

Roving: The causes of interference and re-enabled learning in multi-task visual training

Barbara Anne Doshier

Cognitive Science Department, University of California,
Irvine, Irvine, CA, USA



Jiajuan Liu

Cognitive Science Department, University of California,
Irvine, Irvine, CA, USA



Wilson Chu

Cognitive Science Department, University of California,
Irvine, Irvine, CA, USA
Department of Psychology, Los Angeles Valley College,
Valley Glen, CA, USA



Zhong-Lin Lu

Division of Arts and Sciences, NYU Shanghai, Shanghai,
China; Center for Neural Sciences and Department of
Psychology, New York University, New York, NY, USA



People routinely perform multiple visual judgments in the real world, yet, intermixing tasks or task variants during training can damage or even prevent learning. This paper explores why. We challenged theories of visual perceptual learning focused on plastic retuning of low-level retinotopic cortical representations by placing different task variants in different retinal locations, and tested theories of perceptual learning through reweighting (changes in readout) by varying task similarity. Discriminating different (but equivalent) and similar orientations in separate retinal locations interfered with learning, whereas training either with identical orientations or sufficiently different ones in different locations released rapid learning. This location crosstalk during learning renders it unlikely that the primary substrate of learning is retuning in early retinotopic visual areas; instead, learning likely involves reweighting from location-independent representations to a decision. We developed an Integrated Reweighting Theory (IRT), which has both V1-like location-specific representations and higher level (V4/IT or higher) location-invariant representations, and learns via reweighting the readout to decision, to predict the order of learning rates in different conditions. This model with suitable parameters successfully fit the behavioral data, as well as some microstructure of learning performance in a new trial-by-trial analysis.

Introduction

As humans, our everyday interactions with a complex world often depend on well-practiced visual judgments. Correspondingly, many visual judgments can be substantially improved with training or practice, sometimes from near chance to excellent performance. Examples include judgments of orientation (Crist, Li, & Gilbert, 2001; Doshier & Lu, 1999), motion direction (Ball & Sekuler, 1987), texture pattern (Karni & Sagi, 1991), and many other tasks. Our interactions with the visual world also often require rapid and flexible intermixture (Kuai, Zhang, Klein, Levi, & Yu, 2005) of visual judgments in everyday behavior. Yet, when several visual tasks or task variants have been intermixed (roved) during training, especially those with similar stimuli and judgments, perceptual learning can sometimes be almost completely disrupted (Sagi, Adini, Tsodyks, & Wilkonsky, 2003; Yu, Klein, & Levi, 2004; Kuai et al., 2005; Parkosadze, Otto, Malania, Kezeli, & Herzog, 2008; Zhang et al., 2008).

These so-called roving effects may reveal important properties about how tasks are learned. Two broad theories about how visual perceptual learning occurs have been proposed (Seitz & Watanabe, 2005). We have schematically illustrated them in Figure 1a. According to sensory retuning, learning primarily reflects retuning neurons in early retinotopic cortical areas, as early as V1 (Karni & Sagi, 1991). In reweighting, learning optimizes the readout (reweighting) of evidence

Citation: Doshier, B. A., Liu, J., Chu, W., & Lu, Z.-L. (2020). Roving: The causes of interference and re-enabled learning in multi-task visual training. *Journal of Vision*, 20(6):9, 1–26, <https://doi.org/10.1167/jov.20.6.9>.






Figure 1. (a) A schematic illustration shows the retuning and reweighting theories of visual perceptual learning, with associations to possible cortical substrates. Perceptual learning may change the tuning of neurons in early retinotopic areas (e.g. V1) (turquoise inset), or change the weights connecting sensory representations at several levels (e.g. V1, V4, IT, or higher) to decision (red oval). (b). A diagram showing weights fighting following two similar orientations with opposite responses (CW and CCW). Following the left image, a “CW” response is expected and weights move toward positive values (red); following the right image, a “CCW” response is expected and weights move toward negative values (blue). The overlapping weights for these two orientations, hence fail to improve due to conflicting updates (gray dotted lines).

represented at one or more levels of the visual hierarchy (Doshier & Lu, 1998; Doshier & Lu, 1999; Doshier, Jeter, Liu, & Lu, 2013). One consequence of retuning theory, which has been a dominant proposal in the field, is that learning tasks in separate retinal locations should be largely independent, because they involve separate retinotopic neural populations (Karni & Sagi, 1991). Under reweighting theory, on the other hand, tasks trained in different retinal locations may interact during learning through shared higher-level representations, and (as we will see) this will be especially consequential for tasks in which similar stimuli require different responses. In a neural network model of learning through reweighting, training intermixed tasks can exhibit different levels of interference (Grossberg, 1987;

McCloskey & Cohen, 1989). Of course, learning might, in some circumstances, occur through both retuning and reweighting, as suggested in several integrated reviews of perceptual learning (Seitz & Watanabe, 2005).

Roving effects in perceptual learning

The impact of task intermixture or roving during visual perceptual learning has been documented in a number of visual tasks in experiments that train in a single retinal location. Perhaps the most famous example showed that learning was largely disrupted in two-interval contrast increment detection when base

contrast varied from trial to trial, whereas contrast increment detection otherwise showed robust learning with a fixed base contrast (Adini, Sagi, & Tsodyks, 2002; Kuai et al., 2005; Yu et al., 2004). In bisection tasks, learning has been shown to be disrupted, or at least very slow, when the distance between reference lines or dots was intermixed, or with other certain forms of stimulus variation, but not all (Aberg & Herzog, 2009; Parkosadze et al., 2008). In some studies though, contrast increment learning was “re-enabled” when the base contrasts were cycled in a fixed temporal order (Cong & Zhang, 2014; Kuai et al., 2005; Zhang et al., 2008). Similar damaging effects of roving have also been found in the auditory domain. In one example, learning two interval frequency discrimination was robust with a fixed frequency standard, showed very slow or no learning for roved standards within an auditory band, and for wide roving (Amitay, Hawkey, & Moore, 2005); in another, learning was present in an auditory temporal-interval discrimination task with sequential training of two base intervals but not intermixed training (Banai, Ortiz, Oppenheimer, & Wright, 2010).

The theoretical explanations for these roving effects have attributed them to task variation and some form of interference. First, roving damages learning when the stimuli and tasks are distinct but similar (Tartaglia, Aberg, & Herzog, 2009a). Such combinations would recruit overlapping neuron populations in which training could have interfering effects. As for the form of interference, some explanations have focused on recurrent processing stages, a potential role of variations in position due to fixation fluctuation and nonlinear processes in registration (Otto, Herzog, Fahle, & Zhaoping, 2006; Zhaoping, Herzog, & Dayan, 2003). Other explanations recruited ideas about failure of developing a stable memory trace (Yu et al., 2004), or memory consolidation (Seitz et al., 2005). The theoretical discussion closest to our current modeling of these roving effects (by Tartaglia, Herzog, and colleagues) analyzes stimulus roving effects in the context of network learning models and their selective susceptibility to negative effects, depending on the nature of the learning algorithm and the overlap of representations (Tartaglia, Aberg, & Herzog, 2009b). A computational model can formalize these ideas and be used to test whether they can account for the behavioral learning data.

Modeling approach

Understanding why and under what circumstances learning is disrupted for intermixed tasks has broader implications for the broad theories of brain plasticity (Herzog, Aberg, Frémaux, Gerstner, & Sprekeler, 2012) and may, in turn, be relevant to the design

of real-world training protocols (Kuai et al., 2005; Tartaglia et al., 2009b). In this study, we investigated how training tasks in different retinal locations, together with manipulations of stimulus dissimilarity, impact intermixed learning. From a modeling perspective, we see the disruption of learning when stimuli/tasks are intermixed during training as related to the concept of catastrophic interference in neural networks (Grossberg, 1987). Figure 1b schematically illustrates why stimulus similarity might be critical for interference due to roved tasks or stimuli—essentially when very similar stimuli (with close stimulus representations) require distinct responses.

In order to investigate this theoretical hypothesis, we developed and tested the predictions of a computational model of visual perceptual learning, the integrated reweighting theory (IRT) (Doshier et al., 2013; Liu, Doshier, & Lu, 2015), which accounts for learning by reweighting—improving the readout—from stable sensory evidence using a hybrid (augmented Hebbian) learning rule. It learns by reweighting the evidence in visual representations at several levels to make perceptual decisions. Stimulus images are first processed through a front-end, which simulates early visual cortical responses in orientation and spatial frequency tuned representations (both location-specific and location-invariant); and then weights connecting these representations to a decision unit are updated (reweighted) using augmented Hebbian learning rules on each trial, simulating the outcomes in the actual experiments (Petrov, Doshier, & Lu, 2005; Petrov, Doshier, & Lu, 2006). (See Methods, Simulation Methods, for details of the computational model.) The schematic principle of interference in decision weights for near stimuli is illustrated in Figure 1b, and this idea was implemented computationally in the IRT.

As outlined previously, simple retuning theories of visual perceptual learning, which locate learning primarily in the plasticity of neurons in early retinotopic visual areas, predict little interaction between tasks learned in separate retinal locations—that is, they predict that separating the tasks by location should eliminate effects of roving because the neural populations involved are distinct. In the reweighting model, by contrast, task variants learned in different locations in intermixed training protocols reweights both higher-level location-invariant representations and lower-level location-specific ones. Although roving interference would change the weights on both representation levels if the task variants were trained in the same location, they occur through the weights on location-invariant representations when different task variants are trained in different locations. The interactions between locations, and different amounts of roving interference, were exercised in four groups each of which experienced different task mixtures.

Experimental approach

We compared perceptual learning in four groups of observers who practiced judging the orientations of Gabor patterns ($\theta \pm 12^\circ$) in four different peripheral locations, but with different combinations of reference angles (θ) across the locations in each group. On each trial, observers judged whether the stimulus was rotated clockwise or counter clockwise of the reference angle for the location indicated by a pre-cue. External noise was added on half the trials to constrain the estimates of internal noises in the IRT fits (Doshier et al., 2013; Lu & Doshier, 2008). Figure 2 shows the stimulus layout, sample stimuli with and without external noise, illustrations of the four intermixture (roving) conditions, and a typical trial sequence. The four groups were as follows: in the *All* condition, training intermixed four different reference angles spaced in orientation, one per location. This is a perfect setup for substantial roving interference in learning—if learning in different locations interact—because opposite responses are required for every set of adjacent Gabor stimuli in representations of orientations (see Figure 2b). Two other groups intermixed training of two tasks, each occurring in two locations, either with more similar reference angles (*Near*) or quite dissimilar reference angles (*Far*). Finally, a no-roving condition (*Single*) trained the same reference in all four locations. (See Methods, Simulation Methods, for details.)

The IRT model predicted the observed differences in learning rates for the four groups, on the basis of the interaction of learning in the four locations and the similarity or dissimilarity of the set of tasks being learned. As we will see, it provided an excellent account of the data at both the session level and the microstructure level of trial-by-trial learning.

Methods

Behavioral experiment

Observers

Observers, with normal to corrected-to-normal vision, provided written consent under a protocol approved by the Institutional Review Board of the University of California Irvine. Forty-eight observers were randomly assigned, 12 each, to one of four experimental groups. Each observer participated in 7,680 experimental trials over 8 sessions on different days, usually within a two-week period. Other observers were excluded from analysis if their thresholds were not measurable because they were at ceiling (100% contrast) for more than one session, especially but not exclusively

in the high external noise condition; this occurred more frequently in more difficult *All* ($n = 4$) and *Near* ($n = 4$) conditions, compared to the easier *Far* ($n = 1$) and *Single* ($n = 1$) conditions. (As a consequence, if anything, the results may slightly underestimate the learning rate differences between groups.) Each observer who completed the study performed 7,680 trials over eight training sessions, 960 per session, or 368,640 trials over all observers.

Stimuli and apparatus

A Gabor (windowed sine wave) pattern was presented on each trial at one of the four corners around fixation; its orientation was chosen at random (“clockwise” [CW] or “counter clockwise” [CCW] of the reference angle for each location) and was presented with or without Gaussian noise. The Gabor pattern, defined in a 64×64 pixel patch, is described by: $I(x, y) = I_0 \left(1.0 \pm c \sin(2\pi f(x \sin(\theta) \pm y \cos(\theta))) \exp\left(-\frac{x^2+y^2}{2\sigma^2}\right) \right)$, with $\theta =$ reference angle $\pm 12^\circ$, spatial frequency $f = 1.33$ cpd, and SD of the Gaussian envelope $\sigma = 0.5$ degrees, maximum contrast c , and I_0 is the mid-gray background luminance. Each external noise image, newly generated for each trial, was composed of 2×2 pixel noise elements with contrasts randomly chosen from a Gaussian distribution with mean value θ and SD 0.25. External noise images and signal Gabor images (NNSNN) were displayed sequentially at the frame rate (see Procedure). The 64×64 pixel images subtended $3^\circ \times 3^\circ$ visual angle, located at 5.67° eccentricity, at a viewing distance of 72 cm. Stimuli were generated in MATLAB with PsychToolbox on a Macintosh G4 computer using the internal 10-bit video card, refresh rate of 67 Hz, resolution of 640×480 pixels, and displayed on a 19 in. Viewsonic color monitor in pseudo-monochrome. A lookup table, estimated by a visual calibration procedure (Lu & Doshier, 2013) and validated by photometric measurement, linearized the luminance range into 127 levels from 1 cd/m^2 to 67 cd/m^2 ; the mid-grey background luminance was 34 cd/m^2 . The observer’s head was stabilized using a chin rest.

Design

Observers discriminated the orientation of a Gabor patch tilted $\pm 12^\circ$ (CW or CWW, from a reference angle) in one of the four locations indicated by a pre-cue and a response post-cue (arrows). There were four roving (intermixture) conditions: In the *All* condition, each of the four locations used a different reference angle (i.e. -67.5° , -22.5° , 22.5° , or 67.5° from vertical for the lower left, upper left, upper right, and lower right positions). In the *Near* condition, two closer reference

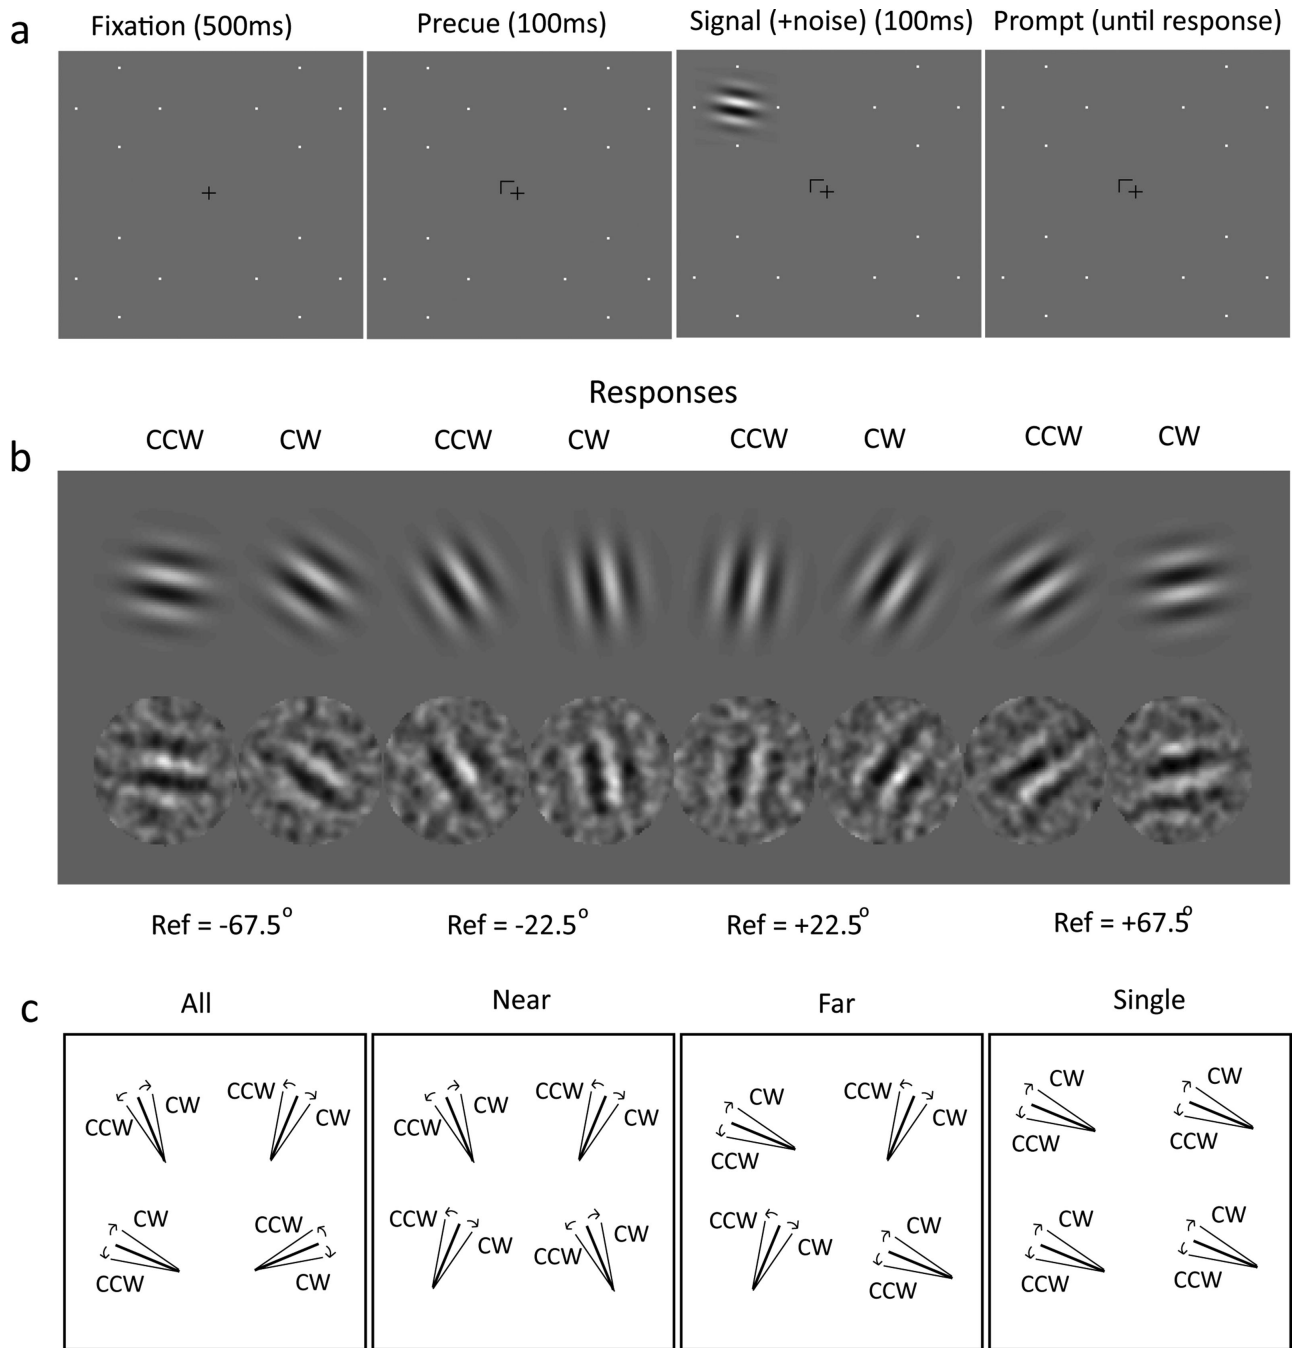


Figure 2. Sample stimuli and responses, and the training task mixtures of orientation judgments in separate locations (“clockwise” (CW), “counterclockwise” (CCW) by $\pm 12^\circ$ of a reference angle). (a) Trial sequence of a fixation (500 ms), a precue (100 ms) marking the location to judge, the stimulus display (100 ms), and a response cue. Adaptive methods estimated contrast thresholds at 75% correct. (b) Sample stimuli with and without external noise (“snow”), along with assigned CW or CWW responses. (c) Illustrations of possible task mixtures in the four roving groups: *All* interleaved different reference angles in each location (-67.5° , -22.5° , 22.5° , or 67.5° from vertical); *Near* interleaved two similar tasks in two locations each (e.g. 22.5° or 67.5°); *Far* interleaved two dissimilar tasks at two locations each (e.g. -22.5° or 67.5°); *Single* trained one reference angle (e.g. -67.5°) in four locations (no roving); showing here only the reference angles, not the actual stimuli. $N = 12$ observers per group performed 7,680 trials each in 8 sessions.

angles were used in two diagonal positions (e.g. -22.5° or 22.5°). In the *Far* condition, two dissimilar reference angles were used in two diagonal positions (e.g. -67.5° or 22.5°). In the *Single* condition, the same reference angle occurred in all locations. Zero and high external noise test conditions were intermixed. There were 8 sessions of 960 trials per session. Adaptive methods (see below) were used to measure contrast thresholds at 75% correct for each location and external noise condition separately (120 trials each within a session).

Procedure

Observers were instructed in the task, shown printed examples of the stimuli, and then participated in a small number of practice trials prior to collecting the experimental data. Each trial of the experiment started with a central fixation mark and four sets of location markers; 500 ms later, the stimulus sequence (external Gaussian noise or blank frames, signal, external Gaussian noise or blank frames) appeared for 2 refresh counts per frame, with a central pre-cue arrow appearing 100 ms prior to the signal Gabor frame. The contrast of the stimulus in a trial was determined by the adaptive procedure. Observers pressed the “j” key for CW or the “f” key for CCW. A feedback tone followed each correct response. Each session included 960 trials, or 120 trials in each of the four locations for each external noise level.

Adaptive threshold measurement

Thresholds were measured with the accelerated stochastic approximation algorithm (Kesten, 1958). The Gabor (signal) contrast on each trial was selected to track a target performance of 75% correct. In the first two trials, contrasts follow the stochastic approximation procedure (Robbins & Monro, 1951): $X_{n+1} = X_n - \frac{s}{n}(Z_n - \phi)$, where n is the trial number, X_n is the stimulus contrast in trial n , $Z_n = 0$ or 1 is the response accuracy in trial n , X_{n+1} is the contrast for the next trial, and s is the pre-chosen step size at the beginning of the trial sequence. From the third trial on, the sequence is “accelerated”: $X_{n+1} = X_n - \frac{s}{2+m_{shift}}(Z_n - \phi)$, where m_{shift} is the number of shifts in response category (from correct response to incorrect response and vice versa). In our application, the method was modified such that while $m_{shift} = 0$ the increased contrast on an error is capped at $0.125s$. See (Treutwein, 1995) for a discussion of this adaptive method, and (Lu & Doshier, 2013) for an analysis of the convergence properties and guidelines for step size and starting values.

Bootstrap methods

Error bars were estimated by bootstrapping. In most cases, this involved (i) generating new pseudo-samples of observers (usually 1,000) by sampling with replacement from the set of actual observers, and then (ii) statistically processing these with the same methods used in the corresponding analysis of human data. For example, session thresholds (symbols in Figure 3) were computed by averaging the contrast thresholds from 12 observers in each condition; and for each observer, the threshold was first averaged from contrasts for the final 30 trials in each testing location, and then averaged these over location. The SD of the mean thresholds of many pseudo-sampled sets of 12 observers per condition was processed to estimate the standard error of these mean values, which gives an estimate of error based on the observed population variability. Similar bootstrapping methods were used as the basis of error estimates for the parameters of power function models fit to the learning curves.

Fitting power functions

Power functions were fit to the contrast threshold learning data (curves in Figure 3a) to estimate the learning rates. The learning curves were fit by power function improvements (Doshier & Lu, 2007; Heathcote, Brown, & Mewhort, 2000): $C(t) = \lambda(t + 1)^{-\beta} + \alpha$, with initial threshold of $\lambda + \alpha$, asymptotic threshold of α , learning rate β , and training block t . The curves for the four roving conditions were tested for significant differences with a lattice of nested F -tests, each of which compares a restricted model to a fuller model of which it is a proper subset. For example, if roving conditions actually differ in learning rate (or any other parameter), constraining the model system to equate that parameter will significantly reduce the quality of fit. The proportion of variance accounted for by a model is r^2 :

$$r^2 = 1.0 - \frac{\sum [x^{theory} - x^{observed}]^2}{\sum [x^{observed} - \bar{x}]^2}.$$

The Σ is over all N observations and \bar{x} is the mean of the observed values. F -tests for nested models compared the fit of the fuller and reduced models:

$F(df_1, df_2) = \frac{(r_{full}^2 - r_{reduced}^2)/df_1}{(1 - r_{full}^2)/df_2}$, where $df_1 = k_{full} - k_{reduced}$, and $df_2 = N - k_{full} - 1$. The k 's are the number of model parameters. The F -test computes the ratio of the improvement in error variance for each additional parameter in the *fuller* model to the (average) error variance per degree of freedom.

The variability in the parameter estimated was computed via bootstrapping as described above.

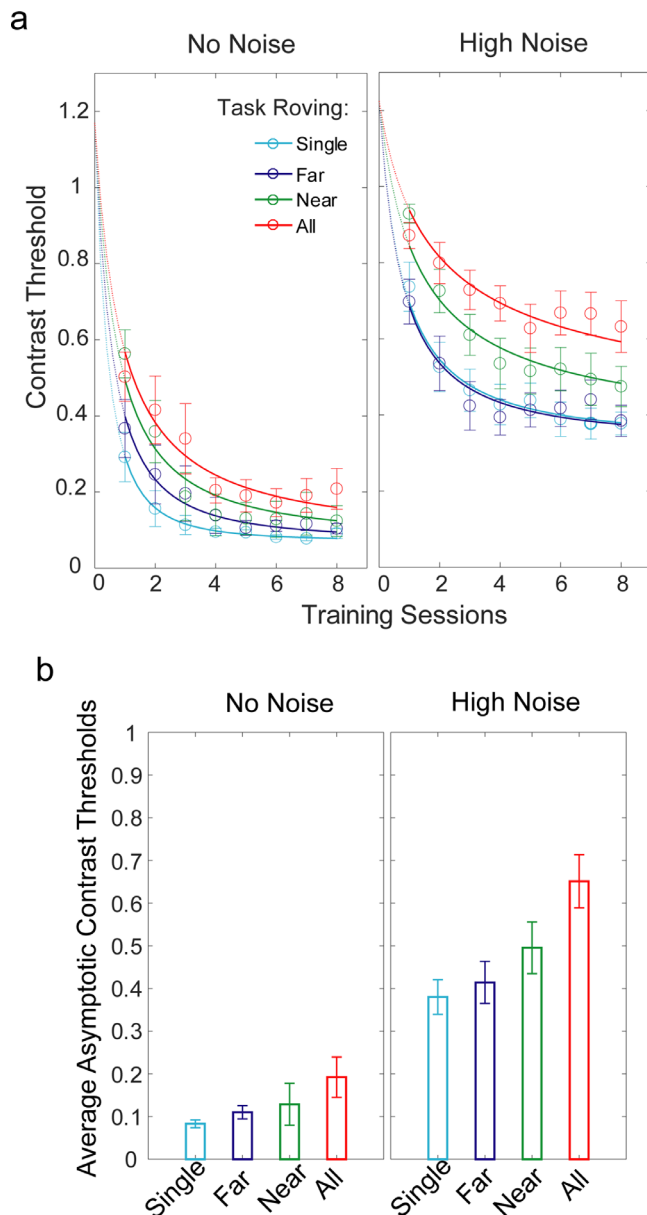


Figure 3. Contrast thresholds as a function of training session for the four roving groups: *All*, *Near*, *Far*, and *Single*, displayed separately for low and high noise test trials. (a) Contrast thresholds at 75% correct from adaptive staircases, averaged over observers. Error bars were bootstrapped ($n = 1000$) from the behavioral data. Smooth curves are the power fits with different learning rates for each group. (b) Post-training contrast thresholds for the four roving groups. Error bars are standard error of the mean.

A set of 1,000 samples of 12 observers per group, drawn randomly with replacement from the data of actual observers, were created, and then power function models were fit to each of these data sets. This was the basis of the estimated mean and SD of the estimated parameter values. Because the SDs were inflated by correlated structures between power

function parameters, we also tabulated the frequency of the ordinal pattern between different rate parameters in the 1,000 fits to bootstrapped data (presented in [Appendix A](#)).

Simulation methods

Integrated reweighting theory

The integrated reweighting theory (IRT) ([Doshier et al., 2013](#)) was implemented in Matlab (The Mathworks, Inc., Santa Clara, CA). It includes a representation module, a decision module, and a learning module. The simulation model is tested by exactly reprising the experimental protocol (i.e. same number of trials, randomization of conditions, etc.), taking stimulus images as input, producing responses as output on each trial, and using exactly the same data analysis as the behavioural experiment. The descriptions of the modules below are similar to those of previous applications ([Doshier et al., 2013](#); [Petrov et al., 2005](#)).

The representation module, inspired by units in early visual cortex, computes the activities in location-specific and location-invariant representations from stimulus images. Signal and external noise images are summed in the model to represent temporal integration by the visual system. This implementation used four sets of location-specific representations and one set of location-invariant (location-independent) representations that responds to inputs from all four locations. The spacing, orientation, and spatial frequency bandwidth parameters, and the spatial summation radius, were all set a priori from earlier applications. There were 5 spatial frequency bands (every $\frac{1}{2}$ octave) centered at (0.7, 1, 1.4, 2, and 2.8 cycle/degree), 12 orientation bands (every 15°) centered at ($0^\circ, \pm 15^\circ, \pm 30^\circ, \pm 45^\circ, \pm 60^\circ, \pm 75^\circ, \text{ and } +90^\circ$ [= -90°]), and four spatial phases ($0^\circ, 90^\circ, 180^\circ, \text{ and } 270^\circ$). The spatial frequency tuning bandwidth was set at $h_f = 1$ octave and the bandwidth of the orientation tuning was set at $h_\theta = 30^\circ$ (half-amplitude full-bandwidth), based on estimated cellular tuning bandwidths in primary visual cortex. The location-invariant representations were more broadly tuned, with bandwidths of $1.6 \cdot$ those of the location-specific units, and also had more internal noise ([Doshier et al., 2013](#)). The descriptions of the representation, decision, and learning modules are similar to those in earlier treatments ([Liu et al., 2015](#); [Petrov et al., 2005](#)).

The representation module computes the activation values $A(\theta, f)$ of the orientation- and frequency-selective representation units, whether location-specific or location-invariant, in response to the stimulus image(s). This measures the normalized spectral energy in those channels. Sets of retinotopic phase-sensitive maps $S(x, y, \theta, f, \phi)$ are applied to the input image $I(x, y)$:

$S(x, y, \theta, f, \phi) = [RF_{\theta, f, \phi}(x, y) \otimes I(x, y)]$, for spatial frequency f , orientation θ , and spatial phase ϕ . The input (stimulus) image $I(x, y)$ is convolved with the filter for each spatial-frequency/orientation unit by fast Fourier transform, followed by half-squaring rectification, to produce phase-sensitive activation maps analogous to “simple cells.” These are pooled over spatial phase: $E(x, y, \theta, f) = \sum S(x, y, \theta, f, \phi) + \varepsilon_1$

and subjected to inhibitory normalization: (Heeger, 1992) $C(x, y, \theta, f) = \frac{aE(x, y, \theta, f)}{k + N(f)}$. The noise ε_1 is

Gaussian-distributed internal noise with mean 0 and standard deviation σ_1 . The normalization pool $N(f)$ is independent of orientation and only modestly tuned for spatial frequency, as suggested by the physiology. The parameter a is a scaling factor, and k is the saturation constant to prevent division by zero at very low contrasts. For this behavioral task where the observer judges orientation, activations were pooled over spatial phase, and then spatially pooled with a Gaussian kernel of radius W_r around the target Gabor. Another Gaussian-distributed noise of mean 0 and SD σ_2 introduces another source of stochastic variability:

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

activations of the representation units are limited within a range by a nonlinear function with gain parameter γ :

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

The decision module takes the weighted sum of activity in the representation units and a bias unit as inputs and generates a predicted response on each trial. The decision variable is: $u = \sum_{i=1}^{60} w_i A(\theta_i, f_i) - w_b b + \varepsilon_d$. The w_i s are the current weights on representation units, b is a bias term integrated with weight w_b , and ε_d (Gaussian, mean 0, SD σ_d) is decision noise. A sigmoidal function with gain γ transforms it into an “early” post-synaptic decision activation o' : $o' = G(u) = \frac{1 - e^{-\gamma u}}{1 + e^{-\gamma u}} A_{max}$, with negative and positive values mapping to a CCW or CW responses, respectively.

The learning module updates the weights from the representation units to the decision unit on every trial using augmented Hebbian learning. Feedback ($F = \pm 1$), when available, moves the “late” post-synaptic activation in the decision unit o toward the correct response: $o = G(u + w_f F)$. If the feedback weight w_f is high, activation of the decision unit approaches the correct positive or negative maximum ($\pm A_{max} = \pm 1$). In the absence of feedback ($F = 0$), learning operates on the early decision activation ($o = o'$), which is often intermediate. Learning occurs by changing the connection strengths w_i from sensory representation units i to the decision unit on each trial. The weight changes depend on the activation at the pre-synaptic connection, $A(\theta, f)$, the post-synaptic

activation compared to its long-term average, ($o - \bar{o}$), the distance of the weight from the minimum or maximum saturation value (w_{min} or w_{max}) and the (system) learning rate (η). So, the change in weight is: $\Delta w_i = (w_i - w_{min})[\delta_i]_- + (w_{max} - w_i)[\delta_i]_+$, where $\delta_i = \eta A(\theta_i, f_i)(o - \bar{o})$, and the time-weighted average of post-synaptic activation is $\bar{o}(t + 1) = \rho o(t) + (1 - \rho)\bar{o}(t)$. This Hebbian learning rule is augmented both by feedback (when it occurs in the behavioral experiment) and by information in the bias control unit b that contributes to the early decision activation. The bias is a time-weighted average of responses $r(t)$, weighted exponentially with a time constant $\rho = 0.02$ (about 50 trials), $r(t + 1) = \rho R(t) + (1 - \rho)r(t)$. The bias serves as a counter to deviations from 50% to 50% response histories (assuming symmetric experimental designs). Bias control tracks the observer’s responses, while feedback tracks the external teaching signals.

Fitting the simulated model to data

Predictions of the IRT were generally based on 1,000 simulated repetitions of the experiment, yielding estimates of the mean and SDs of simulated thresholds. Parameters of the representation module were fixed a priori from the physiology or from prior implementations of an augmented Hebbian reweighting model (the AHRM) and the IRT (Doshier et al., 2013; Petrov et al., 2005). The parameter for the nonlinear activation of the decision unit, γ , was set to 3.5 based on previous estimates. The parameters varied to achieve the best fit of the IRT to the behavioral data of the current experiment were: internal additive noise (σ_1) and internal multiplicative noise (σ_2) for location-specific and for location-independent representations, a decision noise (σ_d), and learning rate (η). A scaling factor (a) matched initial performance of the different randomly assigned observer groups, and could differ slightly between groups if necessary.

The best parameter values to fit the behavioral data were found using successive grid search, followed by more detailed searches in identified regions of the parameter space. These are time consuming due to the computational demands of processing many different Gabor plus noise images through the representation module to create a large image activation cache, and many runs of the simulation for each parameter combination. Several key free parameters, which are listed above, were varied and the best least-squares fit of the model to the average data among those we tested was selected. The quality of the fit was measured by the r^2 between the mean contrast thresholds from the simulation and the average contrast thresholds in the experiment. The fit was also assessed using the statistic Kendall’s τ that measures the consistency in the ordinal

predictions between the model and the data. In this application, Kendall's τ was lower than r^2 because some conditions led to similar predicted outcomes, and so could easily trade ordinal positions in the data (e.g. the Far and Single groups).

Results

Perceptual learning in separate locations interacts, and depends on the task combination

Perceptual learning occurred at quite different learning rates in the four training groups, as seen in the average contrast threshold learning curves, graphed separately for zero and high external noise tests (Figure 3a). Contrast thresholds were measured per session, as is typical of many studies of perceptual learning. The SDs of the average thresholds (error bars) were estimated by bootstrap methods (see Methods, Behavioral experiment). The ultimate differences in learning between groups after 8 sessions were quite substantial (Figure 3b).

The effects of the training (intermixture) group (*All*, *Near*, *Far*, and *Single*) were tested using analysis of variance on the contrast thresholds, with external noise (zero and high) and training block as within-observer factors, roving group as a between-observer factor, and observers as the random factor ($\alpha = 0.05$). Higher contrast was, of course, required to accurately judge the stimuli embedded in external noise ($F(1, 44) = 418.44$; $p < 0.0001$; $\eta_p^2 = 0.905$). Learning reduced contrast thresholds over training sessions ($F(7, 308) = 52.85$; $p < 0.0001$; $\eta_p^2 = 0.5457$). Of central importance for this study, the four training groups showed different amounts of learning ($F(3, 44) = 5.025$; $p < 0.005$; $\eta_p^2 = 0.255$), showing large differences in learning after several sessions. There was also an interaction between external noise and training group ($F(3, 44) = 2.750$; $p < 0.05$; $\eta_p^2 = 0.158$); and among training condition, external noise, and block ($F(21, 308) = 1.48$; $p \approx 0.08$; $\eta_p^2 = 0.092$). The methods of Masson (2011) were used to compute the Bayes information criterion probabilities ($p_{\text{BIC}}(H_1|D)$) (essentially a transformation of $(1 - \eta_p^2)$): $p_{\text{BIC}}(H_1|D) > 0.999$ for the effects of training (roving) group, blocks of training, external noise, and the group by external noise interaction; $p_{\text{BIC}}(H_1|D) < 0.001$ for the interaction among training condition, external noise, and session.

Post hoc tests indicated that the differences between the four intermixture (roving) groups were significant (all $p < 0.0001$ except *All* versus *Near*, $p < 0.01$), although the *Far* and *Single* conditions were statistically equivalent (NS) (Bonferroni correction $\alpha = 0.008$). See Appendix A, 1 for equivalent results of analyses

separately in zero and high external noise conditions. As described later, the IRT reweighting model predicts this same order of learning rates: $All < Near < Far \approx Single$, from least to most.

The rates of learning were estimated from power function learning curves fitted to the average contrast threshold data (smooth curves in Figure 3) (Doshier & Lu, 2007; Heathcote, Brown, & Mewhort, 2000). These functions are described by: $C(t) = \lambda(t + 1)^{-\beta} + \alpha$ where $c(t)$ is the threshold in session t , $\lambda + \alpha$ is the initial threshold, α is the asymptotic value late in training, and β is the learning rate. Power functions provide a good description of average contrast threshold learning functions (Doshier & Lu, 2007). In this case, the training sessions started at $t = 1$ because the thresholds reflect session-end performance, and we tested the equality of pre-training thresholds at $t = 0$ in additional nested model tests (see Appendix A, 1).

The four learning functions differed only in the rate of learning: the $(1\lambda - 4\beta - 1\alpha)$ model provided an excellent fit with rates of: β s = 1.1478, 1.3763, 1.7446, and 2.3077 ($\lambda = 1.0984$, and $\alpha = 0.0713$) for zero noise ($r^2 = 0.9411$) and β s = 0.5538, 0.7936, 1.3242, and 1.2836 ($\lambda = 0.8979$, and $\alpha = 0.3262$) for high external noise ($r^2 = 0.9554$), listed for *All*, *Near*, *Far*, and *Single*, from slower to faster. A lattice of subcase models and nested significance tests rejected more complicated models (see the discussion in Appendix A, 2, and Tables A.1 and A.2). The SDs of the estimated parameters, computed using bootstrap methods, are listed in Table A.3. The parameter SDs are relatively large (reflecting slight threshold level differences between observer groups and parameter correlations; added variance from parameter correlations was partially discounted in SDs of normalized rates). Despite this, the ordinal consistency of the four rates from the bootstrapped methods, which is perhaps more meaningful, was very high. For example, in zero noise, β_{All} was slower than β_{Single} , slower than β_{Far} , and slower than β_{Near} in 998, 949, and 786 fits, respectively, out of 1,000 fits to resampled data sets; and in high noise, β_{All} was slower than β_{Single} in 1,000 fits, slower than β_{Far} in 1,000 fits, and slower than β_{Near} in 950 fits out of 1,000 fits to resampled data sets (ordinal statistics are also listed in Table A.3). Consistent with the ANOVA tests, in high noise, β_{Far} is slower than β_{Single} only in 469 out of 1,000 fits—they are not significantly different from each other.

Since observers were assigned to groups randomly, the threshold performance before training was expected to be equivalent in the four groups. Consistent with this, the difference in contrasts among groups was insignificant at the beginning of the first session ($p > 0.05$), and steadily increased throughout the session ($\eta_p^2 = 0.015$ for trials 40 to 80, $p_{\text{BIC}}(H_1|D) < 0.05$, whereas $\eta_p^2 = 0.221$ for the last 40 trials, $p_{\text{BIC}}(H_1|D) >$

0.999; differences in contrast thresholds between groups emerged as early as 200 to 300 trials of training in the first session p values < 0.05 , uncorrected for multiple tests; see [Appendix A, 3](#) for more detailed analysis for each noise level). Additionally, the contrast thresholds of a subset of observers in the *All* group showed a deterioration in the last few sessions that we believe may have reflected a lack of motivation in this more challenging roving condition.

Differential learning predicted by the IRT

Tasks trained in different retinal locations interacted strongly during learning, which rules out simple forms of the sensory retuning theory, in which learning primarily reflects retuning of separate neural populations in retinotopic representations in early visual cortex. In contrast, the reweighting theory of perceptual learning, implemented computationally in the IRT ([Doshier et al., 2013](#)), predicts different empirical learning rates for the four different task intermixtures, with the same order as seen in the data: $All < Near < Far \approx Single$, least to most. These different learning rates predicted by the model are induced solely through differential training in the different roving conditions.

This computational IRT model processes stimulus images through a visual front end, including normalization and gain control ([Heeger, 1992](#)), producing activities in spatial-frequency and orientation-tuned units that approximately mimic early visual cortical responses. It then weights this evidence (activation scores) in location-specific and location-invariant sensory representations to make a decision (i.e. “counter-clockwise” or “clockwise”). Then, augmented Hebbian learning (with feedback and response bias inputs) changes the readout (e.g. weights on sensory evidence) with experience on every trial. It recapitulates the behavioral experiment exactly in the simulation (e.g. stimuli, training sequence, number of training trials, randomization, and adaptive algorithm to adjust contrast). That is, it takes stimulus images as input, produce responses, and learns by adjusting weights on each trial. The resulting responses are analysed as in the experiment. (See Methods, Simulation Methods, for model specification.)

The different training experiences in the four groups interact differently in neural network weight space. Weights connecting a location-invariant layer to the decision unit are affected by training in all four locations in addition to the location-specific layer trained by tasks in each location separately. The success of learning at the higher level and the amount of training also influences the simultaneously learned weights connecting each location-specific representation to decision. The four groups experience different levels

of interference in the connection weights from the location-invariant representations to decision because of the different task combinations. For an intuition, consider the stimulus-response mappings in the four roving conditions: Any test stimulus in one location of the *All* condition that maps to a “clockwise” response has two stimuli adjacent in orientation (rotation) space that maps to the competing “counter-clockwise” response, leading to lots of interference; each stimulus in the *Near* condition has a competitive response mapping with a stimulus on one side; the stimuli in the *Far* condition are widely separated in orientation space and so have no near stimuli with competing responses; while in the *Single* condition the same stimuli and responses are trained in all locations. Although the focus here is on the weights from the location-invariant representations to decision because the experiment trains different task variants in different locations, task roving in a single location would lead to interference in the weights from both location-invariant and the location-specific levels of representation. These intuitions were validated by computing (nearly) optimal weights for the four training groups by simulating a very large number of training trials in zero noise. Despite different speeds of weight development during early and middle stages of training, with corresponding predicted differences in contrast thresholds, the four conditions nearly converge after very extensive training in zero external noise (this would require an amount of training far far beyond the thousands of trials for each observer in the behavioral experiment).

The IRT provides a theoretical and intuitive understanding of the nature of the interaction between tasks trained in these intermixed training paradigms. It predicts the ordinal properties of the empirical learning rates in the four intermixed training groups, and this is true for many parameter sets—although fitting the data quantitatively requires optimizing parameter values.

The best-fitting parameters were estimated through modified grid search (see Methods, Simulation). The average predicted contrast thresholds (line), and ± 1 SD (shaded areas) are shown in [Figure 4](#), along with the behavioral contrast threshold data (symbols). (Error bars for the behavioural data are in [Figure 3](#)). The parameters free to vary included: internal multiplicative noise σ_1 , additive noise σ_2 , decision noise σ_d , scaling factor a , the weight on feedback w_f , and model learning rate η . Spatial frequency and orientation bandwidths of the sensory representations were selected a priori based on the physiology and some nonlinearity parameters were set from prior model applications, with the orientation and spatial frequency bandwidths of the location-invariant representations slightly broader than the location-specific representations. The location-invariant internal noises were set at twice the location-specific internal noises, based on prior applications of the model. With the model constrained




Figure 4. Predictions of the IRT model of perceptual learning are shown for the four training groups, which differ in the intermixture of trained tasks. (a) A best fitting IRT model for the behavioral data; the model simulations take the stimulus images as input, replay the behavioral experiment, and predict contrast thresholds with exact the same parameters for all groups. Symbols are the behavioral data (see Figure 3 for error bars); the lines are the average prediction of the best-fitting IRT, and the bands are the mean \pm 1 SD of the individual simulations. (b) Selected average weights for the four training conditions (*All*, *Near*, *Far*, and *Single*) are shown at the beginning of training (gray) and at the end of training for units tuned to the relevant orientation channels at two different spatial frequencies (1.4 cycles/degrees, which matches the stimulus; and 2.8 cycles/degrees, control) for both location-dependent and location-invariant representations. The increased ranges reflect higher values of positive and negative weights on the decision-relevant evidence.

Parameters	Parameter values
Parameters set a priori	
Orientation spacing $\Delta\theta$	15°
Spatial frequency spacing Δf	0.5 octave
Maximum activation level A_{\max}	1
Weight bounds w_{\min} w_{\max}	± 1
Location-specific initial weights scaling factor w_{init}	0.0169
Location-independent initial weights scaling factor w_{inlt}	0.0169
Activation function gain γ	3.5
Location-specific orientation bandwidth $\sigma\theta$	30°
Location-independent orientation bandwidth η	48°
Location-specific frequency bandwidth h_f	1 octave
Location-independent frequency bandwidth h_{fi}	1.6 octave
Radial kernel width h_r	2 dva
Parameters adjusted for the data	
Normalization constant k	5e-7
Scaling factor a	0.0667
Location-specific internal noise 1 σ_1	1e-7
Location-specific internal noise 2, σ_2	0.06
Location-independent internal noise 1 $\sigma_{11} = 2^* \sigma_1$	2e-7
Location-independent internal noise 2 $\sigma_{12} = 2^* \sigma_2$	0.12
Decision noise σ_d	0.03
Learning rate η	6e-5

Table 1. Parameters of best fit IRT ($r^2 = 0.919$, $\tau = 0.870$).

from physiology and prior applications, there were 6 parameters (of 20 total parameters) free to vary to optimize the fit to the 64 data points (average contrast thresholds in 8 sessions \times 4 groups \times 2 external noise levels). Table 1 shows the best-fit parameter values.

A single learning rate parameter generated different predicted (“empirical”) learning rates in the *All*, *Near*, *Far*, and *Single* groups. The differences in zero and high external noise thresholds also emerge naturally from the same parameter values. These best-fit model predictions provided a good quantitative fit to the behavioural contrast threshold data (Kendall’s $\tau = 0.870$, $p < 0.000001$; and $r^2 = 0.919$, $p < 0.000001$). We also looked at a more complicated model, one that allowed small differences in internal noise parameters between groups (i.e. small differences between the groups of observers), which slightly but nonsignificantly improved the fit to the data (Kendall’s $\tau = 0.883$; $r^2 = 0.938$) (see Appendix B, 1 and Figure B.1 for the graph of the fit, and Table B.1 for details and significance tests), but even with slight level differences between the randomly assigned groups of observers, it is the intrinsic differences in learning experiences that control the learning rates. In sum, the learning rate differences between the training groups emerge organically from the model based on the intermixture of trial experiences and are qualitatively and quantitatively consistent with the behavioural data in the experiment.

Learning altered the weights connecting the sensory representations to decision; examining how the weights

change in the best-fit model can reveal aspects of the learning process in the model. Initial weights were broadly set for the task, reflecting task instructions and general knowledge of orientation. This accounts for the initial above-chance performance in the behavioural data. The weights then change with training so as to focus on the most useful information in the stimulus. The weights on units tuned to the relevant clockwise or counter-clockwise stimulus orientation and spatial frequency of the Gabor ($sf = 1.4$ cpd, cycles per degree) increased or decreased during training, as appropriate, resulting in increased range in weights (Figure 4b). The weights on units tuned to relevant orientations in other spatial frequency channels (e.g. $sf = 2.8$ cpd) were relatively unchanged. In the *All* condition, changes in weights connecting the location-invariant representations to decision “fight” each other from one trial to the next because adjacent orientation stimuli require opposite correct responses in other locations. This, in turn, forces learning into the location-specific representation weights in this condition. In the *Near* condition, fewer weights on location-invariant representations conflict for the most similar stimuli, whereas in the *Far* condition stimuli are sufficiently dissimilar that weights for the tasks can be nearly independent. In the *Single* condition, the weights on the most relevant location-invariant representation units show the largest increases, due to the consistent training in all locations. The full sets of weight changes in the best-fitting IRT model are shown

in [Appendix B, 2](#), and in [Figure B.2](#). In short, training intermixed tasks with similar stimuli that require different responses sets the conditions for catastrophic interference, a common property in neural network learning models. Furthermore, this interference occurs specifically in weights on higher-level location-invariant representations in this experiment.

Trial-by-trial learning, behavior, and IRT model

For this experiment, it was also possible to evaluate learning based on trial-by-trial data, and this analysis revealed some additional features of learning in the four groups. Trial-by-trial and other more fine-grained analyses are only beginning to be deployed in the literature in perceptual learning ([Zhang, Zhao, Doshier, & Lu, 2019](#)). [Figure 5](#) graphs contrasts and corresponding accuracies from the human data (a) and model-simulated predictions (b). On each trial, the adaptive algorithm—here, the accelerated stochastic approximation staircase ([Kesten, 1958](#))—determined the change in the Gabor stimulus contrast based on the accuracy of the observer’s responses in order to track 75% correct. The figure shows the contrast and proportion correct for every trial for the four roving groups, separately for zero and high external noise, averaged over locations and observers (4 locations and 12 observers per group, for 48 trials per point). The vertical lines indicate the session breaks. (See [Appendix C, 1, Figure C.1](#) for graphs with error bars on the all contrast values; the variability in proportion correct data is seen directly in data.) The adaptive algorithm did a good job of keeping the accuracy within $\pm 1.6\sigma$ of (binomial) variability (horizontal dashed lines) of its target value of 75% correct, except for first few trials in the session where contrast step-sizes are quite large (step sizes rapidly decrease thereafter). The patterns in the contrast thresholds are discussed next.

This trial-by-trial analysis reveals several fine-grained observations about the adaptive method and the learning in different conditions, as well as some within- and between-session micro-patterns that may suggest other processes during learning. To begin with, at a qualitative level, the trial-by-trial data reveals differential rates of learning in the different intermixture (roving) groups. Learning emerges in the first session in the *Single* and *Far* conditions, especially in zero noise. The *Near* condition starts to improve next, somewhere during the second session, and clearly by the third session. The *All* condition is even more delayed in showing improvements. (See [Appendix C](#), for a detailed description of the analysis and results, and contrast error bands for all trials.)

Some within-session micro-patterns are visible in the contrast data. An exploratory analysis of these micro-

patterns suggested several influences working together: the adaptive algorithm, overnight consolidation of learning, and within-session deterioration. The accelerated approximation staircase, like essentially all commonly used adaptive staircases, is designed to estimate an unchanging threshold, whereas the thresholds are changing in learning experiments. (In contrast, the newer quick-Change Detection methods build learning curves into the adaptive measurements; [Zhang et al., 2019](#).) An early dip in average contrast followed by a gradual increase is a consequence of this adaptive algorithm, as the step sizes of contrast changes go from very large to small throughout the session (see [Figure C.2](#)). Essentially, if the starting threshold in each session is a bit high, which causes the first response to be accurate more often, then the next trial takes a large step down in contrast, and then adjusts back up with smaller and smaller step sizes. This dip is clearly visible at the beginning of each session.

The contrast values at the end of the previous session were used as the starting values for the next session; apparently, they were somewhat higher than the true threshold at the beginning of the next session. This could be consistent with the often-claimed consolidation improvements during overnight sleep ([Censor, Karni, & Sagi, 2006](#); [Mednick, Drummond, Boynton, Awh, & Serences, 2008](#); [Mednick, Nakayama, & Stickgold, 2003](#)). This situation is likely to lead to a higher than expected number of correct responses on the first trial, and so to the dip in contrast described above. Additionally, there is a trend for deterioration in performance during each session, especially visible in later training sessions (when the amount of within-session learning is very small in the asymptotic phases). This second pattern of within-session deterioration is also consistent with some reports in the literature ([Censor et al., 2006](#); [Mednick et al., 2008](#)).

The IRT model naturally makes trial-by-trial predictions, which are shown in [Figure 5b](#). Because it exactly reprises the experimental protocols in the simulations, it naturally predicts the within-session microstructure associated with the adaptive method. Despite some small systematic deviations, the IRT model provided a strong account of the data (rank order correlation Kendall’s $\tau = 0.847$; $p < 0.00001$; and proportion variance accounted for $r^2 = 0.929$; $p < 0.00001$). Note that these model predictions used the parameter values estimated from the fits to the session contrast thresholds (e.g. they involved no additional optimization to fit the trial-by-trial data), except for an added lapse rate (i.e. rate of guessing trials) that increased within each session (0.0 to 0.2) to capture within-session deterioration. The same IRT model without the within-session lapse rate provided a reasonable, but somewhat worse fit to the data ($F(1, 7658) = 2372.9$; $p < 0.0001$, with rank order correlation Kendall’s $\tau = 0.835$; $p < 0.00001$; and r^2

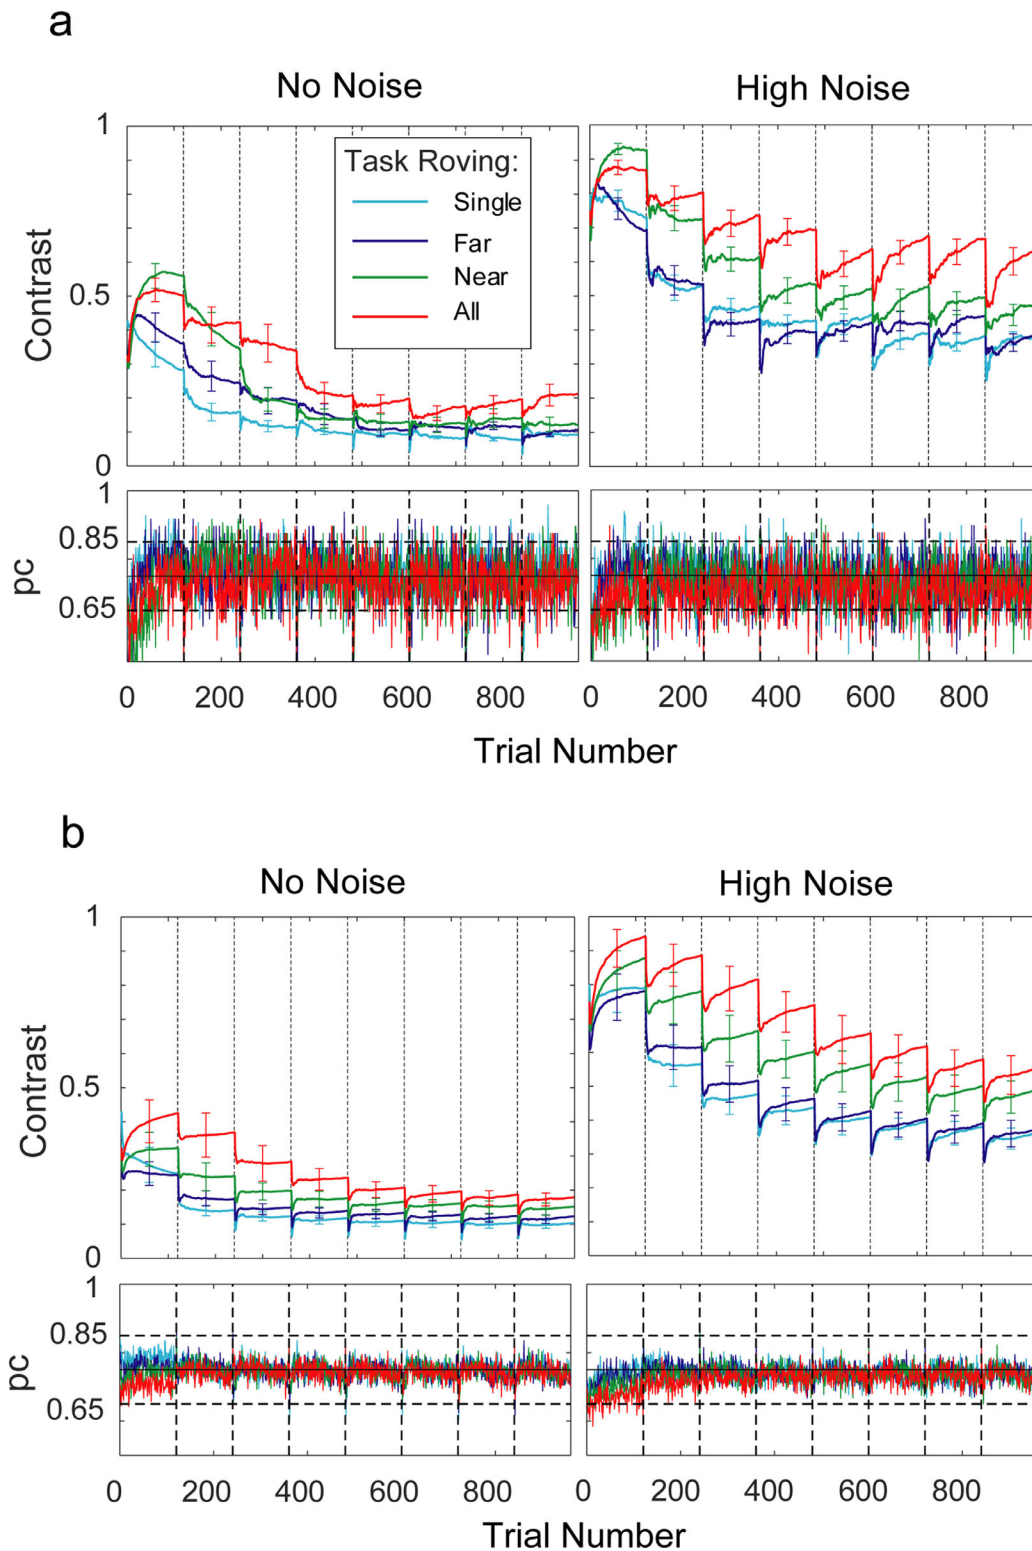


Figure 5. Trial-by-trial contrast and accuracy data averaged over test location and observers (a) and an IRT simulation of these data (b). Trial-by-trial data track the emergence of differential learning between training groups (averaged over four locations and 12 observers per group), and within-session deterioration of performance, especially obvious in later training sessions in which within-session learning is of smaller magnitude. The IRT simulation gives a good account of these trial-by-trial micro-patterns (Kendall's $\tau = 0.847$ and percent variance accounted for $r^2 = 0.929$). Error bars in (a) are from bootstrapping ($n = 1,000$) data, and in (b) are ± 1 SD of the individual simulations. For figure clarity, only error bars in the middle of a session are shown. More detailed error bars are shown in [Figure C.1](#).

= 0.907; $p < 0.00001$). Although these models differed significantly (given the very large number of data points and so degrees of freedom), our choice of the model with an increasing lapse rate was also based on the presence of visible systematic errors in prediction without it, especially at the end of later sessions. A full discussion of trial-by-trial analyses, fits of the model to trial-by-trial data, increasing internal noise as an alternative to the lapse rate, and a discussion of adaptive methods all appear in [Appendix C](#).

Discussion

Summary

This study asked the question: why and under what conditions does training multiple interleaved tasks interfere with perceptual learning? Can we model these patterns of interference? To answer these questions, we manipulated the mixture of tasks trained in four groups of observers. The behaviorally observed learning rates and final performance after thousands of training trials depended on the similarity between intermixed tasks, even when trained in different retinal locations. These interactive effects were substantial, leading to an approximately two-to-one relationship in learning rates (comparing the fastest to the slowest group; e.g. *Single* to *All*).

Training each task in a different retinal location, although it allowed some learning in even the most challenging condition, did not eliminate the damage to learning of intermixture (roving). These results contradicted the predictions of simple forms of pure retuning theory, which attributes perceptual learning to tuning of neurons in the early retinotopic visual cortex, in which plasticity for each retinal location is separate. In contrast, the IRT ([Doshier et al., 2013](#)), using an augmented Hebbian learning rule both qualitatively predicted and quantitatively fitted the substantial differences in learning in the four groups. In the model, tasks trained in different locations interact by shared weights from higher-level location-invariant representations to decision through either destructive interference or cooperative reinforcement. Interference is especially powerful when stimuli that are very similar require opposite responses; or conversely cooperative learning may occur when the weight structures agree. Plasticity must involve both lower-level location-specific and higher-level location-invariant levels in the model to account for the data.

The current explanation, based on the IRT and the general framework of reweighting evidence from a hierarchy of sensory representations, bears some similarities to and some differences from earlier

related network-style explanations of roving (task intermixture) disruption of perceptual learning. In particular, the research of Tartaglia, Herzog, and colleagues also proposed an important role for overlapping representations in inducing loss of learning in roving paradigms ([Tartaglia et al., 2009a](#)). On the other hand, simple neural network learning models were rejected as a class in an earlier paper due to a failure of these models (in some cases) to show disruptions in learning with roving, essentially because the learning mechanisms were too powerful; the conclusion was to discount standard network learning models in favor of reinforcement learning in which roving reflected an inability to separately track a reward expectation for similar stimuli ([Tartaglia et al., 2009b](#)).

Yet, the IRT network model of perceptual learning—together with the central role of higher-level location-invariant representations in the representational architecture—provided a good account of when and how roving challenges learning. These location-invariant representations were originally proposed as the mechanism of transfer across retinal locations, to account for the patterns of transfer when both stimuli and locations might be altered in transfer tests ([Doshier et al., 2013](#)). In more recent work, we have explored other kinds of invariant representations in order to account for other forms of transfer, such as transfer from trained to untrained spatial frequency stimuli in orientation judgments, or transfer from trained to untrained orientations in spatial frequency judgments. All of this suggests the importance of a hierarchy of representations in many visual perceptual learning tasks. In the IRT and its subsequent variants, the location-specific and location-invariant representations have been associated with retinotopic areas of early visual cortex and higher-level areas, respectively ([Doshier et al., 2013](#); [Sotiropoulos, Seitz, & Seriès, 2018](#)).

A similar argument against early retinotopic retuning as the primary basis for perceptual learning was also proposed by [Otto, Ögmen, & Herzog \(2010\)](#). Their study showed that perceptual learning in an illusory Vernier task with moving flankers partially transferred to different retinal locations but not to different orientations or to a standard Vernier stimulus. Furthermore, a series of studies by Yu and colleagues used “double training” and/or “training plus exposure” paradigms to promote the transfer of ordinarily specific training effects to a different retinal location, to other visual features such as orientation, and even to different types of stimuli ([Wang et al., 2016](#); [Xiao, Zhang, Wang, Klein, Levi & Yu, 2008](#); [Zhang & Yang, 2014](#)). These results prompted those researchers to suggest that perceptual learning is “rule-based” and may be mediated through conceptual inference. Another roving study showed that temporal sequencing of the task

variants resulted in a release from roving disruptions in learning (Kuai et al., 2005). Together with the present study, this reveals a complicated set of learning and transfer effects, separate from roving.

Some of the double training and exposure transfer experiments have been successfully modeled by the IRT by us (Liu, Lu, & Doshier, 2011) or by others using a slightly more flexible IRT variant (Talluri et al., 2015). Whether all double-training phenomena can be handled with some version of a reweighting model is an open question. Examples of independent task co-learning of different tasks (e.g. Vernier and bisection) were modeled using distinct decision units and weights for each task (Huang et al., 2012). Release from roving disruption when tasks were temporally sequenced can be similarly modeled if each task uses distinct decision units and weights—although location separation only partially released roving disruptions in the current study. The IRT is an example of a generative model—one that makes predictions for exact experimental designs on a trial-by-trial basis. Any specific experiment might require an extension either to replace the front-end for the stimulus domain or for the task. Further testing of any specific roving or transfer phenomenon would require its own computational modeling study. It remains an open question whether a successful model would utilize additional kinds or layers of representation or, as suggested by some researchers (e.g. Wang et al., 2016), general conceptual learning is involved.

Relation to physiology

These findings demonstrate that retuning of early retinotopic cortex (hypothesized by many researchers) is almost certainly not the only—or even the primary—form of plasticity in perceptual learning. We suggest that this conclusion is broadly consistent with evidence from single-cell recording studies: although small shifts in V1 (Schoups, Vogels, Qian, & Orban, 2001) and V4 (Yang & Maunsell, 2004) tuning have sometimes been reported, they are generally too small to account for the large behavioral improvements associated with perceptual learning (Doshier & Lu, 2017; Law & Gold, 2008), although neural response changes measured while animals are actively performing the task seem to show larger changes in higher levels of visual cortex (in V4 (Raiguel, Vogels, Mysore, & Orban, 2006) or at LIP in motion tasks (Law & Gold, 2008)) that account for a higher portion of the behavioral responses. This pattern of physiological results is consistent with the idea that plasticity must involve representations higher in the visual cortical hierarchy or even in multi-sensory or motor decision areas (Diaz, Queirazza, & Philiastides, 2017). Overall, improved readout (reweighting) appears to be one dominant mode of perceptual learning in low-level or mid-level visual tasks, even though modest

sensory retuning may sometimes occur in certain tasks (Schoups et al., 2001; Seitz & Watanabe, 2005). Indeed, estimates of the influence of changes in V1 or V2 (which have been estimated to account for < 10% of the behavioral improvements) are consistent with the magnitude of learning estimated for retuning within the reweighting models (Petrov et al., 2006). As Figure 1 suggests, even if specific low-level location-specific representations do undergo retuning during learning, the different evidence or activity in these units must still be read out to make a decision and control the motor response (Doshier & Lu, 2017). Although some retuning in early retinotopic areas cannot be ruled out, they are not necessary to account for the behavioral data with the IRT model.

Conclusions

The predictions of the IRT model provided a strong qualitative and quantitative account of the behavioral data in both the session and trial-by-trial measurements. In the model, cross-location interactions reflect learned weight changes for location-specific V1-like representations and for more broadly tuned (and noisier) location-invariant representations, as in higher visual cortex (i.e. IRT or possibly V4-like). Learning is disrupted if the optimized weight structures of the different tasks are in conflict (which occurs when similar sensory stimuli require opposite responses), because updating the weights on one trial may reverse weight changes from other trials—so-called catastrophic interference. The IRT provides a promising framework for predicting the behavioral effects of multiplexed training, but also has been shown to account for many phenomena of transfer (Doshier et al., 2013) and feedback in perceptual learning (Doshier & Lu, 2009; Doshier & Lu, 2017). Visual perceptual plasticity occurs at multiple levels of the visual hierarchy. Further work in physiology or brain imaging may reveal the complex regions underlying this plasticity in particular tasks.

Keywords: perceptual learning, reweighting, readout, orientation discrimination, trial-by-trial

Acknowledgments

Supported by the National Eye Institute (R01EY17491).

The behavioral data and simulated data from this study will be made available upon reasonable request. The code used to generate simulated data from this study will be made available upon reasonable request.

Commercial relationships: Doshier, Liu, and Chu have no competing financial interests. Lu owns intellectual

property rights on qCSF, qPartialReport, and related technologies, and has equity interest in Adaptive Sensory Technology, Inc. (San Diego, CA) and Jiangsu Juehua Medical Technology Co, LTD (Jiangsu, China); these interests are not related to the current research. Corresponding author: Barbara Anne Doshier. Email: bdoshier@uci.edu. Address: Cognitive Science Department, University of California, Irvine, 3151 SSPA, Irvine, CA 92697-5100, USA.

References

- Aberg, K. C., & Herzog, M. H. (2009). Interleaving bisection stimuli—randomly or in sequence—does not disrupt perceptual learning, it just makes it more difficult. *Vision Research*, *49*, 2591–2598.
- Adini, Y., Sagi, D., & Tsodyks, M. (2002). Context-enabled learning in the human visual system. *Nature*, *415*, 790–793.
- Amitay, S., Hawkey, D. J., & Moore, D. R. (2005). Auditory frequency discrimination learning is affected by stimulus variability. *Perception & Psychophysics*, *67*, 691–698.
- Ball, K., & Sekuler, R. (1987). Direction-specific improvement in motion discrimination. *Vision Research*, *27*, 953–965.
- Banai, K., Ortiz, J. A., Oppenheimer, J. D., & Wright, B. A. (2010). Learning two things at once: differential constraints on the acquisition and consolidation of perceptual learning. *Neuroscience*, *165*, 436–444.
- Censor, N., Karni, A., & Sagi, D. (2006). A link between perceptual learning, adaptation and sleep. *Vision Research*, *46*, 4071–4074.
- Chen, P., Engel, S., & Wang, C. (2019). The multivariate adaptive design for efficient estimation of the time course of perceptual adaptation. *Behav Res*, doi:10.3758/s13428-019-01301-6
- Cong, L.-J., & Zhang, J.-Y. (2014). Perceptual learning of contrast discrimination under roving: The role of semantic sequence in stimulus tagging. *Journal of Vision*, *14*, 1.
- Crist, R. E., Li, W., & Gilbert, C. D. (2001). Learning to see: Experience and attention in primary visual cortex. *Nature Neuroscience*, *4*, 519–525.
- Diaz, J. A., Queirazza, F., & Philiastides, M. G. (2017). Perceptual learning alters post-sensory processing in human decision-making. *Nature Human Behaviour*, *1*, Article 0035.
- Doshier, B. A., & Lu, Z.-L. (2009). Hebbian reweighting on stable representations in perceptual learning. *Learning & Perception*, *1*, 37–58.
- Doshier, B., & Lu, Z.-L. (1999). Mechanisms of perceptual learning. *Vision Research*, *39*, 3197–3221.
- Doshier, 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*, *95*, 13988–13993.
- Doshier, B. A., Jeter, P., Liu, J., & Lu, Z.-L. (2013). An integrated reweighting theory of perceptual learning. *Proceedings of the National Academy of Sciences*, *110*, 13678–13683.
- Doshier, B. A., & Lu, Z.-L. (2007). The functional form of performance improvements in perceptual learning: learning rates and transfer. *Psychological Science*, *18*, 531–539.
- Doshier, B., & Lu, Z.-L. (2017). Perceptual learning and models. *Annual Review of Vision Science*, *3*, 343–363.
- Grossberg, S. (1987). Competitive learning: From interactive activation to adaptive resonance. *Cognitive Science*, *11*, 23–63.
- Heathcote, A., Brown, S., & Mewhort, D. J. K. (2000). The power law repealed: The case for an exponential law of practice. *Psychonomic Bulletin & Review*, *7*, 185–207.
- Heeger, D. J. (1992). Normalization of cell responses in cat striate cortex. *Visual Neuroscience*, *9*, 181–197.
- Herzog, M. H., Aberg, K. C., Frémaux, N., Gerstner, W., & Sprekeler, H. (2012). Perceptual learning, roving and the unsupervised bias. *Vision Research*, *61*, 95–99.
- Huang, C. B., Lu, Z.-L., & Doshier, B. A. (2012). Co-learning analysis of two perceptual learning tasks with identical input stimuli supports the reweighting hypothesis. *Vision Research*, *61*, 25–32.
- Jeter, P. E., Doshier, B. A., Liu, S.-H., & Lu, Z.-L. (2010). Specificity of perceptual learning increases with increased training. *Vision Research*, *50*, 1928–1940.
- Jeter, P. E., Doshier, B. A., Petrov, A., & Lu, Z.-L. (2009). Task precision at transfer determines specificity of perceptual learning. *Journal of Vision*, *9*(3), 1.
- 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*, *88*, 4966–4970.
- Kesten, H. (1958). Accelerated stochastic approximation. *The Annals of Mathematical Statistics*, *21*, 41–59.
- Kuai, S.-G., Zhang, J.-Y., Klein, S. A., Levi, D. M., & Yu, C. (2005). The essential role of stimulus temporal patterning in enabling perceptual learning. *Nature Neuroscience*, *8*, 1497–1499.

- Law, C.-T., & Gold, J. I. (2008). Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. *Nature Neuroscience*, *11*, 505–513.
- Liu, J., Lu, Z. L., & Doshier, B. (2011). Multi-location augmented Hebbian reweighting accounts for transfer of perceptual learning following double training. *Journal of Vision*, *11*, 992. (Abstract).
- Liu, J., Doshier, B. A., & Lu, Z.-L. (2015). Augmented Hebbian reweighting accounts for accuracy and induced bias in perceptual learning with reverse feedback. *Journal of Vision*, *15*(10):10, 10–21.
- Lu, Z.-L., & Doshier, B. A. (2008). Characterizing observers using external noise and observer models: assessing internal representations with external noise. *Psychological Review*, *115*, 44.
- Lu, Z.-L., & Doshier, B. (2013). *Visual psychophysics: From laboratory to theory*. MIT Press, Cambridge, MA.
- Massian, M. (2011). A tutorial on a practical Bayesian alternative to null-hypothesis significance testing. *Behavior Research Methods*, *43*, 679–690.
- McCloskey, M., & Cohen, N. J. (1989). Catastrophic interference in connectionist networks: The sequential learning problem. *Psychology of Learning and Motivation*, *24*, 109–165.
- Mednick, S., Drummond, S., Boynton, G., Awh, E., & Serences, J. (2008). Sleep-dependent learning and practice-dependent deterioration in an orientation discrimination task. *Behavioral Neuroscience*, *122*, 267–272.
- Mednick, S., Nakayama, K., & Stickgold, R. (2003). Sleep-dependent learning: a nap is as good as a night. *Nature Neuroscience*, *6*, 697–698.
- Otto, T. U., Herzog, M. H., Fahle, M., & Zhaoping, L. (2006). Perceptual learning with spatial uncertainties. *Vision Research*, *46*, 3223–3233.
- Otto, T. U., Ögmen, H., & Herzog, M. H. (2010). Perceptual learning in a nonretinotopic frame of reference. *Psychological Science*, *21*, 1058–1063.
- Parkosadze, K., Otto, T. U., Malania, M., Kezeli, A., & Herzog, M. H. (2008). Perceptual learning of bisection stimuli under roving: Slow and largely specific. *Journal of Vision*, *8*, 5.
- Petrov, A. A., Doshier, B. A., & Lu, Z.-L. (2005). The dynamics of perceptual learning: an incremental reweighting model. *Psychological Review*, *112*, 715.
- Petrov, A. A., Doshier, B. A., & Lu, Z.-L. (2006). Perceptual learning without feedback in non-stationary contexts: Data and model. *Vision Research*, *46*, 3177–3197.
- Raiguel, S., Vogels, R., Mysore, S. G., & Orban, G. A. (2006). Learning to see the difference specifically alters the most informative V4 neurons. *Journal of Neuroscience*, *26*, 6589–6602.
- Robbins, H., & Monro, S. (1951). A stochastic approximation method. *The Annals of Mathematical Statistics*, *22*, 400–407.
- Sagi, D., Adini, Y., Tsodyks, M., & Wilkowsky, A. (2003). Context dependent learning in contrast discrimination: effects of contrast uncertainty. *Journal of Vision*, *3*, 173. (Abstract).
- Seitz, A., & Watanabe, T. (2005). A unified model for perceptual learning. *Trends in Cognitive Sciences*, *9*, 329–334.
- Seitz, A. R., Yamagishi, N., Werner, B., Goda, N., Kawato, M., & Watanabe, T. (2005). Task-specific disruption of perceptual learning. *Proceedings of the National Academy of Sciences of the United States of America*, *102*, 14895–14900.
- Schoups, A., Vogels, R., Qian, N., & Orban, G. (2001). Practising orientation identification improves orientation coding in V1 neurons. *Nature*, *412*, 549–553.
- Sotiropoulos, G., Seitz, A. R., & Seriès, P. (2011). Perceptual learning in visual hyperacuity: A reweighting model. *Vision Research*, *51*, 585–599.
- Sotiropoulos, G., Seitz, A. R., & Seriès, P. (2018). Performance-monitoring integrated reweighting model of perceptual learning. *Vision Research*, *152*, 17–39.
- Talluri, B. C., Huang, S. C., Seitz, A. R., & Seriès, P. (2015). Confidence-based integrated reweighting model of task-difficulty explains location-based specificity in perceptual learning. *Journal of Vision*, *15*, 17.
- Tartaglia, E. M., Aberg, K. C., & Herzog, M. H. (2009a). Perceptual learning and roving: Stimulus types and overlapping neural populations. *Vision Research*, *49*, 1420–1427.
- Tartaglia, E. M., Aberg, K. C., & Herzog, M. H. (2009b). Modeling perceptual learning: Why mice do not play backgammon. *Learning & Perception*, *1*, 155–163.
- Treutwein, B. (1995). Adaptive psychophysical procedures. *Vision Research*, *35*, 2503–2522.
- Wang, R., Wang, J., Zhang, J. Y., Xie, X. Y., Yang, Y. X., Luo, S. H., Yu, C., . . . Li, W. (2016). Perceptual learning at a conceptual level. *Journal of Neuroscience*, *36*, 2238–2246.
- Xiao, L. Q., Zhang, J. Y., Wang, R., Klein, S. A., Levi, D. M., & Yu, C. (2008). Complete transfer of perceptual learning across retinal locations enabled by double training. *Current Biology*, *18*, 1922–1926.
- Yang, T., & Maunsell, J. H. (2004). The effect of perceptual learning on neuronal responses in

monkey visual area V4. *Journal of Neuroscience*, 24, 1617–1626.

- Yu, C., Klein, S. A., & Levi, D. M. (2004). Perceptual learning in contrast discrimination and the (minimal) role of context. *Journal of Vision*, 4, 4.
- Zhang, J.-Y., Kuai, S.-G., Xiao, L.-Q., Klein, S. A., Levi, D. M., & Yu, C. (2008). Stimulus coding rules for perceptual learning. *PLoS Biol*, 6, e197.
- Zhang, J. Y., & Yang, Y. X. (2014). Perceptual learning of motion direction discrimination transfers to an opposite direction with TPE training. *Vision Research*, 99, 93–98.
- Zhang, P., Zhao, Y., Doshier, B., & Lu, Z.-L. (2019). Assessing the detailed time course of perceptual sensitivity change in perceptual learning. *Journal of Vision*, 19, 9.
- Zhaoping, L., Herzog, M. H., & Dayan, P. (2003). Nonlinear ideal observation and recurrent preprocessing in perceptual learning. *Network: Computation in Neural Systems*, 14, 233–247.

Appendix A

Learning in behavioral data, analysis of variance and post hoc tests on contrast thresholds

The behavioral performance measure for the four training groups (*All*, *Near*, *Far*, and *Single*) was the threshold contrast required to achieve 75% correct from the adaptive staircase procedure. Learning is measured by reduction in thresholds as a function of training. Contrast threshold at the end of each session was estimated by the average contrast in the last 30 trials of the adaptive staircase measured in each condition. (Measuring learning at the scale of large blocks or sessions is typical in the perceptual learning literature; trial-by-trial data are considered below.) The main analysis of variance on the contrast threshold data was described in the main text: 8 sessions for 12 observers in each of four training conditions, in zero and in high external noise (i.e., training session and external noise were within subjects factors and roving or intermixture group was a between subjects factor), or $8 \times 2 \times 12$ or 192 values. Contrast thresholds were, of course, higher in high external noise. Training improved (reduced) the contrast thresholds, and the rate of improvement was different in the four training conditions differed.

Here we report the results of separate analysis of variances in zero and high external noise. As found in previous studies using external noise manipulations (Doshier et al., 2013; Jeter, Doshier, Liu, & Lu, 2010; Jeter, Doshier, Petrov, & Lu, 2009), these differential learning effects were especially clear in high external

noise. In high external noise, there was a significant effect of training ($F(7,308)=53.00$, $p < .0001$; $\eta_p^2 = 0.546$) and training group ($F(3, 44)=5.87$, $p < .005$; $\eta_p^2 = 0.286$), and a marginal interaction ($p \approx .10$; $\eta_p^2 = 0.089$). In zero external noise, there was a significant overall effect of training ($F(7,308)=30.93$, $p < .0001$; $\eta_p^2 = 0.413$) and of training group ($F(3,44)=2.82$, $p \approx .05$; $\eta_p^2 = 0.162$). These separate analyses of variance on zero and high external noise data used 8 sessions and 12 observers per group (96 values). The corresponding values of $p_{BIC}(H_1|D)$ in both high external noise and zero external noise were > 0.999 for training block and training (roving) group, and < 0.001 for the interaction.

Post-hoc t-tests (Bonferroi correction $\alpha = .008$) paralleled those reported for the overall data in the main text, showed the same pattern in both high and low external noise conditions separately: in the high external noise data, condition differences were significant ($p < .0001$), except *All* versus *Near*, $p < .01$, and *Far* versus *Single* (n.s.); in the zero external noise data, all condition differences were significant ($p < .001$ except *All* versus *Near*, $p < .05$, *Far* versus *Single*, $p < .02$, and *Near* versus *Far*, $p < .08$).

Power function fits and learning rates in contrast threshold learning curves

The learning curves of the four training conditions (*All*, *Near*, *Far*, and *Single* groups) differed significantly, as measured by fitted power function learning curves. The power function equation is: $C(t) = \lambda(t + 1)^{-\beta} + \alpha$ where $c(t)$ is the threshold in session t , $\lambda + \alpha$ is the initial threshold, α is the asymptotic threshold late in training, and β is the learning rate (Doshier & Lu, 2007). For this data, the initial threshold estimates performance prior to training at $t=0$; training counts starting at $t=1$ because behavioural performance was measured near the end of the corresponding sessions. (See also [Appendix C](#) for a trial-by-trial analysis.) The four conditions were fit simultaneously by models using different numbers of free parameters in a partial lattice of systems of equations using the Matlab routine *fminsearch*. The null model ($1\lambda - 1\beta - 1\alpha$) assumes that the different training mixture conditions were the same, while a fully saturated model ($4\lambda - 4\beta - 4\alpha$) fits each curve independently with no common parameter values, with a variety of models in between. The data for high external noise and for zero external noise data were fit separately, because the descriptive parameters for threshold learning functions depend on the external noise. Nested sub-models, in which one model is a special restricted case of another, can be tested for significant differences using an F-test (see [Methods](#) for a description of the F test for nested models).

Model	4λ -4 β-4 α	4λ -4β-1α	1λ -4β-4α	1λ -4β-1α	1λ -3β-1α	r ²
4λ -4 β-4α	–	–	–	–	–	0.9852
4λ -4β-1α	F(3,35) = 1.47 ^{ns}	–	–	–	–	0.9851
1λ -4β-4α	F(3,35) = 14.159**	•	–	–	–	0.9677
1λ -4β-1α	F(6,35) = 12.026***	F(3,38) = 25.231***	F(3,38) = 4.852*	–	–	0.9554
1λ -3β-1α	F(7,35) = 10.410***	F(4,38) = 19.112***	F(4,38) = 3.726*	F(1,41) = 0.273 ns	–	0.9551
1λ -1β-1α	F(9,35) = 128.926***	F(6,38) = 205.048***	F(6,38) = 91.379***	F(3,41) = 138.789***	F(2,42) = 211.714***	0.5024

Table A.1. Comparisons of power-function models for the threshold learning curves in high external noise. Note: • Non-nested models; *p < .1, **p < .001, ***p < .00001; ^{ns}p > .1.

Model	4λ -4 β-4 α	4λ -4β-1α	1λ -4β-4α	1λ -4β-1α	1λ -3β-1α	r ²
4λ -4 β-4α	–	–	–	–	–	0.9676
4λ -4β-1α	F(3,35) = 0.255 ^{ns}	–	–	–	–	0.9669
1λ -4β-4α	F(3,35) = 5.691*	•	–	–	–	0.9519
1λ -4β-1α	F(6,35) = 4.779*	F(3,38) = 9.884**	F(3,38) = 2.823 *	–	–	0.9411
1λ -3β-1α	F(7,35) = 7.687**	F(4,38) = 14.089**	F(4,38) = 6.702 **	F(1,41) = 16.183**	–	0.9179
1λ -1β-1α	F(9,35) = 33.404***	F(6,38) = 53.010***	F(6,38) = 34.488***	F(3,41) = 58.370***	F(2,42) = 58.365***	0.6897

Table A.2. Comparisons of power-function models for learning curves in zero external noise. Note: • Non-nested models; *p < .1, **p < .001, ***p < .00001; ^{ns}p > .1.

Tables A.1 and A.2 show r² for the different nested models, and significance values separately for the high external noise and the zero external noise data. Power function learning curves with different learning rates for each training (roving) condition, but equivalent initial and final asymptotic performance levels (1λ - 4β - 1α), provided a significantly better fit than a single learning curve (1λ - 1β - 1α) (F(3,41) = 138.789, p << .00001 in high external noise; F(3,41) = 58.370, p << .00001 in low external noise; see tables for r² values)—demonstrating that the four different training groups differed significantly in learning rate. Although similar, the *Far* and *Single* conditions differed slightly, the (1λ - 4β - 1α was better than a 1λ - 3β - 1α model in low noise, F(1,41)=16.183, p < .001; but the difference is not significant in high noise, F(1,41) = 0.273, p > 0.1).

The random assignment of observers to training condition should lead to approximately equal initial thresholds prior to training (λ + α) (e.g., at t=0 before the first session, although they differ even by the end of the first session of training), corresponding with the (1λ - 4β - 1α) model. There were small, but significant improvements in the quality of fit if the groups were allowed to differ in initial thresholds (4λ - 4β - 1α) (F(3,38) = 25.231, p < .00001 in high external noise; F(3,38) = 9.884, p < 0.001 in low external noise). However, additionally allowing different asymptotic contrast thresholds after extensive training (4λ - 4β - 4 α) did not further improve the quality of the fits

(F(3,35) = 1.470, p = .249 in high external noise; F(3,35) = 0.255, p = 0.839 in zero external noise). Although there may be slight differences between the four randomly assigned groups in their initial levels of performance, we focus in our report on the more parsimonious (1λ - 4β - 1α) fits in the main text. Although significant, the differences in r², with and without the added λ parameters are relatively small and may reflect over fitting unconstrained by data at time t=0, and some of less constrained estimated parameter values are implausible, suggesting that these may be the result of over-fitting.

Table A.3 shows the parameter values for the power function fits to the threshold data from the behavioural experiment, along with the bootstrapped mean and standard deviation (Mean, St. Dev.) of the parameters estimated from the best-fitting power function model. The bootstrap method (n=1000) sampled groups from the original group data sets (selecting n=12 observers at random with replacement for each condition). The standard deviation of rate parameters normalized within fits (St. Dev.*) was also computed; normalization may provide a better estimate of the reliability of the βs because this compensates for the correlations with the estimated values of λ and α for a given fit (e.g., if the estimate of λ is high and of α is low in the fit of a new set of bootstrapped pseudo-data, this will shift all 4 βs to higher values, which contributes to the standard deviation of each β). The rank-orders of the β's for the four conditions in these bootstrapped fits are also

Parameter	Estimate	Mean	St. Dev.	St. Dev.*				
No External Noise								
$\lambda + \alpha$	1.1697	1.2192	0.2511		Ordinal Frequencies			
α	0.0713	0.0659	0.0202		$\beta_{All <}$	$\beta_{Near <}$	$\beta_{Far <}$	$\beta_{Single <}$
β_{All}	1.1478	1.2098	0.3627	0.2231	—	—	—	—
β_{Near}	1.3763	1.4247	0.2937	0.2454	786	—	—	—
β_{Far}	1.7446	1.8394	0.5166	0.2960	949	827	—	—
β_{Single}	2.3077	2.4284	0.5923	0.3438	998	994	875	—
High External Noise								
$\lambda + \alpha$	1.2241	1.2192	0.0796		Ordinal Frequencies			
α	0.3262	0.3048	0.0490		$\beta_{All <}$	$\beta_{Near <}$	$\beta_{Far <}$	$\beta_{Single <}$
β_{All}	0.5538	0.5380	0.1435	0.1193	—	—	—	—
β_{Near}	0.7936	0.7707	0.1571	0.1329	950	—	—	—
β_{Far}	1.3242	1.2885	0.3350	0.2101	1000	979	—	—
β_{Single}	1.2836	1.2570	0.2965	0.1918	1000	981	469	—

Table A.3. Parameter estimates for the $(1\lambda - 4\beta - 1\alpha)$ power function models in zero and high noise. Note: Parameter estimates for the power function fits of the $(1\lambda - 4\beta - 1\alpha)$ model to the observed group data (Estimates), and the resampled mean (Mean), standard deviation (St. Dev.), the normalized standard deviation (St. Dev.*), and the frequency of orders from power function fits to $n=1000$ bootstrapped data sets (12 observers sampled at random with replacement for each condition). The normalized standard deviation partially corrects for rate β_s correlations with λ and α parameters (see text). The ordinal frequencies indicate high consistency in the order of learning rates in fits to bootstrapped data sets; e.g., the count out of 1000 fits to bootstrapped in the four training groups with $\beta_{Far} < \beta_{Single}$ in zero noise was 875.

shown in Table A.3; these frequency statistics indicate a very stable rank ordering of the estimated rates, except for the near-equivalence of the β_{Far} and β_{Single} in high external noise.

Tests for differences between groups at the beginning of the first session

The equivalence of the initial performance of the observers randomly assigned to the four groups (12 observers per group, 48 overall) was verified in an analysis of variance of threshold contrasts in the first session, up to trial 61-70 in zero external noise and trial 41-50 in high external noise ($p_s > .05$) in an analysis across 4 locations of every 5 trials in the first session separately for each noise level (there are four staircases per noise level, one in each location). Excluding the first 5 trials (during which time the differences in contrasts were more affected by initial contrasts given to subjects), the difference among groups steadily increased within the session for both noise levels (for low noise, η_p^2 increased from 0.0113 to 0.174, and $p_{BIC}(H_1/D)$ increased from 0.0011 to 1.000; for high noise, η_p^2 increased from 0.0029 to 0.241, and $p_{BIC}(H_1/D)$ increased from 0.0005 to 1.000; both $p_{BIC}(H_1/D) > 0.99$ from about the first 50 trials in each location, or 200 trials total). This equivalence was tested with the session-level data using the power function fits (above), which fit the data very well with models

assuming equal initial performance before learning, although differences were already seen by the end of the first session at the first data point. We carried out collateral tests of the equivalence between the four groups at the beginning of training—the equivalence of the initial state before learning—from the trial-by-trial data (see Appendix C).

Appendix B

Fits of the integrated reweighting theory (IRT) and a variant

The IRT model presented in the main text accounts well for the human behavioral data in both ordinal predictions (Kendall's $\tau = 0.870$) and in the quantitative fit ($r^2 = 0.919$). The parameters for this simulation were listed in table 1 and the model fits and summaries of the weight changes were shown in figure 4 in the main text, with 6 of 20 parameters free to vary to fit 64 data points: normalization constant k , scaling factor a , location-specific internal noise σ_1 and location-invariant internal noise σ_2 , decision noise σ_d and learning rate η . Location-invariant internal noises were set at twice the location-specific internal noise, and other parameters were set from prior applications of the model.

Parameters	Parameter values
Parameters set a priori	
Orientation spacing $\Delta\theta$	15°
Spatial frequency spacing Δf	0.5 octave
Maximum activation level A_{\max}	1
Weight bounds w_{\min} w_{\max}	± 1
Location-specific initial weights scaling factor w_{init}	0.0169
Location-independent initial weights scaling factor $w_{\text{init}l}$	0.0169
Activation function gain γ	3.5
Location-specific orientation bandwidth σ_d	30°
Location-independent orientation bandwidth η	48°
Location-specific frequency bandwidth h_f	1 octave
Location-independent frequency bandwidth h_{fl}	1.6 octave
Radial kernel width h_r	2 dva
Parameters adjusted for the data	
Normalization constant k	5e-7
Scaling factor a	0.0667
Location-specific internal noise 1 σ_1	5e-8
	1e-7
	1e-7
	8e-8
Location-specific internal noise 2, σ_2	0.05
	0.02
	0.07
	0.08
Location-independent internal noise 1 $\sigma_{1l} = 2^* \sigma_1$	1e-7
	2e-7
	2e-7
	1.6e-7
Location-independent internal noise 2 $\sigma_{2l} = 2^* \sigma_2$	0.10
	0.04
	0.14
	0.16
Decision noise σ_d	0.03
Learning rate η	6e-5

Table B.1. Parameters of the IRT fits with differences between groups ($r^2 = 0.938$, $\tau = 0.883$).

We also fit another version of the IRT in which the four training groups could vary in scaling and internal noise, thus allowing for possible differences between the sets of randomly assigned observers in the four training conditions. This model gives only a slightly and non-significantly better fit to the behavioural data in ordinal predictions (Kendall's $\tau = 0.883$) and quantitative fit ($r^2 = 0.938$). The parameters for this simulated model variant are listed in Table B.1, and the corresponding fit to the data is shown in Figure B.1. This model has 15 varying parameters (4 scaling factors a , 4 σ_1 's and 4 σ_2 's, k , σ_d , and η , (total of 29 parameters). Although this model seems to fit the data somewhat better by eye, the quality of fit was not significantly better ($r^2 = 0.938$ versus $r^2 = 0.919$) after accounting for the increase in the number of free

parameters ($F(9,34) = 1.158$, $p = 0.352$). Even assuming that there were slight differences between the randomly assigned groups of observers, the added parameters primarily adjust the initial level of performance, and it is still the intrinsic differences in learning experiences that largely accounted for the different rates of perceptual learning.

Changes of weight structures in the IRT model

In the IRT, learning changes the weights that connect evidence (activation) in sensory representations to decision, gradually improving the visual task performance over the course of training by changing the “readout” of sensory evidence. Some aspects of the




Figure B.1. A best fitting IRT model with slightly different parameters among groups. (Parameters are listed in Table B.1). With these six additional parameters, Kendall's τ and r^2 were slightly better but the improvement was not statistically different. The line and shaded areas from the model predictions are the mean and ± 1 standard deviation from individual simulations ($n=1000$).

weight changes were illustrated in figure 4b; we provide more complete information about weight changes in the best fitting model in this section.

Changes in the weights for representation units most sensitive to the relevant orientation stimuli in the most relevant and a less relevant spatial frequency were graphed in Figure 4b (main text) as the average magnitude of the weights. Over the course of learning, the weights (positive or negative) on representation units tuned to orientations and spatial frequencies near those of the Gabor stimuli are increased, while the weights on less relevant ones show only small changes. Note that in previous applications of the IRT in different experimental paradigms, weights on less relevant channels decreased over the course of learning (Dosher et al., 2013).

More complete snapshots of the weight structures are included here. Figure B.2 shows the *initial* and *final* weights (blue and red, respectively) for units tuned to all the different orientations at a near-matched spatial frequency ($sf=1.4$ cpd, or cycles per degree) and at a less relevant spatial frequency ($sf=2.8$ cpd) for location-specific and location-invariant representations for the fit listed in the main text. Examination of these patterns shows that learning results in larger positive weights on orientation units best tuned to the “clockwise” stimulus orientations, and larger




Figure B.2. Initial (blue) and final (red) weights from the best-fitting IRT simulation model ($n=1000$ runs) for location-specific and location-invariant representations in the four training groups (*Single*, *Far*, *Near* and *All*) as a function of orientation, shown for (a) the most relevant spatial frequency channel ($sf = 1.4$ cycles/deg) and (b) a less relevant channel ($sf = 2.8$ cycles/deg). (The IRT representations used units tuned to all combinations centered on 5 spatial frequencies and 12 orientations, for 60 channel weights in each location or 300 overall. For visual clarity, each panel shows the weights at the 12 orientations as lines.) For location-specific units, weights increased for those channels tuned to “clockwise” orientations and decreased for “counterclockwise” orientations used in the different tasks for all four groups. When the tasks in the four locations are compatible (*Single*), there is substantial learning in location-invariant weights, while if they are more incompatible (*All*), there is almost no learning in location-invariant weights. In this case, the weights in less relevant channels (b) had little change over the course of learning.

negative weights on orientation units best tuned to the “counter-clockwise” stimulus orientations. For the location-specific representations, the magnitudes of weight change (learning) from location-invariant representations to decision is about the same in the four training conditions. In contrast, the magnitudes of weight change from location-invariant representations to decision differ markedly in the four training conditions; these weights make the most substantial contribution to predicting the condition differences.

The initial weights, as explained in the main text, were set to reflect approximate information about orientation judgments in each training condition, as required to account for above chance initial performance; these initial weights are also thought to embody information that the observers possess based on initial instructions and general information about orientations. In our experiment, observers were shown printed images of the different oriented stimuli to be discriminated in each location as part of the instructions. (This is unlike category learning paradigms in which the required judgments need to be discovered.) In the current simulations, approximate weights were set up only around the instructed reference angle(s) for each group, e.g., positive weights on close angles clockwise and negative on close angles counter clockwise of the reference angle, corresponding to the instructions actually provided. We also investigated several other initial weights patterns. In one, initial weights were set around all four reference angles even when only a subset were used in that training condition; this led to similar ordinal predictions. In another, the initial weights were set randomly around zero; this led to poor fits because it predicts random initial performance and was unable to achieve the accuracy required for the adaptive staircase in the first session. As training goes on, however, each of these simulations (assuming some reasonable learning occurs) makes the same ordinal predictions for the four training groups, predicting behavioural learning rates in the order $All < Near < Far < \approx Single$.

Appendix C

Trial-by-trial data, model-simulations of trial-by-trial data, and within session micro-patterns

Figure 5 shows the empirical trial-by-trial performance, along with the corresponding predictions of an IRT model. Trial-by-trial data, in the experiment and in the model, were measured by taking the contrast and proportion correct for the 1st, 2nd, 3rd, ... etc. trials within each adaptive staircase, averaged over the

separate staircases in the four retinal locations and over observers, separately for the four groups in the two external noise conditions. Vertical dashed lines mark between sessions, carried out on different days. Horizontal dashed lines in the proportion correct graphs are $\pm 1.6\sigma$ (binomial) from the target accuracy of 75% correct of the adaptive algorithm. Error bars were shown in figure 5 only for mid-session values for visibility. Figure C.1 shows bootstrapped error regions for every trial of the trial-by-trial average contrasts.

Some interesting features of the trial-by-trial data, both empirical and simulated, are the within-session patterns in stimulus contrasts. These are an emergent property of the adaptive algorithm. All the relevant adaptive algorithms (in particular, all of those routinely used in perceptual learning studies) were originally designed to estimate a threshold of a stationary, unchanging, psychometric function, whereas learning causes non-stationary improvements, and so the adaptive methods can be sluggish in estimating actual threshold changes within a session. In this study, thresholds were measured using the *accelerated stochastic approximation* algorithm, (Kesten, 1958) which tracks a target accuracy ϕ , here 75% correct (see the [Methods](#) for details). This adaptive algorithm, as most others such as the up/down staircases and others (Lu & Doshier, 2013), either reduces the size of up or down steps (here, in contrast) over trials within a run or session, using big steps early to do range finding, and smaller steps later to converge with more precision or alternatively may focus test conditions in regions determined early in the sequence.

In addition and separately, it appears from the data that the starting values are higher than the actual initial thresholds in each subsequent day (which are taken from the final contrast of the prior session), especially in the first several sessions. This might reflect overnight consolidation that has been widely claimed to occur in perceptual learning (Karni & Sagi, 1991; Mednick et al., 2003; Mednick et al., 2008), or possibly a release from within-session deterioration (Censor et al., 2006; Mednick et al., 2008), as described briefly in the main text. If the starting contrast for the staircases in each session (after the first) is, as seen in the trial-by-trial data, set high relative to the true threshold being estimated, then the observer’s first response is likely to be correct, leading to the undershoot in the accuracy early in the session that is subsequently recovered as the algorithm calibrates, as discussed in the main text.

If the true threshold were also improving during the session due to learning—which is especially apparent in early sessions before learning is approaching asymptote—then the contrasts tend to lag behind. This may be especially true in the first few sessions in which learning is typically more rapid (due to the power function or exponential form of the learning curve) (Doshier & Lu, 2007). If in addition there is

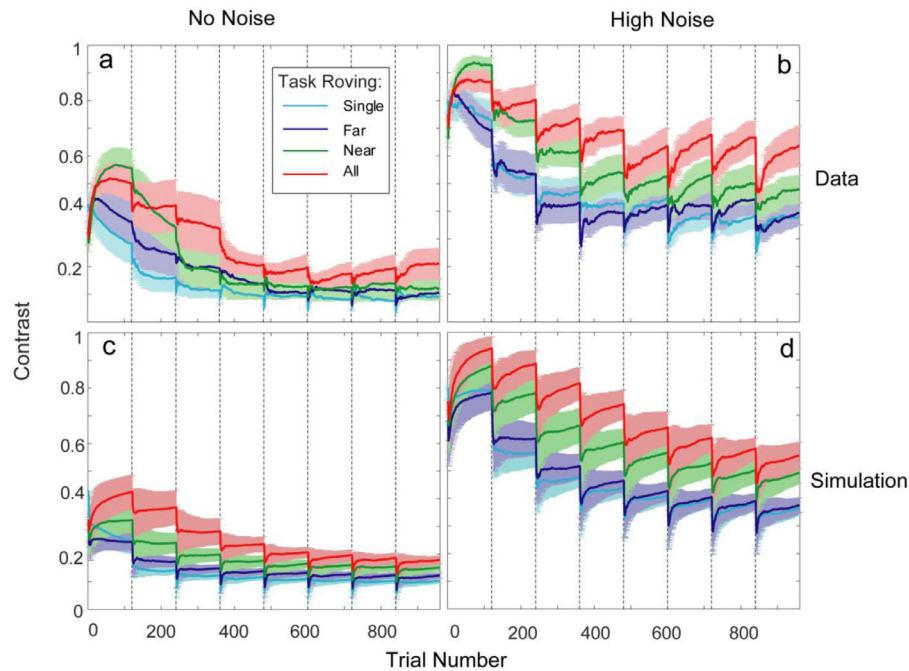


Figure C.1. Trial-by-trial contrast data (top) and simulation (bottom) with shaded error bars. The error bars for the data are ± 1 standard deviation from bootstrapping data ($n=1000$); the error bars for the simulation are ± 1 standard deviation of the individual simulations ($n=1000$).

deterioration in performance over the course of a session, as sometimes claimed, then the thresholds may actually increase within the session, and counteract or even fully overcome the positive effects of within-session improvements due to learning. Both of these features are visible in the trial-by-trial data. These patterns likely would appear in most perceptual learning data if it could be analysed in this way. The majority, indeed nearly all, of current studies of perceptual learning use adaptive methods with similar assumptions. Alternative adaptive methods specifically designed to estimate thresholds that change over trials are currently under development by ourselves (Zhang et al., 2019) and others (Chen, Engel, & Wang, 2019).

To illustrate the within-session micro-patterns that may occur, we simulated how the adaptive staircase would behave given an observer with or without learning (threshold reductions), and with or without within session deterioration (modelled here as an increasing lapse rate corresponding with increased deterioration or fatigue) for three different starting contrast levels c_0 (above, below, and approximately at the “true” threshold). Figure C.2 shows the predicted patterns in the average stimulus contrast as a function of trial within a session that can emerge with the accelerated approximation algorithm for: (a) no changes in threshold, flat lapse rate (stationary performance); (b) decreasing threshold, flat lapse rate; (c) no changes in threshold, increasing lapse rate; and (d) decreasing threshold and increasing lapse rate. Two of these

patterns are typical of those actually seen in the human behavioural data in earlier (d) and later (c) sessions, respectively.

The IRT model generates the predictions for trial-by-trial data shown in figure 5 using the parameters listed in Table 1 in the main text. These trial-by-trial predictions used these parameter values estimated from session-level data, with one elaboration. The model predicts performance improvements throughout training based on its incremental trial-by-trial learning process; it also included an increasing lapse rate throughout each session to account for the apparent deterioration of performance within the testing session. This deterioration is especially visible in the later sessions in which it is not counteracted by the more substantial learned improvements occurring in early sessions. Although we also explored modelling within-session deterioration as increasing internal decision noise (and this may work as well), we settled on using a lapse rate that increased from 0.0 to 0.2 over each session because lapse rates often play a role in staircase threshold algorithms. As described in the main text, the relationship between the predicted and the empirical data is still quite good (though significantly worse) even without including the lapse rate ($r^2=0.907$, Kendall’s $\tau=0.835$). There seems to be a deviation between the model and the data due to slower than expected learning of the *All* condition in the first and second sessions, which is not predicted by the model, with or without lapse rates. We looked at models that

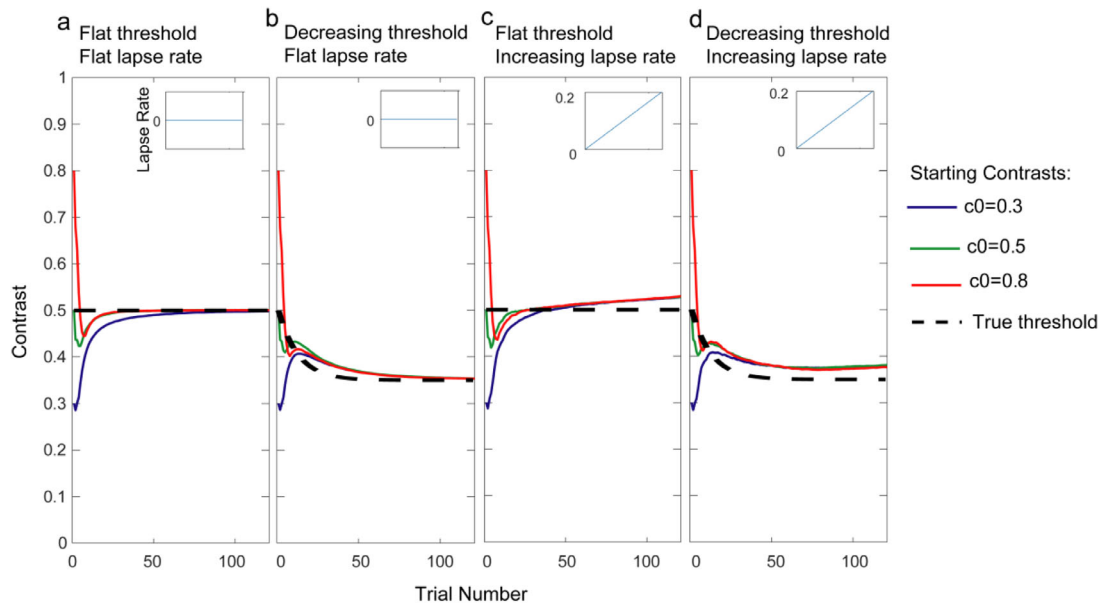


Figure C.2. Simulations illustrating micro-patterns of performance of the adaptive staircase procedure, with or without learning and/or within-session deterioration. In these illustrations, the ASA (accelerated stochastic approximation) staircase is simulated with three different starting contrasts for a “true” threshold starts at a contrast of 0.5: (a) when the threshold is stable – no within-staircase learning or lapsing; (b) with learning (decrease of contrast threshold) but no lapsing; (c) with an increasing lapse rate but without learning; (d) when there are both decrease of threshold and increase of lapse rate. When compared with the trial-by-trial data, the pattern is consistent with within-block learning, especially in early blocks (b or d), while the within-block deterioration is more prominent in late blocks (c).

incorporated different lapse rates for the different task mixture conditions, e.g., larger lapse rates for the *All* condition and smaller for the *Single* condition. While these models visually improved the fits slightly, they did not significantly improve the r^2 enough to outweigh the increase in parameters in F-tests.

In both cases, with or without a lapse rate included, it should be emphasized that these relatively good fits of the model to trial-by-trial data are largely prospective, in the sense that the session-end data were used to estimate the primary values of the model parameters, which were not optimized for the trial by trial data.