Brief Communications

Efficient Visual Recalibration from Either Visual or Haptic Feedback: The Importance of Being Wrong

Wendy J. Adams, Iona S. Kerrigan, and Erich W. Graf
School of Psychology, University of Southampton, Southampton SO17 1BJ, United Kingdom

The human visual system adapts to the changing statistics of its environment. For example, the light-from-above prior, an assumption that aids the interpretation of ambiguous shading information, can be modified by haptic (touch) feedback. Here we investigate the mechanisms that drive this adaptive learning. In particular, we ask whether visual information can be as effective as haptics in driving visual recalibration and whether increased information (feedback from multiple modalities) induces faster learning.

During several hours' training, feedback encouraged observers to modify their existing light-from-above assumption. Feedback was one of the following: (1) haptic only, (2) haptic and stereoscopic (providing binocular shape information), or (3) stereoscopic only. Haptic-only feedback resulted in substantial learning; the perceived shape of shaded objects was modified in accordance with observers' new light priors. However, the addition of continuous visual feedback (condition 2) substantially reduced learning. When visual-only feedback was provided intermittently (condition 3), mimicking the time course of the haptic feedback of conditions 1 and 2, substantial learning returned.

The intermittent nature of conflict information, or feedback, appears critical for learning. It causes an initial, erroneous percept to be corrected. Contrary to previous proposals, we found no particular advantage for cross-modal feedback. Instead, we suggest that an "oops" factor drives efficient learning; recalibration is prioritized when a mismatch exists between sequential representations of an object property. This "oops" factor appears important both across and within sensory modalities, suggesting a general principle for perceptual learning and recalibration.

Introduction

The visual system uses statistical regularities of the environment to bias perception toward the most frequent or likely interpretations of visual input. For example, the "light-from-above" assumption aids the recovery of object shape from otherwise ambiguous shading information (e.g., Metzger, 1936; Kleffner and Ramachandran, 1992) (also see Fig. 1*A*). We have shown (Adams et al., 2004; Champion and Adams, 2007) that this light prior is not hard-wired but adapts in response to visual and haptic (touch) interactions with the environment. However, little is known about the process by which visual priors are acquired or modified. Here we explore visual recalibration to ask a number of questions regarding the nature and specificity of learning with different types of feedback.

It has been proposed that haptics is particularly important in calibrating vision (e.g., Berkeley, 1709) with visual–haptic learning studies often regarding haptic information as "feedback" for the visual system (e.g., Adams et al., 2004; Knill, 2007). Gori et al. (2008) suggest that the developmentally late integration of visual and haptic shape cues (at 8–10 years) enables the different modalities to efficiently recalibrate each other during fast physical growth. Furthermore, in adults, perceptual information may be more closely coupled within than across modalities, such that access to individual visual cue estimates is lost, with access retained only to amalgamated information from vision or from haptics (Hillis et al., 2002). Thus, haptic information might have a privileged role in recalibrating visually based information, beyond that predicted solely from its reliability. However, whether such cross-modal conflict affords more efficient learning is an open question (Jacobs and Shams, 2010). We ask whether visual cues are recalibrated more readily by haptic feedback than by conflicting information from other visual cues.

An alternative proposal is that, given conflicting information, the perceptual system treats any reliable cue as feedback reflecting the true state of the world. More reliable cues, regardless of modality, would thus be used to recalibrate less reliable cues (e.g., Burge et al., 2010). If learning is modulated by relative reliability, then faster learning might occur when the amount and reliability of feedback is increased. In the present study, we compare recalibration of the light prior when either haptics alone or haptics with an additional visual cue (binocular disparity) conflict with shading.

Finally, we investigate whether learned changes in a visual prior (the light prior) are specific to the experimental context, or more generalized. In the latter case, learned effects would quickly diminish when observers return to normal surroundings where illumination is (presumably) most often overhead.

Received May 30, 2010; revised Sept. 2, 2010; accepted Sept. 9, 2010.

This work was funded by Engineering and Physical Sciences Research Council Grant EP/D039916/1 and British Academy Grant SG-49590 to W.J.A., the Royal Society (E.W.G.), and an Economic and Social Research Council studentship (I.S.K.). We thank Samira Bouzit for assistance with data collection, Richard Gramzow for statistical advice, Ed Redhead for advice on the learning literature, and James Elder for modeling discussions.

Correspondence should be addressed to Wendy J. Adams, School of Psychology, University of Southampton, Southampton SO17 1BJ, UK. E-mail: w.adams@soton.ac.uk.

DOI:10.1523/JNEUROSCI.2749-10.2010

0270-6474/10/3014745-05$15.00/0

In summary, we ask the following questions: First, is visual feedback alone effective (and as effective as haptics) in guiding visual recalibration? Second, is learning increased when feedback is more reliable? Third, is the recalibration of visual priors context specific?

Materials and Methods

Observers viewed computer-generated stimuli via a front-silvered mirror while a PHANToM force-feedback device (SensAble Technologies) provided haptic information (Fig. 1*A*). This setup allows simultaneous presentation of visual and haptic information in the same perceived location; the distance from the observers' eyes to the center of both the visual and haptic scenes was 56 cm. Individually molded bite bars ensured head alignment across sessions. Visual scenes were comprised of four shaded disks of diameter 5.6°, each offset from the screen's midpoint by 5.3°. Each disk was consistent with a hemisphere squashed in depth by a factor of 0.5, and illuminated by a distant light source (see Fig. 1*A*). The slant of the light source (the angle between the lighting vector and the screen normal) was held constant at 55°, while tilt (the angle between a vertical axis in the screen's plane and the projection of the lighting vector) varied across trials; changing tilt was equivalent to rotating the shaded disks within the image plane. Within each scene, either two or three objects had identical shading and the remainder were rotated by 180°. Thus, images were consistent with a scene comprising four objects of which one, two, or three were convex (the remainder being concave), all illuminated by a single light source. In general, within a scene, objects with opposite shading gradients were indeed perceived as having opposite curvatures (consistent with a preference for a single light source). However, the individual observer's light prior determined which objects were perceived as convex.

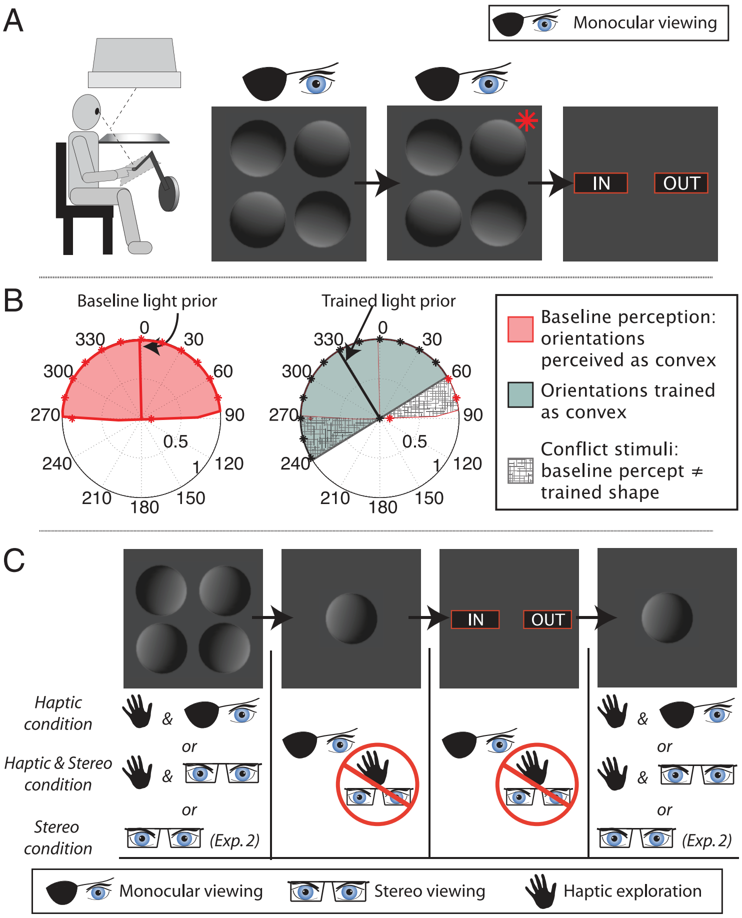

Figure 1. Stimuli and methods. ***A***, The visual– haptic setup and an example baseline trial. The upper-right and lower-left objects in this scene are generally perceived as convex and the others as concave, consistent with overhead lighting. The red asterisk signaled which object should be judged (via buttons labeled "in" and "out"). ***B***, Left plot, Proportion convex judgments as a function of stimulus orientation for one observer: baseline data (red asterisks) and model fit (red line). Right plot, Orientations trained as convex (shaded gray region) and data from haptic trials (black asterisks). ***C***, Schematic of the available shape information (monocular, binocular and haptic) at each stage of a training trial, for each condition.

Visual test trials (including baseline trials). On each test trial, the observer monocularly viewed a scene of four objects, without haptic information (haptic access was blocked by a virtual wall). Halfway through the 1.5 s presentation, an asterisk appeared, indicating the target whose shape should be judged (convex vs concave). A block of 192 visual-only judgments (24 equally spaced target orientations × 8 repetitions) lasting 8–10 min was used to estimate an observer's light prior at baseline and during subsequent test phases. An example of one subject's data can be seen in Figure 1*B*. The peak of "convex" responses occurs for targets that are bright near the top (0°), consistent with overhead lighting. This peak can be interpreted as the mean of the observer's light prior.

Visual–haptic training trials (experiment 1). After completing the baseline visual-only trials, observers underwent training in either a haptic feedback condition (11 observers, 4 male) or a haptic and stereo feedback condition (11 observers, 4 male). In the "Haptic" condition, viewing was monocular; within the visual domain, shape was defined by shading alone. The haptic scene matched the dimensions of the visual scene; an observer could move their finger (fixed to the PHANToM's gimbaled stylus) to "feel" the four objects that they were viewing. However, the haptically defined shape of the objects (convex or concave) was manipulated on a subject-by-subject basis: haptic shape was consistent with a lighting direction within a range that was shifted by either +30° or −30° relative to the observer's baseline light prior distribution (see Fig. 1*B*). The feedback (in all training conditions) disambiguated both lighting direction and object shape; some objects that were perceived as convex in the baseline trials now felt concave and vice versa (hatched "conflict" orientations in Fig. 1*B*).

In the "Haptic&Stereo" condition, training stimuli were viewed via shutter glasses such that each eye received a slightly different view of the scene. As shading gradients provide weak stereoscopic shape information, a vertical contour running across the surface of each object was added (not shown). Disparity information was always consistent with haptic shape (described above) such that in this Haptic&Stereo condition, both haptic and disparity information indicated that the observer should recalibrate their light prior. Test trials were viewed monocularly, regardless of training condition. Thus, in both the Haptic&Stereo and Stereo conditions (experiment 2), only one eye's view was visually rendered during test trials with the other eye's view blocked via the shutter glasses. Vertical contours were rendered only during training trials, preventing observers from using contour curvature as a cue to convexity during test trials.

Training trials (shown schematically in Fig. 1*C*) were procedurally similar to those used previously (e.g., Adams et al., 2004). Each training trial started with a visual–haptic exploration phase. After a minimum of 7 s, and having "touched" all four objects at least once, the observer pressed a virtual button labeled "DONE." After a brief delay, one of the four objects from the previous scene appeared monocularly at the center of the display for 1.25 s. Based on visual appearance alone, the observer judged its shape (responding either convex or concave via virtual buttons). The observer then touched the object while viewing it either monocularly (Haptic condition) or binocularly (Haptic&Stereo condition). Each training block consisted of 112 trials (24 stimulus orientations; 4

repetitions for "non-conflict," 8 repetitions for "conflict" orientations) and lasted 20–30 min.

Visual training trials (experiment 2). In the Stereo feedback condition (experiment 2), no haptic information was provided. Instead, binocular disparity information was provided (see Fig. 1*C*), with one key difference from that used in the Haptic&Stereo training. Rather than presenting disparity information throughout the training trial, its temporal availability mimicked the intermittent presentation of the haptic information in previous conditions. To achieve this, recorded finger movements for an observer in the Haptic condition were used to generate visual feedback for Stereo training trials. On each Stereo feedback trial, rather than touching the scene, the observer was required to visually track a green dot moving over the scene. The motion of this dot was determined by replaying the previously recorded finger movements for the Haptic condition, for the same trial. When the green dot was in contact with one of the four objects, both object and dot were rendered stereoscopically; other objects in the scene remained monocular. In this way, observers gained reliable depth information about the scene in an intermittent manner, with exactly the same time course as observers gained reliable haptic information in previous conditions. Following visual exploration of the scene, an observer judged a single object from the scene, presented monocularly. This was followed by binocular presentation of the same object, such that, as in previous conditions, its shape became unambiguous. Finally, to confirm that observers tracked the dot as required, and thus received the stereo information, they were asked how many of the presented objects were stereoscopically convex (all observers performed well, μ = 93% correct; chance = 33%). Trial order was randomized across the 10 observers (4 male), none of whom had participated in experiment 1.

Test-training schedule. After an initial block of baseline test trials and a short break, observers completed a 3 h train–test session of 8 blocks (4 training blocks, alternated with 4 test blocks). Subsets of observers (6 for the Haptic&Stereo, 8 for the Haptic, and 8 for the Stereo condition) were tested again the following morning before another train–test session. This second day's session lasted ~1.5 h and consisted of 4 blocks (2 training blocks alternated with 2 test blocks). In experiment 1, to investigate decay of the training effects, the 14 observers who completed 2 d of training also completed several more test blocks over the following week. The local ethics committee approved the study and all observers gave informed written consent.

Results

Experiment 1

Pairing ambiguous shading information with unambiguous haptic information is sufficient to drive recalibration of shape-from-shading, by modifying observers' light priors (Adams et al., 2004). Experiment 1 extended this work to ask (1) whether the addition of a second unambiguous cue (binocular disparity) impacts learning rate and (2) whether the resultant effects diminish over time following training.

Data from two example observers and the normalized group data are shown in Figure 2. To quantify the effect of training, each block of training or test trials from each observer was summarized by a single value: the mean of the light prior. Light priors were determined via a simple Bayesian model that provided a good fit to the data (average $r^2 = 0.92$, across participants and conditions), as follows.

The model assumed that a global light source illuminates convex and concave objects to generate two shading orientations (θ_1 and θ_2). The prior for lighting tilt (ϕ) was described by a von Mises distribution whose mean (μ) and concentration (κ) were adjusted to provide the best fit (highest likelihood) to the data; for each condition, responses were predicted by the relative probabilities of convexity and concavity for the judged object (τ_1):

$$\text{proportion\_judged\_convex} = \frac{p(\tau_1 = \text{convex}|\theta_1,\theta_2)}{p(\tau_1 = \text{convex}|\theta_1,\theta_2) + p(\tau_1 = \text{concave}|\theta_1,\theta_2)}.$$

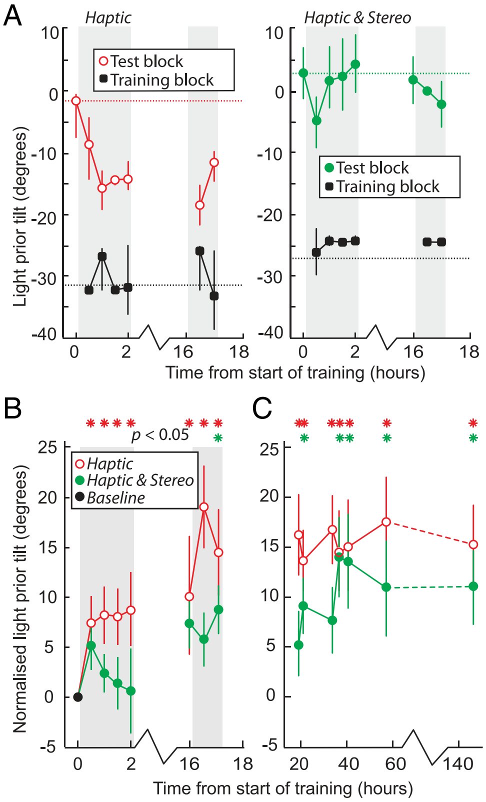

Figure 2. ***A***, Light priors for individual observers for the Haptic condition (left) and Haptic&Stereo condition (right). Circles represent test blocks, while black squares give shape responses on training trials. Error bars show 95% confidence intervals from bootstrapping. Dashed horizontal lines give the observer's baseline light prior (red/green) and the trained light prior direction (black). Shaded gray strips indicate train–test sessions. ***B***, ***C***, Group results for the two feedback conditions during the 2 day training period (***B***) and after training (***C***). Asterisks indicate test blocks that differ significantly from baseline ($p < 0.05$). Error bars give ±1 SE across observers.

Thus, rather than picking the most likely shape (i.e., the maximum a posteriori decision rule), the observer sampled the posterior distribution.

From Bayes' rule, $p(\tau_1|\theta_1,\theta_2) \propto p(\theta_1,\theta_2|\tau_1)p(\tau_1)$. Ignoring measurement noise, shading gradients must be aligned with the lighting direction. Thus the likelihood $p(\theta_1,\theta_2|\tau_1)$ simplifies to $p(\theta_1,\theta_2|\tau_1 = \text{convex}) = p(\phi = \theta_1)$ and $p(\theta_1,\theta_2|\tau_1 = \text{concave}) = p(\phi = \theta_1 + \pi)$, and the above ratio becomes the following:

$$\text{proportion\_judged\_convex} = \frac{p(\phi = \theta_1)}{p(\phi = \theta_1) + p(\phi = \theta_1 + \pi)}.$$

To allow comparisons across observers, each individual's data were normalized by subtracting his or her baseline light prior. For those who trained with a −30° shift, data were flipped around zero.

The effect of training on shape judgments was assessed using a linear mixed model analysis (incorporating light priors from test blocks in the two train–test sessions; i.e., data from Fig. 2*B*, left panel). The Haptic group displayed significantly more learning than the Haptic&Stereo group ($t_{(19)} = 3.92$, $p < 0.01$). For the Haptic group (red circles), significant training effects were seen immediately in the first day's train–test session (regression coefficients differ from zero, $p < 0.05$). These effects were increased in the second day's training. By contrast, the Haptic&Stereo group (filled green circles) showed no significant training effects on day 1, with only the last test block on day 2 differing significantly from

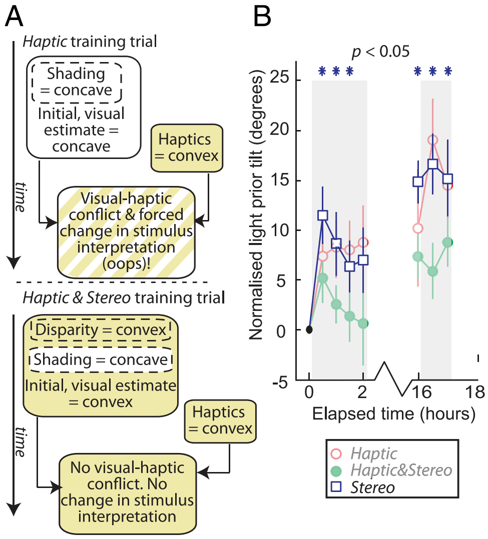

Figure 3. ***A***, The progression of cue availability and the resultant percepts within training trials on a conflict trial where the stimulus is initially perceived as concave, but feedback indicates that it is convex (upper panel: Haptic condition, lower panel: Haptic&Stereo condition). ***B***, Data from all three training conditions. Asterisks indicate test blocks where significant learning was observed in the Stereo feedback condition.

baseline. Both groups retained significant learning effects during the week after training had ceased (one-sample, two-tailed *t* tests against zero, as shown by asterisks on Fig. 2*C*; two points for the Haptic&Stereo group dip below significance).

During training, shape judgments were determined by feedback; fitted light priors did not differ significantly from the trained direction (mean normalized light priors across observers, conditions and training blocks: 29.97°; one-sample, two-tailed *t* tests against 30° for each training block: all $p > 0.05$).

The addition of reliable and unambiguous binocular disparity information, in experiment 1, reduced the rate of learning, contradicting the idea that simply presenting more information should increase learning. However, as well as increasing the amount and reliability of available feedback, the addition of disparity information in the Haptic&Stereo condition changed the nature of the feedback in two potentially important ways: The cross-modal conflict between visual and haptic information was eliminated, and observers were no longer forced to reinterpret the visual stimulus within "conflict" training trials (see Fig. 3*A*). Thus, two plausible explanations exist for the reduced learning: (1) Visual–haptic conflict is more efficacious than visual–visual conflict in driving recalibration. (2) A forced reinterpretation of the visual signal over time following an initial, apparently erroneous interpretation is important in driving recalibration.

When considering the first possibility, we note that in Haptic&Stereo training trials, disparity information immediately disambiguated the shading information. Thus, no conflict remained between estimates of haptic shape and visual shape (defined by both shading, in conjunction with the observer's light prior, and disparity). In contrast, during Haptic training trials, there was a conflict between visually defined shape (from shading) and haptically defined shape. If a difference between shape estimates from these two modalities were a particularly strong cue to recalibrate, then eliminating this cross-modal error signal would result in reduced learning, as observed in the Haptic&Stereo condition.

Importantly, the addition of disparity information in the Haptic&Stereo condition also eliminated the "oops" factor; observers were no longer forced to change their interpretation of the stimulus within training trials. Disparity information was available from the start of the trial, immediately disambiguating the shading information. When observers subsequently touched the stimulus, there was no change in percept: haptic and disparity information were always consistent. (A parallel can be drawn with the phenomenon of "blocking" in associative learning, as described by Kamin (1969), where inhibition of learning is caused by the presence of a concurrent, established cue.) By contrast, in the Haptic condition, the initial stimulus interpretation, formed from visual (shading) information, changed on haptic contact. It may be that this "oops" factor resulted in more learning for the Haptic condition. This second explanation does not confer cross-modal conflict any greater importance for learning than within-modal conflict (e.g., between disparity and shading). Instead it suggests that the visual system prioritizes recalibration when erroneous interpretations are made.

Experiment 2

To distinguish between the two possible explanations for the reduced learning in the Haptic&Stereo condition, experiment 2 investigated the efficacy of visual-only feedback by mimicking the temporal availability of the previous experiment's haptic information. In this way, we reintroduced the forced reinterpretation of the stimulus over time, or "oops" factor, but without any cross-modal error.

Figure 3*B* shows data from the Stereo feedback condition, alongside experiment 1 data, for comparison. A linear mixed model analysis confirmed a significant effect of training in the Stereo condition ($t_{(9)} = 4.18, p < 0.01$). Importantly, the Stereo feedback was as effective as the Haptic feedback in recalibrating observers' light priors: the Stereo and Haptic conditions differed significantly from the Haptic&Stereo condition ($t_{(28)} = -2.11$, $p < 0.05$ and $t_{(28)} = -2.531, p < 0.05$, respectively), but not from each other ($t_{(28)} = 0.07$, $p > 0.05$). As in the Haptic condition, observers in the Stereo condition showed significant training effects from the very first test session [although for reasons that are unclear, the training effects became only marginally significant ($p = 0.05$), in the final test session of day 1]. Significant learning effects were also seen the following day.

For all three training conditions, there is some evidence of overnight consolidation of learning (although this is not statistically significant)—training effects are stronger at the start of day 2 (at 16 h) than at the end of day 1 (at 2 h) without extra training during the intervening period.

Discussion

Two striking conclusions arose from experiment 1. First, learned changes in the interpretation of shading were still apparent several days after training, despite exposure to a range of lighting conditions in the observers' day-to-day activities, presumably with an average direction that was roughly overhead. This is consistent with context-dependent learning, whereby the learned light priors were retained in the laboratory setting, separate from observers' experience in other contexts. [Additional data from author W.J.A. (not shown) demonstrated that, for this observer, learning transferred across different experimental laboratory setups with various stimulus sizes and viewing distances.]

Resistance to learning generalization in humans has been reported using more traditional associative paradigms (e.g., Shanks et al., 1998), suggesting a context specificity that cannot be accounted for by classic, elemental associative learning models (e.g., Rescorla and Wagner, 1972). Interestingly, although our observed effects are context-dependent, a previous study (Adams et al., 2004) suggests that they are not stimulus dependent. In that study, training also affected lightness judgments of a different, untrained stimulus.

Second, in experiment 1 we found that a combination of haptic and stereo feedback reduced the amount of learning relative to haptic feedback alone. Two hypotheses were proposed: Either cross-modal conflict is important in driving recalibration, or the delayed and intermittent nature of feedback is important—increasing recalibration by forcing observers to reinterpret the shaded stimulus. Our second experiment supports the latter hypothesis. Visually driven, visual recalibration occurred readily (i.e., binocular disparity-driven recalibration of the shape-from-shading process) when feedback was intermittent. The perceptual system appears to prioritize recalibration when errors of judgment are made, thus maximizing potential gains in performance. This description contrasts with the notion that discrepancies between cues (and not perceptual or behavioral errors) are the sole predictors of recalibration (e.g., Epstein, 1975). According to Epstein's account, recalibration would readily occur when unreliable (and erroneous) information is completely and immediately dominated by a veridical cue, despite the absence of perceptual errors, or conscious registration of conflict, such as in our Haptic&Stereo condition. Epstein (1975) suggests that "inhibition is the precondition of recalibration," whereas we propose that recalibration is expedited when an initial visual percept is overturned by new evidence.

Visual feedback has been shown to drive both visual–motor and oculomotor recalibration. For example, visual feedback is sufficient to recalibrate visually guided throwing (Martin et al., 1996), while optic flow can recalibrate self-motion toward a target (Bruggeman et al., 2007). Prisms that alter the relationship between convergence and distance cause perceptual aftereffects (e.g., Wallach and Frey, 1972; Wallach et al., 1972), primarily via changes in oculomotor muscle tone (Ebenholtz and Wolfson, 1975).

Recalibration of a purely visual cue (binocular disparity) occurs when observers wear a horizontal magnifier over one eye for an extended period (e.g., Epstein and Morgan, 1970; Epstein, 1971; Adams et al., 2001). In these studies, the full range of visual and visual–motor signals were available to drive learning. However, there is a relative paucity of studies that investigate the effectiveness of visual information alone as the error signal to recalibrate another visual process. Wallach et al. (1963) and Epstein and Morgan-Paap (1974) both reported disparity recalibration resulting from brief viewing of stimuli in which disparity conflicted with other visual cues. However, these observed perceptual aftereffects are largely attributable to changes in tonic vergence (Fisher and Ebenholtz, 1986; Fisher and Ciuffreda, 1990) and "normalization," or "satiation" slant aftereffects (e.g., Köhler and Emery, 1947). The present study provides a clear demonstration of visual recalibration, driven by another visual cue, and moreover offers a description of the conditions under which this recalibration is optimized.

In summary, we provide evidence that the visual system is efficiently recalibrated by either visual or haptic feedback, when that feedback reveals errors in a perceptual estimate. Additional feedback can reduce learning if it entirely dominates erroneous or conflicting information, thus preventing perceptual errors and masking the signal to recalibrate. Should we dismiss the notion that haptics play an important role in recalibrating vision? Certainly not. Although we found that within our paradigm a cross-modal error is not important per se, the temporal dynamics of haptic information as we typically receive it make it particularly useful for calibrating visual information. For example, we usually pick up or touch an object after locating it visually. This delay in haptic relative to visual information means that an "oops" signal, or forced reinterpretation, occurs more often between visual and haptic cues than between multiple visual cues. The present study demonstrates that a similar "oops" signal between two visual cues also drives recalibration. This can be just as efficient as cross-modal feedback when the two visual cues are temporally offset in a way that mimics the intermittent way in which we normally receive haptic information.

References

Adams WJ, Banks MS, van Ee R (2001) Adaptation to three-dimensional distortions in human vision. Nat Neurosci 4:1063–1064.

Adams WJ, Graf EW, Ernst MO (2004) Experience can change the 'light-from-above' prior. Nat Neurosci 7:1057–1058.

Berkeley G (1709) An essay towards a new theory of vision. Dublin: Pepyat. Reprinted in: Works on vision (Turbayne CM, ed), pp 19–97. Indianapolis: Bobbs-Merrill, 1963.

Bruggeman H, Zosh W, Warren WH (2007) Optic flow drives human visuo-locomotor adaptation. Curr Biol 17:2035–2040.

Burge J, Girshick AR, Banks MS (2010) Visual-haptic adaptation is determined by relative reliability. J Neurosci 30:7714–7721.

Champion RA, Adams WJ (2007) Modification of the convexity prior but not the light-from-above prior in visual search with shaded objects. J Vis 7:10.1–10.10.

Ebenholtz SM, Wolfson DM (1975) Perceptual aftereffects of sustained convergence. Percept Psychophys 17:485–491.

Epstein W (1971) Adaptation to uniocular image magnification after varying preadaptation activities. Am J Psychol 84:66–74.

Epstein W (1975) Recalibration by pairing: a process of perceptual learning. Perception 4:59–72.

Epstein W, Morgan CL (1970) Adaptation to uniocular image magnification: modification of the disparity-depth relationship. Am J Psychol 83:322–329.

Epstein W, Morgan-Paap CL (1974) The effect of level of depth processing and degree of informational discrepancy on adaptation to uniocular image magnification. J Exp Psychol 102:585–594.

Fisher SK, Ciuffreda KJ (1990) Adaptation to optically-increased interocular separation under naturalistic viewing conditions. Perception 19:171–180.

Fisher SK, Ebenholtz SM (1986) Does perceptual adaptation to telestereoscopically enhanced depth depend on the recalibration of binocular disparity? Percept Psychophys 40:101–109.

Gori M, Del Viva M, Sandini G, Burr DC (2008) Young children do not integrate visual and haptic information. Curr Biol 18:694–698.

Hillis JM, Ernst MO, Banks MS, Landy MS (2002) Combining sensory information: mandatory fusion within, but not between, senses. Science 298:1627–1630.

Jacobs RA, Shams L (2010) Visual learning in multisensory environments. Topics Cogn Sci 2:217–225.

Kamin LJ (1969) Predictability, surprise, attention, and conditioning. In: Punishment and aversive behavior (Campbell BA, Church RM, eds), pp 242–259. New York: Appleton-Century-Crofts.

Kleffner DA, Ramachandran VS (1992) On the perception of shape from shading. Percept Psychophys 52:18–36.

Knill D (2007) Learning Bayesian priors for depth perception. J Vis 7:1–20.

Köhler W, Emery DA (1947) Figural after-effects in the third dimension of visual space. Am J Psychol 60:159–201.

Martin TA, Keating JG, Goodkin HP, Bastian AJ, Thach WT (1996) Throwing while looking through prisms: II. Specificity and storage of multiple gaze-throw calibrations. Brain 119:1199–1211.

Metzger W (1936) The laws of seeing. Reprint (Spillmann L, translator). Cambridge, MA: MIT Press, 2006.

Rescorla RA, Wagner AR (1972) A theory of Pavlovian conditioning: variations in the effectiveness of reinforcement and nonreinforcement. In: Classical conditioning II: current research and theory (Black AH, Prokasy WF, eds) pp 64–99. New York: Appleton-Century-Crofts.

Shanks DR, Darby RJ, Charles D (1998) Resistance to interference in human associative learning: evidence of configural processing. J Exp Psychol Anim Behav Process 24:136–150.

Wallach H, Frey K (1972) Adaptation in distance perception based on oculomotor cues. Percept Psychophys 11:77–83.

Wallach H, Moore ME, Davidson L (1963) Modification of stereoscopic depth-perception. Am J Psychol 76:191–204.

Wallach H, Frey K, Bode E (1972) The nature of adaptation in distance perception based on oculomotor cues. Percept Psychophys 11:110–116.