

Data Visualization through Graph Drawing ¹

George Michailidis ² and Jan de Leeuw ³

² Department of Statistics, The University of Michigan, Ann Arbor, MI 48109-1285

³ Department of Statistics, University of California, Los Angeles, Los Angeles, CA 90095

Summary

In this paper the problem of visualizing categorical multivariate data sets is considered. By representing the data as the adjacency matrix of an appropriately defined bipartite graph, the problem is transformed to one of graph drawing. A general graph drawing framework is introduced, the corresponding mathematical problem defined and an algorithmic approach for solving the necessary optimization problem discussed. The new approach is illustrated through several examples.

Keywords: Multivariate categorical data, data visualization, graphs, majorization algorithms

1. Introduction

Advances in data collection, computerization of transactions and breakthroughs in storage technology have allowed research and business organizations to collect increasingly large amounts of data. In order to extract useful information from such large data sets, a first step is to be able to *visualize* their global structure and identify patterns, trends and outliers.

¹The work of the first author was supported in part by the National Science Foundation under grant IIS-9988095. The authors would like to thank three anonymous referees for their helpful comments and suggestions.

Graphs are useful entities since they can represent relationships between sets of objects. They are used to model complex systems (e.g. computer and transportation networks, VLSI and Web site layouts, molecules, etc) and to visualize relationships (e.g. social networks, entity-relationship diagrams in database systems, etc). Graphs are also very interesting mathematical objects and a lot of attention has been paid to their properties. In many instances the right picture is the key to understanding. The various ways of visualizing a graph provide different insights, and often hidden relationships and interesting patterns are revealed. An increasing body of literature is considering the problem of how to draw a graph (see for instance the book by [10] on Graph Drawing, the proceedings of the annual conference on Graph Drawing, etc). Also, several problems in distance geometry and in graph theory have their origin in the problem of graph drawing in higher dimensional spaces. Of particular interest in this paper are the representation of data sets through graphs. This bridges the fields of multivariate statistics and graph drawing.

Many of the algorithms that have appeared in the graph drawing literature serve a different purpose; namely, to satisfy some *aesthetic* rules imposed on the final layout such as symmetry, uniformity of edge lengths and distribution of vertices, minimization of edge crossings, etc [5, 10, 16]. However, our main objective is to use graph drawings to visualize data; thus, we are interested in drawings that show important and invariant aspects of the data and whose structure and properties we can examine analytically. In this paper, we provide a rigorous mathematical framework for drawing graphs utilizing the information contained in the adjacency matrix of the underlying graph. At the core of our approach are various loss functions that measure the lack of fit of the resulting representation, that need to be optimized subject to a set of constraints that correspond to different drawing representations. We then establish how the graph drawing problem encompasses problems in multivariate statistics. We develop a set of algorithms based on the theory of majorization to optimize the various loss functions and study their properties (existence of solution, convergence, etc). We demonstrate the usefulness of our approach through a series of examples.

2. Data as Graphs

2.1. Graphs and the Adjacency Matrix. In this paper we consider an undirected graph $G = (V, E)$, where $V = \{v_1, v_2, \dots, v_n\}$ is the set of the n vertices and $E \subset V \times V$ the set of edges. It is assumed that the graph G does not contain either self-loops or multiple edges between any pair of vertices. The set of edges can be represented in matrix form through the *adjacency matrix* $A = \{a_{ij}\}$, $i, j = 1, \dots, n$. Thus, vertices $i, j \in G$ are connected if and only if $a_{ij} > 0$, otherwise $a_{ij} = 0$. If $a_{ij} \in \{0, 1\}$ we are dealing with a *simple* graph, otherwise with a *weighted* graph.

In the following example the graph representation of a familiar data structure from multivariate statistics is given. Consider the following dissimilarity matrix on six objects

$$\begin{bmatrix} 0 & & & & & \\ 7 & 0 & & & & \\ 2 & 3 & 0 & & & \\ 8 & 6 & 1 & 0 & & \\ 5 & 9 & 6 & 3 & 0 & \\ 3 & 1 & 2 & 5 & 12 & 0 \end{bmatrix}$$

It can be represented by a complete weighted graph on 6 vertices as shown in Figure 1.




FIGURE 1. The graph representation of a dissimilarity matrix. The numbered squares correspond to the objects, while the weights on certain edges correspond to the dissimilarities

2.2. Bipartite Graphs. In two recent papers of de Leeuw and Michailidis [9, 18] the problem of representing categorical datasets through *bipartite graphs* is considered. For bipartite graphs, the vertex set V is partitioned into two sets V_1 and V_2 , and the edge set E is defined on $V_1 \times V_2$ and indicates which vertices from V_1 are connected to vertices in V_2 and vice versa. In multidimensional data analysis, the classical data structure where data on J categorical variables (with k_j possible values (categories) per variable) are collected for N objects, can be represented by a bipartite graph. The N objects correspond to the vertices of V_1 , the $K = \sum_j k_j$ categories to the vertices of V_2 and there are $N \times J$ edges in E , since each object is connected to J different categories.

The adjacency matrix of a bipartite graph takes the form

$$A = \begin{bmatrix} 0 & W \\ W' & 0 \end{bmatrix}$$

where W is the familiar superindicator matrix [12] from multiple correspondence analysis (MCA). The superindicator matrix of a small example (discussed in [22] and also [9]) and the corresponding graph representation are given in Table 1 and Figure 2 respectively.

Another data structure that can be represented by a bipartite graph is the contingency table, familiar from elementary statistics [1, 12], where the I categories of the first variable correspond to the vertices in V_1 and the J categories of the second variable to those of V_2 . For this data structure the a_{ij} 's are nonnegative numbers that indicate how many observations fall

		cheap	not expensive	expensive	down fibers	synthetic fibers	good	acceptable	bad
	Sleeping Bag	Price			Fiber		Quality		
1	One Kilo Bag	1	0	0	0	1	1	0	0
2	Sund	1	0	0	0	1	0	0	1
3	Kompakt Basic	1	0	0	0	1	1	0	0
4	Finmark Tour	1	0	0	0	1	0	0	1
5	Interlight Lyx	1	0	0	0	1	0	0	1
6	Kompakt	0	1	0	0	1	0	1	0
7	Touch the Cloud	0	1	0	0	1	0	1	0
8	Cat's Meow	0	1	0	0	1	1	0	0
9	Igloo Super	0	1	0	0	1	0	0	1
10	Donna	0	1	0	0	1	0	1	0
11	Tyin	0	1	0	0	1	0	1	0
11	Travellers Dream	0	1	0	1	0	1	0	0
13	Yeti Light	0	1	0	1	0	1	0	0
14	Climber	0	1	0	1	0	0	1	0
15	Viking	0	1	0	1	0	1	0	0
16	Eiger	0	0	1	1	0	0	1	0
17	Climber light	0	1	0	1	0	1	0	0
18	Cobra	0	0	1	1	0	1	0	0
19	Cobra Comfort	0	1	0	1	0	0	1	0
20	Foxfire	0	0	1	1	0	1	0	0
21	Mont Blanc	0	0	1	1	0	1	0	0

TABLE 1. The superindicator matrix W of the sleeping bags data set

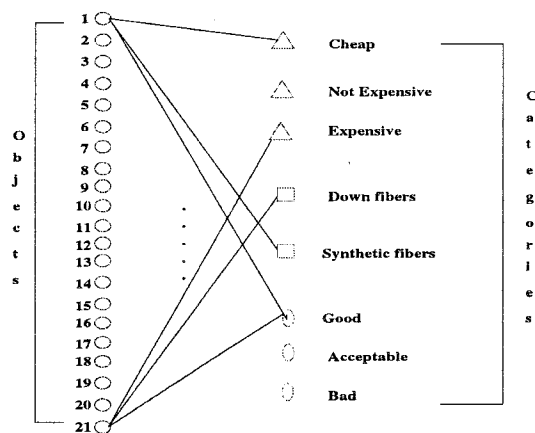


FIGURE 2. The bipartite graph of the sleeping bags example. It can be seen that the first object is cheap, made of synthetic fibres and of good quality, while the last object is expensive, made of down fibres and of good quality

in cell (i, j) in the contingency table; thus, we are dealing with a weighted bipartite graph in this case. Finally, data sets involving measurements on objects organized in a hierarchical structure (e.g. students grouped by class

or school, consumers grouped by geographical areas, etc) can be represented by direct sums of bipartite graphs [19].

3. Graph Drawing

The adjacency matrix represents a useful way to code the information in a data set and moreover contains exactly the same information as the original data set; however, it is hard to use it to uncover patterns and trends in the data. One way of utilizing the information contained in the adjacency matrix is to draw the graph by connecting the appropriate vertices. This goes in the direction of making a picture of the data, and when things work out well, a picture is worth a lot of numbers, especially when these numbers are just zeros and ones. But the "technique" of drawing the coded graph, say in the plane, has a large amount of arbitrariness. Since the graph only contains the qualitative information of which vertices are connected, we can locate them anywhere in the plane and then draw the edges corresponding with the nonzero elements of the adjacency matrix. The trick is to manage to draw the graph in such a way so as the resulting picture becomes as informative as possible.

The general problem of graph drawing discussed in this paper is to represent the edges of a graph as points in \mathbb{R}^s and the vertices as lines connecting the points. Graph drawing is an active area in computer science, and it is very ably reviewed in the recent book by [10]. The choice of \mathbb{R}^s is due to its attractive underlying geometry and the fact that it renders the necessary computations more manageable. In practice, s is usually chosen equal to 2 or 3, in order to be able to visualize the resulting representation.

There are basically two different approaches to make such drawings. In the *metric* or *embedding* approach we use the path-length distance defined between the vertices of the graph and we try to approximate these distance by the Euclidean distance between the points [4, 10]. A data analytic technique in this spirit is Multidimensional Scaling [2, 7].

In this paper, we adopt primarily the *adjacency model*, i.e. we do not emphasize graph-theoretical distances, but we pay special attention to which vertices are adjacent and which vertices are not. Obviously, this is related to distance, but the emphasis is different. We use objective (loss) functions to measure the *quality* of the resulting embedding.

3.1. Force-Directed Techniques. The class of graph drawing techniques we are most interested in here are the *force-directed techniques*. The vertices are bodies that attract and repel each other, for instance because the edges are springs or because the vertices have electric charges. This means that there are forces pulling and pushing the vertices apart, and the optimal graph drawing will be the one in which these forces are in equilibrium. In [10] (Chapter 10) force-directed graph drawing means minimizing a loss function which incorporates both pushing and pulling. A popular choice for pulling is Hooke's law [11], i.e. the force is proportional to the difference between

the distance of the vertices and the zero-energy length of the spring, while the repelling force follows an inverse square law. A Bayesian framework for dynamic graph drawing by using force-directed methods is provided in [3]. Force-directed algorithms are also employed in NicheWorks [27], a visualization tool for the investigation of very large graphs.

In this paper we concentrate on *pulling under constraints*, which means that we do not have explicit pushing components in the loss function. The constraints normalize the drawing in order to prevent trivial solutions in which the whole graph is pulled (collapses) into a single point.

Let us first define the following loss function, that represents in our study the main tool for making the necessary graph drawings,

$$(1) \quad \sigma_\phi(Z|A) = \sum_{i=1}^n \sum_{j=1}^n a_{ij} \phi(d_{ij}(Z)),$$

where $d_{ij}(Z)$ denotes the distance of points with coordinates z_i and z_j in \mathbb{R}^s . We assume that the weights $a_{ij} \in \{0, 1\}$ and that ϕ is an increasing function. Therefore, minimizing σ_ϕ means minimizing the weighted sum of the transformed distances between the points that are connected in the graph. Notice that we do not assume that the distances are Euclidean, they could be ℓ_1 (City Block) or ℓ_∞ (Chebyshev) or general ℓ_p distances. Moreover, the presence of the ϕ function allows us to mitigate the effect of the distances on the graphical representation we are seeking. For example, a convex ϕ function will reinforce large distances by rendering them even larger, thus enabling us to detect unique features in the data. On the other hand, a concave function will dampen the effect of outlier objects (i.e. objects totally dissimilar to other objects in the data) in the final representation and thus provide us with a picture of the most important basic patterns in the data. Interesting choices of ϕ functions include among others: the class of homogeneous power transformations given by $\phi(d) = d^\gamma$, $\gamma > 0$, the logistic function $\phi(d) = \exp(d)/(1 + \exp(d))$, the logarithmic function $\phi(d) = \log(d)$, the Huber function $\phi(d) = d^2/2$ for $d < c$ and $\phi(d) = cd - c/2$ for $d > c$, for some suitably chosen constant $c > 0$, etc. Also, the notation suggests that Z is the only variable that we control; for a given problem both ϕ and A are fixed.

Minimizing the σ_ϕ function without further restrictions does not make much sense. We can minimize it by simply collapsing all the points in the origin of the space ($z_i = 0$ for all i), and thus all corresponding distances become zero. This provides the global minimum of σ_ϕ , but "Indeed, this is not a good drawing!" [10] (page 310). In order to avoid such trivial solutions we need to impose normalization restrictions. Our focus in this work is on imposing constraints on Z , so that the resulting graph drawing problem is well posed. Some of the possible normalizations are:

1. Tutte [25] suggests partitioning Z into (at least 3) fixed points X and free points Y and minimize σ_ϕ over Y only.
2. Impose a constraint on Z (e.g. $Z'Z = I$ or $\text{trace}(Z'Z) = 1$, etc).
3. Impose constraints on the distances (e.g. $\sum_{ij} d^\gamma(z_i, z_j) = 1$, $\gamma > 0$).

An alternative route would be to impose pushing constraints; i.e. vertices in the graph not connected are repelled.

It is important to realize that the choice of normalization is not a trivial matter. It will determine the shape and properties of the drawing. Thus it should be made in an informed way. It is also important that the data influence the drawing through σ_ϕ , while the constraints are usually chosen by considerations of mathematical convenience or global properties of the drawing.

3.2. The special case of squared Euclidean distances. In this section we examine in more depth the special case where $\phi(d_{ij}(Z)) = d_{ij}^2(Z)/2$; that is, we examine the loss function with squared Euclidean distances. This turns out to be the most important case from the algorithmic point of view and many more general cases can be reduced to this one. The following notation is convenient: let $d_{ij}^2(Z) = \text{trace}(Z'B_{ij}Z)$ with $B_{ij} = (e_i - e_j)(e_i - e_j)'$ and where the e_i 's are unit vectors.

Define the matrix O by

$$(2) \quad O = \sum_{ij} a_{ij} B_{ij}.$$

Thus, O has the negative values $-a_{ij}$ as its off-diagonal elements and the row-sums (or column-sums) of the adjacency matrix A as its diagonal elements. Thus, O is doubly-centered, and by construction positive semi-definite. We then have that

$$(3) \quad \sigma_2(Z|A) = \text{trace}(Z'OZ).$$

If in addition we choose $Z'Z = I$ as the normalization, then the coordinates of the points in the optimal drawing correspond to the s eigenvectors corresponding to the smallest s non-zero eigenvalues. Thus, the solution to the problem becomes algorithmically straightforward.

On the other hand, if we decide to partition the coordinates of the solution $Z = [X' Y']'$, according to the two vertex sets (with X containing the coordinates of the vertices of the object points and Y those of the category points) and normalize only X (i.e. $X'X = I$) or only Y (i.e. $Y'Y = I$), then the loss function becomes identical to the one used in multiple correspondence analysis [12, 18]. In particular we have,

$$(4) \quad \sigma_2(X, Y|W) = \sum_{i=1}^N \sum_{k=1}^K w_{ik} d^2(x_i, y_k),$$

subject to the normalization constraint $X'X = I$, or equivalently in matrix form

$$(5) \quad \sigma_2(X, Y|W) = \text{trace}(JX'X + Y'DY - 2Y'WX),$$

where $D = \text{diag}(W'W)$ and contains the univariate marginals of the categories and J corresponds to the number of variables in the data. There are two methods to minimize (5): (i) block relaxation and (ii) projection. In block relaxation the (X, Y) block structure is exploited. We alternate minimization over the variables in block Y by keeping the values in block X fixed, then minimize the variables in block X by keeping Y fixed and iterate between these two steps. This leads to the following well known alternating least squares algorithm [18]:

Step 1: $\hat{Y} = D^{-1}W'X$; Step 2: $\hat{X} = \frac{1}{J}WY$; Step 3: Orthonormalize X using the Gram-Schmidt procedure [12].

In projection methods, we have that for the optimal $Y(X)$, the σ_2 function can be written as

$$(6) \quad \sigma_2(X|A) = \text{trace}(X'(I - P)X),$$

where $P = \frac{1}{J}WD^{-1}W'$ is the average *between categories* projector. Hence, the optimal X corresponds to the solution of the above eigenvalue problem. For both problems the solution at the optimum satisfies $\hat{Y} = D^{-1}W'X$; i.e. category points are in the center of gravity of the objects belonging in a particular category. This is known in the literature as the *centroid principle* [12, 18].

It can be seen that the special case of squared Euclidean distances under the appropriate normalization recovers a well known multivariate technique.

The Sleeping Bag example. Using squared Euclidean distances to draw the graph of the sleeping bag example and using both the $Z'Z = I$ and the $X'X = I$ normalizations (only the coordinates of objects are normalized) we get the representations given in Figures 3 and 4, respectively.




FIGURE 3. The 2-dimensional graph drawing of the sleeping bags using squared Euclidean distances and the $Z'Z = I$ normalization




FIGURE 4. The 2-dimensional graph drawing of the sleeping bags using squared Euclidean distances and the $X'X = I$ normalization

The two representations reveal several things. First of all, the inherent rotational invariance of the solution. Moreover, both solutions capture the presence of good, expensive sleeping bags filled with down fibers and cheap, bad quality sleeping bags filled with synthetic fibers and the absence of bad, expensive sleeping bags. It also shows that there are some intermediate sleeping bags in terms of quality and price filled either with down or synthetic fibers. However, the absence of a centroid principle for the representation based on the $Z'Z = I$ normalization results in placing most vertices (both objects and categories) on the periphery of the graph. On the other hand, the presence of the centroid principle improves the representation aesthetically in Figure 4 and also serves as a guide to the data analyst for uncovering the basic patterns present in the data. It is worth noting that for the sleeping bag example a 3-dimensional representation (i.e. $s = 3$) does not reveal any additional patterns in the data. A classical graph drawing algorithm is Sugiyama's [24], which computes layouts for arbitrary graphs by first converting them into k -level graphs (where k in this example would correspond to the number of variables) and then attempts to reduce the number of crossings. Such a layout is given in Figure 5; however, this representation mostly reveals bivariate relationships in the data (see also Section 5), rather than a global map of both objects and categories.




FIGURE 5. A graph drawing of the sleeping bags using Sugiyama's algorithm

4. A General Optimization Framework.

One of the challenges that our general framework poses is the large number of possible configurations that we may encounter when optimizing the loss function σ_ϕ . The different possibilities come about by the various combinations of parameters we can select, such as the type of distances (e.g. ℓ_1 , ℓ_2 , ℓ_∞), the shape of the ϕ function, the form of the normalization constraint. We propose a ubiquitous optimization scheme based on the principles of majorization that with few alterations can achieve our goal. Majorization is discussed in general terms in a series of papers [6, 8, 15, 17].

Majorization algorithms use *surrogate* functions that render the optimization problem under consideration easier to handle. The goal (in general notation) is to optimize a function $\psi(\theta)$ over $\theta \in \Theta$, with $\Theta \subseteq \mathbb{R}^s$. Suppose that a function $v(\theta, \eta)$ defined on $\Theta \times \Theta$ satisfies

$$(7) \quad \psi(\theta) \geq v(\theta, \eta) \text{ for all } \theta, \eta \in \Theta, \text{ and } \psi(\theta) = v(\theta, \theta) \text{ for all } \theta \in \Theta,$$

Thus, for a fixed η , $v(\bullet, \eta)$ is below ψ and it touches ψ at the point $(\eta, \psi(\eta))$. We then say that $\psi(\theta)$ *majorizes* $v(\theta, \eta)$ or that $v(\theta, \eta)$ *minorizes* $\psi(\theta)$. Two key results in [6] show that: (i) if ψ attains its maximum on Θ at $\hat{\theta}$, then $v(\bullet, \hat{\theta})$ also attains its maximum of Θ at $\hat{\theta}$, and (ii) if $\tilde{\theta} \in \Theta$ and $\hat{\theta}$ maximizes $v(\bullet, \tilde{\theta})$ over Θ , then $\psi(\hat{\theta}) \geq \psi(\tilde{\theta})$. These two results suggest the following algorithm for maximizing $\psi(\theta)$.

- Step 1: Given a value $\theta^{(k)}$ construct a minorizing function $v(\theta^{(k)}, \eta)$.
- Step 2: Maximize $v(\theta^{(k)}, \eta)$ with respect to η and set $\theta^{(k+1)} = \eta^{\max}$.
- Step 3: If $|\psi(\theta^{(k+1)}) - \psi(\theta^{(k)})| < \epsilon$ for some predetermined small ϵ stop; else, go to Step 1.

In order for this algorithm to be of practical use, the minorizing function v should be easily maximized. Obviously, in case we are interested in minimizing ψ a majorizing function v that needs to be minimized in Step 2 needs to be identified.

For the problem at hand it can be shown that (details given in [20]) for almost all choices of transformations that we are interested in, the loss function $\sigma_\phi(Z|A)$ is majorized by $\text{trace}(Z' C(Z, A) Z)$ for an appropriately defined matrix $C(Z, A)$. This latter fact implies that under an orthonormality normalization constraint, at every step of the algorithm we just have to solve an eigenvalue problem.

For example, suppose that we are interested in using $\phi = d^\gamma$ with $\gamma = [1, 2]$. This is a convex function with growth rate slower than the quadratic. It contains as special cases both the σ_2 function and the σ_1 function that deals with Euclidean distances (distances without the square). The corresponding minimization problem can be easily solved by constructing a majorization function based on Young's inequality.

$$(8) \quad d_{ij}^\gamma(Z) \leq \frac{2-\gamma}{2} d_{ij}^\gamma(\tilde{Z}) + \frac{2}{\gamma d_{ij}^{2-\gamma}(\tilde{Z})} d_{ij}^2(Z),$$

which implies that we can construct a quadratic majorizing function, and hence get

$$\sigma_\gamma(Z) \leq \sum_{i=1}^n \sum_{j=1}^n a_{ij} \left(\frac{2-\gamma}{2} d_{ij}^\gamma(\tilde{Z}) + \frac{2}{\gamma d_{ij}^{2-\gamma}(Y)} d_{ij}^2(Z) \right) = \text{trace} \left(\frac{2-\gamma}{2} \tilde{Z}' C_\gamma(\tilde{Z}, A) Z + \frac{2}{\gamma} Z' C_\gamma(\tilde{Z}, A) Z \right),$$

where

$$(9) \quad C_\gamma(\tilde{Z}, A) = \sum_{i=1}^n \sum_{j=1}^n \frac{a_{ij}}{d_{ij}^{2-\gamma}(\tilde{Z})} A_{ij}, \quad \gamma \in [1, 2].$$

Thus, in an iteration we minimize $\text{trace}(Z' C_\gamma(Z^{\text{previous}}, A) Z)$ over normalized Z , where Z^{previous} denotes the value of Z in the previous iteration.

5. Bipartite Graphs and Parallel Coordinate Plots.

In this section it is shown how our graph drawing framework can help improve the presentation of multivariate categorical data using parallel coordinates plots [26]. In such plots, we draw J parallel straight lines, one for each variable. The objects are then plotted on each of the lines, and points corresponding with the same objects are connected by broken line segments (and perhaps colored with different colors). These plots (implemented in many data visualization packages) capture bivariate relationships and can reveal various interesting features in the data. However, the parallel coordinate plot for nominal variables may not be particularly illuminating. In this section we use the mammals dentition data set for illustration purposes (the data are analyzed in [18] and a short description of the variables and their frequencies can be found at <http://www.stat.lsa.umich.edu/gmichail/gd>). The parallel coordinates plot of this data set from XGobi is given in Figure 6. The




FIGURE 6. Parallel coordinates plot of the mammals data set

plot reveals certain interesting features, such as the presence of a group of mammals -the ruminants- that are characterized by the absence of top incisors (TI1) and the presence of a large number of bottom incisors (BI5) (see discussion in [18]), or all mammals with no top canines (TC1) have also zero bottom canines (BC1), etc.

The question is whether some other arrangement of the category points would reduce the number of crossings and thus lead to a cleaner and more informative representation of the data. It has been found in usability studies [23] that reducing the number of crossings is by far the most important aspect in graph drawings. Suppose that we are allowed to place the category points of variables j at arbitrary locations along the j th vertical lines. We propose to use the coordinates for the categories found by the multiple correspondence analysis solution. Notice that each object defines a line through J category points. Let q_{ij} be the induced quantification of object i on variable j , which is the same value as the corresponding quantification of the category of variable j that object i belongs in. We can partition the variance in the induced quantifications, as in Table 2

Source	Sum of Squares	Matrix Expression
Within Objects, Between Variables	$\sum_{i=1}^n \sum_{j=1}^J (q_{ij} - q_{i\bullet})^2$	$q'(D - \frac{1}{J}C)q$
Between Objects	$J \sum_{i=1}^n (q_{i\bullet} - q_{\bullet\bullet})^2$	$\frac{1}{J}y'Cy$
Total Variance	$\sum_{i=1}^n \sum_{j=1}^J (q_{ij} - q_{\bullet\bullet})^2$	$q'Dq$

TABLE 2. Partitioning Quantification Variance

where $C = W'W$, $D = \text{diag}(W'W)$ and q is a column vector containing the quantifications of the category points of all the variables. This measures in how far the lines connecting the objects deviate from horizontal lines, by computing the variance around the best fitting horizontal line, which is given by $q_{i\bullet}$. Minimizing the ratio of the within-object variance $q'(D - \frac{1}{J}C)q$ to the total variance $q'Dq$ amounts to computing the first dimension of an MCA. Observe that this is the same as maximizing the between-object variance $\frac{1}{J}y'Cy$ for a given total variance, i.e. we also want the horizontal lines to be as far apart as possible. This is discussed in more detail in Chapter 3 in [12]. We illustrate the above with our mammals example. It can be seen that in the layout shown in Figure 7 the number of crossings is significantly reduced compared to the original plot (20 crossings vs 30). Moreover, bivariate relationships in the data are uncovered. For example, the almost straight lines between categories TM1-BM1 and TM2-BM2 indicate that most of the objects belong to these combinations of categories (20 and 42) and very few to TM1-BM2 and TM2-BM1 (3 and 1), respectively. A similar conclusion can be reached for the variables TC and BC. On the other hand, the positioning of the categories of variables BC and TP indicates that most of the objects that belong to BC1 also belong to one of the first three categories of variable TP, while the majority of the object in BC2 belong to TP4 and TP5. It is worth pointing out that this representation does not quite reproduce the relationships that one would find by analyzing each pair of variables with correspondence analysis [1], since the quantifications of the categories come from an analysis of all eight variables and not from separate analyses of the bivariate tables. It should also be noted that minimizing the number of crossings is one of the objectives, which can also be achieved by algorithms for hierarchical graphs [14]; the other objective is to obtain information about

bivariate relationships, which standard graph drawing algorithms would not provide due to their different focus.




FIGURE 7. Optimized parallel coordinates plots for the mammals data. On the left axis the quantifications of the category points are given, while the categories are numbered in the plot

6. Choice of Parameters.

In this section we briefly look at the impact of the choice of some of the parameters in our framework such as the normalization constraint and the transformation function on the resulting representation. In Figures 8 and 9 the graph layouts of the mammals dentition data set is given under the σ_2 and $\sigma_{1.6}$ loss functions, respectively.




FIGURE 8. The graph representation of the mammals data using the σ_2 loss function

It can be seen that in the latter case the graph drawing exhibits a much stronger clustering pattern of both the object and category points. This suggests that for large data sets such a solution could focus on the most basic patterns present in the data and ignore the minor ones. However, one should be careful, since under the σ_1 function the resulting representation

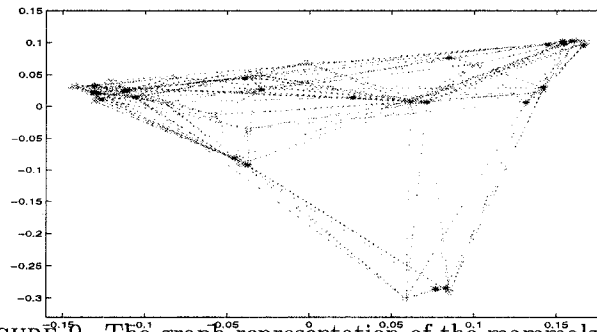


FIGURE 9. The graph representation of the mammals data using the $\sigma_{1.6}$ loss function

consists of $s + 1$ points (s being the embedding dimension) and thus leads to a rather trivial solution. As argued in [21] this is due mostly to the choice of an orthonormality constraint.

On the other hand, the Tutte normalization heavily depends on the choice of the points being fixed as Figures 10-12 indicate. For example, fixing the category points of a single variable leads to a representation where all the other points are clumped near the center of the plot (e.g. Figures 10 and 11), while fixing the coordinates of all the points of one of the vertex sets can lead to informative plots (e.g. Figure 12). In that case the position of the nodes of the other vertex set is guided by the *barycentric* principle [16, 25]. For example, the object points in Figure 12 are in the center of gravity of the categories they belong to. The main issue then becomes of how to choose automatically and in an intelligent way the coordinates of the category points, which is a topic of current research.




FIGURE 10. Tutte solution with the points of variable Price fixed

7. Concluding Remarks.

In this study a general framework of visualizing data sets through graph drawing is considered. The underlying mathematical problem is rigorously defined and a ubiquitous optimization framework based on the concept of majorization is introduced. It is further shown that the familiar technique




FIGURE 11. Tutte solution with the points of variable Quality fixed




FIGURE 12. Tutte solution with the points of all the variables fixed

of Multiple Correspondence Analysis (MCA) is a special case of our general framework. The MCA solution can also be used to improve the presentation of categorical data sets through parallel coordinates plots. Moreover, the choice of normalizations, as well as the type of distances and the transformation function ϕ have a significant impact on the resulting representation. An important future direction is the introduction of pushing constraints in the framework and the analysis of weighted graphs that represent data sets with numerical variables.

REFERENCES

- [1] Benzécri, J.P. (1992), *Correspondence Analysis Handbook*, Marcel Dekker, Inc.
- [2] Borg, I. and Groenen, P.J.F. (1997), *Modern Multidimensional Scaling*, Springer, New York
- [3] Brandes, U. and Wagner, D. (1997), 'A Bayesian paradigm for dynamic graph layout', *Proceedings of the 5th International Symposium on Graph Drawing*, Di Battista (ed), vol 1353 of Lecture Notes in Computer Sciences, 236-247, Springer, New York
- [4] Buja, A., Swayne, D.F., Littman, M. & Dean, N. (1998), 'XGvis: Interactive data visualization with multidimensional scaling', under review *Journal of Computational and Graphical Statistics*
- [5] Coleman, M.K. (1996), 'Aesthetics based graph layout for human consumption', *Software Practice and Experience*, **26**, 1415-1438

- [6] De Leeuw, J. (1994), 'Block relaxation algorithms in statistics', *Information Systems and Data Analysis*, Bock et al. (eds), Springer
- [7] De Leeuw, J. (2001), Multidimensional Scaling, to appear in the *International Encyclopedia of the Social and Behavioral Sciences*, Elsevier
- [8] De Leeuw, J. & Michailidis, G. (2000), 'Majorization methods in statistics', *Journal of Computational and Graphical Statistics*, **9**, 26-31
- [9] De Leeuw, J. & Michailidis, G. (2000), 'Graph layout techniques and multidimensional data analysis', *Game Theory, Optimal Stopping, Probability and Statistics. Papers in honor of Thomas S. Ferguson*, F.T. Bruss and L. Le Cam (eds), IMS Lecture Notes-Monograph Series, 219-248
- [10] Di Battista, G., Eades, P., Tammasia, R. & Tollis, I. (1998), *Graph Drawing: Algorithms for the Geometric Representation of Graphs*, Prentice Hall
- [11] Eades, P. (1984), 'A heuristic for graph drawing', *Congressus Numerantium*, **42**, 149-160
- [12] Gifi, A. (1990), *Nonlinear Multivariate Analysis*, Wiley
- [13] Hartigan, J. (1975), *Clustering Algorithms*, Wiley
- [14] Healy, P. and Kuusik, A. (1999), 'The vertex-exchange graph: a new concept for multi-level crossing minimization', *Proceedings of the 7th International Symposium on Graph Drawing*, Kratochvil (ed), vol 1731 of Lecture Notes in Computer Science, 205-216, Springer, New York
- [15] Heiser, W.J. (1995), 'Convergent computing by iterative majorization: theory and applications in multidimensional data analysis', in *Recent Advances in Descriptive Multivariate Analysis*, Krzasnowski (ed.), Clenredon Press
- [16] Herman, I, Melancon, G. and Marshall, M.S. (2000), 'Graph visualisation and navigation in information visualisation', *IEEE Transactions on Visualization and Computer Graphics*, **6**, 24-43
- [17] Lange, K., Hunter, D.R. & Yang, I (2000), 'Optimization transfer using surrogate objective functions', *Journal of Computational and Graphical Statistics*, **9**, 1-25
- [18] Michailidis, G. & de Leeuw, J. (1998), 'The Gifi system for descriptive multivariate analysis', *Statistical Science*, **13**, 307-336
- [19] Michailidis, G. & de Leeuw, J. (2000), 'Multilevel homogeneity analysis with differential weighting', *Computational Statistics and Data Analysis*, **32**, 411-442
- [20] Michailidis, G. & de Leeuw, J. (2000), 'Multivariate data analysis using constrained pulling', Technical report, Department of Statistics, The University of Michigan
- [21] Michailidis, G. & de Leeuw, J. (2000), 'Homogeneity analysis by alternating least absolute deviations', Technical report, Department of Statistics, The University of Michigan
- [22] Prediger, S. (1997), 'Symbolic objects in formal concept analysis', Proceedings of the Second International Symposium on Knowledge, Retrieval, Use and Storage for Efficiency, G. Mineau & A. Fall (eds)
- [23] Purchase, H.C, Cohen, R.F. and James, M. (1995), 'Validating graph drawing aesthetics', *Proceedings of the 3rd International Symposium on Graph Drawing*, Brandenburg (ed), vol 1027 of Lecture Notes in Computer Science, 435-446, Springer, New York
- [24] Sugiyama, K., Tagawa, S. and Toda, M. (1981), 'Methods for visual understanding of hierarchical systems', *IEEE Transactions on Systems, Man and Cybernetics*, **11**, 109-125
- [25] Tutte, W.T. (1963), 'How to draw a graph', *Proceedings of the London Mathematical Society*, **13**, 743-767
- [26] Wegman, E.J. (1990), 'Hyperdimensional data analysis using parallel coordinates', *Journal of the American Statistical Association*, **85**, 664-675
- [27] Wills, G.J. (1999), 'NicheWorks: Interactive visualization of very large graphs', *Journal of Computational and Graphical Statistics*, **8**, 190-212