Marketing Bulletin, 1998, 9, 16-27, Article 2

Perceptual Modelling of Product Similarities Using Sorting Data

David Bimler and John Kirkland

Perceptual maps and perceptual trees are two graphic techniques for presenting data about the structure of similarities between products or brands. Both do so in a sufficiently simple way to make the data useful in practice as bases for making marketing decisions. This paper considers the special case that arises when the similarity information is collected using the 'Method of Sorting', a convenient procedure when the number of products is large. We argue that the conventional approach of pre-processing 'sorting data' into a form suitable for perceptual mapping algorithms leads to distorted results. Alternative algorithms are described for constructing maps and trees directly from the data, bypassing the preprocessing stage. This combination of perceptual mapping and the Method of Sorting is illustrated with two simple examples

Keywords: perceptual modelling, similarity, scaling, sorting, mapping.

Introduction

One of the broad aims of market research is to inform managers about the position of a product on the market. Managers often want some insight into how consumers perceive a product, relative to competing products. They want to know which attributes of the product distinguish it from its competitors, and whether the product could be marketed more effectively if these attributes were changed to reposition it.

The abstract concept of 'market position' lends itself to a spatial metaphor, in which products are represented by points, with the distance between two points corresponding to the similarity between that pair of products. Thus, a two-stage process of (a) collecting data about the perceived similarities between the products, and (b) representing the data in the form of a two-dimensional map or a three-dimensional spatial model, is one possible means towards the end of pinpointing 'market position'. This is the technique of perceptual mapping. Closely related to it is another common technique of using a 'perceptual tree' (a dendrogram) to represent the data. Both produce models which are clearer than tabulating the raw similarity information, and more likely to be useful as a basis for decision-making.

In the application considered here, the goal is to investigate consumer perceptions of the similarities and differences between brands or varieties within a particular class of products: e.g. pain relievers (Carroll, Clark & DeSarbo 1994); fruit (Clarke, 1976); breakfast snacks (Green & Rao 1972; Young & DeSarbo 1995); cola drinks (Schiffman, Rynolds & Young 1981)**;** political parties (Lingoes & Borg 1978)**.** Such an investigation might be performed as a preliminary to redesigning a product, brand differentiation, or repositioning. In general, it may be used to obtain the knowledge necessary to make a product compete more intensely with rivals, or alternatively, to reduce the intensity of competition, perhaps by finding an unoccupied niche in the market.

When employing these perceptual mapping techniques, it is assumed that some kind of correlation exists between the perceived similarities and consumer behaviour (e.g. purchasing decisions). That is, it is assumed that the maps and trees do not simply summarise the data, but also provide models of the consumers' internal representations of the products – models that are compatible with other data (preferences, purchasing decisions, brand-switching, objective measurements). In other words, it should be possible to interpret the 'similarity' ascribed to a pair of products as 'substitutability': how intensely they compete.

Many market-research applications involve numerous alternatives. For example, there are more than 10 analgesics, to single out one example. As the number of products compared in such an investigation increases, problems arise in the collection of the raw similarity data. It becomes awkward for informants to assess the similarity of every pair. Sometimes the task is even harder because informants are unable to ascribe numerical values to the similarities, so they are asked to rank them instead (e.g. for 10 items, ranking the 45 pairs, from most to least similar, as in Clarke 1976).

In such situations, a useful alternative procedure is the "Method of Sorting" (Rao & Katz 1971; Weller & Romney 1988). Here, informants sort the products into groups with instructions to place similar products in groups together. The burden for each informant is lightened considerably but in compensation a number of data sets must be combined to extract useful results. The question of how this combination should be performed is an open one. This paper summarises a particular way of analysing sorting data which appears to have advantages over more conventional analyses. The method lends itself to the construction of both spatial solutions (*maps*) and *tree* solutions. In this paper we focus on trees of the 'hierarchical' or 'rooted' form, which involve fewer degrees of freedom than 'additive' or 'unrooted' trees.

Two applications are given as examples, with 24 varieties of fruit and vegetable as items in one, and 20 tinned beers in the other. The examples are not intended to imply that the marketing of fruit and vegetables could be improved using our methodology, or that 'cabbage' and 'cauliflower' should be considered as rival products, but to show what could be done with a genuine marketing problem.

Spatial models and Trees

A key distinction between the two forms of model is that, in a map, distances (which model dissimilarities) are measured in the usual spatial way, along a ruler. In a tree, the distance from one item to another is measured along the path one traces getting from one to the other, up a branch and down again.

In a tree the items are located as terminal or 'leaf nodes' (see Figures 1 or 3 for examples). In a tree cluster, one cannot describe any of the items in the cluster as being 'closer' or 'more similar' to an outlying item in a different cluster; they are all equally distant from it. i.e. attached to a particular branch. The concept of 'continuum' is not defined for a tree. However, the various dimensions of a perceptual map are continuous by definition, thus allowing a product to be located relative to others.

A map may be rotated, until its axes are interpretable as broad properties or qualities, thereby locating each product in terms of the value of these qualities, and indicating which ones should be increased or decreased to move a product toward a desired position. A small number of parameters – coordinates on these axes – specify a product's location in a spatial model (of 'the market'), where its close neighbours (competitors) are apparent. But a map fails to inform us which aspects of the product's image are responsible for producing its values on those broad properties. This is where the branches of a perceptual *tree* are useful.

If a number of items are attached as 'leaves' to a particular branch, that branch can be interpreted as corresponding to some feature or attribute which the items have in common. If the goal is to reposition a product, say away from one grouping of brands and into another, then a tree can indicate which feature(s) should be added or subtracted from the product, or from consumers' perceptions of it.

Ghose (1994) provides an insightful comparison of the relative advantages of maps and trees. The point to emphasise is that the two are complementary because of the slightly different ways they present structure abstracted from the raw data.

General Method

Data Collection

The same two-phase sorting procedure was used in each of the applications which follow. Sets of cards containing either the names (fruit and vegetables) or pictures (beers) of the items to be compared were shuffled and given to participants, along with data recording sheets. For the G-(Grouping) phase participants were requested to sort the items, forming between 6 and 10 piles groups according to judged similarity. After these groups were recorded, informants were asked to merge them into larger ones (the Additive or A-phase), again on the basis of perceived similarities. This merging proceeded in a series of steps: in each step, the participants decided which pair of *groups* of items were most similar, and combined them. The process concluded when no further groups could be merged (e.g. if the dissimilarities between them were too great). Once groups had been formed, items could not be redistributed elsewhere.

The pile-merging steps of the additive phase, which distinguish G/A-sorting from the simple sorting procedure, are intended to glean additional data from each informant. On its own the grouping phase provides information about the nearest neighbours of each item, but is deficient in the longer-range or global information (corresponding to the large-scale structure of a map – the positioning of item clusters, relative to each other – and to the higher-level branching structure of a tree).

This G/A-sorting procedure is convenient for up to 30 or 40 items. If a participant is not familiar with any of the items, he or she can simply omit them.

Analysis

The data from each informant can be considered as an incomplete tree. The grouping phase represents a single "slice" through a tree, for the grouping on its own, while the additive phase reveals the branches of the tree at higher levels. One goal is to construct a 'compromise tree', which agrees with each informant's data, as far as possible, minimising the sum over informants of discrepancies between their observed and the predicted sorting decisions. Note that the data are ordinal level rather than interval- or ratio-level: successive stages of group merging do not reflect equal intervals of decreasing similarity.

Young and DeSarbo have published an algorithm (1995) for constructing best-fit trees for 'conditional rankings', a form of data in which the similarities are explicitly ranked in greater-than / less-than comparisons. It is possible to expand an informant's sorting and merging decisions in the G/A task into a sequence of implied similarity comparisons, and once this is done the Young-DeSarbo algorithm generalises to cover them.

A second goal is to produce a perceptual *map* maximising the agreement with each informant's data. For this, sorting and merging decisions are again 'expanded' into their implied ranking of similarities. Then these rankings are subjected to multidimensional scaling (MDS), using an algorithm similar to the MAXSCAL program for triadic and tetradic similarity comparisons (Takane, 1978). The whole process is iterative, since the expansion is performed in the context of an initial estimate of the map, and each time the map is updated to improve its agreement with the implied rankings, the expansion must be performed again. The algorithm for trees has the same iterative aspect.

The workings of this expansion are the innovative component of our methodology (explained in Bimler & Kirkland 1997). They are difficult to summarise in a short paper. Imagine, though, that item C is grouped with A and B by half the participants, and with items X and Y by the other half. It is fair to infer that C is situated about halfway between the groups (A,B) and (X,Y). But the current iteration of the map or the tree affects what conclusions one can draw about the rank-order of the separate similarities between C on one hand and the items A, B, X and Y on the other. This is where our method differs from conventional approaches.

A widely-used approach to MDS with sorting data is to take the number of times a given pair of items occur in the same group (summed over informants and over merging steps), and to use this 'co-occurrence' as an estimate of the similarity for that pair. However, there are methodological problems about summing ordinal data like this, as it seems to introduce artefacts into the solutions. In some cases this causes an exaggerated clustering of the items; in other cases, items are displaced outwards from the centre, to form a kind of annulus. Part of the problem is that co-occurrences conceal the role of *transitivity* in sorting data. If a participant groups A with B because they are similar, and B with C for the same reason, then A and C necessarily co-occur, even if the similarity between them is low. In a series of comparisons in situations where external information was available about the true arrangement of the items (for example, facial expressions), we found that the expansion approach used here – which we have called the "Method of Reconstructed Dyads" – gave superior results (Bimler & Kirkland 1997).

Co-occurrence pre-processing does not explicitly feature in Multiple Correspondence Analysis and Homogeneity Analysis (Tenenhaus & Young 1985), two variants of MDS, which have been applied to sorting data in some studies. But co-occurrences are still used in these analyses, though in a concealed form (Willem van der Kloot, personal correspondence), resulting in the same distortions.

Generally the sorting and merging responses vary from participant to participant. It is tempting to use these variations for the purposes of market segmentation, i.e. to distinguish subgroups of participants (Arabie & Boorman, 1973), but we have met with no success in this.

Application 1: Fruit and vegetables

Procedure

Twenty-four common varieties of fruit and vegetable were selected, and their names written on slips of card for convenient sorting. Thirty-one participants were recruited informally. These participants arranged the labels into an average of 9.7 groups in the first phase (G), and were able to make an average of 6.8 merging steps in the second phase (A). After merging the groups, there was an average of 2.9 groups left. Although some participants omitted items, these were not a problem for the analysis.

Results

A two-dimensional map was adequate for representing the structure of clusters revealed by the tree diagram. The obvious feature of both the map and the tree is a distinction between fruit and vegetables.

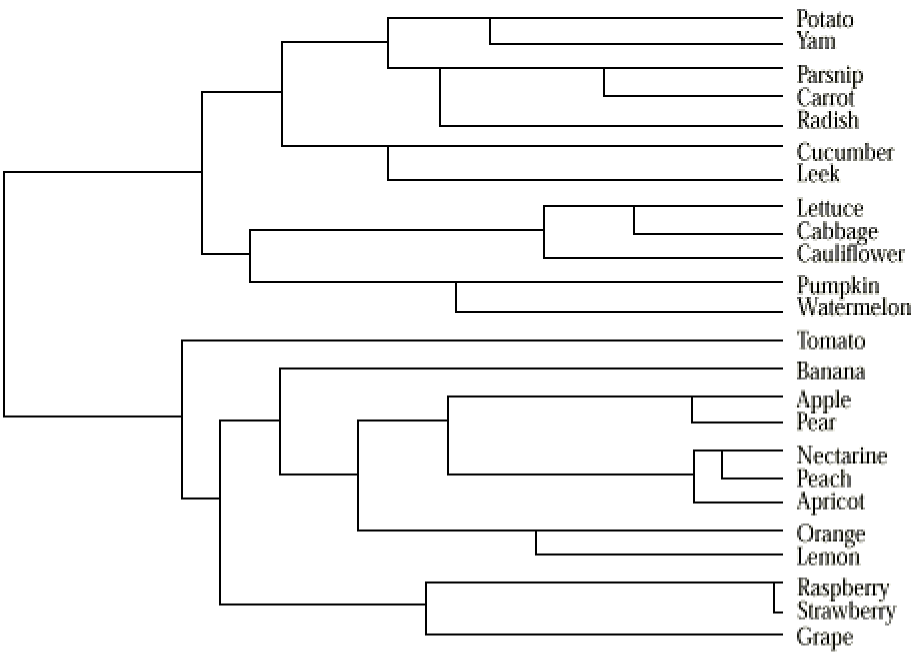

Figure 1. Hierarchical tree for 24 fruit and vegetables

The perceptual tree (see Figure 1) reveals that within the broad 'vegetable' grouping, there are separate clusters of 'underground' vegetables (divided further into tubers and roots) and 'greens'. Within the fruit category, some obvious clusters are the pip fruits; the stone fruits; the citrus fruits; and the berry fruits. The berry fruits, *raspberry* and *strawberry*, were grouped together by all but two informants. *Tomato* occupies a somewhat peripheral position.

In the perceptual map (see Figure 2), rather than forming discrete categories, the fruit and vegetables are separated along a continuous dimension. Some items lie in intermediate positions: such as *tomato* (which has vegetable qualities, despite being generally classified with the fruit), and *watermelon* (which is generally classified as a vegetable, along with *pumpkin*, forming a 'squash'

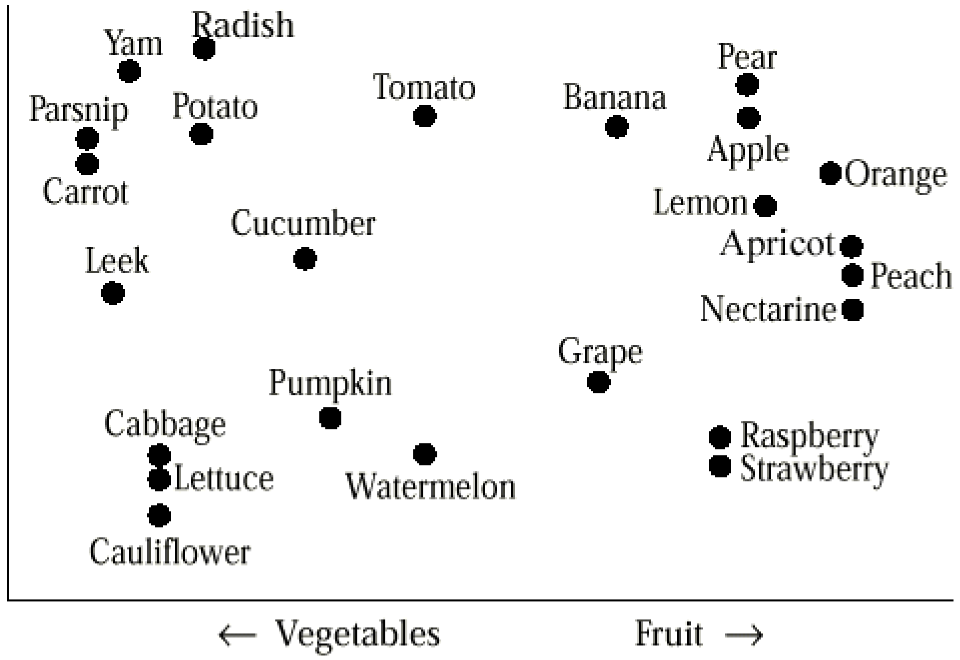

Figure 2. Perceptual map for 24 fruit and vegetables

cluster). One can talk of prototypal vegetables (roots, tubers, and greens) lying at one extreme of this dimension, and prototypal fruit (stone-fruit, then citrus and pip-fruit) at the other end; other items, harder to categorise, are located between the extremes.

This structure uncovered by the map is obscured by the tree diagram, so in this respect the spatial model is preferable. The map also makes it possible to see that 'salad vegetables' (*tomato*, *radish*, *cucumber*) are neighbours, in the sense that all three could be enclosed in a single loop (*lettuce* is further away). In the same sense, 'green vegetables' (*leek* and *lettuce*, as well as *cabbage* and *cauliflower*) are neighbours. The tree separates them into two groups, though it is possible that they would have been grouped together more frequently by participants, and clustered together in the tree, if more 'greens' such as onions and broccoli had been included among the items.

The second dimension in the spatial map does not lend itself to a simple, single interpretation. It may be, given the heterogeneous nature of the items, that this axis has a different interpretation on the left and the right half of the map, since it is illustrating different distinctions being made among fruit and among vegetables.

A similar though more comprehensive selection of fruit and vegetables was mapped by Weller and Romney (1988, p.19). They analysed their sorting data using the 'co-occurrence' approach, producing the annulus effect mentioned above, which has displaced items in their map and arranged them in a figure-8 or S-curve. Apart from this distortion, their map is reassuringly similar to Figure 2.

Application 2: Beers

Procedure

Black-and-white images of beer cans were acquired for 20 varieties of beer, and printed on slips of card. Each can was shown 3 cm. wide and 6 cm. high. Numerical labels from 10 to 29 were assigned randomly to the images, for recording the data. An attempt was made to avoid the products being too heterogeneous. If stouts had been included, for instance, or bottom-brewed ales, the distinctions between them and the present set would have been large and obvious, with no intermediate values – making them difficult to interpret in dimensional terms. For another study of canned beers, concentrating on their sensory properties rather than brand images, see Gains and Thomson (1990).

The slips were sorted by 23 participants, recruited informally from staff and students at Massey University. Participants arranged the labels in an average of 7.1 groups in the first phase (G), and were able to make an average of 3.2 merging steps in the second phase (G). After merging the groups, there was an average of 3.9 groups left.

Some informants were unfamiliar with some of the beers, which were not all widely available or extensively promoted in the Palmerston North area, and omitted them from their sorting. Two imported beers (*Kirin Lager* from Japan, and *Foster's Light*) were sorted by 7 and 10 informants respectively. Three of the regional beers (*Taranaki Draught*, *Coromandel Draught*, and *Vailima Lager*, "The beer of the South Pacific") were sorted 8, 2 and 3 times respectively, so not a lot of weight can be placed on the calculated positions for the last two. Each informant felt knowledgeable enough about 14 of the beers, on average, to sort them.

Results

As in the first study, a two-dimensional map is adequate to contain the clusters formed in the perceptual tree, without distorting them or forcing them to overlap. Within both the tree and the map, there is a broad distinction between the paler, lighter brands ('lagers') and the darker, stronger-flavoured Draught and Bitter brands.

The tree (see Figure 3) makes further distinctions. Within the lagers, the dry varieties (*DB Export Dry*, *Lion Ice* and *Flame Beer*), are separated from the others (*Foster's*, *Rheineck*, *Steinlager*, *Heineken*), with two imported lagers (*Kirin* and *Budweiser*) placed separately. Within the Draught/Bitter grouping, the low-alcohol beers (*Mako*, *Foster's Light* and *Foster's Special*) form a separate cluster (along with *Victoria Bitter*). *Carlsberg Elephant Beer* is an outlier.

From the perceptual map (see Figure 4), it appears that the three low-alcohol beers (*Foster's Light*, *Mako*, *Foster's Special*) lie at the bottom of the second dimension, while the high-alcohol *Elephant Beer* lies near the top; so this second dimension can be tentatively identified as 'alcohol content'. The *first* dimension represents the distinction between pale and dark beers, with particularly light, pale lagers such as *Budweiser*, *Heineken* and *Steinlager* farthest at the right. The calculated position of *Vailima* at the extreme left is based on data from only three informants. Distinctions among the six brands of Draught and Bitter seem to be weaker than between the lagers, but there is a slight tendency for the regional draughts (*Taranaki, Waikato, Coromandel*) to be farther toward the dark pole than the bitters (*DB, Victoria*).

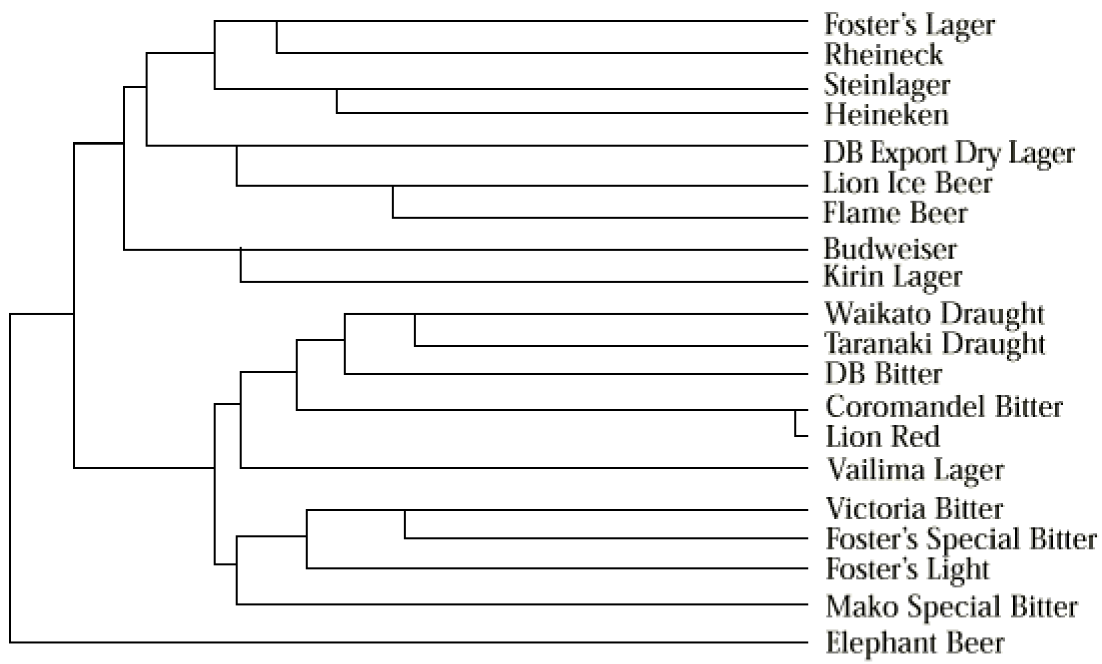

Figure 3. Hierarchical tree for 20 tinned beers

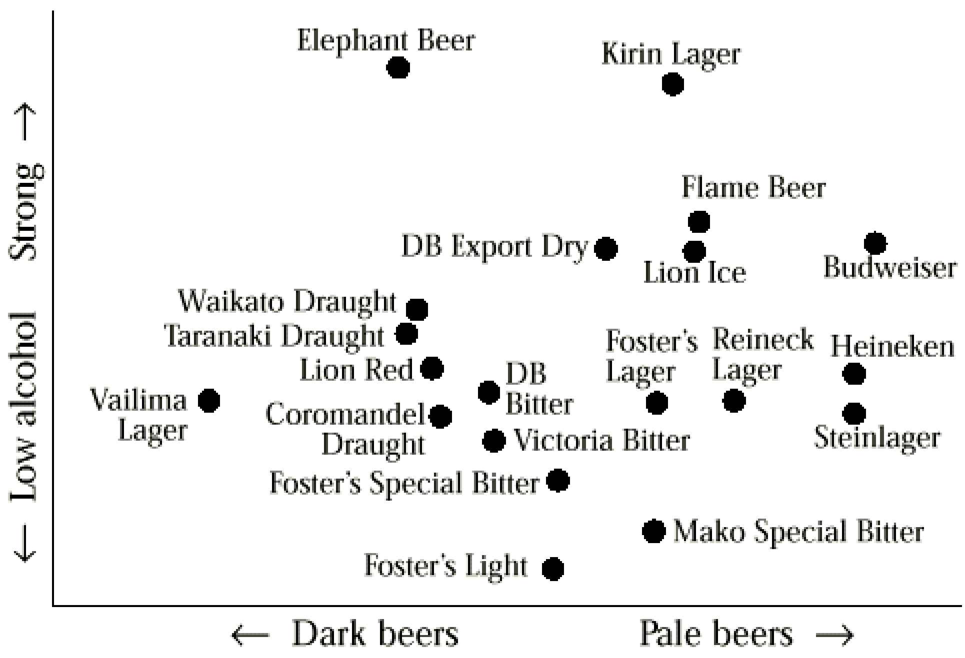

Figure 4. Perceptual map for 20 tinned beers

Conclusions

These results are impressive. First, inspection of both outcomes (tree and map) reveals excellent face validity, and the trees and maps display complementary aspects of the data. Secondly, the sorting tasks are very straightforward. In an unpublished application, the instructions were easily translated into Japanese, while our participants included students as young as 13 years of age. Finally, the whole process is transparent and understandable with easily interpreted results. The MDS procedure we applied to construct the spatial models appears to be free from the artefacts introduced by the usual ‘co-occurrence’ approach, where these artefacts obscure the distinctions between items by forcing them either into tight clusters, or into an annulus around the periphery of the map.

While just two examples were presented here, market research contains other applications of MDS where the number of items is large enough for the Method of Sorting (or its variants) to be an appropriate procedure for collecting data. Rather than using MDS to scale *products*, sometimes it is useful to scale *words* used to describe product ‘image’, thereby constructing a map of the descriptors which could possibly be applied to a product. The result is a ‘semantic space’.

The establishment of a semantic space would be a preliminary to locating an actual product within the space, or rather, to locating the cluster of attributes associated with it. This application could be used to compare products, or to identify the cluster of attributes which one desires to associate with a product. This is more informative than merely knowing where products stand, relative to one another, since different areas of the map are labeled in advance, and interpreting its dimensions is no longer an issue.

Naturally this G/A-sorting procedure is not a ‘magic bullet’; like any other tool, it is not suited for all situations, and the user should be aware of its shortcomings and limitations. First, as we have noted, sorting responses vary from one participant to the next. Any form of MDS is innately probabilistic (we are grateful to an anonymous reviewer for this point), but the use of sorting data exacerbates this. The sample size must be large enough to be sure that one is dealing with a representative sample, bringing the ‘consensus tree’ or ‘consensus map’ close to the group average. Miller (1969) argued that 20 was the minimum acceptable number of informants, but at least twice this is preferable, allowing one to check the results’ stability by analysing random subgroups separately.

A second pitfall is the context-dependence of ‘similarity’; the presence of a given item can call the informants’ attention to a distinction within the item set to which otherwise they might not have attended, and skew the pattern of perceived similarities. This problem arises whenever similarity judgments are collected for MDS: they are affected by the selection of items included. The researcher’s contribution in making this selection is important, though unavoidably subjective. For example, whether informants focus on the features which A and B have in common, or the features which distinguish them, can be affected by how many close neighbours each item has – in turn affecting whether the items are grouped together.

Finally, within the expansion of G/A-sorting data into a sequence of comparisons between similarities, some assumptions are made as to the procedure which informants are following – consciously or unconsciously – as they make their sorting and merging decisions. We assume that they are following a hierarchical clustering algorithm – specifically, the ‘link’ or ‘nearestneighbour' algorithm. However, this assumption does not seem to be crucial, since varying it has little effect on results.

References

Arabie P & Boorman SA (1973). Multidimensional scaling of measures of distance between partitions. *Journal of Mathematical Psychology*, *10*, 148-203.

Bimler DL & Kirkland J (1997). Multidimensional scaling of hierarchical sorting data applied to facial expressions. *Scandinavian Journal of Psychology*, *38*, 349-357.

Carroll JD; Clark LA & DeSarbo WS (1984). The representation of three-way proximity data by single and multiple tree structure models. *Journal of Classification*, *1*, 25-74.

Clarke MJ (1976) *The identification of substitutes among selected fruits using multi-dimensional scaling*. Market Research Centre, Massey University, Technical Report 6.

Gains N & Thomson DMH (1990). Sensory profiling of canned lager beers using consumers in their own homes. *Food Quality and Preference*, *2*, 39-47.

Ghose S (1994). Visually representing consumer perceptions. *European Journal of Marketing*, *28*, 5-18.

Green PE & Rao, V R (1972). *Multidimensional Scaling*. Hinsdale, IL: Dryden Press.

Lingoes JC & Borg I (1978). A direct approach to individual differences scaling using increasingly complex transformations. *Psychometrica*, 43, 491-519.

Miller GA (1969). A psychological method to investigate verbal concepts. *Journal of Mathematical Psychology*, *6*, 169-191.

Rao VR & Katz R (1971). Alternative multidimensional scaling methods for large stimulus sets. *Journal of Marketing Research*, *8*, 488-494.

Schiffman SS; Rynolds ML & Young FW (1981). *Introduction to Multidimensional Scaling*. NY: Academic Press.

Takane Y (1978). A maximum likelihood method for nonmetric multidimensional scaling: I. The case in which all empirical pairwise orderings are independent – Theory. *Japanese Psychological Research*, *20*, 7-17.

Tenenhaus M & Young FW (1995). An analysis and synthesis of multiple correspondence analysis, optimal scaling, dual scaling, homogeneity analysis and other methods for quantifying categorical multivariate data. *Psychometrika*, *50*, 91-119.

Weller SC & Romney AK (1988). *Systematic Data Collection*. Newbury Park, Cal.: SAGE.

Young MR & DeSarbo WS (1995). A parametric procedure for ultrametric tree estimation from conditional rank order proximity data. *Psychometrika*, *60*, 47-75.

Acknowledgments

We are grateful to Donald Kirkland and Robyn Surville who ably assisted with data collection in application 1. We acknowledge too the contributions from various anonymous sorters including high school and university students, as well as friends and staff. We are indebted to the Editor of the Marketing Bulletin, and to two anonymous reviewers, for their constructive and very helpful comments on an earlier version of this paper.

David Bimler is an independent researcher, and John Kirkland is an Associate Professor, in the Department of Education Studies and Community Support, Massey University.