

The Perceptual Organization of Point Constellations

Matthew J. Dry (matt.dry@psy.kuleuven.be)

Department of Psychology, K.U. Leuven,
Tiensestraat 102, B-3000 Leuven

Daniel J. Navarro (daniel.navarro@adelaide.edu.au)

Kym Preiss (kym.preiss@unisa.edu.au)
School of Psychology, University of Adelaide,
Adelaide, SA, 5005 Australia

Michael D. Lee (mdlee@uci.edu)

Department of Cognitive Sciences, 3151 Social Sciences Plaza A,
University of California, Irvine, CA 92697, USA

Abstract

In this paper we present observers with point patterns based on 30 major star constellations and ask them to connect the points to show the structure they perceive. The resulting empirical structures had a high inter-rater reliability and a high degree of overlap with constellation structures recorded in star atlases, suggesting that the perception of structure in point patterns is largely invariant across individuals. Further, we demonstrate that the empirical structures correspond closely with the structures developed in the field of relational geometry. We discuss the results of the experiment in light of previous findings and suggest a number of potential approaches to formally modeling human performance on clustering tasks.

Keywords: Gestalt organizational principles; visual perception; relational structure; Delaunay triangulation; perceptual organization; perceptual modeling.

Introduction

One of the fundamental tasks of early visual perception is the spatial organization of an image (e.g., Marr, 1982; Ullman, 1984). Researchers from the Gestalt school (e.g., Koffka, 1935; Köhler, 1929; Wertheimer, 1938) demonstrated that the organization of many visual stimuli appeared to be dictated by a number of simple principles such as relative proximity, similarity, good continuation, and common fate. Figure 1a is a replication of one of the examples given by Köhler (1929) as demonstration of organizational structure based on the grouping-by-proximity principle: specifically (all other things being equal) elements with a greater relative proximity tend to be grouped together. Hence, Figure 1a is generally seen as being organized into two groups of three objects rather than one group of two objects and one group of four objects (or any other possible configuration).

Köhler suggested that this principle holds equally well for more complex stimuli, citing the constellations in the night sky as an ecologically plausible example. He noted that “If on a clear night we look up at the sky, some stars are immediately seen as belonging together and, as detached

from the environment. The constellation Cassiopeia is an example, the Dipper is another. For ages people have seen the same groups as units, and at the present time children need no instructions to see the same units” (p. 141-142).

As noted by previous researchers (e.g., Compton & Logan, 1993), while there is a large body of experimental and phenomenological evidence suggesting that the Gestalt organizational principles play an important role in early visual perception little effort has been made to develop formal models of the cognitive or perceptual processes underlying these principles. In this paper we outline a formal approach to describing the perception of structure in constellation (and constellation-like) stimuli. The approach is based upon Delaunay triangulation, a powerful measure of relational structure. In the following sections we describe Delaunay triangulation and outline the aims of the paper.

Delaunay Triangulation

Given any group of co-planar points it is possible to obtain a cell structure that defines the regions within the plane that are closer to each point than any other point. This is known as the Voronoi tessellation of the point set, an example of which is shown in Figure 1b. Joining the points in the set that share common Voronoi edges gives the Delaunay triangulation of the set, an example of which can be seen in Figure 1c.

An important property of the Delaunay triangulation is that it is a super-graph of a number of relational structures: The *Gabriel graph* (Gabriel & Sokal, 1969) is the set of the Delaunay edges that intersect with only one Voronoi edge. The *relative neighborhood graph* (Toussaint, 1980) connects points if no other point is closer to both of them than their inter-point distance. A spanning tree is a structure that connects all of the points in a set with $n-1$ edges and contains no circuits. The *minimum spanning tree* is the structure that minimizes the total length of the edges connecting points (Zahn, 1971). Finally, two points are joined as *nearest neighbors* if one of the points lies closer to the other than to any other point within the set. Examples

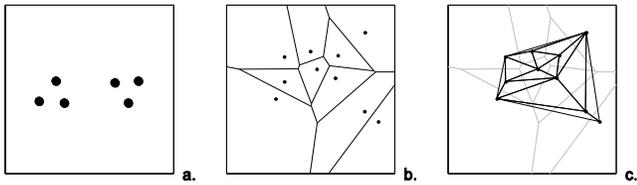


Figure 1. Replication of Köhler’s (1929) demonstration of the proximity principle (a), the Voronoi tessellation of a random point set (b), and the corresponding Delaunay triangulation of the random point set (c).

Examples of these sub-graphs are given in Figure 2. As can be seen, the graphs are hierarchically nested: the nearest neighbors are a subset of the minimum spanning tree, which is a subset of the relative neighborhood graph, and so on, up to the Delaunay triangulation.

Each of these measures has previously been considered in research on structure detection. For example, Zahn (1971) demonstrated that the minimum spanning tree could be used to detect the presence of separate clusters of dots and changes in dot density in random dot textures. Similar demonstrations have also been made by Toussaint (1980) using the relative neighborhood graph, minimum spanning trees and Delaunay triangulation, and by Ahuja and Tuceryan (1989) using Delaunay triangulation. In each of these papers it was demonstrated that algorithms based on the respective relational measures were able to detect the presence of what Toussaint described as “perceptually meaningful” structure.

It is important to note that these relational structures are not merely a convenient geometric measure more suited to computer/artificial vision than human vision. Rather, there is a growing body of psychophysical, physiological and theoretically motivated research suggesting that the human visual system might be generating a Voronoi/Delaunay-like representation at any early stage in visual processing via a spreading activation or ‘grassfire’ process (e.g., Dry, 2008; Kovacs, Feher, & Julesz, 1998; T. Lee, Mumford, Romero, & Lamme, 1998). As such, this form of representation presents a psychologically plausible starting point for developing a formal understanding of the processes underlying perceptual organization.

Aims

The aims of this paper are twofold. First, we are interested in empirically testing Köhler’s suggestion that the perception of structure in constellations is largely invariant across observers. Towards this end we present an experiment in which we asked observers to indicate the structure that they perceive in constellation stimuli. Second, we investigate the degree to which the empirical structures can be described by Delaunay triangulation and its’ sub-graphs.

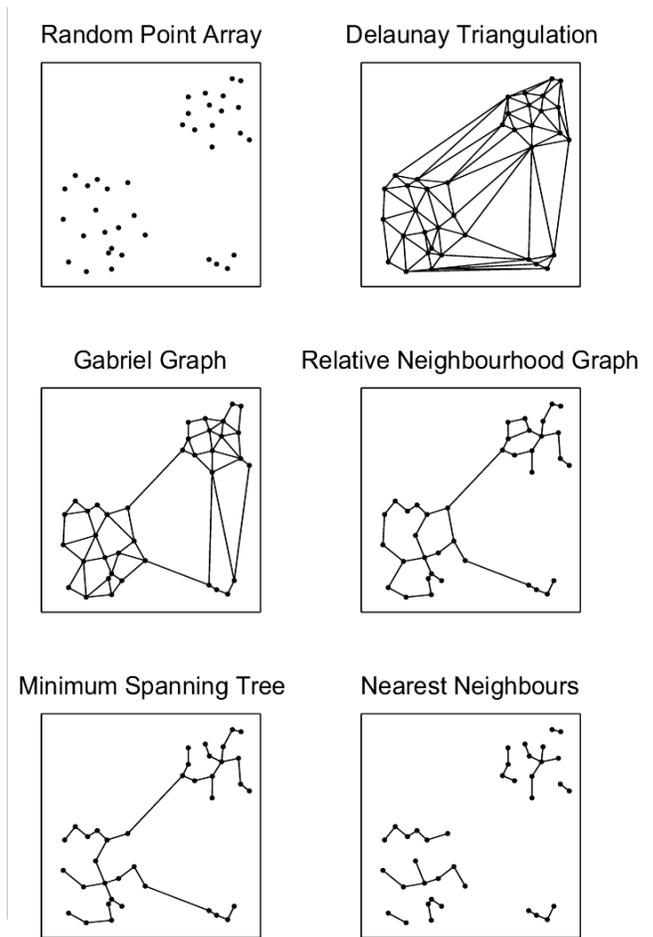


Figure 2. Set of randomly distributed points with its associated Delaunay triangulation and sub-graphs.

Method

Participants

12 observers (six male, six female) participated in the experiment. The mean age of the participants was 28 years. All of the participants were postgraduate psychology students and had normal or corrected-to-normal vision.

Stimuli

30 constellations were selected from the 48 originally identified by Ptolemy (Toomer, 1984). The coordinates of the constellations were taken from the Redshift 3 Desktop Planetarium (RS3). The criteria for selection of a constellation were that it should have 8 or more stars and a structure that was not simply linear. The constellations were flipped across the horizontal axis to minimize the likelihood of the participants recognizing a constellation and reproducing the structure from memory.

Each stimulus was comprised of 0.15 cm diameter black dots presented on a 15 x 15 cm white field.

Procedure

The stimuli were presented on computer monitors. The participants were instructed to connect the points in a stimulus to show the structure they perceive. They were told that they could join any point to any other point that they chose, and make as many or as few links as seemed necessary, with the one provision that the final structure should contain all of the points in the stimulus.

The participants created links between points by left-clicking on a point with the computer mouse. Then, while holding down the mouse button, they drew a path by dragging the mouse cursor to a subsequent point and releasing the button, causing a straight line to be drawn between that point and the previously visited point. By right-clicking on a link to select it and then pressing the 'delete' key on the keyboard, the participants could undo any links they had drawn. The participants were thus free to connect the points in any order, to work alternately from two points, or to work on several separated clusters of points.

The stimuli were presented in a single test session. The order of presentation for the stimuli was randomized across the participants. Prior to the presentation of the experimental stimuli the participants completed three practice stimuli (with 8, 15 and 24 points). Following the experiment the participants were debriefed regarding the aims of the study. None of the participants reported recognizing the stimuli as constellations.

Results

Inter-Rater Reliability

For each of the participants we obtained an n by n matrix of ones and zeros detailing the links between the n points in each stimulus. For example, a link between points 3 and 7 was indicated by a 1 in the third row and seventh column (and seventh row and third column) of the matrix. Points which were not linked had zeros in their corresponding cells.

We measured the reliability of the participants' solutions using split-half correlations between the upper triangles of the link frequency matrices averaged over 10000 random splits. The resulting coefficients gave mean r -values ranging from .89 to .98, suggesting that there is a high degree of overlap between the links chosen by the different participants. This provides support for Köhler's claim that the perception of structure in constellation-like stimuli is largely invariant across individuals.

Overlap Between RS3 And Empirical Structures

The data also indicate that there is a high degree of overlap between the links represented in the RS3 planetarium and the links present in the structures generated by the participants for the constellation stimuli. On average 79% of the empirical links were also present in the RS3 structures, and 80% of the RS3 links were present in the empirical

structures. Figure 3 shows two example stimuli with the RS3 structure indicated by black lines and the empirical structure indicated by white lines. The width of the white lines indicates the frequency with which the participants chose a given link (with thick lines indicating a higher frequency than thin lines). As can be seen, there is a high correspondence between the most frequently chosen empirical links and the links present in the RS3 structure.

We employed a Bayesian approach to assessing the likelihood of finding this degree of overlap by chance (see Navarro, 2008). Briefly, this approach compares the relative likelihoods of four competing explanatory models: M_0 – the empirical and RS3 structures are drawn from populations with different link numerosities and locations; M_1 – the empirical and RS3 structures are drawn from populations with the same link numerosity but different link location; M_2 – the two structures are drawn from populations with the same link numerosity and location; and M_3 – the two structures are drawn from populations with different link numerosities but the same link locations.

The need to consider both link location and numerosity is obvious when one considers the possibility of an empirical structure which connects each node to all other nodes: in this case the overlap between the two structures would be 100%, but the empirical structure would also contain numerous links that are not present in the RS3 structure.

The results of the analyses indicated that for each constellation the most likely model was M_2 , with the next most likely model being *at least* 1.24×10^4 times less likely. In other words, the probability of the empirical and RS3 structures sharing by chance such a high degree of overlap in number and location of structural links is highly unlikely. Again, this result points towards the relative invariance of the perceived organization of this class of stimuli.

Overlap Between Delaunay And Empirical Structures

As mentioned in the Introduction, one of the aims of this study was to investigate the degree to which the empirically produced structures could be described by Delaunay

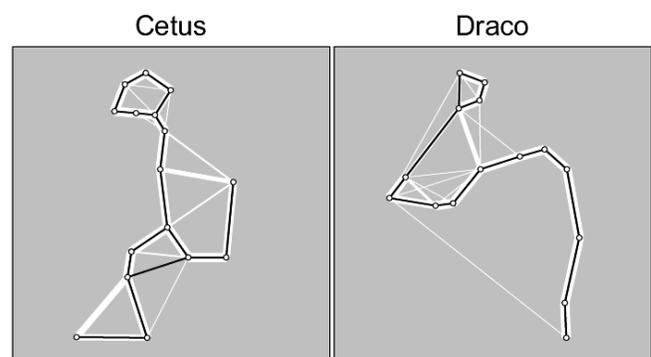


Figure 3. Example constellations with the RS3 structure shown in black, and the aggregated empirical structure shown in white.

triangulation and its subgraphs. Taking an approach analogous to precision/recall analyses in information retrieval research, Figure 4 shows the proportional overlap between the empirical links in the graph structures (y-axis), and the graph links in the empirical structures (x-axis).

For the Delaunay triangulation links the data indicate high recall and low precision: a high proportion of the empirical links are Delaunay triangulation links ($\bar{X} = 98\%$), but the Delaunay structures also contain a high proportion of links that are not present in the empirical structures ($\bar{X} = 58\%$). Conversely, Figure 4 indicates that nearly all of the nearest neighbor links are present in the empirical structures ($\bar{X} = 93\%$), but the empirical structures also contain numerous links that are not nearest neighbors ($\bar{X} = 53\%$).

Figure 4 suggests that the empirical structures are best described by relative neighborhood graph or minimum spanning tree links, with a close correspondence between the proportion of graph links that are empirical links ($\bar{X} = 83\%$ and 86% for RNG and MST respectively) and empirical links that are graph links ($\bar{X} = 84\%$ and 82% for RNG and MST respectively). A qualitative inspection of the empirical structures appears to confirm this result: the participants tended to create structures closely resembling minimum spanning trees, but with additional links employed to close loops.

We employed the previously described Bayesian methodology to determine the likelihood of this degree of overlap occurring by chance. In regards to the overlap between the empirical structures and Delaunay triangulation and nearest neighbor graph structures the analyses indicated that for each stimulus the most likely model was M_3 , indicating that the empirical and graph structures had a high overlap in terms of link locations, but (as could be expected from Figure 4) the link numerosities appeared to be drawn from different populations. In regards to the Gabriel graph link structure the results of the analyses showed that 47% of the stimuli were best described by M_2 indicating a high overlap between both the number and location of structure links, with the remaining 53% of the stimuli best described by M_3 .

Finally, in regards to the overlap between the empirical structures and the relative neighborhood graph and minimum spanning tree graph structures the analyses indicated that for all 30 constellations the correspondence was best described by M_2 , with the next most likely model (which in each case was M_3) being *at least* 3.01×10^{10} and 4.05×10^{10} times less likely for the relative neighborhood graph and minimum spanning tree respectively. In other words, the probability that the empirical structures would by chance share such a high degree of overlap with the minimum spanning tree and relative neighborhood graph structures in terms of both number and location of structure links is extremely low.

Nonetheless, Figure 4 also demonstrates that the links present in the minimum spanning tree and relative neighborhood graph structures are not sufficient to model the empirical structures. Specifically, it can be seen that

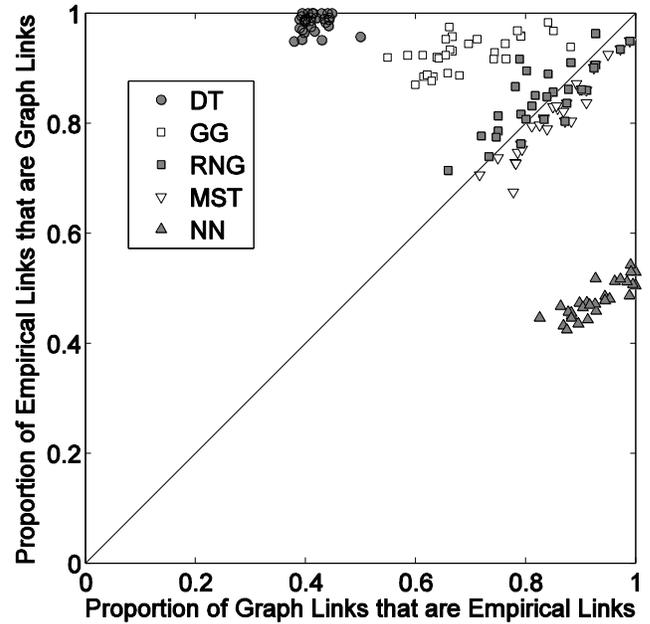


Figure 4. Proportional overlap between graph and empirical structure links for Delaunay triangulation (DT), Gabriel graph (GG), relative neighborhood graph (RNG), minimum spanning tree (MST), and nearest neighbors (NN). Each data point represents one of the 30 stimuli.

there is a wide range in the degree of correspondence between the graph and empirical structures, indicating that in some cases the observers are employing Delaunay links, but are not necessarily employing minimum spanning tree or relative neighborhood graph links. This suggests that in order to model the empirical structures it would be necessary to employ the full Delaunay triangulation, but with some means of restricting or biasing the links to lower-level neighbor relations (i.e., maintaining the recall, but improving the precision).

Figure 5 indicates the plausibility of such an approach. If the Delaunay neighbors are indexed ordinally (i.e., 1st nearest Delaunay neighbor, 2nd nearest Delaunay neighbor, ... k^{th} nearest Delaunay neighbor) it can be seen that the majority of the empirical links are captured by Delaunay neighbors of order ≤ 5 (top panel). Furthermore, the majority of the empirical links are first order neighbors (i.e., nearest neighbors), with the frequency of inclusion decreasing as a function of neighbor order (lower panel).

It should be noted that for uniformly random distributions of points the most prevalent number of Delaunay neighbors is 6, with 5 and 7 neighbors being roughly three quarters as prevalent. Given this it is unsurprising that the proportion of high-order Delaunay neighbors present in the empirical structures is low. Nonetheless, Figure 5 indicates that the empirical links are not chosen with a uniform probability – rather there appears to be a strong bias towards including very low-order neighbors (i.e., 1st to 3rd nearest Delaunay neighbors).

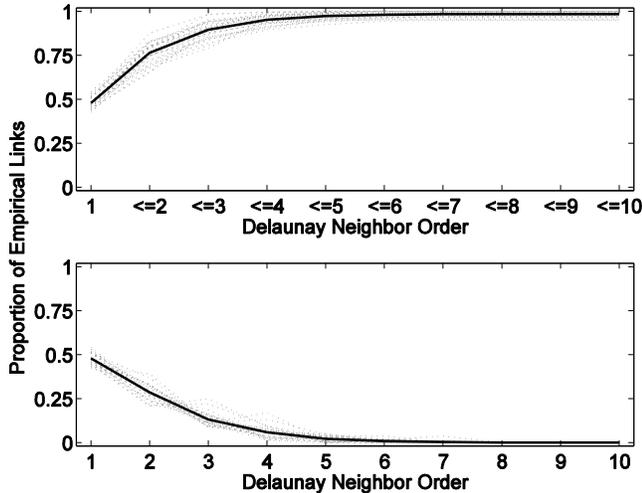


Figure 5. The cumulative distribution function for the proportion of empirical links that are 1st order Delaunay neighbors through to 1st, 2nd, 3rd ... and 10th order neighbors (top panel), and the probability density function for the proportion of empirical links that are 1st order neighbors only, through to 10th order neighbors only (bottom panel). Individual stimuli are shown in gray, with the average shown in black

Controlling For Familiarity

As has been indicated none of the participants reported recognizing the experimental stimuli as constellations. However, in order to control for the possibility that the participants unconsciously reproduced the constellation structures from memory we ran a control study (using the same 12 participants from the current study) which employed 30 random point patterns or ‘pseudo-constellations’ which were matched to one of the 30 constellations in terms of numerosity and degree of clustering as measured by $R = r_o/r_e$, where r_o is the mean nearest neighbor distance for the n points in the stimulus and r_e is the expected mean for a uniformly random distribution of n points. The results of the control study provided the same pattern of results as those reported for the constellation stimuli in terms of inter-rater reliability, proportional overlap with Delaunay and Delaunay sub-graph link structure. Given this we can safely assume that any familiarity effects are either negligible or non-existent. Furthermore, the control study provides an important indication of the generality of the results of the constellation experiment.

Discussion

The high inter-rater reliability and high degree of overlap between the RS3 and empirical constellation structures provide support for Köhler’s suggestion that the perceptual organization of constellation-like stimuli is relatively invariant across individuals. Köhler and colleagues argued that perceptual organization is driven by universal principles or processes, citing phenomenological examples such as

Figure 1a as evidence. In this study we provide quantified empirical evidence pointing towards the same conclusion.

The finding that the empirical constellation structures were best described by minimum spanning tree or relative neighborhood graph structure is analogous to the results of a similar experiment reported in Pomerantz (1981) in which participants were asked to join up the dots in semi-random point patterns to show the structure that they perceived. As with the current study the results indicated that the empirical structures tended to be minimum spanning trees with a small number of additional links closing loops or adding to the overall symmetry of the structure.

Pomerantz noted that the participants’ production of structures corresponding closely to minimum spanning trees could be interpreted as empirical evidence of the law of Prägnanz (the minimum principle). In other words the empirical structures tended towards the simplest possible configuration of the point set in that they had close to the minimum number of links needed to create a tree structure, and came close to minimizing the overall length of the links included in the structure. Furthermore, it was suggested that the fact the participants generated these close-to-minimal structures without prompting could be taken as evidence that much of perceptual organization occurs in a bottom-up fashion without reference to top-down processes based on strategy or learning.

If this is indeed the case, then the results of these experiments might provide insight into human performance on difficult optimization tasks such as the visually presented traveling salesperson problem (TSP). Solving a TSP involves finding the shortest pathway through a set of N cities that begins and ends at the same city. The number of potential solutions to a TSP instance increases factorially as the number of cities in the instance increases, such that for a 5 city instance there are 12 pathways, for a 10 city instance there are 181,400 pathways and for a 15 city instance there are 4×10^{10} pathways. Despite this apparent intractability research has shown that human participants are able to generate near-optimal solutions to TSPs in a timeframe that is a close-to-linear function of problem size (e.g., Dry, M. Lee, Vickers, & Hughes, 2006).

This finding might be explicable in terms of a bottom-up process that is biased towards organizing visual stimuli such that the resulting structure is simple or minimal. If the base representation or initial clustering of a TSP is a minimal structure (by virtue of the bottom-up process employed to generate the representation), then producing a minimal pathway via some form of top-down cluster joining heuristic should be far more efficient than a path-finding heuristic that works entirely from a top-down perspective seeking to actively *impose* minimality on a raw stimulus.

The results of the current study provide some important insights into formally modeling human performance on perceptual clustering tasks. Firstly, the analyses indicate that the empirical structures can be well described by grouping heuristics based upon relative proximity alone. Specifically, it was not necessary to employ more complex heuristics

such as good continuation or symmetry in order to describe the empirical structures. Nonetheless, the fact that the participants were creating near-minimal structures with the addition of extra links to close loops or add some form of balance or symmetry to the structure suggests that some of the remaining variance between the graph and empirical structures might be accounted for by these additional heuristics.

Secondly, the data in Figure 5 suggest that it might be possible to simulate the empirical structures using a model that links together neighboring points in a hierarchical manner by initially forming clusters based on nearest neighbor or low-level Delaunay neighbor links, and then joining these clusters into a single structure. A similar approach has been suggested in relation to modeling human performance on the traveling salesperson problem (e.g., Dry et al., 2006). Preliminary analyses have shown that this form of approach is able to produce structures that have a high degree of overlap with the empirical constellation structures (mean $r = .92$)

There are a number of alternative previously published models that have also been applied to the detection of structure in dot stimuli using spatial filtering (e.g., Smits & Vos, 1986) or some form of relational information (e.g., Caelli, Preston, & Howell, 1978; Pizlo, Salach-Golyska, & Rosenfeld, 1997). Furthermore, Vickers, Navarro and M. Lee (2000) suggested that the visual system might extract structure from point sets by searching for transformations (e.g., rotations, translations, etc) that generate an output that is maximally symmetric with the original image, and demonstrated that such an approach could produce a link structure for the constellation Perseus that closely resembled the structure present in star atlases. It would be highly interesting to compare the performance of these different models on the constellation task to determine which of these approaches provides a better account of the processes underlying human perceptual organization.

Acknowledgements

We would like to thank Douglas Vickers who first drew our attention to the high overlap between constellation structure and the Delaunay triangulation sub-graphs. DJN was supported by an Australian Research Fellowship (ARC grant DP-0773794).

References

Ahuja, N., & Tuceryan, M. (1989). Extraction of early perceptual structure in dot patterns - Integrating region, boundary, and component Gestalt. *Computer Vision Graphics and Image Processing*, 48(3), 304-356.

Caelli, T. M., Preston, G. A. N., & Howell, E. R. (1978). Implications of spatial summation models for processes of contour perception: a geometric perspective. *Vision Research*, 18, 723-734.

Compton, B. J., & Logan, G. D. (1993). Evaluating a computational model of perceptual grouping by proximity. *Perception & Psychophysics*, 53, 403-421.

Dry, M. J. (2008). Using relational structure to detect symmetry: a Voronoi tessellation based model of symmetry perception. *Acta Psychologica*, 128, 75-90.

Dry, M. J., Lee, M. D., Vickers, D., & Hughes, P. (2006). Human performance on visually presented traveling salesperson problems with varying numbers of nodes. *Journal of Problem Solving*, 1(1), 20-32.

Gabriel, K. R., & Sokal, R. R. (1969). A new statistical approach to geographic variation analysis. *Systematic Zoology*, 18, 259-278.

Koffka, K. (1935). *Principles of Gestalt psychology*. New York: Harcourt.

Köhler, W. (1929). *Gestalt Psychology*. New York: Liveright.

Kovacs, I., Feher, A., & Julesz, B. (1998). Medial-point description of shape: a representation for action coding and its psychophysical correlates. *Vision Research*, 38, 2323-2333.

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

Marr, D. (1982). *Vision*. San Francisco: W.H. Freeman and Company.

Navarro, D. J. (2008). *Bayesian tests for correlated binary choices* (Technical Note available online at <http://www.psychology.adelaide.edu.au/personalpages/staff/danielnavarro/papers.html>).

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

Pomerantz, J. R. (1981). Perceptual organization in information processing. In M. Kubovy & J. R. Pomerantz (Eds.), *Perceptual Organization* (pp. 141-180). NJ: Lawrence Erlbaum.

Smits, J. T. S., & Vos, P. G. (1986). A model for the perception of curves in dot figures: The role of local salience of "virtual lines". *Biological Cybernetics*, 54, 407-416.

Toomer, G. J. (1984). *Ptolemy's Almagest*: Springer-Verlag.

Toussaint, G. T. (1980). The relative neighborhood graph of a finite planar set. *Pattern Recognition*, 12(4), 261-268.

Ullman, S. (1984). Visual Routines. *Cognition*, 18, 97-159.

Vickers, D., Navarro, D. J., & Lee, M. D. (2000). Towards a transformational approach to perceptual organisation. In R. J. Howlett & L. C. Jain (Eds.), *KES 2000: Proceedings of the Fourth International Conference on Knowledge-Based Intelligent Engineering Systems & Allied Technologies* (Vol. 1, pp. 325-328). Piscataway, NJ: IEEE.

Wertheimer, M. (1938). Laws of organization in perceptual forms. In W. Ellis (Ed.), *A Source Book of Gestalt Psychology*. New York: Harcourt.

Zahn, C. (1971). Graph-theoretical methods for detecting and describing Gestalt structures. *IEEE Transactions on Systems, Man and Cybernetics*, 20(1), 68-86.