Perceptual learning in contrast discrimination and the (minimal) role of context

Cong Yu
School of Optometry, University of California, Berkeley, CA, USA, & Institute of Neuroscience, Chinese Academy of Sciences, Shanghai, China

Stanley A. Klein
School of Optometry and Helen Wills Neuroscience Institute, University of California, Berkeley, CA, USA

Dennis M. Levi
School of Optometry and Helen Wills Neuroscience Institute, University of California, Berkeley, CA, USA

Unlike most visual tasks, contrast discrimination has been reported to be unchanged by practice (Dorais & Sagi, 1997; Adini, Sagi, & Tsodyks, 2002), unless practice is undertaken in the presence of flankers (context-enabled learning, Adini et al., 2002). Here we show that under experimental conditions nearly identical to those in the no-flanker practice experiment of Adini et al. (2002), practice significantly improved contrast discrimination. Moreover, in a separate experiment, we found that practice without flankers can improve contrast discrimination to a level only reached with flankers in Adini et al. (2002), but further practice with flankers produces no further improvement of contrast discrimination. These results call into question whether the "context-enabled learning" proposed by Adini et al. (2002) is different from regular contrast learning without flankers. In separate experiments, we found that contrast learning is tuned to spatial frequency, orientation, retinal location, and, unexpectedly, contrast. We also replicated Sagi, Adini, Tsodyks, and Wilkonsky's (2003) more recent finding that no regular contrast learning occurs if reference contrasts are randomly interleaved (contrast roving), and further demonstrated that flankers have no effect on contrast learning under contrast roving, another piece of evidence equating "context-enabled learning" to regular contrast learning. The contrast specificity of learning and the lack of learning under contrast roving provide new evidence in favor of a multiple contrast-selective channels model of contrast discrimination, and against saturating transducer models and multiplicative noise models.

Keywords: perceptual learning, contrast discrimination, context, roving

Introduction

A wide range of visual functions, from the activity of single neurons to higher level pattern and brightness perception, are critically dependent on stimulus contrast (Shapley, 1986). One of the fundamental tasks of the human visual system therefore is to detect and discriminate changes in contrast. In a recent letter to *Nature*, Adini, Sagi, and Tsodyks (2002) reported that contrast discrimination for a Gabor stimulus does not improve with practice (their Figure 2b; also see Dorais & Sagi, 1997), unless the training is conducted with the Gabor stimulus flanked by additional pairs of identical Gabors (their Figure 2a). Adini et al. (2002) thus named their newly discovered flanker-induced contrast learning "context-enabled learning," and proposed a context-enabled neural plasticity model on the basis of Hebbian and anti-Hebbian synaptic learning rules to explain this effect.

Numerous studies have reported that practice can improve performance in a variety of visual tasks (Fiorentini & Berardi, 1981; Saarinen & Levi, 1995; Dorais & Sagi, 1997; Fahle, 1997; Matthews & Welch, 1997; Fine & Jacobs, 2000; Sowden, Rose, & Davies, 2002), such as discrimination of orientation, spatial frequency, phase, Vernier offset, etc., even for adults with amblyopia (Levi & Polat, 1996; Levi, Polat, & Hu, 1997; Li & Levi, in press). The failure to learn contrast discrimination (without flankers) appears to be an interesting exception. Adini et al. (2002) interpreted their "nonlearning" phenomenon as an indication that "not all the activations of the primary visual cortex result in long-term modifications." Their results are surprising because Vernier acuity, which is readily improved with learning (Li, Levi & Klein, 2004), is often modeled as a form of contrast discrimination (Hu, Klein, & Carney, 1993; Levi, Klein, & Wang, 1994).

However, a careful inspection of the Adini et al. (2002) data raises questions about their claims of no-learning and context-enabled learning in contrast discrimination. For example, in their Figure 3, which shows contrast discrimination at a reference contrast of 0.50, the mean discrimination threshold before (and after) practice is around 0.20, or a Weber fraction ($\Delta C/C$) of 0.20/0.50 = 0.40. However, the Weber fraction of contrast discrimination is known to be around 0.10–0.20 (Legge, 1981; Legge & Kersten, 1983). In our earlier experiments (Yu & Levi, 1997, 2000), contrast thresholds of experienced observers under stimulus conditions similar to

those in Adini et al. (2002) were about 0.06 to 0.08 at a reference contrast of 0.40, or a Weber fraction of 0.15–0.20, about half that of inexperienced observers in Adini et al. Therefore, we suspected that contrast discrimination could be significantly improved or learned given sufficient practice without the help of flankers. Moreover, we suspected that if contrast discrimination could indeed be learned, improved contrast discrimination with the presence of flankers in Adini et al. (2002) might actually reflect regular learning and have little to do with the flankers.

In this work, we examined both claims (i.e., no-learning and context-enabled learning) made in Adini et al. (2002) and were unable to replicate either. Some of our data also address new issues raised during our communications with Sagi and his colleagues (see Sagi, Adini, Tsodyks, & Wilkonsky, 2003). Experiments I and II are a direct examination of the Adini et al. (2002) main claims of no contrast learning and context-enabled learning. We also examined several new issues that are central to learning. Experiment III probes the specificities of contrast learning to stimulus dimensions, retinal location, and eye of origin. Experiment IV uses contrast-roving methods (randomly interleaved staircases with different base contrasts) to show that contrast learning likely reflects improvement at a decision stage, rather than low-level cortical neural plasticity, regardless of whether the observers practice with or without stimulus context or flankers. Brief reports of results in this work were presented in the 2002 and 2003 annual Vision Sciences Society meetings in Sarasota, Florida.

Methods

Observers and apparatus

More than 30 observers, mostly University of California–Berkeley undergraduate students, with normal or corrected-to-normal vision participated in different phases of the study. Most were new to psychophysical experiments and unaware of the specific purposes of the experiments, though they were informed that the general goal of this study was to investigate whether visual performance could be improved by practice.

The stimuli in Experiments I–III were generated by a VisionWorks program (Vision Research Graphics, Inc., Duham, NH) and presented on a 21-inch Image System Max21L monochrome monitor (1024 x 512 resolution, 0.28 mm (*H*) x 0.41 mm (*V*) pixel size, 117-Hz frame rate, 50 cd/m^2 mean luminance, and 3.8° x 3. 0° screen size at the 5.64-meter foveal viewing distance). Luminance of the monitor was made linear by means of a 15-bit look-up table. The stimuli in Experiments IV were generated by a WinVis program (Neurometrics Institute, Berkeley, CA) and presented on a 19-inch Dell UltraScan P991 color monitor (640 x 480 resolution, 0.59 mm (H) x 0.54 mm (*V*) pixel size, 60-Hz frame rate, 60 cd/m^2 mean luminance, and 4.0°

x 3.0° screen size at the 5-meter foveal viewing distance). Luminance of this monitor was made linear by means of an 8-bit lookup table. Experiments were run in a dimly lit room.

Stimuli and procedure

The test stimulus was a Gaussian windowed sinusoidal grating (Gabor patch). Under most stimulus conditions (foveal viewing), this Gabor patch had a spatial frequency of 6 cycles per degree (cpd), and the standard deviation of the Gaussian envelope was $\sigma = 0.12°$. In Experiments II and IV, additional flanking stimuli were used, either simulated by increasing the length of the pedestal, or by adding additional pairs of Gabor patches. Further details will be provided in "Results."

Contrast thresholds were measured with a temporal two-alternative forced-choice (2AFC) staircase procedure. In Experiments I~III, staircases at different reference contrasts were run non-interleaved. In Experiments IV, staircases at all reference contrasts were run randomly interleaved (contrast roving). Within a staircase, the test and reference stimuli were separately presented in the two stimulus intervals (≈103 msec each) in a random order separated by a 600-msec interstimulus interval. Each stimulus interval was accompanied by an auditory tone of the same duration to reduce temporal uncertainty. The observers' task was to judge which stimulus interval contained the higher contrast Gabor. Each trial was preceded by a 6.3' x 6.3' fixation cross which disappeared 100 msec before the beginning of the trial. Auditory feedback was given on incorrect responses. Each staircase consisted of four preliminary reversals and eight experimental reversals when run non-interleaved, or two preliminary reversals and six experimental reversals when run interleaved. The step size of the staircase was 0.05 log units. A classical 3-down-1-up staircase rule was followed, which resulted in a 79.4% convergence level of the staircase. The geometric mean of the experimental reversals was taken as the contrast threshold for each staircase run.

Experiment I. Perceptual learning of contrast discrimination at multiple contrasts

We first reran the Adini et al. (2002) control experiment (practice without flankers) to check whether their surprising results (no contrast learning, their Figure 2b) could be replicated. We had five inexperienced observers practice contrast discrimination at four reference contrasts from 0 to 0.63. The foveal Gabor (*SF* = 6 cpd, $\sigma = 0.12°$, where σ is the *SD*) and the levels of reference contrasts were identical to those used in the Adini et al. control experiment. The σ used by Adini et al. (2002; see their Figure 1) is sd/√2. Thus their condition $\sigma = \lambda$ was also what we used. Each experimental segment contained four

blocks of trials (or four staircases, each staircase is treated as one block) for four reference contrasts (0, 0.30, 0.47, and 0.63) measured in ascending order. The observers practiced 3-4 experimental segments each session for 2 hours over four to five days.

Four of our five observers, with the exception of VF, showed significant session-by-session improvement of contrast discrimination (reduced thresholds) (Figure 1a), indicating that practice indeed improves contrast discrimination without the help of flankers. The individual data and the means across observers in the first and last sessions are plotted respectively in Figure 1b as TvC (threshold vs. contrast) functions to summarize the learning effects. The ratios of mean pre/post training thresholds (Figure 1b) are 0.44, 0.63, 0.50, and 0.53, respectively, for reference contrasts at 0, 0.3, 0.47, and 0.63, which are

comparable to the mean ratios of pre/post flanker training thresholds in Adini et al. (2002) (1, 0.47, 0.55, & 0.54, respectively, for the same contrasts, calculated from their Figure 2a with thresholds at some contrasts extrapolated from the TvC functions). The exception is at the zero contrast (detection) where no learning was shown in the Adini et al. observers. Therefore, practice without flankers not only improves contrast discrimination, it does this as well as practice with flankers.

Some specifics are worth mentioning. First, Figure 1a shows that learning did not reach an asymptotic level after 4-5 sessions of practice in most cases, suggesting room for further improvement of contrast discrimination. Second, the amount of learning is not dependent on the initial thresholds. For instance, observers CC and CN showed contrast learning as significant as the other two (JC & TG)
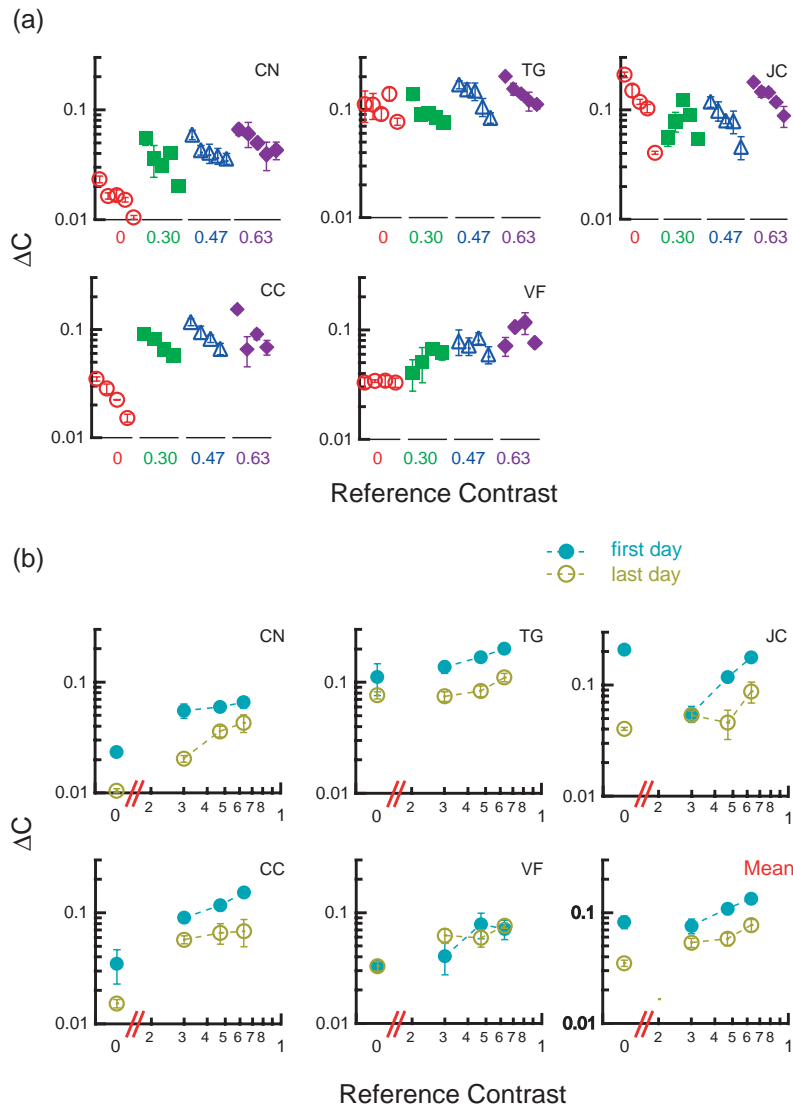

Figure 1. Perceptual learning of contrast discrimination at four stimulus contrasts. (a). The session-by-session data. Data points of each color indicate session-by-session threshold changes (from left to right) at one specific contrast. (b). Individual pre- and post-training (first and final day) TvC functions and their means. Each datum in individual plots is the geometric mean of thresholds within the session.

even though they started with lower initial thresholds. Third, as Figure 1b suggests, there is no clear pattern of slope changes of the TvC functions across observers. Fourth, the mean pre-training thresholds of our observers were 0.09, 0.08, 0.12, and 0.15 for reference contrasts at 0, 0.3, 0.47, and 0.63, respectively, while in Adini et al. (2002) observers' corresponding pre-training thresholds were approximately 0.07, 0.17, 0.22, and 0.24 (extrapolated from their Figure 2a). Our observers actually had lower pre-training thresholds to start with, except at the zero contrast. Therefore it is unlikely that the different learning results between the two studies under otherwise nearly identical conditions are due to the inhomogeneity of observer pools.

For learning at zero reference contrast (contrast detection), similar effects have been reported previously in both the fovea and the periphery (De Valois, 1977; Mayer, 1983; Sowden et al., 2002). Our results extend the findings of contrast learning to suprathreshold contrast discrimination and dispute the no-learning claims made by Dorais and Sagi (1997) and Adini et al. (2002). Because practice produces comparable contrast learning with and without the presence of flankers, and because Adini et al. (2002) observers in the flanker training condition never had previous training without flankers, we suspect that the so-called "context-enabled learning" is at the very least confounded by regular contrast learning. It could even *be* regular contrast learning, and context or flankers may actually be irrelevant, at least under experimental conditions very similar to those used in Adini et al. (2002)!

Experiment II. Context-enabled learning after exclusion of regular contrast learning

To further examine whether "context-enabled learning" (Adini et al., 2002) is confounded by regular contrast learning, we first excluded the effect of regular contrast learning through practice without flankers and then measured any possible further improvement of contrast discrimination under flanker conditions. We reasoned that having "squeezed out" regular contrast learning, any further improvement would reflect pure context-enabled learning effects.

Seven new observers (all inexperienced except AP and YC, who were experienced but had not run similar tasks for a few years) practiced contrast discrimination for the same Gabor as in Experiment I (*SF* = 6 cpd, σ = 0.12°, contrast = 0.5) for 4-6 sessions. This initial practice without flankers greatly reduced contrast thresholds (Figure 2b). The mean threshold ratio of the last and first sessions was 0.33 ± 0.09, a three-fold improvement in threshold! After this phase of practice most regular contrast learning was effectively squeezed out as evidenced by the asymptote in performance. The same observers then practiced contrast discrimination with flankers (Figure 2a) for another four sessions. We

reasoned that this additional practice with flankers should enable mostly pure "context enabled learning" because of the initial exclusion of regular contrast learning. Finally the observers repeated no-flanker contrast discrimination for another session to evaluate the amount of context enabled learning. A comparison of contrast thresholds before and after practice with flankers, however, reveals little threshold reduction or "context-enabled learning" (Figure 2c). The mean threshold ratio of pre- and post-flanker training was 0.94 ± 0.14, suggesting no significant change of contrast discrimination. The only exception was FF (Figure 2c), who shows improvement after flanker training. However, this observer had very high overall thresholds, and her threshold at the end of the initial training stage (Figure 2b) was higher than most other observers' beginning thresholds. It is likely that her threshold reduction (Figure 2c) may have resulted because her regular learning of contrast discrimination was not yet complete.

These results suggest that "context-enabled learning" is inseparable from regular learning. After the observers have learned contrast discrimination to asymptotic performance without flankers, there is no more learning with flankers. Together Experiments 1 and 2 demonstrate that flankers are not necessary for learning contrast discrimination.

Experiment III. The specificities of contrast learning: Stimulus dimensions, retinal location, and eye of origin

In this experiment, we surveyed some basic properties of contrast learning: whether contrast learning is specific to certain stimulus dimensions such as orientation and spatial frequency, and whether it can be transferred to different retinal locations and to the untrained eye. The specificities of learning to stimulus dimensions are often interpreted as indications of V1 involvement in perceptual learning. For example, in a recent case relevant to our study, Sowden et al. (2003) used stimulus specificities to argue that perceptual learning of peripheral contrast detection takes place in a subpopulation of V1 layer 4 cells. Such arguments are open to question because position and orientation specificities can be shown as late as in cells in the inferotemporal cortex (Vogels & Orban, 1994). Mollon and Danilova (1996) also suggested that "The learning may occur at a central site, and what the subject may be learning about are the local idiosyncrasies of his retinal image, of his receptor mosaic, and of the wiring of his visual system."

Specificity to stimulus dimensions

We first had three inexperienced observers practice contrast discrimination for the same Gabor stimulus (*SF* = 6 cpd, orientation = 0°) used in Experiment I at 0.47 contrast, and examined whether contrast learning could be
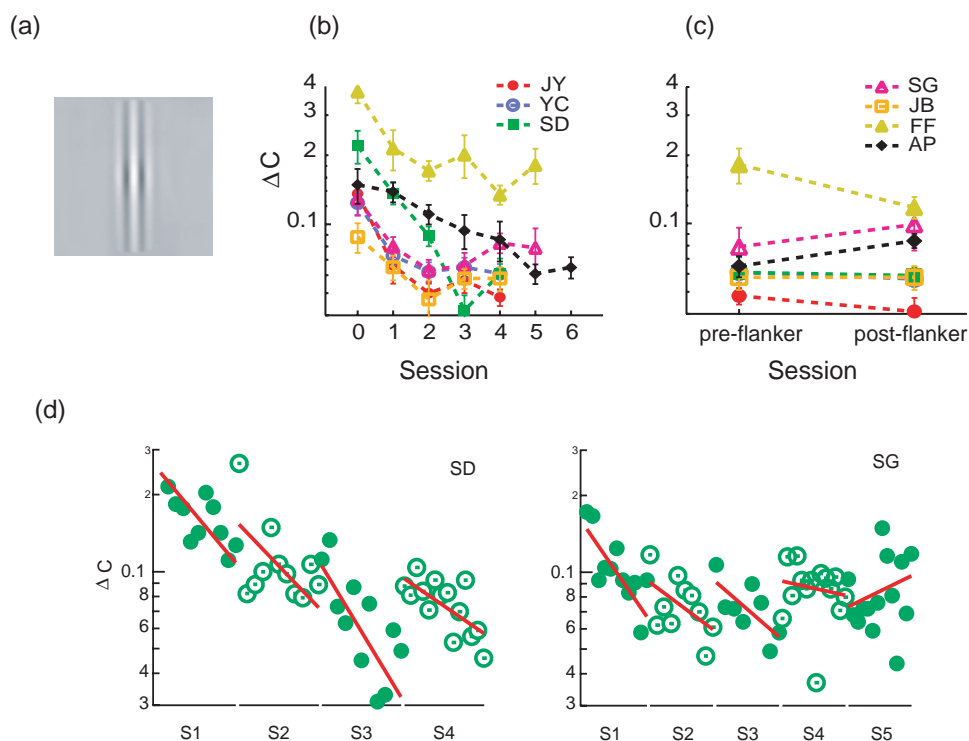
Figure 2. (a). The Gabor stimulus with flankers. The flankers were realized by increasing σ_y of the pedestal ($\sigma_y = 6 \ \sigma_x$). (b). The initial training course of contrast discrimination without flankers. The geometric mean of the first 4 runs of the first session is taken as the pre-practice (Session 0) threshold, and the geometric mean of the last 4 data points of each day is taken as that day's threshold after that day's practice. (c). Contrast discrimination measured with no flankers before and after practice with flankers (pre-flanker and post-flanker). Each observer's pre-flanker threshold is his/her last-session threshold in (b). (d). Examples of detailed training courses in two observers. Each datum represents the threshold from one staircase run. The red linear regression line shows threshold changes with the same training session. It is interesting that most of the learning from the previous training session is lost at the beginning of the new session. This trend is consistent across all our subjects and has also been reported previously in Levi et al. (1997).

transferred to other spatial frequencies (3 & 12 cpd), orientation (90°) and contrasts (0.30 & 0.73, about 0.2 log units above and below the learned contrast). The observers' contrast thresholds at all spatial frequencies, orientations, and contrasts were first measured in one session to set the pre-training baselines. After this, the observers practiced contrast discrimination at 6-cpd spatial frequency, 0° orientation, and 0.47 contrast for 2–3 sessions (approximately 25–35 staircase runs). Finally, the same pre-training measurements were repeated in another session to gather the post-training threshold data.

Figure 3 depicts the pre- and post-training data at the trained (indicated by red arrows in the mean plots) and untrained conditions. A similar pattern is evident across all three specificity measurements, though less consistent in the orientation case. That is, although contrast discrimination in all conditions is more or less improved after training, the trained conditions had the most improvement by approximately a factor of 2 (TW was an exception who showed equal improvement at both trained and untrained orientations). The mean pre/post threshold ratio at the trained condition is about 0.55. The same ratios at untrained spatial frequencies, orientation, and

contrasts are 0.80, 0.72, and 0.74, respectively. These data suggest that about half the learning of contrast discrimination is stimulus-dimension specific and is not transferred to untrained dimensions.

Specificity to retinal location and eye

Another three inexperienced observers were tested in the visual periphery to investigate whether perceptual learning of contrast discrimination is specific to retinal location and eye. Contrast discrimination was practiced for two sessions for a peripheral Gabor (1.5 cpd, $\sigma = 0.68°$, contrast = 0.47) presented in the lower 5° eccentricity of the left visual field of one eye (Figure 4a) (JY and PH's dominant right eyes and EL's nondominant left eye) with the other eye covered by a hand-held white translucent pad. The observers were sitting 1.64 meters from the monitor screen, one-fourth the normal (foveal) viewing distance. Pre- and post-training thresholds (mean of 4-5 staircase runs each in pre- and post-training sessions) were then compared to reveal whether contrast learning at the trained location (or eye) can be transferred to the untrained retinal location (the upper 5° eccentricity) of the same eye, or the same retinal location of the untrained fellow eye.
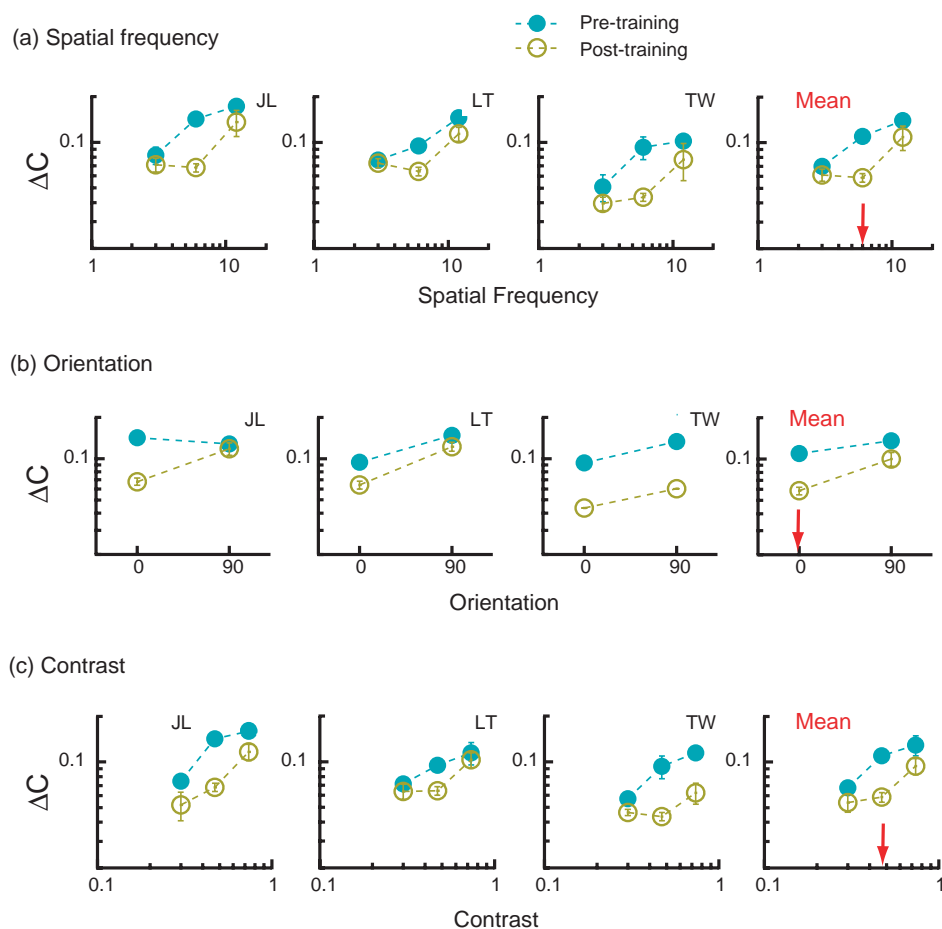
Figure 3. Transfer of contrast learning to untrained stimulus dimensions. Spatial frequency (a); orientation (b); and contrast (c). Contrast thresholds before and after practice at trained and untrained conditions are presented. The trained conditions are indicated by red arrows in the mean plots. Each datum indicates the mean threshold of 3-4 staircase runs.

Figure 4b shows that learning is specific to the practiced retinal location (lower visual field) with little transfer to the unpracticed location (upper visual field) of the same eye. However, Figure 4c indicates that two out of three observers showed complete interocular transfer of learning.

The stimulus specificities and retinotopic properties of contrast learning revealed in Figures 3 and 4 resemble the tuning and retinotopic properties of V1 neurons. Even the surprising specificity of contrast learning to stimulus contrast (Figure 3c) could reflect the limited operating range of V1 neurons (e.g., Sclar, Maunsell, & Lennie, 1990). However, as suggested earlier, a link between V1 neural plasticity and contrast learning need not be the only explanation, and the same data can be equally well explained by higher level visual processes. Therefore, although these data give a more comprehensive description of the contrast-learning phenomenon we are studying, they are not able to effectively limit the locus of contrast learning. The contrast roving method detailed in Experiment IV is aimed at further examining the locus puzzle.

Experiment IV. Contrast learning with roving contrasts

Contrast discrimination is determined by the low-level gain of spatial filters or visual neurons, as well as by high-level decision processes (Legge & Foley, 1980). Practice could enhance contrast discrimination either by increasing the gain (signal to noise) of visual cortical neurons, or by improving the observer's visual decision efficiency, in that the observers learn to attend to the responses of the most relevant neurons or filters. Indeed, after confirming part of our contrast learning data (improvement after training at a single contrast, like those shown in Figure 3), Sagi et al. (2003) suggested that contrast learning at a single contrast may result from optimized discrimination strategies (or improved templates in other words) for stimuli at that specific contrast level.[1] Furthermore, they suggested that this possibility could be tested by manipulation of contrast uncertainty. If contrast learning is interrupted by contrast uncertainty, it would suggest that contrast learning results from optimized (high-level) discrimination strategies, rather
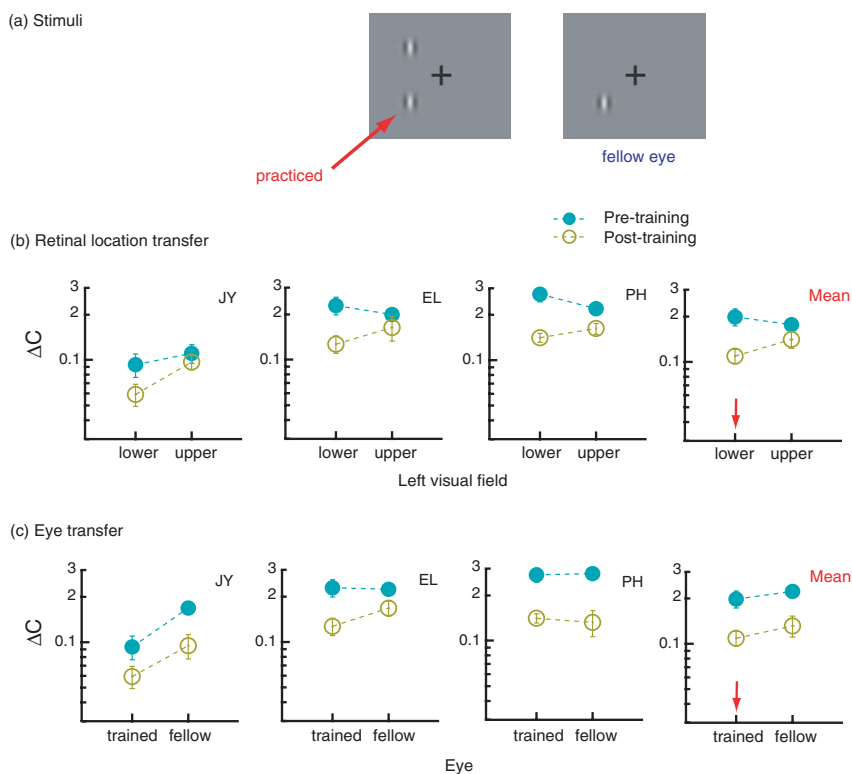
Figure 4. Transfer of learning across retinal location and eye of origin. (a). Stimuli at trained location in the lower left visual field and untrained locations in the upper left visual field and in the fellow eye. (b). Contrast thresholds before and after practice at trained (lower visual field, indicated by the red arrow) and untrained (upper visual field) retinal locations of the same eye. (c). Contrast thresholds before and after practice at trained (indicated by the red arrow) and untrained eyes for the same retinal location (lower visual field).

than improved gain of (low-level) sensory mechanisms (see "Discussion" for clarification).

On the other hand, Adini et al. (2002) and Sagi et al. (2003) suggest that unlike regular contrast learning, contrast learning with flankers or "context-enabled learning" reflects low-level changes in the gain of cortical neurons. This is at odds with our finding in Experiment II that the so-called "context-enabled learning" is inseparable from regular contrast learning. In this experiment we followed the Sagi et al. (2003) logic to investigate whether contrast learning at multiple reference contrasts without flankers (regular contrast learning), as well as with flankers ("context-enabled learning"), are outcomes of changes in the gain of visual cortical neurons or in high-level decision processes. The first part was to replicate the Sagi et al. (2003) experiment, and the second part served as another powerful tool to validate our earlier conclusion that "context-enabled learning" is simply a variation of regular contrast learning.

Contrast learning without flankers under contrast roving

Four inexperienced observers practiced contrast discrimination for otherwise identical Gabor stimuli without flankers at 4 contrasts: 0.20, 0.30, 0.47, and 0.63.

Contrast thresholds were again measured with 2AFC staircases, but 4 staircases were randomly interleaved from trial to trial (contrast roving). Roving is the word first used in auditory studies by Berliner and Durlach (1973). Here it means randomly interleaved staircases for multiple reference contrasts. Under contrast roving, the reference contrast is not predictable from trial to trial (contrast uncertainty). Under contrast roving, observers are required to base their judgments on comparing the two presentations of a single 2AFC trial. They cannot build up a reference template because the template changes from trial to trial.

The red dashed curves and blue solid curves in Figure 5 show the TvC functions in the first and third training sessions, respectively, for each observer. In general, practice with roving contrasts produced no significant learning of contrast discrimination, in sharp contrast to the significant learning effects when staircases were run in blocks (Figure 1). Individually, though no learning is evident for observers SA and CG, IH showed some learning at lower contrasts, and JS had better performance at higher contrasts but worse performance at lower contrasts. These data are consistent with the contrast roving results by Sagi et al. (2003) and support their suggestion that contrast learning probably occurs at a more central decision stage where practice reduces uncertainty.
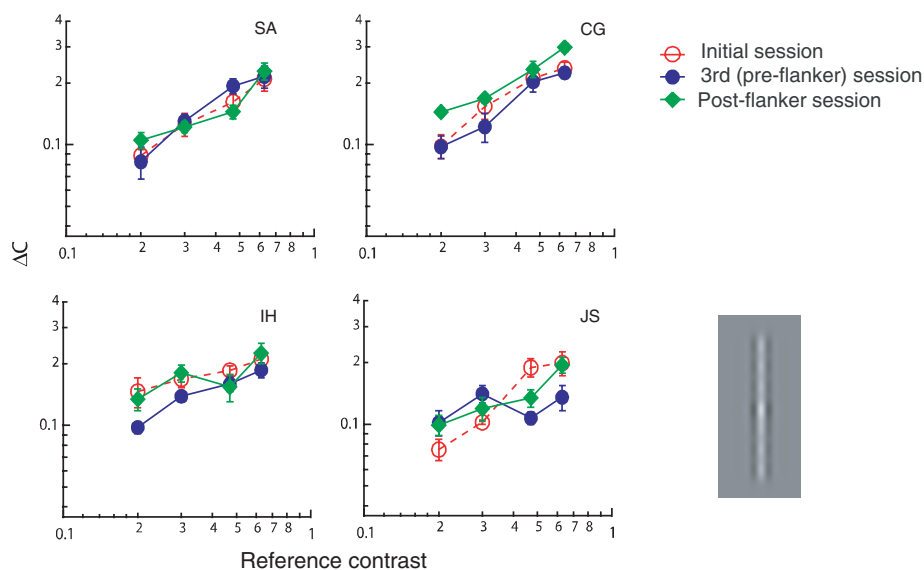
Figure 5. The effects of contrast roving on contrast learning with flankers. The red curves are each observer's initial TvC functions, the blue curves show the third (or pre-flanker) session TvC functions, and the green curves show post-flanker training TvC functions. The stimulus image shows the Gabor test as well as flankers.

Contrast learning with nonmatched flankers under contrast roving

The same four observers continued to practice contrast discrimination for the same Gabor test under contrast roving. However, this time the Gabor test was flanked by three additional pairs of Gabors with Gabor-to-Gabor spacing at 0.19° (Figure 5, the stimulus image).

Two flanker conditions were used. Under the fixed contrast condition, two observers (SA & IH) practiced with the flankers at a fixed contrast (0.40) that was not one of the test contrasts. This condition was similar to the Adini et al. (2002) "context-enabled learning" condition except that the reference contrasts was randomly interleaved. Under the second jittered contrast condition, the other two (CG & JS) practiced with flanker contrasts randomly jittered at one of the four reference contrasts (i.e., randomly set at either 0.2, 0.3, 0.47, or 0.63). We jittered the flanker contrast because we were concerned that a fixed flanker contrast might provide a reference to the reference stimuli even with roving pedestal contrasts. That is, the observers might be able to form a stimulus template at each contrast based on the contrast difference between the reference and the flankers. However, the data show conclusively that our concern was unwarranted because no significant difference was shown between the effects of fixed and jittered flanker contrasts (Figure 5). We did not carry out experiments in which the flank contrast was yoked to the reference because it changes the task to a detection task in which the judgment can be made in a single interval by comparing the target to the yoked flanks.

After three sessions of practice with the presence of flankers at either fixed or jittered contrasts, contrast discrimination for Gabors without flankers was re-tested (green solid curves). The logic for the current experiment is, if contrast learning in the presence of flankers is different from regular contrast learning, and is based on low-level gain changes of visual cortical neurons as proposed by Adini et al. (2002) and Sagi et al. (2003), it should not be disturbed by contrast roving that most likely would only influence visual decision making, and we should see significant contrast learning under contrast roving in these subjects. This, however, is not what the data show.

Figure 5 shows that under contrast roving, practice with flankers produced no improvement of contrast discrimination. In some cases it even made the performance of contrast discrimination worse. For example, practice with flankers raised CG's overall contrast thresholds (the green curve) to be even higher than the initial contrast thresholds (the red curve). It also reversed IH's previous learning at low contrasts obtained in Experiment III (the green curve now overlaps with the red initial curve at low contrasts), as well as partially reversed JS's previous learning at high contrasts (the post-flanker training green curve is now in between the pre-flanker training blue curve and initial session red curve).

These results demonstrate that contrast roving essentially kills contrast learning regardless of the presence of flankers. Thus, we suggest that any context-enabled learning (if it occurs) may not be low level at all and may share the same, more central, mechanisms with regular contrast learning. This is consistent with our results in Experiment II that context-enabled learning is inseparable from regular contrast learning.

Discussion

Our investigations lead to two simple conclusions regarding the Adini et al. (2002) report: First, practice improves contrast discrimination, as it does many other visual tasks. Second, context-enabled learning is probably nothing more than regular contrast learning. These conclusions oppose the two main claims made by Adini et al. (2002). In this discussion, we first examine the differences between our experiments and those of Adini et al. (2002). Then we consider the implications of the two novel findings of the contrast tuning of learning and the difficulty of learning under roving conditions. We also discuss the distinction between early and late learning.

Some differences between our study and the Adini et al. study

We feel it is important to lay out our understanding of the differences between our experimental conditions and results and those of Adini et al. (2002). Some of these differences have been clarified through discussions with Sagi.

First, as we have made clear, our four non-interleaved contrast practice experiment (Experiment I) is nearly identical to the Adini et al. (2002) non-flanker control experiment (same stimuli and 2AFC staircase procedure), but our experiment found significant learning and Adini et al. (2002) did not. Though we are not aware of exactly what caused this difference, we feel it worth pointing out that the two observers' data in the Adini et al. (2002) no-flanker learning experiment (their Figure 2b) show a very consistent but peculiar pattern. That is, both TvC functions are flat at reference contrasts from 0.2 to 0.5, as if two observers already had previous practice at high contrasts.

More recently, Sagi et al. (2003) also reported no learning when discrimination for seven non-interleaved contrasts was practiced (also see Dorais & Sagi, 1997, for practice at 7-8 contrasts). Although practicing at seven contrasts may not be qualitatively different from practicing at four contrasts where significant learning was evident (Figure 1), it may result in insufficient practice per contrast in a given session, and insufficient practice in turn may not be able to produce a sustainable memory trace. For example, in Dorais and Sagi (1997), each contrast was practiced in only one staircase run, which was about 50 trials, in contrast to nearly 200 trials in our experiment (Figure 1). The unsustainable learning due to insufficient trials is further exacerbated by the fact that contrast learning does not transfer much between contrasts (Figure 3c).

Second, in Experiment II, we used an elongated Gabor pedestal to simulate flankers, rather than using additional collinear Gabors. Adini et al. used flankers at a separation of 2.8 Gabor envelope SDs. Could this difference change the experiment outcomes? Possibly, but not very likely, because in our Experiment II the initial practice without

flankers greatly lowered contrast thresholds to be around 0.07–0.08 at a reference contrast of 0.50. This threshold level is even lower than the 0.11–0.12 level at the same reference contrast after context-enabled learning in Adini et al. (2002), consistent with our conclusion that context-enabled learning may actually be part of the regular learning. Moreover, previous studies by Sagi's group and us have reported that contrast discrimination is similarly affected by either the length of the pedestal (Yu & Levi, 1997) or the number of flankers (Adini & Sagi, 2001), indicating that our extended flankers may function similarly to the Gabor flankers.

Third, Sagi et al. (2003) also pointed out that in our Experiment IV, we only used three pairs of Gabor flankers in the stimuli, while the number of flankers used in their experiments gradually increased from session to session (from 2-10). Although Adini et al. (2002) never elaborated why changing the number of flankers is important in "context-enabled learning," it seems to us unlikely that while a fixed number of flankers slightly impairs contrast learning in the contrast roving condition, increasing the number of flankers session by session could significantly improve it. Finally, the timing was slightly different (duration 103 vs. 90 msec and ISI 600 vs. 1000 msec).

We cannot completely rule out the possibility that these differences in the stimuli or experimental methods account for our failure to replicate the Adini et al. (2002) nonlearning data under no-flanker practice conditions. Indeed, although we tried to match what we considered to be the critical experimental conditions between the two labs (see Acknowledgments), Sagi (personal communication) feels that (unknown) subtle experimental details may be crucial. Nonetheless, our results show clearly that contrast discrimination can be improved through practice, and that flankers are not necessary to bring about nonroving contrast learning, nor are they sufficient to bring about learning under contrast roving conditions.

Why does contrast learning not transfer to neighboring contrasts? Three models of contrast discrimination

The dominant feature of contrast discrimination is the power law Weber-like behavior of the TvC function. The TvC function is expressed as $\Delta c = k\,c^n$, with the power n typically between 0.5 and 0.7 and with k typically about 0.15. Three broad hypotheses have been proposed for the mechanisms controlling contrast discrimination: contrast response function saturation (gain control), multiplicative noise, and multiple contrast channels.

The most popular account of the TvC function is in terms of a saturating contrast response function (Figure 6, top row), as would result from a contrast gain control mechanism. Adini, Sagi, and Tsodyks (2002) proposed such a mechanism to account for their finding of context-enabled learning. One difficulty with these contrast response function-based models is that they do not
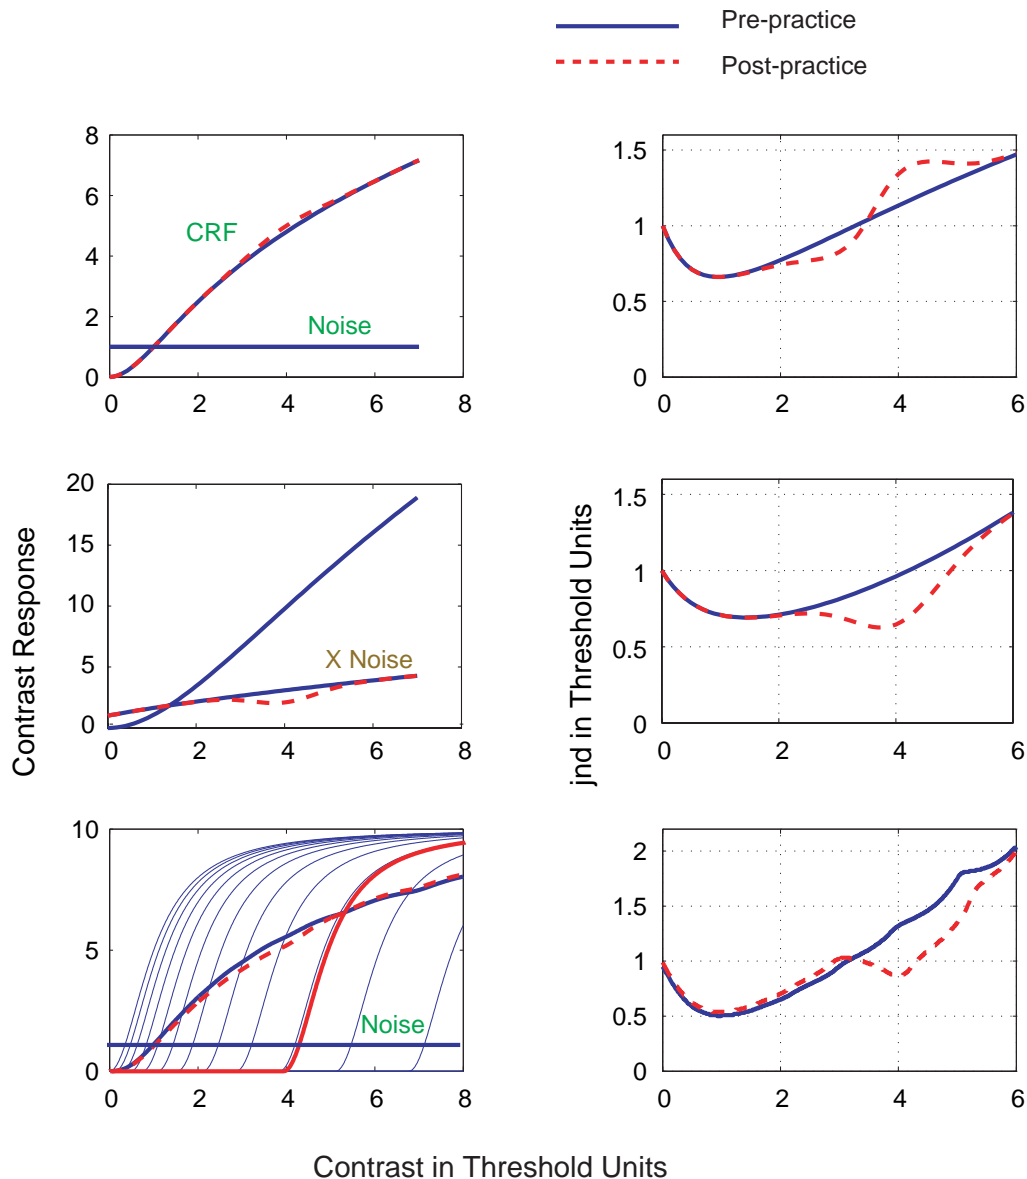
Figure 6. Three models of contrast discrimination and learning: gain control (top row), multiplicative noise (middle row), and multiple contrast-selective channels (bottom row). For each model, the left panels show the contrast response functions (CRF) or multiplicative noise (X Noise) functions pre- (blue) and post (red)-practice. The right panels show the resulting changes of the TvC function. The thin S-shaped blue curves in the left panel of the bottom row are the contrast response functions of individual neurons, each with a limited range of contrast responses. The thick blue curve is the average of individual contrast response functions. Practice at a contrast of 4 threshold units as simulated in the figure induces better attention to the responses of the most relevant "neuron" that we model by adding an extra "neuron" at the practiced contrast (thick solid red curve) and changes the shape of the population contrast response function (dashed red curve). The right panel of the bottom row shows the resulting changes of the TvC function. MATLAB code for this figure is listed in the "Appendix."

naturally account for our finding of non-transfer of learning to nearby contrasts (Figure 3c). One might conceive of a Hebbian model where by repetition the practiced region (near practiced contrast Cp or 4 threshold units in Figure 6) of the contrast response function becomes sensitized. The dashed line (post-training) in the top left panel of Figure 6 is slightly elevated as compared to the solid line (pre-training). What are the implications for the TvC function of facilitation of the contrast response function? We would expect facilitation of discrimination for a small range below Cp and inhibition of discrimination above Cp (Figure 6, top right panel). However, our data show that the facilitation is peaked at the practiced contrast. The full details of the modeling that generated Figure 6 are provided in the Matlab code in the Appendix.

Another account for the near-Weber relationship of the TvC function is the multiplicative noise model (Figure

6, middle row). Signal detection theory reminds us that d' is equal to the change in the contrast response function divided by the noise. If the noise increases as the pedestal increases (multiplicative noise), a Weber-like threshold elevation could result. One possible explanation of learning contrast discrimination is that the noise at the practiced contrast is reduced. This would indeed account for the reduced transfer of learning to nearby contrasts (Figure 6, middle row).

A third account of the TvC function shape is in terms of multiple contrast-selective mechanisms (Figure 6, bottom row). It has been shown by Sclar et al. (1990) and Geisler and Albrecht (1997) that the dynamic response range of cortical neurons is surprisingly limited. There is a factor of about 10 between the contrast at which a neuron begins to fire and the contrast at which it saturates. The full contrast range is spanned by a multiplicity of neurons with different thresholds. As contrast increases, an increasing number of neurons fire. The decision stage could simply count the number of firing neurons. Practice at one contrast level could enhance performance in two ways. First, practice could sharpen the tuning of the mechanisms. In this case that would be done by steepening the slope (narrowing the dynamic range) of the contrast response character of neurons. A more plausible possibility is that because there is always noise in neural systems, it would make sense that with practice at one contrast level, the decision stage could learn to attend to those neurons most sensitive to that range of contrast. That would enhance contrast discrimination in that range and leave unaltered discrimination for contrasts outside the practiced region (Figure 6, bottom row). This hypothesized mechanism is compatible with our finding of the specificity of learning to the region of practiced contrasts.

Early (primary visual cortex) versus late (decision stage) learning

Why is there such a large interest in visual learning? We suspect that this interest stems from the possibility that the learning takes place in early stages of visual processing. Learning in late stages of processing would be less interesting because there are already many examples of cognitive learning tasks. For example, if the visual task involved detecting a subtle pattern with many distracters, we would not be surprised to find strong learning effects as one learns to recognize and discount the distracters. We are more interested when we find learning with simple patterns with aspects such as non-transfer to different locations or orientations that indicate that the learning might take place in early stages of processing. However, as noted above, Vogel and Orban (1994) and Mollon and Danilova (1996) showed that non-transfer of learning, often thought to be early, can be explained by central mechanisms.

The massive interconnectedness of cortex makes it difficult to separate early and late stages of processing even when using brain-imaging techniques. After about 150

msec, it is expected that effects of late decision stages will affect V1 processing through feedback (Lee & Mumford, 2003). In order to be concrete about the subtle distinction between learning occurring in early versus late stages of processing, we propose the following simple operational definition. Learning at an early stage would allow a microelectrode implanted in an early visual area (say in V1) to produce a direct correlate of learning in the first 125 msec of response. For example, if noise reduction (the second model of Figure 6) occurred early, then a microelectrode in a neuron responding to the Gabor patch would show reduced noise in its early firing rate after learning. On the other hand, learning at a late stage would not affect the initial responses of neurons in primary visual cortex. For example, reduced noise in the comparison stage (e.g., by improving the memorized contrast template or by learning a more efficient way to compare the memory to the test stimulus) need not show up in the initial firing rate and would therefore be called late stage learning even if the computations were carried out in V1. Similarly, with our selective mechanism hypothesis (the third model of Figure 6), a decision stage with access to the multiple mechanisms spanning the full contrast range would be needed. Thus for the contrast selective mechanism hypothesis, the main action is carried out at a higher stage of processing. The first two hypotheses (change in the contrast response function or the multiplicative noise in neurons in primary visual cortex) would allow a perceptual decision to be made based on the activity of single cells or cell assemblies in early vision. In an important sense we are using early and late as referring to the temporal domain as well as to whether the learning is top-down. This distinction is relevant to our earlier mention of late versus early mechanisms for learning.

It is worth mentioning that evidence directly linking perceptual learning to neural plasticity has been scarce and inconclusive. Schoups, Vogels, Qian, and Orban (2001) recently reported sharpening of orientation tuning functions after orientation discrimination practice in the primary visual cortex of monkeys, but Ghose, Young, and Maunsell (2002) failed to find evidence for similar physiological correlates of orientation discrimination learning in a separate monkey study and referred the behavioral performance improvement to more central pooling and decision processes (similar to the third model in Figure 6). Recent work suggests that such learning does take place in V4 (Yang & Maunsell, 2004). It would be interesting to know whether more central visual processing is universal in perceptual learning of other visual discrimination tasks, such as phase and spatial frequency discrimination and Vernier acuity, even if we cannot completely exclude the role of early neural plasticity in contrast learning.

Given that high-level (late) learning is well established, the burden of proof in the early versus late argument should, therefore, be on the side of those arguing that learning is done early. By this reasoning one need not

provide evidence against early learning. However, we suggest that our roving experiments do provide evidence against the learning being early.

Why does roving inhibit contrast learning? Perhaps the most surprising result of our study, as well as Sagi et al. (2003), is that roving among 4 reference contrasts, in a 2AFC experiment, inhibits learning (Figure 5). The difference between a roving experiment and a blocked experiment is that in the former the only discrimination cue is the contrast difference between the first and second intervals. In a blocked design experiment, there is the additional cue that after a few trials a long-term memory trace of the reference contrast is built up and that memory trace can be used as a reference for both intervals of the 2AFC trial. It may be useful to illustrate the two cases where the perceived signal strength is represented by a number. Suppose the first interval of the roving trial has a perceived strength of 30±5 and the second interval has strength of 34±5. In this case, the perceived contrast difference would have d'=(34–30)/5, which is below the d'=1 threshold. For the blocked experiment, the perceived strength of the signal relative to the memorized reference would be smaller, more accurate numbers such as 2±3 and 6±3 for the first and second intervals. In this case, the d' of the contrast difference would be (6–2)/3, which is above the d'=1 threshold. In the blocked case, learning could decrease thresholds by either improved memorization of a stable reference template or by learning to more accurately compare the memorized reference to each test.

We now examine how the three models of contrast discrimination discussed in the preceding section would deal with the results of our roving experiment. For both *the contrast response function model* and the *multiplicative noise model*, it is hard to see why roving would make it more difficult to learn contrast discrimination. If practice facilitates the contrast response function or reduces the noise at one contrast level, there is no obvious reason why it should not do the same if the contrasts are roving. (Of course one could always develop post hoc models with assumptions that make it difficult to do learning with roving contrasts.) For the *multiple contrast-selective channels model*, on the other hand, there is a natural explanation for why roving causes a problem. As discussed in the section on transfer of contrast learning, according to this model, learning takes place because the decision stage learns to attend to the optimally sensitive mechanisms. In the presence of roving, this type of selective attention would not be possible because attention would be spread out, as in the pre-practice runs.

Summary

Practice can improve contrast discrimination in the absence of flankers. On the other hand, context (flankers)-enabled learning cannot be separated from regular contrast learning.

No learning under contrast roving suggests that contrast learning may take place at a more central stage.

The contrast specificity of contrast learning and no-learning under contrast roving provide new evidence for a multiple contrast-selective channels model of contrast discrimination, and against saturating transducer models and multiplicative noise models.

Footnotes

[1] Sagi et al. (2003) did not retest the four-contrast learning blocked condition in which Adini et al. (2002) found no effect (their Figure 2b), but we found significant contrast learning (our Figure 1). Instead they studied learning in a seven-contrast blocked training experiment and again found no learning effect.

References

Adini, Y., & Sagi, D. (2001). Recurrent networks in human visual cortex: Psychophysical evidence. *Journal of the Optical Society of America A, 18*, 2228-2236. [PubMed]

Adini, Y., Sagi, D., & Tsodyks, M. (2002). Context-enabled learning in the human visual system. *Nature, 415*, 790-793. [PubMed]

Albrecht, D. G., & Geisler, W. S. (1991). Motion selectivity and the contrast-response function of simple cells in the visual cortex. *Visual Neuroscience, 7*, 531-546. [PubMed]

Berliner, J. E., & Durlach, N. I. (1973). Intensity perception. IV. Resolution in roving-level discrimination. *Journal of Acoustic Society of America, 53*, 1270-1287. [PubMed]

De Valois, K. K. (1977). Spatial frequency adaptation can enhance contrast sensitivity. *Vision Research, 17*, 1057-1065. [PubMed]

Dorais, A., & Sagi, D. (1997). Contrast masking effects change with practice. *Vision Research, 37*, 1725-1733. [PubMed]

Fahle, M. (1997). Specificity of learning curvature, orientation, and vernier discriminations. *Vision Research, 37*, 1885-1895. [PubMed]

Fine, I., & Jacobs, R. A. (2000). Perceptual learning for a pattern discrimination task. *Vision Research, 40*, 3209-3230. [PubMed]

Fiorentini, A., & Berardi, N. (1980). Perceptual learning specific for orientation and spatial frequency. *Nature, 287*, 43-44. [PubMed]

Geisler, W. S., & Albrecht, D. G. (1997). Visual cortex neurons in monkeys and cats: Detection, discrimination, and identification. *Visual Neuroscience, 14*, 897-919. [PubMed]

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

Hu, Q., Klein, S. A., & Carney, T. (1993). Can sinusoidal vernier acuity be predicted by contrast discrimination? *Vision Research*, 33, 1241-1258. [PubMed]

Lee, T. S., & Mumford, D. (2003). Hierarchical Bayesian inference in the visual cortex. *Journal of the Optical Society of America A, 20*, 1434-1448. [PubMed]

Legge, G. E. (1981). A power law for contrast discrimination. *Vision Research, 21*, 457-467. [PubMed]

Legge, G. E., & Foley, J. M. (1980). Contrast masking in human vision. *Journal of the Optical Society of America, 70*, 1458-1471. [PubMed]

Legge, G. E., & Kersten, D. (1983). Light and dark bars: Contrast discrimination. *Vision Research, 23*, 473-484. [PubMed]

Levi, D. M., Klein, S. A., & Wang, H. (1994). Discrimination of position and contrast in amblyopic and peripheral vision. *Vision Research*, 34, 3293-3313. [PubMed]

Levi, D. M., & Polat, U. (1996). Neural plasticity in adults with amblyopia. *Proceedings of the National Academy of Sciences USA, 93*, 6830-6834. [PubMed]

Levi, D. M., Polat, U., & Hu, Y. S. (1997). Improvement in Vernier acuity in adults with amblyopia: Practice makes better. *Investigative Ophthalmology and Visual Science, 38*, 1493-1510. [PubMed]

Li, R. W-H., Levi, D. M. & Klein, S.A. (2004) Perceptual learning improves efficiency by re-tuning the "template" for position discrimination. *Nature Neuroscience, 7*, 178-183. [PubMed]

Li, R. W-H., & Levi, D. M. (in press) Characterizing the mechanisms of improvement for position discrimination in adult amblyopia. *Journal of Vision*.

Matthews, N., & Welch, L. (1997). Velocity-dependent improvements in single-dot direction discrimination. *Perception and Psychophysics, 59*, 60-72. [PubMed]

Mayer, M. J. (1983). Practice improves adults' sensitivity to diagonals. *Vision Research, 23*, 547-550. [PubMed]

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

Saarinen, J., & Levi, D. M. (1995). Perceptual learning in Vernier acuity: What is learned? *Vision Research, 35*, 519-527. [PubMed]

Sagi, D., Adini, Y., Tsodyks, M., & Wilkonsky, A. (2003). Context dependent learning in contrast discrimination: effects of contrast uncertainty [Abstract]. *Journal of Vision* ,*3*(9), 173a, http://journalofvision.org/3/9/173/, doi:10.1167/3.9.173.

Schoups, A., Vogels, R., Qian, N., & Orban, G. (2001). Practicing orientation identification improves orientation coding in V1 neurons. *Nature, 412*, 549-553. [PubMed]

Sclar, G., Maunsell, J. H., & Lennie, P. (1990). Coding of image contrast in central visual pathways of the macaque monkey. *Vision Research, 30*, 1-10. [PubMed]

Shapley, R.(1986). The importance of contrast for the activity of single neurons, the VEP and perception. *Vision Research, 26*, 45-61. [PubMed]

Sowden, P. T., Rose, D., & Davies, I. R. (2002). Perceptual learning of luminance contrast detection: Specific for spatial frequency and retinal location but not orientation. *Vision Research, 42*, 1249-1258. [PubMed]

Vogels, R., & Orban, G. A. (1994) Activity of inferior temporal neurons during orientation discrimination with successively presented gratings. *Journal of Neurophysiology, 71*, 1428–1451. [PubMed]

Yang T., & Maunsell J. H. (2004) The effect of perceptual learning on neuronal responses in monkey visual area V4. *Journal of Neuroscience, 24*, 1617-26. [PubMed]

Yu, C., & Levi, D. M. (1997). End-stopping and length tuning in psychophysical spatial filters. *Journal of the Optical Society of America A, 14*, 2346-2354. [PubMed]

Yu, C., & Levi, D. M. (2000). Surround modulation in human vision unmasked by masking experiments. *Nature Neuroscience, 3*, 724-728. [PubMed]

Appendix

```
% Matlab code with the full details of the modeling that generated Figure 6. Comments are given in green.
clear;
clf;
c=0:.01:8; %the pedestal contrasts being sampled
type=['b- r--']; %blue and red are pre- and post-learning
for ia=0:5 % even and odd numbers for pre- and post-learning
        % ia=0,1 for top; 2,3 for middle; 4,5 for bottom panels of Figure 6
    n=c.^0; %Constant noise [for model 1 (top panels) and model 3 (bottom panels)] ia2=mod(ia,2);ia12=ia-ia2  %for
    plotting conditions
    if ia<2, anum=ia*.2; a=2; %conditions for model 1
    elseif ia<4, anum=0; a=10; n=(1+c).^.7-ia2*exp(-(c-4).^2);%for model 2 (middle panels)
    end
    resp= (a+1)*c.^2./(a+c.^1.5)+anum*exp(-(c-4).^2); %contrast response function for Models 1 & 2
    subplot(3,2,1+ia12); %specify where to place the plot
    if ia<2, plot(c,resp,type(3*ia2+1:3*ia2+3),c,n); hold on %plot left panel of Model 1
    elseif ia<4, plot(c,resp,c,n,type(3*ia2+1:3*ia2+3));hold on %plot left panel of Model 2
    else offset=[.4*1.3.^[0:13]]-.4; %calculate plot 5. All the shifts of black curves
            if ia==5; offset=[offset offset(10)+.05];end %add an extra mechanism post-learning
            for iplot=1:length(offset)
                off=offset(iplot); %offsets for contrast response functions of plot 5
                cshift=(c-off).*(c>off); %contrast of shifted curves
                R(iplot,:)=10*cshift.^2./(1+cshift.^2);%Naka-Rushton type saturation response
            end
            resp=mean(R); %the CRF is the average of all the separate neural responses
            if ia==4; subplot(3,2,5);plot(c,R,'k',c,resp,'b');hold on %plot Model 3 pre-learning
            else subplot(3,2,5);plot(c,R(end,:),'k',c,resp,'r--',[0 8],[1 1],'b');%3 post-learning
            end
    end
    ylabel('CRF and noise');if ia==0, title('contrast response functions and noise');end
    for i=1:length(c)
            [rmin,cmin]=min(abs((resp-resp(i))./n-1));%solves {(CRF(cmin)-CRF(c))/noise = 1} for cmin
            cjnd(i)=cmin/100-c(i);%The jnd contrast. The /100 converts sample units to contrast units
    end
    subplot(3,2,2+ia12);plot(c,cjnd,type(3*ia2+1:3*ia2+3));hold on %plot right panels
    if ia==1, title(' jnd. pre (black) and post (red)');end;
    axis([0,6,0,2.2]); grid on
end
```