

Perceptual learning improves efficiency by re-tuning the decision 'template' for position discrimination

Roger W Li, Dennis M Levi & Stanley A Klein

Visual position discrimination improves with practice; however, the mechanism(s) underlying this improvement are not yet known. We used positional noise to explore the underlying neural mechanisms and found that position discrimination improved with practice over a range of noise levels. This improvement can be largely explained by an increasing efficiency with which observers used positional information in the stimulus. In a second experiment, we tested the hypothesis that the improved efficiency reflects a re-tuning of the observers' perceptual 'template'—the weightings of inputs from basic visual mechanisms—to more closely match the ideal template required to perform the perceptual task. Using a new technique to measure which parts of the stimulus influenced the observer's performance, we were able to record the re-tuning of the decision template across training sessions; we found a robust and steady increase in template efficiency during learning.

In vision, as in life, practice improves performance. This 'perceptual learning' has been shown to improve contrast detection¹, orientation discrimination², bisection³ and Vernier⁴ judgments, depth perception⁵, motion detection⁶, texture segregation⁷ and pattern recognition⁸. The learning effects are often specific to the trained visual tasks, stimulus attributes, and eye and retinal location. Thus, it has been suggested that the adult brain retains a large degree of neural plasticity.

What are the neural changes that occur during repetitious practice of a visual task, and at what level of processing in the nervous system do these changes occur? Many early studies addressed the issue of whether learning takes place early or late in the visual system by examining the specificity of learning (that is, whether improvement transfers to a different task, orientation, eye, and so on). Later studies used luminance noise (for example, adding dark and bright pixels to the picture, like snow on a TV screen) to ask what is being learned^{9–11}.

In the present study we used 'positional noise' (*i.e.*, perturbation of the positions of parts of the stimulus) to investigate what is learned when practicing a position discrimination task. Position discrimination, the ability to distinguish changes in the relative position of two features, is especially susceptible to perceptual learning in normal foveal vision^{4,12–14}, even in the absence of noise. Positional acuity is limited by at least two different mechanisms^{15,16}. For abutting or closely spaced stimuli, position discrimination depends on target visibility, and is thought to be limited by the response properties of contrast-sensitive filters. An earlier study¹⁷ using one-dimensional luminance noise at various orientations showed that perceptual learning shrinks the range of orientations over which noise influences position acuity for abutting lines. The authors conjectured that improvement might occur because observers learn to 'pay attention' to the set of neurons that is most sensitive to the task. For separated

stimuli (like those used in the present study), position discrimination shows little dependence on target visibility (reducing the target contrast from 99% to 50% has no effect on performance), but depends strongly on spatial relations, and is thought to be limited by positional uncertainty^{15,18}. Thus, we used positional noise to mimic the putative limiting internal noise by rendering the positions of the samples uncertain (rather than by obscuring their visibility, as occurs with luminance noise) and to explore the underlying neural mechanisms for position discrimination.

Our first experiment shows that practice improves performance over a wide range of noise levels through more efficient use of the stimulus information (increased efficiency). In the second experiment, we further investigated the improvement in efficiency by using the trial-by-trial effects of noise (termed 'molecular' psychophysics¹⁹) to quantify the influence of different parts of the stimulus on the observers' perceptual decisions and to record changes in the observers' decision template (which reflects the weightings of inputs from basic visual mechanisms). The result is a 'classification image,' which is a map or spatial profile that shows which parts of the stimulus influence the observer's performance. Using an efficient technique for measuring the observers' classification image in positional noise, we were able to record the re-tuning of the decision template across training sessions and found a robust and steady increase in template efficiency during learning.

RESULTS

Experiment 1: position discrimination learning

By measuring performance at different noise levels and using an influential model^{20,21}, it is possible to attribute any changes during learning to two important factors: a change in equivalent input noise and/or an increase in the efficiency with which the stimulus informa-

School of Optometry and Helen Wills Neuroscience Institute, University of California, Berkeley, Berkeley, California 94720-2020, USA. Correspondence should be addressed to D.M.L. (dlevi@uclink.berkeley.edu).

Published online 18 January 2004; doi:10.1038/nn1183

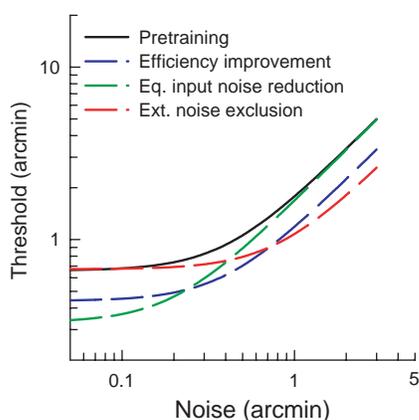


Figure 1 Three possible mechanisms for visual learning. The post-training curve shows the effects of (i) improvement in efficiency, (ii) lowered equivalent input noise levels with fixed efficiency and (iii) external noise exclusion through learning.

tion is used. Equivalent input noise is the noise that must be added to the stimulus to mimic the limiting noise in the visual system. Efficiency reflects the computation underlying the use of the information (samples) of the stimulus²¹.

Three of the possible post-training outcomes (threshold versus noise curves) based on an early noise model^{20,21} are illustrated in **Figure 1**: (i) A pure improvement in efficiency would shift the curve downward. This type of improvement has been reported for learning faces and complex patterns⁹. A mostly downward shift also occurs for learning orientation discrimination in peripheral vision^{10,11}. (ii) A pure decrease in equivalent input noise would shift the ‘knee’ point of curve down and to the left. We have seen this type of change in the fovea of amblyopes during learning (R.W.L. and D.M.L., unpublished data). (iii) Another pattern of learning that has been found^{10,11} for learning simple foveal discrimination tasks is a rightward shift of the curve, produced by a combination of improved efficiency and increased equivalent input noise. This type of learning has also been modeled^{10,11} in terms of improved exclusion of external noise with fixed equivalent late noise.

In the first experiment, observers practiced a position discrimination task in which they were asked to judge which of three pairs of lines was misaligned (top, middle or bottom in **Fig. 2**). We chose this simple task because even three-year-olds can pick the ‘odd man out’²², and we plan to use the same test to study perceptual learning in children. We introduced positional noise by perturbing the positions of the individual patches of each segment according to a Gaussian probability function, with the mean offset of the jittered segments set to zero. Equivalent input noise and efficiency were estimated by systematically manipulating the stimulus positional noise in each training session. A learning session consisted of 750 responses and took about two and a half hours (including rest breaks). Each observer provided more than 6,000 responses over eight learning sessions.

We found that position discrimination gradually improved during practice at all levels of noise (**Fig. 3a**). The most notable change was the downward shift in the curves, which are the model fits to the data; there was also a very subtle leftward shift of the knee points across the measurements for eight of the ten observers. The mean improvement in position discrimination was about 23% for all noise settings. The threshold data are replotted as a function of session in **Figure 3b** to illustrate the time course of learning. A bi-linear regression model

($y = a(x - x_0) + y_0$ for $x < x_0$ and $y = y_0$ for $x > x_0$) was used to fit the data. The position discrimination thresholds for all noise settings decreased (improved) steadily, reaching asymptotic performance after about five sessions. The slopes (a) of the regression lines were found to be significantly different from zero for all five noise levels (see **Fig. 3** legend).

We used a position averaging model²⁰ to parse the improvement into changes in equivalent input noise (**Fig. 4a**) and efficiency (**Fig. 4b**). Since performance reached an asymptote in the last three sessions, the statistical test was performed based on the first and the mean of the last three sessions. We found a small and non-significant decrease (about 11%) in mean equivalent input noise of the last three sessions compared to that of the first session (paired $t = 1.30$; $P = 0.23$), but a substantial and significant increase (about 35%) in mean efficiency (paired $t = 2.68$; $P = 0.03$). The 1.35-fold increase in average efficiency predicts an approximately 1.16-fold ($\sqrt{1.35}$) improvement in average threshold, slightly lower than the 1.23-fold improvement that we observed. There are important individual differences in learning: most observers showed markedly increased efficiency and slightly reduced noise, although two observers showed reduced equivalent input noise but did not show any change in efficiency (**Fig. 4c**).

Experiment 2: classification image

Our aim in the second experiment was to better understand the components that contribute to improved efficiency. For example, improvement in efficiency might reflect improved use of the stimulus samples (patches) through re-tuning the template for position discrimination. We approached this question by computing the classification image during learning. Classification images provide an important tool for measuring the ‘template’ an observer uses to accomplish a task^{23–26}. By keeping track of both the pattern of noise and the observer’s responses on each trial, it is possible to compute the correlation between the noise and the observer’s response, and thus map which parts of the stimulus influence the observer’s performance. In other words, this spatial template reveals the subjective weighting of each stimulus sample for positional judgments. Thus, the classification image may be thought of as a behavioral receptive field²⁷.

Previous studies have used luminance (pixel) noise and reverse correlation to determine the classification image for position discrimination²³. However, these methods are not ideal for studying learning because they require many thousands of trials to compute the classifi-

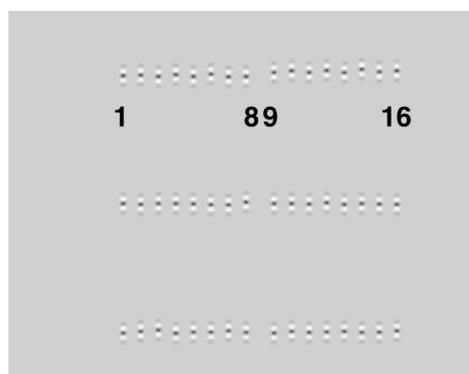


Figure 2 Stimuli with positional noise. The observer’s task was to indicate the position of the ‘test’ stimulus (top, middle or bottom). The top stimulus is misaligned: the right segment is higher than the left segment.

cation image, by which time the learning is complete. Here we used positional noise with relatively few moving parts (only 16 stimulus samples, rather than thousands of pixels) and an efficient linear regression technique, enabling us to measure the classification image with just 750 trials, and to document the changes in the classification image on the time scale of learning.

In this second experiment, we used both low and high levels of external noise. High noise is used because it ‘swamps’ the additive internal input noise. However, high and low noise do not always have the same effect on performance²⁸, and, to the extent that it is possible to obtain classification images in low noise, this provides a closer match to task performance in the absence of external noise and thus may provide some insights into learning with no noise. Two observers (LM and KP) were tested with low positional noise (a standard deviation (s.d.) of 0.17 arcmin), and the other two (LN and SG) were tested with high noise (s.d. = 0.67 arcmin).

Our classification image calculations show that the observer’s spatial templates did indeed change gradually over the four learning sessions (Fig. 5); only the first (s1) and the last (s4) sessions are shown for clarity. The abscissa in Figure 5 shows the patch position (corresponding to the patch positions shown in Fig. 2). Patches 8 and 9 are the inner end of each of the two target segments. The main result is that for each observer the template was broader after practice (blue) than at the start (red). In general, the gains of patches of the left segment and the right segment increased in opposite directions (becoming more negative and more positive respectively) with learning. After practice, more samples were being used by our observers for position judgments.

To provide a quantitative description of template, we fitted the template coefficient c_x with an exponential function: $c_x = A/2 * \text{sign}(x) * \exp(-|x|/w)$ (continuous lines in Fig. 5). In all four observers, the template tuning width w increased during the course of learning (Fig. 6a). Note that the larger the w , the broader is the template tuning.

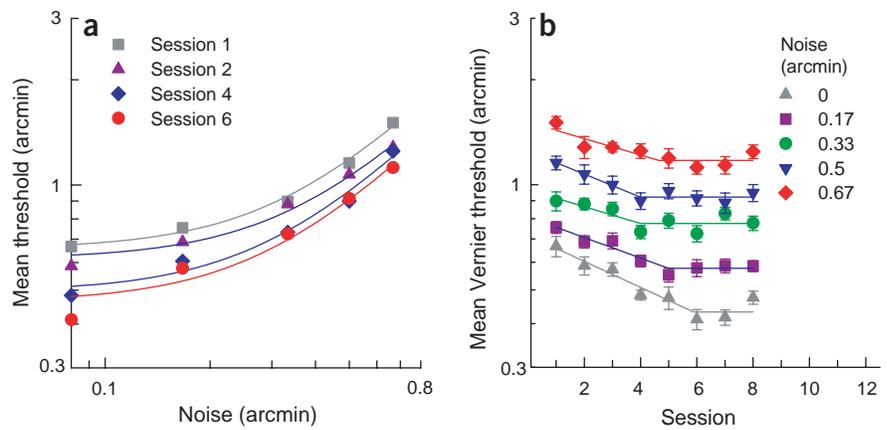


Figure 3 Effect of learning on position discrimination. **(a)** Mean position discrimination thresholds as a function of positional noise across training sessions. **(b)** The threshold data with standard errors are replotted with sessions; note that the ordinate scale and range are the same as that in **a**. A bilinear regression model was used to fit the data. The slopes (standard errors) of the regression lines were -0.037 (0.010), -0.029 (0.007), -0.024 (0.010), -0.033 (0.016) and -0.024 (0.033) for 0–0.67 min noise levels, respectively ($P = 0.003$ for 0.33 min noise; $P < 0.0005$ for the other four cases; the probabilities were obtained from Monte-Carlo simulations based on the mean and standard deviation of each datum using 2,000 randomly generated runs).

We compared our human observers’ templates to that of an ideal observer (*i.e.*, a machine that knows the precise details of the stimulus and task; see inset in Fig. 5) to assess the efficiency of the human template during the course of learning. The ideal observer uses all of the stimulus samples with equal weighting. Note that the templates used by human observers are far from ideal, even after practice; we are not surprised that the outer patches carried much less weight than those patches in the central part of stimulus because of overall orientation uncertainty.

We quantified template efficiency by computing the squared correlation between the human and ideal templates²⁴. Both ‘measured’ (circles and squares in Fig. 5) and ‘fitted’ templates (continuous lines in Fig. 5) were used to calculate template efficiency. We quantified human efficiency by calculating the squared ratio of ideal threshold to human threshold. There was a robust and steady increase in both

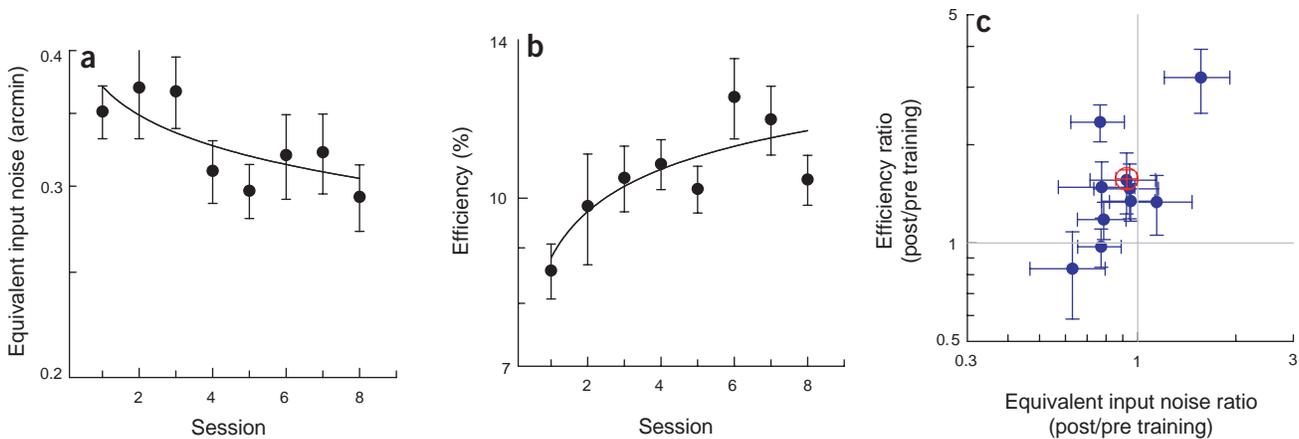


Figure 4 Effect of learning on equivalent noise and efficiency. **(a,b)** Equivalent input noise **(a)** and efficiency **(b)** across training sessions; the curve fitting was based on the mean thresholds in **Figure 3a**. **(c)** The post-training/pretraining ratio for individual observers. The error bars indicate one standard error. The red unfilled circle shows the mean ratios.

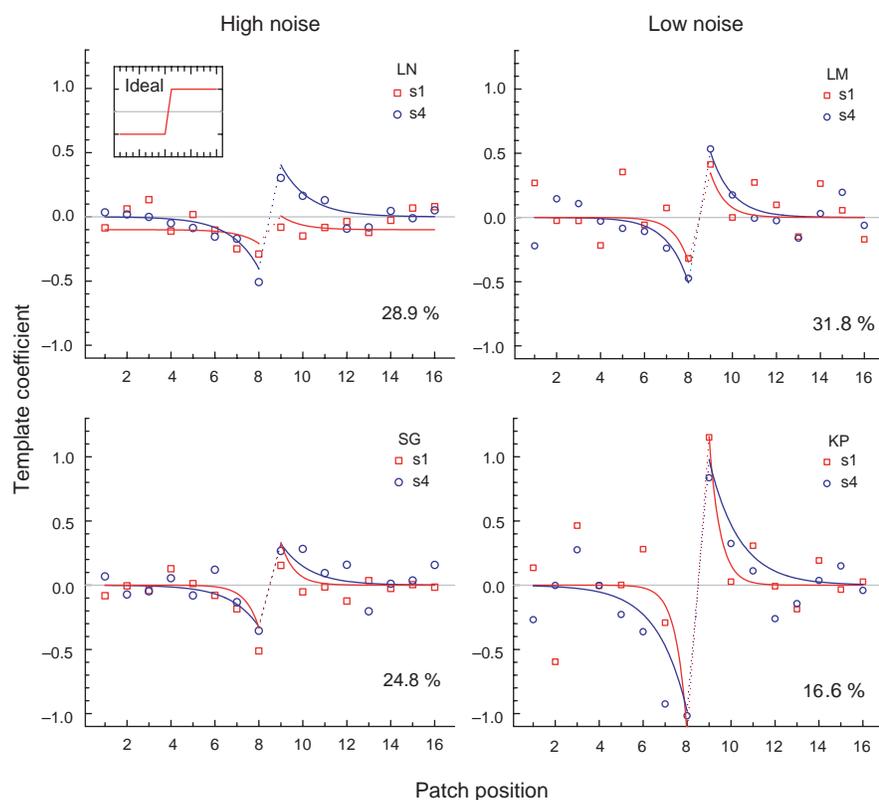


Figure 5 The changing templates across training sessions. The patch position is shown as in **Figure 2**. Note that patches 8 and 9 are the end of each segment. Because our jittered stimulus had zero mean, the calculations were based on the assumption that the mean coefficient of the last four patches in each segment is zero. (We calculated the template of the middle eight patches without zero mean adjustment, and obtained almost the same template as in the calculations with adjustment.) The mean improvement in threshold (%) from the last two sessions is shown at the lower-right corner. The fitting curve *s1* for observer LN was arbitrarily shifted by -0.1 for ease of viewing. For comparison, the template for the ideal observer is included in the top-left panel (inset).

template efficiency and human efficiency (**Fig. 6b**) across training sessions, reflecting the re-tuning of the human templates during learning (*i.e.*, they became more ideal). With practice, the weightings of the patches were adjusted, the tuning width was broader (**Fig. 6a**), and for two observers (LN and LM), the amplitude between patches 8 and 9 was substantially increased (**Fig. 5**). We note that whereas template efficiency was similar at low and high noise levels, human efficiency was much lower with low noise than with high (open triangles, **Fig. 6b**), presumably because of the dominant effects of internal noise.

The results of both experiments are summarized by comparing the improvement in performance (threshold ratio – ordinate) to the improvement in efficiency (efficiency ratio – abscissa) for each observer (**Fig. 6c**). The data (gray open circles) from experiment 1 are replotted in the figure for comparison. The gray line shows the improvement of thresholds predicted by the increased efficiency (the threshold ratio is the square root of the efficiency ratio (equation 1)). For both experiments, the improvement in performance falls close to the prediction. Interestingly, for three of the four observers in experiment 2, the performance improvement is also well predicted by the improvement in their template efficiency (solid symbols). The one clear outlier (shown by the red solid circle) is observer LN, who showed about a 5-fold change in template efficiency (using measured rather than fitted template) and only a 1.3-fold improvement in per-

formance. This observer had the least efficient template initially (see LN red symbols in **Fig. 5**), but surprisingly good performance.

The measured template (circles and squares in **Fig. 5** and red symbols in **Fig. 6b,c**) may give a biased estimate of the true template efficiency. Our simulations show that if the true template varies as a smooth function of separation (as shown by the lines in **Fig. 5**), adding variation on a trial-by-trial basis will bias the estimated template efficiency downward (*i.e.*, using the measured template would underestimate the true template efficiency when the template is extremely noisy, as was the case for LN in the initial data). Using the fitted template (smooth fits to the template in **Fig. 5**) provides a higher estimate of template efficiency, particularly in the initial trials (green symbols in **Fig. 6b,c**) and bring LN's template efficiency close to the prediction line. We note that if the true template is indeed bumpy, the process of smoothing will cause an upward bias in the efficiency. We suspect that the true value of the template efficiency lies between the smooth and the bumpy templates (red and green solid symbols in **Fig. 6b,c**).

DISCUSSION

Our results show that position discrimination can be significantly improved by repetitive training. The improvement in performance can be mainly attributed to increased efficiency; repetitive practice boosts the brain's ability to use the relevant stimulus information more efficiently. In contrast, equivalent input noise remains

almost unchanged (or very little changed) with learning.

The classification image shows that the increased efficiency (**Fig. 6c**) was mainly a consequence of re-tuning the spatial template for position discrimination, which enabled observers to use the stimulus information more efficiently. The increased amplitude and breadth of the template allows extraction of more samples (patches). Our results also provide an explanation for the narrowing of orientation tuning reported previously¹⁷. Our calculations (see Methods) show that the broadening of the template results in a decrease in the orientation tuning width, on average, from 32.0° to 26.5° (17%) over the course of learning. The change in the template shows that practice with feedback enables the observer to use the stimulus information most relevant to the task more efficiently.

Several previous studies have used luminance noise to reveal the template for position discrimination^{23,24}. We combined positional noise with relatively few moving parts with an efficient linear regression technique, enabling us to document changes in the classification image during learning of position discrimination. In contrast, several thousands of trials are needed to compute the classification image using the standard method with luminance noise and reverse correlation²⁹. Interestingly, our explanation for learning in position discrimination may hold across other visual tasks. In a simple Gaussian blob detection task³⁰, there was a marked increase in template correlation

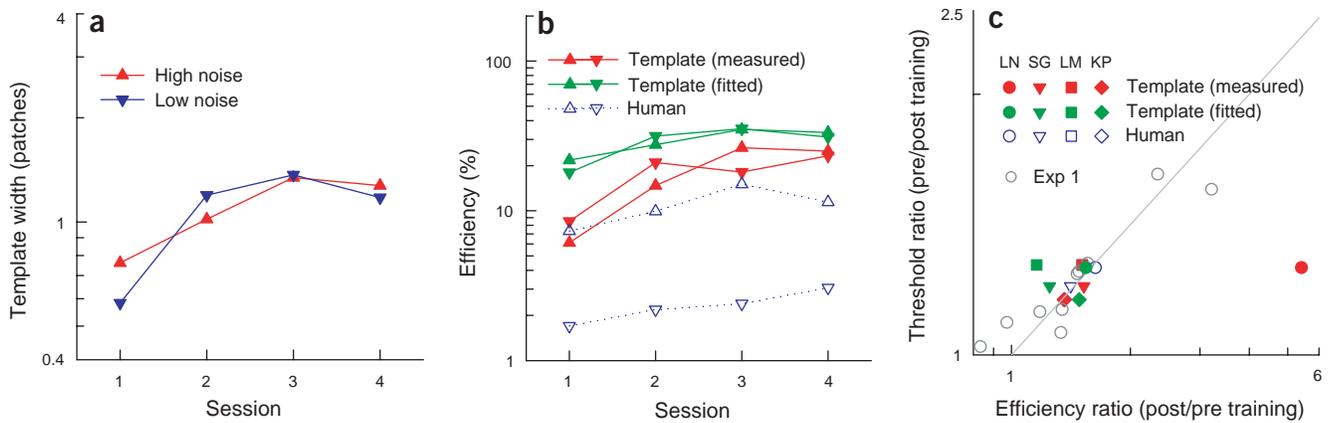


Figure 6 Template retuning. (a) Template tuning width from an exponential fit to the template in each session. (b) Template and human efficiency across sessions. We used both ‘measured’ (circles and squares in Fig. 5) and ‘fitted’ templates (continuous lines in Fig. 5) for the calculations of template efficiency. The true value of the template efficiency is expected to lie between the measured and the fitted templates. The upright and inverted triangles show the data for high and low noise, respectively. (c) The comparison of improvement between threshold and efficiency. The data (gray open circle) from experiment 1 is replotted in the figure; note that the ratios for experiment 1 are based on the first and the mean of the last three sessions, and the threshold ratio represented the mean threshold ratio of three high noise levels (0.33, 0.5 and 0.67 arcmin). The ratios for experiment 2 were based on the means of the first and the last two sessions.

across a block of trials, and it has been reported that the template for face recognition also changes during learning²⁹.

There is evidence for neural plasticity in the adult visual nervous system. For example, receptive fields in cortical area V1 of adult cats reorganize following abnormal visual experience because of retinal injury^{31,32}. Learning effects can be observed in evoked potentials of brain activity in humans^{33,34}. However, the evidence for low-level neural effects in perceptual learning is mixed. For example, it has been suggested that orientation tuning functions are sharpened after orientation discrimination practice in the primary visual cortex of monkeys³⁵. Other investigators, however, have failed to find physiological correlates of orientation discrimination learning and attribute the performance improvement to more central pooling and decision processes³⁶. Brain imaging has been used to demonstrate a decrease in neuronal activity in the task-related cortical regions following practice^{37,38}; this could occur through increased efficiency of processing and the narrowing of neuronal responses to the specific trained tasks but might also reflect top-down influences.

To summarize, we have characterized the limits, time course and mechanisms of improvement in position discrimination in normal foveal vision. It is clear that in vision, as in life, practice is effective in improving performance. Learning operates at all noise levels as an improvement in the efficiency with which the observer uses the positional information in the stimulus. This improved efficiency is a consequence of re-tuning the observer’s perceptual template, such that it becomes closer to the ideal template required by the perceptual task.

METHODS

Experiment 1. The stimulus comprised two line segments with a 17-arcmin gap between the two segments (Fig. 2). Each segment consisted of eight Gabor patches (carrier spatial frequency, 10 c.p.d.), and the patch separation was 10.67 arcmin. The patches were constructed to have an aspect ratio of 1/3: the Gaussian envelope standard deviation was 1.25 and 3.75 arcmin for the horizontal and vertical orientations, respectively. Positional noise was produced by distributing the position of each Gabor patch in the vertical direction according to a Gaussian probability function. The average offset of each jittered segment was made to be zero by uniformly shifting the eight patches. One of the

advantages of using zero mean noise is that no ‘wrong’ feedback would be provided to observers. Previous studies have shown that manipulated feedback can modify the learning patterns³⁹. An offset was produced by randomly shifting the right segment up or down. Observers were tested and trained at five noise levels (including zero); the positional noise was barely visible at 0.17 arcmin. The three neighboring stimuli were separated by 80 arcmin vertically. The stimuli were presented on a flat 21-inch Sony F520 monitor at a 90-Hz refresh rate. The mean center luminance of the stimuli was 55 cd/m², and the contrast of each Gabor patch was 99%. The monitor was viewed directly at a distance of 4 m.

A three-alternative forced-choice protocol was used to measure position discrimination in experiment 1. On each trial, the position of the misaligned stimulus was randomly chosen (top, middle or bottom). The observer’s task was to indicate the position of the ‘test’ stimulus. Stimuli remained on the monitor screen until the observer had given the response. Trial-by-trial feedback was provided. Observers had their heads steadied by a chin rest and forehead bar. A modified interleaved staircase method was used to track the individual thresholds²². The experimental trials were divided into triplets: three correct responses decreased the offset magnitude by a unit step, two correct responses in a triplet made the position offset unchanged and only one or zero correct response increased the offset by two unit steps. The staircases started with about double the predicted threshold and converged to 72%. A Weibull analysis (with a free-floating exponent) was performed to fit the psychometric curve with the response data. The position discrimination threshold was defined as the offset at which 66% correct responses was obtained. A session consisted of 750 responses (150 responses for each noise setting) in 2.5 h (including rest breaks). The task was self-paced and a break was given whenever the subjects requested one.

A positional averaging model²⁰ was used to quantify the effects of external positional noise (σ_e) on the threshold (σ_{th}):

$$\sigma_{th}^2 = 2d^2 \left(\frac{1}{k} - \frac{1}{n} \right) (\sigma_e^2 + \sigma_i^2) \quad (1)$$

where k denotes the number of samples extracted, n is the total number of samples and σ_i is the equivalent input noise. The term $1/n$ is present in equation 1 because of the zero mean adjustment of each line. In this study, each segment consisted of eight Gabor patches and hence n was equal to 8. For 66% correct response probability, the detectability (d') was 1.1 (ref. 40). Sampling efficiency (E) was defined as:

$$E = \frac{k}{n} \cdot 100\% \quad (2)$$

Ten normal adult observers were tested. All of them had corrected to normal visual acuity (20/20 or better) in each eye. Position discrimination was tested in the dominant eye only, the other eye being occluded with a standard eye patch. The data collection for each observer was completed in about 3 weeks. All observers were naive to the purpose of experiment and had no prior experience in psychophysical experiments. The experiments were undertaken with the understanding and written consent of each observer and all procedures were approved via institutional review at the University of California, Berkeley.

Experiment 2. Four new observers participated in this experiment; they were inexperienced with psychophysical experiments and naive to the purpose of this experiment. The stimulus was essentially the 'test' stimulus of experiment 1, briefly presented (200 ms) in the center of monitor. The observers' task was to rate the offset direction with a confidence level. There were four response categories: two 'up' and two 'down'. We used a rating scale technique for the classification image measurements^{24,41}. Two observers were tested with high positional noise (s.d. 0.67 arcmin) and the other two observers with low positional noise (s.d. 0.17 arcmin). We used the method of constant stimuli with five offset levels; the offset levels were chosen to be around the estimated threshold of the individual observer. Runs were paired in each experimental session, with stimuli in the second pass having identical positional noise configurations as the stimuli in the first pass. Each pass consisted of 375 responses, with 75 responses for each offset level. An alternative hypothesis, other than template improvement, could be random noise reduction. With this double-pass design⁹, we found that all four observers showed no significant changes in response consistency with learning.

To keep track of an observer's template across training sessions, we used a linearity assumption to compute the classification image c_x as follows:

$$Y = \sum_{x=1}^{16} c_x N_x + \varepsilon \quad (3)$$

where Y is observer response rating on a given scale, N is the noise on that trial and ε is the residual error. The subscript x specifies the 16 Gabor patches. Because each eight-patch segment had zero mean, the classification image of each segment is defined up to an arbitrary constant offset. Each half of the classification image was shifted vertically so that the outer four patches had zero mean. We calculated the classification image of the middle eight patches without zero mean adjustment, and obtained almost the same template as in the calculations with adjustment.

The orientation tuning function (see Discussion) was obtained by computing the template responses (inner products of the classification image and a Gaussian bar with standard deviation of 10.67 arcmin) at a range of orientations; the orientation at which half the maximum template response was reached represented the tuning width.

ACKNOWLEDGMENTS

This work was supported by research grants R01EY01728 and R01EY04776 from the National Eye Institute. The authors thank E. Bassett, S. Gepshtein, A. Popple, J. Saarinen, K. Young and E. Wong for their comments on an earlier version of the manuscript.

COMPETING INTERESTS STATEMENT

The authors declare that they have no competing financial interests.

Received 29 July; accepted 18 December 2003

Published online at <http://www.nature.com/natureneuroscience/>

1. Fiorentini, A. & Berardi, N. Perceptual learning specific for orientation and spatial frequency. *Nature* **287**, 43–44 (1980).
2. Vogels, R. & Orban, G.A. The effect of practice on the oblique effect in line orientation judgements. *Vision Res.* **25**, 1679–1687 (1985).
3. Fahle, M. & Morgan, M. No transfer of perceptual learning between similar stimuli in

- the same retinal position. *Curr. Biol.* **6**, 292–297 (1996).
4. McKee, S.P. & Westheimer, G.W. Improvement in vernier acuity with practice. *Percept. Psychophys.* **24**, 258–262 (1978).
5. O'Toole, A.J. & Kersten, D.J. Learning to see random-dot stereograms. *Perception* **21**, 227–243 (1992).
6. Ball, K. & Sekuler, R. A specific and enduring improvement in visual motion discrimination. *Science* **218**, 697–698 (1982).
7. Schoups, A.A. & Orban, G.A. Interocular transfer in perceptual learning of a pop-out discrimination task. *Proc. Natl. Acad. Sci. USA* **93**, 7358–7362 (1996).
8. Furmanski, C.S. & Engel, S.A. Perceptual learning in object recognition: object specificity and size invariance. *Vision Res.* **40**, 473–484 (2000).
9. Gold, J., Bennett, P.J. & Sekuler, A.B. Signal but not noise changes with perceptual learning. *Nature* **402** (1999).
10. Doshier, B.A. & Lu, Z.L. Perceptual learning reflects external noise filtering and internal noise reduction through channel reweighting. *Proc. Natl. Acad. Sci. USA* **95**, 13988–13993 (1998).
11. Doshier, B.A. & Lu, Z.L. Mechanisms of perceptual learning. *Vision Res.* **39**, 3197–3221 (1999).
12. Poggio, T., Fahle, M. & Edelman, S. Fast perceptual learning in visual hyperacuity. *Science* **256**, 1018–1024 (1992).
13. Fahle, M., Edelman, S. & Poggio, T. Fast perception learning in hyperacuity. *Vision Res.* **35**, 3003–3013 (1995).
14. Fahle, M. & Edelman, S. Long-term learning in vernier acuity: Effects of stimulus orientation, range and of feedback. *Vision Res.* **33**, 397–412 (1993).
15. Wang, H., Levi, D.M. & Klein, S.A. Intrinsic uncertainty and integration efficiency in bisection acuity. *Vision Res.* **36**, 717–739 (1996).
16. Levi, D.M., Klein, S.A. & Carney, T. Unmasking the mechanisms for Vernier acuity: evidence for a template model for Vernier acuity. *Vision Res.* **40**, 951–972 (2000).
17. Saarinen, J. & Levi, D.M. Perceptual learning in vernier acuity: What is learned? *Vision Res.* **35**, 519–527 (1995).
18. Levi, D. & Klein, S.A. The role of separation and eccentricity in encoding position. *Vision Res.* **30**, 557–585 (1990).
19. Green, D.M. Consistency of auditory detection judgments. *Psychol. Rev.* **71**, 392–407 (1964).
20. Zeevi, Y.Y. & Mangoubi, S.S. Vernier acuity with noisy lines: estimation of relative position uncertainty. *Biol. Cybern.* **50**, 371–376 (1984).
21. Pelli, D.G. The quantum efficiency of vision. In *Vision: Coding and Efficiency* (ed., Blakemore, C.) 3–25 (Cambridge Univ. Press, Cambridge, 1990).
22. Carkeet, A., Levi, D.M. & Manny, R.E. Development of Vernier acuity in childhood. *Optom. Vis. Sci.* **74**, 741–750 (1997).
23. Beard, B.L. & Ahumada, A.J. A technique to extract relevant image features for visual tasks. *Proc. Hum. Vis. Electronic Imaging III SPIE* **3299**, 79–85 (1998).
24. Levi, D.M. & Klein, S.A. Classification images for detection and position discrimination in the fovea and parafovea. *J. Vis.* **2**, 46–65 (2002).
25. Murray, R.F., Bennett, P. & Sekuler, A.B. Optimal methods for calculating classification images: weighted sums. *J. Vis.* **2**, 79–104 (2002).
26. Ahumada, A.J. Classification image weights and internal noise level estimation. *J. Vis.* **2**, 121–131 (2002).
27. Gold, J.M., Murray, R.F., Bennett, P.J. & Sekuler, A. Deriving behavioural receptive fields for visually completed contours. *Curr. Biol.* **10**, 663–666 (2000).
28. Lu, Z.L., Lesmes, L.A. & Doshier, B.A. Spatial attention excludes external noise at the target location. *J. Vis.* **2**, 312–323 (2002).
29. Gold, J.M., Sekuler, A.B. & Bennett, P.J. Visualizing perceptual learning. *Cogn. Sci.* (in press).
30. Solomon, J.A. Noise reveals visual mechanisms of detection and discrimination. *J. Vis.* **2**, 105–120 (2002).
31. Gilbert, C.D. & Wiesel, T.N. Receptive field dynamics in adult primary visual cortex. *Nature* **356**, 150–152 (1992).
32. Chino, Y.M., Kaas, J.H., Smith, E.L.I., Langston, A.L. & Cheng, H. Rapid reorganization of cortical maps in adult cats following restricted deafferentation in retina. *Vision Res.* **32**, 789–796 (1992).
33. Skrandies, W., Lang, G. & Jedynak, A. Sensory thresholds and neurophysiological correlates of human perceptual learning. *Spat. Vis.* **9**, 475–489 (1996).
34. Skrandies, W., Jedynak, A. & Fahle, M. Perceptual learning: psychophysical thresholds and electrical brain topography. *Int. J. Psychophysiol.* **41**, 119–129 (2001).
35. Schoups, A., Vogels, R., Qian, N. & Orban, G. Practising orientation identification improves orientation coding in V1 neurons. *Nature* **412**, 549–553 (2001).
36. Ghose, G.M., Yang, T. & Maunsell, J.H. Physiological correlates of perceptual learning in monkey V1 and V2. *J. Neurophysiol.* **87**, 1867–1888 (2001).
37. Schiltz, C. *et al.* Neuronal mechanisms of perceptual learning: changes in human brain activity with training in orientation discrimination. *Neuroimage* **9**, 46–62 (1999).
38. Kassubek, J., Schmidtke, K., Kimmig, H., Lucking, C.H. & Greenlee, M.W. Changes in cortical activation during mirror reading before and after training: an fMRI study of procedural learning. *Brain Res. Cogn. Brain Res.* **10**, 207–217 (2001).
39. Herzog, M.H. & Fahle, M. The role of feedback in learning a vernier discrimination task. *Vision Res.* **37**, 2133–2141 (1997).
40. Wickens, T.D. *Elementary Signal Detection Theory* (Oxford Univ. Press, New York, 2002).
41. Ahumada, A.J. & Lovell, J. Stimulus features in signal detection. *J. Acoust. Soc. Am.* **49**, 1751–1756 (1971).