PERGAMON

Vision Research 40 (2000) 3209–3230

Vision Research

www.elsevier.com/locate/visres

Perceptual learning for a pattern discrimination task

I. Fine [a,*], Robert A. Jacobs [b]

[a] Department of Psychology, University of California, San Diego, 9500 Gilman Drive Dept. 0109, La Jolla, CA 92093-0109, USA
[b] Center for Visual Science, University of Rochester, Rochester, NY 14627-0268, USA

Received 23 June 1999; received in revised form 20 January 2000

Abstract

Our goal was to differentiate low and mid level perceptual learning. We used a complex grating discrimination task that required observers to combine information across wide ranges of spatial frequency and orientation. Stimuli were 'wicker'-like textures containing two orthogonal signal components of 3 and 9 c/deg. Observers discriminated a 15% spatial frequency shift in these components. Stimuli also contained four noise components, separated from the signal components by at least 45° of orientation or ~ 2 octaves in spatial frequency. In Experiment 1 naive observers were trained for eight sessions with a four-alternative same-different forced choice judgment with feedback. Observers showed significant learning, thresholds dropped to $\sim 1/3$ of their original value. In Experiment 2 we found that observers showed far less learning when the noise components were not present. Experiment 3 found, unlike many other studies, almost complete transfer of learning across orientation. The results of Experiments 2 and 3 suggest that, unlike many other perceptual learning studies, most learning in Experiment 1 occurs at mid to high levels of processing rather than within low level analyzers tuned for spatial frequency and orientation. Experiment 4 found that performance was more severely impaired by spatial frequency shifts in noise components of the same spatial frequency or orientation as the signal components (though there was significant variability between observers). This suggests that after training observers based their responses on mechanisms tuned for selective regions of Fourier space. Experiment 5 examined transfer of learning from a *same-sign* task (the two signal components both increased/decreased in spatial frequency) to an *opposite-sign* task (signal components shifted in opposite directions in frequency space). Transfer of learning from same-sign to opposite-sign tasks and vice versa was complete suggesting that observers combined information from the two signal components independently.

Keywords: Perceptual learning; Pattern discrimination

1. Introduction

A great deal is known about low level visual pattern analyzers and their role in visual perception. In the 1970s and 1980s many psychophysical experiments were carried out using stimuli at or near their own detection threshold (see Graham, 1989, 1992 for a review). Such experiments led to a clear, rigorous and quantitative model of the lowest stages of pattern vision: at early stages of processing retinal input is represented by low level analyzers tuned for spatial frequency (Carter & Henning, 1971; Graham & Nachmias, 1971; DeValois & Switkes, 1980) and orientation (Campbell & Kukilkowski, 1966; Blakemore & Campbell, 1969), with receptive fields of limited spatial extent (Kulikowski & King-Smith, 1973), properties very similar to those of simple cells (Hubel & Wiesel, 1962, 1968, 1977; De Valois & De Valois, 1988). However images of real-world objects contain a wide range of Fourier components, and therefore the combination of information across low level analyzers is necessary to reliably recognize objects. We are interested in how 'mid level' mechanisms might pool information across low level analyzers tuned for a wide range of spatial frequencies and orientations.

Evidence for mid level mechanisms comes from psychophysical experiments showing that the ability to discriminate and identify complex pattern stimuli is usually not predictable from simple probability summation of the outputs of low level analyzers. The independence between different spatial frequencies and orientations found for threshold detection tasks rapidly

* Corresponding author. Fax: +1-858-5347190.

0042-6989/00/$ - see front matter
PII: S0042-6989(00)00163-2

disintegrates for discrimination or identification tasks using suprathreshold stimuli (e.g. Georgeson, 1992; Derrington & Henning, 1989; Burr & Morrone, 1994; Graham & Sutter, 1998; Olzak & Wickens, 1997; Olzak & Thomas, 1986, 1991, 1999). Observers are better at processing particular combinations of spatial frequency and orientation. For example, Olzak and Thomas (1999) described mid level mechanisms sensitive to a wide range of spatial frequencies at a single orientation (as occurs at an edge in a visual scene). They argue that these mid level mechanisms pool information across low level analyzers tuned for a particular orientation but a wide range of spatial frequencies. This selective pooling by mid level mechanisms might be considered the beginning of extracting features commonly found in the natural world.

Studies examining mid level mechanisms have almost always measured performance after considerable practice, and have not addressed the possibility that the characteristics of these mechanisms might change as a function of training. However an increasing number of studies suggest that even relatively early stages of the visual system may change with training (e.g. Ball & Sekuler, 1982, 1987; Fahle & Edelman, 1992; Vidyasagar & Stuart, 1993; Sagi & Tanne, 1994; Ahissar & Hochstein, 1995, 1996; Saarinen & Levi, 1995; Schoups, Vogels, & Orban, 1995; Fahle & Morgan, 1996; Schoups & Orban, 1996; Liu & Vaina, 1998; Liu & Weinshall, 1999). These studies suggest that the changes mediating improvement often, though not always, occur in mechanisms specific for orientation, direction of motion, spatial scale and spatial position. (Learning effects are generally unspecific for eye of origin.) Cells at V1 tend to be tuned for these properties, while cells at higher levels of the visual system tend to be less selective for these properties. This has led experimenters (e.g. Karni & Sagi, 1991; Ahissar & Hochstein, 1995) to suggest that the changes in mechanisms mediating many perceptual learning tasks might take place as early as V1.

In addition, learning effects have been noted (Olzak, personal communication, 1995; Dosher & Lu, 1998, 1999) for tasks involving compound grating discriminations thought to involve mid level mechanisms. Fiorentini and Berardi (1980, 1981) found little learning when observers were asked to discriminate suprathreshold sinusoidal gratings that differed in spatial frequency. However observers did show perceptual learning for a discrimination task based upon the first and third harmonic of a square wave. Observers discriminated the first two harmonics of a square wave from a similar complex waveform differing only in the phase or contrast of the third harmonic. Observers showed improvements in performance over a single session. This learning showed many of the characteristics typical of early learning: it was specific for retinal position and orientation. However learning also showed complete interocular transfer. Fiorentini and Berardi suggested the learning they found might be subserved by mechanisms higher in processing than spatial frequency channels, though still tuned for orientation and spatial frequency.

Although objects tend to contain statistical regularities, such as a wide range of spatial frequencies at the same orientation, they do not exclusively do so. This paper studies the capacity of the visual system to learn to represent relatively arbitrary conjunctions of spatial frequency and orientation. Our signal components were gratings orthogonal to each other and separated by approximately 2 octaves of spatial frequency. Gratings separated by 2 octaves of spatial frequency and 90° of orientation fall within the bandwidth of different low level analyzers. In addition, data suggests that mid level mechanisms selectively tuned for orthogonal gratings widely separated in spatial frequency do not pre-exist training: gratings widely separated in both spatial frequency and orientation seem to be processed independently (Fine & Jacobs, 1998; Olzak & Thomas, 1999). The spatial frequencies and orientations chosen for our stimuli are therefore sensible for studying how the visual system integrates information from arbitrary regions of Fourier space that are not strongly correlated with each other in natural scenes, or processed selectively by the visual system before training.

One difficulty in studying perceptual learning is that tasks must be relatively complicated before observers show significant amounts of learning. (Perceptual learning for simple stimuli is therefore usually examined in the periphery, where there tend to be much larger learning effects). More complicated tasks are difficult to formally model, as usually very little is known about the mechanisms underlying task performance. However complex gratings with additional Fourier noise components are sufficiently complicated to allow a significant amount of perceptual learning while still being simple enough to model. The complex plaids used in this paper are therefore a useful tool: by looking at conjunctions of gratings it is possible to study changes in the mid level mechanisms thought to be responsible for combining information across widely separated regions of Fourier space. These stimuli are also amenable to interesting manipulations, as demonstrated in Experiments 3–5.

It should be noted that although the complex stimuli we use are described in Fourier terms, none of the following experiments are based on any strong assumption that the visual system really can be accurately described as a Fourier system. This paper rests on a ‘very weakly Fourier’ assumption — that at very early stages of visual processing low level analyzers are tuned for both spatial frequency and orientation, and at intermediate stages there exist mid level mechanisms that selectively pool information from these low level analyzers.

Experiment 1 demonstrated that with practice, observers consistently showed significant improvements in discriminating changes in spatial frequency in the underlying gratings of a complex plaid stimulus masked by noise components. Experiment 2 showed that observers show smaller and less consistent learning when asked to discriminate complex plaids without the additional noise components. These experiments indicate that the improvement in performance seen in Experiment 1 is mainly due to changes in mid to high level mechanisms that combine information over wide ranges of spatial frequency and orientation, rather than to changes in low level analyzers tuned for both spatial frequency and orientation.

Experiments 3 and 4 were designed to investigate the selectivity of the mechanisms underlying the learning demonstrated in Experiment 1. Experiment 3 examined transfer of learning across orientation. Observers showed almost complete transfer for three different orientation transfer conditions, suggesting that the improvements in performance shown by observers in the task were not based upon low level analyzers strictly tuned for the orientation and spatial frequency of the signal components. Experiment 4 used a novel technique, based on the use of shifting noise components, to compare how susceptible observers' performances were to shifts in spatial frequency in different regions of Fourier space. Observers were more susceptible to shifts in spatial frequency within noise components that were of the same spatial frequency or orientation as the signal components, or were close to the signal components in Fourier space. These results suggest that observers learned to base their responses on selective regions of Fourier space as a function of practice.

Experiment 5 examined whether the mechanisms underlying discrimination combine information across Fourier space based on an independent combination rule such as probability summation, or whether information is combined based upon a non-independent integration rule. Transfer of learning from a same-sign task (both signal components increased or decreased in spatial frequency) to an opposite-sign task (one component increased in spatial frequency and the other decreased) was tested. Observers showed almost complete transfer from same to opposite-sign tasks, consistent with the use of probability summation, or some other discrimination rule that combined information from the two signal components independently.

2. General methods

2.1. Display

The maximum calibrated luminance of the monitor was 45 cd/m^2 and the minimum luminance was 5 cd/m^2. The mean luminance of the screen remained at 25 cd/m^2 throughout every experiment. Observers were at a distance of 1.5 m from the monitor. The only source of light in the room was the monitor.

2.2. Task

Observers were asked to perform a four alternative forced choice discrimination task (Fig. 1). Four stimuli were presented sequentially in time with the same two-dimensional noise mask after each stimulus. Observers were asked to indicate which of the four stimuli was different from the others using a key press. A short bell marked the onset of each stimulus, and a double bell marked the beginning of the second pair of stimuli. This double bell helped the subjects keep track of the four intervals. Observers were given auditory feedback and were self paced.

There are two advantages of this four alternative forced choice procedure. First, the chance success rate was 25%, thereby providing more information per trial than a two alternative forced choice task. Second, such a task enables subjects to perform a same–different judgment without potential criterion effects (observers showing a bias towards responding same or different).

Before the beginning of the experiment observers were given a small amount of practice ($\sim$ 12 trials) with a very easy, unrelated task. This allowed them to become comfortable with the four alternative forced choice procedure.

Each 40 min session contained five blocks, each containing 50 trials. Observers therefore completed 250 trials per session. Observers were encouraged to take rests between each block of trials. In addition observers could pause the experiment at any time in order to have a rest. Observers never carried out more than a single session in a day, and carried out three to five sessions a week.

The experimental instructions asked observers to pay attention to the stimulus (not the mask) and told observers to look for changes in the texture (rather than the size or contrast of the stimulus). A laboratory assistant sat with observers for the first session, and if the observer seemed to be performing at chance in the first 20 trials the instructions were repeated.

All observers were naive as to the purpose of the experiment. Observers' vision had been tested within the last year and a half, and they wore corrective lenses if necessary. Observers were undergraduate or graduate students and were under 40 years of age. Many observ-

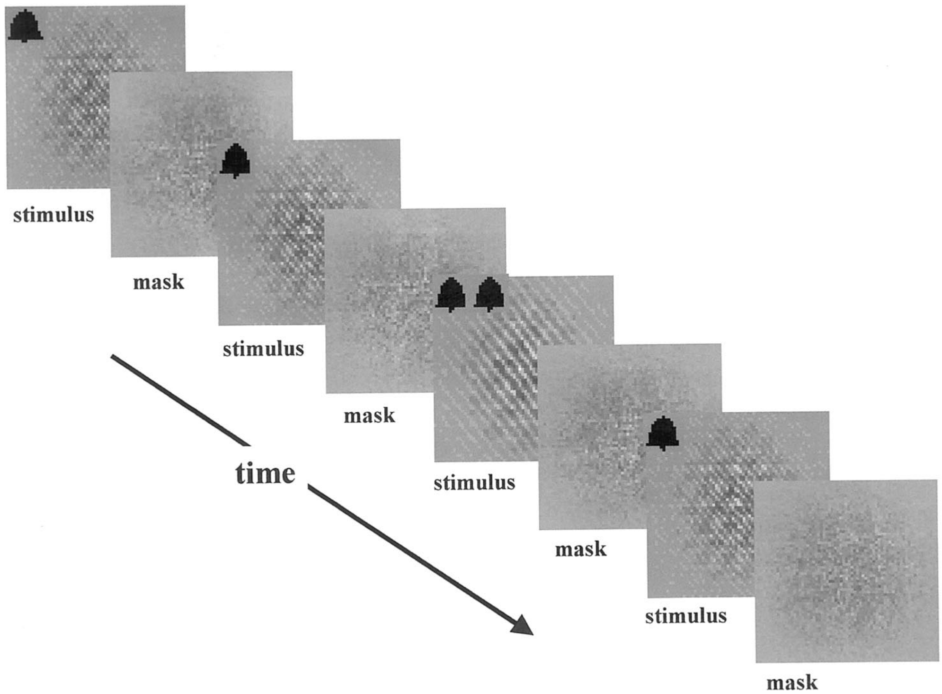

Fig. 1. Diagram of the task used in the experiment. Four stimuli were presented sequentially in time with the same two-dimensional noise mask after each stimulus. Observers were asked to indicate which of the four stimuli was different from the others using a key press.

ers had previously participated in vision experiments using very different tasks and stimuli. Observers who had previous experience as psychophysical observers with tasks using gratings or plaids as stimuli were excluded from the study.

2.3. Stimulus

The stimuli contained two signal components and four noise components. Table 1 provides the spatial frequency, orientation and contrast of the signal and noise components. Fig. 2 represents the stimuli in Fourier space using polar coordinates. The radius represents spatial frequency and the angle represents orientation. The black filled circles represent the two possible signal components. These signal components were widely separated in orientation (at least 90° to each other) and widely separated in spatial frequency (approximately 2 octaves apart). It can be assumed relatively safely that these signal components were processed independently by low level analyzers. (For evidence for independent channels for orthogonally oriented gratings see Campbell & Kukilkowski, 1966; Blakemore & Campbell, 1969; for evidence for independent channels for widely separated spatial frequencies see Carter & Henning, 1971; Graham & Nachmias, 1971; DeValois & Switkes, 1980). There were two possible LOW components, centered on 3 c/deg, only one of which was ever present. The LOW− component was always at 2.55 c/deg (3 c/deg − 15%) and the LOW+ component was always at 3.45 c/deg (3 c/deg + 15%). There were two possible HIGH components, centered on 9 c/deg, only one of which was ever present. The HIGH− component was always at 7.65 c/deg (9 c/deg − 15%) and HIGH+ was always at 10.35 c/deg (9 c/deg + 15%). The black arrows show how the LOW and HIGH components shifted in spatial frequency between LOW− and LOW+ and between HIGH− and HIGH+. The one stimulus in each trial that was different from the other three (the 'odd man out') was always distinguished by a spatial frequency shift in both LOW and HIGH components. The separation between the LOW− and LOW+ components and the

Table 1
Complex plaid used for observers in Experiment 1

Signal/mask	Spatial frequency (c/deg)	Orientation (°)	Contrast (%)
LOW signal	2.55 or 3.45	−45	1.6–12.8
HIGH signal	7.65 or 10.35	45	5.5–44
Noise component 1	9	−45	11
Noise component 2	3	45	3.2
Noise component 3	4.3	0	7.1
Noise component 4	6.2	0	7.1

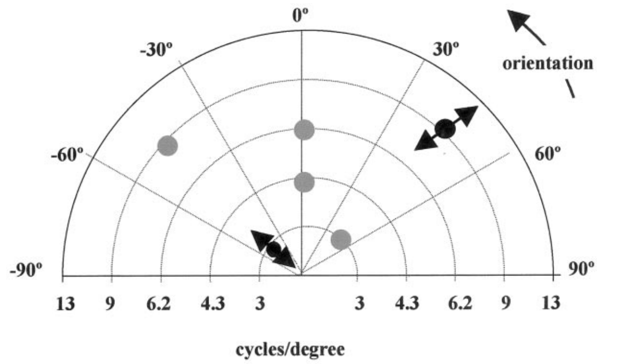

Fig. 2. Fourier representation of typical stimuli. The x-axis represents spatial frequency and the angle represents orientation. One signal component had a spatial frequency of 2.55 or 3.45 c/deg and an orientation of $-45°$ (shown by arrows representing the two possible positions of that component). The other signal component had a spatial frequency of 7.65 or 10.35 c/deg and an orientation of 45° (again shown by arrows representing the two possible positions of that component). There were four noise components, represented by gray circles: (1) spatial frequency of 9 c/deg, $-45°$ orientation, 11% contrast; (2) spatial frequency of 3 c/deg, 45° orientation, 3.2% contrast; (3) spatial frequency of 4.3 c/deg, 0° orientation, 7.1% contrast; (4) spatial frequency of 6.2 c/deg, 0° orientation, 7.1% contrast.

HIGH− and HIGH+ components was fixed to be the same percentage ($\pm 15\%$) of the central frequency in every trial.

The difficulty of the task was manipulated by changing the contrast of the LOW and HIGH components independently. The LOW components had possible contrasts of 1.6, 2.4, 3.2, 6.4 and 12.8% and the HIGH components had possible contrasts of 5.5, 8.3, 11, 22 and 44% (3.44 times the low component). The 15% shift in spatial frequency between LOW− and LOW+ components and the HIGH− and HIGH+ components, and the contrasts used for both the signal and the noise components, were chosen on the basis of pilot data collected with different observers and a QUEST procedure (Watson & Pelli, 1983). The contrast values were chosen to make the HIGH and LOW components approximately equally visible. Examination of the data from Experiment 1 indicates that observers did in fact consistently use both components when making their discriminations.

Four sinusoidal gratings were added as noise components. The gray circles in Fig. 2 represent the noise components. The first noise component had a spatial frequency of 9 c/deg and an orientation of 45° and was at 11% contrast. The second noise component had a spatial frequency of 3 c/deg, an orientation of $-45°$ and was at 3.2% contrast. The third and fourth noise components were both oriented vertically and had a contrast of 7.1%. The spatial frequencies of the third and fourth noise components were 4.3 c/deg and 6.2 c/deg, respectively.

The frequency and orientation of the noise components were chosen to minimize the degree to which they would mask the signal components within low level visual analyzers tuned for spatial frequency and orientation. Every noise component differed from the signal components by at least 45° of orientation or almost 2 octaves of spatial frequency. However they were also designed to maximize the amount of masking within mid level mechanisms responding to a wide range of spatial frequencies or orientations, thereby forcing the observer to selectively combine information over Fourier space.

The noise components were an integral part of the stimulus. As shown in Experiment 2, observers showed almost no learning when there were no noise components. Observers also showed little learning when the noise components were always of lower contrast than the signal components. The contrast levels of the noise components were chosen to be at the median value of the contrast of the stimuli components, i.e. half the time the signal components had a higher contrast than the masking components, and half the time the masking components were of higher contrast than the signal components.

Observers could be asked to discriminate same-sign or opposite-sign spatial frequency shifts in the signal components. In the same-sign discrimination task observers discriminated between stimuli where both underlying components shifted in the same direction in frequency space — i.e. between (LOW− HIGH−) and (LOW+ HIGH+). In the opposite-sign discrimination task observers were asked to discriminate between stimuli where the underlying components shifted in opposite directions in frequency space — i.e. between (LOW+ HIGH−) and (LOW− HIGH+). In all experiments, except Experiment 5, same-sign and opposite-sign discriminations were randomly interleaved. Previous experiments (Olzak & Thomas, 1986; Fine & Jacobs, 1998) have failed to reveal any difference in performance between same and opposite-sign discriminations for the plaids used in these experiments.

The relative phase (with respect to the center of the screen and the origin of their Gaussian envelope) of the signal components was varied randomly between each trial, and remained constant across the four intervals within each trial. Randomizing the phases of the masking components prevented observers from performing the task by learning to identify particular 'beat' patterns or features within the stimulus. The orientation, contrast and spatial frequency of the masking components were fixed across all trials; however the phase of the masking components varied between each trial and across the four intervals of each trial.

The stimuli were modulated spatially by a two dimensional Gaussian envelope with a sigma of 0.5693° centered within a window of 2.76°, and was modulated

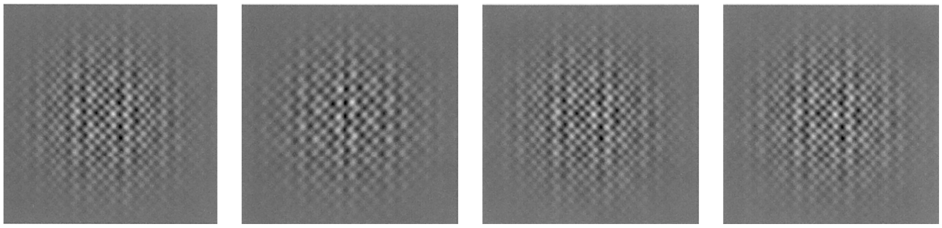

Fig. 3. Illustration of a typical set of stimuli.

temporally by a Gaussian envelope of sigma 0.237 s centered within a 0.67 s temporal window. The monitor had a frame rate of 100 Hz, a pixel depth of 8 bits and a resolution of 1024 by 768 pixels. Stimuli were programmed with the help of routines from the VideoToolbox (Pelli, 1997).

A two-dimensional noise mask (as shown in Fig. 1) presented after each stimulus reduced afterimage interference. The noise mask consisted of low pass filtered white noise. Uniform white noise varying between 10 and 40 cd/m^2 was low pass filtered through convolution with a 5×5 parabolic template[1]. After convolution the mean luminance of the noise mask remained the same as the background luminance. The luminance values of the noise mask after convolution varied in an approximately Gaussian manner between approximately 17 and 33 cd/m^2. This noise mask was then Gaussian windowed spatially with a standard deviation of 0.5693° centered within a window of 4.83°, and was modulated temporally by a Gaussian envelope of standard deviation 0.237 s centered within a 0.67 s temporal window identical to that of the stimulus. Each trial (four stimuli and four masks) therefore took 5.36 s.

3. Experiment 1 — perceptual learning in complex plaids with additional Fourier noise

3.1. Introduction

The purpose of Experiment 1 was to see whether observers improve with practice in performing a complex pattern discrimination task that required combining information over an arbitrary conjunction of spatial frequency and orientation.

3.2. Methods

The complex stimuli presented to the observers in Experiment 1 were as described in the general methods. Each complex stimulus contained six sinusoidal gratings: two signal components and four masking components. The stimuli resembled a 'wicker' pattern, as shown in Fig. 3. Observers were asked to discriminate the 'odd man out' (interval 2 in Fig. 3) using a four-alternative forced choice procedure. Observers were given eight training sessions, over a period of 3 weeks. Further details of the training procedure are described in Section 2.

3.3. Results and conclusions

Our task contained a range of difficulty levels. Fig. 4 shows the probability correct, averaged over all difficulty levels, as a function of session. The x-axis shows the session and the y-axis shows the probability correct. Fig. 4A shows the probability correct averaged across observers, Fig. 4B shows the probability correct for individual observers.

It can be seen that the probability correct increased for every observer as a function of practice. Monte-Carlo simulations showed that the difference between the first and last session and the difference between the second and the last session were significant ($P < 0.01$) for every observer. The average percentage improvement in performance was 31%, and the average drop in threshold (see below) was to 34% of the initial value. These improvements in performance compare very favorably with the amounts of learning shown in other studies such as that of Beard, Levi and Reich (1995), or Ahissar and Hochstein (1995). Our learning effects were surprisingly large for a free fixation task: learning effects tend to be much smaller in the fovea than in the periphery (Mayer, 1983). There was no sign of observers consistently showing a single discrete step in learning suggestive of learning a single cognitive strategy for performing the task; rather learning appeared to improve gradually across a number of sessions for most observers.

One of the advantages of using a range of stimulus difficulty levels, as done in these experiments, is that threshold changes as well as improvement in the probability correct can be calculated. Changes in threshold provide a measure of performance that is more closely related to discrimination ability than percentage im-

[1] The parabolic template used for filtering the white noise was very similar to a Gaussian filter with a sigma of 1.83 pixels (1.52″ of visual angle) along vertical and horizontal dimensions.

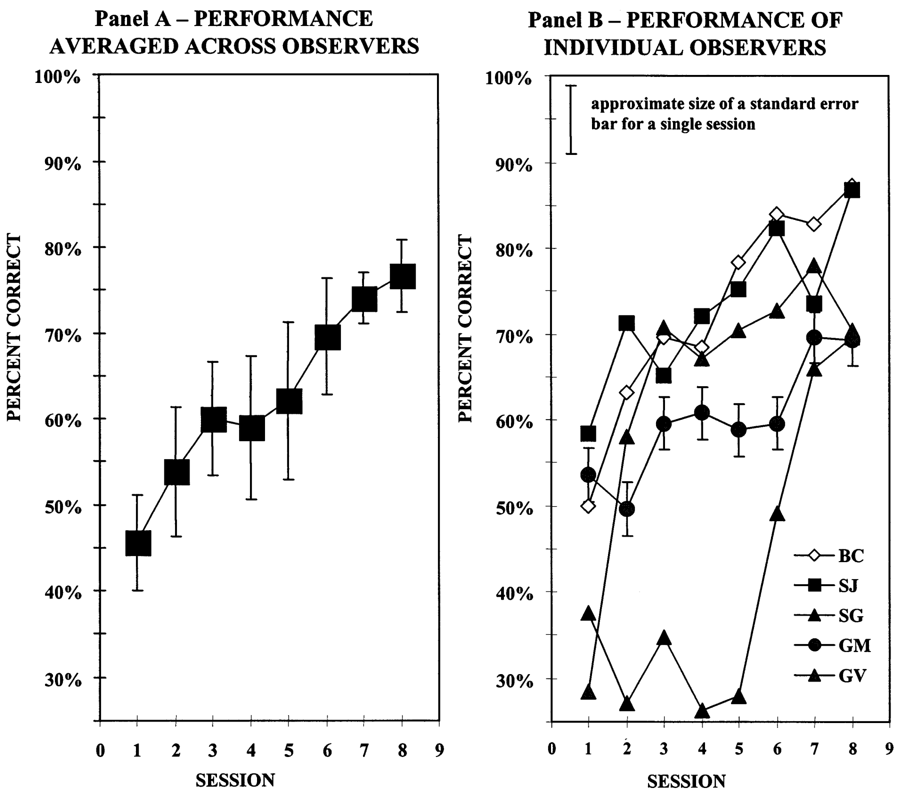

Fig. 4. Panel A shows the probability correct as a function of session averaged across observers in Experiment 1. The x-axis shows the session and the y-axis shows the probability correct. Standard error bars are shown. Panel B shows the probability correct for individual observers

Table 2
Probability correct for observer GM in the fifth session of training

LOWS.F. component contrast	HIGHS.F. component contrast				
	5.5%	8.3%	11%	22%	44%
1.6%	0.2	0.25	0.55	0.85	0.90
2.4%	0.35	0.35	0.50	0.55	0.90
3.2%	0.20	0.15	0.35	0.65	0.85
6.4%	0.45	0.35	0.50	0.70	0.80
12.8%	0.100	0.85	0.75	0.90	0.90

provement. It is possible for large changes in probability correct to be caused by very small changes in observers' thresholds. We wanted to make sure that the improvements in observers' performances were due to significant changes in their discriminative abilities.

Data for each observer can be organized into a 5×5 square. Table 2 shows the probability correct for each of the 25 stimuli presented in a session for a typical observer for a session in the middle of training. The contrast of the low spatial frequency component varies along the rows, and the contrast of the high spatial frequency component varies along the columns. It can be seen that performance is best when either the low and the high spatial frequency component are at high contrast and performance is worst when both the low and high spatial frequency components are at low contrast.

These two-dimensional data were modeled using two assumptions:

1. The probability of detecting a change in spatial frequency of each individual signal component is monotonically related to the contrast of that component, in such a way that it can be fit with a Weibull function. This is a common assumption made when fitting psychophysical data. The particular form of the monotonic function is relatively unimportant and other monotonic functions, such as the logistic function, could also have been used.
2. Discrimination is based upon independent detection of a change in spatial frequency in either component: $P(\mathrm{Discrim}) = P(L_c) + P(H_c) - P(L_c)P(H_c)$.

P(Discrim) is the probability of detecting the shift in spatial frequency in either component, $P(L_c)$ is the probability of detecting the shift in spatial frequency in the low spatial frequency component at a particular contrast and $P(H_c)$ is the probability of detecting the shift in spatial frequency in the high spatial frequency component at a particular contrast. It was assumed that observers always responded correctly when they discriminated the target, i.e. key-press errors were not taken into account.

The data were modeled on the basis of these two assumptions using a parameter optimization procedure. It was impossible to get good consistent fits when fitting Weibulls to both signal components simultaneously[2]. A Weibull was therefore fit to each signal component in turn, and then the Weibull fits were combined using our assumption that discrimination was based upon independent detection of either component (assumption 2). An additional parameter was subtracted from the Weibull for each component. This additional parameter for each component represented the probability of getting the correct answer through chance or through detecting the change in spatial frequency in the other component.

Fig. 5A, shows data for observer GM for session 5. Fig. 5B shows the surface fit by the model. The contrast of the high spatial frequency component, H_c, varies along the x-axis; the contrast of the low spatial frequency component, L_c, varies along the y-axis; the percent correct of the observer or the model varies along the z-axis. The gray dotted line represents the flat plane of 62.5% correct performance. There are several points on the model surface that intersect this plane. We wished to find a single threshold to represent this two dimensional surface. The solid gray line shows the diagonal of the data (from the top left to the bottom right corners of Table 2), where the high frequency component was 3.44 times the contrast of the low frequency component. This diagonal represents stimuli in which the high and low frequency signal components were approximately equally salient. We chose to use the 62.5% threshold that intersected this diagonal.

Because data were noisy, particularly at the beginning of training, it was often difficult to get a good model fit in two dimensions. Data were therefore averaged across the first two and the last two sessions, and the corresponding contrast thresholds were found, as shown in Fig. 6. Fig. 6A shows the contrast threshold for the first and last two sessions averaged across observers and Fig. 6B shows the contrast thresholds for individual observers. Observers' contrast thresholds in sessions 7–8 (white bars) fell on average to a third of their value in sessions 1–2 (black bars). The drop in contrast threshold across the five observers, with a two-tailed paired t-test, was significant ($P < 0.05$).

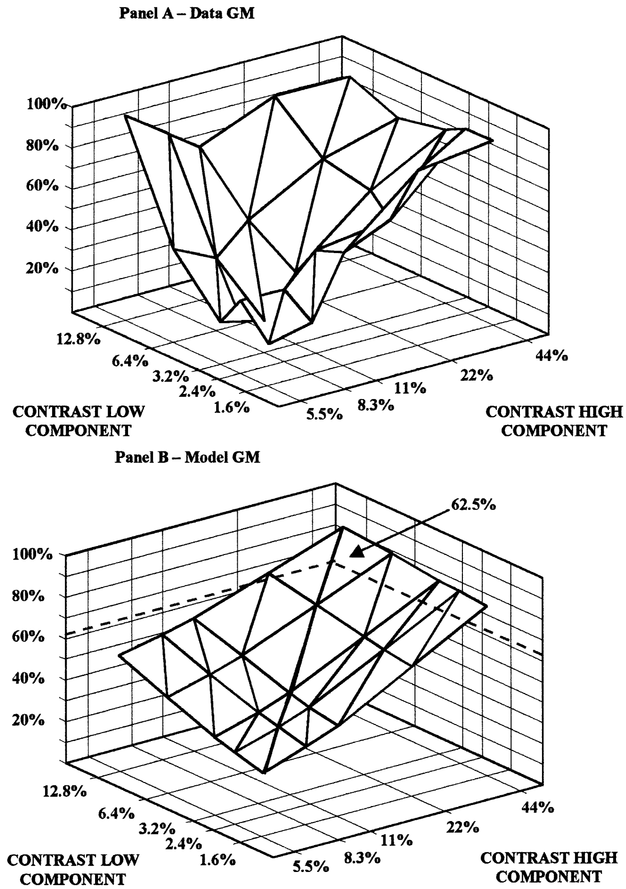

Fig. 5. Panel A shows real data for observer GM in sessions 5–6 of Experiment 1. Panel B shows the model fit for the same observer for the same sessions. The x-axis shows the contrast of the high spatial frequency signal component, the y-axis shows the contrast of the low spatial frequency signal component, the z-axis shows the probability correct. The dotted lines represent the flat plane of 62.5% correct performance. The gray line represents stimuli where the high frequency component was 3.44 times the contrast of the low frequency component (they were approximately equally salient). The threshold was taken to be the intersection of these two lines.

There were two main reasons why the data were fit using such a model. The first was to obtain changes in the discrimination threshold as a function of training for the two components. Fitting a smooth function to our data would have been very difficult without assuming some sort of model or set of constraints. Our model fitted the data reasonably well: correlation coefficients between the real data and the model predictions varied around 0.7–0.95 and were generally above 0.8 when two sessions were averaged together. Our estimates of observers' thresholds can therefore be considered reasonably reliable, regardless of how strictly true the assumptions of our model were. Our other measure of performance, percentage correct, did not depend upon any model assumptions.

[2] Unfortunately, fitting Weibulls to both components simultaneously proved impractical: the data did not sufficiently constrain the model, and parameter optimization procedures tended to fall regularly into local minima.

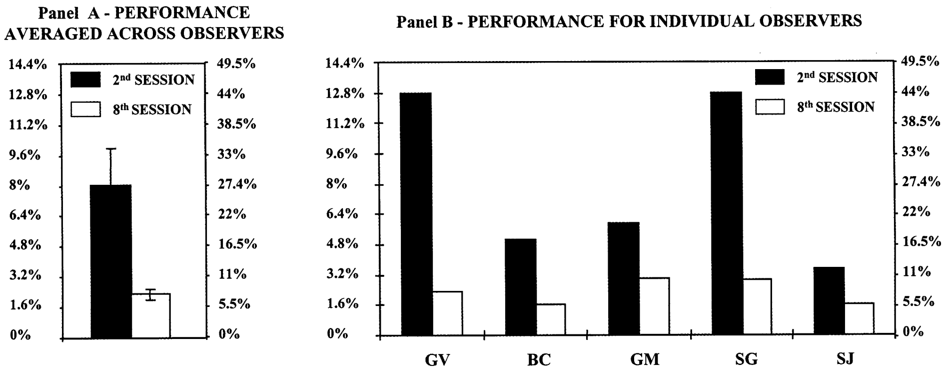

Fig. 6. Panel A shows the contrast threshold in sessions 1–2 and sessions 7–8 averaged across observers. Standard error bars are shown. Panel B shows thresholds for individual observers.

The second motivation for our model was as a 'null hypothesis' for how observers perform the task. The question of whether this model realistically describes observers' performance is further explored in Experiment 5.

Karni and Sagi (1993; also Ahissar & Hochstein, 1997) have noted that performance does not improve uniformly as a function of difficulty level; performance for easier trials improves more rapidly than for more difficult trials. Performance in each session for easier and harder trials was therefore examined, averaged across all five observers. The easy trials were considered to be those where both signal components were of high contrast, as represented by the four bottom right shaded squares of Table 2. The difficult levels were considered to be those where both signal components were of low contrast, represented by the four top left shaded squares of Table 2. In Fig. 7 easy trials are represented by black squares, and difficult trials by gray circles. Standard error bars are shown. As in other studies, observers' performance improved more rapidly for easier stimuli than for more difficult stimuli. In the first four sessions observers' performances for easier stimuli improved on average by 18%, while performances for more difficult stimuli improved by only 0.3%. During the last four sessions this pattern of results was reversed: performances for the easier stimuli only improved by an average of 9%, while performances for the harder stimuli improved by 18.5%. Performance for the more difficult stimuli only began to improve when performances for the easier stimuli were beginning to approach asymptote. It has been shown that having easier trials is often necessary in perceptual learning tasks; observers may 'bootstrap' from easier to harder stimuli. Our data suggest that 'bootstrapping' may also occur in our experiments.

Retention of learning was also examined. Previous studies (Fiorentini & Berardi, 1981) suggest that learning effects tend to be relatively long lasting. Four observers were brought back to complete an additional session more than a month after the end of training. During this month they did not participate in any vision experiments. Observers showed no significant drop in performance[3] after delays of more than a

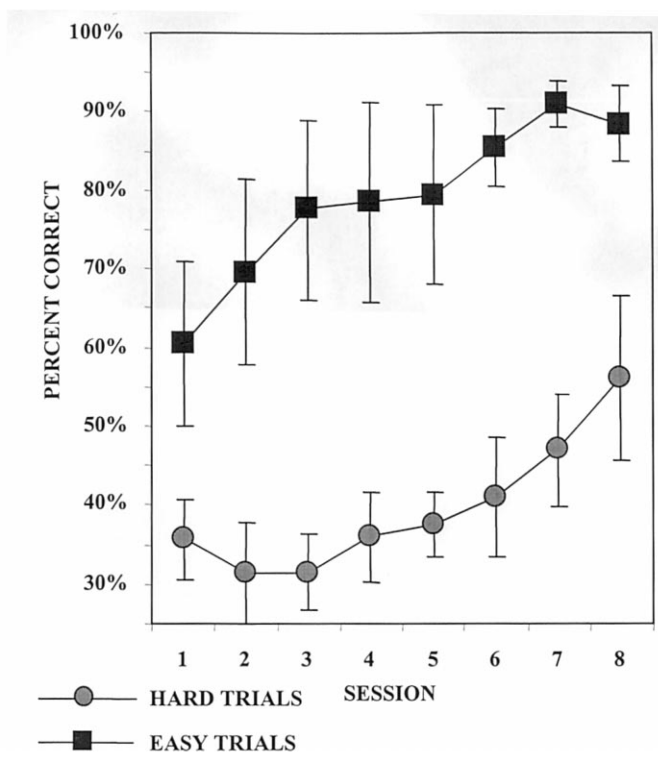

Fig. 7. Performance as a function of session for easier and harder levels of difficulty, averaged across observers. The x-axis shows the session and the y-axis shows the probability correct. Easy trials (signal components were high contrast) are represented by black squares, hard trials (signal components were low contrast) are represented by gray circles.

[3] The performance of one observer (GV) was significantly better ($P < 0.05$) after 4 weeks. This observer reported being less bored with the experiment after the long delay.

month, suggesting that the learning effects of Experiment 1 were relatively long lasting.

To conclude, observers showed a consistent improvement in their performance: indicated by both an increased percent of being correct, and by a decrease in threshold as a function of practice over eight practice sessions. Observers showed faster improvement for easier stimuli than for more difficult stimuli. These effects were relatively long lasting, with no decline in performance more than a month after training.

4. Experiment 2 — perceptual learning in complex plaids without additional Fourier noise

4.1. Introduction

One possibility is that the improvement in performance shown in Experiment 1 might be due to improvement in the ability of low level analyzers to detect or discriminate changes in the spatial frequency of individual gratings. Another possibility is that better performance might be due to an improvement in the ability of mid level mechanisms to selectively combine information over Fourier space.

Observers were trained with the complex plaid stimulus of Experiment 1, without the additional noise components. If the learning demonstrated in Experiment 1 was due to tuning changes within low analyzers tuned for both spatial frequency and orientation then removing the noise components should not affect the amount of learning shown by observers. The noise components were outside the tuning bandwidths of the low level analyzers tuned for the signal components, they were separated from the signal components by at least 45° in orientation or nearly two octaves of spatial frequency. If, however, most of the learning demonstrated in Experiment 1 was due to changes in mid level mechanisms less selectively tuned for spatial frequency and orientation then the removal of the noise components might be expected to affect the amount of perceptual learning.

A number of similar studies (Fiorentini & Berardi, 1980, 1981) suggest that observers show little perceptual learning for discrimination tasks involving single gratings. The amount of learning shown in Experiment 1 is much greater than that shown by Fiorentini and Berardi for single gratings. However there is evidence that tasks thought to involve mechanisms similar to low level analyzers do show some learning in the fovea and parafovea (DeValois, 1977; Mayer, 1983; Beard et al., 1995). Experiment 2 provides a test of how much of the learning demonstrated in Experiment 1 is due to changes in low level analyzers tuned for both spatial frequency and orientation.

4.2. Methods

Display and task were identical to those used in Experiment 1. Only the stimulus differed in Experiment 2, in that the noise components (the gray circles of Fig. 2) were no longer present.

Without the noise components the task would be trivially easy for the contrast levels and spatial frequency shifts used in Experiment 1. If the noise components had simply been removed from the stimulus then observers' performance would have been close to 100% within a few trials. The difficulty of the task was adjusted by reducing the spatial frequency shift to avoid such ceiling effects. The spatial frequency shift was varied between ± 2.5 and $\pm 12.5\%$ to create a range of difficulty levels. The contrast of the signal components was fixed at the median contrast values of Experiment 1: 3.2% for the LOW component and 11% for the HIGH component. The spatial frequency shift rather than the contrast was manipulated to create a range of difficulty levels because detecting low contrast gratings is often modeled as detection of a grating within noise. In Experiment 1 the presence of the masking gratings seems to lead to uncertainty as to what is signal and what is noise. Intuitively, with very low contrast gratings observers also have to deal with signal uncertainty. Minimizing signal uncertainty in Experiment 2 ensured that it was the discriminative powers of low level analyzers that limited performance. Pilot data suggests that we would have obtained similar results had we chosen to manipulate contrast rather than spatial frequency. The range of spatial frequency shifts chosen (based on pilot data), resulted in observers performing at approximately the same mean probability correct in the first sessions of Experiment 2 as they did in the first sessions of Experiment 1.

Three observers were given six sessions of training on the task.

4.3. Results and conclusions

Fig. 8 shows probability correct as a function of session for each of the three observers. Observers showed approximately the same amount of learning between first and second sessions as in Experiment 1. In Experiment 1 observers improved on average by 8.2% between first and second sessions while in Experiment 2 they improved on average by 9.6%. However the amount of learning after the second session was far smaller in Experiment 2. The average amount of learning between the second and sixth session was only 2.9% in Experiment 2 as opposed to 15.8% in Experiment 1, and only observer PCS showed significant learning

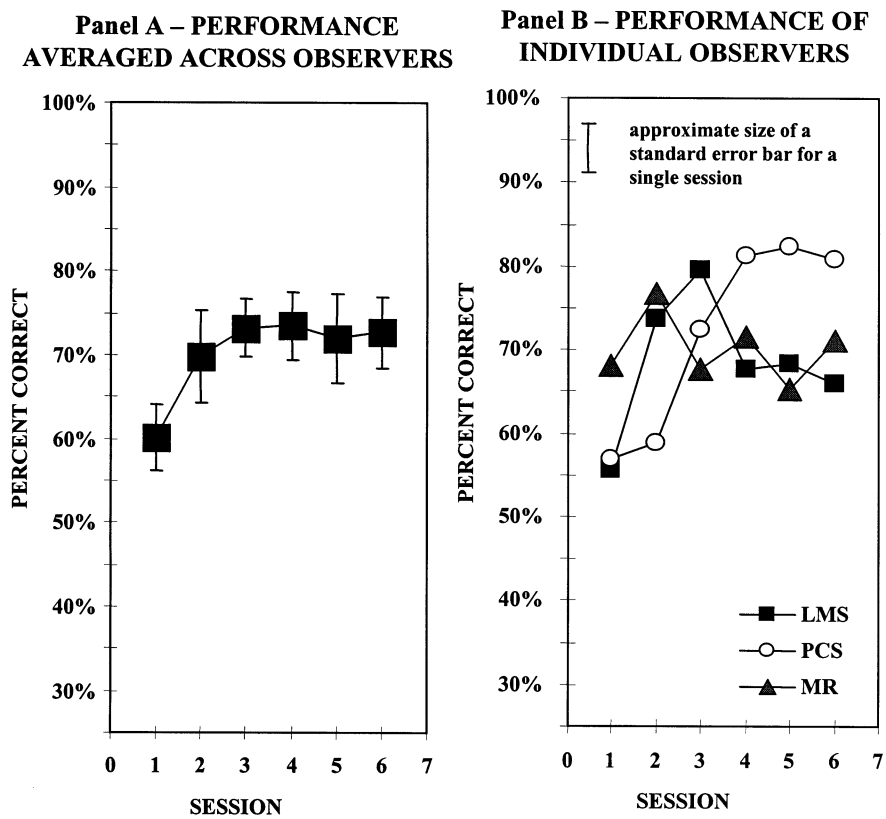

Fig. 8. Panel A shows the probability correct as a function of session averaged across observers for Experiment 2. The x-axis shows the session and the y-axis shows the probability correct. Standard error bars are shown. Panel B shows the probability correct as a function of session for individual observers

between the second and sixth session ($P < 0.01$)[4]. The thresholds in the first two (1–2) and last two sessions (5–6) of Experiment 2 were also compared. Only one of the three observers (PCS) showed a drop in threshold, and there was no significant drop in threshold across the three observers. Observers also failed to show learning in several pilot studies similar to Experiment 2[5].

In comparison, in Experiment 1 five of the six observers (the exception was GM, whose learning was significant at the $P < 0.05$ level) showed learning significant at the $P < 0.01$ level between the second and the sixth session. In total, across a several experiments, we have trained 17 subjects on variants of the ‘wicker’ stimulus described in Experiment 1. Of these 17 subjects, one subject (CG) showed no learning, one subject showed learning significant at the $P < 0.05$ level and the remaining 15 subjects showed learning significant at the $P < 0.01$ level. Given that there was no prescreening of subjects this is evidence of a robust learning effect for the ‘wicker’ task.

Observers in both Experiment 1 and Experiment 2 showed significant learning between the first and second session. This learning between the first and second session, shown in both Experiment 1 and Experiment 2, has two possible causes. Learning between the first and second session might be due to learning in low level analyzers, or alternatively might be due to observers becoming more accustomed to the task and developing better general task strategies (e.g. learning the four-alternative key-press procedure, learning where to fixate, learning when to blink etc.). There is no way to distinguish between these two possibilities with the present data.

[4] Differences in the amount of learning shown in Experiments 1 and 2 are unlikely to be due to ceiling effects — initial performance was closely matched for the majority of subjects. In Experiment 1, three of the five observers performed at 50–60% correct in the first session. In Experiment 2, two of the three observers performed around 55% correct in the first session. None of the observers' performance reached 90% correct by the end of training in either experiment.

[5] In pilot experiments we measured observers' thresholds for a shift in spatial frequency (varied by a QUEST procedure) as a function of practice over five to ten sessions. We found a 10% drop in threshold between the first and second session. We found a further 10% drop in threshold between the second and the final session. However these learning effects were *not* statistically significant — only three of six observers showed consistent learning patterns not swamped by daily variability. We also looked for perceptual learning using gratings much lower in spatial frequency (signal components centered on 0.56 and 1.13 c/deg), and gratings separated by only 1 octave (signal components centered on 3 and 6 c/deg). We did not find a consistent learning effect under either of these conditions.

However these results for Experiment 2 do show that although some learning in Experiment 1 may be due to learning in low level analyzers tuned for both spatial frequency and orientation, most of the learning shown in Experiment 1 *after* the first session is not simply due to learning in low level analyzers. The noise components of Experiment 1 do not fall within the tuning bandwidths of such mechanisms, and similar amounts of learning in Experiment 1 and Experiment 2 would be expected if changes in low level analyzers mediated improvements in performance.

Moreover, the improvements in performance after the second session in Experiment 2 are also unlikely to be due to observers becoming accustomed to the task or developing better non-visual cognitive strategies. It may be that mid level mechanisms become better tuned for the spatial frequencies and orientations relevant for the task, or observers may learn to base their responses on those mechanisms best tuned for the task. Other possibilities include learning to focus attention to the signal components, or some sort of higher level learning, such as template matching.

Experiments 1 and 2 show that most of the learning after the second session is not due to non-visual cognitive strategies, and occurs within mechanisms higher in visual processing than low level analyzers. Experiments 3 and 4 examine the selectivity of the mechanisms underlying these improvements in performance.

5. Experiment 3 — transfer of learning across orientation

As described in the general literature review, transfer of learning has been used in a number of studies to determine the level in visual processing at which changes underlying perceptual learning occur, and the selectivity of the underlying mechanisms (e.g. Ball & Sekuler, 1982, 1987; Fahle & Edelman, 1992; Vidyasagar & Stuart, 1993; Sagi & Tanne, 1994; Ahissar & Hochstein, 1995, 1996; Schoups et al., 1995; Fahle & Morgan, 1996; Schoups & Orban, 1996; Liu & Vaina, 1998). The reasoning behind transfer of learning experiments is that if learning takes place within mechanisms selective for orientation (or direction of motion, retinal position etc.) then learning should not transfer across orientation, as the novel stimulus is processed by different mechanisms from the learned stimulus.

There is evidence that cells in higher visual areas such as V4 or IT are less orientation specific than cells in V1 (D'Zmura & Lennie, 1986; Desimone, Albright, Gross, & Bruce, 1984; Desimone & Schein, 1987; Desimone, Schein, Moran, & Ungerleider, 1985; DeValois, Yund, & Hepler, 1982; Hubel & Wiesel, 1977; Maunsell & Hochstein, 1991; Vogels & Orban, 1994; Wachsmuth & Perret, 1997). A large number of studies have examined transfer of learning across changes in orientation, arguing that failure of transfer across orientation suggests that the changes mediating learning occur at relatively early stages of processing that are orientation specific (e.g. Ahissar & Hochstein, 1995; Schoups et al., 1995). However it is possible that extensive training with stimuli always presented at the same orientation, or to the same eye, leads normally unselective mechanisms (situated higher in visual processing) to become more selective (Mollon & Danilova, 1996). Failure of transfer of learning might demonstrate only that learning has been specific for orientation or the eye of origin, not that the mechanisms themselves are normally selective. Despite these reservations, we examined transfer of learning across changes in orientation. Three different orientation transfer conditions were used to distinguish two different types of sensitivity to orientation, defined as absolute and relative orientation selectivity.

5.1. Methods

The display and task were identical to that used in Experiment 1. Experiment 3 contained two stages. The first training stage was identical to Experiment 1: Experiment 3 was in fact a continuation of Experiment 1. As described in Section 2, observers were trained with a stimulus consisting of two signal components and four noise components for eight sessions.

Fig. 9 shows the stimuli in Fourier space. Fig. 9A shows the stimulus used in Experiment 1. Fig. 9B shows the stimulus in condition *rot all* — the stimulus of Fig. 9A rotated 90° (equivalent to turning the monitor on its side). In condition *rot signal* (Fig. 9C) the signal components and the oblique noise components were rotated 90°, while the vertical noise components remained vertical. Both the absolute orientations of the signal components as well as the relative orientations of the components within the stimulus changed. The oblique noise components were rotated with the signal components to ensure that the signal components remained separated from every noise component by at least 45° orientation or close to 2 octaves of spatial frequency. Rotating the oblique noise components prevented them from being at the same orientation and spatial frequency as the signal components. In Fig. 9D (condition *rot noise*) the signal components and oblique noise components remained at the same orientation and the vertical noise components were rotated 90°. The absolute orientation of the signal components (though not the noise components) remained constant, while the relative orientations of the components within the stimulus changed. Fig. 10 shows what typical stimuli looked like.

It is worth noting (see Fig. 10) that the stimulus in condition *rot signal* (Fig. 10C) appeared most similar to the original stimulus (Fig. 10A). *Rot signal* was the only

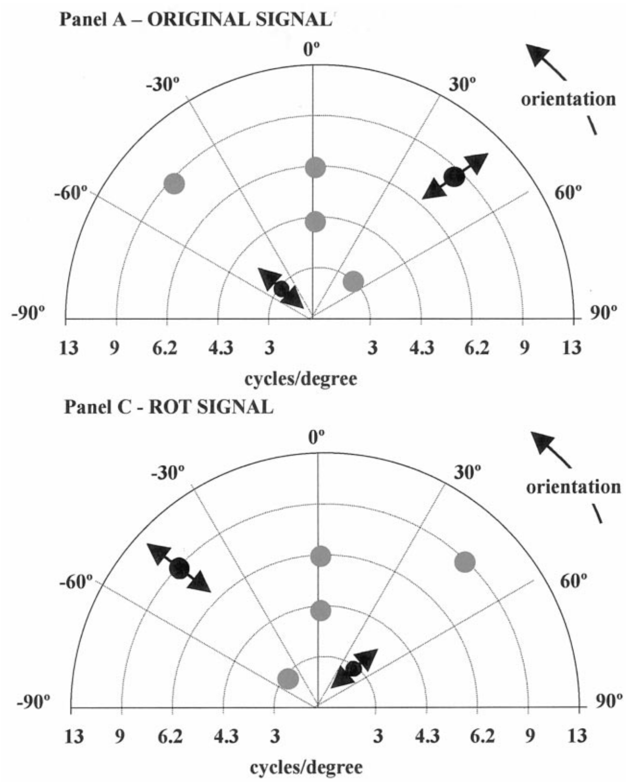

Fig. 9. Panel A shows the stimuli used for the training stage of Experiment 3 in Fourier space. The x-axis represents spatial frequency and the angle represents orientation. Panel B shows stimulus rot all. Panel C shows stimulus rot signal and Panel D shows stimulus rot noise. The x-axis represents spatial frequency and the y-axis represents orientation. Signal components are represented by black circles and noise components are represented by gray circles.

rotation condition in which the stimulus contained the same orientations and spatial frequencies as the original stimulus. In conditions *rot all* and *rot noise* (Fig. 10B and D) the rotation of the noise components resulted in horizontally oriented components not present in the original stimulus.

Five observers were initially trained with the stimulus described in Fig. 9A for eight sessions. In sessions 9, 11 and 13 observers were tested with one of the three transfer conditions. In sessions 10 and 12 observers were presented with the original training stimulus of Fig. 9A. The order in which transfer conditions were presented was based on a Youden square design (an incomplete Latin square).

5.2. Results and conclusions

Fig. 11A shows performance in the various rotation transfer conditions. Standard error bars are shown. The first (black) bar represents performance on the second session and might be considered analogous to a baseline measure: performance dropping to the level of the second day would be considered as almost complete failure of transfer. The second (white) bar represents performance on the last (8th) training session. The remaining bars represent performance for the *rot all* (dark gray), *rot signal* (medium gray) and *rot noise* (light gray) conditions.

As can be seen in Fig. 11 observers showed almost complete transfer of learning across orientation for all three conditions — performance remained very close to performance on the last training session. For all three rotation conditions there tended to be a very slight drop (between 2 and 3%) in performance that was not consistent across observers. There was no significant rise in threshold across observers for any of the three rotation conditions.

In comparison, the average improvement between the second and the eighth session was 22.8% and was significant across observers. It seems that most of the improvement shown by observers was unselective for either absolute or relative orientation.

If improvement in performance were mediated by a ‘template matching’ mechanism sensitive to the appearance (rather than the Fourier content) of the stimulus then a failure of transfer for the *rot noise* condition might be expected, as the appearance of the stimulus changed significantly. Observers showed almost complete transfer to the *rot noise* condition, suggesting that the mechanisms subserving the improvement in performance were not ‘template matchers’.

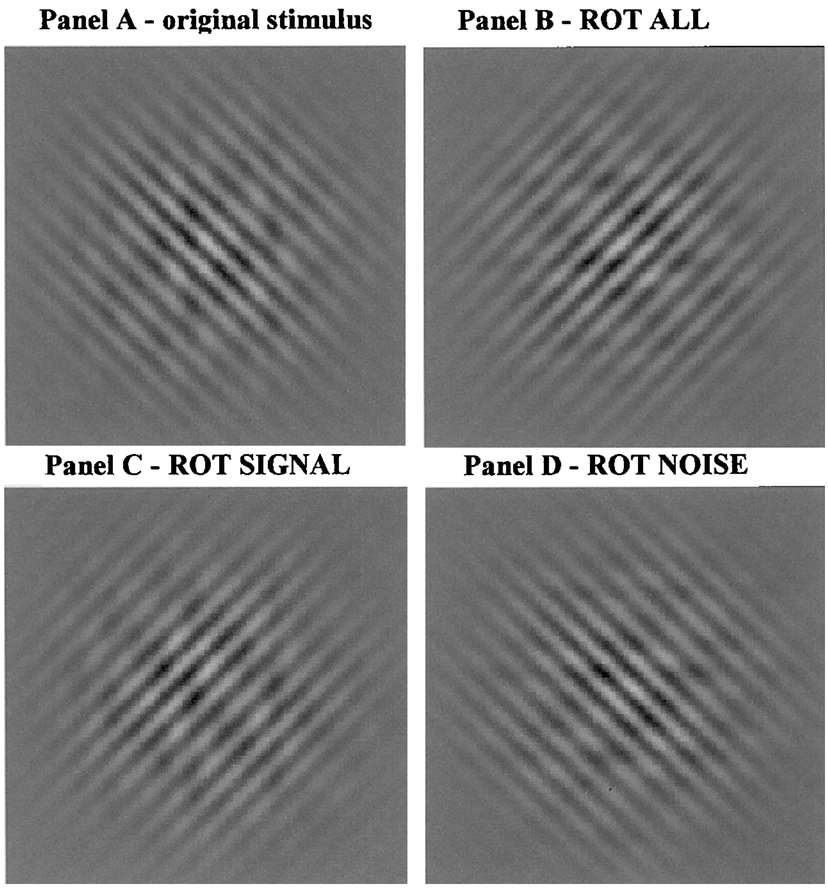

Fig. 10. Illustration of what typical (the easiest) stimuli looked like. Panel A shows the original stimulus used in the training stage of Experiment 1, Panel B shows stimulus rot all, Panel C shows stimulus rot signal and Panel D shows stimulus rot noise.

6. Experiment 4 — using spatial frequency shifts in noise components to examine selectivity

6.1. Introduction

The complex plaids used in our experiments lie somewhere between the simple gratings used in detection tasks and the complex objects used in object recognition and identification tasks. In discussing Experiments 1 and 2 it was assumed that our task, like other similar tasks (Olzak & Thomas, 1999), was mediated by mid level mechanisms combining information over wide ranges of spatial frequency and orientation. Experiment 4 examines whether observers' performance is indeed based on a Fourier representation of the stimulus, rather than a spatial representation as would be expected if performance was based on high level mechanisms (Desimone et al., 1985).

Observers showed significantly more learning when noise components were present (Experiment 1) than they did when only the signal components were present (Experiment 2). This additional learning may be partly due to observers learning which components (the signal components rather than the noise components) of the stimulus are informative. Initially observers have no information as to which regions of Fourier space will prove useful for the task. However, the signal components were always of the same orientation and spatial frequency, so it is possible for observers to learn which regions of Fourier space contain the signal components and are therefore informative. Observers might base

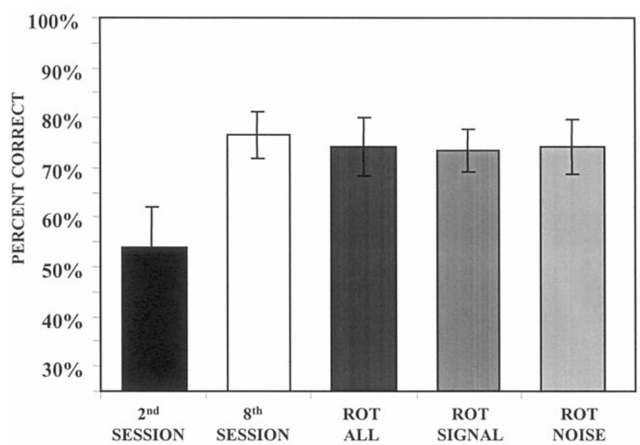

Fig. 11. Performance in the second and eighth sessions of training, and conditions rot all, rot signal and rot noise, averaged across five observers. The x-axis represents the condition and the y-axis represents the probability correct. Standard error bars are shown.

their responses on more selective regions of Fourier space as a function of practice.

Experiment 3 showed almost complete transfer of learning when either the absolute or the relative orientation of the stimuli was changed. These results exclude the possibility that learning is based upon mechanisms tuned for both spatial frequency and orientation, as such mechanisms would fail to show transfer of learning. However it remains possible that mechanisms selective for orientation but broadly tuned for spatial frequency (the 'cigar' mechanisms of Olzak & Thomas, 1999) or mechanisms selective for spatial frequency but broadly tuned for orientation ('doughnuts') subserve improvements in performance.

It is also possible that observers' improvements in performance are not due to basing responses on more selective regions of Fourier space. It may be that practice improves observers' ability to detect shifts in spatial frequency regardless of where they occur. Or it may be that observers develop templates not closely related to the stimulus' Fourier content.

An additional advantage of using compound gratings as our stimuli is that, though they are complicated, their structure can be manipulated in a systematic way to investigate the selectivity of the mechanisms underlying the improvement in performance. Experiment 4 uses a novel 'shifting noise components' technique to examine how performance is affected by shifts in the spatial frequency of the noise components. In all the experiments of this paper except Experiment 4 the spatial frequency and orientation (though not the phase) of all the noise components remained constant through each experiment. In Experiment 4 the effects of allowing particular noise components to randomly shift in spatial frequency between the four intervals of each trial was examined. The degree of masking caused by these shifts in spatial frequency can be used to indicate the regions of Fourier space on which observers based their responses.

Observers might be equally sensitive to shifts in spatial frequency regardless of where they occurred in Fourier space, as one would expect if observers did not base their responses on selective regions of Fourier space. Alternatively observers might be differentially sensitive to shifts within noise components depending on where they were situated in Fourier space.

One difficulty with transfer of learning studies is that mechanisms that are normally unselective for spatial frequency, spatial position, eye of origin etc., may become more selective as a result of extensive training with stimuli that remain constant in regard to these properties (Mollon & Danilova, 1996). One advantage of Experiment 4 is that an increase in selectivity should lead to greater transfer of learning to the shifting noise conditions rather than less: the more selective observers are in Fourier space the less susceptible they will be to shifts in the noise components.

6.2. Methods

The display and task were again identical to those described in Section 2. Three observers were trained for eight sessions. 3/5ths of the trials were the same as in Experiment 1. The remaining trials were noise shift conditions, where selected noise components shifted in spatial frequency randomly across the four intervals. Noise shift trials were randomly interleaved with normal trials.

We examined four noise shift conditions. In condition *vertical shift* (Fig. 12A) the vertical noise components were shifted $\pm 15\%$ in spatial frequency randomly across the four presentations of each trial. The vertical noise components were not of either the same spatial frequency or the same orientation as the signal components.

In condition *oblique shift* (Fig. 12B) the oblique noise components were shifted $\pm 15\%$ in spatial frequency. The oblique shift stimulus would be expected to impair performance if observers based their responses on regions of Fourier space that were either of the same spatial frequency or orientation as the signal components.

Fig. 12C shows the *same orientation shift* condition. The stimulus was identical to the *oblique shift* condition of Fig. 12A except that the oblique noise components were shifted in spatial frequency so that they were intermediate between the LOW and HIGH signal components in log space. The shifting noise components therefore still were of the same orientation as the signal components, but were separated as far as possible in spatial frequency. This condition would be expected to impair performance if observers based their responses on regions of Fourier space of the same orientation as the signal components.

Fig. 12D shows the *same spatial frequency shift* condition. This stimulus was identical to that used in the *oblique shift* condition except that the oblique noise components were shifted in orientation, so that they were both at 90°. The shifting noise components therefore remained at the same spatial frequency as the signal components but were no longer of the same orientation. This condition would impair performance if observers based their responses on regions of Fourier space of the same spatial frequency as the stimulus.

It should be noted that there is only a small difference in contrast masking within low level analyzers between the different conditions: both vertical and oblique noise components were always present, though in conditions *within orientation shift* and *within spatial frequency shift* they are shifted in spatial frequency or orientation. However these noise components remained

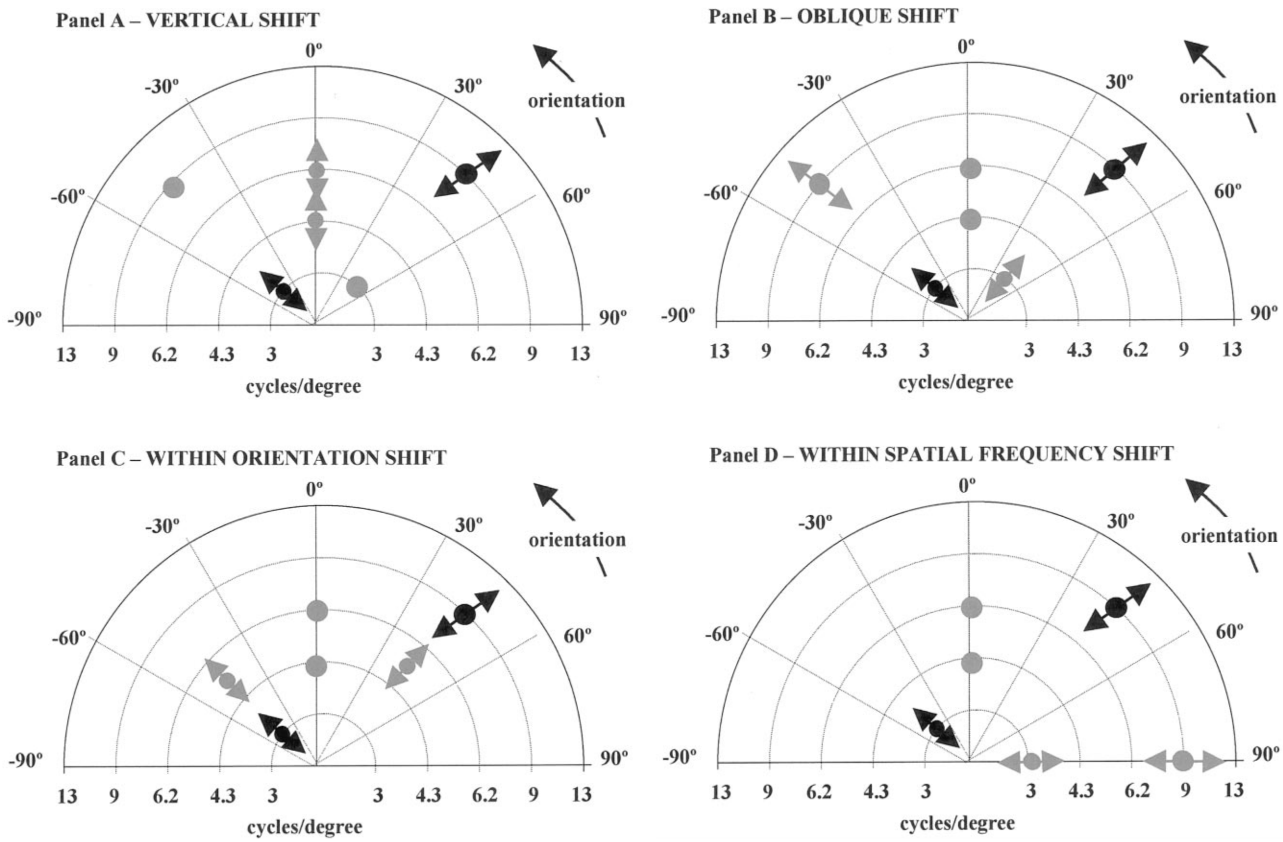

Fig. 12. Representations of the stimuli used for the shifting noise conditions of Experiment 4 in Fourier space. The x-axis represents spatial frequency and the angle represents orientation. Signal components are represented by black circles, noise components are represented by gray circles. Arrows represent the components that shift between trials. When the noise components shifted in spatial frequency they shifted randomly between the four intervals of each trial. Panel A shows the stimulus for the vertical shift condition. Panel B shows the stimulus for the oblique shift condition. Panel C shows the stimulus for the same spatial frequency shift condition and Panel D shows the stimulus for the same orientation shift condition

separated from the signal components by 45° of orientation or approximately an octave of spatial frequency. These shifts in the noise components are not likely to result in radically different amounts of contrast masking: instead it is the irrelevant shift in spatial frequency that makes the task difficult.

6.3. Results and conclusions

Fig. 13 shows performance for easy and hard trials averaged across the three observers for the last three sessions of training. Performance was significantly worse ($P < 0.05$) for hard trials than for easy trials. The x-axis represents whether the trials were easy or hard and the y-axis represents the probability correct averaged across the three observers. The black diamonds represent performance for the original stimulus, where none of the components shifted in spatial frequency.

Observers' performance was affected by the shifting noise components in almost all conditions. For every observer, during the last three sessions there was a significant (at least $P < 0.05$) drop in performance for the shifting noise conditions compared to the trials where the noise components did not shift in spatial frequency.

The effects of the noise components seemed to be greater during hard trials (though interactions between task difficulty and the effects of noise were not significant). It was not surprising that the effects of spatial frequency shifts in the noise components were larger when the noise components were of higher contrast than the signal components. Because each noise shifting condition only contained 10% of the total number of trials we used a slightly broader definition of easy and hard then that used to distinguish easy and hard trials in Experiment 1. We defined easy trials as those where both signal components were greater than the median contrast (11% for the high spatial frequency signal component and 3.2% for the low spatial frequency signal component). We defined hard trials as those where both signal components were of the median contrast or less than median contrast.

After training observers appeared to base their responses on selective regions of Fourier space. In the case of the hard trials observers were more sensitive to shifts in noise components that were either of the same orientation (*oblique* and *within orientation shift* conditions) or of the same spatial frequency (*oblique* and *within spatial frequency shift* conditions) as the signal components. Observers were less sensitive to shifts in the vertical noise components. Observers therefore seemed to base their responses on regions of Fourier space that were of the same orientation or spatial frequency as the signal components.

In the case of the easy trials the pattern of results was less clear. Observers were least sensitive to shifts in noise components of the same spatial frequency and orientation as the signal components (*oblique shift*) and were most sensitive to shifts in the vertical noise components (*vertical shift*).

This pattern of selectivity did differ between observers, and it would be interesting to see whether observers' selectivity maps would become more uniform as a function of more extensive practice. In any case, it seems that observers do become selective in the regions of Fourier space for which they base their discriminations as a function of training. We also carried out a variant of Experiment 4 in which observers were only tested with the shifting noise trials at the end of training[6]. Observers showed a similar pattern of performance under these conditions.

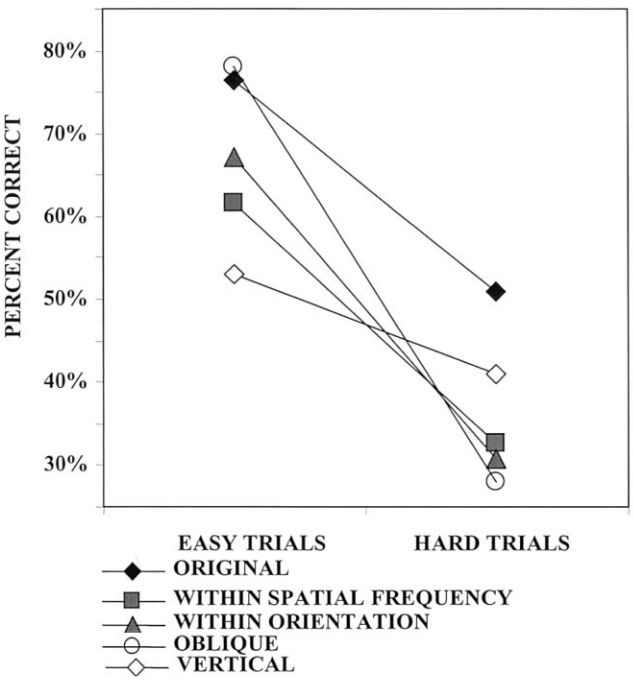

Fig. 13. Performance for the original stimulus, vertical shift, oblique shift, within spatial frequency shift and with orientation shift conditions for easy and hard trials

[6] Performance was pre-tested for all four shifting noise conditions. Observers were then trained using the stimulus and procedure described in Experiment 1 for 6 days. At the end of training performance for the four shifting noise conditions was re-tested. Results were very similar to those described in Experiment 4 — observers did show selectivity in Fourier space, however the pattern of selectivity differed between observers.

There are a number of ways by which observers' responses might become based upon particular regions of Fourier space with practice. One possibility is that observers learn to attend to signals from the relevant regions of Fourier space. This attentional modulation might take place at various stages of processing, attentional effects have been recorded throughout striate and extra-striate cortex (Motter, 1993; Gandhi, Heeger, & Boynton, 1999; Martinez et al., 1999). Another is that mid level mechanisms may become better tuned for the spatial frequencies and orientations relevant for the task. Improvement might also be due to observers learning to base their responses on a more selective pool of pre-existing mid level mechanisms: observers may learn to base their responses on those mechanisms best tuned for the particular task on which they were trained.

7. Experiment 5 — transfer of learning from same-sign to opposite-sign tasks

In the experiments above the data was modeled using two assumptions: (I) that the probability of detecting each component monotonically increased as a function of contrast in a way that could be modeled by a Weibull function; and (II) that detection of each component was based on independent probability summation.

This simple model provided a reasonably good fit to the data (as shown in Fig. 5): correlation coefficients between the real data and the model predictions varied around 0.7–0.95 and were generally above 0.8 when two sessions were averaged together. Differences between the model and the data were averaged across sessions 1–2, 3–4, 5–6 and 7–8 for each observer. There were no consistent deviations between the model fits and the data for any of the five observers. Model fits tended to improve with practice, however there were no systematic tendencies in the data across sessions that were not fit by the model. In addition, visual inspection of the data indicated no strong or consistent linear or nonlinear biases, suggesting that the performance for the two components was indeed separable. Differences between model and data appeared to be due to noise, suggesting that our model might be a reasonable starting place for describing observers' performance.

In previous experiments in half the trials the two components moved in the same direction in frequency space — a same-sign shift. LOW− HIGH− was discriminated from LOW+ HIGH+. The relative spatial frequencies of the two components did not change. In the other half of the trials the two components moved in opposite directions in frequency space

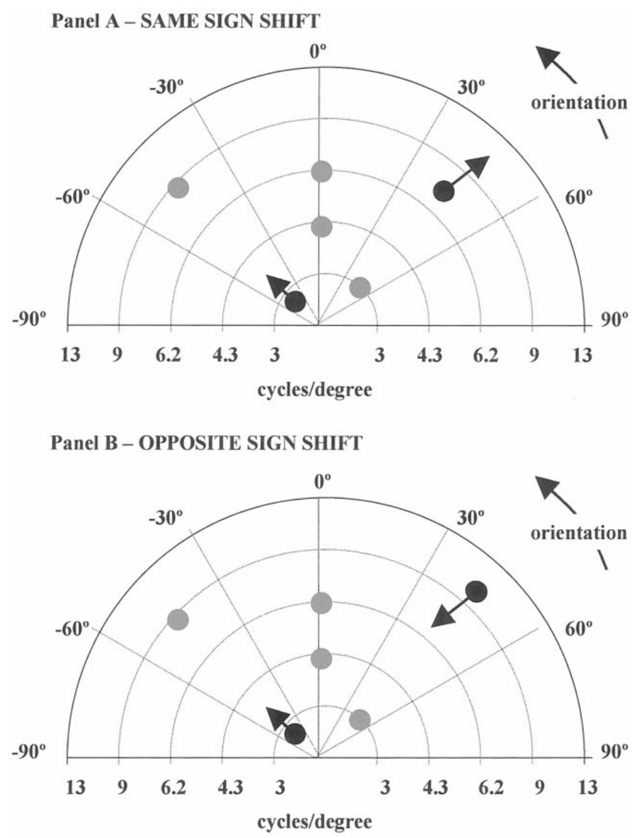

Fig. 14. Polar plot of the stimuli used for the transfer of task experiment. Panel A shows the same-sign shift stimulus and Panel B shows the opposite-sign shift stimulus

— an opposite-sign shift. LOW − HIGH + was discriminated from LOW + HIGH −. Whether the trial was same or opposite-sign was randomized across trials for each observer. Performance for same and opposite-sign shifts was compared for the last three sessions of Experiment 1 in three observers (the necessary information was not recorded for BC and SJ). No systematic difference in performance for same and opposite-sign shifts in spatial frequency was found across these three observers. Only observer SG was significantly better at the opposite-sign task, with a probability of $P < 0.05$. These results are consistent with data showing no difference in performance for same and opposite-sign tasks for components separated in both spatial frequency and orientation, even after extensive training (Fine & Jacobs, 1998; Olzak & Thomas, 1999) under experimental conditions similar to those used in Experiment 2.

The equal performance for same and opposite-sign spatial frequency shifts, and the lack of evidence for any systematic interactions between the two components, is compatible with discrimination being based upon some combination rule, such as probability summation, allowing independent discrimination of each component. As Olzak and Thomas (1991, 1999) have argued, if observers used a non-independent integration rule then there might be differences in performance between same and opposite-sign tasks. For example, if observers computed the average spatial frequency of stimuli performance would be better for the same-sign task, where there is a shift in the average spatial frequency, than for the opposite-sign task, where the shift in average spatial frequency is much smaller (nonexistent if one assumes logarithmic scaling of spatial frequency). However other ways of integrating information non-independently might be compatible with equal performance in same and opposite-sign tasks.

Though Experiment 4 suggests that much of the improvement in performance shown by observers is due to increased selectivity in Fourier space, it is also possible that with practice observers become better at combining information from the two signal components. Experiment 5 tested whether better integration of information from the two components was partially responsible for observers' improvements in performance. In addition Experiment 5 also provided a stronger test of independent detection of each component. Transfer of learning from a same-sign task to an opposite-sign task and vice versa was examined. We were interested in whether observers' performance would continue to look like they were using an independent combination rule such as probability summation even after training with stimuli where the shift in spatial frequency of the underlying signal components was correlated. Complete transfer between same and opposite-sign tasks would suggest that the second stage mechanisms combining information across Fourier space did not distinguish between same and opposite-sign tasks, thereby providing further evidence for an independent combination rule similar to probability matching. Failure of transfer would provide evidence against a combination rule like probability matching, suggesting that different mechanisms were responsible for performance in same and opposite-sign tasks, or that training had modified the combination rule of a single mechanism to favor the trained task over the untrained task.

7.1. Methods

Display and task were similar to that described in Section 2 with the difference that in the training stage observers were only presented with the same-sign task and in the testing stage transfer to the opposite-sign task was tested, or vice versa. The observers were first trained on the same-sign task to asymptote (as in Experiment 1) and then tested with the opposite-sign task, or were trained with the opposite-sign task and tested with the same-sign task. Fig. 14 shows same and opposite-sign tasks in Fourier space. Fig. 15 what the stimuli looked like. The same-sign task required observers to distinguish between the stimuli in Fig. 15 Panels A and D, the opposite-sign task required observers to distinguish between the stimuli in Panels B and C.

In the same-sign task, shown in Fig. 14 Panel A, observers had to distinguish between LOW − HIGH − and LOW + HIGH +. In Fig. 14 Panel A both the arrows point towards higher spatial frequencies to illustrate that both signal components moved in the same direction in spatial frequency space. In the opposite-sign task, observers had to distinguish between LOW + HIGH − and LOW − HIGH +; in Panel B the arrows point in opposite directions in frequency space.

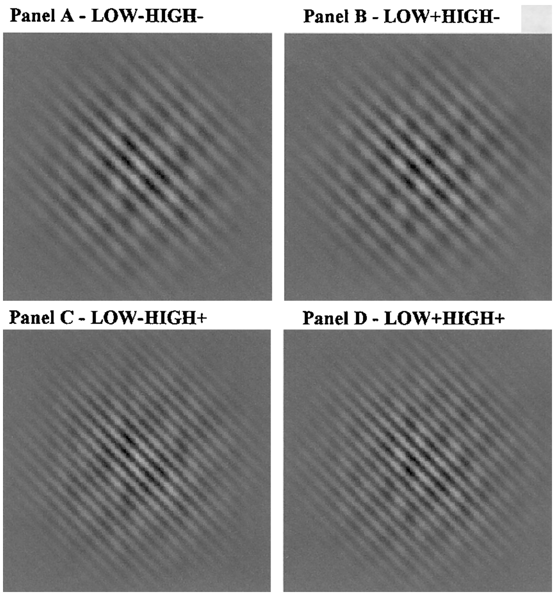

Fig. 15. Illustration of what the easiest stimuli looked like. The signal components are at their maximum contrast. Panel A shows LOW − HIGH −, Panel B shows LOW + HIGH −, Panel C shows LOW − HIGH + and Panel D shows LOW + HIGH +. The same-sign task required observers to distinguish between the stimuli in Panels A and D, the opposite-sign task required observers to distinguish between the stimuli in Panels B and C.

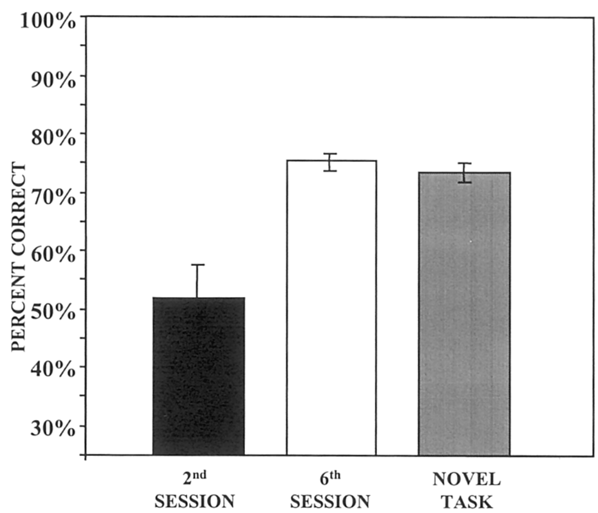

Fig. 16. Performance in the second and sixth sessions of training, and performance in the novel task, averaged across three observers. The black bar represents performance on the 2nd day of training, the white bar represents performance on the sixth day of training. The gray bar represents performance on the novel task. The novel task was the same-sign task for observers trained with the opposite-sign task, and was the opposite sign task for observers trained with the same-sign task.

Two observers were trained with the same-sign task and then transfer of learning to the opposite-sign task was tested, and two other observers were trained with the opposite-sign task, and transfer of learning to the same-sign task was tested.

7.2. Results and conclusions

Like Experiment 1, observers showed significant learning over the six sessions. One observer, CG, who showed no learning throughout the study has been excluded[7]. The average improvement in performance in Experiment 5 was 23.6%. The amount of learning did not differ significantly from Experiment 1, either with or without the excluded observer CG. (However observers showed significant variability in their improvement, which may have hidden a small difference in the amount of learning between Experiments 1 and 5). The differences between the first and last session, and between the second and last session, were highly significant ($P < 0.01$) for the three out of four observers that showed learning.

If different mechanisms mediated performance in same and opposite-sign trials then one might expect improvement in performance in Experiment 5 to be twice as fast as Experiment 1. In Experiment 1 each daily session contained 125 same-sign trials and 125 opposite-sign trials while in Experiment 5 each daily session contained 250 trials of either same-sign or opposite-sign trials. However the rate of learning was not significantly faster in Experiment 5 than in Experiment 1, suggesting that same and opposite-sign trials were mediated by the same mechanisms.

Fig. 16 shows the amount of transfer from same-sign to opposite-sign task or vice versa averaged across observers. The first bar represents performance for the second session of training. The second bar represents performance after six sessions of training with either the same-sign or the opposite sign task. The third bar represents performance in the novel task (the same-sign task if observers were trained with the opposite sign task, and vice versa). If observers showed complete failure of transfer then we would expect performance

[7] In all we ran 17 subjects using ‘wicker’-like stimuli (some experiments are not reported in this paper). CG is the only observer of the 17 who failed to show significant learning, and his performance was near chance in every session. When CG is included in our analysis the average improvement in the probability correct was 17%.

for the novel task to drop to the level of the second session. The almost complete transfer of learning shown in this experiment suggests that observers use an independent combination rule such as probability summation, even when trained with stimuli where the shifts in spatial frequency of the underlying components are correlated.

Olzak and Thomas (1991) have already noted that there is no difference in performance between same and opposite-sign tasks when the underlying components are widely separated in both spatial frequency and orientation. However they did find a difference in performance between same and opposite-sign tasks when the underlying components had similar spatial frequencies or orientations. They argue for the existence of specialized mechanisms pooling over spatial frequency or orientation. However, failure to find a difference in performance for same and opposite-sign mechanisms does not necessarily imply independent discrimination of the two underlying components. Olzak and Thomas used a concurrent response paradigm (observers were asked to make separate judgments about each component, for a fuller description see Olzak and Thomas, 1991) and found non-independent responses for components of similar spatial frequencies or orientations, but independent responses for components separated in both spatial frequency and orientation. Our results support those of Olzak and Thomas — even after a significant amount of training with only same or opposite-sign shift tasks observers showed almost complete transfer between the two tasks, suggesting independent discrimination of each of the two underlying components.

It should be noted that the almost complete transfer shown by observers is a much stronger test of independent discrimination for each component than the equal performance for same-sign and opposite-sign tasks demonstrated in Experiment 1. Experiment 5 shows that learning is not specific for same or opposite-sign discriminations. This transfer between the two tasks suggests that same-sign and opposite-sign tasks are mediated by the same mechanisms. Moreover the mechanisms combining information from low level analyzers are insensitive to the direction of the shift in spatial frequency in the two underlying components, since even with training observers do not become more sensitive to same-sign over opposite-sign tasks or vice versa.

Interestingly, same-sign shifts imply a change in scale (or a change in the distance of the stimulus from the observer), while opposite-sign shifts (resulting in a change in the relative spatial frequencies within the stimulus) imply a change in shape. The idea of scale invariance is strongly ingrained within the literature on high-level object recognition. These data suggest that if scale invariance exists in the visual system it occurs higher in visual processing than our task, as the mechanisms underlying learning in our task did not distinguish between changes in scale and changes in shape.

Although the stimuli used in the experiments above are not particularly complicated compared to the images of real objects, they are significantly more complicated than the single grating, or masking stimuli that are usually modeled with independent probability summation models. It is interesting that our data can be so easily fit with such a model, and that independent discrimination of each component is further supported by the data of Experiment 5. These data, supported by that of Graham and Sutter (1998) and Olzak and Thomas (1999), suggest that pattern processing remains 'roughly Fourier' even for mid level tasks, and that further exploiting many of the techniques developed for studying low level vision may therefore prove productive in studying mid level pattern vision.

8. Summary and general conclusions

One of the difficulties with studies of perceptual learning is that when complicated stimuli are used it is often difficult to determine what has been learned. Although studies of low level learning have been relatively successful in modeling changes in underlying mechanisms as a function of training, experiments using more complicated stimuli have only been able to use relatively crude transfer studies. Our stimuli are relatively complicated, but are also amenable to manipulations allowing the selectivity of underlying mechanisms to be probed in a relatively sophisticated manner. This paper describes several experiments examining how observers learn to put together information across a wide range of Fourier space: both how information is selected and how it is combined.

Experiment 1 demonstrated that observers show a large amount of perceptual learning for relatively complex pattern stimuli consisting of signal and noise gratings. Such stimuli are useful to study in the context of perceptual learning because it is easy to manipulate them in order to examine what is being learned.

In Experiment 2 it was shown that observers showed less learning when presented with simple plaid stimuli without the noise components. These results suggest that the learning of Experiment 1 was not due either to observers developing generalized task strategies (e.g. learning when to blink) or due to changes in low level analyzers tuned for both spatial frequency and orientation.

Experiments 3 and 4 examined the selectivity of mechanisms underlying learning. Experiment 3 suggested that observers do not learn to base their responses on regions of Fourier space localized in both spatial frequency and orientation: observers showed almost complete transfer when the signal components were rotated 90°, thereby changing the position of the signal components in Fourier space. Nor could learning be explained in terms of observers developing 'template matching' mechanisms: observers showed almost complete transfer when the noise components were rotated 90°, thereby changing the appearance of the stimulus. Experiment 4 showed that although observers do not become selective for both spatial frequency and orientation, they do base their responses on select regions of Fourier space. Observers were more sensitive to shifts in the noise components that were of the same spatial frequency or orientation as the stimulus components, or were close to the stimulus components in Fourier space. However observers varied significantly in where in Fourier space they based their responses. The improvement in performance as a function of practice might be due to observers learning to attend to the relevant signal components, developing suitable mid level mechanisms, or to learning to base their responses on the most relevant pre-existing mechanisms. Distinguishing these possibilities is a promising future direction for research.

Experiment 5 examined how these mechanisms combined information from different components — whether each component was discriminated independently or some non-independent integration function was used. Surprisingly, observers showed almost complete transfer between same and opposite-sign tasks, supporting the conclusion that independent probability summation or some other independent function was a reasonable first approximation of observers' performance, even when observers were trained on a task where the components did not vary independently.

By using complex gratings it is possible to carry out experiments to explore how observers select and combine information from low level analyzers, and in particular the effects of practice. We have found that a large amount of learning takes place at a stage of processing, higher than the low level analyzers, that combines information from signal components independently and remains 'roughly Fourier'. Observers' improvement in the task was mainly due to basing their responses on more selective regions of Fourier space as a function of practice. Observers based their responses on regions of Fourier space that were of the same orientation or spatial frequency as the signal components, or were close to the signal components in Fourier space. As our knowledge of the mechanisms underlying mid level tasks increases it should be possible to ask increasingly refined questions about the role of mid level mechanisms, and in particular, the role adaptability plays in allowing such mechanisms to represent an unpredictable world.

Acknowledgements

This work was supported by NIH grants R29-MH54770, P30-EY01319 and EY 01711. We would like to thank R. Aslin, G. Boynton, K. Dobkins, N. Graham, D. MacLeod, W. Makous, L. Olzak and two anonymous reviewers for their advice and comments, and E. Bero, L. O'Brian, A. Pauls and M. Saran for help conducting the experiments.

References

Ahissar, M., & Hochstein, S. (1995). How early is early vision? Evidence from perceptual learning. In T. V. Papathomas, C. Chubb, A. Gorea, & E. Kowler, *Early vision and beyond*. Cambridge, MA: MIT Press.

Ahissar, M., & Hochstein, S. (1996). Learning pop-out detection: specificities to stimulus characteristics. *Vision Research*, *36*(21), 3487–3500.

Ahissar, M., & Hochstein, S. (1997). Task difficulty and the specificity of perceptual learning. *Nature*, *22*, 401–406.

Ball, K., & Sekuler, R. (1982). A specific and enduring improvement in motion discrimination. *Science*, *218*, 697–698.

Ball, K., & Sekuler, R. (1987). Direction-specific improvement in motion discrimination. *Vision Research*, *276*(6), 953–965.

Beard, B. L., Levi, D. L., & Reich, L. N. (1995). Perceptual learning in parafoveal vision. *Vision Research*, *35*(12), 1679–1691.

Blakemore, C., & Campbell, F. W. (1969). On the existence of neurones in the human visual system selectively sensitive to the orientation and size of retinal images. *Journal of Physiology*, *203*, 237–260.

Burr, D. C., & Morrone, M. C. (1994). The role of features in structuring visual images. In M. Morgan, *Higher-order processing in the visual system*. New York, NY: Wiley.

Campbell, F. W., & Kukilkowski, J. J. (1966). Orientation selectivity of the human visual system. *Journal of Physiology*, *187*, 437–445.

Carter, B. E., & Henning, G. B. (1971). Detection of gratings in narrow-band visual noise. *Journal of Physiology*, *219*, 355–365.

D'Zmura, M., & Lennie, P. (1986). Mechanisms of color constancy. *Journal of the Optical Society of America*, *3*(10), 1662–1672.

Derrington, A. M., & Henning, G. B. (1989). Some observations on the masking effects of two-dimensional stimuli. *Vision Research*, *29*(2), 241–246.

Desimone, R., Albright, T. D., Gross, C. G., & Bruce, C. (1984). Stimulus-selective properties of inferior temporal neurons in the macaque. *Journal of Neuroscience*, *4*(8), 2051–2062.

Desimone, R., & Schein, S. J. (1987). Visual properties of neurons in area V4 of the macaque: sensitivity to stimulus form. *Journal of Neurophysiology*, *57*(3), 835–868.

Desimone, R., Schein, S. J., Moran, J., & Ungerleider, L. G. (1985). Contour, color and shape analysis beyond the striate cortex. *Vision Research*, *25*(3), 441–452.

DeValois, K. K. (1977). Spatial frequency adaption can enhance contrast sensitivity. *Vision Research*, *17*(9), 1057–1065.

De Valois, R. L., & De Valois, K. K. (1988). *Spatial vision*. New York: Oxford University Press.

DeValois, K. K., & Switkes, E. (1980). Spatial frequency specific interaction of dot pattern and gratings. *Proceedings of the National Academy of Sciences of the United States of America*, *77*, 662–665.

DeValois, R. L., Yund, E. W., & Hepler, N. (1982). The orientation and direction selectivity of cells in macaque visual cortex. *Vision Research*, *22*(5), 531–544.

Dosher, B. A., & Lu, Z. L. (1998). Perceptual learning reflects external noise filtering and internal noise reduction through channel reweighting. *Proceedings of the National Academy of Sciences USA*, *95*, 13988–13993.

Dosher, B. A., & Lu, Z. L. (1999). Meachanisms of perceptual learning. *Vision Reasearch*, *39* (19), 3197–3221.

Fahle, M., & Edelman, S. (1992). Long-term learning in vernier acuity: effects of stimulus orientation, range and of feedback. *Vision Research*, *33*(3), 397–412.

Fahle, M., & Morgan, M. (1996). No transfer of perceptual learning between similar stimuli in the same retinal position. *Current Biology*, *6*(3), 292–297.

Fine, I., & Jacobs, R. A. (1998). Observers show differential sensitivity to changes in spatial scale and relative frequency in a complex grating discrimination task. *Investigative Ophthalmology and Visual Science*, *39*.

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

Fiorentini, A., & Berardi, N. (1981). Learning in grating waveform discrimination: specificity for orientation and spatial frequency. *Vision Research*, *21*(7), 1149–1158.

Gandhi, S. P., Heeger, D. J., & Boynton, G. M. (1999). Spatial attention affects brain activity in human primary visual cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *96*(6), 3314–3319.

Georgeson, M. A. (1992). Human vision combines oriented filters to compute edges. *Proceedings of the Royal Society of London*, *249B*, 235–245.

Graham, N. (1989). *Visual pattern analysers*. Oxford, UK: Oxford Science Publications.

Graham, N. (1992). Nonlinear processes in spatial-frequency channel models of perceived texture segregation: effects of sign and amount of contrast. *Vision Research*, *32*(4), 719–743.

Graham, N., & Nachmias, J. (1971). Detection of grating patterns containing two spatial frequencies: a comparison of single channel and multichannel models. *Vision Research*, *11*(3), 251–259.

Graham, N., & Sutter, A. (1998). Spatial summation in simple (Fourier) and complex (non-Fourier) channels in texture segregation. *Vision Research*, *38*(2), 231–257.

Hubel, D. H., & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in cat's visual cortex. *Journal of Physiology*, *160*, 106–123.

Hubel, D. H., & Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology*, *195*, 215–243.

Hubel, D. H., & Wiesel, T. N. (1977). Functional architecture of macaque visual cortex. *Proceedings of the Royal Society of London*, *128B*, 1–59.

Karni, A., & Sagi, D. (1991). Where practice makes perfect in texture discrimination: evidence for primary visual cortex plasticity. *Proceedings of the National Academy of Sciences of the United States of America*, *88*, 4966–4970.

Karni, A., & Sagi, D. (1993). The time course of learning a perceptual skill. *Nature*, *365*, 250–252.

Kulikowski, J. J., & King-Smith, P. E. (1973). Spatial arrangement of line, edge and grating detectors revealed by subthreshold summation. *Vision Research*, *13*(8), 1455–1478.

Liu, Z., & Vaina, L. M. (1998). Simultaneous learning of motion discrimination in two directions. *Cognitive Brain Research*, *6*(4), 347–349.

Liu, Z., & Weinshall, D. (1999). Mechanisms of generalization in perceptual learning. In *Advances in neural information processing systems*. Cambridge: MIT Press.

Martinez, A., Anllo-Vento, L., Sereno, M. I., Frank, L. R., Buxton, R. B., Dubowitz, D. J., Wong, E. C., Hinrichs, H., Heinze, H. J., & Hillyard, S. A. (1999). Involvement of striate and extrastriate visual cortical areas in spatial attention. *Nature Neuroscience*, *2*(4), 364–369.

Maunsell, J. H. R., & Hochstein, S. (1991). Effects of behavioural state on the stimulus selectivity of neurons in area V4 of the macaque monkey. In B. Bloom, *Channels in the visual nervous system: neurophysiology, psychophysics and models*. London, UK: Freund.

Mayer, M. J. (1983). Practice improves adult's sensitivity to diagonals. *Vision Research*, *23*(5), 547–550.

Mollon, J. D., & Danilova, M. V. (1996). Three remarks on perceptual learning. *Spatial Vision*, *10*(1), 51–58.

Motter, B. C. (1993). Focal attention produces spatially selective processing in visual cortical areas V1, V2, and V4 in the presence of competing stimuli. *Journal of Neurophysiology*, *70*(3), 909–919.

Olzak, L. A., & Thomas, J. P. (1986). Seeing spatial patterns. In K. R. Boff, L. Kaufman, & J. P. Thomas, *Handbook of perception and human performance*. New York, NY: Wiley.

Olzak, L. A., & Thomas, J. P. (1991). When orthogonal orientations are not processed independently. *Vision Research*, *31*(1), 51–57.

Olzak, L. A., & Thomas, J. P. (1999). Neural recoding in human pattern vision: model and mechanisms. *Vision Research*, *39*(2), 231–256.

Olzak, L. A., & Wickens, T. D. (1997). Discrimination of complex patterns: orientation information is integrated across spatial scale; spatial frequency and contrast information is not. *Perception*, *26*(9), 1101–1120.

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spatial Vision*, *10*, 437–442.

Saarinen, J., & Levi, D. M. (1995). Perceptual learning in vernier acuity: what is learned? *Vision Research*, *35*(4), 519–527.

Sagi, D., & Tanne, D. (1994). Perceptual learning: learning to see. *Current Opinion in Neurobiology*, *4*(2), 195–199.

Schoups, A. A., & Orban, G. A. (1996). Interocular transfer in perceptual learning of a pop-out discrimination task. *Proceedings of the National Academy of Sciences of the United States of America*, *93*(14), 7358–7362.

Schoups, A. A., Vogels, R., & Orban, G. A. (1995). Human perceptual learning in identifying the oblique orientation: retinotopy, orientation specificity and monocularity. *Journal of Physiology*, *483*, 797–810.

Vidyasagar, T. R., & Stuart, G. W. (1993). Perceptual learning in seeing form from motion. *Proceedings of the Royal Society of London*, *254B*, 241–244.

Vogels, R., & Orban, G. A. (1994). Activity of inferior temporal neurons during orientation discrimination with successively presented gratings. *Journal of Neurophysiology*, *71*(4), 1428–1451.

Wachsmuth, E., & Perret, D. I. (1997). The physiology of shape generalization (size and orientation). In V. Walsh, & J. Kulikowski, *Perceptual constancies: why things look as they do*. Cambridge, UK: Cambridge University Press.

Watson, A. B., & Pelli, D. G. (1983). QUEST: a Bayesian adaptive psychometric method. *Perception and Psychophysics*, *33*(2), 113–120.