MASSACHUSETTS INSTITUTE OF TECHNOLOGY
ARTIFICIAL INTELLIGENCE LABORATORY
and
CENTER FOR BIOLOGICAL AND COMPUTATIONAL LEARNING
DEPARTMENT OF BRAIN AND COGNITIVE SCIENCES

A.I. Memo No. 1468 May, 1994
C.B.C.L. Paper No. 98

Trajectory mapping ('TM'): A new non-metric scaling technique

Whitman Richards and Jan J. Koenderink

This publication can be retrieved by anonymous ftp to publications.ai.mit.edu.

Abstract

Trajectory Mapping "TM" is a new scaling technique designed to recover the parameterizations, axes, and paths used to traverse a feature space. Unlike Multidimensional Scaling (MDS), there is no assumption that the space is homogenous or metric. Although some metric ordering information is obtained with TM, the main output is the feature parameterizations that partition the given domain of object samples into different categories. Following an introductory example, the technique is further illustrated using first a set of colors and then a collection of textures taken from Brodatz (1966).

This report describes research done within the Center for Biological and Computational Learning in the Department of Brain and Cognitive Sciences, in association with the Artificial Intelligence Laboratory. Support for the A.I. Laboratory's artificial intelligence research is provided in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract N00014-91-J-4038. Support for the Center's research is provided in part by ONR contracts N00014-93-I-0385 and N00014-92-J-1879; by a grant from the National Science Foundation under contract ASC-9217041 (funds provided by this award include funds from ARPA provided under HPCC). Additional support is provided by the North Atlantic Treaty Organization, ATR Audio and Visual Perception Research Laboratories, Mitsubishi Electric Corporation, Siemens AG., and Sumitomo Metal Industries. Special thanks to Stephen Gilbert for helpful discussion and to Prof. Rosalind Picard for a convenient set of Brodatz samples. WR was supported under NSF-IRI-89 00 267.

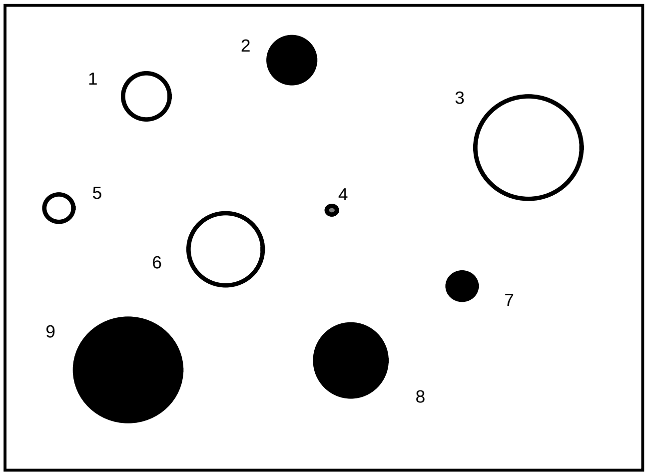

Figure 1: A collection of samples that suggest the domain of objects consists of black and white circles or discs of varying size.

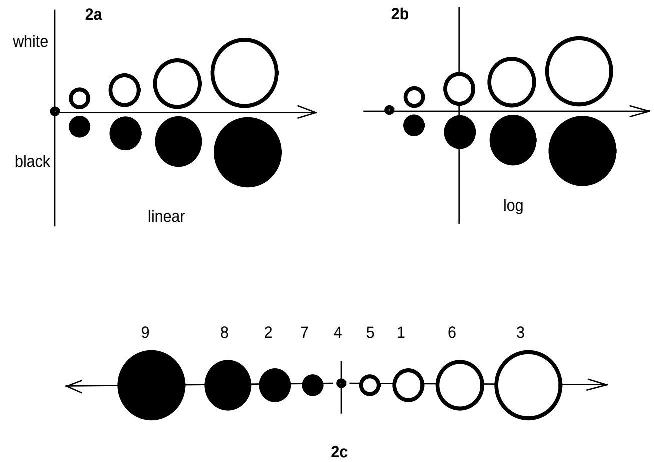

Figure 2: Possible orderings of the black and white objects in a feature space: (A) The white and black objects are separated by a half-line of increasing size from zero. (B) Similar to (A), except now a reference size is chosen for the origin, and size-ratios are evaluated. (C) The black and white objects now take on positive and negative values on the "size" line, with the origin being the point of zero size.

1 Introduction

Object classes are typically considered as a set of properties, seen as particular collection of features. Hence different cognitive classes of objects will occupy different regions (perhaps overlapping) of a multidimensional feature space. The axes and parameterizations of this space are not clear. For example, Figure 1 depicts a domain of colored circles of varying size. Perceptually we immediately group black objects together separately from white objects, suggesting that each group occupy separate half-planes in the feature space. What then is the proper ordering for size? One possibility is to use a half-line, with the point at the origin, increasing from here to infinite size (Figure 2a). Alternately, we could take one (average) size as a unit origin, and then place a circle along the full line according to the log ratio to this reference size (Figure 2b). Still another possibility is to abut two half-lines, one black and the other white, at the degenerate point (Figure 2c). Of course, still other orderings are possible, especially if a continuum through gray is allowed along the black-white dimension. Which of these representations is used by our perceptual system?

One approach to answering such questions is Multidimensional Scaling (MDS) – a procedure introduced in the 50's and 60's by [8, 16] and [18]. (For a brief overview see [6].) This method assumes that the feature space is homogeneous and that distances within the space are measures of the similarity between objects or object classes. For our circle example, we can judge the similarity between all pairs, and then use an MDS algorithm to place an ordering on the samples. Figure 3 illustrates. On a scale of 0 to 10, two circles of the same size but opposite content were given a rank of "5"; two circles differing slightly in size but of the same contrast were ranked "8"; etc. Using these ranks, an MDS algorithm (Systat: Kruskal method) placed the circles on a plane as shown in the upper part of Figure 3. This ordering accounted for about 95% of the variance in a smooth monotonic function relating similarity to distance.

Two observations about the MDS procedure: First, although the circles have been ordered in a feature space, the axes and origin in the derived space are somewhat arbitrary. (Note for example that circles of the same size lie on lines inclined to the x-axis. [1]) Second, the actual relations between the objects in the space have not been made explicit. So, for example, the distance between the largest white circle and the smallest white circle is greater than the distance between the same small white circle and the large black circle. Does it really make sense to infer that the latter are more similar? Furthermore, is it legitimate to conclude that there is a feature (or parameterization) along the path in this space between the small white and large black circle? Nothing in the MDS plot excludes the possibility that there is indeed such a feature common to both, trading whiteness for size, which connects these two objects. The only hint that there is not such a feature (or path) is that all the black objects lie below the $y = 0$ axis, whereas all white objects lie above, excepting the degenerate point. To eliminate some of these ambiguities in an MDS plot, it would be helpful to know exactly which paths through the feature space are legitimate, and which are excluded, such as those that might cross the $y = 0$ axis at an arbitrary angle. This is the primary objective of the Trajectory Mapping, or "TM" procedure.

[1] One might argue that the white objects are interpreted as circles, whereas the black objects are interpreted as solid discs, rather than "circles", and hence the black and white objects differ by more than just their color. Such a proposal would simply strengthen our argument for the value of using the TM procedure.

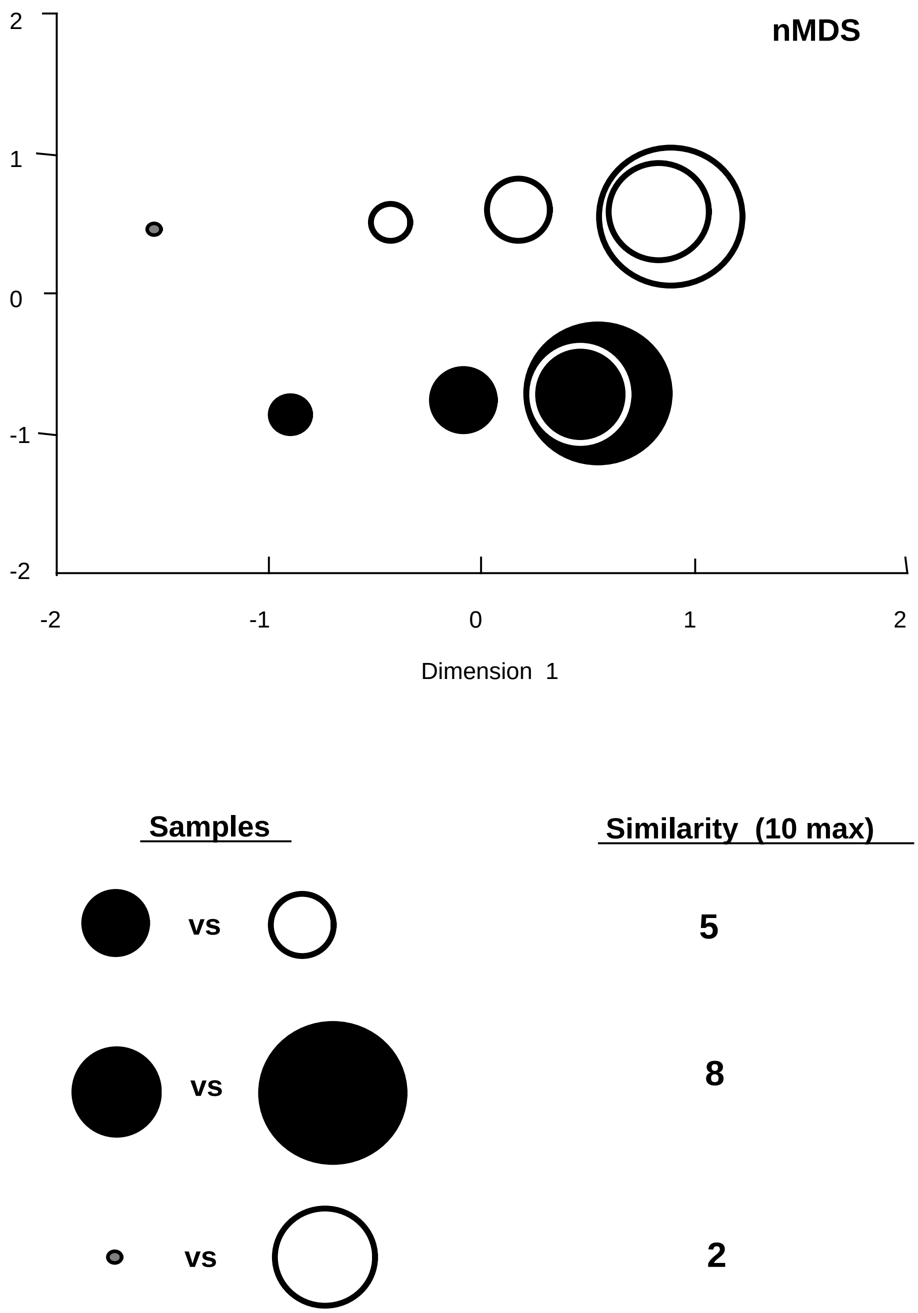

Figure 3: The result of applying a multidimensional scaling procedure to the objects in Figure 1. Typical similarity rankings, on a scale of zero to ten, are illustrated in the lower panel of the figure. The upper panel shows the scaled results. Note that the distance from the second to smallest white circle to the large black circle is less than (i.e. "more similar to") the distance to the large white circle.

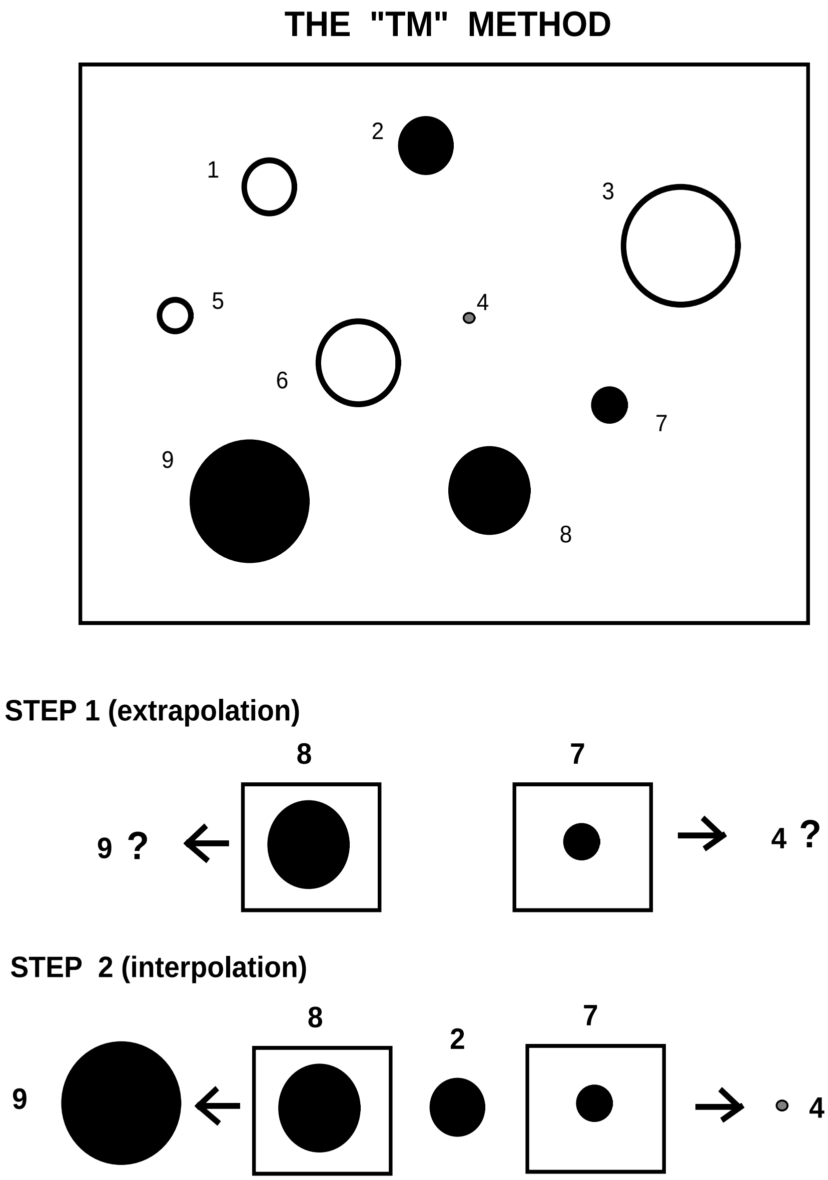

Figure 4: The Trajectory Mapping "TM" procedure. Two samples are picked at random and their extrapolants are found, as illustrated using the pair 8 and 7. Then, continuing to use the same feature for linking the samples, an interpolant is chosen.

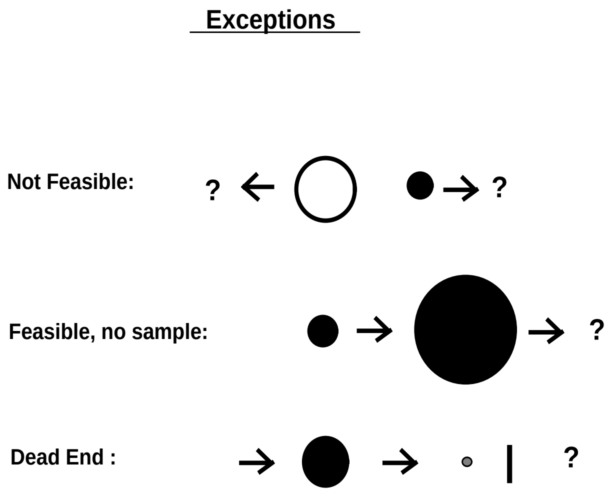

Figure 5: Exceptions to the procedure illustrated in the previous figure. Sometimes two pairs are judged not to have a feature in common. In this case, extrapolation is "not feasible". Second, often extrapolation (or interpolation) is possible, but there is no sample in the given set of objects. Third, either one member of a trial pair, or an extrapolant is at or near the end of the extrapolated feature. All these cases are marked.

2 Trajectory Mapping ("TM")

The "TM" method is conceptually quite simple. Again choose two samples from the given domain, just as was done for MDS. However, rather than asking for a similarity ranking, request that two other samples be chosen from the set which are the best extrapolations from the reference pair. Thus, again for our circle domain, assume the pair are numbers 7 and 8 (see Figure 4). Lay these side by side – or so displayed on a computer screen. Now imagine a feature that takes sample 8 into sample 7, and continue the extrapolation until the feature parameter agrees with one of the other samples (e.g. "4"). Similarly, start with 7 and extrapolate the same feature in the opposite direction through 8 until a sample (9?) is found that is an extention of that chosen feature. Once the two extrapolants have been found, then pick an interpolant, say 2 in this case. (It is important that the interpolant be chosen last so that the reference samples be the primary criteria for extrapolation.) The result of each trial is then a string of five samples linked and ordered by a parameterized feature. Typically, subjects will attempt to choose extrapolants that lie roughly a unit distance from the reference pairs, with the unit set implicitly by the judged similarity of the reference pair. Thus, for our circle example, given the very similar reference pair of 8 and 2, most subjects would pick 7 as the extrapolant and not the limiting case 4 (see Figure 4).

There are three exceptions to the TM procedure as outlined above. First, very often the two reference samples seem incompatible in the context set by the domain of samples. For example, black and white circles are hard to extrapolate. Certainly the subject *could* opt to make a judgment based solely on size, ignoring contrast (see Figure 5). If however, white objects and black objects are considered to be separate classes of objects in the conceptual feature space, then the stronger tendency is to prefer not to forge a path between them. Such "mutually exclusive pairs" are noted because they offer evidence as to boundaries in the subject's feature space.

A second exception to the standard extrapolation protocol is when the the subject feels that an extrapolate (or interpolate) is possible, but the sample is missing from the set. (It is not legitimate to expand the domain, however, say by adding "gray" circles in our illustration.) A simple example would be if the reference pair were the black circles 7 and 9. Clearly the extrapolant is a still larger circle not in the data set. Hence *whenever* extrapolants are feasible, this is noted.

Finally, the last inclusion to the method is to note when a sample (or extrapolant) lies near or at the end of a feature path in the space. These "dead ends" also help identify the boundaries, axes and origins of the TM space. For our example, the degenerate point "4" is a dead end: the circle size can not be further reduced.

3 Circle Example

Using a Mathematica program we applied the TM procedure to the black and white circle set, yielding a collection of $9 \times 8/2 = 36$ quintuples. Of these, half we judged "not feasible", namely the reference pairs of opposite contrast. We then count the frequency of occurence of all feasible pairs and triples (see Table 1). The maximum possible frequency for a pair or triple will thus be 18 – the case where one pair occurs in each feasible (black or white) quintuples.

In order to determine the smooth paths in the feature space, we must use the sets of triples, not the pairs. (If the pairwise data were used, two different feature parameterizations common to one sample could be joined at the common sample – somewhat as if a corner were turned in the space. This will not be the case for triples, assuming the same feature parameterization is used for each trial. Thus, starting with the most frequent triples, we link their overlapping pairs. These linked pairs then become the major pathways in the space, if you will. In Figure 6, we illustrate this construction, with the size of each link indicating the frequency of occurence of that triple. The overlap between the triples thus allows us to follow the feature path. In this construction, we also note the "dead ends" by the vertical bar and the feasible extrapolations beyond the sample set, as indicated by arrows. The result of this procedure is then a web of interconnections between the samples, giving an explicit representation of the paths through the feature space and the boundaries in this space. Of course, we still do not know exactly which feature was varied! However, unlike MDS, the structure of the parameterizations in the space has been made explicit.

To continue this illustrative example, a second step in the analysis is helpful. Often, as shown in the examples that follow, the web of paths in the feature space can be quite complex. Hence it is useful to have some technique for positioning the samples, as well as showing the paths between them. The number of times two samples appear

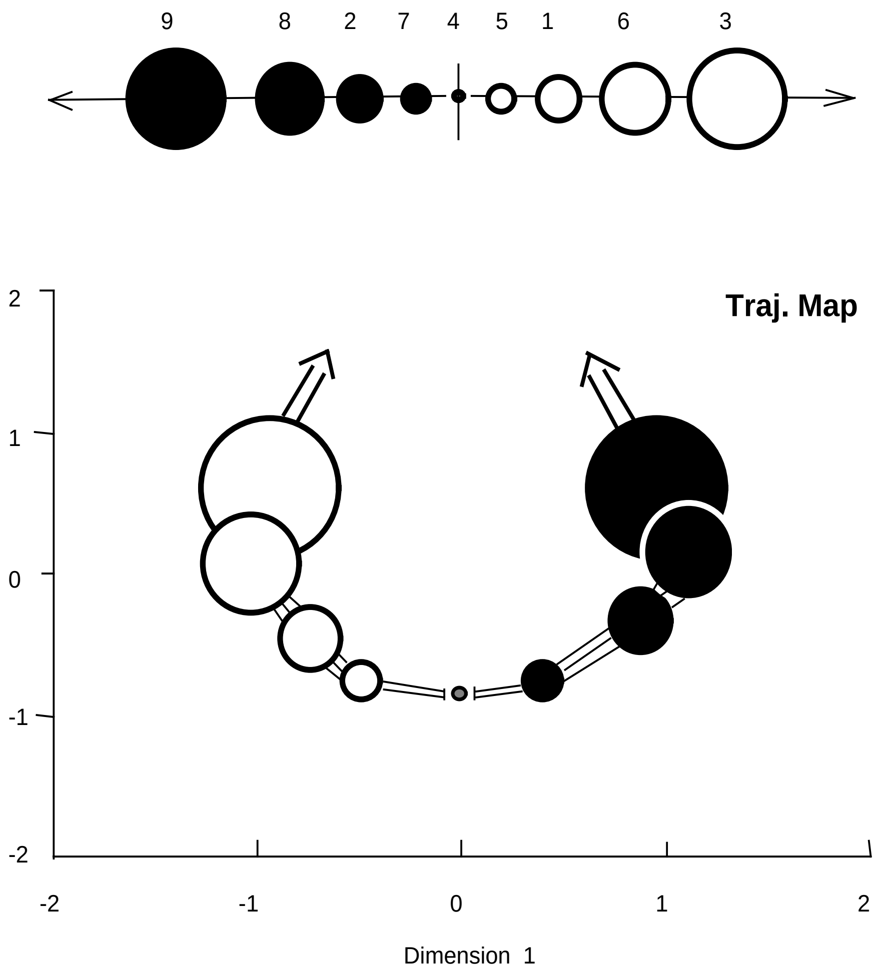

Figure 6: The result of applying the TM procedure to the black and white circles of Figure 1 yields two paths, as illustrated at the top. The black circles can be extrapolated further to the left; the white to the right. There is a "dead end" to each path at the degenerate point (4). The thickness of the lines indicates the "strength" of the path. Below, the linear array has been placed in a space set up by the adjacency counts. The extrapolations appear to converge in this space, but presumably this is an artifact of the use of only two imaginary objects to represent all larger black or white discs.

as neighbors on a path serves this purpose. The rationale is that samples nearby in the space will have a greater likelihood of appearing consistently in the quintuple for each trial than distant samples between which it is difficult (or perhaps impossible) to extrapolate. Furthermore, the adjacency statistic allows us to compare similarities between two samples on two different paths that share a common sample where the two paths cross. Intuitively, these similarities between various crossing points will be determined by the number of shared paths and samples between them. The adjacency frequency then becomes a ranking statistic that can be used in the MDS algorithm. Again Figure 6 (lower) shows the result for our black and white circles. Note that in contrast to the original MDS similarity judgments, now the special "dead-end" point "4" lies at the origin of one of the dimensions in the "adjacency space". Furthermore, there are now clear black-white half planes about this origin, with circles of constant size lying at constant y values. The sole major "artifact" in this adjacency-scaled plotting is the inward curving of the trajectores as the circle size takes on positive values along dimension 2. This is a result of including two "imaginary" samples in the analysis that indicated a feasible extrapolation but not sample in the data set. In any case, with the pathways included, the basic structure of the feature space is clear.

4 The Subway Map

An illustration devised by Stephen Gilbert provides a compelling summary of the major differences between the MDS and TM protocols. If you wish to move about a town using a Metro system, then a map of the routes is needed (Figure 7). These routes are typically identified by color and arranged so that the linking is clear: the topology is preserved but not the metrics. These routes are the threads or features through the space that TM discovers.

In contrast, MDS cares little about these routes, but rather attempts to recover the actual distances between the subway stops – i.e. their true locations on a map, as if a crow were to fly from one to the other. The result is thus an arrangement of the various subway stops without any of the "earth paths" being made explicit. The orientation of the resultant map may also be unclear. If a subject parameterizes the space by, say, the directions "North and East", then the resultant MDS map will have a meaningful orientation. If however, the axes and parameterizations of the space are chosen from the available paths of features i.e. the "subway routes", then there will be no clear relation between these features and the MDS axes. In other words, the MDS procedure requires a homogenous metric space if the cognitive parameterizations of that space are to map into the recovered axes. The TM procedure makes no such assumption. Even if the space is non-metric and not homogenous, the feature parameterizations used to traverse that space can still be recovered.

5 Color Example

As a second example of the TM method, we have begun to explore feature path in color space. This space has received considerable study for over two centuries with the aim of creating a subjectively uniform array of reflectance samples ([2, 4, 7, 9]). Essentially all of this work assumes that the color space is homogenous but non-Euclidean. Classically, with recent support from neurophysiology, the three presumed major axes of the subjective space are: white to black; red to green, and yellow to blue. (These designations can be defined precisely – see [19]). Although movement is allowed in any direction through the space, most theoretical treatments favor radial movement from a neutral achromatic color (i.e. a gray) in a constant lightness plane (corresponding to constant hue of varying saturation) or circumferential movement in this same plane (corresponding to constant saturation but variable hue). We chose for color chips thirty-eight samples from the LJG uniform color space recommended by the Optical Society of America ([9]). All our samples lie on or are adjacent to the equi-lightness plane $L = 0$. The sample chips were 1.5 inch squares placed at tabletop level on a 75% reflecting surface and illuminated with both fluorescent and natural daylight. Because of the large number of pairs ($38 \times 37/2 = 1406$), the space was partitioned into overlapping regions and the judgments were made over four sessions for both JJK and WR and the results pooled.

To illustrate the salient paths in the color space obtained with the TM procedure, we have plotted in Figure 8 only those triples of color samples that occurred three or more times. (The dashed lines indicate additional supporting triples that only occurred twice.) In addition, for clarity, we have omitted those paths which radiate from the neutral gray (square). These data are included in Table 1, however. It is of interest to note that none of these radiating trajectories passed through the neutral gray point. Indeed, in all our 574 trials combined, we only found one triple occurring at least twice that included samples on opposite sides of the gray. Thus, the neutral gray is a "dead end" in the same sense that the point "4" was in our circle example.

The most intriguing finding is that the major pathways – one on the red side, the other on the green side – are not circles about the gray, but rather appear like wings of a butterfly, with the paths folding inward as the yellow-brown colors are approached, giving special status to this hue. It is clear then that the transformations used by the subjects to move along a route in the constant lightness plane must vary in both saturation and hue simultaneously. However, this simultaneous change seems to reach a limiting value when reddish approaches brown or greenish approaches olive – i.e. the yellow-brown-gray boundary. (Note that we also found no strong linkages across the gray-blue axis.) A possible explanation for this effect would be if the paths were based on rates of change of red-versus-green outputs up to but not beyond the equality of red and green, and similarly for the red-blue and green-blue ratios. (See the slope sign condtion of [14].) Rather than analyzing these results in detail at this time, however, our

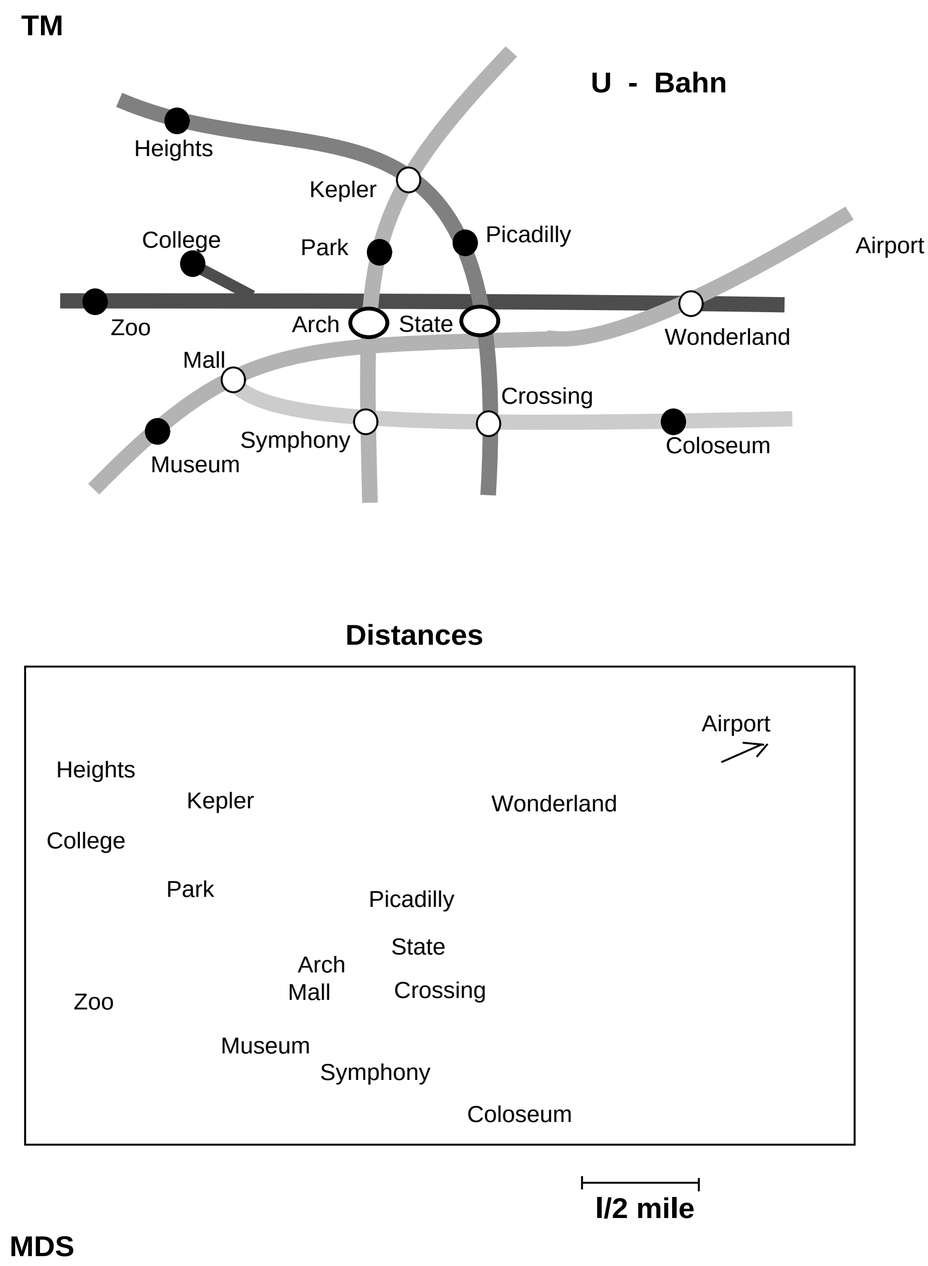

Figure 7: The locations of subway stops can be indicated two ways. The “TM” method shows the topological relationship along different routes; the classical “MDS” method depicts the actual distances.

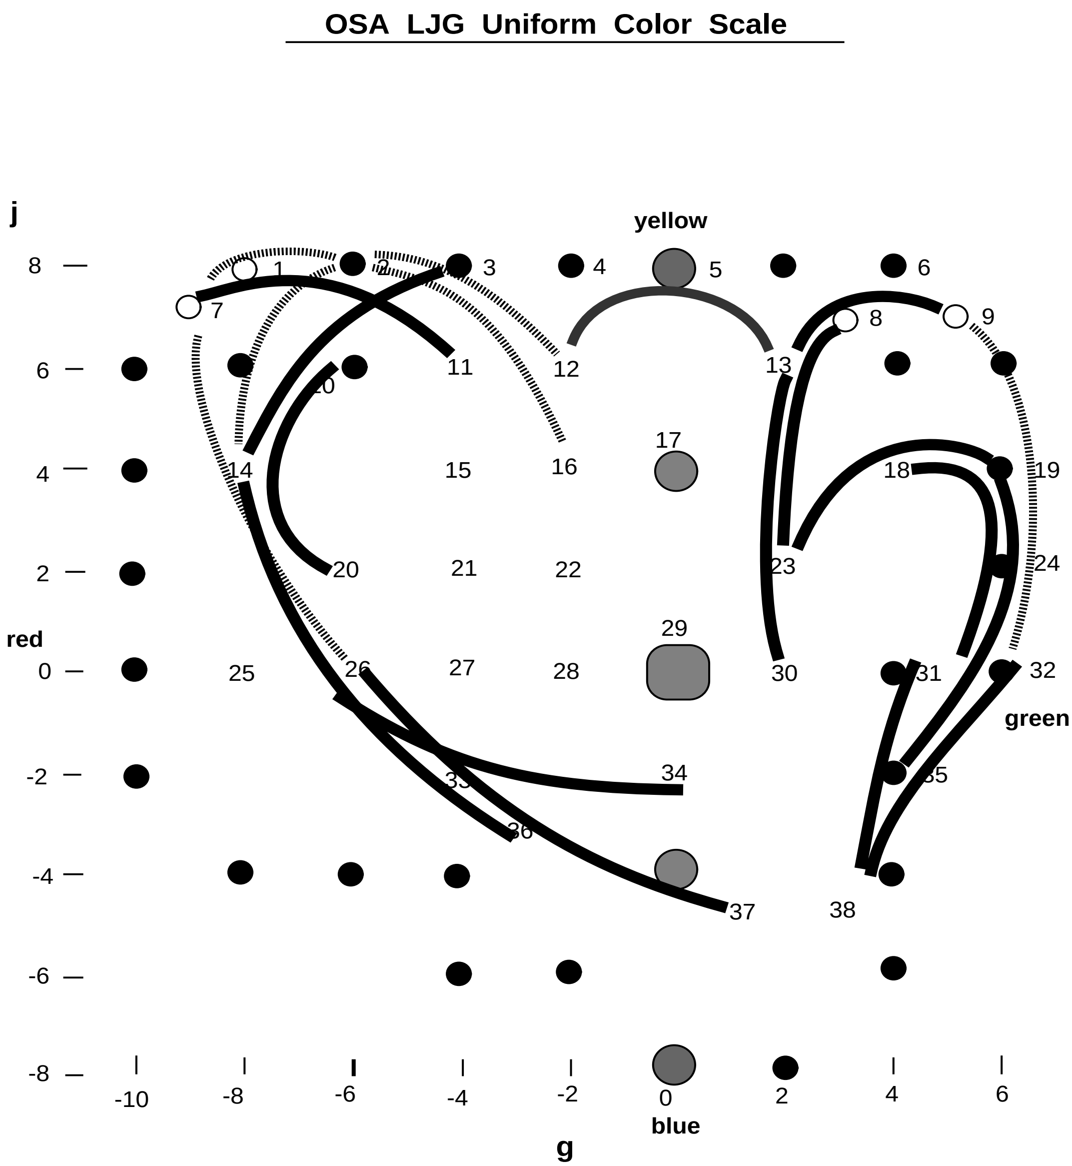

Figure 8: The result of using the TM procedure on 38 color chips taken from a standard (OSA) uniform color space. The TM procedure suggests that the parameterizations of this space vary both in hue and saturation simultaneously, as indicated by the solid arcs. Note that none of these arcs penetrate the neutral gray point; also there appears to be a “reflecting” barrier along the yellow-brown-gray axis. Not shown in the above depiction of our results are the “rays” that emanate from the neutral gray (square).

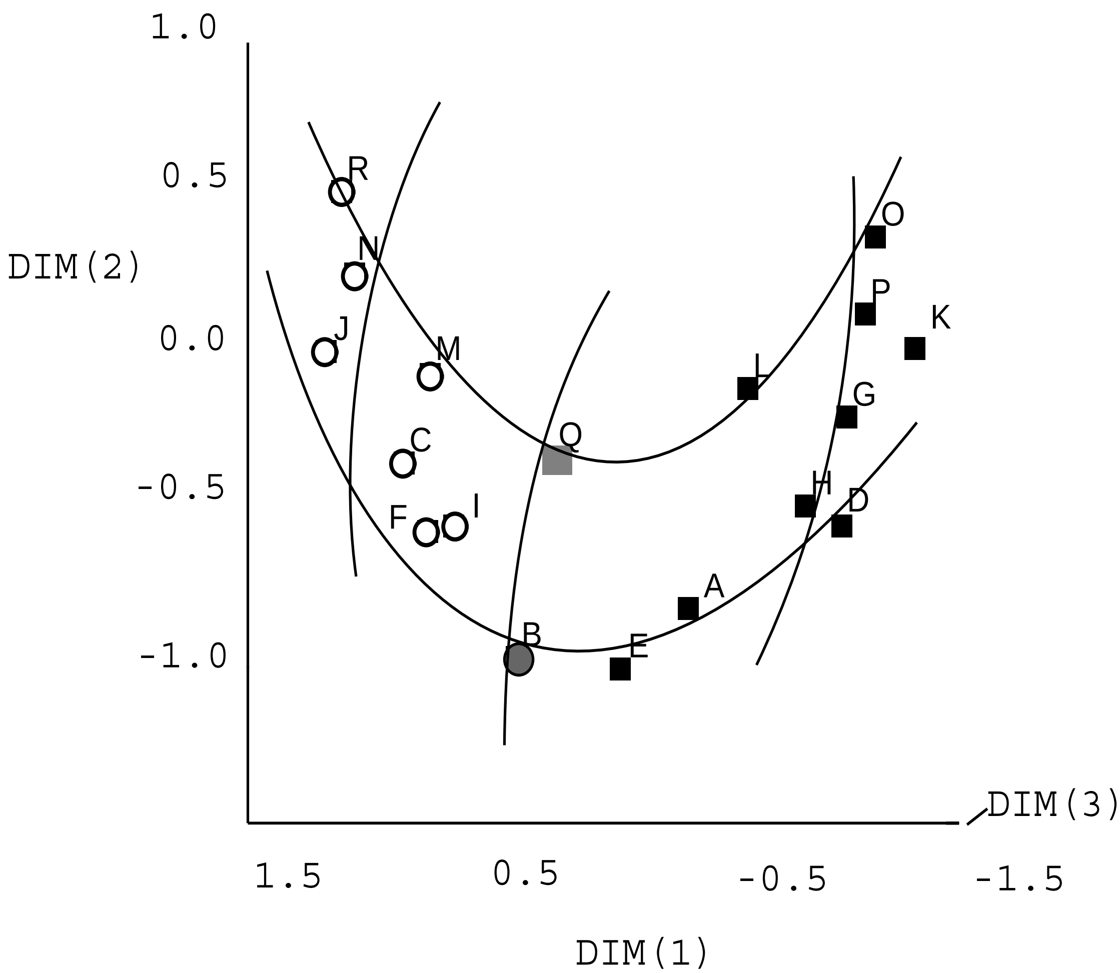

Figure 9: The counts of nearest neighbors along the TM paths can be used as the similarity measure for an MDS algorithm, yielding the array shown. The open circles are the green samples; the filled symbols are reddish. The QB locus is the yellow-blue axis, which lies along a principal curvature on the resultant "saddle surface". The lines indicate constant j or constant g loci.

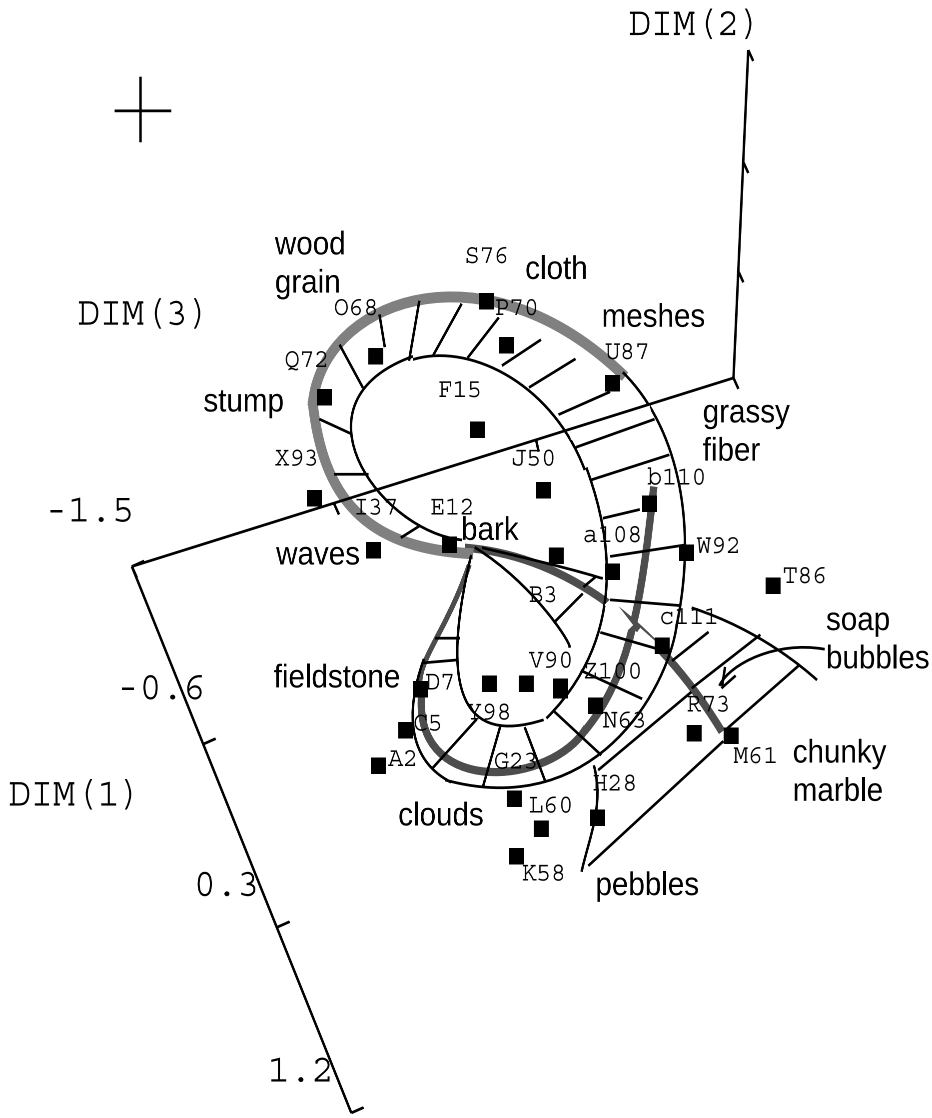

Figure 10: The organization of 30 Brodatz texture samples as revealed by the TM technique (see text).

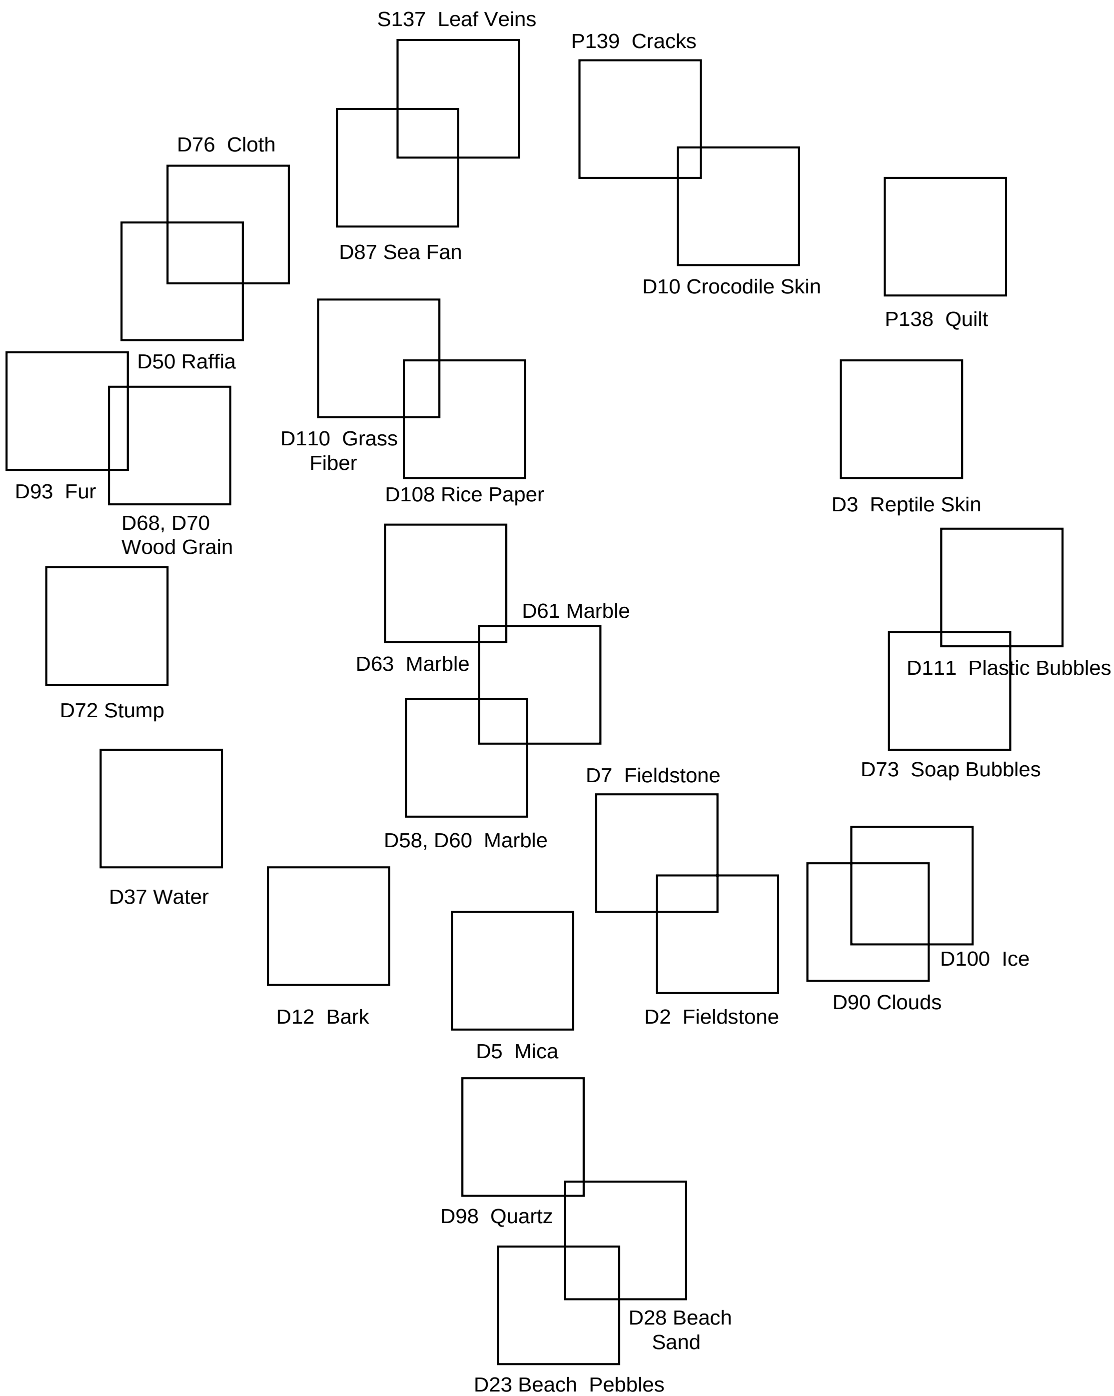

Figure 11: An unfolded projection of the significant paths illustrated in the previous figure. The main loop in this projection is a great circle in three space; the "tail" is probably the beginning of another great circle that completes via the "cracks" pattern P138. (Patterns D15 Straw and D92 Pigskin are not included in the figures: they lie between the depicted loops suggesting still other trajectories are permissible.) The "D" panels are from Brodatz [1]. The "S" and "P" panels are respectively from [17] and [11]. [For a copy of the figure with the texture pictures included, send an email request to burt@psyche.mit.edu.]

point is simply to demonstrate that the TM protocol suggests that there is still another feasible representation for a color space that is quite different from the classical hue/saturation/lightness scheme. A more detailed and theoretical examination of this new representation will be forthcoming.

Before closing this section there is still one more analysis that can be performed using the TM data. Our trajectories assumed that the OSA uniform color space would provide a proper ordering of the space (but no paths). Can we now recover the putative uniform spacing of the samples from the frequencies of adjacent pairs? Figure 9 shows the result of running an MDS algorithm on one data set of eighteen samples using these frequencies as the similarity measure. The open circles represent the greenish samples, the closed the reddish samples. The lines are the horizontal and vertical grid lines transposed from the OSA color space grid of Figure 8 (i.e. indicating constant g and j directions). Although the rough topology is preserved, clearly the spaces are different. The most striking difference is that three dimensions are required to plot the adjacency scaled data because the surface has roughly a saddle-shape. One of the principal curvatures of this surface is the yellow-blue axis. In effect, because the adjacency measure does not require a globally uniform space, we find that the red and green colors each lie on separate surfaces inclined to one another, and intersecting along the yellow-blue axis. (As if we had red and green wings of a butterfly.) The TM procedure, then, is one of the few scaling techniques that actually gave a distinguished geometry to the yellow-blue axis.

6 Texture Example

Our final example is a preliminary exploration of textures. The organization of texture space, unlike the color space, has been extremely resistant to study. Although recently there have been helpful suggestions ([12]), no clear, definitive parameterization of texture space has been found. Even the dimensionality is not clear. Perhaps part of the difficulty is that the space does not have a homogenous metric.

Initially we chose thirty samples from Brodatz ([1]). Later, we added three more samples to the Brodatz set (see Figure 11 and Table 2) and then ran a second TM session using these additions plus ten of the original Brodatz textures. Except for the added set, JJK and WR made independent judgements, and then for clarity we "flattened" the results using Mathematica to obtain the common triples and adjacencies. Table 2 gives these data.

Our initial construction of the pathways by hand gave a complex web of interconnections. (Although E. Todorov has now written an automatic algorithm for unscrambling the simple and robust webs of feature paths, such as data typical of subway routes, this program can not yet handle the complex texture data.) To clarify the organization of the paths, we ran MDS on the adjacency frequencies to obtain a three-dimensional adjacency plot. A Systat package then was used to rotate these MDS adjacency results, allowing us to view the three-dimensional structure of the original 30 points. There was little. The strong impression was a 3D swarm of points.

Next, the major feature paths were added to the swarm, and a hint of structure emerged. The two clearest major paths linking points in the 3-space are shown in Figure 10. Both begin near "wood bark". From here there is an ordered progression through "waves", "wood grain", to "cloth" up to "meshes". Another branch from "bark" ascends through "fieldstone" to "clouds" to "European marble" to "rice paper". There is the suggestion that these two paths might meet, thus completing a loop in the threespace.

To test this prediction, we then selected the new set of thirteen textures, nine taken from our original set plus another three from elsewhere together with one more new sample from Brodatz (see Table 2). Indeed, we were able now to close this loop with the additional textures, and at the same time could clarify the residual paths.

The final result is four major paths –or perhaps more properly transformations in the three-space – roughly lying on a sphere like sections of a great circle.[2] These have been projected onto a plane in Figure 11, and unfolded. The supporting triples (of frequency two or greater) are listed in Table 2 by their Brodatz figure numbers. In addition to the two paths previously mentioned, the data suggest a third rising from "clouds/ice crystals" through "soap" to "plastic bubbles" to "reptile skin" and then into our additional "quilted/cracks" textures, linking up with the "sea fan". The fourth path again emanates from "clouds/fieldstone" to mica and then into 3D structures like "quartz", "beach sand", and "pebbles", perhaps eventually leading to round 3D packed balls such as Brodatz Figures 74 or 66 (which could then collapse back to ice crystals).

One final comment regarding the interpretation of Figure 11. Although the paths have been depicted by texture samples, it may be improper simply to regard this ordering as a space of textures. Rather, the paths reflect how one texture is transformed into another, with each path or segment of a path representing one type of transformation. So, for example, "expanded mica" becomes stretched out to create the linear structure of "wood bark" which is further refined to thin 1D filaments like "fur" or "hair". This texture type then adds orthogonal filaments to create a 2D mesh, etc. Similarly, "expanded mica" can be granularized or pulvarized into ice crystals, or, alternately, have its sharp edges worn smooth to form 3D globules that eventually become a dense packing of pebbles or 3D spheres. Such an interpretation of Figure 11 would suggest, then, that the proper representation for textures is not simply a space of material types, but rather a space of types of transformations. In such a view, the locations of the materials in Figure 11 may be misleading. Rather, the texture classification of materials would be the location of its

[2] There are also other minor paths not shown, such as a loop from reptile skin to soap bubbles to sea fan, and a second loop from leaf veins through rice paper to wood grain. Also not shown are textures D92 Pigskin, D15 Straw and D86 Ceiling tile.

TABLE 1

Color Triples*

TRIPLE			FREQUENCY	EXPT	TRIPLE			FREQUENCY	EXPT
14	26	36	(5)	1	19	18	23	(3)	4
10	14	20	(4)	3	26	36	37	(3)	1
12	5	13	(4)	3	34	36	26	(3)	1
18	19	32	(4)	4	7	14	26	(2)	1
19	32	35	(4,4,3)	1,2,4	8	23	30	(2,2)	1,2
23	**30**	**35**	(4)	2	9	8	23	(2,2)	1,2
30	**35**	**38**	(4,4)	1,2	9	19	32	(2,2)	1,2
31	35	35x	(4)	4	12	2	14	(2)	1
32	35	38	(4)	1	12	4	2	(2)	1
9	8	13	(3,4,3)	1,2,4	13	23	32	(2,2)	1,2
3	10	14	(3)	3	16	2	7	(2)	1
8	13	23	(3,3)	1,2	19	32	30	(2,2)	1,2
11	1	7	(3)	2	30	34	33	(2)	2
13	23	30	(3,3)	1,2	30	34	36	(2)	1
18	**6**	**6x**	(3)	4					

*The color samples are indicated by the numbers in the triple columns, which refer to the LJG locations shown in Figure 8. The parenthetical numbers are the number of times the triple was chosen in the experiments indicated. An "x" following a number indicates extrapolation possible, but the sample was not in the experimental collection. Bold values approximate radial paths and are not plotted in Figure 8. Triples of frequency two are included only from the two experiments that spanned the yellow-blue axis. Some of these less popular triples are indicated as dashed lines on the Figure 8, in order to show evidence for further joins. The data are the intersection of the choices for both JK and WR.

most probable form – say the prototype – together with the extent of its allowable transformations along one of the accepted paths. The texture would thus be represented by a probability distribution along the four paths. For "chunky" mica this distribution seems to have four "arms", whereas for "fur" or regular tesselations such as "reptile skin" or packed "3D pebbles", there may be only two (i.e. the opposite directions of the same path). Perhaps the TM paths in the color space can be viewed in the same light.

7 Conclusion

Trajectory Mapping ("TM") is based on the assumption that objects are categorized by features. The categories become defined by those features which are allowed to vary or be transformed, and those which can not ([3]). Given a domain of objects (and a context), the Trajectory Mapping technique explicitly requires subjects to chose an aspect of a pair of objects that is allowed to vary, thus revealing the parameterizations (and axes) of the relevant feature space (or space of transformations). There is no requirement that the space be Euclidean, homogenous or metric, except along any given feature path. Although the feature attribute or transformation is not made explicit, often it is easy to deduce, as illustrated by our three examples. In this regard, the TM procedure could be augmented by using directed cross-correlation techniques on the sample set to extract the attributes that define any feature parameterization or transformation. That is, of course, assuming an input basis set for the space is known in advance.

References

[1] P. Brodatz. *Textures: a photographic album for artists and designers.* Dover Publications, NY, 1966.

[2] R.W. Burnham. Comparison of color systems with respect to uniform visual spacing. *Jrl. Opt. Soc. Am.*, 39:387-392, 1949.

[3] J. Feldman. Perceptual categories and world regularization. Ph.D. thesis, Dept. Brain & Cog. Sci., Massachusetts Institute of Technology, 1992.

[4] W. Goethe. *Contributions to Optics.* Weimar, 1791. See R. Matthaei, *Goethe's Color Theory*, von Nostrand Reinhold Co., NY, 1971.

[5] L.M. Hurvich and D. Jameson. Some quantitative aspects of an Opponent Colors Theory, *Jrl. Opt. Soc. Am.*, 46:416-421, 1956.

[6] R.A. Johnson and D.W. Wichern. *Multivariate Statistical Analysis.* Prentise Hall, NJ, 1989.

[7] D.B. Judd and G. Wysecki. *Color in Business, Science and Industry*, Wiley, NY, 1975.

[8] J. Kruskal. Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis. *Psychometrika*, 29:1-27, 1964.

[9] D.L. MacAdam. Uniform color scales. *Jrl. Opt. Soc. Am.,* 64:1691-1702, 1974.

[10] D.L. MacAdam. Colorimetric data for samples of OSA uniform color scales. *Jrl. Opt. Soc. Am.*, 68:121-130, 1978.

TABLE 2

Color Triples*

TRIPLE			FREQUENCY	EXPT	TRIPLE			FREQUENCY	EXPT
12	37	68	(9)	1	68	76	87	(3)	1
12	72	68	(6)	1	76	68	93	(3)	1
5	12	72	(5)	1	93	15	108	(3)	1
2	92	12	(4)	2	93	15	110	(3)	1
2	92	3	(4)	1	100	73	111	(3)	2
3	111	73	(4)	1	2	5	12	(2)	1
3	111	87	(4)	1	2	12	72	(2)	1
12	37	72	(4)	1	3	10	12	(2)	2
12	98	23	(4)	1	3	10	76	(2)	2
2	5	23	(3)	1	3	87	12	(2)	1
3	10	87	(3)	2	5	12	37	(2)	1
3	87	12	(3)	2	7	37	70	(2)	1
3	87	111	(3)	1	12	70	93	(2)	2
3	**111**	**76**	(3)	1	**12**	**110**	**100**	(2)	2
3	138	87	(3)	2	23	5	60	(2)	1
3	92	60	(3)	1	37	70	93	(2)	2
3	138	137	(3)	2	60	58	61	(2)	1
5	2	100	(3)	1	60	63	108	(2)	1
12	2	100	(3)	1	68	76	87	(2)	1
12	37	50	(3)	1	63	110	137	(2)	2
12	72	93	(3)	1	68	76	138	(2)	2
15	68	93	(3)	1	73	111	87	(2)	2
28	**73**	**60**	(3)	1	87	139	111	(2)	2
37	50	68	(3)	1	111	87	138	(2)	2

*Numbers in parenthesis indicate the number of times the triple was chosen by either WR or JK, using the greater count. Bold triples were not used in Figure 11. Triples occuring less than two times by both subjects are not included. If a triple occured only twice with one subject, then it is included in the table only if the triple was chosen at least once by the other subject. The triple numbers refer to page numbers in Brodatz or Stevens (#137); or Patterns in the Wild (#138, 139).

[11] National Wildlife Federation. *Patterns in the Wild.* Washington, DC, 1992.

[12] A.R. Rao and G.L. Lohse. Identifying high level features of texture perception. *Comp. Vision, Graphic & Image Proc.*, 55:218-233, 1993.

[13] W. Richards and J.J. Koenderink. Trajectory Mapping: A new scaling technique. European Conference on Visual Perception, Edinburgh, August 1993.

[14] J. Rubin and W. Richards. Color vision: Representing material categories. Chapt. 14 in *Natural Computation*, W. Richards (Ed), MIT Press, Cambridge, MA, 1988.

[15] E. Saund. Putting knowledge into visual shape representation. *Art. Intell.*, 54:71-119.

[16] R.N. Shepard. The analysis of proximities: Multidimensional scaling with an unknown distance function. *Psychometrika*, 27:219-246, 1962.

[17] P.S. Stevens. *Patterns in Nature.* Atlantic Monthly Press, 1974.

[18] W.S. Torgerson. Multidimensional scaling. I. Theory and method. *Psychometrika*, 4:401-419, 1952.

[19] G. Wysecki and W.S. Stiles. *Color Science.* Wiley, NY, 1967.

[20] Wolfram Research. *Mathematica: A System for Doing Mathematics by Computer.* User's Guide for the Macintosh, 1993.