

Perceptual learning without feedback in non-stationary contexts: Data and model

Alexander A. Petrov^{a,*,1}, Barbara Anne Doshier^{a,*}, Zhong-Lin Lu^b

^a Department of Cognitive Sciences, University of California, Irvine, CA 92697, USA

^b Department of Psychology, University of Southern California, Los Angeles, CA 90089, USA

Received 26 July 2005; received in revised form 13 March 2006

Abstract

The role of feedback in perceptual learning is probed in an orientation discrimination experiment under destabilizing non-stationary conditions, and explored in a neural-network model. Experimentally, perceptual learning was examined with periodic alteration of a strong external noise context. The speed of learning, the performance loss at each change in external noise context (switch cost), and the asymptotic accuracy d' without feedback were very similar or identical to those with feedback. However, lack of feedback led to higher decision bias (error responses matching the external noise context). In the model, the stimulus representations are constant, whereas the read-out connections to a decision unit learn by a Hebbian plasticity rule that may be augmented by additional feedback input and criterion control of decision bias.

© 2006 Elsevier Ltd. All rights reserved.

Keywords: Perceptual learning; Neural networks; Orientation discrimination; Plasticity; Hebbian learning; Feedback

1. Introduction

The ability to detect and discriminate simple visual stimuli improves with practice. This perceptual learning is often slow but its effects can last for years without further training. The precision of stimulus control and the detailed knowledge of the physiology of the early sensory areas make perceptual learning a particularly fruitful domain for studying cortical plasticity in general (see Fahle & Poggio, 2002; Gilbert, Sigman, & Crist, 2001; Tsodyks & Gilbert, 2004, for reviews). Slow, long-lasting improvement at the behavioral level suggests incremental, long-lasting

synaptic plasticity at the neural level. The principles governing this plasticity are of great theoretical interest. Neural-network models provide one approach to understanding the computational principles of perceptual learning. Supervised, error-correcting mechanisms imply a reliance on feedback, while unsupervised Hebbian mechanisms can be relatively insensitive to explicit task feedback. The role of explicit feedback thus may reveal the nature of learning in this domain. Previously (Petrov, Doshier, & Lu, 2005), we developed a Hebbian learning model augmented with a feedback unit (equivalent to supervised Hebbian learning when feedback is available) and a criterion control unit to account for a complex and parametrically varied pattern of perceptual learning in alternating external noise contexts. The focus of the current paper is to explicitly evaluate the importance of feedback by testing observers' ability to learn the same challenging non-stationary task without feedback. The alternating external noise contexts force the system to adjust repeatedly to new sets of stimulus statistics, which is especially difficult in the absence of feedback. Despite these challenges, robust perceptual learning

* Corresponding authors. Present address: Department of Psychology, Ohio State University, 225 Psychology Bldg, 1835 Neil Ave, Columbus, OH 43210, USA (A.A. Petrov), 3151 SSPA, University of California, Irvine, CA 92697, USA (B.A. Doshier).

E-mail addresses: apetrov@alexpetrov.com (A.A. Petrov), bdoshier@uci.edu (B.A. Doshier).

¹ All data and software are available on-line at <http://www.socsci.uci.edu/maplab/> or <http://alexpetrov.com/proj/plearn/>.

occurs even without feedback. The augmented Hebbian model, based on selective reweighting, provides a strong framework for understanding the role of feedback in perceptual learning.

1.1. Learning architectures and the role of feedback

A fundamental theoretical distinction exists between supervised and unsupervised learning in neural networks (e.g., Dayan & Abbott, 2001; Hertz, Krogh, & Palmer, 1991; O'Reilly & Munakata, 2000). Supervised learning requires external feedback and updates the connection weights according to plasticity rules that minimize the average error between actual and target activations of the output units. Such error-correcting rules can train the network to perform arbitrary stimulus–response mappings. 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 post-synaptic) units. These rules do not depend on feedback and detect statistical regularities in the training corpus. Hybrid systems are also possible. For example, internal target signals may replace the external teacher in error-correcting learning, or external feedback may be included as additional input in Hebbian learning. This article proposes a hybrid of the latter kind, which we call “augmented Hebbian learning” (see Petrov et al., 2005). The behavioral effects of feedback provide valuable insights into the underlying plasticity rule in perceptual learning. Pure error-correcting systems fail to learn when feedback is not available, whereas pure Hebbian systems fail to benefit from it when it is available (Herzog & Fahle, 1997). Hybrid systems avoid these pitfalls—they do not depend on external feedback but can benefit from it. The behavioral effects of feedback in such systems can be large or small, and a detailed model is necessary to understand how feedback interacts with other environmental variables. The current experiment probes the role of feedback in a challenging non-stationary environment.

1.2. The role of feedback in perceptual learning

Empirically, the role of feedback in perceptual learning reveals a somewhat complex pattern. There are numerous reports of successful learning with no external feedback (Ball & Sekuler, 1987; Crist, Kapadia, Westheimer, & Gilbert, 1997; Fahle & Edelman, 1993; Karni & Sagi, 1991; McKee & Westheimer, 1978; Shiu & Pashler, 1992). Sometimes, however, feedback seems necessary for learning (Shiu & Pashler, 1992, Experiment 1), or learning without feedback succeeds when the stimuli are easy but fails when the stimuli are hard (Shiu & Pashler, 1992, Experiment 3), or succeeds for some observers but not for others (Herzog & Fahle, 1997). Introducing feedback after no-feedback learning had reached asymptote apparently triggers no further improvement (Herzog & Fahle, 1997; McKee & Westheimer, 1978). In aggregate, these data indi-

cate that perceptual learning can occur without feedback, but that feedback can be important for learning under some circumstances.

When there is successful learning without feedback, is it slower than learning with trial-by-trial feedback? The available evidence is inconclusive. One experiment found no significant difference between the slopes of learning curves with and without feedback, whereas a marginal speedup with feedback was found in another (Fahle & Edelman, 1993). In a motion-direction discrimination experiment (Ball & Sekuler, 1987), feedback did not affect the speed of learning or the asymptotic d' for stimuli near the cardinal directions, but did speed up learning near the oblique directions. These observations suggest that feedback may increase the efficiency of learning, especially in difficult conditions.

Herzog and Fahle (1997, p. 109) considered the theoretical implications of the mixed effects of feedback and concluded that they are “not compatible with classical neural-network models”. In their view, both supervised and unsupervised learning schemes are implausible—the former because perceptual learning can occur without feedback, and the latter because feedback can make a difference. Herzog and Fahle (1998) outlined a model in which external feedback, when available, changes the learning rate but does not act as a teaching signal. Here, we propose instead a Hebbian framework with top-down feedback input that may operate as a supervised Hebbian system. This alternative is consistent with Herzog and Fahle’s (1998) overall philosophy, but relies entirely on biologically plausible plasticity mechanisms.

1.3. Non-stationary environments: Testing a perceptual learning model

Petrov et al. (2005) developed a neural-network model to account for perceptual learning in a complex non-stationary environment based on a noisy signal-detection model of the observer (e.g., Lu & Doshier, 1999). Observers identified the orientation (top tilted left or right) of a Gabor target embedded in a larger context of orientation-filtered noisy texture tilted either left or right. This paradigm provides significant challenges for perceptual learning. The target and noise have overlapping spectra, and the task uses the same representations in both noise contexts. Because switches in context are destabilizing, postswitch recovery may be critically dependent on external feedback.

The proposed model learns which features of the early representations are most diagnostic for the task, and strengthens the read-out connections from the units (or “channels”) encoding those features. The representations themselves never change. The channels are modeled on the orientation and frequency tuned units of early visual analysis (see Section 2.1 for details). Although feedback was available in the original study, an error-correcting model was rejected as too powerful in favor of a hybrid model with a Hebbian learning that may operate in a

supervised mode if feedback is available. The Hebbian model (also referred to as “multi-channel reweighting model,” Petrov et al., 2005) takes images as inputs, constructs representations consistent with many principles of the early visual system, and learns by updating the connections from these representations to a decision unit.

The Hebbian learning algorithm incorporates external feedback, when present, simply as another input to the decision unit. This top-down input shifts the postsynaptic activation in the correct direction, which in turn fosters appropriate weight changes. Exactly the same Hebbian learning rule applies in all situations. Without feedback, the weights still move in the correct direction on average because the activation of the decision unit correlates with the correct stimulus classification.

The success of pure Hebbian learning without feedback depends on the initial state of the observer and on task difficulty. When the initial accuracy is sufficiently above chance for a sufficient fraction of the stimuli, Hebbian learning amplifies this positive correlation. When the initial accuracy is low, however, Hebbian learning can be erratic, slow, or even fail altogether. Hebbian learning is most successful when some easy stimuli are included, when the choice is binary, when the decision space admits a linear boundary solution, and when the category structure coincides with the first principle component of the stimulus distribution (cf. Hertz et al., 1991).

In this article, we test the augmented Hebbian learning hypothesis in alternating external noise contexts without feedback. This non-stationary paradigm poses special challenges for statistical learning. Our earlier experiment (Petrov et al., 2005) provides a with-feedback comparison for the current no-feedback test. The previous study showed substantial perceptual learning with feedback (d' improved more than twofold), recurring switch costs (d' dropped temporarily whenever the context changed), and systematic effects of stimulus contrast—all excellently fit with the multi-channel reweighting model.

Will perceptual learning be possible in this non-stationary environment without feedback? Can the system recover from the decrements in performance and biased responses that occur at each context switch? Does feedback improve the asymptotic accuracy level or the rate of learning? Model simulations suggest that: (i) learning will occur without feedback, (ii) the simulated curves will exhibit switch costs, and (iii) feedback should improve performance to some degree. These are strong predictions. It is by no means guaranteed that learning will be successful under context switching without feedback, as a pure Hebbian system may get stuck in its initial response policy and fail to adjust when the environment switches abruptly. We ran a direct experiment to investigate these questions.

2. Experiment

The present experiment is a no-feedback complement to the study of Petrov et al. (2005). Its primary objective is to

test the Hebbian learning principles in a challenging non-stationary environment without feedback. The same task is performed on the same target images in two external noise contexts: orientation discrimination of Gabor patches either -10° (“left”) or $+10^\circ$ from vertical (“right”). The same sensory units, response selection units, and connections are plausibly engaged throughout the experiment (see Petrov et al., 2005, for discussion). The learning dynamics was tracked across numerous switches in the predominant orientation of the filtered-noise context surrounding the Gabor. Three separate target contrast levels were tested to relate performance directly to a parametrically varied physical measure, which constrains estimates of the non-linearity in the perceptual system (Doshier & Lu, 1999; Lu & Doshier, 1999). This contrast manipulation imposes additional strong constraints on the model.

2.1. Method

2.1.1. Stimuli

Each stimulus consists of a Gabor patch $G(x, y)$ embedded in a larger field of filtered visual noise $N(x, y)$ (Eqs. (1) and (2)). The observers are instructed to ignore the background and indicate the orientation θ ($-10^\circ =$ left or $+10^\circ =$ right of vertical) of the Gabor target (Fig. 1). The spectral power of the external noise is concentrated around a predominant orientation $\phi = \pm 15^\circ$, also left or right. In a congruent stimulus, the Gabor and noise texture orientations have the same sign (Fig. 1, bottom left); in an incongruent stimulus they have opposite signs (Fig. 1, top left).

The luminance $L(x, y)$ of each pixel is an additive mixture of a Gabor term $G(x, y)$ and noise $N(x, y)$, where L_0 is the mid-gray value of the monitor

$$L(x, y) = [1 + c_p G(x, y) + c_n N(x, y)] L_0, \quad (1)$$

$$G(x, y) = e^{-(x^2+y^2)/2\sigma^2} \sin [2\pi f(x \cos \theta + y \sin \theta)]. \quad (2)$$

The peak target contrast c_p was set to 0.245, 0.160, or 0.106. The sine-phase Gabor patches have spatial frequency $f = 2$ cyc/deg, and the standard deviation of their Gaussian envelope is $\sigma = 0.4^\circ$. The background $N(x, y)$ is an isotropic field of Gaussian noise, filtered in the Fourier domain with the conical filter defined by Eq. (3)

$$H_\phi(f_x, f_y) = \left[1 + \frac{(f_x \cos \phi + f_y \sin \phi)^2}{\beta^2 (f_x \sin \phi - f_y \cos \phi)^2} \right]^{-1}. \quad (3)$$

Its cross-section at any spatial frequency is a Butterworth bandpass filter of order 1 (Gonzalez & Woods, 1992) with half-amplitude half-bandwidth $\beta = 0.20 \approx \tan 11.3^\circ$. The filter H_ϕ attenuates the spectral power of orientations away from the peak orientation ϕ (Fig. 1, right). To generate a background texture, the algorithm generates a 64×64 sample of iid Gaussian noise and applies a filter tuned for $\phi = -15^\circ$ in context L and $\phi = +15^\circ$ in context R . The noise of each trial was a new sample. The resulting matrix $N(x, y)$

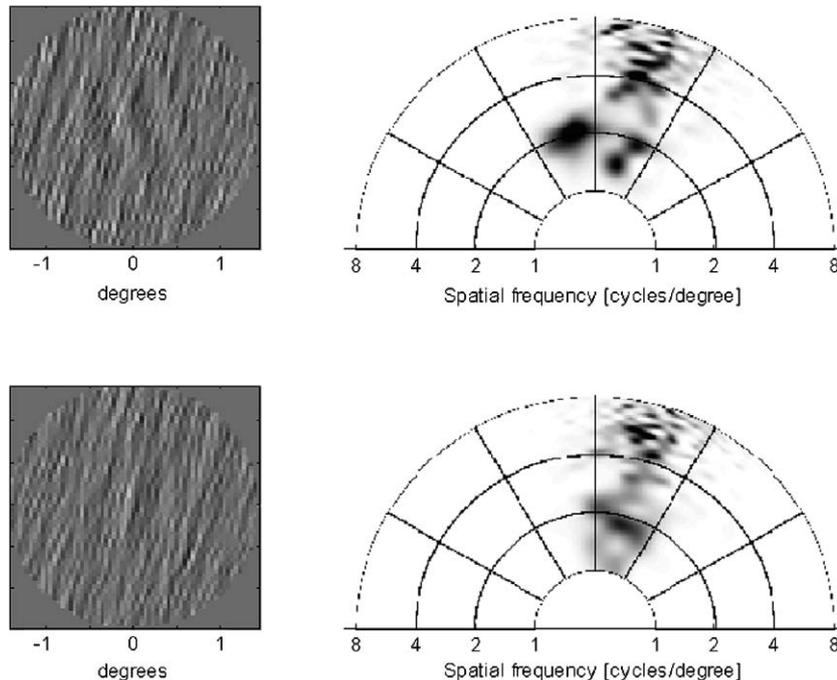


Fig. 1. Examples of a congruent (bottom left) and an incongruent stimulus (top left) and their corresponding power spectra (right, bottom and top). Gabor targets oriented $\pm 10^\circ$ from vertical are embedded in textures of noise filtered in a cone of orientations. The Gabor spatial frequency is 2 cyc/deg. The spectra are plotted in polar coordinates, with log frequency along the radial axis and orientation along the angular axis.

is then rescaled linearly to zero mean (mid-gray) and peak contrast $c_n = 0.667$. The standard deviation of the term $c_n N(x, y)$ in Eq. (1) varies between 0.17 and 0.19 for most noise patches. The Gabor $G(x, y)$ is added to the noise, and the image is quantized to 256 grayscale levels and clipped within a circular window with radius 32 pixels ($\approx 1.44^\circ$ of visual angle).

Fourier analysis verifies that all stimuli occupy the same general region in the frequency domain and thus should activate a common population of orientation-selective neurons. Fig. 1 plots the power spectrum of a congruent (top right) and an incongruent stimulus (bottom right) in context R . Note the peaks at the spatial frequency of the Gabor targets (2 cyc/deg) and near the predominant background orientation for context R ($\phi = +15^\circ$) at all frequencies. The spectra in context L are mirror images.

2.1.2. Apparatus

All stimuli were generated in MATLAB (The MathWorks, 1999) in real time and presented using the Psychophysics Toolbox extensions (Brainard, 1997). They were displayed on a NANA Technology FlexScan 6600 monochrome monitor with P4 phosphor and a refresh rate of 120 frames/s driven by the internal video card of a Power Macintosh 7300. A special circuit combined two 8-bit output channels of the video card to produce 6144 distinct gray levels (Pelli & Zhang, 1991). Luminance calibration was performed both with psychophysical matching judgments (Lu & Sperling, 1999) and by measurement with a Tektronix Lumacolor J17 photometer. A linear lookup table divided the dynamic range of the monitor (from 1

to 22 cd/m^2) into 256 evenly spaced levels, with the background $L_0 = 11 \text{ cd/m}^2$. Displays were viewed binocularly in a dim room with the natural pupil at a viewing distance of $\approx 72 \text{ cm}$. The same apparatus was used in Petrov et al. (2005).

2.1.3. Observers

Eighteen paid volunteers participated in the study. All had normal or corrected-to-normal vision. None of them had been involved in the earlier experiment with feedback.

2.1.4. Design

Two groups of observers differed with respect to the context they first trained on: nine observers began in context L and nine in context R . There were nine sessions on separate days, with four blocks per day and a total of 36 blocks. The presentation schedule (in blocks) was $L-8R-8L-8R-8L-3R$ for group 1 and $R-8L-8R-8L-8R-3L$ for group 2. Mid-session context switches avoided potential confounds from overnight consolidation or forgetting. Each block had 300 trials: two Gabor orientations by three Gabor contrasts by two retinal locations by 25 replications of each stimulus type. The trial sequence was randomized within each block.

2.1.5. Procedure

Each session began with a few demonstration trials (with auditory feedback, ten trials on the first day, two trials on subsequent days) followed by four blocks of 300 experimental trials without feedback. The observers indicated the orientation of the Gabor target by pressing a

key with their left or right hand. The few (less than 0.4%) trials with invalid responses were repeated at the end of the block. Optional rest periods were allowed between blocks. Each trial began with a brief beep and a fixation cross in the middle of the screen. The stimulus appeared for 75 ms at one of two equiprobable locations centered 5° above or below fixation. Then the screen was cleared and the observer's response recorded. The next trial began after a 750 ms inter-trial interval.

2.1.6. Dependent variables

Probability correct is tabulated separately for congruent and incongruent trials at each contrast level in each block (216 data points per observer, each with 50 responses counterbalanced across the two retinal locations). Discrimina-

bility measures (d' 's) are computed by adding the z values for congruent and incongruent stimuli, where $z = \Phi^{-1}(p)$ and extreme frequencies ($p = 50/50$, 160 cases in all) are transformed to $z = 2.33$.

2.2. Results and discussion

The data from the two variants of the presentation schedule are statistically indistinguishable ($F(1, 16) < 1$, n.s.). The two observer groups were combined, and labeled as A-8B-8A-8B-8A-3B.

2.2.1. Robust learning without feedback

Fig. 2A plots the d' learning curves at three target contrast levels, averaged across observers. The corresponding

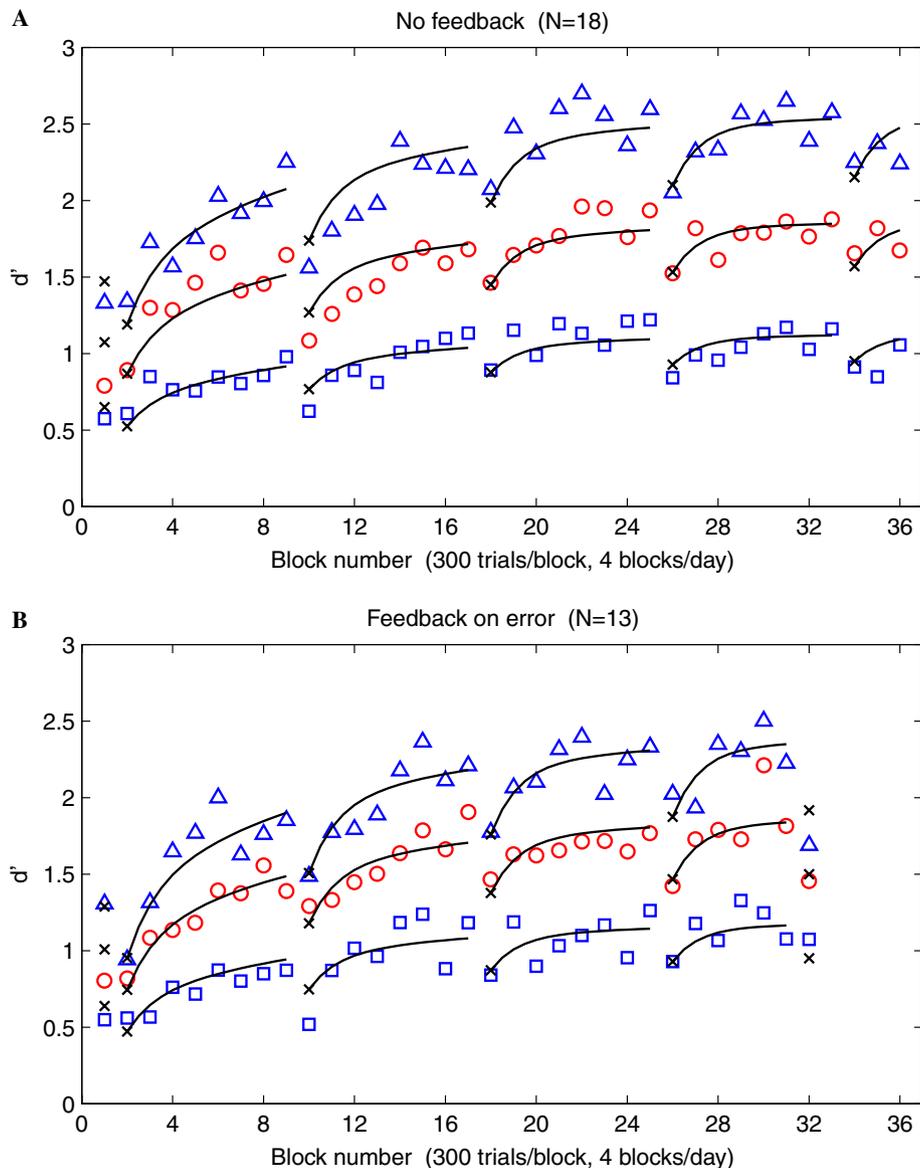


Fig. 2. (A) d' learning curves without feedback, averaged across 18 observers. Target contrast 0.245 (triangles), 0.160 (circles), and 0.106 (squares). Fits of Eq. (4) in the text, $CI_{95} = \pm 0.241$. The connected lines belong to the same context; the discontinuities mark context switches. Note the recurring switch costs. (B) Analogous data with feedback on error trials (Petrov et al., 2005, $N = 13$, $CI_{95} = \pm 0.235$, reproduced with permission).

feedback data from Petrov et al. (2005) are reprinted for comparison (B). There is clear perceptual learning even in the absence of feedback—the d' s increase approximately twofold over the 36 blocks of training. The higher the target contrast c_p , the higher the discriminability.

The connected segments in Fig. 2 mark successive blocks of one external noise context; the discontinuities at blocks 2, 10, 18, 26, and 34 mark changes in external noise context. Switch costs were superimposed on the gradually rising learning curve at each contrast level. The d' s decreased about 15% after each context change, and recovered during the stationary epochs. These recurring switch costs indicate partial stimulus specificity of perceptual learning.

Learning without feedback could have failed initially, or faltered following the disruptive change in external noise context. Instead, the data in Fig. 2 clearly show that perceptual learning does not require external feedback even in non-stationary environments, in excellent agreement with the predictions of the (multi-channel) augmented Hebbian reweighting model. The switch costs are superimposed on the overall upward trend. The same qualitative pattern appears in the feedback data as well (Fig. 2B).

Non-linear regression analyses based on Eq. (4) (see Petrov et al., 2005) decompose the learning dynamics into a context-general and a context-specific component. The general component $(1 - ge^{-T/\tau})$ improves with time (practice) T since the beginning of the experiment. The parameter g measures the overall learning effect—the relative distance between the initial and the asymptotic levels D_i for each target contrast i . The context-specific component $-se^{-t_s/\tau_s}$ accounts for the transient switch costs of magnitude s superimposed on the general trend, where t_s is reset to 0 after each switch. The time constants τ and τ_s quantify the speed of general and specific learning, respectively.

$$d'_i(T, t_s) = D_i(1 - ge^{-T/\tau} - se^{-t_s/\tau_s}). \quad (4)$$

Eq. (4) accounts for 95.5% of the variance of the 108 observations in Fig. 2A with free parameters (rmse = .121), comparable to the fit to the feedback data ($R^2 = .945$, rmse = .118, Petrov et al., 2005). The estimated parameters for the no-feedback condition are shown in Table 1, along with those for the previous feedback data.

2.2.2. Learning dynamics

The learning curves unfold in parallel across all difficulty levels over blocks of practice. Only the asymptotes D_i vary with target contrast i ; everything else scales in proportion (Eq. (4)). The parameters g , s , τ and τ_s are equivalent for all contrasts ($R^2 = .9555$, $df = 15$, $F(6, 92) < 1$, n.s.). The time constant of the general improvement is on the order of days ($\tau \approx 10$ blocks = 2.5 days), while the recovery from switch costs is relatively rapid ($\tau_s \approx 1.2 \ll 10$). The switch cost magnitude does not diminish with successive switches. The seesaw pattern in Fig. 2 indicates that training in one context interferes with earlier training in the other. This interference reflects the fact that our stimuli induce highly overlapping representations in both contexts. A single set of connections is thus forced to cope with successive switches in the noise background. The system fine-tunes its weights for maximal accuracy given the statistics of the prevailing stimuli. It lags behind with suboptimal weights and its accuracy decreases after a switch. The constant cost across successive switches reflects the constant statistical differences between the two contexts.

2.2.3. Congruence effects

The accuracy profiles (Fig. 3, z -transformed probability correct) on incongruent trials mirror the d' profiles:

Table 1
Summary of the empirical results

Behavioral phenomenon	Feedback	No feedback
The d' improves with practice in all conditions, with no need for external feedback	$g = 0.47$	$g = 0.43$
There is no evidence that perceptual learning is any slower without feedback than with it	$\tau = 10$	$\tau = 10^*$
The learning dynamics seems independent of the target contrast; the absolute d' levels are strongly dependent on it, both with and without feedback	$D_{.245} = 2.41$ $D_{.160} = 1.89$ $D_{.106} = 1.20$	$D_{.245} = 2.58$ $D_{.160} = 1.88$ $D_{.106} = 1.14$
Each switch of the background texture incurs a transient switch cost superimposed on the general learning curve	$s = 0.18$ $\tau_s = 1.2$	$s = 0.15$ $\tau_s = 1.2^*$
The switch cost reappears undiminished after each context switch, both with and without feedback	Fig. 2B	Fig. 2A
Stronger target contrast gives rise to slightly lower identification accuracy for congruent stimuli, both with and without feedback	Replicated	Fig. 3
Responses are biased towards the background orientation, especially when there is no feedback	57% vs. 43%	64% vs. 36%

The statistics pertain to Eq. (4) and are explained in the main text. The asterisks indicate no-feedback parameters that are fixed to their corresponding feedback values. Feedback data from Petrov et al. (2005).

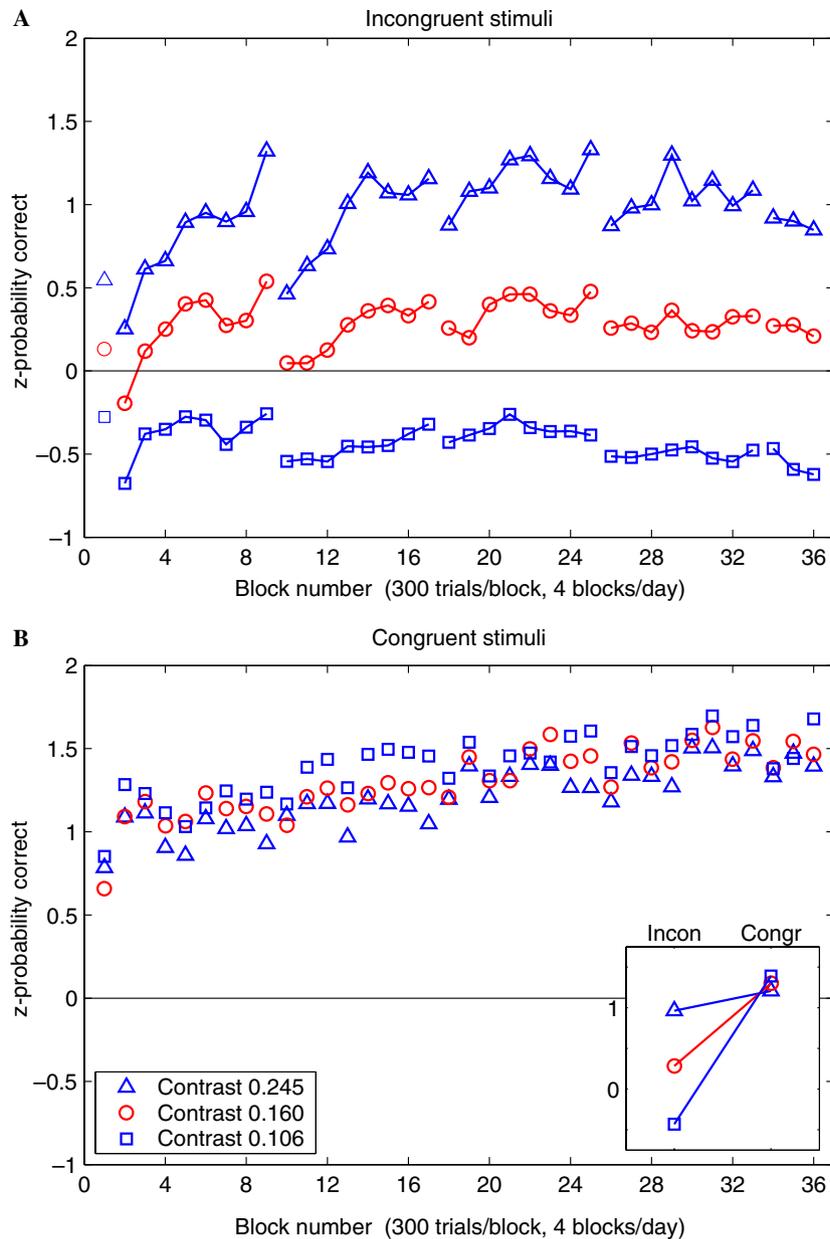


Fig. 3. Accuracy profiles for target oriented against (A) or in the direction (B) of the background noise. Average of 18 observers, no feedback, $CI_{95} = \pm 0.28$. The inset shows the same six curves collapsed in time, $CI_{95} = \pm 0.05$. Note the paradoxical inverse relationship between contrast and accuracy for congruent stimuli.

improvement with practice, switch costs, and strong positive effect of the target contrast. The target contrast effect reverses, however, on congruent trials. When the Gabor orientation matches the primary orientation of the external noise, the accuracy tends to decrease when the target contrast increases! This counterintuitive reversal is small but highly statistically significant ($F(2, 70) = 70$, $p < 10^{-16}$). Table 2 and the inset in Fig. 3 show the strong interaction between target contrast and congruence, averaged over blocks and observers. The group average is representative for all individual observers, and is the same as the pattern in the feedback data (Petrov et al., 2005).

This interaction imposes powerful constraints on models of perceptual learning. The negative relationship between

accuracy and target contrast for congruent trials does not reflect response bias, which would trade off the accuracy of congruent and incongruent stimuli but cannot reverse a contrast effect.

2.2.4. Comparisons of learning with and without feedback

As detailed in the previous sections, the general pattern of learning in the current data set without feedback exhibited all of the major properties of the learning with feedback from Petrov et al. (2005). Subject to individual differences between the two sets of observers, the performance without feedback approaches that with feedback (see parameters in Table 1). Although feedback was accurate and thus potentially useful, the average performance

Table 2
Accuracy (average z -probability across blocks) as a function of target contrast and congruence

Stimulus	Peak target contrast c_p			Total
	0.106	0.160	0.245	
<i>Behavioral data, no-feedback ($CI_{95} = \pm 0.047$ within blocks)</i>				
Incongruent	-0.432	0.284	0.963	0.272
Congruent	1.390	1.300	1.208	1.299
Total (= $d'/2$)	0.479	0.792	1.085	0.786
<i>Behavioral data, feedback on error ($CI_{95} = \pm 0.040$, Petrov et al., 2005)</i>				
Incongruent	-0.070	0.565	1.126	0.541
Congruent	1.031	0.956	0.821	0.936
Total (= $d'/2$)	0.481	0.761	0.974	0.738
<i>Model fits, no feedback</i>				
Incongruent	-0.480	0.188	1.082	0.263
Congruent	1.358	1.306	1.203	1.289
Total (= $d'/2$)	0.439	0.747	1.142	0.776
<i>Model fits, feedback on error</i>				
Incongruent	-0.146	0.481	1.310	0.548
Congruent	0.989	0.948	0.864	0.934
Total (= $d'/2$)	0.421	0.715	1.087	0.741

Feedback data from Petrov et al. (2005).

is actually slightly better, just by chance, without feedback ($d'=1.57$, $SD = .40$) than with it ($d' = 1.48$, $SD = .46$). Learning occurs approximately as quickly without feedback as it does with it ($R^2 = .9545$ with five parameters vs. $R^2 = .9549$ with 7; $F(2, 100) < 1$, n.s.), within the precision² of the time constants. In short, the absence of feedback seemingly had at most small (to no) impact on the key parameters of learning.

Nonetheless, observers were attending to the feedback when it was available because it had a major effect on their overall response bias. The current no-feedback data show a strong bias to respond “right” in context R and “left” in context L —the average probability of a congruent response is almost twice that of an incongruent response (64% vs. 36%). This strong bias was significantly moderated with feedback, reducing the errors in the incongruent condition (57% vs. 43%, in the data of Petrov et al., 2005).

2.3. Summary

In conclusion, with the exception of moderation of response bias, all of the key characteristics of perceptual learning with feedback also occurred without feedback, even in a challenging task that required repeated adjustments and recovery from switch costs. The asymptotic discriminability levels, the speed of learning, and the recurring switch costs were all remarkably similar in the two feedback conditions (Table 1). The correlation of the internal responses with the stimulus category seem sufficient to guide the learning process, in qualitative agreement with the Hebbian learning hypothesis and the predictions of the multi-channel reweighting model (Petrov et al., 2005).

Perceptual learning is statistically driven. The present task engages the same sensory and decision structures in both external noise contexts. The observed persistent switch costs are inconsistent with any model that explains the specificity of perceptual learning solely in terms of non-overlapping stimulus representations or non-overlapping decision structures because all such schemes entail that the switch cost should diminish with repeated switches and eventually disappear entirely. The complex interactions of contrast and congruence, the similarity in overall learning with and without feedback, and the strong feedback effect on response bias present a clear theoretical and modeling challenge.

3. A Hebbian reweighting model of perceptual learning

In this section, we present the model and evaluate its ability to account for the patterns of perceptual learning without feedback, as well as those of the previous experiment with feedback. The multi-channel Hebbian reweighting model is a mechanistically explicit and fully functional neural-network model that instantiates the Hebbian reweighting hypothesis (Petrov et al., 2005). The model builds upon and extends the perceptual template model (PTM), which has been applied to perceptual learning, selective attention, and other domains (Doshier & Lu, 1998, 1999, 2005; Lu, Chu, Doshier, & Lee, 2005; Lu & Doshier, 1998, 1999, 2000). Although inspired by the computational principles of the human brain, the model implementation is deliberately simplified and abstract. It is not intended to replicate the neural pathways of visual processing and response selection. Rather, we have reduced the structure of the model to the bare functional essentials. The objective is to provide an existence proof that the

² Bootstrap standard error is ≈ 4 for τ and ≈ 0.4 for τ_s .

Hebbian reweighting hypothesis is both conceptually coherent and empirically adequate to account for the challenging behavioral results of perceptual learning. Here, we show that the Hebbian learning system augmented with a feedback unit and criterion (bias) control can account for the detailed dynamics of perceptual learning both with and without feedback.

The neural-network model consists of four types of units (Fig. 4). Orientation and frequency tuned representation units encode the input image as a distributed activation pattern, analogous to the function of the early visual system. The orientation discrimination task is carried out by a task-specific decision unit receiving weighted input from the sensory units. An adaptive bias unit accumulates a running average of the response frequencies and provides top-down support for the weaker alternative that helps to balance the frequency of the two responses. Finally, a feedback unit encodes the external feedback when (and if) it is presented, and sends top-down input to the decision unit that supports the correct response.

Learning in this system occurs exclusively through incremental Hebbian modification of the weights between stable representation units and the decision unit—multi-channel reweighting. The reweighting hypothesis (Doshier & Lu, 1998, 1999) is substantively consistent with the available physiological, functional, psychophysical, and computational evidence. 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 feedback (Petrov et al., 2005).

External feedback, when available, acts as just another input to the decision unit. The only effect of feedback is to shift the activation of the decision unit during the late phase of the trial, after the response is emitted but before the weights are updated. The reweighting mechanism operates exactly as in the no-feedback case, but the postsynaptic

activation now reflects a feedback contribution in addition to the sensory and bias contributions.

3.1. Representation subsystem

The representation subsystem encodes the stimulus image as a distributed activation pattern over a population of orientation- and frequency-selective units. The processing principles are consistent with the perceptual template model (Lu & Doshier, 1998) and inspired by the functional organization of the early visual cortex (De Valois & De Valois, 1988; Ferster & Miller, 2000). (See Petrov et al., 2005, for a detailed description.)

In overview, the activation $A(\theta, f)$ of each of the 35 representation units encodes the normalized spectral energy in the corresponding orientation and frequency channel. Fig. 5 summarizes the information flow. Briefly, units tuned to different orientations θ , spatial frequencies f , and spatial phases ϕ compute a set of retinotopic phase-sensitive maps $S(x, y, \theta, f, \phi)$ of the input image $I(x, y)$. Weight matrices defined by two-dimensional Gabor receptive fields $\text{RF}_{\theta, f, \phi}$ split the processing stream into 35 channels spanning seven orientations and five frequencies (Fig. 5B; $\theta \in \{0^\circ, \pm 15^\circ, \pm 30^\circ, \pm 45^\circ\}$, $f \in \{1, 1.4, 2, 2.8, 4\}$ cyc/deg). The exact orientation of our target stimuli ($\pm 10^\circ$) is not explicitly included; representations are distributed. The \otimes symbol in Eq. (5) denotes the convolution operator; $[\cdot]_+^2$ denotes half-squaring rectification (Fig. 5C; Heeger, 1992a). Receptive field sizes correspond to bandwidth parameters representative of the tuning of parafoveal simple cells in macaque striate cortex (half-amplitude full-bandwidth $h_\theta = 30^\circ$ for orientation, $h_f = 1$ octave for spatial frequency; De Valois, Yund, & Hepler, 1982). Sensitivity analyses show that these values, while typical of physiological reports, are not critical to the model predictions (Petrov et al., 2005).

$$S(x, y, \theta, f, \phi) = [\text{RF}_{\theta, f, \phi}(x, y) \otimes I(x, y)]_+^2, \quad (5)$$

$$E(x, y, \theta, f) = \sum_{\phi} S(x, y, \theta, f, \phi), \quad (6)$$

$$C(x, y, \theta, f) = E(x, y, \theta, f) / (s^2 + N(f)). \quad (7)$$

The phase-sensitive maps are then combined into phase-invariant maps $C(x, y, \theta, f)$ by summing across phases ($\phi \in \{0, 90, 180, 270\}$, Eq. (6), Fig. 5E and F). Phase invariance is characteristic of V1 complex cells (De Valois et al., 1982; Movshon, Thompson, & Tolhurst, 1978) and is used in energy models of texture and motion perception (Adelson & Bergen, 1985; Knutsson & Granlund, 1983; Pollen & Ronner, 1981). Non-linear divisive normalization approximates the shunting inhibition in visual cortex (Eq. (7); Heeger, 1992b). In line with physiological and psychophysical evidence, the normalization pool $N(f)$ is assumed to be essentially independent of orientation and modestly tuned for spatial frequency (Cannon & Fullenkamp, 1991; Carandini, Heeger, & Movshon, 1997; Chubb, Sperling, & Solomon, 1989; Graham & Sutter, 2000; Lu &

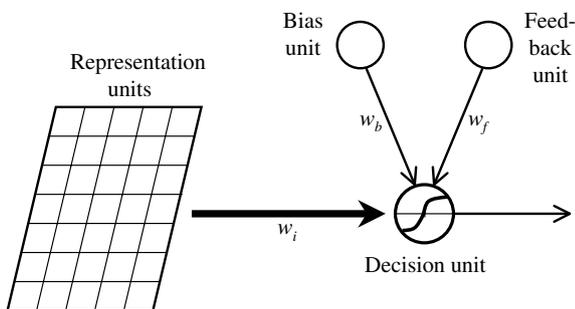


Fig. 4. The multi-channel reweighting model encodes the input image as a distributed activation pattern over 35 representation units. A single decision unit carries out the orientation discrimination task. The model learns through incremental Hebbian reweighting of the connections between representation and decision (bold arrow). A feedback unit and an adaptive bias unit augment the Hebbian framework with modulatory top-down inputs.

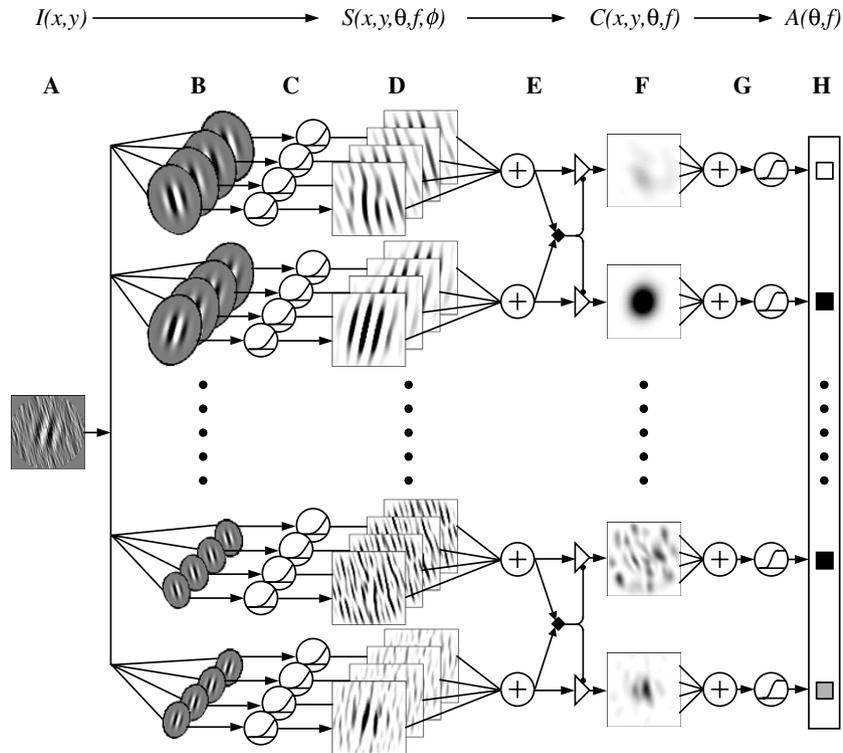


Fig. 5. Schematic outline of the model representation subsystem. Top row: the stimulus $I(x,y)$ is processed by units tuned for different orientations θ , spatial frequencies f , and spatial phases ϕ . The resulting phase-sensitive maps $S(x,y,\theta,f,\phi)$ are processed further into phase-invariant maps $C(x,y,\theta,f)$. Finally, a third population of units $A(\theta,f)$ pools the information across space. The spatial extent of the receptive fields at this stage is commensurate with the diameter of the stimuli ($\approx 2.8^\circ$ of visual angle). Main diagram: (A) stimulus image. (B) Schematic rendition of the weight matrices of the various channels. The receptive fields are drawn to scale and the elliptical aspect ratio is accurate. (C) Rectifying non-linearity (half-squaring). (D) Phase-sensitive maps, maximal activations are depicted in black. (E) Contrast gain control implemented via divisive normalization. This is equivalent to shunting inhibition, depicted by triangles and diamonds. (F) Phase-invariant maps. (G) Spatial pooling followed by a saturating non-linearity. (H) Final representation $A(\theta,f)$. Note: From “The Dynamics of perceptual learning: An incremental reweighting model,” by A. Petrov, B. Doshier and Z.-L. Lu, 2005, *Psychological Review*. Copyright 2005 of the American Psychological Association. Reprinted with permission.

Doshier, 1999; see Petrov et al., 2005, for details). The semi-saturation constant s^2 is relevant only at near-threshold contrasts.

Because spatial phase is not relevant in this task (and to simplify the representations) the model aggregates each retinotopic phase-invariant map into the activation level $A(\theta,f)$ of a single representation unit per channel (Fig. 5H). Eq. (8) pools the energy maps across space,³ weighted by a radially symmetric Gaussian kernel W_r with full-width at half-height $h_r = 2.0^\circ$, commensurate with the size of the target Gabor. The representation units have non-negative activations and limited dynamic range and their responses saturate for high inputs, with gain parameter γ_r (Eq. (9) and Fig. 5G).

$$A'(\theta,f) = \sum_{x,y} W_r(x,y)C(x,y,\theta,f) + \varepsilon_{\theta,f}, \quad (8)$$

$$A(\theta,f) = \begin{cases} \frac{1-e^{-\gamma_r A'}}{1+e^{-\gamma_r A'}} A_{\max} & \text{if } A' \geq 0, \\ 0 & \text{otherwise.} \end{cases} \quad (9)$$

The representation noise $\varepsilon_{\theta,f}$ in Eq. (8) models various inefficiencies in the visual system and any variability that is not related to the current stimulus (Ahumada & Watson, 1985; Burgess, Wagner, Jennings, & Barlow, 1981; Lu & Doshier, 1999). In combination with the external stimulus noise and the internal decision noise in the task-specific subsystem, this representation noise limits the accuracy of the overall performance. In the current experiment, the variability in the representations is dominated by external rather than internal noise.

Fig. 6 illustrates the activity representations of two sample stimuli: one incongruent (A, target $\theta = -10^\circ$, noise $\phi = +15^\circ$) and one congruent (B, $\theta = +10^\circ$, $\phi = +15^\circ$). The representations are distributed, noisy, and overlap significantly for the two stimulus classes. The two Gabor targets have very similar activation patterns. The task-specific subsystem must identify which units have diagnostic value for the particular task at hand.

³ No positional uncertainty is implemented because the high-contrast noise patch identifies the retinal position of each stimulus presentation. Future extensions could simply duplicate the current model and have two independent pools of units with receptive fields above and below the horizontal meridian.

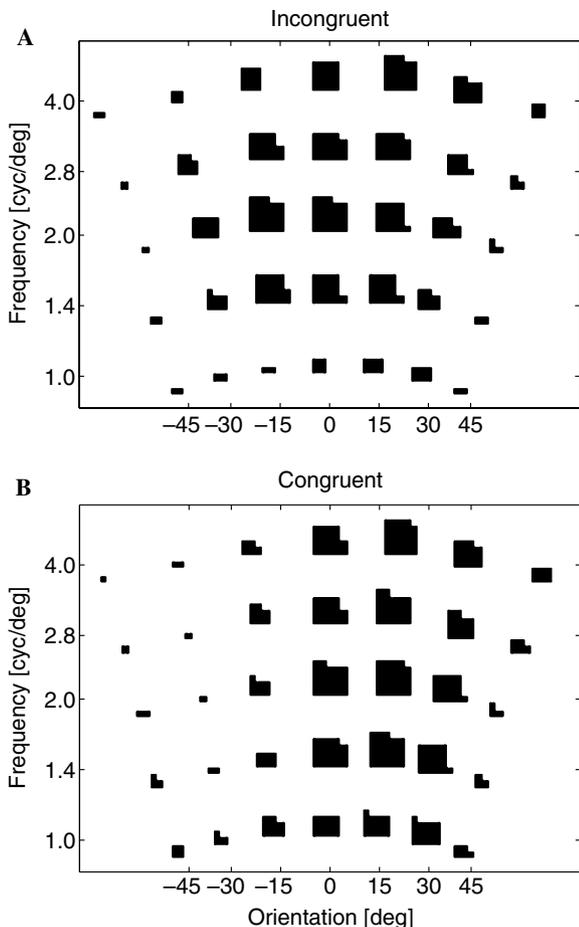


Fig. 6. Internal representations of a congruent (B) and incongruent stimulus (A). There are 35 units tuned for seven orientations (rays) and five spatial frequencies (arcs). Darker blobs depict higher activations. Compare with the power spectra of the same images in Fig. 1.

3.2. Task-specific reweighting subsystem

The stimuli are categorized as having either “left” or “right” orientation by a task-specific reweighting subsystem based on Hebbian learning (Fig. 4). A single layer of connections is sufficient for our simple orientation discrimination task.

The decision unit aggregates the sensory information using the current weights w_i and the current top-down bias b (Eq. (10)). Gaussian noise ε with mean 0 and standard deviation σ_d models the random fluctuations in the decision-making process.

$$u = \sum_{i=1}^{35} w_i a_i - w_b b + \varepsilon. \quad (10)$$

By convention, negative activation of the decision unit denotes “left,” and positive activation denotes “right.” The activations $a_i = A(\theta_i, f_i)$ of the representation units are always positive (Eq. (9)). The noisy sum u in Eq. (10) defines the net early input to the decision unit. The early activation o' of the unit is a sigmoidal function of the early input u (Eqs. (11) and (12), where γ is the gain parameter of the

decision unit). Note that the sign of o' always equals the sign of u because the activation function G is symmetric. The network implements a commonly used linear classifier with an adjustable decision boundary whose orientation in representation space is set by the current weight vector (Eq. (10)), which achieves near-optimal accuracy for our stimulus environment (Petrov et al., 2005).

$$G(u) = \frac{1 - e^{-\gamma u}}{1 + e^{-\gamma u}} A_{\max}, \quad (11)$$

$$o' = G(u) \quad (\text{early}), \quad (12)$$

$$o = G(u + w_f F) \quad (\text{late}). \quad (13)$$

The initial weights are set in proportion to the preferred orientation of the units: $w_i = (\theta_i/30^\circ)w_{\text{init}}$, reflecting general prior knowledge about orientation that supports above chance performance at the very beginning of practice. The initial weights single out neither the frequency nor the orientation of the Gabor targets; these are factors that must be learned. Step-like initial weights yield very similar results (Petrov et al., 2005).

3.3. Supervised Hebbian learning with feedback inputs

Following the response of the task-dependent reweighting system, feedback—if present—is encoded by the feedback unit and sent as top-down input to the decision unit (Fig. 4). This arrangement is known as “supervised Hebbian learning with soft clamping” in the neural-network literature (e.g., Dayan & Abbott, 2001; O'Reilly & Munakata, 2000). The feedback input F adds to the early input u driving the decision unit, which changes its activation to a new, late level o according to Eq. (13). All learning happens during this late phase (cf. O'Reilly & Munakata, 2000). (A pure Hebbian learning rule would involve the early postsynaptic activation o' as defined by Eq. (12).) The impact of feedback depends upon the weight w_f on the feedback input. The late activation is driven to $\pm A_{\max} = \pm 0.5$ when feedback $F = \pm 1$ is present and the feedback weight is relatively high (hard clamping). Lower feedback weights may simply shift the activation slightly (soft clamping). When no-feedback signal is present ($F = 0$), as in the current experiment, the late decision activation is the same as the early decision activation ($o = o'$), which typically is in the intermediate range.

In the multi-channel Hebbian reweighting model, the only mechanism for long-term changes operates on the synaptic strengths w_i of the connections between the sensory units i and the decision unit. The Hebbian rule (Eq. (14)) is exactly the same both 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} . This associative property makes learning both stimulus and task specific. The low learning rate $\eta = 0.0016$ ensures slow accumulation of statistical information over many trials.

$$\delta_i = \eta a_i (o - \bar{o}), \quad (14)$$

$$\Delta w_i = (w_i - w_{\min})[\delta_i]_- + (w_{\max} - w_i)[\delta_i]_+, \quad (15)$$

$$\bar{o}(t+1) = \rho o(t) + (1 - \rho)\bar{o}(t). \quad (16)$$

Eq. (14) corrects the postsynaptic activation o by its long-term average \bar{o} (Eq. (16) with rate parameter ρ ; t is the trial number). Thus, the Hebbian term δ_i tracks systematic stimulus–response correlations rather than mere response bias. Such sliding baselines in Hebbian models appear well grounded physiologically (e.g., Bienenstock, Cooper, & Munro, 1982). Eq. (15) keeps the weights within bounds by scaling δ_i down in proportion to the remaining range⁴ (O’Reilly & Munakata, 2000). The non-linearity in Eq. (15) sensitizes the model to second-order (variance) statistics. The weights converge to near-optimal values, driven by the signal-to-noise ratios of the sensory units (see Petrov et al., 2005, for detailed analysis).

3.4. 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 top-down input from a bias unit b with weight w_b (Eq. (10)). This mechanism counterbalances the biasing effect of the background texture in the current experiment. Adaptive criterion control may be especially useful in stabilizing performance without feedback in non-stationary environments. Observers are assumed to approximately equalize the frequencies of “left” and “right” responses to match the presentation probabilities in the experiment. The bias $b(t+1)$ on each successive trial equals the current weighted running average $r(t)$ (Eqs. (17) and (18)), discounting the distant past exponentially with a time constant of 50 trials ($\rho = 0.02$, Eq. (17)). Note the similarity between Eq. (16), which tracks the internal activation $o(t)$ of the decision unit, and Eq. (17), which tracks the overt thresholded responses $R(t) = \pm 1$. For parsimony, the same parameter ρ sets the time constant of both equations.

$$r(t+1) = \rho R(t) + (1 - \rho)r(t), \quad (17)$$

$$b(t+1) = r(t). \quad (18)$$

The sensory input provides strong bottom-up bias for the “left” response in context L and the “right” response in context R . This bias propagates into the overt responses and deflects the running average $r(t)$ in the direction of the current context. The bias term b (Eq. (10)) supports the weaker (incongruent) alternative. The resulting negative-feedback loop stabilizes the learning dynamics and promotes balanced response frequencies. The correction

is reactive and always lags behind, leaving a residual response bias favoring the predominant background noise orientation. This residual bias is consistent with the human data. 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. Based on the model fits below, one prominent role of external feedback apparently is to promote stricter criterion control.

In summary, the model has two augmentations relative to canonical unsupervised Hebbian networks: a feedback unit and a bias unit. Both provide top-down inputs to the decision unit and contribute to the accuracy and stability of the model in complementary ways, as discussed later.

4. Model fits

The multi-channel Hebbian reweighting model is implemented in a MATLAB program¹ that takes grayscale images as inputs, produces binary (left/right) responses as outputs, and learns on a trial-by-trial basis. Model performance is directly comparable to the behavioral data and is tested in a simulation experiment that replicates the psychophysical experiment. Excellent quantitative fits demonstrate that an augmented Hebbian system can learn without feedback even in a challenging non-stationary environment. The model accounts for over two hundred data points with four free parameters and explains all the phenomena listed in Table 1.

4.1. Simulation experiment method

Stimulus sequences in the simulation were generated by the same MATLAB program as in the experiment. There are 2000 runs through the 36-block protocol of the present no-feedback study and 2000 feedback runs through the 32-block protocol of the original with-feedback study (Petrov et al., 2005). Half of the runs begin with context L and half begin in context R . On model runs with feedback, a binary top-down signal F is provided only after incorrect responses where explicit feedback was given. Scaled by weight w_f , it changes the late activation of the decision unit. On all other trials the decision activation remains unchanged. The no-feedback runs never provide any feedback. The weights are updated at the end of each trial (Eqs. (14)–(16)).

The model parameters are listed in Table 3. Four parameters are optimized⁵ to fit the no-feedback data set. All other parameters are set a priori or taken verbatim from our original fits to the feedback data set (Petrov et al., 2005). The model is applied to both the no-feedback data from the current experiment and the with-feedback data⁶ of Petrov et al. (2005) to allow a direct comparison.

⁴ The $[x]_+$ operator returns x if $x > 0$ and 0 otherwise; $[x]_-$ returns x if $x < 0$ and 0 otherwise. Thus, repeated potentiation ($\delta_i > 0$) drives the weight up, exponentially approaching the upper bound w_{\max} ; repeated depression ($\delta_i < 0$) drives it down towards the lower bound w_{\min} .

⁵ An iterative parameter search algorithm minimizes the summed squared error. See Petrov et al. (2005) for details.

⁶ The implementation of the task-specific subsystem was modified here to explicitly implement the sliding baseline in the Hebbian Eq. (14).

Table 3
Parameters of the multi-channel reweighting model

Parameter	Value	
<i>Newly estimated parameters</i>		
Decision noise (Eq. (10))	Feedback $\sigma_d = 0.170$	No feedback $\sigma_d = 0.156$
Bias weight (Eq. (10))	$w_b = 2.20$	$w_b = 0.95$
Feedback weight (Eq. (13))	$w_f = 1.80$	Irrelevant
Learning rate (Eq. (14))		$\eta = 0.0016$
Activation gain of the decision unit (Eq. (11))		$\gamma = 5.0$
<i>Parameters taken verbatim from Petrov et al. (2005)</i>		
Maximum activation level (Eqs. (9) and (11))		$A_{\max} = 0.5$
Weight bounds (Eq. (15))		$w_{\min/\max} = \pm 1$
Orientation tuning bandwidth		$h_\theta = 30^\circ$
Frequency tuning bandwidth		$h_f = 1.0$ oct
Radial kernel width (Eq. (8))		$h_r = 2.0$ dva
Running-average rate (Eqs. (16) and (17))		$\rho = 0.02$
Activation gain of the sensory units (Eq. (9))		$\gamma_r = 0.80$
Representation noise (Eq. (8))		$\sigma_r = 0.100$
Initial weight scaling factor		$w_{\text{init}} = 0.17$

Only four parameters are fitted to the current no-feedback data; the rest are set a priori or taken verbatim from Petrov et al. (2005).

4.2. Simulation results and discussion

The model provides an excellent account of both the no-feedback and with-feedback data. Fig. 7 plots the d' learning curves averaged across model runs and compared with the behavioral data. Quantitatively, the model accounts for 92% of the variance of the 108 no-feedback points in Fig. 7A with four free parameters ($R^2 = .915$, $\text{rmse} = .182$). The fit to the feedback data in (B) is also very good ($R^2 = .861$, $\text{rmse} = .224$, replicating Petrov et al., 2005). The slight R^2 difference between the two conditions reflects observer variability and the smaller sample size in the feedback experiment. The model fits approach the fits of the regression (Eq. (4)) discussed earlier ($R^2 = .955$ without and $.945$ with feedback), which are purely descriptive and allow many more factors to vary independently. In contrast, the reweighting model provides a principled, mechanistic explanation with strong ties to direct physical manipulations of factors such as signal contrast.

Of the four estimated parameters, only two parameters differ between the feedback and no-feedback conditions: the bias weight and the decision noise. The bias weight is much stronger in the feedback condition and leads to more stringent criterion control, probably because feedback makes errors more salient. A 10% decision noise difference ($\sigma_d = 0.156$ without and 0.170 with feedback) accommodates the slightly higher asymptotic accuracy in the no-feedback data, which we believe on logical grounds must be due to inter-observer variability.

4.2.1. Robust learning without feedback

The Hebbian reweighting model improves performance without the need of any external feedback by detecting statistical regularities in the stimulus stream and adjusting the weights to improve the classification accuracy. The model predicts the slow but steady general performance

improvement, recurring switch costs superimposed on the general trend, and parallelism at the three contrast levels. The switch costs do not diminish with successive switching because they reflect the same differences between the stimulus statistics in the two noise contexts. Additional model simulations confirm that the switch costs persist long after the general improvement has reached asymptote. The d' curves for the three target contrast levels unfold in parallel because they reflect the adjustment of a common weight vector. The magnitude of the separation between the three target contrast levels in the model d' is driven entirely by the differential signal-to-noise ratio of the input images. Unlike the descriptive regression model, the multi-channel Hebbian reweighting model has no adjustable parameters to control the vertical placement of the individual curves. The good parameter-free fit of this aspect of the data suggests that the model representations are generally appropriate.

4.2.2. Congruence effects

Fig. 8 plots the accuracy predictions for incongruent and congruent stimuli averaged across the 2000 model runs without feedback. The three incongruent curves (A) are widely spaced whereas the congruent ones are almost on top of each other (B), in close agreement with their empirical counterparts in Fig. 3. The model accounts for the counterintuitive reversal of the target contrast effect for congruent stimuli. Accuracy is slightly lower for high-contrast congruent targets than for low-contrast ones, regardless of feedback (see Table 2 and the inset of Fig. 8). The activation of the congruent channels scales little with the target contrast because of background noise saturation. The slight reversal arises from the “leakage” of activation into the incongruent channels when a congruent stimulus is presented. Incongruent channels operate in the steep region of the sigmoidal activation function, amplifying this

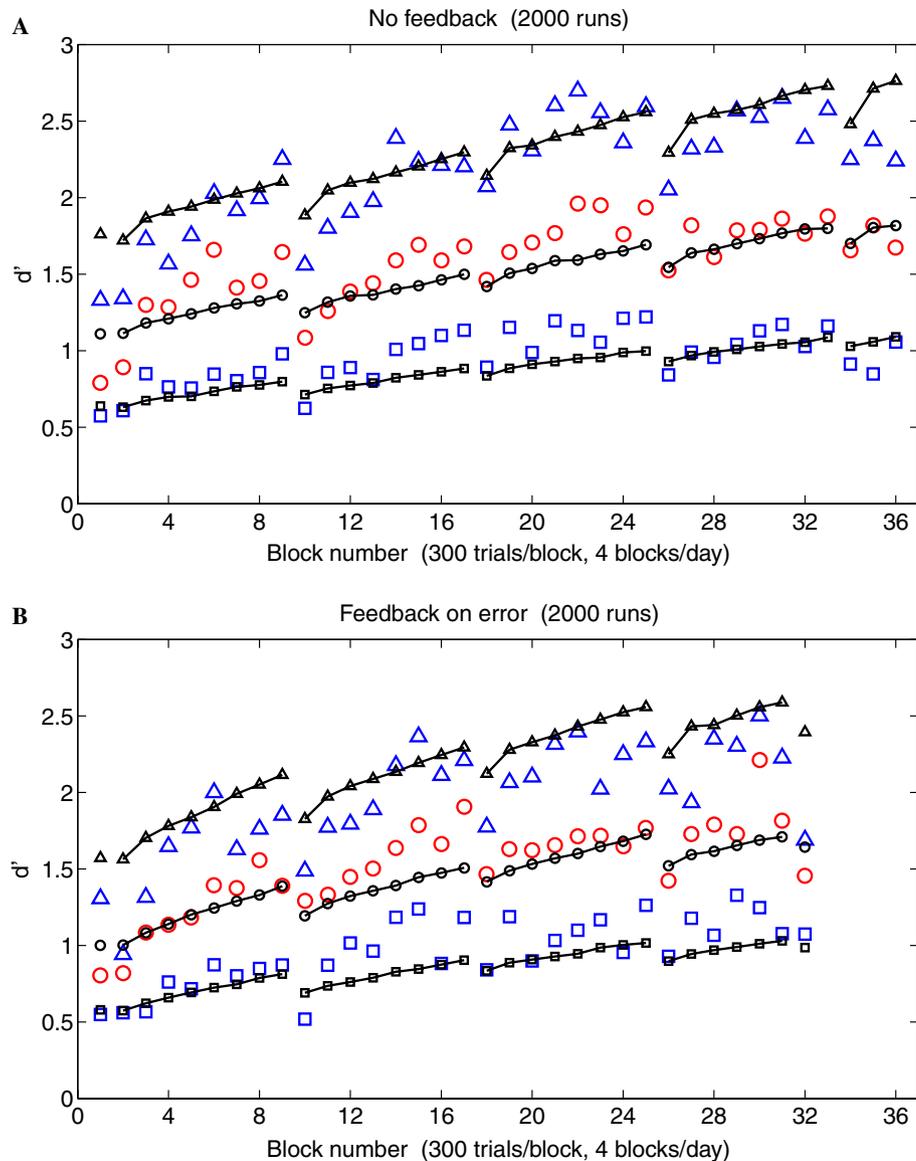


Fig. 7. The model d' learning curves (solid lines) reproduce all patterns in the behavioral data (open symbols) both with (B) and without feedback (A). The connected lines belong to the same context; the discontinuities mark context switches. Target contrast 0.245 (triangles), 0.160 (circles), and 0.106 (squares). Compare with Fig. 2.

leakage. As increased activation in the incongruent channels is interpreted as evidence against the congruent response, a negative correlation between contrast and accuracy emerges for congruent targets.

4.2.3. Weight dynamics

The weight dynamics is qualitatively the same with and without feedback (Fig. 9). The Hebbian learning rule drives the weights somewhat farther away from their initial values during the feedback runs because then the postsynaptic term in Eq. (14) is more consistent. However, the relative weights are almost the same with and without feedback, which also explains the similarity of the switch cost patterns in the observable d' curves.

The initial weights carry no information about which spatial frequency channel contains the target. With

practice, the weights of the different channels are adjusted to match the statistical structure of the stimulus environment. Fig. 9 plots two spatial frequency bands—one centered on the target frequency ($f = 2$ cyc/deg, left) and another a full octave away (4 cyc/deg, right). The weights of the diagnostic units increase with training, whereas the weights of the non-diagnostic units decrease toward zero. The weights of different orientation channels are ordered from -45° (bottom line) to $+45^\circ$ (top line) within each panel. Although the weights on channels at $\theta = \pm 10^\circ$ (thick lines) nearest the target orientation increased, even higher weights develop for the channels tuned to $\theta = \pm 30^\circ$, which turn out to be more diagnostic given the presence of masking noise. This off-channel strategy is consistent with behavioral and physiological data (e.g., Reagan & Beverley, 1985; Schoups, Vogels, Qian, & Orban, 2001),

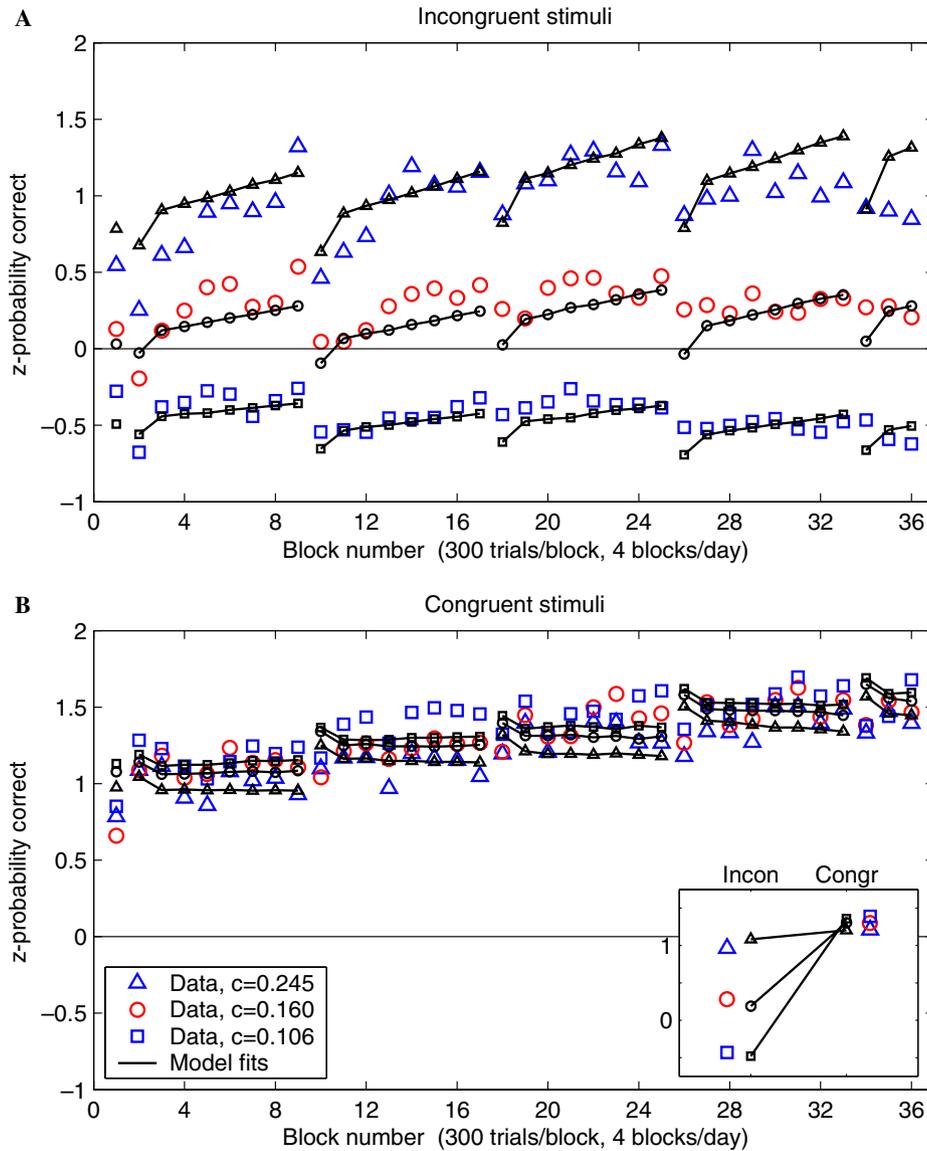


Fig. 8. No-feedback accuracy profiles for targets oriented against (A) or in the direction (B) of the background noise. The inset shows the same six curves collapsed in time. Compare with Fig. 3.

as well as with the negative correlation between contrast and accuracy for our congruent stimuli. Although highly active, the weight of the midline unit ($\theta = 0^\circ$) remains close to zero precisely because it is activated on every trial. The Hebbian learning rule strengthens only units that are correlated with the response states.

Some small but important weight asymmetries contribute to performance costs when the noise backgrounds are switched. Practice in context *L* (Fig. 9, trial 2700) leads the right (incongruent) channels to develop stronger weights than the left (congruent) channels whose reliability is compromised by the background noise—the model learns to “look” for incongruent targets standing out of the noise. When the context switches abruptly, the current weight vector is no longer optimal and the macroscopic *d'* drops until the system adapts to the new conditions.

4.2.4. Adaptive criterion control

The most salient difference between the no-feedback and with-feedback conditions involves the response bias. The bias favoring the congruent response is two times stronger without feedback than with it (compare 64% vs. 36% without feedback and 57% vs. 43% with feedback in Table 1). The error beeps apparently alert the observers that they use the congruent response too often, causing them to apply their bias control strategy more robustly. The estimated bias weight is higher in the feedback condition ($w_b = 2.20$) than in the no-feedback condition ($w_b = 0.95$). The simulated response frequencies are 65% vs. 35% without feedback and 57% vs. 43% with feedback—in excellent agreement with the behavioral data.

Adaptive criterion (bias) control is important for facilitating the recovery after a context switch. At the end of

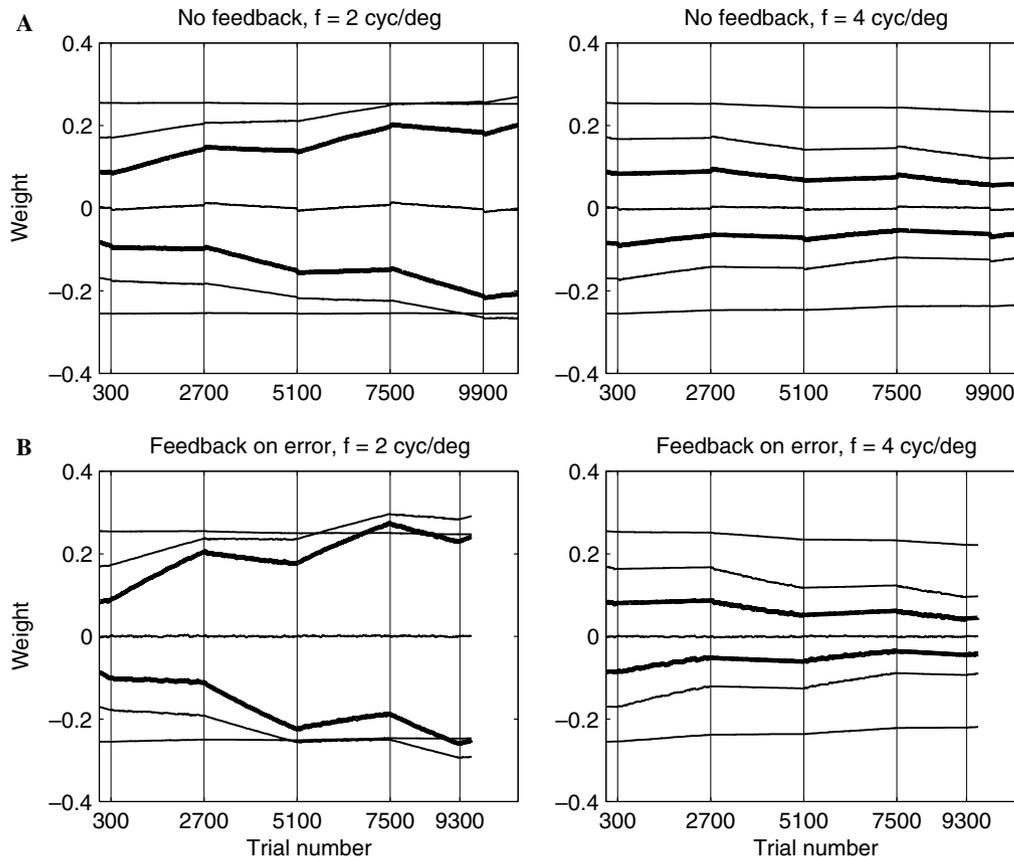


Fig. 9. Weight dynamics of a typical model run under schedule $R-8L-8R-8L-8R-3L$ with (B) and without feedback (A). Each trace corresponds to a particular orientation, with $\pm 15^\circ$ plotted in bold. Hebbian learning strengthens the frequency band of the Gabor targets (2.0 cyc/deg, left) and tunes out irrelevant bands (4.0 cyc/deg, right). Weight asymmetries in the two contexts give rise to switch costs. See text for details.

training in one context, the weight structure emphasizes activation in the “incongruent” channels, which exaggerates the bias once that external noise has switched. The bias unit provides compensatory top-down input. Simulations eliminating adaptive criterion control ($w_b = 0$) exhibited slower initial learning and much slower recovery from switch costs.

4.2.5. Supervised Hebbian learning with feedback

Although generally consistent with a number of reports in the literature, a surprising aspect of the comparison between our feedback and no-feedback learning data is the near equivalence of the rates of learning. Hebbian models augmented by feedback with reasonable weight ($w_f > 0$) predict that accurate feedback should benefit learning at least to some extent. The similarity of the data in the two conditions suggests either that the feedback weight is very low or that feedback is not effective on every trial. We selected an implementation in which feedback is incorporated only on error trials—that is, only when an explicit error buzz occurred in the experiment.⁷ Alternatively, feed-

back could be incorporated on a random subset of trials, or on all trials but with very low weight. Various other implementations were considered, but led to counter-intuitive compensatory changes in other parameters between the two conditions.

4.2.6. Relationship between feedback and bias

There are similarities and differences between the feedback and bias units. Both provide top-down inputs to the decision unit that tend to counteract the bottom-up preponderance of activity in the congruent channels. Increasing the bias weight can compensate for decreased feedback weight and vice versa, within limits. Despite these similarities, however, the bias and the feedback inputs are structurally and functionally distinct. Their respective temporal dynamics is different, as is their impact on the overall processing. The timing of the feedback unit introduces the distinction between “early” and “late” postsynaptic activations (Eqs. (12) and (13), respectively; cf. contrastive Hebbian learning, O’Reilly & Munakata, 2000). The stronger the feedback, the weaker the impact of the system’s own responses on the learned weights. The bias unit, by contrast, changes very slowly according to the running-averaging Eq. (17). The stronger the bias, the stronger the impact

⁷ This differs from our original simulations (Petrov et al., 2005) where feedback was introduced on every trial. See Section 5.1 for details.

of the system's own responses on the learned weights. The adaptive bias also introduces implicit competition between the two responses.

5. General discussion

In this article, we demonstrated perceptual learning without feedback in a challenging non-stationary environment. The experiment generated a complex and articulated data set, including: overall improvements in d' without feedback, higher accuracies for higher target contrasts, with parallel improvements with practice, persistent switch-costs at each change of the external noise context, interactions between target contrast and congruence, and an overall response bias favoring the context-congruent orientation (Table 1). The context alternations could have caused significant problems for non-error-correcting learning systems, yet perceptual learning succeeded without feedback.

A theoretical framework for perceptual learning based on selective reweighting (Doshier & Lu, 1998, 1999) and instantiated in a fully implemented computational model (Petrov et al., 2005) was tested in this new no-feedback environment. Without feedback, the multi-channel reweighting model learns the statistical correlations between the stimulus representation and its own internal responses using Hebbian learning augmented by a response-bias unit and capitalizing on initial above-chance accuracy. With only four free parameters, the model provides an elegant and accurate account for the learning dynamics across thousands of trials and for various interlocking patterns in a complex data set.

The model embodies a number of established principles of human perception and cognition: orientation- and frequency-tuned representations, contrast gain control, weighted decision units, incremental associative learning, and intrinsic variability. Its representation subsystem is based on fairly standard assumptions and its tuning parameters are representative of the tuning properties of V1 neurons. It operates on the stimulus images themselves. The task-specific subsystem embodies the principle of learning through incremental selective reweighting in a deliberately simplified, one-layer neural network. This principle is consistent with the available physiological, functional, psychophysical, and computational evidence (Petrov et al., 2005). The Hebbian learning hypothesis—augmented with feedback inputs and adaptive criterion control—arises naturally within the selective reweighting framework and is examined on the solid foundation of the mechanistically explicit and fully functional multi-channel model.

5.1. Feedback in perceptual learning

The primary result of the current empirical study concerns the robustness of perceptual learning without feedback, even in a non-stationary environment alternating the dominant orientation of the external noise context from

epoch to epoch. This study, then, joins other reports of successful learning without external feedback (Ball & Sekuler, 1987; Crist et al., 1997; Fahle & Edelman, 1993; Karni & Sagi, 1991; McKee & Westheimer, 1978; Shiu & Pashler, 1992; but see Vallabha & McClelland, in press). In conjunction with our previous study with feedback (Petrov et al., 2005), the present results go beyond a simple demonstration of the ability to learn without feedback. They document a close similarity of most qualitative and quantitative patterns of learning in the two conditions.

The primary role of feedback in the current experiments was to induce more stringent criterion (bias) control. Feedback may have had its primary influence on error trials when an explicit signal was given (Petrov et al., 2005), and did not substantially influence the learning rate. Yet the potential of the Hebbian model to be influenced (supervised) by feedback is important because the literature suggests that feedback can be quite important in some situations. Although it is possible to learn without feedback, in some cases feedback seems necessary for perceptual learning, especially for difficult stimuli (Herzog & Fahle, 1997; Shiu & Pashler, 1992). In other cases, feedback improves the rate or extent of learning (Fahle & Edelman, 1993; Ball & Sekuler, 1987). Overall, the cumulative evidence indicates that feedback can be important in a range of circumstances, particularly when the task difficulty is high and the initial performance is low.

The augmented Hebbian reweighting model is compatible not just with the current data, but also with this broad empirical literature. Unsupervised Hebbian learning with no feedback inputs is sufficient to account for the possibility of learning without feedback. It fails, however, to account for the positive role of feedback in other circumstances. The supervised Hebbian model explains why feedback may even be “necessary” in some cases (at least over the time scale of a perceptual learning experiment). For near-threshold stimulus blocks (e.g., Herzog & Fahle, 1997) where the initial response accuracy is low, unsupervised Hebbian learning may need a very long time to discover the weak statistical correlations. Similarly, the Hebbian framework explains why late introduction of feedback may fail to improve performance further (Herzog & Fahle, 1997; McKee & Westheimer, 1978). If the model finds an optimal or near-optimal weight vector before external feedback becomes available, then augmentation with feedback at that point should be unimportant.

In summary, the augmented Hebbian learning hypothesis and the multi-channel reweighting model that instantiates it are consistent with both the current data and the broader literature. The varied effects of feedback on perceptual learning can be understood in this simple and elegant framework.

5.2. Adaptive criterion control

The adaptive criterion control mechanism is a second augmentation of the Hebbian learning model. The bias

weight w_b is estimated to be much stronger in the feedback condition than in no-feedback condition. Indeed, this model parameter accounts for one of the largest differences between the two data sets—the bias toward the context-congruent response was twice as large without feedback as with it. Apparently, the feedback drew the observer's attention to errors associated with a particular response. The adaptive criterion control strategy implemented here measures deviations from the 50–50 response frequencies appropriate for the counterbalanced experiment. Some form of differential criterion control appears necessary to account for the different biases in the two feedback conditions.

Adaptive criterion control is particularly important in non-stationary environments without feedback—the ones that approximate the real world. Additional simulations indicated that without criterion control the model failed to recover after the first context switch in the no-feedback condition. Despite very strong initial response bias, the model improves initially by increasing the weights of the channels for the incongruent orientation at the target frequency. These weight changes, however, only exacerbate the response bias when the context switches. When the adaptive criterion control mechanism is operational, it quickly moves the system (correctly) back in the other direction. Without this mechanism, the d' in many cases actually becomes worse than the initial unpracticed d' and never fully recovers because the system tends to be out of phase with the optimal solution. We are not aware of other cases in the literature where adaptive criterion control has been proposed as an important function in perceptual learning.

Finally, we speculate that the criterion control mechanism might provide at least a partial explanation for other phenomena such as the recovery from false or misleading feedback (Herzog & Fahle, 1997, 1999). Also, blocked (as opposed to trial-by-trial) feedback may be sufficient to inform (increase the weight of) the criterion control mechanism, despite the limited utility of block feedback for the reweighting mechanism. This is a possible explanation of some counterintuitive effects of block feedback (Herzog & Fahle, 1997; Shiu & Pashler, 1992).

5.3. Hebbian learning and task structure

The augmented Hebbian hypothesis proposes a hybrid form of learning that introduces feedback- and criterion-related inputs into an intrinsically unsupervised model. Despite this hybrid quality, the Hebbian core suggests that perceptual learning will have many of the qualities of pure Hebbian learning. It may be more successful when some easy stimuli are included (e.g., Rubin, Nakayama, & Shapley, 1997). Also, perceptual learning will probably be best suited to binary choices in simpler decision spaces with approximately linear boundary solutions, and where the first principle component of the stimulus distribution is decision-relevant (cf. Hertz et al., 1991). The current task

environment has all these properties (see Petrov et al., 2005, for details), and this seems true for most tasks in the visual perceptual learning literature. Perceptual learning might be difficult or impossible in tasks that violate these boundary conditions. In the final analysis, perceptual learning without feedback in discrimination derives from the statistical simplicity of these tasks.

5.4. Selective reweighting versus representation modification

The multi-channel reweighting model (Petrov et al., 2005) instantiates, illustrates, and proves the empirical adequacy of the selective reweighting hypothesis (Doshier & Lu, 1998, 1999; Mollon & Danilova, 1996). Perceptual learning is hypothesized to occur in the connections between the representation and task-specific subsystems. The representations themselves need not change. This view is different from (though not incompatible with) the oft-cited, and perhaps dominant, *representation modification hypothesis* that focuses perceptual learning within the representation subsystem (e.g., Karni & Sagi, 1991; Teich & Qian, 2003). This alternative hypothesis attributes the behavioral improvement to sharper, more coherent and/or reliable stimulus representations after practice. Various mechanisms for representation modification are sketched in the literature, but few are formalized explicitly and none is implemented in sufficient detail to handle actual images or provide a model of incremental learning (see Gilbert et al., 2001; Tsodyks & Gilbert, 2004, for reviews). The stimulus specificity of perceptual learning (e.g., Ahissar & Hochstein, 1996; Ball & Sekuler, 1987; Crist et al., 1997; Fahle, 1997; Fiorentini & Berardi, 1980; Shiu & Pashler, 1992), often cited as conclusive evidence for representation modification, is equally consistent with both explanations. The representation units are selectively tuned for various stimulus features and any plasticity mechanism involving these units—in either presynaptic or postsynaptic position—gives rise to behavioral effects that inherit the same specificity.

Petrov et al. (2005) reviewed all available physiological, functional, psychophysical, and computational evidence in the visual modality and concluded that selective reweighting plays a primary role in visual perceptual learning, although representation modification may also contribute to some extent. Psychophysical demonstrations of stimulus and task specificity of perceptual learning (Ahissar & Hochstein, 1993; Crist et al., 1997; Fahle, 1997; Fahle & Morgan, 1996; Saffell & Matthews, 2003; Shiu & Pashler, 1992) are also consistent with changes in the connections between representation and decision. A new task analysis (Petrov et al., 2005) distinguished between transfer experiments in which two tasks share: (i) input representations, (ii) read-out connections to a decision unit, (iii) both or (iv) neither. Virtually all existing behavioral demonstrations of stimulus specificity plausibly involve distinct stimulus representations and hence neither test nor support representation modification. Direct evidence from

single-cell recordings from intact visual cortex in adult behaving monkeys is also broadly consistent with the reweighting claim (Doshier & Lu, 1998, 1999). Four recent studies found little or no practice-induced changes in V1, V2, and/or V4 (Crist, Li, & Gilbert, 2001; Ghose, Yang, & Maunsell, 2002; Schoups et al., 2001; Yang & Maunsell, 2004). Small representation changes in V4 (barely reaching statistical significance, Yang & Maunsell, 2004) may reflect reweighting from V1 and V2, whereas the modest changes in V1 (Schoups et al., 2001) are hard to replicate (Ghose et al., 2002) and cannot account for the massive behavioral improvement in these studies (according to an ideal observer model of population activity, Schoups et al., 2001). Furthermore, there are tight functional constraints on the primary visual areas given their involvement in a wide range of closely calibrated computations. The representations themselves may need to change only to compensate for damage. Additionally, sharpening the tuning curves of the representation units in the multi-channel reweighting model accounted for less than 10% of the overall learning effect in our experimental paradigm (Petrov et al., 2005). This replicates the inability of sharpening to account for the data of Schoups et al. (2001). By contrast, Hebbian reweighting reproduces the full magnitude of the learning effect, the switch costs, and several other interlocking patterns. All this converging evidence justifies the selective reweighting scheme instantiated by our model. On the other hand, the tight functional constraints on early vision make it extremely challenging to implement a representation modification scheme. The present simulations and our earlier analyses indicate that the latter scheme is neither necessary nor sufficient (Petrov et al., 2005).

5.5. Comparison with other models

The full task analysis of the behavioral literature in visual perceptual learning, as well as physiological reports (Doshier & Lu, 1998, 1999; Petrov et al., 2005) motivated the reweighting hypothesis. An analysis of the role of feedback motivated the Hebbian learning model. This model is related to a number of general models of learning and earlier models of perceptual learning. The Hebbian learning mechanism is related to general competitive learning models (e.g., Carpenter & Grossberg, 1987; Kohonen, 1995; O'Reilly & Munakata, 2000; Petrov & Anderson, 2005; Rumelhart & Zipser, 1985), although in our model competition occurs only indirectly through criterion control. In the domain of perceptual learning, related models have been applied to vernier hyperacuity (Herzog & Fahle, 1998; Weiss, Edelman, & Fahle, 1993), motion perception (Vaina, Sundaeswaran, & Harris, 1995), bisection (Zhaoping, Herzog, & Dayan, 2003), and phoneme categorization (Vallabha & McClelland, in press) (see Tsodyks & Gilbert, 2004, for an excellent review). Each model assumes an appropriate stimulus representation, reflecting the modeled task, and none proposes systematic changes in representation. All models postulate incremental learning, consistent

with perceptual learning data. However the actual mechanisms of this learning differ considerably.

Our research follows the pioneering work of Herzog and Fahle (1997, 1998) in concluding that neither purely supervised nor purely unsupervised learning rules can account for the multi-faceted effects of feedback on perceptual learning. Weiss et al. (1993) proposed an error-correcting system based on a Widrow-Hoff rule in a radial basis network, in which learning without feedback occurs only on a small subset of very easy trials. Vaina et al. (1995) report simulations with a self-supervised learning rule that involves a non-local term and a threshold on the presynaptic activation that promotes weight sparseness that ignores noisy sensory channels, but does not directly incorporate external feedback. The limited-capacity sampling model of Liu and Weinsall (2000) implements a similar idea (see also Ahissar & Hochstein, 2004). The model of Herzog and Fahle (1998) uses a more complex multi-level representation with hidden units in which feedforward connections are gated by learnable top-down (attention) connections; feedback is used to control learning rate, but not to provide response targets. Their system evaluates a given weight configuration for some time, and then discards the changes if they do not improve performance—a notion of neurological reversibility. Other models also emphasize more complex or non-linear problem spaces (Zhaoping et al., 2003). The current proposal builds upon this earlier work, particularly that of the closely similar model of Vaina et al. (1995). Our model differs in its focus on standard representations, biological plausibility, explicit top-down inputs within a (supervised or unsupervised) Hebbian framework augmented by a mechanism of criterion control. There are differences of application and development as well. Our approach identified a relatively rare set of test designs that are capable of contrasting reweighting and representation change. Finally, our model was fully implemented and quantitatively tested against a complex and demanding data set with internal checks for physical manipulations such as contrast, and switch costs on classification of identical input representations.

Further research is needed to assess the adequacy of the augmented Hebbian approach more fully. The exact conditions in which learning without feedback slows down or fails must be identified. The overall task difficulty and the initial accuracy seem two important factors, as well as the investigation of more complex statistical structure, integrating the perceptual learning literature with the non-perceptual categorization literature (e.g., Ashby, Alfonso-Reese, Turken, & Waldron, 1998). Studying the effects of various feedback manipulations is of obvious significance too. Aggregated block feedback seems a particularly interesting form of feedback that can be as efficient as trial-by-trial feedback in conditions when no-feedback learning fails (Herzog & Fahle, 1997; Shiu & Pashler, 1992). Block feedback may have its effect through triggering more rigorous criterion control. Although additional research is needed to test the various forms of experimental feedback, the

mechanism of either unsupervised or supervised Hebbian learning plus criterion control advocated here seems entirely consistent with all currently available evidence.

6. Conclusion

Behavioral studies of the role of feedback in perceptual learning can provide valuable insights about the underlying neural plasticity. In this article, we propose a Hebbian framework augmented with top-down feedback input and bias control. Even when no external feedback is available as a supervisory signal, the statistical distribution of the sensory evidence across multiple channels is sufficient to guide the learning process. Although a number of previous studies report learning without feedback, the present study is the first to demonstrate that human perceptual learning does not require feedback even in challenging non-stationary environments, to directly compare performance in this environment to with-feedback performance, and to provide a quantitative model for the observed perceptual learning. A multi-channel reweighting model accounts for hundreds of no-feedback data points with only four free parameters, only two of which differ between the no-feedback and with-feedback fits. The representations themselves never change. External feedback, when present, acts as top-down input to the decision unit. An adaptive criterion control mechanism stabilizes the learning dynamics and balances the response frequencies. This parsimonious, biologically plausible model establishes the empirical adequacy of the feedback-supervised or unsupervised Hebbian reweighting hypotheses, augmented by criterion control.

Acknowledgments

This research is supported by grants from the National Institute of Mental Health and the National Science Foundation.

References

- Adelson, E. H., & Bergen, J. R. (1985). Spatiotemporal energy models for the perception of motion. *Journal of the Optical Society of America, A*, 2(2), 284–299.
- Ahissar, M., & Hochstein, S. (1993). Attentional control in early perceptual learning. *Proceedings of the National Academy of Sciences, USA*, 90, 5718–5722.
- Ahissar, M., & Hochstein, S. (1996). Learning pop-out detection: Specificities to stimulus characteristics. *Vision Research*, 36(21), 3487–3500.
- Ahissar, M., & Hochstein, S. (2004). The reverse hierarchy theory of visual perceptual learning. *Trends in Cognitive Sciences*, 8(10), 457–464.
- Ahumada, A. J., & Watson, A. B. (1985). Equivalent-noise model for contrast detection and discrimination. *Journal of the Optical Society of America, A2*, 1133–1139.
- Ashby, F. G., Alfonso-Reese, L. A., Turken, A. U., & Waldron, E. M. (1998). A neuropsychological theory of multiple systems in category learning. *Psychological Review*, 105(3), 442–481.
- Ball, K., & Sekuler, R. (1987). Direction-specific improvement in motion discrimination. *Vision Research*, 27, 953–965.
- Bienenstock, E. L., Cooper, L. N., & Munro, P. W. (1982). Theory for the development of neuron selectivity: Orientation specificity and binocular interaction in visual cortex. *Journal of Neuroscience*, 2, 32–48.
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, 10, 433–436.
- Burgess, A. E., Wagner, R. F., Jennings, R. J., & Barlow, H. B. (1981). Efficiency of human visual signal discrimination. *Science*, 214, 93–94.
- Cannon, M. W., & Fullenkamp, S. C. (1991). Spatial interactions in apparent contrast: Inhibitory effects among grating patterns of different spatial frequencies, spatial positions, and orientations. *Vision Research*, 31(11), 1985–1998.
- Carandini, M., Heeger, D. J., & Movshon, J. A. (1997). Linearity and normalization in simple cells of the macaque primary visual cortex. *Journal of Neuroscience*, 17(21), 8621–8644.
- Carpenter, G. A., & Grossberg, S. (1987). A massively parallel architecture for a self-organizing neural pattern recognition machine. *Computer Vision, Graphics, and Image Processing*, 37, 54–115.
- Chubb, C., Sperling, G., & Solomon, J. A. (1989). Texture interactions determine perceived contrast. *Proceedings of the National Academy of Sciences, USA*, 8, R33–R76.
- Crist, R. B., Kapadia, M. K., Westheimer, G., & Gilbert, C. D. (1997). Perceptual learning of spatial location: Specificity for orientation, position, and context. *Journal of Physiology*, 78, 2889–2894.
- Crist, R. B., Li, W., & Gilbert, C. D. (2001). Learning to see: Experience and attention in primary visual cortex. *Nature Neuroscience*, 4(5), 519–525.
- Dayan, P., & Abbott, L. F. (2001). *Theoretical neuroscience: Computational and mathematical modeling of neural systems*. Cambridge, MA: MIT Press.
- De Valois, R. L., & De Valois, K. K. (1988). *Spatial vision*. New York: Oxford University Press.
- De Valois, R. L., Yund, E. W., & Hepler, H. (1982). The orientation and direction selectivity of cells in macaque visual cortex. *Vision Research*, 22, 531–544.
- 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, USA*, 95(23), 13988–13993.
- Dosher, B. A., & Lu, Z.-L. (1999). Mechanisms of perceptual learning. *Vision Research*, 39, 3197–3221.
- Dosher, B., & Lu, Z.-L. (2005). Perceptual learning in clear displays optimizes perceptual expertise: Learning the limiting process. *Proceedings of the National Academy of Sciences, USA*, 102(14), 5286–5290.
- Fahle, M. (1997). Specificity of learning curvature, orientation, and vernier discriminations. *Vision Research*, 37(14), 1885–1895.
- Fahle, M., & Edelman, S. (1993). Long-term learning in vernier acuity: Effects of stimulus orientation, range, and feedback. *Vision Research*, 33(3), 397–412.
- Fahle, M., & Morgan, M. (1996). No transfer of perceptual learning between similar stimuli in the same retinal position. *Current Biology*, 6(3), 292–297.
- Fahle, M., & Poggio, T. (Eds.). (2002). *Perceptual learning*. Cambridge, MA: MIT Press.
- Ferster, D., & Miller, K. D. (2000). Neural mechanisms of orientation selectivity in the visual cortex. *Annual Review of Neuroscience*, 23, 441–471.
- Fiorentini, A., & Berardi, N. (1980). Perceptual learning specific for orientation and spatial frequency. *Nature*, 287(4), 43–44.
- 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.
- Gilbert, C. D., Sigman, M., & Crist, R. E. (2001). The neural basis of perceptual learning. *Neuron*, 31(5), 681–697.
- Gonzalez, R. C., & Woods, R. E. (1992). *Digital image processing*. Reading, MA: Addison-Wesley.
- Graham, N., & Sutter, A. (2000). Normalization: Contrast-gain control in simple (Fourier) and complex (non-Fourier) pathways of pattern vision. *Vision Research*, 40(20), 2737–2761.

- Heeger, D. J. (1992a). Half-squaring in responses of cat striate cells. *Visual Neuroscience*, 9, 427–443.
- Heeger, D. J. (1992b). Normalization of cell responses in cat striate cortex. *Visual Neuroscience*, 9, 181–197.
- 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(15), 2133–2141.
- Herzog, M. H., & Fahle, M. (1998). Modeling perceptual learning: Difficulties and how they can be overcome. *Biological Cybernetics*, 78, 107–171.
- Herzog, M. H., & Fahle, M. (1999). Effects of biased feedback on learning and deciding in a vernier discrimination task. *Vision Research*, 39, 4232–4243.
- 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, USA*, 88, 4966–4970.
- Knutsson, H., & Granlund, G. H. (1983). Texture analysis using two-dimensional quadrature filters. In 1983 IEEE computer society workshop on computer architecture for pattern analysis and image database management (pp. 206–213). Silver Spring, MD: IEEE Computer Society.
- Kohonen, T. (1995). *Self-organizing maps* (3rd ed.). Berlin, Germany: Springer-Verlag.
- Liu, Z., & Weinshall, D. (2000). Mechanisms of generalization in perceptual learning. *Vision Research*, 40, 97–109.
- Lu, Z.-L., Chu, W., Doshier, B. A., & Lee, S. (2005). Independent perceptual learning in monocular and binocular motion systems. *Proceedings of the National Academy of Sciences, USA*, 102, 5624–5629.
- Lu, Z.-L., & Doshier, B. A. (1998). External noise distinguishes attention mechanisms. *Vision Research*, 38(9), 1183–1198.
- Lu, Z.-L., & Doshier, B. A. (1999). Characterizing human perceptual inefficiencies with equivalent internal noise. *Journal of the Optical Society of America*, 16(3), 764–778.
- Lu, Z.-L., & Doshier, B. A. (2000). Spatial attention: Different mechanisms for central and peripheral temporal precues? *Journal of Experimental Psychology: Human Perception and Performance*, 26(5), 1534–1548.
- Lu, Z.-L., & Sperling, G. (1999). Second-order reversed phi. *Perception and Psychophysics*, 61, 1075–1088.
- The MathWorks. (1999). *MATLAB user's guide*. Natick, MA: The MathWorks, Inc.
- McKee, S. P., & Westheimer, G. (1978). Improvement in vernier acuity with practice. *Perception and Psychophysics*, 24(3), 258–262.
- Mollon, J. D., & Danilova, M. V. (1996). Three remarks on perceptual learning. *Spatial Vision*, 10(1), 51–58.
- Movshon, J. A., Thompson, I. D., & Tolhurst, D. J. (1978). Spatial and temporal contrast sensitivity of neurones in areas 17 and 18 of the cat's visual cortex. *Journal of Neurophysiology*, 283, 101–120.
- 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., & Zhang, L. (1991). Accurate control of contrast on microcomputer displays. *Vision Research*, 31, 1337–1350.
- Petrov, A. A., & Anderson, J. R. (2005). The dynamics of scaling: A memory-based anchor model of category rating and absolute identification. *Psychological Review*, 112(2), 383–416.
- Petrov, A. A., Doshier, B. A., & Lu, Z.-L. (2005). The dynamics of perceptual learning: An incremental reweighting model. *Psychological Review*, 112(4), 715–743.
- Pollen, D. A., & Ronner, S. (1981). Phase relationship between adjacent simple cells in the visual cortex. *Science*, 212, 1409–1411.
- Reagan, D., & Beverley, K. I. (1985). Postadaptation orientation discrimination. *Journal of the Optical Society of America*, 2(2), 147–155.
- Rubin, N., Nakayama, K., & Shapley, R. (1997). Abrupt learning and retinal specificity in illusory-contour perception. *Current Biology*, 7, 461–467.
- Rumelhart, D. E., & Zipser, D. (1985). Feature discovery by competitive learning. *Cognitive Science*, 9, 75–112.
- Saffell, T., & Matthews, N. (2003). Task-specific perceptual learning on speed and direction discrimination. *Vision Research*, 43, 1365–1374.
- Schoups, A. A., Vogels, R., Qian, N., & Orban, G. A. (2001). Practising orientation identification improves orientation coding in V1 neurons. *Nature*, 412, 549–553.
- Shiu, L.-P., & Pashler, H. (1992). Improvement in line orientation discrimination is retinally local but dependent on cognitive set. *Perception and Psychophysics*, 52(5), 582–588.
- Teich, A. F., & Qian, N. (2003). Learning and adaptation in a recurrent model of V1 orientation selectivity. *Journal of Neurophysiology*, 89(4), 2086–2100.
- Tsodyks, M., & Gilbert, C. (2004). Neural networks and perceptual learning. *Nature*, 431, 775–781.
- Vaina, L. M., Sundaeswaran, 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.
- Vallabha, G.K., & McClelland, J.L. (in press). Success and failure of new speech category learning in adulthood: Consequences of learned Hebbian attractors in topographic maps. *Cognitive, Affective & Behavioral Neuroscience*.
- 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(7), 1617–1626.
- Zhaoping, L., Herzog, M. H., & Dayan, P. (2003). Quadratic ideal observation and recurrent preprocessing in perceptual learning. *Network: Computation in Neural Systems*, 14, 233–247.