Sensorimotor adaptation compensates for distortions of 3D shape information

Evan Cesanek[1], Jordan A. Taylor[2], and Fulvio Domini[1]

[1] Department of Cognitive, Linguistic, & Psychological Sciences
Brown University
Providence, RI, 02906

[2] Department of Psychology
Princeton University
Princeton, NJ, 08544

Abstract

Visual perception often fails to recover the veridical 3D shape of objects in the environment due to ambiguity and variability in the available depth cues. However, we rely heavily on 3D shape estimates when planning movements, for example reaching to pick up an object from a slanted surface. Given the wide variety of distortions that can affect 3D perception, how do our actions remain accurate across different environments? One hypothesis is that the visuomotor system performs selective filtering of 3D information to minimize distortions. Indeed, some studies have found that actions appear to preferentially target stereo information when it is put in conflict with texture information. However, since these studies analyze averages over multiple trials, this apparent preference could be produced by sensorimotor adaptation. In Experiment 1, we create a set of cue-conflict stimuli where one available depth cue is affected by a constant bias. Sensory feedback rapidly aligns the motor output with physical reality in just a few trials, which can make it seem as if action planning selectively relies on the reinforced cue—yet no change in the relative influences of the cues is necessary to eliminate the constant errors. In contrast, when one depth cue becomes less correlated with physical reality, variable movement errors will occur, causing canonical adaptation to fail as the opposite error corrections cancel out. As a result, canonical adaptation cannot explain the preference for stereo found in studies with variable errors. However, Experiment 2 shows that persistent errors can produce a novel form of adaptation that gradually reduces the relative influence of an unreliable depth cue. These findings show that grasp control processes are continuously modified based on sensory feedback to compensate for both biases and noise in 3D visual processing, rather than having a hardwired preference for one type of depth information.

1. Introduction

Every day, we use our hands to perform highly skilled interactions with a wide variety of physical objects. At a glance, we can recognize a target object's 3D structure from available sources of visual information (depth cues) and use this information to guide our movements, for example shaping and orienting the grip in a suitable way to grasp the object. However, visual perception of 3D object properties is known to be systematically distorted, even when viewing real, fully illuminated objects from within reaching distance (Heine, 1900; Norman, Todd, Perotti, & Tittle, 1996; Norman, Lappin, & Norman, 2000). In this study, we investigate how the motor system manages to produce accurate movements despite the serious challenge posed by distortions of 3D shape perception.

Perhaps the most well-known explanation of this paradox is that perception and action are supported by independent visual systems (Goodale & Milner, 1992), with the purported "vision-for-action" system capable of processing 3D information in a way that somehow recovers accurate metric estimates of physical shape. A major cornerstone of support for this theory is the literature on the effects of visual illusions on motor behavior. On average, the motor responses in such studies tend to be quite accurate, apparently resisting the perceptual distortions induced by the illusions. In one representative example, Bruggeman, Yonas, and Konczak (2007) presented participants with the Ames window illusion, comparing perceptual judgments of surface slant with bimanual pointing movements targeting the left and right edges of the display. When viewing an Ames window, texture cues to the slant of the depicted window are put in conflict with stereo cues specifying the actual flat surface of the display; in general, the resulting perceived slant lies somewhere in between. Bruggeman *et al*. (2007) found that the average separation between the left and right hands at the end of the pointing movements was less than the reported perceived depth of the display, suggesting a preference for the flat surface specified by stereo.

However, the availability of sensory feedback in this study (and others like it) suggests another, more parsimonious account of results. By this account, actions are guided by the same visual processing of 3D information as perception, but dissociations can arise rapidly as a result of short-term motor learning based on sensory feedback. In particular, if the initial reaches were aimed at the edges of the perceived surface, one hand would have stopped short of its intended alignment with the edge of the display, while the other hand would have passed beyond the display's other edge. At the end of this movement, the observed relative disparities between the fingertips and the display edges would provide a visual error signal that could be used to adjust planned reach distances on the next trial. These trial-by-trial corrections could occur without modifying the relative influences of stereo and texture information—all that is required is a constant reduction of the separation between the hands at the end of the movement. In other words, action planning could always rely on the visually perceived slant of the surface, yet the *average* separation of the hands across many trials would nevertheless show a preference for stereo information, simply because it was reinforced by sensory feedback.

The canonical form of sensorimotor adaptation described above, where motor outputs are shifted in one direction, is a powerful way to bring actions into alignment with the physical world, in just a handful of movements. However, it can only succeed when movement errors are relatively constant, as in the situation described above, where one of the available sources of 3D information (texture) was affected by a constant bias while the other (stereo) was consistently accurate. In contrast, if movement errors are highly variable, which is exactly what occurs when one of the available sources of 3D information becomes noisy, this form of adaptation will fail. As a result, canonical sensorimotor adaptation cannot explain the preference for stereo

information found in other studies, where *opposing* movement errors were experienced from one trial to the next.

For example, Knill (2005) rendered slanted surfaces using conflicting stereo and texture information in a virtual reality rig with haptic feedback from a real physical surface, and asked participants to place a handheld cylinder so that its bottom would be parallel with the surface at contact. On some trials, one of the cues was perturbed by ±4° from the underlying physical slant, causing the perceived slant to be biased in that direction. Thus, unlike the study of Bruggeman *et al*. (2007), which involved a constant difference between stereo and texture slant on each trial, and therefore a relatively constant perceptual bias, the perceptual bias in Knill's (2005) study changed sign randomly throughout the block. Yet the results were similar to those of Bruggeman *et al*. (2007): the average cylinder orientation at first contact was found to be slightly closer to the stereo slant. This finding has been interpreted as evidence that "vision-for-action" selectively relies on stereo information (*cf.* Goodale, 2011), enabling the motor system to avoid making mistakes based on faulty processing of other, typically less reliable cues. This idea is motivated by the observation that binocular disparities are a straightforward function of an object's depth, the viewing distance, and the separation between the two eyes, thus requiring only an estimate of viewing distance from ocular convergence to recover a metric depth estimate.

However, we will show that it is also possible to explain the stereo preference in Knill's (2005) study without positing a separate "vision-for-action" system, but simply as a result of a more complex form of sensorimotor adaptation, which we have termed *cue-selective adaptation*. We hypothesized that when unidirectional shifts of the motor output do not suffice to eliminate errors, the relative influences of stereo and texture information can be changed based on error feedback, a previously unexplored idea in sensorimotor adaptation research. In theory, the availability of stereoscopic visual feedback in Knill's (2005) experiment could have driven gradual reinforcement of the stereo cue in movement planning: when the planned cylinder orientation did not match the stereo slant, the resulting movement would have *looked* incorrect, thus producing an error signal (similar to the misalignment of the fingertips with the physical edges of the Ames window in Bruggeman *et al*., 2007). Over time, these errors could drive movement planning to increase its reliance on stereo information in order to eliminate these errors in visual feedback.

In summary, we propose two, non-mutually exclusive ways in which sensorimotor adaptation could overcome errors produced by distorted 3D perception, while maintaining the parsimonious assumption that actions depend on the same combination of stereo and texture information as perception. In developing models of these two mechanisms, we looked to the adaptation literature, which has produced an elegant theory explaining how observed errors drive changes in motor output. First, we apply a standard computational model of adaptation to account for cases where one depth cue is systematically biased to appear shallower than the physical slant; this reflects the design of Bruggeman *et al*. (2007), where a constant bias was present in one of two depth cues while the other was reinforced by sensory feedback. Second, we explore a more complex scenario in which the slant estimates from one cue are sometimes larger and sometimes smaller than the true surface slant (*i.e.*, the cue is noisy), while the other cue is reinforced through sensory feedback, as in Knill's (2005) study. This produces conflicting positive and negative movement errors that cannot be resolved by simple shifts of the motor output. We refer to these two types of distortion, respectively, as cue-specific bias and cue-specific noise. Our results show that the two sensorimotor adaptation mechanisms operate as predicted in these two situations, neatly explaining why motor responses often seem to prefer stereo information in studies that provide stereo-consistent sensory feedback.

2. Results

2.1. Model of adaptation to a biased slant cue

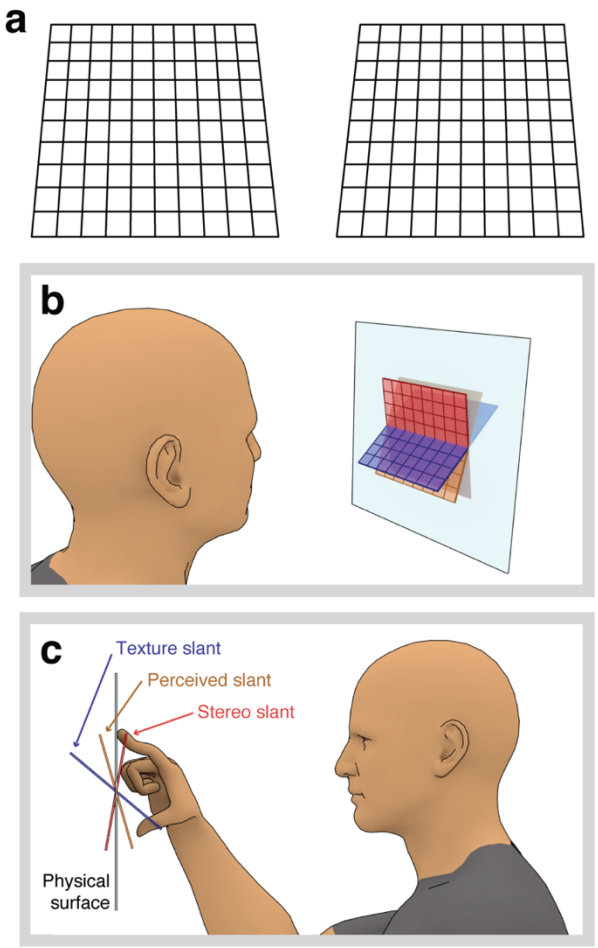

Figure 1. Interacting with slanted surfaces defined by conflicting stereo and texture cues. (a) Example stimulus (cross-fuse). Stereo information specifies a surface with the top edge nearer to the viewer than the bottom edge, while texture information specifies an opposite slant direction. Observers typically perceive a slant that is between the two component slants. (b) An observer viewing the cue-conflict slant stimulus rendered in panel a. The physical surface of the page is depicted by the large, transparent frame. The component stereo and texture surfaces are shown in conflict, with the perceived slant in the middle. In this example of viewing a stereogram printed on a flat page, the physical surface does not match the perceived slant, stereo slant, or the texture slant. (c) In our experiments, the observer attempts to place the index finger and thumb simultaneously on the displayed surface, as shown. When the planned grip orientation is not appropriate for the physical slant, an error signal is registered as one of the fingers bumps into the surface earlier than anticipated. We hypothesize that during this Grip Placement task, the grip orientation will initially target the perceived slant, but gradually come to target the haptically reinforced slant of the underlying physical surface. Unlike the example depicted here, the physical surfaces in our experiments were made to be consistent with either texture or stereo information, depending on the condition.

To motivate our model of sensorimotor adaptation to a biased slant cue, it will be helpful to first walk through a short example. Consider the cue-conflict stereogram in Figure 1a, which has been constructed so that the perceived surface should appear to have a deeper slant than the flat plane of the document it is printed on (*i.e.*, the top edge should appear slightly beyond the page while the bottom edge is slightly protruding). As depicted in Figure 1b, this perceived slant (yellow) is due to perceptual combination of a stereo slant (red) and a conflicting texture slant (blue) printed on the flat surface of the page (transparent frame). Now, imagine trying to simultaneously place your index finger on the top edge and your thumb on the bottom edge of this surface. If perceived slant guides action planning, you will reach with your grip angled slightly forward so the index finger leads the thumb. However, this means you will bump the physical page sooner than expected with your index finger (Fig. 1c), giving rise to an error signal. This error will trigger sensorimotor adaptation mechanisms that adjust movement control by changing the mapping that links visually perceived slant to the appropriate motor commands for this grip placement task. Thus, on the next trial, while the visual percept may remain the same, the mapping will now produce commands that are more appropriate for the physical slant of the surface. Within a few trials, all of your reaches will tend to be accurate.

Adaptation to a constant bias in slant perception can be formalized using a state-space model (Thoroughman & Shadmehr, 2000; Cheng & Sabes, 2006). On a given trial n, the

observer perceives some slant $\hat{s}$ that is a function of the available texture and stereo information $\Psi(s_T, s_S)$:

1.1. $\hat{s} = \Psi(s_T, s_S)$

The planned grip orientation y_n for that trial is the combination of the perceived slant $\hat{s}$ and an adjustable internal state x_n:

1.2. $y_n = \hat{s} + x_n$

When the planned grip orientation y_n does not match the physical surface slant s_Φ, haptic feedback will produce an error signal ε_n, which we can approximate as the difference between the planned grip orientation and the physical surface slant:

1.3. $\varepsilon_n = y_n - s_\Phi$

Note that Equation 1.3 is meant only as an approximation of the actual error signal, and not as a mechanistic claim that the planned movement parameter is compared directly to the physical slant encountered at the end of the movement. The error signal then causes an update of the internal state for the next trial x_{n+1} by extracting some of the error ε_n according to the learning rate b:

1.4. $x_{n+1} = x_n - b\varepsilon_n$

So, with repeated reaches, registered error signals produce cumulative adjustments to the visuomotor mapping. Through these adjustments, actions become increasingly accurate: the grip orientation becomes progressively flatter, while the visual percept remains stable. If the conflict between the stereo and texture cues is removed, then reaches should still be biased by the now-adapted internal state and give rise to an aftereffect, often considered the hallmark of adaptation.

In sum, we propose that this canonical adaptation process is triggered by observed movement errors, rapidly updating an internal state that compensates for constant perceptual errors arising from cue-specific bias (*i.e.*, consistent underestimation or overestimation of depth from a given cue). This hypothesis provides an alternative account of some dissociations between perceptual judgments and visually guided actions without supposing a fixed preference for stereo information in motor planning. Rather, actions are guided by perceptual processing, but the resulting action is also shaped by the contributions of a flexible visuomotor mapping. This hypothesis is the focus of Experiment 1, where we examine adaptation of the grip orientation when reaching to place the fingers on surfaces defined by texture and stereo cues. During Adaptation, the underlying physical slant was consistent with the texture slant while the stereo cue consistently underestimated the slant by 30°.

2.2. Model of adaptation to a noisy slant cue

The standard adaptation process described above assumes that the slant estimates provided by the individual cues cannot be accessed. Instead, adjustments can only be applied after the combined estimate $\hat{s} = \Psi(s_T, s_S)$ has been processed. However, sensorimotor adaptation may be more flexible than this, with the capability to operate on an encoding that distinguishes different values of the individual cues. To illustrate, consider a set of physical surfaces with varying slants. An observer receives texture and stereo cues about surface slant, but one of the two cues is noisy, showing a poor correlation with physical slant. As a result, the movements performed in this situation will show both positive and negative errors, depending on whether the noisy cue has indicated a spuriously large or a spuriously small slant value. In this case, there is no consistent bias for motor adaptation to simply adjust the motor output, so the canonical adaptation process modeled in the previous section will fail due to interference, producing unhelpful oscillations of the internal state.

In contrast, the hypothesis of cue-selective adaptation suggests that the individual cues (not just the combined estimate) may remain encoded in the adapted mapping, which could allow gradual tuning to relevant visual cues for a more specific form of adaptation. We can approximate this mapping as a linear function of the individual cues:

2.1. $y_n = k_{T_n} s_T + k_{S_n} s_S + x_n$

The proposed process of cue-selective adaptation would involve adjustments of the coefficients k_{T_n} and k_{S_n} of Equation 2.1 in order to reduce the influence of unreliable cues and increase the influence of reliable ones. In Experiment 2, we use an interference design to test whether cue-selective adaptation occurs when simpler changes, like the uniform shifts of motor output described above, are ineffective in reducing error.

2.3. Experiment 1: Adapting to a constant underestimation of slant from stereo

Our first experiment examined adaptation to a constant bias in stereo processing, for which our model includes two hypotheses: (1) action planning relies on perceived slant, without regard for the specific mixture of slant cues, and (2) when sensory feedback is incompatible with the perceived slant, action planning changes as a result of an adapted internal state. To set up a straightforward test of the first hypothesis in our visuomotor task, we began the experiment with a Matching task, asking participants to produce two sets of stimuli that were perceived to have the same slants but composed of different combinations of stereo and texture information. This was done by adjusting the slant of a cue-consistent stimulus ($s_T = s_S$; yellow line in Figs. 2a, 2b) in order to match each of three fixed cue-conflicts ($s_T = s_S + 30°$; blue and red lines indicating texture and stereo slants in Figs. 2a, 2b). Having obtained these matches, we predicted that equivalent motor responses would be produced when we suddenly switched from a cue-consistent stimulus to its perceptually matched cue-conflict, or vice versa, in the visuomotor task.

Following the Matching task, participants performed a Grip Placement task on slanted surfaces in a standard three-phase adaptation design (Fig. 2c). Participants reached forward with the hand held in a precision grip, orienting the grip in order to contact the surface with their index finger and thumb simultaneously (Fig. 2d). In the initial Baseline phase (Fig 2c, left), each participant interacted with their personalized set of cue-consistent stimuli indicated during the Matching task. This phase was meant to emulate normal visuomotor coordination—the physical slant encountered at the end of each movement matched both cues. After establishing Baseline performance, we transitioned into the main Adaptation phase (Fig 2c, middle), suddenly replacing the three cue-consistent stimuli with the fixed set of three cue-conflicts. We predicted that grip orientations would not change across this transition because the perceived slant remained the same—notice how the black curve in Fig 2c is at the same level as Baseline on the very first Adaptation trial. During Adaptation, the physical surface reinforced the texture slant, which was consistently 30° greater than the stereo slant (and therefore a relatively constant amount greater than the perceived slants). We predicted that the grip orientation would gradually tend toward these deeper physical surface slants as the internal state was modified to reduce error signals. Finally, we switched back to the cue-consistent surfaces in a short Washout phase (Fig 2c, right)—again we predicted no sudden change across the transition due to the perceptual matching, followed by rapid convergence back on Baseline performance.

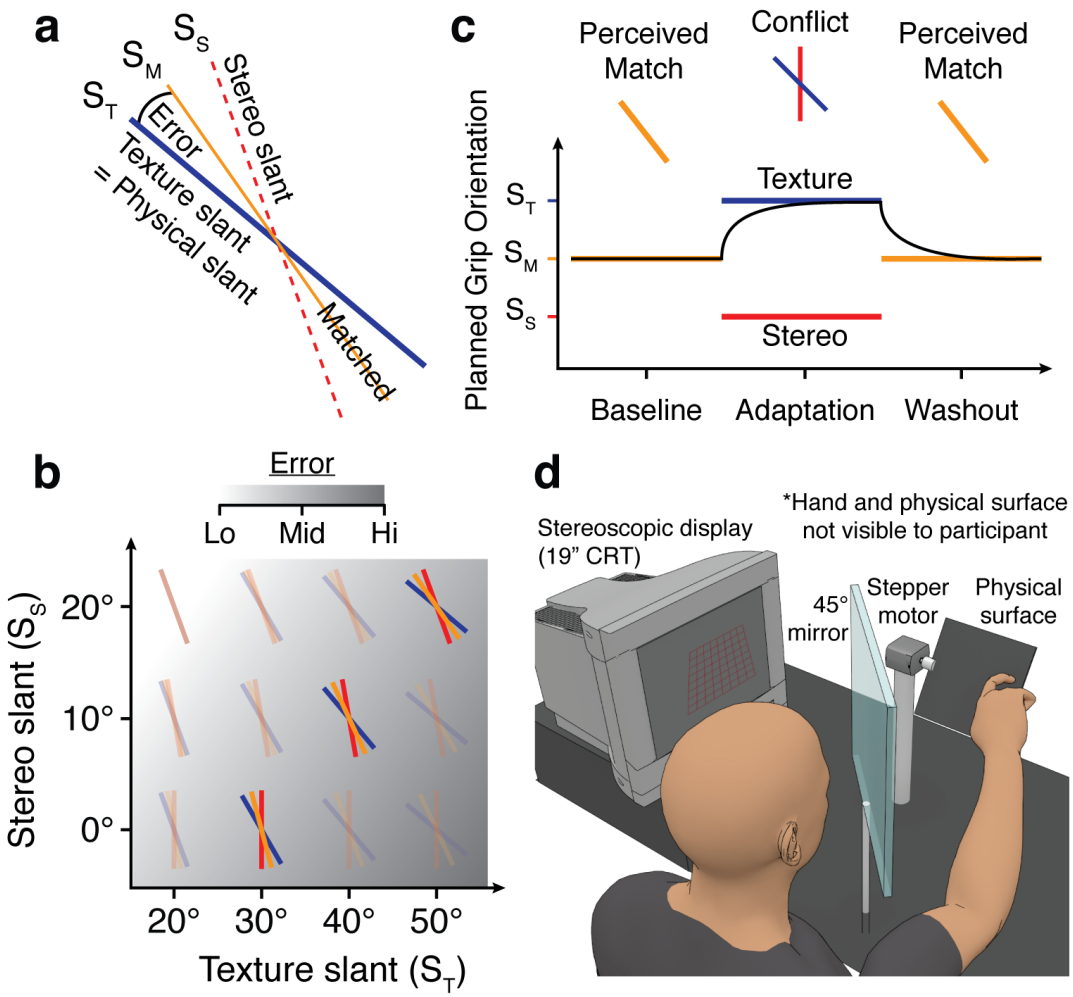

Figure 2. Experiment 1: Adaptation to a constant bias in stereo. (a) Magnified view of an individual cue-conflict surface. Texture (blue) correctly specifies the physical slant, while stereo (red) shows an underestimation bias, producing an intermediate perceived slant (yellow). Aiming at the perceived surface will produce a movement error. (b) Stereo was consistently 30° shallower than texture slant for all three conflict surfaces (highlighted), so errors are relatively constant (background gradient). (c) Grip Placement task. In Baseline, cue-consistent surfaces perceived to have the same slants as the conflict surfaces (yellow line) were presented to establish normal movement coordination (black line). In Adaptation, the cue-consistent surfaces were replaced by the conflict surfaces (blue and red lines). On the very first trial of Adaptation, the planned grip orientation (black curve) should be identical to Baseline, mirroring the perceptual equivalence of the two sets of stimuli. Thereafter, planned grip orientations should adapt toward the texture-consistent physical slants. At the transition to the Washout phase, where we reintroduce the cue-consistent surfaces, again we should observe no sudden change in the motor response due to their perceptual equivalence. Thereafter, we predict rapid convergence on Baseline performance. (d) The multisensory virtual reality rig used in our experiments. Planar surfaces are rendered with stereo and texture cues on a CRT monitor. The image is reflected by the oblique mirror, making it appear beyond the mirror in the workspace. A physical surface in the workspace is precisely aligned with the visual stimulus to provide haptic feedback; its slant is controlled by a stepper motor. On each trial, the participant reaches with the right hand in a precision grip, orienting the hand in order to place the index finger and thumb simultaneously on the observed surface. Optotrak markers (not shown) are worn on the index finger and thumb nails of the right hand to track the grip orientation throughout the movement. The room was completely dark, and an opaque back panel was placed on the mirror to prevent vision of the physical surface or of the hand.

The results of the Matching task are presented in Figure 3a, with mean cue-consistent slants of 15.6°, 28.8°, and 39.3° for the three fixed cue-conflicts (s_T/s_S= 0°/30°, 10°/40°, 20°/50°). These data correspond to a relative weight on texture information w_T of 0.60 (SEM = 0.08) according to $s_{match} = w_T s_T + (1 - w_T)s_S$. We computed a relative weighting of texture and stereo information here, rather than a freely varying slope parameter, because our Matching task did not provide an estimate of absolute perceived slant; it only indicated the relative influences of the two conflicting cues. w_T is relatively constant in the tested range, though there was a slight trend toward an increased influence of texture information at greater slants, consistent with previous findings (Hillis *et al*., 2004).

Figure 3b depicts the average planned grip orientations for each of the three cue-consistent match slants during Baseline (small circles), where the physical slants matched the perceived slants, as well as during the final 30 trials of Adaptation (large circles), after exposure to the perceptually matched cue-conflicts with texture-consistent haptic feedback. We found that the planned grip orientations from the final 30 trials of Adaptation closely matched the reinforced physical slants, though on average slightly exceeding them.

Most importantly, the timeline of grip orientations depicted in Figure 3c shows that the results of Experiment 1 are consistent with our model of adaptation to a constant bias (in this case, an underestimation of slant from stereo). At the transition from Baseline to Adaptation, the average grip orientation did not change—the model accurately predicted the planned grip orientation on the very first trial with zero free parameters (the error-correction rate, our only free parameter, does not affect the first trial of Adaptation or Washout). Following this initial trial, grip orientations rapidly shifted toward the feedback-reinforced texture slants and stabilized. The model fit produced an error-correction parameter b of 0.21, which was robust to random variation of the initial guess in the fitting procedure. On the first trial of Washout, the planned grip orientation again matched the model prediction almost perfectly, with no sudden shift in response to the altered values of texture and stereo information. Thereafter, the planned grip orientation converged back to the same point as Baseline, responding to haptic feedback from the now-shallower physical slants. These findings support our claim that visually guided actions are based on the same processing of 3D information as perception. If a separate visual encoding of 3D shape were used to drive action, we should have observed some shift in the planned grip orientation on the first trials of Adaptation and Washout, indicating different processing of the two cues.

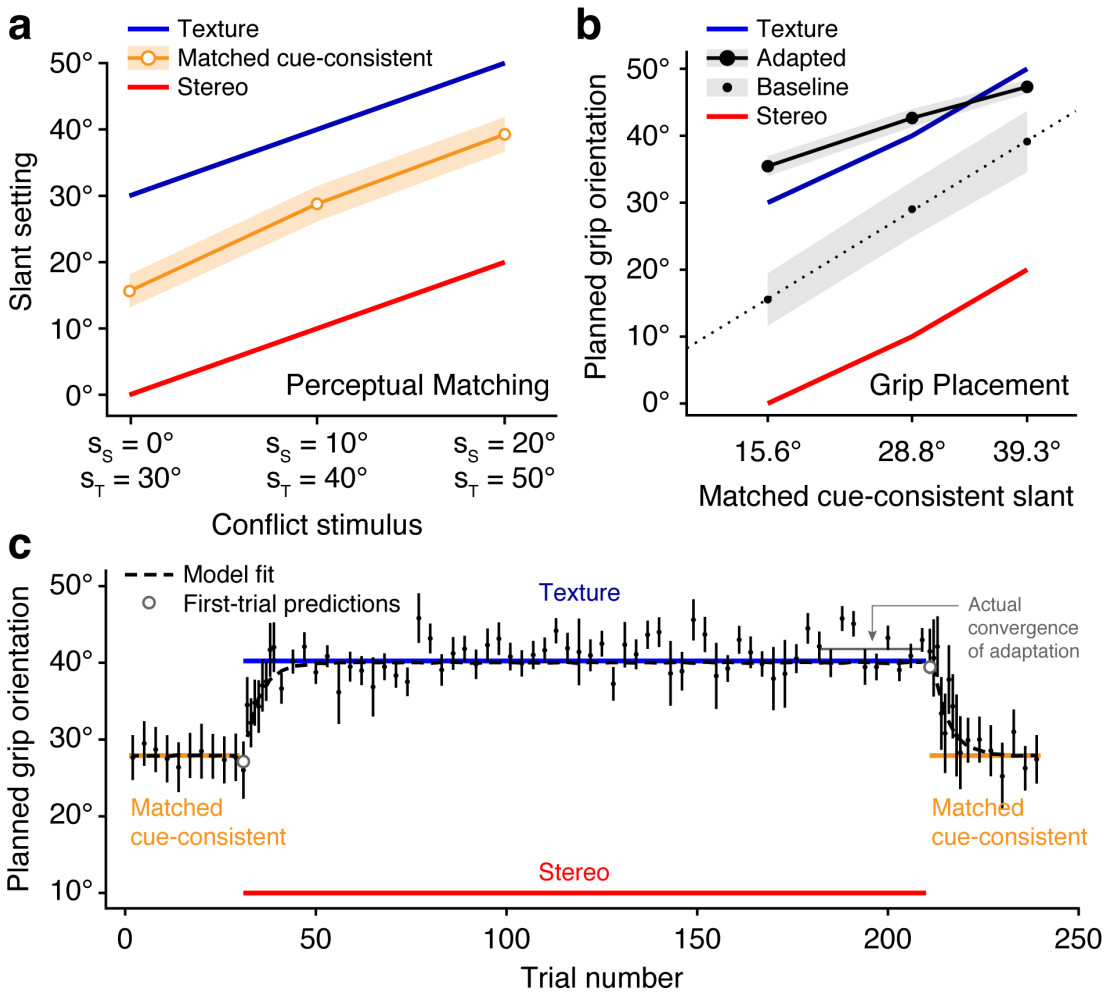

Figure 3. Experiment 1 results. (a) Matching: On average, participants set the cue-consistent slant midway between the component stereo and texture slants of the cue-conflict stimuli they were asked to match. (b) Grip Placement: Average grip orientations as a function of perceived slant, in Baseline and during the final 30 trials of Adaptation ("Adapted"). Note that in this graph we are plotting the dependent variable from panel a on the x-axis; this is why the solid red and blue lines, which show the stereo and texture slants of the three cue-conflict stimuli, are not straight. In Baseline, the planned grip orientations were the physical slant values (see Fig. S1 & Methods for more details); this is why these data points fall on the unity line. By the end of Adaptation, however, the planned grip orientations shifted to track with the texture slants, which were reinforced by haptic feedback. (c) Grip Placement: Timeline of planned grip orientations. Our model accurately predicted the absence of any sudden change in grip orientation in the first trials of Adaptation and Washout, and the time course of adaptation was closely fit with an error-correction rate of 0.21. There was a slight underestimation of the actual convergence point of adaptation (compare model fit with black line labeled Adapted; this difference is also noticeable in panel b, Texture Slant versus Adapted)—we made a slight adjustment of the internal state of the model on the first trial of Washout to account for this discrepancy.

In the Introduction, we argued that the preferential reliance on stereo information found in some previous studies could have been a result of averaging over the rapid sensorimotor adaptation that occurs in response to sensory feedback. This argument is based on the fact that, in most cases, available feedback signals tend to be consistent with stereo. This is always true of haptic feedback in real-world settings, because while surface markings on an object can be drawn to mislead the viewer, binocular disparities are fully determined by physical structure.

Visual feedback of the hand also tends to reinforce stereo, in real-world and virtual settings, because we expect relative disparities to be nulled at object contact. Either of these sources of sensory feedback would therefore drive motor outputs toward the stereo-specified 3D shape of objects. However, based on the data of Experiment 1, this argument can only be applied to situations involving an approximately constant degree of bias in the unreliable cue, across all target objects. This is because the canonical adaptation mechanism operates by simply shifting motor outputs to minimize movement errors—it does not strengthen the influence of the reinforced cue in the planning process. In Experiment 2, we tested whether such reweighting is possible when shifting the motor output is insufficient to reduce movement errors.

2.4. Experiment 2: Adapting to an unreliable slant cue

The form of sensorimotor adaptation explored in Experiment 1 increased the planned grip orientation, independent of the specific values of stereo and texture information that were being presented, which was sufficient to compensate for the constant bias in stereo information. In Experiment 2, we consider the more complex pattern of perceptual inaccuracy that occurs when the signals obtained from a given cue suddenly become noisier. Cue-specific noise occurs naturally when processing depth cues in real-world situations, but it is also produced in many experimental situations using cue-conflict stimuli, due the randomized order of stimulus presentation. During exposure to cue-specific noise, accurate performance cannot be achieved by uniformly shifting motor outputs (*i.e.*, by changing the "intercept" x_n, Eqn. 1.2). Instead, selective adjustment of the influence of each individual cue (*i.e.*, changes to the slopes k_{T_n} and k_{S_n}, Eqn. 2.1) is needed.

In order to determine whether cue-selective adaptation is possible, in Experiment 2 we created a set of stimuli with independently varying stereo and texture information ($s_T \in \{15°, 30°, 45°\} \times s_S \in \{15°, 30°, 45°\}$) and asked separate groups of participants to interact with these objects in either a haptic-for-texture condition or a haptic-for-stereo condition. Figure 4a illustrates this stimulus set: the three stimuli located along the identity line were cue-consistent surfaces, whereas the other six stimuli (off-diagonal entries in Fig 4a) were rendered with different degrees of cue conflict. Critically, stimuli on the opposite sides of the main diagonal will bias perception in opposite directions with respect to the haptically reinforced cue, creating interference in adaptation due to conflicting error signals. To illustrate this point, consider performing a precision-grip placement on the stimulus in the top-left of Fig. 4a (blue = texture, red = stereo, yellow = perceived), where texture specifies a flatter surface than stereo, and imagine haptic feedback is consistent with the texture slant. As before, we assume that the initial reach will be directed at the perceived slant, and as a result the index finger will bump into the physical surface, registering as a *positive* error. In contrast, on a reach to the stimulus in the lower-right of Fig. 4a, the *thumb* will bump too soon, registering as a *negative* error. This interference renders ineffective the canonical form of adaptation that uniformly shifts the motor output. However, if sensorimotor adaptation also operates on encodings that distinguish between the different values of the stereo cue and the texture cue (as opposed to a singular, fused estimate of slant; see Hillis, Ernst, Banks, & Landy, 2002), then it should be possible to elicit a change in the relative influences of the two cues over time. Since the non-reinforced cue is completely uncorrelated with the slants of the physical surfaces in our stimulus set, the optimal solution would be to eliminate that cue's influence and to increase the influence of the reinforced cue to match Baseline performance.

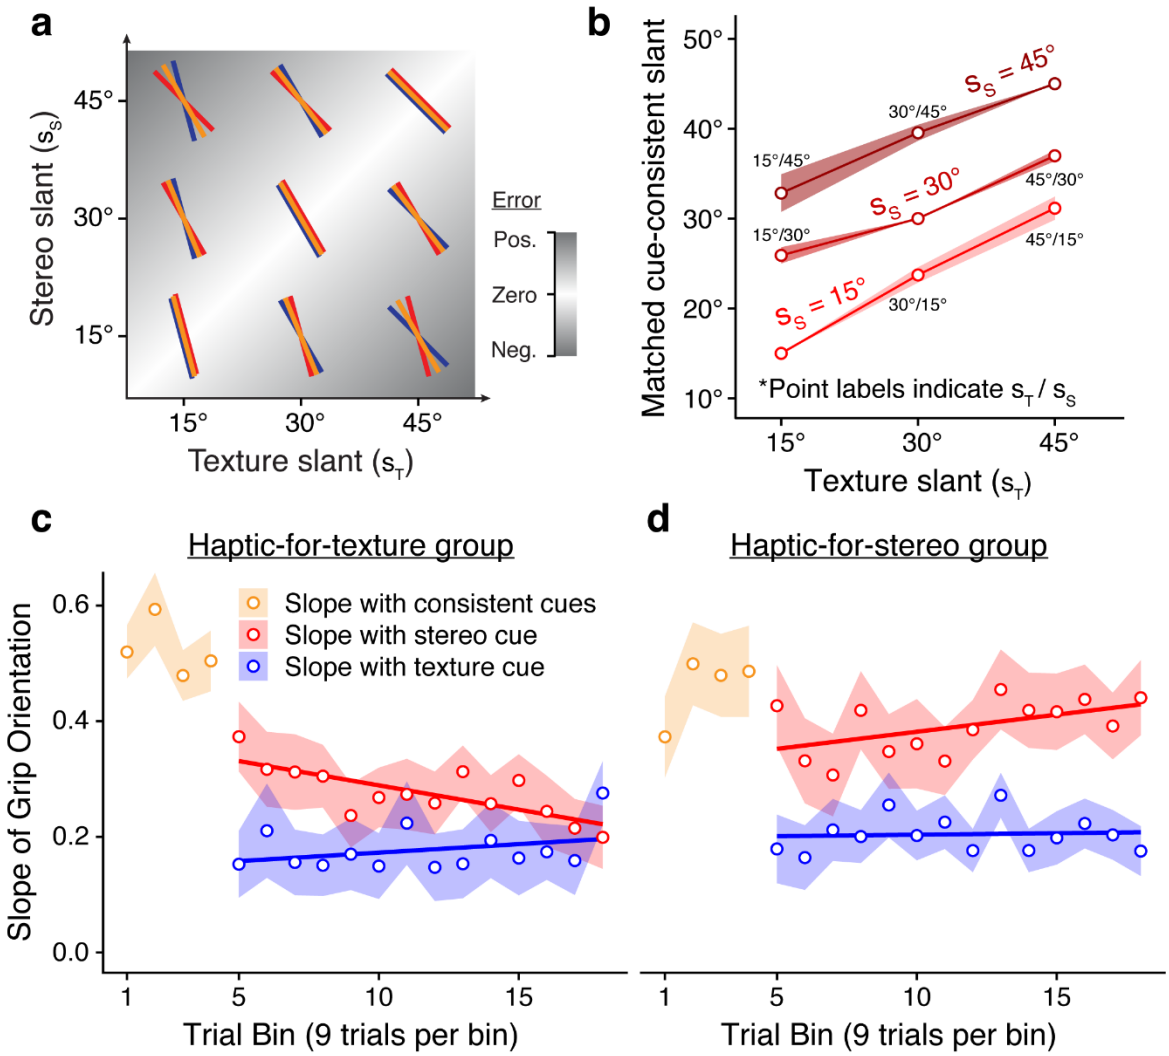

Figure 4. Experiment 2: "Interference" design and results. (a) Nine fixed target stimuli were obtained from a factorial combination of three texture slants (blue) and three stereo slants (red). The main diagonal of this array contains three cue-consistent stimuli, while the six off-diagonal stimuli are cue-conflicts. Notice that for half of the cue-conflicts (bottom-right region), the expected cue-consistent match slant (yellow) is shallower than texture and deeper than stereo, while for the other half (top-left region), it is deeper than texture and shallower than stereo. When haptic feedback reinforces one cue and is uncorrelated with the other, conflicting positive and negative movement errors will be experienced, creating interference in standard adaptation processes (Eqns. 1.1–1.4). To reduce these errors, cue-selective adaptation is required. (b) Matching results: for each of the six cue-conflict stimuli, matched cue-consistent slant settings (y-axis) were midway between the two component cues of the conflicts. The nine data points correspond to the nine fixed target stimuli, with texture slant indicated by position along the x-axis and stereo slant indicated by the three groupings (different shades of red). (c, d) Grip Placement task: In each bin, participants aimed at each of the nine stimuli once. In Baseline, we presented the cue-consistent match stimuli, and fit simple linear regressions to the grip orientations in each bin; yellow data indicate the slope coefficients of these regressions. In Adaptation, we presented the nine target stimuli with physical slant reinforcing texture (panel c, haptic-for-texture group) or stereo (panel d, haptic-for-stereo group). The red and blue data indicate the stereo and texture slopes estimated via 2D linear regression. The lines plotted over these data depict the results of a bin-wise linear regression on the slopes.

The results of the initial Matching task are presented in Figure 4b for the six cue-conflict stimuli from the target set. The cue-consistent match slants were set, on average, about halfway between the component stereo and texture slants of the conflict stimuli. Similar to the relative weight on texture information w_T of 0.60 found in Experiment 1, these data correspond to a w_T of 0.56 (SEM = 0.04).

The linear regression coefficients (slopes) estimated in each bin of the Grip Placement task are depicted in Figures 4c (haptic-for-texture group) and 4d (haptic-for-stereo group). During the Baseline phase, participants' planned grip orientations were related to cue-consistent slants by a slope of 0.50 (SEM = 0.03), comparable to the slope of 0.60 found in the Baseline phase of Experiment 1. In the Adaptation phase, we obtained independent slopes for stereo and texture information in each bin via multiple linear regression; these indicate the influence of each cue on the motor response. To measure changes in the influence of each cue over time, we fit an additional linear regression on these slopes as a function of bin number (solid red and blue lines in Figs 4c & 4d). The slope parameters of these temporal regressions indicate the bin-wise rate of change in the influence of each of each cue. Analyzing this rate-of-change measure using a mixed-design ANOVA (Feedback Group × Depth Cue), we found a significant interaction of the two factors ($F(1, 45) = 5.13$, $p = 0.028$), confirming our main hypothesis that the two feedback conditions elicited changes in the relative influences of the two cues.

Follow-up analyses revealed a significant difference between feedback conditions in the rate of change of the influence of stereo (one-tailed, two-sample *t*-test; $t(39.79) = 3.47$, $p < 0.001$), but not in the rate of change of the influence of texture ($t(44.06) = 0.44$, $p = 0.33$). Additional follow-up tests demonstrated that changes in the influence of stereo were observed in both conditions: we found a significantly increasing influence of stereo during the haptic-for-stereo condition (one-tailed *t*-test; mean = +0.0059 per bin, $t(18) = 1.86$, $p = 0.040$) and a significantly decreasing influence of stereo during the haptic-for-texture condition (mean = −0.0084 per bin, $t(27) = 3.09$, $p = 0.0023$).

3. Discussion

Consistent with our hypothesis of cue-selective sensorimotor adaptation, Experiment 2 demonstrated that haptic feedback can alter the relative influence of stereo and texture cues in the motor planning processes that control grip orientation. Specifically, we found that the influence of stereo information was enhanced over time when it was reliably correlated with physical surface slant, and dramatically reduced over time when it was uncorrelated with physical slant. Meanwhile, the influence of texture information remained relatively constant throughout Adaptation, regardless of the feedback condition. This indicates that different cues to 3D shape remain separably encoded at the level of the visuomotor mapping where sensorimotor adaptation operates. Critically, we did not observe proportionally equivalent changes in the influences of stereo and texture information, as would be expected if the adaptation mechanism were simply changing the scaling of the planned grip orientation with respect to a singular, fused estimate of slant, like that suggested by Hillis *et al*. (2002).

However, the observed changes fell quite short of the optimal form of cue-selective adaptation that might have been achieved. Ideally, participants would have increased the influence of the reliable cue to roughly match their Baseline slope, while decreasing the influence of the unreliable cue to zero. While the slope with stereo slant appears to be converging on the average Baseline slope in the haptic-for-stereo group, the influence of texture information did not show any noticeable decrease in this condition. Similarly, the slope with stereo slant was reliably suppressed over time in the haptic-for-texture condition, but the corresponding slope with texture slant remained surprisingly constant. Thus, the present data are

sufficient to conclude that cue-selective adaptation takes place when one slant cue suddenly becomes unreliable, potentially creating a dissociation between perception and action (assuming perception has not also changed). However, we are unable to identify the specific constraints on this process that produced sub-optimality in this experiment. In the final section, we will speculate about what these constraints may be.

3.1. Depth cues as feature dimensions supporting selectivity in motor adaptation

To gain additional insight into the mechanism of cue-selective adaptation, we again looked to the adaptation literature, focusing on how interference between conflicting error signals might be resolved. One way to do this would be to selectively apply error corrections from erroneous movements so that they primarily affect contextually similar movements. For instance, imagine you are presented with two different targets in a reaching task, a square target and a circular target, both of which are displayed in the same spatial location. When reaching to the square target, visual feedback is shifted to the right, but when reaching to the circular target, it is shifted to the left. Ideally, you would learn to apply opposite movement corrections according to the visual identity of the target, but adaptation studies applying these sorts of cue-dependent perturbations have generally found that explicit cues like shape and color are not sufficient to overcome the interference between the two perturbations (Jordan: best citations?). Instead of being selectively applied to each target, error corrections applied to a square target will automatically generalize to a circular target, and vice versa, canceling out one another.

However, previous adaptation studies have also identified some target dimensions that do in fact provide a basis for selective application of error corrections. In the reach adaptation literature, the most well-known of these dimensions is spatial location (Jordan: best citations?). When opposing perturbations of visual feedback are systematically linked to leftward and rightward reach targets, a sufficient separation between the targets makes it possible to simultaneously learn the opposite shifts of planned reach direction that are needed to perform accurately. We recently reported a similar result for grasp adaptation: participants simultaneously learned to increase the planned grip aperture for targets in one location and decrease it in another, thus reducing errors caused by opposing visual size perturbations (Cesanek & Domini, 2017). This location selectivity is a critical feature of motor adaptation because it provides a basis for learning to interact with different objects in different ways, which is just as important as generalizing previous motor learning to interactions with novel objects. Meanwhile, the search for other task dimensions that support selective adaptation, whereby shielding new skills from the potentially deleterious effects of generalization, has yielded few candidates.

Remarkably, the results of Experiment 2 reveal that the visual feature dimensions established by individual depth cues are capable of supporting selectivity in motor adaptation. All of the surfaces in our experiment appeared in the same spatial location, thus resulting in considerable trial-by-trial interference in the motor output shifts produced by the canonical adaptation process. Still, participants succeeded in partially compensating for the opposing visual perturbations caused by the unreliable depth cue. For example, in the haptic-for-texture condition, participants markedly reduced their planned grip orientations over time for the subset of objects with misleadingly deep stereo slants, while increasing their planned grip orientations for objects with misleadingly shallow stereo slants. It would have been impossible for participants to explicitly adopt these modifications, as they had no conscious access to the component values of stereo slant and texture slant that comprised each of the conflict stimuli; the percept is always a unified, coherent surface. Notice, however, that the analysis of individual depth cues is integral to the stream of processing that produces a grasping movement—stereo and texture *must* be analyzed to extract the slant or depth estimate that guides the movement. In

contrast, the sorts of explicit cues that have failed to mediate selective motor adaptation in previous studies generally play no functional role in the planning process—shape and color are not involved in the computation of a reach target's location. Tentatively, we suspect that the persistence of movement errors, due to interference in the canonical adaptation process, allows the corrective effects of these errors to propagate upstream, gradually producing changes at deeper layers of the visuomotor network where individual depth cues are encoded.

3.2. A potential mechanism for cue-selective adaptation

Converging behavioral and neurophysiological evidence from reach adaptation paradigms suggest that location-selective adaptation relies on a form of local function approximation: the distributed neural encoding constituting a visuomotor mapping provides the flexibility needed to selectively apply error corrections (Pouget & Snyder, 2000; Donchin, Francis, & Shadmehr, 2003; Hwang, Donchin, Smith, & Shadmehr, 2003). Let us assume that some network units respond maximally when planning a reach to a leftward target, while others respond maximally to a rightward target, and that the activations of all units for a given target direction are weighted and summed to produce the motor output. In this formulation, the adaptation process modeled by Equations 1.1–1.4 would be equivalent to adjusting the weights of all units equally based on each observed movement error. However, with a learning rule where the amount of change in a given unit's weight is directly proportional to its influence in planning the erroneous movement (*e.g.*, the delta rule), the network will gradually learn opposite compensations at the two target locations, due to the location selectivity of the individual units.

Considering the present results, we speculate that the local function approximation framework might be usefully extended to consider units that are tuned to different combinations of depth cues within the visuomotor mapping. This would be one way that opposite compensations could be learned for the opposite perceptual biases found across our stimulus set (Fig. 4a). This idea remains highly speculative, but reasoning about our findings in this way provides one possible explanation for the noisiness and relatively slow rate of the observed cue-selective adaptation, compared to the standard adaptation observed in Experiment 1. Some of this is due to the increased complexity of our analysis—regression slopes estimated with only nine data points are bound to be noisy. However, noisy trial-by-trial fluctuations in the planned motor output are an inevitable occurrence when a network model with broadly tuned units encounter large errors in opposite directions. If neurons tuned to specific combinations of stereo and texture cues do in fact constitute part of the visuomotor mapping for actions guided by 3D spatial properties, then the current results suggest broad tuning along the component depth cue dimensions. While this would make cue-selective adaptation slower, it provides an economical benefit by reducing the number of neurons required to encode the full range of co-variability across the large number of cue dimensions occurring in natural viewing. In particular, our data suggest that the tuning of these units may be broader along the texture dimension than along the stereo dimension—this would lead to increased interference among the three texture slant values we tested, helping to explain why the influence of texture information did not change significantly in either condition.

4. Conclusion

In two experiments, we examined changes in the performance of visually guided reaching movements in response to haptic feedback, with the goal of understanding whether sensorimotor adaptation actively contributes to the reduction of movement errors that arise when 3D shape perception is inaccurate. In particular, we focused on two types of distortion—cue-specific bias and cue-specific noise—that could be expected to produce inaccurate perception of slant from

binocular disparity (stereo) and texture gradients, even in natural viewing conditions. Experiment 1 provides clear support for a parsimonious model of perception and action, in which movement planning is based on the same visual encodings of 3D shape as conscious perceptual judgments. Experiments 1 and 2 both show that sensorimotor adaptation fine-tunes motor performance when the processing of a particular cue becomes biased or unreliable, helping to compensate in the short term for fluctuations in perceptual processing, but in very different ways. In Experiment 1, we observed a constant shift of the motor output, while in Experiment 2, we found more gradual changes in the relative influences of stereo and texture information, which we take as evidence that these upstream visual dimensions can also be targeted by motor adaptation processes. Our findings have important methodological implications as well: the operation of standard and cue-selective adaptation mechanisms over relatively short timescales provides a simple explanation of why some studies find greater reliance on stereo information in visuomotor tasks involving cue-conflict stimuli. Notably, this account avoids postulating redundant encodings of spatial properties for perception and for action. In contrast to a dissociated view of perception and action, we take these findings to be indicative of a close connection between the two functions, as the action-relevant process of sensorimotor adaptation was found to be sufficient for correcting certain inaccuracies known to afflict 3D shape perception.

4. Methods

4.1. Participants

Participants were between 18 and 35 years old, right-handed, and with normal or corrected-to-normal vision. They were either granted course credit or paid $8/hour as compensation. Informed consent was obtained from all participants prior to any participation, in accordance with protocol approved by the Brown University Institutional Review Board and performed in accordance with the ethical standards set forth in the Declaration of Helsinki.

Fifteen participants were recruited for Experiment 1. Forty-eight participants were recruited for Experiment 2. Twenty-eight were exposed to the haptic-for-texture condition, and twenty were exposed to the haptic-for-stereo condition. One participant from the latter condition was excluded from analysis because more than half of their Grip Placement trials were marked for exclusion by the criteria indicated in Section 4.1.1.

4.2. Apparatus

The general layout of the tabletop virtual reality apparatus used in this experiment is depicted in Figure 2d. Participants were seated in a height-adjustable chair so that the chin rested comfortably in a chinrest. Movements of the right hand were tracked using an Optotrak Certus motion-capture system. Small, lightweight posts containing three infrared emitting diodes were attached to the fingernails of the index finger and thumb, and the system was calibrated prior to the experiment to track the extreme tips of the distal phalanxes of each finger. This motion-capture system was coupled to a tabletop virtual reality environment: participants looked into a half-silvered mirror slanted at 45° relative to the sagittal body midline, which reflected the image displayed on a 19” CRT monitor placed directly to the left of the mirror at the correct distance to provide consistent accommodative and vergence information.

Participants viewed stereoscopic renderings of slanted 3D surfaces in this tabletop virtual reality setup, where stereo and texture information were controlled independently using a 3D graphics technique known as backprojection (Hillis, Watt, Landy, & Banks, 2004). The surfaces were slanted by rotating around a transverse axis through the middle of the object, which was positioned directly in front of the eyes at a viewing distance of 38 cm. This arrangement caused

the rendered 3D surfaces to appear to be floating in space beyond the mirror. When fully frontoparallel, the surface diagonally subtended 13° of visual angle. Stereoscopic presentation was achieved with a frame interlacing technique in conjunction with liquid-crystal goggles synchronized to the frame rate. Additional linear positioning stages and stepper motors in the workspace were used to position and change the orientation of the physical stimulus with respect to the visual stimuli.

4.3. Procedure

Experiment 1 involved two tasks. First, in the Matching task, participants matched through a psychometric procedure the perceived slants of three stereo-texture conflict surfaces (s_{T_i}, s_{S_i}) with cue-consistent surfaces $(s_{match_i}, s_{match_i})$ where stereo and texture specified the same slant, such that $\Psi(s_{T_i}, s_{S_i}) = \Psi(s_{match_i}, s_{match_i})$. The three stimuli were chosen such that, for each surface, stereo slant and texture slant differed by a constant conflict angle of 30° ($s_{T_i} = s_{S_i} + 30°$). These combinations of stereo and texture slant are illustrated in Fig. 2b. On each trial, we allowed participants to switch freely between the fixed cue-conflict stimulus and an adjustable cue-consistent stimulus, using keypresses to make incremental changes to the slant of the cue-consistent stimulus until it appeared to match the slant of the cue-conflict stimulus. To prevent the use of motion information, we displayed a blank screen with a small fixation dot for an inter-stimulus interval of 750 ms whenever the stimulus was changed. Participants performed five repetitions of matching for each of the three fixed cue-conflicts, for a total of 15 trials.

The resulting set of six stimuli (three pairs of matched cue-conflict and cue-consistent surfaces) were then presented as stimuli in the Grip Placement task. With the hand shaped into a precision grip, participants reached toward the displayed surface with the goal of making the index finger and thumb hit the surface at the same time. They were instructed to use only the wrist and arm to change the grip orientation while holding their fingers in fixed precision grip posture. This task is similar to handheld object placement tasks where cylinders are placed so that their bottoms are parallel with the surface at contact (Knill, 2005), but also engages the fingers' tactile sensitivity. A standard, three-phase "ABA" design was adopted for the visuomotor task. In the Baseline phase, participants reached toward their personalized set of cue-consistent surfaces for 30 trials. At the transition from Baseline to the Adaptation phase, the cue-consistent surfaces were suddenly replaced by the perceptually matched cue-conflict surfaces, with the underlying physical surface reinforcing the texture slant. Following 180 trials of exposure to these conflict surfaces, the experiment concluded with a 30-trial Washout phase, identical to Baseline. Throughout this task, we used a binned trial order such that each of the three surfaces was presented once before any one was repeated, facilitating local averaging of trials.

In Experiment 2, participants again performed the Matching and Grip Placement tasks, but there were nine target stimuli presented during Adaptation (Fig 4a): three cue-consistent (main diagonal) and six cue-conflict (off-diagonal cells). The six cue-conflict surfaces were perceptually matched with adjustable cue-consistent stimuli in the first block using the same psychometric procedure as Experiment 1. We eliminated the Washout phase from this experiment in order to reduce the overall duration. During Baseline, the targets were the nine cue-consistent perceptual matches, specific to each subject. These nine targets were presented four times each in a binned trial order for a total of 36 Baseline trials. During Adaptation, the visual stimuli were the nine target stimuli, and the physical surface slants were consistent either with the texture slants (haptic-for-texture group) or the stereo slants (haptic-for-stereo group). These were presented 14 times each in a binned trial order for a total of 126 trials.

4.4. Analysis

Raw motion-capture position data was processed and analyzed offline using custom software. Missing frames due to marker dropout were linearly interpolated, and the 85-Hz raw data was smoothed with a 20-Hz low-pass filter. We excluded from analysis all trials where (a) the proportion of missing frames exceeded 90%, (b) fewer than 5 frames were not missing, or (c) the center of the grip traveled less than 25 mm. In Experiment 1, these criteria resulted in the exclusion of 79 individual trials out of a total 3,600 trials (approx. 2%).

In-flight grip orientation was calculated as the declination of the line joining the fingertips in the sagittal plane. To infer the planned grip orientation at the end of the movement, we used a snapshot of the grip orientation at a moment 10 mm prior to contact (see Fig. S1a). This was done in order to avoid any contamination caused by adjustments made after one of the fingers contacted the surface. To extract this *terminal* grip orientation, we first processed the entire trajectory, then searched for the first motion-capture frame where either one of the fingers contacted the physical surface (minimum orthogonal distance from each finger to the surface plane), then scanned 10 mm backward along the trajectory.

However, this terminal grip orientation does not represent the planned grip orientation because the movement is still in progress. For the data displayed in Fig. 3c (Experiment 1), we used each participant's Baseline performance to determine how their measured terminal grip orientations related to the actual physical slants they intended to place the fingers on, which were presumably their planned grip orientations (Fig. S1b). Accordingly, we fit a linear regression to this Baseline data and used the inverse of the estimated function to transform terminal grip orientations, our raw dependent variable, into the planned grip orientations (Fig. S1c). This does not affect the statistical analysis or the modeling of our data, but simply helps to present the data in a more understandable format, with the dependent variable sharing the same metric as the rendered slant information.

In Experiment 1, we fit the error-correction parameter b of our first adaptation model to minimize the root mean squared error of the model with respect to the average trial-by-trial grip orientations. To do so, we used the constrained optimization by linear approximation (COBYLA;

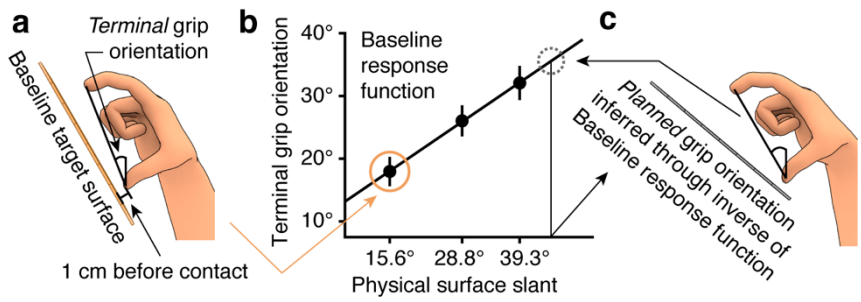

Figure S1. Computation of the planned grip orientation. (a) Our raw dependent variable was the terminal grip orientation, the angle between the index finger and thumb when the fingers were 10 mm from their first contact with the physical surface. Even in the Baseline phase, where we presented cue-consistent slants, this raw dependent variable tended to be slightly less than the actual slants of the physical surfaces, as the movement was still in progress. For ease of visualization, we sought to transform these values into planned grip orientations that would be consistent with rendered slant values. (b) The Baseline terminal grip orientations were presumed to be aimed at the physically presented slant values, since participants performed enough trials to find a comfortable movement patterning. We fit a linear regression on this data in order to estimate the relationship between planned grip orientations (represented by the physical slants in Baseline) and our raw dependent variable. (c) We then used this estimated function to transform our raw terminal grip apertures from the whole experiment into planned grip orientations, which are on the same scale as the rendered slants.

Powell, 1998) algorithm of the nloptr package (Johnson, n.d.) in R (R Core Team, 2014), constraining the fit so that *b* was bound between 0 and 1. After fitting the model, we observed that participants' planned grip orientations converged on a slightly greater value than did the model, as the model was bound to converge on the actual physical slants that were presented. In order to incorporate this empirical measure of the fully adapted state into the model before making continuing to make predictions for the Washout phase, we manually set the internal state of the model on the first trial of Washout to reflect the average change in grip orientations from Baseline to the final 30 trials of Adaptation.

The factorial design of the conflict stimuli in Experiment 2 allowed us to measure the relative influence of stereo and texture information in the Grip Placement task by estimating coefficients (slopes) for each cue via multiple linear regression according to Eqn. 2.1, with the terminal grip orientation as the response variable. Unlike in Experiment 1, we did not transform the terminal grip orientations into planned grip orientations—this was not necessary in Experiment 2 because we were interested only in how the slope coefficients from the multiple regression changed over time. Just as before, terminal grip orientations were measured by finding the moment at which one of the fingers first touches the surfaces, then scanning backwards by 10 mm of hand movement. A regression was computed for each bin of nine trials within the Adaptation phase, producing a fine-grained (though somewhat noisy) timeline of the influence of each cue on terminal grip orientation. In Baseline, we computed the slope of the terminal grip orientation with respect to the perceptually matched slant values using simple linear regression in each bin.

In Experiment 2, since there were no differences in procedure between the two feedback conditions until the Adaptation phase of the Grip Placement task, and since no significant

differences were found in Matching performance between the two groups for any of the six cue-conflict stimuli, the two groups were combined when analyzing the Matching task.

Acknowledgments

This work was supported by a National Science Foundation grant to F. D. and J. A. T. (BCS 1827550).

References

Bruggeman, H., Yonas, A., & Konczak, J. (2007). The processing of linear perspective and binocular information for action and perception. *Neuropsychologia*, *45*(7), 1420–1426.

Cesanek, E. & Domini, F. (2017). Error correction and spatial generalization in human grasp control. *Neuropsychologia*, *106*, 112–122.

Cesanek, E., Campagnoli, C., Taylor, J. A., & Domini, F. (2018). Does visuomotor adaptation contribute to illusion-resistant grasping? *Psychonomic Bulletin & Review*, *25*(2), 827–845.

Cheng, S. & Sabes, P. N. (2006). Modeling sensorimotor learning with linear dynamical systems. *Neural Computation*, *18*(4), 760–793.

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. *Journal of Neuroscience*, *23*(27), 9032–9045.

Heine, L. (1900). Ueber orthoskopie oder ueber die abhängigkeit relativer entfernungsschätzungen von der vorstellung absoluter entfemung [On "orthoscopy" or on the dependence of relative distance judgments on the representation of absolute distance]. *Albrecht von Græfe's Archiv für Ophthalmologie*, *51*, 563–572.

Hillis, J. M., Ernst, M. O., Banks, M. S., & Landy, M. S. (2002). Combining sensory information: mandatory fusion within, but not between, senses. *Science*, *298*(5598), 1627–1630.

Hillis, J. M., Watt, S. J., Landy, M. S., & Banks, M. S. (2004). Slant from texture and disparity cues: Optimal cue combination. *Journal of Vision*, *4*(12), 967–992.

Hwang, E. J., Donchin, O., Smith, M. A., & Shadmehr, R. (2003). A gain-field encoding of limb position and velocity in the internal model of arm dynamics. *PLoS Biology*, *1*(2), 209–220.

Johnson, S.G. (n.d.). The NLopt nonlinear-optimization package. <http://ab-initio.mit.edu/nlopt>.

Knill, D. C. (2005). Reaching for visual cues to depth: The brain combines depth cues differently for motor control and perception. *Journal of Vision*, *5*(2), 103–115.

Kopiske, K. K., Cesanek, E., Campagnoli, C., & Domini, F. (2017). Adaptation effects in grasping the Müller-Lyer illusion. *Vision Research*, *136*, 21–31.

Norman, J. F., Lappin, J. S., & Norman, H. F. (2000). The perception of length on curved and flat surfaces. *Perception & Psychophysics*, *62*(6), 1133–1145.

Norman, J. F., Todd, J. T., Perotti, V. J., & Tittle, J. S. (1996). The visual perception of three-dimensional length. *Journal of Experimental Psychology: Human Perception and Performance*, *22*(1), 173–186.

Pouget, A. & Snyder, L. H. (2000). Computational approaches to sensorimotor transformations. *Nature Neuroscience*, *3*(11s), 1192–1198.

Powell, M. J. D. (1998). Direct search algorithms for optimization calculations. *Acta Numerica*, *7*, 287–336.

R Core Team. (2014). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. <http://www.R-project.org/>.

Thoroughman, K. A. & Shadmehr, R. (2000). Learning of action through adaptive combination of motor primitives. *Nature*, *407*(6805), 742–747.

Todd, J. T. (2004). The visual perception of 3D shape. *Trends in Cognitive Sciences*, *8*(3), 115–121.

Todd, J. T., Oomes, A. H., Koenderink, J. J., & Kappers, A. M. (2004). The perception of doubly curved surfaces from anisotropic textures. *Psychological Science*, *15*(1), 40–46.