Noise Induced Perceptual Completion. [1,2]

Paul R. Schrater [a] Trinh Khuu [b] Daniel Kersten [a]

[a] *Department of Psychology,University of Minnesota,Minneapolis,MN,55455*

[b] *Carleton College, Northfield, MN 55057 (Now at: State University of New York State College of Optometry)*

Abstract

When an observer is faced with an ambiguous, occluded or degraded image, the image will frequently be perceptually completed. In perceptual completion observers report seeing the presence of features or objects which are actually absent that are the complement or "completion" of features or objects in the image. We report a novel completion effect resulting from adding half of a bilaterally symmetric image to visual noise, which results in an illusory appearance of the missing half-image in the noise. In this paper we investigate the role of prior knowledge and symmetry in this perceptual completion stimulus using a signal detection task. The task allows us to obtain separate measures of the magnitude of the completion effect (via false alarm rate) and the processing involved in the completion effect (via sensitivity, d'). With these separate measures, we can begin to distinguish between low-level and high-level theories of perceptual completion. We contrast low-level theories that produce differences in sensitivity between completion and no completion conditions with high-level theories that predict no change in detection sensitivity. Our results are most consistent with completion in our stimuli occuring as a result of a high-level perceptual inference or judgement rather than low-level, bottom-up processing. In addition, we show that additional symmetry information can augment the completion effect without changing sensitivity.

Key words: perceptual completion,symmetry,surface,ideal observer,detection

[1] Supported by NIH RO1 EY02857, RO1 EY11507, and The Presidential Fellowship of Carleton College.
[2] Draft submitted to *Vision Research*

1 Introduction

1.1 Perceptual Completion

The visual system is constantly faced with the task of interpreting image data that is ambiguous, corrupted, or missing (occluded). In the presence of these difficulties, the visual system must bring to bear prior knowledge to successfully interpret the image, whether explicitly represented or implicitly built into the visual system's architecture (Kersten, 1990; Yuille and Bulthoff, 1996; Kersten and Schrater, 2000). One of the most striking demonstrations of the visual system's use of prior knowledge is the phenomenon of perceptual completion. In perceptual completion observers report seeing the presence of absent features that are the complement or "completion" of features or objects present in the image (Pessoa et al., 1998).

Common examples of perceptual completion include the Kanizsa Triangle (Kanizsa, 1979), abutted grating illusory contours (Paradiso et al., 1989) and Neon Coloring spreading (van Tuijl, 1975). It is not clear how the visual system effects completion, or whether it is the same phenomenon in each example. On the other hand completion phenomena like illusory contours and neon-coloring spreading are frequently attributed to a side-effect of the image processing, rather than the result of a perceptual inference (Grossberg and Mingolla, 1985; Lesher, 1995; Spillman and Dresp, 1995). Many attempts at determining the conditions for completion have been made, beginning with the Gestaltists (Wertheimer, 1923; Kofka, 1935). While several attempts have been centered around formulating general laws or principles like proximity, good continuation, and common fate (Ullman, 1976; Parent, 1989; Kellman and Shipley, 1991; Heitger, 1993; Yin et al., 1997), many recent ideas revolve around completion as a result of object or surface inference or fitting (Nakayama et al., 1989; Nakayama and Shimojo, 1992; Sekuler et al., 1994; Tse, 1999a,b; van Lier, 1999). Stimuli like the Kanizsa Triangle suggest that one of the key requirements for generating perceptual completion is an adequate explanation for the missing or corrupted data (Rock and Anson, 1979; Dennett, 1992). For example, in the triangle, the presence of the "pac-men" can be explained by an occluding triangle, but only if the occluding surface has the same color as the background.

However, if completion occurs when the visual system finds an adequate explanation for the missing or corrupted data then we should be able to create novel stimuli that produce completion based on this principle. Below we describe a novel class of stimuli that uses a partial image placed in an ambiguous context to induce a kind of perceptual completion.

One of the simplest contexts that indicates the possibility of missing data is visual noise. Because all patterns exist on average in white noise, it is possible that the visual system can be biased toward preferentially extracting a particular signal from the noise. We have found that adding half of a bi-laterally symmetric image to uncorrelated noise produces an illusory completion for all observers tested. In this case, familiarity with symmetric objects suggests that the other half should be present, while the noise provides an explanation for why the data is missing. Examples are shown in fig. 1a and fig. 1b, half of a face and half of an image of bull's-eye rings. The two patterns were used as stimuli in the experiments which follow, and were selected because of their symmetry and potential familiarity. When additional symmetry information was provided in the form of symmetric noise, the subjective reports of completion are much higher. Examples of this are shown in 1c and fig. 1d. These completion effects are to our knowledge novel, and it is important to point out that there is no reason the visual system should necessarily complete these images. On the contrary, there is substantial evidence in the image for the absence of the right half.

The main goals of this paper are to quantify this completion effect and to investigate how it arises. We distinguish three basic possibilities. First, completion could be the result of an automatic low-level "filling-in" process, in which missing image data is filled-in via image-based extrapolation that assumes implicit, generic prior knowledge about mirror symmetry. Second, completion could be the result of fitting a global template to the image data (e.g. fitting a familiar object, surface or global pattern model to the image data). Third, completion could be the result of a high-level decision that a particular global pattern is present, despite the lack of local image evidence. In the next section we explain how these basic possibilities can be distinguished within the context of signal detection theory, where low-level and model-fitting processes predict changes in *sensitivity* (d'), while high-level processes predict no change in sensitivity, but instead a change

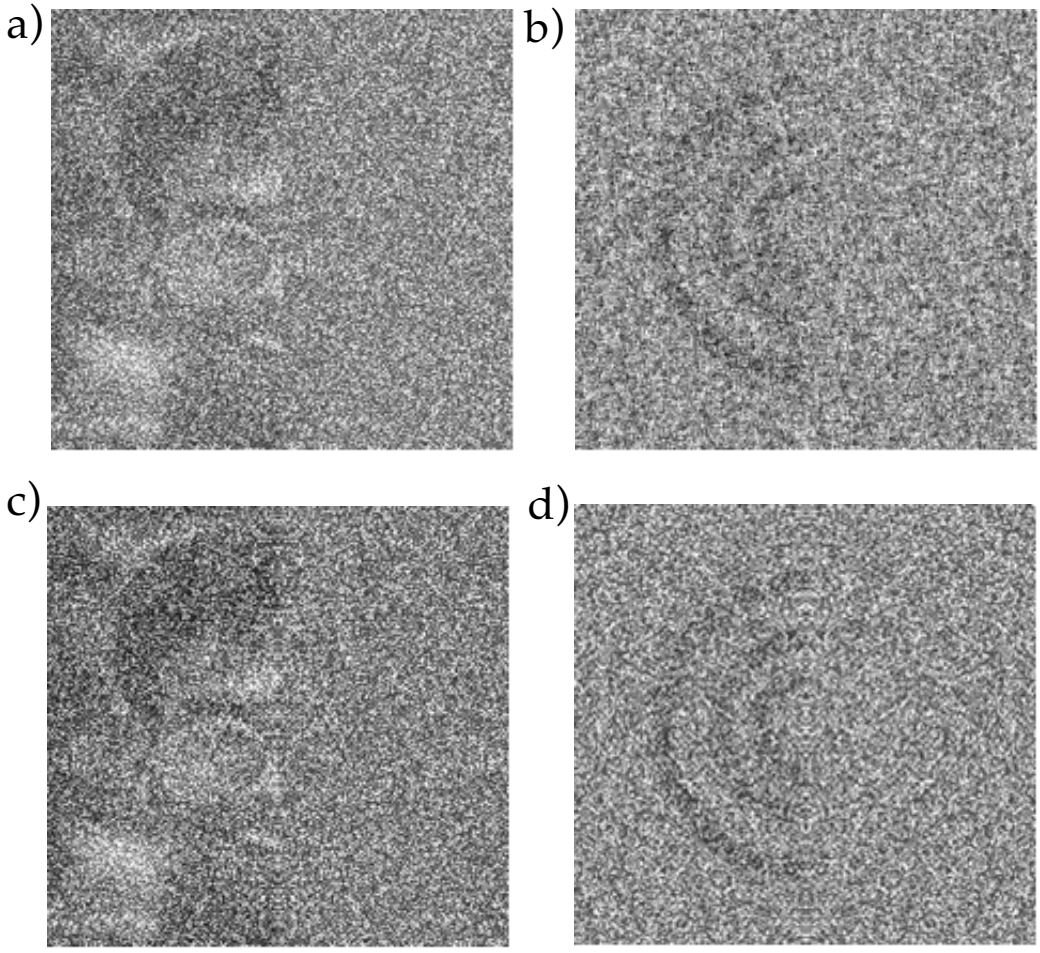

Fig. 1. Does the detection of symmetry induce illusory figural completion? The presence of symmetry is cued by symmetric noise. **a)** Half image of a face added to noise. Many observers report seeing illusory completions of the face, particularly close to the midline.**b)** Half image of a set of concentric rings added to noise. **c)** The same face half image as in a), but added to symmetric noise produced by reflecting the right half of the noise image into the left. **d)** The half rings image added to symmetric noise. Most observers report an increase in completion in the presence of symmetric noise.

in the *decision criterion*.

1.2 Task and Stimuli

In order to study completion within a signal detection framework, we have developed a measure of completion using a detection task. The benefit of using the signal detection paradigm is that the

magnitude of the perceptual completion can be measured separately from processing sensitivity.

The idea is straightforward: we add the left-half of a symmetric image to noise. On each trial, observers were asked to detect the presence of the right-half image added to noise (Green and Swets, 1974). The right-half was presented 50% of the time at just detectable contrasts, and was absent half the time. A natural measure of completion is the proportion of times the observer reports the presence of the right-half of the pattern when the right half-is absent. This measure is simply the false alarm rate for the detection task. On the other hand, signal detection theory allows us to simultaneously study sensitivity through the d' measure of sensitivity. d' takes into account both hit and false alarm rates, and remains constant for fixed distributions of the underlying decision variable. Thus using the detection task allows us to separate the magnitude of the completion effect from the effect of processing on the stimulus on the right hand.

We performed a set of experiments that addressed three issues: 1) to measure the magnitude of the completion effect; 2) investigate models of completion and 3) assess the role of additional symmetry information on completion. The conditions for the experiment are schematically presented in figures 2&3. The conditions are arranged in a table, columns showing the type of reference image that was placed in the left hand side, and rows showing the two noise types used: symmetric and non-symmetric.

The three reference types, "Noise Alone", "Completion", and "Signal-Copy" allow us to quantify the completion effect and control for non-completion changes in false-alarm rate. The difference between false alarm rates in "Noise Alone" vs. "Completion" reference conditions provide a natural measure of the magnitude of the completion effect. On the other hand, an increase in false-alarms could be due to having information about the signal in the left hand side rather than completion per se. Comparing false alarm rates between the "Completion" and "Signal-Copy" reference conditions controls for this possibility. The use of symmetric and non-symmetric noise allow us to investigate the effect of the addition of signal-uncorrelated symmetry information on completion. Another non-completion use of the signal buried in the left hand side is to decrease the observer's uncertainty about which signal is present. In order to minimize this effect, the signal to be detected was always presented in full contrast above the noise image on the right

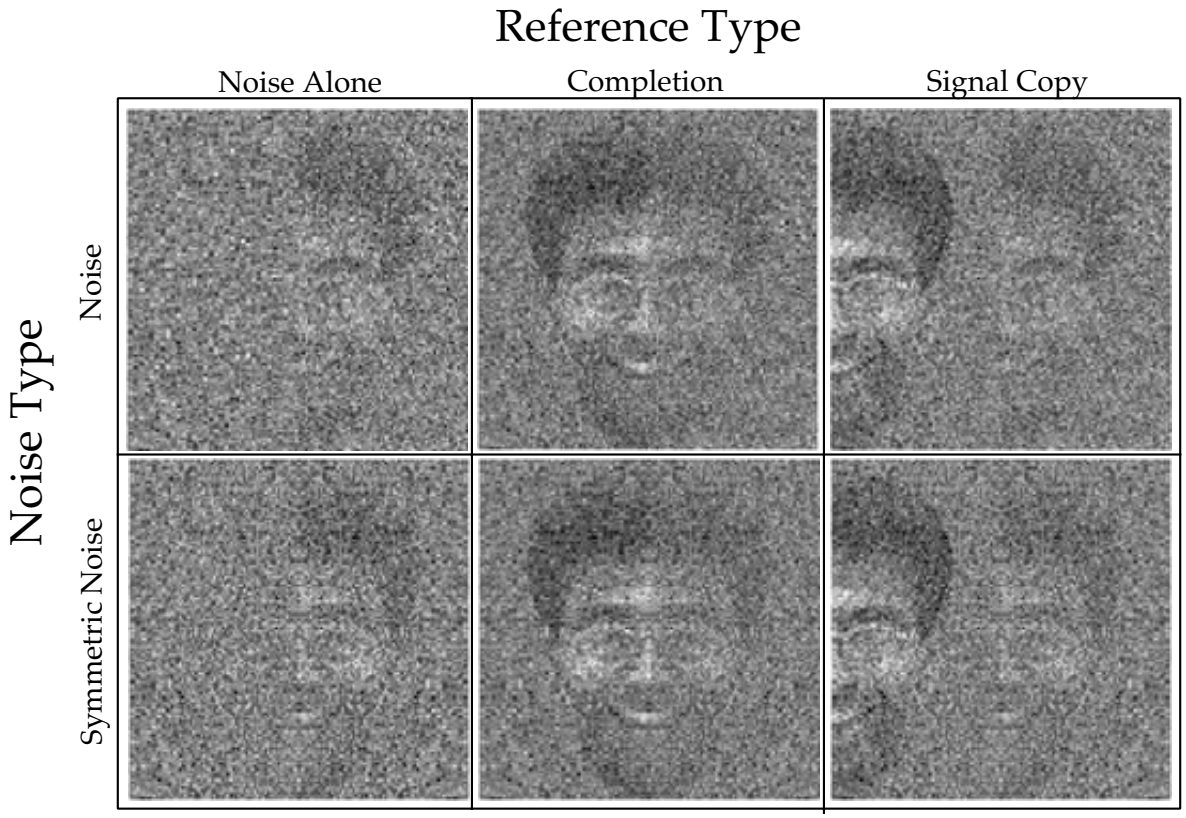

Fig. 2. Experimental stimuli and conditions. Face images.

hand side. This control is based on the results of Burgess (Burgess and Ghandeharian, 1984) who showed that subjects were able to use this kind of information to reduce the effect of signal uncertainty.

1.3 Predictions

Low-Level image-based extrapolation

The idea that completion results as a result of low-level processing predicts a change in d', that in general will produce a decrease. This is true because processing theories of completion modify the image data to produce completion. In general it is impossible to modify the input data without affecting sensitivity (d').

Intermediate-level Model-based extrapolation

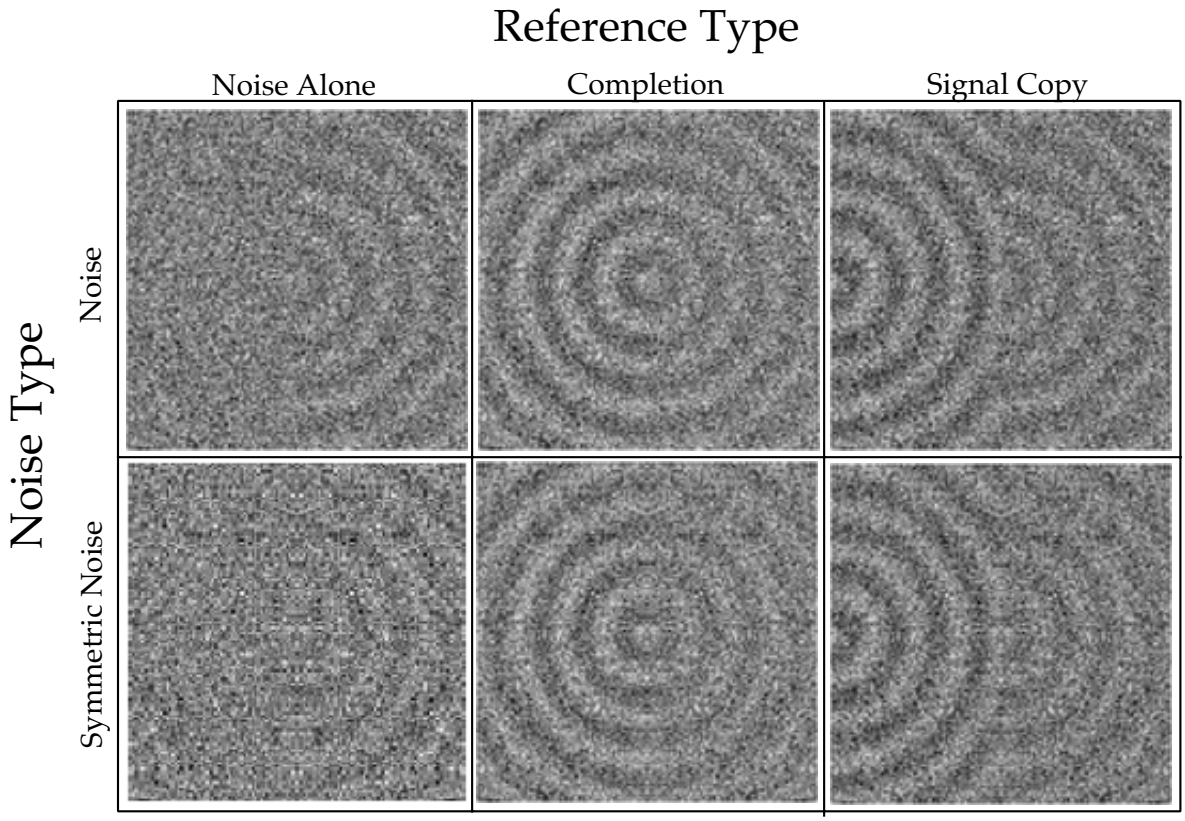

Fig. 3. Experimental stimuli and conditions. Rings images.

The idea that completion results from a model-fitting process in general predicts a decrease in d'. This decrease is due to the fact that fitting a complete model or template will result in the inclusion of more background noise into the decision variable.

High-Level decision

The idea that completion results from a high-level decision process predicts no change in d'. Here familiarity with the complete image of a face or rings can increase the prior on viewing the complete image, which in turn biases the decision and changes the false alarm rate. However, such changes in the decision criterion do not affect d'.

2 Methods

2.1 Subjects

Twelve subjects (five males and seven females) participated in four sessions of 800 trials each. All subjects were naive as to the purposes of the experiment except TK (author) and CS. Observers viewed the stimuli binocularly and had normal or corrected-to-normal vision.

2.2 Stimuli

Stimuli consisted of: a signal image presented in the right half of the display, a reference image presented in the left half, and both of theses images were added to Gaussian white luminance noise. The signal image was either the left half of a face or rings, presented at 1.3% contrast. The reference image was either a translated copy of the signal image, denoted "Signal Copy" and presented at 20% contrast, a reflected copy (denoted "Completion") also presented at 20%, or no image (denoted "Noise Alone"). Noise was either symmetric or non-symmetric and had a residual mean square contrast of 0.2. Symmetric noise was created by copying the reflected left half noise image into the right half.

2.3 Apparatus

Stimuli were presented on on a gamma calibrated color monitor in 8 bit grayscale, using *Psych-Toolbox for Matlab* display software written by David Brainard and Denis Pelli for the Macintosh. Viewing distance was 50 cm and screen pixels were 0.6 millimeters in height and width. Stimuli were 128x128 pixels, and subtended 17.3 degrees visual angle.

2.4 Procedure

The twelve experimental conditions were randomly presented in 4 sessions of 800 trials each. The task was to determine whether the right half of the noise image contained the target (either a face or concentric rings). The targets were presented at a contrast levels of zero or 1.3% randomly 50% of the time. The signal contrast level was chosen to produce a d' of about 1 in observers TK and CS. A full contrast copy of the signal was presented immediately above the target (right) side of the noise image.

The twelve experimental conditions consisted of two different target stimuli x 3 reference types x 2 types of noise. Subjects responded by pressing the right arrow key for "yes," if they saw the target and the left arrow key for "no," if they did not see the target. Subjects were given no feedback, but were given as much time as needed to respond to each trial. and were given breaks between trial runs when needed. Before data collection began, subjects performed 128 practice trials of to provide familiarity with the experimental task.

2.5 Data Analysis

For each condition, false-alarm rates and d' were measured. False Alarm rates are presented as *normal values*,(e.g. $FA = \Phi^{-1}(p(respondsignal|nosignal))$, where $\Phi^{-1}()$ is the inverse cumulative normal function). d' was computed via the relation $d' = \Phi^{-1}(p(respondsignal|signal)) - \Phi^{-1}(p(respondsignal|nosignal))$.

The data were analyzed both across and within subjects. Standard statistical analyses (T-tests, one-way ANOVAs) were performed across subjects on the raw d' and false-alarm normal values. Individual analyses were performed via parametric bootstrap (equivalent to Monte-Carlo simulation) by assuming the subject's response data are binomially distributed random variables. Standard errors were estimated from the standard deviation of 1000 bootstrap replications. Bootstrap T-tests follow Efron & TibshiraniEfron and Tibshirani (1993).

In order to test perform ANOVAs across the individual data, we extended the idea for the boot-

strap T-test to perform a bootstrap ANOVA. Given several conditions (groups), the bootstrap ANOVA tests whether the within condition means μ_i are significantly different from the group mean $\mu_g = \frac{1}{N_g}\sum_{i=1}^{N_g}\mu_i$, where N_g is the number of groups. The bootstrap ANOVA statistic computes:

$$F_b^j = \frac{1}{N_g}\sum_{i=1}^{N_g}(s_i^j - \mu_i)^2 - \frac{1}{N_g}\sum_{i=1}^{N_g}(s_i^j - \mu_g)^2$$

. This statistic is the deviation of bootstrap sample means from grand mean minus the deviation of the bootstrap sample means from the sample means, where s_i^j are the jth bootstrap replicate from the ith group, and N is the number of bootstrap replicates. If the sample means equal the group mean, then this quantity is distributed around zero. A test of the significance is given by the proportion of bootstrap replicates that are less than zero. Essentially, the test uses the deviations of the bootstrap sample means around the group mean to generate a distribution for the Null Hypothesis, that the sample means are all equal to the group mean.

3 Results

The results will be presented first pooled across subjects to highlight the main results followed by a presentation of the individual results.

3.1 Group Results

3.1.1 Completion Effect

A bar graph of the mean false alarm rate averaged across the 12 subjects for faces and rings is shown in fig. 4. Note that the false-alarm rates are presented as normal values. The false-alarm rates are highest for the completion condition in both faces and rings in both symmetric and non-symmetric noise. The fact that the false alarm rates are significantly greater for the 'Completion' reference condition than the 'Signal Copy' or 'Noise Alone' conditions (T-test, t=7.96, df=22, $p < 0.001$) indicates the presence of the completion effect. This result also holds in a within subjects data analysis, where all subject's data show the same effect at the $p < 0.005$ level (see

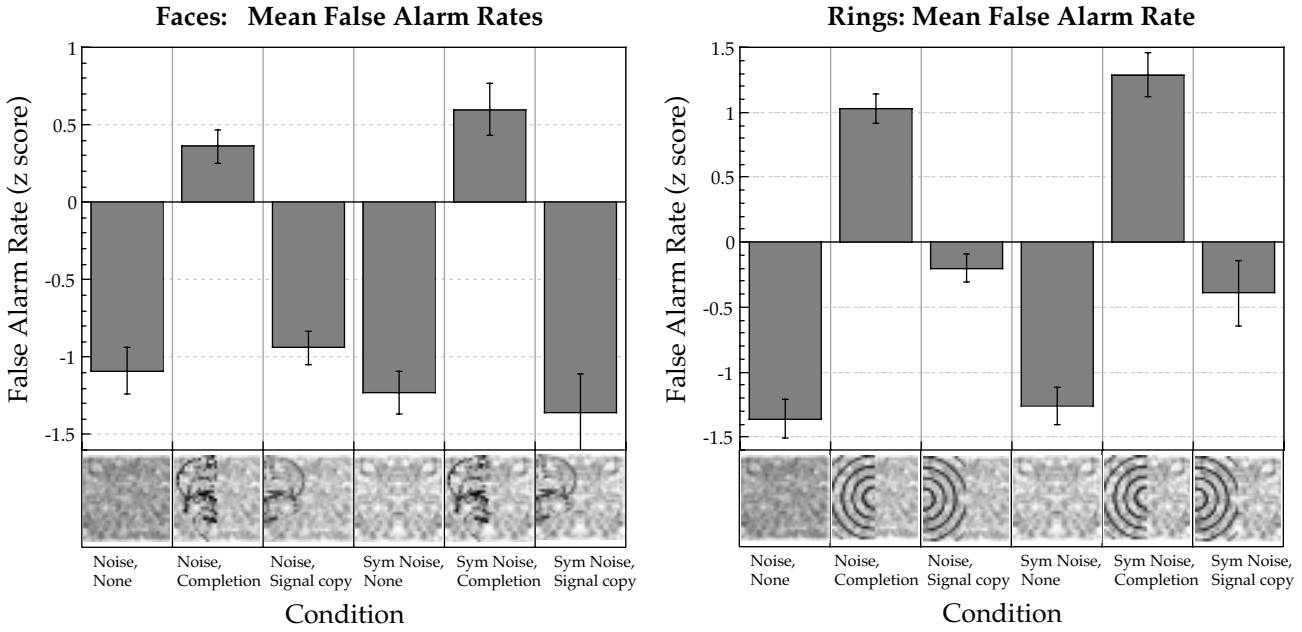

Fig. 4. False-alarm normal values averaged across the 12 subjects are shown for the faces and rings stimuli. Error bars represent the standard error of the mean.

Methods section for the details of individual analysis).

For faces stimuli, there is no effect of the presence of the signal copy reference on false-alarm rates over Noise Alone. However, for rings stimuli there is a significant increase for signal copy over no reference, suggesting that there may be some completion in the this condition as well. Subjectively, the rings in the signal copy condition look that ripples or waves in the presence of the signal in the right hand side (see figure 3). These waves have a regular appearance that may allow a kind of completion that is not possible in the signal copy condition for faces stimuli.

3.1.2 Effect of Symmetric Noise

The main effect of symmetric noise on false-rates is to significantly increase the magnitude of the completion effect for both faces and rings (T-test: $p < 0.05$). However, this increase is small, and is only significant in 7 out of 12 subjects analyzed individually for faces and 6 out of 12 subjects for rings. Thus, symmetry information which is unrelated to the signal can still promote completion. However, symmetric noise alone does not significantly increase false-alarm rates for 9 of the 12 subjects .

3.1.3 Effects of completion on the d′ measure of performance.

A bar graph of d' averaged across the 12 subjects shows no significant change in d' for any of the conditions for the face stimulus, and only one significant effect for the rings stimulus: a decrease in performance for the "signal-copy" reference in non-symmetric noise. The condition that gave rise to the highest false alarm rate, the "Completion" reference in symmetric noise, gave the highest d' rate but the results were not statistically significant.This lack of significance could be due to averaging across subjects with large differences in d' ranges, which could mask within subject differences. However, performing the within subjects analysis shows no significant difference in d' across the reference types and noise types for all but two subjects at the $p < 0.05$ level and no subjects at the $p < 0.01$ level. These results suggest that the completion effect is due to a high-level inference or judgement rather than a result of bottom-up processing.

Given the results of the d' rate, the evidence supports a theory of perceptual completion that produces no change in performance. We will discuss the kinds of completion mechanisms that could produce these results in the discussion section.

3.1.4 Effects of symmetric noise on performance.

The presence of symmetric noise has no significant effect on performance. For the completion stimulus, symmetric noise seems to enhance FA rate but not d'.The simplest explanation of this phenomena is that the presence of symmetric noise acts as a flag which increases the prior probability of the full face or rings image occurring.

3.2 Individual results

Individual results are presented in tables 1-4.

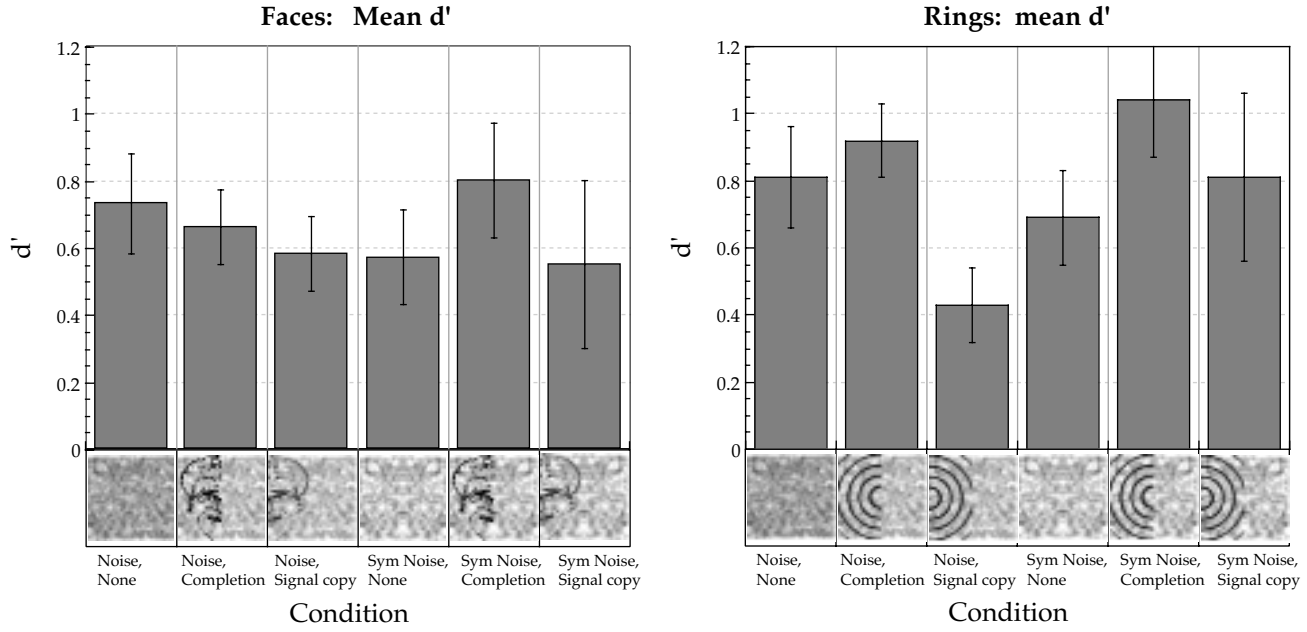

Fig. 5. False-alarm normal values averaged across the 12 subjects are shown for the faces and rings stimuli. Error bars represent the standard error of the mean.

Conditions		**Subjects**											
Noise Type	*Reference*	CE	CY	JP	ML	TC	CS	TK	PC	VR	WS	LL	JL
Noise	Noise Alone	-1.1	-1.0	-1.7	-1.5	-0.9	-0.1	-1.5	-0.1	-1.5	-1.05	-1.43	-1.38
	Completion	0.1	0.9	0.8	0.2	-0.2	0.3	0.7	0.6	0.0	0.8	-0.2	0.2
	Signal Copy	-0.9	-0.6	-1.0	-1.2	-0.7	-0.8	-1.4	-0.3	-0.7	-1.7	-1.2	-1.0
Sym. Noise	Noise Alone	-1.0	-1.1	-1.9	-1.3	-1.2	-1.2	-1.4	-0.7	-1.4	-1.2	-2.1	-0.3
	Completion	0.5	0.8	1.1	1.0	0.2	0.1	0.0	1.2	0.3	0.8	-0.4	1.7
	Signal Copy	-1.3	-0.8	-1.3	-1.3	-0.9	-1.4	-3.1	-0.9	-1.2	-0.8	-3.1	-0.4

Table 1

Faces stimuli: Table of FA rates expressed as normal values for all subjects.

Conditions		**Subjects**											
Noise Type	*Reference*	CE	CY	JP	ML	TC	CS	TK	PC	VR	WS	LL	JL
Noise	Noise Alone	-0.9	-1.0	-1.9	-1.1	-1.3	-0.5	-0.6	-1.8	-3.1	-1.3	-1.8	-1.0
	Completion	1.0	1.1	1.1	1.3	-0.2	0.3	1.7	1.8	0.5	0.8	1.7	1.3
	Signal Copy	-0.1	-0.1	-0.9	-1.1	-0.9	0.4	0.6	-0.6	-0.3	-1.4	0.0	1.9
Sym. Noise	Noise Alone	-0.7	-0.9	-1.7	-1.4	-1.4	-0.5	-0.7	-1.5	-2.1	-1.6	-1.8	-0.7
	Completion	0.9	1.0	1.1	1.3	-0.0	0.5	3.1	1.6	0.8	3.1	0.7	1.2
	Signal Copy	-0.1	0.2	-1.0	-1.1	-1.0	0.1	0.5	-1.0	-0.7	-1.4	-0.2	0.9

Table 2

Rings stimuli: Table of FA rates expressed as normal values for all subjects.

Conditions		**Subjects**											
Noise Type	*Reference*	CE	CY	JP	ML	TC	CS	TK	PC	VR	WS	LL	JL
Noise	Noise Alone	-0.0	0.2	0.8	0.5	0.2	3.2	1.1	0.6	1.6	0.0	0.4	0.3
	Completion	0.1	0.5	0.9	0.8	0.1	1.1	0.6	0.5	1.0	0.3	1.3	0.7
	Signal Copy	-0.1	0.1	0.4	0.4	-0.0	3.9	0.7	0.9	0.3	0.5	-0.2	0.1
Sym. Noise	Noise Alone	-0.0	0.1	0.5	0.3	0.2	1.9	0.8	0.4	1.1	0.8	0.7	0.1
	Completion	0.1	0.5	0.7	0.8	0.2	1.3	1.7	1.9	0.8	0.9	0.7	0.0
	Signal Copy	0.3	0.2	0.6	0.5	0.0	1.4	2.1	0.9	0.8	-0.0	0.0	-0.2

Table 3

Faces stimuli: Table of d' values for all subjects.

4 Discussion

4.1 Completion effect

We have presented a novel completion effect in which one half of a pattern appears to be completed when embedded in visual noise. The fact that we found a completion effect for these

Conditions		**Subjects**											
Noise Type	*Reference*	CE	CY	JP	ML	TC	CS	TK	PC	VR	WS	LL	JL
Noise	Noise Alone	0.2	0.3	1.0	0.4	0.4	0.9	0.9	1.3	2.6	-0.4	1.0	1.1
	Completion	0.2	0.7	1.3	0.7	0.7	1.6	1.4	-0.0	2.6	0.5	-0.5	1.8
	Signal Copy	0.1	0.2	1.0	1.1	0.2	-0.1	0.9	0.3	0.6	0.7	0.5	-0.2
Sym. Noise	Noise Alone	0.2	0.2	0.5	0.9	0.3	1.0	0.6	0.9	1.3	0.7	1.1	0.5
	Completion	1.4	0.8	2.0	1.7	0.8	1.2	0.0	1.5	0.5	0.0	0.6	1.9
	Signal Copy	0.2	0.2	0.9	1.0	0.1	0.3	0.8	0.7	1.5	0.7	1.0	2.2

Table 4

Rings stimuli: Table of d' values for all subjects.

stimuli supports the general idea that perceptual completion is the result of using prior knowledge to form an interpretation of ambiguous images with missing data. Although our data support a high-level implementation of the completion for our stimuli, there may be several ways that prior knowledge can be implemented, and other completion effects may differ in their implementation.

However, the noise completion effect fits the general pattern of completion stimuli: 1) due to missing or corrupted data the stimulus is ambiguous in the sense that the image information is not sufficient to form a unique interpretation; 2) the stimulus admits an explanation for the missing/corrupted data that is consistent with the observer's prior knowledge.

For our stimuli, the white noise satisfies both conditions. Condition 1 is satisfied because white noise can be considered as carrying infinitely many signals and hence is fundamentally ambiguous. Thus, for any sample of white noise, there will always be some evidence for the existence of the signal in the right half noise. However, the evidence for this signal is no better than the evidence for any other signal. On the other hand, the presence of noise constitutes positive evidence that the data has been corrupted/destroyed, hence satisfying condition 2.

For condition 2 to be completely satisfied, however, there must be some prior knowledge that the

left half signals normally occurs as part of a complete image. In this regard, the signal images we used can be considered as eliciting the use of two kinds of prior information indicating the left-half signal usually occurs with its left-half complement: knowledge of the complete signal's symmetry and a familiarity with the complete stimuli. We will discuss the use of symmetry and familiarity to promote completion below.

4.1.1 Symmetry

The term symmetry is generally used to describe the mapping of a pattern onto itself or onto other patterns (Tyler and Hardage, 1996), and hence connotes a characteristic redundancy of information. Of the possible symmetries, mirror symmetry plays a special role in vision. Mirror symmetry is ubiquitous in biological forms, and there is abundant evidence that visual systems make use of it in many ways, including discrimination and recognition (Delius and Nowak, 1979; Tyler and Hardage, 1996) and mate selection (Moeller, 1992; Swaddle and Cuthill, 1993), among others. In keeping with the notion of the biological relevance of mirror symmetry, there is abundant evidence that mirror symmetric images are specially processed by the human visual system (Wagemans, 1996). Ernst Mach (Mach, 1979), in one of the first experiments on symmetry perception, found that people can easily perceive bilateral symmetries of simple amorphous shapes. Since then, a large number of studies have looked at the effects of different factors on the speed and efficiency of symmetry detection. In general, symmetry detection is best for symmetry close to a vertical axis (Wagemans, 1996; Barlow and Reeves, 1979; Jenkins, 1982; Tyler et al., 1995).

In addition, symmetry has a long history as a potential factor in figural completion (Kofka, 1935; Dinnersten and Wertheimer, 1957; Barlow and Reeves, 1979; Kanizsa, 1979; Buffart and Leeuwenberg, 1981; Buffart et al., 1983; Sekuler, 1994), beginning with Gestalt theory, where symmetry was one of the factors postulated to affect the organizational Gestalt principles of closure and good continuation. The common thread in all of this work is that if symmetry is detected, the characteristic redundancies of symmetric patterns can be exploited to infer missing or corrupted information in the image. Because symmetry is efficiently detected (Barlow and

Reeves, 1979), it is plausible that detected symmetry can be used to automatically generate a completed figure.

4.1.2 Familiarity

Familiarity denotes the recognition that the image includes features that could belong to a familiar pattern. If the visual system fits object or surface models to the image data, then familiarity can bias the selection of which model best interprets the data . The reason for studying familiarity is that it makes clearer the distinction between completion that could arise due to generic image processing, independent of the objects in the scene, and completion that could arise from fitting familiar objects to a scene.

5 Theories of completion

The phenomena of completion forms strong empirical evidence that vision involves an inferential process. There are many ways the visual system might perform completion, but existing theories fall into three broad categories based on the kind of knowledge the visual system brings to bear on the completion process. One possibility is that high-level, domain specific knowledge is used to infer the image should be complete (Gregory, 1972; Rock and Anson, 1979; Durgin et al., 1995). Another broad possibility is that completion is the result of automatic processes in which the early visual system constructively "fills in" or interpolates the missing pieces based on generic principles (e.g. smoothing, or Gestalt mechanisms). The constraints for filling-in could be implemented using low-level image-based operations (Kellman and Shipley, 1992; Wouterlood and Boselie, 1992). Alternatively, the filling-in could involve an intermediate-level process that first indexes a candidate (complete) familiar pattern, object, or surface (Nakayama et al., 1989; Paradiso et al., 1989; Shimojo and Nakayama, 1990). In this section we investigate and quantify specific versions of these three completion schemes and compare their predictions to our results.

Low-level completion: Image-based data modification The most common account of completion is that it is the result of automatic image processing that is not tied to a model of the scene. The

basic idea is that there are a set of heuristics which the visual system applies to all incoming images which extract and represent information the organism needs. These heuristics, however, sometimes generate representations of the image data which have "filled-in" information not in the original image.

Heuristics used to explain phenomena like the filling-in of the blind-spot and neon color spreading involve competitions between neurons with inhibitory lateral connections (van Tuijl, 1975; Grossberg and Mingolla, 1985; Fiorani Junior et al., 1992; Grossberg, 1997). Heuristics for stimuli like the Kanisza triangle involve the construction of "illusory contour detectors" (von der Heydt et al., 1984), that embody simple rules for detecting line configurations which are likely to generate the percept of illusory contours. These rules are largely based on Gestalt principles like good continuation and symmetry, and on the knowledge of the kinds of line configurations (e.g. T-junction) which would result from occlusion (Ullman, 1976; Parent, 1989; Kellman and Shipley, 1992; Wouterlood and Boselie, 1992; Heitger, 1993; Williams and Hanson, 1994; Williams, 1994).

The basic idea is that the incoming data is automatically modified, and this modified data forms the substrate of a perception of completion.

Intermediate-level Model Fitting: Indexing a familiar global pattern

A second basic possibility is that completion is the by-product of the visual system fitting a global pattern, such as an object or surface model, to the image data. Such a process could involve an "indexing" stage similar to that in models of recognition. Evidence from the left-image is used to select (the indexing stage) a global template from a library of familiar patterns. This global template is then applied to the whole image. When there is a constellation of image features that are consistent with a common or high probability model, the mismatch is ignored and the visual system makes the inference that the complete surface/object must be there.

High-level: decision-based

Finally, completion could be a consequence of a high level process, like a perceptual judgement or high-level feedback that "fills-in" the missing information, but only *after* the decision as to

the pattern is made. In this view it is a perceptual inference that the completed object is in the scene that creates the percept of completion.

5.1 Signal detection models of completion

In this section we use a signal-detection approach to quantify and illustrate the predictions of several general classes of completion model and differentiate between them in light of our results.

The generative model for the task is simple:

$$\vec{I} = \vec{r} + \vec{n} \tag{1}$$

A fixed signal $\vec{r}$ is added to a sample of gaussian white noise to form an image $\vec{I}$. In the detection task, $\vec{r}$ has a 50% probability of being a signal in the right half image (e.g. face or rings) denoted by $\vec{s}$ and a 50% probability of being zero.

The likelihoods for the task are given by:

$$P(\vec{I}|\vec{r} = \vec{s}) = \prod_{i=1}^{N} \frac{1}{\sqrt{2\pi\sigma^2}} \exp(-\frac{(I_i - s_i - B_i)^2}{2\sigma^2}) \tag{2}$$

$$P(\vec{I}|\vec{r} = 0) = \prod_{i=1}^{N} \frac{1}{\sqrt{2\pi\sigma^2}} \exp(-\frac{(I_i - B_i)^2}{2\sigma^2}) \tag{3}$$

where $\vec{B}$ is the background luminance that is identical at each pixel, σ^2 is the variance of the image noise and N is the number of pixels in the image.

It is well-known that the log-likelihood ratio comparison is equivalent to comparing $\vec{s} \cdot \vec{I} = \sum_{i=1}^{N} s_i I_i$ with a fixed criterion T. This is called the *matched template* observer. To compute the sensitivity (d') and false-alarm rates of the matched filter observer, we need to compute the distribution of $\vec{s} \cdot \vec{I}$ assuming $\vec{I}$ is generated by either by the signal plus noise or the noise alone. Because $\vec{I}$ is gaussian distributed and linear combinations of gaussian variables are gaussian distributed, $\vec{s} \cdot \vec{I}$ is gaussian distributed and hence is completely described by its mean and variance (see Appendix).

Because it is the simplest, we first describe the high-level decision-based model.

5.1.1 Model 1: High-level: Decision-based

We are interested in modelling an observer that performs perceptual completion via a high-level decision event. In other words, completion results purely from the *decision* that something should be in the right half given the presence of the left-half of the image. Because this model involves only a perceptual decision, we can model the observer completely through the decision threshold criterion T and assume that data processing is done optimally via the matched template.

One of the simplest assumptions is that the presence of the left-half image effectively increases the observer's prior probability for the signal being present by some factor $c_L (c_L > 1)$: $P(signal|\text{left-half}) = P(signal)c_L$. Thus

$$T_L = \frac{\mu_s - \mu_n}{2} + \log(\frac{P(noise)/c_L}{P(signal)c_L}) = T - 2\log(c_L)$$

Predictions

Under these conditions, it is shown in the appendix:

$$\Delta d' = 0 \tag{4}$$

$$\Delta FA = \Phi^{-1}(F_L(T)) - \Phi^{-1}(F(T)) = \frac{2\log(c_L)}{\sqrt{\sigma^2 \vec{s} \cdot \vec{s}}} \tag{5}$$

where FA is the false-alarm rate, and $\Phi()$ is the inverse cumulative gaussian, and σ^2 is the variance of the background additive noise. Thus we expect no change in d' with this model, but an increase in the false-alarm rate.

5.1.2 Model 2: Intermediate-level, Full-template model-fitting

One simple model of completion is that the presence of the left-half image triggers the observer's use of a complete template to detect the right-half signal. In the appendix we show how for this model increases in false-alarm rate are accompanied by decreases in sensitivity.

In this model, we assume that the observer in the presence of the left-half is a matched template

observer with its template composed of both the right-half signal image $\vec{s}$ and the left-half reflected image $\vec{l}$ to form the template $\vec{t_f} = \vec{s} + \vec{l}$. We assume that this observer uses a fixed threshold T for all its decisions, in order to isolate the effect of using the full-template on false-alarm rates and sensitivity. Thus the full-template observer computes $x = \vec{t_f} \cdot \vec{I}$ and decides signal or noise by comparing x to T. We can determine the properties of this observer by computing the mean and variance of x on the signal and noise conditions.

Predictions

In the appendix, we show that if we assume $T = 2\vec{s} \cdot \vec{B} + \vec{s} \cdot \vec{s}/2$ in the presence of the left-half, then we can derive a simple formula for the change in d' as a function of the change in false-alarm rate.

This model predicts:

$$\Delta FA = \frac{2\log(T)}{\sqrt{\sigma^2 \vec{s} \cdot \vec{s}}} \tag{6}$$

$$\Delta d' = -\sqrt{2}\Delta FA \tag{7}$$

Thus we expect an increase in the false-alarm rate in the presence accompanied by a decrease in d' with this model.

5.1.3 Model 3: Low-Level, Image-based Data Modification

A third basic possibility is that completion results from a low-level process that automatically completes or fills-in the right-hand side based on the left-half image. Many such strategies have been proposed () that differ widely in terms of implementation. The essence of the idea, however, is that the information on the left-hand side is used to complete or fill-in the 'missing' information on the right-hand side. A reasonable, simple, and mathematically tractable model of such a process assumes the completion process effectively adds a reflected copy of the left hand image to the right hand image to form a new internal representation of the image:

$$\vec{I_n} = \vec{I} + bR(\vec{I_l} - \vec{B})$$

where $R()$ is the reflection function, I_l is the left subimage minus the background luminance and b is a scale factor.

Again, in the presence of the left-half, we assume that the observer is a matched template observer, and that this observer uses a threshold $T = \vec{s} \cdot \vec{B} + \vec{s} \cdot \vec{s}/2$ for all its decisions. Thus the full-template observer computes $x = \vec{s} \cdot \vec{I_n}$ and decides signal or noise by comparing x to T. We can determine the properties of this observer by computing the mean and variance of x on the signal and noise conditions.

This model predicts:

$$\Delta F = \left(\frac{b - \frac{1}{2}}{\sqrt{b^2 + 1}} + \frac{1}{2} \right) \frac{\sqrt{\vec{s} \cdot \vec{s}}}{\sigma} \tag{8}$$

$$\Delta d' = \frac{2(1 - \sqrt{b^2 + 1})}{\sqrt{b^2 + 1} + 2b - 1} \Delta F \tag{9}$$

Because $b > 0$, we expect an increase in the false-alarm rate in the presence of the left half image accompanied by a decrease in d' with this model. In particular, if $b = 1$ (copied without attenuation), then:

$$\Delta d' = -0.34 \Delta F$$

This value is used in figure 6.

5.1.4 Comparing predictions to data

The three models of completion make simple predictions about how sensitivity changes with increases in the false alarm rate. In figure 6 we plot the change in d' between the 'noise alone' and 'completion' conditions as a function of the change in false-alarm rate.

Of the three models, the best overall fit is generated by the high-level model which predicts no change in d'. The $\Delta d'$ for all subjects are clustered around zero except one (CE) for face stimuli and two (CS,WS) for rings. Although 4 subjects appear to match the low-level (dashed-line) predictions for faces, none of these subjects show a similar agreement for rings. Thus, parsimony suggests completion is the result of a high-level decision procedure.

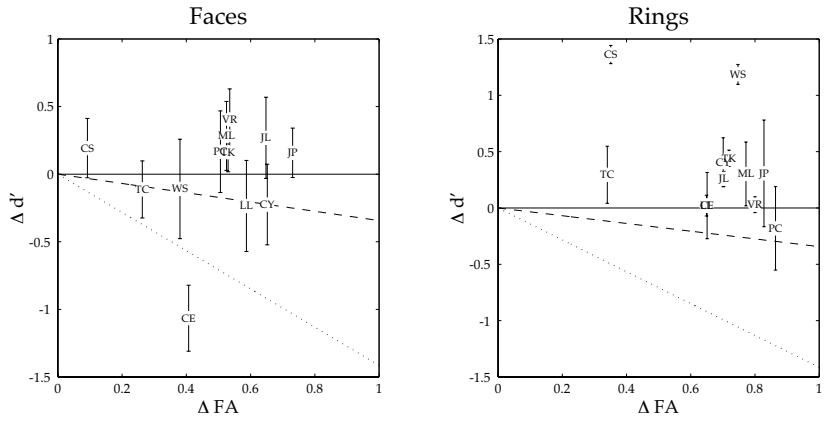

Fig. 6. Change in d' between the 'completion' and 'noise alone' conditions is plotted as a function of the change in false-alarm rate for the 12 subjects for faces (left) and rings (right). Subject labels are used as data points and the error bars represent bootstrap estimates of the standard errors of the estimate. The solid line represents the prediction of the high level model (model 1), the dotted line represents the complete template model (model 2), and the dashed line represents the low-level (model 3) predictions.

However, for faces and especially for rings, $\Delta d'$ is greater than zero for several subjects, whereas all three predict models predict no change or decreases in d' with increases in ΔFA. What could produce an increase in d'?

5.1.5 Spatial Uncertainty

One of the simplest ways that $\Delta d'$ could be greater than zero is that the observer has substantial spatial uncertainty about the location of the signal in the image. Spatial uncertainty would be expected to reduce sensitivity in the 'noise alone' over the 'completion' condition because in the 'completion' condition the left-half forms a visual landmark that specifies where the signal will spatially appear.

If spatial uncertainty is driving $\Delta d'$ toward values greater than zero, then a real decrease in $\Delta d'$ with increases in ΔFA could be present but masked, which confounds our conclusion that the completion effect is likely caused by a high-level decision.

If real decreases in $\Delta d'$ are being masked by spatial uncertainty in the 'noise alone' condition, then we would expect that observers with the least spatial uncertainty should show the largest

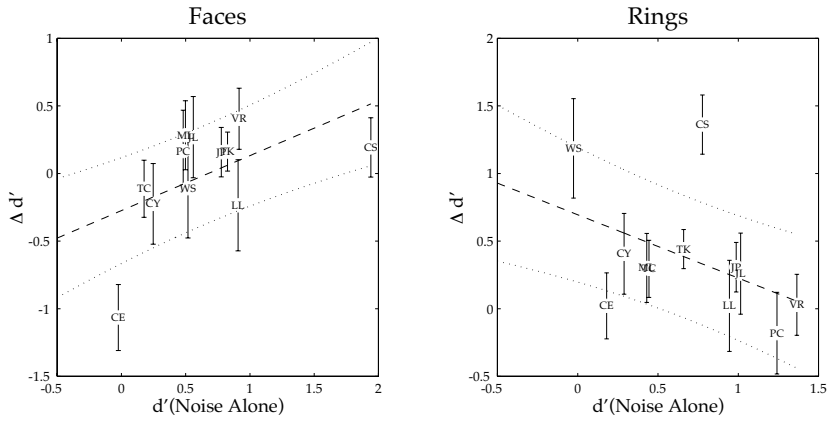

Fig. 7. Change in d' between the 'completion' and 'noise alone' conditions is plotted as a function of 'noise alone' d' for the 12 subjects for faces (left) and rings (right). Subject labels are used as data points and the error bars represent bootstrap estimates of the standard errors of the estimate. The dashed line represents the best straight line fit and dotted lines represent the 50% confidence intervals of the straight line fit.

decrease in $\Delta d'$. However, observers with the least spatial uncertainty should have the highest d' in the 'noise alone' condition. Thus, if this masking is occurring, we would expect to see a negative trend in a plot of $\Delta d'$ vs. 'noise alone' d', which is shown in figure **??**. Instead we observe no significant trends, suggesting that any masking should not affect our conclusions.

5.2 Symmetry and processing

While there is abundant evidence for the special processing and early detection of bilateral symmetry, we found little effect of noise symmetry on sensitivity. We found this result somewhat surprising because there are many ways in which this symmetry could have been used that would alter sensitivity.

In fact, error-free performance can be achieved on the symmetric noise trials. This is due to the fact that given one part of the image and the knowledge that the image is symmetric, the other part can be produced by copying and reflecting. Thus the symmetric noise can be entirely removed by subtracting off the reflected left-hand image, to perfectly reveal the presence or absence of the signal (e.g. by setting $b = -1$ in Model 3).

In general there are two basic ways that symmetry can be used to improve visual performance and efficiency, by allowing more efficient codes and (as mentioned above) providing a means to do noise removal.

5.2.1 Efficient Coding

The redundancy in symmetric images could be used to construct an efficient coding scheme, such that the redundant data is discarded. The notion that symmetry may be used for efficient encoding was put forward by Barlow (Barlow and Reeves, 1979). In the context of our experiments, such encoding should decrease the salience of a signal added to the symmetric pattern. If symmetry is used to compress the encoding of the image, then a small signal added to the symmetric noise would also get compressed, and hence less detectable than a signal added to non-symmetric noise. Thus, this coding idea predicts a decrease in d' in the presence of symmetric noise which is inconsistent with our results.

5.2.2 Denoising

Alternately, given the presence of symmetry, the two halves of the image can be compared for discrepancies. This comparison can be used for denoising corrupted images or the completion of occluded images. Because the two halves of a symmetric noise image are reflected copies of each other, a signal buried in one half can be perfectly detected by reflecting and subtracting one side of the noise image from the other. Thus an observer that uses symmetry to denoise an image should have substantially better performance detecting a signal added to symmetric noise than a signal added to non-symmetric noise. Again, this prediction is not consistent with our results.

5.2.3 Filling in missing/occluded data

On the other hand symmetry can be used to fill in missing data. In the presence of symmetry, unaffected areas of the image can be copied into areas that have missing or occluded data. In the context of the experiment, the presence of the symmetric noise could cause the copying of the left hand image noise into the right, obscuring the signal and decreasing sensitivity. This idea is

not supported by our results.

5.2.4 Changing priors

Another possible use for symmetry information is to adjust prior expectations and/or uncertainty about the objects in the scene. The presence of symmetry in an image constitutes a substantial restriction on the kind of scene that could have caused the image. Simply stated, we may want to detect symmetry because the probability of a symmetric image given a relevant biological cause (e.g. predator or prey) is high compared the probability of a symmetric image given a generic cause. Thus the presence of symmetry can be used to indicate the likely presence of a biological form in the scene. In the context of the experiment, the presence of symmetry could increase the prior expectation of observing a symmetric image, which in the case of the face and rings would mean a complete image. This possibility can produce increases in false-alarm rates without changing sensitivity, consistent with our results.

A similar conclusion has been reached by Vetter et al. ((Vetter et al., 1994)). They found that human observers can use a priori information about symmetry to recognize 3-D objects. From an accidental view of a bilateral symmetric object, humans can recognize novel views. They explained this effect as a reduction in the required number of model views for the recognition of symmetrical objects.

6 Conclusions

We have investigated a novel completion phenomenon in which partial patterns are perceptually completed in visual noise. The phenomenon is consistent with the general idea that completion is the result of using symmetry to find a viable interpretation of ambiguous image data, irregardless of whether completion is implemented by high-level global processes or low-level local processes. However, signal detection theory analysis took us one step further in suggesting that the completion is not due to low-level image-based processes, nor due to intermediate-level model-fitting processes. Rather, it may be best explained by a change in an observer's prior probability of a

whole pattern being present. This change in prior affects the decision criterion, and thus the false alarm rate for reporting whole patterns. In contrast to most signal detection tasks, the change in prior probability is manifest perceptually–rather than as a pure response bias.

7 Conclusions

References

[1] Barlow, H. B., Reeves, B. C., 1979. The versatility and absolute efficiency of detecting, mirror symmetry in random dot displays. Vision Research 19, 783–793.

[2] Buffart, H., Leeuwenberg, E., 1981. Coding theory of visual pattern completion. J Exp Psychol Hum Percept Perform 7 (2), 241–74.

[3] Buffart, H., Leeuwenberg, E., Restle, F., 1983. Analysis of ambiguity in visual pattern completion. J Exp Psychol Hum Percept Perform 9 (6), 980–1000.

[4] Burgess, A. E., Ghandeharian, H., 1984. Visual signal detection. ii. signal-location identification. J Opt Soc Am [A] 1 (8), 906–10.

[5] Delius, J. D., Nowak, B., 1979. Visual symmetry recognition in pigeons. Psychol. Res. 44, 199–212.

[6] Dennett, D., 1992. "filling in" versus finding out: a ubiquitous confusion in cognitive science. In: H.L. Pick, Jr., P. v. d. B. . D. K. (Ed.), Cognition: Conceptual and Methodological Issues. American Psychological Association, Washington, D.C.

[7] Dinnersten, D., Wertheimer, M., 1957. Some determinants of phenomenal overlapping. American Journal of Psychology 70, 21–37.

[8] Durgin, F. H., Tripathy, S. P., Levi, D. M., 1995. On the filling in of the visual blind spot: some rules of thumb. Perception 24 (7), 827–40, using Smart Source Parsing.

[9] Efron, B., Tibshirani, R. J., 1993. An introduction to the bootstrap. Chapman & Hall, New York, NY.

[10] Fiorani Junior, M., Rosa, M. G., Gattass, R., Rocha-Miranda, C. E., 1992. Dynamic surrounds of receptive fields in primate striate cortex: a physiological basis for perceptual completion?

Proc Natl Acad Sci U S A 89 (18), 8547–51.

[11] Green, D. M., Swets, J., 1974. Signal detection theory and psychophysics. Wiley, New York.

[12] Gregory, R., 1972. Cognitive contours. Nature 238, 51–52.

[13] Grossberg, S., 1997. Cortical dynamics of three-dimensional figure-ground perception of two-dimensional pictures. Psychol Rev 104 (3), 618–58.

[14] Grossberg, S., Mingolla, E., 1985. Neural dynamics of form perception: boundary completion, illusory figures, and neon color spreading. Psychol Rev 92 (2), 173–211.

[15] Heitger, R. & von der Heydt, R., 1993. A computational model of neural contour processing, figure-ground segregation and illusory contours. In: Internal Conference Computer Vision , 32–40.

[16] Jenkins, B., 1982. Redundancy in the perception of bilateral symmetry in dot textures. Percept Psychophys 32 (2), 171–7.

[17] Kanizsa, G., 1979. Organization in Vision. Praeger, New York.

[18] Kellman, P. J., Shipley, T. F., 1991. A theory of visual interpolation in object perception. Cognit Psychol 23 (2), 141–221.

[19] Kellman, P. J., Shipley, T. F., 1992. A theory of visual interpolation in object perception. Cognitive Psychology 23, 141–221.

[20] Kersten, D., 1990. Statistical limits to image understanding. In: Blakemore, C. (Ed.), Vision: Coding and Efficiency. Cambridge University Press, Cambridge, UK, pp. 32–44, chapter 3.

[21] Kersten, D., Schrater, P., 2000. Pattern inference theory: A probabilistic approach to vision. In: Mausfeld, R., Heyer, D. (Eds.), Perception and the Physical World. John Wiley & Sons, Ltd., Chichester.

[22] Kofka, K., 1935. Principles of Gestalt Psychology. Harcourt, New York.

[23] Lesher, G., 1995. Illusory contours: toward a neurally based perceptual theory. Psychonomic Bulletin & Review 2, 279–321.

[24] Mach, E., 1979. Analysis of the Sensations. Open Court Publishing House, Chicago.

[25] Moeller, A. P., 1992. Female swallow preference for symmetrical male sexual ornaments. Nature 357, 238–240.

[26] Nakayama, K., Shimojo, S., 1992. Experiencing and perceiving visual surfaces. Science 257, 1357–

1363.

[27] Nakayama, K., Shimojo, S., Silverman, 1989. Stereoscopic depth: Its relation to image fragmentation, grouping and the recognition of occluded objects. Perception 18, 55–68.

[28] Paradiso, M., Shimojo, S., Nakayama, K., 1989. Subjective contours, tilt-aftereffects, and visual cortical organization. Vision Research 29, 1205–1213.

[29] Parent, P. & Zucker, S., 1989. Trace inference, curvature consistency and curve detection. IEEE Transactions on Pattern Analysis and Machine Intelligence 11, 823–839.

[30] Pessoa, L., Thompson, E., Noe, A., 1998. Finding out about filling-in: a guide to perceptual completion for visual science and the philosophy of perception. Behav Brain Sci. 21, 723–48.

[31] Rock, I., Anson, R., 1979. Illusory contours as a solution to a problem. Perception 8, 665–681.

[32] Sekuler, A. B., 1994. Local and global minima in visual completion: effects of symmetry and orientation. Perception 23, 529–545.

[33] Sekuler, A. B., Palmer, S. E., Flynn, C., 1994. Local and global processes in visual completion. Psychological Science 5, 260–267.

[34] Shimojo, S., Nakayama, K., 1990. Amodal representation of occluded surfaces: role of invisible stimuli in apparent motion correspondence. Perception 19 (3), 285–99, using Smart Source Parsing.

[35] Spillman, L., Dresp, B., 1995. Phenomena of illusory form: can we bridge the gap between levels of explanation? Perception 24, 1333–1364.

[36] Swaddle, J. P., Cuthill, I. C., 1993. Preference for symmetric males by female zebra finches. Nature 367, 165–166.

[37] Tse, P. U., 1999a. Complete mergeability and amodal completion. Acta Psychol (Amst) 102 (2-3), 165–201.

[38] Tse, P. U., 1999b. Volume completion. Cognit Psychol 39 (1), 37–68.

[39] Tyler, C., Hardage, L., 1996. Mirror symmetry detection: predominance of second-order pattern processing throughout the visual field. In: Tyler, C. (Ed.), Human Symmetry Perception. Springer-Verlag, New York, pp. 157–171.

[40] Tyler, C., Hardage, L., Miller, R. T., 1995. Multiple mechanisms for the detection of mirror symmetry. Spatial Vision 9, 79–100.

[41] Ullman, S., 1976. Filling-in the gaps: the shape of subjective con-tours and a model for their generation. Biological Cybernetics 25, 1–6.

[42] van Lier, R., 1999. Investigating global effects in visual occlusion: from a partly occluded square to the back of a tree-trunk. Acta Psychol (Amst) 102 (2-3), 203–20.

[43] van Tuijl, H., 1975. A new visual illusion: neonlike color spreading and complementary color induction between subjective contours. Acta Psychologica 39, 441–445.

[44] Vetter, T., Poggio, T., Bulthoff, H. H., 1994. The importance of symmetry and virtual views in three-dimensional object recognition. Curr Biol 4 (1), 18–23.

[45] von der Heydt, R., Peterhans, E., Baumgartner, G., 1984. Illusory contours and cortical neuron responses. Science 224 (4654), 1260–2.

[46] Wagemans, J., 1996. The role of pattern outline in bilateral symmetry detection with briefly flashed dot patterns. Spatial Vision 9, 9–32.

[47] Wertheimer, M., 1923. Untersuchungen zur lehre von der gestalt,ii. Psychol. Forsch. 4, 301–350.

[48] Williams, L. R., 1994. Perceptual completion of occluded surfaces. Ph.D. thesis, University of Massachusetts, Amherst, Mass.

[49] Williams, L. R., Hanson, A. R., Jun. 1994. Perceptual completion of occluded surfaces. In: Proceedings of the Conference on Computer Vision and Pattern Recognition. IEEE Computer Society Press, Los Alamitos, CA, USA.

[50] Wouterlood, D., Boselie, F., 1992. A good continuation model of some occlusion phenomena. Psychological Research 54, 267–277.

[51] Yin, C., Kellman, P. J., Shipley, T. F., 1997. Surface completion complements boundary interpolation in the visual integration of partly occluded objects. Perception 26 (11), 1459–79, using Smart Source Parsing.

[52] Yuille, A. L., Bulthoff, H. H., 1996. Bayesian decision theory and psychophysics. In: D.C., K., W., R. (Eds.), Perception as Bayesian Inference. Cambridge University Press, Cambridge, U.K., competitive priors.

Appendix

The ideal observer for the detection task is the simple and well-known matched template observer).

The generative model for the task is simple:

$$\vec{I} = \vec{r} + \vec{n} \tag{10}$$

A fixed signal $\vec{r}$ is added to a sample of gaussian white noise to form an image $\vec{I}$. In the detection task, $\vec{r}$ has a 50% probability of being a signal in the right half image (e.g. face or rings) denoted by $\vec{s}$ and a 50% probability of being zero.

The likelihoods for the task are given by:

$$P(\vec{I}|\vec{r}=\vec{s}) = \prod_{i=1}^{N} \frac{1}{\sqrt{2\pi\sigma^2}} \exp(-\frac{(I_i - s_i - B_i)^2}{2\sigma^2}) \tag{11}$$

$$P(\vec{I}|\vec{r}=0) = \prod_{i=1}^{N} \frac{1}{\sqrt{2\pi\sigma^2}} \exp(-\frac{(I_i - B_i)^2}{2\sigma^2}) \tag{12}$$

where $\vec{B}$ is the background luminance that is identical at each pixel, σ^2 is the variance of the image noise and N is the number of pixels in the image.

The unbiased Bayes optimal decision is to choose signal present ($\vec{r} = \vec{s}$) is present if $\frac{P(\vec{I}|\vec{r}=\vec{s})}{P(\vec{I}|\vec{r}=0)} > \frac{P(\vec{r}=0)}{P(\vec{r}=\vec{s})}$ and otherwise choose no signal. This decision rule is invariant to monotonic transforms like the natural logarithm, hence the decision rule can be written: $\log(\frac{P(\vec{I}|\vec{r}=\vec{s})}{P(\vec{I}|\vec{r}=0)}) > T$ where $T = \log(\frac{P(\vec{r}=0)}{P(\vec{r}=\vec{s})})$.

It is well-known that the log-likelihood ratio comparison is equivalent to comparing $\vec{s} \cdot \vec{I} = \sum_{i=1}^{N} s_i I_i$ with a fixed criterion T. This is called the *matched template* observer. To compute the performance and false-alarm rates of the matched filter observer, we need to compute the distribution of $\vec{s} \cdot \vec{I}$ assuming $\vec{I}$ is generated by either by the signal plus noise or the noise alone. Because $\vec{I}$ is gaussian distributed and linear combinations of gaussian variables are gaussian distributed, $\vec{s} \cdot \vec{I}$ is gaussian distributed and hence is completely discribed by its mean and

variance. Let $x = \vec{s} \cdot \vec{I}$. Then:

$$P(x|signal) = \frac{\exp(-\frac{(x-\mu_s))^2}{2\sigma_n^2})}{\sqrt{2\pi\sigma_n^2}} = \frac{1}{\sqrt{2\pi\sigma^2\vec{s}\cdot\vec{s}}}\exp(-\frac{(x-\vec{s}\cdot(\vec{s}+\vec{B}))^2}{2\sigma^2\vec{s}\cdot\vec{s}}) \tag{13}$$

$$P(x|noise) = \frac{\exp(-\frac{(x-\mu_n))^2}{2\sigma_n^2})}{\sqrt{2\pi\sigma_n^2}} = \frac{1}{\sqrt{2\pi\sigma^2\vec{s}\cdot\vec{s}}}\exp(-\frac{(x-\vec{s}\cdot\vec{B})^2}{2\sigma^2\vec{s}\cdot\vec{s}}) \tag{14}$$

The threshold T is given by $T = \frac{\mu_s - \mu_n}{2} + \log(\frac{P(noise)}{P(signal)})$.

The formula for d' and the false-alarm rate are easily derived from the expressions:

$$d' = \frac{\mu_s - \mu_n}{\sigma_n} = \frac{\sqrt{\vec{s}\cdot\vec{s}}}{\sigma} \tag{15}$$

$$F(T) = \int_T^\infty \frac{\exp(-\frac{(x-\mu_n))^2}{2\sigma_n^2})}{\sqrt{2\pi\sigma_n^2}} = \Phi(\frac{\mu_n - T}{\sigma_n}) = \Phi(\frac{\vec{s}\cdot\vec{B} - T}{\sqrt{\sigma^2\vec{s}\cdot\vec{s}}}) \tag{16}$$

where Φ is the cumulative normal function.

7.1 Model 1: High Level-Decision based

In this model we assume that the presence of the left-half image increases the prior probability for the signal being present by some factor c_L: $P(signal|\text{left-half}) = P(signal)c_L$. Thus

$$T_L = \frac{\mu_s - \mu_n}{2} + \log(\frac{P(noise)/c_L}{P(signal)c_L}) = T - 2\log(c_L)$$

No Left-Half

In the no left-half condition:

$$d' = \frac{\sqrt{\vec{s}\cdot\vec{s}}}{\sigma} \tag{17}$$

$$F(T) = \Phi(\frac{\vec{s}\cdot\vec{B} - T}{\sqrt{\sigma^2\vec{s}\cdot\vec{s}}}) \tag{18}$$

Left-Half present

In the left-half present condition the template ignores the left-half, thus:

$$d'_L = \frac{\sqrt{\vec{s} \cdot \vec{s}}}{\sigma} \tag{19}$$

$$F_L(T) = \Phi(\frac{\vec{s} \cdot \vec{B} - T_L}{\sqrt{\sigma^2 \vec{s} \cdot \vec{s}}}) \tag{20}$$

Predictions

This model predicts:

$$d'_L - d' = 0 \tag{21}$$

$$\Phi^{-1}(F_L(T)) - \Phi^{-1}(F(T)) = \frac{2\log(c_L)}{\sqrt{\sigma^2 \vec{s} \cdot \vec{s}}} \tag{22}$$

Thus we expect no change in d' with this model, but an increase in the false-alarm rate.

7.2 Model 2: Full-template model

In this model, we assume that the observer in the presence of the left-half is a matched template observer with its template is composed of both the right-half signal image $\vec{s}$ and the left-half reflected image $\vec{l}$ to form the template $\vec{t_f} = \vec{s} + \vec{l}$. We assume that this observer uses a fixed threshold T for all its decisions, in order to isolate the effect of using the full-template on false-alarm rates and performance. Thus the full-template observer computes $x = \vec{t_f} \cdot \vec{I}$ and decides signal or noise by comparing x to T. We can determine the properties of this observer by computing the mean and variance of x on the signal and noise conditions.

No Left-Half

In the no left-half condition, $x = \vec{s} \cdot \vec{I}$ is gaussian distributed with means and variance as before. Thus:

$$d' = \frac{\vec{s} \cdot \vec{s}}{\sigma} \tag{23}$$

$$F(T) = \Phi(\frac{\vec{s} \cdot \vec{B} - T}{\sqrt{\sigma^2 \vec{s} \cdot \vec{s}}}) \tag{24}$$

Left-Half present

Similarly, in the left-half present condition:

$$\mu_x|signal = \vec{t_f} \cdot (\vec{t_f} + \vec{B}) = (\vec{l} + \vec{s}) \cdot (\vec{l} + \vec{s} + \vec{B}) = 2\vec{s} \cdot \vec{s} + 2\vec{s} \cdot \vec{B} \tag{25}$$

$$\mu_x|noise = \vec{t_f} \cdot (\vec{l} + \vec{B}) = \vec{l} \cdot \vec{l} + (\vec{l} + \vec{s}) \cdot \vec{B} = \vec{s} \cdot \vec{s} + 2\vec{s} \cdot \vec{B} \tag{26}$$

$$\sigma_x^2 = \vec{t_f} \cdot \vec{t_f}\sigma^2 = (\vec{l} + \vec{s}) \cdot (\vec{l} + \vec{s})\sigma^2 = 2\vec{s} \cdot \vec{s}\sigma^2 \tag{27}$$

where we have used $\vec{s} \cdot \vec{l} = 0$, $\vec{l} \cdot \vec{l} = \vec{s} \cdot \vec{s}$, and $\vec{l} \cdot \vec{B} = \vec{s} \cdot \vec{B}$.

Thus:

$$d'_L = \frac{\sqrt{\vec{s} \cdot \vec{s}}}{\sqrt{2}\sigma} \tag{28}$$

$$F_L(T) = \Phi(\frac{\vec{s} \cdot \vec{s} + 2\vec{s} \cdot \vec{B} - T}{\sqrt{2\sigma^2 \vec{s} \cdot \vec{s}}}) \tag{29}$$

Predictions

If we assume $T = 2\vec{s} \cdot \vec{B} + \vec{s} \cdot \vec{s}/2$ in the presence of the left-half, then we can derive a simple formula for the change in d' as a function of the change in false-alarm rate.

This model predicts:

$$\Delta F = \Phi^{-1}(F_L(T)) - \Phi^{-1}(F(T)) = \frac{(1 + \sqrt{2})\sqrt{\vec{s} \cdot \vec{s}}}{2\sqrt{2}\sigma} \tag{30}$$

$$d'_L - d' = \frac{(1 - \sqrt{2})\sqrt{\vec{s} \cdot \vec{s}}}{\sqrt{2}\sigma} = -\sqrt{2}\Delta F \tag{31}$$

Thus we expect a decrease in d' with this model, but an increase in the false-alarm rate in the presence of the left half image.

7.3 Model 3: Low-Level (Data Modification)

In this model, a reflected copy of the left hand image is added to the right hand image to form a new internal representation of the image:

$$\vec{I_n} = \vec{I} + bR(\vec{I_l} - \vec{B})$$

where $R()$ is the reflection function, I_l is the left subimage, $\vec{B}$ is the background luminance and b is a scale factor.

Again, in the presence of the left-half, we assume that the observer is a matched template observer, and that this observer uses a threshold $T = \vec{s} \cdot \vec{B} + \vec{s} \cdot \vec{s}/2$ for all its decisions. Thus the full-template observer computes $x = \vec{s} \cdot \vec{I_n}$ and decides signal or noise by comparing x to T. We can determine the properties of this observer by computing the mean and variance of x on the signal and noise conditions.

No Left-Half

In the no left-half condition, $x = \vec{s} \cdot \vec{I_n}$ is gaussian distributed with means and variance as before. Thus:

$$d' = \frac{\vec{s} \cdot \vec{s}}{\sigma} \tag{32}$$

$$F(T) = \Phi(\frac{\vec{s} \cdot \vec{B} - T}{\sqrt{\sigma^2 \vec{s} \cdot \vec{s}}}) \tag{33}$$

Left-Half present

Similarly, in the left-half present condition:

$$\mu_x|signal = \vec{s} \cdot ((b+1)\vec{s} + \vec{B}) \tag{34}$$

$$\mu_x|noise = \vec{s} \cdot (b\vec{s} + \vec{B}) \tag{35}$$

$$\sigma_x^2 = (b^2+1)\vec{s} \cdot \vec{s}\sigma^2 \tag{36}$$

Thus:

$$d'_L = \frac{\sqrt{\vec{s}\cdot\vec{s}}}{\sqrt{b^2+1}\sigma} \tag{37}$$

$$F_L(T) = \Phi(\frac{\vec{s}\cdot\vec{s}+2\vec{s}\cdot\vec{B}-T}{\sqrt{2\sigma^2\vec{s}\cdot\vec{s}}}) \tag{38}$$

Predictions

This model predicts:

$$\Delta F = \Phi^{-1}(F_L(T)) - \Phi^{-1}(F(T)) = \left(\frac{b-\frac{1}{2}}{\sqrt{b^2+1}}+\frac{1}{2}\right)\frac{\sqrt{\vec{s}\cdot\vec{s}}}{\sigma} \tag{39}$$

$$d'_L - d' = \left(\frac{1}{\sqrt{b^2+1}}-1\right)\frac{\sqrt{\vec{s}\cdot\vec{s}}}{\sigma} = \frac{2(1-\sqrt{b^2+1})}{\sqrt{b^2+1}+2b-1}\Delta F \tag{40}$$

Because $b > 0$, we expect a decrease in d' with this model, but an increase in the false-alarm rate in the presence of the left half image.