



Good continuation in dot patterns: A quantitative approach based on local symmetry and non-accidentalness



José Lezama^{a,b,*}, Gregory Randall^b, Jean-Michel Morel^a, Rafael Grompone von Gioi^a

^a CMLA, ENS Cachan, France

^b IIE, Universidad de la República, Uruguay

ARTICLE INFO

Article history:

Received 30 January 2015

Received in revised form 18 June 2015

Accepted 18 September 2015

Available online 1 October 2015

Keywords:

Gestalt

Good continuation

Dots

Non-accidentalness

Local symmetry

ABSTRACT

We propose a novel approach to the grouping of dot patterns by the good continuation law. Our model is based on local symmetries, and the non-accidentalness principle to determine perceptually relevant configurations. A quantitative measure of non-accidentalness is proposed, showing a good correlation with the visibility of a curve of dots. A robust, unsupervised and scale-invariant algorithm for the detection of good continuation of dots is derived. The results of the proposed method are illustrated on various datasets, including data from classic psychophysical studies. An online demonstration of the algorithm allows the reader to directly evaluate the method.

© 2015 Elsevier Ltd. All rights reserved.

1. Introduction

The Gestalt school of psychology (Wertheimer, 1923; Metzger, 1975; Kanizsa, 1979; Wagemans, Elder, et al., 2012; Wagemans, Feldman, et al., 2012) proposed the existence of a short list of grouping laws governing visual perception. Among them, the law of *good continuation* can be stated as “All else being equal, elements that can be seen as smooth continuations of each other tend to be grouped together” (Palmer, 1999, p.259). Fig. 1 exemplifies this law; a perceptual organization of this image would result in a three part configuration: a line, an arc of circle, and a zigzag, all formed by dots. Unfortunately, this law, as the other Gestalt laws, was enunciated only qualitatively, without a formalization into a predictive framework.

Since it was first enunciated by Wertheimer (1923), the Gestalt law of good continuation has been extensively studied in vision science. Various aspects of this law have been examined, including amodal completion and contour integration of basic oriented and unoriented elements. In this work we concentrate on the case of unoriented elements.

The advantage of working with unoriented elements (such as dots), is that the grouping and masking processes are not influenced by the appearance of the basic elements. Experiments using dot patterns have been explored by many works in psychophysics.

Notably in the early works of French (1954), Uttal, Bunnell, and Corwin (1970) and Uttal (1973) but also in more recent works (Mussap & Levi, 2000). Uttal studied the influence of length, dot spacing, curvature and outlier noise in the perceptual grouping of dot structures. Regularity in dot patterns was analyzed by Feldman (1997b) and Kubovy, Holcombe, and Wagemans (1998). In the famous work of Glass (1969), superposed random dot patterns were used to show the importance of local interactions in the building of global percepts.

The particular case of perceptual grouping by good continuation in dot patterns is also covered in a vast literature. Prinzmetal and Banks (1977) use dot patterns to prove the existence of the good continuation phenomenon. Early algorithmic proposals for modeling the phenomenon (Caelli, Preston, & Howell, 1978; van Oeffelen & Vos, 1983; Smits, Vos, & Van Oeffelen, 1984; Smits & Vos, 1986) consisted in convolving the dot patterns with a Gaussian kernel to detect groups of dots by thresholding the result of this convolution. This idea, formalized in the CODE algorithm of van Oeffelen and Vos (1983), found relative success and was further extended in psychophysics (Compton & Logan, 1993; Logan, 1996) and also in computer vision, as will be discussed below. More recent approaches analyze the curvature of the curves generated by successive dots. In Feldman (1997a) the probabilistic properties of successive angles in a perceived chain of dots was studied. Pizlo, Salach-Golyska, and Rosenfeld (1997) proposed a clever pyramidal system to account for local–global interactivity.

* Corresponding author at: CMLA, ENS Cachan, France.

E-mail address: lezama@cmla.ens-cachan.fr (J. Lezama).

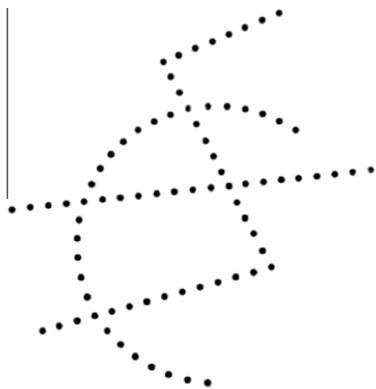


Fig. 1. Good Continuation law: human perception tends to group elements on a smooth, continuous order. Image extracted from Kanizsa (1980).

Another approach to the study of the good continuation law is through the following simple experiment: dots are arranged along a virtual circular contour that produces a linear interpolation when the number of points is small, and a curvilinear interpolation when the number is large (van Assen & Vos, 1999). The exact number of dots needed to pass from one type of interpolation to the other is studied with psychophysical experiments. In a similar line, Gori and Spillmann (2010) take a collinear arrangement of dots, and modify the spacing between given pairs, studying the boundary between the perception of an irregular alignment and its splitting into multiple segments.

The law of good continuation applied to oriented elements has also been extensively studied in psychophysics, notably in Field, Hayes, and Hess (1993), who use Gabor patterns (Machilsen & Wagemans, 2011; Demeyer & Machilsen, 2012) or line segments (Feldman, 2007) as the oriented elements.

In the computational domain, many algorithms inspired by the good continuation law have been developed. Early proposals tried to define a global vision mechanism detecting multiple Gestalts (Grossberg & Mingolla, 1987; Carpenter & Grossberg, 1987). Sha'asua and Ullman (1988) applied, in a remarkable early algorithm, the good continuation law to identify salient image features. A saliency map is obtained by iterative local computations on the image edge elements that minimize an energy favoring smooth and long curves. Parent and Zucker (1989) proposed a rigorous and clever approach to inferring curves as a labeling problem, based on local interactions that favor co-circularity (which is a form of local symmetry, a notion that our method also exploits). The image elements are convolved with oriented filters formed with Gaussians, to determine tentative tangent orientations. In Gigus and Malik (1991), the convolution with oriented Gaussians is also exploited, simply taking the maximum filter responses. This has the algorithmic advantage of being non-iterative, although performance might be affected. In Herault and Horaud (1993) the problem of figure-ground segmentation is posed from a combinatorial optimization perspective, and it is solved with simulated annealing. Another groundbreaking work is the Tensor Voting approach introduced by Guy and Medioni (1993), which proposes a saliency measure that involves the summation of vector votes emitted by each element. The votes encourage co-circularity and proximity of the elements. Unlike the other approaches, this framework specifically incorporates the detection of curves of unoriented elements. Perona and Freeman (1998) pose the figure-ground segmentation as a factorization of a matrix representing the affinity between elements. Williams and Thorner (1999) provide a very good review of existing approaches and introduce a new one, where the saliency measure is given by the number of times a random walk passes by an edge, and where the transition

probability matrix is also given by an affinity matrix. In general, all of these methods are based on finding combinations of local interactions that favor curve smoothness, length and elements proximity.

In this work, we propose a new model and algorithm for the grouping by good continuation (restricted to unoriented elements) using a simple model that favors local symmetries, and with a detection control based on the non-accidentalness principle. This allows the method to be general in the sense that it can capture smooth curves of any shape and scale, and is robust to the presence of outliers and noise. It is also unsupervised because detections are given by their statistical significance, which requires only a single parameter, namely the number of false detections that would be allowed in an image of random noise. In an earlier work, a preliminary version of the algorithm was introduced to the image processing community (Lezama, Grompone von Gioi, Randall, & Morel, 2014). Here we present a reformulated theory to achieve scale invariance and we establish a link with psychophysics.

The proposed algorithm consists of two main steps: building candidate chains of points, and validating them. Candidate chains of points are built by considering triplets of points formed by joining nearest neighbors. Once valid triplets have been obtained, a graph representation is produced where each node corresponds to a triplet. A classic path finding algorithm is run on this graph to obtain paths between all pairs of triplets. Finally, the paths found are validated as non-accidental or rejected using thresholds obtained with the *a contrario* approach (Desolneux, Moisan, & Morel, 2008). It will be shown that the number of false alarms (NFA) defined in that theory provides an effective measure of the meaningfulness of a good continuation configuration. The potential use of the method is evaluated on data from classic psychophysical experiments and image processing applications.

This article is organized as follows: Section 2 presents our proposed mathematical model of good continuation chains. Section 3 describes an efficient algorithm for detecting good continuation configurations in dot patterns. The mathematical model and the algorithm are then evaluated in Section 4. Finally, Section 5 presents the conclusions of this study.

2. Mathematical model

Let us consider a set of N planar points. The aim is to find a mathematical model that can predict when an ordered subset of points lie on a smooth curve that is salient relative to the background of the other points, see Fig. 2(a). Each ordered subset of points (a sequence of points) will be called a *chain*; each set of three consecutive points in a chain will be called a *triplet*. The proposed model is based on the simple idea that the better the symmetry of the triplets, the better the saliency of the sequence. Ideally, the third point in a triplet should be symmetric to the first, relative to the middle point, in the position marked with an X on Fig. 2(b). Symmetric triplets also enforce a second Gestalt grouping law: *proximity* (Wagemans, Elder, et al., 2012). A regular spacing of the points along the curve favors their perceptual grouping; inversely, an irregular spacing (Wertheimer, 1923; Gori & Spillmann, 2010) would tend to stop the curve at larger gaps.

To obtain a perceptually plausible model, the evaluation of a chain of points will be based on the *non-accidentalness principle*, proposed as the rationale underlying perceptual thresholds (Witkin & Tenenbaum, 1983; Rock, 1983; Lowe & Binford, 1981; Albert & Hoffman, 1995). In a nutshell, an observed structure is relevant if it would rarely occur by chance. Quoting D. Lowe, “we need to determine the probability that each relation in the image could have arisen by accident, $P(a)$. Naturally, the smaller that this value is, the more likely the relation is to have a causal

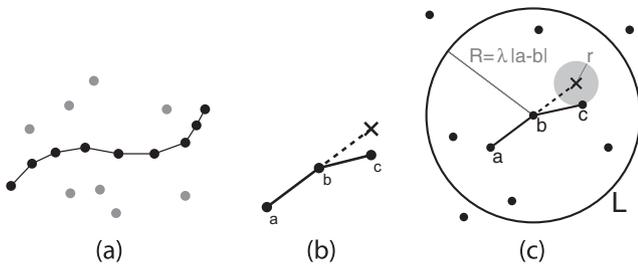


Fig. 2. Definition of the *good continuation* event. (a) A candidate *chain* is defined by an ordered sequence of points. (b) Three consecutive points in a chain define a *triplet*, (a, b, c) in this case. Ideally, the triplet should be symmetric. That is, the third point c should be symmetric to the first point a , relative to the middle point b . The ideal third point is represented by X . (c) The symmetry precision of a triplet is measured by the distance r from the third point c to the ideal point X . The scale-invariant error is expressed as the probability that, among the n points in the local window of radius R , the nearest point to X be at most at a distance r .

interpretation” (Lowe, 1985, p.39). The *a contrario* framework (Desolneux et al., 2008), a formalization of this principle, is used to provide automatic detection thresholds, compatible with perception, and to handle noise points. Given a random model for the data, the *a contrario* methodology consists in evaluating the expectation of the occurrence of an error as small as the one observed, relative to an ideal structure. If this expectation is small, the event is termed non accidental and thus perceptually meaningful. This methodology has already been applied in several computer vision problems. Among them, there was a previous attempt at detecting *good continuations* of image edges (Cao, 2004); even if related, the main source of information for this detector was the orientation of the edges, while no orientation is associated to points in our case.

We will evaluate the probability of observing a chain of points where all of its triplets have a given degree of symmetry. This probability is evaluated in a random background model assuming that the points in the image were randomly distributed. The imperfection of a triplet translates into the distance r between the observed third point and its ideal symmetric position X , giving the local context. To provide scale-invariance, this error will be evaluated relatively to the context contained in a circular local window L with radius R , where $R = \lambda \cdot |a - b|$ is proportional to the triplet size, see Fig. 2(c).

Given that n points were observed in L (not counting the first two of the triplet, as they define the local window of the first triplet), our random or *a contrario* model H_0 , used to evaluate accidentalness, is that these points are independent and uniformly distributed in L . In other words, our *a contrario* model assumes that the n points result from a spatial uniform Poisson process in L . This *a contrario* model H_0 is not intended to model the statistics of the sought structure; quite the opposite, it models random data where the sought structure is not present, and is used to calibrate rejection thresholds.

Under these assumptions, we would like to translate the error of each triplet into probabilistic terms. Let us call ρ the distance between the ideal point X and its nearest point in L under H_0 . We will evaluate the precision of a triplet by the probability $\mathbb{P}(\rho \leq r)$, for the observed radius r . It is simpler to compute its complement, $\mathbb{P}(\rho > r)$, which implies that all the n points in L fall outside the disk of radius r ; given that L is a disk of radius R ,

$$\mathbb{P}(\rho > r) = \left(1 - \frac{\pi r^2}{\pi R^2}\right)^n. \text{ Finally, the error associated to a triplet is}$$

$$e = \mathbb{P}(\rho \leq r) = 1 - \left(1 - \frac{r^2}{R^2}\right)^n. \tag{1}$$

This quantity would be zero only for the ideal symmetric triplet ($r = 0$), and a small value corresponds to a near symmetric triplet. If only one point is observed in the local window L (not counting

the points a and b that define the triplet), this error measures the deviation of the triplet from the ideal symmetric configuration. When more points are present, the error grows, reflecting the lower relevance of the triplet due to crowding. Note that r only appears in the ratio $\frac{r}{R}$, thus making the measure independent of its absolute value and ensuring the scale-invariance of the criterion.

Consider a chain C of k points a_1, a_2, \dots, a_k . The error e_i of each of the $k - 2$ triplets (a_i, a_{i+1}, a_{i+2}) can be evaluated by Eq. (1), and the worst case value, $e_{\max} = \max\{e_1, e_2, \dots, e_{k-2}\}$, is associated to the whole chain. The event we are considering is a chain C of $k - 2$ triplets, each with error e_{\max} or less relative to the ideal symmetric triplet. We will evaluate now the probability of this event under H_0 .

Let us consider a random triplet under H_0 , with the third point selected as the nearest one to the ideal symmetric point. By construction of the error measure, when Eq. (1) is evaluated on points following the same model H_0 , the error defined by Eq. (1) becomes a random variable E with uniform distribution in $[0, 1]$.¹ Thus, $\mathbb{P}(E \leq \alpha) = \alpha$ for any $\alpha \in [0, 1]$. We are now in a position to evaluate a random chain of $k - 2$ triplets under H_0 . The corresponding errors are E_1, E_2, \dots, E_{k-2} and the worst error is $E_{\max} = \max\{E_1, E_2, \dots, E_{k-2}\}$. Now, under the *a contrario* Poisson assumption, the probability of each triplet is independent from the previous ones, so the probability of all errors being lower than e_{\max} is

$$\begin{aligned} \mathbb{P}(E_{\max} \leq e_{\max}) &= \mathbb{P}(E_1 \leq e_{\max} \cap E_2 \leq e_{\max} \cap \dots \cap E_{k-2} \leq e_{\max}) \\ &= \mathbb{P}(E_1 \leq e_{\max}) \times \mathbb{P}(E_2 \leq e_{\max}) \times \dots \times \mathbb{P}(E_{k-2} \leq e_{\max}) \\ &= e_{\max}^{k-2}. \end{aligned} \tag{2}$$

Notice that this is not the probability of observing the exact chain C , but the probability of observing, under H_0 , chains whose triplets have all error e_{\max} or less relative to ideal symmetric triplets. This term can take the value zero only for straight chains with equally-spaced points, where each one of the triplets would be perfectly symmetric. Inversely, irregular curves necessarily have values near to one.

The fundamental quantity in the *a contrario* methodology is the number of false alarms (NFA) of an event, defined as the number of tests, times the probability of the event (Desolneux et al., 2008). The NFA for a chain of points in *good continuation* is computed as

$$NFA(C) = N_{\text{tests}} \cdot \mathbb{P}(E_{\max} \leq e_{\max}). \tag{3}$$

The NFA is an upper bound on the expected number of chains with the same error as C or smaller, to be observed by chance in the *a contrario* model H_0 . A large NFA means that such an event is to be expected under the *a contrario* model and therefore is irrelevant. On the other hand, a small NFA corresponds to a rare event and therefore arguably a meaningful one. Geometrically, a perfectly aligned and equally-spaced curve would have $NFA = 0$ and small values correspond to regular curves. On the other hand, irregular curves necessarily have large NFAs, near N_{tests} .

The number of tests N_{tests} counts the chains considered as potential good continuations. The proposed method will generate candidates starting at each of the N points and every one of its b nearest neighbors will be tried as a second point. Using these two points, the ideal point X is constructed (see Fig. 2(b)) and the closest point to it is selected as the third point of the triplet. Iteratively, the last two points of the previous triplet are used to construct the ideal point X for the next one. This process is repeated until a maximal chain length of \sqrt{N} is reached. The assumption here is that a smooth 1D subset of a 2D set of N points would be typically limited to \sqrt{N} points. Each of the intermediate chains is also tested as a potential meaningful chain. Thus, the

¹ This is due to the fact that, given a continuous random variable X with CDF F_X , the random variable $Y = F_X(X)$ is uniform in $[0, 1]$.

number of tests is $bN\sqrt{N}$. Finally, the NFA of the event “having k points in *good continuation* configuration with error e_{\max} or less” is
$$\text{NFA} = bN\sqrt{N} \cdot e_{\max}^{k-2}. \quad (4)$$

Given an observed set of N points and a candidate chain of k points, we will consider the latter event as an ε -meaningful *good continuation* when the corresponding NFA is lower than ε . It can be shown (Desolneux et al., 2008) that the expected number of events with $\text{NFA} < \varepsilon$ is bounded by ε in the *a contrario* model H_0 . This justifies the definition and name of the NFA, as it controls the average number of accidental (thus false) detections. Following Desolneux et al. (2008), we will always fix $\varepsilon = 1$ to ensure an unsupervised detection, as having less than one false detection on average is tolerable in our case. (See the evaluation section for a confirmation of this fact.)

3. Algorithm

This section describes an efficient but heuristic algorithm for searching meaningful chains of dots using the model presented in the previous section. Given an input of N planar points, the candidate chains are obtained by exploring the b nearest neighbors of each point to construct candidate triplets and then by connecting paths between every two triplets. To find the paths, a graph representation of the triplets is constructed and the Floyd–Warshall algorithm is used. The Floyd–Warshall algorithm is a classic graph analysis algorithm that finds the shortest paths between all pairs of nodes in a graph. Each path found is a candidate chain that is finally evaluated using the NFA, Eq. (4), and the most significant chains are kept using a redundancy reduction step, described below. This process is described in Algorithm 1.

Algorithm 1. *Good continuation* detection

Input: A list of N planar points, b the number of nearest neighbors used for exploration and λ the proportion of the local window radius w.r.t. the triplet size.
Output: A list of *good continuation* chains \mathcal{GC}

```

1  $\mathcal{T} \leftarrow \emptyset$ 
2 for point  $i = 1$  to  $N$  do
3   for point  $j \in \text{NearestNeighbors}_b(i)$  do
4     Count  $n$ , number of points in the local window  $L$ 
       centered in  $j$  with radius  $R = \lambda \cdot \text{dist}(i, j)$ 
5     Compute point  $s_j^i$ , symmetric to  $i$  w.r.t.  $j$ 
6     for point  $k \in \text{NearestNeighbors}_b(s_j^i)$  do
7       Compute  $e = \mathbb{P}(\rho < \text{dist}(s_j^i, k) | R, n)$  (Eq. (1))
8        $\mathcal{T} \leftarrow (i, j, k; e)$ 
9     end
10  end
11 end
12  $\mathcal{D} \leftarrow a |\mathcal{T}| \times |\mathcal{T}|$  matrix with values  $\infty$ 
13 for triplet  $s \in \mathcal{T}$  do
14   for triplet  $t \in \mathcal{T}$  do
15     if  $s$  is adjacent to  $t$  then
16        $\mathcal{D}_{s,t} \leftarrow e_s + e_t$ 
17     end
18   end
19 end
20  $\mathcal{P} \leftarrow \text{Floyd-Warshall}(\mathcal{D})$ 
21 for path  $p \in \mathcal{P}$  do
22   if  $\text{NFA}(p) < 1$  (Eq. (4)) then
23      $\mathcal{GC} \leftarrow p$ 
24   end
25 end
```

The algorithm requires two parameters: the number of nearest neighbors b used for exploration, and λ , the proportion of the local window size to a triplet’s size² (see Section 2). Lines 1 to 11 of Algorithm 1 build the list \mathcal{T} of triplets to be considered. Note that each triplet is stored with its error e .

A pair of triplets will be called *adjacent* when they share two points in such a way that they can form a chain of four points. (Triplets that share two points but form a “Y” shape are not adjacent.) We define a graph where triplets are the vertices and adjacent triplets s and t share an edge with value $e_s + e_t$, the sum of their errors (lines 12 to 19). This is a heuristic step, without any probabilistic interpretation, justified on what follows. Once the adjacencies are determined, the Floyd–Warshall algorithm is used to compute the path with shortest distances between every two vertices (line 20). Finally, all the candidate chains provided by Floyd–Warshall are evaluated for significance using Eq. (4) and the ones with $\text{NFA} < \varepsilon$ are kept (lines 21 to 25). The heuristic is necessary to be able to benefit from the efficiency of the Floyd–Warshall algorithm, which works with additive distances.

The resulting paths are the best in the sense of the smallest sum $\sum_i e_i$ along the path. Notice that the “distance” used is not Euclidean, but the sum of the triplet errors (thus the common bias toward small objects in shortest path methods is not present here). There is no theoretical guarantee that paths with minimal NFA (smallest e_{\max}^{k-2}) are all contained in minimal paths for the Floyd–Warshall algorithm. However, our simulations show that this approximation is acceptable.

The computational complexity of the Floyd–Warshall algorithm is $O(|V|^3)$, where $|V|$ is the number of vertices in the graph, i.e., the number of triplets. In terms of computation time, this is the bottleneck of the proposed algorithm. The result is a non-linear algorithm, but fast enough to perform example simulations with some thousand points in tens of seconds on a laptop computer.

Once all the *good continuation* events are found, we are interested in keeping only non-redundant detections. Note that a *good continuation* event might mask another smaller event contained in itself (e.g. a subset of the points in a meaningful chain can be also meaningful). We shall say that an event A masks an event B , if $\text{NFA}_A < \text{NFA}_B$ and the chains share at least two points. The latter is just a simple criterion to allow crossing chains, which share one point; future work will focus on this point to develop a principled criterion. Obtaining a list of only the most meaningful events, defined as those that are not masked by any other event, can be done by following the simple steps: First, the meaningful chains are ordered by their NFA (lowest first). A second list is created which in the beginning contains only the most meaningful chain. Then, the first list is traversed, checking if each chain is masked with any of the chains in the second list. If a chain is not masked, it is added to the second list. The resulting second list gives the non-redundant *good continuation* chains.

4. Evaluation

In this section we will perform a detailed analysis of the NFA obtained with our model through multiple examples, showing its applicability as a quantitative perceptual measure. Next, we will present some results using the heuristic algorithm for dot patterns taken from the *good continuation* literature. Finally, we will show its application to image analysis. The reader is invited to try the online demo of this algorithm to evaluate the method directly.³

² All the results shown in this article use $b = 5$ and $\lambda = 4$.

³ http://dev.ipol.im/jlezama/ipol_demo/lgrm_good_continuation_matlab/
 user: demo password: demo

4.1. Evaluation of the NFA on selected chains

To analyze the effect of curvature, dot density and dot regularity in the NFA of a chain, we considered the dot pattern examples of Uttal (1973). We analyzed the NFA obtained from each curve using Eq. (4). To isolate the NFA evaluation from the candidate search heuristic, the chains were specified manually; in the next subsection we shall analyze the result of the complete detection process. To demonstrate the perceptual plausibility of the NFA, we added two types of noise: a surrounding outlier noise and inlier noise implemented as a random jitter on the position of the points.

The left side of Fig. 3 shows the original figures scanned from Uttal (1973). We will refer to the first row of the original curves as dataset 1, and to the second and third rows as datasets 2 and 3, respectively. In each dataset, curves are numbered from #1 to #6. To recover the position of the dots we used a Harris corner detector (Harris & Stephens, 1988) on the scanned images. On the right side of the figure, the NFA obtained for each of the curves is shown. Note how the NFA increases as the curvature increases or as the length of the curves decreases, correctly indicating their lower meaningfulness. Minor variations in the NFA can be due to small errors in the dot positions from the scanning and dot

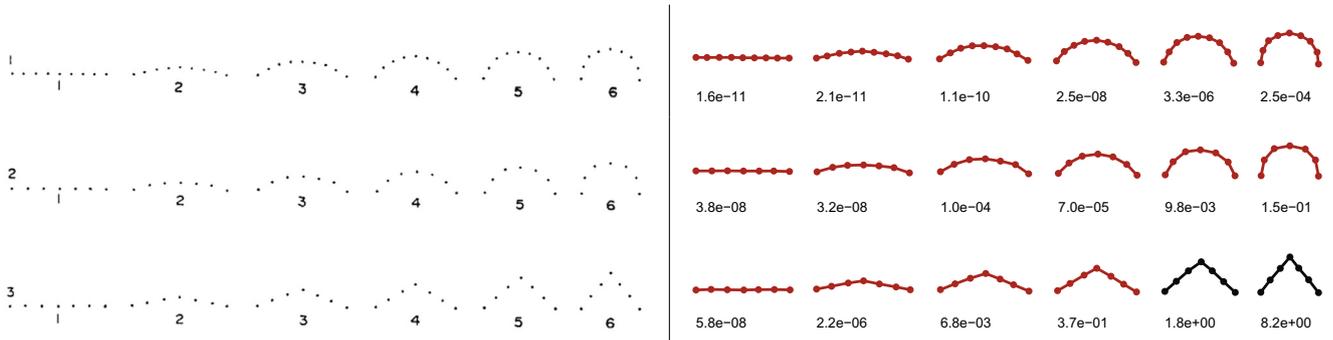


Fig. 3. NFAs for the curves used in Uttal (1973). The points have been obtained by scanning the figure and running a Harris corner detector (Harris & Stephens, 1988). On the right side, the lines indicate the manually selected path; the NFA obtained for each curve is shown in E notation. Curves with NFA value lower than 1 (red) are considered meaningful, while the ones with NFA value higher than 1 (black) are not considered meaningful. Note that the NFA associated to each curve increases as it contains less dots or the curvature of the underlying curve is bigger.

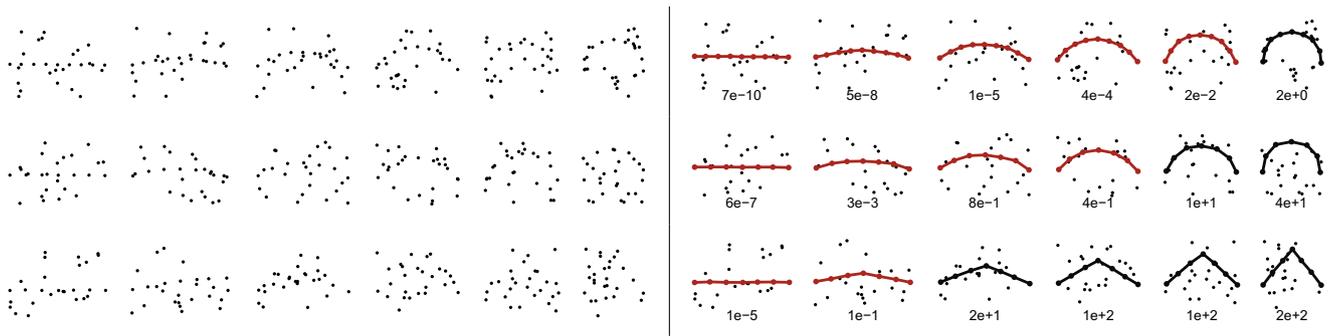


Fig. 4. NFAs for the curves of Fig. 3 plus 20 outlier noise points. On the right side, the lines indicate the manually selected path; the NFA obtained for each curve is shown in E notation. Curves with NFA value lower than 1 (red) are considered meaningful, while the ones with NFA value higher than 1 (black) are not considered meaningful. Note how the NFA increases (less meaningfulness) in the presence of noise, sometimes above the meaningfulness threshold. The NFA as a quantitative predictive measure is consistent with the perception of the curves in the figures on the left.

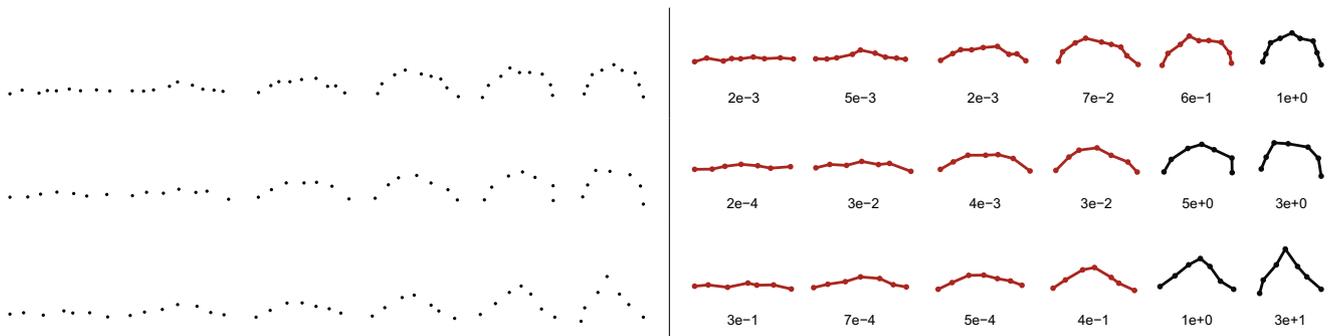


Fig. 5. NFAs for the curves of Fig. 3 with Gaussian jitter ($\sigma = 3\%$ of image width) on the points position. On the right side, the lines indicate the manually selected path; the NFA obtained for each curve is shown in E notation. Curves with NFA value lower than 1 (red) are considered meaningful, while the ones with NFA value higher than 1 (black) are not considered meaningful. The NFA is degraded because the local symmetries are deteriorated. In some cases a human observer might prefer to split the curve in two. In general, for those cases the NFA is above the meaningfulness threshold.

detection processes. If we set the meaningfulness threshold to $\varepsilon = 1$ (one false alarm in average in noise), curves #5 and #6 of dataset 3 are not meaningful. (On Figs. 3–5, the curves with $NFA < 1$ are drawn in red, while the curves with $NFA \geq 1$ are in black.) This is because the angle in the middle (which will determine the precision of the event because it forms the worst triplet) is too acute. In this case, our model would prefer to split the curve into two straight segments. This can be seen in the actual result of the algorithm in Fig. 6, which shows the most significant curves. Still, for a human observer, it is possible that a higher-level grouping process produces the junction of both segments.

Fig. 4 shows the same curves of Fig. 3 with 20 random points added to each one. The aim of this extended dataset is to show how the NFA correctly models the masking/unmasking perception process. The NFA of the curves is less meaningful in the presence of noise, but it is still meaningful where the structure can still be

perceived. On the other hand, when the noise is sufficient to mask the structure, (notably in datasets 2 and 3 where the structure has fewer points), the structure becomes statistically as well as perceptually indistinguishable from noise, and the NFA goes above 1. This observation is in line with the original conclusions of Uttal (1973), although the experiment setup is by no means the same.

The datasets of Fig. 5 aim at showing the effect on the NFA of the irregular placement of dots. Starting with the original curves of Fig. 3, we added isotropic random displacements to each dot. The random displacements are Gaussian distributed, centered at each dot and with a 4 pixels standard deviation (as a reference, inter-dot distance is approx. 20 pixels). The results show a degradation of the NFA with respect to the original curves. This was to be expected because the local symmetry is strongly violated. In particular, the results of the dataset 1, curve #6 and dataset 2,

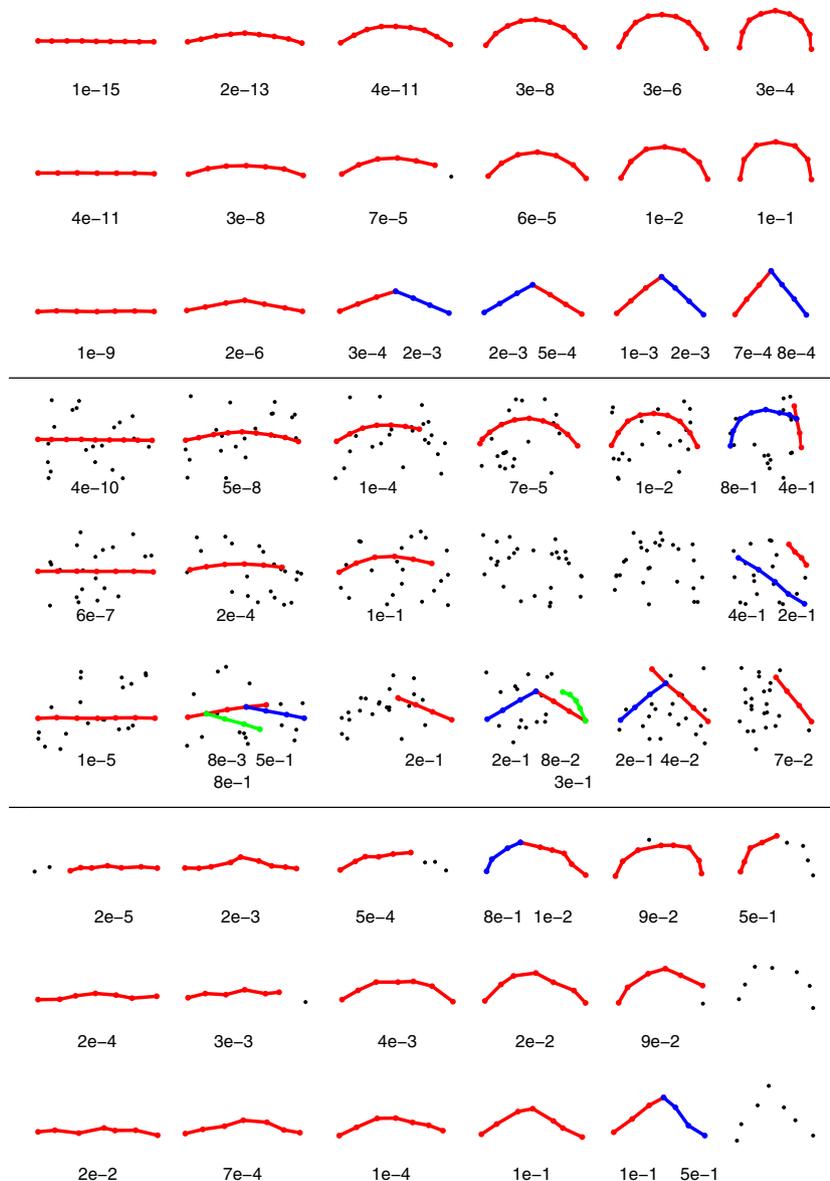


Fig. 6. Result of the algorithm for the dot patterns of Figs. 3–5. The examples are grouped by the type of noise: none, outlier and inlier noise. When the algorithm detects meaningful curves ($NFA < 1$), the NFA value is printed for each detection. In red are shown the most meaningful detections, in blue the second most meaningful ones, and in green the third most meaningful one, when exist. Note that in some cases the algorithm detects no meaningful curve. The value of the NFA (or the meaningfulness of the detected curves) are in accordance with the perceptual analysis of the dots: the bigger the (inlier or outlier) noise, the less meaningful the curve.

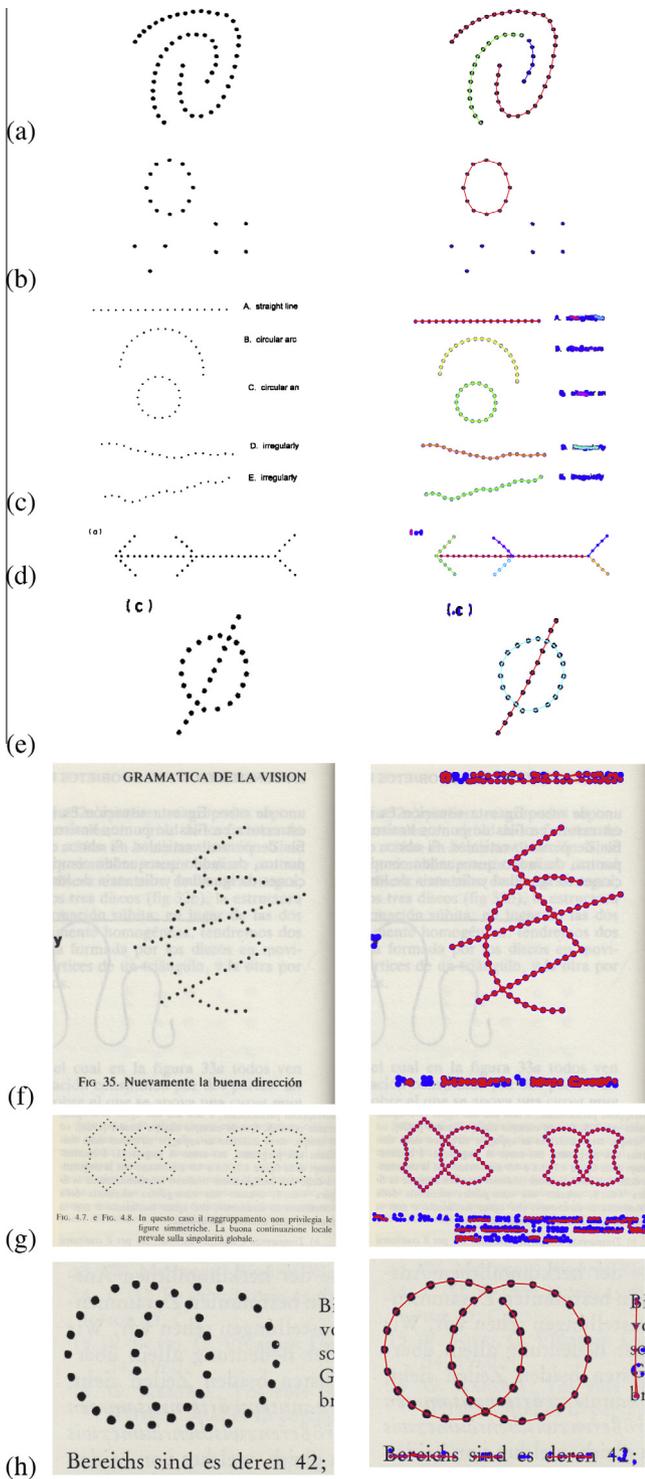


Fig. 7. Result of the good continuation detection algorithm for images scanned from the following articles: (a) & (b) van Oeffelen and Vos (1983); (c) Pizlo et al. (1997); (d) & (e) Caelli et al. (1978); (f) Kanizsa (1980); (g) Kanizsa (1991); (h) Metzger (1975). On each scanned image in the left column, a Harris corner detector was run to find the dots. The right column shows the Harris detections in blue, and the good continuation configurations found by our method as curves connecting a set of dots. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

curves #5 and #6 have now NFA values above the meaningfulness threshold of 1. When looking at those curves, a human observer might find it easier to interpret them as two separate pieces of curves instead of a single one.

4.2. Results of the heuristic algorithm

In this subsection we shall discuss the results of the complete algorithm, which includes the heuristic for searching for candidate chains and their evaluation using the NFA. Fig. 6 shows the result of the good continuation detection algorithm described in Section 3 for the dot patterns of Figs. 3–5. The figure is divided in three groups according to the type of noise: no noise, outlier noise, and inlier noise. Only the curves with $NFA < 1$ are displayed.

In the noise-free examples, the algorithm finds the original curve in most cases. One exception is dataset 2, curve #3, where the algorithm prefers the curve that is formed by leaving out the last dot. Note that points have some position noise due to the scanning and dot detection. What is happening is that the precision of the last triplet is such that the NFA would increase rather than decrease if it were included, so the most significant curve leaves that dot out. Indeed, the detection algorithm keeps only the most significant curves. In dataset 3, when the angle is too strong, the individual segments are more meaningful than the entire curve, resulting in two curves (red and blue).

In the second group of examples, where outlier noise is present, the algorithm tends to find the curve when it is still perceived. Otherwise, there are two reasons for a curve not to be detected. The first trivial reason is when the NFA is simply too large and therefore not meaningful, because the triplet's probabilities increase as more points are present in the local window (see Fig. 2(c)). The second reason is that due to the algorithm's heuristic of searching among the nearest neighbors, the curve may never be considered as a candidate, and never be evaluated. This is why for some curves the NFA as it would be calculated by an ideal observer is meaningful, but they are not detected by the algorithm. An example of this case can be found in dataset 2, curve #4. This effect is actually perceptually plausible: when there are many points in the image, the complexity of evaluating every possible combination

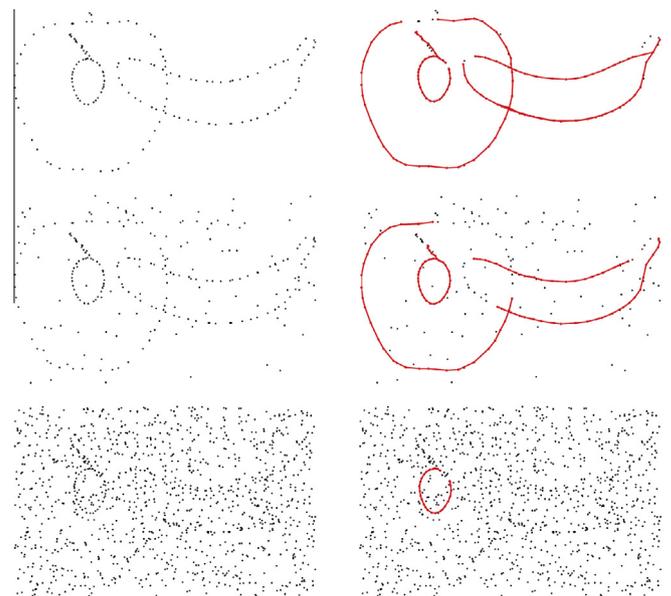


Fig. 8. Result of the algorithm using dot patterns adapted from the “yellow apple”, “banana” and “tamarillo” edge patterns from Williams and Thornber (1999). The first row shows the superposition of the three original patterns, where the orientation information was dropped and the figures were rescaled and displaced. In the second and third row 100 and 1000 random points are added, respectively. The red curves are the groupings of dots found by the algorithm. Note how the detection decreases as the noise increases, as expected. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

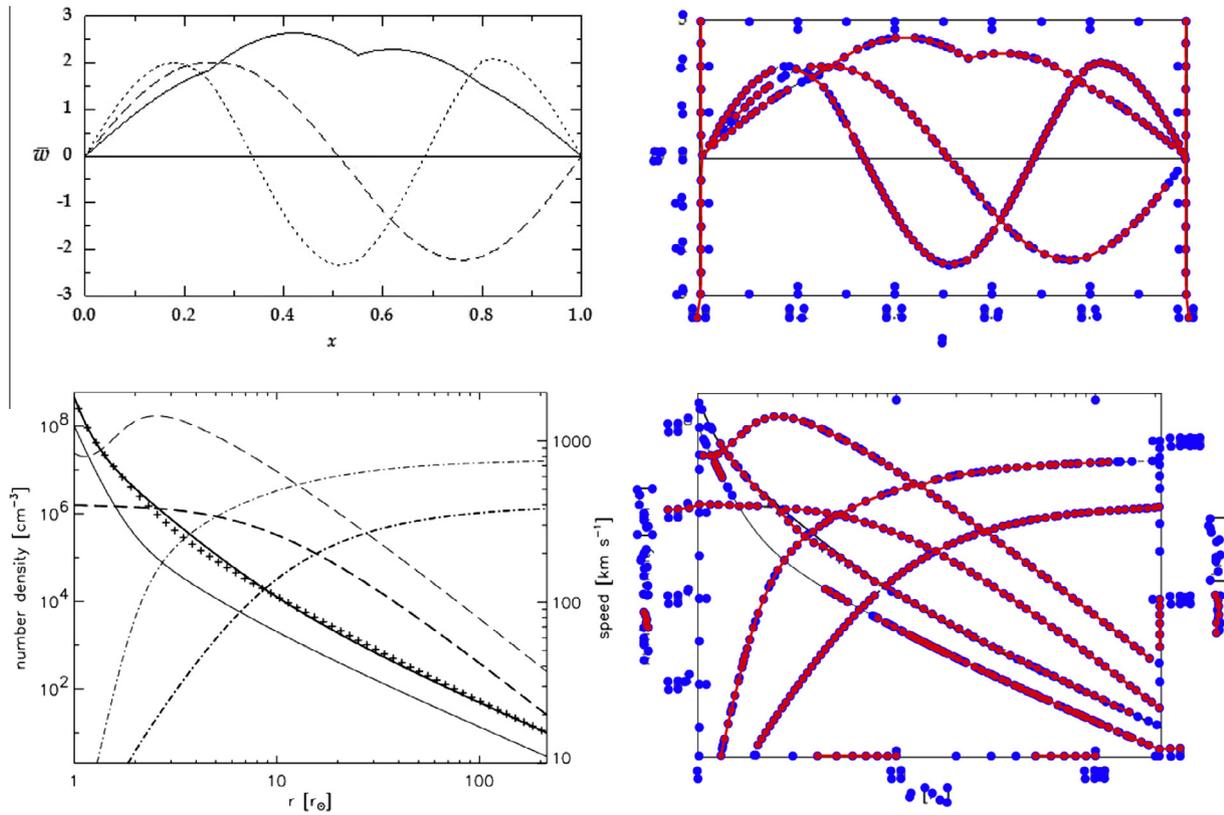


Fig. 9. A possible application for the method is the processing of scanned documents. In the left column we show graphs scanned from real documents. In the right column, Harris corner detections are represented by blue dots and curve detections in red. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

is arguably intractable for human perception. Small detections are also produced that are not part of the original curves, in dataset 1 curve #6, dataset 2 curve #6 and dataset 3 curves #2 and #4. These small detections are arguably perceptually salient when looking at the figures, but their NFA is close to 1, meaning that they are just slightly meaningful.

In the third group of examples, where random jitter is added to the points positions, some curves were split at triplets with large error. In those cases, separate segments of the curve are individually more meaningful than the whole curve. Such splittings are perceptually plausible. Another interesting case is dataset 1, curve #5, where the algorithm prefers to leave one point out of the curve. Again, looking at the input dots this interpretation seems natural.

4.3. Applicability to image analysis

This subsection illustrates the applicability of the proposed method to image and data analysis by showing results on some typical data.

Fig. 8 presents examples of detection results on shape analysis. We took three images from the fruit and vegetables dataset of Williams and Thornber (1999) and kept only the position of the oriented elements. The examples are silhouettes of an apple, a banana and a tamarillo. To demonstrate the scale invariance of our approach, we scaled each image, so they are at 1/3 and 1/6 scales, respectively. To produce a second dot pattern, we added 100 random points to the figure, to test the robustness of our approach in differentiating perceptually relevant structure from noise (second row of Fig. 8). For a final dot pattern, we added 1000 random points, which visually mask the two largest figures (third row of Fig. 8). In this case, only the contour of the tamarillo is still detected by the algorithm. The detected structures show a

good match with the perceived structures. The method handled the masking by noise automatically, as well as the different scale of the shapes.

Finally, Figs. 7 and 9 show a possible application of the method: the automatic processing of graphs from scanned documents. The input points to the algorithm were once again the points obtained by running a Harris corner detector. Spurious corner detections are also part of the input, which causes reasonable detections where text is present. Fig. 7 shows results of the algorithm in dot patterns obtained from scanning figures from popular articles in the literature. The examples in Fig. 9 illustrate the potential application to the automatic vectorization of dotted lines, by combining our detection algorithm with a Harris point detector.

5. Conclusions

We introduced a new quantitative model for the grouping of dots under the good continuation Gestalt. Our approach is a formalization of the non-accidentalness principle, based on a very simple model that favors local symmetries. This makes the model prefer smooth and long curves, where the dots are equally spaced. The model also accounts for the masking effect produced by surrounding noise points. We presented an algorithm for the detection of good continuation groupings under this model. In the evaluation section we presented through examples the theoretical limits of the model as well as the results obtained with the algorithm. These informal tests show a good match between the quantitative measure introduced and human perception.

We envision three different lines of future work. One is the incorporation of the closure Gestalt to the model. The second is the utilization of more complex models for the local interactions,

modeling curvature explicitly, and in line with local interaction fields studied in the literature. A third line is the formulation of a coarse-to-fine version of the algorithm that would model a hierarchical process of perception and reduce the required computations.

Acknowledgments

Work partly founded by the Centre National d'Etudes Spatiales (CNES, MISS Project), the European Research Council (advanced grant Twelve Labours), the Office of Naval Research (ONR Grant N00014-141-0023), DGA project Stéréo, and ANR-DGA project ANR-12-ASTR-0035.

References

- Albert, M. K., & Hoffman, D. D. (1995). Genericity in spatial vision. In D. Luce, K. Romney, D. Hoffman, & M. D'Zmura (Eds.), *Geometric representations of perceptual phenomena: Articles in honor of Tarow Indow's 70th birthday* (pp. 95–112). Erlbaum.
- Caelli, T. M., Preston, G., & Howell, E. (1978). Implications of spatial summation models for processes of contour perception: A geometric perspective. *Vision Research*, 18(6), 723–734.
- Cao, F. (2004). Application of the Gestalt principles to the detection of good continuations and corners in image level lines. *Computing and Visualization in Science*, 7, 3–13.
- Carpenter, G. A., & Grossberg, S. (1987). A massively parallel architecture for a self-organizing neural pattern recognition machine. *Computer Vision, Graphics, and Image Processing*, 37(1), 54–115.
- Compton, B. J., & Logan, G. D. (1993). Evaluating a computational model of perceptual grouping by proximity. *Perception & Psychophysics*, 53(4), 403–421.
- Demeyer, M., & Machilsen, B. (2012). The construction of perceptual grouping displays using gert. *Behavior Research Methods*, 44(2), 439–446.
- Desolneux, A., Moisan, L., & Morel, J. (2008). *From Gestalt theory to image analysis, a probabilistic approach*. Springer.
- Ellis, W. (Ed.). (1967 originally 1938). *A source book of Gestalt psychology*. Humanities Press.
- Feldman, J. (1997a). Curvilinearity, covariance, and regularity in perceptual groups. *Vision Research*, 37(20), 2835–2848.
- Feldman, J. (1997b). Regularity-based perceptual grouping. *Computational Intelligence*, 13(4), 582–623.
- Feldman, J. (2007). Formation of visual “objects” in the early computation of spatial relations. *Perception & Psychophysics*, 69(5), 816–827.
- Field, D. J., Hayes, A., & Hess, R. F. (1993). Contour integration by the human visual system: Evidence for a local “association field”. *Vision Research*, 33(2), 173–193.
- French, R. S. (1954). Pattern recognition in the presence of visual noise. *Journal of Experimental Psychology*, 47(1), 27.
- Gigus, Z., & Malik, J. (1991). *Detecting curvilinear structure in images*. University of California, Berkeley, Computer Science Division.
- Glass, L. et al. (1969). Moiré effect from random dots. *Nature*, 223(5206), 578–580.
- Gori, S., & Spillmann, L. (2010). Detection vs. grouping thresholds for elements differing in spacing, size and luminance. An alternative approach towards the psychophysics of gestalten. *Vision Research*, 50(12), 1194–1202.
- Grossberg, S., & Mingolla, E. (1987). Neural dynamics of perceptual grouping: Textures, boundaries, and emergent segmentations. *Advances in Psychology*, 43, 144–210.
- Guy, G., Medioni, G. (1993). Inferring global perceptual contours from local features. In: *Computer vision and pattern recognition, proceedings CVPR'93, 1993 IEEE computer society conference on IEEE* (pp. 786–787).
- Harris, C., Stephens, M. (1988). A combined corner and edge detector. In: *Alvey vision conference Vol. 15* (p. 50). Manchester, UK.
- Herault, L., & Horaud, R. (1993). Figure-ground discrimination: A combinatorial optimization approach. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(9), 899–914.
- Kanizsa, G. (1979). *Organization in vision: Essays on Gestalt perception*. New York: Praeger.
- Kanizsa, G. (1980). *Grammatica del vedere*. Il Mulino.
- Kanizsa, G. (1991). *Vedere e pensare*. Il Mulino.
- Kubovy, M., Holcombe, A. O., & Wagemans, J. (1998). On the lawfulness of grouping by proximity. *Cognitive Psychology*, 35(1), 71–98.
- Lezama, J., Grompone von Gioi, R., Randall, G., Morel, J.-M. (2014). A contrario detection of good continuation of points. In: *IEEE international conference on image processing*.
- Logan, G. D. (1996). The code theory of visual attention: An integration of space-based and object-based attention. *Psychological Review*, 103(4), 603.
- Lowe, D. (1985). *Perceptual organization and visual recognition*. Kluwer Academic Publishers.
- Lowe, D., Binford, T. O. (1981). The interpretation of three-dimensional structure from image curves. In: *Proceedings of IJCAI-7*, pp. 613–618.
- Machilsen, B., & Wagemans, J. (2011). Integration of contour and surface information in shape detection. *Vision Research*, 51(1), 179–186.
- Metzger, W. (1975). *Gesetze des Sehens* (3rd ed.). Frankfurt am Main: Verlag Waldemar Kramer.
- Mussap, A. J., & Levi, D. M. (2000). Amblyopic deficits in detecting a dotted line in noise. *Vision Research*, 40(23), 3297–3307.
- Palmer, S. E. (1999). *Vision science: Photons to phenomenology*. The MIT Press.
- Parent, P., & Zucker, S. W. (1989). Trace inference, curvature consistency, and curve detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(8), 823–839.
- Perona, P., & Freeman, W. (1998). A factorization approach to grouping. In *Computer vision—ECCV'98* (pp. 655–670). Springer.
- Pizlo, Z., Salach-Golyska, M., & Rosenfeld, A. (1997). Curve detection in a noisy image. *Vision Research*, 37(9), 1217–1241.
- Prinzmetal, W., & Banks, W. P. (1977). Good continuation affects visual detection. *Perception & Psychophysics*, 21(5), 389–395.
- Rock, I. (1983). *The logic of perception*. The MIT Press.
- Sha'asua, A., & Ullman, S. (1988). Structural saliency: The detection of globally salient structures using a locally connected network. In *Second International Conference on Computer Vision* (pp. 321–327). IEEE.
- Smits, J., & Vos, P. (1986). A model for the perception of curves in dot figures: The role of local salience of “virtual lines”. *Biological Cybernetics*, 54(6), 407–416.
- Smits, J., Vos, P. G., & Van Oeffelen, M. P. (1984). The perception of a dotted line in noise: A model of good continuation and some experimental results. *Spatial Vision*, 1(2), 163–177.
- Uttal, W. R. (1973). The effect of deviations from linearity on the detection of dotted line patterns. *Vision Research*, 13(11), 2155–2163.
- Uttal, W. R., Bunnell, L. M., & Corwin, S. (1970). On the detectability of straight lines in visual noise: An extension of French's paradigm into the millisecond domain. *Perception & Psychophysics*, 8(6), 385–388.
- van Assen, M. A., & Vos, P. G. (1999). Evidence for curvilinear interpolation from dot alignment judgements. *Vision Research*, 39(26), 4378–4392.
- van Oeffelen, M. P., & Vos, P. G. (1983). An algorithm for pattern description on the level of relative proximity. *Pattern Recognition*, 16(3), 341–348.
- Wagemans, J., Elder, J. H., Kubovy, M., Palmer, S. E., Peterson, M. A., Singh, M., et al. (2012). A century of gestalt psychology in visual perception: I. Perceptual grouping and figure-ground organization. *Psychological Bulletin*, 138(6), 1172–1217.
- Wagemans, J., Feldman, J., Gepshtein, S., Kimchi, R., Pomerantz, J. R., van der Helm, P. A., et al. (2012). A century of gestalt psychology in visual perception: II. Conceptual and theoretical foundations. *Psychological Bulletin*, 138(6), 1218–1252.
- Wertheimer, M. (1923). Untersuchungen zur Lehre von der Gestalt. II. *Psychologische Forschung*, 4, 301–350. an abridged translation to English is included in Ellis (1967 (originally 1938)).
- Williams, L. R., & Thornber, K. K. (1999). A comparison of measures for detecting natural shapes in cluttered backgrounds. *International Journal of Computer Vision*, 34(2–3), 81–96.
- Witkin, A. P., & Tenenbaum, J. M. (1983). On the role of structure in vision. In J. Beck, B. Hope, & A. Rosenfeld (Eds.), *Human and machine vision* (pp. 481–543). Academic Press.