



# A taxonomy of glyph placement strategies for multidimensional data visualization

Matthew O Ward<sup>1</sup>

<sup>1</sup>Computer Science Department, Worcester Polytechnic Institute, 100 Institute Road, Worcester, MA 01609 U.S.A.

**Correspondence:**

Matthew O Ward, Computer Science Department, Worcester Polytechnic Institute, 100 Institute Road, Worcester, MA 01609 U.S.A.  
Tel: (508) 831-5671; Fax: (508) 831-5776;  
E-mail: matt@cs.wpi.edu

**Abstract**

Glyphs are graphical entities that convey one or more data values via attributes such as shape, size, color, and position. They have been widely used in the visualization of data and information, and are especially well suited for displaying complex, multivariate data sets. The placement or layout of glyphs on a display can communicate significant information regarding the data values themselves as well as relationships between data points, and a wide assortment of placement strategies have been developed to date. Methods range from simply using data dimensions as positional attributes to basing placement on implicit or explicit structure within the data set. This paper presents an overview of multivariate glyphs, a list of issues regarding the layout of glyphs, and a comprehensive taxonomy of placement strategies to assist the visualization designer in selecting the technique most suitable to his or her data and task. Examples, strengths, weaknesses, and design considerations are given for each category of technique. We conclude with some general guidelines for selecting a placement strategy, along with a brief description of some of our future research directions.

*Information Visualization* (2002) 1, 194–210. doi:10.1057/palgrave.ivs.9500025

**Keywords:** Multivariate data visualization; information visualization; glyphs; layout algorithms.

**Introduction**

It is well known that we are experiencing an explosion in the volume and complexity of data to which we have ready access. Scientific measurements, business transactions, surveys, simulations, and a host of other activities contribute to this growing mass, and our ability to effectively analyze the data has become increasingly strained. Much of this data is multivariate in nature, whether it be from tables, spreadsheets, multispectral sensors, or complex computations. Each row or data entry may have anywhere from two to several thousand entries, and the larger the number and size of the entries, the harder it is to detect, classify, and measure features and relations of interest.

Visualization has emerged as a powerful tool for the exploration of such large and complex data sets. While algorithmic analysis can be used to quickly and accurately process data to identify patterns and outliers, it is dependent on having a computational model of the phenomena of interest. The problem is that one may not know what one is looking for, or may not be able to set fixed parameters and thresholds to effectively guide the analysis. Visualization, on the other hand, uses the human perceptual system to extract meaning from the data, focus attention, and reveal structure and patterns. The author's Visual Information Seeking Mantra, I'll Know it When I See it (IKIWISI), which is a follow up to Shneiderman's Mantra<sup>1</sup> overview first, zoom and filter, then details on demand, is a reflection of the notion that we often do not know what we are looking for when presented with a data set, and we rely on our visual pattern recogni-

tion system to help us in gaining knowledge of the data. It is also relevant to the concept that visualization is a powerful mechanism for confirming hypotheses about data phenomena; quantitative analysis may indicate that some relation exists, but a supporting visualization can improve our confidence in the results and our ability to remember them.

Glyphs are a popular method for conveying information visually. Individual dimensions/variables for a given data point are mapped to attributes of a particular shape or symbol, and variations, clusterings, and anomalies among these graphical entities may be readily perceived. Glyphs are a powerful communication mechanism in that a large number of data dimensions can be incorporated into the attributes of a single shape or symbol. The particular mappings may also be customized to reflect semantics relevant to specific domains, which can greatly facilitate the interpretation process.

Once a glyph has been designed and generated from a data entry, it must be placed at a location in 2D or 3D display space. The position attribute can be very effective in communicating data attributes or improving the detection of similarities, differences, clustering, outliers, or relations in the data. Many strategies exist for setting the position attribute; some are based on the raw data or information derived from the data, while others are based on the structure of the data set (e.g., ordered or hierarchical). Methods also vary as to whether they allow overlapping glyphs, whether they are space-filling, or whether they employ empty space to highlight differences.

The goal of this paper is to present a comprehensive taxonomy of glyph placement strategies to support the design of effective visualizations. We first present a formal definition of glyphs, how they are formed, examples of glyphs that have been used in the past, and some limitations of glyphs for visualizing multivariate data. We then describe a categorization of methods for positioning glyphs, present several examples, and describe limitations and strengths of each. We conclude with some general strategies for selecting a placement strategy, and point to some directions for future research.

## Glyph fundamentals

*Multivariate data*, also called multidimensional or  $n$ -dimensional, consists of some number of points,  $m$ , each of which is defined by an  $n$ -vector of values. Such data can be viewed as an  $m \times n$  matrix, where each row represents a data point and each column represents an observation, variable or dimension. An observation may be scalar or vector, nominal or ordinal, and may or may not have a distance metric, ordering relation, or absolute zero. Each variable/dimension may be independent or dependent. For this paper, we assume that each dimension is an independent bounded numeric scalar value that we can normalize to the range  $[0 \rightarrow 1]$ .

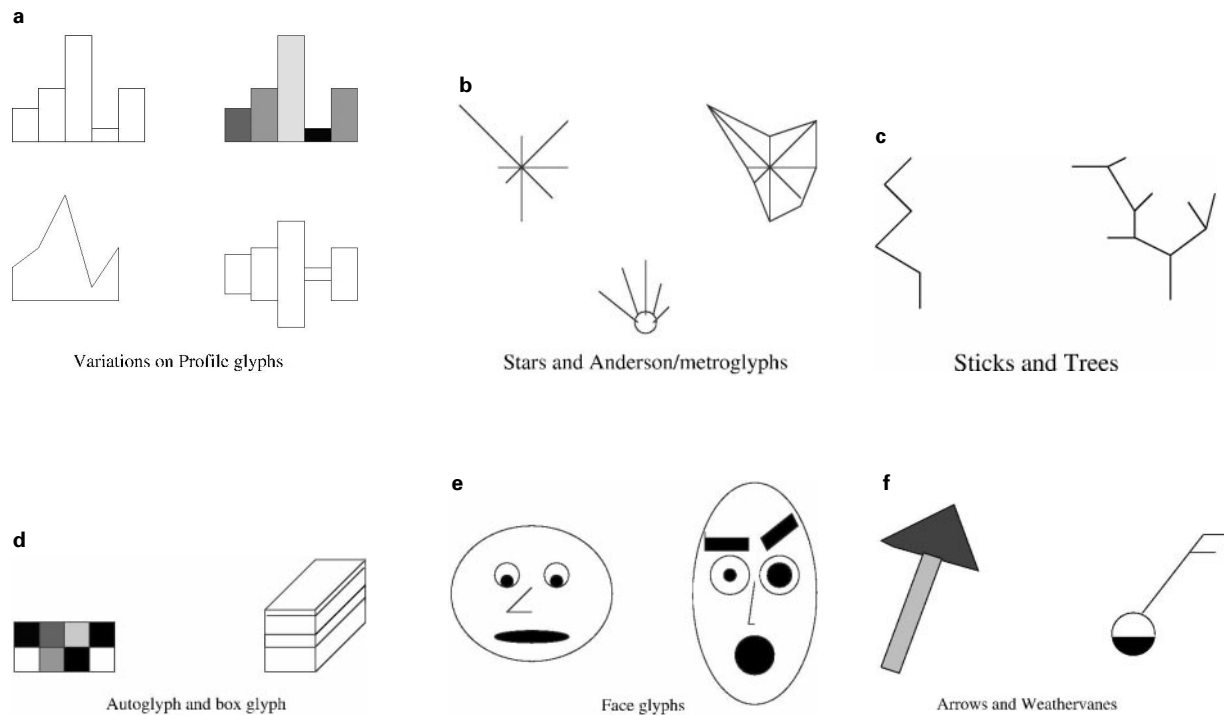
A *glyph* consists of a graphical entity with  $p$  components, each of which may have  $r$  geometric attributes

and  $s$  appearance attributes. Typical geometric attributes include shape, size, orientation, position, and direction/magnitude of motion, while appearance attributes include color, texture, and transparency. Attributes can be discrete or continuous, scalar or vector, and may or may not have a distance metric, ordering relation, or absolute zero.

The process of creating a glyph thus becomes one of *mapping* one or more data dimensions for a data point to one or more geometric and/or appearance attributes of one or more components of a graphical entity. A glyph may also contain components or attributes that are independent of the data being used in its formation. For example, in a scatterplot the size and color of the plotting symbol may be data-independent, such as using color to indicate the data point has been selected.

Figure 1 shows a variety of examples of glyphs that have been proposed and used in the past. The list below describes these and other glyphs from the literature in terms of the graphical entities and attributes controlled by the multivariate data point.

- profiles<sup>2</sup>: height and color of bars (Figure 1a).
- stars<sup>3</sup>: length of evenly spaced rays emanating from center (Figure 1b).
- Anderson/metroglyphs<sup>4,5</sup>: length of rays (Figure 1b).
- stick figures<sup>6</sup>: length, angle, color of limbs (Figure 1c).
- trees<sup>7</sup>: length, thickness, angles of branches; branch structure derived from analyzing relations between dimensions (Figure 1c).
- autoglyph<sup>8</sup>: color of boxes (Figure 1d).
- boxes<sup>9</sup>: height, width, depth of first box; height of successive boxes (Figure 1d).
- hedgehogs<sup>10</sup>: spikes on a vector field, with variation in orientation, thickness, and taper.
- faces<sup>11</sup>: size and position of eyes, nose, mouth; curvature of mouth; angle of eyebrows (Figure 1e).
- arrows<sup>12</sup>: length, width, taper, and color of base and head (Figure 1f).
- polygons<sup>13</sup>: conveying local deformation in a vector field via orientation and shape changes.
- dashtubes<sup>14</sup>: texture and opacity to convey vector field data.
- weathervanes<sup>15</sup>: level in bulb, length of flags (Figure 1f).
- circular profiles<sup>16</sup>: distance from center to vertices at equal angles.
- color glyphs<sup>17</sup>: colored lines across a box.
- bugs<sup>18</sup>: wing shapes controlled by time series; length of head spikes (antennae); size and color of tail; size of body markings.
- wheels<sup>18</sup>: time wheels create ring of time series plots, value controls distance from base ring; 3D wheel maps time to height, variable value to radius.
- boids<sup>19</sup>: shape and orientation of primitives moving through a time-varying field.
- procedural shapes<sup>20,21</sup>: blobby objects controlled by up to 14 dimensions.



**Figure 1** Examples of glyphs. Top row: (a) variations on profiles; (b) stars/metroglyphs; and (c) stick figures and trees. Bottom row: (d) autoglyphs and boxes; (e) faces; and (f) arrows and weathervanes.

- Glyphmaker<sup>22</sup>: user-controlled mappings.
- Icon Modeling Language<sup>23</sup>: attributes of a 2D contour and the parameters that extrude it to 3D and further transform/deform it.

### Glyph limitations

Glyphs are not without limitations in the communication of multivariate data. Most, if not all, mappings introduce biases in the process of interpreting relationships between dimensions. Some relations are much easier to perceive than others, such as data dimensions mapped to adjacent components. Some recent research<sup>24</sup> has focused on clustering and reordering of data dimensions to improve perception of relationships. The accuracy with which humans perceive different graphical attributes varies tremendously as well.<sup>25</sup> For example, our ability to accurately measure length is superior to attributes such as orientation and color. In addition, the accuracy can vary significantly between individuals and even for a single observer in different contexts. Color perception has been shown to be extremely sensitive to context;<sup>26,27</sup> background color as well as the color of adjacent entities can modify our judgment of the color of a glyph.

There are also limitations based on the media being used to communicate the information. Screen space and resolution are limited, and displaying too many glyphs at once can lead to either overlaps that can hinder accurate discernment of individual dimensions or very

small glyphs (though, as we shall see, dense packing can form texture patterns for global analysis). Dynamic scaling control for glyphs can thus be a critical component for an effective interface to glyph displays. Finally, having too many data dimensions in the mapping can make it hard to discriminate or interpret the individual dimensions.

### Glyph placement issues

The first consideration when selecting a placement strategy is whether the placement will be *data-driven*, eg, based on two or more data dimensions, or *structure-driven*, such as methods based on an explicit or implicit order or other relationship between data points.

A second consideration is whether overlaps between glyphs will be allowed. This can have a significant impact on the size of the data set that can be displayed, the size of the glyphs used, and the interpretability of the resulting images.

A third consideration is the trade-off between optimized screen utilization, such as found in space-filling algorithms, versus the use of white space to reinforce distances between data points.

A fourth consideration is whether the glyph positions can be adjusted after initial placement to improve visibility at the cost of distorting the computed position. Overlapping glyphs can be difficult to interpret, but any movement alters the accuracy of the visual depiction. We need to know, for the given domain, what are the trade-offs between accuracy and clarity.

### Glyph placement strategies

Figure 2 shows a taxonomy of glyph placement strategies for multivariate data sets. In the next two sections we describe each class of strategies and give examples of their use. To formalize the different algorithms we assume the following. Let  $V=(v_1, v_2, \dots, v_N)$  correspond to the  $N$  variables or dimensions of the data set, and  $D=(d_1, d_2, \dots, d_M)$  correspond to the  $M$  data records. Each record  $d_i=(o_{i1}, o_{i2}, \dots, o_{iN})$  contains  $N$  numeric observations normalized to the range  $[0.0 \rightarrow 1.0]$ . Finally, the set  $L=(l_1, l_2, \dots, l_M)$  contains the locations in normalized display space for the  $M$  data points. For the examples and algorithms in this paper, we assume the display space to be two dimensional, though many of the techniques and concepts readily extend to other display dimensions.

### Data-driven glyph placement

In *data-driven* placement, the data are used to compute or specify the location parameters for the glyph. The two categories of this strategy class are *raw* and *derived*. The placement algorithm can be characterized as a function that maps  $D \rightarrow L$ , and may involve either local knowledge, such as when  $l_i$  is computed based solely on  $d_i$ , or global knowledge, as when  $l_i$  is computed using the entirety of  $D$  or both  $D$  and  $L$ . Each category of data-driven placement strategies can be further divided based on whether or not the computed positions are subjected to a modification/distortion process in order to reduce visual clutter or overlap. Examples of strategies, both for original placement and distortion, are described below.

### Raw data strategies

In raw data placement strategies, one, two or three of the data dimensions are used as positional components;

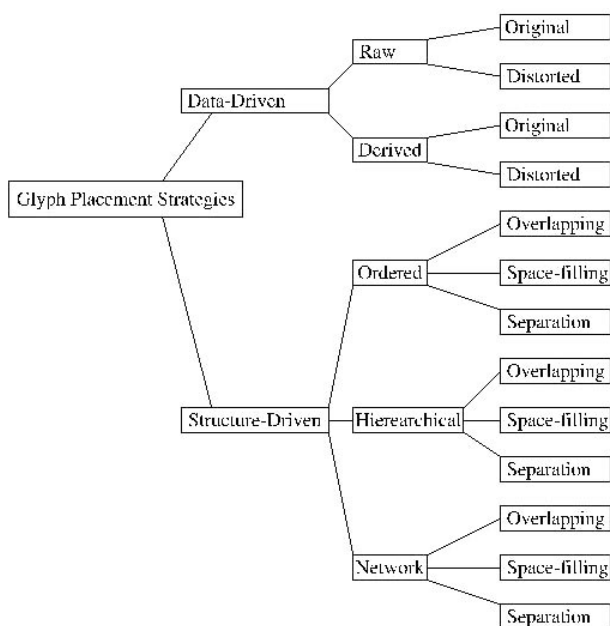


Figure 2 Taxonomy of glyph placement strategies.


these dimensions may also be redundantly represented in the glyph itself. Thus for two dimensional display space,  $l_i=(o_{ia}, o_{ib})$ , where  $0 < a, b \leq N$ . This method conveys detailed relationships between the dimensions selected, as our ability to detect positional, as opposed to shape, patterns is quite strong. Many systems and techniques have been developed using this strategy. IVEE, now known commercially as Spotfire, allows users to select the mappings of data dimensions to both position and graphical attribute.<sup>28</sup> Becker *et al*<sup>29</sup> used geographically located glyphs to convey communication network traffic attributes. Many researchers<sup>30,31</sup> have employed the cityscape metaphor to place 3D box glyphs at data driven positions. Figure 3 shows all pairwise mappings of the Iris data set, which consists of 150 points and four dimensions: petal and sepal length and width. Note that while each view separates one iris family (sailboat shape) from the other two (kite shape), varying degrees of separation can be seen within the large cluster. Also, some views reveal a number of outliers.

Since for  $N$ -dimensional data and a two-dimensional display there are  $N(N-1)/2$  possible mappings (we consider only orthogonal views in the  $N$ -dimensional space, and do not include rotational variations), it is critical that this mapping be done carefully. An ineffective mapping can result in substantial cluttering and poor screen utilization. For example, if one of the dimensions used for the mapping only takes on a small number of distinct values, a significant amount of overlap is possible. Also, depending on the semantics of the data, some mappings may be more meaningful than others to the person interpreting the display. Analytic techniques, such as the measures of interestingness used in projection pursuit algorithms,<sup>32,33</sup> can be used to identify potentially informative views, such as those exhibiting clustering or significant separation.

The major strengths of this placement strategy are its ease of interpretation, low computational overhead, and effectiveness at uncovering correlations between the dimensions controlling the position, such as in a scatterplot. Among the limitations of this approach is the fact that bias is given to the dimensions involved in the mapping, and thus conveys only pairwise (or three-way, for 3-D) relations between the selected dimensions. Structure that is not orthogonal to the viewing direction may be difficult to discern. The issue of occlusion is also a significant concern when using this strategy; this is discussed in the section below on distortion methods.

### Derived data strategies

A derived data placement technique uses an analytic process to generate positions using the data values, and perhaps other information, as input. Thus, instead of a location reflecting only one, two, or three of the data dimensions, it reflects some combination of all the dimensions in an attempt to convey  $N$ -dimensional relational information in a smaller number of display dimensions. Some commonly used dimensionality reduc-



**Figure 3** All pairwise raw data-driven views (star glyphs) of the Iris data set: (a) sepal length vs. sepal width, (b) sepal length vs. petal length, (c) sepal length vs. petal width, (d) sepal width vs. petal length, (e) sepal width vs. petal width, and (f) petal length vs. petal width.



tion techniques include Principal Component Analysis (PCA),<sup>34</sup> Multidimensional Scaling (MDS),<sup>35–38</sup> Self-Organizing Maps (SOMs),<sup>39,40</sup> simulated annealing,<sup>41</sup> spring-based models,<sup>42</sup> and other optimization techniques.<sup>43</sup> PCA attempts to find linear combinations of the dimensions that explain the largest variation in the multivariate data set. The first two principal components can be used to specify the position of a glyph. SOMs and MDS are iterative refinement/optimization processes that attempt to adjust weights or positions until a certain criteria is met; in our case, the distances or similarities between points in 2-D is a good approximation of the N-D distances/similarities. Variations in the starting conditions and distance or similarity measure can have significant impact on the results of these methods.

As an example of this sort of derived placement, Figure 4 displays the iris and cars data sets, using PCA to determine glyph positions. One can see the first two principal components were reasonably successful at separating the three classes of iris, as depicted by the small sailboat shapes and large and medium size kite shapes. Likewise, clusters of similar shapes can be seen in the cars data set.

Another type of derived placement algorithm places objects in terms of their relationships with one or more reference objects, which may or may not be part of the original set. These techniques, often referred to as barycentric methods,<sup>44</sup> start with  $B$  reference objects ( $r_1, r_2, \dots, r_B$ ) located at the vertices of a convex shape. All data points are then positioned within the shape based on their relations to these reference points. The barycentric coordinates of point  $P$  form a linear combination with the reference points, such that  $P = a_1 r_1 + a_2 r_2 + \dots + a_B r_B$ , where  $a_1 + a_2 + \dots + a_B = 1$ . The coordinates ( $a_1, a_2, \dots, a_B$ ) define a unique location within the shape that can be derived from a data set. Each choice regarding the size and content of the reference set generates a different view; thus the setting of these parameters is crucial in creating an informative visualization. A variation on this approach, described in,<sup>45</sup> uses  $A$  equally spaced anchors

around a circle, with each data point attached to the anchors by springs with constants proportional to the data values corresponding to the dimensions/anchors. Another related technique, the *bull's eye* display,<sup>46</sup> places objects in concentric rings relative to their distance from a selected node of interest. Yet another technique based on relations to anchor positions was presented by Spoerri,<sup>47</sup> where the corners of a polygon represent concepts and distinct locations within the polygon represent all the possible boolean relations among those concepts. The *Sunflower* metaphor presented by Rose<sup>48</sup> has a similar goal, using overlapping circles to convey all possible relationships between one or more documents and a set of domains.

In general, the techniques described above reinforce the perception of similar shapes by attempting to place them close to each other. However, the techniques do have their limitations. The resulting display coordinates, unlike in the raw data-driven techniques, have no semantic meaning. Instead they can be viewed as supplemental dimensions to the original data set. PCA assumes that the majority of the variation in a data set will be well embodied in the first few principal components, which is not always the case. MDS and SOMs, like all iterative optimization procedures, are not guaranteed to be optimal, and the results are generally not unique. Iterative strategies can also be computationally demanding, although several efforts have been recently made to improve the processing time.<sup>36,49</sup>

A category of data-driven placement strategies that can be perceived as a hybrid of raw and derived techniques can be found in 2-D and 3-D field data, where uniform or non-uniform samples are intended to represent a continuous field of scalar, vector, or tensor values. In these strategies, positions are first computed, and then the data are resampled to generate the values corresponding to the positions. These positions are often specified on a user-controlled regular grid, but significant research has been performed on using the data to compute positions that are most effective for



Figure 4 Star glyphs of the Iris and cars data sets with position based on the first two principal components.

conveying information. For example, while most flow visualizations place arrow glyphs on a regular grid, Globus *et al*<sup>50</sup> and Helman and Hesselink<sup>51</sup> describe experiments in placing glyphs at critical points in the flow field. This means that significant structures are less likely to be overlooked. Another approach to placing glyphs on a vector field, described by Dovey,<sup>52</sup> allows resampling either in physical or parameter space, thus enabling the technique to be applied to curvilinear and unstructured grids. Finally, Fuhrmann and Groller<sup>14</sup> describe a method to place 3-D textured streamlines, called *dashtubes*, to optimize uniformity in spacing. Because these methods compute positions based on the data, we have grouped them with the data-driven, derived techniques. However, one could argue that they are not so much placement strategies as filtering strategies, as a continuum of possible glyphs is being filtered to select the subset to be shown.


#### Distortion techniques for data-driven placement

Once the positions for the glyphs have been computed, a possible post-processing step for data-driven techniques would involve distorting these positions to reduce clutter and overlap. For example, random jitter has been employed in statistical graphics when data-driven positioning is being used with discrete (as opposed to continuous) dimensions.<sup>53</sup> Other approaches try to shift positions to minimize or avoid overlaps. Keim and Hermann<sup>54</sup> use a quadtree structure to relocate points to maintain proximity to their original positions while eliminating overlaps. One concern of this approach is the level of distortion that is introduced. While a minor shifting probably will not significantly change the interpretation of the display, in cases where the data density is extremely high, a point's final position may be far from its initially assigned location. Another variation on position shifting was presented by Chuah *et al*,<sup>30</sup> where elements of interest on a planar city-scape could be elevated above the plane and brought closer to the viewer

for inspection. Woodruff *et al*<sup>55</sup> employ several strategies commonly used in map generalization to maintain a constant density of glyphs on the screen. These strategies include *aggregation* (combining low-level glyphs into high-level ones), *deletion* (removing glyphs based on proximity and similarity to their neighbors), and *replacement* (using smaller, less informative glyphs in areas of high density).

Distortion techniques have also been used to selectively vary the level of detail shown in the visualization.<sup>56</sup> In this way, regions of dense glyphs can be assigned a larger segment of the display space, with either position or position and scaling modified to emphasize or de-emphasize subsets of the data. These methods generally act on a local, rather than global, scale, requiring one or more focus points and the specification of a region of influence for the distortion process. Thus it is critical to the effectiveness of these methods to provide users interactive control of the transformation to facilitate variation of the focus and maintenance of context, ie, enabling users to retain a global view of the data while expanding the screen space dedicated to the data subset of interest. Smooth animation between the original and distorted views is also a powerful mechanism in this regard.

A simple distortion technique for reducing glyph overlap is outlined below. Note that no constraint is placed on how far a glyph can move from its original position. Also, if the maximum distance a glyph can move per iteration is set too high, the procedure may increase, rather than decrease, the number of overlaps. Finally, no provision is made for oscillating behaviors, such as a glyph moving back and forth between two or more other glyphs. A small random perturbation to the glyph offsets could solve this problem. Figure 5 shows the breakfast cereal data set with data-driven positioning, without and with overlap reduction. Note that some data features that had been occluded in the original placement strategy are now visible.



**Figure 5** Star glyphs of 7 dimensions of the breakfast cereal data set with position based on fiber and sugar levels. The left image is the original mapping, and the right was distorted to reduce the level of overlap.

Let *Tolerance*=the maximum percentage of overlapping glyphs.

Thus  $\text{MaxOverlaps} = M * \text{Tolerance}$ .

Let *Size*=the maximum size for a glyph.

Let  $D(i,j)$ =the distance between the centers of glyphs *i* and *j*.

Thus  $\text{NumberOverlaps}$ =number of distinct glyph pairs for which  $D(i,j) < \text{Size}$ .

Let *Speed*=the maximum distance a glyph can move per iteration (we suggest this be a fraction of *Size*).

Let *Iterations*=the maximum number of iterations of adjusting positions.

Main Loop:

```

for i=0; i < Iterations AND NumberOverlaps
  < MaxOverlaps; i++ {
  for j=0; j < M; j++ { for k=j+1; k < M; k++ {
    Overlaps=0;
    PushVector=(0, 0)
    if  $D(j, k) < \text{Size}$  { /* j and k overlap */
      Direction=normalized vector from k to j
      if Direction=(0, 0), set Direction to a random
        displacement
      Magnitude=1 -  $D(j,k)/\text{Size}$ 
      PushVector=PushVector + Magnitude *
        Direction
      Overlaps=Overlaps + 1
    }
  }
  if Overlaps > 0 {
    if PushVector=(0, 0), set PushVector to a random
      displacement
    PushVector=PushVector * Speed / Overlaps
    Offset for glyph j=PushVector
  }
}
Offset all glyphs
}

```

**Algorithm 1:** Distorting glyph positions in derived placement strategies to reduce overlaps.

Distortion techniques can be very effective at reducing the level of overlap while maintaining a reasonable degree of data integrity, a quality metric based on the difference between the perceived and actual data values and inter-relationships. The permissible variation in glyph position is often domain and context dependent. For example, the placement of a glyph representing a particular city should reside on the appropriate side of a state border. Indeed, the more constraints placed on the type and degree of distortion allowed, the more complex the algorithm becomes. The major trade-off thus becomes one of ease-of-implementation, such as random jitter and lens techniques, vs. accuracy of the results (e.g., through constraint optimization).

## Structure-driven glyph placement

Structure-driven algorithms assume there exists some implicit or explicit connectivity or relationship between data points, which may or may not be derivable from the data values. For example, there might be a temporal ordering or hierarchical relationship that might not be included as one of the data fields. We categorize strategies of this form as either ordered, hierarchical, or generalized network/graph structures. While ordered and hierarchical structures can be considered special cases of the generalized graph structure, it is possible to take advantage of the constraints imposed on certain types of structure to derive layout strategies not applicable to the more general structure. Algorithms for structure-driven placement are characterized as functions that map  $D \rightarrow S$ , where *S* is some generic structure space, followed by a mapping of  $S \rightarrow L$ . Each category of structure-driven technique can be further divided based on whether overlaps are permitted or not, and whether extra space is added to emphasize relationships. Overlap reduction and space padding will be discussed under distortion techniques for structure-driven methods.


**Ordered structure** Ordered structure may be linear (1-D) or grid-based (N-D). By sorting the data on one or more dimensions, and using this ordering to specify the glyph placement, we facilitate the detection of changes in the dimensions used in the sorting as well as provide clues into potential dependencies in the other variables (detecting if the patterns of change coincide with those of the sorted dimensions). Several placement algorithms have been developed in the past based on linear orderings, including raster scan,<sup>57</sup> circular,<sup>58,59</sup> and recursive space-filling patterns.<sup>60</sup> These are illustrated in Figure 6.

Figure 7 shows a time series of economic data (Dow Jones index, Standard and Poors 500 index, retail sales, and unemployment over a three year period) positioned in a circular pattern. Data for December radiate straight up. Expected patterns include high sales and low unemployment in January and the inverse in June. Abnormally high unemployment is visible through much of the first year shown, which improves as other market indicators rise in value.


Figure 8 displays 5-dimensional remote sensing data (Spot, magnetics, and three channels of radiometrics) using star glyphs and a 2-D raster ordering. Because of the number of glyphs involved, individual data points are difficult to interpret. However, clearly discernable regions of uniform texture are visible, including a zoomed in region from the lower right of the original data identifying an area with particularly high uranium content (data courtesy of CSIRO Division of Exploration and Mining).

The algorithm for a raster-based placement of a 1-D sorted glyph set is as follows: Let  $C = \sqrt{M+0.5}$  be the number of glyphs one can fit into a row of the display canvas. Thus, the maximum size of a glyph is given by the equation  $S_{\max} = W/C$ , where *W* is the display canvas

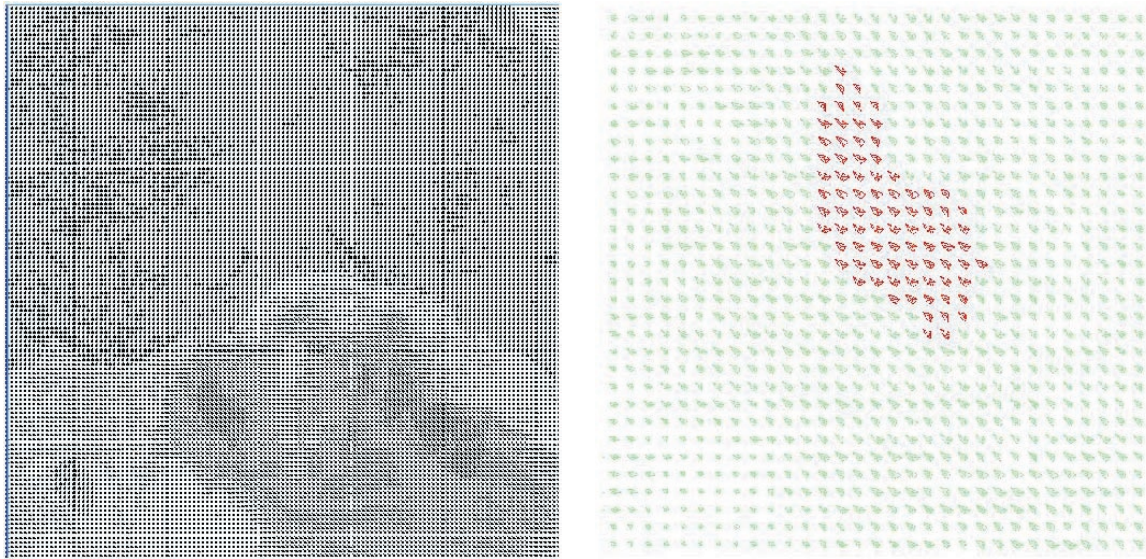





**Figure 6** Various placement patterns for linearly structured data (from<sup>58,60</sup>), including raster, radial, and recursive raster.



**Figure 7** Profile glyphs of monthly economic time series data using linear ordering and spiral layout (from SpiralGlyphics<sup>59,81</sup>). Dimensions are (from left to right) Dow Jones average, Standard and Poors 500 index, retail sales, and unemployment.



**Figure 8** Star glyphs of remote sensing data from the Grant's Pass region of Western Australia (5 channels) using 2-D raster ordering (from XmdvTool<sup>57</sup>). Second image shows a zoomed in region from the lower right of the original data highlighting an area with high Uranium values.



**Figure 9** Profile glyphs of average values from hierarchical partitioning of the Iris data using the DySect partitioning algorithm<sup>82</sup>.

width. Finally the location of the  $i^{th}$  glyph ( $l_i$ ) is computed as  $l_{ix} = S_{max} * (i \text{ modulo } C)$  and  $l_{iy} = S_{max} * (i/C)$ . Note that if  $\sqrt{M}$  has a fractional component, the final row will contain some empty space.

The strengths of order-based positioning are its ease of implementation, ease of interpretation, and effective utilization of screen space. It also provides a number of potentially revealing views of the data, as significant clusterings and discontinuities within a given dimension can become visually apparent. The main limitation of the strategy is that it assumes significant features in N-D will be perceptible in 1-D or 2-D snapshots, which often is not the case. It could be difficult, if not impossible, to correlate the changes seen in one ordering with those seen in another to create a composite description of an overall pattern.

### Hierarchical structure

Hierarchical structure in data sets can be explicit or implicit. Explicit hierarchies are those in which each level of the hierarchy is associated with a single data dimension, and the branches emanating from this level correspond to some number of distinct ranges for that data dimension. For example, sales data may have dimensions associated with particular time periods, geographical locations, sales personnel, and products. Different hierarchies are generated depending on the order in which the dimensions are processed. Other examples of explicit hierarchies are file systems and organizational charts. Implicit hierarchies are generated algorithmically using clustering or partitioning algorithms in conjunction with some N-dimensional distance or similarity metric.<sup>2</sup> These methods differ

from the derived data-driven techniques in that the goal is not to generate screen positions, but a hierarchical relationship structure that then can be conveyed to a placement algorithm customized to display this relational information.

Given a hierarchical structure, the task is to position glyphs on the display in such a way as to convey the relationships inherent in the structure. In most methods proposed to date, these relationships are reinforced through the use of additional graphics, using lines to connect parent and child nodes and/or to separate generations. White space can also be used to support this task. Many algorithms have been developed for drawing node-link graphs of trees; see<sup>61</sup> for an excellent survey. These algorithms vary according to features such as where the root node is relative to the rest of the tree (eg, centered, top-most), and the relative direction between a node and its children (e.g., radially outward, horizontal, vertical, or alternating horizontal and vertical). Different criteria for layout can result in configurations that reveal different aspects of the data. For example, Wilson and Bergeron<sup>62</sup> describe a system for visualizing hierarchies that implements four layout algorithms: leaf-based, subnode-based, range-based, and density-based. Tree layout algorithms also vary depending on whether they are 2-D or 3-D (as in Cone Trees,<sup>63</sup> Disc Trees,<sup>64</sup> and others<sup>65</sup>). Each of these methods can be augmented to display glyphs as opposed to simple circles, boxes, or spheres. Since the multivariate data may only represent the terminal nodes of the tree, one could either represent non-terminal nodes with aggregation-based glyphs, such as using the average of the values in the subgraph for each dimension, or simply use the original node representations above.

A simple algorithm for placement and drawing of hierarchically structured glyphs is given below. To create a more aesthetically pleasing graph, extra space on rows could be used to evenly separate all nodes at a particular level. Additionally, as will be described later, free space can be used to further separate non-sibling nodes at the same level. Finally, for unbalanced trees, space could be

used beneath branches that terminate prior to other branches to reduce edges crossing over glyphs. Figure 9 shows the result of this technique on a hierarchical representation of the iris data, where the non-terminal glyphs show the average values for the glyphs below them. Glyphs at a given level have been evenly spaced.


Let  $H$ =the height of the tree  
Let  $W$ =the maximum number of nodes in one level of the hierarchy.  
Let  $C$ =the canvas size (assumed a square)  
Compute  $S=C / \max(H, W)$  as the maximum size of a glyph

Assume glyph  $i$  is at position  $(l, j)$  of the hierarchy, where  $l$  is the level number and  $j$  is the offspring number ( $i$  is the  $j$ th child of its parent)

For each node in the hierarchy  
Set its  $y$  position (row) to be  $l * S$   
Compute its offset  $o$  in the row as the sum of the child counts for the nodes to left of its parent  
Set its  $x$  position (column) to be  $(j + o) * S$   
Draw the node  
If it is not the root node, draw a link to its parent

Algorithm 2: Hierarchical glyph placement.

An alternative to node-link based visualizations are the filled region methods, which have mostly been used for conveying univariate data. One of the earliest of these, the *Tree-Map*,<sup>66</sup> recursively divides the screen space, alternating horizontal and vertical, based on the number and populations of the subtrees, filling the resulting rectangular region with a color based on the dependent variable. A radial version, with aggregations also being depicted, was described in.<sup>67</sup> Wills<sup>68</sup> describes a hierarchical clustering mechanism that allows varying levels of detail to be conveyed with a tree-map structure, thus making it amenable to very large hierarchies. Ward and Keim<sup>69</sup> proposed a variant on tree-maps for positioning glyphs. This method, instead of being space-filling as in the original tree-map, allocates fixed size regions for each glyph,



**Figure 10** Star glyphs of Iris data set using recursive hierarchical placement based on the DySect<sup>8</sup> data partitioning algorithm and the Tree-map<sup>66</sup> recursive space partitioning technique<sup>69</sup>.



and thus may produce significant amounts of white space. A high-level description of the placement algorithm is outlined below.

If we assume a unit-sized glyph, the space required for a 2-level hierarchy consisting of a root node  $R_{0,0}$  with  $C_{0,0}$  children (all terminal nodes) is either  $(1 \times C_{0,0})$  or  $(C_{0,0} \times 1)$ . The two element subscript indicates the level of the hierarchy and the index of the node at the level specified. For a 3-level hierarchy, where the root has  $C_{0,0}$  non-terminal nodes  $(R_{1,0}, R_{1,1}, \dots, R_{1,C_{0,0}-1})$ , each of which has child counts given by  $(C_{0,0}, C_{1,1}, \dots, C_{1,C_{0,0}-1})$ , the space required will be  $(C_{0,0} \times \max(C_{1,i}))$ , where  $0 \leq i < C_{0,0}$ , or alternately the same dimensions rotated 90°C. Thus we can envision the screen first being divided in one direction according to the number of branches of the root, followed by the terminal nodes laid out in the opposite direction. The space needed for the terminal nodes, assuming we maintain a rectangular area, is the maximum of the number of children of the level one nodes.

Continuing this process, we discover the width and height needed for an arbitrary hierarchy can be computed by alternately multiplying the sum or the maximum of the children counts for nodes as we recurse down the tree. At any given level, either the height or the width will use the sum, while the other uses the maximum. To draw the glyphs then requires each node to be placed relative to its parent's base location (eg, the lower left corner), using offsets based on its earlier siblings' space requirements. To visually emphasize the family relationships, we draw rectangles to separate siblings, with the width of the rectangle sides decreasing as we get deeper into the tree. Finally, glyphs are drawn as the terminal nodes, with their bounding rectangles encompassing their immediate siblings. An example using the iris data set is shown in Figure 10. From the figure it is clear that the clustering algorithm performed relatively well, as most glyphs within a given block have clear shape and size similarities. There are, however, a few misplaced glyphs, such as the sibling pair in the third row from the bottom of the lower left region.

Another filled region method, called *dimensional stacking*,<sup>70</sup> recursively divides the display space based on binning two dimensions at a time, with inner subimages being divided by subsequent pairs of dimensions until all dimensions are used. The final rectangle is colored according to the dependent variable or an aggregation of all points mapping to the corresponding bin. Again, this technique, in conjunction with the constraint of using a uniform-sized region, could be employed to convey multivariate data with glyphs.

Placement strategies for hierarchical structures can be evaluated on a number of aspects, including the limitations on the width and depth of trees that can be effectively accommodated, the screen utilization, and the ease at which relationships between nodes can be perceived. This in turn depends on the degree of clutter caused by the display of links and the amount of

control the user has on the view. Most of the techniques work best for balanced trees; however, most real data does not behave in this fashion. Thus evaluation must also include tests with trees of varying degrees of complexity and balance. Also, since most hierarchy visualization techniques concentrate predominantly on the display of the relations, rather than the data at the node, tests must be performed to determine how easy it is to interpret the values depicted at the glyphs positioned at the nodes of the trees. Clearly, link placement becomes more problematic in the presence of glyphs on the display.

### Graph/Network structure

A generalization of hierarchical structure is that of a graph or network, which consists of a set of nodes (in our case, the data points) and a finite set of directed or undirected links/connections, each of which represents a relationship between a pair of nodes. As with hierarchical techniques, we can replace the simple node representation from an existing method with our glyph to convey the data. However, unlike hierarchical techniques it is not possible, in general, to avoid drawing some representation of the links, as implying connectivity via positioning can lead to errors in interpretation.

Network visualization has become a very popular research area in recent years. Reichenberger *et al*<sup>71</sup> developed a theory of diagram design and showed how different relationships can be mapped to different graph styles and the attributes of the nodes and links. Force-based techniques, such as described in,<sup>72</sup> are commonly used to automatically place nodes on the screen to convey relations and their strengths. However, these methods can lead to cluttered displays, depending on the data characteristics.

Determining positions for graph structured data is a complex process, as many, often conflicting, factors must be considered.<sup>44,64</sup> For example, rules of aesthetics, such as minimizing edge crossings or maintaining a relatively uniform distribution of nodes, are common.<sup>73</sup> Drawing conventions, such as using only straight lines or 90° bends in the links, also must be adhered to for consistency and ease of interpretation. Constraints may also be imposed on subsets of the graph, such as centering or clustering the elements of the subgraph or ensuring that a particular node is on the outer edge of the subgraph.

As in hierarchically-structured glyphs, perhaps the greatest concern when designing graph-structured placement algorithms is the scalability to large and complex networks. Since the links themselves may convey other information besides the connectivity of the nodes (e.g., traffic volume), it is important to allocate sufficient screen space and user control to these information-rich components of the display. Indeed, the links can be considered a secondary form of glyph, as their style, size, color, and shape can all be controlled by different aspects of the relationship. Graph-struc-

tured glyphs also have a disadvantage over hierarchies in that there is not a clear method to navigate through the structure. Thus, there must be support for elevating arbitrary nodes to the center of focus and examining neighborhoods of this focus, rather than simply providing drill-down and roll-up mechanisms as found in hierarchical structures.

### Distortion techniques for structure-driven layouts

Given an initial position assignment for structure-based techniques, we can choose to distort the results to attain other desirable properties. Distortion techniques have long been used in visualization to emphasize subsets of data while maintaining context and to deal with clutter resulting from large data sets. Examples include fish-eye and other lens techniques,<sup>56,74,75</sup> pliable surfaces,<sup>76</sup> and mapping to hyperbolic space.<sup>77</sup> Shape distortion has also been used to convey additional information, such as the area-based methods found in some cartographic algorithms. Position distortion has been employed in cartography to avoid overlaps and congestion, with the understanding that in many circumstances, the degree of accuracy needed varies considerably. As mentioned in the discussion of data-driven techniques, random jitter and algorithms to remove overlaps may also be applied. Luders and Ernst<sup>78</sup> used a combination of node scaling and layout reorganization to improve the readability of graphs. Beaudoin *et al*<sup>79</sup> use controlled overlapping to compress large hierarchies, with a variety of tools to support navigation and exploration of the stacked space.

An alternate goal of position distortion is to emphasize differences. Space padding can be used to separate glyphs that have data-specific differences which are not sufficiently reflected in the screen positions. For example, Figure 11 shows an ordered raster placement with spacing inserted between glyphs in a given row proportional to their n-

dimensional distance. The algorithm for this is given below.

*Let  $G$  be the largest gap the user wishes between glyphs, with gap units equal to the maximum glyph size*

*Compute  $D_i$  as the N-D distance from glyph  $i$  to glyph  $i+1$*   
*Compute  $D_m$  as the maximum N-D distance between adjacent glyphs*

*For each glyph  $i$ , compute the spacing to the next glyph as  $1 + G * D_i / D_m$*

*Compute the sum  $S$  of all spacings, which includes the space for drawing the glyphs*

*Compute  $R$ , the number of glyph spaces per row, which is  $\text{sqrt}(S) + 0.5$*

*Compute the maximum size for a glyph, which is  $W / R$  ( $W$  is canvas width)*

*position=0*

*For  $i=0$ ;  $i < M$ ;  $i++$  {*

*draw glyph  $i$  at row=(position/ $R$ ) and*


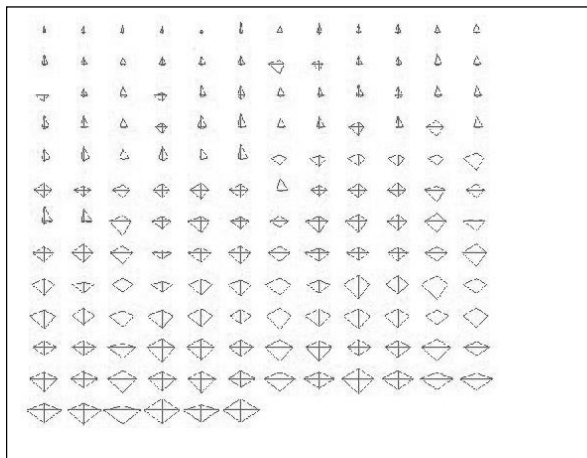
*column=(position modulo  $R$ )*

*add spacing  $i$  to position*

*}*


**Algorithm 3:** Inserting white space to emphasize differences in adjacent glyphs.

As another example, a tree drawing algorithm often must put nodes from different subtrees adjacent to each other. By adjusting positions such that inter-sibling distances are significantly smaller than between less-related nodes sharing a level of the tree, node relationships become clearer (see Figure 12). However, this reduces the number of glyphs one may be able to display



**Figure 11** Star glyphs of the iris data set, sorted by sepal length, with spacing inserted at positions of significant variation between adjacent glyphs.





**Figure 12** Profile glyphs of average values from hierarchical partitioning of the Iris data using the DySect partitioning algorithm<sup>82</sup>, with added space placed between adjacent non-siblings to emphasize sibling relationships.

without overlap. Other ways to emphasize these differences would be via small vertical offsets (eg, alternating between two y-values at each non-sibling transition along a row of the hierarchy) or coloring the nodes, links, or background between non-sibling adjacencies.

The concept of an acceptable level of overlapping is a crucial one; in many cases, minor, or even major, overlaps don't detract from the message being conveyed in the visualization. The ideal solution might be to allow user control of the overlapping, with smooth animation of positions between overlapped and non-overlapped states to allow easy tracking of relationships between the glyphs. This has been demonstrated in several 3-D tree visualization systems, such as reconfigurable disk trees.<sup>64</sup> Distortion techniques that permit compression and expansion of arbitrary sections of the structure, while maintaining context via animated transitions, can allow users to build a mental model of the regions of interest within the data without being overly constrained by the display resolution. While ideally the system itself would analyze the congestion on the screen and attempt to reduce overlap via distortion, to enable the system to understand the user's task in sufficient detail to automatically perform this process reliably has not yet been achieved.

### Summary of placement strategies

In this section we attempt to get a feel for the number of possible layout algorithms that can be applied to a given data set, and then give some guidelines on how one might go about selecting one or more algorithms to test. It should be noted that, to date, no formal, extensive evaluations and/or user studies have been performed on this problem, although it is clearly a desirable future goal. Some studies focusing on subsets of the techniques have been reported (eg,<sup>65,75</sup>), and evaluation metrics for information visualization have been proposed that involve issues such as density, dimensionality, occlusion, and number of identifiable points.<sup>79</sup> However, we have not found a comprehensive study of glyph layout strategies in the literature.

### The size of the design space

One might wonder what the total number of distinct ways a set of  $N$ -D glyphs can be positioned on a 2-D data space. The rather simplistic computations below cover a significant subset of the methods described in this paper. Clearly adding the third display dimension, or even the time dimension, would increase this number considerably, and visually exploring the entire space of alternatives would be a daunting task.

**data-driven (raw):**  $(1+D)(N^2 - N)$ , where  $D$  is the number of distortion techniques (e.g., random jitter, relaxation)

**data-driven (derived):**  $(1+D) \sum_{j=1}^k p_j$ , where  $k$  is the number of distinct derivation algorithms and  $p_j$  is the number of variations within algorithm  $j$  (e.g., different distance metrics).

**structure-driven (linear ordered):**  $N \times FP(1 + SP + OL)$ , where  $FP$  is the number of filling patterns,  $SP$  is the number of methods which introduce white space for separation (space-padding), and  $OL$  is the number of methods which distort the placement to allow overlaps.

### Selecting a placement algorithm

In selecting an appropriate method, one must consider several factors, including the characteristics of the data (size, distribution), the purpose of the visualization (presentation, confirmation, exploration) and the specific task(s) at hand (e.g., detection, classification, or measurement of patterns or outliers). An additional critical component in the selection process is the skills of the prospective user of the visualization. Some visualization methods require significant training to interpret correctly.

Domain knowledge also plays a significant role in selecting meaningful strategies. In data driven approaches, there are usually mappings of data dimensions to position that make more intuitive sense than others. Similarly, the dimensions used to control structure-driven techniques, whether it be the key for sorting

or the order of dimensions used for layering a hierarchy, are usually best left to selection by the user.

Placement techniques often force the user to trade off between efficient screen utilization, the degree of occlusion, and distortion of the values being used to position the glyphs. Ideally, the user should be able to dynamically adjust the mapping to strengthen or weaken the relative importance of each of these factors. Thus rather than selecting a single visualization technique, he or she would have a continuum of interrelated methods. This is a relatively unexplored area in the development of glyph placement strategies, though static constraint satisfaction algorithms have been applied to graph drawing tasks in the past.

Another question the user must ask is whether to impose a structure on the data in the event that no explicit structure already exists. While the structure provides relational information that might not be readily perceivable in data-driven approaches, the results are often not unique and may convey misleading or erroneous information. For example, clustering algorithms require the specification of a distance or similarity metric to process the data, and it is often difficult to capture the relative interdependencies between data dimensions or their contextual importance within this metric. Again, user control of the process is vital, and this in turn requires users to be knowledgeable not only of the semantics of their data but also of the structuring algorithm being utilized.

As an initial, though not exhaustive or definitive, set of guidelines for selecting a layout strategy, we look at two of the important design criteria: data characteristics and user task. The table in Figure 13 displays several points within this design space, and a recommendation for a starting layout. Each can be augmented with distortion depending on the degree of overlap present and amount of empty space on the screen, as well as the desired accuracy of placement for data-driven methods. Note that these recommendations are not infallible, but reasonable starting points. Other methods, such as imposing structure on unstructured data or ignoring the structure in a structured set, may also prove insightful. The main point is that no single placement strategy is good for all tasks

and data characteristics, which implies that for other than very limited circumstances, support should be provided for a variety of layout methods.

### Summary and future work

The visualization of multivariate data is a task of growing importance in a wide range of disciplines, and numerous techniques have been proposed and employed. In this paper we have focused on  $N$ -dimensional glyphs as a visualization component, and have presented a taxonomy of techniques for positioning the glyphs on the screen to better convey relationships between data points. Methods were divided between data-driven and structure-driven approaches, and a variety of post-processing distortions were described. Several of the techniques presented are novel in the context of glyph-based visualization. These include:

- force-based distortion to reduce overlap, which is similar to techniques employed in digital cartography for placing symbols on maps;
- insertion of white space in linear and raster layouts to emphasize transitions between adjacent glyphs; and
- space-filling layout of hierarchically structured data, which was first proposed by the author in<sup>69</sup> but not fully implemented.

In addition, we proposed a number of considerations and criteria for designing glyph layout techniques and for evaluating and selecting strategies based on the characteristics of the data and the user's task. While not an exhaustive list of criteria and evaluation/selection parameters, we feel they are reasonable starting points for continued development.

Our future work will involve consolidating the techniques described in this paper into a single application to facilitate a comprehensive analysis of the applicability and limitations of each placement algorithm. A wide assortment of customizable glyphs will be supported, similar to the Glyphmaker program.<sup>22</sup> Usability and perceptual testing will also be conducted. Finally, we are interested in exploring the use of animation to help overcome some of the deficiencies found in the static

Data Characteristics	User Task	Recommended Strategy
Small to moderate size	univariate analysis	sorted structure-driven
Small to moderate size	bivariate analysis	raw data-driven
Small to large size	outlier detection	raw or derived data-driven
Moderate to large size	cluster analysis	derived data-driven, e.g., MDS
Small to moderate size	cluster analysis	hierarchical structure-driven after imposing a hierarchy
Small to moderate size, relational	link analysis	network structure-driven

**Figure 13** Sample points in a design space of glyph layout strategies.

presentation of multivariate data with glyphs. For example, smoothly animating the transformation of one placement strategy to another might convey relational information not readily discernible in either static display.

## References

- Shneiderman B. *The eyes have it: a task by data type taxonomy for information visualization*. Proceedings 1996 IEEE Symposium on Visual Languages (1996); 336–343.
- du Toit S, Steyn A, Stumpf R. *Graphical Exploratory Data Analysis*. Springer-Verlag: Berlin, 1986.
- Siegel J, Farrell E, Goldwyn R, Friedman H. The surgical implication of physiologic patterns in myocardial infarction shock. *Surgery* 1972; **72**: 126–141.
- Anderson E. A semigraphical method for the analysis of complex problems. *Proceedings of the National Academy of Science* 1957; **13**: 923–927.
- Gnanadesikan R. *Methods of Statistical Data Analysis of Multivariate Observations*. John Wiley and Sons: New York, 1997; 384 pp.
- Pickett R, Grinstein R. *Iconographic displays for visualizing multidimensional data*. Proceedings IEEE Conference on Systems, Man and Cybernetics 1988; 164–170.
- Kleiner B, Hartigan J. Representing points in many dimension by trees and castles. *Journal of the American Statistical Association* 1981; **76**: 260–269.
- Beddow J. *Shape coding of multidimensional data on a micro-computer display*. Proceedings Visualization 90 (San Francisco, CA). Los Alamitos, CA: IEEE Computer Society Press, 1990; 238–246.
- Hartigan J. Printer graphics for clustering. *Journal of Statistical Computing and Simulation* 1975; **4**: 187–213.
- Klassen R, Harrington S. *Shadowed hedgehogs: a technique for visualizing 2D slices of 3D vector fields*. Proceedings Visualization 91 (San Diego, CA). Los Alamitos, CA: IEEE Computer Society Press, 1991; 148–153.
- Chernoff H. The use of faces to represent points in k-dimensional space graphically. *Journal of the American Statistical Association* 1973; **68**: 361–368.
- Wittenbrink C, Pang A, Lodha S. Glyphs for visualizing uncertainty in vector fields. *IEEE Transactions on Visualization and Computer Graphics* 1996; **2**: 266–279.
- Schroeder W, Volpe C, Lorensen W. *The Stream Polygon: a technique for 3D vector field visualization*. Proceedings Visualization 91 (San Diego, CA). Los Alamitos, CA: IEEE Computer Society Press, 1991; 126–132.
- Fuhrmann A, Groller E. *Real-time techniques for 3D flow visualization*. Proceedings Visualization 98 (Research Triangle Park, NC). Los Alamitos, CA: IEEE Computer Society Press, 1998; 305–312.
- Friedman J, Farrell E, Goldwyn R, Miller M, Sigel J. *A graphic way of describing changing multivariate patterns*. Proceedings Sixth Interface Symposium on Computer Science and Statistics 1972; 56–59.
- Mezzich J, Worthington D. A comparison of graphical representations of multidimensional psychiatric diagnostic data. In: Wang P (Ed.) *Graphical Representation of Multivariate Data*. Academic Press: New York, 1978; pp 123–141.
- Levkowitz H. *Color icons: merging color and texture perception for integrated visualization of multiple parameters*. Proceedings Visualization 91 (San Diego, CA). Los Alamitos, CA: IEEE Computer Society Press, 1991; 164–170.
- Chuah M, Eick S. Information rich glyphs for software management data. *IEEE Computer Graphics and Applications* 1998; **18**: 24–29.
- Kerlick G. *Moving iconic objects in scientific visualization*. Proceedings Visualization 90 (San Francisco, CA). Los Alamitos, CA: IEEE Computer Society Press, 1990; 124–130.
- Rohrer R, Ebert D, Sibert J. *The shape of Shakespeare: visualizing text using implicit surfaces*. Proceedings Visualization 98 (Research Triangle Park, NC). Los Alamitos, CA: IEEE Computer Society Press, 1998; 121–129.
- Ebert D, Rohrer R, Shaw C, Panda P, Kukla J, Roberts D. *Procedural shape generation for multi-dimensional data visualization*. Proceedings Data Visualization 99. Springer-Verlag: Berlin, 1999; 3–12.
- Ribarsky W, Ayers E, Eble J, Mukherjee S. Glyphmaker: creating customized visualizations of complex data. *Computer* 1994; **27**: 57–64.
- Post F, Walsum T, Post F, Silver D. *Iconic techniques for feature visualization*. Proceedings Visualization 95 (Atlanta, GA). Los Alamitos, CA: IEEE Computer Society Press, 1995; 288–295.
- Ankerst M, Berchtold S, Keim D. *Similarity clustering of dimensions for an enhanced visualization of multidimensional data*. Proceedings Information Visualization 98 (Research Triangle Park, NC). Los Alamitos, CA: IEEE Computer Society Press, 1998; 52–60.
- Cleveland W, McGill R. Graphical perception: theory, experimentation and application to the development of graphical methods. *Journal of the American Statistical Association* 1984; **79**: 531–554.
- Levkowitz H. *Color Theory and Modeling for Computer Graphics, Visualization, and Multimedia Applications*. Kluwer: Boston, 1997.
- Gershon N. From perception to visualization. In: Rosenblum L, et al. (Eds) *Scientific Visualization: Advances and Challenges*. Academic Press: San Diego, CA, 1994; 000 pp.
- Ahlberg C, Wistrand E. *IVEE: an information visualization and exploration environment*. Proceedings Information Visualization 95 (Atlanta, GA). Los Alamitos, CA: IEEE Computer Society Press, 1995; 66–73.
- Becker R, Eick S, Wilks A. Visualizing network data. *IEEE Transactions on Visualization and Computer Graphics* 1995; **1**: 16–28.
- Chuah M, Roth S, Mattis J, Kolojechick J. *SDM: malleable information graphics*. Proceedings Information Visualization 95 (Atlanta, GA). Los Alamitos, CA: IEEE Computer Society Press, 1995; 36–42.
- Rossi A, Varga M. *Visualization of massive retrieved newsfeeds in interactive 3D*. Proceedings International Conference on Information Visualization 99 (London, UK). Los Alamitos, CA: IEEE Computer Society Press, 1999; 12–17.
- Friedman J, Tukey J. A projection pursuit algorithm for exploratory data analysis. *IEEE Transactions on Computers* 1974; **C-23**: 881–889.
- Huber P. Projection pursuit. *Annals of Statistics* 1985; **13**: 435–475.
- Jolliffe J. *Principal Component Analysis*. Springer Series in Statistics. Springer: Berlin, 1986.
- Kruskal J, Wish M. *Multidimensional Scaling*. Sage Publications: London, 1978.
- Bentley C, Ward M. *Animating multidimensional scaling to visualize n-dimensional data sets*. Proceedings Information Visualization 96 (San Francisco, CA). Los Alamitos, CA: IEEE Computer Society Press, 1996; 72–73.
- Brodbeck D, Chalmers M, Lunzer A, Cotture P. *Domesticating Bead: adapting an information visualization system to a financial institution*. Proceedings Information Visualization 97 (Phoenix, AZ). Los Alamitos, CA: IEEE Computer Society Press, 1997; 73–80.
- Wong P, Bergeron R. *Multivariate visualization using metric scaling*. Proceedings Visualization 97 (Phoenix, AZ). Los Alamitos, CA: IEEE Computer Society Press, 1997; 111–118.
- Kohonen T. *Self-Organizing Maps* Springer Series on Information Sciences, Vol. 30. Springer: Berlin, 2001; 501 pp.
- Rushmeier H, Lawrence R, Almasi G. *Visualizing customer segmentation produced by self organizing maps*. Proceedings Visualization 97 (Phoenix, AZ). Los Alamitos, CA: IEEE Computer Society Press, 1997; 463–466.

- 41 Assa J, Kohen-Or D, Milo T. *Displaying data in multidimensional relevance space with 2D visualization maps*. Proceedings Visualization 97 (Phoenix, AZ). Los Alamitos, CA: IEEE Computer Society Press, 1997; 127–134.
- 42 Gross M, Sprenger T, Finger J. *Visualizing information on a sphere*. Proceedings Information Visualization 97 (Phoenix, AZ). Los Alamitos, CA: IEEE Computer Society Press, 1997; 11–16.
- 43 Eick S, Wills G. *Navigating large networks with hierarchies*. Proceedings Visualization 93 (San Jose, CA). Los Alamitos, CA: IEEE Computer Society Press, 1993; 204–210.
- 44 Herman I, Melancon G, Marshall M. Graph visualization and navigation in information visualization. *IEEE Transactions on Visualization and Computer Graphics* 2000; 6: 24–43.
- 45 Hoffman P, Grinstein G, Marx K, Grosse I, Stanley E. *DNA visual and analytic data mining*. Proceedings Visualization 97 (Phoenix, AZ). Los Alamitos, CA: IEEE Computer Society Press, 1997; 437–441.
- 46 Mukherjee S, Hara Y. *Visualizing world wide web search engine results*. Proceedings International Conference on Information Visualization 99 (London, UK). Los Alamitos, CA: IEEE Computer Society Press, 1999; 400–405.
- 47 Spoerri A. *InfoCrystal: a visual tool for information retrieval*. Proceedings Visualization 93 (San Jose, CA). Los Alamitos, CA: IEEE Computer Society Press, 1993; 150–157.
- 48 Rose S. *The Sunflower visual metaphor, a new paradigm for dimensional compression*. Proceedings Information Visualization 99 (San Francisco, CA). Los Alamitos, CA: IEEE Computer Society Press, 1999; 128–131.
- 49 Chalmers M. *A linear iteration time layout algorithm for visualizing high-dimensional data*. Proceedings Visualization 96 (San Francisco, CA). Los Alamitos, CA: IEEE Computer Society Press, 1996; 127–132.
- 50 Globus A, Levit C, Lasinski T. *A tool for visualizing the topology of three-dimensional vector fields*. Proceedings Visualization 91 (San Diego, CA). Los Alamitos, CA: IEEE Computer Society Press, 1991; 33–40.
- 51 Helman J, Hesselink L. Visualizing vector field topology in fluid flows. *Computer Graphics and Applications* 1991; 11: 36–46.
- 52 Dovey D. *Vector plots for irregular grids*. Proceedings Visualization 95 (Atlanta, GA). Los Alamitos, CA: IEEE Computer Society Press, 1995; 248–253.
- 53 Cleveland W. *Visualizing Data*. Hobart Press: Summit, NJ, 1993.
- 54 Keim D, Hermann A. *The Gridfit algorithm: an efficient and effective approach to visualizing large amounts of spatial data*. Proceedings Visualization 98 (Research Triangle Park, NC). Los Alamitos, CA: IEEE Computer Society Press, 1998; 181–188.
- 55 Woodruff A, Landay J, Stonebraker M. *Constant density visualization of non-uniform distributions of data*. Proceedings ACM Symposium on User Interface Software and Technology 1998; 19–28.
- 56 Leung Y, Apperley M. A review and taxonomy of distortion-oriented presentation techniques. *ACM Transactions on Computer-Human Interaction* 1994; 1: 126–160.
- 57 Ward M. *XmdvTool: integrating multiple methods for visualizing multivariate data*. Proceedings Visualization 94 (Washington, DC). Los Alamitos, CA: IEEE Computer Society Press, 1994; 326–333.
- 58 Keim D, Kriegel H. VisDB: database exploration using multidimensional visualization. *Computer Graphics and Applications* 1994; 40–49.
- 59 Lipchak B, Ward M. *Visualization of cyclic multivariate data*. Proceedings Visualization 97 Late Breaking Hot Topics (Phoenix, AZ). Los Alamitos, CA: IEEE Computer Society Press, 1997; 61–64.
- 60 Keim D, Kriegel H, Ankerst M. *Recursive pattern: a technique for visualizing very large amounts of data*. Proceedings Visualization 95 (Atlanta, GA). Los Alamitos, CA: IEEE Computer Society Press, 1995; 279–286.
- 61 Di Battista G, Eades P, Tamassia R, Tollis I. *Graph Drawing: Algorithms for the Visualization of Graphs*. Prentice Hall: Upper Saddle River, NJ, 1999; 397 pp.
- 62 Wilson R, Bergeron R. *Dynamic hierarchy specification and visualization*. Proceedings Information Visualization 99 (San Francisco, CA). Los Alamitos, CA: IEEE Computer Society Press, 1999; 65–72.
- 63 Robertson G, Mackinlay J, Card S. *Cone trees: animated 3D visualizations of hierarchical information*. Proceedings ACM SIGCHI Conference on Human Factors in Computing Systems, 1991; 331–337.
- 64 Jeong C, Pang A. *Reconfigurable disc trees for visualizing large hierarchical information space*. Proceedings Information Visualization 98 (Research Triangle Park, NC). Los Alamitos, CA: IEEE Computer Society Press, 1998; 19–25.
- 65 Carriere J, Kazman R. *Interacting with huge hierarchies: beyond cone trees*. Proceedings Information Visualization 95 (Atlanta, GA). Los Alamitos, CA: IEEE Computer Society Press, 1995; 74–81.
- 66 Johnson B, Shneiderman B. *Tree-maps: a space-filling approach to the visualization of hierarchical information structures*. Proceedings Visualization 91 (San Diego, CA). Los Alamitos, CA: IEEE Computer Society Press, 1991; 189–194.
- 67 Chuah M. *Dynamic aggregation with circular visual designs*. Proceedings Information Visualization 98 (Research Triangle Park, NC). Los Alamitos, CA: IEEE Computer Society Press, 1998; 35–43.
- 68 Wills G. *An interactive view for hierarchical clustering*. Proceedings Information Visualization 98 (Research Triangle Park, NC). Los Alamitos, CA: IEEE Computer Society Press, 1998; 26–31.
- 69 Ward M, Keim D. *Screen layout methods for multidimensional visualization*. Proceedings CODATA Euro-American Workshop on Visualization of Information and Data. (Paris, France); 2 pp.
- 70 LeBlanc J, Ward M, Wittels N. *Exploring n-dimensional databases*. Proceedings Visualization 90 (San Francisco, CA). Los Alamitos, CA: IEEE Computer Society Press, 1990; 230–237.
- 71 Reichenberger K, Kamps T, Golovchinsky G. *Towards a generative theory of diagram design*. Proceedings Information Visualization 95 (Atlanta, GA). Los Alamitos, CA: IEEE Computer Society Press, 1995; 11–18.
- 72 Hendley R, Drew N, Wood A, Beale R. *Narcissus: visualizing information*. Proceedings Information Visualization 95 (Atlanta, GA). Los Alamitos, CA: IEEE Computer Society Press, 1995; 90–96.
- 73 Davidson R, Harel D. Drawing graphs nicely using simulated annealing. *ACM Transactions on Graphics* 1996; 15: 301–331.
- 74 Schaffer D, Zuo Z, Greenberg S, Bartram L, Dill J, Dubs S, Roseman M. Navigating hierarchically clustered networks through fisheye and full-zoom methods. *ACM Transactions on Computer-Human Interactions* 1996; 3: 162–188.
- 75 Storey M, Wong L, Fracchia F, Muller H. *On integrating visualization techniques for effective software exploration*. Proceedings Information Visualization 97 (Phoenix, AZ). Los Alamitos, CA: IEEE Computer Society Press, 1997; 38–45.
- 76 Carpendale T, Cowperthwaite D, Fracchia F. *3-dimensional pliable surfaces for the effective presentation of visual information*. Proceedings ACM Symposium on User interface Software and Technology, 1995; 217–226.
- 77 Munzner T. *H3: laying out large directed graphs in 3D hyperbolic space*. Proceedings Information Visualization 97 (Phoenix, AZ). Los Alamitos, CA: IEEE Computer Society Press, 1997; 2–10.
- 78 Luders P, Ernst R. *Improving browsing in information by the automatic display layout*. Proceedings Information Visualization 95 (Atlanta, GA). Los Alamitos, CA: IEEE Computer Society Press, 1995; 26–33.
- 79 Beaudoin L, Parent M-A, Vroomen L. *Cheops: a compact explorer for complex hierarchies*. Proceedings Visualization 96 (San Francisco, CA). Los Alamitos, CA: IEEE Computer Society Press, 1996; 87–92.
- 80 Brach R. *Metrics for effective information visualization*. Proceedings Information Visualization 97 (Phoenix, AZ). Los Alamitos, CA: IEEE Computer Society Press, 1997; 108–111.
- 81 Ward M, Lipchak B. A visualization tool for exploratory analysis of cyclic multivariate data. *Metrika* 2000; 51: 27–37.
- 82 Andrae P, Dawkins B, O'Connor P. DySect: an incremental clustering algorithm. [Document included with public-domain software] retrieved from StatLib at CMU (<http://lib.stat.cmu.edu/general/>) (accessed 21 September 2002).