Title
Modeling trial by trial and block feedback in perceptual learning

Authors
Liu, Jiajuan
Dosher, Barbara
Lu, Zhong-Lin

Publication Date
2014-06-01

DOI
10.1016/j.visres.2014.01.001

Vision Res. Author manuscript; available in PMC 2015 June 01.

Published in final edited form as:
Vision Res. 2014 June ; 99: 46–56. doi:10.1016/j.visres.2014.01.001.

Modeling trial by trial and block feedback in perceptual learning

Jiajuan Liu[1], **Barbara Dosher**[1], and **Zhong-Lin Lu**[2]

[1]Department of Cognitive Sciences, University of California, Irvine, CA 92697 USA

[2]Department of Psychology, The Ohio State University, Columbus, OH 43210 USA

Abstract

Feedback has been shown to play a complex role in visual perceptual learning. It is necessary for performance improvement in some conditions while not others. Different forms of feedback, such as trial-by-trial feedback or block feedback, may both facilitate learning, but with different mechanisms. False feedback can abolish learning. We account for all these results with the Augmented Hebbian Reweight Model (AHRM). Specifically, three major factors in the model advance performance improvement: the external trial-by-trial feedback when available, the self-generated output as an internal feedback when no external feedback is available, and the adaptive criterion control based on the block feedback. Through simulating a comprehensive feedback study (Herzog & Fahle 1997, *Vision Research*, *37* (15), 2133–2141), we show that the model predictions account for the pattern of learning in seven major feedback conditions. The AHRM can fully explain the complex empirical results on the role of feedback in visual perceptual learning.

Introduction

Perceptual learning – performance improvements through training or practice – has been demonstrated in a wide range of perceptual tasks in the adult population (Fahle & Poggio, 2002; Lu & Dosher, 2012; Lu, Hua, Huang, Zhou & Dosher, 2011). One important factor in perceptual learning is the availability of feedback. The availability of feedback on performance accuracy can be consequential in determining how quickly or whether learning occurs and in what task circumstances (Herzog & Fahle 1997; Petrov et al 2006; Liu, Lu & Dosher, 2010, 2012). Understanding how perceptual learning is achieved in the adult perceptual system may both reveal the nature of brain plasticity and suggest more effective training paradigms for treating diseases such as amblyopia (Huang, Lu & Zhou, 2008; Levi & Li, 2009; Polat, Sagi & Norcia, 1997). In this study, we considered the complex effect of feedback in perceptual learning and simulated a comprehensive feedback study (Herzog & Fahle, 1997) using the augmented Hebbian reweighting model (AHRM, Petrov, Dosher, & Lu, 2005, 2006). In doing so, we aim to gain a more systematic understanding of the role of

Address correspondence to BD: bdosher@uci.edu.

feedback in perceptual learning and shed light on how different feedback can be used to promote perceptual learning in practice.

Feedback plays an interesting role in perceptual learning (see Dosher & Lu, 2009 for a review). Two main types of feedback have been used in perceptual learning studies: trial-by-trial feedback when a correct/incorrect signal was provided after each trial and block feedback when only a single proportion correct was provided after a whole block (often contains dozens of trials). Or, alternatively, feedback may not be provided in some tasks. Although trial-by-trial feedback is used in most perceptual learning experiments and is associated with performance improvement, significant perceptual learning has also been observed using tasks without any external feedback (Ball & Sekuler, 1987; Crist, Kapadia, Westheimer & Gilbert, 1997; Fahle & Edelman, 1993; Karni & Sagi, 1991; McKee & Westheimer, 1978; Shiu & Pashler, 1992; Petrov, Dosher & Lu, 2006), with just block feedback (Herzog & Fahle, 1997; Shiu & Pashler, 1992; Shibata, Yamagishi, Ishii & Kawato, 2009), or with temporally coincident feedback from an unrelated task (Seitz & Watanabe, 2003; Watanabe, Nanez & Sasaki, 2001; Watanabe, Nanez, Koyama, Mukai, Liederman & Sasaki, 2002; Seitz, Nanez, Holloway, Tsushima & Watanabe, 2006). Two studies found that, after achieving asymptotic performance through training without feedback, the addition of external feedback had little effect (Herzog & Fahle, 1997; McKee & Westheimer, 1978). On the other hand, in other cases it has been documented that feedback improved the rate or extent of learning (Ball & Sekuler, 1987; Fahle & Edelman, 1993; Vallabha & McClelland, 2007), and was necessary for perceptual learning, especially for difficult stimuli (Herzog & Fahle, 1997; Shiu & Pashler, 1992; Seitz et al., 2006). Perceptual learning was found to be absent with false feedback, but performance rebound occurred with subsequent correct feedback (Herzog & Fahle, 1997). Or performance may be affected by details of misleading block feedback (Shibata, Yamagi, Ishii & Kawato, 2009). We also documented that feedback is necessary for perceptual learning when the training accuracy is low, and that its presence is not important when training accuracy is high (Liu, Lu & Dosher, 2010). Furthermore, when trials of high and low training accuracy are mixed, feedback is no longer necessary for the performance improvement in low training accuracy trials (Liu, Lu & Dosher 2012).

The complex pattern of empirical results for feedback in perceptual learning creates a challenge for models of learning and of feedback (Dosher & Lu, 2009). The ability to learn perceptual tasks without feedback in some circumstances and the relevance of feedback in others rules out both a pure supervised model (Hertz, Krogh & Palmer, 1991) and a pure unsupervised model of learning (Polat & Sagi, 1994; Vaina, Sundareswaran & Harris, 1995, Weiss; Edelman & Fahle, 1993), and inspired the development of a hybrid model – the augmented Hebbian reweighting model (AHRM) of perceptual learning (Petrov et al 2005, 2006). In this model, when feedback is absent, it is similar to an unsupervised Hebbian learning system; when feedback is present, the feedback acts as another input to the system, and helps to shift weights in the correct direction. This model naturally predicts and explains several results regarding the effect of trial-by-trial feedback in perceptual learning: if present, the trial-by-trial feedback shifts the post-synaptic activation in the correct direction and in turn fosters appropriate weight changes, hence its beneficial effect. On the other hand, when there is no feedback but the task is relatively easy, the weights still move in the correct direction on average because the activation of the decision unit strongly correlates with the correct stimulus classification. However, when the task is very difficult, such correlation is weak, and so the update of the weights may be very slow or even erroneous and learning is not manifest. In this paper, we examine the ability of this model to quantitatively predict several variants of trial-by-trial feedback.

Even if the AHRM explains the effect of trial-by-trial feedback well, the effects of block feedback on perceptual learning may still pose a challenge. Providing only a proportion correct as feedback after a whole block, trial-by-trial feedback is no longer available; hence the system learns as an unsupervised Hebbian system within a block. How a single proportion correct input at the end of each block facilitates learning is worth further investigation. Two studies provide insights related to this question. Firstly, in a study about perceptual learning in a non-stationary learning environment in which the characteristics of a misleading external noise background switched back and forth every several sessions, Petrov et al. (2005, 2006) found a smaller response bias toward the orientation of the background noise with feedback. Although discrimination improved at approximately the same rate with and without feedback, the presence of feedback allows observers to achieve a more balanced response profile and improves learning in a changing stimulus environment. Secondly, Herzog & Fahle (1997) showed that when there is no feedback, the performance of individual subjects is highly varied: some improved as much as subjects with trial-by-trial feedback, some zigzagged and showed no overall learning, while others deteriorated significantly. Feedback seems to reduce variation in the learning over subjects. Given these results, we hypothesized that block feedback, though having no information about the correctness of every single trial, may help reduce response bias and/or performance variance and in turn enhance perceptual learning. In contrast, Herzog and Fahle (1998) have proposed a model in which block feedback affects the rate of perceptual learning. Though we focus on the bias correction mechanism of block feedback in this study, possible similarities and differences between these two models are also discussed.

To test the potential role of response bias, and to develop and test the theoretical framework, we used the AHRM to simulate all main results in Herzog and Fahle (1997), a comprehensive study about the effect of different trial-by-trial (complete, partial, uncorrelated and reversed) and block (true and manipulated) feedback in a two-alternative forced choice vernier task. The goal is to see whether the AHRM accounts for the role of feedback in perceptual learning with an internally consistent set of mechanisms and parameters. We select the Herzog and Fahle (1997) paper because it is both representative and inclusive of most results about feedback in the literature: In the Herzog and Fahle data, significant performance improvement was shown with complete and partial trial-by-trial feedback and block feedback, but not with other feedback conditions. The no feedback group showed great variance in performance. By following the exact stimulus and training procedure in their experiments, we were able to test the ability of the AHRM model framework to account for all the main results in the study. We were also able to reproduce the varied individual learning process in the group with no feedback. This new application of the AHRM extended our previous studies of the model on various tasks of identifying Gabor stimuli with different contrast and/or orientation, furthered our understanding of feedback and may point to new experimental directions concerning the mode of perceptual learning.

Simulating the Effects of Feedback with the AHRM

The AHRM model for perceptual learning consists of a *representation subsystem*, or visual front end, and a *learning module* that uses augmented Hebbian reweighting. It augments learning by using feedback, when present, and information from a *criterion-control unit* as inputs. Figure 1 shows a schematic diagram of the AHRM (Petrov, Dosher, & Lu, 2005). Model simulations replicate the stimulus and test sequences of the target experiments in generating the behavioral predictions of the AHRM model. First we briefly describe all the subsystems of the AHRM along with the experimental stimuli and procedure in Herzog & Fahle (1997). A detailed description of the model can be found in previous studies (Petrov et al 2005, 2006; Lu et al 2010; Liu et al 2010, 2012). We provide a mathematical description of the model in Appendix A. A related theoretical framework, the integrated reweighting theory (IRT, Dosher, Jeter, Liu & Lu., 2013) extends the reweighting model to learning and transfer over retinal locations.

Representation Subsystem

Herzog and Fahle (1997) reported a series of experiments using a vernier task in which observers discriminated “left” and “right” displacement between an upper and lower line segment (Figure 2). The model front-end computes the activation profile of the vernier stimuli with different displacements. In the modeling study, the size of the displacement is taken from the data reported for each experiment in Herzog and Fahle (1997), which is a threshold obtained using PEST. Following prior observations of the similarity between vernier and orientation judgments (Saarinen & Levi, 1995), the model makes the vernier judgments based on activations of representational units tuned to different orientations. The front-end uses activation channels that are spatial frequency and orientation selective, spanning 5 spatial frequencies and 7 orientations. Each unit incorporates standard normalization processing and stochastic internal noise in the response. This front-end is the same as that used in prior AHRM simulations (Petrov et al., 2005, 2006; Liu et al., 2010, 2012; Lu et al., 2011). The representation system is summarized in the Appendix. The activations in different orientation and frequency bands reflect the processing of the input image through the respective filters; activations on different trials differ due to the incorporation of internal noises.

In the Herzog and Fahle (1997) experiments, the appropriate displacement size for each observer was determined by an adaptive procedure (PEST) in a pretest phase. Then in the main experiment the vernier stimulus was rotated by 90 degree to avoid possible training effects from this pre-testing. In simulation, we first selected model parameters to achieve the initial performance level in Herzog and Fahle (1997), and then simulated only the training part of the study. Simulations on both vertical and horizontal vernier stimuli are essentially the same; we report the results from a vertical simulation.

Decision Subsystem

The decision subsystem of the AHRM takes the weighted sum of activation outputs from the representation subsystem on each trial, adding the bias term and the decision noise to classify a specific stimulus:

$$u=\sum_{i=1}^{nfilter} w_i A\left(\theta_i, f_i\right) - w_b b + \varepsilon_d. \quad (1)$$

The Appendix includes additional details about the decision subsystem.

Adaptive Criterion Subsystem

A top-down bias control system is also employed to balance the response frequency and augment learning. The bias correction term ("*b*" in eq.1) is the recency-weighted aggregate response bias (left as −1, right as +1) of the simulated observer. It is computed from a history of approximately the last 50 trials with the early trials exponentially discounted, and reflects deviations from the expected proportion of left and right responses in the sequence of trial decisions. The bias correction term shifts the response criterion to counterbalance the bias effect. For example, when *b* is negative, which means the observer has produced more "left" responses, the decision unit is shifted upward by $-w_b b$, a positive term which makes a "right" response more likely. This is equivalent to shifting the criterion downward by the same amount. Bias correction reduces future bias and stabilizes the system. The higher the bias weight (w_b), the stronger is the correction effect and the smaller the bias. Figure 3 shows this relationship for some standard model parameters. The appendix provides additional details about the bias control system.

In the current implementation of the AHRM, we introduce a relationship between the accuracy in the last block of trials—either provided in block feedback conditions or estimated in trial feedback conditions—and the bias weight (w_b in eq.1). If the block feedback indicates high discrimination accuracy, then the system has more confidence in the bias information and sets a high bias weight. If the block feedback indicates low accuracy then there is less confidence in bias information and the system sets a low bias weight. With high bias weights, the response frequency tends to be closer to a balanced "50-50" in these experiments where left and right test instances are balanced. In the implementation of the AHRM, the bias weight only changes after every block in the block feedback conditions. We selected a linear relationship between the bias weight for the upcoming block and the accuracy indicated by block feedback for the prior block. The minimum and maximum of the bias weight is at 0 and 1 for performance accuracies (proportion correct or *pc*) between chance at 0.50 and perfect performance at 1.0:

$$w_b = 2 * pc - 1. \quad (2)$$

Slightly different monotonic functions relating the bias weight to performance accuracy account for the empirical data similarly well. Performance accuracy and so block feedback in the Herzog and Fahle data tend to range between 0.65 and 0.85, and so do not constrain the bottom or top of this function, so we cannot rule out a function relating w_b to proportion correct with a minimum above zero or a maximum less than one from these data. If proportion correct is very near chance it is not clear that introducing high bias weight would improve optimal weight discovery.

Augmented Hebbian Learning

Following the response of the decision system, the trial-by-trial feedback, if present, is sent as a top-down input to the decision unit. It forms a late input together with the early input (u in eq. 1) and drives learning. If there is no trial-by-trial feedback, only the early input u is used to drive Hebbian learning. See the appendix for a description of the Hebbian learning mechanism and the related equations.

Simulation methods

The AHRM was implemented in a MATLAB program. The program takes grayscale images as inputs, produces binary ("left"/"right") responses as outputs, and learns on a trial-by-trial basis by adjusting weights on the activations in different representation units. Varying parameters of the model fit the improvements in performance with simulated training with different forms of feedback. Just as in the experiment, there are seven feedback conditions: trial-by-trial feedback, block feedback, partial (50%) trial-by-trial feedback, no feedback, manipulated feedback (where a fake block feedback of 65±3% is provided regardless of observers' actual performance), uncorrelated feedback (a random trial-by-trial feedback) and reversed trial-by-trial feedback.

The Augmented Hebbian Reweighting Model (AHRM) of learning was fit to the behavioral data in different feedback conditions reported in Herzog and Fahle (1997) using modified grid search methods. The parameters that control the front end were set a-priori as in Petrov et al. (2005, 2006), or were fixed based on model fits to experimental data in a number of other applications (Dosher et al., 2013; Lu et al 2010; Liu et al 2010, 2012) (see the Appendix for a discussion). Similarly, the initial weights before learning were set in proportion to the preferred orientation of the units: $w_i = (\theta_i/30)w_{init}$, reflecting general prior knowledge about orientation given initial task instructions in the target experiments. Five basic parameters were varied to optimize the fits of the model to the data: internal multiplicative noise σ_m, internal additive decision noise σ_d, scaling factor a, the weight on feedback w_f, and learning rate η. Four of the five parameters—all but a —were constrained to be equal in all seven learning conditions, although w_f is operational only in the five trial-by-trial feedback conditions. The seven scaling factors, a, one for each group, accommodate small random differences in performance level for these randomly assigned groups. This led to a total of 11 free parameters to fit all 55 data points in the model over 7 different conditions.

The adjustment of parameters to best fit the model to the data was based on Least Squared error:

$$L = \sum \left[\log\left(p_\tau^{theoretical}\right) - \log\left(p_\tau^{measured}\right) \right]^2 \quad (3)$$

where $p_{\tau}^{measured}$ and $p_{\tau}^{theoretical}$ represent measured and model-generated proportion correct, and Σ represents summation over all data points across all seven experimental conditions.

Optimization of the model fits were carried out in two stages: In the first step, the internal noises and scaling factors were adjusted so that the model performance approximately matched the initial performance levels of the groups of human observers in the beginning of the experiment before learning. In the second step, we evaluate the differential effects of learning under different feedback conditions by simulating the human experiment in the model on a trial-by-trial basis. The weights from the stimulus representations to the task decision unit changed dynamically throughout the learning phase, corresponding to learning on each trial of the simulated experiment with the Hebbian mechanism. The output of the decision unit and/or the external feedback was used to update the weights depending on the specific feedback condition. The model performance was then compared to that of the human observers using Least Squared error defined in Eq. 3. The two steps were repeated until the model predictions were reasonably matched to the data using elaborated grid search methods.

For every experimental condition, the model, just as the human observers, ran 7~8 blocks with 80 trials/block. Each simulated experiment was repeated 1000 times. A bootstrap procedure was used to generate confidence intervals on model performance. In each bootstrap step, we sampled performance curves from the same number of simulations as the number of real observers in the experiment to generate the average performance curve. This was repeated 1000 times. Following standard bootstrap procedures, we computed the mean and standard deviations of the proportion correct of the learning curves of the model from the 1000 learning curves. Analysis of variance on model performance was also performed based on the mean and standard deviations of the model curves.

Results

Herzog and Fahle's study aimed to compare the rate of learning in seven feedback conditions: trial-by-trial feedback, partial (50%) trial-by-trial feedback, no feedback, block feedback, manipulated block feedback (always 65%±3%), uncorrelated feedback, and reversed feedback. We simulate all these conditions, with the parameters summarized in Table 1.

Critically, a single learning rate was used to model the learning curves in all different feedback conditions, and the predicted differences between conditions entirely reflect differential effectiveness of Hebbian learning and feedback in these conditions[1]. The internal multiplicative noise, decision noise, and feedback weights are also constant in all feedback conditions. The augmented Hebbian learning model (AHRM) makes straightforward predictions for different variants of trial-by-trial feedback. Block feedback effects, as described earlier, are implemented by incorporating different weights on the bias term in learning. This weighted bias term reflects the observer's sensitivity to changes in the balance of the two responses in the trial history, a property of the AHRM that can lead to

[1]The initial performance level in a group of observers also may in some cases contribute to the level of learning, see Liu et al. (2010).

improved performance. In order to model the effect of block feedback, we assume that the higher the feedback-based proportion correct – whether this is accurate or not – the higher the weight on the bias unit. The bias weight is updated after every block when block feedback is provided. For consistency, we also updated the bias weight in trial-by-trial feedback conditions, reflecting an approximate impression about performance from trial-by-trial feedback. In practice, however, changes in the bias weight were relatively unimportant in conditions with trial-by-trial feedback, and so the trial-by-trial fits can also default to unchanging bias weights. Our previous fits of the model to data focused on conditions with consistent trial-to-trial feedback, either on every trial (Liu, Lu, & Dosher, 2010, 2012), or on errors only or no feedback (Petrov, Dosher, & Lu, 2005, 2006). This test extends results to partial trial-by-trial feedback, and forms of block feedback. Also, this is the first time the AHRM has been applied to vernier task learning for any feedback condition.

For purposes of discussion, we divided the seven conditions into three categories: trial-by-trial “real” feedback, block feedback, and trial-by-trial “irrelevant” feedback. The first category includes the first three conditions of the experiment: trial-by-trial, partial trial-by-trial, and no feedback groups. In these groups, the trial-by-trial feedback, if present, is real and correct. The second category includes the block feedback and manipulated block feedback groups, in which only block feedback was provided. The third category includes the uncorrelated feedback and reversed feedback groups, in which the trial-by-trial feedback is totally irrelevant or wrong.

Trial-by-trial feedback, partial trial-by-trial feedback, and no feedback

In the experiment, when the trial-by-trial feedback is present, even only half of the time, observers improved over blocks of practice. The data and the fits of the AHRM model to these three conditions are shown in Figure 4. The model captures this pattern as shown by the red lines and shaded areas. In the no feedback condition, on average the observers did not show significant learning, and this again was consistent with the results simulated by the model. The amount of performance improvement for both the experiment and model is summarized in Table 2. The quality of the fits of the model to the data was excellent. The scaling factors (a) in the model were used to adjust for apparent slight level differences between the groups at the beginning of training; other model parameters set the level and general speed of perceptual learning. Group differences in learning were solely a consequence of the different feedback protocols. The findings of little learning in the absence of feedback with about 70% correct staircases[2] are generally consistent with prior reports and the predictions of the Hebbian model that feedback (supervised learning) is necessary for perceptual learning when training tracks lower accuracies (Liu, Lu, & Dosher, 2010). In that experiment, learning at 85% correct performance levels did not depend upon the availability of feedback, while learning at 65% correct did. The no-feedback data of Herzog and Fahle (1997) started with slightly higher initial performance than the feedback conditions. A higher starting level would be expected to make learning more possible not less possible; consistent with this, a simulation of the no-feedback condition with starting

[2]The slight uptrend and final downtrend in the model for no-feedback data were the result of variability in performance over time; continuing training the model led to what are essentially stochastic fluctuations in performance.

level equated to the trial-by-trial condition also generated no perceptible learning (2.6±1.8% and 1.1±1.7% improvement for starting level of complete and partial trial-by-trial feedback respectively; compare with AHRM values in Table 2). So the model predictions are consistent with no perceptible learning in the no-feedback condition under a range of initial levels. The AHRM predicts that learning should be possible even in the absence of feedback if the initial performance level is high enough.

Block feedback and manipulated block feedback

In Herzog and Fahle (1997), when (accurate) block feedback was available, observers did almost as well as when trial-by-trial feedback was available. On the other hand, misleading or fake block feedback prevented learning when block feedback was set at 65±3%. However, the AHRM with the same parameters predicts that if the misleading or fake block feedback is set at 85±3%, learning can be reinstated at these higher levels. Herzog & Fahle (1997) did not test this specific condition, but a related experiment by Shibata et al. (2009) shows exactly such results (see discussion for more details). In this simulation, the AHRM was extended to account for block feedback by introducing a relationship between the level of block feedback and the weight placed on bias control; the higher the block feedback, the higher the weight on the bias unit and hence the smaller the response bias. In turn, this improves the opportunity to learn the correct weights. The AHRM model predictions and data for these conditions are shown in Figure 5.

In contrast, Herzog and Fahle (1998) suggested that the increase or decrease in block feedback from one block to the next directly alters the learning rate. In particular, the learning rate for the next block is increased when the estimated accuracy of the previous block times the magnitude of the internal decision signal was an improvement over the one before and decreased if it is less than the one before. We carried out a supplementary model fit that held bias weight constant and instead varied learning rate in the block feedback condition. In the context of the AHRM, performance sensitive modification of learning rate without altering the criterion control did not fit the data as well. Learning rate changes underperformed the learning in the data (or in the AHRM model) (learning slope is 0.77±0.31; learning amount is 5.4±2.2%; compare to AHRM model values in Table 2). This resulted in moderately large instability in learning rates from one block to the next, and consequently increased the variability in the learning as well. And, in the context of trail-by-trial feedback, the feedback dominated learning.

Uncorrelated feedback and reversed feedback

With totally random feedback, Herzog and Fahle’s observers did not improve significantly over time. This is captured by the AHRM as well. For the condition with uncorrelated feedback, the model predicts that the performance would approach chance level (50%). For the condition with reversed feedback, the model predicts that the performance would generally decrease toward 0% if given enough time. Herzog and Fahle (1997) report the data from a single observer in the reversed feedback condition that has quite a bit of variability from block to block and may not be very informative. The model predicts reduction in performance from the reversed feedback. However, it may also be that observers recognize that the feedback is not veridical and choose to ignore it. However, the current model implementation did not require an explicit discounting system but instead naturally incorporates relative values of the internal response and the external feedback at the learning stage. The internal response and the weighted feedback jointly determine the direction and extent of response used during the learning phase based on their relative values; opposition of the two reduces or may eliminate learning.

In Table 2, we summarize both the slope and net extent of performance improvements from the model simulation, and compared them to those from the original experiment. Consistent with Figures 4–6, the trial-by-trial, block and partial trial-by-trial feedback groups showed significant learning while other groups did not. This is true in the experiment and in the simulation. There is no significant difference between the experimental and simulation results. In the groups with significant performance improvement (true and partial trial-by-trial feedback, and block feedback), the model accounts for 88.2% variance of the experimental results. In the groups with no significant improvement, the model generally shows no learning effect (a flat line) and hence explains little variance of the behavioural results (35.8%). In the absence of learning, the mean of a condition is the best prediction; the model does slightly better as a result of the scaling factors a and may also capture small but non-significant learning (the r^2 is benchmarked to the global mean over different groups as the predictor, which by definition would lead to 0% variance over the grand mean). The learning slopes and learning amounts in the experiments are remarkably consistent with the learning slopes and learning amounts generated by model-simulated learning.

Dynamics of learning: The change of weights

How was learning achieved through the AHRM? A look at the change of weights from the representation units to the decision unit revealed possible dynamics of perceptual learning in these conditions. Figure 7 shows how practice in the model alters the weights on different orientation sensitive channels in the representation on average, while figure 8 shows the weight change of a single trace of simulation as an example. Here only orientation channels of one spatial frequency (1.41 c/d) are shown, since channels of other spatial frequencies show a similar pattern. As seen in figure 7, most groups except the uncorrelated and reversed feedback groups show a pattern of weight change in the same direction as the best trial-by-trial feedback condition – the absolute values of the weights of the most relevant channels increased over the process of training, which should support better discrimination between “left” and “right” stimuli. This is puzzling because the proportion correct performance did not follow the same pattern – only the first three groups, the trial-by-trial, block and partial trial-by-trial feedback groups, improved vernier judgment accuracy. The single trace weight dynamics (figure 8) shed light on why the accuracy predictions of the model differ between these conditions.

Similar to figure 7, in the conditions with significant learning (top panels), the absolute values of the weights for the most relevant channels (at ±15 deg) increased with practice, while the weights on the irrelevant channels stayed about the same or decreased slightly. For the “no feedback” and “manipulated feedback” group, however, although the weights on the most relevant channels did differentiate to some degree, response bias developed and moved the weights in one direction or another. Behaviorally, bias refers to a departure from balanced Left and Right responses over sets of trials; the correspondence to this in the weight traces occurs when the weights become asymmetric in favor of either left or right (as can be seen for example in the drifts of the turquoise line above or below the zero baseline). For different traces of the simulation, developing bias could be either negative or positive or negligible; or, sometimes bias changed direction in the middle of training. Even if the weights on the most informative inputs improve, the aggregate variability and bias in the remaining weights dominate the process and result in no performance improvements. The traces of individual simulation learning histories may show significant learning, while others show none, or even reductions in performance. However, the overall effect of bias and variability from many trials can render the predicted learning insignificant.

To better understand the effects of bias on weight variability, we calculated the variability of the weights in the training course, as shown in figure 9. We selected one representative channel, orientation of 15° and spatial frequency of 1.4 cycle/deg, the closest channel to the experimental stimulus, and calculated its standard deviation from all 1000 simulations in the left panel, and the ratio between the standard deviation and the mean in the right panel. As shown in figure 9, the no-feedback and manipulated feedback groups showed the biggest standard deviation, i.e. some traces may improve but others may deteriorate; overall there was no performance improvement. The block feedback has a slightly bigger standard deviation than the trial-by-trial feedback groups, but it was not severe enough to prevent learning. The scalloped structure in the no feedback and manipulated feedback groups, reflects the fact that the bias from the criterion control unit is set at zero at the beginning of each block (reflecting the absence of history within the training block, regardless of the bias weight) and bias correction gradually increases during the block as evidence accumulates. The partial and complete trial-by-trial feedback groups showed a small standard deviation of the weights, representing a less variable performance in these groups. Improved weights in the presence of small variability yields improved predicted performance accuracy.

For the uncorrelated feedback and reversed feedback groups, there are no systematic improvements of the weights, and hence no learning. Specifically, for the uncorrelated feedback group, weights generally move toward zero, since the feedback is not informative at all, and hence there is no performance improvement. For the reversed feedback group, weights may actually move in the direction that is opposite to the optimal weights, and theoretically performance could be significantly damaged during training by inducing observers to learn an opposite response. This did not happen in the experiment.

Discussion

In this study, we reviewed the effect of feedback in visual perceptual learning and simulated a comprehensive feedback study (Herzog & Fahle, 1997) by extending the augmented Hebbian reweighting model (AHRM) (Petrov et al., 2005, 2006) to account for the effects of different forms of trial-by-trial feedback and the facilitatory effects of block feedback. We successfully modelled all the results in Herzog & Fahle (1997): both true trial-by-trial and block feedback facilitate learning; false trial-by-trial feedback (uncorrelated and reversed feedback) abolishes learning; no feedback leads to highly variable performance with no average learning; manipulated block feedback, when low (65 ± 3%), also rendered learning negligible. We also showed that if falsely high (85 ± 3%), a manipulated block feedback in some circumstances may actually facilitate learning. This particular experiment was not done in Herzog & Fahle (1997), but the prediction is consistent with results from another study, Shibata et. al (2009), where exaggerated positive block feedback was shown to benefit learning.

The positive effect of block feedback has been shown in multiple studies (Herzog & Fahle, 1997; Shiu & Pashler, 1992; Shibata et al, 2009). We simulated this effect by recognizing that, in comparison with no feedback (which was shown to be far more biased than with trial-by-trial feedback, Petrov et al., 2006), the response bias and/or performance variance is smaller with block feedback (Herzog & Fahle, 1997). We predicted the experimental results for block feedback in Herzog & Fahle (1997) by employing a straightforward relationship between bias weight and the block feedback: the higher the block feedback, the bigger the bias weight – hence the smaller the bias.

In a related study about block feedback in perceptual learning, Shibata et al (2009) showed that false block feedback, if more positive than the actual performance, can facilitate learning. They also found no effect of false feedback that was more negative than the actual performance. Shibata et al (2009) developed a computational model in which a performance gradient together with performance variance alters the learning rate. The AHRM simulation treatment of the Shibata study, which uses a task with a different two-interval same-different judgment rule, is considered in a separate development.

Herzog and Fahle (1998, 1999) proposed a model for perceptual learning in vernier experiments based on task-dependent top-down reduction of the effective connections from an input layer of a network that represents orientations. In this model, learning rate is altered by trial feedback and block feedback. Properties of the Herzog and Fahle (1997) data were cited as inspiration for the model. It is possible that an implemented form of the Herzog and Fahle model could also provide a competitive quantitative fit to the broad set of feedback data; however, they did not provide fits of the model to the Herzog and Fahle (1997) data. Instead, response shifts resulting from asymmetric training of left and right vernier stimuli were qualitatively predicted and experimentally confirmed in Herzog and Fahle (1998).

These asymmetric training effects were the topic of a series of experiment in Herzog and Fahle (1999 in Herzog and Fahle (2006), with the same testing stimuli used in Aberg and Herzog (2012). The asymmetric set included offsets of −15″, −10″, −5″, +10″, and +15″ (arc s), tested with different probabilities and in some conditions the feedback for −5″ was replaced on some or all trials with false feedback indicating a “right” stimulus. Quantitative fitting of the data for the varied training schedules in Herzog and Fahle (1999) and Aberg and Herzog (2012) by the AHRM would require a very extensive new modeling project. Additionally, Herzog & Fahle (1999), report performance only for left offset conditions, and effective model fitting would require more complete data sets. However, without quantitative fitting, simulations of the currently implemented AHRM using the training protocols in asymmetric exposure and feedback reversal conditions predict data patterns (Figure 10) that are qualitatively akin to those shown in Herzog and Fahle (1999)—biased training with false feedback on the smallest left offset leads to decreases in correct labelling of the negative offsets, essentially shifting responses to "right". The model also appears to be qualitatively consistent with the results in Aberg and Herzog (2012). In the AHRM, these response shifts primarily reflect shifts (biases) in learned weights towards "right" and only secondarily the operation of the bias control unit.

Aberg and Herzog (2012), who also used the asymmetric design, argued that block feedback left the decision criterion across blocks unaltered in a line vernier task. In our model, block feedback changes bias weight, and the decision criterion changes slightly on every trial by the product of bias (recency-weighted aggregate response bias) and bias weight (eq. 1). Our current results and those of Aberg and Herzog are not necessarily in contradiction. The decision criterion is changed on each trial by the product of the bias and bias weight. When bias weight increases, however, average bias in the AHRM model generally decreases (see figure 3). Therefore, the magnitude of the effect on responses can be small and may not be detectable. Also, bias may be negative or positive from trial to trial, with no accumulating effect of criterion change in a specific direction absent an asymmetric or false feedback design.

The current augmented Hebbian reweighting model (AHRM) naturally makes predictions about different forms of trial-by-trial feedback in perceptual learning, as well as most effects of random and false trial-by-trial feedback. It was extended to account for the effects of true and manipulated block feedback by using adaptive criterion control that depends upon block feedback or a sense of overall accuracy and its interactions with learning rate. The current implementation of the AHRM accounted for a wide range of data patterns of feedback in perceptual learning without the introduction of other more complex functions of feedback discounting or of complex criterion control. Several reasonable complications seem intuitive and might be required to account for learning and performance in other paradigms. For example, it seems likely that observers could note the existence of false feedback when it is applied to a range of easy stimuli—essentially noting that the internal response and the feedback are either randomly related or negatively correlated. Feedback monitoring could be incorporated in the model, and then used to lower or set to zero the weight on feedback in learning if the observer decides that it is misleading. The current criterion control unit does not use information from the feedback to weight the inputs in estimating the bias. If feedback is reliable, it could be integrated to more heavily weight errors in the bias history. There may be some circumstances in which separate criterion control tracking should apply to distinct stimulus conditions. These more complex rules for augmentation of the Hebbian learning via feedback or bias control might assist in predictions in some circumstances.

We conclude that the reweighting model and framework of the AHRM provide a successful account of the impact of major variants of feedback and their effects on perceptual learning.

References

Aberg KC, Herzog MH. Different types of feedback change decision criterion and sensitivity differently in perceptual learning. Journal of Vision. 2012; 12(3):3, 1–11. [PubMed: 22396463]

Ball K, Sekuler R. Direction-specific improvement in motion discrimination. Vision Research. 1987; 27(6):953–965. [PubMed: 3660656]

Crist RB, Kapadia MK, Westheimer G, Gilbert CD. Perceptual Learning of Spatial Location: Specificity for Orientation, Position, and Context. Journal of Physiology. 1997; 78:2889–2894.

Dosher BA, Lu ZL. Hebbian reweighting on stable representations in perceptual learning. Learning & Perception. 2009; 1:37–58. [PubMed: 20305755]

Dosher BA, Jeter P, Liu J, Lu Z-L. An integrated reweighting theory of perceptual learning. Proc Natl Acad Sci. 2013; 110(33):13678–83. [PubMed: 23898204]

Fahle M, Edelman S. Long-term learning in vernier acuity: Effects of stimulus orientation, range and of feedback. Vision Research. 1993; 33(3):397–412. [PubMed: 8447110]

Fahle, M.; Poggio, T. Perceptual learning. Cambridge, Mass: MIT Press; 2002.

Hertz, J.; Krogh, A.; Palmer, RG. Santa Fe Institute studies in the sciences of complexity. Redwood City, California: Addison-Wesley Pub. Co; 1991. Introduction to the theory of neural compuation.

Herzog MH, Fahle M. The role of feedback in learning a vernier discrimination task. Vision Research. 1997; 37(15):2133–2141. [PubMed: 9327060]

Herzog MH, Fahle M. Modeling perceptual learning: difficulties and how they can be overcome. Biol Cybern. 1998; 78:107–117. [PubMed: 9525037]

Herzog MH, Fahle M. Effects of biased feedback on learning and deciding in a vernier discrimination task. Vision Research. 1999; 39(25):4232–4243. [PubMed: 10755160]

Herzog MH, Eward KRF, Hermens F, Fahle M. Reverse feedback induces position and orientation specific changes. Vision Research. 2006; 46(22):3761–3770. [PubMed: 16844180]

Huang CB, Lu Z-L, Zhou YF. Broad bandwidth of perceptual learning in the visual system of adults with anisometropic amblyopia. Proc Natl Acad Sci U S A. 2008; 105(10):4068–4073. [PubMed: 18316716]

Karni A, Sagi D. Where practice makes perfect in texture discrimination: evidence for primary visual cortex plasticity. Proc Natl Acad Sci U S A. 1991; 88(11):4966–4970. [PubMed: 2052578]

Levi DM, Li RW. Perceptual learning as a potential treatment for amblyopia: a mini review. Vision Research. 2009; 49:2535–3549. [PubMed: 19250947]

Liu J, Lu Z-L, Dosher BA. Augmented Hebbian reweighting: Interactions between feedback and training accuracy in perceptual learning. Journal of Vision. 2010; 10(10):29, 1–14. [PubMed: 20884494]

Liu J, Lu ZL, Dosher BA. Mixture of training at high and low accuracy levels facilitates learning at low training accuracy level. Vision Research, Special Issue on Perceptual Learning. 2012; 61:15–24.

Lu ZL, Liu J, Dosher BA. Modeling mechanisms of perceptual learning with augmented Hebbian re-weighting. Vision Research. 2010; 50:375–390. [PubMed: 19732786]

Lu, Z-L.; Dosher, BA. Visual perceptual learning. In: Seel, Norbert M., editor. Encyclopedia of the Sciences of Learning. Springer; 2012. p. 3415-3418.

Lu Z-L, Hua T, Huang C-B, Zhou Y, Dosher BA. Visual perceptual learning. Neurobiology of Learning and Memory. 95:145–151. [PubMed: 20870024]

McKee SP, Westheimer G. Improvement in vernier acuity with practice. Perception & Psychophysics. 1978; 24(3):258–262. [PubMed: 704286]

Petrov A, Dosher BA, Lu Z-L. Perceptual learning through incremental channel reweighting. Psychological Review. 2005; 112(4):715–743. [PubMed: 16262466]

Petrov AA, Dosher BA, Lu ZL. Perceptual learning without feedback in non-stationary contexts: Data and model. Vision Research. 2006; 46(19):3177–3197. [PubMed: 16697434]

Polat U, Sagi D. The architecture of perceptual spatial interactions. Vision Research. 1994; 34(1):73–78. [PubMed: 8116270]

Polat U, Sagi D, Norcia AM. Abnormal long-range spatial interactions in amblyopia. Vision Research. 1997; 37(6):737–744. [PubMed: 9156218]

Saarinen J, Levi DM. Perceptual learning in vernier acuity: what is learned? Vision Research. 1995; 35(4):519–27. [PubMed: 7900292]

Seitz AR, Nanez JE, Holloway S, Tsushima Y, Watanabe T. Two cases requiring external reinforcement in perceptual learning. Journal of Vision. 2006; 6(9):966–973. [PubMed: 17083288]

Seitz AR, Watanabe T. Is subliminal learning really passive? Nature. 2003; 422(6927):36–36. [PubMed: 12621425]

Shibata K, Yamagishi N, Ishii S, Kawato M. Boosting perceptual learning by fake feedback. Vision Research. 2009; 49:2574–2585. [PubMed: 19531366]

Shiu, L-p; Pashler, H. Improvement in line orientation discrimination is retinally local but dependent on cognitive set. Perception & Psychophysics. 1992; 52(5):582–588. [PubMed: 1437491]

Vaina LM, Sundareswaran V, Harris JG. Learning to ignore: Psychophysics and computational modeling of fast learning of direction in noisy motion stimuli. Cognitive Brain Research. 1995; 2(3):155–163. [PubMed: 7580397]

Vallabha GK, McClelland JL. Success and failure of new speech category learning in adulthood: consequences of learned Hebbian attractors in topographic maps. Cognitive, affective & behavioral neuroscience. 2007; 7(1):53–73.

Watanabe T, Nanez JE, Sasaki Y. Perceptual learning without perception. Nature. 2001; 413(6858): 844–848. [PubMed: 11677607]

Watanabe T, Nanez JES, Koyama S, Mukai I, Liederman J, Sasaki Y. Greater plasticity in lower-level than higher-level visual motion processing in a passive perceptual learning task. Nature Neuroscience. 2002; 5(10):1003–1009.

Weiss Y, Edelman S, Fahle M. Models of perceptual learning in vernier hyperacuity. Neural Computation. 1993; 5(5):695–718.

Appendix

The Augmented Hebbian Reweighting Model (AHRM) simulates a multi-channel network model that takes stimulus images as input, produces a task response, and updates weights from stimulus information to decision reflecting learned improvement in task performance. A Hebbian model of learning is augmented by inputs from feedback and from a criterion control unit. Learning occurs through channel reweighting (Dosher & Lu, 1998, 1999; Petrov et al., 2005, 2006). This Appendix provides a brief description of the model.

The *representation subsystem* applied in this paper consists of orientation- and frequency-selective units. This system has previously been used for tasks based on discrimination of the orientation of Gabor patches; here the same representations are used to discriminate vernier lines based on their orientation evidence. The representation units compute the activation value $A(\theta, f)$ of the stimulus image—the normalized spectral energy in each channel.

Retinotopic *phase-sensitive maps* $S(x, y, \theta, f, \varphi)$ are computed for the input image $I(x, y)$:

$$S(x,y,\theta,f,\phi)=[RF_{\theta,f,\phi}(x,y)\otimes I(x,y)]^2 \quad (4)$$

These units at location (x, y) are tuned to spatial frequency f, orientation θ, and spatial phase φ. There were 5 spatial frequencies {1, 1.4, 2, 2.8, 4 c/d}, 7 orientations {0°, ±15°, ±30°, ±45°}, and four spatial phases {0°,90°,180°, 270°}. The bandwidth of spatial frequency tuning and of orientation tuning were set at $h_f = 1$ octaves and $h_\theta = 30°$ (half-amplitude full-bandwidth). These values were based on estimates cellular tuning bandwidths in primary visual cortex and are the same ones used in other applications of the AHRM (Petrov et al., 2005, 2006; Liu et al., 2010, 2012; Dosher, Jeter, Liu & Lu, 2013).

The input image $I(x, y)$ is convolved with each unit filter using fast Fourier transform, followed by a half-squaring rectification operation, followed by spatial phase pooling and then inhibitory normalization (Heeger, 1992), respectively:

$$E(x,y,\theta,f)=\sum_{\phi}S(x,y,\theta,f,\phi)+\varepsilon_1 \quad (5)$$

and

$$C(x,y,\theta,f)=\frac{aE(x,y,\theta,f)}{k+N(f)} \quad (6)$$

The normalization pool N_f is weakly tuned for spatial frequency and independent of orientation (see Petrov et al., 2005 for discussion). *a* is a scaling factor; *k* is a saturation constant relevant for extremely small contrasts. Spatial phase is pooled in this application where phase does not distinguish stimuli; but pooling could be omitted for phase-sensitive tasks or stimuli. The Gaussian kernel of radius W_r determines the spatial pooling for the region of the stimulus.

There are two internal noises. The internal noise term ε_1 has mean 0, standard deviation σ_1, with a Gaussian distribution. Another internal noise ε_2 of mean 0 and standard deviation σ_2 introduces another source of stochastic variability:

$$A^{'}(\theta,f)=\sum_{x,y}W_r(x,y)C(x,y,\theta,f)+\varepsilon_2 \quad (7)$$

An activation function with gain parameter γ range-limits the activation of the representation units:

$$A(\theta,f)=\begin{cases}\frac{1-e^{-\gamma A^{'}}}{1+e^{-\gamma A^{'}}}A_{\max}, & if A^{'}\geq 0\\ 0, & otherwise\end{cases} \quad (8)$$

The activation pattern over the representation units is combined to yield a decision that weights these inputs by w_i, including a top-down bias factor *b* with weight w_b: $u=\sum_{i=1}^{65}w_iA(\theta_i,f_i)-w_bb+\varepsilon_d$, and includes random decision noise ε_d (Gaussian with mean 0 and standard deviation σ_d). The "early" activation of the decision unit o' is a sigmoid function of the weighted activations *u* with gain γ:

$$o=G(u)=\frac{1-e^{-\gamma u}}{1+e^{-\gamma u}}A_{\max} \quad (9)$$

A negative o' maps to one response ("left"), while a positive o' maps to the other response ("right").

The weight structure is learned through updating on every trial. When feedback is available, the feedback pushes the decision unit to a late level *o*:

$$o' = G(u) = \frac{1 - e^{-\gamma u}}{1 + e^{-\gamma u}} A_{\max} \ (\text{late}) \quad (10)$$

Learning occurs during this late phase. The late activation will go to its maximum ($\pm A_{max} = \pm 1$) with feedback ($F = \pm 1$) with high feedback weight, while lower feedback weight will only slightly shift activation in the direction of the correct response. If feedback is not present, learning operates without benefit of this shift towards a correct response ($o = o'$). Except for very low accuracy conditions, the learned weights tend to move towards a more optimum weight distribution because o' tends to correlate with the correct response.

Learning occurs by updating the synaptic connection weights from sensory representation units to the decision unit. The change in each weight, w_i, depends on the learning rate, η, the presynaptic activation $A(\theta, f)$, how far the post-synaptic activation is from its long-term average, $(o - \bar{o})$, and how far the weights are from their saturation values, w_{min} or w_{max}. Weights are learned as:

$$\Delta w_i = (w_i - w_{\min}) [\delta_i]_- + (w_{\max} - w_i) [\delta_i]_+ \quad (11)$$

where

$$\delta_i = \eta A(\theta_i f_i)(o - \overline{o}), \quad (12)$$

and the average of post-synaptic activation is

$$\overline{o}(t+1) = \rho o(t) + (1-\rho)\overline{o}(t). \quad (13)$$

The Hebbian learning process is augmented not just by feedback (when it occurs), but also by a criterion control unit that tracks deviations of the recent response frequencies from 50% or the instructed presentation probabilities in the experiment. Top-down input b weighted by w_b is input to the decision unit. The bias on each trial is an exponentially weighted average of the responses with a time constant of 50 trials ($\rho = 0.02$):

$$r(t+1) = \rho R(t) + (1-\rho) r(t) \quad (14)$$

$$b(t+1) = r(t) \quad (15)$$

Here, $R(t)$ is the response for the current trial (−1 for “Left” and +1 for “Right”), and $r(t)$ is the response running average which exponentially discounts past trials. Prior studies found more pervasive response biases, and correspondingly lower weights on adaptive criterion control, in the absence of feedback (Petrov et al., 2005, 2006). Bias control tracks responses, while feedback tracks external teaching signals. Bias control is more important to learning in the absence of trial-by-trial external feedback (Petrov et al., 2006).

- We model the varied effects of different kinds of feedback on perceptual learning.
- We extend an Hebbian reweighting model to consider different kinds of feedback.
- The model fits seven conditions of feedback in the data of Herzog and Fahle 1997.
- Block feedback is modeled through adaptive criterion setting.
- The study provides an integrated account of a full range of feedback phenomena.

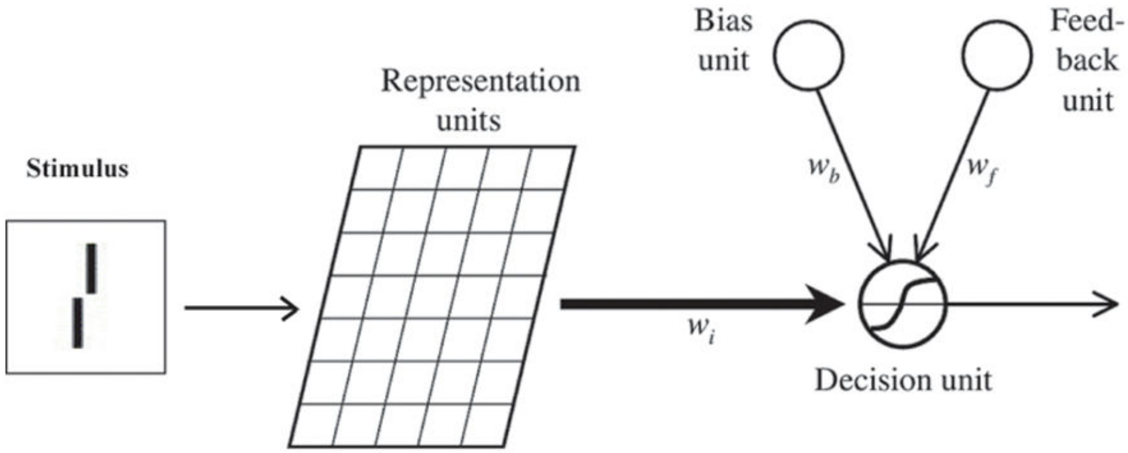

Figure 1.

An illustration of the AHRM model framework. The model takes gray-scale image as the input, encodes the stimulus as an activation pattern through representation units, calculates weighted sum of the representation together with the bias control and makes a decision about the stimulus. The feedback, if present, shifts the output of the decision unit for learning (reweighting of the *wi* in the figure). This figure is modified from figure 4 of Petrov, Dosher, & Lu (2006).

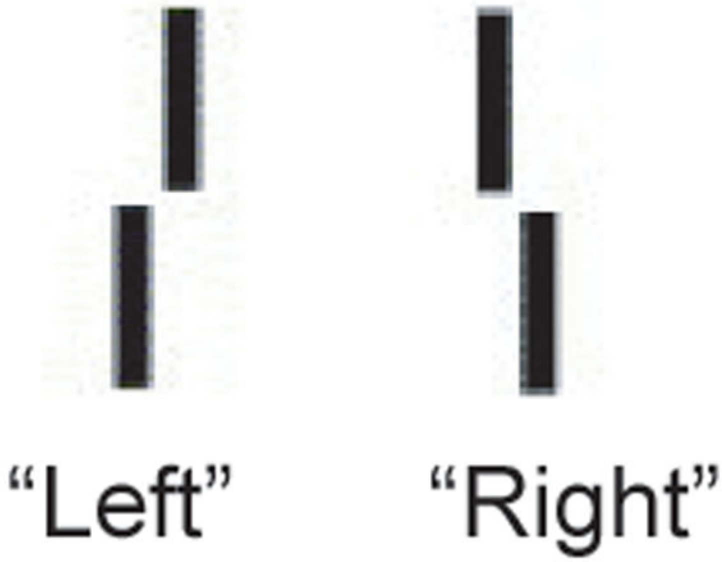

Figure 2.
The line vernier stimuli used in Herzog & Fahle (1997).

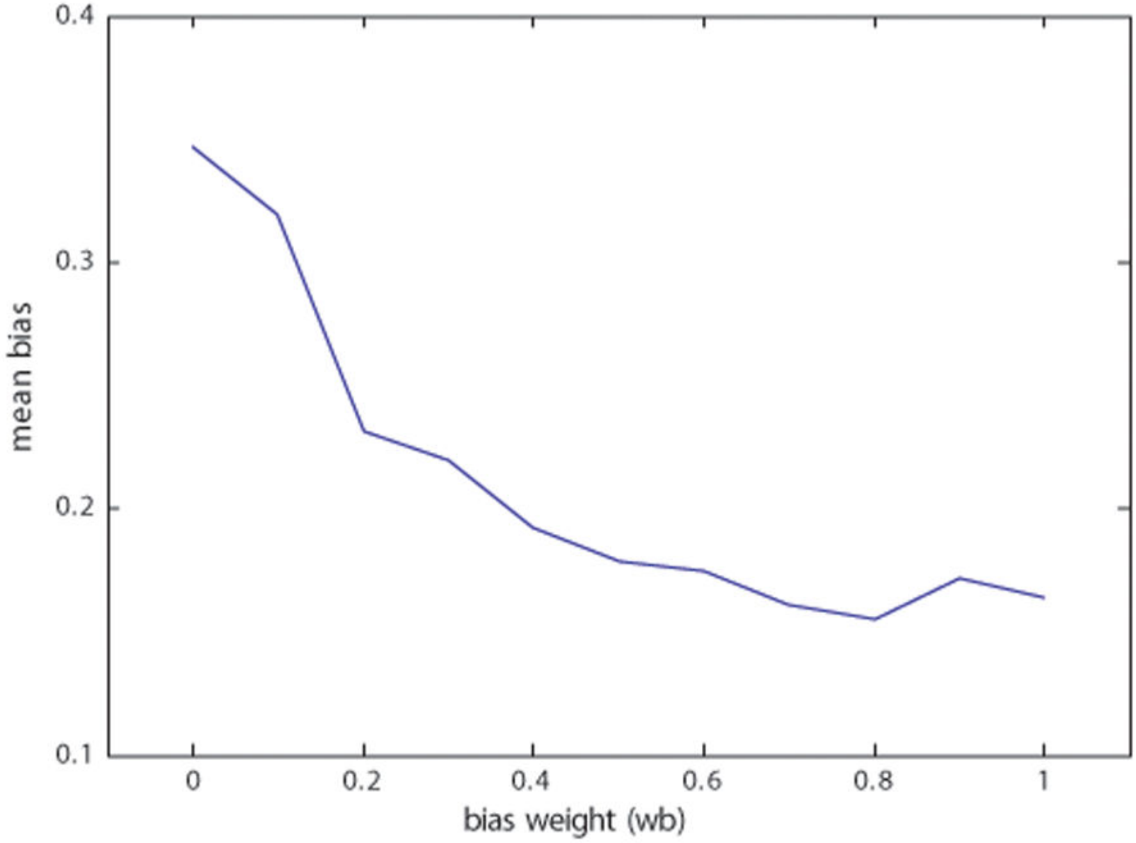

Figure 3.

The mean of absolute value of bias from the whole experiment as a function of the bias weight. Since the bias can be either positive or negative, a direct mean may not indicate the amplitude of bias correctly. An absolute value is taken for bias from each trial and then averaged over all the trials. As the bias weight (wb) increases, the overall bias is reduced.

NIH-PA Author Manuscript
NIH-PA Author Manuscript

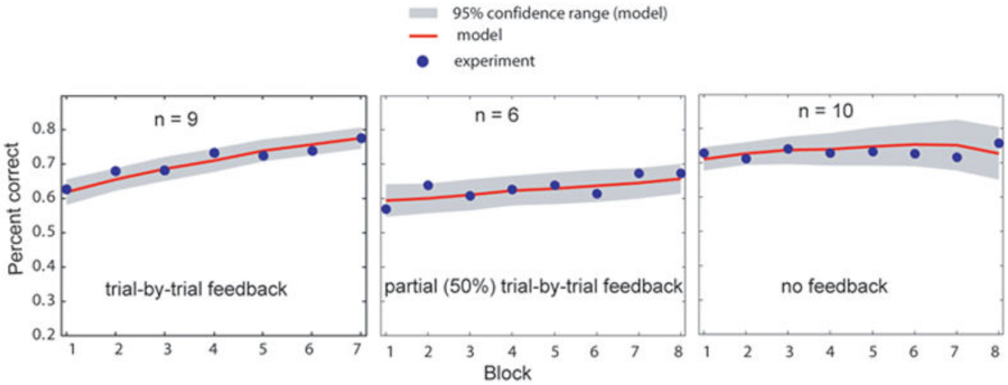

Figure 4.
Data and model fits for the three trial-by-trial “real” feedback conditions. Just as shown in Herzog & Fahle (1997), correct trial-by-trial feedback, even present only half of the time, facilitated learning; while on average no significant performance improvement was shown in the no-feedback group.

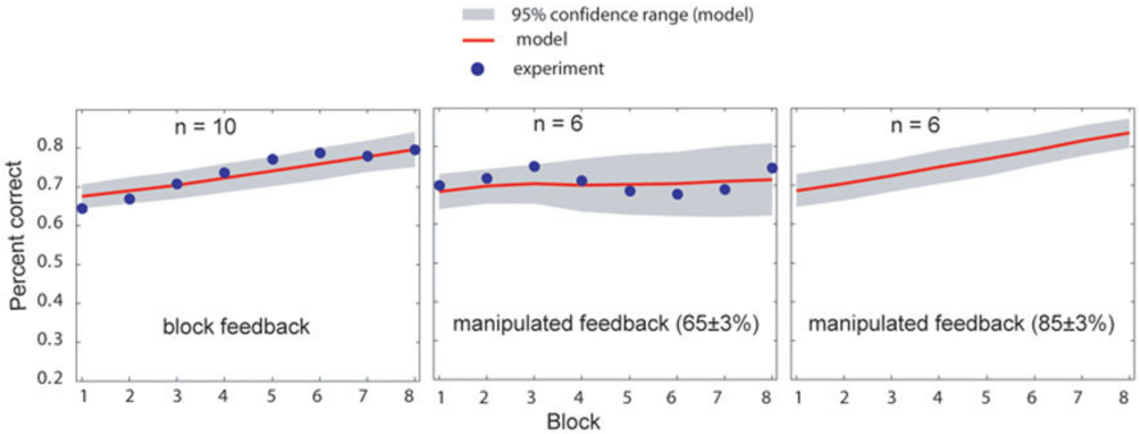

Figure 5.

Model fits to the data in the block feedback conditions. Significant performance improvement was present in the block feedback condition, but not in the pre-set low block feedback (65±3%) group, just as shown in Herzog & Fahle (1997). The right panel shows that the AHRM predicts significant learning with a hypothetical high block feedback (85±3%).

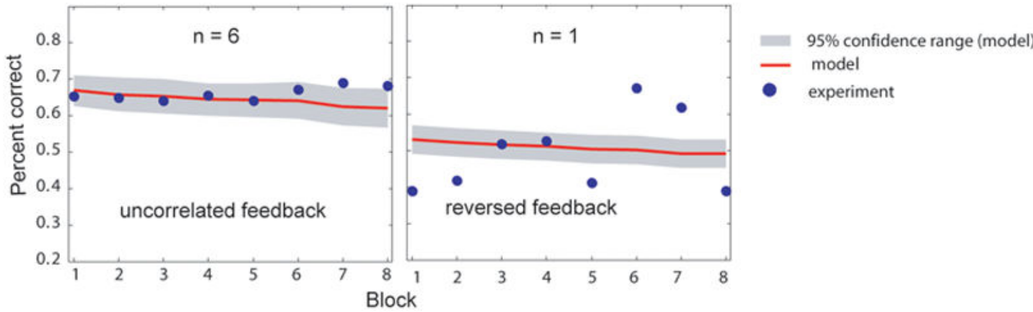

Figure 6.
Model fits in the uncorrelated feedback and reversed feedback conditions. The AHRM predicts slow performance deterioration over time and is in agreement with no performance improvement with the amount of training.

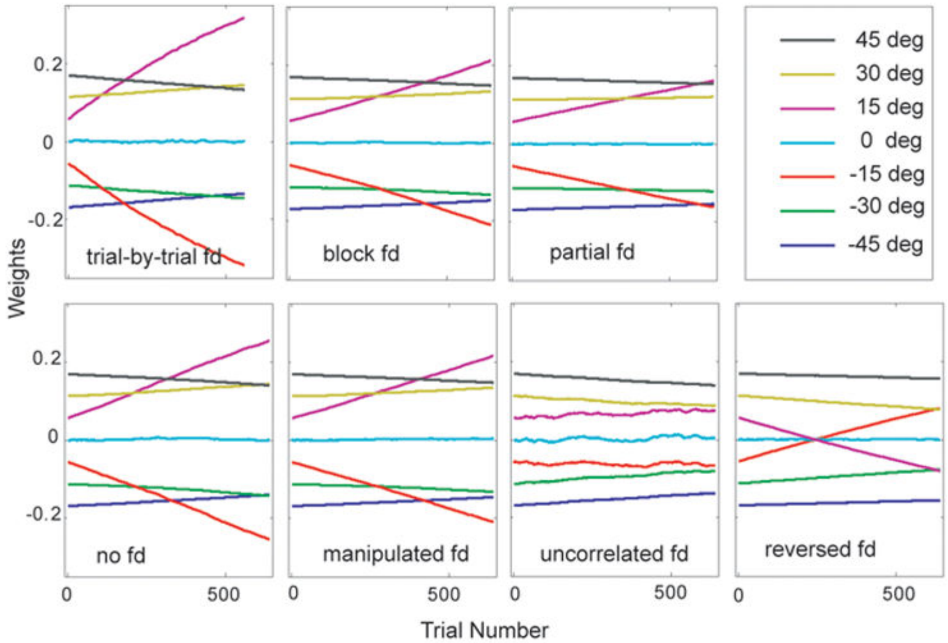

Figure 7.
ARHM channel weights over the training period, averaged over 1000 simulations. Most except the last two groups showed similar patterns: the absolute values of the weights of most relevant channels increased over training.

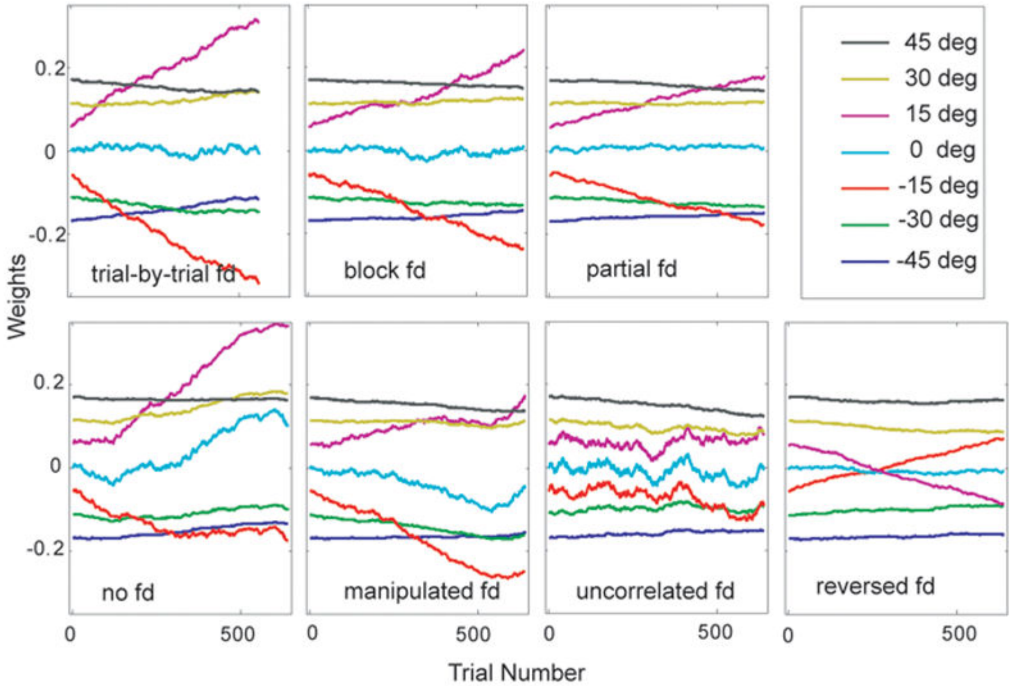

Figure 8.

ARHM channel weights over the training period from a single trace of simulation. For the groups with performance improvements (top panels), the single-trace weights changed in the same way as the averaged weights (Figure 7). For the no-feedback and manipulated feedback groups, the weights changed somewhat irregularly and became biased.

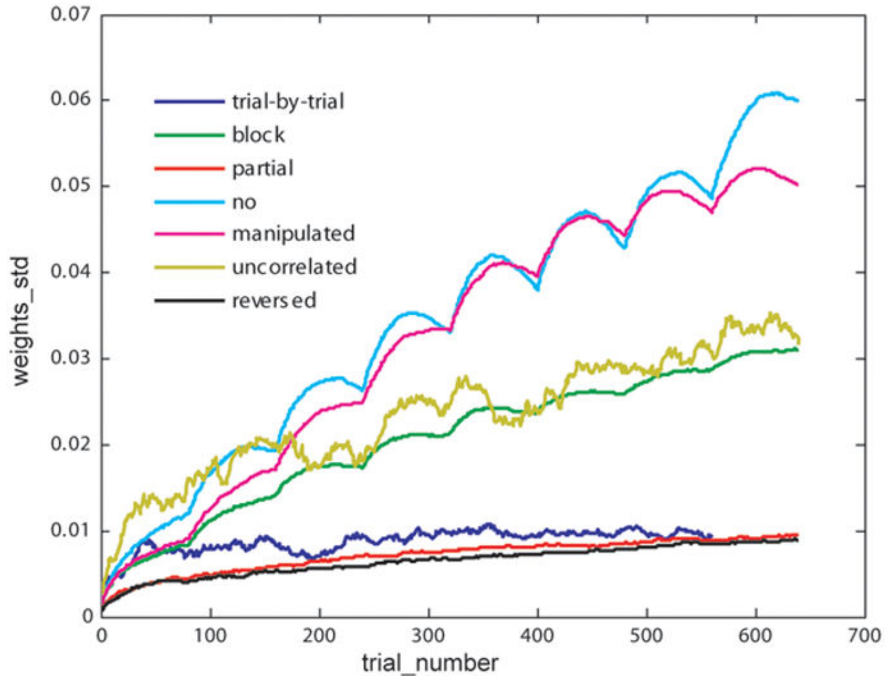

Figure 9.

The standard deviation of weights from the most relevant channel (15°, 1.4 c/deg). The no-feedback and manipulated feedback groups have larger standard deviations than block feedback groups, corresponding with variable performance. The partial and complete trial-by-trial feedback groups had small standard deviations, representing a less variable performance and improvement in accuracy with training. The uncorrelated and reversed feedback groups had small standard deviation, but also had small amplitudes of the weights (fig 7 and 8).

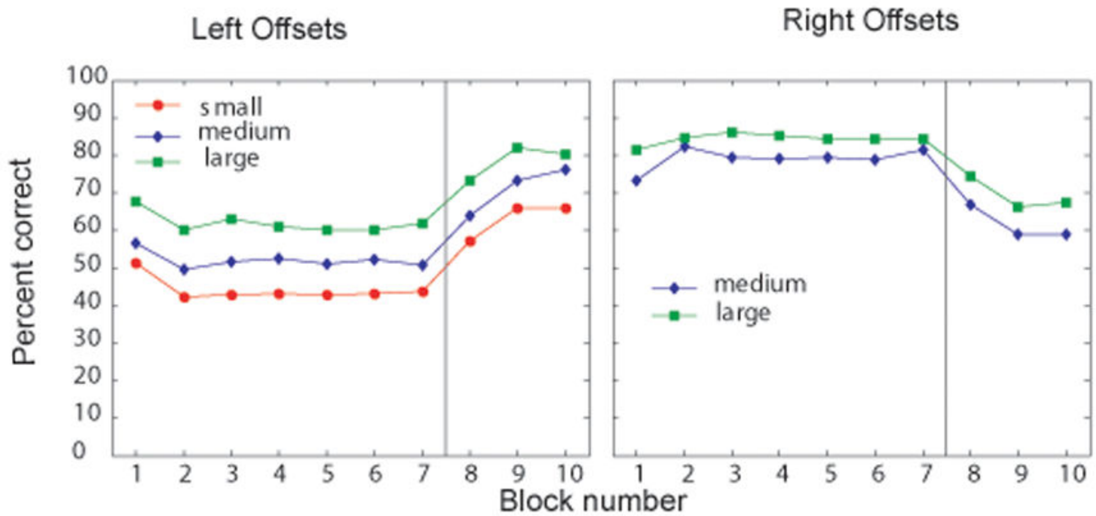

Figure 10.

The simulated results of a veriner task with biased feedback (exp.2 from Herzog & Fahle(1999); see fig. 7 & 8 in their paper). With the reversed feedback for the smallest “left” offset, performance of all left offsets dropped (left panel) while performance of right offsets increased (right panel). After the introduction of correct feedback (black line after 7^{th} block), performance of left offsets rebounded (left panel) while performance of right offsets dropped (right panel). Simulation of other experiments in the study (exp1 through exp5) show a similar pattern.

Table 1

Model parameters

	Parameter	Value						
Parameters set a priori	Orientation spacing	$\Delta\theta = 15°$						
	Spatial frequency spacing	$\Delta f = 0.5$ oct						
	Maximum activation level	$A_{max} = 1$						
	Weight bounds	$w_{min/max} = \pm 1$						
	Running average rate	$\rho = 0.02$						
	Activation function gain	$\gamma = 0.8$						
	Bias weight	$w_b = 2*pc - 1$						
	Normalization constant	$k = 0$						
	Internal additive noise	$\sigma_1 = 0$						
	Initial weight scaling factor	$w_{ini} = 0.169$						
Parameters constrained by published data	Orientation tuning bandwidth	$h_\theta = 30°$						
	Frequency tuning bandwidth	$h_f = 1.0$ oct						
	Radial Kernel width	$h_r = 2.0$ degrees of visual angle						
Parameters optimized to fit the present data		tbt fd	blk fd	par fd	no fd	man fd	unc fd	rev fd
	Representation scaling factor	a=0.26	0.44	0.23	0.52	0.46	0.46	0.10
	Internal multiplicative noise	$\sigma_m = 0.15$						
	Decision noise	$\sigma_d = 0.18$						
	Learning rate	$\eta = 0.01$						
	Feedback weight	$w_f = 0.4$						

Table 2

The summary of learning slope and amount from both experiment and model. The experimental results came from Herzog & Fahle (1997). All numbers are regression-based.

Feedback condition	Learning slope		Learning amount	
	Experiment	Model	Experiment	model
trial-by-trial*	2.13 ± 0.38	2.61 ± 0.31	14.7%	15.7 ± 1.8%
Block*	2.38 ± 0.4	1.74 ± 0.35	16.5%	12.2 ± 2.4%
Partial*	1.1 ± 0.32	0.88 ± 0.35	10.5%	6.2 ± 2.4%
no	0.23 ± 0.39	0.38 ± 0.54	2.9%	2.7 ± 3.8%
manipulated block	−0.1 ± 0.22	0.36 ± 0.72	4.3%	2.5 ± 5.0%
uncorrelated	0.02 ± 0.82	−0.82 ± 0.43	2.8%	−5.7 ± 3.0%
reversed	Not reported	−0.55 ± 0.86	Not reported	−3.8 ± 6.0%