Preexposure disrupts learning of location-contingent perceptual biases for ambiguous stimuli

Loes C. J. van Dam Max Planck Institute for Biological Cybernetics, Tübingen, Germany

Marc O. Ernst Max Planck Institute for Biological Cybernetics, Tübingen, Germany

The perception of a bistable stimulus as one or the other interpretation can be biased by prior presentations of that stimulus. Such learning effects have been found to be long lasting even after small amounts of training. The effectiveness of training may be influenced by preexposure to the ambiguous stimulus. Here we investigate the role of preexposure for learning a position-dependent perceptual bias. We used rotating Necker Cubes as the bistable stimuli, which were presented at two locations: above or below fixation. On training trials, additional depth cues disambiguated the rotation direction contingent on the location. On test trials, the rotating cube was presented without disambiguation cues. Without preexposure to the ambiguous stimulus, subjects learned to perceive the cube to be rotating in the trained direction for both locations. However, subjects that were preexposed to the ambiguous stimulus did not learn the trained percept–location contingency, even though the preexposure was very short compared to the subsequent training. Preexposure to the disambiguated stimulus did not interfere with learning. This indicates a fundamental difference between ambiguous test and disambiguated training trials for learning a perceptual bias. In short, small variations in paradigm can have huge effects for the learning of perceptual biases for ambiguous stimuli.

Keywords: perceptual bistability, learning, perceptual memory, learning mechanisms

Citation: van Dam, L. C. J., & Ernst, M. O. (2010). Preexposure disrupts learning of location-contingent perceptual biases for ambiguous stimuli. *Journal of Vision, 10*(8):15, 1–17, http://www.journalofvision.org/content/10/8/15, doi:10.1167/10.8.15.

Introduction

Human perception and behavior largely depends on prior knowledge about relationships between the world around us and the sensory inputs resulting from it. Most of this prior knowledge has been learned and updated since our infancy, but even in our adulthood our perceptual system is continuously updating our knowledge (learning) as we encounter new correlations. Learning these new correlations can lead to changes so that the same sensory input is perceived differently. Such learning effects appear to occur on several different timescales. For instance, depending on the statistics of the environment, adaptation to a new visuomotor offset only needs a few trials of exposure (e.g., Burge, Ernst, & Banks, 2008) and thus can occur in about a minute. A longer exposure time, in the order of 1 h, was needed for humans to learn to integrate previously uncorrelated signals from vision and touch like the luminance (visual signal) and stiffness (haptic signal) of an object (Ernst, 2007). In addition, Adams, Graf, and Ernst (2004) trained subjects for about 1.5 h to induce a shift in the "light-from-above" prior, an example of a previously learned relationship between the environment and the sensory input. However, to completely reverse a previously learned relationship between the size of an object and its estimated weight when compared to similar objects (thus leading also to a reversal of the size–weight illusion), Flanagan, Bittner, and Johansson (2008) needed to train their subjects for at least 11 days. Here we show just how vulnerable learning of a new correlation can be when brief preexposure to the testing stimulus interferes with training on a longer timescale.

To do this, we used a perceptually bistable stimulus for which several different learning timescales are known to be involved. Perceptual bistability can occur when one and the same stimulus can be interpreted in two different ways. A good example is the well-known Necker Cube (Necker, 1832), which can be interpreted as if seen from above or as if seen from below. At any given time, only one of the two possible interpretations is perceived, but, if the figure is viewed continuously over a longer period of time, the percept will be experienced to alternate between the two interpretations. Such stimuli have been of great interest to psychologists and neuroscientists because the visual input alone contains insufficient information to drive the percept to one interpretation or the other and therefore the perceptual conflict must be entirely resolved within the perceptual system.

It is known for quite some time that one of the factors involved in resolving the ambiguous input for a bistable stimulus are in fact prior presentations of that same stimulus. These influences usually manifest themselves as short-term interactions, like negative aftereffects when

doi: 10.1167/10.8.15 Received February 1, 2010; published July 21, 2010 ISSN 1534-7362 © ARVO

subjects are presented with a disambiguated form of the stimulus for at least a few seconds immediately prior to being presented the ambiguous stimulus (e.g., for various types of perceptually bistable stimuli: Brascamp, Knapen, Kanai, van Ee, & van den Berg, 2007; Hochberg, 1950; Kanai & Verstraten, 2005; Long, Toppino, & Mondin, 1992; Nawrot & Blake, 1989; Virsu, 1975; Wolfe, 1984). In addition, short-term positive priming effects are known to occur when a disambiguated form of the stimulus is either presented very briefly, for just a few milliseconds, before presenting the ambiguous stimulus or when there is a pause of a few seconds between presenting the disambiguated priming stimulus and the ambiguous stimulus itself (e.g., for various types of bistable stimuli: Brascamp et al., 2007; Chong & Blake, 2006; Kanai & Verstraten, 2005; Long et al., 1992; Mitchell, Stoner, & Reynolds, 2004). Positive priming also occurs when viewing the ambiguous stimulus intermittently, with interstimulus intervals of about 1 s or more, in which case the percept tends to be the same across the separate stimulus presentations (e.g., Brascamp, Pearson, Blake, & van den Berg, 2009; Leopold, Wilke, Maier, & Logothetis, 2002; Maier, Wilke, Logothetis, & Leopold, 2003; Orbach, Ehrlich, & Heath, 1963). The latter effect is often called perceptual stabilization although it has been shown that perceptual alternations do still occur in this case, only at increasingly larger intervals with increasing interstimulus intervals (Brascamp et al., 2009). Evidence has also been found that the interactions across several stimulus intervals can act on multiple independent timescales on the order of seconds to even minutes (Brascamp et al., 2008; Pastukhov & Braun, 2008). For example, Brascamp et al. (2008) intermixed intermittent presentation with periods of continuous viewing and found that the percept during successive intermittent presentation periods also tended to be the same, regardless of the last percept seen during the intermediate continuous viewing period. In order for this longer term stabilization to occur, the continuous viewing period could take up to approximately a minute. Taken together, the effects described so far span the range of prior experience timescales from milliseconds to the order of several minutes.

Interestingly, coming from a somewhat different direction, it has also been shown that learned biases for bistable figures can last for days, and that reversing them requires extensive retraining (Haijiang, Saunders, Stone, & Backus, 2006). Haijiang et al. (2006) used a rotating version of the Necker Cube, in which case the cube can be perceived as rotating to the left (i.e., the front face moves leftward) or rotating to the right. During a training phase, they added depth cues in order to disambiguate the bistable Necker Cube. Using such disambiguated stimuli, they created a novel contingency by presenting subjects with leftward motion on one location (e.g., above fixation) and the rightward motion on another location (below fixation). In test trials in which the stimulus was presented without the disambiguation cues, it was found that subjects had learned the contingency between stimulus location and rotation direction. Interestingly, when subjects came back the next day they still perceived the same location–rotation direction contingency and to reverse the contingency needed extensive retraining. In short, they showed that a relatively brief exposure to an entirely new contingency can result in persistent and relatively long-lasting perceptual effects.

Haijiang et al. (2006) attributed these effects to a form of associative learning: Pavlovian cue recruitment and, consistent with the literature on Pavlovian learning, in their paradigm, the training started immediately. However, in most of the perceptual learning literature[1] the methods commonly include a pretest prior to training and a posttest after training. Such pretest–training–posttest paradigms are often used to obtain a more quantifiable measure on how the training relates to the change in perception. To obtain such a quantifiable measure for learning a bias for a bistable figure, we revisited the paradigm of Haijiang et al. (2006) and included pre- and posttests to get a true measure of the change in perception brought about by the training. We realized, based on the prior presentation effects described above, that there was an off chance that the pretest might somewhat interfere with the training. To our surprise, however, as will be clear from the results, we found a much stronger pretest interference than anticipated and this interference was even to the extent that subjects did not or hardly learned the trained percept–location contingency at all. Here, we will describe the experiments in which we found this strange effect and tested certain hypotheses as to what aspect of the pretest may have been causing the interference with subsequent training.

Another interesting question is to what extent the learned biases are specific to the training conditions used. More specifically, if the rotating Necker Cube is used to learn the biases, are these biases only applied when presented with a Necker Cube? Or do these learned biases generalize to other objects, like ambiguous spheres? This question strongly bears on which level the learning occurs, i.e., if high level object recognition is involved or if the learning occurs more on a low level. If object recognition is involved, the cube somehow has to be recognized as such before it can influence which of the two possible percepts is going to dominate. In other words, learning should be specific for cubes alone and no transfer of the learned biases to spheres is to be expected in this case. However, if the learning occurs more on a low level, on features that are common across different objects, the learned biases should also be more commonly applied. In this case, we would expect the biases learned explicitly for cubes to fully transfer also to spheres that were not included in the training. To investigate this question, in the experiments described here, the pre- and posttests include both cube and sphere trials to test for such transfer effects of learning.

Experiment 1

In Experiment 1, we compared two different paradigms to investigate perceptual learning of a location-dependent perceptual bias: a training–posttest paradigm versus a pretest–training–posttest paradigm. The first is similar to the paradigm of Haijiang et al. (2006); except for the posttest, and, therefore, serves as a control whether their results can be replicated using our setup. The reasoning for the second paradigm is that usually in learning experiments posttest results are compared to pretest results to better quantify the effect of training and to correct for a bias that may have preexisted even without training. For each of these two paradigms, a separate group of observers participated and since for both groups the training schema itself is exactly the same we initially did not expect to find a very big difference in learning between the two groups.

Methods

Apparatus

Stimuli were displayed stereoscopically using a haploscope. The half-images were displayed on separate CRT monitors (60 Hz) and viewed via mirrors, one for each eye. The resulting viewing distance was 50.5 cm. The display area on the monitors was limited to approximately 20 by 20 degrees by physical apertures in front of the monitors. A chin and head rest restricted the subjects' head movements. Stimuli were created using OpenGl and the C programming language.

Stimuli and task

At the beginning of each trial, subjects were presented with a fixation square (4.5×4.5 arcmin) at the center of the display area. Subjects initiated stimulus onset by button press. The stimuli were 2-s animations of a cube, of which the sides were covered with small squares, rotating around a vertical axis (67 deg/s). For a stationary example of the stimulus, see Figure 1A. The edges of the cube subtended 2.4 deg and a single square on the cube's surface subtended 8.8 arcmin. Square density on the cube surface was approximately 8.9 squares/deg^2. The cube was tilted with respect to the rotation axis such that one of the cube's diagonals was perpendicular to the rotation axis. Cube starting rotation angle was randomized across trials. Orthogonal projection was used to display the cube in order to prevent perspective cues from influencing the percept.

The cube was displayed 4.3 deg either above or below the fixation square. On the other side of the fixation square, also at a vertical distance of 4.3 deg, a single white dot (diameter of 10 arcmin) moved at 53 arcmin/s either leftward or rightward (see Figure 1B). The motion direction (left or right) of the separate dot was randomized across trials. The subject's task was to indicate by pressing one of two buttons whether this white dot moved in the same or in the opposite direction as the cube's front face. This task was chosen to prevent a direct link between the perceived rotation direction of the cube and the required response and to make subjects less aware that the perceived rotation direction of the cube was the actual parameter of interest.

Without additional cues, the rotation direction of the cube is ambiguous and can be perceived as rotating either leftward or rightward. On training trials, the cube's rotation direction was disambiguated by presenting a cylinder (radius of 34 arcmin; length of 9.1 deg) going through the center of the cube along the rotation axis, and by adding veridical stereoscopic depth information (binocular disparity). The cylinder occluded part of the backside of the cube and together with the stereoscopic depth cue this results in strong disambiguation of the rotation direction. For each subject, the rotation direction on training trials was always the same for each location but different across locations. That is, trained contingencies could either be leftward when presented above fixation and rightward below or, vice versa leftward rotation when presented below fixation and rightward above. This was done in order to be able to distinguish learning position as a cue from developing a general perceptual bias that does not depend on location. The two possible rotation direction/location contingencies for the training were counterbalanced across subjects.

On test trials, the cube was shown without occlusion and stereoscopic cues and the rotation direction of the cube was ambiguous. The perceived rotation direction of the cube on test trials will be informative whether or not the subjects learned the rotation direction vs. cube location contingencies. To completely remove any stereoscopic cues, the cube was presented monocularly, only to the left eye, during test trials.

After each trial, subjects received feedback about their response in the form of a smiling or a sad cartoon face for a correct and an incorrect response, respectively. For an incorrect response, subjects received an additional penalty in the sense that they had to wait 6 s before they could start the next trial. For correct responses, the feedback was shown for 1 s. The feedback, on the ambiguous test trials, was always positive, i.e., indicating a correct answer.

Training procedure

For a schematic of the training procedure, see Figure 1C. During the training phase, the trials were divided into blocks of 12 trials. The first 10 trials were always training

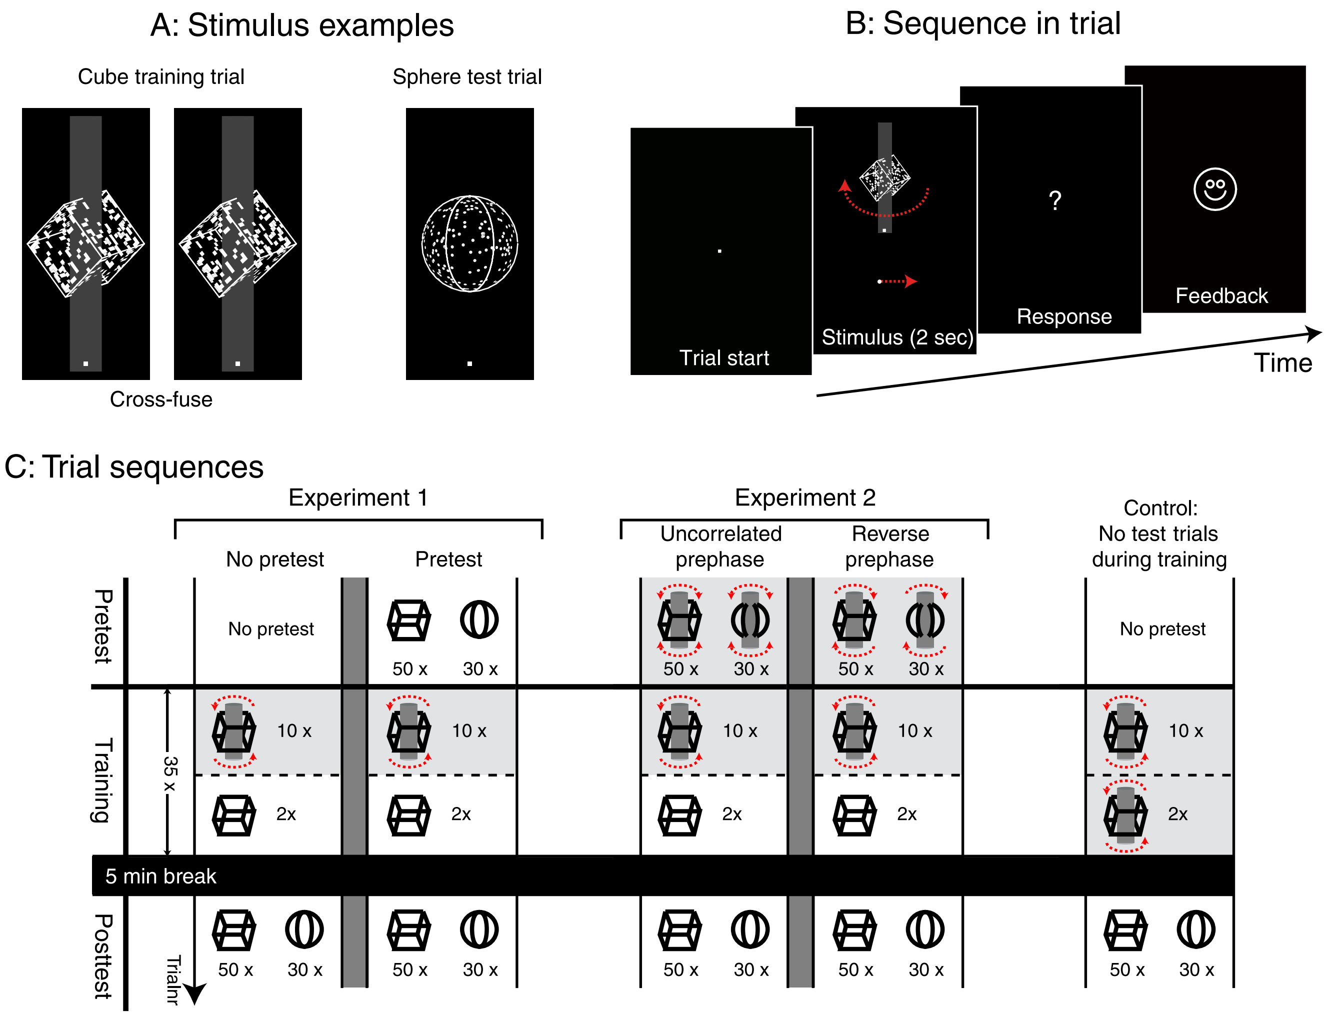

Figure 1. (A) Examples of the used stimuli. On the left, a stereogram is shown, depicting the stimulus for a training trial with added stereo- and occlusion cues (cross-fuse for 3D representation). For cube test trials (not shown), the gray bar was removed and the stimulus was only presented to the left eye. On the right, the sphere stimulus, as used in the pre- and posttests, is depicted. (B) Schema for one trial. Subjects initiated stimulus onset and a rotating cube and a moving dot were displayed. Subjects responded whether the dot moved in the same or opposite direction as the front face of the cube, upon which they received feedback on the correctness of their response in the form of a cartoon face. On test trials, subjects always received positive feedback. (C) Training schedule for the separate experimental conditions. Gray blocks represent training trials; white blocks represent test trials. The red arrows indicate the trained contingencies (e.g., during training above fixation is trained for leftward rotation to be perceived and below fixation is trained for rightward rotation to be perceived). Double-sided arrows indicate no correlation with the stimulus location. The separate conditions in Experiments 1 and 2 differ only in the trials shown before the actual training starts.

trials and in 5 of these the cube was presented above and in 5 below fixation in random order. The last two trials in each training block were 2 test trials, one for each location. Subjects performed 35 of these training blocks (420 trials). Directly after the training phase, the word "break" appeared on the screen and the subjects were required to take a 5-min break before continuing to the posttest phase. The break was inserted to prevent any direct short-term aftereffects from contributing to the posttest results. Subjects were told to do what ever they want during the break except to close their eyes for the whole period of the break. After the break, the fixation square appeared again and subjects could continue the experiment.

The posttest consisted of test trials only. For 50 of these posttest trials, a cube was shown as before and cube presentations were equally divided across the above and below fixation locations of the cube. To test for transfer effects to other objects, there were 30 additional posttest trials in which a sphere (see Figure 1A) appeared instead of the cube, again equally divided across the above and below fixation locations. The presented sphere had a radius of 1.7 deg and was covered in dots with a radius of 3.4 arcmin. Dot density on the sphere surface was

5.5 dots/(deg visual angle)^2. Six additional, equally spaced, semicircular lines from the top to the bottom of the sphere added contour information. All cube and sphere trials in the posttest were presented in random order.

In Experiment 1, the subjects were divided in two groups. For each of these groups, the training, break, and posttest schemas were exactly the same and the only difference between the groups is what happened before the training started. One group followed the procedure exactly as described above, i.e., directly start the training phase and afterward the posttest ("no pretest" group). The second group of subjects ("pretest" group) was preexposed to the ambiguous stimulus in a pretest before the training phase started. The procedure for the pretest was exactly the same as for the posttest (50 ambiguous cube trials and 30 sphere trials in random order and equally divided across the two locations).

Participants

Participants were recruited from a general population and for each condition 10 observers participated. Observers had normal or corrected-to-normal vision and stereo-acuity was assessed using the Stereo Fly Test (Stereo Optical, Chicago IL, USA). Subjects without good stereo-vision were excluded from participation.

Results

The task is fairly easy to perform and for training trials subjects can be expected to have a high accuracy. Most subjects only made a few mistakes (average error rate was 2.7%, which is equivalent to about 10 trials out of the 350 training trials) and reported afterward that they pressed the wrong button by accident on these cases. However, there were also a few subjects who were not very accurate on training trials and whose proportion of correct responses was, although still above chance, below 87.5%. For these subjects, it is not clear how to interpret the results for the test trials. Thus, subjects with a proportion of correct responses below 87.5% on training trials were excluded from the analysis and replaced by a new participant.

Figure 2A displays the results for test trials for the two different subject groups. The *x*-axis represents the phase within the session, i.e., pretest (if included), training, and posttest. The *y*-axis represents the percentage of all test trials for each phase on which the perceived rotation direction was in agreement with the trained direction; 100% means that subjects for both locations always saw the cube rotating in the trained direction, indicating a complete learning effect, and 0% means that subjects always saw the cube rotating in the opposite direction than the direction that was presented on training trials. To track learning during the training phase, we divided the training phase in 7 chronological blocks, each consisting of 50 training and 10 test trials. For the training phase, the proportion of test trials within such a block, for which the perceived rotation direction of the cube was in the trained direction, is displayed against the cumulative number of training trials in the session.

The different colors depict the different subject groups. Blue symbols and lines show the results for the no-pretest group and red symbols and lines show the results for the pretest group. The different symbols stand for the different objects displayed on different test trials. Squares display the results for cube trials and circles the results for sphere trials. The sphere trials only occurred in the pre- and posttests of the session. The error bars show 95% confidence intervals for binomially distributed data across all test trials for each phase (note that this means that these estimates are based on several test trials per subject rather than on a single representative value per subject). The dashed line indicates chance level.

Figure 2A shows the results across subjects. Figures 2B and 2C show the individual subject results for the no-pretest and pretest groups, respectively. Each line in Figures 2B and 2C represents one subject and only results for cube trials are shown (for sphere trials, posttest results are shown in Figure 4).

The results for the no-pretest group show clear learning effects both during training and in the posttest: the perceived rotation direction conform with the trained direction on 80% of the cube test trials or more on average. For 7 out of the 10 individual subjects, the proportion perceived as trained are even very close to 100%. In addition, note that the training effect appears to be virtually immediate, i.e., even after as few as 50 training trials the proportion of test trials seen as trained is already about 80% and significantly above chance. This could mean that only a few trials are necessary to learn a location contingent bias, but note that during training also short-term priming effects (e.g., Kanai, Knapen, van Ee, & Verstraten, 2007; Kanai & Verstraten, 2005; Long et al., 1992) additionally cause the perceived rotation direction for test trials to be in the trained direction. Thus from the no-pretest group results during training alone it is unclear how many trials are sufficient to ensure a long-term retention of the bias. That the bias is really learned becomes clear only from the posttest results. Short-term priming effects should have decayed during the 5-min break and therefore cannot be responsible for the significant learning effect there. Furthermore, the results in the posttest for cube and sphere test trials are quite similar and also for the sphere the proportion of rotation direction seen as trained is significantly different from chance ($p = 0.006$ in a signed rank test across proportions per subject). This indicates that the learning effect generalizes to a different object class.

Such learning effects, as reported here for the no-pretest group, have previously been demonstrated by Haijiang et al. (2006), whose experimental design differed in the

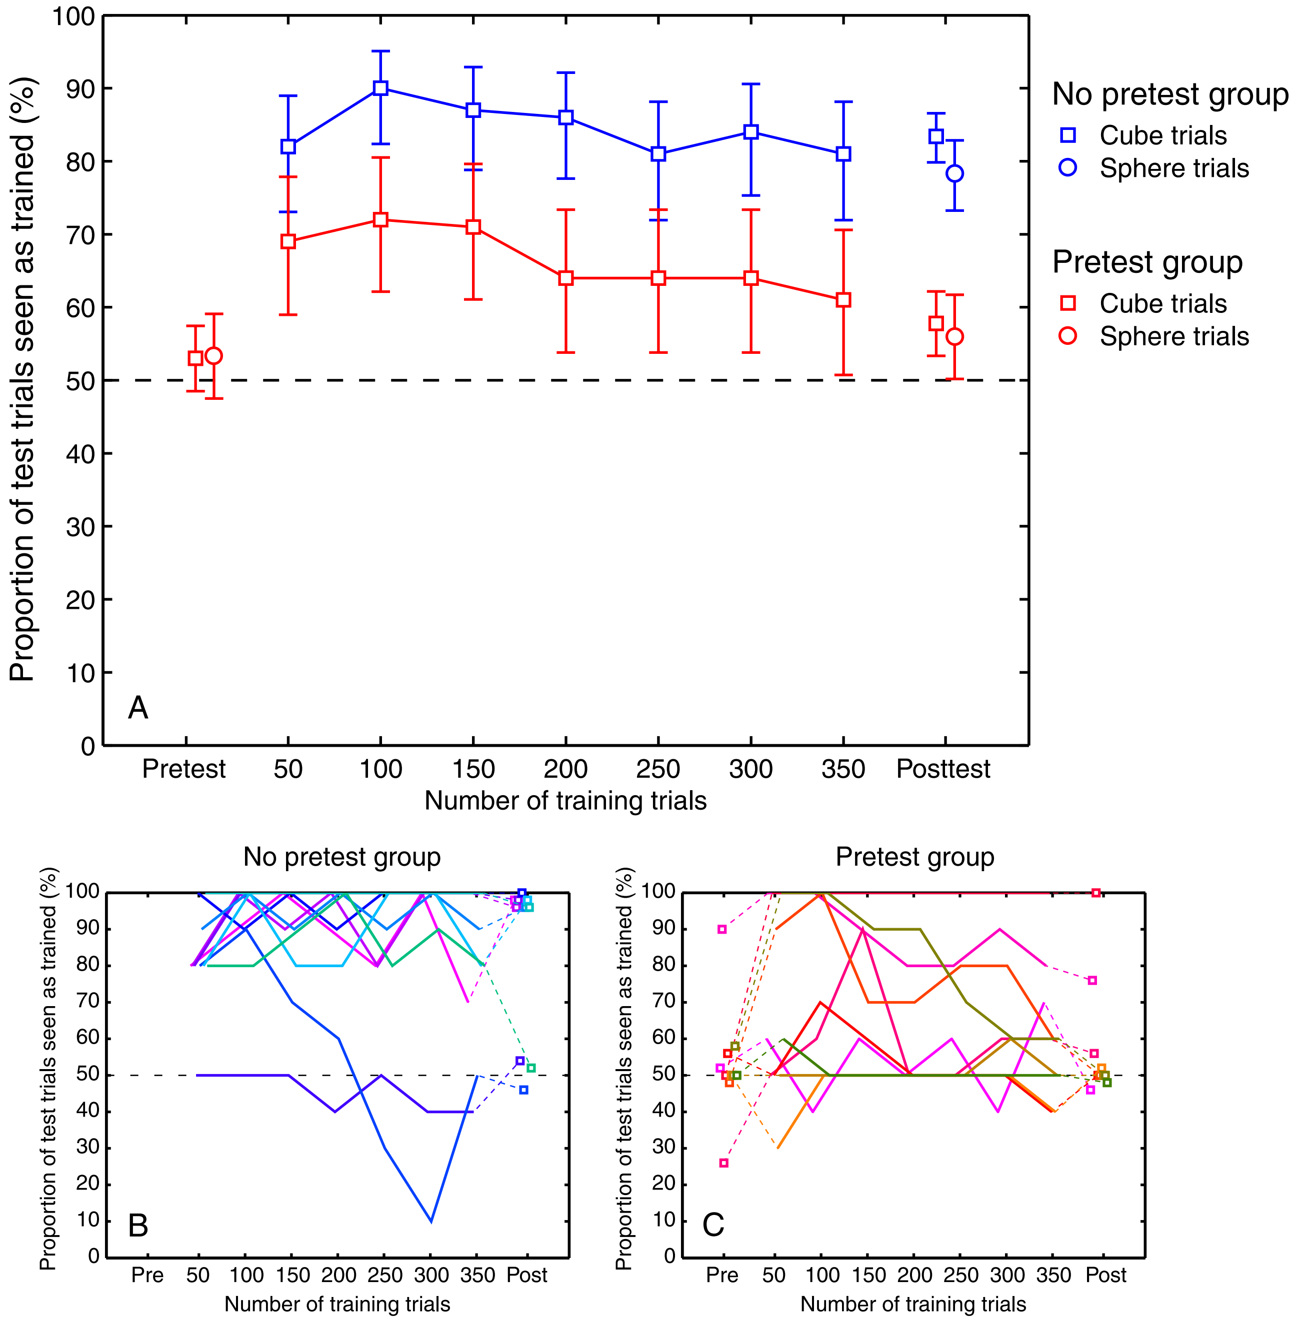

Figure 2. Results for Experiment 1. (A) The percentage trained direction perceived on test trials is shown versus the phase in the experimental session (pretest, training phase in blocks of 50 training trials (10 test trials), and posttest). The two colors represent the subject groups, blue: no-pretest group; red: pretest group. Squares show results for cube trials and circles for sphere trials. Error bars represent 95% confidence intervals. The dashed line represents chance level. For the pretest group, training appears to be less effective compared to the no-pretest group. (B, C) Each line represents the results for an individual subject for the no-pretest and pretest groups, respectively.

size of the cubes (14 deg versus 2.4 deg in our experiment), distance between fixation and cube locations (9.4 deg vs. 4.3 deg), stimulus duration (8.33 s vs. 2 s), the method of stereoscopic display (red–green anaglyphs vs. haploscope), and viewing distance (200 cm vs. 50.5 cm). So the results for the no-pretest group replicate the results of Haijiang et al. (2006) and demonstrate that the previously reported learning effect is relatively robust across small variations in experimental conditions and methods. Note though that our results appear to be slightly more variable across subjects than was the case in Haijiang et al.'s (2006) study, i.e., in our experiment there were 3 subjects that did not learn the location–rotation direction contingency, whereas in the study of Haijiang et al. (2006) all the subjects learned. In principle, all of the small variations in experimental condition mentioned above could have contributed to this slight difference in results. However, for the short-term perceptual stabilization effect it has been shown that for two separate retinotopic locations the short-term memory traces become more and more independent with increasing distance between the locations (Knapen, Brascamp, Adams, & Graf, 2009). If the long-term perceptual memory studied here operates on similar retinotopic principles, the increase in variability across subjects in our study could partially be due to the smaller distance between the two training locations.

As noted before, we did not expect to find a big difference between the pretest group and the no-pretest group since in principle the training scheme is the same for both groups. Instead, we assumed that the pretest would not or hardly interfere with the training. The advantage of including the pretest would have been that posttest results could have been compared with an initial bias, were it not for the fact that we were very much mistaken in the above assumption. As can be seen in the results, the difference between the pretest group and the no-pretest group is quite drastic and subjects in the pretest group hardly appear to have learned (only 2 subjects in the pretest group had a perceptual bias during the posttest that was consistent with the trained direction and a signed rank test across subjects' individual proportions did not show a significant difference from chance; $p = 0.13$).

To see in what way the pretest is interfering with the later training, it is useful to look in more detail at individual subject data for the two training locations separately (see Figure 3). Here it is important to note that a proportion of 50% test trials seen as trained for an individual subject in Figure 2C does not necessarily mean that the average percept for each of the two separate locations was completely unbiased. The 50% proportion could also be a result of subjects having a bias completely in the trained direction for one of the two locations (100%) and for the other location a bias completely opposite to training (0%). In other words, the 50% results for subjects in the pretest group could mean that subjects did have strong perceptual biases for each separate location but that these biases just did not correspond to the trained contingency. We wondered whether this was the case and if any biases that subjects might have obtained during the pretest could be responsible for the interference with the training afterward. Figure 3 shows the individual subject data separately for the two training locations. In Figure 3, each column represents a single subject (the order in which the results for individual subjects are displayed is such to more easily detect commonalities across the individual subject results). The top row shows the results for the above fixation location and the bottom row shows the results for the below fixation location. Results are displayed as proportion of test trials seen as trained separately for test trials in the pretest (light gray bars), during training (intermediate gray bars), and in the posttest (dark gray bars). What can be seen is that indeed the individual subject pretest biases for each separate location (light gray bars) were relatively strong: only one subject had an initial bias that was close to chance for one location (subject 2, top graph), all the other separate biases per location were significantly below or above chance and most even close to either 0% or 100%. Comparing the biases observed in the posttest to the ones observed in the pretest, we see that most subjects (7 out of 10) stick to the initial biases they had observed in the pretest for both locations. So for 7 out 10 subjects the pretest biases seem to dominate in the posttest. In contrast, only for two subjects (subjects 9 and 10) the posttest results are more or less consistent with the trained directions for both locations but note that for one of these subjects that already happened to be the case in the pretest.

In short, what can be seen in Figure 3 is that particularly during the posttest the subjects' perceptual bias corresponds more to their initial perceptual bias, observed during the pretest, than the contingencies they are exposed to during training.[2] This indicates that perceptual biases

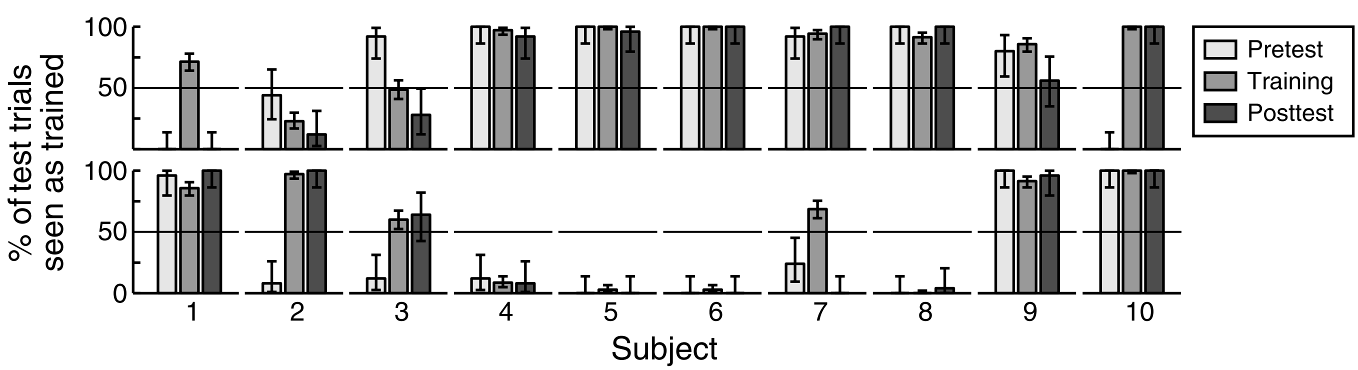

Figure 3. Individual subject results separated for the two training locations for the pretest group only. Each column represents a single subject (the order in which the results are displayed is such to more easily detect commonalities between individual subjects). The top row shows the results for the above fixation location and the bottom row shows the results for below fixation. Results are displayed as proportion of test trials seen as trained for test trials in the pretest (light gray bars), during training (intermediate gray bars), and in the posttest (dark gray bars). For individual subjects, perceived rotation directions for test trials in the posttest appear to correspond more to the initial biases observed in the pretest rather than to the trained biases.

may be fixed very early on and later training is not as effective as the initially obtained bias. To examine if, and if so how, the initial perceptual history can influence later learning mechanisms, we conducted Experiment 2.

Experiment 2

From Experiment 1, it is evident that presenting an observer with a pretest interferes with later perceptual learning in this paradigm. The perceived rotation direction during the pretest was not experimentally controlled in Experiment 1 and therefore accidental correlations may have existed and already have been learned at this stage. Therefore, in Experiment 2, we tested whether any such previous correlation between location and perceived rotation direction observed during the pretest may have been the cause for the interference with learning a different correlation during training. We tested two hypotheses concerning possible accidental correlations.

The first hypothesis is that the results in Experiment 1 could be due to a "learned irrelevance." During the pretest where the percept for the rotating cube was ambiguous, and therefore not controlled experimentally, subjects may have learned that the location cue is irrelevant and therefore fail to learn the location cue when it does become relevant during the training phase. This hypothesis closely corresponds to Pavlovian latent inhibition, where preexposure to the conditioned stimulus results in slower Pavlovian learning. For example, the analogous experiment in the animal learning literature is that if you expose a rat to a tone that is uncorrelated with food, it will take longer for the rat to learn, subsequently, that a tone predicts reward, as compared to not presenting the uninformative tone before training. To test this "learned irrelevance" hypothesis in our paradigm, we simulated a pretest, i.e., instead of showing ambiguous cubes in the pretest we showed disambiguated cubes (with the same occlusion and stereo-cues as described for Experiment 1) and in this way controlled the observed perceptual history. During this simulated pretest, 50% of the cubes and spheres were presented rotating leftward and 50% were presented rotating rightward on both locations and in random order, thus ensuring that during this pretraining phase there was no correlation whatsoever between location and perceived rotation direction. After this initial phase, the remaining training, break, and posttest procedures were the same as in Experiment 1. If the initial phase of the experiment is important for the system to determine whether a certain source of information is relevant, in this case the location cue, then no learning should occur in this condition since there is no correlation between perception and location at the beginning of the experiment. We will refer to this condition as the "uncorrelated prephase" condition.

The second hypothesis concerns the results in Figure 3, which indicates that an initial bias, obtained during the pretest, may be stronger than the bias that is subsequently being trained on. To test this hypothesis again, a pretraining phase with perceptually controlled disambiguated cubes and spheres was conducted before training started. During this prephase, the rotation direction for each location was 100% opposite to the subsequently trained direction. Training, break, and posttest procedures were again the same as in Experiment 1. This condition will be referred to as the "reverse prephase" condition and if the initial perceptual history is stronger than the later training trials, training is expected to have no or at least less of an effect on posttest results. So again no learning due to the training phase is expected.

Two groups of 10 observers each participated in this experiment, none of whom had participated in the previous conditions. All observers had normal or corrected-to-normal vision and good stereo-vision.

Results of Experiment 2

The posttest results for Experiment 2 together with the previous results of Experiment 1 are displayed in Figure 4. From left to right, the graphs in Figure 4 show the results for Experiment 1 (no-pretest vs. pretest condition), Experiment 2 with prephases consisting of disambiguated trials (either uncorrelated or 100% reversed from training direction) and a further control condition, which will be described below. The top row displays the proportion of test trials for which the perceived rotation direction was in the same direction as trained across subjects. Green bars represent the results for cube trials and brown bars represent the results for sphere trials. Error bars depict 95% confidence intervals. The black dashed line represents chance level.

The middle row of the graphs displays the posttest results for individual subjects in each group. For each subject, the proportion of test trials seen as trained is shown for cube trials (green squares) as well as for sphere trials (brown circles). For each subject, cube and sphere results are connected by a gray line. The black dashed line indicates chance level and the green and brown dashed lines represent the one-sided significance boundaries for individual subject cube and sphere results, respectively. The bottom row shows the proportion of subjects that individually showed a significant learning effect. Subjects were determined to have a significant learning effect if their posttest results for cube trials (green bars) were significantly above 50% chance level and the perceptual biases for both locations were in the trained direction. Brown bars show the results of the same analysis for sphere trials. Error bars show the 95% confidence intervals for the proportions. The black dashed line depicts chance level, which in this case is 25% since the bias at both separate locations needed

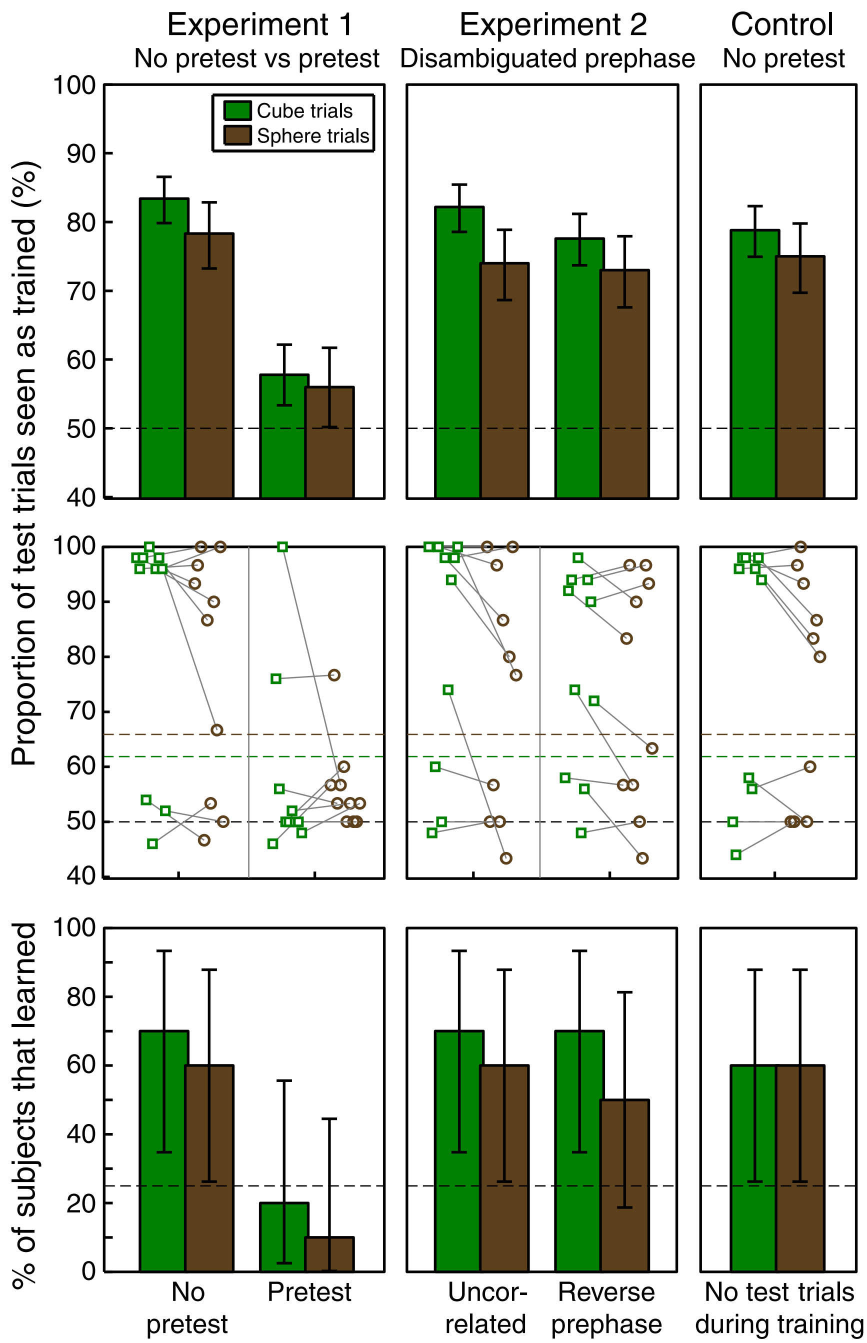

Figure 4. Posttest results for Experiment 1, Experiment 2, and the control condition. From left to right, the graphs show the results for Experiment 1 (no-pretest vs. pretest condition), Experiment 2 with prephases consisting of disambiguated trials (either uncorrelated or reversed training direction), and the no-test-trials-during-training condition. (Top) The proportion of test trials for which the perceived rotation direction was in the same direction as trained for cube trials (green bars) and sphere trials (brown bars). Error bars represent 95% confidence intervals; the dashed line represents chance level. (Middle) Posttest results for individual subjects in each group. Green squares indicate results for cube trials, brown circles for sphere trials. Cube and sphere results for individual subjects are connected by gray lines. The black dashed line represents chance level and the green and brown dashed lines represent the 0.05 one-sided significance boundaries for cube trials and sphere trials, respectively. (Bottom) The proportion of subjects that individually showed a significant learning effect for cube trials (green bars) and sphere trials (brown bars). Error bars represent the 95% confidence intervals; the dashed line represents chance level. For cube trials, all groups except the pretest group in Experiment 1 show very strong learning effects.

to correspond to the trained direction in order for the subject to be counted as having learned.

What can be seen in these graphs is that all groups except the pretest group from Experiment 1 show very strong learning effects. This is somewhat surprising especially if the result for the pretest condition of Experiment 1 is compared with the results from Experiment 2. As noted before, the accidental perceptual biases observed during the pretest in Experiment 1 were relatively strong and for most subjects did not correspond to the subsequently trained directions, at least not for both locations. This means that subjects were initially exposed to a different correlation than that which they were subsequently trained on. Therefore, in Experiment 2 we especially expected to find similar results for the reverse prephase condition in which subjects also are initially exposed to a different correlation, this time experimentally controlled. Instead, in Experiment 2 we find a strong learning effect in both the disambiguated prephase conditions when compared to the pretest condition of Experiment 1 for which the only difference is that ambiguous test trials were used in the pretest. This suggests a fundamental difference between ambiguous test and disambiguated training trials in terms of the contribution to the learning process. A bias obtained from viewing ambiguous test trials appears to be much stronger than a bias obtained from viewing disambiguated training trials.

Control condition

The difference in results between the pretest condition in Experiment 1 and the disambiguated prephase conditions of Experiment 2 suggests that there may be a strong difference between training trials and test trials in terms of learning the bias. Furthermore, in the pretest condition of Experiment 1 just 80 test trials were sufficient to seriously hamper any learning from 350 subsequent training trials, more than 4 times the same amount. Does this mean that any correlation learning is stronger for test trials than for training trials?

In our paradigm, as well as the paradigms of Backus and Haijiang (2007) and Haijiang et al. (2006) the training trials were always interleaved with test trials during the training phase. So from these earlier results it is not clear whether it is important for the training phase to include test trials besides the training trials and whether the learning effect hinges on these test trials rather than the training trials. In order to verify whether training trials by themselves can elicit the learning effect we repeated the no-pretest condition of Experiment 1, but this time the interleaved test trials during training, i.e., the 2 test trials after every 10 training trials (see Methods section and Figure 1C), were replaced with 2 additional training trials, one for each location. After training, the procedures for the break and posttest were the same as before. If learning hinges completely on the interleaved test trials during training, then no learning should occur in this case, whereas if training trials do contribute to the learning effect we should find similar results to the no-pretest condition of Experiment 1.

Again 10 observers who had normal or corrected-to-normal vision, good stereovision, and who had not participated in any of the previous conditions participated. The results for this control condition are also depicted in Figure 4, the rightmost graph in each row. The results show that learning occurs for most subjects in this condition as well, indicating that training trials do contribute to the learning effect.

Experiment 3

The interference effect of Experiment 1 raises another question: namely, is the learning process based on acquiring knowledge about a global correlation between location and rotation direction? This would be equivalent to learning the function $f(x)$, where f represents the percept dependent on the retinal location x based on the global history. Alternatively, the learning process could involve multiple independent, retinally localized learners, with each only taking into account the local history. So far, the paradigm used lead to the same predictions for these two possibilities. However, with the interference effect we can piece the two apart by pretesting on one location only and examining the effect of subsequent training on several locations. If the underlying mechanism learns the global correlation between location and perceived rotation direction, this one location pretest should still interfere with learning over all locations, since it interferes with learning $f(x)$. If instead the learning mechanism involves separate localized learners, the pretest should only interfere with learning at the same location at which it was performed and not on other locations.

Methods

The paradigm for Experiment 3 was similar to the pretest condition in Experiment 1. The only difference is that the pretest was performed on only one location, i.e., either above or below fixation. This pretest contained 25 cube test trials and 15 sphere trials, which is consistent with the pretest conditions in Experiment 1 for a single location. The pretest location was counterbalanced across subjects to be above or below fixation.

In the pretest, a bias was obtained and the rotation direction for the subsequent training trials was in the opposite direction as this bias on both locations. For example, if a subject showed a rightward bias during the pretest on the below fixation location, the subsequent training trials would show leftward rotation on both the above and below fixation locations. Note that we can only be certain about the initial bias on the location that the pretest was performed at. However, in the pretest of Experiment 1, 80% of the subjects exhibited the same bias in perceived rotation direction on both locations. Additional data from pilot experiments also suggest that subjects generally have the same directional bias on both the above and below fixation locations. Therefore, we can safely assume that the initial bias, if also measured on the non-pretested location, would be similar for both locations. Note that this is consistent with the study by Carter and Cavanagh (2007) who investigated existing location-dependent biases at several different retinal locations for a bistable stimulus and whose results indicate that although biases can differ across locations there usually is one percept that dominates most locations. Therefore, their

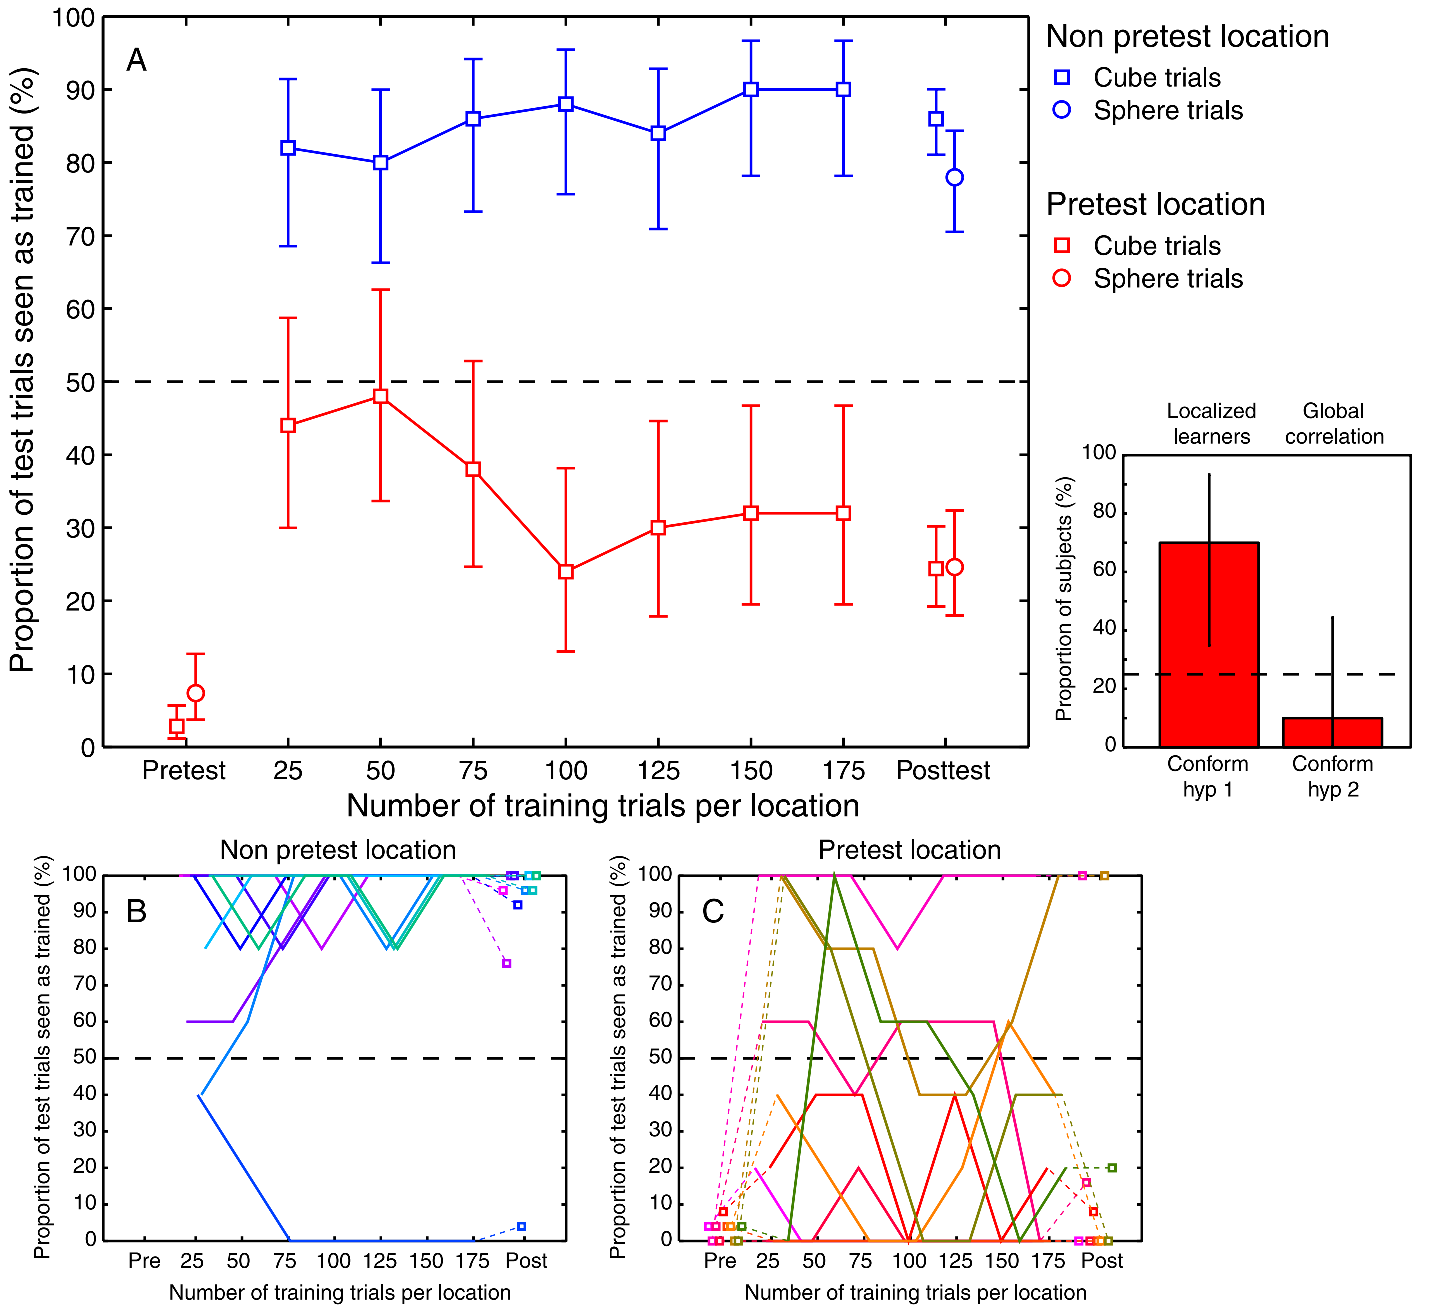

Figure 5. Similar to Figure 2 but now for Experiment 3. (A) Results for the non-pretest location (blue line and symbols) and the pretest location (red line and symbols) are shown separately. The bar chart inset shows the proportion of subjects for which the results conform with the separate localized learners hypothesis (left bar) or conform with a global correlation learning hypothesis (right bar). (B, C) Individual subject data for the non-pretest location and pretest location, respectively. Results show that the pretest only interferes at the location it was performed at, which is in agreement with a separate localized learners hypothesis.

results also indicate that the chance of the bias being the same at a subsample of two specific locations is relatively high.

The schemes for training and posttest were the same as for Experiment 1, except that, as noted above, the trained rotation directions for the above and below fixation locations were in the same direction rather than in opposite directions. As before, 10 observers who had normal or corrected-to-normal vision, good stereovision, and who had not participated in any of the previous conditions participated.

Results

Figure 5 shows the results for Experiment 3. In Figure 5A, the proportion of test trials seen as trained, across all subjects, is shown versus phase within a session, i.e., pretest, training blocks, and posttest. Note that on the *x*-axis we now plot the number of training trials per location since the results are shown separately for the non-pretest location (blue line and symbols) and the pretest location (red line and symbols). Cube trial results are depicted by squares, sphere trial results by circles. In Figures 5B and 5C, each line represents the results for cube trials for a single subject for the non-pretest location and the pretest location, respectively.

For the pretest itself, the results are very close to zero, both on average (Figure 5A) and for each individual subject (Figure 5C). This means that subjects show a very strong initial preference for seeing the cube rotate in a specific direction. Furthermore, for the non-pretest location the pretest hardly interferes, i.e., there is strong learning trend for this location (Figures 5A, blue line, and 5B). However, for the location where the pretest was performed (red line, Figures 5A and 5C), the training hardly seems to have any effect and in the posttest eight out of the ten subjects still have a bias in the same direction as in the pretest. Thus, the results show that the pretest on one location only interferes with training for that particular location and does not globally affect the learning.

The bar chart inset shows the extent to which these results are consistent with each of the two global learning or localized learners hypotheses. The left bar shows the proportion of subjects that in the posttest showed the result of interference on only the pretest location, which is according to the localized learners hypothesis. The right bar shows the proportion of subjects that showed interference on both training locations according to the prediction based on learning more global correlations between location and rotation direction. This bar chart clearly shows that the results favor the separate localized learners hypothesis. Note that the two bars in the graph sum up to 80% of the subjects. For the remaining two subjects, the training was effective on both locations.

Discussion

Previously, it has been shown, by exposing human subjects to new correlations with the world they operate in, that their perceptual biases and a priori knowledge are continuously being updated (e.g., Adams et al., 2004; Ernst, 2007; Flanagan et al., 2008). Recently, in one particular study it was shown that for a bistable stimulus, the rotating Necker Cube, a new correlation between perceived rotation direction and a location cue can be learned in a very short amount of time (Haijiang et al., 2006). The resulting location-dependent biases needed extensive retraining in order to reverse. Such strong and long-lasting effects could be due to subjects not having encountered the specific stimulus before training starts and therefore there being no previous bias that the new correlation has to compete against. In this study, we examined the role of preexposure to the ambiguous stimulus on learning a perceptual bias dependent on location. We found that preexposing subjects to the ambiguous stimulus in a pretest hugely interfered with the subsequent training. This interference effect was not due to having already obtained a perceptual history, as was shown by the results of Experiment 2. Rather, the combined results of Experiments 1 and 2 indicate a fundamental difference between training and test trials. Biases obtained using test trials may be stronger and more long-lived than ones obtained using training trials. Last but not least, the results for Experiment 3 show that the learning mechanism does not focus on a global correlation between location and perceived rotation direction but rather involves independent separately localized learners.

Intermittent presentation of bistable figures

In this study, we used an ambiguous stimulus, the rotating Necker Cube, to test for learning of location-contingent biases. The perception of such figures, when presented continuously, usually switches between two different states over time. However, with intermittent presentation, like in our pre- and posttests, this perceptual cycle tends to slow down (Brascamp et al., 2009; Leopold et al., 2002; Maier et al., 2003; Orbach et al., 1963). This means that observers usually have the same percept for a longer period of time when viewing the ambiguous figure intermittently compared to continuous viewing. This effect is often referred to as perceptual stabilization. Apart from the percept surviving interleaved blanks with intermittent viewing, there have also been reports about more long-lasting effects of stabilization (Brascamp et al., 2008; Pastukhov & Braun, 2008). Especially, Brascamp et al. (2008) showed that the percept seen during a previous stabilization period tends to also be observed in the next stabilization period. This occurs even when the two subsequent stabilization periods are separated by a period of time for which the stimulus is presented continuously and long enough for the observer to have experienced several perceptual reversals. Even when the last percept during this continuous viewing period is the opposite compared to the last winner during stabilization, the next percept to win when the stimulus is shown intermittently again is most likely the last stabilization winner and not the most recent percept. According to Brascamp et al. (2008), this effect takes about a minute of continuous viewing to decay completely.

This stabilization effect will also have influenced the stability of the percept in our pre- and posttests for which the ambiguous stimulus was viewed in an intermittent fashion. That is, the fact that in our pre- and posttest the subjects' biases are almost completely in one direction will, at least partially, be due to this stabilization effect. In addition, the stabilization could in this way be responsible for allowing a small initial perceptual bias to be reinforced during the pretest in Experiment 1. However, the stabilization effect cannot, by itself, explain the sign of these biases. For the pretest, there is nothing to direct the bias in a particular direction that is consistent across subjects. For the posttest, subjects were obliged to take a 5-min break before continuing with the posttest procedure. Although blank periods have been shown to be less effective for decay of priming effects than a neutral stimulus (Kanai et al., 2007), we expect that during this time any relatively short-term effect should have decayed and therefore is unlikely to be responsible for the results.

Furthermore, the short-term stabilization effect cannot possibly account for the differences in posttest results between our experimental groups. Any tentative explanation based on short-term adaption or priming effects would weigh recent trials more heavily than trials that happened at the very beginning of the experiment and therefore predict similar results in the posttest for all conditions given that the training conditions were the same. In that sense, especially the results for the pretest group in Experiment 1 present a curious puzzle when compared to the other experimental groups. For this condition, the posttest results indicate that the biases obtained after only 80 trials in the pretest completely override the more than 4 times as extensive subsequent training. Interestingly enough, as can be seen in Figures 2C, 3, and 5C, there are some subjects for whom the training after the pretest initially does appear to be effective at the start of training. If their perceptual biases were governed by stabilization effects alone, these subjects should have stuck to the trained percepts from then on. Instead they switch back to their original biases, obtained in the pretest, usually already during training, and especially in the posttest, after the 5-min break, the biases are consistent again with the ones observed in the pretest. So it seems that there is a more long-term perceptual bias at play here on which disambiguated trials have less of an impact. Even taking into account that adaption and priming for bistable stimuli is likely to occur at several timescales (Brascamp et al., 2008, 2009; Pearson & Brascamp, 2008), it will be very hard to explain this particular result especially when compared to the results for the other conditions.

Test trials vs. training trials

In Experiment 2, we mimicked the pretest condition of Experiment 1 using disambiguated training trials in the preexposure phase instead of ambiguous test trials. In this way, we simulated two possible causes for the interference effect in two separate conditions. In one condition, the disambiguated preexposure phase indicated the location cue to be irrelevant; in the second condition, the correlation in the preexposure phase was the complete opposite with respect to the subsequent training. However, in both cases learning still occurred (see Figure 4), i.e., the percepts during training and in the posttest were consistent with the trained direction rather than being random or opposite. This shows that training trials themselves do not, or hardly, interfere with later training trials, whereas from Experiment 1 we know that test trials do. This is also consistent with recent work that shows that correlation between test and training trials decays about after 4 trials (Fuller, Backus, van Dam, & Ernst, 2009).

The fact that preexposure using training trials does not interfere with training, indicates that the interference due to the pretest in Experiment 1 is not due to having previously obtained a perceptual history, since the percepts for training and test trials are relatively similar. Rather, the comparison between the results of Experiments 1 and 2 indicate that there must be a more fundamental difference between training trials and test trials. For short-term aftereffects, differences between ambiguous and unambiguous figures have been found before. For instance measuring aftereffects after viewing an ambiguous stimulus or an unambiguous stimulus for a brief interval can be in opposite directions (Kanai et al., 2007; Kanai & Verstraten, 2005; Pearson & Clifford, 2005). Here we show for the first time that ambiguous and disambiguated figures can have different effects on a longer timescale as well.

So what is causing the difference in results between test and training trials in our paradigm? In terms of cube percepts, training and test trials are relatively similar but the difference in the way of how to obtain the percept is obvious. In training trials, there are additional disparity and occlusion cues specifying how the cube should be perceived. A perceptual bias will have little to add in this case since it will be overwhelmed by the available evidence in the stimulus that the cube is rotating in a particular direction. For test trials, however, there is no signal specifying how the cube should be perceived, and the resulting perceptual ambiguity is completely resolved within the perceptual system. For test trials, a perceptual bias will therefore play a huge role. For this reason, if for test trials the visual system regards the resulting percept as a correct interpretation it will also regard this percept as more informative about in what direction to change (learn) the bias. The next time the same stimulus is presented the observer will thus also be more likely to have the same percept. For training trials, the resulting percept is more related to the accurate interpretation of the additional cues, rather than being informative for the perceptual bias, and therefore, the perceptual bias is updated to a lesser extent. This difference in effectiveness in updating a bias might be the reason that a bias obtained using test trials is able to survive a much larger amount of training trials. The reasoning here is consistent with results from previous studies on learning a new light-from-above prior. Adams et al. (2004) showed that subjects can learn a new light-from-above prior when haptic feedback was provided about the correct interpretation of the visually ambiguous training stimuli. However, when on top of the haptic feedback visual disambiguation cues were added to the training stimuli the learning effect was greatly reduced (Graf, Adams, & Bouzit, 2007).

Another possibility to explain the difference between training and test trials is if the learning process takes place entirely within the process of resolving the ambiguous input, i.e., the rivalry process. For training trials, due to the additional cues, the competition for dominance between the two different interpretations is more or less bypassed, since dominance is already specified in the sensory input. So if somehow for the learning to occur competition for dominance is necessary, this could also explain why an initial bias obtained with training trials is less strong than an initial bias obtained for test trials.

Transfer of learning to new objects

The transfer effects of training using cubes to testing with spheres has not yet been discussed in great detail so far. Although Figure 4 hints that there is a strong correlation between cube and sphere trials in the posttest, we performed a separate analysis addressing this question more specifically. In order to verify whether there is such a correlation, Figure 6 shows the percentage that spheres in the posttest were perceived in the trained direction versus the percentage that cubes in the posttest were perceived as trained. The results are shown together for all six conditions in this study and separate conditions can be distinguished by the symbols used. For each subject, 2 points are shown, one for each separate location. Since there is a large amount of overlap between the data points, the data points have been jittered slightly about their actual values for reasons of clarity. What is clear from this figure is that posttest biases for spheres are very similar to biases for the cubes on which the training was explicitly focused. The correlation between the sphere and cube biases is highly significant (Pearson's $r = 0.93$; $p \ll 0.001$). This indicates that learning the rotation direction for cubes, to a large extent, transfers to different object shapes. This, in turn, strongly suggests that the underlying learning mechanism focuses on aspects that are common for both shapes, for instance local motion patterns, rather than involving more high-level object recognition. Although, here it is important to note that the spheres were presented in the same block of trials as the cubes, so stabilization effects might have been responsible for the high correlation (Maier et al., 2003). To separate the transfer from stabilization effects and to find out what exactly the common aspects are requires further study.

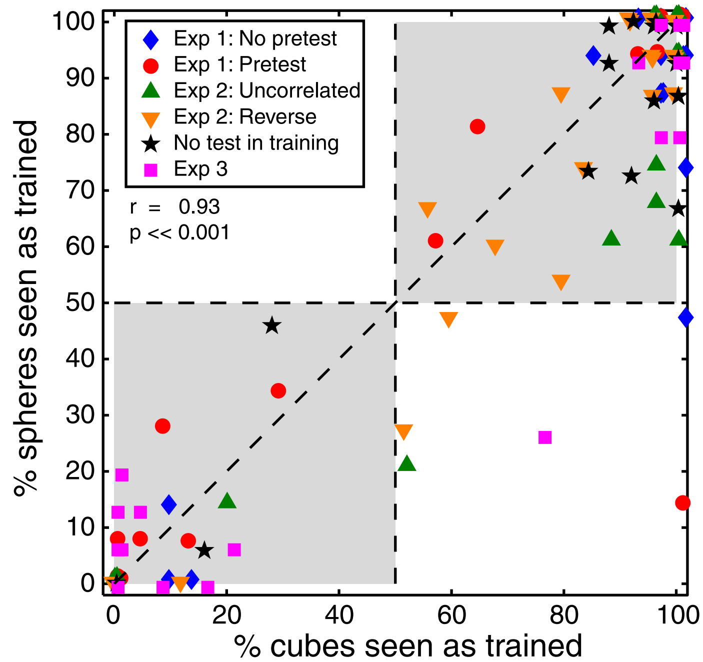

Figure 6. Individual subject posttest results for the proportion of spheres perceived to be rotating in the same direction as trained versus the proportion of cubes perceived to be rotating in the same direction as trained. The results for each individual subject are shown as 2 points, one for each training location, in the graph. The different symbols represent the different experimental groups. Because results from different conditions largely overlap, the symbols are jittered about their actual locations for clarity. Across all groups, there is a highly significant correlation between rotation directions perceived for spheres and cubes.

Long-term learning vs. short-term priming

One question that one might ask is if this long-term learning effect investigated here is not just an extremely long timescale of the same mechanism responsible for the short-term priming effects, given that the pattern of observed results seems quite similar; for example, long-term as well as and short-term "learning" effects, both transfer to different objects, indicating that common features from the objects are involved in the learning process (Maier et al., 2003, and see results here). In addition, both effects appear to be mostly based on a retinotopic organization rather than outside world coordinates (Backus & Haijiang, 2007; Knapen et al., 2009). Thus, they seem to involve retinally localized learners (although note that short-term stabilization has been shown to extend at least somewhat beyond the exact stimulation location; Knapen et al., 2009).

Short-term priming has been explained as the result of perceptual memory within the more general process of bistable figure perception and perceptual alternations. That is, the current understanding for bistable figure perception includes a "memory trace" that more or less tracks the prior perceptual history (for theoretical models, see, e.g., Brascamp et al., 2008, 2009; Gigante, Mattia, Braun, & Del Giudice, 2009; Noest, van Ee, Nijs, & van Wezel, 2007). However, in light of the current results there are two things that these models, at least in their current form, can hardly explain. The first observation that cannot be explained is the difference in results between test and training trials in the pretest. In terms of strength of the perceptual memory trace, the models do not strongly differentiate between the percept being the result of the competition when viewing an ambiguous stimulus (i.e., for test trials) or the percept being the result of a disambiguated stimulus (for training trials). Therefore, these models would predict similar results for both preexposure conditions that either use ambiguous or disambiguated preexposure trials. This is not what we found. Second, these models cannot explain why the more extensive and controlled perceptual history during training fails to override the relatively short history during pretest. In short, it is likely that the memory trace for the long-term learning effects investigated here is governed by a different kind of process than the ones described in the models mentioned above.

The preexposure interference suggests a more conventional learning mechanism that keeps track of prior information, weighting each new input relative to existing information. That is, if the prior information is still relatively sparse at the beginning of the experiment, this would explain why initial trials have a relatively large effect, compared to later trials. Furthermore, such a mechanism is more likely to weigh ambiguous and disambiguated stimuli differently based on their respective salience (see Test trials vs. training trials section above). In addition, such a mechanism would explain the long-term retention of the obtained biases (Carter & Cavanagh, 2007; Haijiang et al., 2006), since there is no decay of the bias if the system is not acquiring new information.

In conclusion, an updated model for the perception of ambiguous figures would ideally incorporate both a competition/percept adaptation mechanism (as in, e.g., Noest et al., 2007) to explain the short-term effects (including negative aftereffects) as well as a prior information/bias updating mechanism to explain the long-term learning effects. Given that both short- and long-term effects seem to depend on similar retinal location conditions, these two mechanisms are likely to be closely tied and coupled retinotopically. Such a closely coupled hybrid model would also explain why a previous attempt to train subjects to obtain a sound-contingent bias for a bistable stimulus failed (Haijiang et al., 2006). It also makes it unlikely that other cues that do not in any way relate to the retinotopic organization will work this easily. However, future research will have to verify the full merit of such a hybrid model.

Conclusion

Here we investigated the role of preexposure for learning location-contingent biases for an ambiguous stimulus. We found that presenting the subjects with the ambiguous stimulus in a pretest prior to training strongly interferes with learning the trained biases. When simulating a pretest using disambiguated training trials, no such interference occurs. Therefore, we concluded that the interference does not depend on the perceptual history alone but rather that there is a fundamental difference between test and training trials in terms of learning a perceptual bias. The results further indicate that the learning mechanism does not necessarily seek out global correlations between location and rotation direction but involves multiple localized perceptual learners. Further, we found that the learned biases for the rotating Necker Cubes mostly transfer to ambiguous rotating spheres. This suggests that learning the perceptual bias occurs on features that are common in both figures, rather than involving higher level object-based learning. A new framework for an updated model for the perception of ambiguous figures is suggested, which would tie in these new results with existing findings.

Acknowledgments

The authors would like to thank Massimiliano Di Luca, Benjamin T. Backus, and two anonymous reviewers for helpful discussions and insightful comments and Jasmin Steinwender and Verena Eikmeier for running subjects. This research was supported by the Human Frontier Science Program and EU Grant 27141 "ImmerSence."

Commercial relationships: none.
Corresponding author: Loes C. J. van Dam.
Email: loes.van.dam@tuebingen.mpg.de.
Address: Spemannstrasse 41, 72076 Tübingen, Germany.

Footnotes

[1]Here it is important to note that "perceptual learning" can have two different meanings in the existing literature. Most perceptual learning studies concern cases where observers get better at a perceptual task due to training. Other studies define "perceptual learning" as a change in perception or appearance due to training. The current study concerns the latter.

[2]Note that this cannot be due to the response feedback always being positive during the pretest: a separate control experiment in which no feedback about the correctness of the response was provided led to similar results.

References

Adams, W. J., Graf, E. W., & Ernst, M. O. (2004). Experience can change the "light-from-above" prior. *Nature Neuroscience, 7,* 1057–1058. [PubMed]

Backus, B. T., & Haijiang, Q. (2007). Competition between newly recruited and pre-existing visual cues during the construction of visual appearance. *Vision Research, 47,* 919–924. [PubMed]

Brascamp, J. W., Knapen, T. H. J., Kanai, R., Noest, A. J., van Ee, R., & van den Berg, A. V. (2008). Multi-timescale perceptual history resolves visual ambiguity. *PLoS One, 3,* e1497. [PubMed] [Article]

Brascamp, J. W., Knapen, T. H. J., Kanai, R., van Ee, R., & van den Berg, A. V. (2007). Flash suppression and flash facilitation in binocular rivalry. *Journal of Vision, 7*(12):12, 1–12, http://www.journalofvision.org/content/7/12/12, doi:10.1167/7.12.12. [PubMed] [Article]

Brascamp, J. W., Pearson, J., Blake, R., & van den Berg, A. V. (2009). Intermittent ambiguous stimuli: Implicit memory causes periodic perceptual alternations. *Journal of Vision, 9*(3):3, 1–23, http://www.journalofvision.org/content/9/3/3, doi:10.1167/9.3.3. [PubMed] [Article]

Burge, J., Ernst, M. O., & Banks, M. S. (2008). The statistical determinants of adaptation rate in human reaching. *Journal of Vision, 8*(4):20, 1–19, http://www.journalofvision.org/content/8/4/20, doi:10.1167/8.4.20. [PubMed] [Article]

Carter, O., & Cavanagh, P. (2007). Onset rivalry: Brief presentation isolates an early independent phase of perceptual competition. *PLoS One, 2,* e343. [PubMed] [Article]

Chong, S. C., & Blake, R. (2006). Exogenous attention and endogenous attention influence initial dominance in binocular rivalry. *Vision Research, 46,* 1794–1803. [PubMed]

Ernst, M. O. (2007). Learning to integrate arbitrary signals from vision and touch. *Journal of Vision, 7*(5):7, 1–14, http://www.journalofvision.org/content/7/5/7, doi:10.1167/7.5.7. [PubMed] [Article]

Flanagan, J. R., Bittner, J. P, & Johansson, R. S. (2008). Experience can change distinct size-weight priors engaged in lifting objects and judging their weights. *Current Biology, 18,* 1742–1747. [PubMed]

Fuller, S., Backus, B., van Dam, L., & Ernst, M. (2009). Short-term dynamics of perceptual bias for bistable stimuli [Abstract]. *Journal of Vision, 9*(8):41, 41, http://www.journalofvision.org/content/9/8/41, doi:10.1167/9.8.41.

Gigante, G., Mattia, M., Braun, J., & Del Giudice, P. (2009). Bistable perception modeled as competing stochastic integrations at two levels. *PLoS Computational Biology, 5,* e1000430. [PubMed] [Article]

Graf, E. W., Adams, W. J., & Bouzit, S. (2007). Light priors, learning and feedback [Abstract]. *Journal of Vision, 7*(9):309, 309, http://www.journalofvision.org/content/7/9/309, doi:10.1167/7.9.309.

Haijiang, Q., Saunders, J. A., Stone, R. W., & Backus, B. T. (2006). Demonstration of cue recruitment: Change in visual appearance by means of Pavlovian conditioning. *Proceedings of the National Academy of Sciences of the United States of America, 103,* 483–488. [PubMed]

Hochberg, J. E. (1950). Figure–ground reversal as a function of visual satiation. *Journal of Experimental Psychology, 40,* 682–686.

Kanai, R., Knapen, T. H. J., van Ee, R., & Verstraten, F. A. J. (2007). Disruption of implicit perceptual memory by intervening neutral stimuli. *Vision Research, 47,* 2675–2683. [PubMed]

Kanai, R., & Verstraten, F. A. J. (2005). Perceptual manifestations of fast neural plasticity: Motion priming, rapid motion aftereffect and perceptual sensitization. *Vision Research, 45,* 3109–3116. [PubMed]

Knapen, T., Brascamp, J., Adams, W. J., & Graf, E. W. (2009). The spatial scale of perceptual memory in ambiguous figure perception. *Journal of Vision, 9*(13):16, 1–12, http://www.journalofvision.org/content/9/13/16, doi:10.1167/9.13.16. [PubMed] [Article]

Leopold, D. A., Wilke, M., Maier, A., & Logothetis, N. K. (2002). Stable perception of visually ambiguous patterns. *Nature Neuroscience, 5,* 605–609. [PubMed]

Long, G. M, Toppino, T. C., & Mondin, G. W. (1992). Prime time: Fatigue and set effects in the perception of reversible figures. *Perception & Psychophysics, 52,* 609–616. [PubMed]

Maier, A., Wilke, M., Logothetis, N. K., & Leopold, D. A. (2003). Perception of temporally interleaved ambiguous patterns. *Current Biology, 13,* 1076–1085. [PubMed]

Mitchell, J. F., Stoner, G. R., & Reynolds, J. H. (2004). Object-based attention determines dominance in binocular rivalry. *Nature, 429,* 410–413. [PubMed]

Nawrot, M., & Blake, R. (1989). Neural integration of information specifying structure from stereopsis and motion. *Science, 244,* 716–718. [PubMed]

Necker, L. A. (1832). Observations on some remarkable optical phenomena seen in Switzerland; and on an optical phenomenon which occurs on viewing a figure of a crystal of geometrical solid. *London and Edinburgh Philosophical Magazine and Journal of Science, 1,* 329–337.

Noest, A. J., van Ee, R., Nijs, M. M., & van Wezel, R. J. A. (2007). Percept–choice sequences driven by interrupted ambiguous stimuli: A low-level neural model. *Journal of Vision, 7*(8):10, 1–14, http://www.journalofvision.org/content/7/8/10, doi:10.1167/7.8.10. [PubMed] [Article]

Orbach, J., Ehrlich, D., & Heath, H. A. (1963). Reversibility of the Necker Cube: I. an examination of the concept of "satiation of orientation". *Perceptual and Motor Skills, 17,* 439–458. [PubMed]

Pastukhov, A., & Braun, J. (2008). A short-term memory of multi-stable perception. *Journal of Vision, 8*(13):7, 1–14, http://www.journalofvision.org/content/8/13/7, doi:10.1167/8.13.7. [PubMed] [Article]

Pearson, J., & Brascamp, J. (2008). Sensory memory for ambiguous vision. *Trends in Cognitive Sciences, 12,* 334–341. [PubMed]

Pearson, J., & Clifford, C. W. G. (2005). Mechanisms selectively engaged in rivalry: Normal vision habituates, rivalrous vision primes. *Vision Research, 45,* 707–714. [PubMed]

Virsu, V. (1975). Determination of perspective reversals. *Nature, 257,* 786–787. [PubMed]

Wolfe, J. M. (1984). Reversing ocular dominance and suppression in a single flash. *Vision Research, 24,* 471–478. [PubMed]