Vision Research 126 (2016) 220–231

Contents lists available at ScienceDirect

Vision Research

journal homepage: www.elsevier.com/locate/visres

The role of shape complexity in the detection of closed contours

John Wilder [a,*], Jacob Feldman [b], Manish Singh [b]

[a] Department of Computer Science, University of Toronto, Toronto, Canada
[b] Department of Psychology, Center for Cognitive Science, Rutgers University – New Brunswick, USA

ARTICLE INFO

Article history:
Received 2 February 2015
Received in revised form 15 October 2015
Accepted 17 October 2015
Available online 7 November 2015

Keywords:
Shape perception
Contour detection
Gestalt perception
Bayesian inference
Information theory
Pattern detection

ABSTRACT

The detection of contours in noise has been extensively studied, but the detection of *closed* contours, such as the boundaries of whole objects, has received relatively little attention. Closed contours pose substantial challenges not present in the simple (open) case, because they form the outlines of whole shapes and thus take on a range of potentially important configural properties. In this paper we consider the detection of closed contours in noise as a probabilistic decision problem. Previous work on open contours suggests that contour complexity, quantified as the negative log probability (Description Length, DL) of the contour under a suitably chosen statistical model, impairs contour detectability; more complex (statistically surprising) contours are harder to detect. In this study we extended this result to closed contours, developing a suitable probabilistic model of whole shapes that gives rise to several distinct though interrelated measures of shape complexity. We asked subjects to detect either natural shapes (Exp. 1) or experimentally manipulated shapes (Exp. 2) embedded in noise fields. We found systematic effects of global shape complexity on detection performance, demonstrating how aspects of global shape and form influence the basic process of object detection.

© 2015 Elsevier Ltd. All rights reserved.

1. Introduction

A central function of the visual system is the ability to segregate visual objects from cluttered backgrounds. A simple laboratory approximation to this natural task is contour detection, in which subjects are asked to detect relatively collinear chains of oriented elements amid fields of randomly oriented elements. The large literature on this task (Field, Hayes, & Hess, 1993; Hess & Field, 1995; Pettet, McKee, & Grzywacz, 1998; Hess & Field, 1999; Field, Hayes, & Hess, 2000; Li & Gilbert, 2002; Geisler, Perry, Super, & Gallogly, 2001; Yuille, Fang, Schrater, & Kersten, 2004; Wilder, Singh, & Feldman, 2015) has shed light on the mechanisms underlying what the Gestaltists called "good continuation" (meaning grouping into smoothly elongated patterns), and has proved revealing about the organization of primary visual cortex. Because this literature has primarily focused on basic processes of orientation selectivity and contour integration, most studies have used targets consisting of simple open contours ("open" meaning that they lack self-intersections, i.e. do not loop).

By comparison, the detection of closed contours (contours that meet at their ends) has been studied relatively rarely (e.g. Pizlo, Salach-Golyska, & Rosenfeld, 1997). A closed contour defines not only the contour itself but also a bounded interior region (Koffka, 1935). As a result, closed contours pose a number of theoretical and experimental problems not present with open ones. The shape of the enclosed region can be parameterized in a multitude of ways, leaving it extremely unclear exactly what shape properties might be important in the process of detection. It is unclear, for example, whether global aspects of the shape of the enclosed region play a role in detection above and beyond the local properties of the bounding contour. More broadly, it is unclear whether known principles of open contour detection can be simply extended to the closed case, or whether new principles specific to whole forms will have to be introduced.

Contour closure is known to convey a detection advantage above and beyond that of the combined local relations among constituent contour elements (Kovacs & Julesz, 1993; Braun, 1999; Mathes & Fahle, 2007). However (Tversky, Geisler, & Perry, 2004) argued that the detection advantage for closed contours can simply be explained via probability summation (the transitivity rule of Geisler & Super (2000)). It is not clear, however, how transitivity would explain other benefits conferred on closed contours. For example, judgments about the aspect ratio of a contour are more accurate when the contour is closed than when it is open (Saarinen & Levi, 1999). Closed contours are processed more efficiently than open ones (Elder & Zucker, 1993), which gives rise to an advantage in remembering and recognizing them (Garrigan,

* Corresponding author.
E-mail address: jdwilder@cs.toronto.edu (J. Wilder).

http://dx.doi.org/10.1016/j.visres.2015.10.011
0042-6989/© 2015 Elsevier Ltd. All rights reserved.

2012). Altmann, Bülthoff, and Kourtzi (2003) found that the visual system gives special treatment to a chain of elements that could be perceived as a shape (i.e. a closed contour) over individual elements or an open contour. Others have found that the advantage for closed contours is driven by a preference for explanations of visual data that involve fewer objects over those that treat each element as an object unto itself (Murray, Kersten, Olshausen, Schrater, & Woods, 2002; Murray, Schrater, & Kersten, 2004; Fang, Kersten, & Murray, 2008; He, Kersten, & Fang, 2012).

Computational procedures for detecting closed contours in natural images have been extensively studied, primarily in applied contexts such medical imaging analysis, for example to analyze the shape of the interior of arteries and the heart ventricles (Guttman, McVeigh, & Prince, 1992; Guttman, Prince, & McVeigh, 1994), or to detect the boundaries of blood vessels (Yuan, Lin, Millard, & Hwang, 1999). Kass, Witkin, and Terzopoulos (1988)'s active contour model (sometimes known as the snake algorithm), has been employed successfully to find closed contours in natural images. While active contours often work with an edge map output by standard edge detection algorithms, newer algorithms for active contours work directly on the original image without an edge detection step. For example, the edgeless active contour algorithm of Chan and Vese (1999) and Chan and Vese (2001) has found success at finding contours on a variety of images, even very noisy images. Additionally, this algorithm can be used to find the boundaries between fuzzy objects without edges, or where the boundaries are not defined by image gradients. This relates directly to the stimuli in the experiments below, which contain closed contours embedded in random pixel noise. Standard edge detection algorithms fail to detect the target contours our stimuli at all, while human observers are able to detect them with ease, suggesting a substantial discrepancy between known algorithms and the mechanisms of the human visual system.

Previous work on the human visual system's thresholds for detecting shapes on blank backgrounds and the role of the shape's complexity has given mixed results (see Zusne, 1970 for a review). Much of this work measured detection thresholds for small shapes of uniform luminance over a dark background. Cheatham (1952) was unable to find an effect of complexity on detection thresholds. Other authors (Bitterman, Krauskopf, & Hochberg, 1954; Engstrand & Moeller, 1962; Hochberg, Gleitman, & Macbride, 1948; Krauskopf, Duryea, & Bitterman, 1954) used compactness (the ratio of the perimeter squared to the area) as a measure of complexity, and found that it was correlated with detection thresholds. For example, a five-pointed star is less easily detected than a circle. Kincaid, Blackwell, and Kristofferson (1960) describe a model that can account for the data connecting compactness to detection. Their model is a neural model in which the visual system responds to the presentation of a shape with a propagation of excitation, resulting in "fronts" of excitation meeting at a point resulting in an even larger amount of activiation which facilitates in the detection of the shape—anticipating the grassfire procedure described later by Blum (1973) and the more recent shock graphs of Siddiqi, Shokoufandeh, Dickinson, and Zucker (1998). However none of these early studies resolved the question of whether shape complexity influences detection of shapes in noise, nor developed the principled connection between detection and complexity that we propose below.

1.1. Statistical properties of contours

Many modern accounts of perceptual processes involve the idea that the visual system is optimized to the statistics of natural stimuli (Barlow, 1961; Geisler, 2008). Along these lines, a number of approaches to contour detection exploit statistical properties of natural contours. Several studies have shown that the visual system's implicit assumptions about the statistics of contour structure mirror those of contours in natural images (Elder & Goldberg, 2002; Geisler et al., 2001), and that contour integration can be understood as optimal or near-optimal probabilistic inference (Claessens & Wagemans, 2008; Feldman, 2001; Ernst et al., 2012). More specifically, a number of studies have suggested that the visual system's implicit probabilistic contour model assumes that "smooth" contours exhibit turning angles that are approximately von Mises distributed (Feldman, 1997; Feldman & Singh, 2005; Singh & Fulvio, 2005; Singh & Fulvio, 2007; Wilder et al., 2015). The turning angle α is the deviation of the contour from its tangent direction, and can be thought of as a discretization of curvature.[1] The von Mises distribution is the angular analog of the Gaussian (see Mardia, 1972). This model assumes that along a smooth curve turning angles are distributed approximately as

$$p(\alpha) \propto \exp\cos\beta\alpha, \tag{1}$$

where β is the parameter of the von Mises model analogous to the precision $1/\sigma^2$ of a Gaussian.

Eq. (1) provides a probabilistic generative model of smooth contours, and allows a number of predictions to be formulated about the characteristics of contour detection performance. In particular, Wilder et al. (2015) showed how it gives rise to a natural definition of *contour complexity*. Shannon (1948) defines Description Length (DL) as

$$\mathrm{DL} = -\log p(M), \tag{2}$$

where M is a quantity being measured, and $p(M)$ is the probability of obtaining that measurement (Cover & Thomas, 1991). The DL expresses the idea that a measurement is informative to the extent that it is "surprising" under a given probability model, and is a natural measure of complexity because it is the length of the shortest expression of M in an optimal code. Using this definition, if we think of a contour as a series $[\alpha_i]$ of turning angles generated independently and i.i.d. under the von Mises model $p(\alpha)$, then it follows that the DL of the contour is expressed by the integrated DL along the curve

$$\begin{aligned}\mathrm{DL}(\mathrm{CONTOUR}) &= -\sum_i \log p(\alpha_i) \\ &\propto -\beta\sum_i \cos\alpha_i,\end{aligned} \tag{3}$$

(see Feldman & Singh, 2005). This DL measure quantifies how unpredictable the contour is—how much it "zigs and zags" relative to the smooth expectation expressed by the von Mises model. Wilder et al. (2015) showed that contours with higher complexity in this sense are more difficult for subjects to detect, with performance declining with increasing DL.

In what follows, we extend this reasoning to the case of closed contours. The mathematical argument is similar in the closed case in that again complexity is defined as $-\log p$ for a suitable probabilistic model. But closed contours present a more substantial challenge because of the difficulty in defining a suitable probability model for them. Such a model would need to incorporate not only the local structure of the bounding contour, but also probabilistic regularities in the shape of the bounded region, which are much harder to quantify. Below we introduce such a model, and show how it gives rise to a set of interconnected shape and contour complexity measures. For simple open contours, the integrated DL is closely related to conventional (non-probabilistic) measures of contour complexity; for example the contour DL is related to the total squared contour curvature, a factor previously known to affect human performance. In contrast, the closed contour DL

[1] More technically, $\alpha \approx \Delta s\kappa$, where κ is the local curvature and Δs is the stepsize in arclength s; see Singh and Feldman, 2013.

measures that we introduce below have no simple counterparts in the conventional literature, as they quantify aspects of global form that are notoriously difficult to capture mathematically. The main goal of the experiments is then to test which, if any, of these shape complexity measures influences human detection of closed contours in noise.

In the experiments below, we asked subjects to detect closed contours in fields of pixel noise. On each trial, the subject saw (in random order) a target image containing a field of pixel noise with an embedded contour, and a noise image containing only pixel noise, and were asked to indicate with a button press which image (first or second) contained the target (2-interval forced choice). In Exp. 1, targets were the bounding contours of real shapes (animals and leaves drawn from real images), in order to investigate the system's sensitivity to naturally-occurring shapes. In Exp. 2, we use artificial contours whose structure has been manipulated in order to investigate more systematically the complexity factors that influence detection performance.

2. Experiment 1: Natural shapes

This study investigates the sensitivity of human observers to the structure of natural shapes and contours in a detection task. In this experiment subjects detected the closed contours of natural shapes that were embedded in a noise field. Closed contours appeared in an image displayed in one of two intervals and the subject responded, with a key-press, which interval contained the image with the contour.

2.1. Method

2.1.1. Subjects

Fifteen subjects participated in the experiment in return for credit in an introductory Psychology course. The subjects were naive to the purpose of the experiment. All subjects gave informed consent prior to participation. The experiment was carried out in accordance with the Code of Ethics of the World Medical Association (Declaration of Helsinki).

2.1.2. Stimuli

Stimuli were displayed on a gamma corrected iMac computer running Mac OS X 10.4 and Matlab 7.5, using the Psychophysics toolbox (Brainard, 1997; Pelli, 1997). We created two types of experimental images: non-target images, which contained only pixel noise; and target images, which contained pixel noise plus an embedded contour. The contours were the bounding contours of natural leaf or animal shapes drawn from standard databases. (We have previously used these shapes to investigate subjects' classification of shapes into superordinate shape categories; see Wilder, Feldman, & Singh, 2011.) Examples of shapes are shown in Fig. 1 and examples of a target image with shape embedded is shown in Figs. 2 and 3.

To create the noise images, we first created a 225 by 225 matrix of intensity values ranging from 0 and 1, with each intensity value drawn independently from a uniform distribution. When displayed an intensity of 0 corresponded to 0 cd/m^2 and an intensity value of 1 corresponded to 98 cd/m^2. Next the image was increased in size, to 550 by 550 pixels, using nearest-neighbor interpolation so that each pixel in the original image corresponded to a 2×2 pixel area in the resulting image. Target images were created identically, except that in addition to the noise a target contour with a constant intensity value was embedded at a random location and random orientation somewhere in the image. Leaf and animal databases were sampled with equal probability, and each image within each database was selected with equal probability. The leaf database was larger ($N = 424$) than the animal database ($N = 341$), but the probability of encountering a leaf or animal on each trial was 0.5. The bounding contour of the selected shape was extracted and scaled so that the contour would be 500 ± 25 pixels long. This means that there was, on average, no difference in mean luminance between target images with different shapes. This also means that each contour had 500 ± 25 turning angles, one for each pixel in the contour. Each turning angle could be estimated from three neighboring pixels, or we could estimate the turning angle at each pixel using an arbitrarily large window of integration (additional details will be provided in the "Modeling" section of this paper). Once the contour was properly sized it was dilated using the dilation mask ($[0\,1\,0;\ 1\,1\,0;\ 0\,0\,0]$). The dilation was performed so that a perfectly straight contour oriented horizontally or vertically was detected approximately as easily as one oriented at an arbitrary angle in pilot studies (see Wilder et al., 2015). Finally, the contour was embedded in the image at 56% contrast[2], which corresponds to a luminance of 21.6 cd/m^2. Because the noise images were doubled in size the dilated contours matched the scale of the noise in the images.

A sample target image containing an animal shape can be seen in Fig. 2. In Fig. 3 there are many examples of the contours shown in the experiment. (In Fig. 3 the contours are shown at maximum contrast for illustration, rather than at the lower level contrast used in the experimental stimuli.) The noise images are identical except for the lack of embedded contour.

The two experimental images were interleaved with masking images. The masks were noise images consisting of random black and white pixels. The masks were created by randomly sampling a binary matrix of size 225 by 225, and then scaling in a manner identical to experimental images.

2.1.3. Design and procedure

Each subject participated in 300 trials, with breaks after trial 100 and 200. Each trial began with a central fixation mark. After fixating the subject pressed a key to begin the trial. A sequence of five images was rapidly presented to the subject. The first, third, and fifth images were the masks. The second and fourth images were the experimental images (target and non-target), in random order. Each mask and experimental image was presented for 500 ms, with a 100 ms gap between images. The subject's task was to decide whether the contour was present in the first or second of the two experimental images (2IFC), and respond by a key press.

Head position was not constrained, but the head was placed at 66 cm from the display and subjects were instructed to remain as still as possible. During the three breaks, if necessary the head was repositioned to be 66 cm from the display. Images were presented at 6.5° of visual angle, measuring 15×15 cm on the display.

2.2. Results

The majority of the subjects were able to perform the task very well. Two of the 15 subjects were excluded from analysis because they performed below 55% correct. The remaining 13 subjects performed between 65% and 90.3% correct, with a mean of 77.7% correct. Subjects were significantly better ($t = 2.96$, $p = 0.0068$) at detecting leaf shapes (82.8% correct) than animals shapes (73% correct). All 13 subjects had a higher proportion of correct answers for leaves than for animals; this trend was significant for 8 of the 13 subjects using a two-proportion z-test with $\alpha = .05$.

[2] Contrast was measured as the difference of the mean luminance of the image (i.e. the background) and the luminance of the contour divided by the mean luminance of the image.

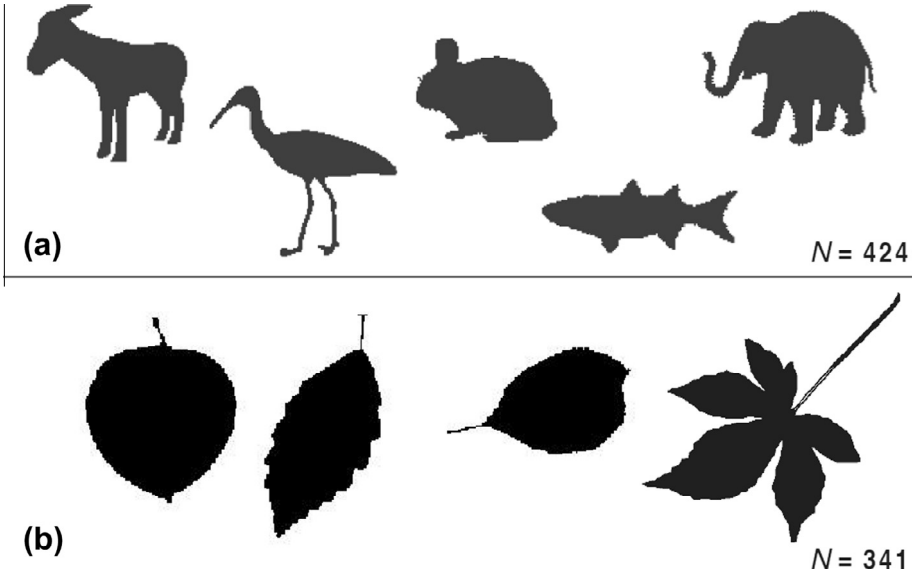

Fig. 1. Examples of the natural shapes, from which bounding contours would be extracted and embedded in noisy images. The top row (a) shows a sample of the animal shapes, and the bottom row (b) shows a sample of leaf shapes.

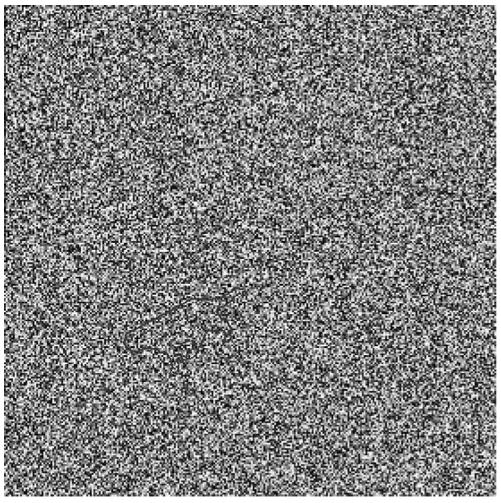

Fig. 2. An example stimulus image containing a contour, in this case an animal embedded slightly down and to the left of center. This image is best viewed or printed at a zoom of 100%, otherwise the image will be distorted due to pixel averaging or interpolation.

Beyond these basic analyses we would like to know what properties of the contours, if any, influenced their detectability. To this end, we now introduce complexity measures suitable for closed shapes. As explained above, doing so requires developing a suitable probabilistic generative model for closed contours.

3. Modeling

3.1. Contour complexity

As discussed above, in Wilder et al. (2015) we showed how subjects' ability to detect a contour in a noise field is influenced by a measure of contour complexity based upon the integrated information along the contour. Here we extend this model to closed contours by introducing a probability model for closed contours, and developing complexity measures from it.

We begin by expressing the contour detection problem as a decision problem, in which the observer must decide if each experimental image is more likely to be a target image or a noise image—that is, whether or not it contains a contour that is likely to be sample from an assumed probabilistic generative model. Above, we introduced the von Mises model which quantifies the likelihood of any given chain of elements under the "smooth contour" hypothesis. For mathematical convenience, here we adopt a (numerically almost identical) Gaussian model[3] in which each turning angle has probability

$$p(\alpha|\text{CONTOUR}) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left[-\frac{(\mu-\alpha)^2}{2\sigma^2}\right], \tag{4}$$

where μ and σ are respectively the mean and standard deviation of the smooth turning angle distribution. A contour consists of a chain of N turning angles $\{\alpha_i\}$, assumed to be drawn i.i.d. from this model, so the likelihood of a contour under this model is

$$\begin{aligned} L_c &= p\left(\{\alpha_i\}_{i=1}^N|\text{CONTOUR}\right) = \prod_{i=1}^N p(\alpha_i|\text{CONTOUR}) \\ &= \left(\frac{1}{\sigma\sqrt{2\pi}}\right)^N \exp\left[-\frac{1}{2\sigma^2}\sum_{i=1}^N(\mu-\alpha_i)^2\right]. \end{aligned} \tag{5}$$

Now, alternatively, consider seeing a sequence of turning angles under the hypothesis that this sequence is simply noise. Under the noise assumption there is no reason to expect one turning angle over another, so the turning angle distribution is uniform, where each angle has some small probability ϵ. Since again the sequence of turning angles is i.i.d., the probability of the entire sequence is the product of the individual probabilities, yielding

$$L_0 = p\left(\{\alpha_i\}_{i=1}^N|\text{NOISE}\right) = \epsilon^N. \tag{6}$$

The sequence $[\alpha_i]$ can now be classified as smooth contour rather than noise if

$$p\left(\text{CONTOUR}|[\alpha_i]_{i=1}^N\right) > p\left(\text{NOISE}|[\alpha_i]_{i=1}^N\right). \tag{7}$$

[3] Only the von Mises distribution is technically appropriate for angles, because it has domain $[-\pi, \pi]$, whereas the Gaussian has domain $(-\infty, \infty)$. We use the Gaussian here simply because it makes the relationship between complexity and curvature more transparent. In any case the two distributions are virtually identical over the range of angles considered here; see Feldman and Singh (2006).

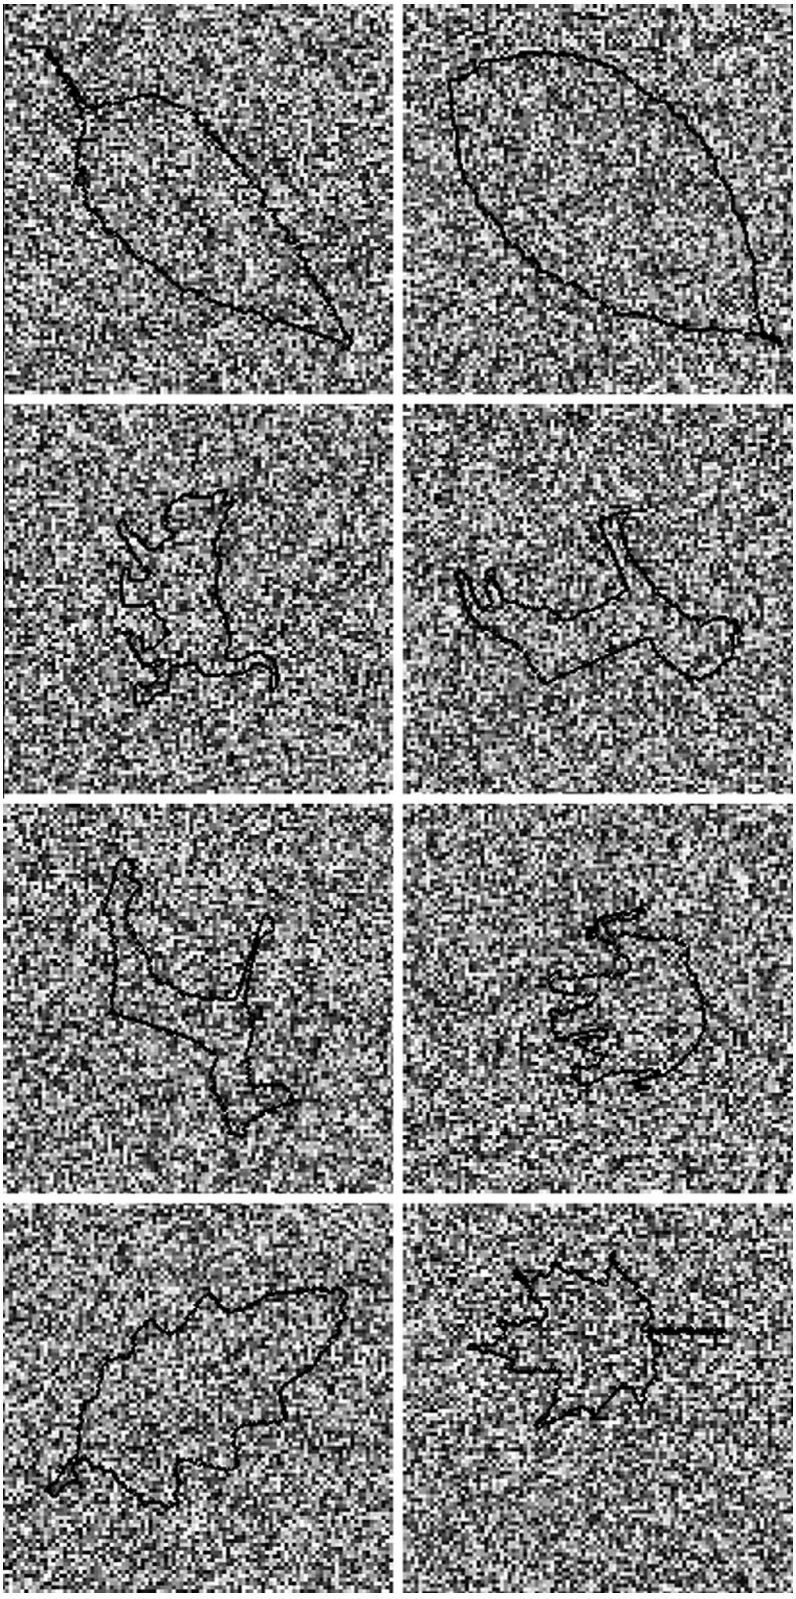

Fig. 3. Several example contours embedded in noise. Only the portion of the image containing the contour is shown. The contours are shown at the lowest luminance for the ease of the reader, though in the actual experiment the images were much closer to the mean luminance (64% contrast, or a luminance of 21.6 cd/m^2). This image is best viewed or printed at a zoom of 100%, otherwise the image will be distorted due to pixel averaging or interpolation.

Assuming equal priors (as in the experiment), this reduces to a criterion on the likelihoods: the observer should decide "Contour" if $L_C > L_0$. Equivalently, we can rewrite this criterion in terms of the weight of evidence (WOE) in favor of the smooth contour model, that is, the log likelihood ratio:

$$w = \log(L_c/L_0) = \log L_c - \log L_0. \tag{8}$$

The turning angles should be classified as a contour if the WOE is greater than zero, $w > 0$.

As discussed above, the DL ($-\log p$) can be thought of as the complexity of a given observation given a model. With this in mind, we can rewrite the classification criterion in terms of the DL of the contour. We define the complexity of the contour under the smooth model,

$$\mathrm{DL(CONTOUR)} = -\log p\left([\alpha_i]_{i=1}^N|\mathrm{CONTOUR}\right) = -\log L_C. \tag{9}$$

Note that this measure expresses the complexity of the contour under the smooth contour model specifically; below we will define analogous measures of complexity that use alternative contour models. With this in mind, the decision variable w can be expressed as

$$w = \log L_c - \log L_0 = -\mathrm{DL(CONTOUR)} - \log L_0. \tag{10}$$

As complexity DL(CONTOUR) increases, the decision variable w decreases, and the contour is more likely to be classified as noise. As mentioned above, the DL is closely related to contour curvature, because it depends monotonically on the squared turning angle, though in our view the decision-theoretic motivation here gives a more principled explanation of the role played by curvature than do conventional accounts.[4]

In Wilder et al. (2015), we considered the case of open contours, for which the expectation (mean) of the turning angle μ is 0°, meaning straight continuation, because they are assumed to have no bias to turn in either direction. For closed contours, the mean is instead $2\pi/N$, where N is the number of angles in the sequence (Feldman & Singh, 2005), because the sum of the exterior angles of any closed non-self-intersecting contour must equal 2π radians (360°). Intuitively, the contour must, on average, bend slightly to the right (moving clockwise around the contour) so that it will eventually close on itself. The σ of the Gaussian (or the β of the von Mises) is fixed for the shapes under consideration. This parameter can be fixed to any suitable value without affecting the relative DLs of the contours. Changing this value will keep the ordinal ranking of a set of contours constant, only changing the magnitude of the difference between the DL of two contours. In what follows, the notation DL(CONTOUR) will indicate DL under this smooth closed contour model.

[4] To better appreciate the role of curvature in the decision, plug L_c (Eq. (5)) into the DL equation (Eq. (9))

$$\mathrm{DL(CONTOUR)} = -\log L_C = N\ln\sigma\sqrt{2\pi} + \frac{1}{2\sigma^2}\sum_{i=1}^{N}(\mu-\alpha)^2. \tag{11}$$

Plugging Eq. (11) into Eq. (10) leads to a more comprehensible form of the decision variable

$$w = -\mathrm{DL(CONTOUR)} - \log L_0, \tag{12}$$

$$= -N\ln\sigma\sqrt{2\pi} - \frac{1}{2\sigma^2}\sum_{i=1}^{N}(\mu-\alpha)^2 - N\ln\epsilon, \tag{13}$$

$$= -c_1 - c_2\sum_{i=1}^{N}(\mu-\alpha)^2 - c_3, \tag{14}$$

$$\propto -\sum_{i=1}^{N}(\mu-\alpha)^2 + c_4. \tag{15}$$

The WOE is proportional to the negative of the squared difference between each observed angle and the mean of the turning angle distribution. In these expressions, c_4, where c_1, c_2, c_3 and c_4 are constants that do not depend on the turning angle.

DL(CONTOUR) has one free parameter: the scale at which the turning angles are measured. In Wilder et al. (2015) each turning angle was defined by three neighboring pixels, getting turning angles by looking forward and backward one pixel from each pixel on the contour, i.e. the finest available scale. In what follows we instead optimize the scale, i.e. we chose the scale at which the contour complexity was most closely related to detection performance. The best scale (for both Exp. 1 and the following experiment, Exp. 2) was to look forward and backward 32 pixels from each pixel. Note that this does not decrease the number of turning angles in the contour, because the turning angle filter is a sliding window 65 pixels wide that is evaluated at each point along the contour.

Fig. 4 shows detection performance as a function of DL(CONTOUR) at the optimal scale (which we use henceforth). The data show a decrease in performance as contour complexity increases; a 2-AFC logistic regression shows that DL(CONTOUR) is a significant predictor of performance, $p = 0.024$. Hence we conclude that the complexity of the contour in this particular sense, as in Wilder et al. (2015), does have an influence on detection. Next, we consider alternative shape complexity factors that are more global in nature.

3.2. Shape complexity measures

While a closed contour is simply a contour that closes, it is perceptually much more than that. A closed contour defines a bounded region whose shape may be defined in a variety of ways. In this section we introduce several new complexity measures to quantify the shape of the bounded region. As before, each complexity measure has the form $-\log p$, but with probability defined via a suitable region-based generative model.

We adopt a skeleton-based shape-generating model introduced in Feldman and Singh (2006) (see also Wilder et al., 2011; Feldman et al., 2013). This generative model builds on the extensive literature on medial representations of shape that began with Blum (1973)'s medial-axis transform (MAT) and grassfire algorithm, and has been extended in numerous ways since (Siddiqi et al., 1998; August, Siddiqi, & Zucker, 1999; Leymarie, Kimia, & Giblin, 2004). Medial shape representations have known psychological consequences (Psotka, 1978; Kovacs & Julesz, 1994; Burbeck & Pizer, 1995; Harrison & Feldman, 2009 and neural correlates Lescroart & Biederman, 2012; Hung, Carlson, & Connor, 2012).

In our Bayesian skeleton framework, we assume that the bounding contour is generated by a process of stochastic "growth" from a branching axial structure, the *generating skeleton* (Feldman & Singh, 2006). We assume some underlying prior probability distribution of skeletons, $p(\mathrm{SKEL})$, which favors skeletons with relatively few branches and relatively straight component axes. Given a skeleton, random vectors (called *ribs*) sprout laterally on both sides of each axis, at angles and distances drawn from probability distributions, and their endpoints define the bounding contour; this defines the likelihood function, $p(\mathrm{SHAPE}|\mathrm{SKEL})$, which assigns probability to specific shapes given a skeleton. In this framework, the observer interprets a given shape by estimating the skeleton *most likely* to have generated it, called the maximum a posterior (MAP) skeleton. That is, the MAP skeleton is the skeleton that maximizes the posterior

$$p(\mathrm{SKEL}|\mathrm{SHAPE}) \propto p(\mathrm{SKEL})p(\mathrm{SHAPE}|\mathrm{SKEL}). \tag{16}$$

The MAP skeleton has a number of desirable properties as a representation of global shape. It is an optimal representation of shape given the assumed generative model, balancing the complexity of the shape model (the MAP skeleton) with the degree of fit between the shape and the model (see Fig. 5). As a result, distinct axes of the MAP skeleton tend to correspond to distinct intuitive parts of the

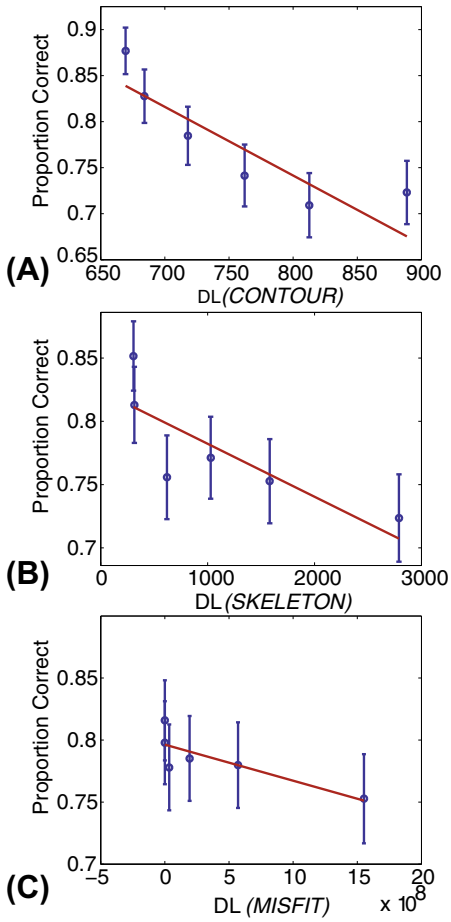

Fig. 4. Performance (proportion correct) as a function of (A) DL(CONTOUR), (B) DL (SKELETON), or (C) DL(MISFIT) for the combined subjects' data. Error bars show one standard error. Data was binned into 6 bins of equal size for visualization; for all analysis described in the text no binning was used. The data is noisy, but shows the trend in the predicted direction for each complexity; detection was best for low complexities and worse for higher complexities. The largest effect appears to be with (B) DL(SKELETON).

shape (Feldman & Singh, 2006; Feldman et al., 2013). Superordinate shape classification based on the estimated skeleton closely tracks human judgments (Wilder et al., 2011). Broadly speaking, the MAP skeleton is a useful low-parameter representation of shape that effectively quantifies properties of global form. Here, we investigate the complexity measures that arise from the MAP skeleton, and in particular their influence on detection performance.

The MAP skeleton has several parameters that can be adjusted, but for this paper the parameters are not adjusted to improve the fit with the data, but are simply those used in Feldman and Singh (2006) and Wilder et al., 2011.

As above, we introduce a complexity measure by taking the negative log probability (DL) of the contour under a given probability model, here the skeletal model. Notice that the DL of the posterior

$$DL(\text{SKELETON}|\text{SHAPE}) = -\log p(\text{SKEL}|\text{SHAPE}), \tag{17}$$

decomposes into two additive parts,

$$\begin{aligned}\text{DL}(\text{SKELETON}|\text{SHAPE}) &= -\log p(\text{SKEL}) - \log p(\text{SHAPE}|\text{SKEL}) + c, \\ &= DL(\text{SKELETON}) + DL(\text{SHAPE}|\text{SKELETON}) + c,\end{aligned} \tag{18}$$

arising respectively from the prior and likelihood (plus a constant c that corresponds to the proportionality constant in Eq. (16)). These can each be interpreted as distinct complexity measures, which we refer to respectively as DL(SKELETON) and DL(MISFIT):

$$\text{DL}(\text{SKELETON}) = -\log p(\text{SKEL}), \tag{19}$$

$$\text{DL}(\text{MISFIT}) = \log p(\text{SHAPE}|\text{SKEL}). \tag{20}$$

DL(SKELETON) expresses the complexity of the estimated skeleton itself, penalizing for the multiplicity and curvature of the component branches. It is larger for shapes whose skeletons branch in more complicated ways (see Fig. 6, top row). DL(MISFIT) expresses the discrepancy between the estimated skeleton and the shape itself, penalizing shapes that are not well-described by their own best skeleton. It is high for shapes that have relatively noisy contours relative to the typical shape generated from the given skeleton (see Fig. 6, bottom row). Our working hypothesis is that each of these measures quantifies a distinct sense in which shapes can be globally "complex".

Fig. 4(C) shows the influence of DL(MISFIT) on detection performance. The effect of DL(MISFIT) on detection is similar to that of DL (CONTOUR): weak, but in the predicted direction. A 2-AFC logistic regression shows no significant effect of DL(MISFIT) when aggregating over subjects ($p = 0.403$; Fig. 4(C)). The decreasing trend was present in about half of the subjects.

A slightly clearer effect on detection is found when looking at DL(SKELETON). Fig. 4(B) shows a slightly larger decrease in detection with an increase in complexity. The 2-AFC logistic regression reveals that DL(SKELETON) is a significant predictor of performance ($t = -4.555$, $p = 0.0000054$). All but one of the individual subjects showed a decrease in detection as DL(SKELETON) increased.

The above analyses test for the effect of the three different types of complexity in isolation. Now we will look at multiple factors together. First, we will simply look at the correlation between the separate factors, and then at regression models involving multiple factors.

Table 1 shows that the two shape complexities, DL(SKELETON) and DL(MISFIT), are very weakly correlated with each other and with DL(CONTOUR). All of the correlations were significant ($p < 0.001$, df = 3877).

To understand how these three factors might interact, we entered all three factors into 2-AFC logistic regression model using a stepwise function (the "step" function in R) that includes the factors and interactions that increased proportion of explained variance, while penalizing for model complexity using the Akaike-Information Criterion (AIC). In this analysis, both DL(CONTOUR) and DL(SKELETON) were part of the winning (lowest AIC) model, meaning that both factors were required to best explain the data. The complexity of the MAP skeleton DL(SKELETON) significantly decreased detection ($z = -2.33$, $p = 0.0197$). The complexity of the contour DL(CONTOUR) significantly decreased detection ($z = -2.974$, $p = 0.003$).

The regression analysis supports the idea that in order to account for our observers' data we need to include local information, DL(CONTOUR), and non-local information, DL(SKELETON). Further support for this idea comes from an additional analysis using leave-one-out cross validation. We use a model with factors DL(CONTOUR) and DL(SKELETON) as well as their interaction, and compare this model to simpler models without the interaction (the winning model from the regression analysis) or with only one of the individual factors. For each of those models we fit the regression model using leave-one-out cross validation. Cross validation error was the smallest in the model including both DL(CONTOUR) and DL(SKELETON) as factors, which is the same "winning" model as in the stepwise regression.

These results suggests that detection of the closed contours cannot be simply explained by local contour effects. The subjects'

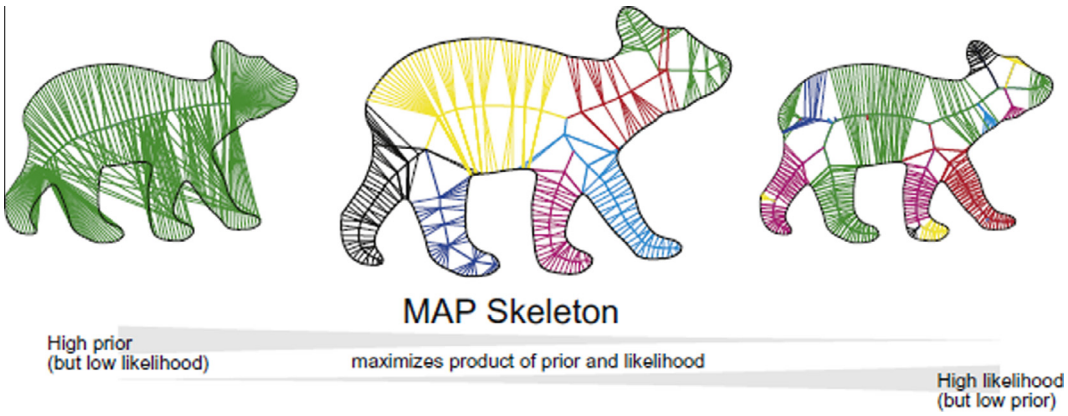

Fig. 5. On the left is a bear shape with a simple skeleton (high prior probability). This skeleton would not be likely to generate the bear shape (low likelihood of the shape given the skeleton). On the right there is a bear shape with a complex skeleton (low prior probability) that is very likely to generate the bear shape (high likelihood of the shape given the skeleton). In the middle is a bear shape with the MAP skeleton. The skeleton is a balance of skeletal complexity (medium prior probability) and likelihood of generating the shape (medium likelihood) that fits "just right". Figure adapted from Fig. 17 of Wilder et al. (2011).

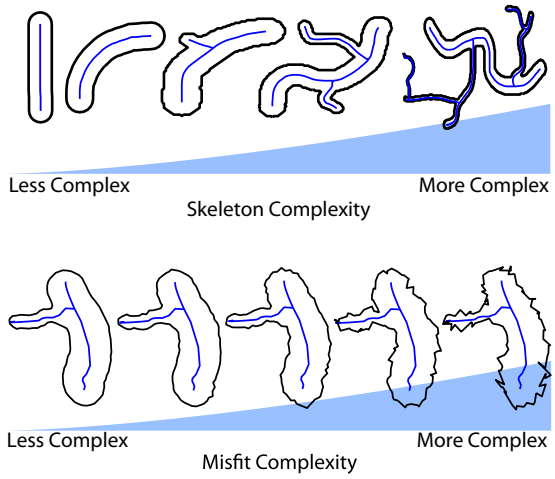

Fig. 6. Illustration of the two shape complexity measures, DL(SKELETON) (top) and DL(MISFIT) (bottom). Straight, single-axis score low on DL(SKELETON) (top row) because they have high prior probabilities, while multi-axis skeletons with high axial curvature score high. Shapes with low noise relative to their skeletons score low on DL(MISFIT), because they have high likelihood, while noisy shapes score high.

detection performance is influenced by *global* form as well, in that it is best explained by a combination of DL(CONTOUR) and DL(SKELETON). Closed contours that bound more complex regions are more difficult to detect, at least as they arise in natural shapes. This result must be interpreted with caution, however, because the use of natural shapes meant that potentially interactions among shape factors were uncontrolled. In the next experiment, we use artificially generated shapes whose form has been manipulated in order to more carefully distinguish the various potential shape complexity measures.

4. Experiment 2: Experimentally manipulated shapes

Exp. 2 follows the same basic method as Exp. 1, except here we create shape artificially and manipulate their structure in order to study the effects of shape complexity more systematically. More specifically, natural shapes do not cover complexity space uniformly, meaning that in Exp. 1 some complexity values were overrepresented in the stimulus set relative to others. Moreover, correlations among the various complexity measures were uncontrolled, making it difficult to tease apart the effects of the separate factors. In Exp. 2 shapes, we attempted to manipulate DL(MISFIT) and DL(SKELETON) somewhat more uniformly and more independently by independently manipulating the *axial structure* (which primarily affects DL(SKELETON)) and the *noise along the contour* (which primarily affects DL(MISFIT)). The goal is to create a shape space that samples complexity space more uniformly and that minimizes correlation among complexity measures. More broadly, the experiment serves to confirm that the results of Exp. 1 were not due to any other unknown artifacts of the shapes included in the database.

Table 1
Correlations of the individual complexities from Exp. 1

Factor 1	Factor 2	Pearson correlation
DL(*SKELETON*)	DL(*MISFIT*)	−0.072
DL(*SKELETON*)	DL(*CONTOUR*)	0.287
DL(*MISFIT*)	DL(*CONTOUR*)	0.106

As in Exp. 1, subjects were required to decide which of two intervals contained the image containing the closed contour, and which contained the noise image. The effect of the three types of complexity, DL(CONTOUR), DL(SKELETON), and DL(MISFIT), were first analyzed in isolation, and then in combination, to find the factors that best describe the subjects' data without over-fitting.

4.1. Methods

4.1.1. Subjects

Twenty (new) subjects participated in the experiment in return for credit in an introductory Psychology course. The subjects were naive to the purpose of the experiment and gave informed consent prior to participating in the experiment.

4.1.2. Stimuli

As in Exp. 1, there were two experimental images, one containing the target embedded in noise, and the other just noise, presented in random order. The subject's task was to indicate with a keypress which image contained the target. As in Exp. 1, the experimental images were interspersed with mask images.

To create the target shape, a shape skeleton was generated and then a shape was "grown" outward using the generative model described above. There were four different shape skeleton conditions: single part shapes, two part shapes, three part shapes where two parts are on opposite sides of a main part, and three part

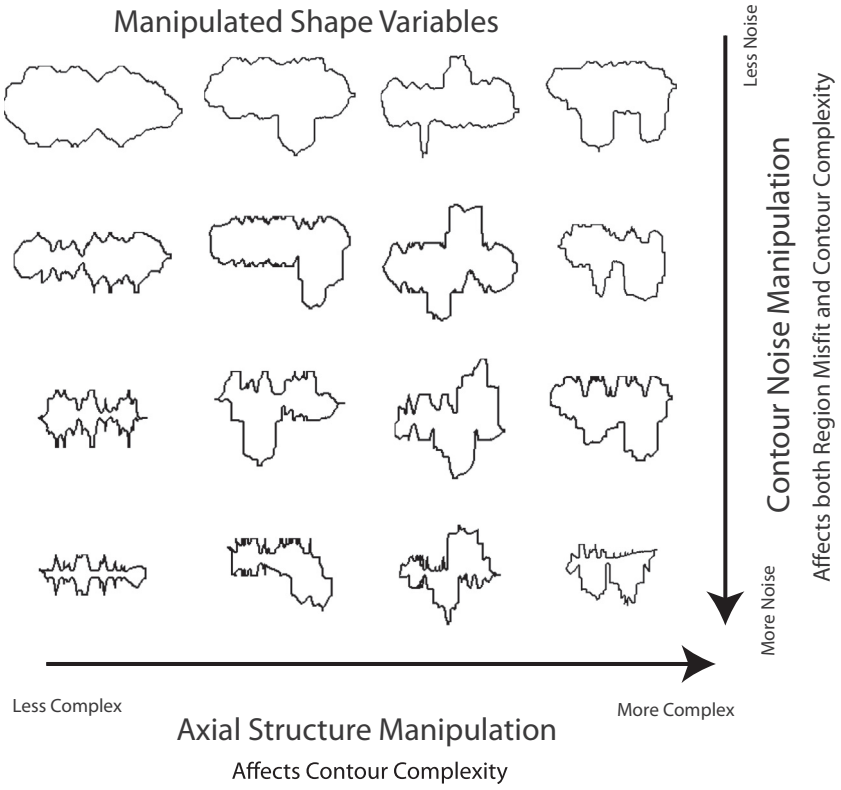

Fig. 7. Illustration of the shape manipulations used in Exp. 2, intended to cover two complexity measures more uniformly. The horizontal axis illustrates the effect of the manipulation of axial structure, which primarily affects DL(SKELETON). Moving left to right, shapes have respectively a single axis, then two axes, then three axes, first with two axes on opposite sides of the shape and then two axes on the same side of the shape. The vertical axis illustrates the manipulation of contour noise, which primarily affects DL(MISFIT). At the top, shapes have a small amount of variability in the lengths of the ribs that generated the contour, while shapes towards the bottom have larger amounts of noise

shapes where two parts extend off the same side of a main part. The first condition began with a single straight axis. The second condition used a straight main axis with a second skeletal axis added, perpendicular to the main axis, at a random location, yielding a skeleton with a higher complexity; this automatically yields a lower prior and thus higher DL(SKELETON). The third condition used a straight main axis with a second axis emanating from one side of the main axis, and a third axis emanating from the other side of the main axis. Both the second and third axis were perpendicular to the main axis and were at random locations. This again results in a skeleton that has higher DL. The fourth axis condition has a straight main axis with a second and third axis emanating from the same side of the main axis. Both are perpendicular to the main axis and at a random location, with the constraint that the third axis must be more than $2 \times 15 = 30$ units away from the second axis, where 15 is the length of the ribs for the second and third axis (explained in more detail below). This condition results in skeletons with similar complexities as in the third axis condition, but higher complexities than skeletons in the first or second conditions. The overall result of was a variety of skeletal structures varying in DL(SKELETON).

Once a skeleton was generated, a shape was grown from the skeleton via the likelihood function described above, in which "ribs" (random deviates) sprout laterally from the skeleton to form the shape. Beginning at one end of the first axis a rib was grown perpendicularly left and right from the skeleton. A contour point was placed at the end of the rib. The length of the ribs was sampled from a Gaussian distribution centered at 25 units away from the skeleton. The contour noise condition set the standard deviation of the Gaussian; the contour noise conditions ranged from lowest noise to highest noise, and had standard deviations of 1, 3, 5, or 7 units. After obtaining the first contour point, the second point on the skeletal axis is considered. Again we sample a rib length from a Gaussian, but this time the Gaussian is centered at the previous rib length. The contour was sampled in this Markovian fashion around the entire axis. Once a child axis is reached the ribs are generated in the same manner, but the Gaussian from which rib lengths are sampled is initially centered at 15 instead of 25. This process yields a variety of shapes for each skeletal structure differing in DL(MISFIT).

Fig. 7 shows examples of shapes with each combination of factors.

4.1.3. Design and procedure

The design was a 4×4 factorial design (four skeletal axis conditions and four contour noise conditions). Subjects saw 25 trials from each condition, in a random order, for a total of 400 trials. The stimulus was randomly generated on each trial, so no subject saw the same stimulus as another, and no stimulus was repeated. Each trial was identical to a trial of Exp. 1, except that the stimuli were generated as described above.

5. Results

The results from Exp. 2 were generally noisier than those from Exp. 1, perhaps because the unnatural shapes induced more uncertainty in the observers. We found no significant effects of our manipulated variables (see Fig. 8), though the skeletal axis manipulation was close to reaching significance ($p = 0.053$). Note though that our main focus was not on the manipulated variables per se, but rather on the effects of the various complexity measures,

which are related but different. The goal of the manipulations was to induce a more uniform sampling of complexity space, and the main focus of our analyses is on any effects of complexity that this sampling might uncover.

Both complexity measures, DL(SKELETON) and DL(MISFIT), influenced detection in the expected direction, with increasing complexity leading to decreased detection (Fig. 9). Neither DL (CONTOUR) nor DL(MISFIT) had a significant effect on detection when considered in isolation (respectively $p = 0.274$ and $p = 0.243$), similarly to Exp. 1. The effect of DL(SKELETON) in isolation was marginally significant ($p = 0.045$).

As in Exp. 1 we measured the correlations between the different complexities (Table 2). The correlation coefficients between the different complexities were small (but all were significant, $p < 0.0001$, df = 6390), and in fact DL(SKELETON) and DL(CONTOUR) are very weakly anti-correlated.

As in Exp. 1, in addition to looking for effects of the three complexity measures in isolation, we next considered models of their interaction. We again used a stepwise minimization procedure (the step function in R) to find the 2-AFC logistic regression model that resulted in the lowest AIC (best fit to the data after taking into account the number of degrees of freedom). Here we find that a combination of multiple factors was not better at explaining the data than a model using only the effect of DL(SKELETON).

Again, as in Exp. 1, in addition to the regression analysis, we can confirm this result using leave-one-out cross validation. We fit models using the single factor DL(CONTOUR), the single factor DL (SKELETON), both as factors, and both as factors with their interaction included. We again find that the cross validation analysis results in the same model as obtained in the regression analysis; the model consisting only of DL(SKELETON) gives the best cross-validated fit to the data.

In summary, the best fitting model of the effect of complexity of experimentally manipulated shapes (Exp. 2) shows that non-local factors influence detection, just as in Exp. 1, even though detection does not in principle require perceiving the whole. The effects of Exp. 2 are small, presumably because the range of skeletal complexities is smaller than the range found in natural shapes.

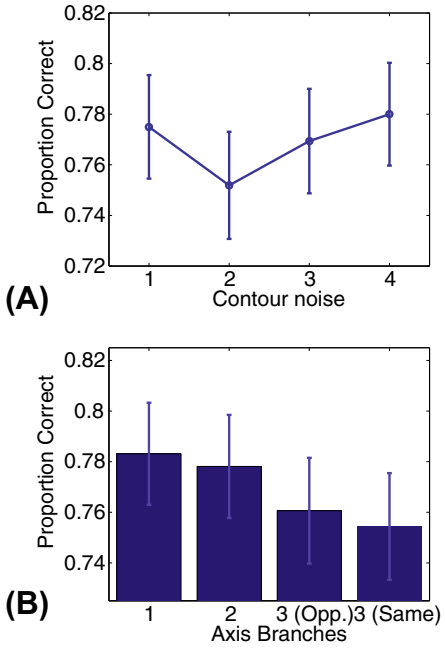

Fig. 8. Performance in Experiment 2 (proportion correct) is shown as a function of the manipulated variables – (A) contour noise, and (B) axial structure. The axial structure manipulation has a non-significant effect, but the trend is in the expected direction; an increase in the number of branches tends to decrease the performance. The contour noise manipulation appears not to have affected detection performance.

6. Discussion

Our results show that human subjects' detection of closed contours in noise is impaired, slightly but reliably, by the complexity of the shape of the bounded region. Specifically, we found that the complexity of the global form of a target, as measured by DL (SKELETON), influenced performance. These results cannot be explained by local contour complexity by itself; DL(CONTOUR) in isolation had a negligible affect on performance in Exp. 2, and a significant effect in Exp. 1 with the caveat that a model that includes DL(SKELETON) explains the data better, even though the same measure of contour complexity *does* strongly impair performance with simple open contours (Wilder et al., 2015). The closure of the contours in the current displays seems to invoke involuntary processes of form representation that in turn influence performance.

To fully appreciate our results, it is important to keep in mind that our detection task could, in principle, be accomplished by a completely local process, such as one that simply detects short smooth sections of contour. There is no inherent requirement that the system be sensitive to the holistic form of the contour, or for the fact that the target contours are closed to elicit different processes from those at work with open contours. Hence it is all the more remarkable that we found consistent, albeit small, effects of global form on contour detection.

Moreover our results cannot be explained by standard computer vision techniques, which are generally unable to detect the contours in our displays at all. A more complete understanding of the nature of the complexities affecting human subjects, and the global shape representation processes that these complexity effects uncover, would take us closer to achieving computational models that can achieve human-like performance, though at the moment such models are unavailable, and in any case are not the main goal of our research.

Our results are broadly consistent with "global precedence" effects suggesting that global stimulus properties are computed early and dominate the eventual percept (Hochstein & Ahissar, 2002; Navon, 1977; Lee, Mumford, Romero, & Lamme, 1998; Fulvio, Jacqueline, & Singh, 2006; Singh, 2015). Though object detection effects such as those at work in our displays presumably begin locally as information flows through visual cortex, representations of whole shapes that arise later in the visual stream, including some that specifically involve skeletal representations of shape (Hung et al., 2012; Lescroart & Biederman, 2012), quickly begin to convey a decisive influence (Kovacs & Julesz, 1994; Burbeck & Pizer, 1995; Lee et al., 1998; Harrison & Feldman, 2009). Our results add to the weight of this viewpoint, suggesting that those later representations influence basic detection mechanisms even when they are not required to accomplish the experimental task.

While many of the closed contour studies mentioned earlier suggested that closed contours are treated differently than open contours, Tversky et al. (2004) were unable to find an effect of closure on detection that could not be explained by a simple grouping rule. The transitivity rule (Geisler & Super, 2000) simply says that if elements *A* and *B* are grouped, and *B* and *C* are grouped, then *A* will also be grouped with *C*. Subjects in Tversky et al. (2004)'s study were able to detect a sequence of edge elements that closed to form a circle using fewer edge elements than were necessary to detect an open sequence of edge elements. The performance difference was predicted by an application of the transitivity rule, which increased the grouping strength between all of the elements in the

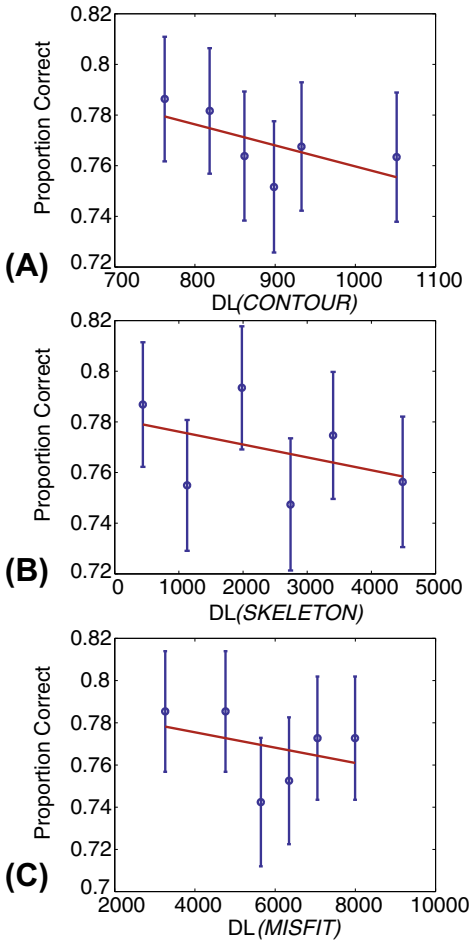

Fig. 9. Performance (proportion correct) is shown as a function of (A) DL (CONTOUR), (B) DL(SKELETON), and (C) DL(MISFIT) for the combined subjects' data from Experiment 2. The data is binned into six bins with an equal number of trials per bin for visualization purposes. Statistical analyses described in the text do not use binning. Error bars show one standard error. There is a weak trend in the predicted direction for each type of complexity; detection was best for low complexities, and worse for higher complexities.

circle; no additional closure mechanism was required. Our results, in contrast, suggest that the complexity of a closed contour cannot be reduced to the complexity of the constituent contours. It would be too strong, however, to infer a contradiction between our results and these earlier results, in light of the many differences between the stimuli and tasks.

Finally, our results further corroborate the Bayesian skeletal model of Feldman and Singh (2006), in that the complexity measures that arise from it have some correlation with performance in this task. Certainly, other shape complexity measures more narrowly tuned to this task might also have found support in our detection data. But the DL measures we introduced here arise naturally out of this broad, mathematically unified framework which has empirical support from a variety of other tasks (see Feldman et al., 2013). In this light, the current results provide converging evidence for a probabilistic account of shape representation and perceptual grouping.

The findings reported here raise several additional questions, such as the nature of the perceptual transition between open and closed contours. Our results suggest that the system "explains" a contour via a different perceptual model when it is closed than when it is open, but the criteria for deciding which model applies to any given stimulus are unclear. Bayesian theory can provide criteria for choosing among distinct probabilistic models (called *model selection*, see Myung, 2000), but how such criteria would apply to the problem at hand is speculative, especially because intermediate cases were not the focus of the experiments. A contour in the shape of a broken circle or C shape, for example, might be treated as a simple closed contour (a circle) with missing data, or as a fairly complex (curved) contour with complete data, or as something in between. An understanding of how the visual system treats such cases would require further experimentation.

Table 2
Correlations of the individual complexity measures from Exp. 2

Factor 1	Factor 2	Pearson correlation
DL(*SKELETON*)	DL(*MISFIT*)	0.109
DL(*SKELETON*)	DL(*CONTOUR*)	−0.058
DL(*MISFIT*)	DL(*CONTOUR*)	0.059

7. Conclusions

When a contour closes, a shape is formed. Our results show that when observers seek a closed contour amid noise, they are influenced by the complexity of the bounded shape. This result is not predicted by conventional account of shape representation or contour integration, but is a natural consequence of a Bayesian model of shape representation as applied to the detection problem. The key mathematical idea is the quantification of complexity as the negative logarithm of probability under a suitable shape model (Eqs. 19, 20), which as we have shown appears as the critical term in a natural criterion for the presence of a contour (see Wilder et al., 2015). The contribution of complexity to performance is not best explained by local contour complexity by itself, but requires the inclusion of the complexity of the shape skeleton.

Our broader hope is that our mathematical treatment of shape complexity will contribute to a better understanding of the role of closure and the representation of whole shapes. The Gestaltists had the right intuition that whole shapes are represented in a manner that goes beyond their constituent contours. But while the representation of simple contours in visual cortex is now fairly well understood, the representation of whole shapes is not. Our results corroborate the Gestalt view that a complete understanding of contours is not, in and of itself, sufficient to understand whole shapes. The quantification of the complexity of whole shapes is merely one component of a much larger story, but we hope that more attention to the problem of representing whole shapes, and specifically to skeletal representations, will spur development of a more comprehensive account.

Acknowledgments

J.W. was supported by the Rutgers NSF IGERT program in Perceptual Science, NSF DGE 0549115 and NSERC CREATE 397882-2011. J.F. and M. S. were supported by NIH EY021494.

References

Altmann, C. F., Bülthoff, H. H., & Kourtzi, Z. (2003). Perceptual organization of local elements into global shapes in the human visual cortex. *Current Biology, 13*(4), 342–349.

August, J., Siddiqi, K., & Zucker, S. (1999). Ligature instabilities in the perceptual organization of shape. *Computer Vision and Image Understanding, 76*(3), 231–243.

Barlow, H. B. (1961). Possible principles underlying the transformation of sensory messages. In: W. A. Rosenblith (Ed.), (pp. 217–234). Cambridge: M.I.T. Press.

Bitterman, M., Krauskopf, J., & Hochberg, J. (1954). Threshold for visual form: A diffusion model. *The American Journal of Psychology, 67*(2), 205–219.

Blum, H. (1973). Biological shape and visual science (part 1). *Journal of Theoretical Biology, 38*, 205–287.

Brainard, D. (1997). The psychophysics toolbox. *Spatial Vision, 10*, 433–436.

Braun, J. (1999). On the detection of salient contours. *Spatial Vision, 12*, 211–225.

Burbeck, C., & Pizer, S. (1995). Object representation by cores: Identifying and representing primitive spatial regions. *Vision Research, 35*, 1917–1930.

Chan, T., & Vese, L. (1999). An active contour model without edges. In: M. Nielsen, P. Johansen, O. Olsen, & J. Weickert (Eds.), *Scale-space theories in computer vision*, Vol. 1682 (pp. 141–151). Berlin Heidelberg: Springer.

Chan, T., & Vese, L. (2001). Active contours without edges. *IEEE Transactions on Image Processing, 10*(2), 266–277.

Cheatham, P. (1952). Visual perceptual latency as a function of stimulus brightness and contour shape. *Journal of Experimental Psychology, 43*(5), 369–380.

Claessens, P. M. E., & Wagemans, J. (2008). A Bayesian framework for cue integration in multistable grouping: Proximity, collinearity, and orientation priors in zigzag lattices. *Journal of Vision, 8*(7), 1–23.

Cover, T. M., & Thomas, J. A. (1991). *Elements of information theory*. New York: John Wiley.

Elder, J., & Goldberg, R. M. (2002). Ecological statistics of Gestalt laws for the perceptual organization of contours. *Journal of Vision, 2*, 324–353.

Elder, J., & Zucker, S. (1993). The effect of contour closure on the rapid discrimination of two-dimensional shapes. *Vision Research, 33*, 981–991.

Engstrand, R., & Moeller, G. (1962). The relative legibility of 10 simple geometric-figures. *American Psychologist, 17*(6), 386.

Ernst, U. A., Mandon, S., Schinkel-Bielefeld, N., Neitzel, S. D., Kreiter, A. K., & Pawelzik, K. R. (2012). Optimality of human contour integration. *PLoS Computational Biology, 8*(5), e1002520.

Fang, F., Kersten, D., & Murray, S. O. (2008). Perceptual grouping and inverse fMRI activity patterns in human visual cortex. *Journal of Vision, 8*(7).

Feldman, J. (1997). Curvilinearity, covariance, and regularity in perceptual groups. *Vision Research, 37*(20), 2835–2848.

Feldman, J. (2001). Bayesian contour integration. *Perception, & Psychophysics, 63*, 1171–1182.

Feldman, J., & Singh, M. (2005). Information along contours and object boundaries. *Psychological Review, 112*(1), 243–252.

Feldman, J., & Singh, M. (2006). Bayesian estimation of the shape skeleton. *Proceedings of the National Academy of Sciences, 103*(47), 18014–18019.

Feldman, J., Singh, M., Briscoe, E., Froyen, V., Kim, S., & Wilder, J. (2013). An integrated Bayesian approach to shape representation and perceptual organization. In S. Dickinson & Z. Pizlo (Eds.), *Shape perception in human and computer vision: An interdisciplinary perspective*. Springer Verlag.

Field, D., Hayes, A., & Hess, R. (1993). Contour integration by the human visual system: Evidence for a local association field. *Vision Research, 33*(2), 173–193.

Field, D., Hayes, A., & Hess, R. (2000). The roles of polarity and symmetry in the perceptual grouping of contour fragments. *Spatial Vision, 13*(1), 51–66.

Fulvio, M., Jacqueline & Singh, M. (2006). Surface geometry influences the shape of illusory contours. *Acta Psychologica, 123*, 20–40.

Garrigan, P. (2012). The effect of contour closure on shape recognition. *Perception, 41*(2), 221–235.

Geisler, W. S. (2008). Visual perception and the statistical properties of natural scenes. *Annual Review of Psychology, 59*, 167–192.

Geisler, W. S., Perry, J. S., Super, B. J., & Gallogly, D. P. (2001). Edge co-occurrence in natural images predicts contour grouping performance. *Vision Research, 41*, 711–724.

Geisler, W. S., & Super, B. J. (2000). Perceptual organization of two-dimensional patterns. *Psychological Review, 107*(4), 677–708.

Guttman, M. A., McVeigh, E. R., & Prince, J. (1992). Contour estimation in tagged cardiac magnetic resonance images. In: *Engineering in medicine and biology society, 1992 14th annual international conference of the IEEE*, Vol. 5 (pp. 1856–1857).

Guttman, M. A., Prince, J. L., & McVeigh, E. R. (1994). Tag and contour detection in tagged mr images of the left ventricle. *IEEE Transactions on Medical Imaging, 13*(1), 74–88.

Harrison, S. J., & Feldman, J. (2009). The influence of shape and skeletal axis structure on texture perception. *Journal of Vision, 9*(6).

He, D., Kersten, D., & Fang, F. (2012). Opposite modulation of high- and low-level visual aftereffects by perceptual grouping. *Current Biology, 22*(1), 1040–1045.

Hess, R., & Field, D. (1995). Contour integration across depth. *Vision Research, 35*(12), 1699–1711.

Hess, R., & Field, D. (1999). Integration of contours: New insights. *Trends in Cognitive Sciences, 3*(12), 480–486.

Hochberg, J., Gleitman, H., & Macbride, P. (1948). Visual threshold as a function of simplicity of form. *American Psychologist, 3*, 341–342.

Hochstein, S., & Ahissar, M. (2002). View from the top: Hierarchies and reverse hierarchies in the visual system. *Neuron, 36*, 791–804.

Hung, C.-C., Carlson, E. T., & Connor, C. E. (2012). Medial axis shape coding in macaque inferotemporal cortex. *Neuron, 74*(6), 1099–1113.

Kass, M., Witkin, A., & Terzopoulos, D. (1988). Snakes: Active contour models. *International Journal of Computer Vision, 1*(4), 321–331.

Kincaid, W., Blackwell, H., & Kristofferson, A. (1960). Neural formulation of the effects of target size and shape upon visual detection. *Journal of the Optical Society of America, 50*(2), 143–146.

Koffka, K. (1935). *Principles of gestalt psychology*. New York: Harcourt, Brace, & World.

Kovacs, I., & Julesz, B. (1993). A closed curve is much more than an incomplete one: Effect of closure in figure-ground segmentation. *Proceedings of the National Academy of Sciences, 90*(16), 7495–7497.

Kovacs, I., & Julesz, B. (1994). Perceptual sensitivity maps within globally defined visual shapes. *Nature, 370*, 644–646.

Krauskopf, J., Duryea, R., & Bitterman, M. (1954). Threshold for visual form: Further experiments. *The American Journal of Psychology*, 427–440.

Lee, T. S., Mumford, D., Romero, R., & Lamme, V. A. (1998). The role of the primary visual cortex in higher level vision. *Vision Research, 38*, 2429–2454.

Lescroart, M. D., & Biederman, I. (2012). Cortical representation of medial axis structure. *Cerebral Cortex, 23*(3), 629–637.

Leymarie, F., Kimia, B., & Giblin, P. (2004). Towards surface regularization via medial axis transitions. In: *Proceedings of the 17th international conference on pattern recognition, 2004. ICPR 2004*, Vol. 3 (pp. 123–126).

Li, W., & Gilbert, C. D. (2002). Global contour saliency and local colinear interactions. *Journal of Neurophysiology, 88*(5), 2846–2856. http://dx.doi.org/10.1152/jn.00289.2002.

Mardia, K. V. (1972). *Statistics of directional data*. London: Academic Press.

Mathes, B., & Fahle, M. (2007). Closure facilitates contour integration. *Vision Research, 47*(6), 818–827.

Murray, S. O., Kersten, D., Olshausen, B. A., Schrater, P., & Woods, D. L. (2002). Shape perception reduces activity in human primary visual cortex. *Proceedings of the National Academy of Sciences, 99*(23), 15164–15169.

Murray, S. O., Schrater, P. R., & Kersten, D. (2004). Perceptual grouping and the interactions between visual cortical areas. *Neural Networks, 17*(5–6), 695–705.

Myung, I. J. (2000). Importance of complexity in model selection. *Journal of Mathematical Psychology, 44*, 190–204.

Navon, D. (1977). Forest before trees: The precedence of global features in visual perception. *Cognitive Psychology, 9*, 353–383.

Pelli, D. (1997). The videotoolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision, 10*, 437–442.

Pettet, M. W., McKee, S. P., & Grzywacz, N. M. (1998). Constraints on long range interactions mediating contour detection. *Vision Research, 38*(6), 865–879.

Pizlo, Z., Salach-Golyska, M., & Rosenfeld, A. (1997). Curve detection in a noisy image. *Vision Research, 37*(9), 1217–1241.

Psotka, J. (1978). Perceptual processes that may create stick figures and balance. *Journal of Experimental Psychology, 4*(1), 101–111.

Saarinen, J., & Levi, D. M. (1999). The effect of contour closure on shape perception. *Spatial Vision, 12*(2), 227–238.

Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal, 27*, 379–423, pp. 623–656.

Siddiqi, K., Shokoufandeh, A., Dickinson, S., & Zucker, S. (1998). Shock graphs and shape matching. In: *Sixth international conference on computer vision, 1998* (pp. 222–229).

Singh, M., & Feldman, J. (2013). Principles of contour information: A response to Lim and Leek (2012). *Psychological Review, 119*, 678–683.

Singh, M., & Fulvio, J. M. (2005). Visual extrapolation of contour geometry. *Proceedings of the National Academy of Sciences of the United States of America, 102*(3), 939–944.

Singh, M., & Fulvio, J. M. (2007). Bayesian contour extrapolation: Geometric determinants of good continuation. *Vision Research, 47*(6), 783–798.

Singh, M. (2015). Visual representation of contour and shape. In J. Wagemans (Ed.), *Oxford handbook of perceptual organization*. Oxford: Oxford University Press.

Tversky, T., Geisler, W. S., & Perry, J. (2004). Contour grouping: Closure effects are explained by good continuation and proximity. *Vision Research, 44*, 2769–2777.

Wilder, J., Feldman, J., & Singh, M. (2011). Superordinate shape classification using natural shape statistics. *Cognition, 119*(3), 325–340.

Wilder, J., Singh, M., & Feldman, J. (2015). Contour complexity and contour detection. *Journal of Vision, 6*(6), 1–16.

Yuan, C., Lin, E., Millard, J., & Hwang, J.-N. (1999). Closed contour edge detection of blood vessel lumen and outer wall boundaries in black-blood MR images. *Magnetic Resonance Imaging, 17*(2), 257–266.

Yuille, A. L., Fang, F., Schrater, P. R., & Kersten, D. (2004). Human and ideal observers for detecting image curves. In: *Neural information processing systems (NIPS)*.

Zusne, L. (1970). *Visual perception of form*. New York: Academic Press.