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

A R T I C L E   I N F O

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

A B S T R A C T

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.

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/

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

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

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

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

(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.