The psychophysics of stereopsis can be explained without invoking independent ON and OFF channels

Jenny C. A. Read

Bruce G. Cumming

Institute of Neuroscience, Newcastle University, Newcastle upon Tyne, UK

Laboratory of Sensorimotor Research, National Eye Institute, National Institutes of Health, Bethesda, MD, USA

Early vision proceeds through distinct ON and OFF channels, which encode luminance increments and decrements respectively. It has been argued that these channels also contribute separately to stereoscopic vision. This is based on the fact that observers perform better on a noisy disparity discrimination task when the stimulus is a random-dot pattern consisting of equal numbers of black and white dots (a "mixed-polarity stimulus," argued to activate both ON and OFF stereo channels), than when it consists of all-white or all-black dots ("same-polarity," argued to activate only one). However, it is not clear how this theory can be reconciled with our current understanding of disparity encoding. Recently, a binocular convolutional neural network was able to replicate the mixed-polarity advantage shown by human observers, even though it was based on linear filters and contained no mechanisms which would respond separately to black or white dots. Here, we show that a subtle feature of the way the stimuli were constructed in all these experiments can explain the results. The interocular correlation between left and right images is actually lower for the samepolarity stimuli than for mixed-polarity stimuli with the same amount of disparity noise applied to the dots. Because our current theories suggest stereopsis is based on a correlation-like computation in primary visual cortex, this postulate can explain why performance was better for the mixed-polarity stimuli. We conclude that there is currently no evidence supporting separate ON and OFF channels in stereopsis.

Introduction

Harris and Parker (1995) made a striking claim about stereoscopic vision. They argued for the existence of neural mechanisms for bright and dark information that make independent contributions to stereopsis. Their evidence came from the performance of their observers on a depth discrimination task made challenging by the addition of disparity noise. The stimulus was a random-dot stereogram with a vertical step edge. The mean disparity was opposite on either side of the edge, but each dot was given a noise disparity drawn from a Gaussian distribution with a given standard deviation. The task was to detect which half of the stereogram was closer. Observers performed better when the stimulus was made up of equal numbers of black and white dots on a gray background (Figure 1A) than when all dots were black (Figure 1B) or white (Figure 1C).

One can imagine an ideal observer solving this task by computing the mean disparity of N dots on either side of the step-edge. Despite the noise, if observers averaged enough dots on each side of the boundary, they would correctly judge the sign of the step. Harris and Parker (1995) worked out what N would have to be for this ideal observer to match the performance of their observers. They found that the implied number of dots was around twice as large for mixed-polarity stimuli as for same-polarity stimuli, representing a doubling of statistical efficiency. Harris and Parker related this to the problem of stereo correspondence. Before the disparity of a dot can be identified, it has to be successfully matched up with the corresponding dot in the other eye. When all the dots are the same color, each dot in one eye could potentially be the correct match for any dot in the other eye. But if the stereo system only matches black dots with black dots and white dots with white ones, the number of false matches is halved for mixed-polarity stimuli. This could enable more dots to be successfully matched up, increasing the number N of disparities which can be averaged on each side of the step-edge and so improving performance. They suggested that the independent processing of black and white dots could be mediated by the separate ON and OFF channels that are well-established early in

Citation: Read, J. C. A., & Cumming, B. G. (2019). The psychophysics of stereopsis can be explained without invoking independent ON and OFF channels. *Journal of Vision*, *19*(6):7, 1–14, https://doi.org/10.1167/19.6.7.

Received October 31, 2018; published June 7, 2019

ISSN 1534-7362 Copyright 2019 The Authors

 \sim





Figure 1. Mixed and same-polarity stimuli. The stimuli are stereograms suitable for free fusion. (A) Mixed-polarity dot stimuli (black and white dots on a gray background); (B) and (C) same-polarity stimuli with either black-only or white-only dots. All three examples are for the no-overlap condition, where dots are not allowed to overlap/occlude one another, as in Harris and Parker (1995). The images are 100×100 pixels and the dots are 4×4 pixels with density d = 0.4, corresponding to an average of 250 dots per image). The images have zero mean disparity and Gaussian disparity noise with standard deviation equal to the dot size. The images have been normalized such that they all have the same mean and standard deviation of luminance. Psychophysically, the mixed-polarity advantage is unaffected by this normalization, showing that it cannot be explained by a contrast artefact.

the visual system (Jiang, Purushothaman, & Casagrande, 2015; Schiller, 1992, 2010).

Harris and Parker's (1995) mixed-polarity advantage was replicated by Read, Vaz, and Serrano-Pedraza (2011). These authors explored a different way of making the task hard: They made a certain proportion of the dots uncorrelated. That is, some dots had the disparity associated with the step-edge, while other dots were removed and then replaced at random, independently in each eye. Read et al. argued that this presents even more of a challenge to stereo correspondence. In the disparity-noise version of the task, each dot does have a definable disparity, even if this difference may be hard for the visual system to extract. But in the decorrelated version of Read et al., the uncorrelated dots carried no useful disparity signal, although local false matches may contribute noise. Read et al. found that the mixed-polarity advantage was even more pronounced in the decorrelation version of the task, with an implied efficiency ratio of up to five. Although this is not as neat as the original doubling of efficiency, mapped on to distinct ON/OFF channels, it is still consistent with Harris and Parker's argument that stereo correspondence is easier in mixed-polarity stimuli.

Read & Cumming

The problem is that it is not at all clear how to reconcile this with the understanding of disparity encoding that has developed since Harris and Parker's 1995 paper was published. Physiologically, stereo correspondence is believed to begin in primary visual cortex (V1). Many V1 neurons are tuned to disparity in random-dot patterns like those shown in Figure 1. These responses can readily be explained by the binocular energy model (Ohzawa, DeAngelis, & Freeman, 1990) which has since become the canonical description of the early stages of disparity encoding. It has been extensively tested against real V1 neurons and although modifications are certainly required, the basic principle underlying the energy model has been vindicated (Cumming & DeAngelis, 2001; Henriksen, Tanabe, & Cumming, 2016; Ohzawa, 1998; Prince, Cumming, & Parker, 2002; Prince, Pointon, Cumming, & Parker, 2002; Read, 2005); notably, the model has made successful predictions about the response of real neurons to impossible stimuli (Cumming & Parker, 1997). Stereopsis is widely used as a model system to study how perceptual experience relates to early cortical encoding (Parker, 2007; Read, 2014; Roe, Parker, Born, & DeAngelis, 2007).

However, the energy model does not recognise image features such as dots. Rather, it effectively computes the interocular cross-correlation between left and right images after filtering with filters that are bandpass for orientation and spatial frequency (Allenmark & Read, 2011; Oian & Zhu, 1997). For the energy model, the dark spaces represented by the background in an allwhite-dot stimulus should serve just as well as black dots. Recently, Ichiro Fujita and colleagues have presented psychophysical evidence which they argue implies a separate "matching" computation, which is sensitive to contrast polarity and will not match black dots with white (Doi & Fujita, 2014; Doi, Takano, & Fujita, 2013; Doi, Tanabe, & Fujita, 2011). In principle, such a matching computation is consistent with Harris and Parker's (1995) theory, but its neural substrate is unknown. We have recently shown that the psychophysical data can be equally well explained by a slight modification (an additional output nonlinearity) to the energy model (Henriksen, Cumming, & Read, 2016). This model makes specific predictions regarding the effect of dot size and dot density on performance, which were tested and verified, and is also consistent with the properties of V1 neurons (Henriksen, Read, & Cumming, 2016). Thus, Doi and Fujita's observations do not require a dot-matching process.

Therefore, it is not currently clear how one could build a model of V1 neurons that would provide a neuronal basis for Harris and Parker's (1995) independent mechanisms. One might consider the modified version of the energy model proposed by Read, Parker, and Cumming (2002), where monocular inputs are halfwave rectified before binocular combination. This enables a binocular neuron to receive input only from an ON or an OFF channel. However, Read et al. (2011) showed that in this model, as for the original energy model, disparity tuning curves have the same amplitude for mixed- as for same-polarity stimuli. The reason, again, is that the energy model and variants sense interocular correlation, but do not look for specific image features. Similarly, there is not a simple mapping between white/black dots and ON/OFF channels. "OFF" detectors based on linear filters are activated by light decrements in the receptive field centre. But they are also activated by light increments in the surround. Consequently, a single-polarity pattern containing only bright dots does not stimulate only ON channels, since the OFF channels are stimulated by the dark regions between the bright dots. Similarly, in mixed polarity patterns the OFF channel is not blind to the white dots. So, even if ON and OFF channels are treated separately, as in the stereo correspondence algorithm of Glennerster (1998), they would not provide independent estimates of disparity. This is what makes the observation by Harris and Parker so striking, and so hard to reconcile with current views about disparity processing.

Very recently, Goncalves and Welchman (2017) put forward the first simulation that reproduces the mixed-polarity advantage observed psychophysically. Their model was a convolutional neural network, which they called a Binocular Neural Network. They explained its mixed-polarity advantage by noting that "the network depends on the activity of the simple units moderated by readout weights. Presenting mixed versus single-polarity stimuli increases the simple unit activity, in turn changing the excitatory and suppressive drives to complex units. We found that mixed stimuli produce greater excitation for the preferred output unit and increased suppression to the nonpreferred unit". This is puzzling, because Goncalves and Welchman's Binocular Neural Network is simply a generalized version of the original energy model with more subunits and a half-linear instead of halfsquaring output nonlinearity (discussed in Read & Cumming, 2017). These additional ingredients could not enable the model to work as envisaged by Harris and Parker (1996), matching dot to dot in a way sensitive to contrast sign rather than simply responding to the interocular correlation between left and

right images. Despite this, the model in Goncalves and Welchman clearly did reproduce the effect reported by Harris and Parker. Even more striking, their model reproduced a second subtle feature of human psychophysics. In constructing their stimuli, Harris and Parker did not permit dots to overlap one another. Read et al. (2011) reproduced the results of Harris and Parker, but also showed that the mixed-polarity advantage is no longer present when dots were allowed to overlap. Goncalves and Welchman reproduced this behaviour also. Here, we try to understand what explains it.

It seems remarkable that a relatively simple model based on initial linear filtering can reproduce these psychophysical phenomena, which seem to depend on specific image features (dots), and subtle rules about feature placement (no overlap). We will show that avoiding dot-overlap has a subtle effect on the binocular cross-correlation in the images, and that this is different for same-polarity and mixed-polarity stimuli. As a result, models much simpler than that of Goncalves and Welchman can also explain these phenomena. In these simple models, it is clear that there are not separate ON and OFF channels. As a result, the existing evidence does not support the conclusion that human stereopsis uses separate ON and OFF channels. Of course, our analysis does not prove that humans do not use separate channels eitherfurther experimental work will be required to determine that.

Methods

Pearson correlation coefficient

We measure the correlation between left and righteye images with the standard formula for the sample Pearson correlation coefficient, r. L_j , R_j is the value of the *j*th pixel in the left, right eye. Then the sample Pearson correlation coefficient is

$$r = \frac{\langle LR \rangle - \langle L \rangle \langle R \rangle}{\sqrt{\langle L^2 \rangle - \langle L \rangle^2} \sqrt{\langle R^2 \rangle - \langle R \rangle^2}} \quad (1)$$

where $\langle \rangle$ indicates the average over all pixels *j*. Equation 1 describes the correlation at zero disparity, but a similar expression holds for any uniform disparity if $\langle LR \rangle$ is computed between appropriately displaced pixels. In Figure 2 and Figure 3, we plot this correlation coefficient for zero-disparity images. In Figure 6, we plot it as a function of image displacement for disparate images.



Figure 2. Disparity noise is more disruptive for same-polarity patterns. Interocular image correlation plotted as a function of disparity noise (the standard deviation of the Gaussian noise distribution as a fraction of the dot size in the pattern). Lines show the mean Pearson correlation coefficient between left and right images, averaged over 100 different random-dot patterns; shaded regions show \pm *SD*. Note that the vertical axis is logarithmic and all images have zero mean disparity. Insets show example mixed-polarity images. Same-polarity images are the same except all dots are black or all white. In the simulations, the mixed-polarity images were generated first, and then the dot colors were manipulated to produce same-polarity images. Thus, the same dot patterns were used for both conditions. (A) and (C) No-overlap condition (dots are placed only in empty regions of the stimulus); (B) and (D) Overlap condition (dots are placed at random, occluding one another when they overlap). (A) and (B) Low-density (d=0.1); (C) and (D) High-density (d=0.4). The same number of dots was drawn in each case, so the area unoccupied by dots (the background) is a little higher for the no-overlap condition. Insets show example regions of mixed-polarity images. Images were 200 × 200 pixels; dots were 4 × 4 pixels.

Image generation

In each run of a simulation, a mixed-polarity imagepair was generated first, with a background of 0 and white and black dots of ± 1 . The absolute value was then taken to produce a same-polarity stereogram with white dots on a gray background, and then the sign was inverted to produce a same-polarity stereogram with black dots on a gray background (cf. Figure 1). Dots were square, and the term "dot size" in this paper refers to the side of the square. To generate the pattern, first the number of dots was specified, and then dots were added one after another. First, the luminance of the dot was chosen, either black or white with equal probability. The x and y coordinates of each dot were chosen from a uniform random distribution across the image.



Figure 3. For patterns with no dot overlap, correlation is lower in same-polarity dot patterns even where the dots have two different contrasts. As Figure 2, but the cyan/dotted curves show results for same-polarity patterns with equal numbers of black and dark gray dots, for comparison with same-polarity patterns with all-black or all-white dots (black/dashed curves). There is no significant difference between the three same-polarity patterns. (A) Low-density (d = 0.1); (B) High-density (d = 0.4); no dot overlap throughout. Other details as in Figure 2.



Figure 4. Patterns can be divided into (A) and (C) "matched" pixel-pairs, where the pixel-values in the two eyes are necessarily the same, and (B) and (D) "unmatched pairs" disrupted by noise. Graphs show the different possible pixel values in the two eyes (L) and (R): (A) and (B), for mixed-polarity images, and (C) and (D), for same-polarity. The size of each blob (not to scale) represents the probability of that particular combination. These probabilities are also labeled on the diagram; see section on formal derivation for definitions. All blobs of a given color have the same probability in each panel; thus only one is labeled.



Figure 5. Probability tree showing the possible situations. Each location in the image can be a matched pixel-pair, with probability m, or an unmatched pair (i.e., affected by noise), with probability u = 1 - m. Given that the location is matched, the probability that the left-eye pixel is covered by a dot is d_m (by definition), in which case the right-eye must also contain a dot (probability 1). The probability that a matched location has background in the left eye is $(1 - d_m)$, in which case the right eye must also be background. If the location is unmatched, the probability that the left-eye pixel is covered by a dot is d_u (by definition). The right-eye pixel may contain either dot or background. Conversely the probability that the left-eye pixel is background is $(1 - d_u)$. In that case, the right eye must contain a dot, since under our definition unmatched pixel-pairs must contain a dot in at least one eye.



Figure 6. Cross-correlation functions between left and right images, for different types of disparity noise. Pink solid curves: results for mixed-polarity dot patterns like those shown on the left; black dashed curves: for same-polarity patterns, where the dots are either all black or all white. Top row: for random-dot patterns where dots are scattered at random; bottom row: where dots are not allowed to overlap. (B) and (F) "No noise": all dots have the same disparity, 2 pixels. This thus has the same shape as the auto-correlation function of the monocular images, shifted to the stimulus disparity. (C) and (G) "Transparent planes": half the dots (at random) have disparity 0 pixels and half have disparity 4 pixels. (D) and (H) "Gaussian noise": every dot has random disparity noise drawn from a Gaussian of standard deviation 2 pixels. The images are 100×100 pixels and the dot size is 4 pixels. The curves show the mean cross-correlation function averaged over 100 different random images; shaded regions show $\pm SD$.

The appropriate disparity was then applied to the x coordinate. In the Overlap condition, the dot was then simply drawn at the resulting location, overwriting any pixels belonging to existing dots. In the No-overlap condition, we checked to see if this dot would overwrite any pixels belonging to existing dots. If it did, the dot was abandoned and a new one was chosen. This process was repeated until the desired number of dots had been placed.

Read & Cumming

When comparing Overlap and No Overlap conditions in Figure 2 through Figure 5, we set the number of dots such that each pixel in a monocular image had a specified probability d of being covered by a dot. If none of the dots overlap, then the number of dots required is

$$N_{NoOv} = d \frac{A_{im}}{A_{dot}}$$
 when dot overlap is forbidden

where A_{im}/A_{dot} is the ratio of the area of the whole image to the area of a single dot. If dots are allowed to overlap, obviously more dots are required in order to achieve the same probability. It can be shown that

$$N_{Ov} = \frac{\ln(1-d)}{\ln\left(1 - \frac{A_{dot}}{A_{im}}\right)} \qquad \text{when dot overlap is allowed.}$$

The monocular images were then normalised to have zero mean luminance and unit variance. This is a simple way of representing low-level luminance and contrast adaptation in the visual system. As Read et al. (2011) showed, if the mean luminance is not zero, energy model neurons have much lower amplitude disparity tuning to same-polarity images, even in the absence of noise. (This artefact cannot explain the psychophysical results, both because early adaptation removes such changes in overall luminance before binocular combination, and because empirically the mixed-polarity advantage persists even if the psychophysical stimuli are adjusted to have the same luminance.) The same normalisation is implied in the definition of the Pearson correlation coefficient r (Equation 1), which is unchanged when a constant luminance offset is added to both images or when both images are scaled. The normalisation of variance ensures that all images have the same contrast energy. Without this step, samepolarity stimuli would have lower contrast than mixedpolarity. We wished to study the response of model neurons to correlation differences in the stimuli, unconfounded by changes in contrast.

For clarity, image parameters are given in each figure legend. In most figures, the images were 100×100 pixels, the mean disparity was 0 pixels, and the dot size was 4 pixels. For the neuronal simulations shown in Figure 7, we aimed to reproduce the stimulus of Figure 7 of Read et al. (2011). In that paper, the dots were circles with an area of 8^2 arcmin and, scattered without

overlap, occupied 0.28 of the stimulus area; the step size was \sim 3 arcmin and the noise \sim 2 arcmin (precise values differed between observers, depending on what was needed to bring their performance to around 75% correct on average). In our simulation, the images were 241 × 241 pixels, and the dot density was the same as in the experiments of Read et al. (2011), i.e., d = 0.28 in the No Overlap condition and correspondingly lower in the Overlap condition. We took 1 pixel to represent 0.5 arcmin and made the dot size 6 pixels, stimulus disparity 6 pixels, and disparity noise 4 pixels.

Simulating response of binocular neurons

Our model receptive fields were even or odd Gabor functions:

$$f_e(x, y) = \exp\left(-\frac{x^2 + y^2}{2\sigma^2}\right) \cos(2\pi x/\lambda)$$
$$f_o(x, y) = \exp\left(-\frac{x^2 + y^2}{2\sigma^2}\right) \sin(2\pi x/\lambda)$$

The receptive field standard deviation was $\sigma = 32$ pixels, representing 16 arcmin or 0.27° since we are representing 1 arcmin with 2 pixels, and the carrier spatial period λ was 128 pixels or 1.1°. For a model neuron with position disparity x_0 , we computed the inner product of each of these receptive fields, appropriately shifted, with the monocular images:

$$v_{Le}(x_0) = \sum_{j=1}^n f_e(x_j - x_0/2, y_j) L_j$$
$$v_{Re}(x_0) = \sum_{j=1}^n f_e(x_j + x_0/2, y_j) R_j$$

and similarly for v_{Lo} . We considered various possible model V1 neurons, as follows:

ODF TE: $R = (v_{Le} + v_{Re})^2$ ODF TI: $R = (v_{Le} - v_{Re})^2$ ODF ODD: $R = (v_{Le} + v_{Ro})^2$ RPC TE: $R = (\lfloor v_{Le} \rfloor + \lfloor v_{Re} \rfloor)^2$ RPC ODD: $R = \lfloor \lfloor v_{Lo} \rfloor - \lfloor v_{Re} \rfloor \rfloor^2$

where the symbol $\lfloor \rfloor$ denotes half-wave rectification, i.e., $\lfloor x \rfloor = x$ if x > 0 and is 0 otherwise. Within each neuron class we simulated a population of neurons with different position disparities x_0 (from -20 to +20 pixels in steps of 1 pixel). The tuning curves shown in Figure 7 represent each neuron's mean response to 10,000 different random-dot patterns with the same disparity, normalised by that neuron's mean response to binocularly uncorrelated stimuli.

8



Figure 7. Population response for five different types of model V1 simple cells, for mixed-polarity (solid pink curve) and same-polarity (dashed, black) random-dot stereograms. The stimulus mean disparity is 6 pixels. In the top row, the stimuli have no disparity noise; i.e., all dots have a disparity of 6 pixels. In the bottom two rows, each dot additionally has a noise disparity drawn from a Gaussian with a standard deviation of 12 pixels. In the top two rows, dots did not overlap; in the bottom row, they did, as shown by the inset images of mixed-polarity stimuli. The horizontal axis shows the position disparity of the neurons. Curves show the average response of a cell with the position disparity indicated on the horizontal axis to 10,000 random-dot stereograms with the specified parameters. The vertical lines show the disparity of the stimulus (solid blue line) and its opposite (dashed). The response of each neuron is normalised by its response to uncorrelated stereograms, marked with the horizontal gray line. The images were 241×241 pixels, dot size was 6 pixels, and images were normalized to have zero mean luminance and unit variance.

"ODF" refers to the original form of the binocular energy model introduced by Ohzawa, DeAngelis, and Freeman (1990). "RPC" refers to the modified version introduced by Read, Parker, and Cumming (2002), in which the monocular inputs are half wave-rectified prior to binocular combination. "TE" denotes tunedexcitatory, i.e., a cell whose monocular receptive fields have the same phase and whose disparity tuning curve is therefore symmetric about a central peak (Read & Cumming, 2004). "TI" denotes tuned-inhibitory, i.e., a cell whose monocular receptive fields have opposite phase and whose disparity tuning curve is therefore symmetric about a central trough. "ODD" denotes a cell whose monocular receptive fields are $\pi/2$ out of phase. In the ODF model, such a cell has oddsymmetric tuning.

For simplicity, the example tuning curves in Figure 7 are all for simple cells. The same average tuning curves are obtained if we use phase-invariant complex cells, as for example $R = (v_{Le} + v_{Re})^2 + (v_{Lo} + v_{Ro})^2$, but complex cells show less variability in their response across different random-dot patterns.

Results

Where dots do not overlap, same-polarity stereograms have lower correlation than mixed-polarity. Figure 2 shows how the interocular image correlation declines as noise is added to the stimulus. The plots show the mean Pearson correlation coefficient between left and right images of random-dot stereograms, averaged over different random-dot patterns. The pink/solid lines show results for mixed-polarity stimuli, with equal numbers of black and white dots; the gray/dashed lines show results for same-polarity stimuli, where the dots are all black or all white. In the left-hand column, the patterns were generated such that dots are not allowed to overlap one another, as was done in Harris and Parker (1995); in the right-hand column, dots are scattered at random, with dots occluding others where they overlap. The top row is for a low dot density, and the bottom row for a high dot density.

Unsurprisingly, image correlation falls as disparity noise increases. Where dots are scattered at random (Overlap condition, Figure 2B and D), the decline is equal for mixed-polarity and same-polarity stimuli: The same amount of disparity noise results in the same decrease in correlation for both. However, where dots are placed so as to avoid overlap, the decline is steeper for same-polarity stimuli. This means that for a given amount of disparity noise, image correlation is lower for all-black or all-white dots. This effect becomes stronger as dot density increases (compare Figure 2C vs. Figure 2A). Of course, where dot density is low enough, overlap will be rare anyway and so results must tend to be the same whether or not overlap is forbidden. The same effect applies for two shades of the same polarity.

Harris and Parker (1995) also examined stimuli containing two dot colors but both with the same contrast polarity relative to the background, e.g., black and dark-gray dots on a gray background. They found that performance in this case was no better than for patterns containing dots of only one contrast. This result is also predicted by the image correlation. Figure 3 shows how image correlation is disrupted by disparity noise for the mixed- and same-polarity patterns considered before, and for "dark and darker" same-polarity patterns. Clearly the addition of "dark and darker" dots makes little difference; disparity noise is just as disruptive as for the other same-polarity patterns, and mixed-polarity patterns still have an advantage.

Why forbidding dot overlap produces lower correlation for same-polarity patterns

Why does this happen? In brief, because preventing dot overlap changes the pairwise statistics of pixels in

the stereo images. Stereo images are made up of pairs of pixels, one in the left eye and one in the right. We imagine starting with a blank stimulus in both eyes. All the pixel-pairs contain background in both eyes, so they are "matched" in this sense. Now we add dots. Correlated, zero-disparity dots by definition change both left- and right-eye pixels to the same value, so the pixel-pair remains matched. However, disparity noise and/or uncorrelated dots produce "unmatched" pixelpairs. By definition, these either contain a dot in one eye but background in the other, or where both eyes' pixels are dots, this sameness is just by chance; it is not the left and right images of "the same" dot.

Figure 4 shows the pairwise distribution of dot contrasts for matched (A,C) and unmatched (B,D) pixel-pairs. Note that under the description developed in the previous paragraph, all the unmatched pixelpairs have a dot in at least one eye, whereas only matched pixel-pairs can have background in both eyes. Figure 4A and B is for mixed-polarity images, where there are three possible values of luminance in each eye and thus nine possible pairs (L, R). For matched pixelpairs (A), only three of the nine are possible: Pixels may be background in both eyes (color-coded gray in the figure), and pairs that have dots of the same polarity in both eyes (green). For unmatched pixel-pairs (B), we find pixel-pairs that have a dot in one eye but not the other (pink) and pairs that have dots of opposite polarity (blue) as well as pairs that have dots of the same polarity by chance (green). By inspection, the matched pixels have correlation 1, and the unmatched pixels have correlation 0.

Figure 4B and D shows what happens if we convert the pattern into a same-polarity stereogram by turning all the black pixels white. Now, pixel values of -1become +1, effectively folding the other three quadrants of the space onto the first quadrant. Again by inspection, the matched pixel-pairs still have correlation 1, even though there are now only two points in the space. However, now the unmatched pixel-pairs have a negative correlation. This immediately provides an intuitive reason why correlation is lower for samepolarity images: The overall correlation is dragged down by the negative correlation of the unmatched pairs.

It is clear from Figure 4D that this negative correlation is driven by the "monocular dots": That is, unmatched pixel-pairs which have a dot in one eye and background in the other. These monocular dots occur less often when overlap is allowed. This is because when overlap is allowed, a noise dot in the left eye can land on a pixel that was previously a matched pair with a dot in both eyes. As a result, this pixel-pair becomes unmatched, with unrelated dots in both eyes (the noise dot in the left eye and the previously matched dot in the right eye). This situation shifts area from the pink blobs to the green blob in Figure 4D, making the correlation less negative. As a result, as we show below, when dot overlap is allowed the overall interocular Pearson correlation ends up being the same for same- as for mixed-polarity images. When overlap is forbidden, the situation is different. Now a noise dot in the left eye can only land on background. A previous noise dot may have landed at that location in the right eye, resulting in unmatched pixel-pairs with dots in both eyes, but previous signal dots cannot contribute to this situation. This means that less area goes to the green and more to the anticorrelating pink blobs. This results in lower correlation for same-polarity stimuli when overlap is forbidden.

Formal derivation

We now provide formal proofs of the above statements.

Basic definitions

Above, we defined d to be the probability that any given pixel in a monocular image is covered by a dot, as opposed to being part of the background. In general, this will be different for matching versus nonmatching pixel pairs. We thus define d_m to be the probability that a given pixel in the left eye is covered by a dot, given that this pixel belongs to a matched pair (n.b. a pixel covered by the background in both eyes is a matched pair). Similarly, we define d_u to be the probability that a given pixel in the left eye is covered by a dot, given that this pixel belongs to an unmatched pair. We define m to be the probability that any given pixel is matched, and u = 1 - m the converse probability that it is unmatched. By the law of total probability,

 $d = md_m + ud_u$

Now let us think through the probability tree for any pixel-pair (Figure 5). With probability m, the pixel-pair is matched; and then with probability d_m , it contains a dot in the left eye. Then—since it is matched—it also contains a dot in the right eye. Thus on average a fraction md_m pixel-pairs are matched pairs with dots in both eyes. Similarly a fraction $m(1 - d_m)$ are matched pairs with background in both eyes.

Probability tree

For unmatched pairs, by definition the probability that there is a dot in the left eye is d_u . Let x be the probability that, given this, there is also a dot in the right eye. Then $d_u x$ is the probability that an unmatched pair has a dot in both eyes, and $d_u(1 - x)$ is the probability that an unmatched pair has a dot in the left eye and background in the right. By symmetry, this must also be the probability that an unmatched pair has a dot in the right eye and background in the left. These are the only three possibilities, so their probabilities must sum to 1: $d_ux + 2d_u(1 - x) = 1$ and thus $xd_u = 2d_u - 1$. So, a fraction $u(2d_u - 1)$ pixel-pairs are unmatched pairs with dots in both eyes; $2u (1 - d_u)$ are unmatched pairs with a dot in one eye and background in the other. The precise values of d_m, d_u depend on how the patterns are generated, but note that $d_u \ge 0.5$ to avoid negative probabilities.

Expected correlation for mixed-polarity patterns

We can now work out expressions for the expected correlation. Without loss of generality, we'll do this for the case where black dots are -1, and white are +1 and background is 0 (the correlation coefficient is unchanged by shifts or rescalings). Let us first consider mixed-polarity patterns. Since black and white dots are equally likely, the mean luminance of the pattern in both eyes is zero: $\langle L \rangle = \langle R \rangle = 0$. The mean of the square depends on the dot density: $\langle L^2 \rangle = \langle R^2 \rangle = d$. To compute $\langle LR \rangle$, pairs where either pixel is background don't contribute to the sum, so we need only consider the situation where both pairs are covered by dots. And for unmatched pairs, the dots are as often oppositeluminance as same, so these also contribute nothing on average. We need only consider the matched pairs, so $\langle LR \rangle = md_m$. Putting all these into Equation 1, then, we find

$$r_{mixed} = \frac{md_m}{d}$$

Expected correlation for same-polarity patterns

For same-polarity patterns, all the dots have the same luminance, so now $\langle L \rangle = \langle R \rangle = \langle L^2 \rangle = \langle R^2 \rangle = d$, and when computing $\langle LR \rangle$ we now also need to consider the unmatched pairs which have a dot in both eyes, so $\langle LR \rangle = md_m + u(2d_u - 1)$. From Equation 1 we find

$$r_{same} = \frac{md_m + u(2d_u - 1) - d^2}{d(1 - d)}$$

Expected correlation for mixed- and same- polarity patterns when dots overlap

We now work out expressions for the dot probabilities d_u, d_m and thus for the correlations when dots are scattered at random, occluding other dots where they overlap. Consider an unmatched pair. The probability that it has a dot in the left eye is d_u . But when dots are scattered truly at random, the probability that the corresponding pixel in the right eye is also a dot is just d, the overall dot probability. Thus the probability that an unmatched pixel-pair has a dot in both eyes, which earlier we saw was $(2d_u - 1)$, is just $d_u d$. This means that

$$d_u = \frac{1}{2-d} \quad (2)$$

and thus

$$d_m = \frac{m - (1 - d)^2}{m(2 - d)}$$

Substituting this into the above expressions for r_{mixed} and r_{same} , we find

$$r_{mixed,Ov} = r_{same,Ov} = 1 - \frac{u}{d(2-d)} \quad (3)$$

Thus, when dots are scattered entirely at random, the average interocular correlation is the same for same- as for mixed-polarity patterns.

Expected correlation for mixed- and same- polarity patterns when overlap is forbidden

When overlap is forbidden, the locations of dots in the two eyes are not independent. A noise dot in the left eye cannot have been placed on top of a previous correlated dot, and so unmatched pixel-pairs are less likely to have a dot in both eyes than when overlap is allowed; see Equation 2. Thus, when overlap is forbidden, $d_u < \frac{1}{2-d}$. The smallest value for d_u is determined by the extreme case where noise dots are only placed where there is background in both eyes, when $d_u = 0.5$ and $d_m = (d - u/2)/m$. Then $r_{mixed,limNoOV} = 1 - \frac{u}{2d}$ and $r_{same,limNoOv} = 1 - \frac{u}{2d(1-d)}$ (only for the limiting case where unmatched pixels never have dots in both eyes) and so $r_{same,limNoOv} <$ $r_{mixed,limNoOy}$. Note that these expressions cannot be directly compared with those for Overlap, Equation 3, since the way the pattern is generated may change the proportion of unmatched pixel-pairs, u, even if the overall dot density is held constant.

In general for a pattern where overlap is forbidden,

$$\frac{1}{2} \le d_u < \frac{1}{2-d}$$

and so $r_{same,NoOv} < r_{mixed,NoOv}$.

Spatial properties

The above discussion simply considered images as a string of pixel values, neglecting the spatial properties except as they influence the pairwise statistics. However spatial properties could be important in real images, which undergo filtering before their interocular correlation is evaluated. We now present an alternative way of considering the problem which places more emphasis on the spatial structure. This is particularly relevant when noise is introduced by adding disparity noise, rather than by making some dots uncorrelated (as in Read et al., 2011).

Consider a gray pixel adjacent to an existing dot. and consider all the possible positions of the next random dot that could cover this pixel. Half of these random dot placements would overlap the existing dot, and hence are not permitted. As a result, the probability that this particular pixel is a dot is lower than the mean value across the image, d. This is reflected in the autocorrelation of monocular images. When overlap is allowed, the autocorrelation function is a triangle function (reflecting the autocorrelation function of a single dot), as shown in Figure 6B. But when overlap is not allowed, there are regions of negative correlation at displacements just larger than the dot width, caused by the reduced probability of dots there, as shown by the black dashed lines in Figure 6F. Note however that in mixed-polarity images the auto-correlation function is still a triangle function (pink curves in Figure 6F). This is because, while pixels adjacent to an existing dot are more likely to be gray, the probability of their being white or black is reduced equally, so that the mean product is unaffected.

The auto-correlation function, or the binocular cross-correlation where there is no noise, has a peak at 1.0 for all images (Figure 6B and F). To understand what happens when noise is added, it is useful to consider a simplified case where the dots are given one of two disparities, depicting two transparent planes. Figure 6C and G shows the cross-correlation for this situation; the blue dotted lines mark the disparities of the two planes. The images are the sum of two noisefree stereograms, for each of which the cross-correlation would be as shown in Figure 6B and F, but shifted. Although the cross-correlation of the combined stereograms is not necessarily the mean of the crosscorrelations for the individual stereograms, it closely resembles this. As illustrated in Figure 6G, for samepolarity stereograms the negative side-lobes in the autocorrelation reduce the peak cross-correlation, when there are multiple disparities.

Figure 6D and H shows the cross-correlation function for stereograms with Gaussian disparity noise. The same effect applies, meaning that for No-Overlap random-dot patterns, the cross-correlation function peaks at a lower value for same-polarity images than for mixed-polarity. This is the effect we saw in Figure 2C, where we plotted this peak correlation as a function of disparity noise. In summary, preventing dot overlap introduces systematic negative troughs in the binocular cross correlation. In the presence of disparity noise, troughs from one disparity can reduce the peaks produced by other disparities.

Finally, when considering spatial properties, it is useful to recognize that we have limited ourselves to just two image forms—random dot patterns where overlap is forbidden, and patterns where dots occlude one another and their placement is unrestricted. Patterns could be constructed in a variety of other ways (dot placement could reduce overlap but not forbid it; dot luminance could be added rather than allowing occlusion). Were psychophysical studies of stereopsis to use such stimuli, it would be important to examine their binocular correlation properties. Similarly, studies that attempted to use more naturalistic images to measure stereo performance would benefit from studying binocular correlation in the images used.

Standard models can explain the psychophysics

As we have seen, with the conventional way of generating no-overlap dot patterns, same-polarity stereograms end up with lower interocular correlation than mixed-polarity images affected by the same disparity noise. Human stereo performance declines as interocular correlation falls (Cormack, Stevenson, & Schor, 1991; Tyler & Julesz, 1978). It is therefore not surprising that humans perform better for noisy mixedpolarity stereograms than for same-polarity stereograms with the same amount of noise. Standard models of disparity encoding, such as the binocular energy model (Allenmark & Read, 2011; Filippini & Banks, 2009; Ohzawa et al., 1990; Qian & Zhu, 1997), are also based on interocular correlation, so we would expect these neurons to show stronger disparity tuning for noisy same-polarity than for mixed-polarity stimuli.

Figure 7 confirms this expectation. It shows the mean response of a population of model V1 neurons to mixed-polarity (pink, solid curves) and same-polarity (black, dashed). In the top row, the stimuli are noisefree stereograms with a uniform disparity of 6 pixels, and the disparity tuning curves are identical for mixedand same-polarity stimuli. This is the result reported by Read et al. (2011) in their figure 12. However, in the bottom two rows, we consider noisy stereograms, with the same mean disparity of 6 pixels but now disparity noise of 12 pixels. The noise effectively decorrelates the stimuli seen by the receptive fields, so the amplitude of disparity tuning falls for both stimulus types. In the middle row, the dots in the stimuli are not allowed to overlap so, as we have seen, the effective image correlation is lower for the same-polarity stimuli. Accordingly, the amplitude of disparity tuning for same-polarity stimuli is approximately half that for mixed-polarity.

Most obvious ways of decoding this population will therefore predict a mixed-polarity advantage in performance. For example, to simulate a front/back discrimination task, we could assume that the observer answers correctly when the response of the neuron tuned to the stimulus disparity of +6 pixels exceeds that of the "anti-neuron" tuned to -6 pixels (these neurons are marked with blue lines in Figure 7A and F). The quantitative neurometric performance depends on the size of the receptive fields and the number of subunits. For our ODF TE simple cells, whose Gabor receptive fields have a standard deviation of 32 pixels or 5.3 times the dot size, we obtain a performance of 73% correct for mixed-polarity stimuli and 63% for same-polarity stimuli, giving an efficiency ratio of 3.6 (equation 4 of Read et al., 2011). For a pair of ODF TE complex cells, which sum input from pairs of subunits in quadrature phase, performance is better since the phase-independence reduces the image-dependent variability: 83% correct for mixed-polarity stimuli and 70% for same-

correct for mixed-polarity stimuli and 70% for samepolarity stimuli, an efficiency ratio of 3.3. These values are comparable to those for human observers. This similarity confirms that the mixed-polarity advantage can be explained very well even by a single pair of classic energy-model simple cells. A complex neural network is not required.

The bottom row of Figure 7 is the same as the middle row except that stimulus dots are allowed to overlap. Now, the amplitude of disparity tuning is the same regardless of dot polarity. This confirms the key role of "dot repulsion" in generating the difference in interocular correlation and hence in model cell response.

Discussion

We have long been puzzled by the evidence apparently supporting independent neural mechanisms for bright and dark information-ON and OFF channels-in human stereopsis. This evidence seemed to show that disparity was easier to extract in mixedpolarity random-dot stereograms-made up of black and white dots on a gray background-than in samepolarity stereograms, where all dots are black or all white. This conclusion is puzzling for two reasons. First, linear filtering in the early visual system means that even single-polarity random-dot stereograms stimulate both ON and OFF channels. Second, by the time disparity is encoded in the primary visual cortex, the evidence suggests that ON and OFF channels have been combined. Simulations confirmed that our current models predicted the same amplitude of disparity tuning for mixed-polarity as for same-polarity stereograms (Read et al., 2011).

In this paper, we have shown that all previous results can be explained by a subtle consequence of the way the stimuli were generated. Given a value of disparity noise or of decorrelation, preventing dot overlap has a greater effect on the correlation of the entire image when dots all have the same contrast polarity relative to the background. We have shown that this effect produces a good quantitative account of human performance. This result means that the mixed-polarity advantage is entirely consistent with our current understanding of disparity encoding and with other aspects of stereo psychophysics. Stereo correspondence is not, after all, intrinsically easier in mixed-polarity stereograms; it is simply easier in stereograms with higher interocular correlation.

If the stimuli are generated differently, with dots scattered at random and occluding one another where they overlap, then the image correlation ends up being the same for mixed- and for same-polarity stimuli. Significantly, the mixed-polarity advantage is then not observed (Read et al., 2011). This strongly suggests that differences in correlation explain the mixed-polarity advantage in the original stimuli.

One of us (JCAR) has the dubious distinction of having examined the evidence for independent channels previously without having noticed this explanation. Read et al. (2011) understood Harris and Parker's result as implying that disparity is intrinsically easier to detect in mixed-polarity stimuli, and that noise functioned solely to remove the ceiling effect, bringing performance down to a level where this difference was observable. Accordingly, we examined the response of model neurons only to 100% correlated, noise-free stereograms, and concluded-wrongly-that these models could not account for the psychophysics. Goncalves and Welchman (2017) examined the response of their model to the actual, noisy stereograms used in experiments, and so revealed the effect that Read et al. (2011) missed. Amusingly, Read et al. (2011) even drew attention to the lack of dot overlap in Harris and Parker, and showed that the psychophysical mixed-polarity advantage was abolished if dots were allowed to overlap. However, we failed to appreciate the significance of this observation, attributing it to interference from the occlusion cue present in mixedpolarity stimuli.

We have now demonstrated that the psychophysical advantage for noisy mixed-polarity stereograms can be reproduced just as well with the standard binocular energy model as with the Binocular Neural Network of Goncalves and Welchman (2017). Both approaches judge depth based on which of two neurons, one tuned to near and one to far disparities, is responding most strongly. However, whereas Goncalves and Welchman's decision neurons receive input from a convolutional network made up of thousands of binocular subunits, ours each receive input from just a single binocular subunit. This demonstrates that computation specific to Goncalves and Welchman's Binocular Neural Network is not required for this result. We conclude that the differences in image correlation shown in Figure 2 are responsible for the mixedpolarity advantage in both models.

The mixed-polarity advantage was originally put forward as evidence that bright and dark information is processed separately in stereopsis, via distinct ON and OFF channels. We have now shown that the apparent advantage can be entirely explained by changes in binocular correlation in the stimulus. Of course, this result does not disprove the existence of independent ON and OFF channels in stereopsis, but it certainly undermines the current evidence for them.

Keywords: binocular vision, stereopsis, energy model

Acknowledgments

Read & Cumming

We are very grateful to Andrew Glennerster for extensive discussions on why correlation is lower for same-polarity stereograms, which greatly improved our understanding of our result and resulted in Figure 4.

Commercial relationships: none. Corresponding author: Jenny C. A. Read. Email: jenny.read@newcastle.ac.uk. Address: Institute of Neuroscience, Newcastle University, Newcastle upon Tyne, UK.

References

- Allenmark, F., & Read, J. C. A. (2011). Spatial stereoresolution for depth corrugations may be set in primary visual cortex. *PLoS Computational Biology*, 7(8): e1002142.
- Cormack, L. K., Stevenson, S. B., & Schor, C. M. (1991). Interocular correlation, luminance contrast and cyclopean processing. *Vision Research*, 31, 2195–2207.
- Cumming, B. G., & DeAngelis, G. C. (2001). The physiology of stereopsis. *Annual Review of Neuroscience*, 24, 203–238.
- Cumming, B. G., & Parker, A. J. (1997, September 18). Responses of primary visual cortical neurons to binocular disparity without depth perception. *Nature*, 389, 280–283.
- Doi, T., & Fujita, I. (2014). Cross-matching: A modified cross-correlation underlying threshold energy model and match-based depth perception.

Frontiers in Computational Neuroscience, 8, 127, https://doi.org/10.3389/fncom.2014.00127.

- Doi, T., Takano, M., & Fujita, I. (2013). Temporal channels and disparity representations in stereoscopic depth perception. *Journal of Vision*, *13*(13): 26, 1–25, https://doi.org/10.1167/13.13.26. [PubMed] [Article]
- Doi, T., Tanabe, S., & Fujita, I. (2011). Matching and correlation computations in stereoscopic depth perception. *Journal of Vision*, 11(3):1, 1–16, https:// doi.org/10.1167/11.3.1. [PubMed] [Article]
- Filippini, H. R., & Banks, M. S. (2009). Limits of stereopsis explained by local cross-correlation. *Journal of Vision*, 9(1):8, 1–18, https://doi.org/10. 1167/9.1.8. [PubMed] [Article]
- Glennerster, A. (1998) d_{max} for stereopsis and motion in random dot displays. *Vision Research*, 38(6), 925–935.
- Goncalves, N. R., & Welchman, A. E. (2017). "What not" detectors help the brain see in depth. *Current Biology*, 27(10), 1403–1412. e8, https://doi.org/10. 1016/j.cub.2017.03.074.
- Henriksen, S., Cumming, B. G., & Read, J. C. A. (2016). A single mechanism can account for human perception of depth in mixed correlation random dot stereograms. *PLoS Computational Biology*, *12*(5), e1004906, https://doi.org/10.1371/journal. pcbi.1004906.
- Henriksen, S., Read, J. C. A., & Cumming, B. G. (2016). Neurons in striate cortex signal disparity in half-matched random-dot stereograms. *Journal of Neuroscience*, 36(34), 8967–8976.
- Henriksen, S., Tanabe, S., & Cumming, B. (2016). Disparity processing in primary visual cortex. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 371(1697):20150255, 1–12, https://doi.org/10.1098/rstb.2015.0255.
- Jiang, Y., Purushothaman, G., & Casagrande, V. A. (2015). The functional asymmetry of ON and OFF channels in the perception of contrast. *Journal of Neurophysiology*, 114(5), 2816–2829, https://doi. org/10.1152/jn.00560.2015.
- Ohzawa, I. (1998). Mechanisms of stereoscopic vision: The disparity energy model. Current Opinion in Neurobiology, 8(4), 509–515.
- Ohzawa, I., DeAngelis, G. C., & Freeman, R. D. (1990, August 31). Stereoscopic depth discrimination in the visual cortex: Neurons ideally suited as disparity detectors. *Science*, 249, 1037–1041.
- Parker, A. J. (2007). Binocular depth perception and the cerebral cortex. *Nature Reviews Neuroscencei*, 8(5), 379–391.

- Prince, S. J., Cumming, B. G., & Parker, A. J. (2002). Range and mechanism of encoding of horizontal disparity in macaque V1. *Journal of Neurophysiol*ogy, 87, 209–221.
- Prince, S. J., Pointon, A. D., Cumming, B. G., & Parker, A. J. (2002). Quantitative analysis of the responses of V1 neurons to horizontal disparity in dynamic random-dot stereograms. *Journal of Neu*rophysiology, 87, 191–208.
- Qian, N., & Zhu, Y. (1997). Physiological computation of binocular disparity. *Vision Research*, 37(13), 1811–1827.
- Read, J. C. (2014). The place of human psychophysics in modern neuroscience. *Neuroscience*, https://doi. org/10.1016/j.neuroscience.2014.05.036.
- Read, J. C. A. (2005). Early computational processing in binocular vision and depth perception. *Progress in Biophysics and Molecular Biology*, 87(1 Special issue), 77–108.
- Read, J. C. A., & Cumming, B. G. (2004). Ocular dominance predicts neither strength nor class of disparity selectivity with random-dot stimuli in primate V1. *Journal of Neurophysiology*, 91, 1271– 1281.
- Read, J. C. A., & Cumming, B. G. (2017). Visual perception: Neural networks for stereopsis. *Current Biology: CB*, 27(12), R594–R596, https://doi.org/ 10.1016/j.cub.2017.05.013.
- Read, J. C. A., Parker, A. J., & Cumming, B. G. (2002). A simple model accounts for the reduced response of disparity-tuned V1 neurons to anti-correlated images. *Visual Neuroscience*, 19, 735–753.
- Read, J. C. A., Vaz, X. A., & Serrano-Pedraza, I. (2011). Independent mechanisms for bright and dark image features in a stereo correspondence task. *Journal of Vision*, 11(12):4, 1–14, https://doi. org/10.1167/11.12.4. [PubMed] [Article]
- Roe, A. W., Parker, A. J., Born, R. T., & DeAngelis, G. C. (2007). Disparity channels in early vision. *The Journal of Neuroscience*, 27(44), 11820–11831.
- Schiller, P. H. (1992). The ON and OFF channels of the visual system. *Trends in Neurosciences*, 15(3), 86– 92, https://doi.org/10.1016/0166-2236(92)90017-3.
- Schiller, P. H. (2010). Parallel information processing channels created in the retina. *Proceedings of the National Academy of Sciences*, USA, 107(40), 17087–17094, https://doi.org/10.1073/pnas. 1011782107.
- Tyler, C. W., & Julesz, B. (1978). Binocular crosscorrelation in time and space. *Vision Research*, *18*(1), 101–105, https://doi.org/10.1016/0042-6989(78)90083-4.