Available online at www.sciencedirect.com
SCIENCE DIRECT
Vision Research 46 (2006) 1493–1506
Vision Research
www.elsevier.com/locate/visres

Noise masking reveals channels for second-order letters

İpek Oruç [a,*], Michael S. Landy [b], Denis G. Pelli [b]

[a] Department of Psychology, University of British Columbia, 3008-2136 West Mall, Vancouver, BC, Canada V6T 1Z4
[b] Department of Psychology and Center for Neural Science, New York University, 6 Washington Place, New York, NY 10003, USA

Received 12 April 2005; received in revised form 5 July 2005

Abstract

We investigate the channels underlying identification of second-order letters using a critical-band masking paradigm. We find that observers use a single 1–1.5 octave-wide channel for this task. This channel's best spatial frequency (c/letter) did not change across different noise conditions (indicating the inability of observers to switch channels to improve signal-to-noise ratio) or across different letter sizes (indicating scale invariance), for a fixed carrier frequency (c/letter). However, the channel's best spatial frequency *does* change with stimulus carrier frequency (both in c/letter); one is proportional to the other. Following Majaj et al. (Majaj, N. J., Pelli, D. G., Kurshan, P., & Palomares, M. (2002). The role of spatial frequency channels in letter identification. *Vision Research, 42*, 1165–1184), we define "stroke frequency" as the line frequency (strokes/deg) in the luminance image. That is, for luminance-defined letters, stroke frequency is the number of lines (strokes) across each letter divided by letter width. For second-order letters, letter *texture* stroke frequency is the number of carrier cycles (luminance lines) within the letter ink area divided by the letter width. Unlike the nonlinear dependence found for first-order letters (implying scale-dependent processing), for second-order letters the channel frequency is half the letter texture stroke frequency (suggesting scale-invariant processing).

Keywords: Letter identification; Second-order vision; Critical-band masking; Scale invariance; Channel switching

1. Introduction

The mechanisms for detecting and localizing patterns defined by variations of luminance have been studied extensively over the past forty years. It is widely accepted that such patterns are encoded using a bank of linear filters that are spatially localized and selective for spatial frequency. Their properties have been elucidated using such methods as summation (e.g., Graham & Nachmias, 1971), adaptation (e.g., Blakemore & Campbell, 1969), and masking (e.g., Stromeyer & Julesz, 1972; Wilson, McFarlane, & Phillips, 1983).

However, the linear filters, also commonly known as spatial frequency channels, cannot detect all types of patterns. For example, consider Fig. 1. The texture-defined edges in this image, although apparent, would not be detected by a such a linear channel and subsequent thresholding or peak detection. This type of visual pattern is an example of a *second-order* pattern. The name is derived from a generic model consisting of (1) a bank of first-order linear filters that gives enhanced responses to either the foreground or background texture, but not both, (2) a nonlinearity, and (3) a second bank of linear filters, called second-order filters. Models of this form have been proposed to account for a wide variety of texture discrimination data (Bergen, 1991; Chubb & Landy, 1991; Landy & Graham, 2004).

Much of the work used to measure the properties of both the first- and second-order filters involves detection and discrimination tasks. However, one presumes that this same machinery is used for the broad variety of tasks the

* Corresponding author. Tel.: +1 604 822 0069; fax: +1 604 822 6923.
E-mail addresses: ipek@psych.ubc.ca (İ. Oruç), landy@nyu.edu (M.S. Landy), denis.pelli@nyu.edu (D.G. Pelli).
URLs: http://bar.psych.ubc.ca/People/Ipek.html (İ. Oruç), http://www.cns.nyu.edu/~msl (M.S. Landy), http://psych.nyu.edu/pelli (D.G. Pelli).

0042-6989/$ - see front matter © 2005 Elsevier Ltd. All rights reserved.
doi:10.1016/j.visres.2005.08.016

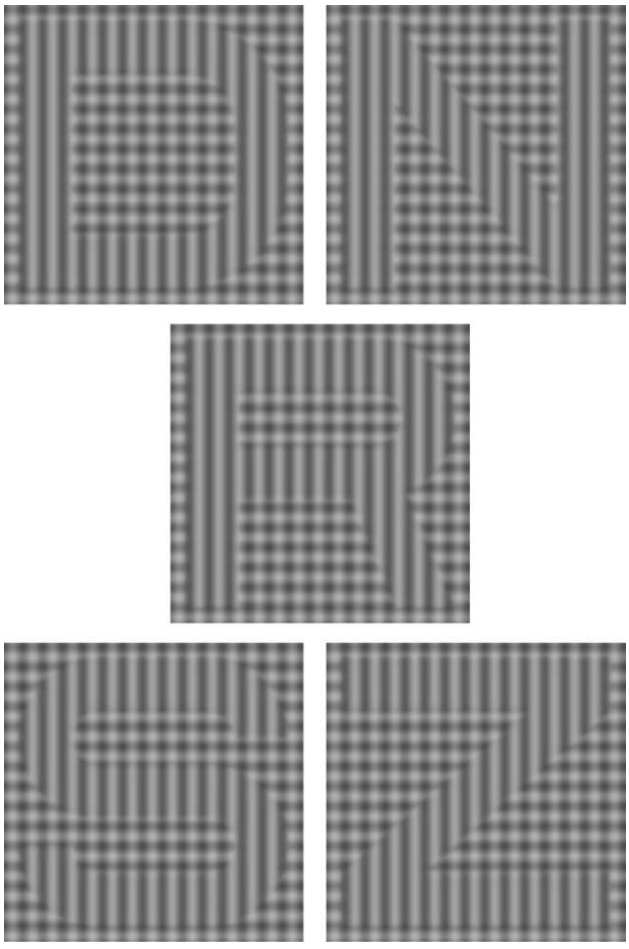

Fig. 1. Example stimuli. All five letters used in the experiment are shown with high modulation contrast and no noise added. For display purposes here, the carrier frequency (relative to the letter size) has been lowered by 2 octaves.

visual system faces in daily life. One such task is pattern identification. In this paper, we concentrate on the task of identifying letters. Others have considered which first-order channels are used for the identification of luminance-defined letters (e.g., Majaj, Pelli, Kurshan, & Palomares, 2002). In this paper, we characterize the second-order channels that observers use to identify texture-defined letters.

Studies of second-order vision have typically used methods that are second-order analogues of methods developed to study sensitivity to luminance-defined form. Typically, the experiment is translated to the second-order domain by delivering a sine wave grating, or other desired pattern, to the second-order channels. This is usually done by using the desired second-order signal to modulate the contrast of a luminance-defined *carrier* pattern. Some first-order channels are sensitive to the carrier. However, no first-order channels respond directly to the second-order signal or *modulator* because it is not present in the Fourier spectrum as power at the modulation frequency. Rather, the first two stages of the model (first-order linear filter and subsequent nonlinearity) reveal (i.e., demodulate) the signal, presenting an imperfect version of the modulator to the second-order channels.

Second-order contrast sensitivity is a measure of the observer's sensitivity to texture modulations, and is defined as the reciprocal of threshold second-order modulation contrast. A number of studies have measured second-order contrast sensitivity using a variety of second-order stimuli including contrast-modulated noise (Jamar & Koenderink, 1985; Schofield & Georgeson, 1999; Sutter, Sperling, & Chubb, 1995), orientation modulation of a texture pattern (Kingdom, Keeble, & Moulden, 1995) and modulation between two differently oriented noise textures (Landy & Oruç, 2002). Unlike the first-order case, where numerous independent measurements of the contrast sensitivity function (CSF, i.e., contrast sensitivity as a function of spatial frequency) yield similar results, there is some disagreement among second-order findings. All studies find the second-order CSF to be broadly tuned, but the shape of this broad second-order CSF has been found to be low pass (Jamar & Koenderink, 1985), band pass (Kingdom et al., 1995), or flat (all pass) (Landy & Oruç, 2002) in different studies (using different stimuli). All studies have found the second-order CSF to be *scale invariant*: second-order grating visibility is unaffected by a change in viewing distance (to a fixed stimulus on the monitor) as long as the carrier is visible. The first-order CSF is not scale-invariant; contrast sensitivity depends on spatial frequency in *retinal* units (e.g., in c/deg; Robson, 1966). However, second-order contrast sensitivity depends on spatial frequency in units relative to the carrier frequency.

Several recent studies point out intriguing similarities between letter identification and grating detection. Solomon and Pelli (1994) found that luminance-defined letters, just like gratings, are detected and identified by a single spatial frequency channel. Majaj et al. (2002) confirmed this finding. Despite the broad spectral frequency content of letters, observers use only a narrow (1.6-octave-wide) portion of this spectrum to identify them. In addition, they repeated the experiment using various letter sizes, expecting to find scale invariance. They expected to find that the observer used the same portion of the letter spectrum at all sizes. To their surprise, they found scale-dependence: observers used a slightly higher frequency band when identifying large letters and similarly a lower frequency band when identifying small letters.

Letters may be defined by differences in texture, as well as differences in luminance. Fig. 1 shows some second-order letters. When demodulated, a second-order letter looks much like a luminance-defined letter. We wondered whether second-order letters are identified using second-order channels in ways that mimic the findings for first-order letters. Here, we estimate the channel used for second-order letter identification as a function of letter size and carrier frequency.

To investigate the channels underlying second-order letter identification, we used a noise-masking paradigm (Majaj et al., 2002; Solomon & Pelli, 1994). Our results support

the hypothesis that second-order letter identification is mediated by a single second-order spatial frequency channel. The channel center frequency (in c/letter) was constant across all letter sizes, i.e., second-order letter identification is scale invariant. We also show that threshold elevation by and large is proportional to second-order noise spectral density, in parallel with the finding for luminance-defined signals (see, e.g., Pelli & Farell, 1999).

2. Experiment: Second-order channels for letter identification

2.1. Introduction

In this experiment, we employed the masking paradigm described by Solomon and Pelli (1994) to determine the power gain of the channel mediating second-order letter identification as a function of spatial frequency. Identification thresholds were measured without added noise. This no-noise threshold serves as a baseline, relative to which we calculate the threshold *elevation* caused by low- or high-pass noises with varying cut-off frequencies. To isolate and measure the second-order channel, second-order noise was employed (Section 2.2.3).

Consider first the case of low-pass masking noise. As the cut-off frequency increases, the masking effect gradually increases and saturates as the noise cut-off frequency passes through and then passes beyond the channel pass band. The slope of the threshold elevation curve is proportional to the channel power gain at the cut-off frequency. If high-pass noise is used, the logic is the same. Thus, the slopes of the high- and low-pass measurements provide independent estimates of the channel gain as a function of frequency.

If the observer uses a fixed channel, then low- and high-pass noise sweeps should yield similar estimates of the channel's tuning curve. However, the observer might use a different channel in different noise conditions. For example, with a low-pass noise the observer might switch to a channel centered on a higher frequency to improve signal-to-noise ratio. When the stimulus has a well-defined frequency (e.g., a grating) this channel switch is called *off-frequency looking* by analogy to off-frequency listening that has been reported in the auditory psychophysics literature (Lutfi, 1983; Patterson & Nimmo-Smith, 1980). Pelli (1981) and Perkins and Landy (1991) have demonstrated that channel switching occurs in vision as well. In principle, channel switching seems even more advantageous for discriminating broad-band stimuli such as letters than for detecting a sine wave grating, since broad-band stimuli provide useful information over a broad range of spatial frequencies. For luminance-defined letters, Gold, Bennett, and Sekuler (1999) found that letter identification performance for band-pass filtered letters is relatively independent of the center frequency of the filter. Parish and Sperling (1991) also found that observers could utilize a wide range of spatial frequencies for letter identification. Thus, if the default channel is masked by noise and there is information available to the observer elsewhere in the spectrum, then observers might improve performance by switching to another channel.

To discover whether our observers switch channels while performing the second-order letter identification task, we measured threshold elevation using both low- and high-pass noise sweeps. If the observer used a fixed channel, the sum of threshold elevations for low- and high-pass noise for any given cut-off frequency, f_c, should equal the threshold elevation caused by white noise. This prediction follows from the assumption that threshold elevation is proportional to the total noise power passed by the channel. If the observers were able switch channels, they would use a higher-frequency channel in low-pass noise and a lower-frequency channel in high-pass noise. By doing this, the observers would improve signal-to-noise ratio for the low- and high-pass noises, making the effect of white noise appear superadditive: greater than the sum of the threshold elevations caused by the low- and high-pass noises.

2.2. Methods

2.2.1. Stimuli

In many studies of second-order processing, stimuli are designed so as to deliver a particular stimulus to the second-order linear spatial filters by using that stimulus as a modulator of a carrier texture. In the present study, our modulator was a letter corrupted by additive, low- or high-pass filtered noise, and our carrier was a pair of sine waves oriented vertically and horizontally. Example no-noise stimuli are shown in Fig. 1.

The stimuli were constructed as follows. The modulator was a letter (D, N, R, S or Z, in the Sloan font,[1] see Fig. 6), possibly with added noise. The carrier sine waves were either 53.2 c/letter (4 c/deg at the standard 87 cm viewing distance) or 106.4 c/letter. (Note: the letter width at the standard viewing distance was 13.3°, and by c/letter we always mean cycles per letter width.) Each stimulus was a weighted combination of the two carrier patterns, where a noisy letter acted as the weight function. The vertical and horizontal sine wave carriers, S_V and S_H, were treated as contrast images (i.e., mean luminance was represented by the value zero). The noise-free letter, M, had a value of 1 inside the letter area, and 0 in the background. A stimulus L was defined as

$$L(x,y) = L_0 + A\Big([m_V(x,y)]^{1/2} S_V(x,y) + [m_H(x,y)]^{1/2} S_H(x,y)\Big), \tag{1}$$

where

$$\begin{aligned} m(x,y) &= k_{lt} M(x,y) + k_{bg}(1 - M(x,y)) + I_N(x,y), \\ m_V(x,y) &= \lceil \lfloor m(x,y) \rfloor_0 \rceil^1, \\ m_H(x,y) &= 1 - m_V(x,y), \end{aligned} \tag{2}$$

[1] The Sloan font is available at http://psych.nyu.edu/pelli/software.html.

where m_V and m_H are weights applied to the vertical and horizontal carriers S_V and S_H, respectively. The vertical carrier's weight m_V takes a value of k_{lt} in the letter and k_{bg} in the background. k_{bg} was fixed at 0.5 so that the background was a plaid (with equal vertical and horizontal energy). k_{lt} was varied to determine identification threshold. A was a fixed amplitude, set so that stimulus peak contrast used the full range of available luminance values. L_0 was the mean luminance of the display. I_N was a noise image produced by applying a low- or high-pass filter to zero-mean, white noise.

The weight functions m_V and m_H were clipped to range from 0 to 1 (Eq. (2)). This guaranteed that the square roots were well defined in Eq. (1). The phases of the carrier patterns were chosen randomly for each stimulus. Thus, $S_V(x,y)$ and $S_H(x,y)$ were independent, and their variances add. The square root in Eq. (1) ensured that the variance (and hence expected root-mean-squared contrast) was constant across the stimulus, and was not a useful cue to the task (following Watson & Eckert, 1994).

Example stimuli are shown in Fig. 1. As you can see, the background is plaid, and inside the letters' "ink" area the vertical grating is more prominent.

2.2.2. Second-order stimulus contrast

For each stimulus, the background area was a plaid. As second-order letter contrast was increased, the relative weight of the vertical component was increased in locations corresponding to the ink area of the letter. For the analysis that follows, we need a definition of second-order stimulus contrast. Any given pixel value is computed as a weighted sum of the vertical and horizontal carrier gratings, with weights of $\sqrt{m_V}$ and $\sqrt{1-m_V}$, respectively. Thus, the proportion of weight applied to the vertical carrier is

$$f(m_V) = \frac{\sqrt{m_V}}{\sqrt{m_V} + \sqrt{1-m_V}}. \quad (3)$$

We specify second-order contrast as the contrast between the values of this proportion in the letter and background. Since the letters we used cover approximately half of the stimulus area (Fig. 4), we use Michelson contrast c,

$$c = \frac{f(k_{lt}) - f(k_{bg})}{f(k_{lt}) + f(k_{bg})}, \quad (4)$$

where k_{lt} and k_{bg} are the values of m_V in the letter and background, respectively. This is analogous to defining the vertical component as "white" and the horizontal as "black" in the standard luminance contrast formulation. Note that alternative definitions of second-order contrast are possible. One could consider only the weight on the vertical component, $f(m_V) = \sqrt{m_V}$. Also, one could use the component power rather than the weight, e.g., $f(m_V) = m_V$. These alternative formulations are nearly linearly related across the range of conditions tested here, and thus using them would not change any of the conclusions made in this paper.

2.2.3. Noise

Each noise image, I_N, was low- or high-pass filtered, zero mean, Gaussian white noise. The filters had 100% modulation transfer on one side of the cut-off, and 0% on the other. The sum of a letter template and a noise pattern was then used to modulate between the two carrier patterns (Eqs. (1) and (2)). The noise did not appear as luminance noise in the stimulus, but rather as a variation in carrier texture weights. Thus, it was second-order noise. Demodulation of the stimulus yields the letter template corrupted by the noise.

For each condition we compute threshold elevation signal-to-noise ratio by dividing threshold elevation signal energy by noise spectral density. Within a condition, identification thresholds were determined by varying the signal contrast while second-order noise spectral density was fixed. Across conditions and observers, second-order noise spectral density was varied over the range 0.04–17.24 deg^2 (in an effort to keep identification thresholds within achievable bounds).

2.2.4. Artifactual cues

To ensure that the task was purely second order, it was important to eliminate all artifactual cues. The edge that separates the letter and the background is such a cue. It is true that average luminance, and luminance contrast on both sides of the edge are the same. But along the edge there is a sharp luminance discontinuity that could serve as an additional cue to the task. Therefore, we smoothed the edges by low-pass filtering the letter templates using a Butterworth filter

$$B(f) = \frac{1}{1 + (\frac{f}{f_c})^2}, \quad (5)$$

where f denotes frequency, and f_c denotes the cut-off frequency.

In addition to mitigating the edge artifacts, there was another reason to low-pass filter the letter templates: The carrier cannot effectively alternate between "features" of the modulator that are shorter than one period of the carrier itself. The carrier frequency is equivalent to the "rate of sampling" of the modulator. And so, according to the sampling theorem (Bracewell, 1978) we low-pass filtered the modulator with a cut-off frequency equal to half the carrier frequency. For example, the cut-off frequency f_c was selected to be 26.6 c/letter (30 c/stimulus image) for the lower-carrier-frequency stimuli (53.2 c/letter carrier).

We also confirmed the absence of any luminance-defined features that would allow identification of the letters. We applied a standard edge detector (Shaw, 1979) to our stimuli. Any luminance edges that outline the letters should be visible after this operation. Fig. 2 shows a luminance-defined letter and an example of our second-order letters (top row). The results of edge detection are shown in the bottom row. As expected, the edges of the luminance-defined letter were clearly detected. In the second-order case, the edges of the individual carrier elements have been

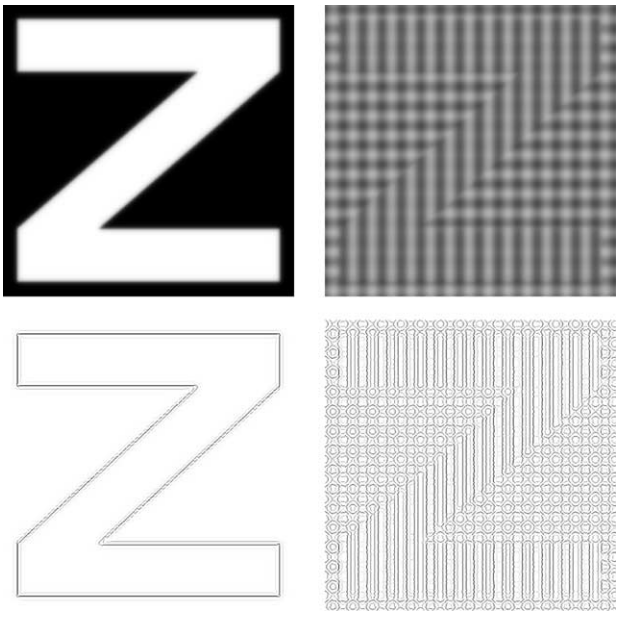

Fig. 2. Results of an edge detection algorithm (Shaw, 1979) applied to a first- and a second-order letter. In the luminance case (left column), the edges outline the letter. In the second-order case (right column), the edges of the individual carrier elements have been detected. However, no edges were found running along the letter borders.

detected. However, no continuous line segments outlining the borders of the letter were found.

One type of first-order information that might be used to perform the letter discrimination task is a difference in the image power spectra between the different letters. Fig. 3 shows two second-order letters as well as two phase-randomized versions of each (preserving their power spectra). It is clear that differences in the power spectrum per se would have been insufficient for the task, as the two right-hand images are not readily identifiable as R or S.

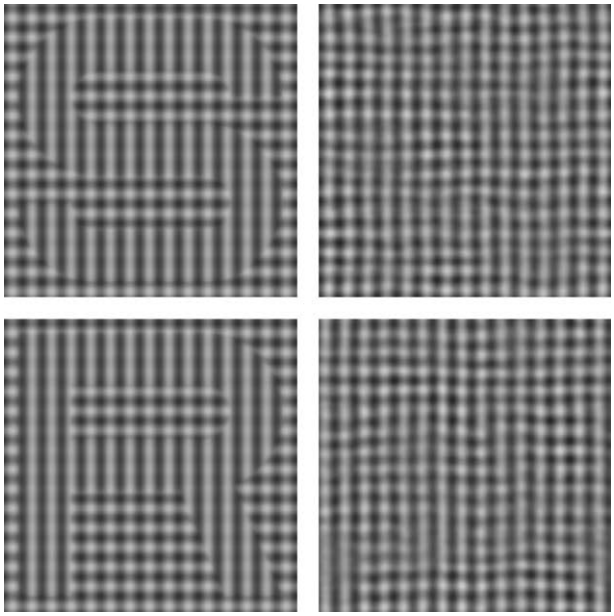

Fig. 3. Effects of randomizing the Fourier phase coefficients. On the left are two intact second-order letters. On the right are phase-randomized versions of these stimuli (a process that leaves the Fourier power spectrum of the image unchanged). The two right-hand images are not identifiable as R or S, indicating that Fourier spectral content alone was insufficient to carry out our task.

Another possible artifactual cue to the task is the "ink area" of the letter. If the ink area varied substantially, one could potentially identify letters by the total spectral power in the vertical orientation (this is again a difference in the power spectrum), regardless of the shape of the letters. Fig. 4 shows the ink area for the 10 letters of the Sloan font. We selected five letters for our experiment that have approximately equal ink areas: D N R S Z. In addition, Sloan (1951) has shown these letters to have roughly equal visibility.

2.2.5. Conditions

The first part of the experiment consisted of a total of 14 conditions: letter identification in high- and low-pass second-order noises, at six cut-off frequencies ranging from 1.66 to 20.6 c/letter, in addition to the white noise and no-noise conditions. Example stimuli with low-pass noise are shown in Fig. 5. Later, the experiment was repeated using a higher-frequency carrier (double that of the original) to see whether second-order letter channels vary with changes in the carrier. This time, 4–5 cut-off frequencies were used ranging from 6.65 to 39.9 c/letter (a total of 10–12 noise conditions including the no-noise and white noise conditions).

2.2.6. Apparatus

The stimuli were computed using the HIPS image processing software (Landy, Cohen, & Sperling, 1984). Each stimulus was 512×512 pixels. The stimuli were displayed on a SONY Trinitron GDM-G500 monitor at 848×646 resolution using a computer equipped with a CRS VSG 2/3 frame buffer, and all displays used a linearized lookup table. The mean luminance was 40 cd/m^2. The standard viewing distance was 87 cm. At this distance the whole

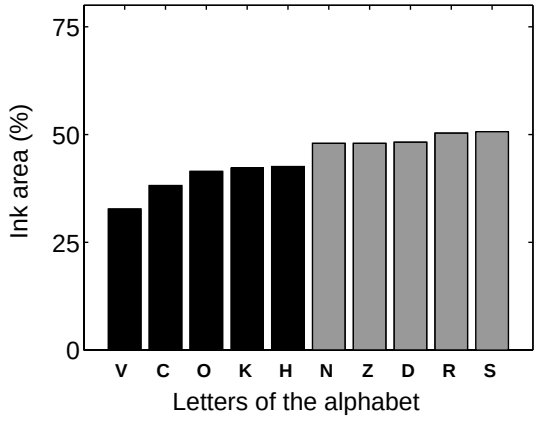

Fig. 4. Ink areas. Ink areas for all 10 letters in the Sloan font. The letters used in this study (S, R, D, Z and N, displayed in gray) have approximately equal values, all covering about half the stimulus image.

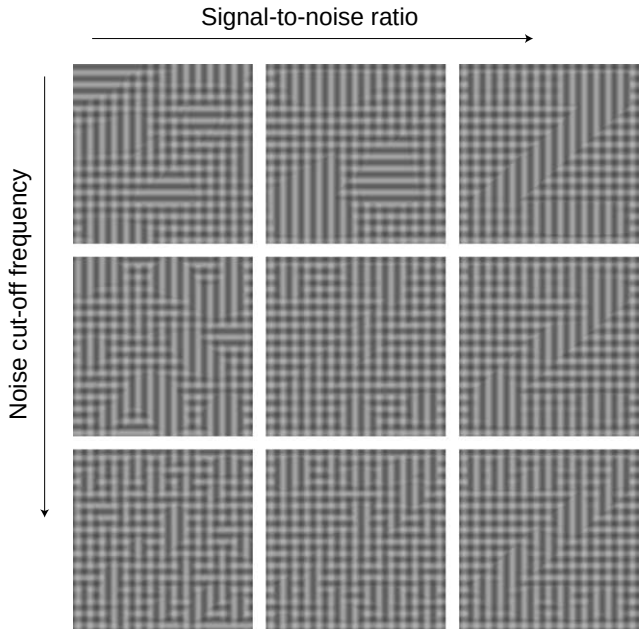

Fig. 5. Illustration of the low-pass noise conditions. Noise cut-off frequency increases from top to bottom. Signal-to-noise ratio increases from left to right. In the experiments, signal-to-noise ratio thresholds for letter identification were determined for various noise cut-offs. For display purposes, the carrier frequency (relative to letter size) has been lowered by 2 octaves.

stimulus subtended $15 \times 15\,\mathrm{deg}^2$, and the letters were approximately 13.3 deg wide. In some experiments, the same stimuli were viewed from half, double, and triple this viewing distance. Each block of trials used three versions of the five letters (using different noise images) for each contrast level. New stimuli were computed for each block of trials.

2.2.7. Observers

Four observers ran in this experiment. Observers MSL and IO are authors of this paper. JT and HB were experienced psychophysical observers who were naive to the purpose of the experiment. All four observers had normal, or corrected-to-normal vision.

2.2.8. Procedure

The task was letter identification (1-interval 5-alternative forced choice). In each trial observers viewed a fixation cue for 500 ms, followed by a 250 ms stimulus interval containing one of five possible letters. Finally, the choices screen (Fig. 6) was displayed until the observer selected a letter using the mouse. Auditory feedback indicated whether the choice was correct.

Each block consisted of 100 trials. Within each block the second-order letter contrast was controlled by two interleaved staircases, each of which ran for 50 trials. One was a 1-up-1-down staircase converging to 50% correct, and the other was a 1-up-2-down staircase converging to 71% correct. One fifth of the trials were full-contrast, no-noise trials to keep subjects motivated and remind them of what the stimuli looked like. These easy trials were not included in the analysis.

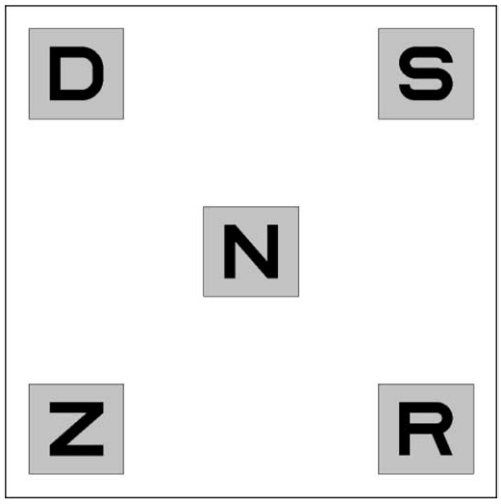

Fig. 6. The letters used in the experiment. At the end of each trial the observers viewed this screen and used the mouse to indicate which letter they thought had been displayed.

Each observer ran 2–4 blocks of each condition in random order. For the lower-frequency carrier (53.2 c/letter), two observers (MSL and IO) repeated the experiment at half the standard 87 cm distance. One observer (MSL) repeated the experiment at double the standard distance, and another (IO) at triple the standard distance. For the higher-frequency carrier (106.4 c/letter), all observers ran at the standard distance. In addition to that, three of the four observers (HB, MSL, IO) repeated the experiment at half and double the standard viewing distance. The viewing conditions are summarized in Table 1.

2.2.9. Data analysis

2.2.9.1. Contrast threshold estimates. For each condition, second-order contrast threshold was calculated by fitting a cumulative Gaussian curve to the plot of proportion correct identification (corrected for chance performance, which was 0.2) as a function of second-order contrast, using a maximum likelihood criterion. Data from each block were fitted separately and the threshold estimates were averaged. Error bars in the plots are standard errors of the threshold estimates across the 2–4 experimental blocks run per condition.

Table 1
Viewing conditions

	Distance condition			
	Half	Standard	Double	Triple
Distance (cm)	43.5	87	174	261
Letter size (deg)	26.6	13.3	6.65	4.43
Low-frequency carrier (c/deg)	2	4	8	12
High-frequency carrier (c/deg)	4	8	16	—

2.2.9.2. Channel estimates. We follow the notation of Pelli and Farell (1999) to describe our experimental design and data analysis. Contrast energy E is the squared contrast integrated over the signal. For letter stimuli, this is squared contrast multiplied by the letter area. A straightforward way to do this experiment would be to use a constant noise spectral density N across all conditions (i.e., all cut-off frequencies). This way, threshold elevation depends only on the channel's tuning function and the noise cut-off frequency, and can be used directly to derive it. Threshold elevation is defined as $E^+ = E - E_0$, where E_0 is the threshold contrast energy in the absence of noise, and E is the threshold contrast energy for a given condition with added noise.

In fact, we used various noise densities in different conditions (noise cut-off frequencies). When noise densities are not constant, threshold elevation depends not only on the tuning function of the channel and the noise cut-off frequency, but also on the density of the noise. Therefore, we normalized the threshold elevations by the noise densities, and report the threshold elevation signal-to-noise ratio

$$D^+ = \frac{E^+}{N}. \tag{6}$$

We use the symbol D^+ in analogy to d' of signal detection theory, but use a capital letter to indicate that the signal-to-noise ratio we report is in units of signal energy. The channel tuning curve for each observer was computed by differentiating the D^+ curves with respect to cut-off frequency for both the low-pass and the high-pass noise sweeps. The low- and high-pass threshold elevation D^+ data were fitted (least squares) by a sigmoid function

$$y = \frac{a}{1 + e^{-b(x - f_{\text{channel}})}} + d, \tag{7}$$

where x is log noise cut-off frequency, a is a scaling factor, b determines steepness of the curve (i.e., bandwidth of the channel), f_{channel} shifts the function along the x axis (i.e., f_{channel} is the center frequency of the channel), and d is the shift along the y axis.

If the two derived channels are identical then the observer did not switch channels. To check this, we computed *noise additivity* graphs Majaj et al. (2002). We compared the D^+ caused by white noise to the sum of D^+'s caused by a pair of low- and high-pass noises with the same cut-off frequency. If the observer used the same channel throughout, the sum of the two should equal the D^+ caused by white noise. For each estimated channel, we plot noise additivity graphs that show $D^+_{\text{lowpass}}(f_c) + D^+_{\text{highpass}}(f_c)$ as a function of the noise cut-off frequency, f_c. If the observer used a fixed channel, this curve should be flat.

2.3. Results

2.3.1. Channel switching

Noise additivity graphs for two observers and viewing distances with the lower (53.2 c/letter) and higher (106.4 c/letter) frequency carriers are plotted in Figs. 7A and B, respectively. Data for other subjects and viewing distances were similar. If there is no channel switching, the data should fall on the flat line that is the threshold elevation for white noise. The noise additivity curves are close to the predicted flat line. A consistent deviation from this prediction (especially a dip near the peak of the estimated channel) might indicate channel switching. No such consistent pattern is visible in our results. We conclude that observers used the same channel in all noise conditions.

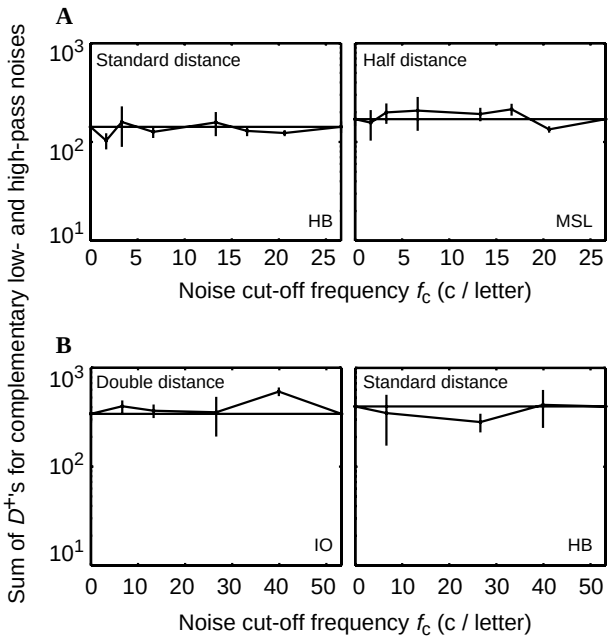

Fig. 7. Noise additivity graphs. (A) Noise additivity graphs for the lower frequency carrier (53.2 c/letter), for two observers and viewing distances. (B) Noise additivity graphs for the higher-frequency carrier (106.4 c/letter), for two observers and viewing distances. The sum of D^+'s for low- and high-pass noise masks at a given cut-off frequency is plotted as a function of noise cut-off frequency. Data were similar for all subjects and viewing distances. The flat line shows the threshold elevation for white noise.

2.3.2. Threshold elevation curves

D^+ is shown as a function of noise cut-off frequency for two observers and viewing distances for the lower (53.2 c/letter) and higher (106.4 c/letter) frequency carriers in Figs. 8A and B, respectively. Similar data were obtained for all subjects and viewing distances. The data for both the low-pass and the high-pass noise conditions are shown. For low-pass noise, a zero cut-off frequency indicates the no-noise condition and the highest cut-off frequency is white noise, or the all-pass condition. As expected, threshold elevation increased by and large monotonically with cut-off frequency in low-pass noise conditions (and decreased in the high-pass noise conditions). Note that the same sigmoid function (i.e., two identical sigmoids with opposite signs) was used to fit both the low- and the high-pass data on a given graph. The fits are quite good, as expected from the flat noise-additivity results.

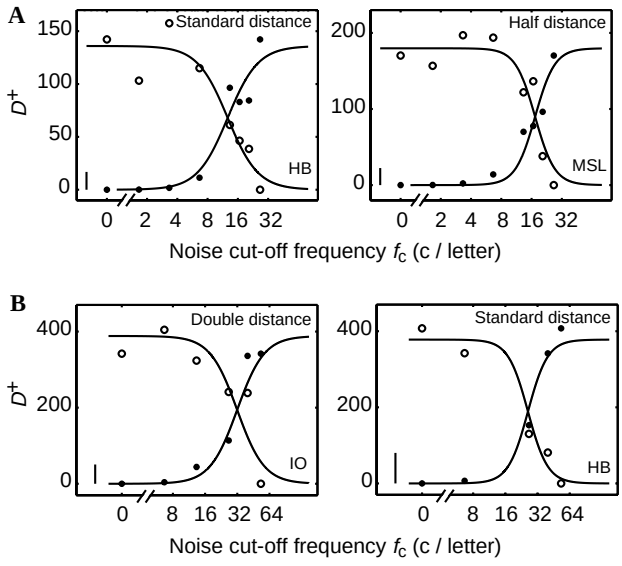

Fig. 8. Threshold elevation signal-to-noise ratio graphs. (A) Threshold elevation signal-to-noise ratio D^+ graphs for the lower-frequency carrier (53.2 c/letter) and (B) for the higher-frequency carrier (106.4 c/letter), each for two observers and viewing distances. D^+'s for low- (filled symbols) and high-pass (open symbols) noises are plotted as a function of noise cut-off frequency. The average standard error for each data graph is shown on the left of each plot. Data were similar for all subjects and viewing distances. Note the varying y-axis scales in these plots.

2.3.3. Estimated channels

The derivative of the sigmoid curve fitted to the threshold elevation data yields an estimate of the channel tuning function. The derived channels for two observers and viewing distances for the lower-frequency carrier (53.2 c/letter) are shown in Fig. 9A. At the standard viewing distance, the estimated channel center frequency averaged over all observers is 14 ± 1.2 c/letter (geometric mean $\pm$ standard deviation). There is reasonable consistency across observers.

What happened when viewing distance was changed? Note that different viewing distances were implemented by seating observers at different distances in front of the same display. Thus, both letter size and carrier frequency in c/deg changed together, leaving the ratio (carrier cycles/letter) constant across different viewing conditions for the given stimulus. If observers used the same channel (in c/letter) across all viewing distances, that would indicate scale-invariant letter recognition. That is what we found: the estimated channel frequency averaged over all observers and viewing distances was 15.2 ± 1.1 c/letter. The half-height bandwidth of the channels' power gains, where bandwidth in octaves was averaged over all observers and viewing distances, was 1.4 ± 0.6 octaves.

Fig. 9B shows the derived channels for the higher-frequency carrier (106.4 c/letter) for two subjects and viewing distances. Again we found evidence for scale invariance: At the standard viewing distance, the estimated channel center frequency averaged over all observers was 31.4 ± 1.1 c/letter. The estimated channel center frequency averaged over all viewing distances and observers was 30.9 ± 1.1 c/letter. The half-height bandwidth of the channels' power gains, averaged over all observers and viewing distances, was 1.2 ± 0.4 octaves.

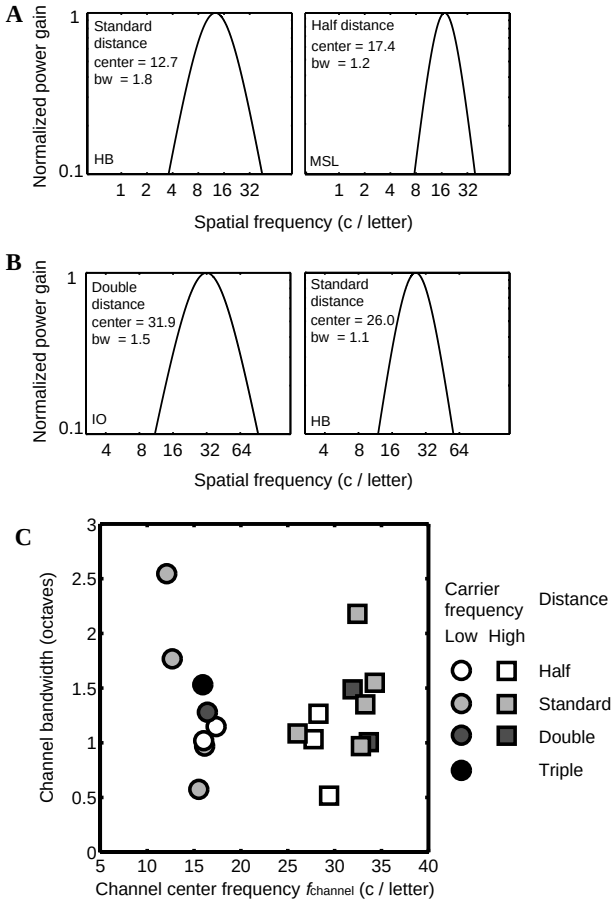

Fig. 9. Derived channels. (A) Derived channels for the lower-frequency carrier (53.2 c/letter) and (B) for the higher-frequency carrier (106.4 c/letter), each for two observers and viewing distances. The derivative of a D^+ curve yields an estimate of the channel tuning function. All derived gain functions were normalized to have a peak of 1. The center frequencies (in c/letter) and half-height bandwidths (in octaves) of the derived channels are noted on each graph. (C) Scatterplot of center frequency versus bandwidth for all channels. Circles and squares denote low- and high-frequency carrier respectively. Viewing distance is denoted by marker face color (half distance: white, standard distance: gray, double distance: dark gray, triple distance: black).

Fig. 9C shows derived channel parameters for all subjects and viewing distances, at both low (circle) and high (square) carrier frequencies. Viewing distance is denoted by marker face color: white for half, gray for standard, dark gray for double, and black for triple viewing distance. There is reasonable consistency among subjects in terms of channel center frequency, and no systematic change in bandwidth with viewing distance.

For luminance-defined letters, Majaj et al. (2002) tested whether letter identification was scale-invariant, using a wide variety of letter fonts. They anticipated that the channel used for identification would be a function of *letter frequency*. But, since letters are broad-band stimuli, one needs

a method for determining a characteristic frequency for a given font and letter size. They introduced *stroke frequency* f_{letter} for this purpose. It was defined as the number of lines a horizontal slice at half the letter height intersects, averaged over all letters of an alphabet, divided by the average width of a letter. This turned out to be a useful definition: they found that the center frequency of the channel used for letter identification depends solely on stroke frequency over a wide variety of fonts and sizes.

In Fig. 10 we plot our results as well as a summary of those of Majaj et al. (2002). Channel frequency f_{channel} is plotted as a function of the stroke frequency f_{letter}. If the frequency of the channel was identical to the nominal frequency of the stimulus, then the data would fall on the dotted identity line. One data point is plotted from a previous study (Landy & Oruç, 2002), showing the second-order channel used to detect a second-order, orientation-modulated sine wave grating. It lies on the identity line as expected. Scale invariance requires the channel frequency to be proportional to the nominal letter frequency, resulting in a line of unit slope in this plot. The solid lines are least-square fits of lines of unit slope to the data for each of the carrier frequencies we used. The fits are excellent, indicating scale invariance. The dashed line summarizes the data of Majaj et al. (2002) for first-order letter identification. The slope of this line (0.7) is less than one, indicating that first-order letter identification is not scale invariant. Like our data point for detecting a second-order sine wave, their data (not shown) for detecting first-order sine waves fall on the identity line.

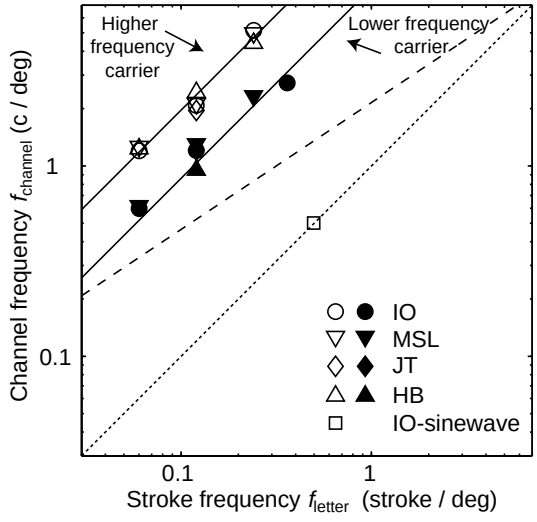

Fig. 10. Scale invariance. Channel frequency (in c/deg) is plotted as a function of stroke frequency for all observers and viewing distances for both carrier frequencies. Filled and open symbols are data for the 53.2 and 106.4 c/letter carriers, respectively. Solid lines are lines of unit slope fit to the data. Doubling the carrier frequency resulted in a doubling of the channel frequency. The open square indicates the channel frequency used to detect a second-order sine wave grating (Landy & Oruç, 2002); it lies on the dotted identity line. The dashed line summarizes the data of Majaj et al. (2002) for identification of first-order letters. Its slope is less than 1, indicating a failure of scale invariance.

When the carrier frequency was doubled, the center frequency of the channel used for letter identification doubled as well (rising from 15.2 ± 1.1 to 30.9 ± 1.1 c/letter). For the lower-frequency carrier, the channel center frequency $f_{\text{channel}} = 9f_{\text{letter}}$, where f_{letter} is the stroke frequency. For the higher-frequency carrier, $f_{\text{channel}} = 20f_{\text{letter}}$. Thus, we found that the channel frequency, relative to letter size, is independent of letter size, but does depend on carrier frequency f_{carrier} (Fig. 11). The ratio of the carrier frequency to the frequency of the letter identification channel remained constant at approx. 3:1. This suggests a link between the scales of first- and second-order channels. This is consistent with Graham (1994), who estimated a 3–4:1 ratio, although others have suggested larger ratios (Sagi, 1990; Kingdom et al., 1995).

The paradigm used in this experiment depends on the assumption that the contrast energy of a signal at threshold is linearly related to the power density of the noise passed through the channel used in the task. In Appendix A, we test this assumption and find that it is true over most of the noise contrast range, but fails at very low noise contrasts. In Appendix A, we argue that this nonlinearity has no impact on the conclusions of this paper.

3. Experiment 2: Second-order template matching

The results show that observers use one fixed channel to identify second-order letters at a fixed size and carrier frequency, and that if the letter stimulus is scaled, the channel scales with it, so the observer's performance is scale invariant. In many studies of vision, there is a concern as to whether a given result really tells us something about the visual system itself, rather than merely revealing a property that follows inevitably from the information available for the task with the particular stimuli used. This question is

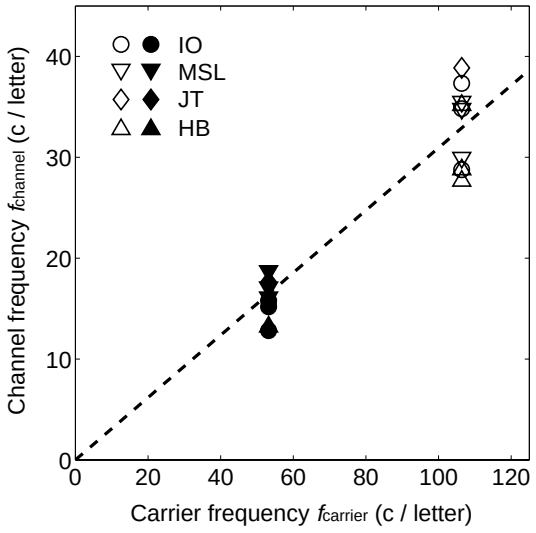

Fig. 11. Channel frequency depends on carrier frequency. Estimated channel peak frequencies are plotted for the two carrier frequencies used. The dotted line shows the line fitted to the data: $f_{\text{channel}} = 0.31f_{\text{carrier}}$. Thus, channel frequency changes proportionally with carrier frequency.

generally addressed by reference to an ideal (maximum likelihood) observer performing the same task. If human sensitivity is a constant, high fraction of ideal sensitivity, then one concludes that performance is mainly limited by the information content of the stimuli. Such an ideal observer would not replicate the scale-dependence of letter identification found by Majaj et al. (2002).

Chung, Legge, and Tjan (2002) describe a sub-ideal observer for letter identification that initially filters the stimulus with a filter matching the human contrast sensitivity function for luminance modulations, and then adds white noise. The letter identification task is then carried out on the resulting noisy, filtered images by an ideal observer. This is equivalent to having an ideal observer view the stimulus masked by additive nonwhite noise (i.e., white noise filtered by the inverse of the CSF). That noise is the only scale-dependent part of their model; the subsequent ideal observer would otherwise operate equally well at all letter sizes. They point out that such a sub-ideal observer displays scale dependence when tested with displays of filtered letters with no added stimulus noise. Yet, any critical-band experiment that succeeds in measuring large threshold elevations must be swamping the observer's equivalent noise. Thus, the additive noise assumed in their model is negligible in the critical-noise-masking paradigm, so that their model would, in that paradigm, predict scale-invariance, contrary to what is found.

The ideal observer for our task (identification of a second-order letter in low- or high-pass second-order noise) is well-defined and requires a calculation of likelihood that integrates over possible carrier phases and modulator noise images. Instead of the ideal observer, we have implemented a somewhat simpler, sub-ideal model that demodulated the stimulus, followed by template matching. Although this is not an ideal observer, we believe its threshold elevation curves are reasonably similar in shape to that of the ideal, and that it is an appropriate reference against which to compare human observers' data.

In the model, the stimulus image was convolved with a filter tuned to one of the two constituent textures. We used a filter with a Gaussian power gain function centered on 60 c/image (i.e., 53.2 c/letter, the same as the lower carrier frequency) and 90° (i.e., vertical) with a 0.5 octave frequency bandwidth and 15 deg orientation bandwidth (the impulse response was a Gabor). After the convolution, a point-wise nonlinearity (x^2) was applied to demodulate the signal. Finally, the resulting image was cross-correlated with each of five candidate letters and the letter with the highest correlation value was chosen.

The model performed the letter identification task using the same stimuli as the human observers in Experiment 1. The model was tested on every stimulus image with the lower-frequency carrier (all conditions, all contrasts, etc.) in the stimulus set for one observer. The percentage correct was computed for the model for each contrast level. A psychometric function was fit and thresholds were calculated in the same way as for the human observers.

D^+s for the model are plotted in Fig. 12A. As expected, the model performance degrades with increasing noise bandwidth in both the low- and high-pass conditions. As with the human observers, the derivative of these plots yields an estimated "channel" for the model observer, even though the model has no second-order filter (it cross-correlates with the full letter templates). These estimated channels (Fig. 12B) are quite different from the band-pass curves found for human observers. For comparison, we plot a channel (dashed curve) that represent average human data (15.2 c/letter center frequency, and 1.4-octave bandwidth), along with the model results. The model seems to peak at no higher than 2–4 c/letter, a factor of four or more below the 15.2 c/letter peak frequency for the human data. We conclude that the channels found for human letter identification measured in Experiment 1 are real prop-

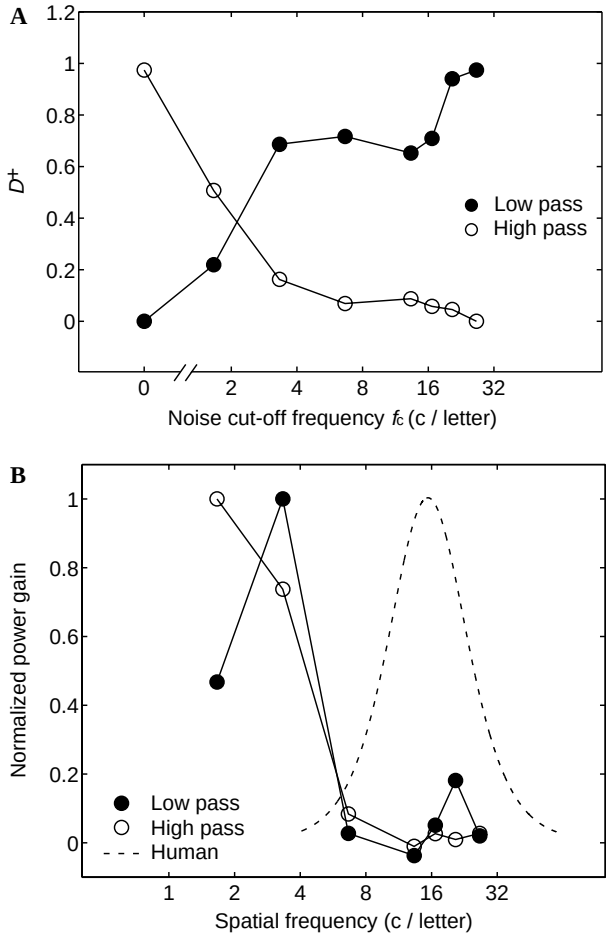

Fig. 12. Model results. (A) Threshold elevation signal-to-noise ratio D^+ is plotted as a function of noise cut-off frequency. Results are plotted for simulations of the model discussed in the text. Low- (filled symbols), and high-pass (open symbols) noise data are plotted. (B) "Channel estimates" for the model for the lower-frequency carrier stimuli. Derivatives of the model's D^+ curves are plotted for low- (filled symbols) and high-pass (open symbols) conditions. The dashed curve represents the average channel for the human observers. The shapes of the derived channels from the model are unlike those derived for the human observers. At this carrier frequency, the channel peak frequencies are around 15 c/letter for human observers.

erties of the human observers, and not inevitable consequences of the task and stimuli.

In addition, the model performed the letter identification task for different noise levels as in Experiment 3 (see Appendix A). Unlike the human observers, the model's threshold energy increased linearly over the entire range as noise density increased.

4. Discussion

It is usually assumed that observers employ the channel that is most sensitive to the stimulus in detection tasks, and most sensitive to stimulus differences in discrimination tasks. In sine wave grating detection there is clear evidence for this. Both first- and second-order studies have found channels centered at the frequency of the target sine wave. In addition, if a first-order stimulus is presented in noise, observers shift to a channel less affected by the noise (Pelli, 1981; Perkins & Landy, 1991). In other words, they use a channel with better signal-to-noise ratio.

Unlike sine wave gratings, where all power is at a single spatial frequency, letters are broad-band signals. Identification efficiencies for band-pass filtered letters are equally good for a wide range of filter center frequencies (Gold et al., 1999). When identifying broad-band signals, such as letters, it is reasonable to expect observers to utilize information at all available frequencies. At the very least, one would expect an observer to switch channels when a, perhaps preferred, portion of the spectrum is corrupted with noise. Surprisingly, we find this not to be the case. Observers identify second-order letters using a single, second-order spatial frequency channel and do not switch channels with changes in the noise spectrum. The channel used by the subject depends on the stroke frequency and carrier frequency of the letter. For the lower-frequency carrier (53.2 c/letter) observers used a channel with average peak frequency of 15.2 ± 1.1 c/letter and a 1.4 ± 0.6 octave bandwidth at all viewing distances. For the higher-frequency carrier (106.4 c/letter) average peak channel frequency was 30.9 ± 1.1 c/letter, and average channel bandwidth was 1.2 ± 0.4 at all viewing distances. Observers were unable to switch channels to avoid noise. These results are consistent with the findings of Majaj et al. (2002) for first-order letter identification.

The channel used for second-order letter identification scales with letter size, unlike the first-order case. Why does one find scale dependence in first-order letter identification, and not in second order? The ideal observer uses information in all useful frequency bands of the letter spectrum. But human observers are not limited *only* by the available letter identity information in the spectrum. We are not equally sensitive to all spatial frequencies. In the absence of added stimulus noise, one would expect that the channel used by an observer would be biased toward the peak of the CSF. This is essentially the suggestion of Chung et al. (2002) that the channels found for letter identification are determined by the combination of the human CSF and the available letter identity information in the spectrum. Perhaps this scale-dependent behavior persists, though no longer adaptive, with added stimulus noise, as in critical-band masking experiments. The second-order CSF is scale invariant (Jamar & Koenderink, 1985; Kingdom et al., 1995; Landy & Oruç, 2002; Schofield & Georgeson, 1999; Sutter et al., 1995). Thus, by a similar argument, one would expect second-order letter identification to be scale invariant as well, which is precisely what we found.

How does the channel frequency compare for first-order and second-order letter identification? We found that channel frequency depended on carrier frequency (Fig. 10) using a definition of stroke frequency based on the modulator (the Sloan font letter template). However, suppose we considered the carrier texture, although it is random, to be an intrinsic part of the second-order letter font which increases its complexity. This would increase the nominal stroke frequency much as the extra curlicues on a fancy script font increase stroke frequency. Fig. 13 replots the data of Fig. 10 using such a definition. In particular, we define letter *texture* stroke frequency f_{stroke} for second-order letters as follows. Along a horizontal slice through the letter we count the number of carrier cycles within the letter ink area and divide by the letter width. With this revised definition of stroke frequency, the data for the two carriers now

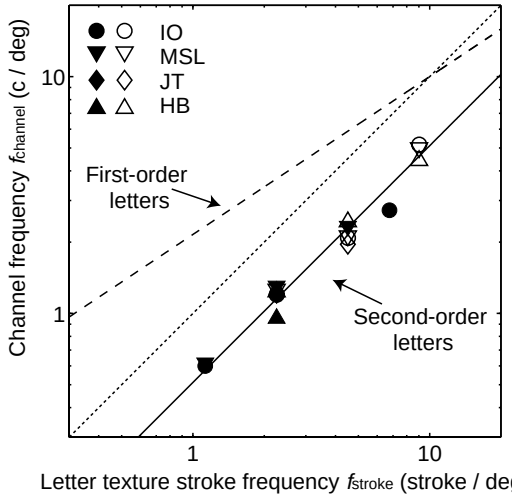

Fig. 13. Letter texture stroke frequency. The data of Fig. 10 are replotted using a revised definition of stroke frequency for the second-order letters. Here, we treat the carrier texture as if it were, although random, an intrinsic characteristic of the font, thus computing letter texture stroke frequency f_{stroke} as follows: Along a horizontal slice through the letter we count the number of carrier cycles within the letter ink area and divide by the letter width in degrees (and then average over vertical positions). With this definition, the data for both carrier frequencies (lower: filled symbols, higher: open symbols) superimpose so that, as in the first-order letter identification of Majaj et al. (2002), channel frequency depends solely on stroke frequency. Second-order channels generally have lower peak frequencies compared to first order (dashed line). The dotted identity line is shown for reference.

superimpose, so that now second-order letter identification may be described as a function solely of letter texture stroke frequency. However, the channels found for second-order letters have lower peak frequency than those for first order for similar letter stroke frequencies.

5. Conclusions

1. Threshold contrast energy to identify a second-order letter increased linearly with noise power spectral density of nonwhite second-order noise masks at all suprathreshold noise levels tested. However, to our surprise, threshold increased slightly when the noise was removed entirely (see Appendix A). Even so, these results validate the linearity assumption that underlies the procedure for estimating channel gain from threshold measurements as a function of noise spatial frequency cut-off. Each channel's gain as a function of frequency is summarized by its center frequency, and we identify the channel by its center frequency.
2. As previously found for first-order letters, the channel used by observers to identify second-order letters is independent of the noise spectrum. This implies that the channel is determined by the signal and task; the observer cannot switch to another channel to avoid the noise.
3. As previously found for first-order letters, for second-order letters the channel frequency is wholly determined by the stroke frequency of the signal.
4. For second-order letters the channel frequency is half the letter texture stroke frequency, $f_{\text{channel}} = 0.5 f_{\text{stroke}}$. Thus, identification of second-order letters is scale invariant. This is unlike first-order letters for which the channel frequency is greater than stroke frequency, following a nonlinear power law, $f_{\text{channel}} \propto f_{\text{stroke}}^{2/3}$, which is scale dependent.

Acknowledgments

The work reported here was supported by Grants EY08266 (MSL) and EY04432 (DGP) from the National Institutes of Health. We thank our anonymous reviewers for their useful comments.

Appendix A. Experiment 3: Dependence of D^+ on noise level

A.1. Introduction

The paradigm used in Experiment 1 depends on the assumption that the contrast energy of a signal at threshold is related linearly to the power density of the noise passed through the channel used in the task. In other words

$$E = kN + l, \tag{8}$$

where E is the threshold contrast energy, N is the noise power density, and k and l are constants. In our experiments, we first measure threshold in the absence of external noise ($N = 0$), so that $E_0 = l$. Threshold elevation is $E^+ = E - E_0 = kN$, and therefore threshold elevation signal-to-noise ratio is

$$D^+ = \frac{E^+}{N} = k. \tag{9}$$

In other words, if the linearity assumption is true, then D^+ is independent of the noise density at which it is measured. Since we varied noise power across conditions, this assumption is crucial, as the estimated channel shape could otherwise change if we repeated the experiment using a different set of noise densities.

Pelli (1981) showed that the contrast energy of a first-order signal at threshold is related linearly to the power spectral density of white noise. Solomon (2000) tested the assumption of linearity for high- and low-pass noise masks with a wide range of cut-off frequencies and found that it holds for most cases tested. In Experiment 1, most conditions involved nonwhite noise masks. So, in Experiment 3 we test whether threshold contrast energy is a linear function of noise power for one choice of letter size and low-pass noise.

A.2. Methods

The stimuli and task were similar to those of Expt. 1. Here, we used only one low-pass noise with cut-off at 6.65 c/letter (0.5 c/deg at the standard viewing distance). There were six conditions with noise densities ranging from 0.068 to 0.44 deg^2, plus the no-noise condition. The lower-frequency carrier was used, viewed from the standard viewing distance (87 cm). Three observers participated in this experiment.

A.3. Results

Fig. 14 shows the results for all observers. Threshold contrast energy, E, is plotted as a function of noise density, N. The results are surprising: for all subjects, the addition of a small amount of noise to a noise-free stimulus *reduces* second-order letter identification contrast threshold. Additional noise results in a steady, linear increase in threshold. We do not have a good explanation as to how noise aids identificaiton. A little bit of noise improves second-order letter identification, but that effect is overwhelmed by noise masking as noise power is increased.

The results in Fig. 14 may remind the reader of the "dipper functions" that result from increment threshold experiments (e.g., Nachmias & Sansbury, 1974), wherein contrast discrimination of a test pattern is easier on a pedestal of that same pattern having near-threshold contrast than on a uniform background. However, the results in Fig. 14 show an improvement in letter identification on a background of noise, unrelated to the pattern to be identified. Our result is analogous to the finding of stochastic resonance in some biological systems (e.g., Bezrukov &

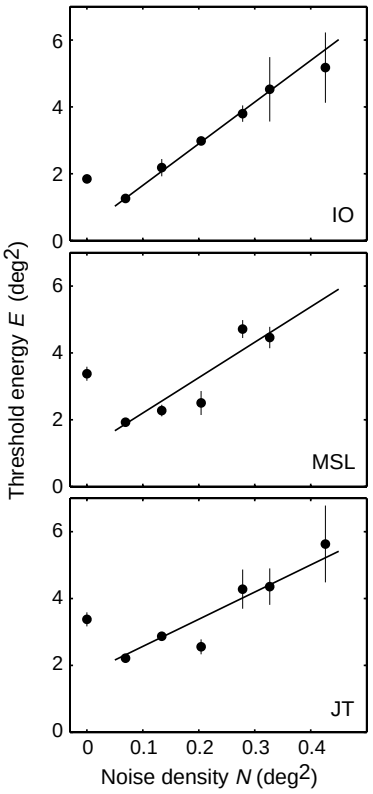

Fig. 14. Test of linearity. Threshold signal contrast energy is plotted as a function of the power spectral density of a low-pass noise (cutoff: 6.65 c/letter, 0.5 c/deg at the standard viewing distance). The solid line is a linear fit to the data excluding the no-noise point.

Vodyanoy, 1997) but has *not* been found for detection of luminance stimuli (Pelli & Farell, 1999).

The slope of the linear portion of these curves corresponds to signal-to-noise ratio (D^+). However, in Experiment 1, signal-to-noise ratio was computed as a ratio of threshold contrast energy at a fairly high noise level and E_0. Given the increased threshold for the no-noise case relative to low-noise stimuli, this should result in an underestimate of D^+. We argue that this effect is small and has no consequence for the results reported in Experiment 1. First, in all noise conditions of Experiment 1 we used noise densities that produced thresholds as high as those of the highest noise density we used in Expt. 3, so that the slope estimated using E_0 should differ little from the true slope (D^+). Also note that the noise densities used for the high-pass noises were much lower than those used for the low-pass noises. If the nonlinearity in Fig. 14 had a consequence for our results, we would expect all high-pass D^+ estimates to be systematically lower than those in the low-pass conditions. This was not the case. Finally, we can compare the slopes of the lines in Fig. 14 with the estimates of D^+ for corresponding conditions in Expt. 1. The estimates from Experiments 1 and 3 for observer IO were 10.96 and 12.46, for MSL: 9.81 and 10.59, and for JT: 11.35 and 8.13, indicating close agreement (within the standard errors indicated by the error bars in Fig. 8).

References

Bergen, J. R. (1991). Theories of visual texture perception. In D. Regan (Ed.). *Vision and visual dysfunction* (Vol. 10, pp. 114–134). New York: Macmillan.

Bezrukov, S. M., & Vodyanoy, I. (1997). Stochastic resonance in non-dynamical systems without response thresholds. *Nature, 385*, 319–321.

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 (London), 203*, 237–260.

Bracewell, R. N. (1978). *The Fourier Transform and its Applications*. New York: McGraw-Hill.

Chubb, C., & Landy, M. S. (1991). Orthogonal distribution analysis: A new approach to the study of texture perception. In M. S. Landy & J. A. Movshon (Eds.), *Computational Models of Visual Processing* (pp. 291–301). Cambridge, MA: MIT Press.

Chung, S. T. L., Legge, G. E., & Tjan, B. S. (2002). Spatial-frequency characteristics of letter identification in central and peripheral vision. *Vision Research, 42*, 2137–2152.

Gold, J., Bennett, P. J., & Sekuler, A. B. (1999). Identification of band-pass filtered letters and faces by human and ideal observers. *Vision Research, 39*, 3537–3560.

Graham, N. (1994). Non-linearities in texture segregation. In G. R. Bock & J. A. Goode (Eds.). *Proceedings of the CIBA Foundation symposium* (Vol. 184, pp. 309–329). New York: Wiley.

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

Jamar, J. H. T., & Koenderink, J. J. (1985). Contrast detection and detection of contrast modulation for noise gratings. *Vision Research, 25*, 511–521.

Kingdom, F. A. A., Keeble, D. R. T., & Moulden, B. (1995). Sensitivity to orientation modulation in micropattern-based texture. *Vision Research, 35*, 79–91.

Landy, M. S., Cohen, Y., & Sperling, G. (1984). Hips: A unix-based image processing system. *Computer Vision Graphics and Image Processing, 25*, 331–347.

Landy, M. S., & Graham, N. (2004). Visual perception of texture. In L. M. Chalupa & J. S. Werner (Eds.), *The visual neurosciences* (pp. 1106–1118). Cambridge, MA: MIT Press.

Landy, M. S., & Oruç, I. (2002). Properties of second-order spatial frequency channels. *Vision Research, 42*, 2311–2329.

Lutfi, R. A. (1983). Additivity of simultaneous masking. *Journal of the Acoustical Society of America, 73*, 262–267.

Nachmias, J., & Sansbury, R. V. (1974). Grating contrast: Discrimination may be better than detection. *Vision Research, 14*, 1039–1042.

Majaj, N. J., Pelli, D. G., Kurshan, P., & Palomares, M. (2002). The role of spatial frequency channels in letter identification. *Vision Research, 42*, 1165–1184.

Parish, D. H., & Sperling, G. (1991). Object spatial frequencies, retinal spatial frequencies, noise, and the efficiency of letter discrimination. *Vision Research, 31*, 1399–1415.

Patterson, R. D., & Nimmo-Smith, I. (1980). Off-frequency listening and auditory-filter asymmetry. *Journal of the Acoustical Society of America, 67*, 229–245.

Pelli, D.G. (1981). Effects of visual noise. PhD thesis. Cambridge University, Cambridge, England.

Pelli, D. G., & Farell, B. (1999). Why use noise. *Journal of Optical Society of America A, 16*, 647–653.

Perkins, M. E., & Landy, M. S. (1991). Nonadditivity of masking by narrow-band noises. *Vision Research, 31*, 1053–1065.

Robson, J. G. (1966). Spatial and temporal contrast-sensitivity functions of the visual system. *Journal of Optical Society of America, 56*, 1141–1142.

Sagi, D. (1990). Detection of an orientation singularity in gabor textures: Effect of signal density and spatial-frequency. *Vision Research, 30*, 1377–1388.

Schofield, A. J., & Georgeson, M. A. (1999). Sensitivity to modulations of luminance and contrast in visual white noise: separate mechanisms with similar behavior. *Vision Research, 39*, 2697–2716.

Shaw, G. B. (1979). Local and regional edge detectors: some comparisons. *Computer Graphics and Image Processing, 9*, 135–149.

Sloan, L. L. (1951). Measurement of visual acuity: A critical review. *Archives of Ophthalmology, 45*, 704–725.

Solomon, J. A. (2000). Channel selection with non-white-noise masks. *Journal of Optical Society of America A, 17*, 986–993.

Solomon, J. A., & Pelli, D. G. (1994). The visual filter mediating letter identification. *Nature (London), 369*, 395–397.

Stromeyer, C. F., & Julesz, B. (1972). Spatial-frequency masking in vision: Critical bands and spread of masking. *Journal of Optical Society of America, 62*, 1221–1232.

Sutter, A., Sperling, G., & Chubb, C. (1995). Measuring the spatial frequency selectivity of second-order texture mechanisms. *Vision Research, 35*, 915–924.

Watson, A. B., & Eckert, M. P. (1994). Motion-contrast sensitivity: visibility of motion gradients of various spatial frequencies. *Journal of Optical Society of America A, 11*, 496–505.

Wilson, H. R., McFarlane, D. K., & Phillips, G. C. (1983). Spatial frequency tuning of orientation selective units estimated by oblique masking. *Vision Research, 23*, 873–882.