Vision Research 126 (2016) 207–219

Contents lists available at ScienceDirect

Vision Research

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

Quantifying density cues in grouping displays

Bart Machilsen, Johan Wagemans, Maarten Demeyer *

Laboratory of Experimental Psychology, Brain & Cognition, University of Leuven, Belgium

ARTICLE INFO

Article history:
Received 6 December 2014
Received in revised form 1 May 2015
Available online 9 July 2015

Keywords:
Contour integration
Proximity
Local density
Ideal observer
Gabor displays
GERT

ABSTRACT

Perceptual grouping processes are typically studied using sparse displays of spatially separated elements. Unless the grouping cue of interest is a proximity cue, researchers will want to ascertain that such a cue is absent from the display. Various solutions to this problem have been employed in the literature; however, no validation of these methods exists. Here, we test a number of local density metrics both through their performance as constrained ideal observer models, and through a comparison with a large dataset of human detection trials. We conclude that for the selection of stimuli without a density cue, the *Voronoi* density metric is preferable, especially if combined with a measurement of the distance to each element's nearest neighbor. We offer the entirety of the dataset as a benchmark for the evaluation of future, possibly improved, metrics. With regard to human processes of grouping by proximity, we found observers to be insensitive to target groupings that are more sparse than the surrounding distractor elements, and less sensitive to regularity cues in element positioning than to local clusterings of target elements.

1. Introduction

The human visual system is capable of grouping spatially separated image features in the retinal input into coherent perceptual units. Empirical studies into the mechanisms underlying this process often employ stimuli composed of a large number of small image elements, that together elicit a global interpretation of the stimulus. For instance, in a typical contour integration stimulus (Field, Hayes, & Hess, 1993) a subset of collinearly aligned Gabor elements is perceived as a continuous path surrounded by randomly oriented distractor elements. However, since grouping by *collinearity* is the subject of such a study, the placement of contour and background elements should be precisely controlled, in order to avoid an additional positional cue to the location of the embedded contour.

The problem is illustrated in Fig. 1. The left panel shows an array of oriented Gabor patches with an embedded open–ended contour. The percept of a smooth contour against a randomly oriented background appears to be due to the collinearity of neighboring Gabor patches. But, despite appearances, an additional cue besides the local alignment could contribute to the detection of the smooth contour here. In the right panel all local orientation information has been removed, and the path remains detectable through proximity or *local density* cues to the location of the contour – either via explicit grouping mechanisms, or via a simple low-pass spatial frequency filter connecting nearby image elements. To eliminate the unwanted cue, the distribution of element spacings around contour and background elements needs to be controlled. In the literature, various methods have been proposed.

The **grid method** distributes all display elements evenly according to a latent grid. Field, Hayes, and Hess (1993), for instance, divided the stimulus display in squares corresponding to the desired eventual spacing between the contour elements. After the contour elements have been placed, background elements are inserted at a random location within each still empty grid cell. However, the use of such a grid does not by itself prevent systematic differences in the density profiles of contour and background elements (Braun, 1999). Refinements of the grid method have been suggested to remove the residual density differences. For instance, Nygård, Van Looy and Wagemans (2009) sampled the contour and background element locations from a shared distribution, and Braun (1999) perturbs the initial grid configuration by a process simulating the diffusion of particles in a liquid.

The grid method of element placement has been criticized by Dakin and Baruch (2009), who argued that more uniform element densities can be obtained through pseudo-random placement of the elements. In this approach, contour elements are first placed as desired and background elements are then randomly added with a fixed minimum distance from previously placed elements (e.g., Kovács & Julesz, 1994; Mijović et al., 2014; Sassi et al., 2014; Watt, Ledgeway, & Dakin, 2008). We have implemented this **minimal distance method** of element placement in our previously

* Corresponding author at: Tiensestraat 102, 3000 Leuven, Belgium.
E-mail address: maarten.demeyer@ppw.kuleuven.be (M. Demeyer).
URL: https://github.com/mpjdem/densitycue (code/data repository).

http://dx.doi.org/10.1016/j.visres.2015.06.004
0042-6989/

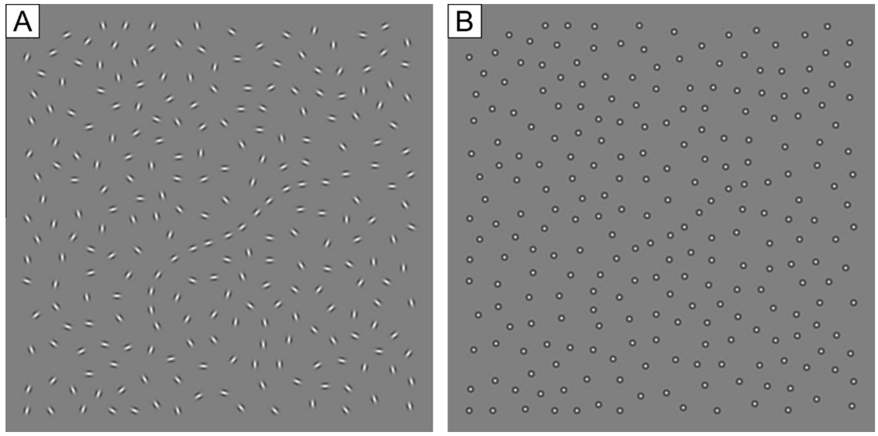

Fig. 1. Local density cues in perceptual grouping displays. (A) The alignment of neighboring oriented Gabor elements gives rise to the percept of a smooth contour. (B) However, without orientation information the contour is still visible, because the neighboring element distances are distributed differently between contour and background elements.

released toolbox GERT (Grouping Elements Rendering Toolbox; Demeyer & Machilsen, 2012).

However, the element placement method is only the first step of local density control, since even the best random placement method might by chance result in a display containing a local density cue. This is where we found the literature to be lacking a consistent quantitative approach to the problem. Many researchers in the past have relied on visual inspection of the stimuli, or on visual inspection of density statistics computed from the displays (Braun, 1999). This lack of a stringent control for local density cues is surprising given the well-known strength of proximity as a grouping cue (Elder & Goldberg, 2002; Kubovy & Wagemans, 1995; Wertheimer, 1923), and the human sensitivity to local density variations (Kovács, 2000; Tripathy, Mussap, & Barlow, 1999).

One systematic approach can be found in the work of Kovács (2000), Kovács and Julesz (1993, 1994)and Kovács et al. (1999), where the ratio D of the average background to average contour element distance is related to contour integration performance. These authors illustrate that density differences become irrelevant to human observers when $D < 1$, that is, where the contour is more sparse in element density than the surrounding background elements.

In the GERT toolbox (Demeyer & Machilsen, 2012), we instead implemented methods to completely eliminate the density cue, not only to human observers but also to ideal observers implemented in a computer algorithm. This generic framework offers quantitative decision criteria, applicable to various density metrics found in the literature. However, until now these metrics lacked validation. In the present study, we aimed to develop procedures for the quantitative evaluation of these and other local density metrics, and formulate recommendations for researchers in the field that are supported by empirical evidence. Three types of metrics are implemented in GERT.

First, the *Voronoi* metric, inspired by the methods used in Dakin and Baruch (2009), starts from a tessellation of the image space in polygon regions such that (a) each polygon contains only one element, and (b) all points within a polygon are closer to that element than to any other element in the display (Fig. 2(A)). The surface area of a polygon is then inversely related to the local density of the element contained within that polygon. To detect the presence of a density cue in the display, we statistically compare the distribution of contour element local densities to those of background elements (see Section 2).

Second, the *AvgDist* metric computes the average distance from each element to its n nearest neighbors. Below, we evaluate the performance of this metric for different values of n (see 2.6). In Fig. 2(B) the *AvgDist* metric is illustrated for the 'natural' neighbors of an element, as defined by a Delaunay triangulation of the image space (neighboring Voronoi cells; see Mathes & Fahle, 2007). This can be considered a parameter-free implementation of the *AvgDist* metric.

Third, the *RadCount* metric counts the number of neighboring elements within a circle of radius r centered on each element of the display (Fig. 2(C)). This number, normalized by the area of the circle, is then taken as the local density measurement of each element. GERT (1.3+) also implements a parameter-free variant of the *RadCount* metric, where a range of radii is tested (Braun, 1999). The sum of the absolute mean differences between contour and background, across all radii, is then used as the local density measure (see Section 2.6).

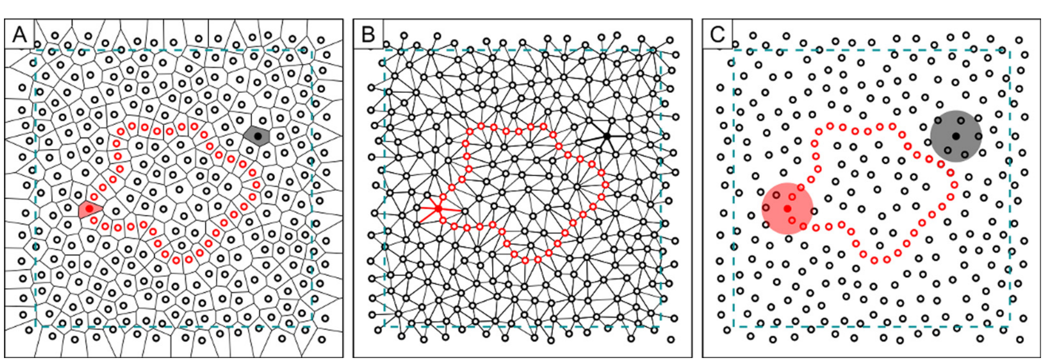

Fig. 2. Schematic of the three types of local density metrics evaluated in this study. (A) *Voronoi*. (B) *AvgDist*. (C) *RadCount*. See main text for more details.