Measuring the spatial extent of texture pooling using reverse correlation

Daniel H. Baker [a,b,*], Tim S. Meese [b]

[a] Department of Psychology, University of York, Heslington, York YO10 5DD, UK
[b] School of Life & Health Sciences, Aston University, Birmingham B4 7ET, UK

ARTICLE INFO

ABSTRACT

The local image representation produced by early stages of visual analysis is uninformative regarding spatially extensive textures and surfaces. We know little about the cortical algorithm used to combine local information over space, and still less about the area over which it can operate. But such operations are vital to support perception of real-world objects and scenes. Here, we deploy a novel reverse-correlation technique to measure the extent of spatial pooling for target regions of different areas placed either in the central visual field, or more peripherally. Stimuli were large arrays of micropatterns, with their contrasts perturbed individually on an interval-by-interval basis. By comparing trial-by-trial observer responses with the predictions of computational models, we show that substantial regions (up to 13 carrier cycles) of a stimulus can be monitored in parallel by summing contrast over area. This summing strategy is very different from the more widely assumed signal selection strategy (a MAX operation), and suggests that neural mechanisms representing extensive visual textures can be recruited by attention. We also demonstrate that template resolution is much less precise in the parafovea than in the fovea, consistent with recent accounts of crowding.

© 2014 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/3.0/).

1. Introduction

The human visual system is structured hierarchically, with spatially local analyses at early stages feeding into representations of extensive textures, objects and surfaces at later stages. But despite extensive work focussing on local processes, e.g. in primary visual cortex (V1), we know relatively little about the later stages of representation. In particular, the limits of contrast integration across space, and the pooling strategy involved.

For several decades, the psychophysics literature has favoured a probability summation rule for pooling contrast beyond the classical receptive fields typically found in V1 (e.g. Mayer & Tyler, 1986; Robson & Graham, 1981) and contemporary accounts implement this with a MAX operator (Meese & Summers, 2012; Pelli, 1985; Tyler & Chen, 2000). This detection strategy is sometimes referred to as *signal selection* (Meese & Baker, 2011), since the MAX operator chooses one signal over several others. An alternative strategy is *signal combination*, in which many signals are combined to generate an overall response (Meese & Baker, 2011). In fact, a recent body of work supports the signal combination account of spatial pooling over the signal selection account (Baker & Meese, 2011; Meese, 2010; Meese & Baker, 2011; Meese & Summers, 2007, 2012; Morgenstern & Elder, 2012). According to this work, the signal combination strategy operates across various visual dimensions such as space, time, orientation and eye (Meese & Baker, 2013). However, although several attempts have been made (Baker & Meese, 2011; Meese, 2010), it has been difficult to firmly establish the spatial extent of the signal combination process using conventional contrast detection techniques (Baker & Meese, 2011).

We address this problem here by tailoring a psychophysical reverse correlation procedure to the problem, and developing novel analysis techniques. The general approach involves comparing the contrasts of discrete stimulus elements with observer responses in a contrast increment detection task, so as to build up a 'map' of the elements that contribute to the observer's decisions over many trials (see Figs. 2 and 3). By comparison with simulated observers using different decision rules, we infer the strategies used by the human observers in our study. Our analysis reveals that observers sum contrast over large areas (9–13 carrier cycles) using a signal combination strategy, rather than a signal selection (MAX) strategy, and that observers can be very poor at ignoring visual data at fixation, even when it is uninformative.

1.1. Simulated observers

We first consider the behaviour of two canonical model observers in a two-interval-forced-choice (2IFC) contrast increment

* Corresponding author at: Department of Psychology, University of York, Heslington, York YO10 5DD, UK.
E-mail address: daniel.baker@york.ac.uk (D.H. Baker).
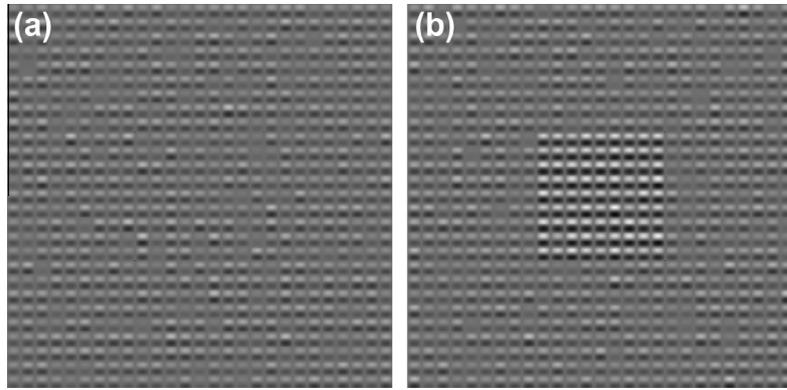
Fig. 1. Example stimulus for null (a) and target (b) intervals of a contrast increment detection task. The contrast of each element was determined by a Gaussian distribution with a mean of 32% (i.e. it was a fixed contrast pedestal, with zero-mean noise added). In the target interval (b), a contrast increment was applied to elements in the target region, here a 9 × 9 element square in the centre of the display. In the experiments, the target increment was either 0% or near threshold, so was less salient than in the example above.
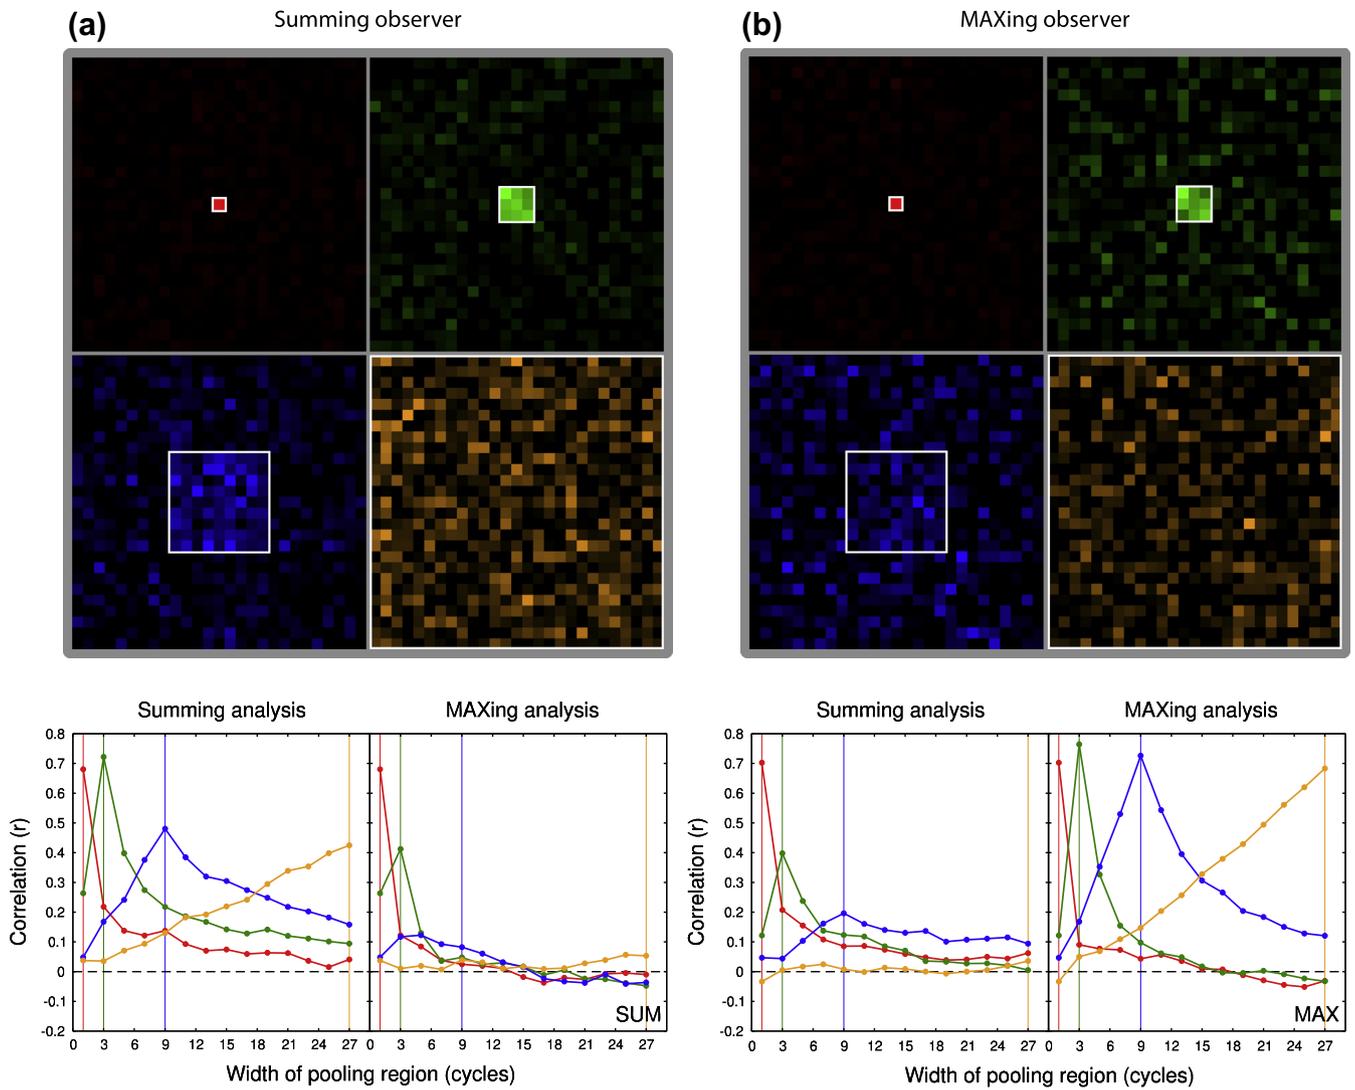
Fig. 2. Data from two model observers, who either summed (a) or MAXed (b) over the target region. The maps treat each element independently, and correlate the contrast difference across intervals with the observer's responses, over 2000 trials. Each map was peak-normalised, with the target regions indicated by white squares. Luminance at each location indicates how well the element there predicted responses, with correlations ⩽0 shown in black. In the lower graphs in each column, correlations were obtained by either summing or MAXing over a range of square windows to infer the width of the observer's pooling region and the decision rule used (see text).
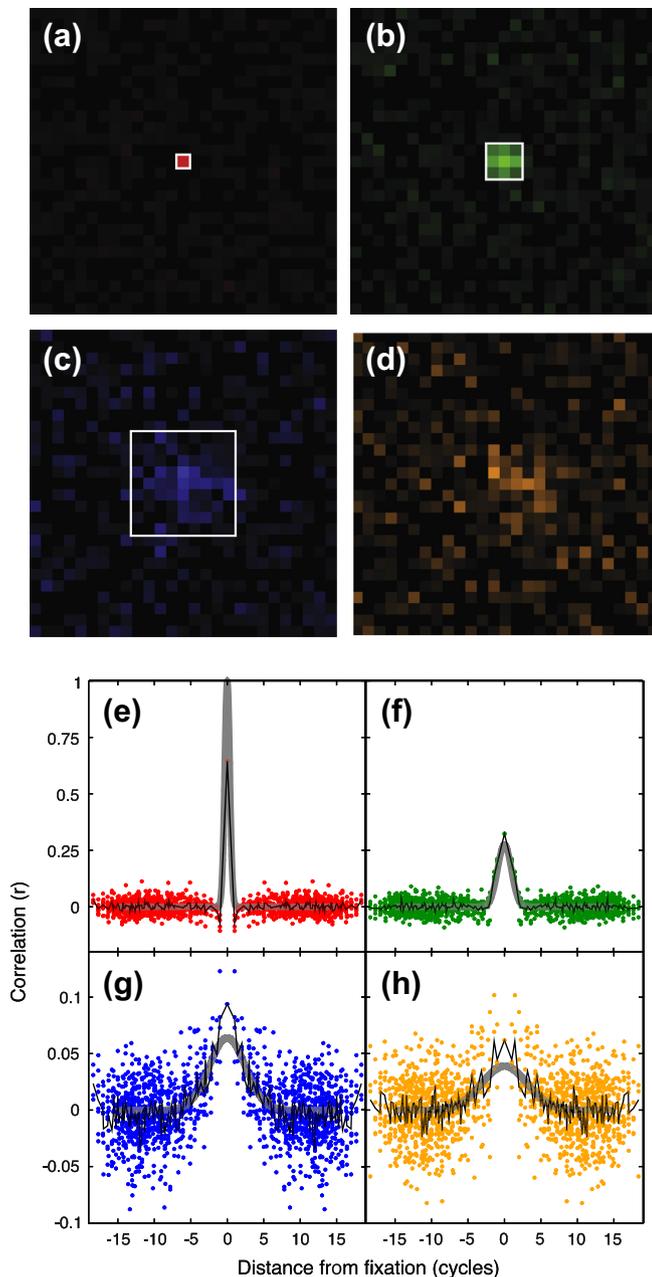
Fig. 3. Correlation maps (a–d) and cross-section traces (e–h) for observer DHB for four target sizes (indicated by the white squares in panels a–d). The maps (a–d) are normalised to the maximum value for each map, with correlation coefficients ⩽0 shown as black. The trace plots (e–h) show absolute values (note the different scales for the ordinate across the rows). Coloured points are correlation values for individual elements, plotted as a function of absolute distance from the central fixation element (and mirrored about zero). The black trace is the average of the individual correlations at each location, and the grey curves are fitted Gaussian functions with two free parameters. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

detection task (see Materials & Methods): the 'summing observer' and the 'MAXing observer'. Each model observer monitored a square region of the stimulus (e.g. Fig. 1), with a width of 1, 3, 9 or 27 elements (defined by the white squares in Fig. 2). On each trial, the summing observer adds the contrast values within the target region linearly, and selects the interval with the largest total as being the one most likely to contain the target. The MAXing observer selects the interval with the highest single contrast element in the target region. This process was repeated for 2000 trials per condition and observer. To mimic the non-determinacy of human

observers, we added zero-mean Gaussian noise to the contrast of each element on every interval of every trial.

The behaviour of the model observers was analysed in two ways. First, we performed reverse correlation for each individual element on the contrast difference between the null and target interval and the interval selected by the observer. The correlation coefficients are plotted in Fig. 2 (top row) as correlation maps, and are similar to classification images (Ahumada, 2002). For both simulated observers, correlations were concentrated in the target area for small target regions (red, green), and became more diffuse as targets grew larger (blue, orange). Using this technique, the two model observers produced similar maps, so there was no way to distinguish between the two very different decision rules.

A more informative analysis combines information across multiple elements, rather than treating each element independently. Pooling regions of different widths were assessed, within which the sum or the MAX was correlated with the responses of the simulated observer. The strongest correlations occurred when the observer rule and the analysis rule matched (i.e. the. summing analysis, for the summing observer; the MAXing analysis, for the MAXing observer), and the pooling region equalled the target region, as shown by the graphs in Fig. 2 (bottom row). When the rules were mismatched, weaker correlations were observed, particularly for the summing observer paired with the MAXing analysis. Thus, by applying this type of analysis to human results, we can determine which of the two pooling strategies they use, and also derive an estimate of the size of the stimulus region over which pooling takes place (which may be sub-optimal owing to physiological limitations). The results of this study have previously been reported in abstract form (Baker & Meese, 2013).

2. Methods

2.1. Apparatus and stimuli

All stimuli were presented on a gamma corrected NEC MultiSync Pro monitor running at 75 Hz. Stimuli were generated in Matlab running on an Apple computer, and presented at 14-bit greyscale resolution by a BITS++ box (Cambridge Research Systems, Kent, UK). The monitor was viewed from 91 cm, such that 48 monitor pixels subtended one degree of visual angle. Throughout, we define contrast as Michelson contrast in percent ($C_\% = (L_{max} - L_{min})/(L_{max} + L_{min})$, where L is luminance), often expressed in decibels ($C_{dB} = 20\log_{10}(C_\%)$).

Stimuli were square arrays of 27×27 'Battenberg' micropatterns (see Meese, 2010) with a spatial frequency of 2c/deg. In brief, these are constructed from a horizontal sinusoidal grating, which is contrast modulated by a full-wave rectified vertical sinusoidal grating at half the carrier spatial frequency. This segments the stimulus into vertical columns of horizontal stripes. Horizontal segmentation occurs naturally at the zero crossings of the carrier grating, and is accentuated by the contrast differences between the elements (carrier cycles). The contrast of each element was sampled independently on each interval of every trial from a Gaussian distribution (in linear contrast units) with a standard deviation of 10% (20 dB), and a mean of 32% (30 dB) (this is equivalent to a pedestal of 30 dB with contrast jitter ("0D noise", see Baker & Meese, 2012) added). Example stimuli are shown in Fig. 1. On the very rare occasions (<0.15% of elements) when an element's contrast exceeded 100% or fell below 0% they were fixed at these limiting values.

2.2. Procedures

Observers viewed the display from a head-and-chin rest. The task was a two interval forced choice (2IFC) contrast increment

detection between two noisy Battenberg textures (see above). One interval contained the target, the other did not, and they were presented in random order for 100 ms (with an interstimulus interval of 400 ms). The observer's task was to indicate using the buttons of a computer trackball which interval they believed contained the target. Before beginning the main experiments, we used a staircase procedure to estimate increment thresholds for the various conditions. The thresholds from this procedure guided our choice of contrasts used in the main experiments. No feedback was given in any experiment.

In Experiment I there were four target sizes, all of which were square, with widths of 1, 3, 9 and 27 elements. The target regions were centrally located. Observers were explicitly informed of target spatial extent by a quad of continuously present dark dots that framed the target area. For each target array size, observers completed 20 blocks of 100 trials using the method of constant stimuli. There were two target contrast levels: 0% and a near-threshold level informed by the staircase procedure described above. For all observers, these near-threshold contrasts were 22 dB (12.6%) for the 1x1 target, and 12 dB (4%) for the other target sizes. We included the non-zero target contrast to keep observers on task, and the 0% contrast because this produces data that is uncontaminated by the presence of a physical target (e.g. the observer is comparing two statistically identical noise fields in each trial). The blocks for different target sizes were run separately in a random order, and each block lasted around 3 min. Each observer completed 2000 trials for each target size, split between the two target contrast levels.

In Experiment II there were two conditions. In the first, a one-element target was offset below fixation by three elements. In the second, there were two target locations, equidistant above and below fixation; the target appeared in both locations in every trial. Quads of dark dots indicated the locations of the target and fixation elements. The target contrast levels were 0% and either 20% (26 dB; DHB and SAW) or 32% (30 dB; TSM), based on pilot staircase data. These pilot data can be considered to be practice sessions, with observers completing around 160 trials for each target size and location.

We calculated correlations between the observer responses and the contrast difference across intervals for each trial. To avoid bias, the target contrast increments were not included in the calculation of contrast difference between the null and target intervals. Positive differences indicate a higher contrast in the target interval and, on average, should correspond to 'correct' observer responses. Because the response data are binary, and the contrast differences continuous, the appropriate statistic is the point biserial correlation. This has an effective maximum limit of $r \approx 0.8$, because the binary response data can never fully predict the continuous contrast difference data. The calculations were performed on an element-by-element basis to produce the correlation maps in Figs. 2 and 3, and on the sum or MAX over groups of elements for the more elaborate analysis (e.g. Fig. 4). We initially calculated correlations for the two target contrast levels separately but, since these produced very similar results, we pooled the data across contrast levels to give us 2000 trials per correlation.

2.3. Observers

Three observers completed all conditions. These were the two authors (DHB, TSM), and a psychophysically experienced postdoctoral researcher (SAW) who was naïve regarding the specific expectations of the study. All observers had normal or corrected-to-normal visual acuity.
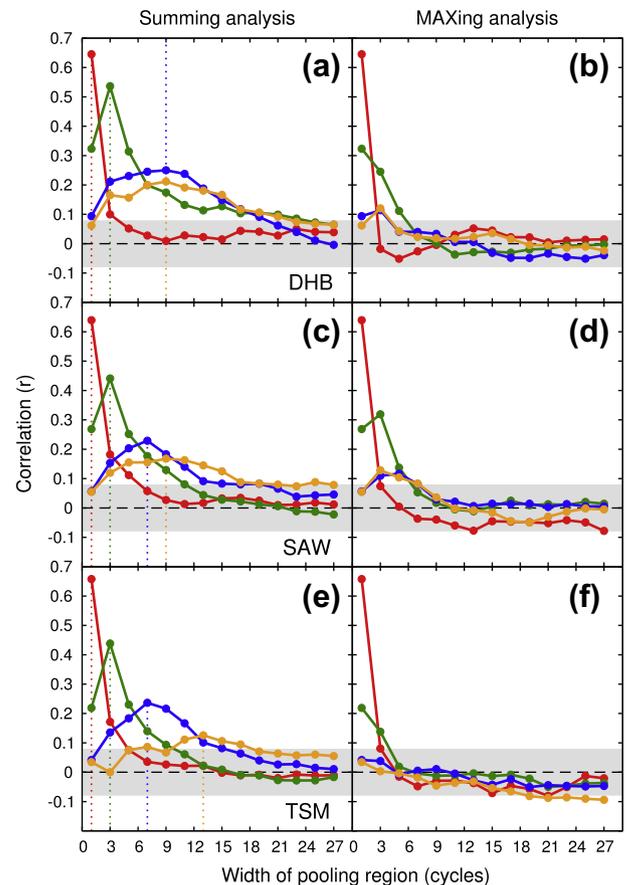
Fig. 4. Correlation coefficients for 3 observers, calculated for different pooling widths monitored by the model observer using either a summing (a, c, e) or MAXing (b, d, f) rule. Dotted lines indicate the location of the peak of each function for the summing analysis. The grey shaded region indicates the range of r values that are not significant at $p < 0.05$ for 2000 observations (Bonferroni corrected for 336 multiple comparisons (14 widths * 4 target sizes * 2 analysis methods * 3 observers)). Points outside of the shaded region indicate statistically significant correlations.

3. Results

3.1. Experiment I: monitoring targets of different sizes

Example correlation maps for a representative human observer (DHB) are shown in Fig. 3a–d, with fits to the data in Fig. 3e–h. For the small target regions (red, green), correlations were strong in the expected locations, and weak outside of them, just as for the simulated observers (Fig. 2). However, for the two larger target regions (blue, orange[1]) there was a clear clustering of correlation coefficients in the centre of the stimulus. This suggests that contrast integration occurs over a limited range, and is non-uniform over space.

To uncover the decision rule used by human observers, we calculated correlation coefficients for all three observers for each of several pooling windows (various widths) and the two decision rules described above. As shown in Fig. 4, correlations were stronger for the summing rule (left panels) than the MAXing rule (right panels) in all cases. This indicates the existence of summing mechanisms that are either pre-wired (in size), or constructed according to prevailing demands. It argues against peak-picking (akin to probability summation) from a population of local mechanisms.

---

[1] For interpretation of color in Fig. 2, the reader is referred to the web version of this article.

Next we asked how spatially extensive the pooling is. We estimated this from the location of the peak in the correlation functions (Fig. 4a, c, and e) for each observer for each stimulus size. For the smaller target sizes (red, green), the correlations peak at windows of 1 and 3 cycles wide, as for the model observers (Fig. 2), suggesting that the human observers selected a pooling region matched to the size of the stimulus in these two conditions. For the 9 × 9 element region (blue), one observer (DHB) appears to have monitored the full target region, whereas the other two (SAW, TSM) monitored a slightly smaller region, 7 elements wide. For the largest target, observers DHB and SAW based their responses only on the central 9 × 9 elements, whereas TSM was able to sum over 13 × 13 elements. However, none of the three observers could uniformly monitor the entire 27 × 27 element region. Also, note that the correlation functions for the largest target are much less sharply peaked than those for smaller regions or for the model observers (Fig. 2). This might derive from trial-to-trial variability in the size of the pooling regions that observers used, perhaps caused by switching across pooling mechanisms of different sizes and positions, or perhaps variations in fixation or attention.

3.2. Experiment II: monitoring targets away from fixation

We then asked if the stimulus region monitored by the observer changed when attention was directed to targets away from fixation. In Fig. 5 we present data from two conditions, in which the target was a single element displaced three cycles below fixation (purple), or a pair of elements offset above and below fixation by the same amount (turquoise). For comparison, the equivalent function from the first experiment for a single central target element is shown in red. We confirmed that there were no significant correlations for the elements in adjacent horizontal locations, and so plot correlations only for the vertical column of elements within which the target(s) were placed.

For a single displaced element (purple), all observers produced maximum correlations at the target location (dotted line). However, whereas for a centrally placed element (red) the adjacent elements showed no (or sometimes negative) correlation with performance, pooling occurred over a larger area for a peripheral element (purple). For observer DHB, this broader spatial footprint is approximately symmetrical about the target element. For SAW there is a hint of an additional peak at the fixated location, whereas TSM shows an inhibitory trough at fixation, and at the adjacent element on the far side to the target. These individual differences might imply differences in strategy, or in spatial uncertainty, between observers.

When targets were placed on both sides of fixation, there was even greater variation between observers. TSM shows a bimodal function (Fig. 5f), though the correlations are weak (~0.1), and more widely distributed in space (turquoise) than for a single central element (red). DHB showed a similar pattern with stronger correlations, but also showed substantial contribution from the central fixated location, even though this element was uninformative for the task (Fig. 5b). Observer SAW (Fig. 5d) based his responses on a broad region centred on the fixation point, with by far the strongest correlations occurring at and around the uninformative element at fixation.

These variations in peripheral strategy are surprising, and might suggest that some observers are very poor at dividing their attention consistently across two spatial locations, or at suppressing the influence from uninformative locations. Alternatively, it may be that small or spatially localised detectors are not available in the periphery, and responses are based on mechanisms that pool over a larger region of space (e.g. for SAW). We note that all observers were able to perform the detection task effectively, achieving
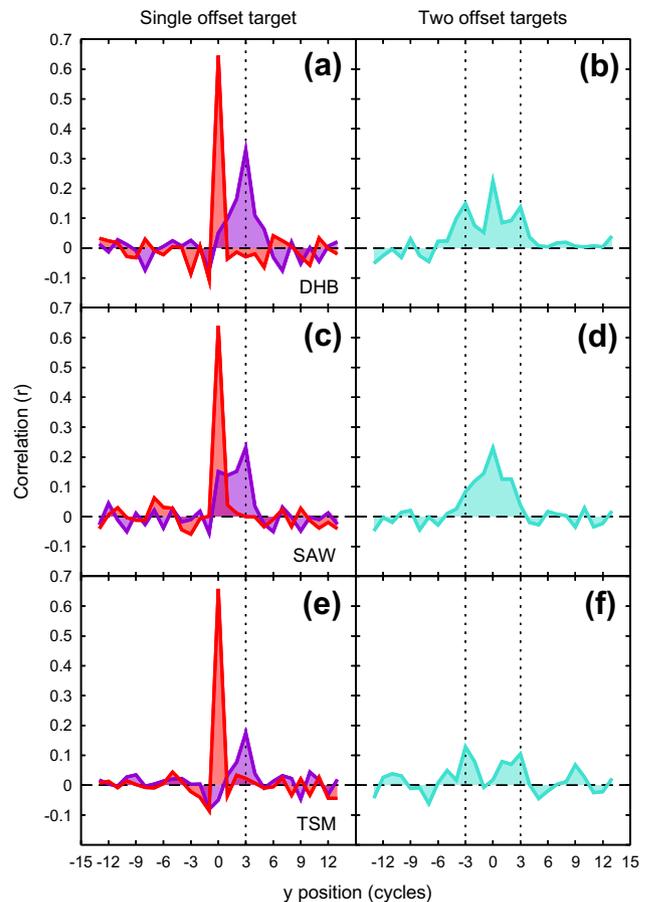
Fig. 5. Correlations between element state and observer response for targets at fixation (red) or displaced from fixation (purple, turquoise). Dotted lines indicate the target locations for the displaced conditions. It is clear that inappropriate pooling occurred when the target was away from fixation. Some observers (DHB, SAW) even produced stronger correlations at (e.g. based more of their decisions on) the fixated element than at the target elements (panels b and d). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

75% (DHB), 67% (SAW) and 79% (TSM) correct in the target-present trials. However, as detailed in the Methods section, observer TSM required a factor of 1.58 (4 dB) more target contrast relative to the other two observers to achieve this level of performance.

4. Discussion

We used a reverse correlation technique to estimate the maximum area over which observers can combine contrast, and demonstrated that this occurs by summing linearly over space, rather than merely selecting the region of highest contrast response. We also show that contrast increment detection becomes markedly less spatially precise when dissociated from fixation, and that some observers are unable to ignore an uninformative region around the fixation point. We now discuss the implications of these findings for our understanding of area summation of luminance contrast, attentional processing, and the classification image technique.

4.1. Area summation of contrast involves signal combination

A long-standing account (Robson & Graham, 1981) of the increase in contrast sensitivity with stimulus area (area summation) is that the improvement in performance owes to probability summation over multiple independent, spatially localised

detectors (see Meese & Summers, 2012; Tyler & Chen, 2000). An alternative explanation supposes that local detectors are summed at a later stage of processing producing mechanisms with a spatially extensive footprint (see Baker & Meese, 2011; Meese, 2010; Meese & Summers, 2007, 2012; Morgenstern & Elder, 2012). We have compared the predictions of these two models in several studies of contrast sensitivity (e.g. Baker & Meese, 2011; Meese & Summers, 2012, 2007; Meese, 2010; Morgenstern & Elder, 2012), all of which have favoured the linear summation account. The results here provide strong evidence that this signal combination strategy also predicts observer responses on a trial-by-trial basis better than a signal-selection (MAXing) strategy (see also Morgenstern & Elder, 2012). This result does not necessarily exclude the possibility that observers can use a MAXing strategy when it is appropriate for the task. In paradigms such as visual search, this may very well be the preferred option. However, our results indicate that pooling mechanisms that sum contrast are available to perception, and can be used in this type of experiment.

Previous estimates of the largest available size of pooling region have largely come from detailed computational modelling of psychophysical detection data. Meese and Summers (2007) concluded that their observers must have been pooling over at least 7 cycles of the carrier grating to produce the observed levels of empirical area summation. Meese (2010) cautiously extended this estimate to 16 cycles using so-called Battenberg stimuli. Using stimuli similar to Meese and Summers (2007), but a wider range of spatial frequencies, Baker and Meese (2011) estimated pooling regions of more than 12 carrier cycles.

The present results permit a more direct estimate of maximum pooling widths. Assuming a square integration region, our observers behaved in a way consistent with pooling over widths of up to 9 (DHB, SAW) or 13 (TSM) cycles (Fig. 4) for the largest stimuli. We obtained further estimates by fitting isotropic 2D Gaussian functions to the correlation maps (see Fig. 3e–h). For the largest target size, best fitting functions indicate pooling over a full-width-at-half-height (2.35 ∗ SD) of between 7 (DHB, SAW) and 11.4 (TSM) grating cycles, broadly consistent with the estimates assuming a hard-edged square summation field.

Thus, overall, two very different approaches (our previous studies above, and the one here) both lead to the conclusion that summation can extend over a substantial portion of the stimulus when the task requires it. Furthermore, we note also that the individual differences between DHB, SAW and TSM are similar across the two studies in which these three observers took part. In Baker and Meese (Baker & Meese, 2011) and here, TSM summed over a larger region than did DHB and SAW. This is also consistent with other informal observations in our laboratory. What remains less clear is why observers are unable to extend the summation field even further so as to improve performance for the larger stimuli.

4.2. Is pooling the same as attention?

In our experiments, observers monitored a large array of elements, and were instructed to base their responses on some subset of those elements. In 50% of trials, no contrast increment was applied to the elements designated as 'target', so the only difference in behaviour was due to the instructions. Thus, observers can deploy their spatial attention according to instructions. This could involve attending to multiple V1-type mechanisms spread across the stimulus and summing their responses (i.e. constructing a pooling mechanism by demand), or attending to an appropriately sized pre-wired pooling mechanism. So what implications might this have for our understanding of spatial attention?

The widespread notion of an attention 'beam' that can be directed around a stimulus at will comes largely from work on visual search (reviewed in Carrasco, 2011) and is usually

conceptualised as monitoring local mechanisms at multiple spatial locations. But as suggested above, if a range of different sized pooling mechanisms were available to the observer, one can conceive of attention as deploying the mechanism most appropriate to the task, e.g. a single large mechanism to monitor a wide area. In other words, the 'beam' becomes 'defocussed' for large stimuli, rather than moving around in space. Our observation that peripheral stimuli are poorly resolved (Fig. 5) might imply a minimum mechanism size in the parafovea, consistent with poorer peripheral resolution (Baldwin, Meese, & Baker, 2012), increased positional uncertainty (Levi, Klein, & Yap, 1987; Michel & Geisler, 2011), or some explanations of crowding phenomena (e.g. Parkes et al., 2001). The variation in observers' success in dividing their attention between two locations (and ignoring intermediate ones) is consistent with the lack of consensus on human ability to do this successfully (Jans, Peters, & De Weerd, 2010). However, it could be that with training and/or feedback, observers might improve at this task.

Throughout, we have discussed the width of pooling (or attention) in terms of cycles of the carrier grating. Although here we used only a single spatial frequency, our previous work (Baker & Meese, 2011) has indicated that area summation could be invariant of the carrier frequency when expressed in terms of cycles (see also Howell & Hess, 1978), as is also the case for retinal inhomogeneity (Baldwin, Meese, & Baker, 2012; Pointer & Hess, 1989; Robson & Graham, 1981). Since natural scenes are broadband (Field, 1987), predicting which combination of mechanisms will govern performance in everyday environments is not straightforward. We anticipate that advances in this area will require combining multiscale filter models (e.g. Georgeson et al., 2007) with detailed formal models of attention (e.g. Gobell, Tseng, & Sperling, 2004).

4.3. Comparison with classification image (CI) studies

The reverse correlation technique used here to produce the correlation maps (e.g. Fig. 3) is related to the CI technique (Ahumada, 2002; Morgenstern & Elder, 2012; Murray, 2011). We also calculated CIs for our experiment by averaging the contrasts of the intervals selected by the observers as containing the target, and subtracting the averaged contrasts from the other (nonselected) intervals. When peak-normalised, these were almost indistinguishable from our correlation maps (i.e. Fig. 3), revealing a close similarity between the two methods (not shown). However, calculating correlation coefficients is more flexible, as it can be easily extended to compare different models and decision rules, as we have done (e.g. Fig. 4).

We think that our approach here is valuable for two reasons. First, our use of contrast jitter instead of white pixel noise means that larger templates can be measured without using extremely large pixel sizes (or requiring implausible numbers of trials for small pixel sizes). Second, the testing of model hypotheses on a trial-by-trial basis is very powerful (see also Morgenstern & Elder, 2012; Neri, 2011), and offers important insights beyond the visual representation of the observer's template produced by standard CI techniques (reviewed in Murray, 2011). We note that a recent study (Morgenstern & Elder, 2012) also used a classification image method to ask related questions about spatial pooling strategies. This work, which used traditional white pixel noise, also found evidence for signal combination over signal selection, and provided estimates of the local filters used for detection, but did not attempt to estimate of the size of the pooling region.

5. Conclusions

We have presented a new multivariate technique for measuring the extent of spatial pooling. Reverse-correlation shows that pool-

ing extends to around 9–13 carrier cycles, and can be precisely limited to small target areas in the central visual field. However, spatial precision is much poorer for larger stimuli, and for small stimuli placed in the parafovea. These findings prompt new metaphors for spatial attention, and indicate that large aggregating mechanisms are available to top-down monitoring in basic detection tasks.

Acknowledgment

References

Ahumada, A. J. (2002). Classification image weights and internal noise level estimation. *Journal of Vision, 2*, 121–131.

Baker, D. H., & Meese, T. S. (2011). Contrast integration over area is extensive: A three-stage model of spatial summation. *Journal of Vision, 11*(14), 1–16. art. 14.

Baker, D. H., & Meese, T. S. (2012). Zero-dimensional noise: The best mask you never saw. *Journal of Vision, 12*(10), 1–12. art. 20.

Baker, D. H., & Meese, T. S. (2013). Using psychophysical reverse correlation to measure the extent of spatial pooling of luminance contrast. *Perception, 41*, 1512.

Baldwin, A. S., Meese, T. S., & Baker, D. H. (2012). The attenuation surface for contrast sensitivity has the form of a witch's hat within the central visual field. *Journal of Vision, 12*(11), 1–17. art. 23.

Carrasco, M. (2011). Visual attention: The past 25 years. *Vision Research, 51*, 1484–1525.

Field, D. J. (1987). Relations between the statistics of natural images and the response properties of cortical cells. *Journal of the Optical Society of America A: Optics, Image Science, and Vision, 4*, 2379–2394.

Georgeson, M. A., May, K. A., Freeman, T. C. A., & Hesse, G. S. (2007). From filters to features: Scale-space analysis of edge and blur coding in human vision. *Journal of Vision, 7*(13), 1–21. art. 7.

Gobell, J. L., Tseng, C., & Sperling, G. (2004). The spatial distribution of visual attention. *Vision Research, 44*, 1273–1296.

Howell, E. R., & Hess, R. F. (1978). The functional area for summation to threshold for sinusoidal gratings. *Vision Research, 18*, 369–374.

Jans, B., Peters, J. C., & De Weerd, P. (2010). Visual spatial attention to multiple locations at once: the jury is still out. *Psychological Review, 117*, 637–684.

Levi, D. M., Klein, S. A., & Yap, Y. L. (1987). Positional uncertainty in peripheral and amblyopic vision. *Vision Research, 27*, 581–597.

Mayer, M. J., & Tyler, C. W. (1986). Invariance of the slope of the psychometric function with spatial summation. *Journal of the Optical Society of America A: Optics, Image Science, and Vision, 3*, 1166–1172.

Meese, T. S. (2010). Spatially extensive summation of contrast energy is revealed by contrast detection of micro-pattern textures. *Journal of Vision, 10*(8), 1–21. art. 14.

Meese, T. S., & Baker, D. H. (2011). Contrast summation across eyes and space is revealed along the entire dipper function by a "Swiss cheese" stimulus. *Journal of Vision, 11*(1), 1–23. art. 23.

Meese, T. S., & Baker, D. H. (2013). A common rule for integration and suppression of luminance contrast across eyes, space, time, and pattern. *i-Perception, 4*, 1–16.

Meese, T. S., & Summers, R. J. (2007). Area summation in human vision at and above detection threshold. *Proceedings of the Royal Society of London. Series B: Biological Sciences, 274*, 2891–2900.

Meese, T. S., & Summers, R. J. (2012). Theory and data for area summation of contrast with and without uncertainty: Evidence for a noisy energy model. *Journal of Vision, 12*(11), 1–28. art. 9.

Michel, M., & Geisler, W. S. (2011). Intrinsic position uncertainty explains detection and localization performance in peripheral vision. *Journal of Vision, 11*(1), 1–18. art. 18.

Morgenstern, Y., & Elder, J. H. (2012). Local visual energy mechanisms revealed by detection of global patterns. *Journal of Neuroscience, 32*, 3679–3696.

Murray, R. F. (2011). Classification images: A review. *Journal of Vision, 11*(5), 1–25. art. 2.

Neri, P. (2011). Coarse to fine dynamics of monocular and binocular processing in human pattern vision. *Proceedings of the National Academy of Sciences of the United States of America, 108*, 10726–10731.

Parkes, L., Lund, J., Angelucci, A., Solomon, J. A., & Morgan, M. (2001). Compulsory averaging of crowded orientation signals in human vision. *Nature Neuroscience, 4*, 739–744.

Pelli, D. G. (1985). Uncertainty explains many aspects of visual contrast detection and discrimination. *Journal of the Optical Society of America A: Optics, Image Science, and Vision, 2*, 1508–1532.

Pointer, J. S., & Hess, R. F. (1989). The contrast sensitivity gradient across the human visual field: With emphasis on the low spatial frequency range. *Vision Research, 29*, 1133–1151.

Robson, J. G., & Graham, N. (1981). Probability summation and regional variation in contrast sensitivity across the visual field. *Vision Research, 21*, 409–418.

Tyler, C. W., & Chen, C. C. (2000). Signal detection theory in the 2AFC paradigm: Attention, channel uncertainty and probability summation. *Vision Research, 40*, 3121–3144.