PERGAMON

Vision Research 39 (1999) 4076–4084

Vision Research

www.elsevier.com/locate/visres

Contrast inconstancy across changes in polarity

Christopher P. Benton *, Alan Johnston

Department of Psychology, University College London, Gower Street, London WC1E 6BT, UK

Received 28 August 1998; received in revised form 12 March 1999

Abstract

In this study, we show that negative polarity noise patterns appear to have a higher contrast than positive polarity noise patterns with identical expected Fourier amplitude spectra. This demonstrates a failure of contrast constancy over changes in pattern polarity. An examination of local contrast measures shows that negative polarity noise has a wider distribution of local contrast values than positive polarity noise. We propose that the difference in apparent contrast between the two patterns may be based upon spatial non-linearities in the combination of local contrast measures. © 1999 Elsevier Science Ltd. All rights reserved.

Keywords: Texture; Contrast constancy; Contrast perception; Suprathreshold; Visual modelling

1. Introduction

The notion that the visual world is analysed through spatial frequency and orientation tuned channels has become widely accepted (Hubel & Wiesel, 1962; Campbell & Robson, 1968; Hubel & Wiesel, 1968; Pantle & Sekular, 1968; Movshon & Blakemore, 1973). The examination of suprathreshold contrast perception has sought to uncover the mechanisms through which information is conveyed within channels, and how that information is combined across channels.

A recent model for contrast perception across spatial frequency and orientation tuned channels has been provided by Georgeson and Shackleton (1994). This model is useful because it encapsulates a number of findings within the perceived contrast literature. The model utilises the quadratic summation across channels proposed by Quick, Hamerly and Reichert (1976), but responses within channels are non-linear. The particular form of the non-linearity is a power law applied to the thresholded input. The need for a threshold was indicated by Kulikowski (1976), who presented evidence that perceived contrasts of sine wave gratings of differing spatial frequencies were equal when the differences between the physical contrasts and detection thresholds were equal. Also, a number of studies have postulated the need for a within channel non-linearity (e.g. Gottesman, Rubin & Legge, 1981; Cannon & Fullenkamp, 1991). The model also incorporates the notion of gain control or normalisation within channels to account for contrast constancy (Georgeson & Sullivan, 1975).

This study examines the perception of contrast in positive and negative polarity patterns. These patterns are simply refections of one another through mean luminance. In the Fourier domain, a pattern may be viewed as a DC signal (mean luminance) plus a string of Fourier components, each with an associated amplitude and phase. Reflecting a pattern through mean luminance can be achieved by phase shifting each of the components by half a cycle. Change in polarity can therefore be viewed as a manipulation of phase.

If there is no interaction between Fourier components, then two patterns with identical Fourier amplitude spectra (but different phase spectra) should have the same perceived contrast. Because each pattern contains the same components, the response of any single channel should be the same for each of the two patterns. If the within channels responses are the same for two different patterns, then the subsequent combination of the outputs across channels should give the same measure of perceived contrast. An account of phase dependencies might well require the incorpora-

* Corresponding author. Tel.: +44-171-5045301; fax: +44-171-4364276.

E-mail address: c.benton@ucl.ac.uk (C.P. Benton)

0042-6989/99/$ - see front matter © 1999 Elsevier Science Ltd. All rights reserved.

PII: S0042-6989(99)00110-8

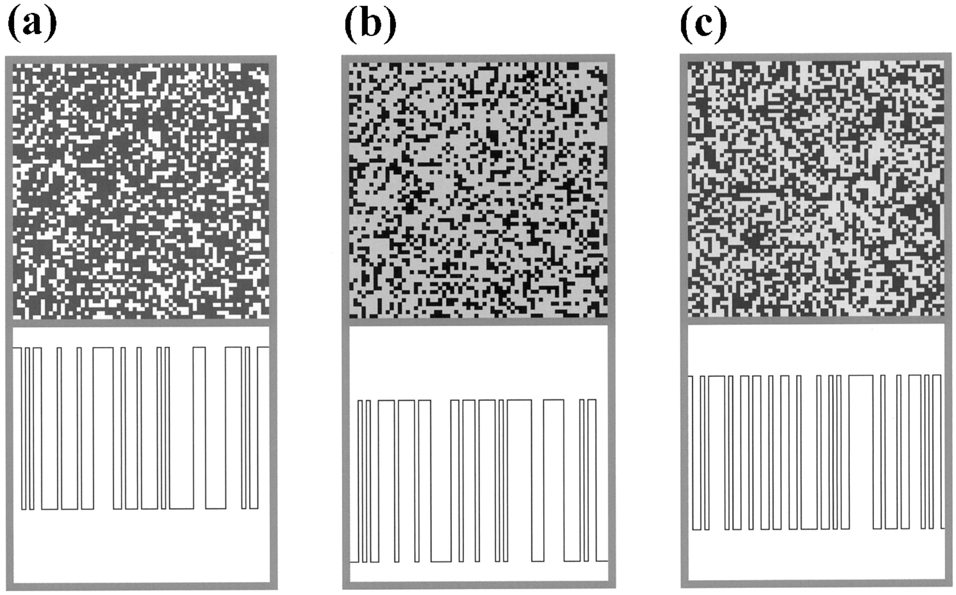

Fig. 1. Top panels show examples of (a) positive polarity noise; (b) negative polarity noise; and (c) standard binary noise. Bottom panels show spatial luminance profiles from the bottom-most line of each of the corresponding examples. The patterns in (a) and (b) are reflections of one another through mean luminance. The patterns were generated with equal mean luminance and RMS contrasts.

tion of phase dependent non-linearities into models of contrast perception.

Here we describe an effect of phase manipulation on perceived contrast. We use asymmetric noise patterns in which the greater part of the variance within the pattern is either carried by pixels that are brighter than mean luminance (positive polarity) or darker than mean luminance (negative polarity). The positive polarity pattern appears to be light random dots on a darker grey background, whilst the negative polarity noise appears to be dark random dots on a lighter grey background. In spite of their different appearance, these two patterns have identical expected Fourier amplitude spectra (if their RMS contrasts are equal). However, the negative polarity noise patterns appear to have the greater contrast. The difference in perceived contrast presents an example of a failure of contrast constancy over a change in pattern polarity.

2. General methods

In the experiments described in this paper, subjects compared the contrasts of two simultaneously presented patterns and indicated which appeared to have the higher contrast. Three types of noise pattern were employed; positive polarity binary noise, negative polarity binary noise and standard binary noise (examples are shown in Fig. 1). Each noise type is defined by its composition in terms of the direction and relative displacement from mean luminance of its pixel luminances, and by the probability of one of those pixels occurring. This is best illustrated by example. Standard binary noise contains two pixel luminances, one at a displacement $+d$ from mean luminance and one at a displacement $-d$ from mean luminance (where $d>0$). The probability of the $+d$ pixel occurring is 1/2 which of course means that the probability of the $-d$ pixel occurring is also 1/2. Standard binary noise can therefore be represented by the following list of proportional deviations and associated probabilities:

Standard binary noise: ([+1, 1/2], [−1, 1/2]).

Using this terminology the other noise types are defined as follows:

Positive polarity binary noise: ([+2, 1/3], [−1, 2/3]).

Negative polarity binary noise: ([+1, 2/3], [−2, 1/3]).

Note that the sum of the products of each pair equals zero. Given the mean luminance, the contrast and the noise type, the precise luminance levels can be calculated. In the present study contrast is always the root mean square contrast and is defined as follows:

$$C_{\mathrm{RMS}} = \frac{\sqrt{\sum_{n=1}^{h} P(I_n)\,(I_0 - I_n)^2}}{I_0}, \tag{1}$$

where there are h possible pixel values of intensity I_n each with a probability $P(I_n)$ of occurring and where I_0 is the mean luminance. This equation gives the *expected* RMS contrast of a noise pattern.

The spectral density of noise is described by the following equation:

$$N(f_x,f_y) = C^2_{RMS} b_x b_y \left(\frac{\sin(\pi f_x b_x)}{\pi f_x b_x}\right)^2 \left(\frac{\sin(\pi f_y b_y)}{\pi f_y b_y}\right)^2, \quad (2)$$

where $N(f_x,f_y)$ is the expected power at coordinates (f_x,f_y) within frequency space, and b_x and b_y are the spatial dimensions of the noise elements (Legge, Kersten & Burgess, 1987; Kukkonen, Rovamo & Näsänen, 1995). The shape of the expected Fourier amplitude spectrum is dependent upon the size of the noise elements, which was the same for all patterns. Noise patterns with the same RMS contrasts and noise element sizes have identical expected Fourier amplitude spectra.

Stimuli were generated and displayed on a Sun Sparcstation LX. The system is capable of displaying 256 grey levels. Mean luminance (I_0) was 37.4 cd/m^2. The displayable area of the screen was 26.0° wide by 20.4° high and was comprised of 1152 × 896 pixels. Linearisation was achieved by measuring the resultant luminance with a Graseby (UDT) model 265 luminance probe to a number of different grey levels. Measurements were taken from a window measuring 512 pixels wide and 256 pixels high presented in the centre of the screen. The remainder of the screen was set to minimum luminance (0.4 cd/m^2). A cubic equation of the form

$$I = b_0 + b_1 p + b_2 p^2 + b_3 p^3 \quad (3)$$

was fitted (where I is luminance, p is the grey level value and b_0, b_1, b_2 and b_3 are constants) and for all the values of p from 0 to 255 a list of the respective luminances was calculated. Discreet output levels were assigned as follows. Suppose a luminance of 20.4 cd/m^2 was required. The nearest grey levels on either side would be determined. A grey level of 135 gives a luminance value of 20.27 cd/m^2 whilst a grey level of 136 gives a luminance value of 20.57 cd/m^2. A random procedure is then used to choose between the two possible grey levels. This is done in such a way that, with an infinite number of random choices, the mean calculated luminance from all the grey levels chosen by the procedure would be equal to the required luminance.[1] The standard procedure for gamma correction involves the use of a look-up table (LUT), a procedure which reduces the number of available grey levels. Although the procedure employed in this study is more computationally expensive, all 256 grey levels are available for use within a stimulus.

We checked for adjacent pixel non-linearity (Naiman & Makous, 1992) by measuring the luminances of a number of instantiations of positive and negative polarity noise. This was carried out with the stimuli used in the experiments. There was no consistent difference between the mean luminances of the two noise types. It should be noted that this procedure does not guarantee that there is no effect of adjacent pixel non-linearity on pixel luminance variance. We address this issue in Experiment 3.

3. Experiment 1

In this experiment the perceived contrasts of asymmetric noise stimuli are measured using a contrast matching paradigm. The experiments took place in a darkened room where the stimulus was the only major source of illumination. The viewing distance was 75 cm, subjects had normal or corrected to normal vision. The task was to indicate which of two noise patterns had the higher contrast by pressing a mouse button. During the experiment subjects were allowed to move their eyes freely. The stimulus consisted of a rectangle with a square of target noise on one side and a square of comparison noise on the other. The two squares were 5.89° (256 pixels) on each side, and were separated from one another by a region of mean luminance with a contrast of zero and width of 0.23° (10 pixels). Noise element size was 2.76 arc minutes (two pixels). Each stimulus was presented for 2 s, minimum interstimulus interval was 1 s. During the experiment the screen around the display rectangle was set to minimum luminance. In the interstimulus interval the display rectangle reverted to mean luminance.

Each run contained 64 trials. For each run, target type and target contrast remained constant. The side on which the target pattern was presented was randomised for each trial. Comparison contrast was determined through an adaptive method of constants procedure (Watt & Andrews, 1981; Treutwein, 1995). Probit analysis was applied to the data from each run. This was done to estimate the point of subjective equality (PSE) at which the comparison noise had the same apparent contrast as the target noise (Finney, 1971). Target type was either positive polarity binary noise or negative polarity binary noise. Comparison type was always standard binary noise. Subjects were tested over three target contrast levels, 0.6, 0.5 and 0.4. For each level and each type, four PSEs were extracted, giving a total of 24 runs. The order of presentation of the runs was randomised. Fig. 2 shows the results for three subjects (the first author, CB, and two näive subjects) in terms of percentage contrast difference as a function of target contrast. If C_{PSE} is the estimated contrast of standard binary noise necessary to match the apparent contrast of a target with contrast C_T, then percentage contrast difference is

$$100\left(\frac{C_{PSE} - C_T}{C_T}\right). \quad (4)$$

[1] With this procedure, the contrast of a pattern produced on the display may differ from the required contrast. At the high contrasts utilised in these experiments, this effect is negligible.

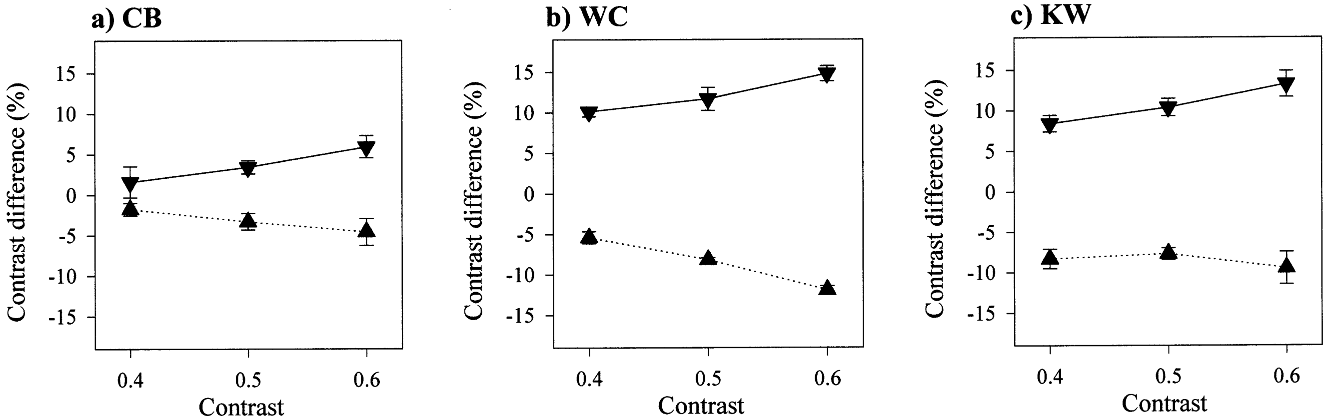

Fig. 2. Contrast difference (see Eq. (4)) as a function of target contrast for three subjects. Matches of standard binary noise to positive polarity noise are indicated by ▲ symbols, matches to negative polarity noise are indicated by ▼ symbols. Error bars indicate ±1 standard error.

These results show that negative polarity noise is judged to have a higher apparent contrast than positive polarity noise. For subjects WC and KW the difference between the asymmetric patterns is as high as 25%, for subject CB the magnitude of the difference appears to be about half of this. For all subjects the difference lessens with decreasing target contrast. A similar polarity bias has been described by Nam and Chubb (1998).

4. Experiment 2

In the previous experiment, results from a contrast matching task showed that negative polarity noise appeared to have a higher apparent contrast than positive polarity noise with the same RMS contrast. One possible objection to this finding is that judgement of contrast may be biased by cues other than contrast, such as the perceived brightness of the component pixels. To control for this factor we need to devise a task in which patterns of differing polarities can have contrast dependent effects without the necessity of comparing patterns of different luminance distributions. To this end, we utilised the contrast simultaneous contrast effect. In this, the perceived contrast of a noise pattern is affected by the contrast of the surrounding pattern. As the contrast of the surround increases, the perceived contrast of the inner patch decreases (Chubb, Sperling & Solomon, 1989).

A schematic diagram of the stimulus is shown in Fig. 3a. The square target and comparison patterns were 2.57° on each side. Both patterns were standard binary noise. The target pattern was surrounded by a narrow mean luminance band of width 11 arc min, and was positioned in the centre of a square (5.89° on each side) containing the surround pattern. The latter was either a positive polarity pattern or negative polarity pattern. The comparison pattern occupied the centre of an adjacent square (also of width 5.89°) the rest of which was filled with a mean luminance field. The contrast of the target pattern was identical to that of the surround pattern, and was set to either 0.6, 0.5 or 0.4. The task of the subject was to indicate whether the comparison pattern, or the target pattern, appeared to have the greater contrast. Each stimulus was presented for as long as it took the subject to provide a response. During the 1 s interstimulus interval, the window in which the stimulus was presented was set to mean luminance. Other than the differences in stimulus configuration, the procedure was identical to that described in the previous experiment. Results for three subjects (the two authors and a näive subject, DW) are shown in Fig. 3(b, c and d).

All subjects show a clear effect of noise type. The negative polarity noise has a greater contrast reducing influence on the perceived contrast of the centre patch than the positive polarity noise. This finding is fully consonant with the notion that negative polarity noise has a greater apparent contrast than positive polarity noise. Fig. 3e shows a repetition of the experiment with a single subject (CB) in which target, comparison and surround patterns were standard binary noise. The contrast of the target pattern was 0.5 and the contrast of the surround pattern was set to 0.4, 0.45, 0.5, 0.55 or 0.6. The results confirm that, as surround contrast is increased, the perceived contrast of the target pattern decreases.

The effects detailed in this experiment appear at odds with those of another study that examined polarity dependence in simultaneous contrast using regular textures. Solomon, Sperling and Chubb's (1993) positive polarity textures were evenly spaced bright pixels on a grey background. Each bright pixel was surrounded by eight grey pixels. The negative polarity texture was a reflection of the positive through mean luminance. Using all four combinations of surrounds and centres, and matching comparison patterns to centre patterns of the same type, no effect of polarity was obtained. In a

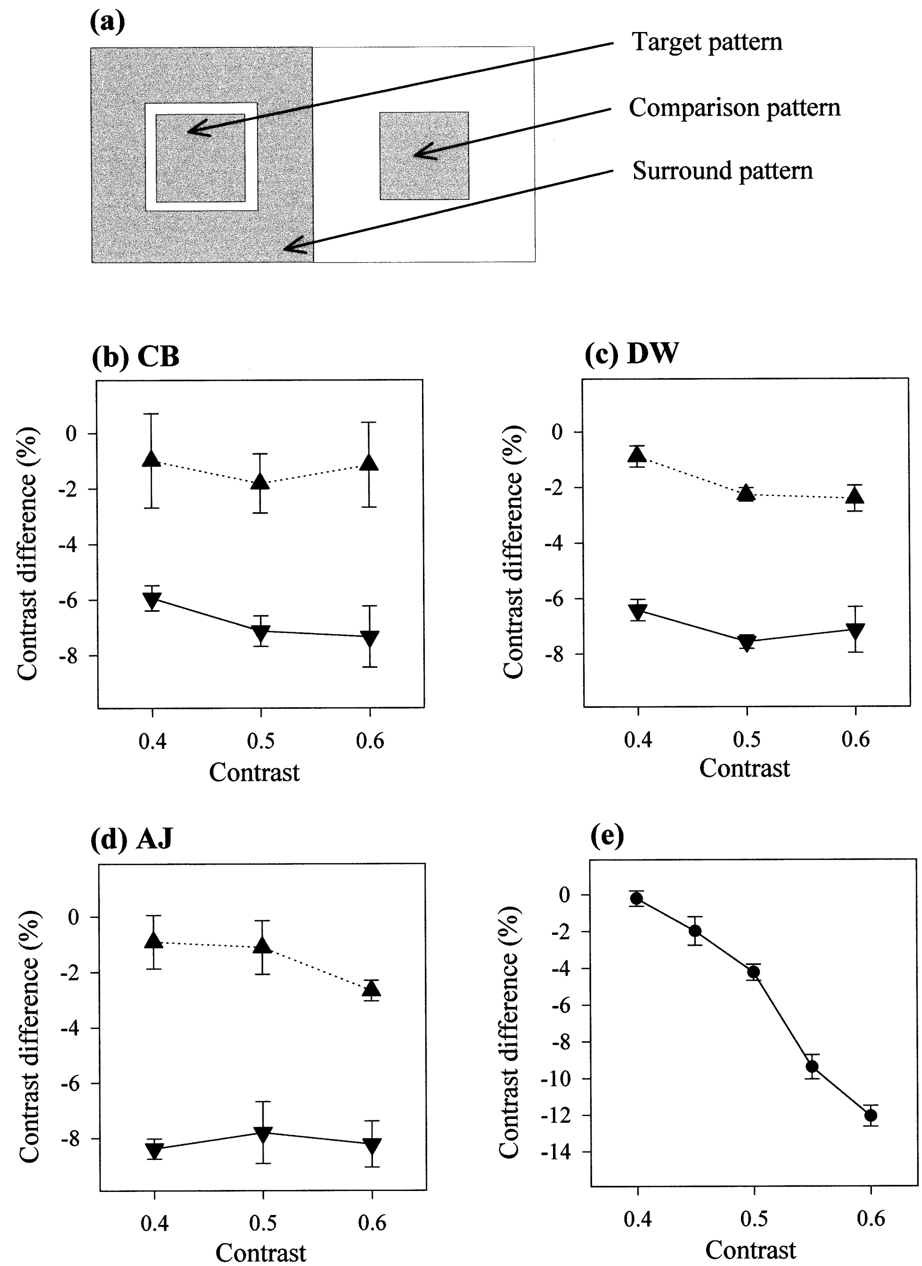

Fig. 3. (a) Schematic diagram of the stimulus employed in Experiment 2. (b, c and d) Data for three subjects showing the contrast reducing effect of positive polarity (▲) and negative polarity (▼) surround patterns on a standard binary noise target pattern. Results are plotted as a function of target and surround contrast. (e) Contrast reducing effect of a standard binary noise surround on a standard binary noise target. Results are plotted as function of surround contrast, target contrast was always 0.5. For all graphs, error bars indicate ± 1 standard error.

contrast matching task employing the same regular textures as Solomon et al., but using the methodology described in Experiment 1, we found only a small difference in apparent contrast. One subject (CB), matched positive polarity regular patterns to negative polarity regular patterns with a contrast of 0.3.[2] The mean contrast difference was -1.05% with a standard error of 0.86. It is unlikely that such a small difference in perceived contrast would translate into any noticeable difference in simultaneous contrast.

5. Experiment 3

In the previous two experiments we provided evidence for an effect of pattern polarity on perceived contrast. In these experiments we were careful to check that the mean luminances of the different polarity patterns did not differ to any significant degree. It

[2] If positive and negative polarity regular patterns are to have the same mean luminance, then the maximum possible C_{RMS} is 0.35.

should be noted however that this procedure does not necessarily control for any distortions of variance in noise element luminance. In this experiment we check for deviations in both luminance and luminance variance. This experiment is a repetition of experiment 1 but at only one contrast level (0.6). The experiment was carried out on a Silicon Graphics Octane workstation (eight bit resolution per colour plane). From the viewing distance of 75 cm, the screen had a width of 26.4° and a height of 21.0° (1280 × 1024 pixels). The squares in which the patterns were displayed measured 8.05° (384 pixels) on each side and were separated by a region of mean luminance (27.7 cd/m^2) with a width of 0.67° (32 pixels). Noise element size was 3.8 arc min (three pixels) on each side. The remainder of the screen was set to mean luminance. In this experiment, we used only the green gun of the monitor. The gamma correction procedure was identical to that described in the methods section. Results for three subjects are shown in Fig. 4, each point shows the mean and standard error of three PSEs.

To check for distortions of mean luminance and luminance variance, we took 200 readings (with a Graseby UDT model 265 luminance probe) for each of the asymmetric noise patterns. For each pattern, half of the readings were taken from the left side of stimulus display area, and half from the right side. The mean luminance for the positive polarity noise was 26.8 cd/m^2 with a standard deviation of 1.54 cd/m^2. The mean luminance for the negative polarity noise was 27.0 cd/m^2 with a standard deviation of 1.55 cd/m^2. For each test sequence the same randomising seed value was used. Therefore, within the set of positive polarity patterns, each pattern was generated with a different set of random values. However, each positive polarity pattern had a counterpart within the set of negative polarity patterns that was generated with exactly the same set of random values. The negative polarity counterpart would therefore be a reflection of the positive polarity pattern through mean luminance (as in (a) and (b) in Fig. 1).

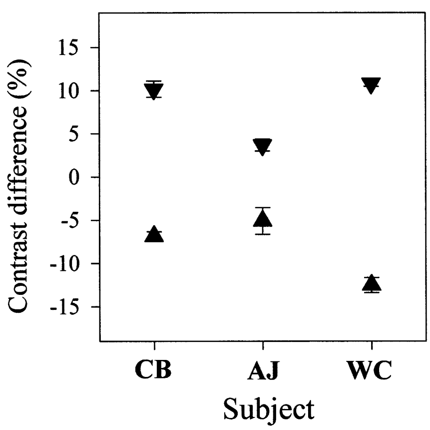

Fig. 4. Contrast differences for three subjects (AJ, CB and WC). Matches of standard binary noise to positive polarity noise are indicated by ▲ symbols, matches to negative polarity noise are indicated by ▼ symbols. Error bars indicate ±1 standard error.

Our measurements show only a small difference in mean luminance and standard deviation (and therefore variance) between the two patterns. Although there does appear to be some effect of adjacent pixel non-linearity, this effect is far too small to account for the magnitude of the difference in perceived contrast between the positive and negative polarity patterns.

6. Discussion

We have presented evidence to show that two noise patterns with identical RMS contrasts, Fourier amplitude spectra and mean luminances can appear to have radically different apparent contrasts. We examined the perception of contrast in noise patterns that were asymmetrically distributed around mean luminance. In our positive polarity patterns the greater part of the pattern's variance was carried by those pixels with luminances greater than mean luminance. In our negative polarity patterns, the greater part of the variance was carried by those pixels with luminance values below mean luminance. When the perceived contrast of a standard binary noise pattern was matched to that of the asymmetric patterns, negative polarity noise appeared to have the greater contrast then positive polarity noise. The positive and negative polarity noise patterns are reflections of one another through mean luminance. A change in polarity can be achieved by phase shifting all of a pattern's Fourier components by half a cycle. This study therefore describes an effect of phase manipulation on perceived contrast.

One obvious way to explain a distortion in contrast perception is to propose that some early non-linearity changes the patterns so that they appear to have different contrasts. It is useful to make the distinction between a *luminance* non-linearity and a *contrast* non-linearity. For the purposes of the present discussion, a luminance non-linearity is some function that is applied either to the luminance signal (or to some quantity that is monotonically related to luminance). A contrast non-linearity is one that is applied after mean luminance has been extracted from the signal and the signal has been split into positive and negative halfwave rectified components (for example after bandpass filtering). Whilst a luminance non-linearity can introduce distortions of both effective mean luminance and contrast, a contrast non-linearity cannot introduce distortions of effective mean luminance.

Turning first to contrast non-linearities. In order to obtain a difference between opposite polarity signals, different non-linearities must be applied to positive and negative filter outputs. The most effective patterns at producing differences in response to these non-linearities will be those with the greatest dissimilarity between the image structure above mean luminance and the structure below mean luminance. The positive polarity regular patterns utilised by Solomon et al. (1993) have a ratio of one high luminance pixel to eight lower luminance pixels. In the noise patterns utilised in this study, the ratio is one to two. Textures such as those used by Solomon et al should be far more effective at uncovering polarity dependent non-linearities than those utilised in the present study. The fact that they do not provides good evidence against a polarity dependent contrast non-linearity.

If a contrast non-linearity cannot explain the difference in perceived contrast between the two types of pattern, can a luminance non-linearity account for the effect? The output of photodetectors is held to be adequately modelled by a compressive non-linearity (Tomita, 1968), and non-linearities in the cone response have been shown to exist by projecting interference fringes onto human retinae (MacLeod, Williams & Makous, 1992; MacLeod & He, 1993). When two interference fringes of different spatial frequencies, each of which cannot be subjectively resolved, are projected onto the retina, a beat pattern is seen. So we do have good reason to believe that some luminance non-linearities do exist in the human visual system.[3]

If there is a strong influence of luminance non-linearities then we might well expect to find differences in the perceived contrasts of opposite polarity patterns. However, this observation must apply not only to the patterns described in this study but also to the regular patterns described by Solomon et al. (1993). In a simultaneous contrast task they found no evidence for a difference in perceived contrast between opposite polarity patterns. With their stimuli, we also found no evidence for a difference in perceived contrast. More generally, luminance non-linearities can potentially effect the perceived contrast of any stimulus in which the luminance profile is drastically altered by manipulations of phase. In gratings composed of f and $3f$ components, judgement of perceived contrast appears unaffected by whether the peaks of the components add or subtract (Quick et al., 1976; Arende & Lange, 1980). If these patterns were distorted by a luminance non-linearity then we might well expect a difference between the two conditions. Whilst we cannot discount the contribution of a luminance non-linearity, it is difficult to see why evidence for such an effect has not been obtained in the studies described above.

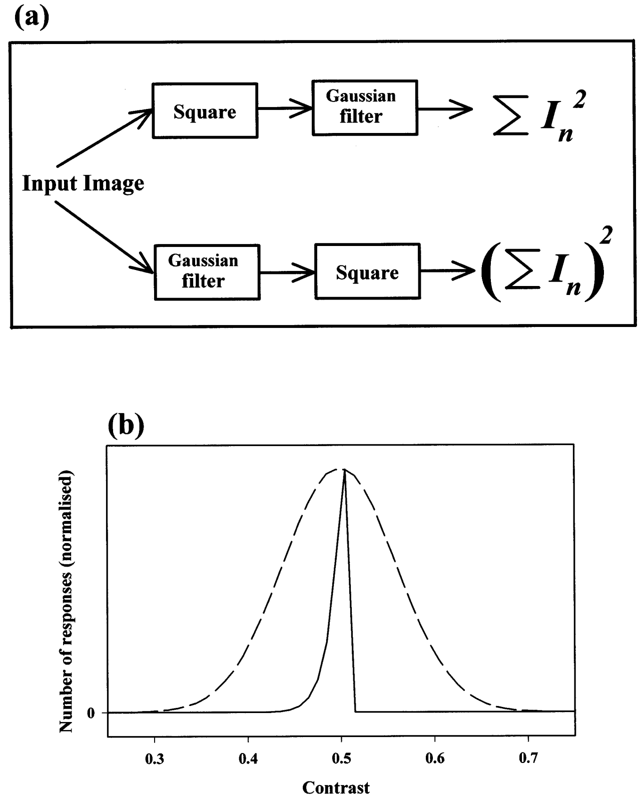

Fig. 5. (a) Schematic diagram showing procedure to extract Gaussian weighted square and square Gaussian weighted image values. These may be combined as described in Eq. (5) to give a measure of local RMS contrast. (b) Probability distribution of local contrast for noise patterns with expected RMS contrasts of 0.5. The unbroken line shows data for positive polarity noise, the broken line shows data for negative polarity noise. Data have been normalised to peak at the same value.

The fact that the contrast difference appears to be present with the opposite polarity patterns employed in this study, but not with the opposite polarity patterns used by Solomon et al., implies that it is not polarity per se that underlies the effect. One obvious difference between the two studies lies in the randomness of their stimuli. In the regular textures of Solomon et al., there is little variation over the image. Measures of local contrast are likely to be the same at whatever point in the image the measurement is taken.

However in the patterns employed in this study, the manner in which local contrast can vary over space is different for the two patterns. The RMS contrast of a given region (as opposed to expected contrast—Eq. (1)) may be calculated by the following equation,[4]

[3] We do not propose that luminance non-linearities are confined to photoreceptor outputs. The term may apply, for example, to retinal ganglion cells with high spontaneous firing rates.

[4] The equation gives the normalised standard deviation. It is derived from the standard computational formula for standard deviation (see Howell, 1992, p. 41).

$$C_{RMS} = \sqrt{\frac{h\sum_{n=1}^{h} I_n^2 - \left(\sum_{n=1}^{h} I_n\right)^2}{\left(\sum_{n=1}^{h} I_n\right)}}, \quad (5)$$

where there are h pixel values of intensity I_n. This allows us to develop a simple computational method to investigate variations in local contrast (Fig. 5a, also see Appendix). It should be emphasised that this is simply a tool to examine the way that local contrast may vary, it is not an attempt to directly model information processing within the human visual system. We calculate the Gaussian weighted sum of the squares of the image intensity, and the square of the Gaussian weighted sum of the image values. From these two values, a local Gaussian weighted measure of RMS contrast can be derived. Note that the volume of the Gaussian filter was set to one.

Results of this calculation for 128 instantiations of positive polarity noise and for the same number of instantiations of negative polarity noise were collected. Input images were 128×128 pixels, each noise element occupied one pixel. The standard deviation of the Gaussian was two pixels.[5] Input patterns had contrasts of 0.5 and notional mean luminances of 34.7. The range of contrast between 0 and 1 was divided into 100 bins, results from each noise pattern were collected into bins to derive a measure of the probability of occurrence of a range of local contrasts. These results are shown in Fig. 5b. Data have been normalised so that both curves have the same maximum value.

It is clear that the negative polarity noise has a broader distribution of possible local contrast values than the positive polarity noise. The way that this occurs can readily be shown by example. Let us suppose that we have a positive polarity noise pattern with a mean luminance of 50 cd/m^2. The pattern contains two luminance levels, bright pixels with a luminance of 90 cd/m^2 and dark pixels with a luminance of 30 cd/m^2. The variance of the pattern is 800 and the mean luminance is 50 cd/m^2, giving a contrast of 0.57. Because the stimulus is random, there may well be some local area in which there are equal numbers of bright and dark pixels. In this region the mean luminance is greater, it is 60 cd/m^2. The local variance is also greater, it is now 900, the local contrast is therefore 0.6. If we now consider a negative polarity pattern which has the same mean luminance and expected contrast as the positive polarity pattern; this contains two luminance levels, 10 and 70 cd/m^2. This pattern is simply a reflection of the positive polarity pattern through global mean luminance. In an area where there are equal numbers of light and dark pixels the negative polarity pattern will also have a local variance of 900 but will have a local mean luminance of 40 cd/m^2 which gives a local contrast of 0.75. This demonstrates how the difference in distribution of local contrast shown in Fig. 5b is a reflection of the manner in which local mean luminance and local variance covary within each of the patterns.

As an alternative to the notion of a luminance non-linearity, we propose the following: that the difference in apparent contrast between positive and negative polarity patterns is based on spatial non-linearities in the combination of local contrast measures into a pattern-wide global contrast measure. For example, in making judgements of contrast it is possible that subjects see high contrast regions as more salient in relation to their decision of contrast. Higher local contrasts would therefore have a disproportionately greater effect in determining global contrasts. This would readily account for the higher apparent contrast of the negative polarity noise.

Appendix A

We wish to calculate the RMS contrast over an area containing h pixels, each with luminance I_n. We also wish to apply a set of weights so that, for example, pixels at the center of the area have more influence in the calculation of local contrast than those at the edges. To this end, each pixel is given a weight W_n. For the sake of analytical simplicity, the weights are scaled so that

$$\sum_{n=1}^{h} W_n = 1. \quad (A1)$$

Weighted luminance (I_0) is then given by

$$I_0 = \sum_{n=1}^{h} W_n I_n, \quad (A2)$$

and weighted mean variance (V_0) is given by

$$V_0 = \sum_{n=1}^{h} W_n (I_n - I_0)^2. \quad (A3)$$

The weighted Rms contrast can then be calculated as

$$C_{\mathrm{RMS}} = \sqrt{\frac{V_0}{I_0^2}} = \sqrt{\frac{\sum_{n=1}^{h} W_n (I_n - I_0)^2}{\left(\sum_{n=1}^{h} W_n I_n\right)^2}}. \quad (A4)$$

Note that, to avoid the possibility of an imaginary contrast, all weights must be positive. With some simple reorganization, C_{RMS} can be written

$$C_{\mathrm{RMS}} = \sqrt{\frac{\sum_{n=1}^{h} W_n I_n^2 - \left(\sum_{n=1}^{h} W_n I_n\right)^2}{\left(\sum_{n=1}^{h} W_n I_n\right)^2}}. \quad (A5)$$

[5] The choice of a standard deviation of two was arbitrary.

Eq. (A5) presents us with a simple computational formula for calculating local weighted RMS contrast around each point in an image. In the present study, the weights are derived from a two dimensional circularly symmetric Gaussian function centered in the middle of a 20×20 pixel square.

References

Arende, L. E. Jr., & Lange, R. V. (1980). Narrow-band spatial mechanisms in apparent contrast matching. *Vision Research*, *20*, 143–147.

Campbell, F. W., & Robson, J. G. (1968). Application of Fourier analysis to the visibility of gratings. *Journal of Physiology*, *197*, 551–566.

Cannon, M. W., & Fullenkamp, S. C. (1991). A transducer model for contrast perception. *Vision Research*, *31*, 983–998.

Chubb, C., Sperling, G., & Solomon, J. A. (1989). Texture interactions determine perceived contrast. *Proceedings of the National Academy of Sciences USA*, *86*, 9631–9635.

Finney, D. J. (1971). *Probit analysis*. Cambridge: Cambridge University.

Georgeson, M. A., & Shackleton, T. M. (1994). Perceived contrast of gratings and plaids: non-linear summation across oriented filters. *Vision Research*, *34*, 1061–1075.

Georgeson, M. A., & Sullivan, G. D. (1975). Contrast constancy: deblurring in human vision by spatial frequency channels. *Journal of Physiology*, *252*, 627–656.

Gottesman, J., Rubin, G. S., & Legge, G. E. (1981). A power law for perceived contrast in human vision. *Vision Research*, *21*, 791–799.

Howell, D. C. (1992). *Statistical methods for psychology*. Belmont, CA: Duxbury.

Hubel, D. H., & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *Journal of Physiology*, *160*, 106–154.

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

Kukkonen, H., Rovamo, J., & Näsänen, R. (1995). Masking potency and whiteness of noise at various check sizes. *Investigative Ophthalmology & Visual Science*, *36*, 513–518.

Kulikowski, J. J. (1976). Effective contrast constancy and linearity of contrast sensation. *Vision Research*, *16*, 1419–1431.

Legge, G. E., Kersten, D., & Burgess, A. E. (1987). Contrast discrimination in noise. *Journal of the Optical Society of America A*, *4*, 391–404.

MacLeod, D. I. A., Williams, D. R., & Makous, W. (1992). A visual nonlinearity fed by single cones. *Vision Research*, *32*, 347–363.

MacLeod, D. I. A., & He, S. (1993). Visible flicker from invisible patterns. *Nature*, *361*, 256–258.

Movshon, J. A., & Blakemore, C. (1973). Orientation specificity and spatial selectivity in human vision. *Perception*, *2*, 53–60.

Naiman, A. C., & Makous, W. (1992). Spatial non–linearities of grayscale CRT pixels. *SPIE Proceedings*: *Human Vision*, *Visual Processing and Digital Displays III*, *1666*, 41–56.

Nam, J.-H., & Chubb, C. (1998). Perceived texture contrast is determined by a negative half-wave rectifying mechanism. *Investigative Ophthalmology & Visual Science* (*Supplement*), *39*, 649.

Pantle, A., & Sekular, R. W. (1968). Size-detecting mechanisms in human vision. *Science*, *162*, 1146–1148.

Quick, R. F. Jr., Hamerly, J. R., & Reichert, T. A. (1976). The absence of a measure of critical band at low suprathreshold contrasts. *Vision Research*, *16*, 351–355.

Solomon, J. A., Sperling, G., & Chubb, C. (1993). The lateral inhibition of perceived contrast is indifferent to on-centre/off-centre segregation, but specific to orientation. *Vision Research*, *33*, 2671–2683.

Tomita, T. (1968). Electrical responses of single photoreceptors. *Proceedings of the IEEE*, *56*, 1015–1023.

Treutwein, B. (1995). Adaptive psychophysical procedures. *Vision Research*, *35*, 2503–2522.

Watt, R. J., & Andrews, D. P. (1981). APE: adaptive probit estimation of psychometric functions. *Current Psychological Reviews*, *1*, 205–214.