Interactions between sensory prediction error and task error during implicit motor learning

Short title: Interactions between errors during implicit motor learning

Jonathan S. Tsay[1,2], Adrian M. Haith[3], Richard B. Ivry[1,2], Hyosub E. Kim[4, 5]

[1]*Department of Psychology, University of California, Berkeley*
[2]*Helen Wills Neuroscience Institute, University of California, Berkeley*
[3]*Department of Neurology, Johns Hopkins University*
[4]*Department of Physical Therapy, University of Delaware, Newark*
[5]*Department of Psychological and Brain Sciences, University of Delaware*

Corresponding author Information:

Name: Jonathan Tsay
Email: xiaotsay2015@berkeley.edu
Address: 2121 Berkeley Way, Berkeley, CA 94704

Key words: Motor learning, Visuomotor Adaptation, Error Based Learning, Target Jump

Abstract

Implicit motor recalibration allows us to flexibly move in novel and changing environments. Conventionally, implicit recalibration is thought to be driven by errors in predicting the sensory outcome of movement (i.e., sensory prediction errors). However, recent studies have shown that implicit recalibration is also influenced by errors in achieving the movement goal (i.e., task errors). Exactly how sensory prediction errors and task errors interact to drive implicit recalibration and, in particular, whether task errors alone might be sufficient to drive implicit recalibration remain unknown. To test this, we induced task errors in the absence of sensory prediction errors by displacing the target mid-movement. We found that task errors alone failed to induce implicit recalibration. In additional experiments, we simultaneously varied the size of sensory prediction errors and task errors. We found that implicit recalibration driven by sensory prediction errors could be continuously modulated by task errors, revealing an unappreciated dependency between these two sources of error. Moreover, implicit recalibration was attenuated when the target was simply flickered in its original location, even though this manipulation did not affect task error – an effect likely attributed to attention being directed away from the feedback cursor. Taken as a whole, the results were accounted for by a computational model in which sensory prediction errors and task errors, modulated by attention, interact to determine the extent of implicit recalibration.

Author's summary

What information does the brain use to maintain precise calibration of the sensorimotor system? Using a reaching task paired with computational modeling, we find that movements are implicitly recalibrated by errors in predicting both the sensory outcome of movement (i.e., sensory prediction errors) as well as errors in achieving the movement goal (i.e., task errors). Even though task errors alone do not elicit implicit recalibration, they nonetheless modulate implicit recalibration when sensory prediction error is present. The results elucidate an unappreciated interaction between these two sources of error in driving implicit recalibration.

Introduction

Sensorimotor adaptation is an essential feature of human competence, allowing us to flexibly move in novel and changing environments (Kim, Avraham, & Ivry, 2020; J. Krakauer, Hadjiosif, Xu, Wong, & Haith, 2019; Ryan Morehead & de Xivry, 2021; Shadmehr, Smith, & Krakauer, 2010). Multiple learning processes have been shown to contribute to the performance changes observed in adaptation tasks, including an aiming process which is explicit, volitional, and learns rapidly and a recalibration process which is implicit, automatic, and learns slowly (Haith, Huberdeau, & Krakauer, 2015; Hegele & Heuer, 2010; McDougle, Ivry, & Taylor, 2016; Taylor & Ivry, 2011; Taylor, Krakauer, & Ivry, 2014; Werner et al., 2015). Recent work has focused on how these two learning processes may be driven by distinct error signals: Whereas explicit aiming responds to task error (TE) – a signal reflecting task performance (Day, Roemmich, Taylor, & Bastian, 2016; Taylor & Ivry, 2011) – implicit recalibration responds to sensory prediction error (SPE) – an error reflecting the difference between predicted and actual feedback (Donchin, Francis, & Shadmehr, 2003; Kim, Morehead, Parvin, Moazzezi, & Ivry, 2018; Lee, Oh, Izawa, & Schweighofer, 2018; Mazzoni & Krakauer, 2006; Morehead, Taylor, Parvin, & Ivry, 2017; Shadmehr et al., 2010; Wolpert, Miall, & Kawato, 1998). Moreover, these two learning processes are thought to rely on distinct neural modules, with explicit aiming requiring more prefrontal control (Anguera, Reuter-Lorenz, Willingham, & Seidler, 2010; Benson, Anguera, & Seidler, 2011; Taylor & Ivry, 2014) and implicit recalibration requiring more cerebellar control (Butcher et al., 2017; Hadjiosif et al., 2014; Izawa, Criscimagna-Hemminger, & Shadmehr, 2012; Schlerf, Xu, Klemfuss, Griffiths, & Ivry, 2013; Taylor, Klemfuss, & Ivry, 2010; Tseng, Diedrichsen, Krakauer, Shadmehr, & Bastian, 2007).

However, recent results from visuomotor rotation tasks have motivated a broader perspective of implicit recalibration, and in particular, led to the proposal that implicit recalibration is sensitive not only to sensory prediction error, but also to task outcome. Empirically, the evidence supporting this hypothesis comes from studies in which perturbed visual feedback (the source of SPE) is combined with a manipulation of target size or target jumps (Cameron, Franks, Inglis, & Chua, 2010a, 2010b; Magescas & Prablanc, 2006) to create a condition in which the visual feedback "hits" the target (Figure 1). Adaptation in such situations is attenuated by about ~20% compared to that observed in control conditions with a similar SPE (Kim, Parvin, & Ivry, 2019; Leow, Marinovic, de Rugy, & Carroll, 2018). The hypothesis that implicit recalibration is sensitive to both SPE and task outcome is consistent with recent neurophysiological observations of reward-related activity in the cerebellum (Heffley & Hull, 2019; Hull, 2020; Ohmae & Medina, 2015; Sendhilnathan, Ipata, & Goldberg, 2020; Wagner, Kim, Savall, Schnitzer, & Luo, 2017).

But how exactly are SPE and task outcome combined to drive implicit recalibration? One possibility is that behavior reflects the operation of two independent learning processes, one sensitive to SPE and the other sensitive to task outcome (Kim et al., 2019; Leow et al., 2018). While this dual-error model is consistent with existing findings, it is unknown whether this reflects the operation of two learning processes that operate independently. For example, it remains to be seen if TE-only would be sufficient to drive adaptation, as would be predicted by such a dual-error model.

Alternatively, SPE and task outcome may interact. For example, the strength of the SPE might be modulated by task outcome; if the displaced cursor still manages to intersect the target, a reward signal linked with task success could weaken the system's sensitivity to SPE, reducing the rate of recalibration (Gonzalez Castro, Monsen, & Smith, 2011; Shmuelof et al., 2012). A different form of interaction might arise from processes tangential to recalibration. For example, displacement of the target, as is commonly used to manipulate TE, might capture attention and weaken the salience of the SPE. In principle, the interaction between TE and SPE could also be a combination of multiple effects.

To examine how SPE and TE collectively shape implicit recalibration, we performed a series of visuomotor experiments that systematically varied the size of these two errors. We also compared participants' performance to a series of computational models designed to catalogue potential ways in which SPE and TE may interact. To control the size of SPE (i.e., operationalized as the difference between the cursor feedback and the original target location), we used clamped visual feedback (Morehead et al., 2017), in which the timing and extent of cursor motion is linked to hand motion, but the cursor trajectory is offset by a fixed angle relative to the target, and thus independent of the hand trajectory. To control the size of TE (i.e., operationalized as the difference between the cursor feedback and the new target location), we jumped the target by a variable amount soon after movement initiation. In all cases, these manipulations were coupled with instructions to ignore the visual feedback and always reach straight towards the original target – an approach which has been shown to reliably elicit implicit recalibration without contamination from explicit strategies (Leow et al., 2018; Tsay, Parvin, & Ivry, 2020). These experiments, coupled with computational models, allow us to precisely characterize the effects of SPE and TE on implicit recalibration.

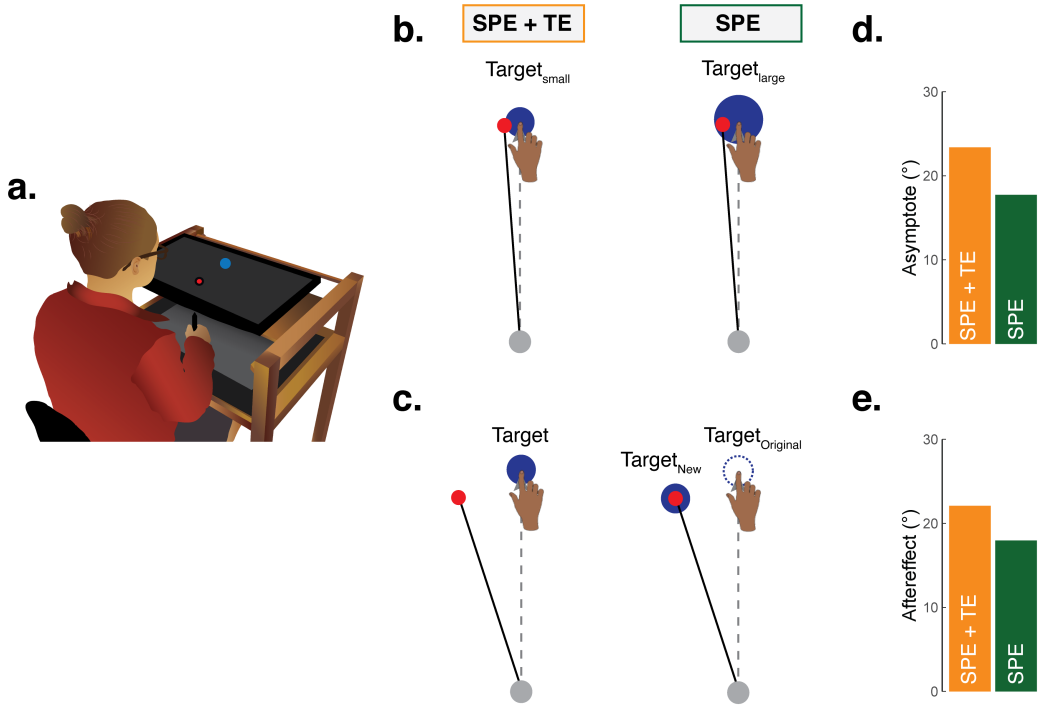

Figure 1: Implicit recalibration elicited by SPE + TE and SPE-only. (a) Illustration of experimental apparatus. **(b-c)** Task outcome was manipulated by either varying the size of the target (Kim et al, 2019) or varying the size of the target jump (Leow et al, 2018). Both SPE and TE are present when the cursor feedback straddles or misses the target, and only SPE is present when the cursor "hits" the target. **(d-e)** Implicit recalibration, as measured by the asymptote of hand angle in a clamped feedback design in Kim et al 2019 or during no-feedback aftereffect trials in a standard visuomotor rotation design, was reduced when TE was removed.

98

Results

Experiment 1-2: TE alone is not sufficient to drive implicit adaptation.

We first examined whether TE-only perturbations would elicit implicit recalibration (Figure 2). We induced TEs by jumping the target a varying amount, between ±16° from trial to trial, while pairing all of these conditions with a clamped cursor that always moved through the original target (i.e., 0° clamp). If TE alone is sufficient to elicit implicit recalibration, the participant's movement would be expected to shift in the direction of the jumped target on the subsequent trial. As a point of comparison, we also tested a condition in which both SPE and TE varied together, through separate blocks in which cursor feedback was clamped between ±16° while the target remained stationary (SPE+TE). We expected the participant's movement would be shifted in the opposite direction of the cursor for these conditions. To ensure that learning was implicit, participants were instructed to always move directly to the original target location, ignoring both the cursor feedback and the target jump.

In trials when both SPE and TE were present, all participants exhibited robust changes in hand angle to (partially) counter the imposed error, a key signature of implicit recalibration (Figure 2c; Mean slope ± SEM: $\beta = -0.1 \pm 0.0$; $F_{(1,212)} = 136.0$, $p = 1.3 \times 10^{-24}, \eta^2 = 0.2$). The change in hand angle as a function of error size appeared to be sublinear, composed of a linear zone for smaller perturbations (0° – 4°) and a saturated region for larger perturbations (4° – 16°), consistent with previous reports of saturated learning across a wide range of error sizes (Hayashi, Kato, & Nozaki, 2020; Kasuga, Hirashima, & Nozaki, 2013; Kim et al., 2018; Morehead et al., 2017; Wei & Körding, 2009).

A very different picture was observed in the TE-only blocks. Here participants exhibited no reliable change in hand angle in response to the TE (Figure 2d; $\beta = 0.0 \pm 0.0$; $t_{(212)} = 0.6, p = 0.69, D = 0.1$). Critically, there was a striking interaction between perturbation size and perturbation type ($\beta = 1.2 \pm 0.1$; $F_{(1,212)} = 61.1, p = 2.5 \times 10^{-13}, \eta^2 = 0.2$), where robust implicit recalibration was observed when both SPE + TE were present, but not when TE-only was provided.

As a test of generality, we examined whether TE-only perturbations would elicit implicit recalibration when all perturbation conditions were scheduled in a random manner, rather than providing TE-only and SPE+TE trials in separate blocks (see

Table 3 in the Methods section). Again, robust sign-dependent changes in hand angle were observed for all participants in the SPE + TE condition (Set A: $\beta = -0.4 \pm 0.0$; $F_{(1,196)} = 138.7, p = 1.5 \times 10^{-24}, \eta^2 = 0.2$; Set B: $\beta = -0.4 \pm 0.0$; $F_{(1,196)} = 128.9, p = 2.8 \times 10^{-23}$, $\eta^2 = 0.1$; Figure 2d & f). In contrast, TE-only trials again failed to elicit any sign-dependent changes in hand angle (Set A: $\beta = 0.0 \pm 0.1$; $t_{(196)} = 0.5, p = 0.62$, $D = 0.1$; Set B: $\beta = 0.0 \pm 0.1$; $t_{(196)} = -0.4, p = 0.72, D = 0.1$; Figure 2f & h). The interaction between perturbation type and size was replicated (Set A: $\beta = 0.4 \pm 0.0$; $F_{(1,196)} = 67.5, p = 2.9 \times 10^{-14}, \eta^2 = 0.2$; Set B: $\beta = 0.4 \pm 0.0$; $F_{(1,196)} = 93.3, p = 2.7 \times 10^{-18}, \eta^2 = 0.3$), showing robust implicit recalibration when both SPE + TE were present, but not when TE-only was provided.

Together, these results indicate that TE alone is not sufficient to drive implicit recalibration. This stands in contrast to SPE, which leads to implicit recalibration whether or not TE is present (Kim et al., 2019; Leow et al., 2018; Leow, Marinovic, de Rugy, & Carroll, 2020). Moreover, these results challenge the hypothesis that SPE and TE operate *strictly* in an independent manner.

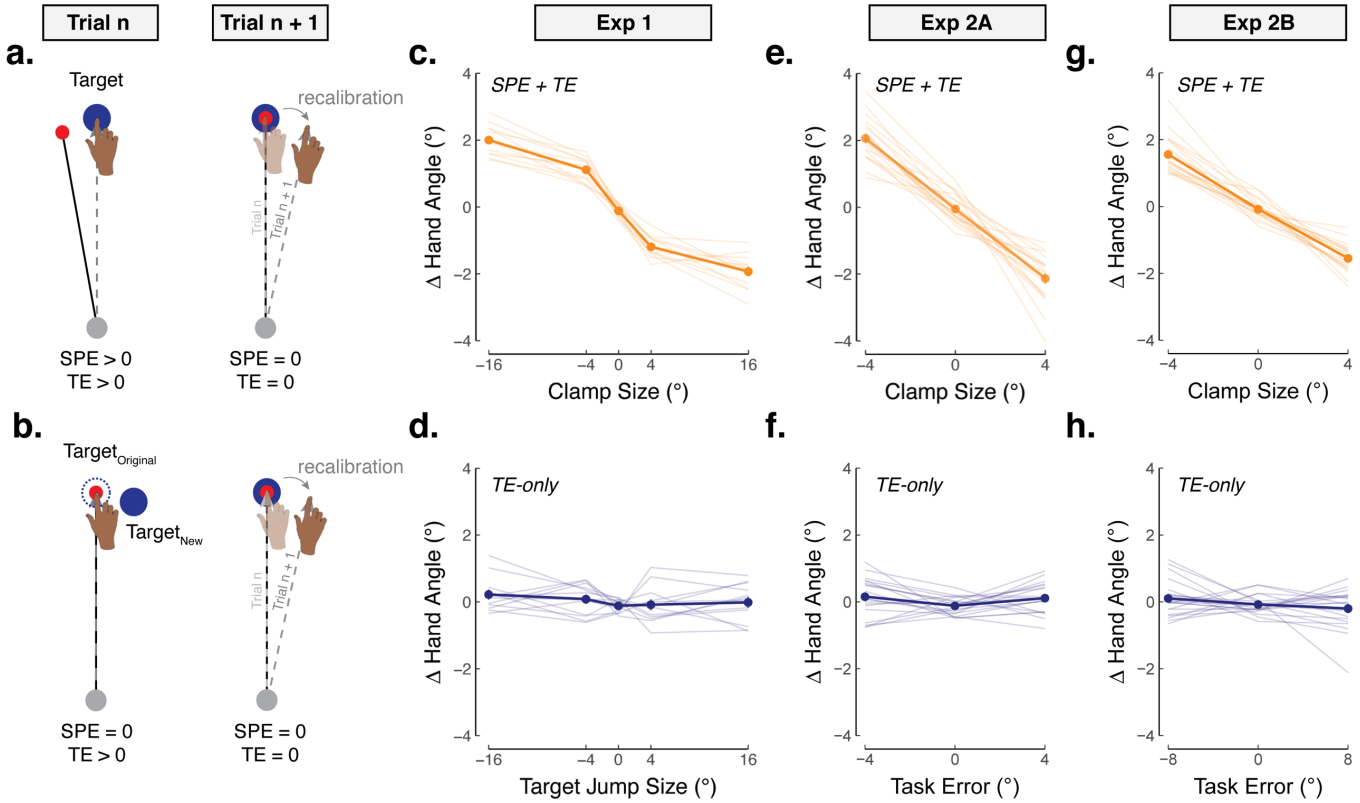

Figure 2: Task error alone does not elicit implicit recalibration (Exp 1 – 2). Using clamped visual feedback for testing implicit responses to: **(a)** SPE + TE, induced by offsetting the cursor trajectory at a fixed angle relative to the target, and **(b)** task error (TE), induced by jumping the target immediately after movement initiation, with the cursor clamped to 0° (the original target location). **(c) – (d)** In Exp 1, participants experienced 4 alternating blocks of target jumps and clamped feedback (201 trials/block). The perturbation sizes within a given block were randomized in order to prevent accumulated learning. Adaptation was quantified by measuring how much the hand angle changed on trial n + 1 in response to the perturbation on trial n. **(e) – (h)** In Exp 2, participants experienced a fully randomized (mixed) schedule of target jumps (Set A with ±4° TE-only perturbations and Set B with ±8° TE-only perturbations) and clamped feedback (both sets with ±4° SPE + TE perturbations). Dots connected with thick line represent the across participant average; thin lines represent individual data. Note that the x- and y-axes are not drawn on the same scale.

Modeling the potential ways in which TE and SPE may interact to drive implicit recalibration.

Although TE alone may not induce recalibration, previous work has shown that the presence or absence of TE will modulate the response to SPE (Kim et al., 2019; Leow et al., 2018, 2020). To understand the potential ways in which SPE and TE may interact to drive learning, we considered several models that encapsulated a variety of possible mechanisms. Figure 3 shows these models with their predicted responses to a fixed clamp size (i.e., fixed SPE) and varying TE size.

We first consider two simple base models, both of which cannot account for previously established results (including Experiment 1) but will serve as a foundation and a contrast for more elaborated models. The first model is one in which TE does not contribute to implicit recalibration. By this Invariant SPE model, we would expect recalibration to be invariant to

the size of target jumps and thus the size of TE (Figure 3a). As noted above, this model is insufficient given the demonstrations in the literature where the response to feedback involving only SPE is attenuated compared to feedback in which there is both SPE and TE (Figure 1).

The second base model is one in which TE and SPE make independent contributions to implicit recalibration (Dual-Error model), with their respective contributions simply being summed. Consequently, jumping the target in the same signed direction as the clamped cursor (e.g., clockwise target jump and clockwise clamp) will decrease the absolute magnitude of TE (as long as the target displacement is less than twice the magnitude of the cursor perturbation). This ought to decrease recalibration since SPE and TE make opposing contributions to the behavioral change. Conversely, jumping the target away from the cursor will increase TE, and thus increase recalibration (Figure 3b). This model, however, cannot account for the failure of TE-only to elicit recalibration (see Figure 2d, e, f).

Building on the failure of these base models, we considered potential ways in which task outcome might influence recalibration to SPE in an interactive manner. One possible way is based on the hypothesis that recalibration is attenuated by a scalar intrinsic reward signal that simply indicates whether or not the movement goal was achieved (i.e., whether or not the cursor "hits" the target) (Cashaback, McGregor, Mohatarem, & Gribble, 2017; Galea, Mallia, Rothwell, & Diedrichsen, 2015; Izawa & Shadmehr, 2011; Kim et al., 2019; Konrad Paul Körding & Wolpert, 2004; Nikooyan & Ahmed, 2015). The intrinsic reward signal can be interpreted as a gain controller, similar to previous efforts to model the effect of explicit rewards and punishments on recalibration (Galea et al., 2015). That is, when the movement goal is achieved, the drive to recalibrate the motor system is reduced. This Rewarded SPE model predicts a transient drop in recalibration only for a narrow range of target jumps corresponding to the cursor hitting the target (Figure 3c).

An alternative model is that recalibration might be modulated by the distracting presence of the target jump, perhaps due to attention being directed away from processing the feedback signal (Song, 2019; Taylor & Thoroughman, 2007). This Distracted SPE model is grounded in a rich history of visual psychophysics revealing worse accuracy at detecting, discriminating, and processing visual stimuli in unattended regions of visual space (Dosher, Sperling, & Wurst, 1986; Kinchla, 1992). For the model, we assume that displacing the target distracts attention away from the feedback cursor, and thus decreases the efficacy of recalibration. As a first approximation, we model this as a gaussian gain function in which

the attentional cost increases with the magnitude of the target jump (variable cost depicted in Figure 3e), an assumption we will test in Experiment 4.

This attentional hypothesis highlights that jumping the target has two effects: In addition to modifying the size of a putative TE signal, the standard motivation for this manipulation (Leow et al., 2018, 2020), it is a source of attentional distraction. One way to separate these factors is to transiently turn off the target while keeping its position fixed. Assuming the flicker serves to distract attention, this "jump-in-place" condition would identify an attentional cost that is independent of the change in TE, an assumption we will test in Experiment 3. This attenuating effect is shown in Figure 3e as a fixed attentional cost, that is, implicit recalibration when the target flickers in the same place during the trial (jump-in-place) would be attenuated compared to a condition when the target remains stationary and visible throughout the trial (no-jump). This fixed cost rides on top of a variable attentional cost that is dependent on the distance of the target displacement.

The Rewarded SPE (Figure 3c) and Distracted SPE (Figure 3e) models consider the modulatory effects of intrinsic reward and attention on a base model in which TE does not directly influence implicit recalibration (the Invariant SPE model). We also considered how the modulatory effects of reward and attention might influence implicit recalibration if both SPE and TE drive learning (Dual-Error model). The predictions of these hybrid, dual-error models are presented in Figures 3d (Rewarded SPE+TE) and 3f (Distracted SPE + TE), both of which predict an asymmetrical effect of target jumps.

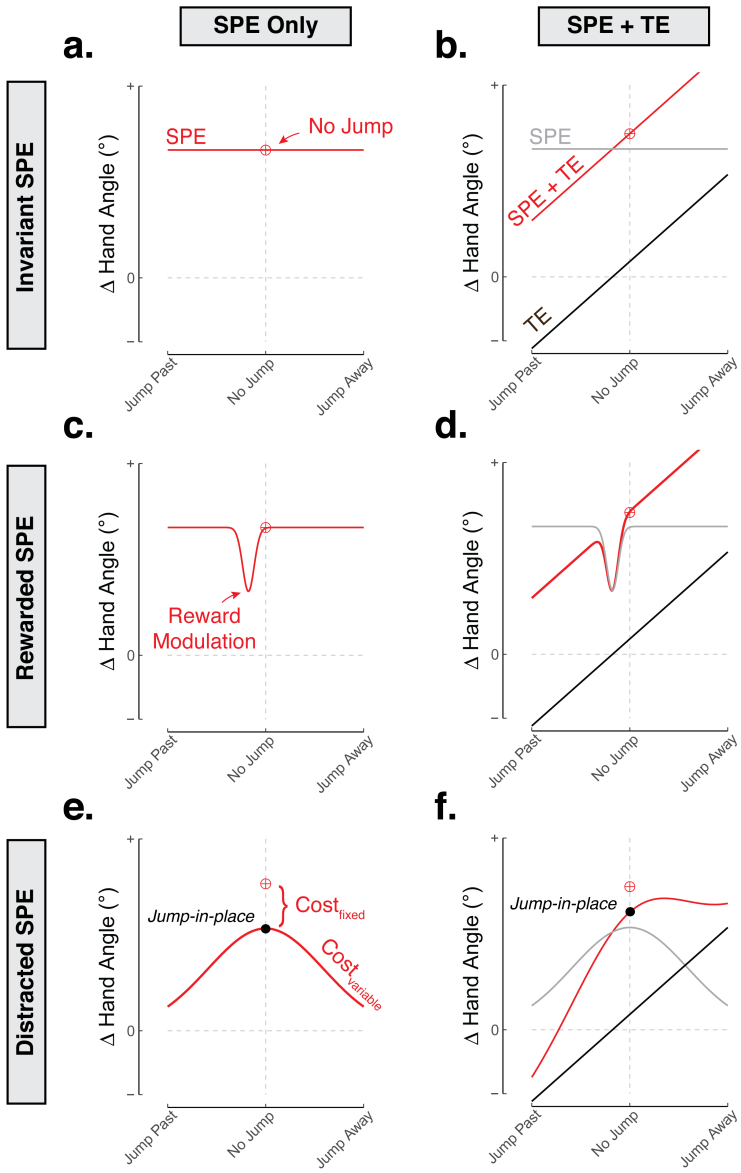

Figure 3: Modeling the influence of target jumps on adaptation to TE and SPE. Given a constant SPE magnitude, SPE may be **(a)** impervious to target jumps **(b)** attenuated when the cursor "hits" the target (modulated by intrinsic reward), or **(c)** attenuated due to the motion of the jumping target diverting attention away from computing a SPE. The attenuation is assumed to be driven by the mere presence of a target jump (fixed cost – an effect isolated by flickering the target, also known as the jump-in-place condition) and varied with target jump size (variable cost). Right column **(b)**, **(d)**, **(f)**: Adaptation may also be driven by a TE-based learning process, assumed here to be a linear function of the distance between the feedback and new position of the target. The red indicates expected behavior, which is the composite of the SPE process (grey) and TE-based process (black).

Experiment 2: TE modulates implicit recalibration in the presence of SPE.

To empirically examine the interactions between SPE and TE, and evaluate the models described above, we performed a second experiment in which we varied the size of target jumps in the context of an SPE, induced by non-zero clamped feedback (see Table 3). To vary the size of TE, we jumped the target between ±8° away from the original target location.

For the non-zero SPE, we clamped the cursor at ±4° from the original target, randomizing the direction of the feedback cursor from trial to trial.

In response to a stationary target (i.e., no jump), participants adapted 1.5° in response to a 4° clamp (Figure 4a). When the target jumped towards the cursor, implicit recalibration was reduced in a roughly stepwise, linear manner (Table 3): Jump-to (i.e., target jumps to the cursor) reduced implicit recalibration by 13% and jump-past (i.e., target jumps in the direction of *and* beyond the cursor) reduced implicit recalibration by 33%. The fact that jumping the target influenced behavior argues against the Invariant SPE model; task outcome indeed influences behavior in the presence of SPE. The graded effect is also not compatible with the Rewarded SPE models (Figure 3c, d), as these models predict a modulating effect of target jumps only when the target intersects the cursor feedback, providing a putative intrinsic reward.

Implicit recalibration was greater when the target jumped away from the cursor compared to when it jumped past (0.3 ± 0.1; $t_{(77)} = 2.6, p = 0.01, D = 0.8$) (Figure 4a; Table 1). This pattern is most consistent with the unique, asymmetrical function predicted by the Distracted SPE + TE model (Figure 3e) and refutes the symmetrical function predicted by the Distracted SPE-only model (Figure 3f). That is, implicit recalibration may be dependent on both SPE and TE (conditioned on the presence of SPE), although the act of manipulating TE via target jumps may have a distracting effect that reduces sensitivity to SPE.

Experiment 3: Target jumps vary the size of TE but also attenuate implicit recalibration.

Exp 3 was designed to provide a strong test of the assumption that jumping the target distracts attention: Namely we predict that recalibration in response to an SPE will be attenuated by distraction, even if the distracting event does not influence TE (or SPE). To test this prediction, we introduced a condition in which the target was perturbed *without* changing locations, disappearing upon movement initiation and then reappearing in its original location on the next screen refresh (jump-in-place; i.e., flickering the target). The difference between implicit recalibration for no-jump (i.e., stationary target) and jump-in-place should indicate the effect of distraction. By varying the size of the SPE, we can ask if the magnitude of the distraction effect is independent of SPE magnitude. To test this prediction, we used two clamp sizes (±3° and ±7°). Because the experiment was conducted online (in response to COVID pandemic restrictions regarding in-person testing), we were able to increase the sample size.

On average, participants adapted 1.1° and 1.5° in response to 3° and 7° clamps, respectively (no jump; Figure 4b-c; Table 1). Strikingly, the response was attenuated in the jump-in-place conditions despite the fact that the SPE and TE were identical to that in the corresponding no-jump conditions. Moreover, the magnitude of this effect, which represents the fixed attentional cost on recalibration, was similar for the two clamp sizes, ~40% (no interaction: $= 0.4 \pm 0.0$; $F_{(3,294)} = 0.1, p = 0.96, \eta^2 = 0$). In addition to the fixed attentional cost due to the flicker of the target, we observed an approximately linear effect of TE on implicit recalibration. For instance, in Exp 3A recalibration was larger by approximately 0.5° in the jump-in-place condition compared to the jump-to condition, and increased by another 0.5° in the jump-away condition (Table 1). This linear effect of TE is uniquely predicted by the Distracted SPE + TE model.

In summary, the results indicate that perturbing the target yields 1) an asymmetrical hand angle function, 2) a fixed cost most clearly evident in the jump-in-place condition, and 3) a linear effect of TE *after* accounting for this fixed attentional cost. These three effects, in aggregate, provide strong support that jumping the target not only modulates the size of the TE, but also attenuates the extent of recalibration from SPE (Distracted SPE+TE model). Notably, these results were replicated both in-lab and online settings.

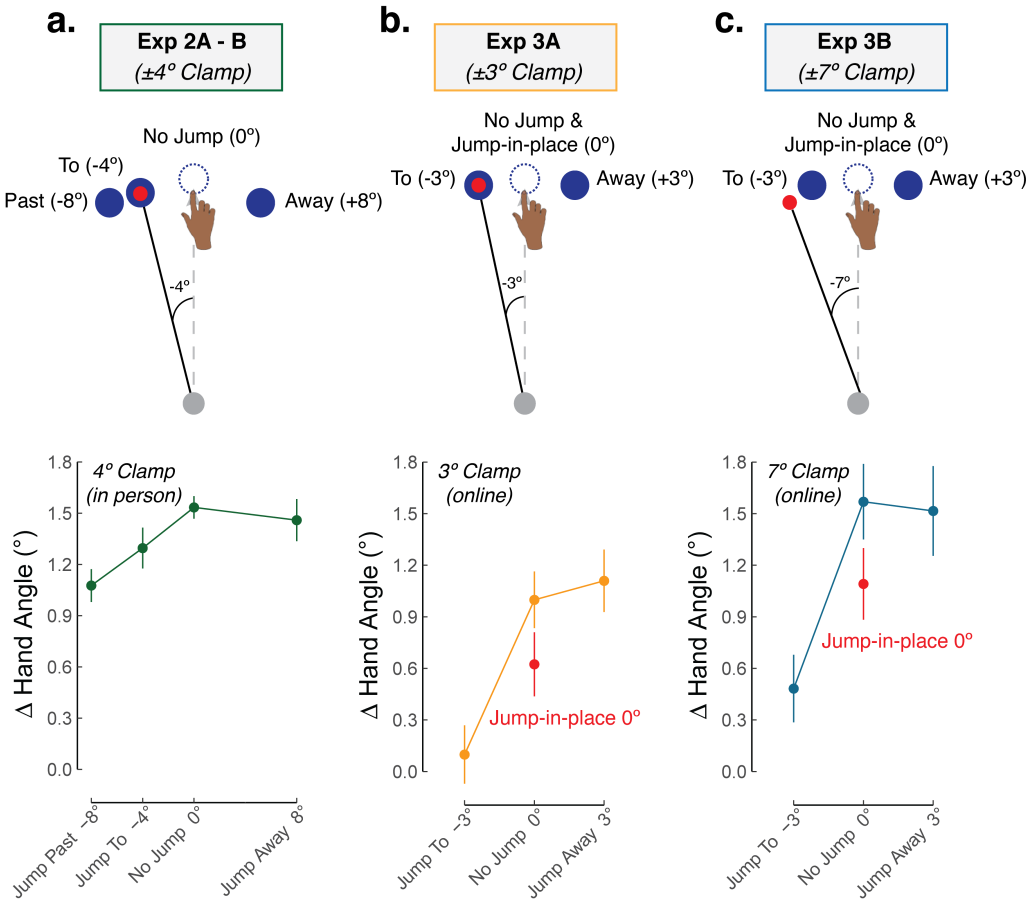

Figure 4: Implicit recalibration is modulated by TE in the presence of SPE (Exp 2 – 3). (a) – (c) Participants experienced a randomized zero-mean perturbation schedule where both clamp size (Exp 2, in-person: ±4° clamp; Exp 3, online: ±3° or ±7°) and target jump size (Exp 2 range: -8 to 8; Exp 3 range: -3 to 3) were varied. A positive change in hand angle signified recalibration in the expected direction, by flipping the sign of hand angles in response to counterclockwise (+) clamped feedback and clockwise (-) target jumps. Dots represent mean and vertical lines represent SEM.

Table 1	**Exp 2A-B (–4° Clamp)**				**Exp 3A (–3° Clamp)**				**Exp 3B (–7° Clamp)**			
Fixed Effects	Past	To	No Jump	Away	To	Jump-in-place	No Jump	Away	Near	Jump-in-place	No Jump	Away
Target Jump Size	–8°	–4°	0°	+8°	–3°	0°	0°	+3°	–3°	0°	0°	+3°
Mean (*SEM*)	1.1 (0.1)	1.3 (0.1)	1.5 (0.1)	1.4 (0.1)	0.1 (0.2)	0.6 (0.2)	1.0 (0.2)	1.1 (0.2)	0.5 (0.2)	1.1 (0.2)	1.5 (0.2)	1.5 (0.2)
Mean – No Jump (*SEM*)	–0.5 (0.1)	–0.2 (0.1)		–0.1 (0.1)	–1.0 (0.2)	–0.4 (0.2)		0.0 (0.2)	–1.0 (0.2)	–0.4 (0.2)		0.0 (0.2)
D	–1.0	–0.4		–0.2	–0.8	–0.3		0.0	0.3	0.2		0.1
P	<0.001	0.14		0.42	<0.001	0.02		0.88	<0.001	0.02		0.88

Table 1: Summary of model-free results. Mean estimates (SEM) from the linear mixed effect model for each target jump condition. Changes in hand angle in response to counterclockwise (+) clamped feedback were flipped to clockwise (-), such that a positive change in hand angle always signify adaptation in the expected direction (i.e., away from the clamped feedback). Contrasts between no jump and other target jump conditions are also shown, with Cohens' D and P values provided. Significant contrasts ($P < 0.05$) are highlighted in a shaded light-grey box.

Experiment 4: Implicit recalibration reflects the joint contribution of TE, SPE, and the distractive effects of target jumps

To further probe how the distracting effect of target jumps interacts with the magnitude of TE, we sampled a wide range of target jump sizes in Experiment 4 (Figure 5a). As shown in Figure 3f, we assume that the attenuating effect of distraction will increase with the size of the target jump due to attention being further displaced from the feedback cursor. As such, the inclusion of a larger range of target jumps should produce a marked asymmetrical function.

This prediction was confirmed (Figure 4b): Implicit recalibration decreased when the target jumped towards the cursor and remained relatively invariant when the target jumped away from the cursor, even as far as 30° (jump-away). This phenomenon could be attributed to the contribution of a TE process that offsets the attentional costs of target jumps on a SPE-based implicit recalibration process.

Sampling a wider range of target jumps also allowed us to fit our candidate models to the data (see formalization in Table 4 of the Methods section). In doing so, we could quantitatively evaluate how well our six candidate models fit the data while taking into account model complexity. Consistent with the qualitative assessments described above, the Distracted SPE + TE model provided the best fit, having the highest R^2 and lowest AIC (Table 2).

The modeling work also allowed us to evaluate the best fitting parameters of the Distracted SPE + TE model. The parameter values suggest that TE may contribute to learning. The slope (β_{TE}) of the TE function was 0.02 ± 0.003, suggesting that of the 0.5° change in hand angle observed for the 3° no-jump condition (where both SPE and TE are present), ~12% of the change came from TE. Similarly, when the error increased to 7° in a no-jump condition, ~16% of the 0.09° change in hand angle came from TE. These results indicate that SPE has a much larger impact on implicit adaptation compared to TE.

The Distracted SPE + TE model has two parameters to capture the effects of perturbing the target. First there is a fixed effect arising from the transient changes that occur when the target is perturbed. The estimate of this parameter (C_J) in the best fitting model was 0.84 ± 0.13. Thus, the mere perturbation of the target, even if it was not spatially displaced reduced recalibration by 15%. Second there is a variable cost (σ_d^2) due to SPE-based learning being attenuated as the target jump

distance increased. The estimate of this parameter was 11.8 ± 2.3. From this value, SPE would no longer be effective in driving implicit recalibration for target jumps greater than 35°.

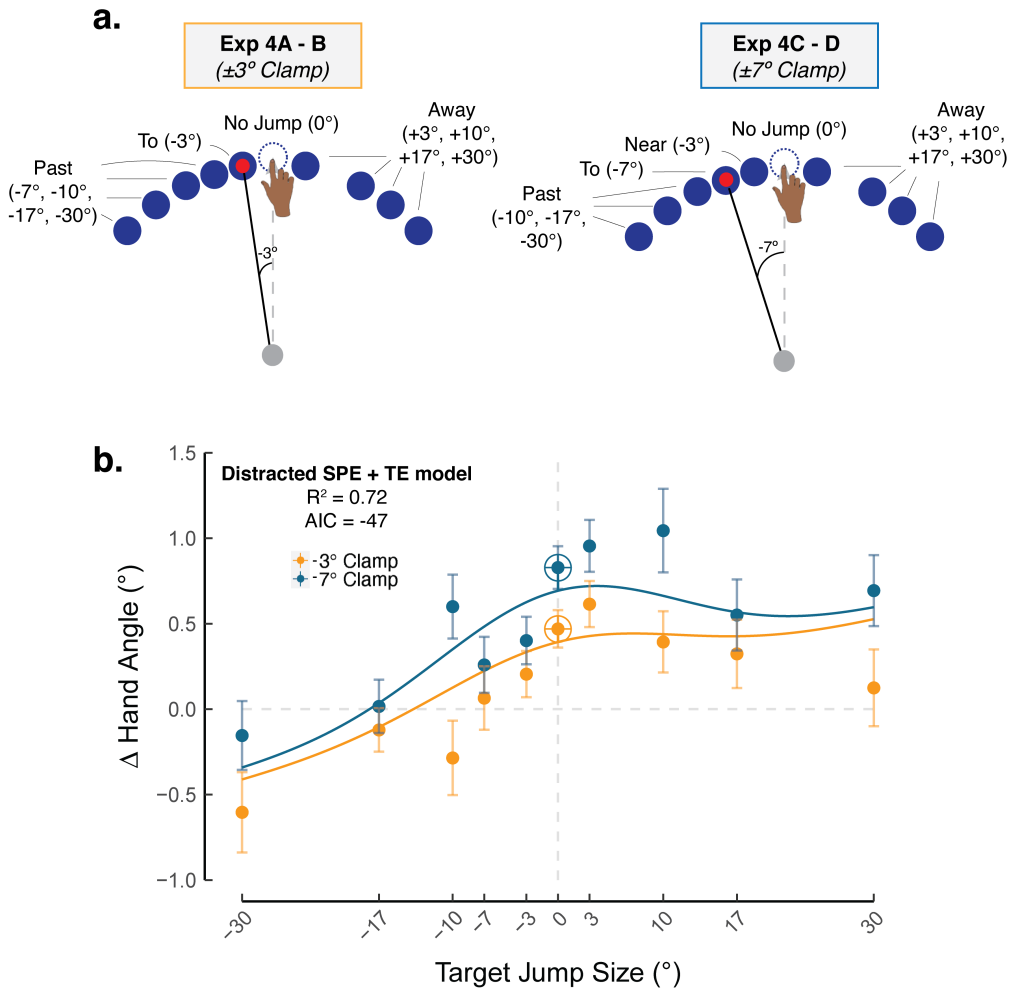

Figure 5: Implicit recalibration reflects the contribution of learning from task error and sensory prediction error, with the latter sensitive to distraction from target jumps (Exp 4). (a-b) Participants experienced a randomized zero-mean perturbation schedule with clamp sizes (-3° shown in orange; -7° shown in blue) × target jumps (x axis, -30° through 30°). The Distracted SPE + TE model was the winning model.

Table 2	SPE Only			SPE + TE		
	# of free param	R^2	AIC	# of free param	R^2	AIC
Invariant SPE	0	-0.54	-18	1	-0.54	-18
Rewarded SPE	2	0.15	-28	3	0.53	-37
Distracted SPE	2	0.33	-33	3	0.72	-47

Table 2: Summary of model-based results.

Discussion

Although it is widely recognized that implicit sensorimotor recalibration serves to minimize motor execution errors, the error signals that drive this learning process remain the subject of considerable debate (Kim et al., 2020; J. Krakauer et al., 2019; Shadmehr et al., 2010). In particular, the idea that sensory prediction error (SPE), the mismatch between the expected and actual feedback, is the sole learning signal has been challenged by recent evidence demonstrating that task error (TE), the mismatch between the target location and feedback may also impact implicit recalibration (Albert et al., 2020; Kim et al., 2019; Leow et al., 2018; Miyamoto, Wang, & Smith, 2020). Whether these two types of error drive implicit recalibration independently or interactively remains unknown.

In traditional sensorimotor adaptation tasks, SPE and TE are confounded. Displacing the hand in a force field or perturbing the feedback in a visuomotor rotation task introduces both SPE and TE. To unconfound these signals, researchers have developed methods that selectively influence one signal or the other. For example, by making the angular trajectory of the feedback cursor independent of the movement, an SPE of a fixed size may either be accompanied by TE (when the target is small, and the cursor misses the target) or occur without TE (when the target is large, and the cursor hits the target). Conversely, displacing the target (i.e., target jump) selectively modulates TE given the assumption that the expected location of the feedback remains at the original target location.

Building on these methodological advances, we designed a series of experiments to systematically manipulate SPE and TE and used the data to test a set of computational models. We first considered a model in which these two types of error make independent contributions to implicit recalibration, with the resultant behavior being the composite operation of two distinct learning processes (Figure 3a, b). This idea takes inspiration from the work of Mazzoni and Krakauer (2006) who showed that implicit recalibration continued to operate even in the absence of task error, a result that suggests SPE-dependent learning is modular. A natural extension of this modular, dual-error model would posit that TE alone should also be sufficient to drive implicit recalibration. In three experiments, we failed to find support for this hypothesis. When the clamped feedback moved directly to the original target location (no SPE), hand angle remained unchanged in response to target jumps. That is, TE in the absence of SPE failed to induce implicit recalibration, arguing against models in which SPE-dependent and TE-dependent learning processes operate in a strictly independent manner.

Given the failure of this simple model and the dependency of TE on SPE, we considered different ways in which SPE-dependent and TE-dependent processes might interact. We varied task outcome in a continuous manner by jumping the target, either away from the perturbed cursor (increasing TE), towards the perturbed cursor (reducing TE), or to the location of the perturbed cursor (nullifying TE; i.e., SPE only). In stark juxtaposition to the failure of TE only to elicit implicit recalibration, SPE only reliably elicited implicit recalibration, which is a finding consistent with previous literature (Kim et al., 2019; Leow et al., 2018, 2020).

By modulating TE in a fine-grained, continuous manner, we revealed an unexpected, asymmetrical effect on implicit recalibration: Implicit recalibration decreased when TE decreased yet remained largely unaffected when TE increased. These results are at odds with the hypothesis that task outcome provides a *binary* reward signal (Figure 3c, d), with TE being present when the cursor misses the target and TE being absent when the cursor hits the target (Cashaback et al., 2017; Galea et al., 2015; Kim et al., 2019; Konrad Paul Körding & Wolpert, 2004; Nikooyan & Ahmed, 2015). This asymmetrical function is also at odds with the hypothesis where SPE-learning, the sole process driving implicit recalibration, is attenuated by a generic distractor effect of displacing the target (Figure 3e).

Instead, the pattern of results supports a hybrid model, where implicit recalibration is driven by both TE and SPE, with each error signal having a modulatory effect on the other error signal (Figure 3f). Implicit recalibration scales with the size of TE, but only when SPE is also present. Implicit recalibration also scales with the size of SPE but is attenuated when the target is perturbed. We hypothesize that the modulation of SPE-based learning occurs because attention is directed away from the feedback cursor and towards the (displaced) target, an effect that increases with the size of the displacement. Taken together, this hybrid perspective underscores the rich, dynamic interplay between two distinct error signals that drive implicit recalibration in an interactive manner.

We recognize that at this stage of development, the models are largely descriptive, intended to provide a qualitative sense of the behavioral changes that would be expected given different ways in which sensory prediction error and task error might interact. Future work will be required to develop stronger theoretical foundations and more rigorous experimental tests for the different assumptions underlying the models; for example, to ask if the signals follow normative principles such as optimal integration (Burge, Ernst, & Banks, 2008; Ernst & Banks, 2002) or relevance estimation (Wei & Körding, 2009).

While implicit recalibration seems to scale with TE in the presence of SPE, it remains unclear why TE alone fails to elicit recalibration. One possibility is that SPE serves as a gating signal, with recalibration only engaged in the presence of SPE and then responding to all sources of error information. Alternatively, the lack of an SPE may allow the brain to correctly attribute the target jump to an external cause (Konrad P. Körding et al., 2007; Shams & Beierholm, 2010; Wei & Körding, 2009). In contrast, when a TE occurs in the presence of an SPE, the brain may be less confident in attributing the error to an external cause and, as such, use both sources of information to recalibrate the sensorimotor system. This latter hypothesis might account for the recent findings of Ranjan and Smith (2020) who observed robust adaptation in response to TE alone, a result at odds with the current study (Ranjan & Smith, 2020). Whereas our participants were told to ignore the manipulations of the cursor and the target, Ranjan and Smith instructed participants to hit the target with their cursor. These instructions may have rendered the position of the displaced target and the cursor relevant, thereby motivating participants to recruit more explicit re-aiming strategies to reduce TE (Oza, Kumar, & Mutha, 2020).

The current study also highlights an important methodological issue. Similar to the way error clamps have provided a tool to isolate implicit recalibration, target jumps have been viewed as a way to provide a "pure" manipulation of TE. However, our results show an attenuated effect on implicit recalibration from the transient effects associated with perturbing the target, a result made salient by the conditions in which the target briefly disappeared and then reappeared at its original location. The transient sensory events associated with a target jump or flash might siphon attention away from the visual feedback, thereby weakening the overall learning signal. Alternatively, a transient distraction may have increased the likelihood that visual feedback is mis-localized, thus attenuating the motor system's reliance on this uncertain feedback (Burge et al., 2008; Konrad P. Körding & Wolpert, 2004; Tsay, Avraham, et al., 2020; Tsay, Kim, Parvin, Stover, & Ivry, 2021; Wei & Körding, 2010). Regardless of the mechanism, our results underscore the importance of considering the distractive effect of a target jump manipulation and the consequences of this on implicit recalibration.

Methods:

Participants and Apparatus

All participants were between 18 – 30 years old and right-handed, as determined by either the Edinburgh handedness inventory (Oldfield, 1971) or through self-report. The protocol was approved by the IRB at University of Delaware and UC Berkeley.

In-person participants (Exp 1 – 2): Undergraduate students were recruited from the University of Delaware community, receiving financial compensation for their participation at a rate of $10/hour. Participants were seated in front of a custom tabletop setup and held the handle of a robot manipulandum (KinArm: BKIN Technologies, sampling rate 200 Hz) that was positioned below a mirror. Visual feedback was projected by a monitor placed directly above onto the mirror, which occluded vision of the participant's hand during the experiment. Peripheral vision of the arm was minimized by extinguishing the room lights. Participants completed the task by moving the robot manipulandum, which was constrained to a horizontal 2D plane.

Online participants (Exp 3 – 4): Participants were recruited via Amazon Mechanical Turk, receiving financial compensation for their participation at a rate of $8/hour. Participants used their own laptop computer to access a customized webpage (Tsay, Lee, Ivry, & Avraham, 2021) hosted on Google Firebase (sampling rate typically ~60 Hz) (Anwyl-Irvine, Dalmaijer, Hodges, & Evershed, 2020; Bridges, Pitiot, MacAskill, & Peirce, 2020). Recruitment was restricted to trackpad users to minimize variability from different response devices. Participants completed the task by swiping their index finger on the trackpad.

Reaching Task Procedure

In-person procedure: Reaches were made from a start location to one target (90° location, straight ahead). The start location was indicated by a white ring (6 mm diameter) and the target by a blue circle (6 mm diameter), with the radial distance between the start location and target fixed at 10 cm. To initiate a trial, the robot arm moved the participant's hand to the start location. Visual feedback of the hand position was given via a cursor (white circle 3.5 mm diameter) only when the hand was within 1 cm of the start position. Once the hand remained within the start location for 500 ms, the target appeared,

serving as a cue to indicate the location of the target and an imperative to initiate the reach. To discourage on-line corrections, participants were instructed to perform fast, 'shooting' movements through the target as soon as the target appeared.

Reaction time (RT) was defined as the time from initial target presentation to the start of movement (defined as when the hand first exceeded 5 cm/s for at least 50 milliseconds). Movement time (MT) was defined as the time between the start of movement and when the hand crossed the radial target distance of 10 centimeters. To ensure that participants moved at a fast speed that excluded online feedback corrections, the message "Too Slow" appeared on the screen at the end of the trial when MT was < 40 cm/s at peak velocity. We also presented the message "Too Fast" if MT was > 70 cm/s at peak velocity to ensure that participants did not make simple ballistic movements in the general direction of the target (this criterion was rarely exceeded). After completing the reach, the participant was instructed to keep the arm and shoulder relaxed as the robot moved the hand back to the starting position.

Online procedure: The reaching task was adapted for an online study. We did not obtain information concerning the monitors used by each participant; as such, we cannot specify the exact size of the stimuli. However, from our experience in subsequent studies, we assume that most online participants used a laptop computer. To provide a rough sense of the stimulation conditions, we assume that the typical monitor had a 13" screen with a width of 1366 pixels and height of 768 pixel (Anwyl-Irvine et al., 2020). The center position was indicated by a white circle (0.5 cm in diameter) and the target location was indicated by a blue circle (also 0.5 cm in diameter). To ensure that reaches remain in the trackpad, we reduced the radial distance of the target to 6 cm and positioned the target at the 45° target (upper right quadrant).

The participant made center-out planar movements by moving the computer cursor with her trackpad to a visual target. To initiate each trial, the participant moved their hand to the start location. Visual feedback of the hand position was given via a cursor (white circle 0.5 cm diameter) when the hand was within 1 cm of the start position. Once the hand remained within the start location for 500 ms, the target appeared, serving as a cue to indicate the location of the target and an imperative to initiate the reach. To discourage on-line corrections, participants were instructed to perform fast, 'shooting' movements through the target as soon as the target appeared.

RT was defined as the time from initial target presentation to the start of movement (i.e., when the hand movement exceeded 1 cm from the start location). Due to the lower sampling rate of standard computer monitors compared to in-person setup, we opted to define RT in terms of movement distance (requiring fewer samples) rather than movement velocity (requiring more samples to adequately estimate). There were no constraints on RT. MT was defined as the time between the start of the movement and when the radial distance of the movement reached 6 cm. To ensure that the movements were made quickly, the computer displayed a "too slow" message if MT exceeded 300 ms. We did not include a "too fast" message since participants recruited online, based on our pilot results, err on the side of moving too slowly.

There were three types of cursor feedback trials used throughout the in-person and online experiments: On veridical feedback trials, the cursor corresponded to the position of the hand. On clamped feedback trials, the cursor followed an invariant path along a constant angle with respect to the target. The radial distance of the cursor, relative to the start position, was yoked to the participant's hand. In both types of feedback trials, the radial position of the cursor matched the radial position of the hand until the movement amplitude reached the radial distance of the target, at which point the cursor froze for 50 ms. On no-feedback trials, the cursor was blanked when the target appeared, and did not re-appear until the participant had completed the reach and returned to the start location for the next trial.

There were also target jump trials, where upon movement initiation (i.e., in-person: velocity > 5 cm/s; online: radial distance > 1 cm), the original target was blanked and immediately re-positioned at a new target location (i.e., one screen refresh between offset of original target to onset of new target; in-person: within 1 ms; online: <15 ms, accounting for the delay in the monitor system (Anwyl-Irvine et al., 2020; Bridges et al., 2020)). We varied the size of the target jump and categorized these based on the relative position of the new target location to the clamped cursor position: jump-past, jump-to, jump-near, jump-away, and jump-in-place. When the target jumps in the direction of the the clamped cursor feedback, the size of the target jump could either be greater than (jump-past), equal to (jump-to), or less than (jump-near) the clamped angle. On jump-away trials, the target was repositioned in the direction opposite to the clamped feedback. On jump-in-place trials, the target disappeared upon movement initiation (1 refresh) and then reappeared (1 refresh) in the same (original) location (<30 ms, accounting for delays in the system). While jump-in-place has a longer interval between successive displays of the target compared to other target jump conditions, this interval ensured that jump-in-place trials elicited a detectable disturbance to the visual display, something that was obvious in the other target jump conditions.

Table 3	N	Setting	Perturbation Conditions			
			Set	Clamp size (°)	Target jump (°)	Figure
Exp 1	12	In-Person	––	0, ±4, ±16	0	2c
				0	0, ±4, ±16	2d
Exp 2	40	In-Person	A	–4	0, –4, –8	2e, 3a
				+4	0, +4, +8	
				0	0, ±4	2f
			B	±4	0, ±8	2g, 3a
				0	0, ±8	2h, 3a
Exp 3	100	Online	A	±3	±3, 0, $0_{\text{jump-in-place}}$	3b
			B	±7	±3, 0, $0_{\text{jump-in-place}}$	3c
Exp 4	210	Online	A	+3	–10, –3, 0, +3, +7, +10, +17	4b
				–3	+10, +3, 0, –3, –7, –10, –17	
			B	+7	–10, –3, 0, +3, +7, +10, +17	
				–7	+10, +3, 0, –3, –7, –10, –17	
			C	±3	±0, ±10, ±17, ±30	
			D	±7	±0, ±10, ±17, ±30	

Table 3: Summary of experiments.

Experiment 1 – 2, In-person Experiments

Reaching trials were performed to the 90º target (straight ahead). The experiment began with 100 baseline reaching trials with veridical feedback, provided to familiarize the participants with the reaching task. These trials were used to emphasize that movements should "shoot" through the target and demonstrate that the feedback and target would disappear soon after the movement amplitude exceeded the radial distance of the target.

The participant then completed a block of perturbation trials. Just before the start of this block, the error clamp and target jump manipulations were described to the participant, and she was told to ignore the cursor "feedback" as well as any change in the position of the target, always attempting to reach directly to the original target. To help the participant understand the task irrelevant nature of the clamped feedback and target jump, three demonstration trials were provided. The target appeared straight ahead at 90º and the participant was told to reach to the left (demo 1), to the right (demo 2), and backward (demo 3). The cursor moved in a straight line with a 45º offset from the original target in all three trials, and the target "jumped" upon movement initiation 0° (demo 1), 45° (demo 2), and 90° (demo 3) away from the original target.

In Exp 1, the perturbation block was composed of mini-blocks (Table 3; 804 perturbation trials = 4 mini-blocks x 201 trials/mini-block) of either SPE + TE perturbations (i.e., when clamped feedback is paired with a stationary target) or TE-only perturbations (i.e., when a 0° clamp is paired with a target jump). We opted to keep these perturbation conditions separate to minimize any interference or generalization of learning from one trial type to another (Dang, Parvin, & Ivry, 2019; J. W. Krakauer, Ghilardi, & Ghez, 1999; Lerner et al., 2020). SPE + TE and TE-only mini-blocks were interleaved, with the order counterbalanced across individuals. Within each mini-block, there were 20 trials per condition provided in a random, zero-mean order (with the exception of 21 trials for 0° clamp x 0° target jump). This resulted in 80 trials per clamp size x target jump combination across the entire experiment (84 trials in the 0° clamp, 0° target jump condition).

The perturbation block in Exp 2 was not composed of mini-blocks. Instead, we opted to randomize all perturbation conditions across the entire experiment (724 trials) to evaluate whether our results from Exp 1 would hold under another perturbation schedule. To sample a wider range of clamp size x target jump combinations while keeping the experiment within 1 hour to minimize fatigue, participants experienced different sets of perturbations (Set A or Set B). In Set A, the target always jumped in the same direction as the error clamp, while in Set B, the target either jumped in the same or in the opposite direction of the error clamp (Table 3). There were 80 trials per clamp size x target jump combination (84 trials for the 0° clamp, 0° target jump condition).

Experiment 3 – 4, Online Experiments

Due to the onset of the pandemic, Exp 3 – 4 were conducted online. With this approach, we were able to increase our sample size in an efficient manner, providing greater power to detect subtle differences between target jump conditions. We used an motor learning platform (OnPoint) (Tsay, Lee, et al., 2020, 2021) and recruited participants using Amazon Mechanical Turk. Despite substantial differences between in-person and online sensorimotor learning experiments (e.g., in-person: dark room to occlude vision of the hand; online: full visibility of the hand for trackpad users), we have found that the results obtained online are quite similar to those obtained in-person (Tsay, Lee, et al., 2021).

We made several additional changes to the experiment. We included "attention checks" to verify whether participants attended to the task. Specifically, during the inter-trial interval, participants occasionally were instructed to make an arbitrary response (e.g., "Press the letter "b" to proceed."). If participants failed the make the specified keypress, the experiment was

terminated. These attention checks were randomly introduced within the first 50 trials of the experiment. We also included "instruction checks" after our three demo trials to assess whether participants understood the nature of the error clamp and target jump manipulations: "Identify the correct statement. Press 'a': I will aim away from the original target. I will ignore the white dot. Press 'b': I will aim directly towards the original target location and ignore the white dot." The experiment was terminated if participants failed to make an accurate keypress (i.e., "b").

The block structure in Exp 3 and 4 were the same, composed of a baseline block with veridical feedback (28 trials) and a perturbation block with clamp feedback paired with target jumps (Exp 3: 120 trials; Exp 4: 252 trials). All perturbation conditions were randomized in a zero-mean manner throughout the experiment. The perturbation conditions were again divided into sets (See Table 1; Exp 3: Sets A—B; Exp 4: Sets A—D) to sample a wider range of clamp size x target jump combinations, while keeping the experiment within 1 hour. There were 30 trials per clamp size x target jump combination in Exp 3 and 18 trials per combination in Exp 4.

Data analysis, Model Free

The primary dependent variable of reach performance was the hand angle, defined as the hand position relative to the target when the movement amplitude reached the target distance (i.e., angle between the lines connecting start position to target and start position to hand).

Outlier responses were defined as trials in which the hand angle deviated by more than 3 standard deviations from a moving 5-trial window. These outlier trials were excluded from further analysis, since behavior on these trials could reflect attentional lapses or anticipatory movements to another target location (average percent of trials removed per participant ± SD: Exp 1: 0.2 ± 0.2%; Exp 2: 0.1 ± 0.2%; Exp 3: 0.8% ± 0.8%; Exp 4: 1.1 ± 0.1%).

As a measure of trial-by-trial implicit recalibration, we evaluated each participant's median change in hand angle on trial n + 1, as a function of the perturbation condition (clamp size x target jump) on trial n (Δ Hand Angle).

We sought to determine whether SPE + TE and TE-only perturbations elicit robust sign-dependent changes in hand angle (Exp 1 and 2). Specifically, in the SPE + TE condition, we expect implicit recalibration to result in a change in hand angle

in the opposite direction of the error clamp (e.g., a CW clamp eliciting a CCW change in hand angle). In contrast, in the TE-only condition, we expect implicit recalibration to be in the same direction as the target jump (e.g., a CW target jump eliciting a CW change in hand angle). To better visualize the difference between SPE + TE and TE-only conditions, the sign of the target jump was flipped, such that the expected change in hand angle would also be in the opposite direction of the perturbation (i.e., a negative target jump would elicit a positive change in hand angle). Each participants' data were submitted to a linear regression with perturbation size (Exp 1: 0, ±4°, ±16°; Exp 2, Set A: 0, ±4°; Exp 2, Set B: 0, ±4°, ±8°) and perturbation type (clamp vs target jump) as main effects. The mean regression slopes (β) ± SEM across participants were provided.

To ask whether the effect of TE would be conditional on the presence of SPE, we submitted each participants' data in Exps 2 and 3 to a linear regression with target jump size and task set as main effects. Post-hoc contrasts were performed using two tailed t-tests, and P values were Bonferroni corrected. The mean regression values (β) ± SEM across participants were provided.

Data analysis, Model Based

In this section, we formalize the six models justified in the Results section titled: "*Modeling the potential ways in which TE and SPE may interact to drive implicit recalibration*." The development of these models was based on different assumptions about how the size of target jumps (θ_j) and the size of the error clamp (θ_c) impact the processing of SPE and TE.

The first set of models posit that the motor system responds only to SPE (Table 4: SPE only column): First, SPE may be impervious to target jumps (Invariant SPE), where motor updates are not affected by target jumps ($U_{\theta_j=0}$, or the motor update during no-jump). Second, SPE may be attenuated when the cursor lands in the target, modulated by intrinsic reward (Rewarded SPE). The amount of reward modulation could vary with γ_r, a gain value determining the amount of attenuation, and σ_r, the standard deviation of reward function determining the scope of attenuation. Third, SPE may be attenuated due to a distracting effect of target jumps, which may siphon attention away from processing feedback and/or the movement goal (Distracted SPE). The attenuation may be due to the presence of a target jump (a fixed cost, C_J) and the size of the target jump (variable cost, modeled as a gaussian decay with standard deviation σ_d).

We recognize that the distracted SPE hypothesis may take on a different form, where there may only be a fixed cost or only be a variable cost (or a different type of variable cost, like an inverted gaussian). However, these models fail to qualitatively capture our results, and therefore, we opted not to include these models in our formal analysis. We also recognize that, at present, we only consider how target jump impacts learning from SPE, whereas target jumps may also impact learning from TE.

Alternatively, implicit recalibration may also be driven by both SPE and TE-based learning processes (Table 4: SPE + TE). The contribution of TE was assumed to vary with the distance between the cursor feedback and the new target position in a linear fashion. β_{TE} captures the slope of this function, and the $\theta_c - \theta_j$ term constrains implicit recalibration from TE to 0 when TE is 0 (i.e., when the target jumps onto the cursor feedback). This model assumes the net motor update (U_{Total}) to be the sum of a SPE-based learning process (U_{SPE}) and a TE based learning process (U_{TE}).

Table 4	SPE Only	SPE + TE
Invariant SPE	$U_{SPE} = U_{\theta_j = 0}$	$U_{Total} = U_{TE} + U_{SPE}$ $U_{TE} = \beta_{TE}(\theta_c - \theta_j)$
Rewarded SPE	$U_{SPE} = U_{\theta_j=0} - \gamma_r e^{-(\theta_c - \theta_j)^2 / 2\sigma_r^2}$	
Distracted SPE	$U_{SPE} = (U_{\theta_j=0} - C_J)\, e^{(\theta_j)^2 / 2\sigma_d^2}$	

Table 4: Summary of models. Parameters could either be free (red) or fixed (black, based on empirical data in Exp 4).

We evaluated the six models by simultaneously fitting group-averaged data for both ±3° and ±7° clamp groups in Exp 4. To quantify model performance, we calculated R^2 and AIC (Akaike Information Criterion) scores. The winning model was the model with the largest R^2 and the smallest AIC. In order to calculate confidence intervals for the parameter estimates, we applied standard bootstrapping techniques, constructing group-averaged hand angle data 1000 times by randomly resampling with replacement from the pool of participants within each group. We started with 10 different initial sets of parameter values and estimated parameter values that minimized the least squared error between the bootstrapped data and the model output.

Funding: RBI is funded by the NIH (NINDS: NS116883; NIDCD: DC0170941). HEK is funded by NIH K12 (HD055931). JST is funded by the PODSII scholarship from the Foundation for Physical Therapy Research.

Acknowledgements: We thank Joie Tang for her assistance with data collection (Exp 1a/b).

Disclosures and Competing interests: None to disclose.

References

Albert, S. T., Jang, J., Haith, A. M., Lerner, G., Della-Maggiore, V., Krakauer, J. W., & Shadmehr, R. (2020). Competition between parallel sensorimotor learning systems (p. 2020.12.01.406777). doi:10.1101/2020.12.01.406777

Anguera, J. A., Reuter-Lorenz, P. A., Willingham, D. T., & Seidler, R. D. (2010). Contributions of spatial working memory to visuomotor learning. *Journal of Cognitive Neuroscience*, *22*(9), 1917–1930.

Anwyl-Irvine, A. L., Dalmaijer, E. S., Hodges, N., & Evershed, J. (2020). *Online Timing Accuracy and Precision: A comparison of platforms, browsers, and participant's devices*. doi:10.31234/osf.io/jfeca

Benson, B. L., Anguera, J. A., & Seidler, R. D. (2011). A spatial explicit strategy reduces error but interferes with sensorimotor adaptation. *Journal of Neurophysiology*, *105*(6), 2843–2851.

Bridges, D., Pitiot, A., MacAskill, M. R., & Peirce, J. W. (2020). The timing mega-study: comparing a range of experiment generators, both lab-based and online. *PeerJ*, *8*, e9414.

Burge, J., Ernst, M. O., & Banks, M. S. (2008). The statistical determinants of adaptation rate in human reaching. *Journal of Vision*, *8*(4), 20.1-19.

Butcher, P. A., Ivry, R. B., Kuo, S.-H., Rydz, D., Krakauer, J. W., & Taylor, J. A. (2017). The cerebellum does more than sensory prediction error-based learning in sensorimotor adaptation tasks. *Journal of Neurophysiology*, *118*(3), 1622–1636.

Cameron, B. D., Franks, I. M., Inglis, J. T., & Chua, R. (2010a). Implicit motor learning from target error during explicit reach control. *Experimental Brain Research*, *206*(1), 99–104.

Cameron, B. D., Franks, I. M., Inglis, J. T., & Chua, R. (2010b). Reach adaptation to explicit vs. implicit target error. *Experimental Brain Research*, *203*(2), 367–380.

Cashaback, J. G. A., McGregor, H. R., Mohatarem, A., & Gribble, P. L. (2017). Dissociating error-based and reinforcement-based loss functions during sensorimotor learning. *PLoS Computational Biology*, *13*(7), e1005623.

Dang, K. V., Parvin, D. E., & Ivry, R. B. (2019). *Exploring Contextual Interference in Implicit and Explicit Motor Learning* (p. 644211). doi:10.1101/644211

Day, K. A., Roemmich, R. T., Taylor, J. A., & Bastian, A. J. (2016). Visuomotor Learning Generalizes Around the Intended Movement. *ENeuro*, *3*(2). doi:10.1523/ENEURO.0005-16.2016

Donchin, O., Francis, J. T., & Shadmehr, R. (2003). Quantifying generalization from trial-by-trial behavior of adaptive systems that learn with basis functions: theory and experiments in human motor control. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *23*(27), 9032–9045.

Dosher, B. A., Sperling, G., & Wurst, S. A. (1986). Tradeoffs between stereopsis and proximity luminance covariance as determinants of perceived 3D structure. *Vision Research*, *26*(6), 973–990.

Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, *415*(6870), 429–433.

Galea, J. M., Mallia, E., Rothwell, J., & Diedrichsen, J. (2015). The dissociable effects of punishment and reward on motor learning. *Nature Neuroscience*, *18*(4), 597–602.

Gonzalez Castro, L. N., Monsen, C. B., & Smith, M. A. (2011). The binding of learning to action in motor adaptation. *PLoS Computational Biology*, *7*(6), e1002052.

Hadjiosif, A. M., Criscimagna-Hemminger, S. E., Gibo, T. L., Okamura, A. M., Shadmehr, R., Bastian, A. J., & Smith, M. A. (2014). Cerebellar damage reduces the stability of motor memories. *Proceeding of the Translational and Computational Motor Control*.

Haith, A. M., Huberdeau, D. M., & Krakauer, J. W. (2015). The influence of movement preparation time on the expression of visuomotor learning and savings. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *35*(13), 5109–5117.

Hayashi, T., Kato, Y., & Nozaki, D. (2020). Divisively Normalized Integration of Multisensory Error Information Develops Motor Memories Specific to Vision and Proprioception. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *40*(7), 1560–1570.

Heffley, W., & Hull, C. (2019). Classical conditioning drives learned reward prediction signals in climbing fibers across the lateral cerebellum. *ELife*, *8*. doi:10.7554/eLife.46764

Hegele, M., & Heuer, H. (2010). Implicit and explicit components of dual adaptation to visuomotor rotations. *Consciousness and Cognition*, *19*(4), 906–917.

Hull, C. (2020). Prediction signals in the cerebellum: beyond supervised motor learning. *ELife*, *9*. doi:10.7554/eLife.54073

Izawa, J., Criscimagna-Hemminger, S. E., & Shadmehr, R. (2012). Cerebellar contributions to reach adaptation and learning sensory consequences of action. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *32*(12), 4230–4239.

Izawa, J., & Shadmehr, R. (2011). Learning from sensory and reward prediction errors during motor adaptation. *PLoS Computational Biology*, *7*(3), e1002012.

Kasuga, S., Hirashima, M., & Nozaki, D. (2013). Simultaneous processing of information on multiple errors in visuomotor learning. *PloS One*, *8*(8), e72741.

Kim, H. E., Avraham, G., & Ivry, R. B. (2020). The Psychology of Reaching: Action Selection, Movement Implementation, and Sensorimotor Learning. *Annual Review of Psychology*. doi:10.1146/annurev-psych-010419-051053

Kim, H. E., Morehead, J. R., Parvin, D. E., Moazzezi, R., & Ivry, R. B. (2018). Invariant errors reveal limitations in motor correction rather than constraints on error sensitivity. *Communications Biology*, *1*, 19.

Kim, H. E., Parvin, D. E., & Ivry, R. B. (2019). The influence of task outcome on implicit motor learning. *ELife*, *8*. doi:10.7554/eLife.39882

Kinchla, R. (1992). Attention. *Annual Review of Psychology*, *43*(1), 711–742.

Körding, Konrad P., Beierholm, U., Ma, W. J., Quartz, S., Tenenbaum, J. B., & Shams, L. (2007). Causal inference in multisensory perception. *PloS One*, *2*(9), e943.

Körding, Konrad P., & Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature*, *427*(6971), 244–247.

Körding, Konrad Paul, & Wolpert, D. M. (2004). The loss function of sensorimotor learning. *Proceedings of the National Academy of Sciences of the United States of America*, *101*(26), 9839–9842.

Krakauer, J., Hadjiosif, A. M., Xu, J., Wong, A. L., & Haith, A. M. (2019). Motor Learning. *Comprehensive Physiology*, *9*(2), 613–663.

Krakauer, J. W., Ghilardi, M. F., & Ghez, C. (1999). Independent learning of internal models for kinematic and dynamic control of reaching. *Nature Neuroscience*, *2*(11), 1026–1031.

Lee, K., Oh, Y., Izawa, J., & Schweighofer, N. (2018). Sensory prediction errors, not performance errors, update memories in visuomotor adaptation. *Scientific Reports*, *8*(1), 16483.

Leow, L.-A., Marinovic, W., de Rugy, A., & Carroll, T. J. (2018). Task errors contribute to implicit aftereffects in sensorimotor adaptation. *The European Journal of Neuroscience*, *48*(11), 3397–3409.

Leow, L.-A., Marinovic, W., de Rugy, A., & Carroll, T. J. (2020). Task Errors Drive Memories That Improve Sensorimotor Adaptation (pp. 3075–3088). doi:10.1523/JNEUROSCI.1506-19.2020

Lerner, G., Albert, S., Caffaro, P. A., Villalta, J. I., Jacobacci, F., Shadmehr, R., & Della-Maggiore, V. (2020). The Origins of Anterograde Interference in Vxisuomotor Adaptation. *Cerebral Cortex* . doi:10.1093/cercor/bhaa016

Magescas, F., & Prablanc, C. (2006). Automatic drive of limb motor plasticity. *Journal of Cognitive Neuroscience*, *18*(1), 75–83.

Mazzoni, P., & Krakauer, J. W. (2006). An implicit plan overrides an explicit strategy during visuomotor adaptation. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *26*(14), 3642–3645.

McDougle, S. D., Ivry, R. B., & Taylor, J. A. (2016). Taking Aim at the Cognitive Side of Learning in Sensorimotor Adaptation Tasks. *Trends in Cognitive Sciences*, *20*(7), 535–544.

Miyamoto, Y. R., Wang, S., & Smith, M. A. (2020). Implicit adaptation compensates for erratic explicit strategy in human motor learning. *Nature Neuroscience*, *23*(3), 443–455.

Morehead, J. R., Taylor, J. A., Parvin, D. E., & Ivry, R. B. (2017). Characteristics of Implicit Sensorimotor Adaptation Revealed by Task-irrelevant Clamped Feedback. *Journal of Cognitive Neuroscience*, *29*(6), 1061–1074.

Nikooyan, A. A., & Ahmed, A. A. (2015). Reward feedback accelerates motor learning. *Journal of Neurophysiology*, *113*(2), 633–646.

Ohmae, S., & Medina, J. F. (2015). Climbing fibers encode a temporal-difference prediction error during cerebellar learning in mice. *Nature Neuroscience*, *18*(12), 1798–1803.

Oldfield, R. C. (1971). The assessment and analysis of handedness: the Edinburgh inventory. *Neuropsychologia*, *9*(1), 97–113.

Oza, A., Kumar, A., & Mutha, P. K. (2020). Task Errors Do Not Induce Implicit Sensorimotor Learning (p. 2020.11.13.381285). doi:10.1101/2020.11.13.381285

Ranjan, T., & Smith, M. (2020, October 21). *Implicit motor adaptation is driven by motor performance prediction error rather than sensory prediction error*. Presented at the Motor Control and Motor Learning.

Ryan Morehead, J., & de Xivry, J.-J. O. (2021). A Synthesis of the Many Errors and Learning Processes of Visuomotor Adaptation. *BioRxiv*. doi:10.1101/2021.03.14.435278

Schlerf, J. E., Xu, J., Klemfuss, N. M., Griffiths, T. L., & Ivry, R. B. (2013). Individuals with cerebellar degeneration show similar adaptation deficits with large and small visuomotor errors. *Journal of Neurophysiology*, *109*(4), 1164–1173.

Sendhilnathan, N., Ipata, A. E., & Goldberg, M. E. (2020). Neural Correlates of Reinforcement Learning in Mid-lateral Cerebellum. *Neuron*. doi:10.1016/j.neuron.2019.12.032

Shadmehr, R., Smith, M. A., & Krakauer, J. W. (2010). Error correction, sensory prediction, and adaptation in motor control. *Annual Review of Neuroscience*, *33*, 89–108.

Shams, L., & Beierholm, U. R. (2010). Causal inference in perception. *Trends in Cognitive Sciences*, *14*(9), 425–432.

Shmuelof, L., Huang, V. S., Haith, A. M., Delnicki, R. J., Mazzoni, P., & Krakauer, J. W. (2012). Overcoming motor "forgetting" through reinforcement of learned actions. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *32*(42), 14617–14621.

Song, J.-H. (2019). The role of attention in motor control and learning. *Current Opinion in Psychology*, *29*, 261–265.

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

Taylor, J. A., & Ivry, R. B. (2014). Cerebellar and prefrontal cortex contributions to adaptation, strategies, and reinforcement learning. *Progress in Brain Research*, *210*, 217–253.

Taylor, J. A., Klemfuss, N. M., & Ivry, R. B. (2010). An explicit strategy prevails when the cerebellum fails to compute movement errors. *Cerebellum* , *9*(4), 580–586.

Taylor, J. A., Krakauer, J. W., & Ivry, R. B. (2014). Explicit and implicit contributions to learning in a sensorimotor adaptation task. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *34*(8), 3023–3032.

Taylor, J. A., & Thoroughman, K. A. (2007). Divided attention impairs human motor adaptation but not feedback control. *Journal of Neurophysiology*, *98*(1), 317–326.

Tsay, J. S., Avraham, G., Kim, H. E., Parvin, D. E., Wang, Z., & Ivry, R. B. (2020). The Effect of Visual Uncertainty on Implicit Motor Adaptation. *Journal of Neurophysiology*. doi:10.1152/jn.00493.2020

Tsay, J. S., Kim, H. E., Parvin, D. E., Stover, A. R., & Ivry, R. B. (2021). Individual differences in proprioception predict the extent of implicit sensorimotor adaptation. *Journal of Neurophysiology*. doi:10.1152/jn.00585.2020

Tsay, J. S., Lee, A. S., Avraham, G., Parvin, D. E., Ho, J., Boggess, M., … Ivry, R. (2020). OnPoint: A package for online experiments in motor control and motor learning. doi:10.31234/osf.io/hwmpy

Tsay, J. S., Lee, A. S., Ivry, R. B., & Avraham, G. (2021). Moving outside the lab: The viability of conducting sensorimotor learning studies online (p. 2021.01.30.181370). doi:10.1101/2021.01.30.181370

Tsay, J. S., Parvin, D. E., & Ivry, R. B. (2020). Continuous reports of sensed hand position during sensorimotor adaptation. *Journal of Neurophysiology*, *124*(4), 1122–1130.

Tseng, Y.-W., Diedrichsen, J., Krakauer, J. W., Shadmehr, R., & Bastian, A. J. (2007). Sensory prediction errors drive cerebellum-dependent adaptation of reaching. *Journal of Neurophysiology*, *98*(1), 54–62.

Wagner, M. J., Kim, T. H., Savall, J., Schnitzer, M. J., & Luo, L. (2017). Cerebellar granule cells encode the expectation of reward. *Nature*, *544*(7648), 96–100.

Wei, K., & Körding, K. (2009). Relevance of error: what drives motor adaptation? *Journal of Neurophysiology*, *101*(2), 655–664.

Wei, K., & Körding, K. (2010). Uncertainty of feedback and state estimation determines the speed of motor adaptation. *Frontiers in Computational Neuroscience*, *4*, 11.

Werner, S., van Aken, B. C., Hulst, T., Frens, M. A., van der Geest, J. N., Strüder, H. K., & Donchin, O. (2015). Awareness of sensorimotor adaptation to visual rotations of different size. *PloS One*, *10*(4), e0123321.

Wolpert, D. M., Miall, R. C., & Kawato, M. (1998). Internal models in the cerebellum. *Trends in Cognitive Sciences*, *2*(9), 338–347.