

MATCHING THREE-DIMENSIONAL OBJECTS USING A RELATIONAL PARADIGM

LINDA G. SHAPIRO, JOHN D. MORIARTY*, ROBERT M. HARALICK
and PRASANNA G. MULGAONKAR

Department of Computer Science, Virginia Polytechnic Institute and State University, Blacksburg,
VA 24061, U.S.A.

(Received 7 September 1983; in revised form 11 November 1983; received for publication 1 December 1983)

Abstract—A relational model for describing three-dimensional objects has been designed and implemented. The model, which provides a rough description to be used at the top level of a hierarchy for describing objects, was designed for initial matching attempts on an unknown object. Each description is in terms of the set of simple parts of an object. Simple parts can be sticks (long, thin parts), plates (flat, wide parts) and blobs (parts that have three significant dimensions). The relations include an attribute-value table for global properties of the object, the properties of the simple parts, binary connection and support relationships, ternary connection relationships, parallel relationships, perpendicular relationships and binary constraints.

An important use of the model is to characterize the similarity and differences between three-dimensional objects. Toward this end, we have defined a measure of relational similarity between three-dimensional object models and a measure of feature similarity, based only on Euclidean distance between attribute-value tables. In a series of computer tests, we compare the results of using the two different similarity measures and conclude that the relational similarity is much more powerful than the feature similarity and should be used when grouping the objects in the database for fast access.

Matching Relational models Three-dimensional modeling Similarity measures

I. INTRODUCTION

In scene analysis, we are given one or more views of a three-dimensional scene. As part of the analysis, we must identify the three-dimensional objects using only the two-dimensional views. The data are arrays of numbers representing light intensities or distances or other measurable quantities, depending on the sensor. Noise, distortion and sampling error are common. Segmentations of the data into objects or surfaces or cylinders are far from perfect.

In this kind of environment, it seems reasonable that the first attempts at matching should involve only very rough three-dimensional object models. Exact dimensions and exact geometric specification will not be useful until the analysis procedure has narrowed down the choice of models. Instead, rough models that characterize the structure of the object can be used.

In Section III, we motivate and define our relational model for three-dimensional objects. The object models are intended for use in a scene analysis system. A database of such models has been set up, as described in Section V. However, our first experiments with these models are not concerned with scene analysis. Instead, we look at the more fundamental question, "what makes two three-dimensional objects similar?" Surely,

if we cannot answer this question, we cannot hope to answer the harder question, "to which three-dimensional objects is this two-dimensional projection most similar?" In Section IV, we define a new kind of two-way relational matching to be used to compare two relational descriptions. In Section VI we describe several kinds of experiments comparing two three-dimensional models using (1) only their global features and (2) the entire relational structure.

II. RELATED LITERATURE

We have divided the relevant literature into two categories: three-dimensional object representation and matching.

II.1. Three-dimensional object representation

We have chosen to use relational models to represent three-dimensional objects. In this section, we survey these and other 3D object representations. There are several categories of applications that require modeling of three-dimensional objects. These include mechanical design and manufacturing, computer graphics and computer vision. We will restrict our discussion to the representations used in computer vision. Other representations can be found in Badler and Bajcsy⁽¹⁾ and the proceedings of the Workshop on the Representation of Three-Dimensional Objects.⁽²⁾

Surface-edge-vertex models. All of the early work in vision and much of the present work used

* Currently employed at The Analytical Sciences Corporation.

This research was supported by the National Science Foundation under grants MCS-7923827 and MCS-7919741.

surface-edge-vertex models of three-dimensional objects. Roberts⁽³⁾ model included points, lines and planar surfaces in three-space implemented as a ring structure. The object of his work was to find junction points in a given line drawing that fit a transformation of some stored model. Huffman,⁽⁴⁾ Clowes,⁽⁵⁾ and Waltz⁽⁶⁾ labeled the line segments of a line drawing as corresponding to concave, convex, boundary and, in Waltz's work, crack or shadow edges in three-space. Regions delimited by line segments could be labeled as background or as a surface of one of the objects in the scene. There are no stored object models; what is stored is knowledge, in the form of all two-dimensional junction labelings that can correspond to trihedral blocks world objects.

Analysis by line labeling has been extended to curved surfaces, as in Shapira and Freeman,⁽⁷⁾ Turner⁽⁸⁾ and Chakravarty.⁽⁹⁾ The model that Shapira and Freeman employ allows either quadratic or planar surfaces, edges which are all or part of the intersection of two surfaces and vertices which are the intersection of three or more edges. In their model a boundary is a closed chain of edges, a face is a bounded portion of a surface and a body is a closed connected part of three-space, delimited by a finite number of faces. They used multiple views of objects in their analyses and were able to validate junctions that correspond to real vertices, connect some pairs of junctions by empty lines where no line appeared in the line drawing and create synthetic junctions where the real junction was hidden. Their program finds the faces of each body and the corresponding region in each view.

Chakravarty worked with planar-faced or curved-surface solid bodies having vertices formed by at most three surfaces. Junctions are labeled with respect to the number of regions at the junction, the junction type (based on the arrangement of the lines) and the number of regions associated with the line leaving the junction. Lines are labeled as limb, non-occluding, occluding, partial limb, partial non-occluding, partial occluding and concave. He developed a junction transition graph where a cycle having consistent line labels represents the traversal of a region's boundaries.

Another extension is to the "origami world" where objects are created by folding paper. Origami world objects can have surfaces that are not part of a solid. Kanade⁽¹⁰⁾ extends line labeling to the origami objects where he will usually get several legal labelings per drawing. He then maps geometric properties of the line drawing, such as parallel lines and skewed symmetry, into gradient space where a gradient represents how a plane is slanted relative to the line of view. His problem is to uniquely determine the gradients of the surface of each object and he has succeeded with several simple objects.

Surface-edge-vertex models have also been used by Nagao *et al.*⁽¹¹⁾ whose method was to estimate defects in the two-dimensional description, produce imperfect models from the perfect models and match to the

imperfect models, McKee and Aggarwal⁽¹²⁾ who performed recognition on partial views of known objects, Richard and Hemami⁽¹³⁾ who used Fourier descriptors of the silhouettes of objects stored as wire-frame models and numerous others.

Relational models. Relational or graph models have become very popular since it was discovered^(6,14) that the relational matching problem (described by Barrow *et al.*⁽¹⁵⁾) can be greatly reduced by using relaxation processes. Two recent studies are of particular interest to our work. Chien and Selander⁽¹⁶⁾ use object models that include networks of surface, edge and vertex atoms, each having several properties and connected by surface-edge and edge-vertex arcs. The surfaces in their models may be planar, cylindrical or spherical. A complete object model consists of several networks representing several views. Matching is from an image graph, extracted from the input image into a library of object models. The network matching utilizes a cost function and tries to find a low-cost association that pairs image parts with object parts.

Schneier⁽¹⁷⁾ represents objects by primitives and relations, but with the special feature that common primitives and relations are shared across models and within models. His program produces a scene graph from several views of range data of an object. It tries to find an isomorphism between the scene graph and a structure derived from the graph of models where all three-dimensional models are represented. The matching process utilizes fast indexing: (a) primitives and relation schemata index all models in which they occur; (b) models index all primitives and relation schemata within them. The main advantage is the elimination of the need to match against every one of a library of stored models.

Generalized cylinders. The second major type of three-dimensional model used for computer vision is the generalized cylinder model suggested by Binford⁽¹⁸⁾ and first used with laser range data to produce descriptions of curved objects (Agin and Binford⁽¹⁹⁾). A generalized cylinder is a volume defined by a space curve axis and the cross section function at each point of the axis. In Nevatia's work,⁽²⁰⁾ the three-dimensional models consist of generalized cylinders with normal cross sections for primitives, plus connectivity relations and global properties. Cylinders are described by length of axis, average cross-section width, ratio of the two and cone angle. Global properties of an object include number of pieces, number of elongated pieces and symmetry of the connections. In the matching phase, an indexing scheme is used to access objects that are likely to match an unknown. Each object has a three bit code describing each of its distinguished pieces. Encoded are (1) connectivity (one end or both), (2) type (long or wide) and (3) conical (true or false). Objects with the same code are grouped together and the correct group is found before full matching is started.

Marr and Nishihara⁽²¹⁾ think of objects as stick

figures, where each stick is the axis in one or more generalized cylinders. They advocate hierarchical models; at the top level a hand may be represented by a single cylinder which is broken down further at subsequent levels. In order to describe the connections between cylinders they employ two vectors: \$AXIS, which can be placed along the axis of a cylinder whose connection is to be described, and \$SPASAR, which can be used to describe the rotation of a second cylinder about the first. The relationship of two touching cylinders is described by a triple (p, i, g) , where p is the position at which \$SPASAR attaches to \$AXIS, i is the inclination of \$SPASAR to \$AXIS and g is the girdle angle describing the rotation of \$SPASAR about \$AXIS. If the cylinders do not touch directly, then the description uses the pair (d, e) , where d is the perpendicular distance from \$AXIS to the beginning of \$SPASAR and e is the girdle angle.

Marr and Nishihara also believe in the use of indexing in recognition. They distinguish between indexing clues that can be used before there is a guess at the three-dimensional configuration (for example, connectivity and some length comparisons) and those that cannot. Their matching scheme uses relaxation to rotate the model into the appropriate view to match the description obtained from the two-dimensional image.

Hollerbach⁽²²⁾ used generalized cylinders in his hierarchical models of pottery vases. His model of a vase is a main cylinder segmented into possible parts: foot, body, neck and lip. Parts can be described by a general shape (i.e. ovoid) plus modifiers (i.e. protrusions, size, position). Soroka⁽²³⁾ used generalized cylinders with elliptical cross sections to model three-dimensional biological data obtained from tomographic data. In other recent work, Agin⁽²⁴⁾ has developed a new system where objects are modeled by generalized cylinders and arbitrary spatial relationships. Relationships include snakes (several cylinders grouped along a single axis), attachment points and arbitrary transforms. Users can code S-expression descriptions such as (CUBE2 (ATTACH CUBE2 TOP) CUBE 1) to describe objects to the system.

General knowledge models. The models discussed so far have specific primitives (surface, edge, vertex or generalized cylinder), some description of the properties of those primitives and often some kind of connection relation. Several more general models have been proposed. Minsky⁽²⁵⁾ has defined a "frame" as a data-structure for representing a stereotyped situation. A frame is like a network of nodes and relations where the top levels are fixed and represent things that are always fixed about the situation and the lower levels have slots that must be filled with specific information. For a stereotyped scene of three-dimensional objects, Minsky's model is a set of frames describing the scene from different viewpoints plus the transformations between pairs of these frames representing the effect of moving the camera. A related model has been pro-

posed by Ballard *et al.*⁽²⁶⁾ Their model is a semantic network, where nodes represent primitive and complex objects and concepts such as assertions or procedures. Given this model, an image and a query pertaining to a particular object, their system would construct a sketch map (an instantiation of part of the model that matches the scene) and use it to answer the query.

II.2. Matching and constraint satisfaction

Our three-dimensional models are relational structures. Relational matching, the process of finding relational homomorphisms between two structures is an NP-complete problem; in the worst case, its behavior is expected to be exponential. However, it has been shown that the use of look-ahead or relaxation operators can speed up the tree search used for finding a match. Since our relational matching will require some form of relaxation, we will survey some of the recent work in this important area.

Discrete relaxation. In Haralick and Shapiro⁽²⁷⁾ we defined a general network constraint analysis problem, called the *consistent labeling problem*, which was a generalization of specific problems from several different specialty areas. In the general problem, we are given a compatibility model (U, L, T, R) where $U = \{1, \dots, M\}$ is a set of M objects called *units*, L is a set of names for the units called *labels*, $T \subset U^{**}N$ specifies N -tuples of units that constrain one another and $R \subseteq (U \times L)^{**}N$ specifies N -tuples of unit-label pairs $((u1, l1), (u2, l2), \dots, (uN, lN))$, where unit $u1$ can have label $l1$, unit $u2$ can have label $l2, \dots$, and unit uN can have label lN , all at the same time. A labeling of U is a mapping $f: U \rightarrow L$ that assigns a label to each unit. The consistent labeling problem is to find all labelings f that satisfy $(u1, \dots, uN) \in T$ implies $(u1, f(u1), \dots, uN, f(uN)) \in R$. We have shown that the relational homomorphism problem is a consistent labeling problem.

Consistent labeling problems have traditionally been attacked by a depth first search where the search procedure assigns labels to units as long as it can find a label for each new unit that is compatible according to R with the labels already fixed to previous units. Whenever the procedure cannot find a label for a new unit, it backtracks. Such a procedure suffers from thrashing; a poor choice of labels for one of the first units can cause failure of all paths stemming from that choice.

Ullman⁽¹⁴⁾ first tried to avert this thrashing behavior in a matching application. Waltz⁽⁶⁾ popularized discrete relaxation by using it in a program to label the edges of a line drawing as concave, convex, boundary, shadow or crack. His 'filtering' program was applied prior to the tree search and it removed so many possible labelings that frequently the tree search became unnecessary. Rosenfeld *et al.*⁽²⁸⁾ formalized the relaxation operator used by Waltz. Using our consistent labeling notation, U is the set of edges, L is the set of edge labels, $T = \{(\text{line } 1, \text{line } 2) \mid \text{line } 1 \text{ connects to line } 2\}$ and $R = \{((\text{line } 1, \text{label } 1), (\text{line } 2, \text{label } 2)) \mid$

(line 1, line 2) $\in T$ and there is some physically possible labeling of the junction where line 1 and line 2 meet in which line 1 can take label 1 while line 2 takes label 2}.

In the Rosenfeld *et al.*⁽²⁸⁾ formalism a labeling is an N -tuple of sets $Lk = (L(1, k), \dots, L(M, k))$, where $L(i, k)$ is a set of labels that are still allowed for line i at iteration k . $L(i, 0)$ is the set {label | ((line i , label), (line j , label j)) $\in R$ for some j }. $L(k + 1)$ is obtained from $L(k)$ by discarding from each $L(i, k)$ any label l such that there exists a j with $\{((line\ i, l), (line\ j, lj)) \in r \mid lj \in L(j, k)\} = \Phi$. The relaxation procedure obtains $L(1)$ from $L(0)$, $L(2)$ from $L(1)$, ..., until some $L(k + 1) = L(k)$, whereupon it halts. If $L(k) = (\Phi, \Phi, \dots, \Phi)$, there is no legal labeling; if $L(k)$ is single-valued, the single consistent labeling has been found; and otherwise a tree search in a reduced tree is necessary.

Discrete relaxation operators have also been proposed in Haralick and Shapiro⁽²⁷⁾ and by Ullman,⁽²⁹⁾ Montanari,⁽³⁰⁾ Haralick and Kartus,⁽³¹⁾ Mackworth,⁽³²⁾ Freuder,⁽³³⁾ Gaschnig,⁽³⁴⁾ Davis,⁽³⁵⁾ and others. In a recent paper, Haralick and Elliot⁽³⁶⁾ compared several discrete relaxation operators by constructing random binary compatibility models on which to test the operators. It was found that a very simple operator that they call *forward checking* performed best, i.e. it had fewer operations and smaller execution time. The more powerful operators searched less nodes of the tree than forward checking, but took many more operations to do so. Haralick and Elliot also found that a strategy of ordering the units, by always taking the next unit having fewest possible labels left, tended to cause less backtracking to occur and thus reduced search time.

Inexact discrete relaxation. In our past work we developed a structural model for describing two-dimensional shapes and a corresponding procedure for matching an unknown shape to a stored model.⁽³⁷⁾ Since the unknown shapes could be distorted or noisy, their decompositions into simple pieces and intrusions and the corresponding relational descriptions were generally not identical to the stored models. To deal with this problem, we define an inexact match or an ϵ -consistent labeling as a function $f: U \rightarrow L$ that satisfies

$$\sum_{\substack{t \in T \\ f(t) \neq R}} w(t) \leq \epsilon$$

where w is a function assigning a weight to each N -tuple t in T . In subsequent work,⁽³⁸⁾ we further generalized the concept of an inexact match and developed relaxation operators for inexact matching corresponding to the forward checking and lookahead-by-one operators used in exact matching. We found again that the forward checking operator used less operations and less time than the lookahead-by-one or the backtracking tree search.

Continuous relaxation. Rosenfeld *et al.*⁽²⁸⁾ proposed a continuous version of the binary consistent labeling problem. In the continuous version, each possible label l for unit i has an attached weight ($pi(l)$ indicating the

certainty with which label l is attached to unit i). The weights for each label are between 0 and 1 and the sum of the weights of all labels of a given unit must be 1. The constraint relation R also has a weight attached to each tuple. In their notation this amounts to a set of coefficients $\{rij, i, j, \in U\}$ where $rij(l, l')$ denotes the compatibility of label l on unit i with label l' on unit j . The rij 's range from -1 to 1 . The job of the relaxation operator is (a) to increase $pi(l)$, if other units' labels that have high weights are highly compatible with l at unit i , and (b) to decrease $pi(l)$, if other highly weighted labels are incompatible with l at unit i . The relaxation operator that updates the pi 's at iteration $k + 1$ is given by

$$pi\langle k + 1 \rangle(l) = \frac{pi\langle k \rangle(l)[1 + qi\langle k \rangle(l)]}{\sum_i pi\langle k \rangle(l)[1 + qi\langle k \rangle(l)]}$$

where

$$qi\langle k \rangle(l) = \sum_j dij \sum_{l'} rij(l, l') pj\langle k \rangle(l')$$

and the dij 's are coefficients that weight the total interaction between units i and j .

Continuous relaxation has been used by Hanson and Riseman⁽³⁹⁾ for edge enhancement, by Barrow and Tenebaum⁽⁴⁰⁾ for scene analysis, by Zucker and Hummel⁽⁴¹⁾ for clustering, and by others. Current work in the area deals with theoretical analyses of the continuous relaxation process to determine exactly what problem it is actually solving. To this end, Zucker *et al.*^(42,43) have shown that under certain restrictions, continuous relaxation is equivalent to local maxima selection. Haralick *et al.*⁽⁴⁴⁾ have derived necessary conditions for a fixed point. Faugeras *et al.*^(45,46) have also been active in both theory and practice in the area of continuous relaxation. Haralick⁽⁴⁷⁾ has shown under what conditions probabilistic relaxation results in probabilities which can be given a Bayesian interpretation.

III. A RELATIONAL MODEL

In this section we first describe a relational model that provides a rough description of the structure of three-dimensional objects. This model is to be used at the top level of a hierarchy for describing objects. Lower levels will be more precise descriptions, including finer details. The rough descriptions will be used for initial matching attempts and as input to a clustering procedure that will group similar objects together.

The parts of an object: sticks, plates and blobs. The objects that we work with are complex man-made objects, such as office furniture and industrial manufacturing parts. These objects are physically built from parts. The parts can have flat or curved surfaces and they exist in a large variety. Instead of trying to describe each of these many parts, at the top level we classify each part as either a *stick*, a *plate* or a *blob*.




Fig. 1. Illustrates several examples each of sticks, plates and blobs.

Sticks are long, thin parts that have only one significant dimension. Plates are flatish, wide parts with two nearly flat surfaces connected by a thin edge between them. Plates have two significant dimensions. Blobs are parts that have all three significant dimensions. All three kinds of parts are “near-convex”, i.e. a stick cannot bend very much, the surfaces of a plate cannot fold very much and a blob can be bumpy, but cannot have large concavities. Figure 1 shows several examples of sticks, plates, and blobs.

Because we wish to analyze the structure of objects, we need to define sticks, plates and blobs more precisely. Formally, a stick is a 4-tuple $ST = (En, I, Cm, L)$, where En is the set of two end points of the stick, I is the set of interior points of the stick, Cm is its center of mass and L is its length. Since straight line segments have each of the components of a stick, we will be able to informally represent all sticks by straight line segments to simplify our thinking about them.

A plate is a 4-tuple $PL = (Eg, S, Cm, A)$, where Eg is the set of edge points, $S = \{S1, S2\}$ is the set of surface points of the plate, partitioned into the two surfaces, Cm is the center of mass and A is the area. Again, to simplify analyses, we can informally represent all plates by circles.

A blob is a triple $BL = (S, Cm, V)$, where S is the set of surface points, Cm is the center of mass and V is the volume of the blob. We can informally represent all

blobs as spheres. We choose line segments, circles and spheres because they have no corners that we might be tempted to use in our descriptions. At the top level, the descriptions are to be as general and as rough as possible.

Constraints on assembling the parts. Our three-dimensional models must describe how the sticks, plates and blobs are put together. These descriptions will also be rough; they cannot specify the physical points where two parts join. The stick has two logical end points, a logical set of interior points and a logical center of mass that can be specified as connection points. The plate has a set of edge points, a set of surface points and a center of mass. The blob has a set of surface points and a center of mass. We will now discuss how to use such minimal information in the models.

Binary connections. Clearly the connections between pairs of parts are an integral part of any three-dimensional relational model. We specify a connection between two parts by specifying (1) the type of connection and (2) the constrained angles of the connection. The type of connection describes which distinguished entity of the first part touches which distinguished entity of the second part. Thus possible connection types are end-end, end-interior, end-center, end-edge and so on. For each type of connection, there is a corresponding set of angles which, when specified as single values or as ranges,




Fig. 2. Illustrates three examples of the constrained connections of two simple parts.

constrain the binary connection further. For example, when two sticks join end-end or interior-center (as illustrated in Fig. 2), a single angle constrains their connection. Figure 2 specifies the full range of this angle. If the angle is restricted to a single value, the connection is restricted to an exact form. If the angle is specified as an allowable range of values, then the form of connection is more flexible. At most, three angle ranges are required to uniquely describe a binary connection between two arbitrary parts.


Ternary connections. The binary connections are not sufficient to entirely describe a three-dimensional model, since they do not place any global constraints on the resulting object. For example, two sticks each connected end-end to the same third stick might be at the same or opposite ends and, if at the same end, might coincide in space or not coincide, but might still be at the same angle with respect to the third stick. If we were trying to describe the object very precisely, we might need to specify *N*-ary connections for arbitrary *N*. However, we have determined that we can add powerful constraints to the model by considering triples of simple parts. Since at this level, our descriptions are rough anyway, we will only go as far as ternary relations to describe connections.

Let (*s*1, *s*2, *s*3) be a triple of simple parts satisfying that *s*1 and *s*3 both touch *s*2. The description of the


spatial relationship between *s*1 and *s*3 with respect to *s*2 has two components. The first component specifies whether *s*1 and *s*3 meet *s*2 on the same end (or surface). The second component constrains the angle subtended by the centers of mass of *s*1 and *s*3 at the center of mass of *s*2. This angle constraint can also be a single value or an allowable range. Figure 3 illustrates several connections between the parts and the full range of the angles to be specified.

Support structure. Another important aspect of a multi-part three-dimensional object is its support structure. The legs of a chair not only touch the seat, but they also support the seat. The support structure is related to the function of the object and its parts. For example, objects that have four upright sticks supported by the ground and supporting a horizontally-oriented plate tend to be stationary objects and tend to be used to set another object (book, vase, person) on. In addition to its importance in the three-dimensional description, the support structure can be useful in helping to identify two-dimensional perspective projections of an object, since much of the support structure of an object is often evident in a right-side-up two-dimensional view. Thus the support structure seems to be an essential component in a three-dimensional object model.


Additional constraints. The binary and ternary con-



PARTS	SIDE	ANGLE CONSTRAINT
(b1, b2, b3)	OPPOSITE	$90^\circ < \theta < 180^\circ$



PARTS	SIDE	ANGLE CONSTRAINT
(b1, b2, b3)	OPPOSITE	$0^\circ < \theta < 90^\circ$



PARTS	SIDE	ANGLE CONSTRAINT
(b1, b2, b3)	SAME	$0^\circ < \theta < 90^\circ$
(b3, b2, b4)	OPPOSITE	$90^\circ < \phi < 180^\circ$

Fig. 3. Illustrates three examples of constrained connections among three simple parts.


nection relations plus the support structure provide a basic description of the structure of an object. Additional constraints on groups of parts that do not touch each other may also be necessary to completely characterize some objects. Consider the typical table shown in Fig. 4. The four sticks are all the same length, all parallel and their centers of mass are equidistant from the center of mass of the plate. Furthermore, there are constraints on the angles between adjacent pairs of line segments from stick centers of mass to plate center of mass. These constraints rule out certain three-dimensional objects that are not typical tables, although they may have four sticks connected to a plate. The constraints are illustrated in Fig. 4. For a given class of objects and a given application, there will be a set of relations that are important additional constraints.

The relational data structure. In Shapiro and Haralick⁽⁴⁸⁾ we discussed a general relational data structure. The structure, which has also been called a spatial data structure or a structural description, can be formally defined as follows. A relational data structure

D is a set $D = \{R_1, \dots, R_K\}$ of relations. For each relation $R_k, k = 1, \dots, K$, there is a positive integer N_k and a sequence of domain sets $S(1, k), \dots, S(N_k, k)$ such that $R_k \subseteq S(1, k) \times \dots \times S(N_k, k)$. The elements of the domain sets may themselves be atoms (nondecomposable) or relational data structures. In most relational data structures, one of the relations is an attribute-value table (A/V) and contains the values of global properties of the object being represented by the structure.

The relational data structure for a three-dimensional object will consist of an attribute-value table plus nine other relations. The unary SIMPLE PARTS relation is a list of the parts of the object. Each part is represented by a relational data structure consisting of an attribute-value table. The attributes of a simple part consist of TYPE (stick, plate or blob), RELATIVE LENGTH, RELATIVE AREA and RELATIVE VOLUME. The length, area and volume values may be real numbers or may be marked "don't-care" when they are unimportant or inappropriate.

The CONNECTS/SUPPORTS relation contains



Constraints

- parallel (S1,S2,S3,S4)
- equi-length (S1,S2,S3,S4)
- equi-length (L1,L2,L3,L4)
- equi-angle (α, α')
- equi-angle (β, β')

Fig. 4. Illustrates a standard table made of 4 sticks and a plate and some of the constraints on how the parts fit together.

some of the most important information on the structure of the object. It consists of 10-tuples of the form (s1, s2, SUPPORTS, HOW, v11, vh1, v12, vh2, v13, vh3). The components s1 and s2 are simple parts, SUPPORTS is true if s1 supports s2 and false otherwise and HOW describes the connection type of s1 and s2. The values in the HOW field are elements of the set {end-end, end-interior, end-center, end-edge, interior-center, center-center} where 'end' refers to an end of a stick, 'interior' refers to the interior of a stick or surface of a plate or blob, 'edge' refers to the edge of a plate and 'center' refers to the center of mass of any part. The field pairs (v11, vh1), (v12, vh2), and (v13, vh3) hold the low and high values for the allowed angle ranges for the (at most) three angles that can be specified for a binary connection.

The other eight relations express constraints. The TRIPLE CONSTRAINT relation has 6-tuples of the form (s1, s2, s3, SAME, vl, vh) where simple part s2 touches both s1 and s3, SAME is true if s1 and s3 touch s2 on the same end (or surface) of s2 and false otherwise and vl and vh specify the permissible low and high values for the constrained angle as shown in Fig. 3. The

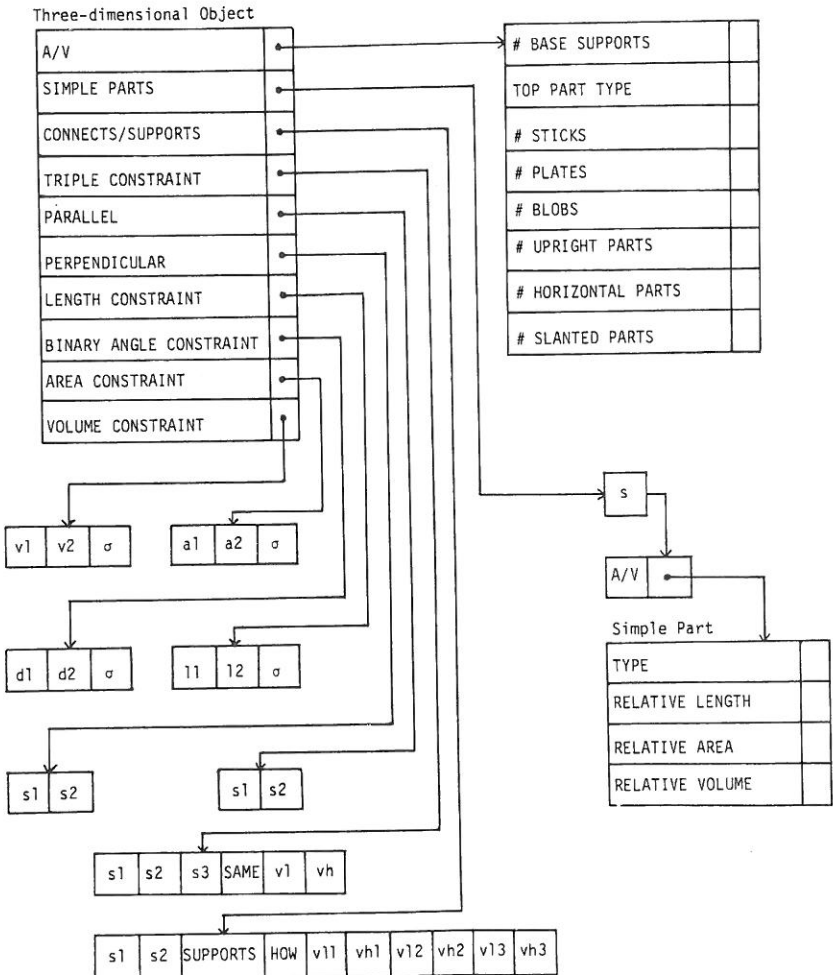


Fig. 5. Illustrates the logical structure of a relational data structure for a three-dimensional object model.

PARALLEL relation and the PERPENDICULAR relation have pairs of the form $(s1, s2)$ where simple parts $s1$ and $s2$ are parallel (or perpendicular) in the model. The LENGTH CONSTRAINT relation has triples of the form $(l1, l2, o)$ where $l1$ and $l2$ refer to specific line segments or part lengths and $o \in \{>, \geq, <, \leq, =, \neq\}$. The components $l1$ and $l2$ of this relation must be powerful enough to express such concepts as the line segment joining the centers of mass of two parts or the length of a stick. The BINARY ANGLE CONSTRAINT relation has triples of the form $(d1, d2, o)$ where $d1$ and $d2$ specify angles and again $o \in \{>, \geq, <, \leq, =, \neq\}$. Finally, the AREA and VOLUME relations have triples of the form $(a1, a2, o)$ and $(v1, v2, o)$, respectively, where $a1$ and $a2$ refer to areas, $v1$ and $v2$ refer to volumes and $o \in \{>, \geq, <, \leq, =, \neq\}$.

Note that specification of areas and volumes in the constraint relations may be redundant due to the fact that each part has a RELATIVE AREA and RELATIVE VOLUME attribute. The constraint relations will only be used when the constraint is meant to be emphasized as important to the object. It is expected that the inexact matching process to be used will be quite lenient about relative area and volume requirements in general. (See Shapiro and Haralick⁽³⁸⁾ for the definition of inexact matching and for some fast algorithms to do it.) When a few constraints on area or volume need to be more severe, the thresholds that the constraint relations have to satisfy can be increased. Similarly, the PERPENDICULAR relation is redundant and will only be used for emphasis or severe constraints.

The attribute-value table of a three-dimensional object contains its global properties. Our intention is to include as many properties as possible while

keeping the description at a gross level. The set of attributes currently planned are number of base supports, type of topmost part, number of sticks, number of plates, number of blobs, number of upright parts, number of horizontal parts and number of slanted parts, number of support levels and position of topmost part. The logical structure of a three-dimensional object is illustrated in Fig. 5. Notice that most of the information in the relational structure is invariant with respect to orientation of the object. Only the SUPPORTS field of the CONNECTS/SUPPORTS relation and the attributes that mention position (upright, horizontal, slanted) or support are related to orientation. These were included in the model because we expect to analyze scenes where objects are usually in their normal upright position. When this is not the case, these attributes can simply be ignored.

IV. RELATIONAL MATCHING

Going from relational descriptions to three-dimensional objects is essentially a matching problem: matching the relational description from the image to the relational description of the object. Doing this matching in a database of one hundred or one thousand separate objects would be computationally expensive. Given an unknown object (three-dimensional object or two-dimensional view), we do not want to compare its description with every object in the database. A solution to this problem which has been used by Salton⁽⁴⁹⁾ in an information storage and retrieval system is to cluster the objects in the database, represent each cluster by a profile description and compare the description of an unknown object only to each profile description. If the unknown object is




Fig. 6. Illustrates two similar chairs and their binary connection relations. Two shapes match when their structural error is sufficiently low.




Fig. 7. Illustrates our 57 objects whose attribute-value tables are in the database.




Fig. 7 continued

judged sufficiently close to one or more clusters, then it is compared only to objects in those groups.

Comparing relational descriptions. Suppose we are given two relational descriptions $D = \{R_1, R_2, \dots, R_K\}$ and $D' = \{R'_1, R'_2, \dots, R'_K\}$, where for each $k =$

$1, \dots, K, R_k \subseteq S(1, k) \times \dots \times S(Nk, k)$ and $R'_k \subseteq S'(1, k) \times \dots \times S'(Nk, k)$. Intuitively, description D is similar to description D' if relation R_k is similar to relation R'_k for $k = 1, \dots, K$. Thus to measure the distance between two relational descriptions, we must

first be able to measure the distance between two relations.

Let $R \subseteq S^N$ and $R' \subseteq T^N$ be two N -ary relations and let f be a binary relation $f \subseteq S \times T$ that associates an element of S with an element of T . We will define a measure of the error of the association f . We define the composition $R \circ f$ of N -ary relation R with binary relation f by

$$R \circ f = \{(t_1, \dots, t_N) \in T^N \mid \text{there exists } (s_1, \dots, s_N) \in R \text{ with } (s_n, t_n) \in f \text{ for } n = 1, \dots, N\}.$$

There are four sets of N -tuple that can be used to describe the error of the association f .

- (1) $R \circ f - R'$. This set consists of N -tuples that arise when an N -tuple of relation R is transformed by f to an N -tuple of T^N , but this new N -tuple is not a part of R' .
- (2) $R' \circ f^{-1} - R$. This set consists of N -tuples that arise when an N -tuple of relation R' is transformed by f^{-1} to an N -tuple of S^N , but this new N -tuple is not a part of the relation R . This set is the symmetric equivalent of set (1) and is used here because we are interested in two-way matching.
- (3) $R - R' \circ f^{-1}$. This set consists of N -tuples of R that are not included in the group of N -tuples obtained by applying f^{-1} to each N -tuple of R' .
- (4) $R' - R \circ f$. This set consists of N -tuples of R' that are not included in the group of N -tuples obtained by applying f to each N -tuple of R .

Example. Consider the two chairs C and C' shown in Fig. 6 and the corresponding simplified binary connection relations

$$R = \{(1, 2), (2, 3), (2, 4), (2, 5), (2, 6)\}$$

and

$$R' = \{(A, B), (B, C), (B, D), (B, E), (B, F), \\ (C, G), (D, G), (E, H), (F, H)\}.$$

Suppose we wish to measure the error of the association f given by

$$f = \{(1, A), (2, B), (3, C), (4, G), (6, F)\}.$$

Then the two compositions are given by

$$R \circ f = \{(A, B), (B, C), (B, G), (B, F)\}$$

and

$$R' \circ f^{-1} = \{(1, 2), (2, 3), (2, 6), (3, 4)\},$$

and the four sets of interest are

	Set	Number elements
$R \circ f - R'$	$\{(B, G)\}$	1
$R' \circ f^{-1} - R$	$\{(3, 4)\}$	1
$R - R' \circ f^{-1}$	$\{(2, 4), (2, 5)\}$	2
$R' - R \circ f$	$\{(B, D), (B, E), (C, G), \\ (D, G), (E, H), (F, H)\}$	6

One method of defining the error of a mapping f is by a weighted sum of the number of elements in each of the

four sets. In particular, we define the *structural error* to be

$$E_s(f) = |R \circ f - R'| + |R' \circ f^{-1} - R|,$$

the *completeness error* to be

$$E_c(f) = |R - R' \circ f^{-1}| + |R' - R \circ f|,$$

and the *total error* to be

$$E_{R,R}(f) = c_1 E_s(f) + c_2 E_c(f)$$

where c_1 and c_2 are non-negative constants.

The total error will be 0 when R' is an isomorphic image of R and f is the isomorphism; and if normalized by dividing by a factor proportional to $|R| + |R'|$ it will be 1 in the worst possible case when $R \circ f \cap R' = R' \circ f^{-1} \cap R = \Phi$. It has the advantage of simplicity and the disadvantage of counting all N -tuples of a relation equally when some relationships may be more important than others.

Once such a mapping error has been defined, we can define the general distance $GD(D, D')$ of two relational descriptions D and D' by

$$GD(D, D') = \min_f \sum_{k=1}^K w_k E_{R_k, R_k}(f)$$

where w_k is the weight assigned to relation R_k and the minimization is taken over all binary relations f associating elements of S with elements of T .

Use of global attributes. The global attributes of a three-dimensional object are stored in the attribute-value table $A/V = \{(a, v) \mid a \text{ is an attribute and } v \text{ is its value}\}$. The attribute-value table is essentially a feature vector and by itself cannot fully describe an object. Yet it is an important aspect of the total description. A human, when asked to describe a chair might answer, "it is an object having four legs, a back and a seat. The legs are long, thin and vertically oriented, the seat is flat, wide and horizontally oriented. The legs connect to and support the seat which connects to and supports the back." Note that in this description, it is very natural to mention the parts and their features before coming to the relational structure. Similarly, the human, when asked the difference between a chair and a table, might say, "the table has no back." Here the presence or absence of a part is important. This all suggests that when comparing two objects, we should first compare their attribute-value tables and only if these are judged similar enough should we continue with the full relational matching.

V. THE EXPERIMENTAL DATABASE SYSTEM

Organization. Each object model is accessed by a unique integer. An object model consists of an attribute-value table plus five relations: SIMPLE PARTS, CONNECTS/SUPPORTS, TRIPLES, PARALLEL and PERPENDICULAR. The relations contain the fields described in Section III, but the angle and size data has been omitted at this stage of

Cluster	Objects
1	1, 2, 3, 4, 5, 6, 7, 8, 23, 24, 25, 26, 28, 30, 34, 35, 36, 38, 39, 40, 42, 45, 48, 52
2	9, 10, 11, 13, 14, 21, 22, 33, 42, 44, 46, 47, 49, 51, 53
3	12, 15, 16, 17, 18
4	19
5	20, 37, 47
6	23, 24, 25, 26, 27, 29, 31, 40, 41, 52, 54, 55, 56, 57
7	32
8	43
9	50

Fig. 8. Gives the clusters of the 57 objects of Fig. 7.

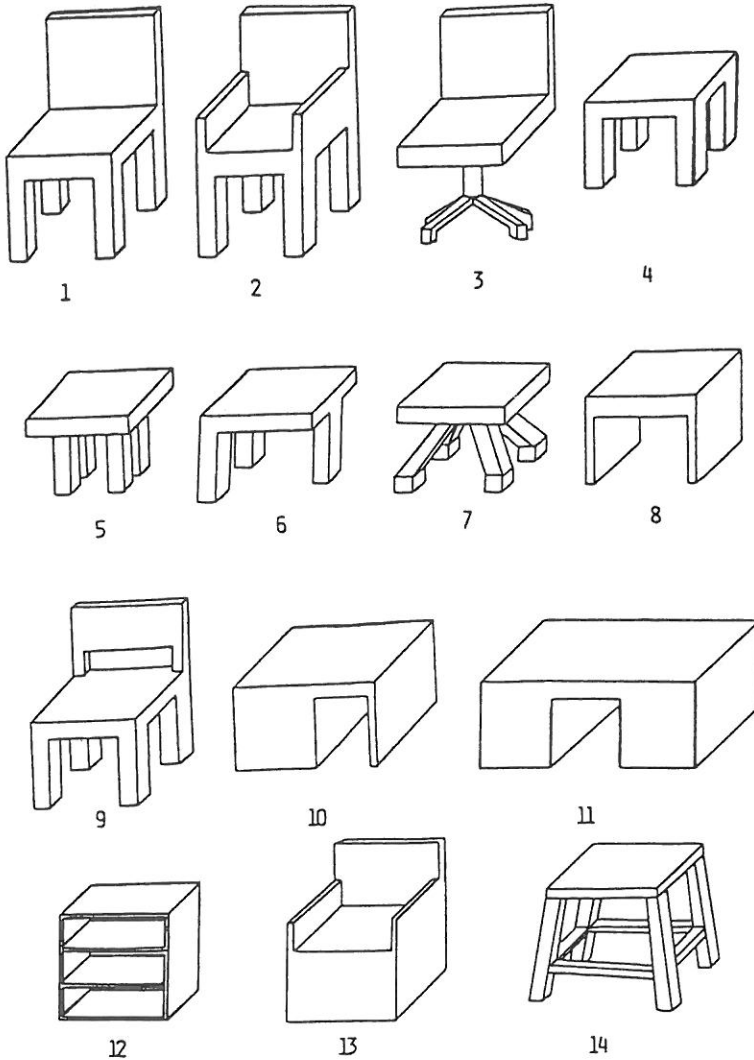


Fig. 9. Illustrates the 14 objects whose full relational models are in the database.

experimentation and the supports component for a pair of parts (i, j) is true if i supports j or j supports i and false if no support is involved. Thus, in the objects on which the experiments were performed, we have five relations $R1, R2, R3, R4$ and $R5$ where

$$R1 = \text{SIMPLE-PARTS} \subseteq \text{parts} \\ \times \{\text{stick, plate, blob}\},$$

$$R2 = \text{CONNECTS-SUPPORTS} \subseteq \text{parts} \\ \times \text{parts} \times \{(\text{supports, how})\} \\ \text{supports} \in \{\text{true false}\} \text{ and} \\ \text{how} \in \{\text{end to end, end to edge,} \\ \text{end to interior, edge to} \\ \text{interior, interior to} \\ \text{interior}\},$$

$$R3 = \text{TRIPLES} \subseteq \text{parts} \times \text{parts} \times \text{parts} \\ \times \{\text{same, opposite}\},$$

$$R4 = \text{PARALLEL} \subseteq \text{parts} \times \text{parts, and}$$

$$R5 = \text{PERPENDICULAR} \subseteq \text{parts} \times \text{parts.}$$

The attribute-value tables of 57 objects have been entered into the database so far. The full 5-relation structure has only been entered for 14 objects, due to the difficulty of constructing the models by hand. The 57 objects are illustrated in Fig. 7.

The database is organized as a set of possibly overlapping clusters of similar objects. In the current set-up, only the attribute-value tables of each pair of objects were used in judging their similarity. With the attributes "number of upright pieces", "number of horizontal pieces" and "number of slanted pieces" weighted by 1 and the other seven attributes weighted by 10, the Euclidean distance between each pair of attribute-value tables was determined. A binary relation, including each pair of objects whose Euclidean distance was below a given threshold, was constructed. The relation was the input to the graph-theoretic clustering procedure which has been described in Shapiro and Haralick.⁽⁵⁰⁾ The set of resulting clusters currently being used in the database are shown in Fig. 8.

In the standard mode of operation, an unknown object may be entered into the database or merely compared to some of the models without being entered. For object entry, the attribute-value table and five relations are input by the user, the attribute-value table is compared to the centroid attribute-value tables of each cluster and the object is added to those clusters that it is most similar to. For matching, instead of being added, the unknown object is compared to each object in the best clusters and the results displayed to the experimenter.

Matching. The relational matching is performed using a treearch with lookahead. Let $U = \{S_1, S_2, S_3, S_4, S_5\}$ be the unknown object, represented by its five relations and $M = \{T_1, T_2, T_3, T_4, T_5\}$ be the model. For a given association f , the total structural error is given by

$$E_s(f) = \sum_{i=1}^5 (\#(Si \text{ of } - Ti) + \#(Ti \text{ of } f^{-1} - Si)) \quad (1)$$

and the total completeness error is given by

$$E_c(f) = \sum_{i=1}^5 (\#(Ti - Si \text{ of } f) + \#(Si - Ti \text{ of } f^{-1})) \quad (2)$$

where $\#$ denotes cardinality. In our matching experiments, we used a combined, weighted measure of total error

$$E(f) = 4 * E_s(f) + E_c(f), \quad (3)$$

which stems from our intuitive feelings that structural error is more important than completeness error. The goal of a matching experiment between two objects U and M with parts $P(U)$ and $P(M)$, respectively, is to find that association $f \subseteq P(U) \times P(M)$ with minimum total error. In the experiments reported in this paper, we restricted the mapping f to being single-valued and one-one, to reduce search time.

Find the best association is achieved with the help of two tree searches. Treearch 1, the "super-quick" search, follows only one path from the root of the tree to the bottom. At each level, it chooses that pair (p, p') , $p \in P(U)$, $p' \in P(M)$, with least accumulated error in the lookahead tables and performs forward checking^(36,38) with respect to the new pair and the so-far-uninstantiated parts. The forward checking operation updates the lookahead tables and determines if this pair can be instantiated. If so, it is added to the association being constructed.

The association f obtained from Treearch I falls into one of three categories:

- (1) exact match: $E_s(f) = 0$ and $E_c(f) = 0$;
- (2) subset match: $E_s(f) = 0$ and $E_c(f) \neq 0$;
- (3) approximate match: $E_s(f) \neq 0$.

In case (1), clearly the best association has already been found. In case (2), one object is contained in the other and the similarity of their attribute-value tables guarantees that not too many parts are missing. In case (3), there is no guarantee that f has minimal error and Treearch II is called.

Treearch II is similar to a branch and bound search using forward checking. Its job is to find an association g such that:

- (1) $E(g) \leq E(f)$
[g 's total error is not greater than f 's];
- (2) $\#proj_1(g) \geq t * (\#P(U))$
 $\#proj_2(g) \geq t * (\#P(M))$

[The projection of g onto its first (second, respectively) coordinate gives a set whose cardinality is at least the percentage specified by parameter t of the number of parts in U (respectively, M).];

- (3) $E(g') \leq k * \sum_{i=1}^5 (\#Si + \#Ti)$, $g' \subseteq g$.

Table 1. Distance matrix for the objects of Fig. 9

	1	2	3	4	5	6	7	8	9	10	11	12	13	14
1	∅	3.46	5.20	5.57	5.57	7.21	6.24	9.64	5.66	9.70	10.49	11.66	10.15	8.66
2		∅	6.24	6.56	6.56	8	7.81	8.72	6.32	10.30	11.26	11.31	9.75	9.33
3			∅	7.48	7.48	8.66	7.21	10.68	4.36	10.72	11.66	11.70	11.22	8.94
4				∅	∅	4.58	2.83	9.01	7.94	9.11	9.17	12.04	10.58	6.63
5					∅	4.58	2.83	9.01	7.94	9.11	9.17	12.04	10.58	6.63
6						∅	5.20	7.81	9.17	7.87	7.94	11.14	9.54	8.06
7							∅	9.27	9	9.22	9.17	12.21	10.86	7.21
8								∅	11.18	4.58	6.48	7.94	5.66	11.22
9									∅	11.22	12.12	12.17	11.62	8.19
10										∅	4.58	9.17	5.74	11.27
11											∅	10.25	7.35	11.31
12												∅	8.66	13.53
13													∅	12.49
14														∅

At each stage of Treesearch II, the partial association g' must satisfy the requirement that its total error is not greater than the percentage specified by parameter k of the sum of the number of N -tuple in all of the relations involved. The parameter k allows the user to control the size of the tree searched at the risk of not finding a best mapping whose error is very high near the top of the tree and very low near the bottom.

VI. EXPERIMENTS

We have run several kinds of experiments using the database of relational models. The purpose of these experiments was to study the relationship between the Euclidean distance between a pair of objects obtained only from their attribute-value tables and the total relational error obtained from the tree search. Figure 9 illustrates the fourteen objects whose relational descriptions and attribute-value tables were used in these experiments.

Table 1 gives the Euclidean distances for each pair of the fourteen objects of Fig. 9. Notice that, as far as the grouping in the database which was obtained using a distance threshold of 8, objects 1-7 and 9 fell into cluster 1, objects 8, 10, 11 and 13 fell into cluster 2, object 12 fell into cluster 3 and object 14 fell into cluster 4.

Table 2 gives the structural and completeness errors for each pair of the same fourteen objects as obtained from the full matching process—Treesearch I, followed by Treesearch II if necessary. Table 2 also gives the k -parameter used in Treesearch II for those matches where Treesearch II was required. An asterisk (*) next to the k -parameter indicates that Treesearch II

ran out of time after 3 min and the best mapping found so far is reported rather than the true best mapping. The t -parameter used in these experiments was 75%. Again, notice that the matrix is not symmetric, although it ideally should be. This is due to our restriction on Treesearch II, which forces it to find a single-valued function from one object to another, instead of an unrestricted binary association. In the cases where they differ, the smaller of the two total errors may be used to estimate the relational distance between the two objects.

In Fig. 10, we graphed the Euclidean distance between attribute-value tables versus total error (3) as determined by TREESEARCH II. As can be seen from the figure, there is some correlation, but the graph is far from a straight line. Part of the reason for this is that in the normal mode of operation of the database system, an object would only be matched against those objects that are in clusters whose centroid is deemed similar to the object. However, in these experiments, we allowed every object to be matched against every other. Comparing two objects that are extremely different can result in meaningless mappings. On the other hand, we did not expect the graph to be a straight line, since this would have indicated that structural descriptions are unnecessary and feature vectors are sufficient to distinguish between objects!

To analyze the situation further, we again constructed a binary relation consisting of those pairs of the fourteen objects whose total error was less than a threshold. The threshold (59) was chosen so that this binary relation had the same number of pairs as the binary relation previously derived from the Euclidean distances with distance threshold 8. Clustering the relational distance binary relation, using the same

Table 2. Completeness and structural errors along with KVAL/100 after TREE_SEARCHII for objects of Fig. 9

	1	2	3	4	5	6	7	8	9	10	11	12	13	14
1	0	0	6	0	0	0	3	17	0	22	20	32	24	0
	0	27	57	7	7	16	35	33	29	29	39	63	39	47
			1				1.2	1	1	1	1	2	2	1
2	0	0	8	0	0	0	3	9	33	17	20	27	24	50
	27	0	80	34	34	43	62	56	66	60	66	98	66	114
			1				1.2	1	.99*	1	1	1*	1	1
3	6	8	0	6	6	7	3	2	10	2	12	29	15	10
	57	80	0	52	52	43	42	36	70	36	38	74	44	88
	1	1		1.5	1.5	1.5	1.5	1	1	1	1	1	1	1
4	0	0	6	0	0	0	3	12	0	10	12	24	24	0
	7	34	52	0	0	9	28	30	22	28	30	60	32	40
			1				1	1	1	1	1	1	1	
5	0	0	6	0	0	0	3	12	0	10	12	24	24	0
	7	34	52	0	0	9	28	30	22	28	30	60	32	40
			1				1	1	1	1	1	1	1	
6	0	0	6	0	0	0	3	12	0	10	12	24	24	0
	16	43	43	9	9	0	19	21	31	19	21	51	23	49
			1.5				1	3	1	1	1	1	1	
7	3	3	3	3	3	3	0	12	1	12	13	26	26	0
	35	62	42	28	28	19	0	22	56	20	22	58	28	74
	1.2	1.2	1	1	1.2	1		1	1.2	1	1	1	1	1
8	2	0	2	12	12	12	12	0	1	2	12	0	0	10
	35	62	36	30	30	21	22	0	52	12	14	48	20	70
	1.2	1	1.2	3	3	3	3		1	1	1	1	1	1
9	0	63	10	0	0	0	3	1	0	2	12	30	15	36
	29	98	70	22	22	31	50	52	0	52	54	86	60	96
	1	1*	1	1	1	1	1	1		2	2	1	1	1*
10	2	2	2	10	10	10	11	2	2	0	2	2	2	10
	35	62	36	28	28	19	20	12	52	0	12	48	20	68
	1	1	1	3	3	3	3	1	2		1	1	1	2
11	12	12	12	12	12	12	13	12	12	2	0	12	1	12
	37	64	38	30	30	21	22	14	54	12	0	50	20	70
	3	3	3	3	3	3	3	1	1	2		2	2	1
12	32	15	29	24	24	24	26	0	30	3	13	0	0	33
	63	78	74	60	60	51	58	48	86	50	52	0	44	102
	2	2	1	1	1	1	1	.99	1	1	2		1	1
13	17	4	15	24	24	24	26	0	15	2	1	0	0	23
	41	62	44	32	32	23	28	20	60	22	20	44	0	78
	3	2	1	1	1	1	1	1	2	1	1	1		1
14	0	43	18	0	0	0	3	12	42	10	12	41	24	0
	47	110	88	40	40	49	68	70	98	68	70	104	72	0
	1	1*	1*				1	1	1	1.5*	1	1*	1	

* These errors are estimates, because the search exceeded its time limit.

graph-theoretic clustering procedure with the same parameters as used previously, gave the following results.

Cluster	Objects
1	1, 2, 4, 5, 6, 7, 9, 14
2	3
3	1, 8, 10, 11, 12, 13

Cluster	Objects
1	1, 2, 3, 4, 5, 6, 7, 9
2	8, 10, 11, 13
3	12
4	14

The main differences are:

Recall that the clusters obtained from the Euclidean distance binary relation were as follows.

- (1) Object 3, which grouped with 1, 2, 4, 5, 6, 7 and 9 using Euclidean distance, did not group with any





Fig. 10. Illustrates the graph of total error after TREE_SEARCHII versus distance from Table 1.

Object 1



ATTRIBUTE VALUE TABLE

BASE SUPPORTS	4
TOP TYPE	2
NO. STICKS	4
NO. PLATES	2
NO. BLOBS	0
NO. UPRIGHTS	5
HORIZONTALS	1
SLANTEDS	0
NO. LEVELS	3
TOP PCS. POS.	2

SIMPT	PARTS TYPE	RELATION LENGTH	AREA	VOLUME
1	1	1.00	0.0	0.0
2	1	1.00	0.0	0.0
3	1	1.00	0.0	0.0
4	1	1.00	0.0	0.0
5	2	1.00	1.00	0.0
6	2	1.00	1.00	0.0

SP1	SP2	SUPPORTS	HOW
1	5	TRUE	12
2	5	TRUE	12
3	5	TRUE	12
4	5	TRUE	12
5	6	TRUE	23


SP1	SP2	SP3	RELATION SAME
1	5	2	TRUE
1	5	3	TRUE
1	5	4	TRUE
1	5	6	FALSE
2	5	3	TRUE
2	5	4	TRUE
2	5	6	FALSE
3	5	4	TRUE
3	5	6	FALSE
4	5	6	FALSE

SP1	SP2
1	2
1	3
1	4
2	3
2	4
3	4

SP1	SP2
1	5
2	5
3	5
4	5
5	6

Fig. 11. Illustrates the full relational structures of objects 1 and 3 of Fig. 9. (Continued over.)

Object 3



ATTRIBUTE VALUE TABLE

BASE SUPPORTS	4
TOP TYPE	2
NO. STICKS	5
NO. PLATES	2
NO. BLOBS	0
NO. UPRIGHTS	2
HORIZONTALS	1
SLANTEDS	4
NO. LEVELS	4
TOP PCS. POS.	2

SIMPLE PARTS RELATION					CONNECTS-SUPPORTS RELATION			
SIMPT	TYPE	LENGTH	AREA	VOLUME	SP1	SP2	SUPPORTS	HOW
1	2	6.00	28.26	0.0	1	2	TRUE	23
2	2	6.00	28.26	0.0	2	3	TRUE	21
3	1	2.00	0.0	0.0	3	4	TRUE	11
4	1	4.00	0.0	0.0	3	5	TRUE	11
5	1	4.00	0.0	0.0	3	6	TRUE	11
6	1	4.00	0.0	0.0	3	7	TRUE	11
7	1	4.00	0.0	0.0	4	5	TRUE	11
					4	6	TRUE	11
					4	7	TRUE	11
					5	6	TRUE	11
					5	7	TRUE	11
					6	7	TRUE	11

TRIPLES RELATION			
SP1	SP2	SP3	SAME
1	2	3	FALSE
2	3	4	FALSE
2	3	5	FALSE
2	3	6	FALSE
2	3	7	FALSE
4	3	5	TRUE
4	3	6	TRUE
4	3	7	TRUE
5	3	6	TRUE
5	3	7	TRUE
6	3	7	TRUE

PARALLEL RELATION	
SP1	SP2
1	3

PERPENDICULAR RELATION	
SP1	SP2
1	2
2	3

Fig. 11 continued.

objects using relational distance.

- (2) Objects 12 and 14, which did not group with any of the other twelve objects using Euclidean distance, each found a (different) group using relational distance.
- (3) Object 1 falls into two different clusters using relational distance, but only one using Euclidean distance.

What caused the differences? Consider object 1 and object 3 whose relational models are shown in Fig. 11. As far as the attribute value tables, they differ in number of sticks, number of uprights, number of slanted pieces and number of levels. In the highly-weighted attributes (number of sticks and number of levels) they differ only by one. In the low-weighted attributes (number of uprights and number of slanted pieces) they differ by 3 and 4, respectively. Thus the total difference was relatively small.

The mapping used in Table 2 from object 1 to object 3 had a structural error of 6 and completeness error of

57. It was a reasonable mapping that sent parts 1, 2, 3 and 6 of object 1 to parts 4, 5, 6 and 1, respectively, of object 3. The structural error of 6 stemmed from 3 connects/supports errors and 3 parallel errors. The completeness error was due to parts 4 and 5 of object 1 and parts 2, 3 and 7 of object 3 mapping to no part at all. The main problems are the lack of connectivity between the legs and seat of object 3 and its slanted legs. The information in the attribute-value table is insufficient to detect all the structural differences.

In the attribute-value table comparisons, object 12 had several more plates than the other objects and object 14 had several more sticks than the others. Thus they were too dissimilar in heavily weighted attributes to the other objects to cluster with them. In relational matching, however, object 14 shares with objects 1, 2, 4, 5, 7 and 9 a seat and four legs in the same connection and triples relationships. In this case, the relational matching is more powerful than the attribute-value table matching. Object 12 was considered similar to

objects 8, 10 and 13 in the relational matching. This is really the case of a subset of object 12 having structure similar to objects 8, 10 and 13, since partial matches were allowed.

As far as object 1, it was deemed relationally similar to objects 2, 4, 5, 6, 7, 8, 9, 10, 13 and 14. Again we find that a subset of object 1 (the seat and back) has the same simple parts and same structure as a subset of objects 8, 10 and 13. Another subset of object 1 (the seat and four legs) has the same simple parts and same structure as a subset of all of objects 2, 4, 5, 6, 7, 9 and 14. This accounts for object 1 clustering with two different groups and makes intuitive sense also.

One criticism of these results might be that the total relational error, as used, was dependent on the number of N -tuples in the relations of an object. Objects with more N -tuples would necessarily generate more errors. To study the effects of this problem, we 'normalized' the total errors by dividing the error for object i vs object j by the total number of N -tuples in object i plus the total number of N -tuples in object j

The results were, again, three clusters.

Cluster	Objects
1	1, 2, 4, 5, 6, 7, 9, 14
2	3
3	1, 2, 8, 10, 11, 13
4	12

This as opposed to the former relationally obtained clusters.

Cluster	Objects
1	1, 2, 4, 5, 6, 7, 9, 14
2	3
3	1, 8, 10, 11, 12, 13.

These results are, of course, somewhat dependent on the threshold that produced the binary relation to be clustered. The threshold was again chosen so that the binary relation would include the same number of pairs as the previous two binary relations. Object 12 was still deemed similar to objects 8, 10 and 13, but the addition of object 2 to cluster 3 forced object 12 out.

VII. CONCLUSIONS

We have defined a relational model to be used as a rough description of three-dimensional objects, such as furniture. A database of such models has been constructed and used in a preliminary set of matching experiments. The database is currently organized into clusters of objects with similar features, based on the Euclidean distance between their attribute-value tables. When an unknown object is analyzed, its attribute-value table is compared with the centroid of each such cluster and relational matching takes place against the objects in those clusters deemed similar enough. The relational matching produces a mapping from the unknown object to the model, plus a measure of their relational distance.

Our experiments comparing relational distance to Euclidean 'feature' distance showed that they are related, but not similar enough to trust the attribute value clusters. Clustering is a viable alternative, but it should be based on relational clusters. This relational grouping introduces some important new problems to study. In particular, how can we define the centroid of a relational cluster and how do we match an unknown object against the centroid description? These and other methods of reducing the number of models that participate in full relational matching are crucial to the use of a large object database. Thus, the results of the experiments reported here serve to define important new work for the future.

REFERENCES

1. N. Badler and R. Bajcsy, Three-dimensional representation for computer graphics and computer vision, *Representation of Three-Dimensional Objects*. Springer-Verlag, Berlin (1983).
2. Workshop on the representation of three-dimensional objects, R. Bajcsy, Director, University of Pennsylvania, Philadelphia (1979).
3. L. G. Roberts, Machine perception of three-dimensional solids, *Optical and Electro-optical Information Processing*, pp. 159-197, J. T. Tippett *et al.*, eds. MIT Press, Cambridge, MA (1965).
4. D. A. Huffman, Impossible objects as nonsense sentences, *Machine Intelligence 6*, pp. 295-323. Edinburgh University Press (1978).
5. M. B. Clowes, On seeing things, *Artif. Intell. 2*, 79-116 (1971).
6. D. Waltz, Understanding line drawings of scenes with shadows, *The Psychology of Computer Vision*, pp. 19-91, P. Winston, ed. McGraw-Hill, New York (1975).
7. R. Shapira and H. Freeman, Computer description of bodies bounded by quadratic surfaces from a set of imperfect projections, *IEEE Trans. Comput. C-27*, 841-854 (1978).
8. K. J. Turner, Computer perception of curved objects using a television camera, Ph.D., dissertation, School of Artificial Intelligence, Edinburgh University (1974).
9. I. Chakravarty, A generalized line and junction labeling scheme with applications to scene analysis, *IEEE Trans. Pattern Anal. Mach. Intell. PAMI-1*, 202-205 (1979).
10. T. Kanade, Recovery of the three-dimensional shape of an object from a single view, CMU U-CS-79-153, Dept. of Computer Science, Carnegie-Mellon University (1979).
11. M. Nagao, S. Hachimoto and T. Sakai, Automatic model generations and recognition of simple three-dimensional bodies, *Comput. Graphics Image Process. 2*, 272-280 (1973).
12. J. W. McKee and J. K. Aggarwal, Computer recognition of partial views of three-dimensional curved objects, Technical Report No. 171, Information Systems Research Lab., University of Texas at Austin (1975).
13. C. W. Richard and H. Hemami, Identification of three-dimensional objects using Fourier descriptors of the boundary curve, *IEEE Trans. Syst. Man Cybernet. SMC-4*, 371-378 (1974).
14. J. R. Ullman, Associating parts of patterns, *Inf. Control 9*, 583-601 (1966).
15. H. G. Barrow, A. P. Ambler and R. M. Burstall, Some techniques for recognizing structures in pictures, *Frontiers of Pattern Recognition*, pp. 1-29, S. Watanabe, ed. Academic Press, New York (1972).
16. R. T. Chien and J. M. Selander, On the use of graph models for representation and object recognition, *Pro-*

- ceedings of the Workshop on the Representation of Three-Dimensional Objects, pp. G1–G15 (1979).
17. M. Schneider, A compact relational structure representation. *Proceedings of the Workshop on the Representation of Three-Dimensional Objects*, pp. O1–O30 (1979).
 18. T. O. Binford, Visual perception by computer, *IEEE Conference on Systems and Control*, Miami (1971).
 19. G. J. Agin and T. O. Binford, Computer description of curved objects, Third International Joint Conference on Artificial Intelligence, Stanford (1973).
 20. R. Nevatia and T. O. Binford, Description and recognition of curved objects, *Artif. Intell.* 8, 77–90 (1977).
 21. D. Marr and H. K. Nishihara, Spatial disposition of axes in a generalized cylinder representation of objects that do not encompass the viewer, MIT AI Lab., Memo No. 341 (1975).
 22. J. M. Hollerbach, Hierarchical shape descriptions of objects by selection and modification of prototypes, M.I.T. AI Lab (1975).
 23. B. I. Soroka, Generalized cylinders from parallel slices, *Proceedings of the IEEE Conference on Pattern Recognition and Image Processing*, Chicago, pp. 421–426 (1979).
 24. G. J. Agin, Hierarchical representation of three-dimensional objects using semantic models, *Proceedings of the Workshop on the Representation of Three-Dimensional Objects*, pp. A1–A23 (1979).
 25. M. Minsky, A framework for representing knowledge, *The Psychology of Computer Vision*, pp. 211–277, P. H. Winston, ed., McGraw-Hill, New York (1975).
 26. D. H. Ballard, C. M. Brown and J. A. Feldman, *Outline of a Query-Driven Vision System*. Computer Science Dept., University of Rochester (1976).
 27. R. M. Haralick and L. G. Shapiro, The consistent labeling problem: Part II, *IEEE Trans. Pattern Anal. Mach. Intell.* PAMI-1, 173–184 (1979).
 28. A. Rosenfeld, R. A. Hummel and S. W. Zucker, Scene labeling by relaxation operations, *IEEE Trans. Syst. Man. Cybernet.* SMC-6, 420–433 (1976).
 29. J. R. Ullman, An algorithm for subgraph homomorphisms, *J. Ass. Comput. Mach.* 23, 31–42 (1976).
 30. U. Montanari, Networks of constraints: fundamental properties and applications to picture processing, *Inf. Sci.* 7, 95–132 (1974).
 31. R. M. Haralick and J. Kartus, Arrangements, homomorphisms, and discrete relaxation, *IEEE Trans. Syst. Man Cybernet.* SMC-8, 600–612 (1978).
 32. A. Mackworth, Consistency in network of relations, *Artif. Intell.* 8, 99–118 (1977).
 33. E. C. Freuder, Synthesizing constraint expression, *Commun. Ass. comput. Mach.* 21, (1978).
 34. J. Gaschnig, A general backtrack algorithm that eliminates most redundant tests, *Proceedings of the 5th International Joint Conference on Artificial Intelligence* p. 457 (1972).
 35. L. S. Davis, Shape matching using relaxation techniques, *IEEE Trans. Pattern Anal. Mach. Intell.* PAMI-1, 60–72 (1979).
 36. R. M. Haralick and G. Elliot, Increasing tree search efficiency for constraint satisfaction problems, *Proceeding of the 6th International Joint Conference on Artificial Intelligence*, pp. 356–364 (1979).
 37. L. G. Shapiro, A structural model of shape, *IEEE Trans. Pattern Anal. Mach. Intell.* PAMI-2, 111–126 (1980).
 38. L. G. Shapiro and R. M. Haralick, Structural descriptions and inexact matching, *IEEE Trans. Pattern Anal. Mach. Intell.* (1981).
 39. A. Hanson and E. Riseman, Segmentation of natural scenes, *Computer Vision Systems*, pp. 129–163, A. Hanson and E. Riseman, eds. Academic Press, New York (1978).
 40. H. G. Barrow and J. M. Tenenbaum, MSYS: a system for reasoning about scenes, SRI AI Technical Report. 121 (1976).
 41. A. W. Zucker and R. A. Hummel, Computing the shape of dot clusters I: labeling edge interior, and noise points, Technical Report TR-543, University of Maryland (1977).
 42. S. W. Zucker, Y. G. Leclerc and I. L. Mohammed, Continuous relaxation and local maxima selection: conditions for equivalence, Technical Report No. 78-15R, Computer Vision and Graphics Laboratory, McGill University (1978).
 43. S. W. Zucker and J. L. Mohammed, Analysis of probabilistic relaxation labelling processes, *Proceedings of the IEEE Conference on Pattern Recognition and Image Processing*, Chicago, pp. 307–312 (1978).
 44. R. M. Haralick, J. Mohammed and S. Zucker, Compatibilities and the fixed points of arithmetic relaxation processes, *Comput. Graphics Image Process.* 13, 242–256 (1980).
 45. O. D. Faugeras and K. E. Price, Semantic description of aerial images using stochastic labeling, *IEEE Trans. Pattern Anal. Mach. Intell.* PAMI-3, 633–642 (1981).
 46. O. D. Faugeras and M. Berthod, Improving consistency and reducing ambiguity in stochastic labeling: an optimization approach, *IEEE Trans. Pattern Anal. Mach. Intell.* PAMI-3, 412–423 (1981).
 47. R. M. Haralick, An interpretation for probabilistic relaxation, *Comput. Graphics Image Process.* 22, 388–395 (1983).
 48. L. G. Shapiro and R. M. Haralick, A general spatial data structure, *IEEE Conference on Pattern Recognition and Image Processing*, pp. 238–249 (1978).
 49. G. Salton, *Dynamic Information and Library Processing*. Prentice Hall, Englewood Cliffs (1975).
 50. L. G. Shapiro and R. M. Haralick, Decomposition of two-dimensional shapes by graph-theoretic clustering, *IEEE Trans. Pattern Anal. Mach. Intell.* PAMI-1, 10–20 (1979).

About the Author—LINDA G. SHAPIRO was born in Chicago, Illinois, in 1949. She received a B.S. degree in Mathematics from the University of Illinois at Urbana Champaign in 1970 and M.S. and Ph.D. degrees in Computer Science from the University of Iowa, Iowa City, in 1972 and 1974, respectively.

She was an Assistant Professor of Computer Science at Kansas State University, Manhattan, from 1974 to 1978. She is currently an Associate Professor at Virginia Polytechnic Institute and State University, Blacksburg. Her research interests include computer vision, pattern recognition, intelligent spatial information systems, computer graphics and data structures. She has completed an undergraduate textbook on data structures with R. Baron.

Dr. Shapiro is currently editor of *Computer Vision, Graphics, and Image Processing*. She is a senior member of the IEEE Computer Society and a member of the Association for Computing Machinery, the Pattern Recognition Society and the American Association for Artificial Intelligence. She is also co-editor of the newsletter for the IEEE Technical Committee on Machine Intelligence and Pattern Analysis.

About the Author—ROBERT M. HARALICK received a Ph.D. from the University of Kansas in 1969 and is

presently a Professor in the departments of Electrical Engineering and Computer Science and director of the Spatial Data Analysis Laboratory at Virginia Polytechnic Institute and State University. Dr. Haralick is the Computer Vision and Image Processing Editor of *Communications of the Association of Computing Machinery* and an associate editor of the *IEEE Transactions on Systems, Man, and Cybernetics*, *Pattern Recognition*, *Computer Vision, Graphics and Image Processing* and the *IEEE Transactions on Pattern Analysis and Machine Intelligence*. He is the Chairman of the IEEE Computer Society Technical Committee on Pattern Analysis and Machine Intelligence and the chairman of the 1983 IEEE Computer Society Computer Vision Workshop.

Professor Haralick has done research in pattern recognition, multi-image processing, remote sensing, texture analysis, image data compression, clustering, artificial intelligence and general system theory. He is a senior member of the Institute of Electrical and Electronic Engineers and a member of the Association for fellow of the Institute of Electrical and Electronic Engineers and a member of the Association for Computer Machinery and the Pattern Recognition Society.

About the Author—JOHN D. MORIARTY was born in Passaic, New Jersey, in 1956. He received a B.A. in Mathematics from Fairleigh Dickinson University (1978) and a M.S. in Computer Science and Applications from Virginia Polytechnic Institute and State University (1980). He was employed at Science Applications Inc., Falls Church, Virginia, working on operations and applications of sensor systems, from 1980 to 1983. Since 1983, Mr. Moriarty has worked for The Analytic Sciences Corporation, McLean, Virginia, on the design of image analysis systems.

About the Author—PRASANNA G. MULGAONKAR was born in New Delhi, India, on 20 January 1958. He received his B.Tech. degree in Electrical Engineering from the Indian Institute of Technology, Kanpur, India, in 1979, and his M.S. in Computer Science from Virginia Polytechnic Institute and State University in 1981. He is currently a Ph.D. candidate at VPI & SU. His research interests include computer vision, high level matching, modelling of three-dimensional objects, developing knowledge sources for matching and vision and computer graphics. He is a student member of the IEEE Computer Society and the Computer Society of India.