Unity units (au) and wider aspect ratios featuring more of a task-irrelevant background texture. The experiments could only be completed in full-screen mode with the desktop cursor locked and hidden, with raw mouse or trackpad input used to perform in-game movements.

Participants first filled in a form to provide details on age, gender, handedness, computer type, and input device (mouse or trackpad), and also clicked a button to ensure game audio was audible. During this stage, participants used their desktop cursor to navigate the form, and the movements of this cursor were tracked in pixel and Unity game units to provide an initial calibration for the in-game cursor. Participants were shown a cut-scene providing exposition for the purpose of the study (popping non-sentient space bubbles), before completing a tutorial that introduced aspects of the experimental task sequentially and interactively to ensure participants fully understood how to complete the experiment. Following this, participants were required to practice using both contexts used in the experiment (Point and Look for Experiments 1-3, or Point and Inverted Look for one group in Experiment 3), where they could switch contexts and adjust cursor sensitivity with key presses. Once participants had completed at least 20 trials in each context, they could progress to the main task.

Across all experiments, participants used their mouse or trackpad to move an FPS style cursor towards a presented target and shoot it (Figure 1a). Throughout a trial, participants saw a dark space themed background. On top of this, the target plane was shown, which had a diameter of 3au, a black background, and a thin ring around it which was white for Experiments 1 and 2 and either yellow or cyan for Experiment 3. When a trial was ready to start, the start-point (0.1au diameter circle centered on the target plane) was colored orange. Participants needed to move their in-game cursor, a small white circle surrounded by a thin white ring of diameter 0.15au (to mimic FPS style reticles), to the start-point and left-click it to start a trial. If participants were within 0.05au of the start-point during this homing movement, the cursor snapped to the start-point to speed up this centering (mimicking FPS-style auto-aim features). Once the start-point had been clicked, a target (0.2au diameter magenta circle) immediately appeared on the target plane, and the start-point turned green to indicate participants could attempt to move to and left-click the target to shoot it. Potential targets were located in increments of 90° on an imaginary circle of diameter 2au. For Experiments 1 and 2, participants had a time limit within which to shoot the target. If the time limit elapsed without a successful shot being registered, the target immediately vanished and a whooshing sound was played. A shot was only successful if, at the time of the left-click, the center of the cursor intercepted the target circle. If the shot was successful, the target exploded and a shooting sound was played. Unsuccessful shots had no effect upon the game, and an unlimited number of attempted shots were allowed per trial. Feedback of a trial outcome, either a successful shot or a lapsed time limit, was provided for 300ms, during which the start-point was colored grey. Upon feedback finishing, the start-point turned orange, indicating that a new trial could begin. Participants had online feedback of the cursor

throughout a trial, apart from Experiment 3, where certain trials had no feedback when returning to the start-point.

Cursor movements could be performed in two main contexts, either the Point or Look (Figure 1a, with an Inverted Look context added for Experiment 3, Figure 4a). In the Point context, input movements of the mouse or trackpad translated the cursor while the in-game camera remained static, whereas in the Look and Inverted Look contexts, input movements rotated the camera about a fixed point to change the camera's view of the scene while the cursor was fixed central to the camera's view. Critically, the two conditions were equated to ensure identical movements were required to reach a given target in either context. Details of the implementation are described in a previous paper [15]. Following completion of the experimental trials, participants filled in a questionnaire probing their perception of the experiment's performance (e.g. lag, difficulty, technical issues), as well as questions about strategy use to overcome the imposed perturbation.

Experiment 1

Participants (n = 68; mean age = 36, SD age = 11, age range = 20 – 65; 24 male, 42 female) were assigned to complete the experiment in either the Point (n = 32) or Look (n = 36) context. During the main task, participants completed 680 experimental trials. Participants first completed a baseline period of 100 experimental trials where no perturbation was applied to cursor movements. Following this, participants completed 480 experimental trials with a 30° visuomotor rotation, where the cursor movement was rotated away from the actual input movement in either the clockwise or counterclockwise direction (counterbalanced across participants within each group). Finally, participants completed 100 experimental trials where the perturbation was turned off to assess after-effects of learning from the rotation period. Throughout the experiment, trials were tested in cycles of four trials, where each target (0°, 90°, 180° or 270°) was tested in a random order. Participants were given breaks of unrestricted length every 20 cycles through the experiment.

To induce the type of time pressure that might be experienced in a typical FPS-style game, the time limit within which participants had to show a target was continuously staircased throughout the experiment. A pair of staircases, initiated at 450ms and 1050ms for mouse users or 780ms and 1380ms for trackpad users (determined through previous work [15]), were interleaved and used equally within each block to set the time limit for the upcoming trial. Following a 1-up 1-down procedure, the tested staircase's time limit was decreased by 30ms following a successful trial, and increased by 30ms following an unsuccessful trial, to approximately give a 50% success rate over the course of the experiment.

Experiment 2

Participants were assigned to either Experiment 2a (n = 62; mean age = 37, SD age = 22, age range = 18-67; 24 male, 38 female) or Experiment 2b (n = 66; mean age = 36, SD age = 11, age range = 20-75; 20 male, 46 female), and were further assigned to complete the rotation block in either the Point (Experiment 2a, n = 32, Experiment 2b, n = 35) or Look (Experiment 2a, n = 30, Experiment 2b, n = 31) context. The experiment consisted of 800 trials. Participants first completed 100 baseline trials in the 'untrained' context (the Look context for the Point Trained group), and then 100 baseline trials in the 'trained' context (the Point context for the Point trained group). During these trials, no perturbation was applied to the cursor. Following this, participants completed 480 trials in the trained context with a 30° visuomotor rotation (direction counterbalanced within each group). For Experiment 2a the perturbation was then

turned off, and participants completed 100 trials in the untrained context and a further 20 trials in the trained context, to assess the transfer of learning from the trained to untrained context. Participants in Experiment 2b completed the same pair of blocks after the rotation, but the perturbation was left on during these blocks, as a second method of assessing transfer. As in experiment 1, trials were tested in cycles of the four target directions, and the time limit was continuously staircased throughout to induce time pressure. Participants were given untimed breaks after cycles 23 and 45, and thereafter every 20 cycles.

In Experiment 1, it was observed that participants sometimes initiated movements to an uncued target (hereafter 'jump starts'). To combat this, a further trial outcome was added. If a participant's movement at a radius of 0.2au was more than 60° in either direction from the ideal movement path between the start-point and target, the trial immediately stopped, and participants were shown an error message for 5 seconds instructing them to make sure they waited to see the target before moving. This check was performed early enough that participants only received a limited sample of their movement on these failed trials. The trial was then repeated, to ensure that participants all received the same number of trials to assess learning and transfer.

Experiment 3

Participants (n = 76; mean age = 36, SD age = 13, age range = 19-67; 26 male, 49 female) were assigned to use either the Look (n = 40) or Inverted Look (n = 36) context during probe trials. In total, participants completed 760 trials during this experiment. In contrast to the previous experiments, no time limit was used, as pilot testing found this caused a substantial number of errors on the critical probe trials. Further, the ring around the target plane changed color as participants returned their cursor to the start-point to indicate the context of the upcoming trial.

Participants first completed 60 baseline trials to the top and bottom targets (90° and 270°). The first 40 trials had 20 contiguous trials for each tested context (Point and either Look or Inverted Look, in a random order), followed by 20 trials where the two contexts were interleaved to get participants used to switching between contexts. Participants then completed 20 trials in the Point context to the top target, where the no-feedback return was introduced. Here, upon clicking the target, the cursor disappeared, and participants were required to return their hand to a comfortable position within 1.5 seconds of outcome feedback extinguishing. Tones were played after 500ms, 1000ms and 1500ms to ensure participants knew when their hand had to be back to a comfortable position. After the final tone, the cursor appeared in the middle of the start-point. From this trial onwards, Point trials never had return feedback (whereas Look and Inverted Look trials always had return feedback) and were always directed to the 90° (top) target.

Participants then completed the baseline generalization block, an unperturbed version of the generalization block. The baseline generalization block consisted of 15 generalization cycles of 18 trials. The first three trials were in the probed context, either Look or Inverted Look, and probed participants movements to either the left or right target (0° or 180°) in the first trial and both the top and bottom targets in trials two and three (arranged so the top target was tested first either 7 or 8 times, in a random order across cycles). The first movement to a lateral target served as an additional cue, as well as the ring color, that the context had switched before the critical vertical targets were probed. The last 15 trials per cycle were in the Point context to the top target. During the baseline generalization block, movements were unperturbed compared to their nominal mapping. Participants then performed 120 trials in the Point

context with a 30° visuomotor rotation applied (rotation direction counterbalanced within each group). Because there was no return feedback, participants only sampled learning during movements to the top target, so should show only local generalization. Participants then completed the generalization block, which had the same structure as the baseline generalization block except the 15 Point trials per cycle had the same visuomotor rotation as in the rotation block (while the probe trials were never perturbed). Because the probe trials had return feedback and probed both trained and untrained directions, unlearning could occur during outward and return reaches, so these 15 trials 'topped up' learning to ensure the probes were made after learning again reached asymptotic levels, as determined during pilot testing. While generalization studies typically preclude feedback to avoid this issue, the feedback during Look trials arose because the camera panned and tilted, so clamping the camera's rotation during Look or Inverted Look probes would give no indication a Look style movement had theoretically occurred. Participants finally completed 20 Point movements with unperturbed feedback, to assess after-effects. Participants were given breaks of unrestricted length during the 8th probe cycle in both the baseline generalization and generalization blocks (after the probe trials but before the top-up trials), and after 100 trials in the rotation block.

The same online check for jump-starts used in Experiments 2a and 2b was used here for trials during the first baseline block for both contexts, and thereafter only for the Point context. As we expected some level of 'slips' for the Inverted Look group (i.e. expressions of the non-inverted mapping), we did not check for jump-starts during the probe trials.

Data Analysis

Kinematic data (cursor position in x and y axis) was sampled at the participant's screen refresh rate (often approximately 60Hz, 120Hz, or 144Hz) and uploaded to a remote database during the experiments. Offline processing was performed in R (version 4.2.2). Technical issues meant a small number of trials were not uploaded, which removed 0.02% of trials for Experiment 1, 0.05% of trials in Experiment 2, and 0.02% of trials in Experiment 3. Further, a small number of trials were affected by an issue that caused a large single-frame mouse input, causing participant's cursor to move a large distance suddenly and unexpectedly. This appears to be an issue with WebGL games more broadly, so we could only filter out affected trials post-hoc by removing trials where cursor speed changed by more than 40au/s from a previous or following sample. This removed 0.4% of trials in Experiment 1, 0.04% of trials in Experiment 2, and 0.3% of trials in Experiment 3.

To ensure a common analysis procedure was applied for all participants, movement data was resampled to a consistent 100Hz using linear interpolation and filtered using a second order, zero-phase, low-pass Butterworth filter with a 15Hz cut-off, with the start and end of each trial's time series padded to avoid transient effects of the filter affecting the movement data. Raw hand paths were visualized (e.g. Figure 2b, rotated to a common target and flipped to show a consistent rotation direction) to ensure the actual movements performed by participants were consistent with canonical observations for these task manipulations. To perform quantitative analysis, key measures were extracted from the time series of cursor movements. The main measure used throughout was hand angle at peak speed. To extract this, the peak radial speed reached during the trial was found and the input position at this time was identified (unaffected by perturbations). Hand angle was then defined as the difference in angle between straight lines connecting the start-point to either the target or input position at peak speed. For participants who experienced a counterclockwise perturbation, the sign of the hand angle was flipped so for both

perturbation directions, positive hand angles reflected a movement of the hand in the direction opposite to the perturbation. For Experiment 3, hand angles on Inverted Look trials also had an 180° rotation applied so that they reflected the angle from the ideal aim location to make Look and Inverted Look trials directly comparable. Trials in Experiments 1 and 2 were grouped into cycles of four movements (one to each target), with mean hand angles calculated per cycle, whereas for Experiment 3 hand angles were left per-trial.

In Experiment 1 we noticed participants occasionally initiated movements to an uncued target, likely because the time pressure applied induced some guessing. For example, at extremely low reaction times in forced response paradigms, participants guess uniformly which target will be probed [73]. To alleviate the effect of this on our data, we removed trials where the absolute hand angle at both peak speed and take-off (hand angle at a radius of 0.2au) were more than 60° from the ideal angle, to ensure the included movements were not directed to an uncued target, which removed 10.5% of trials for the first experiment (*inclusion of these trials did not change any statistical analyses*). To counteract this, Experiments 2 and 3 had online detection of 'jump starts' to ensure any such trials were repeated. This led to 5.3% of trials being repeated in Experiment 2, and 1.6% of trials in Experiment 3 (*note the reduction in repeated trials when the time limit was removed for Experiment 3, with a rate comparable to other online adaptation experiments* [74,75]). Despite this, some trials did not pass the filter at peak speed, removing 1.2% of trials in Experiment 2 (*including these outliers produced identical analyses besides the transfer analysis for Experiment 2b where, despite similar point estimates, there was no statistical difference in transfer between contexts, nor a reduction in learning for the Point-trained group*), and 1.6% of trials in Experiment 3 (*inclusion of these trials did not change any statistical analyses*).

To perform analyses based on periods of the experiment, average hand angles over pre-specified regions were calculated. For Experiment 1, where we wanted to understand how hand angle compared between the contexts across the experiment, the baseline period comprised the last 10 cycles of the baseline block, and early and late rotation and washout were defined as the first and last 10 cycles of the rotation and washout blocks respectively. We also calculated a measure for the after-effect that accounted for potential differences in asymptotic learning by subtracting the hand angle during late learning from that during early washout. For Experiment 2, where we were primarily concerned with transfer, we calculated early and late rotation as in Experiment 1, and additionally calculated transfer as late learning subtracted from the first 10 cycles of the transfer block. For Experiment 3, early and late learning was assessed during the first and last 10 trials during the rotation block respectively. Analysis of probe trials used baseline-corrected values, where the average hand angle per target in the baseline generalization block was subtracted from average hand angle per target in the generalization block for each participant, to correct for any intrinsic biases (analysis was identical with or without baseline correction).

Statistical Analysis

All statistical analyses were performed in R (version 4.2.2). For all ANOVAs assessing the progression of learning in Experiments 1-3 (performed using the *afex*, *BayesFactor*, and *bayestestR* packages), in addition to reporting the F-statistic and *p-values*, we included Bayes Factors (BF_{10}) as a measure of the evidence for the alternative hypothesis (H_1) over the null hypothesis (H_0) to supplement the p-values, and partial eta squared (η_p^2) as a measure of effect size. For all follow-up t-tests using the estimated marginal means (performed using the *emmeans* package) and all independent t-tests, we use two-tailed tests and report the mean and 95% confidence interval of the differences as well as *p-value* and Bayes Factor. Finally, for

all mixed-effect linear regressions (featuring a random intercept for participant, performed using the *lme4* package with p-values from the *lmerTest* package calculated according to the Satterthwaite approximation), we report the mean and 95% confidence interval of the regression coefficient. All statistical tests were assessed against a significance threshold of $p < .05$.

Data availability

All data and scripts to analyse the data will be made available on the Open Science Framework upon publication of the manuscript.

Acknowledgements

Authors F.M and M.M-W were supported by Fellowships from the Alan Turing Institute.

Author contributions

Matthew Warburton: Conceptualization, Methodology, Software, Formal analysis, Investigation, Visualization, Investigation, Writing – Original draft, Writing – Review & Editing. **Carlo Campagnoli**: Methodology, Writing – Review & Editing, Supervision. **Mark Mon-Williams**: Writing – Review & Editing, Funding acquisition. **Faisal Mushtaq**: Writing – Review & Editing, Funding acquisition, Supervision. **J. Ryan Morehead**: Conceptualization, Methodology, Writing - Review & Editing, Supervision, Funding acquisition.

References

1. McNamee, D. & Wolpert, D. M. Internal Models in Biological Control. *Annu. Rev. Control Robot. Auton. Syst.* **2**, 339–364 (2019).
2. Shadmehr, R. & Krakauer, J. W. A computational neuroanatomy for motor control. *Exp. Brain Res.* **185**, 359–381 (2008).
3. Franklin, D. W. & Wolpert, D. M. Computational Mechanisms of Sensorimotor Control. *Neuron* **72**, 425–442 (2011).
4. Blakemore, S.-J., Wolpert, D. & Frith, C. Why can't you tickle yourself? *NeuroReport* **11**, R11–R16 (2000).
5. Sawtell, N. B. Neural Mechanisms for Predicting the Sensory Consequences of Behavior: Insights from Electrosensory Systems. *Annu. Rev. Physiol.* **79**, 381–399 (2017).
6. Miall, R. C. & Wolpert, D. M. Forward Models for Physiological Motor Control. *Neural Netw.* **9**, 1265–1279 (1996).
7. Mehta, B. & Schaal, S. Forward Models in Visuomotor Control. *J. Neurophysiol.* **88**, 942–953 (2002).
8. Shadmehr, R., Smith, M. A. & Krakauer, J. W. Error Correction, Sensory Prediction, and Adaptation in Motor Control. *Annu. Rev. Neurosci.* **33**, 89–108 (2010).
9. Desmurget, M. & Grafton, S. Forward modeling allows feedback control for fast reaching movements. *Trends Cogn. Sci.* **4**, 423–431 (2000).
10. Wolpert, D. M. & Kawato, M. Multiple paired forward and inverse models for motor control. *Neural Netw.* **11**, 1317–1329 (1998).
11. Haruno, M., Wolpert, D. M. & Kawato, M. MOSAIC Model for Sensorimotor Learning and Control. *Neural Comput.* **13**, 2201–2220 (2001).
12. Yon, D., Gilbert, S. J., De Lange, F. P. & Press, C. Action sharpens sensory representations of expected outcomes. *Nat. Commun.* **9**, 4288 (2018).
13. Maunsell, J. H. & Van Essen, D. C. Functional properties of neurons in middle temporal visual area of the macaque monkey. I. Selectivity for stimulus direction, speed, and orientation. *J. Neurophysiol.* **49**, 1127–1147 (1983).
14. Cohen, Y. E. & Andersen, R. A. A common reference frame for movement plans in the posterior parietal cortex. *Nat. Rev. Neurosci.* **3**, 553–562 (2002).
15. Warburton, M., Campagnoli, C., Mon-Williams, M., Mushtaq, F. & Morehead, J. R. Kinematic markers of skill in first-person shooter video games. *PNAS Nexus* **2**, pgad249 (2023).
16. Poggio, T. & Bizzi, E. Generalization in vision and motor control. *Nature* **431**, 768–774 (2004).

17. Ghahramani, Z., Wolpert, D. M. & Jordan, M. I. Generalization to Local Remappings of the Visuomotor Coordinate Transformation. *J. Neurosci.* **16**, 7085–7096 (1996).
18. Thoroughman, K. A. & Shadmehr, R. Learning of action through adaptive combination of motor primitives. *Nature* **407**, 742–747 (2000).
19. Mazzoni, P. An Implicit Plan Overrides an Explicit Strategy during Visuomotor Adaptation. *J. Neurosci.* **26**, 3642–3645 (2006).
20. Tseng, Y., Diedrichsen, J., Krakauer, J. W., Shadmehr, R. & Bastian, A. J. Sensory Prediction Errors Drive Cerebellum-Dependent Adaptation of Reaching. *J. Neurophysiol.* **98**, 54–62 (2007).
21. Krakauer, J. W. & Mazzoni, P. Human sensorimotor learning: adaptation, skill, and beyond. *Curr. Opin. Neurobiol.* **21**, 636–644 (2011).
22. Hadjiosif, A. M., Krakauer, J. W. & Haith, A. M. Did We Get Sensorimotor Adaptation Wrong? Implicit Adaptation as Direct Policy Updating Rather than Forward-Model-Based Learning. *J. Neurosci.* **41**, 2747–2761 (2021).
23. Raibert, M. H. A model for sensorimotor control and learning. *Biol. Cybern.* **29**, 29–36 (1978).
24. Cunningham, H. A. Aiming error under transformed spatial mappings suggests a structure for visual-motor maps. *J. Exp. Psychol. Hum. Percept. Perform.* **15**, 493 (1989).
25. Krakauer, J. W., Pine, Z. M., Ghilardi, M.-F. & Ghez, C. Learning of Visuomotor Transformations for Vectorial Planning of Reaching Trajectories. *J. Neurosci.* **20**, 8916–8924 (2000).
26. Krakauer, J. W. Motor Learning and Consolidation: The Case of Visuomotor Rotation. in *Progress in Motor Control* (ed. Sternad, D.) vol. 629 405–421 (Springer US, 2009).
27. Morehead, J. R. & de Xivry, J.-J. O. A synthesis of the many errors and learning processes of visuomotor adaptation. *bioRxiv* (2021) doi:10.1101/2021.03.14.435278.
28. Krakauer, J. W., Hadjiosif, A. M., Xu, J., Wong, A. L. & Haith, A. M. Motor Learning. in *Comprehensive Physiology* (ed. Terjung, R.) 613–663 (Wiley, 2019). doi:10.1002/cphy.c170043.
29. Day, K. A., Roemmich, R. T., Taylor, J. A. & Bastian, A. J. Visuomotor Learning Generalizes Around the Intended Movement. *eneuro* **3**, ENEURO.0005-16.2016 (2016).
30. McDougle, S. D., Bond, K. M. & Taylor, J. A. Implications of plan-based generalization in sensorimotor adaptation. *J. Neurophysiol.* **118**, 383–393 (2017).
31. Ghilardi, M. F., Gordon, J. & Ghez, C. Learning a visuomotor transformation in a local area of work space produces directional biases in other areas. *J. Neurophysiol.* **73**, 2535–2539 (1995).
32. Cullen, K. E. Sensory signals during active versus passive movement. *Curr. Opin. Neurobiol.* **14**, 698–706 (2004).
33. Wolpert, D. M., Ghahramani, Z. & Jordan, M. I. An Internal Model for Sensorimotor Integration. *Science* **269**, 1880–1882 (1995).
34. Miall, R. C., Weir, D. J., Wolpert, D. M. & Stein, J. F. Is the Cerebellum a Smith Predictor? *J. Mot. Behav.* **25**, 203–216 (1993).
35. Jordan, M. I. & Rumelhart, D. E. Forward models: Supervised learning with a distal teacher. *Cogn. Sci.* **16**, 307–354 (1992).
36. Haith, A. M., Yang, C. S., Pakpoor, J. & Kita, K. De novo motor learning of a bimanual control task over multiple days of practice. *J. Neurophysiol.* **128**, 982–993 (2022).
37. Wolpert, D. M. & Flanagan, J. R. Motor prediction. *Curr. Biol.* **11**, R729–R732 (2001).
38. Morasso, P. Spatial control of arm movements. *Exp. Brain Res.* **42**, (1981).
39. Abend, W., Bizzi, E. & Morasso, P. Human arm trajectory formation. *Brain J. Neurol.* **105**, 331–348 (1982).
40. Georgopoulos, A. P., Kalaska, J. F. & Massey, J. T. Spatial trajectories and reaction times of aimed movements: effects of practice, uncertainty, and change in target location. *J. Neurophysiol.* **46**, 725–743 (1981).
41. Flash, T. & Hogan, N. The coordination of arm movements: an experimentally confirmed mathematical model. *J. Neurosci.* **5**, 1688–1703 (1985).
42. Rosenbaum, D. A. Human movement initiation: specification of arm, direction, and extent. *J. Exp. Psychol. Gen.* **109**, 444 (1980).
43. Favilla, M., Hening, W. & Ghez, C. Trajectory control in targeted force impulses: VI. Independent specification of response amplitude and direction. *Exp. Brain Res.* **75**, (1989).
44. Gordon, J., Ghilardi, M. F. & Ghez, C. Accuracy of planar reaching movements: I. Independence of direction and extent variability. *Exp. Brain Res.* **99**, 97–111 (1994).
45. Vindras, P. & Viviani, P. Frames of Reference and Control Parameters in Visuomanual Pointing. *J. Exp. Psychol. Hum. Percept. Perform.* **24**, 569–591 (1998).
46. Wang, J. & Sainburg, R. L. Adaptation to Visuomotor Rotations Remaps Movement Vectors, Not Final Positions. *J. Neurosci.* **25**, 4024–4030 (2005).

47. Wu, H. G. & Smith, M. A. The Generalization of Visuomotor Learning to Untrained Movements and Movement Sequences Based on Movement Vector and Goal Location Remapping. *J. Neurosci.* **33**, 10772–10789 (2013).
48. Kim, O. A., Forrence, A. D. & McDougle, S. D. Motor learning without movement. *Proc. Natl. Acad. Sci.* **119**, e2204379119 (2022).
49. Gandolfo, F., Mussa-Ivaldi, F. A. & Bizzi, E. Motor learning by field approximation. *Proc. Natl. Acad. Sci.* **93**, 3843–3846 (1996).
50. Sheahan, H. R., Franklin, D. W. & Wolpert, D. M. Motor Planning, Not Execution, Separates Motor Memories. *Neuron* **92**, 773–779 (2016).
51. Howard, I. S., Wolpert, D. M. & Franklin, D. W. The effect of contextual cues on the encoding of motor memories. *J. Neurophysiol.* **109**, 2632–2644 (2013).
52. Hwang, E. J., Smith, M. A. & Shadmehr, R. Dissociable effects of the implicit and explicit memory systems on learning control of reaching. *Exp. Brain Res.* **173**, 425–437 (2006).
53. Hirashima, M. & Nozaki, D. Distinct Motor Plans Form and Retrieve Distinct Motor Memories for Physically Identical Movements. *Curr. Biol.* **22**, 432–436 (2012).
54. Dawidowicz, G., Shaine, Y. & Mawase, F. Separation of multiple motor memories through implicit and explicit processes. *J. Neurophysiol.* **127**, 329–340 (2022).
55. Howard, I. S., Ingram, J. N., Franklin, D. W. & Wolpert, D. M. Gone in 0.6 Seconds: The Encoding of Motor Memories Depends on Recent Sensorimotor States. *J. Neurosci.* **32**, 12756–12768 (2012).
56. Howard, I. S., Wolpert, D. M. & Franklin, D. W. The Value of the Follow-Through Derives from Motor Learning Depending on Future Actions. *Curr. Biol.* **25**, 397–401 (2015).
57. Abekawa, N., Ito, S. & Gomi, H. Gaze-specific motor memories for hand-reaching. *Curr. Biol.* S0960982222007011 (2022) doi:10.1016/j.cub.2022.04.065.
58. Forano, M., Schween, R., Taylor, J. A., Hegele, M. & Franklin, D. W. Direct and indirect cues can enable dual adaptation, but through different learning processes. *J. Neurophysiol.* **126**, 1490–1506 (2021).
59. Listman, J. B., Tsay, J. S., Kim, H. E., Mackey, W. E. & Heeger, D. J. Long-Term Motor Learning in the "Wild" With High Volume Video Game Data. *Front. Hum. Neurosci.* **15**, 777779 (2021).
60. Donovan, I. *et al.* Assessment of human expertise and movement kinematics in first-person shooter games. *Front. Hum. Neurosci.* **16**, (2022).
61. Park, E. *et al.* Secrets of Gosu: Understanding Physical Combat Skills of Professional Players in First-Person Shooters. in *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* 1–14 (ACM, 2021). doi:10.1145/3411764.3445217.
62. Stafford, T. & Dewar, M. Tracing the Trajectory of Skill Learning With a Very Large Sample of Online Game Players. *Psychol. Sci.* **25**, 511–518 (2014).
63. Stafford, T. & Vaci, N. Maximizing the potential of digital games for understanding skill acquisition. *Curr. Dir. Psychol. Sci.* **31**, 49–55 (2022).
64. Baddeley, A. D. & Longman, D. J. A. The Influence of Length and Frequency of Training Session on the Rate of Learning to Type. *Ergonomics* **21**, 627–635 (1978).
65. Crossman, E. R. F. W. A theory of the acquisition of speed-skill. *Ergonomics* **2**, 153–166 (1959).
66. Song, J.-H. & Nakayama, K. Hidden cognitive states revealed in choice reaching tasks. *Trends Cogn. Sci.* **13**, 360–366 (2009).
67. Freeman, J. B., Dale, R. & Farmer, T. A. Hand in Motion Reveals Mind in Motion. *Front. Psychol.* **2**, (2011).
68. Spivey, M. J. & Dale, R. Continuous Dynamics in Real-Time Cognition. *Curr. Dir. Psychol. Sci.* **15**, 207–211 (2006).
69. Dotan, D., Pinheiro-Chagas, P., Al Roumi, F. & Dehaene, S. Track It to Crack It: Dissecting Processing Stages with Finger Tracking. *Trends Cogn. Sci.* **23**, 1058–1070 (2019).
70. Green, C. S. & Bavelier, D. Learning, Attentional Control, and Action Video Games. *Curr. Biol.* **22**, R197–R206 (2012).
71. 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* 1–22 (2021) doi:10.51628/001c.26985.
72. Brookes, J., Warburton, M., Alghadier, M., Mon-Williams, M. & Mushtaq, F. Studying human behavior with virtual reality: The Unity Experiment Framework. *Behav. Res. Methods* **52**, 455–463 (2019).
73. Haith, A. M., Pakpoor, J. & Krakauer, J. W. Independence of Movement Preparation and Movement Initiation. *J. Neurosci.* **36**, 3007–3015 (2016).
74. Avraham, G., Morehead, J. R., Kim, H. E. & Ivry, R. B. Reexposure to a sensorimotor perturbation produces opposite effects on explicit and implicit learning processes. *PLOS Biol.* **19**, e3001147 (2021).
75. Tsay, J. S., Irving, C. & Ivry, R. B. Signatures of contextual interference in implicit sensorimotor adaptation. *Proc. R. Soc. B Biol. Sci.* **290**, 1–8 (2023).