A Contrario 2D Point Alignment Detection

José Lezama, Rafael Grompone von Gioi, Jean-Michel Morel, Grégory Randall

► To cite this version:

José Lezama, Rafael Grompone von Gioi, Jean-Michel Morel, Grégory Randall. A Contrario 2D Point Alignment Detection. 2014. hal-00956596

HAL Id: hal-00956596

<https://hal.inria.fr/hal-00956596>

Preprint submitted on 6 Mar 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Point Alignment Detection

José Lezama, Rafael Grompone von Gioi, Jean-Michel Morel, Gregory Randall

Abstract—In spite of many interesting attempts, the problem of automatically finding alignments in a 2D set of points seems to be still open. This paper tries to explain why it is so difficult, and to elaborate a solution. We show that a correct alignment detection depends on not less than four interlaced criteria, namely the amount of masking in texture, the relative bilateral local density of the alignment, its internal regularity, and finally a redundancy reduction step. Extending tools of the *a contrario* detection theory, we show that all of these detection criteria can be naturally embedded in a single probabilistic *a contrario* model with a single user parameter, the number of false alarms. Our contribution to the *a contrario* theory is the use of sophisticated conditional events on random point sets, for which expectation we nevertheless find easy bounds. By these bounds the mathematical consistency of our detection model receives a simple proof. Our final algorithm also includes a new formulation of the *exclusion principle* in Gestalt theory to avoid redundant detections. Aiming at reproducibility, a source code and an online demo open to any data point set are provided. The method is carefully compared to three state-of-the-art algorithms and counterexamples to the final methods are also explained.

Index Terms—point alignment detection, clustering, a contrario methods, Poisson point process

The code of the final algorithm used to generate all the experiments presented here is freely available at the companion webpage to this submission.¹ An online demo is also available and readers are invited to try out the method on their own data (the website allows users to upload or draw directly their own point patterns).

1 INTRODUCTION

We will consider the problem of finding collinear subsets within a planar set of points. This problem arises in many contexts of data analysis: Alignments are among the simplest structures observable in a point set. They constitute a classic example in statistical shape analysis [40]. Alignment detection is relevant in geology, where the alignment of features, for example earthquake epicenters, reflects underlying faults and joints [32], [18], [19]. In archaeology, geometric configurations of post holes, in particular alignments, often reveal the disposition of buildings even in presence of overlaps from different time periods [40], [31], [6]. The computer vision applications include the detection of

Fig. 1. Exactly the same set of aligned dots is present in the three images, but it is only perceived as such in the first one. The second one is a classic “masking by texture” case and the third shows a “masking by structure”, often called “Gestalt conflict”.

grids [11], calibration patterns [14] or vanishing points [21], [53], and the interpretation of high resolution remote sensing images [48].

Dot patterns are often used in the study of visual perception. Several psychophysical studies led by Utal have investigated the effect of direction, quantity and spacing in dot alignment perception [47], [44], [45], [46]. The detection of collinear dots in noise was the target of other studies attempting to assess quantitatively the masking effect of the background noise [41], [28], [43], [38]. A recent work by Preiss analyzes various perceptual tasks on dot patterns from a psychophysical and computational perspective [39]. An interesting computational approach to detect gestalts in dot patterns is presented in [1], although the study is limited to very regularly sampled patterns. The work described here was initiated in the context of a psychophysical research; here we concentrate, however, on the general problem of detection of point alignments in noise.

While it may seem that point alignments are simple structures, Fig. 1 shows how complex an alignment event can be. From a purely factual point of view, the same alignment is present in the three figures. However, it is only perceived as such by most viewers in the first one. The second and the third figures illustrate two occurrences of the *masking phenomenon* discovered by gestaltists [26]: the *masking by texture*, which occurs when a geometric structure is surrounded by a clutter of randomly distributed similar objects or *distractors*, and the *masking by structure*, which happens when the structure is masked by other

- José Lezama, Rafael Grompone von Gioi, and Jean-Michel Morel are with CMLA, ENS Cachan, France.
- José Lezama and Gregory Randall are with the IIE, Universidad de la República, Uruguay.

1. http://bit.ly/point_alignments

perceptually more relevant structures, a phenomenon also called *perceptual conflict* by gestaltists [34], [35], [25]. The magic disappearance of the alignment in the second and third figures can be accounted for in two very different ways. As for the first one, a probabilistic *a contrario* model [10] is relevant and can lead to a quantitative prediction. As for the second disappearance, it requires the intervention of another more powerful grouping law, the *good continuation* [24] in the perceptual conflict.

These examples show that a mathematical definition of point alignment perception is required before even starting to discuss how to detect them. A purely geometric-physical description is clearly not sufficient to account for the masking phenomenon. Indeed, an objective observer making use of a ruler would be able to state the existence of the very same alignment at the same precision on all three figures. But this statement would contradict our perception, as it would contradict any reasonable computational (definition and) theory of alignment detection.

This experiment also shows that alignment detection is highly dependent on the context of the alignment. It is therefore a complex question, and must be decided by building mathematical definitions and detection algorithms, and confronting them to human perception. As the patterns of Fig. 1 already suggest, simple computational definitions with increasing complexity will nevertheless find perceptual counterexamples.

A classic approach to this problem uses the Hough transform [23], [12], first used for the detection of subatomic particles in bubble chamber pictures [22]. To compute the Hough transform, each point votes in a parameter space for the lines that pass through it. After accumulation of the votes of all points, the lines that correspond to local maxima in the parameter space are selected as detections. Several variations of the basic method were proposed; in particular, the methods proposed in [42], [29], [30] are robust to errors in the point positions. When the Hough transform is applied to a random set of points, it will still find a local maximum which does not correspond to a significant collinear subset. A threshold on the number of votes is usually imposed to cope with this problem. Even if the Hough transform methods provide successful solutions in many applications [49], a sound criterion for setting this threshold is missing.

Other approaches use point clustering methods especially adapted to clusters of linear shape [37]. Using a particular distance between points and clusters [8], general clustering algorithms can be used to detect collinear subsets [16], [15]. The same problem can be approached using a parametric model fitting [7].

We are particularly interested in methods that provide an evaluation of the statistical significance of the detected aligned structures. An example in astronomy may illustrate the importance of such evaluation: In

1980 the discovery of several very precise alignments of quasars in the sky raised the question of a theory explaining this presence [5]. These alignments, however, were later dismissed by a statistical analysis, first by simulations [13] and then analytically [54], showing that alignments of such precision could easily occur just by chance.

The expected number of events where k among n random points are to be found in some rectangle of a given shape was already computed in 1950 using a Poisson random model [33]. This could be the origin of the *strip* method for defining alignments as a large number of points covered by a thin rectangle (the thinner, the more precise). The same random model was used in [6], now explicitly used for detecting point alignments. But the alignment was defined differently: three points are considered aligned when the triangle formed by them is flat enough. Alignments of more points are evaluated by all the possible triangles observed among the points. Various theoretical results about the flat triangles methods are described in [27], where Poisson as well as Gaussian distributions are considered in different domain shapes.

Since then, many different algorithms have been proposed, most of them variations of the strip method. Monte Carlo simulations of random points provide the estimate of the significance in [52] while a binomial model is used in [3]. A set of heuristics are added in [4]. The method in [18] also applies the strip method with a Poisson model, but the density is estimated locally. A refined statistical test, including angular statistics, is proposed in [19].

Here we develop a method deriving from the *a contrario* methodology proposed by Desolneux, Moisan and Morel [9], [10]. It is a mathematical formalization of the so-called *non-accidentalness principle* proposed for perception [51], [2], [50] (sometimes called *Helmholtz principle*). In a nutshell, an observed structure is relevant if it would rarely occur by chance. The result is a statistical framework that provides estimates of significance similar in spirit to the methods mentioned before.

As a simple example to introduce the methodology, Desolneux, Moisan and Morel showed how to handle point alignment detection using a simple strip method with a Poisson model [10, Sect. 3.2]. We shall nevertheless show that this initial method (and the similar ones we mentioned above) is far from sufficient for detecting perceptual alignments. To cope with obvious objections and counterexamples, we shall prove that three new features are necessary to cope with the variety of alignments. We shall show that a reliable algorithm requires: a) a *local* Poisson density estimation, b) an evaluation of the *regularity of the spacing* of the points in the alignment, and c) a criterion to *select the best interpretation among redundant detections*.

The rest of this article is organized as follows:

Fig. 2. A schematic representation of the evaluated rectangle. In a domain with N points, there are $\frac{N(N-1)}{2}W$ possible rectangles. In this example, $N = 47$ and $k(r, \mathbf{x}) = 8$ among them are inside the rectangle r .

Section 2 introduces the basic concepts and techniques used in the state of the art point alignment detectors, and describes the classic strip method. Sections 3, 4 and 5 improve this basic method by incorporating local point density estimation, lateral estimation, and measurements of the regularity of the point spacing. Section 6 discusses how to cope with the redundancy of detections. Section 7 shows some experiments and comparisons and Sect. 8 concludes the paper.

2 BASIC POINT ALIGNMENT DETECTOR

The key point in a statistical detection method is how to define *accidental* detections. This requires a stochastic model H_0 , the so-called *a contrario* model, characterizing unstructured or random data in which the sought structure could only be observed by chance. This basic idea of comparing the detection to the *expected* number of detections in a random model in an *a contrario* model is inevitable [6], [27], [19].

Consider a set of N points defined in a domain D with total area S_D , see Fig. 2. We are interested in detecting groups of points that are well aligned. A reasonable *a contrario* hypothesis H_0 for this problem is to suppose that the N points are the result of a random process where points are independent and uniformly distributed in the domain. The question is then to evaluate whether the presence of aligned points contradicts the *a contrario* model or not.

Given an observed set of N points $\mathbf{x} = \{x_i\}_{i=1\dots N}$ and a rectangle r (a candidate for alignment), we will denote by $k(r, \mathbf{x})$ the number of those points observed inside r . The decision of whether to keep this candidate or not is based on two principles: a good candidate should be non-accidental, and any equivalent or better candidate should be kept as well. The degree of non-accidentalness of an observed rectangle r can be measured by how small the probability $\mathbb{P}[k(r, \mathbf{X}) \geq k(r, \mathbf{x})]$ is, where $\mathbf{X}$ denotes a random set of N points following H_0 . In the same vein, a rectangle r' will be considered at least as good as r given the observation $\mathbf{x}$, if $\mathbb{P}[k(r', \mathbf{X}) \geq k(r', \mathbf{x})] \leq \mathbb{P}[k(r, \mathbf{X}) \geq k(r, \mathbf{x})]$.

The question is now to control the expected number of accidental detections. Given that N_{tests} candidates will be tested, the expected number of rectangles which are as good as r under H_0 is, following [10], less than

$$N_{tests} \cdot \mathbb{P}[k(r, \mathbf{X}) \geq k(r, \mathbf{x})]. \quad (1)$$

The H_0 stochastic model fixes the probability law of the random number of points in the rectangle, $k(r, \mathbf{X})$. The discrete nature of this law implies that (1) is not actually the expected value but an upper bound of it [10], [17]. Let us now analyze the two factors in (1).

The *a contrario* model H_0 assumes that the N points are i.i.d. with uniform density on the domain. (H_0 is a planar Poisson process [36].) Under the *a contrario* hypothesis H_0 , the probability that one point falls into the rectangle r is $p = \frac{S_r}{S_D}$, where S_r is the area of the rectangle and S_D the area of the domain. As a consequence of the independence of the random points, $k(r, \mathbf{X})$ follows a binomial distribution. Thus, the probability term $\mathbb{P}[k(r, \mathbf{X}) \geq k(r, \mathbf{x})]$ is given by

$$\mathbb{P}[k(r, \mathbf{X}) \geq k(r, \mathbf{x})] = \mathcal{B}(N, k(r, \mathbf{x}), p) \quad (2)$$

where $\mathcal{B}(n, k, p)$ is the tail of the binomial distribution

$$\mathcal{B}(n, k, p) = \sum_{j=k}^n \binom{n}{j} p^j (1-p)^{n-j}. \quad (3)$$

The number of tests N_{tests} corresponds to the total number of rectangles that could contain an alignment, which in turn is related to the number of pairs of points defining such rectangles. With a set of N points this gives $\frac{N(N-1)}{2}$ different pairs of points. The set of rectangle widths to be tested must be specified *a priori* as well. In the *a contrario* approach, a compromise must be found between the number of tests and the precision of the structures that are being sought for. The larger the number of tests, the lower the statistical relevance of detections. However, if the set of tests is chosen wisely, structures fitting accurately the tests will have a very low probability of occurrence under H_0 and will therefore be more significant.

For a particular problem, one may have reasons to restrict the shape of rectangles. Nevertheless, this inquiry is not aimed at any particular application. Thus, we will rely on the following general criteria. An alignment should be an elongated structure, so a minimal ratio between the length and the width of the rectangle must be fixed. Then, a fixed number of widths must be tested, decreasing geometrically the maximal width. (The choice of a geometric series is justified by the obvious scale invariance of the detection problem.) Our implementation uses a length/width_{max} ratio of 10 and a geometric series of 8 width values with a factor $1/\sqrt{2}$. The total number of widths to be tested will be denoted by W . Then the

total number of tested rectangles is

$$N_{tests} = \frac{N(N-1)}{2} W. \quad (4)$$

The fundamental quantity of an *a contrario* approach is the Number of False Alarms (NFA) associated with a rectangle r and a set of points $\mathbf{x}$,

$$\begin{aligned} \text{NFA}_1(r, \mathbf{x}) &= N_{tests} \cdot \mathbb{P}[k(r, \mathbf{X}) \geq k(r, \mathbf{x})] \\ &= \frac{N(N-1)}{2} W \cdot \mathcal{B}(N, k(r, \mathbf{x}), p). \end{aligned} \quad (5)$$

This quantity gives a precise meaning to Eq. (1). It will be interpreted as a bound of the expected number of rectangles containing enough *points* to be as rare as r under H_0 . When the NFA associated with a rectangle is large, this means that such an event is to be expected under the *a contrario* model and therefore is not relevant. On the other hand, when the NFA is small, the event is rare and probably meaningful. A rarity threshold ε must nevertheless be fixed for each application. Rectangles with $\text{NFA}_1(r, \mathbf{x}) \leq \varepsilon$ will be called ε -*meaningful rectangles* [10], constituting the detection result of the algorithm. We will refer to this method as Algorithm 1.

Theorem 1 ([10]).

$$\mathbb{E} \left[\sum_{r \in \mathcal{R}} \mathbb{1}_{\text{NFA}_1(r, \mathbf{x}) \leq \varepsilon} \right] \leq \varepsilon$$

where $\mathbb{E}$ is the expectation operator, $\mathbb{1}$ is the indicator function, $\mathcal{R}$ is the set of test rectangles, and $\mathbf{X}$ is a random set of points under H_0 .

The theorem states that the average number of ε -*meaningful rectangles* under the *a contrario* model H_0 is bounded by ε . Thus, the number of detections in noise is controlled by ε and it can be made as small as desired. In other words, this detector satisfies the non-accidentalness principle.

As shown in [10], the detection result is not very sensitive to the value of ε . Following Desolneux, Moisan, and Morel [9], [10], we shall therefore fix $\varepsilon = 1$ for our experiments. This corresponds to accepting on average at most one false detection per data set in the *a contrario* model.

Figure 3 shows the results of the basic algorithm in two simple cases. The results are as expected: the visible alignment in the first example is detected, while no detection is produced in the second. Actually, the points in the first example are also present in the second one, but the addition of random points masks the alignment to our perception. The first example produces many redundant detections; this issue will be addressed in Sect. 6.

3 LOCAL DENSITY ESTIMATION

The basic point alignment detector of section 2 takes as *a contrario* assumption a uniform point density

Fig. 3. Results from the basic point alignment detector (Algorithm 1). (a) and (c) are the input data, and (b) and (d) are the corresponding results. Each detection is represented by a rectangle. In (b) the algorithm correctly detects the obvious alignment. Notice that multiple and redundant rectangles were detected; this issue will be dealt with in Sect. 6. The data set (c) contains the same set of points in (a) plus added noise points. The aligned points are still present but hardly perceptible. The algorithm handles correctly this masking phenomenon and produces no detection.

in the whole domain and evaluates alignments as a local excess with respect to this global density. This comparison is nevertheless too restrictive, because alignments are in fact *local violations of uniformity*. Consider a configuration of points with two zones of different point density, like in Fig. 4 (a). Applying the basic alignment detector yields an unexpected detection shown in Fig. 4 (b). Each of the detected rectangles certainly has a non-accidental excess of points in the rectangles with respect to the global density, but this is definitely not what we are looking for. This example shows that we are actually interested in non-accidental events with an excess of points conditioned to the observation of a local density (which may well be lower or higher than the global density). Such local density estimations for the random point models have been used in [3], [18], [19]. In the interpretation proposed here, a more sophisticated definition of the alignment event should not measure the non-accidentalness by an unusually small probability, but by an unusually small *conditional* probability.

The local density is estimated by counting the points in a rectangular local window, with the same length as the alignment and a given width. To account for the scale invariance of the detection, the width

Fig. 4. Local vs. global density estimation. **(a)** The set of points. **(b)** Alignments found using global density estimation (Algorithm 1). The many detected rectangles indeed have a high point density compared to the average image density used as background model. **(c)** Alignments found using local density estimation (Algorithm 2). The local density is in that case lower on the border, hence the deceptive detection. **(d)** No alignment is found when the local density is estimated by the maximum density on both sides of the alignment (Algorithm 3).

of the local window is proportional to the length of the alignment. For every alignment, a number of proportion ratios or scales are tried. The largest window is square, its width equal to the length of the alignment. Then a fixed number $L = 8$ of widths in geometric series are also tried. The choice for a geometric series with factor $1/\sqrt{2}$ is again justified by the scale invariance of the detection problem. The number of tests N_{tests} corresponds to the total number of observations performed, which in turn is related to the number of rectangles and the different local windows evaluated for each rectangle. With a set of N points and L different sizes of local windows, this gives $\frac{N(N-1)}{2}WL$ different tests.

When the rectangle to be tested lies near the border of the domain, the local window may be partly outside it, where no point information is available, leading to a wrong density estimation. This also happens when the rectangle covers the diagonal of the domain. In accordance with perception theory and the *good continuation* principle, a sound solution is obtained by creating a mirror-symmetric extension of the point set outside the domain. The extended point set is only used for the local density estimation, but not for selecting candidates, and its points are never counted as part of an alignment.

Let R be the local window surrounding the alignment r , as shown in Fig. 5(a). The probability of one point in R falling in r is $p = \frac{S_r}{S_R}$ where S_r and S_R are the areas of r and R respectively. The degree of non-accidentalness of an observation will be measured by the probability that a rectangle has a higher density than its surroundings, conditioned by the observation of the surrounding density. The NFA for the new

Fig. 5. **(a)** The point density is estimated in the local window R surrounding the alignment r . **(b)** In a refined version of the algorithm, the density of points is measured on each side of the evaluated rectangle. The maximum of the densities in R_1 and R_3 is taken as an estimation of the point density in both R_1 and R_3 .

detector is accordingly defined as

$$\begin{aligned} \text{NFA}_2(r, R, \mathbf{x}) &= \\ N_{tests} \cdot \mathbb{P} \left[k(r, \mathbf{X}) \geq k(r, \mathbf{x}) \mid n(R, \mathbf{X}) = n(R, \mathbf{x}) \right] \\ &= \frac{N(N-1)}{2} WL \cdot \mathcal{B}(n(R, \mathbf{x}), k(r, \mathbf{x}), p), \quad (6) \end{aligned}$$

where $n(R, \mathbf{x})$ is the number of points observed in $\mathbf{x}$ inside R . We will call this method Algorithm 2. The next theorem shows that the method controls the number of false detections in H_0 ; its proof is very similar to the proof of the more general Th. 4 and will therefore be omitted.

Theorem 2.

$$\mathbb{E} \left[\sum_{r \in \mathcal{R}} \sum_{R \in \mathcal{R}'(r)} \mathbb{1}_{\text{NFA}_2(r, R, \mathbf{X}) \leq \varepsilon} \right] \leq \varepsilon$$

where $\mathbb{E}$ is the expectation operator, $\mathbb{1}$ is the indicator function, $\mathcal{R}$ is the set of rectangles considered, $\mathcal{R}'(r)$ is the set of surrounding local windows for each rectangle r , and $\mathbf{X}$ is a random set of points under H_0 .

4 LATERAL DENSITY ESTIMATION

While the local density estimation can provide a more adjusted background model, it can also introduce new problems such as a “border effect”, as shown in Fig. 4 (c). Indeed, the density estimation is lower on the border of the left half of the image than inside it. Thus, the previous algorithm (Algorithm 2) detects alignments on the border with non-accidental, meaningful excess with respect to the local density.

In order to avoid this effect, the more sophisticated Algorithm 3 used in Fig. 4 (d) takes, as a conservative estimation of the background density, the *maximum of the densities* measured on both sides of the alignment. In short, to be detected, an alignment must show a higher point density than in both regions immediately on its left and right. This local alignment detector is therefore similar to a classic second order Gabor filter where an elongated excitatory region is surrounded by two inhibitory regions. The local density estimation is calculated as illustrated in Fig. 5 (b): The local window is divided in three parts. R_1 is the rectangle formed by the area of the local window on the left of the alignment. R_3 is the area of the local window on the right of the alignment, and R_2 is the rectangle which forms the candidate alignment. Note that the length of the local window is the same as the alignment and that we can consider any arbitrary orientation for it. Next, the algorithm counts the numbers of points M_1 , M_2 , and M_3 in R_1 , R_2 and R_3 , respectively, and defines the conservative estimate of the local number of points as

$$n^*(R, \mathbf{x}) = 2 \max(M_1, M_3) + M_2. \quad (7)$$

We then define the NFA of the event “the density in R_2 has a significant excess with respect to the density estimated in R ” by

$$\begin{aligned} \text{NFA}_3(r, R, \mathbf{x}) = \\ N_{\text{tests}} \cdot \mathbb{P} \left[k(r, \mathbf{X}) \geq k(r, \mathbf{x}) \mid n(R, \mathbf{X}) = n^*(R, \mathbf{x}) \right] \\ = \frac{N(N-1)}{2} WL \cdot \mathcal{B} \left(n^*(R, \mathbf{x}), k(r, \mathbf{x}), p \right). \end{aligned} \quad (8)$$

Indeed, conditioned to the fact that we assume $n(R, \mathbf{X}) = n^*(R, \mathbf{x})$ under the model H_0 , the $n^*(R, \mathbf{x})$ points in R are still uniformly and independently distributed. We call this method Algorithm 3. The following theorem shows that the method controls the number of false detections in H_0 . Again, its proof is omitted because it is very similar to the proof of Th. 4.

Theorem 3.

$$\mathbb{E} \left[\sum_{r \in \mathcal{R}} \sum_{R \in \mathcal{R}'(r)} \mathbb{1}_{\text{NFA}_3(r, R, \mathbf{X}) \leq \varepsilon} \right] \leq \varepsilon$$

where $\mathbb{E}$ is the expectation operator, $\mathbb{1}$ is the indicator function, $\mathcal{R}$ is the set of rectangles considered, $\mathcal{R}'(r)$ is the set of surrounding local windows for each rectangle r , and $\mathbf{X}$ is a random set of points under H_0 .

5 ALIGNMENT REGULARITY

We proceed incrementally toward a more complex algorithm, by showing that each new more sophisticated version of an alignment detection is still not sufficient to cope with all features of this problem. There is indeed still an objection to Algorithm 3: One can stir wrong detections by introducing small point clusters as shown in Fig. 6 (a). The detected alignment in Fig. 6 (b) seems clearly wrong. It is nevertheless explainable in the setting of Algorithm 3: there is indeed a meaningful point density excess inside the red rectangle. But this excess is caused by the clusters, not by what could be termed an alignment. While the algorithm counted every point, human perception seems to group the small clusters into a single entity, and to count them only once. This unwanted result is a consequence of the fact that Algorithm 3 is searching for elongated clusters of higher density without any cluster regularity requirement. As suggested in other studies [39], [43], [44], the density is not the only property that makes an alignment perceptually meaningful; another characteristic to consider is the uniform spacing or regularity of the points in it, which the gestaltists call the law of *constant spacing*. This observation forces to a still more sophisticated version of the alignment detector. To cope with both issues (avoiding small clusters and favoring regular spacing) a more advanced version of the alignment detector divides each candidate rectangle into equal *boxes*. Instead of counting the total number of points, the algorithm counts the number of boxes that are occupied by at least one point. We call them *occupied boxes*. In this way, the minimal NFA is attained when the points are perfectly distributed along the alignment. In addition, a concentrated cluster in the alignment has no more influence on the alignment detection than a single point in the same position.

We want to estimate the expected number of occupied boxes in the background model H_0 . The probability of one point falling in one of the boxes is $p_0 = \frac{S_B}{S_L}$, where S_B and S_L are the areas of the boxes and the local window respectively. Then, the probability of having one box occupied by at least one of the $n^*(R, \mathbf{x})$ points (i.e., of an *occupied box*) is

$$p_1(R, c) = 1 - (1 - p_0)^{n^*(R, \mathbf{x})}. \quad (9)$$

We will denote by $b(r, c, \mathbf{x})$ the observed number of occupied boxes in the rectangle r when divided into c boxes. Finally, the probability of having at least $b(r, c, \mathbf{x})$ of the c boxes occupied is

$$\mathcal{B}(c, b(r, c, \mathbf{x}), p_1(R, c)). \quad (10)$$

A set $\mathcal{C}$ of different values are tried for the number of boxes c into which the rectangle is divided, and the one producing the lowest NFA is taken. Thus, the number of tests must be multiplied by its cardinality

Fig. 6. Counting occupied boxes to avoid false detections due to the presence of clusters. The dot pattern shown in image **(a)** presents two point clusters but no alignment. However, Algorithm 3 finds a thin rectangle with a high point density, hence a false detection, as shown in **(b)**. Algorithm 4 divides the rectangle into boxes and counts the occupied ones, avoiding this misleading cluster effect, as seen in **(c)**, where the occupied boxes are marked in red and no alignment is actually detected.

Fig. 7. Redundant detections. **Left:** point pattern. **Center:** all significant alignments found by the refined point alignment detector (Algorithm 4) described in Sect. 5. The color represents the relative NFA value, where red is the most significant (smallest NFA value) and blue the least (highest NFA value). **Right:** Result of the masking process.

Fig. 8. Examples of two alternative formulations of the masking process. **Left:** Set of points. **Center:** The Exclusion Principle as defined in [10], a validated gestalt prevents others from using its points. The vertical alignments (evaluated first) mask the horizontal ones. **Right:** The Masking Principle, described in the text, which solves the ambiguities without forbidding basic elements to participate of two different structures. In this example, no individual alignment can mask an individual one in another direction. Thus we get oblique, horizontal and vertical meaningful alignments.

$\#\mathcal{C} = C$. In practice we set $C = \sqrt{N}$ and that leads to

$$N_{tests} = \frac{N(N-1)}{2} WLC = \frac{N(N-1)}{2} WL\sqrt{N}.$$

The NFA of the new event definition is then

$$\begin{aligned} \text{NFA}_4(r, R, c, \mathbf{x}) = \\ N_{tests} \cdot \mathbb{P}\left[b(r, c, \mathbf{X}) \geq b(r, c, \mathbf{x}) \mid n(R, \mathbf{X}) = n^*(R, \mathbf{x})\right] \\ = \frac{N(N-1)}{2} WLC \cdot \mathcal{B}(c, b(r, c, \mathbf{x}), p_1(R, c)). \quad (11) \end{aligned}$$

Fig. 6 (c) shows an example of the resulting algorithm and we will show some more after discussing the masking problem, in section 6. Algorithm 4 presents the pseudo-code for this final refined version of the alignment detector.

Theorem 4.

$$\mathbb{E} \left[\sum_{r \in \mathcal{R}} \sum_{R \in \mathcal{R}'(r)} \sum_{c \in \mathcal{C}} \mathbb{1}_{\text{NFA}_4(r, R, c, \mathbf{X}) \leq \varepsilon} \right] \leq \varepsilon$$

where $\mathbb{E}$ is the expectation operator, $\mathbb{1}$ is the indicator function, $\mathcal{R}$ is the set of rectangles considered, $\mathcal{R}'(r)$ is the set of surrounding local windows for each rectangle r , $\mathcal{C}$ is the set of number of boxes tested, and $\mathbf{X}$ is a random set of points under H_0 .

Proof: We define $\hat{b}(r, R, c, M)$ as

$$\begin{aligned} \hat{b}(r, R, c, M) = \min \left\{ \beta \in \mathbb{N}, \right. \\ \left. \mathbb{P}[b(r, c, \mathbf{X}) \geq \beta \mid n(R, \mathbf{X}) = M] \leq \frac{\varepsilon}{\frac{N(N-1)}{2} WLC} \right\}. \end{aligned}$$

R determines the domain of the local window and M the number of points in it. The probabilistic model inside R , conditioned to the fact that the number of observed points is M , is still uniform and independent, and the conditional law of the number of points inside any subset of R follows a binomial law. Then, $\text{NFA}_4(r, R, c, \mathbf{X}) \leq \varepsilon$ is equivalent to $b(r, c, \mathbf{X}) \geq \hat{b}(r, R, c, M)$ when $M = n^*(R, \mathbf{X})$. Now,

$$\begin{aligned} \mathbb{E} \left[\sum_{r \in \mathcal{R}} \sum_{R \in \mathcal{R}'(r)} \sum_{c \in \mathcal{C}} \mathbb{1}_{\text{NFA}_4(r, R, c, \mathbf{X}) \leq \varepsilon} \right] = \\ \sum_{r \in \mathcal{R}} \sum_{R \in \mathcal{R}'(r)} \sum_{c \in \mathcal{C}} \mathbb{P}[\text{NFA}_4(r, R, c, \mathbf{X}) \leq \varepsilon] = \\ \sum_{r \in \mathcal{R}} \sum_{R \in \mathcal{R}'(r)} \sum_{c \in \mathcal{C}} \sum_{M=0}^{2N} \mathbb{P}[\text{NFA}_4(r, R, c, \mathbf{X}) \leq \varepsilon \mid \\ n(R, \mathbf{X}) = M] \cdot \mathbb{P}[n(R, \mathbf{X}) = M] = \\ \sum_{r \in \mathcal{R}} \sum_{R \in \mathcal{R}'(r)} \sum_{c \in \mathcal{C}} \sum_{M=0}^{2N} \mathbb{P}[b(r, c, \mathbf{X}) \geq \hat{b}(r, R, c, M) \mid \\ n(R, \mathbf{X}) = M] \cdot \mathbb{P}[n(R, \mathbf{X}) = M]. \quad (12) \end{aligned}$$

Algorithm 4: Point alignment detector with boxes

input : A set $\mathbf{x}$ of N points [$W = 8, L = 8, \varepsilon = 1$]
output: A list **out** of point alignments

```

for  $i = 1$  to  $N$  and  $j = 1$  to  $i - 1$  do
     $l \leftarrow \text{distance}(x_i, x_j)$ 
     $w \leftarrow l/10$ 
    for  $1$  to  $W$  do
         $r \leftarrow \text{rect}(x_i, x_j, w)$ 
         $w_L \leftarrow l$ 
        for  $1$  to  $L$  do
             $R_1 \leftarrow \text{local-window-left}(x_i, x_j, w_L)$ 
             $R_3 \leftarrow \text{local-window-right}(x_i, x_j, w_L)$ 
            for  $c \in \mathcal{C}$  do
                Compute  $\text{NFA}_4(r, R, c, \mathbf{x})$  [Eq. 11]
                if  $\text{NFA}_4(r, R, c, \mathbf{x}) \leq \varepsilon$  then
                    out  $\leftarrow r$ 
                end
            end
             $w_L \leftarrow w_L / \sqrt{2}$ 
        end
         $w \leftarrow w / \sqrt{2}$ 
    end
end

```

Note that, because of the maximum density estimation $n^*(R, \mathbf{x})$, the estimated number of points inside a rectangle can theoretically be as large as $2N$, and thus the range for M . By definition of $\hat{b}(r, R, c, M)$ we know that

$$\mathbb{P}[b(r, c, \mathbf{X}) \geq \hat{b}(r, R, c, M) \mid n(R, \mathbf{X}) = M] \leq \frac{\varepsilon}{\frac{N(N-1)}{2} WLC},$$

and using $\#\mathcal{R} = \frac{N(N-1)}{2} W$, $\#\mathcal{R}'(r) = L$, $\#\mathcal{C} = C$ and

$$\sum_{M=0}^{2N} \mathbb{P}[n(R, \mathbf{X}) = M] = 1$$

$$\begin{aligned} \text{we get } \mathbb{E} \left[\sum_{r \in \mathcal{R}} \sum_{R \in \mathcal{R}'(r)} \sum_{c \in \mathcal{C}} \mathbb{1}_{\text{NFA}_4(r, R, c, \mathbf{X}) \leq \varepsilon} \right] \leq \\ \sum_{r \in \mathcal{R}} \sum_{R \in \mathcal{R}'(r)} \sum_{c \in \mathcal{C}} \frac{\varepsilon}{\frac{N(N-1)}{2} WLC} \sum_{M=0}^{2N} \mathbb{P}[n(R, \mathbf{X}) = M] \\ = \varepsilon, \end{aligned}$$

which concludes the proof. $\square$

6 REDUNDANCY

As was observed in Fig. 3, all the described alignment detectors may produce redundant detections. In a very meaningful alignment many smaller or larger rectangles overlapping the main alignment are also meaningful. This redundancy phenomenon can involve points that belong to the real alignment as well as background points near the alignment, as

illustrated in Fig. 7. In such cases, humans usually perceive only one alignment. The question is how to detect this best rectangle, both explaining and masking the redundant detections.

A simple model for this masking process was proposed by Desolneux et al. [10] under the name of “exclusion principle”. The main idea is that each *basic element* (a point in our case) cannot contribute to more than one perceived group or gestalt. This leads to a masking algorithm proceeding as follows: The most meaningful observed structure (the one with smallest NFA) is kept as a valid detection. Then, all the basic elements (the points) that were part of that validated group are assigned to it and the remaining candidate structures cannot use them anymore. The NFA of the remaining candidates is re-computed without counting the excluded elements. In that way, redundant structures lose most of their supporting elements and are no longer meaningful. On the other hand, a candidate that corresponds to a different structure keeps most or all of its supporting basic elements and remains meaningful. The most meaningful candidate among the remaining ones is then validated and the process is iterated until there are no more meaningful candidates.

This formulation of the masking process often leads to good results, removing redundant detections while keeping the good ones. But it may also lead to unsatisfactory results as illustrated in Fig. 8. The problem arises when various valid alignments have many elements in common. As one alignment is evaluated after the other, it may happen that all of its elements have been removed, even if the alignment is in fact not redundant with any of the other ones. In the example of Fig. 8, individual horizontal and vertical alignments are not redundant, but if all the vertical ones have been detected first, the remaining horizontal ones will be (incorrectly) masked. This example shows a fundamental flaw of the exclusion principle: it is not sound to impose that a basic element belongs to a single perceptually valid structure. There must be a global explanation of the organization of the basic elements in visible structures which is at the same time coherent with each individual structure (eliminating local redundancy) and with the general explanation of the scene in such a way that some basic elements can participate of several structures without contradiction. The solution seems to be in a sort of relaxation of the exclusion principle. The following definitions sketch a possible solution.

Definition 1 (Building Elements). *We call building element any atomic component that can be a constituent element of several structures. An example of building elements are points that can be recursively grouped in alignments.*

Definition 2 (Masking Principle). *A meaningful structure B will be said “masked by a structure A ” if B is*

no longer meaningful when evaluated without counting its building elements belonging to A . In such a situation, the structure B is not retained as detected.

In short, a meaningful structure will be detected if it is not masked by any other detected structure. The difference with the former exclusion principle is that here a structure can only be masked by another *individual* structure and not by the union of several structures. A procedural way to attain this result is to validate alignments one by one, starting by the one with smallest NFA. Before accepting a new alignment, it is checked that it is not masked by any one of the previously detected alignments.

Figures 13 and 14 show some point alignment detection results when combining Algorithm 4 with the masking principle. The results obtained in these examples are as expected and this masking procedure was applied to all experiments below. For simplicity, we shall still refer to it as “Algorithm 4”.

7 COMPARATIVE EXPERIMENTS

We will first compare the proposed algorithm with two methods described in the literature designed for point alignment detection [18], [19]. These two methods were selected because they were introduced recently and include statistical significance tests. As we shall see, these methods depend critically on user parameters; we will describe their result for three data sets, commenting on the results while varying the parameters. Next, we will show the results, for the same data, of a well-known clustering method [15], which, as our algorithm, requires no parameter tuning. Finally, the results of our method for these and other point sets will be shown and commented. More experiments can be performed by the reader using the freely available online demo and source code.²

The method introduced in 2006 by Hall, Tajvidi and Malin [18] makes measurements similar to Algorithm 3: the alignment is evaluated as a thin strip and two lateral rectangles are used to estimate the point density; the set of candidates and the statistical test are different. The results presented here were computed using our own implementation of the method; we reproduced the experiments in the original publication to verify the correctness of our code.

The statistical test of the method by Hall et al. is designed to reject alignments in a uniform Poisson random point model. The method works well when the intensity of the point process is high. Indeed, the authors showed that the method approaches optimality as the intensity increases [18]. The method is less efficient when the density of points is low relative to the size of the operator; in such conditions the sampling is not suitable, the point density estimation is poor, and the statistical test is not able to reject

2. http://bit.ly/point_alignments

Fig. 9. Result of Hall, Tajvidi, Malin 2006 [18] on a set of 100 uniform and independent random points. Each detection is represented by a thin rectangle, surrounded by the local window. **Left:** The same parameters as in [18, Sect. 3.1] were used: 10×10 grid, 5 degree angle step, $a = 0.1$, $b = 0.6$, $c = 0.01$, $u = 6$ $v = 2$. As in the original publication, about 3 alignments were detected (2 in this example). **Right:** Result with a slightly different candidate set: 20×20 grid and $b = 0.3$, producing 47 detections. A similar behavior is observed with sets of 1000 random points.

random configurations. The test depends on two parameters, u and v , to be set manually. The first parameter controls the statistical level. In extreme cases of wrong density estimation (e.g. no point is observed in the local window), the statistical test fails. The second parameter, v , is a threshold imposed on the number of points in the strip. It is necessary to cope with cases of density undersampling. The method assumes that the domain is the unit square and tests candidates centered in an $n \times n$ grid, at regular orientations with an angle step θ ; three parameters define the strip: the length b , the strip width c , and the local window width a . Thus seven parameters must be provided by the user: n , θ , a , b , c , u , and v .

Figure 9 shows two detection results in a set of 100 points generated according to a Poisson model. The first result (left) is using the same parameters as in [18, Sect. 3.1]; in accordance with the results of the original article, in these conditions is observed an average of 3 detections per data set (2 in the example shown). However, when the shape parameters are modified, the statistical test is no longer able to control the number of false detections, see Fig. 9 (right). In the second experiment, the number of candidates is larger, (the grid is 20×20 instead of 10×10), and the density estimation is worse because the candidates are half as long (thus the local window is half as big). Under the new conditions, the same u and v values lead to 47 detections in the same point set. This experiment shows the need to set manually the statistical significance parameters (u , v) to produce reliable results. This behavior was also observed for sets of 1000 random points.

We will see now how the method handles data

sets that do contain point alignments. Figure 10(a) shows a set of 186 points; perceptually one can see three alignments, three clusters, and random points. We first adjusted the statistical test parameters u and v so as to obtain very few detections with the same number of random points. The candidate shape parameters (a , b , c) were then adjusted to obtain the best result, see Fig. 10(b). As one can see, the three alignments were detected. Nevertheless, all detections only partially cover the perceived alignment; this is of course due to the selected length ($b = 0.4$), but longer candidates produced less complete results. Also, there is some redundancy in the detections; no redundancy reduction step is included in Hall et al.'s method. Finally, one can observe that one of the clusters led to a false detection. When the shape parameters are changed to less optimal values, Figs. 10(c) and (d), we obtain less useful results: some alignments or parts of them are missing, and many spurious detections were produced. Some are due indeed to deviations from the random model as is the case of the clusters, and this shows the need for a more complex event definition. Others, as in Fig. 10(d), reveal a failure of the statistical test. This experiment shows that the candidate shape parameters must be carefully adjusted to produce good results.

The second point alignment detection method we used for comparison was introduced by Hammer in 2009 [19]. This method requires an alignment length parameter (defined as a radius). Each point of the input set defines a candidate. The distribution of angles from the center point to each of the points inside the radius is evaluated. When the circular-uniform distribution is rejected, using a Rayleigh test, the candidate produces a detection. The last decision requires a significance level α .

The experiments presented here were done using the author implementation of the algorithm, included in the software package PAST [20]. Figure 11 shows the results for the three data sets considered in this comparison and for four parameter sets (no detection was produced for $l = 0.6$ and $\alpha = 0.00001$). As one can see in the first row, the default significance level used in PAST is not satisfactory: it produces many detections on random sets of points (left and middle). Using this significance level one gets some of the expected alignments in the data set containing alignments (right). But, as in the previous method, redundancy is observed and many false detections, mainly caused by the presence of the clusters. When the significance level is changed to $\alpha = 0.00001$ (second row), the number of false detections in noise is reduced significantly; unfortunately, the true alignments also disappear, leaving only two detections due to a cluster. Using a longer alignment length (third row) produced no better results: false detections in random points and unexpected detections in the structured data set. Increasing the significance level

Fig. 10. Result of Hall, Tajvidi, Malin 2006 [18]. Each detection is represented by a thin rectangle, surrounded by the local window. In the following results a 20×20 candidate center grid was used, a 5 degree angle step, $u = 7$ and $v = 3$. **(a)** Input set of 186 points, containing 3 alignments, 3 clusters, and random points. **(b)** Result of the method for $a = 0.2$, $b = 0.4$, and $c = 0.02$. **(c)** Result of the method for $a = 0.3$, $b = 0.6$, and $c = 0.02$. **(d)** Result of the method for $a = 0.05$, $b = 0.6$, and $c = 0.005$.

Fig. 11. Result of Hammer 2009 [19] for three point sets in a normalized unit square domain. Two length values were tested: 0.3 (radius 0.15) and 0.6 (radius 0.3); and two significance levels: $\alpha = 0.1$ and $\alpha = 0.00001$. No detection was produced with $l = 0.6$ and $\alpha = 0.00001$. **Left:** The same set of 100 random points used in Fig. 9. **Middle:** 1000 points drawn independently with uniform distribution in a unit square. **Right:** The same point set as in Fig. 10.

Fig. 12. Unsupervised clustering of the sets of points in Figs. 9, 10, 11 obtained by the Figueiredo and Jain method [15]. Each ellipse represents a detected cluster.

for long alignments led to no detection.

As discussed in the introduction, general clustering methods can provide point alignments when a criterion is added to select elongated clusters. We will show results of this approach using the well-known algorithm by Figueiredo and Jain of 2002 [15]. This algorithm adds an important aspect to our comparison: like our algorithm it is unsupervised. Figure 12 shows the results for the same point sets used before. Being a randomized algorithm, different results are obtained at each run; the best results obtained in our tests are presented. The method was used in its standard form, fitting Gaussian mixtures. Each ellipse in the figure corresponds to a Gaussian cluster. The results obtained for the structured point sets are surprisingly good: each one of the perceived alignments and clusters is well represented. The middle result on random point sets is far less satisfactory as it includes many elongated clusters interpretable as alignment detections.

Figure 13 shows the results of the proposed algorithm for the same data sets. As one can see, no detection is produced in the random points, and the three alignments were found. Two of the detections are however shorter than expected and the top vertex of the “A” is missing. Note how the method correctly handled the redundant detections.

Some further results of our method, with increasingly difficulty, are shown in Figs. 14 and 15. The first four are correctly solved. Notice how the very low relative density alignment in Fig. 14(d) was correctly detected. Figure 15 shows some failures and requires a longer comment because they show the limitations of our algorithm. All the alignments in Fig. 15(a) were found; however, the redundancy reduction step did not select the candidates best covering the alignments from a global gestaltic viewpoint. The set of points in Fig. 15(b) is the same already shown in Fig. 1; the alignment found by the algorithm is correct, but as discussed before, does not correspond to the most common interpretation by a human observer. A natural way of handling this problem would be to detect the “curves” by good continuation and then forcing a

global interpretation by methods similar to our masking methodology, that would discard the detected alignment as “masked” by the curves. In Fig. 15(c) the presence of a large cluster masks an alignment: the large number of points causes a large number of tests, raising the detection threshold and imposing a very restrictive test on the alignments. The seven points alignment is detected but the six points one is missing. Handling this example probably implies a round cluster detector but also a recursive approach: once a cluster is detected and removed, our current algorithm would easily detect both alignments. Finally, two of the structures in Fig. 15(d) are slightly curved, rendering them inaccurate as alignments; they are therefore not detected. Again, the detection of good continuation may provide a solution. In short, no alignment detection algorithm can be fully satisfactory *per se*; it requires the interaction (and conflicts) with other feature detectors, namely cluster and curve detectors.

8 CONCLUSION

In this work we have presented a series of algorithms with growing efficiency for detecting alignments of points in a point pattern. The two key aspects of the alignment have been shown to be its local density and its regularity. Our final method combines both criteria into a single coherent detection. We have also introduced a new procedure to resolve the problem of redundant detections. As future work, the same research methodology can be used to address the detection of related structures, such as the *Good Continuation* or curves of points. This will in particular be necessary to resolve the still unsolved *gestalt conflict* between good continuation and alignment pointed out in Fig. 1.

REFERENCES

- [1] N. Ahuja and M. Tuceryan. Extraction of early perceptual structure in dot patterns: Integrating region, boundary, and component Gestalt. *Computer Vision, Graphics, and Image Processing*, 48:304–356, 1989.

Fig. 13. Result of the proposed algorithm for the sets of points in Figs. 9, 10, 11 and 12. Each detection is represented by a thin rectangle divided into boxes, and surrounded by the local window.

Fig. 14. Results illustrating the issues that are solved by the proposed algorithm. For better visualization, the local window and the boxes are not drawn.

Fig. 15. Issues that are still unsolved by the proposed algorithm. For a better visualization, the local window and the boxes are not drawn.

- [2] M. K. Albert and D. D. Hoffman. Genericity in spatial vision. In D. Luce, K. Romney, D. Hoffman, and D'Zmura M., editors, *Geometric Representations of Perceptual Phenomena: Articles in Honor of Tarow Indow's 70th Birthday*, pages 95–112. Erlbaum, 1995.
- [3] D. Amorese, J. L. Lagarde, and E. Laville. A point pattern analysis of the distribution of earthquakes in Normandy (France). *Bulletin of the Seismological Society of America*, 89(3):742–749, 1999.
- [4] A. Arcasoya, V. Toprakb, and N. Kaymakçı. Comprehensive strip based lineament detection method (COSBALID) from point-like features: a GIS approach. *Computers & Geosciences*, 30:45–57, 2004.
- [5] H. Arp and C. Hazard. Peculiar configurations of quasars in two adjacent areas of the sky. *The Astrophysical Journal*, 240(3):726–736, 1980.
- [6] S. Broadbent. Simulating the Ley Hunter. *Journal of the Royal Statistical Society. Series A*, 143(2):109–140, 1980.
- [7] G. Danuser and M. Stricker. Parametric model fitting: From inlier characterization to outlier detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(2):263–280, 1998.
- [8] R. N. Davé. Use of the adaptive fuzzy clustering algorithm to detect lines in digital images. *Intelligent Robots and Computer Vision VIII*, 1192:600–611, 1989.
- [9] A. Desolneux, L. Moisan, and J.-M. Morel. Meaningful alignments. *International Journal of Computer Vision*, 40(1):7–23, 2000.
- [10] A. Desolneux, L. Moisan, and J.-M. Morel. *From Gestalt Theory to Image Analysis, a Probabilistic Approach*. Springer, 2008.
- [11] M. Dubská, A. Herout, and J. Havel. Real-time precise detection of regular grids and matrix codes. *Journal of Real-Time Image Processing*, 2013.
- [12] R. O. Duda and P. E. Hart. Use of the Hough transformation to detect lines and curves in pictures. *Communications of the ACM*, 15(1):11–15, 1972.
- [13] M. G. Edmunds and George G. H. Random alignment of quasars. *Nature*, 290:481–483, 1981.
- [14] A. de la Escalera and J. M. Armingol. Automatic chessboard detection for intrinsic and extrinsic camera parameter calibration. *Sensors*, 10:2027–2044, 2010.
- [15] M. A. T. Figueiredo and A. K. Jain. Unsupervised learning of finite mixture models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(3):381–396, 2002.
- [16] H. Frigui and R. Krishnapuram. A robust competitive clustering algorithm with applications in computer vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(5):450–465, 1999.
- [17] R. Grompone von Gioi and J. Jakubowicz. On computational Gestalt detection thresholds. *Journal of Physiology – Paris*, 103(1–2):4–17, 2009.
- [18] P. Hall, N. Tajvidi, and Malin P. E. Locating lines among scattered points. *Bernoulli*, 12(5):821–839, 2006.
- [19] Ø. Hammer. New statistical methods for detecting point alignments. *Computers & Geosciences*, 35:659–666, 2009.
- [20] Ø. Hammer, D. A. T. Harper, and P. D. Ryan. PAST: Paleontological statistics software package for education and data analysis. *Palaeontologia Electronica*, 4(1), 2001.
- [21] J. Havel, A. Herout, and M. Dubská. Vanishing points in point-to-line mappings and other line parameterizations. *Pattern Recognition Letters*, 34:703–708, 2013.
- [22] P. V. C. Hough. Machine analysis of bubble chamber pictures. *International Conference on High Energy Accelerators and Instrumentation*, 1:554–556, 1959.
- [23] P. V. C. Hough. *Method and Means for Recognizing Complex Patterns*. U.S. Patent 3,069,654, Dec. 18, 1962.
- [24] G. Kanizsa. *Organization in Vision*. Praeger, 1979.
- [25] G. Kanizsa. *Grammatica del vedere*. Il Mulino, 1980.
- [26] G. Kanizsa. *Vedere e pensare*. Il Mulino, 1991.
- [27] D. G. Kendall and W. S. Kendall. Alignments in two-dimensional random sets of points. *Advances in Applied probability*, 12:380–424, 1980.
- [28] N. Kiryati, O. Henricsson, and L. Rosenthaler. On the perception of dotted lines. Technical Report BIWI-TR-137, Institute for Communications Technology, Image Science Laboratory, 1992.
- [29] N. Kiryati and Bruckstein A. M. On navigating between friends and foes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(6):602–606, 1991.
- [30] N. Kiryati and Bruckstein A. M. What's in a set of points? *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(4):496–500, 1992.
- [31] C. D. Litton and J. Restorick. Computer analysis of post hole distributions. *Computer Applications in Archaeology*, pages 85–92, 1983.
- [32] T. M. Lutz. An analysis of the orientations of largescale crustal structures: A statistical approach based on areal distributions of pointlike features. *Journal of Geophysical Research*, 91(B1):421–434, 1986.
- [33] C. Mack. The expected number of aggregates in a random distribution of n points. *Mathematical Proceedings of the Cambridge Philosophical Society*, 46(2):285–292, 1950.
- [34] W. Metzger. *Gesetze des Sehens*. Verlag Waldemar Kramer, Frankfurt am Main, third edition, 1975.
- [35] W. Metzger. *Laws of Seeing*. The MIT Press, 2006 (originally 1936). English translation of the first edition of [34].
- [36] R. E. Miles. On the homogeneous planar Poisson point process. *Mathematical Biosciences*, 6:85–127, 1970.
- [37] F. Murtagh and A. E. Raftery. Fitting straight lines to point patterns. *Pattern Recognition*, 17(5):479–483, 1984.
- [38] A. J. Mussap and D. M. Levi. Amblyopic deficits in detecting a dotted line in noise. *Vision Research*, 40:3297–3307, 2000.
- [39] A. K. Preiss. *A Theoretical and Computational Investigation into Aspects of Human Visual Perception: Proximity and Transformations in Pattern Detection and Discrimination*. PhD thesis, University of Adelaide, 2006.
- [40] C. G. Small. Techniques of shape analysis on sets of points. *International Statistical Review*, 56(3):243–257, 1988.
- [41] J. T. S. Smits, P. G. Vos, and M. P. van Oeffelen. The perception of a dotted line in noise: a model of good continuation and some experimental results. *Spatial Vision*, 1(2):163–177, 1985.
- [42] P. R. Thrift and S. M. Dunn. Approximating point-set images by line segments using a variation of the Hough transform. *Computer Vision, Graphics, and Image Processing*, 21:383–394, 1983.
- [43] S. P. Tripathy, A. J. Mussap, and H. B. Barlow. Detecting collinear dots in noise. *Vision Research*, 39:4161–4171, 1999.
- [44] W. R. Uttal. The effect of deviations from linearity on the detection of dotted line patterns. *Vision Research*, 13(11):2155–2163, 1973.
- [45] W. R. Uttal. *An Autocorrelation Theory of Form Detection*. John Wiley & Sons, 1975.
- [46] W. R. Uttal. *The Perception of Dotted Forms*. Lawrence Erlbaum Associates, 1987.
- [47] W. R. Uttal, L. M. Bunnell, and S. Corwin. 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, 1970.
- [48] M. C. Vanegas, I. Bloch, and J. Inglada. Detection of aligned objects for high resolution image understanding. *IGARSS*, pages 464–467, 2010.
- [49] G. Wadge and A. Cross. Quantitative methods for detecting aligned points: An application to the volcanic vents of the Michoacan-Guanajuato volcanic field, Mexico. *Geology*, 16:815–818, 1988.
- [50] J. Wagemans. Perceptual use of nonaccidental properties. *Canadian Journal of Psychology*, 46(2):236–279, 1992.
- [51] A. P. Witkin and J. M. Tenenbaum. On the role of structure in vision. In J. Beck, B. Hope, and A. Rosenfeld, editors, *Human and Machine Vision*, pages 481–543. Academic Press, 1983.
- [52] D. Zhang and T. Lutz. Structural control of igneous complexes and kimberlites: a new statistical method. *Tectonophysics*, 159:137–148, 1989.
- [53] Y. G. Zhao, X. Wang, L. B. Feng, G. Chen, T. P. Wu, and C. K. Tang. Calculating vanishing points in dual space. In J. Yang, F. Fang, and C. Sun, editors, *Intelligent Science and Intelligent Data Engineering, Third Sino-foreign-interchange Workshop*, pages 522–530. Springer, 2013.
- [54] E. J. Zuiderwijk. Alignment of randomly distributed objects. *Nature*, 295:577–578, 1982.