Augmented Hebbian reweighting: Interactions between feedback and training accuracy in perceptual learning

Jiajuan Liu
Neuroscience Graduate Program, Department of Biological Sciences, University of Southern California, Los Angeles, CA, USA

Zhong-Lin Lu
Laboratory of Brain Processes (LOBES), Dana and David Dornsife Cognitive Neuroscience Imaging Center, Department of Psychology and Neuroscience Graduate Program, University of Southern California, Los Angeles, CA, USA

Barbara A. Dosher
Department of Cognitive Sciences, University of California, Irvine, CA, USA

Feedback plays an interesting role in perceptual learning. The complex pattern of empirical results concerning the role of feedback in perceptual learning rules out both a pure supervised mode and a pure unsupervised mode of learning and leads some researchers to the proposal that feedback may change the learning rate through top-down control but does not act as a teaching signal in perceptual learning (M. H. Herzog & M. Fahle, 1998). In this study, we tested the predictions of an augmented Hebbian reweighting model (AHRM) of perceptual learning (A. Petrov, B. A. Dosher, & Z.-L. Lu, 2005), in which feedback influences the effective rate of learning by serving as an additional input and not as a direct teaching signal. We investigated the interactions between feedback and training accuracy in a Gabor orientation identification task over six training days. The accelerated stochastic approximation method was used to track threshold contrasts at particular performance accuracy levels throughout training. Subjects were divided into 4 groups: high training accuracy (85% correct) with and without feedback, and low training accuracy (65%) with and without feedback. Contrast thresholds improved in the high training accuracy condition, independent of the feedback condition. However, thresholds improved in the low training accuracy condition only in the presence of feedback but not in the absence of feedback. The results are both qualitatively and quantitatively consistent with the predictions of the augmented Hebbian learning model and are not consistent with pure supervised error correction or pure Hebbian learning models.

Keywords: learning, plasticity, detection/discrimination

Citation: Liu, J., Lu, Z.-L., & Dosher, B. A. (2010). Augmented Hebbian reweighting: Interactions between feedback and training accuracy in perceptual learning. *Journal of Vision, 10*(10):29, 1–14, http://www.journalofvision.org/content/10/10/29, doi:10.1167/10.10.29.

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). Understanding how perceptual learning is achieved in the adult perceptual system might reveal the nature of plasticity in the brain and may also have profound implications for remediation of perceptual functions in visual clinical populations (Huang, Lu, & Zhou, 2008; Levi & Li, 2009; Polat, Sagi, & Norcia, 1997). In this study, we investigated the mode of perceptual learning by examining the interactions between feedback and training accuracy.

Feedback plays an interesting role in perceptual learning (see Dosher & Lu, 2009 for a review). Whereas trial-by-trial feedback is used in most perceptual learning experiments, significant perceptual learning has been observed using tasks without any external feedback (Ball & Sekuler, 1987; Crist, Kapadia, Westheimer, & Gilbert, 1997; Fahle & Edelman, 1993; Herzog & Fahle, 1997; Karni & Sagi, 1991; McKee & Westheimer, 1978; Petrov, Dosher, & Lu, 2006; Shiu & Pashler, 1992), with block feedback (Herzog & Fahle, 1997; Shiu & Pashler, 1992), or with temporally coincident feedback to an unrelated task (Seitz, Nanez, Holloway, Tsushima, & Watanabe, 2006; Seitz & Watanabe, 2003; Watanabe, Nanez, & Sasaki, 2001; Watanabe et al., 2002). 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;

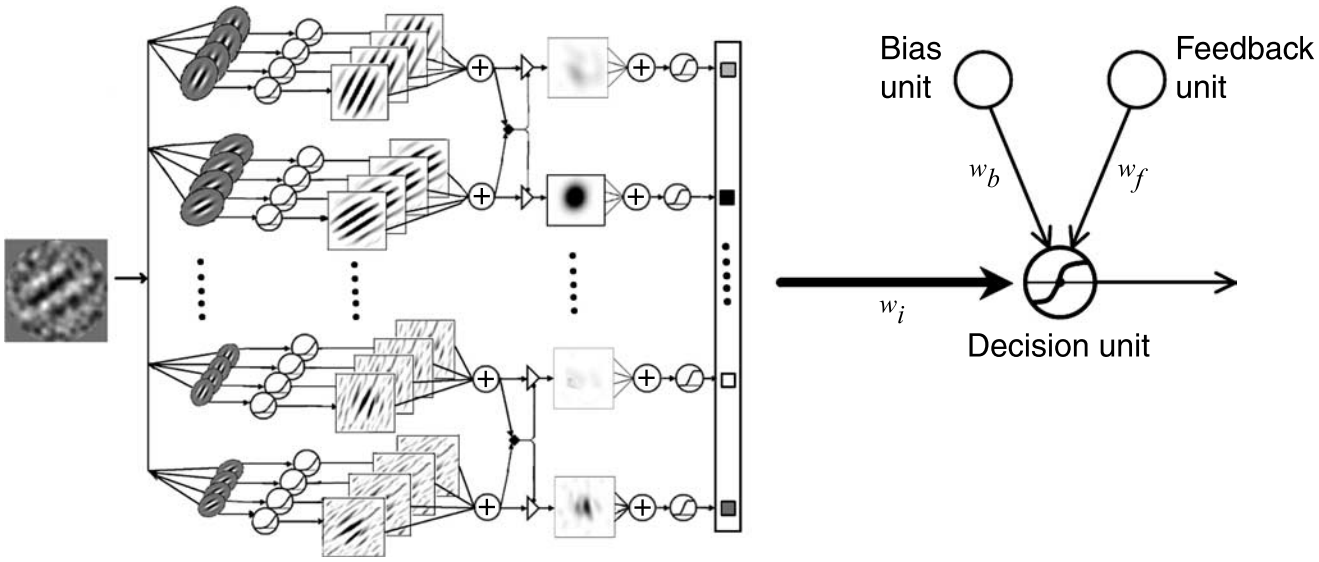

Figure 1. The augmented Hebbian reweighting model (AHRM). (adopted from Petrov et al., 2005, 2006).

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). A recent study (Shibata, Yamagishi, Ishii, & Kawato, 2009) also found that fake block feedback, if more positive than observers' actual performance, enhanced learning; however, if this fake feedback underestimated observers' actual performance, learning was not affected. The complex pattern of empirical results concerning the role of feedback in perceptual learning rules out both a pure supervised mode (Hertz, Krogh, & Palmer, 1991) and a pure unsupervised mode of learning (Polat & Sagi, 1994; Vaina, Sundareswaran, & Harris, 1995, Weiss, Edelman, & Fahle, 1993). In the pure supervised mode, explicit trial-by-trial feedback serves as the teaching signal to update the connection weights according to plasticity rules that minimize the average error between actual and target activations of the output units. No learning is possible without trial-by-trial feedback. In contrast, unsupervised learning is associated with Hebbian learning rules that update the connection weights on the basis of co-activation of input and output (pre- and postsynaptic) units. These rules do not depend on feedback and detect statistical regularities in the training corpus. Herzog and Fahle (1997) concluded that "both supervised and unsupervised (feed-forward) neural networks are unable to explain the observed phenomena and that straightforward *ad hoc* extensions also fail." Herzog and Fahle (1998) suggested that feedback changes the learning rate through (unspecified) top-down control but does not act as a teaching signal in perceptual learning (also see Shibata et al., 2009).

Inspired by Herzog and Fahle (1998), and to serve as an implementation of the reweighting hypothesis (Dosher & Lu, 1998, 1999), Petrov, Dosher, and Lu (2005; Petrov et al., 2006) proposed an augmented Hebbian reweighting model (AHRM) of perceptual learning (Figure 1). The model is a hybrid system where external feedback is included as an additional input in Hebbian learning (leading to the term "augmented Hebbian"). The augmented Hebbian learning algorithm incorporates external feedback, when present, simply as another input to the decision unit that shifts the postsynaptic activation in the correct direction and fosters appropriate weight changes. Without feedback, the model uses the internal response of the observer to update the weights: the weights still move in the correct direction on average because the activation of the decision unit correlates with the correct stimulus classification. Consistent with the general idea outlined in Herzog and Fahle (1998) but implemented with biologically plausible plasticity mechanisms, feedback changes the effective learning rate but does not directly serve as teaching signals in the AHRM.

The AHRM was applied by Petrov et al. (2006) to successfully model perceptual learning with and without trial-by-trial feedback in an orientation discrimination experiment under destabilizing non-stationary task conditions. Using an experimental procedure that is identical to Petrov et al. (2005) except without trial-by-trial feedback, Petrov et al. (2006) measured perceptual learning in a peripheral Gabor orientation discrimination task (±10 deg) in two filtered noise "contexts" with predominate orientations at either ±15 deg that alternated every two training days. Three target contrast levels (and therefore three initial performance levels) were tested in a mixed design. Eighteen observers received no feedback, yet improved both discriminability and speed within and across blocks. The initial and asymptotic d' levels and learning dynamics were comparable to those obtained for observers with feedback (Petrov et al., 2005). Petrov et al. (2006) found that the AHRM, which tracks the external feedback when available, or else reinforces the model's

own response, along with a criterion control bias input, can account for all the experimental results.

In this study, we tested another major prediction of the AHRM—a strong interaction between feedback and training accuracy in perceptual learning. Without feedback, when the training accuracy is above chance for a sufficient fraction of the training trials, Hebbian learning capitalizes on this positive correlation and predicts learned performance improvements. When the training accuracy is low, however, Hebbian learning can be erratic, slow, or even fail altogether. On the other hand, when trial-by-trial feedback is available, the reliance on training accuracy is greatly reduced because external feedback shifts the postsynaptic activation in the correct direction and supports learning. We tested these predictions in this study. The critical manipulation is the level of training accuracy combined with the presence or absence of feedback.

Adult observers practiced in a Gabor orientation identification task in fovea over six training days. They were divided into four equal-size groups: high training accuracy (85% correct) with and without feedback, and low training accuracy (65%) with and without feedback. The accelerated stochastic approximation method (Kesten, 1958) was used to track threshold contrasts at particular performance accuracy levels and keep performance within a narrow range of the target training accuracy. Training accuracy and feedback were jointly manipulated. The augmented Hebbian reweighting model (AHRM) of perceptual learning (Petrov et al., 2005, 2006) was used to account for the experimental results.

Methods

Observers

Twenty-four naive observers, all undergraduate or graduate students at the University of Southern California, participated in the study. They were randomly assigned into four experimental groups: high training accuracy with and without feedback, and low training accuracy with and without feedback. All observers had normal or corrected-to-normal vision and provided written consent under the USC Institutional Review Board protocol.

Apparatus

All the experiments were conducted on a Macintosh Power PC G4 computer running Matlab with Psychtoolbox extensions (Brainard, 1997; Pelli, 1997). Visual stimuli were displayed on a Nanao Technology Flexscan 6600 monitor with a P4 phosphor, a 480 × 640 spatial resolution, and a refresh rate of 120 Hz. Fine control of luminance levels was achieved through a special circuit, which combined two 8-bit outputs of a video card to produce 6144 (12.6 bit) distinct gray levels (Pelli & Zhang, 1991). A lookup table, generated with a psychophysical procedure, provided a linear transformation of pixel value to display luminance (Li, Lu, Xu, Jin, & Zhou, 2003). All displays were viewed binocularly with natural pupil at a viewing distance of approximately 72 cm in a dimly lit room. A chin rest was used to help observers maintain their head positions.

Stimuli

The signal stimuli were Gaussian-windowed sinusoidal gratings, oriented ±10° from 45°. The luminance profile of the Gabor stimulus is described by

$$L(x,y) = L_0\left(1.0 + c\sin(2\pi f(x\cos\theta + y\sin\theta))e^{-\frac{x^2+y^2}{2\sigma^2}}\right), \quad (1)$$

where c is the signal contrast, L_0 is the background luminance, set in the middle of the dynamic range of the display (L_{min} = 1 cd/m^2; L_{max} = 53 cd/m^2), f = 1.29 c/d is the center spatial frequency of the Gabor, and σ = 0.77 deg is the standard deviation of the Gaussian window. The Gabors were rendered on a 64 × 64 pixel grid, extending 3.09° × 3.09° of visual angle (Figure 2).

External noise images, 64 × 64 pixels (3.09° × 3.09°), were constructed using 2 × 2 pixel elements (0.097° × 0.097°); 2 × 2 pixel elements were used in external noise images to increase the overlap of the Fourier energy spectra of the signal and external noise stimuli and therefore the impact of external noise on observer performance. In every trial, the contrasts of all the noise elements were drawn randomly and independently from the same Gaussian distribution with mean 0 and standard deviation 0.25. The external noise images were then filtered through a second-order Butterworth band-pass filter with cutoff frequencies of 0.70 c/d and 2.82 c/d. Because the display contrast ranged from −1.0 to 1.0, a sample with standard deviation of 0.33 conforms reasonably well to a Gaussian distribution. Both signal and noise images were centered at fixation and combined via temporal integration.[1]

Design and procedure

A two-by-two between-subject design was used in this study. Both training accuracy (85% and 65% correct) and feedback (trial-by-trial and no feedback) were manipulated across four independent experimental groups.

Following a key press, each trial started with a fixation display that lasted 150 ms, followed by an external noise, a signal, and an independent external noise image, each lasting 33 ms, and a blank window that lasted until the end

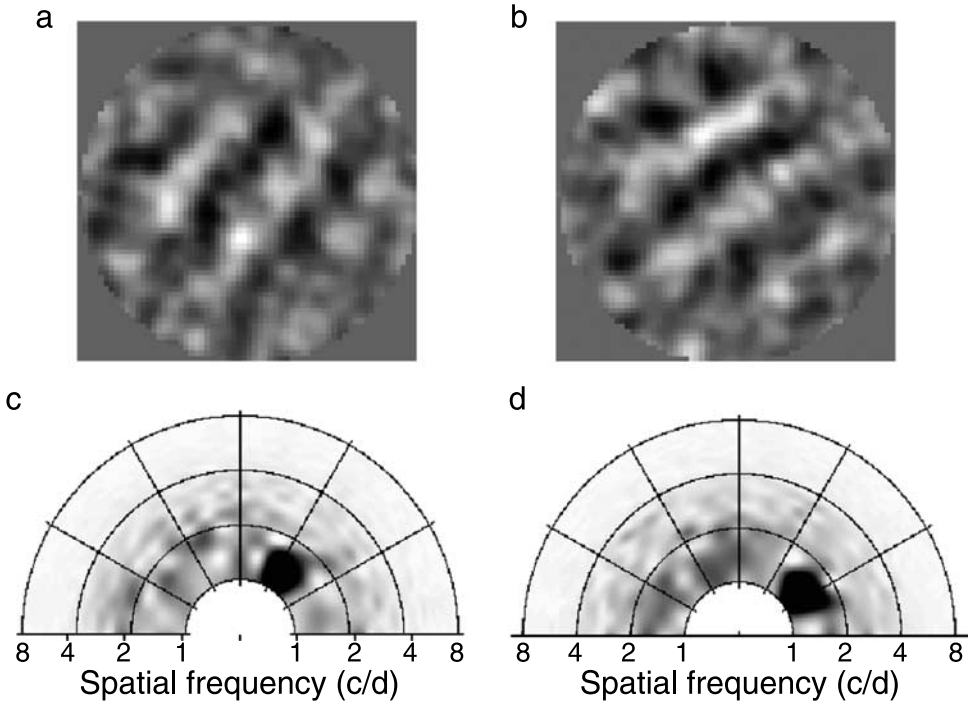

Figure 2. Examples of a "Left" (top left) and a "Right"(top right) stimulus, and their corresponding power spectra (bottom, left and right). Gabor targets oriented ±10° from 45° direction are embedded in external Gaussian noise. The Gabor spatial frequency is 1.29 cyc/deg. The spectra are plotted in polar coordinates, with log frequency along the radial axis and orientation along the angular axis.

of the trial. Observers responded by pressing different keys on the computer keyboard to indicate whether the Gabor was tilted clockwise ("Right") or counterclockwise ("Left") relative to the 45° reference angle. An auditory beep followed each correct response in the trial-by-trial feedback groups. No feedback was provided to the other two groups.

Because training accuracy is a critical experimental variable in this study, we used the accelerated stochastic approximation method (Kesten, 1958), an adaptive procedure that converges to any specified accuracy level, to control performance level throughout the course of training. The stimulus contrast in every trial was determined by the stochastic approximation procedure (Robbins & Monro, 1951). To converge to a target performance level ϕ, stimulus contrasts in the first two trials are given by

$$c_{n+1} = c_n - \frac{s}{n}(Z_n - \phi), \quad (2)$$

where n is the trial number, c_n is the stimulus contrast in trial n, Z_n = 0 (incorrect) or 1 (correct) is the response accuracy in trial n, c_{n+1} is the stimulus contrast for the next trial, and s is the pre-chosen step size at the beginning of the trial sequence. From the third trial onward, the sequence is "accelerated":

$$c_{n+1} = c_n - \frac{s}{2 + m_{\text{shift}}}(Z_n - \phi), \quad (3)$$

where m_{shift} is the number of shifts in response category (switches from consecutive correct responses to incorrect responses and vice versa). In an influential review of adaptive psychophysical procedures, Treutwein (1995) recommended the accelerated stochastic approximation procedure as the best available procedure for measuring thresholds. Simulation studies prior to the experiments suggest that the optimal s is the same as c_1 and the optimal c_1 is the threshold, which can be estimated from a QUEST procedure (see below).

Each 20-min experimental session consisted of four blocks of 80 trials each. Each observer ran six sessions in six training days. In the beginning of the study, several instruction trials using Gabor stimuli without external noise were used to familiarize the observer with the task. A QUEST procedure (Watson & Pelli, 1983) with 60–80 trials was then used to estimate the starting contrast for each subject. After that, an accelerated stochastic staircase method (Kesten, 1958) was used to estimate subjects' threshold for every block of 80 trials. The initial contrast in every block was set to be that of the last trial in the previous block.

The augmented Hebbian model

The AHRM is a multi-channel neural network model implementation of the channel reweighting hypothesis

outlined in Dosher and Lu (1998), originally developed to model the detailed learning dynamics and recurring switch costs of perceptual learning in non-stationary contexts (Petrov et al., 2005). It consists of four types of units (Figure 1): representation units that encode input images as activation patterns, a task-specific decision unit that receives weighted inputs from the representation units, an adaptive bias unit that accumulates a running average of the response frequencies and works to balance the frequency of the two responses, and a feedback unit that makes use of external feedback when (and if) it is presented. Learning in the model occurs exclusively through incremental Hebbian modification of the weights between representation units and the decision unit; while the early processing pathway that constructs representations from the retinal image remains fixed throughout training. This simple and powerful reweighting mechanism accounts for the detailed learning dynamics and the recurring switch costs in the d' curves with and without feedback (Petrov et al., 2005, 2006). Detailed descriptions of the augmented Hebbian reweighting model can be found in Petrov et al. (2005, 2006). Here, we briefly review the major computations in the model.

Representation subsystem

The representation subsystem consists of 7 × 5 orientation- and frequency-selective units (Figure 1). The activation $A(\theta, f)$ of each of the 35 representation units encodes the normalized spectral energy in the corresponding orientation and frequency channel. First, units tuned to different orientations ($\theta \in \{0°, 15°, 30°, 45°, 60°, 75°, 90°\}$; half-amplitude full-bandwidth $h_\theta = 30°$), spatial frequencies ($f \in \{0.7, 1, 1.4, 2, 2.8$ c/d$\}$; $h_f = 1$ octave for spatial frequency), and spatial phases ($\phi \in \{0°, 90°, 180°, 270°\}$) compute a set of retinotopic *phase-sensitive maps* $S(x, y, \theta, f, \phi)$ of the input image $I(x, y)$:

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

Specifically, the input image $I(x, y)$ is convolved with all the 140 tuned units via fast Fourier transformation. The output images are then rectified by a half-squaring operator to generate phase-sensitive maps, which can be interpreted as activation patterns across a large retinotopic population of "simple cells" (Heeger, 1992). Because spatial phase is not relevant in this task (and to simplify the representations), the model aggregates across phases in channel at each spatial location and uses shunting inhibition to obtain normalized outputs (Heeger, 1992):

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

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

where ε_1 represents a Gaussian-distributed internal noise source with mean 0 and standard deviation σ_1, the normalization pool $N(f)$ is assumed to be essentially independent of orientation and modestly tuned for spatial frequency, a is a scaling factor, and k is the saturation constant that is relevant only at near-zero contrasts.

We further pooled the energy maps across space, weighted by a radially symmetric Gaussian kernel W_r with full-width at half-height $h_r = 2.0$ degrees, commensurate with the size of the target Gabor:

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

where ε_2 represents another Gaussian-distributed noise source with mean 0 and standard deviation σ_2.

Finally, an activation function with gain parameter γ was used to limit the dynamic range of the representation units:

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

Task-specific decision subsystem

The decision unit aggregates the sensory information using the current weights w_i and the current top-down bias b:

$$u = \sum_{i=1}^{35} w_i A(\theta_i,f_i) - w_b b + \varepsilon_d, \quad (9)$$

where Gaussian noise ε_d with mean 0 and standard deviation σ_d models the random fluctuations in the decision-making process.

The early activation o' of the unit is a sigmoidal function of the early input u with gain γ:

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

$$o' = G(u) \text{ (early)}. \quad (11)$$

The model generates a "left" response if o' is negative, and a "right" response if o' is positive.

Augmented Hebbian learning

Following the response of the task-dependent reweighting system, feedback—if present—is encoded by the

feedback unit and sent as a top-down input to the decision unit. This new input F adds to the early input u driving the decision unit, which changes its activation to a new, late level o according to the following equation:

$$o = G(u + w_f F) \text{ (late)}. \quad (12)$$

All learning happens during this late phase (O'Reilly & Munakata, 2000).

The impact of feedback depends upon the weight w_f on the feedback input. The late activation is driven to $\pm A_{\max} = \pm 1$ when feedback $F = \pm 1$ is present and the feedback weight is relatively high. Lower feedback weights may simply shift the activation slightly. When no feedback signal is present ($F = 0$), the late decision activation is the same as the early decision activation ($o = o'$), which typically is in the intermediate range. In other words, the model uses the internal response of the observer to update the weights: the weights still move in the correct direction on average because the activation of the decision unit correlates with the correct stimulus classification.

In the AHRM, the only mechanism for long-term changes due to learning operates on the synaptic strengths w_i of the connections between the sensory units i and the decision unit. The Hebbian rule is exactly the same with and without feedback. Each weight change depends on the activation a_i of the presynaptic sensory unit and the activation o of the postsynaptic decision unit relative to the baseline $\bar{o}$:

$$\delta_i = \eta A(\theta_i, f_i)(o - \bar{o}), \quad (13)$$

$$\Delta w_i = (w_i - w_{\min})[\delta_i]_- + (w_{\max} - w_i)[\delta_i]_+, \quad (14)$$

$$\bar{o}(t+1) = \rho o(t) + (1-\rho)\bar{o}(t). \quad (15)$$

Equation 13 corrects the postsynaptic activation o by its long-term average $\bar{o}$. Thus, the Hebbian term δ_i tracks systematic stimulus-response correlations rather than mere response bias.

Adaptive criterion control

The Hebbian learning process is also augmented by a mechanism for adaptive criterion control based on self-monitoring of relative response frequencies and implemented as a top-down input from a bias unit b with weight w_b (Equation 9). Observers are assumed to approximately equalize the frequencies of recent "Left" and "Right" responses to approximately match the presentation probabilities in the experiment. The bias $b(t + 1)$ on each successive trial equals the current weighted running average $r(t)$, discounting the distant past exponentially with a time constant of 50 trials ($\rho = 0.02$):

$$r(t+1) = \rho R(t) + (1-\rho)r(t), \quad (16)$$

$$b(t+1) = r(t). \quad (17)$$

The adaptive criterion control may be more or less prominent in various circumstances, which is captured by a parametric bias weight w_b in the model. For example, Petrov et al. (2006) found a higher weight on adaptive criterion control in the presence of feedback. Although their effects may interact in subtle ways, the bias and the feedback inputs are structurally separate—the former tracks the model's own response R, whereas the latter tracks the external feedback signal F.

The adaptive criterion control unit is crucial for learning in certain circumstances when trial-by-trial external feedback is not available. Petrov et al. (2006) showed in simulation that the unit was necessary to reduce performance bias induced by the learning context and could be especially critical during task switching. The unit is important in the current study because trial-by-trial feedback was absent in the "no feedback" conditions.

Model fits

The augmented Hebbian reweighting model (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. The model parameters are listed in Table 1. Seven parameters, including four scaling factors (a), one for each group, internal multiplicative noise (σ_2), decision noise (σ_d), and learning rate η, were adjusted to fit the experimental data. Three parameters, orientation tuning bandwidth (h_θ), frequency tuning bandwidth (h_f), and radial kernel width (h_r) were set a priori based on publications in the literature. All other parameters were set a priori based on Petrov et al. (2005). The scaling factors and internal noises allow a match to the overall performance shown in the initial performance measure. Critically, a single learning rate η was used to model the learning curves in all four experimental conditions, and the predicted differences between conditions entirely reflect differential effectiveness of Hebbian learning and feedback in these conditions.

Each iteration of parameter adjustment takes place in two steps. First, the initial weights were set in proportion to the preferred orientation of the units: $w_i = (\theta_i/30°)w_{\text{init}}$, reflecting general prior knowledge about orientation. The weights were fixed. Six out of the seven "free-to-vary" parameters, four as, σ_2, and σ_d, were adjusted, with two of

	Parameter	Value			
Parameters set a priori	Orientation spacing	$\Delta\theta$ = 15°			
	Spatial frequency spacing	Δf = 0.5 oct			
	Maximum activation level	A_{max} = 1			
	Weight bounds	$w_{min/max}$ = ±1			
	Running average rate	ρ = 0.02			
	Activation function gain	γ = 5			
	Bias weight	w_b = 2.2			
	Normalization constant	k = 0			
	Internal additive noise	σ_1 = 0			
	Initial weight scaling factor	w_{ini} = 0.169			
	Feedback weight	w_f = 1.0			
Parameters constrained by published data	Orientation tuning bandwidth	h_θ = 30°			
	Frequency tuning bandwidth	h_f = 1.0 oct			
	Radial kernel width	h_r = 2.0 dva			
Parameters optimized to fit the present data		65% w	65% wo	85% w	85% wo
	Representation scaling factor	a = 0.090	0.085	0.14	0.18
	Internal multiplicative noise	σ_2 = 0.16			
	Decision noise	σ_d = 0.20			
	Learning rate	η = 0.00025			

Table 1. Model parameters.

them (internal multiplicative noise and decision noise) restricted to be the same for all 4 groups. The scaling factor was free to vary among the groups such that the model performance can match that of the human observers in the beginning of the experiment, and so reflects small random differences in performance level for these randomly assigned groups.

The adjustment was based on least squared error:

$$L = \sum [\log(c_\tau^{theoretical}) - \log(c_\tau^{measured})]^2, \quad (18)$$

where c_τ^{measured} and $c_\tau^{\text{theoretical}}$ represent measured and model-generated contrast thresholds, and $\sum$ represents summation over the first data point across all four (training accuracy × feedback) experimental conditions.

In the second step, the weights were allowed to change. The output of the decision unit and/or the external feedback was used to update the weights depending on the specific feedback condition. The accelerated stochastic approximation procedure was used to generate the training sequence. The model performance was then compared to that of the human observers using least squared error defined in Equation 18 with summation over all the training sessions in the four experimental conditions. The two steps were repeated until the model predictions were reasonably matched to the data.

For every experimental condition, the model, just as the human observers, ran 24 blocks with 80 trials/block in each simulated experiment. This 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 six simulations to generate the average performance curve of six simulated observers. This was repeated 1000 times. Following standard practice in bootstrap, we computed the mean and standard deviations of the slope 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

Learning curves

The accelerated stochastic procedure did an excellent job in tracking constant target training accuracy levels during the experiment. The average performance accuracy for every forty trials is plotted in Figure 3 for the four groups of observers. All the data points were within ±5% of the target accuracy.

Learning curves, log(contrast threshold) versus log (training block), are plotted in Figure 4. Observers in the high training accuracy (85% correct) condition, regardless of their feedback condition, exhibited significant threshold reduction; the slopes of their learning curves were significantly less than zero ($p < 0.01$). However, for observers in the low training accuracy (65% correct) condition, perceptual learning depended on the availability of trial-by-trial feedback—only observers in the trial-by-trial feedback condition exhibited significant threshold reduction ($p < 0.01$); observers in the no-feedback condition did not show significant threshold reduction ($p > 0.10$).

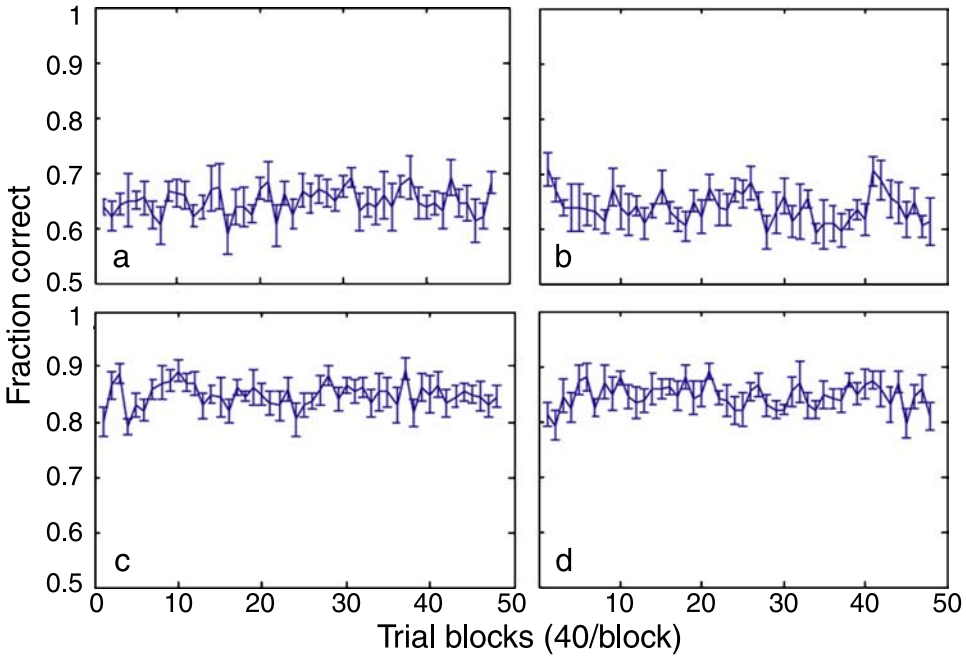

Figure 3. The proportional correct of every 40 trials for all 4 groups. Every point in a graph is the average of the proportion correct of 6 subjects in a group and the bar shows the standard error of the mean. (a) 65% Correct (low accuracy) with feedback. (b) 65% Correct (low accuracy) without feedback. (c) 85% Correct (high accuracy) with feedback. (d) 85% Correct (high accuracy) without feedback.

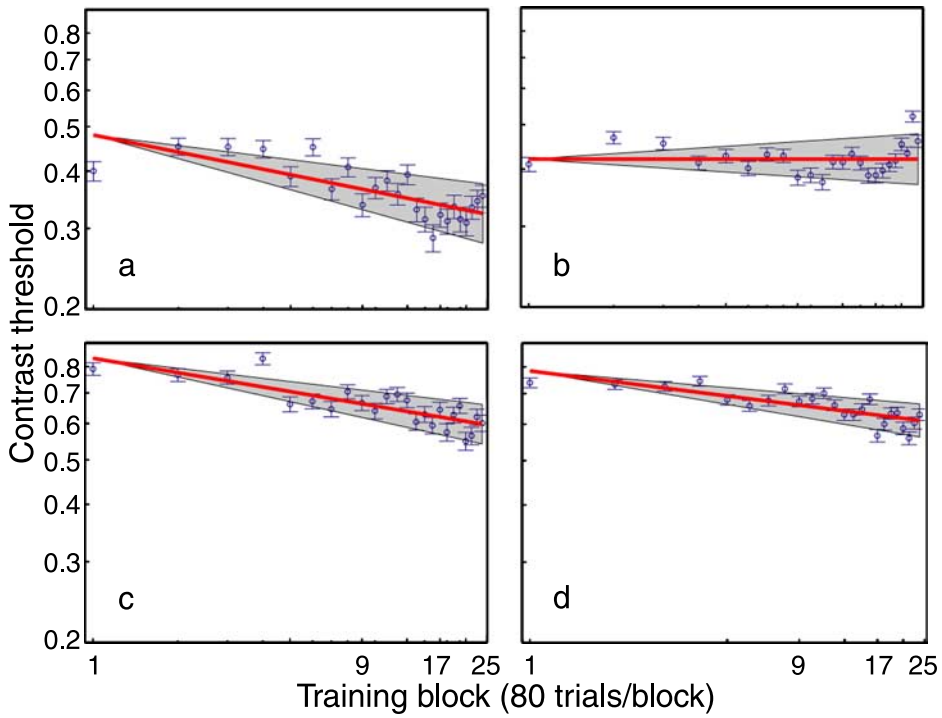

Figure 4. Average learning curves for all four groups. Every point in a graph is the average of the contrast thresholds of 6 subjects in a group and the error bar shows the standard error of the mean. The red line is the linear regression of the contrast threshold over the training blocks. A negative slope indicates the existence of learning, showing the drop of contrast threshold over time. The gray area in each panel covers the 95% range of confidence of the regressed slope. (a) 65% Correct (low accuracy) with feedback. (b) 65% Correct (low accuracy) without feedback. (c) 85% Correct (high accuracy) with feedback. (d) 85% Correct (high accuracy) without feedback.

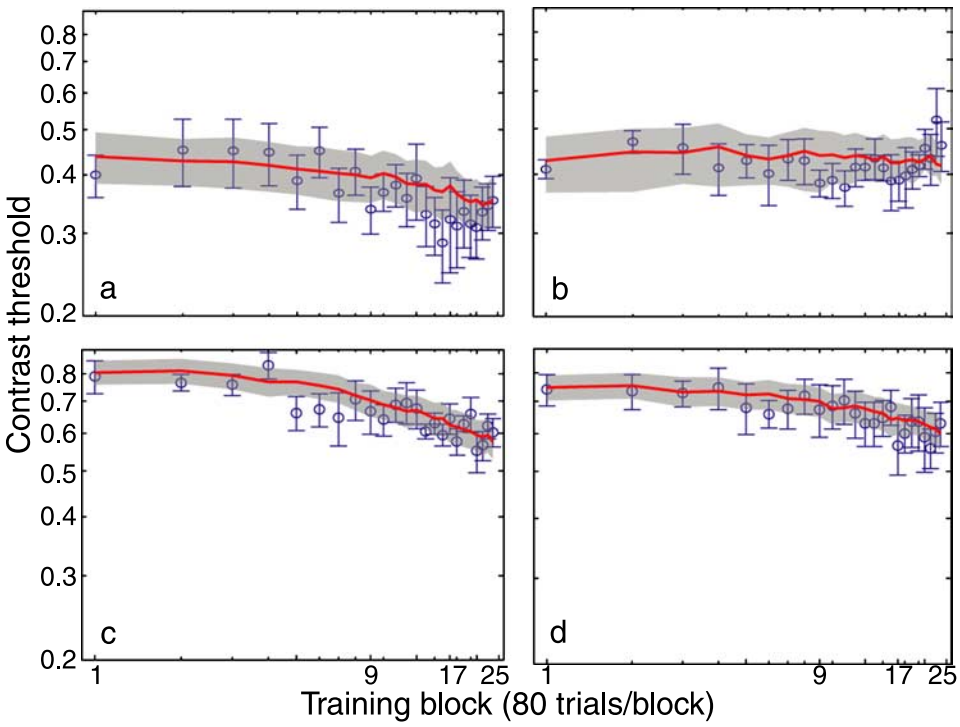

Figure 5. Model fits to the behavioral results. The points and error bars represent the average contrast thresholds and the standard error of the mean in each training condition. The red lines represent performance of the simulated AHRM. Shaded areas represent ±1 *SD* of the model performance. (a) 65% Correct (low accuracy) with feedback. (b) 65% Correct (low accuracy) without feedback. (c) 85% Correct (high accuracy) with feedback. (d) 85% Correct (high accuracy) without feedback.

The average magnitudes of learning, calculated as percent threshold reduction between the last and first training blocks, were 33.2 ± 8.0%, 0.7 ± 14.0%, 27.5 ± 9.0%, and 23.2 ± 8.0%, for the 65% training accuracy with feedback and no feedback, and 85% training accuracy with and without feedback groups, respectively. The slopes of the learning curves for the four groups were: −0.13 ± 0.04 (*SD*), 0.00 ± 0.04, −0.10 ± 0.04, −0.09 ± 0.04. Analysis of variance on the slopes found significant differences among the four different groups ($p < 0.001$). Tukey post-hoc analysis ($\alpha = 0.05$) identified a significant difference only between the 65% training accuracy without feedback group and all other three groups; there was no significant difference among those three. Consistent with Herzog and Fahle (1998) and Petrov et al. (2006), we conclude that the patterns of results are not consistent with either a pure supervised or a pure unsupervised learning model.

AHRM fits

The AHRM provided an excellent account of the data in all four experimental conditions. The predicted learning curves of the Hebbian reweighting model are plotted in Figure 5 along with the behavioral contrast threshold data. Quantitatively, the model accounted for 93.1% of the variance. The pattern of model performance was essentially the same as that of the human observers: In the model, the average magnitudes of learning, calculated as percent threshold reduction between the last and first training blocks, were 21.8 ± 9.0%, 1.9 ± 12.0%, 29.3 ± 5.0%, and 19.9 ± 5.0%, for the 65% training accuracy with feedback and no feedback, and the 85% training accuracy with and without feedback conditions, respectively. The slopes of the learning curves for the four groups were: 0.09 ± 0.04 (*SD*), −0.02 ± 0.04, −0.12 ± 0.02, −0.08 ± 0.02. Analysis of variance on the simulation results identified significant differences among the four different groups ($p < 0.001$). Tukey post-hoc analysis ($\alpha = 0.05$) identified a significant difference only between the 65% training accuracy group without feedback and all other three groups; there was no significant difference among those three.

The weight dynamics were qualitatively similar with and without feedback (Figure 6). The initial weights embody general directional information, yet carry very little information about which spatial frequency and orientation channel contains the target. With practice, the weights of the different channels are adjusted to match the statistical structure of the stimulus environment. The most significant weight optimization occurred in the 85% correct training accuracy with feedback condition. The weight changes in the 65% correct with feedback condition are very similar to those in the 85% correct without feedback group, reflecting a complementary interaction between feedback and training accuracy. The

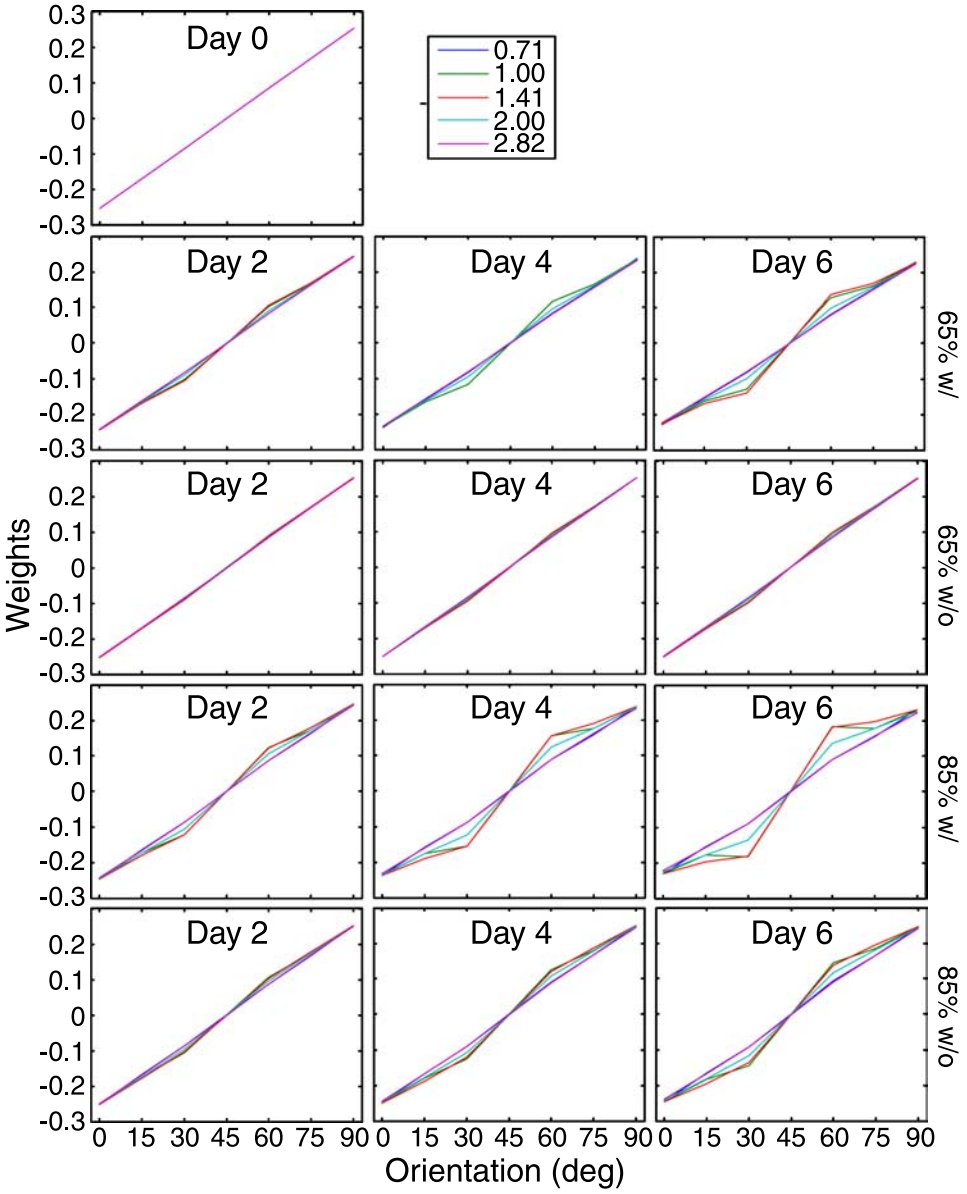

Figure 6. Weight dynamics of the AHRM. The first row is the initial weights before any training, which are the same for all the groups. The second to fifth rows show the weights after 2, 4, and 6 days of training for all the groups, with each group labeled on the right side of the figure. The different colors of the curves represent the channels with different spatial frequencies as shown in the inset.

differences between the 85% correct with feedback condition and the 65% correct with feedback condition, and with the 85% correct without feedback conditions, are quite small and were not statistically distinguishable in our experiments.

Upon request of a reviewer, we also performed an ideal observer analysis on the Gabor orientation identification task. Contrast thresholds of the ideal observer were calculated from simulations in which the ideal observer decided the orientation of the Gabor based on the outputs of two templates that were matched to the two target Gabors. Specifically, we computed the dot products of the two templates with input stimulus (signal Gabor + external noise) in a range of Gabor contrast conditions and extracted contrast thresholds of the ideal observer, $c_{\text{ideal}}(Pc)$, at both 65% and 85% correct performance levels. The sampling efficiency of the human observer was then calculated as (Barlow, 1956)

$$\upsilon(Pc) = \frac{c^2_{\text{ideal}}(Pc)}{c^2_{\text{human}}(Pc)}. \tag{19}$$

Prior to perceptual learning, the sampling efficiencies of the human observers were 6.3%, 8.2%, 6.3%, and 7.2% in the 65% correct with and without feedback, and 85% correct with and without feedback groups, respectively.

The different sampling efficiencies reflect small average differences between the observers randomly assigned to the four groups. At the end of training, the sampling efficiencies became 14%, 8.2%, 12%, and 12% in the four groups. Consistent with the pattern of results based on contrast thresholds, all except the 65% correct without feedback group significantly improved their sampling efficiencies.

Discussion

In this study, we found a significant interaction between the presence of feedback and criterion or target training accuracy in perceptual learning: When the training accuracy is sufficiently high, performance accuracy is relatively high, and external feedback is not critical for perceptual learning. However, when the training accuracy is low, external feedback is essential for perceptual learning. The pattern of results cannot be explained by either a pure supervised learning model (Poggio, Fahle, & Edelman, 1992) or a pure unsupervised learning model (Polat & Sagi, 1994; Vaina et al., 1995; Weiss et al., 1993). Instead, performance is qualitatively and quantitatively consistent with the predictions of an augmented Hebbian reweighting model (Petrov et al., 2006), which, as suggested by Herzog and Fahle (1998), uses feedback not as a teaching signal but a positive influence on the weight changes. (Note, however, that the specific form of influence of feedback in the AHRM is related but not identical to the proposals of Herzog and Fahle.)

Like all its predecessor computational models on perceptual learning (Hertz et al., 1991, Herzog & Fahle, 1998; Polat & Sagi, 1994, Vaina et al., 1995, Weiss et al., 1993), the AHRM learns to improve its performance by maintaining stable early representations of visual inputs and incrementally changing the connection weights between the representation units and the decision unit. Although perceptual learning in the visual domain has been widely claimed to reflect long-lasting plasticity of sensory representations in early visual cortex (Gilbert, Sigman, & Crist, 2001; Seitz & Watanabe, 2005; Tsodyks & Gilbert, 2004), there is increasing evidence supporting the proposal (Dosher & Lu, 1998, 1999; Jacobs, 2009; Law & Gold, 2008; Mollon & Danilova, 1996) that the behavioral expression of specificity in perceptual learning in visual system reflects reweighted decisions, or changed readout, from stationary visual sensory representations. Physiological analysis using single cell recording has documented remarkable robustness and stability of early visual representations in the face of training (Ghose, 2004; Ghose, Yang, & Maunsell, 2002; Law & Gold, 2008; Yang & Maunsell, 2004; but see Hua et al., 2010). In the visual domain as distinct from other sensory or sensory motor domains, reweighting or altered readout at higher levels of visual system is perhaps the dominant mode of perceptual learning (see Dosher & Lu, 2009; Petrov et al., 2005 for a review).

In addition to its application in this study, the AHRM has been used to successfully account for complex patterns of perceptual learning in an orientation discrimination experiment under destabilizing non-stationary manipulations both with and without trial-to-trial feedback (Petrov et al., 2005, 2006). The model has recently been used to account for a large number of data patterns in external noise studies of perceptual learning (Lu, Liu, & Dosher, 2010). It also predicts the Eureka effect in perceptual learning, i.e., mixing easy trials in a training task facilitates learning (Ahissar & Hochstein, 1997; Rubin, Nakayama, & Shapley, 1997), a prediction that has been tested and confirmed in a recent experiment (Liu, Lu, & Dosher, 2009).

We believe that the AHRM is broadly compatible with the complex empirical literature on the role of feedback in perceptual learning. As argued in Dosher and Lu (2009): unsupervised Hebbian learning with possible input from feedback allows for learning without feedback but also accounts for the positive role of feedback in some circumstances. The model provides an explanation and suggests conditions under which feedback may even be "necessary" (at least over the time scale of a perceptual learning experiment). Where the initial response accuracy is targeted to be low, such as in near-threshold tasks (e.g., Herzog & Fahle, 1997), unsupervised Hebbian learning may take a much longer time to discover the weak statistical correlations and may never succeed. Accurate feedback should be most beneficial in those cases where initial accuracies are low, or where accuracy is held low in adaptive staircase methods. Similarly, the Hebbian framework explains why introduction of feedback late in training may fail to improve performance further (Herzog & Fahle, 1997; McKee & Westheimer, 1978). If an optimal or near-optimal weight vector, corresponding to (near) asymptotic performance, has been found by the model before external feedback becomes available, then introduction of feedback at that point should be unimportant. The augmented Hebbian model may also explain why block-level feedback may have some positive effects on learning by modulating the weights of the criterion control (bias) unit based on block feedback—more stringent criterion (bias) control is used when the feedback is more positive (Dosher & Lu, 2009; Liu, Lu, & Dosher, 2010).

Recently, Shibata et al. (2009) found that fake block feedback on every 40 trials indicating a larger performance improvement facilitated learning compared with genuine block feedback in a same/different task using two complex gratings. Modulating the variance of the fake block feedback also affected learning. On the other hand, the fake block feedback generated based on a smaller performance gradient had no effect on learning compared with the genuine feedback. Because the observers reported that they were not aware of the fake feedback, the authors

concluded that perceptual learning in their experiment was implicitly influenced by the fake block feedback and proposed a computational model that successfully accounted for the empirical results.

The computational model proposed by Shibata et al. (2009) consists of three systems: visual, decision, and learning rate control systems. The visual system is a modified version of the hyper basis function network (Weiss et al., 1993), in which the output unit calculates a weighted sum of the responses of the basis functions. The decision system performs the same/different task for two successive grating presentations. The model learns through reweighting by increasing the weights of the task-relevant basis functions and reducing the weights of the task-irrelevant basis functions. Critically, the rate of learning is determined by the learning rate control system based on two processes: Bayesian estimation of the performance gradient and optimistic bias against worse feedback compared to expectation. Using a Kalman filter, the model estimates performance gradient by considering mechanisms underlying changes in the performance gradient over time (state transition process) and how the performance feedback contains noise (observation process). The learning rate increases linearly as a function of the estimated accuracy gradient. However, if the estimated accuracy gradient is less than a lower bound, the learning rate control system sets the learning rate at the corresponding lower bound. This optimistic "self-serving bias" weakens the effect of the smaller gradient fake feedback on the learning rate of the visual system.

Both the computational model proposed by Shibata et al. (2009) and the AHRM originally proposed by Petrov et al. (2005) are reweighting models of perceptual learning that consist of a sensory representation unit, a weight structure, and a decision unit. Both models learn by adjusting the weights between the representation unit and the decision unit. The two models are however different in terms of the learning algorithm. In the AHRM, an adaptive criterion control unit is used to monitor observers' performance history and adjust the bias in the decision process. Trial-by-trial feedback is included as an additional input in Hebbian learning (leading to the term "augmented Hebbian"); block feedback is used to modulate the weight of the criterion control (bias) unit. In the Shibata et al. (2009) model, an exposure-dependent learning rule is used, and the learning rate is determined by Bayesian estimation of the performance gradient and optimistic bias; block feedback is used in estimating the performance gradient.

One natural question is whether the AHRM can explain the empirical data in Shibata et al. (2009), and whether the Shibata et al. (2009) model can account for the range of data, including the data in this paper, that has been successfully modeled by the AHRM. Liu et al. (2010) simulated the Shibata et al. (2009) experiment in a computational study based on the AHRM. Instead of explicitly changing the learning rate as that in Shibata et al. (2009), block feedback was used to modify the weight of the criterion control unit in the AHRM. The simulation results are virtually the same as those generated by Shibata et al.'s computation model (Liu, Lu, & Dosher, in preparation). On the other hand, it is not clear how to adapt the Shibata et al. (2009) model to simulate the experiments in the current study. Critically, we measured contrast thresholds at fixed performance accuracy levels in all experimental conditions. Because the Shibata et al. (2009) model uses performance accuracy to estimate the performance gradient and determine learning rate, a direct application of the Shibata et al. (2009) model to our experiment would result in a zero performance gradient in all four training accuracy × feedback conditions in the current study and the same learning rate, inconsistent with the observed training accuracy × feedback interaction in this study. It is also not entirely clear to us how to modify the Shibata et al. (2009) model to compute a performance gradient and learning rate from contrast thresholds rather than performance accuracy.

In summary, the observed pattern of interactions between feedback and training accuracy provides further support for the augmented Hebbian reweighting model of perceptual learning. The mode of perceptual learning is neither pure supervised, nor pure unsupervised, but rather, augmented Hebbian learning that makes use of feedback when it is available.

Acknowledgments

This research was supported by NEI. There are no competing interests on the research.

Commercial relationships: none.
Corresponding author: Zhong-Lin Lu.
Email: zhonglin@usc.edu.
Address: Laboratory of Brain Processes (LOBES), Dana and David Dornsife Cognitive Neuroscience Imaging Center, Department of Psychology and Neuroscience Graduate Program, University of Southern California, 3620 McClintock Avenue, Los Angeles, CA 90089, USA.

Footnote

[1]Lu, Jeon, and Dosher (2004) measured the window of temporal integration in an external noise study and concluded that the full width of the perceptual temporal window at half-maximum height is 120 ms. Therefore, the three signal and external noise images, each lasting 33 ms, were merged into a single image by the perceptual system through temporal integration.

References

Ahissar, M., & Hochstein, S. (1997). Task difficulty and the specificity of perceptual learning. *Nature, 387,* 401–406. [PubMed]

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

Barlow, H. B. (1956). Retinal noise and absolute threshold. *Journal of the Optical Society of America, 46,* 634–639. [PubMed]

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision, 10,* 433–436. [PubMed]

Crist, R. B., Kapadia, M. K., Westheimer, G., & Gilbert, C. D. (1997). Perceptual learning of spatial location: Specificity for orientation, position, and context. *The Journal of Physiology, 78,* 2889–2894. [PubMed]

Dosher, B. A., & Lu, Z.-L. (1998). Perceptual learning reflects external noise filtering and internal noise reduction through channel reweighting. *Proceedings of the National Academy of Sciences of the United States of America, 95,* 13988–13993. [PubMed]

Dosher, B. A., & Lu, Z.-L. (1999). Mechanisms of perceptual learning. *Vision Research, 39,* 3197–3221. [PubMed]

Dosher, B. A., & Lu, Z.-L. (2009). Hebbian reweighting on stable representations in perceptual learning. *Learning & Perception, 1,* 37–58. [PubMed]

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

Fahle, M., & Poggio, T. (2002). *Perceptual learning*. Cambridge, MA: MIT Press.

Ghose, G. M. (2004). Learning in mammalian sensory cortex. *Current Opinion in Neurobiology, 14,* 513–518. [PubMed]

Ghose, G. M., Yang, T., & Maunsell, J. H. R. (2002). Physiological correlates of perceptual learning in monkey V1 and V2. *Journal of Neurophysiology, 87,* 1867–1888. [PubMed]

Gilbert, C. D., Sigman, M., & Crist, R. E. (2001). The neural basis of perceptual learning. *Neuron, 31,* 681–697. [PubMed]

Heeger, D. J. (1992). Normalization of cell responses in cat striate cortex. *Visual Neuroscience, 9,* 17. [PubMed]

Hertz, J., Krogh, A., & Palmer, R. G. (1991). *Introduction to the theory of neural computation*. Reading, MA: Addison-Wesley.

Herzog, M. H., & Fahle, M. (1997). The role of feedback in learning a vernier discrimination task. *Vision Research, 37,* 2133–2141. [PubMed]

Herzog, M. H., & Fahle, M. (1998). Modeling perceptual learning: Difficulties and how they can be overcome. *Biological Cybernetics, 78,* 107–117. [PubMed]

Hua, T., Bao, P., Huang, C. B., Wang, Z., Xu, J., Zhou, Y., et al. (2010). Perceptual learning improves contrast sensitivity of V1 neurons in cats. *Current Biology, 20,* 887–894. [PubMed]

Huang, C. B., Lu, Z.-L., & Zhou, Y. F. (2008). Broad bandwidth of perceptual learning in the visual system of adults with anisometropic amblyopia. *Proceedings of the National Academy of Science of United States of America, 105,* 4068–4073. [PubMed]

Jacobs, R. A. (2009). Adaptive precision pooling of model neuron activities predicts the efficiency of human visual learning. *Journal of Vision, 9*(4):22, 1–15, http://www.journalofvision.org/content/9/4/22, doi:10.1167/9.4.22. [PubMed] [Article]

Karni, A., & Sagi, D. (1991). Where practice makes perfect in texture discrimination: Evidence for primary visual cortex plasticity. *Proceedings of the National Academy of Sciences of the United States of America, 88,* 4966–4970. [PubMed]

Kesten, H. (1958). Accelerated stochastic approximation. *Annals of Mathematical Statistics, 29,* 41–59.

Law, C.-T., & Gold, J. (2008). Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. *Nature Neuroscience, 11,* 505–513. [PubMed]

Levi, D. M., & Li, R. W. (2009). Perceptual learning as a potential treatment for amblyopia: A mini review. *Vision Research, 49,* 2535–3549. [PubMed]

Li, X., Lu, Z.-L., Xu, P., Jin, J., & Zhou, Y. (2003). Generating high gray-level resolution monochrome displays with conventional computer graphics cards and color monitors. *Journal of Neuroscience Methods, 130,* 9–18. [PubMed]

Liu, J., Lu, Z.-L., & Dosher, B. A. (2009). Augmented Hebbian learning accounts for the Eureka effect in perceptual learning [Abstract]. *Journal of Vision, 9*(8):851, 851a, http://www.journalofvision.org/content/9/8/851, doi:10.1167/9.8.851.

Liu, J., Lu, Z.-L., & Dosher, B. A. (2010). Augmented Hebbian learning accounts for the complex pattern of effects of feedback in perceptual learning [Abstract]. *Journal of Vision, 10*(7):1115, 1115a, http://www.journalofvision.org/content/10/7/1115, doi:10.1167/10.7.1115.

Liu, J., Lu, Z.-L., & Dosher, B. A. (in preparation). Augmented Hebbian learning accounts for the complex pattern of effects of feedback in perceptual learning.

Lu, Z.-L., Jeon, S.-T., & Dosher, B. A. (2004). Temporal tuning characteristics of the perceptual template and

endogenous cuing of spatial attention. *Vision Research, 44,* 1333–1350. [PubMed]

Lu, Z.-L., Liu, J., & Dosher, B. A. (2010). Modeling mechanisms of perceptual learning with augmented Hebbian re-weighting. *Vision Research, 50,* 375–390. [PubMed]

McKee, S. P., & Westheimer, G. (1978). Improvement in vernier acuity with practice. *Perception & Psychophysics, 24,* 258–262. [PubMed]

Mollon, J. D., & Danilova, M. V. (1996). Three remarks on perceptual learning. *Spatial Vision, 10,* 51–58. [PubMed]

O'Reilly, R. C., & Munakata, Y. (2000). *Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain*. Cambridge, MA: MIT Press.

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision, 10,* 437–442. [PubMed]

Pelli, D. G., & Zhang, L. (1991). Accurate control of contrast on microcomputer displays. *Vision Research, 31,* 1337–1350. [PubMed]

Petrov, A., Dosher, B. A., & Lu, Z.-L. (2005). Perceptual learning through incremental channel reweighting. *Psychological Review, 112,* 715–743. [PubMed]

Petrov, A. A., Dosher, B. A., & Lu, Z. L. (2006). Perceptual learning without feedback in non-stationary contexts: Data and model. *Vision Research, 46,* 3177–3197. [PubMed]

Poggio, T., Fahle, M., & Edelman, S. (1992). Fast perceptual learning in visual hyperacuity. *Science, 256,* 1018–1021. [PubMed]

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

Polat, U., Sagi, D., & Norcia, A. M. (1997). Abnormal long-range spatial interactions in amblyopia. *Vision Research, 37,* 737–744. [PubMed]

Robbins, H., & Monro, S. (1951). A stochastic approximation method. *Annals of Mathematical Statistics, 22,* 400–407.

Rubin, N., Nakayama, K., & Shapley, R. (1997). Abrupt learning and retinal size specificity in illusory-contour perception. *Current Biology, 7,* 461–467. [PubMed]

Seitz, A. R., Nanez, J. E., Holloway, S., Tsushima, Y., & Watanabe, T. (2006). Two cases requiring external reinforcement in perceptual learning. *Journal of Vision, 6*(9):9, 966–973, http://www.journalofvision.org/content/6/9/9, doi:10.1167/6.9.9. [PubMed] [Article]

Seitz, A. R., & Watanabe, T. (2003). Is subliminal learning really passive? *Nature, 422,* 36. [PubMed]

Seitz, A. R., & Watanabe, T. (2005). A unified model for perceptual learning. *Trends in Cognitive Science, 9,* 329–334. [PubMed]

Shibata, K., Yamagishi, N., Ishii, S., & Kawato M. (2009). Boosting perceptual learning by fake feedback. *Vision Research, 49,* 2574–2585. [PubMed]

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

Treutwein, B. (1995). Adaptive psychophysical procedures. *Vision Research, 35,* 2503–2522. [PubMed]

Tsodyks, M., & Gilbert, C. (2004). Neural networks and perceptual learning. *Nature, 431,* 775–781. [PubMed]

Vaina, L. M., Sundareswaran, V., & Harris, J. G. (1995). Learning to ignore: Psychophysics and computational modeling of fast learning of direction in noisy motion stimuli. *Cognitive Brain Research, 2,* 155–163. [PubMed]

Vallabha, G. K., & McClelland, J. L. (2007). Success and failure of new speech category learning in adulthood: Consequences of learned Hebbian attractors in topographic maps. *Cognitive, Affective & Behavioral Neuroscience, 7,* 53–73. [PubMed]

Watanabe, T., Nanez, J. E., & Sasaki, Y. (2001). Perceptual learning without perception. *Nature, 413,* 844–848. [PubMed]

Watanabe, T., Nanez, J. E. S., Koyama, S., Mukai, I., Liederman, J., & Sasaki, Y. (2002). Greater plasticity in lower-level than higher-level visual motion processing in a passive perceptual learning task. *Nature Neuroscience, 5,* 1003–1009. [PubMed]

Watson, A. B., & Pelli, D. G. (1983). QUEST: A Bayesian adaptive psychometric method. *Perception & Psychophysics, 33,* 113–120. [PubMed]

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

Yang, T., & Maunsell, J. H. R. (2004). The effect of perceptual learning on neuronal responses in monkey visual area V4. *Journal of Neuroscience, 24,* 1617–1626. [PubMed]