Supporting Information

Dosher et al. 10.1073/pnas.1312552110

SI Materials and Methods

Methods. ***Observers.*** Observers had normal to corrected-to-normal vision and provided written consent under a University of California, Irvine institutional review board protocol. They participated in 9,984 experimental trials per observer. Observers were randomly assigned to experimental groups (P, O, or OP, with 12, 11, and 10 completed observers, respectively).

Design. Observers discriminated the orientations of a Gabor (windowed sinewaves) tilted ±5° top clockwise ("right") or counterclockwise ("left") from a reference angle R of either –35° or +55°. Test locations were on one or another diagonal (NW/SE or NE/SW) at 5.67° eccentricity. Zero noise and high noise test conditions were intermixed within each testing block. Observers were randomly assigned to switch group, initial reference angle, and training positions. They trained for eight blocks over four sessions in an initial task, switched tasks, and trained for another eight blocks or four sessions. Contrast thresholds were measured using adaptive staircase methods (*SI Materials and Methods, Staircases*) tracking 75% correct.

Apparatus. A Macintosh G4 using the internal 10-bit video card with a refresh rate of 67 Hz and resolution of 640 × 480 pixels displayed stimuli on a 19" Viewsonic color monitor. Luminance was calibrated with psychophysical matching judgments and measurement with photometer. A lookup table linearized the luminance range into 127 levels from 1 cd/m^2 to 67 cd/m^2. A chin rest stabilized the observer's head.

Stimuli. The signal Gabor appeared within a 64 × 64 pixel patch (3° × 3° visual angle at a viewing distance of 72 cm): $l(x,y) = l_0\left(1.0 \pm c\sin(2\pi f(y\sin(\theta) \pm x\cos(\theta)) \times \exp\left(\frac{x^2+y^2}{2\sigma^2}\right)\right)$, with angle θ of –35° ± 5° or +55° ± 5°, spatial frequency f = 2 cycles per degree (cpd), and SD of the Gaussian envelope σ = 0.4°, maximum contrast c, and midgray luminance l_0 is the midgray. Each 64 × 64 noise image had 2 × 2 pixel noise elements with Gaussian-distributed contrasts with mean value 0 and SD 0.33, newly generated for each frame and trial. Signal and noise images were combined by alternation at the frame rate (*SI Materials and Methods*, *Procedure*). A single Gabor appeared on any trial, preceded by 150 ms by a brief arrow indicating the test location. Stimuli were generated using MATLAB with PsychToolbox extensions.

Sample images and power spectra for the eight stimuli (four angles, with and without external noise) are shown in Fig. S1. The signal Gabors and the external noise images were integrated via temporal integration of frames shown in rapid succession, as described next.

Procedure. Observers trained at one reference orientation and layout for four sessions and then switched to a new task for another four sessions. Each session was 1,248 trials of four blocks of 312 trials plus two practice trials. Sessions 1 and 5 began with a brief preexposure to the locations and timing, followed by 10 practice trials. A beep occurred 250 ms after the central fixation cross; 250 ms later, the stimulus sequence (two white noise frames + two Gabor frames + two new noise frames) appeared for a total of 90 ms (15 ms/frame), with a central precue arrow appearing 100 ms before the oriented Gabor, which prevented eye movements. Staircases set the contrasts of the Gabors, as described in the next section. Observers responded with the "j" ("clockwise/right") or "f" key ("counterclockwise/left"). A negative feedback tone followed each error. There was a 750 ms intertrial interval.

Staircases. Adaptive staircases (1) were used to track threshold Gabor contrasts. The 3/1 and 2/1 staircases track accuracies of 79.4% and 70.7% correct, respectively, by reducing contrast by 10% after either three or two consecutive correct responses, and increasing contrast 10% after each incorrect response. Staircases were interleaved, one for each external noise level, target accuracy, and testing location. There were 1,248 trials per session, 672 in 3/1 staircases and 576 in 2/1 staircases organized into four blocks (84 and 72 trials per block per staircase for 3/1 and 2/1 staircases, respectively). Threshold contrast levels averaged an even number of reversals for each staircase sequence, excluding the first four or five, to limit bias. Averaging the 3/1 and 2/1 staircases tracks an overall performance level of ∼75% correct. These were averaged over two blocks, to yield two larger block thresholds per session.

Model Implementation of the Integrated Reweighting Theory. The integrated reweighting theory (IRT) is a multichannel, multilevel network model to account for transfer of perceptual learning over location. Learning occurs through channel reweighting (2). The IRT architecture includes location-independent as well as location-specific representations. It incorporates computational modules from the prior augmented Hebbian reweighting model (AHRM) that modeled perceptual learning in a single location (3). The IRT makes predictions about transfer of perceptual learning, and also has implications (not tested in this study) concerning training task mixtures. Here we provide a description and equations for the representation subsystem and the learning system, following refs. 3, 4.

The representation subsystem consists of orientation- and frequency-selective units. The activation value $A(\theta, f)$ of the representation units corresponds to the normalized spectral energy in each channel. A set of retinotopic phase-sensitive maps $S(x, y,\theta, f, \phi)$ is computed for the input image $I(x, y)$:

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

for units tuned to different spatial frequencies f, orientations θ, and spatial phases ϕ. We assumed five spatial frequency bands at {1, 1.4, 2, 2.8, 4 c/d}, orientations sampled every 15°, or {0°, ±15°, ±30°, ±45°, ±60°, ±75°, +90°(=–90°)}, and four spatial phases {0°, 90°,180°, 270°}. Based on estimates of cellular tuning bandwidths in primary visual cortex, the bandwidth of the spatial frequency tuning was set at h_f = 1 octave, and the bandwidth of the orientation tuning was set at h_θ = 30° (half-amplitude full-bandwidth). The location-specific representations used these values, whereas the more broadly tuned location-invariant representations used spatial frequency and orientation bandwidths of 1.6× those of the location-specific units.

The input image $I(x, y)$ is convolved with the filter for each unit by fast Fourier transform. This is followed by a half-squaring rectification operation that generates phase-sensitive activation maps analogous to "simple cells," followed by spatial phase pooling and inhibitory normalization (5): $E(x,y,\theta,f) = \sum_\phi S(x,y,\theta,f,\phi) + \varepsilon_1$ and $C(x,y,\theta,f) = \frac{aE(x,y,\theta,f)}{k+N(f)}$, respectively. The noise term ε_1 is Gaussian-distributed internal noise source with mean 0 and SD σ_1. The normalization pool $N(f)$, consistent with physiological reports, is essentially independent of orientation and only modestly tuned for spatial frequency. The value a is a scaling factor, and k is the saturation constant that operates at near-zero contrasts. Spatial phase is pooled for the orientation judgment task, where phase is not relevant, but it could be retained for judgments that are contingent on phase. The spatial activation maps are

combined and weighted through a Gaussian kernel of radius W_r commensurate with the size of the target Gabor. A second Gaussian-distributed noise of mean 0 and SD σ_2 reflects a second source of stochastic variability: $A'(\theta,f) = \sum_{x,y} W_r(x,y) C(x,y,\theta,f) + \varepsilon_2$. These responses are then passed through an activation function with gain parameter γ to limit the range of the activation of the representation units: $A(\theta,f) = \begin{cases} \frac{1-e^{-\gamma A'}}{1+e^{-\gamma A'}} A_{\max}, & \text{if} \quad A' \geq 0 \\ 0, \text{ otherwise} \end{cases}$.

The system makes a behavioral *decision* on each trial by combining the sensory information in the activation pattern over the representation units in a single decision unit. The decision variable weights these inputs by the current weights w_i, including a top-down bias factor b with weight w_b : $u = \sum_{i=1}^{65} w_i A(\theta_i, f_i) - w_b b + \varepsilon_d$, and includes random decision noise ε_d (Gaussian with mean 0 and SD σ_d). The "early" activation of the decision unit o' is a sigmoidal function of the weighted activations u with gain γ: $o' = G(u) = \frac{1-e^{-\gamma u}}{1+e^{-\gamma u}} A_{\max}$. A negative o' maps to one response (left), whereas a positive o' maps to the other response (right).

The weight structure is learned through updating on every trial. When feedback is available, the feedback F is used to compute a new, late level o: $o = G(u + w_f F)$ (late). Learning occurs during this late phase (6). The effect of feedback depends upon its weight w_f. The late activation will go to its maximum ($\pm A_{\max} = \pm 1$) with feedback ($F = \pm 1$), and the feedback weight is relatively high, whereas low feedback weight will only slightly shift activation in the direction of the correct response. If no feedback signal is present ($F = 0$), learning operates on the early decision activation ($o = o'$), which is usually in the intermediate range. For all but the lowest accuracy conditions, the weights still move statistically toward a more optimum weight distribution because the activations at decision are correlated with the correct response.

Learning occurs by updating the synaptic connection strengths w_i from sensory representation units i to the decision unit. The change in each weight, w_i, depends on the learning rate, η, the activation at the presynaptic connection, $A(\theta,f)$, the postsynaptic activation compared with its long-term average, $(o - \overline{o})$, and the distance of the weight from the minimum or maximum saturation value, $w_{\min}$ or $w_{\max}$. So the change in weight is: $\Delta w_i = (w_i - w_{\min})[\delta_i]_- + (w_{\max} - w_i)[\delta_i]_+$, where $\delta_i = \eta A(\theta_i f_i)(o - \overline{o})$, and the average of postsynaptic activation is $\overline{o}(t+1) = \rho o(t) + (1-\rho)\overline{o}(t)$. The Hebbian term δ_i tracks systematic stimulus–response correlations rather than mere response bias. The Hebbian learning process is augmented not just by feedback (when it occurs), but also by a criterion control that tracks deviations of the recent response frequencies from 50% or the instructed presentation probabilities in the experiment. Top-down input from a bias unit b feeds into the decision unit with weight w_b. The bias $b(t+1)$ on each successive trial is a running average of responses $r(t)$, weighted exponentially with a time constant of 50 trials ($\rho = 0.02$), $r(t+1) = \rho R(t) + (1-\rho) r(t)$. Prior studies found more pervasive response biases, and correspondingly lower weights on adaptive criterion control, in the absence of feedback (4). Both bias control and feedback augment pure Hebbian reweighting. Bias control tracks the model's responses, whereas feedback tracks external teaching signals. Bias control is more important to learning in the absence of trial-by-trial external feedback (4).

Simulation Methods

The IRT simulation, implemented in MatLab, takes grayscale images, computes activity in location-specific and location-independent representation units, generates a response, and then updates the weights. The simulated experiment exactly replayed the procedure in the human experiment. Filter spacing and orientation and spatial frequency bandwidth parameters, the radius of spatial summation, and several other parameters were set a priori for the location-specific representations (3, 4). The bandwidths of location-specific representations were 30° for orientation and 1 octave in spatial frequency (full bandwidth at half height), every 15° in orientation and every one-half octave in spatial frequency. The spatial frequency and orientation bandwidths of location-independent representations were set at 1.6 times the location-specific values based on preliminary fits to the data.

Typical activation patterns in location-specific and location-independent representations for the four oriented Gabor stimuli with and without external noise are shown in Fig. S2, with the patterns for the to-be-discriminated angles shown in the same panels for ease of comparison. The patterns vary for each sample of external noise and for different Gabor contrasts.

The performance levels and the learning curves from experimental data were fit with a scaling factor (a), two parameters for internal additive noise (σ_1) and for internal multiplicative noise (σ_2), one each for location-specific and one for location-independent representations, a decision noise (σ_d), and learning rates (η). The model was optimized for the experimental data through successive grid search of the parameter values, as described below. Scaling parameters, noise parameters, and learning rate were adjusted to yield the best least-squares fit of the model to the average data. The parameter for the activation function of the decision unit, γ, was set to 3.5 after some initial testing. The predicted performance curves were based on 1,000 iterations of the model experiment.

The implementation of the IRT (*SI Materials and Methods*, *Model Implementation of the Integrated Reweighting Theory*) was fit to the data by grid search. To speed the simulation, large caches of the representation activations for stimulus images with different samples of external noise were computed for many Gabor contrasts. This image processing stage is very time-consuming, requiring up to several days of computation. Using precomputed caches in simulated learning runs avoided duplicating these time-consuming steps. Location-independent activations were computed by passing the images through the representation subsystem with broader bandwidths. The grid search of model parameters was also computationally intensive.

SI Results

Learning over practice is shown in the weight structures for the location-specific and location-independent representations. Fig. S3 displays the changes in weights from the beginning of the experiment (Fig. S3, *Left*) to the end of the first phase of training (Fig. S3, *Center*), and the end of the second phase of training (Fig. S3, *Right*). Learning shows increased weights for relevant spatial frequency and orientation channels and decreased weights for irrelevant channels for both the relevant location-specific units and the location-independent units. It also shows reduced reliance on location-independent representations. Subsequent training on the new task further trains the location-independent weights, and retrains the same location-dependent weights for orientation changes, or newly trains the second location weights for changes of just position, or position and orientation.

Examination of the model predictions over the many parameter combinations tested during optimization of the fit of the IRT model to the data essentially all showed the same ordinal pattern at transfer for the P, O, and OP conditions. So the ordinal predictions for the three switch conditions are quite general. Finding optimal parameters largely served to match overall level of performance and speed of learning. The weight structures showed similar qualitative patterns with learning.

Search grids were programmed that examined factorial combinations of the manipulated parameters, and looked for regions of higher values of fit. Finer grids within these regions were tested. Although we searched in most major regions of parameter combinations through these means, no search grid of this type can be exhaustive, so it is possible that some additional adjustment of

parameters would have further optimized the fit of the model to the data.

For example, increasing noise terms generally increases thresholds, but differentially increasing additive noise increases thresholds in low noise more than that in high noise, whereas increasing multiplicative noise affects thresholds in high noise more. Increasing the scaling factor normally improves performance; however, if the scaling factor is too big, the saturation of activation makes the initial performance worse. Finally, if learning rate is increased, learning is faster. The profile of initial weights can also be manipulated. Such manipulation mainly changes initial performance of a training or transfer session. We finally chose one with moderate tuning to orientation and no tuning to spatial frequency, because it provided a better fit to the data. With a square-wave profile of initial weights—that is, less orientation tuning—the discrepancy of performance between zero and high noise is larger in the simulation than in the behavioral data. Specifically, a broad-tuned profile of initial weights worsens performance in high noise more by channeling more external noise into the decision unit. The amplitude of initial weights can also affect the initial performance.

1. Levitt H (1971) Transformed up-down methods in psychoacoustics. *J Acoust Soc Am* 49(2):2, 467.
2. Dosher BA, Lu ZL (1998) Perceptual learning reflects external noise filtering and internal noise reduction through channel reweighting. *Proc Natl Acad Sci USA* 95(23):13988–13993.
3. Petrov AA, Dosher BA, Lu ZL (2005) The dynamics of perceptual learning: An incremental reweighting model. *Psychol Rev* 112(4):715–743.
4. Petrov AA, Dosher BA, Lu ZL (2006) Perceptual learning without feedback in non-stationary contexts: Data and model. *Vision Res* 46(19):3177–3197.
5. Heeger DJ (1992) Normalization of cell responses in cat striate cortex. *Vis Neurosci* 9(2): 181–197.
6. O'Reilly RC, Munakata Y (2000) *Computational Explorations in Cognitive Neuroscience: Understanding the Mind by Simulating the Brain* (The MIT Press, Cambridge, MA).

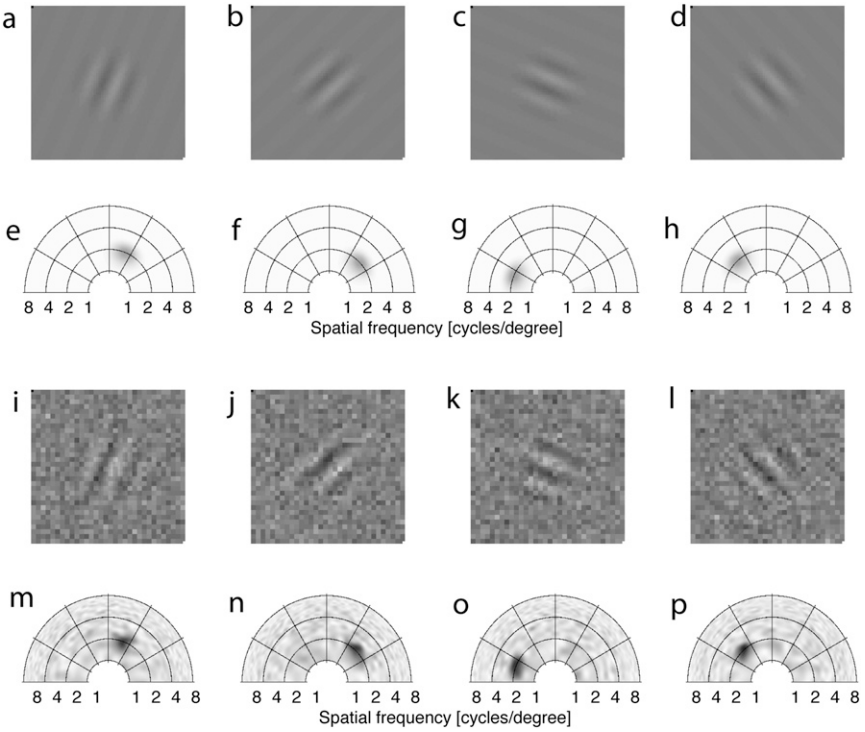

Fig. S1. Sample images and power spectra from the experiment. *A*, *B*, *C*, and *D* show images corresponding with intermediate contrast Gabors without external noise tilted top left and right of the relevant implicit reference angle, with *E*, *F*, *G*, and *H* corresponding power spectra. *I*, *J*, *K*, and *L* show sample high noise tests with high contrast Gabors, with corresponding power spectra in *M*, *N*, *O*, and *P*. Power spectra show orientation on the angular axis (left tilting to right tilting) and spatial frequency (low to high) on the radial axis in polar coordinates.

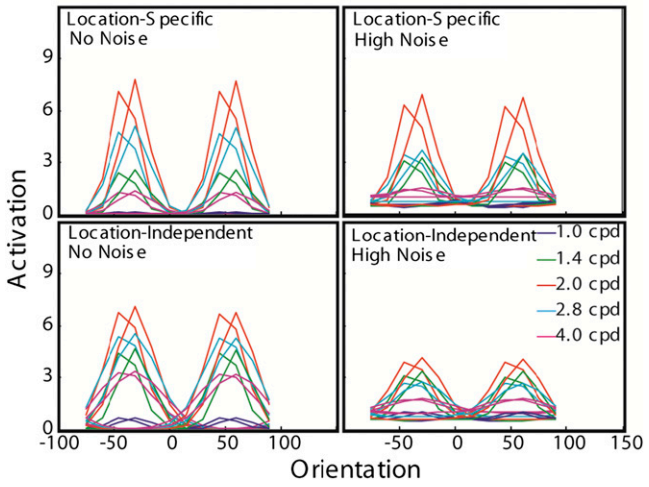

Fig. S2. Average activation profiles in the location-specific and location-independent representations for the four oriented Gabor stimuli with and without external noise, with the two orientations to be discriminated shown in the same panels, with center orientation of filters on the *x*-axis, center spatial frequency of filters indicated by line color, and arbitrary activation units on the *y*-axis. The activation profiles show the largest discriminative information in critical orientations in the spatial frequency band closest to the stimulus, and the two neighboring bands.

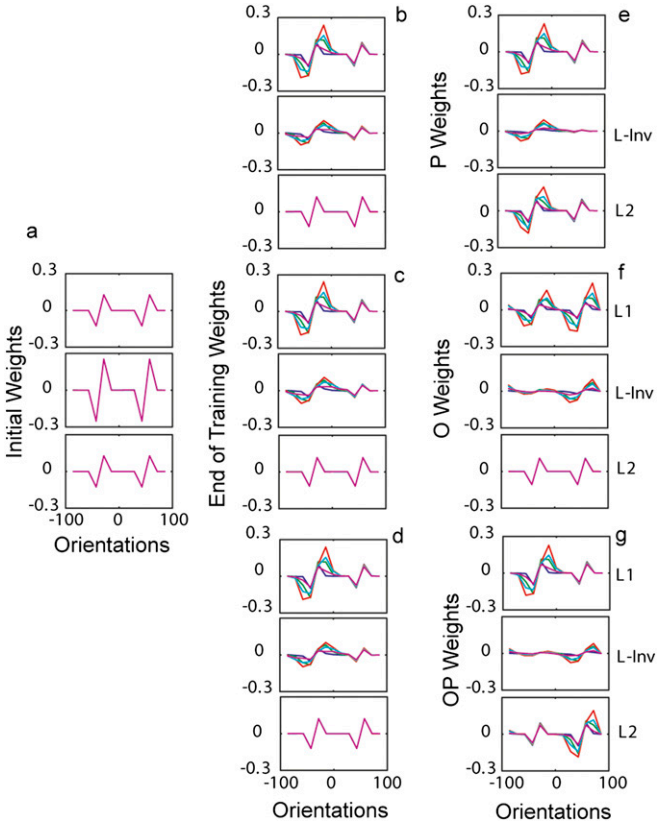

Fig. S3. The connection weights on location-specific and location-independent representations over the course of learning: (*A*) initial weights; (*B–D*) weights at the end of the initial training phase for switches of position P, orientation O, and both orientation and position OP; (*E*) weights after the transfer phase for switches of position P; (*F*) weights after the transfer phase for switches of orientation O; and (*G*) weights after the transfer phase for switches of both position and orientation OP. Activity in spatial frequency bands 2.0 cpd (red), 1.0 cpd (blue), 1.4 cpd (green), 2.8 cpd (cyan), and 4.0 cpd (magenta).