TR/02/85 January 1985

A prototypical approach to machine learning.

by

R. I. Phelps and P. B. Musgrove

w9259389

A prototypical approach to machine learning.

R. I. Phelps, Brunel University
P. B. Musgrove, Polytechnic of the South Bank.

Keywords: Learning, categorization, prototypes, knowledge representation.

Abstract

This paper presents an overview of a research programme on machine learning which is based on the fundamental process of categorization. This work draws upon the psychological theory of prototypical concepts. This theory is that concepts learnt naturally from interaction with the environment (basic categories) are not structured or defined in logical terms but are clustered in accordance with their similaritry to a central prototype, representing the "most typical" member.

A structure of a computer model designed to achieve categorization is outlined and the knowledge representational forms and developmental learning associated with this approach are discussed.

Introduction

This paper gives an overview of the approach being taken to learning within a continuing research project. The project is concerned with developing an operational computer based model of learning drawing on a psychological theory of categorization proposed by Rosch et al. [1]

Learning can occur in several forms.

Taught learning occurs when the system is tuned to give correct responses on a training set. It can be achieved on two levels. Superficially, by changing parameter values within a set procedure. No attempt is made to discover structure in the input and response is linked directly to the input. This is often called adaptive learning. More deeply, by representing its world in a structured way, e.g. a tree or set of classes so that input is firstly assigned to one of a number of categories and then response is inked to category membership.

In untaught learning no training set is used to set the system on the right lines. Again, learning can be achieved in two ways. Superficially, by updating parameter values within a set procedure, e.g. the use of Bayes' rule to update a probability estimate thus 'learning' about a likelihood. More deeply, by finding 'natural' structure in the data, i.e. categorizing input into naturally occurring classes. Teaching, in the sense of feedback, is necessary to associate response with these classes, but the fundamental structuring of the system's world has been achieved.

The type of learning that has attracted most attention in the A.I. literature is the learning of propositional rules to discriminate between categories [2]. Rule learning is clearly an important component of decision making systems and is well positioned to give exploitable results in working systems. Thus commerjcial systems have been produced e.g. Expertease [3]. Such systems, however, exhibit an important drawback of tackling learning at this level. In order to learn new rules the system must be given not only positive examples of the rule (and perhaps examples of near misses) but also the features from which the rule will be constructed.

But the most difficult part of the work has already been accomplished when the relevant features for rule formation have been found. Further, these systems only deal efficiently with concepts naturally described in terms of conjunctive features and do not respond kindly to 'noisy' data. Because of these considerations we feel that such propositional rule approaches, although undoubtedly useful in applications, do not provide a firm basis from which to model learning.

In addition to there being several sorts of learning, there are several sorts of knowledge that can be learned, including factual knowledge, planning knowledge involving the ability to construct sequences of actions and categorization knowledge involving the abi1ity to discriminate between classes of objects. Most present A.I. learning work concentrates on categorization, and there are reasons for considering this the fundamental type of learned knowledge. All other types depend on the pre-existence of a set of concepts which can be related, extended analyzed or combined; but before any of this can be accomplished it is firstly necessary to develop the initial set of concepts. A baby, before Other types of learning can start, has first to make sense of its 'buzzing, booming world'. This is done by recognizing sensory input, i.e. by categorizing its perceptions and relating input to those categories. Only when this is done and recognition is possible can higher levels of learning begin to operate based upon this initial structuring of knowledge.

Following from this discussion we feel it is necessary to develop a model of learning which addresses not only the questions of how to discriminate between categories using a given feature set but also the question of how such features are found. Efficient learning algorithms will depend upon the knowledge representation scheme upon which they operate: learning of logical rules fits naturally with a prepositional knowledge base but not so well with an associative network knowledge base. We therefore decided to model the entire categorization process of knowledge representation, feature selection and discrimination as one integrated system.

Because learning of categories influences the structure of knowledge in -memory and this structure in turn affects higher levels of learning operating on it, studying categorization and related representations also has implications for the operation of other types of learning.

The development of a model for such an integrated system is a complex task and so a restricted domain of classification was chosen for this work. This is the domain of 2—D silhouettes of visual objects. Reasons for this choice are the primacy of visual perception in humans and the importance of recognition. Further, Marr [4] has shown that if the difference between convex and concave parts of a silhouette represents properties of the 3- D surface and where the surface looks continuous in 2-D it really is continuous in 3-D, then the 3-D surface can often be successfully inferred from the silhouette. In limiting ourselves to this domain we are ignoring various sources of visual information, e.g. colour, texture, depth, motion. However, these properties could easily be incorporated within the learning scheme to be proposed, and the key properties of shape and parts will remain the same if we are judicious in our choice of projections. For these reasons it was thought reasonable to explore learning and representation in the 2-D domain which is much easier to handle as regards object description and segmentation. We consider unoccluded silhouettes; occlusion can be treated within this approach but we do not treat this aspect within this paper.

In the main body of this paper we outline the following aspects of this research: firstly, the psychological background relevant to this work; secondly the structure of the knowledge representation scheme used; thirdly the grouping of objects into potential members of the same category ; fourthly the use of clustering to discover categories and prototypes; fifthly the updating of categories as new objects are perceived and finally we present some illustrative examples.

PSYCHOLOGICAL CONSIDERTIONS

Some psychological evidence is available about the structure of categories, Rosch et al [1] identified three levels of categorization: superordinate, basic and subordinate . Evidence was found to support the 'prototype' theory of concepts in which membership of a category is determined by the typicality of a particular object to an ideal member of the category which has the average attributes of all class members. This theory implies that most natural concepts are ill-defined, that is, there is no rule that can determine membership for all members of a category. Furthermore, not all members of a category have equal status. Members judged to be typical of a category (e.g. apples for the ca tegory ' fruit') can be categorized faster and more accurately than members judged less typical (e.g. tomato). This is not the whole story, however, as any object may be categorized at each of several different levels, higher levels being abstract and lower levels more detailed and specific, e.g. a chair may be classified as an inorganic object, a piece of furniture, a chair or a kitchen chair. These psychologists have argued that the most cognitively efficient and therefore most basic level of categorization is that level at which the categories produced provide the most distinct clusters, i.e. the level which maximizes the similarity of objects within a category and maximises the differences between objects in different categories .

Thus of the classifications suggested for a chair, the basic level category is 'chair' because chairs are quite similar amongst themselves and distinct from tables, tomatoes, etc., whereas items of furniture are not very similar amongst themselves and kitchen chairs do not differ sharply from other chairs. Rosen et al [1] and Tversky and Hemingway [5] provide evidence that this basic level of categorization has the following properties:

1) it is the most abstract level at which instances have similar shapes.

2) it is the most abstract level at which instances have similar parts

3) it is most abstract level at which a mental image can reflect the appearance of the entire category.

4) objects are recognized more quickly as members of basic level categories than as members of categories at other levels.

5) it is the level at which humans spontaneously name an object.

The overall intention of the work of which this paper forms a part is to construct a model of the categorization process which can learn basic level "natural' categories from simple visual data without instruction. The approach adopted is to produce a model that operationalizes above psychological findings.

KNOWLEDGE REPRESENTATION

The knowledge representation scheme adopted for objects and categories has been explicitly designed to fit in with the prototype theory, and is explained in greater detail in [6] . Its application to representing silhouettes is outlined here.

Shapes and parts seem a good starting point for visual categorization, both on the psychological grounds advanced previously and intuitively. They are therefore taken as the fundamental descriptors of visual object perceptions. This raises the question of precisely what is meant by 'parts' . Marr's research in machine vision has taken the view that objects are most naturally segmented into convex parts and we have followed this 1ine of thought but refined it so that parts need only be 'psuedo-convex' in the sense that further dividing them into more convex subparts does not significantly increase the measure of convexity . This approach is reported elsewhere in more detail [7] . The effect is illustrated by the resulting 'parts' of a horse in FIG 1.

The description of the visual image of a horse is achieved in stages. Firstly, the horse is described holistically by a set of descriptors including such measures as principal axis, axis extension ratio, compactness (perimeter/area), size, etc., applied to the whole image [8] . The precise set of descriptors used is unimportant as long as it contains a rough description of the shape. Only a rough description is necessary as more accurate descriptions are provided by successive stages. At the second stage the horse is decomposed *as* in FIG 1 into its primary subparts . Each subpart is now described by the same set of descriptors, and the relative position of each subpart is also stored. This process is now repeated to any desired number of stages, the subparts being successively divided and described in increasing detail. The result is a hierarchial description as in FIG 2.

ESTABLISHING POTENTIAL CATEGORIES

We initially consider finding categories within a fixed set of objects. The objects are seen and represented in the form above.

The prototypical approach implies that at basic level objects fall naturally into distinct clusters around prototypical centres. The approach we adopt is therefore to extract measures from the representational descriptions and to use these measures as axes of a representational space wherein the objects are examined for the existence of clusters. Clustering of this type is only sensibly considered if the objects under consideration can all be represented in the same space, i.e. if they are all roughly of the same type. Thus a cow and a horse both have bodies, necks, heads and four legs and so a space with measures along these axes is concievable. However, a cow and a tomato are so different that to attempt to devise a space capable of representing them both is pointless. Hence there is an initial need to determine which of a set of objects may possibly cluster together and which are definitely in seperate categories with their own (separate) representational spaces.

The hierarchical descript ion of objects gives the possibility of extracting features at different level of detail . The first features to be considered are those at the holistic level. These will often by themselves be sufficient to rule out two objects from membership of the same category, because of an extreme difference in one or more of the measures. At this stage we only sort out obiously unlike objects: we do not wish to regard a man with arms raised as totally different from a man with his arms at his sides , but he should be differentiated from a bus. Comparisons are pairwise and so each object creats its own similarity class of potential same category objects. These classes are then reduced by deleting any two objects which on their own pairwise comparison have been found to be totally different. Under this procedure it is possible to have non-disjoint sets of objects as the potential categories. Although unlikely, it is therefore possible that an object may end up as a member of more than one category. This is not a drawback, to the learning process as its purpose is to discover feature clusters representing natural structure in its input, not to uniquely assign objects to categories. Such non-uniqueness is to be expected in borderline cases, e.g. a large stool might be equally well classified as a smal1 table; the important aspect for learning is that the two categories 'table' and 'stool' should be found.

The next stage of the categorization process is the identification of potential category membership at the level of primary part descriptions. This stage operates separately on each set of objects identified as a potential category at the prevous stage. Extraction of parts from an image (FIG 1) is not always a clear-cut process. Often there will be more than one way of dividing a shape into convex parts and the measures used to decide the best division will not significantly favour any on division. In FIG 1 extracting the horse hooves a separate parts is another satisfactory option. All possible 'good' divisions into parts are therefore considered. For each pair of objects, A, B, all primary part descriptions of A are compared with each of B's. A match is obtained if each description has the same number of parts, the description of each part of A is similar to that if a corresponding part in B and the relationships between parts of A are similar to those between corresponding parts of B. The relationship between parts may be encoded by a simple device such as the points of contact or direction between centres of adjoining parts. If a match cannot be achieved the two objects are not members of the same category.

This is too restrictive, however, as can be seen by considering a chair with or without arms, or a man with his arms bent giving rise to two extra parts (forearms). We therefore allow objects A and B to be potential members of the same category if we can find a subset of parts A and a subset of parts of B which match and which account for most of the area of images A and B. When such matched subsets are found, any configuous sets of parts not in the subset is fused into one 'lumped' part. It may then be possible to match lumped parts of A and B, in which case these matched lumped parts are added the matched subsets. Within each potential category each object is represented only by its matched subset. The process of finding such matching subsets operate by attempting to match the largest parts of each object first, followed by other parts roughly in order of size but subject to parts being chosen so that the relationships between each set of parts chosen so far are similar for both objects.

If potential categories of at least eight members have been identified than a first attempt is made to discover actual categories within each potential set by means of cluster analysis.

CLUSTER ANALYSIS

Each object is represented (in part subset form) as a point in a representational space $(p_{11} \ldots p_{1n} \ldots . p_{k1} \ . \ . \ . \ p_{kn} \ r_{12} \ \ldots r_{k-1k})$

Where P_{ij} is the result of measure j applied to part i and r_{ij} is the relationship between parts i and j. A clustering algorithm has been developed which will seek any convex clusters present among these points. In particular this algorithm allows clusters to intersect, thus allowing for the fact that some natural categories may not have sharp dividing lines between them. Further, it does not require the number of clusters present to be set not does it require the alteration of parameters to give good results on different data sets. In these ways it represents an advance in automatic cluster detection over other existing algorithms. Details of this method, which is designed to emulate human performance in detecting dot clusters, are described elsewhere [7] . This algorithm is used to explore the cluster structures found using different subsets of the part and relationship measures as the axes of the space. The objective is to find a minimal set of these measures which provides a 'good' cluster structure . Good in this sense means that the clusters located should be well separated from one another relative to the dispersion of points within each cluster, convex and should account for a high proportion of the objects, leaving a minimum of unclustered 'noise' points.

If the object vectors display clustering then in most cases this will be due at least in part to the different clusters displaying differences in their parts. Thus the search for object clusters can in fact be largely carried out by searching for clusters within corresponding parts of the objects (p_{i1} ... p_{in}). This is a much easier search owing to the reduced dimens iorial ity of the space . The subsets of the measures finally used to examine clustering of the whole objects are then limited to subsets chosen from those part measures which gave rise to part clusters and the relationship measures . If , however, no clustering is apparent within the parts, clustering must be carried out using the whole object vectors.

It may be that distinct subclusters are contained within the clusters and potential clusters found, which are not shown-up by the set of measures giving the best overall breakdown of this set into clusters. Therefore each actual and potential cluster set goes iteratively through the clustering stage of the process until no further clusters can be found. This also deals with any clusters based on subtle interactions between parts which are not detected when limiting measures to those which gave rise to part clusters.

If at least one cluster is found, the objects within it are members of the same category. Any noise points are still potential members of the same category. The outcome of the process at this or further stages of processing is a partitioning of the object set into (possibly overlapping) classes , each containing category members or candidates for membership of a category [FIG 3] . However, the category structure formation we seek is only contained in clusters of objects found via the clustering algorithm. It is the set of measures used to find each cluster and the position of the cluster within this space which provide the operational notion of categories. Objects in a class where no cluster has been found or which have been classified as 'noise' points are uncategorized at this stage .

The above stage of processing from the primary parts description can be repeated on the next level of the description hierarchy, the secondary parts, and again repeated at more detailed levels until the descript ion hierarchy ends. At each more detailed stage more emphasis is planed on close matching of object descript ions and so categories differing in smaller details may be distinguished. However we are searching for sets of features to form spaces wherein naturally well separated and relatively noise free clusters occur. As we descend to consider more detailed levels, it is expected on the basis of Rosch et al's results [1] that natural clusters will not form so readily, the similarities between objects being reduced by detailed differences, giving less tightly knit clusters and less sharp boundaries between clusters and their surroundings. Thus in general we would expect basic categories to be formed at the primary parts level and that clustering at level of greater detail would result in a less 'good' cluster structure. The processing is continued, however to further levels of detail as long as clusters can be located, in order to find subordinate categories within the basic ones. Different classes my therefore be processed to different levels of detail.

PROTOTYPE FORMATION

For each cluster that has been found a prototype is formed. A prototype consists of the matched subset of parts and relations common to all objects in the set in which the cluster was found, with the measures of the parts and relations set to the average values of all objects in the cluster. If the cluster was formed at a detailed level of the description hierarchy then it has descended from potential or actual categories formed at earlier levels of less detail. Not only the prototype but also the matched subsets of parts common to objects in the sets at previous stages from which this cluster has descended are stored with average values for this cluster. This has advantages in recognizing new instances of the category [6]. Once a prototype has been constructed it is no longer necessary to retain the original object descriptions; all that is necessary is the prototype plus a measure of the variation within the cluster. The measure of variation is necessary for two reasons: firstly, in order to decide whether a new object is sufficiently close to the category centre to be a member; secondly, as new objects are encountered and assigned to new or existing categories, it is necessary to check that clusters do not merge with one another, removing the natural distinction between them which was their raison d'etre.

DYNAMICS OF LEARNING

We consider the development of categories as more objects are encountered. As objects are perceived their representations are stored. Since several objects are required to reognize a category, a large memory of instances must be accumulated and regularly checked for new categories. For a finite memory a FIFO discipline seems natural where a new object removes the oldest uncategorized object from memory. This enables the system to concentrate on up to date input and gradually eliminate any incorrectly represented objects.

New objects identified as members of an existing category are not stored but may be used to update the average measures and the measure of variation in the category. The proportions of objects falling within a cluster in a given representational space is also stored to check that cluster density remains significantly above noise density.

If the cluster structure breaks down because clusters merge or because the proportion of noise points increases reducing the relative density of the clusters, the objects in this space must be regarded as uncategorized and the categorization process reapplied to a representative set of these objects. For this reason it is desirable to store a small number of the more recent members of each category.

In the short term the categories formed will be influenced by the order of presentation of objects, since once a prototype has been formed it will be maintained and the system will attempt to fit later objects into the existing prototypical framework until the cluster structure so formed breaks down. In the long term, with sufficient storage available, any order of presentation should lead to convergence on the same category structure. Short term order dependence is a necessary corollary of seeking to extract the maximum structure from limited data and the "importance of first impressions" mirrors a common facet of human performance.

ILLUSTRATIVE ANALYSIS

As an example of the scheme outlined above, consider a set of four silhouettes: two horses, a cow and a bird (wings closed). These have been broken down into their primary convex parts, but these do not necessarily correspond to the parts which we normally consider these animals to have. For example, one leg may occlude another, so in the cow the two front legs are considered as one convex part. In the horses, the area at which the two legs merge into the body and into each other has been separated out as a convex part. The back legs and tail of the cow have merged and been split into three convex regions. The labels given to the regions are for illustrative purposes only — it is not known to the algorithm at this stage that the regions of the two horses here labelled 'body' both correspond to the same part of the concept 'horse'. [TABLE 1].

The bird may be immediately differentiated from the others because of the great difference in the number of parts found. The cow and horses cannot be immediately distinguished and so an initial attempt must be made to match their part descriptions. The best match of corresponding subsets can be made by matching the largest parts (bodies), the next largest (heads) and one of the next largest (necks), which triples preserve roughly the same relationships between their parts and account for around 0.75 of their areas. With these subsets matched the front junctions and legs of a horse are contiguous and would be treated as one lumped part, as would the rear junctions and legs. Similarly, the front legs and hooves of the cow would be lumped as would the 3 rear leg parts and hooves.

The three matched parts are now considered individually to look for clustering within each part . For the bodies, the major difference is in area between cow and horses, which might form a clustering characteristic with a larger sample. There is little evidence of clustering from the head measures. For the necks there is evidence of higher compactness, lower elongation and higher area measures for the horses which again might form clustering characteristics. The triples would now be examined in a space with axes chosen from body area , neck compactness, neck elongation, neck area plus interpart relations, such as direction between part centres.

In fact for these three animals, using the above four part meaksures , there is a clear indication of difference between the horses and the cow, indicating a probable cluster structure .

Average values of the objects in each cluster go to form prototypes. In this case the prototypical horse silhouette would consist of the simplified representation: body, neck, head, frontlegs, hindlegs, tail, together with their average measures and relations.

CONCLUSONS

We have presented an outline of the approach we are taking to machine learing . Although it is presently limited to simple visual categorization it is hoped that by studying learning at this fundamental level we will provide a foundation for models of higher level learning, and that principles of organizations that emerge in categorization will also be incorporated there.

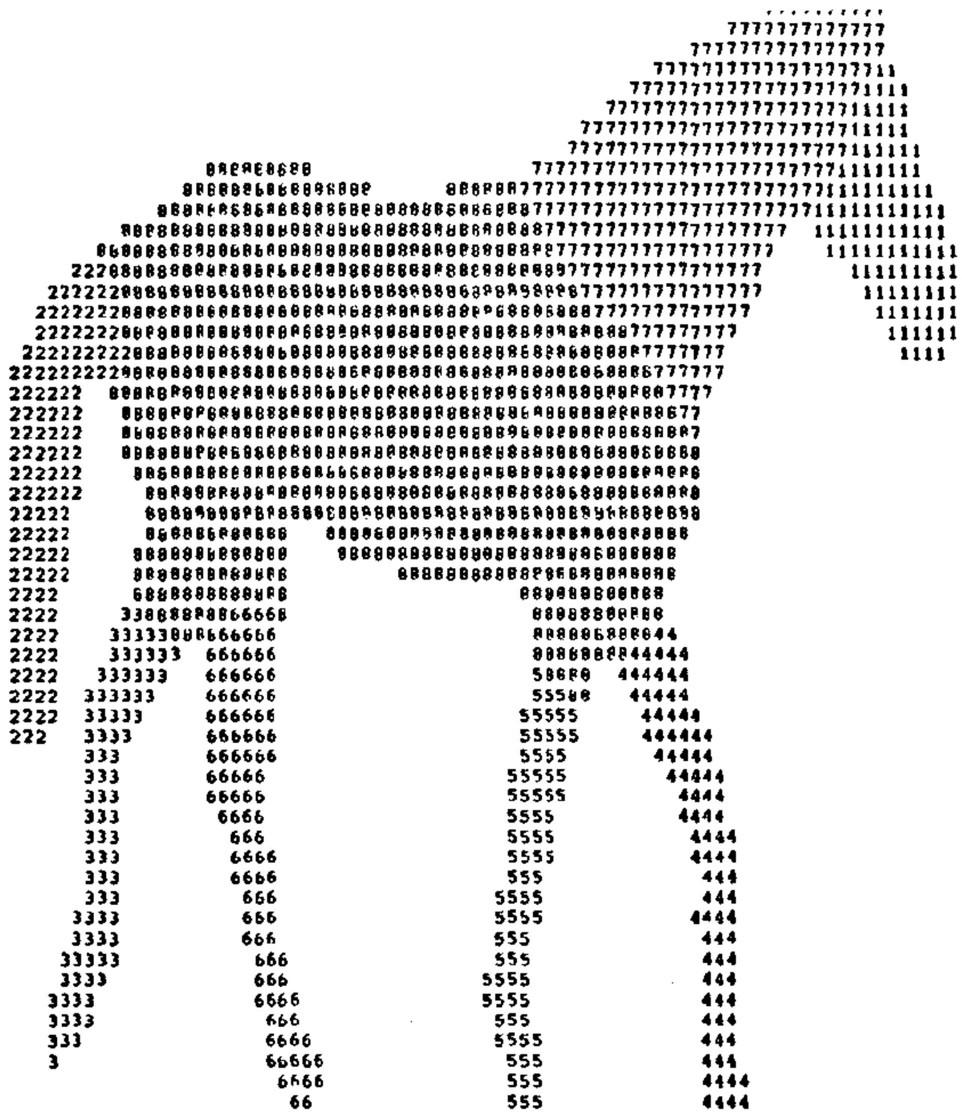

FIG. 1

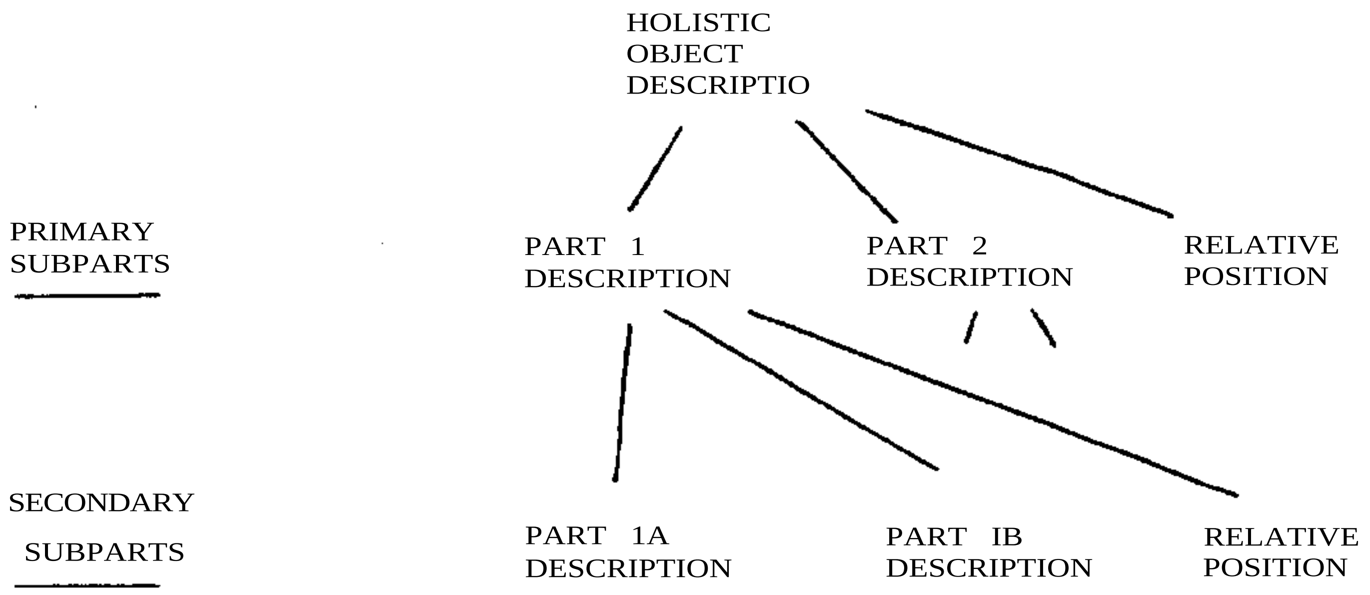
HOLISTIC
OBJECT
DESCRIPTIO
PRIMARY
SUBPARTS
PART 1
DESCRIPTION
PART 2
DESCRIPTION
RELATIVE
POSITION
SECONDARY
SUBPARTS
PART 1A
DESCRIPTION
PART IB
DESCRIPTION
RELATIVE
POSITION

FIG 2.

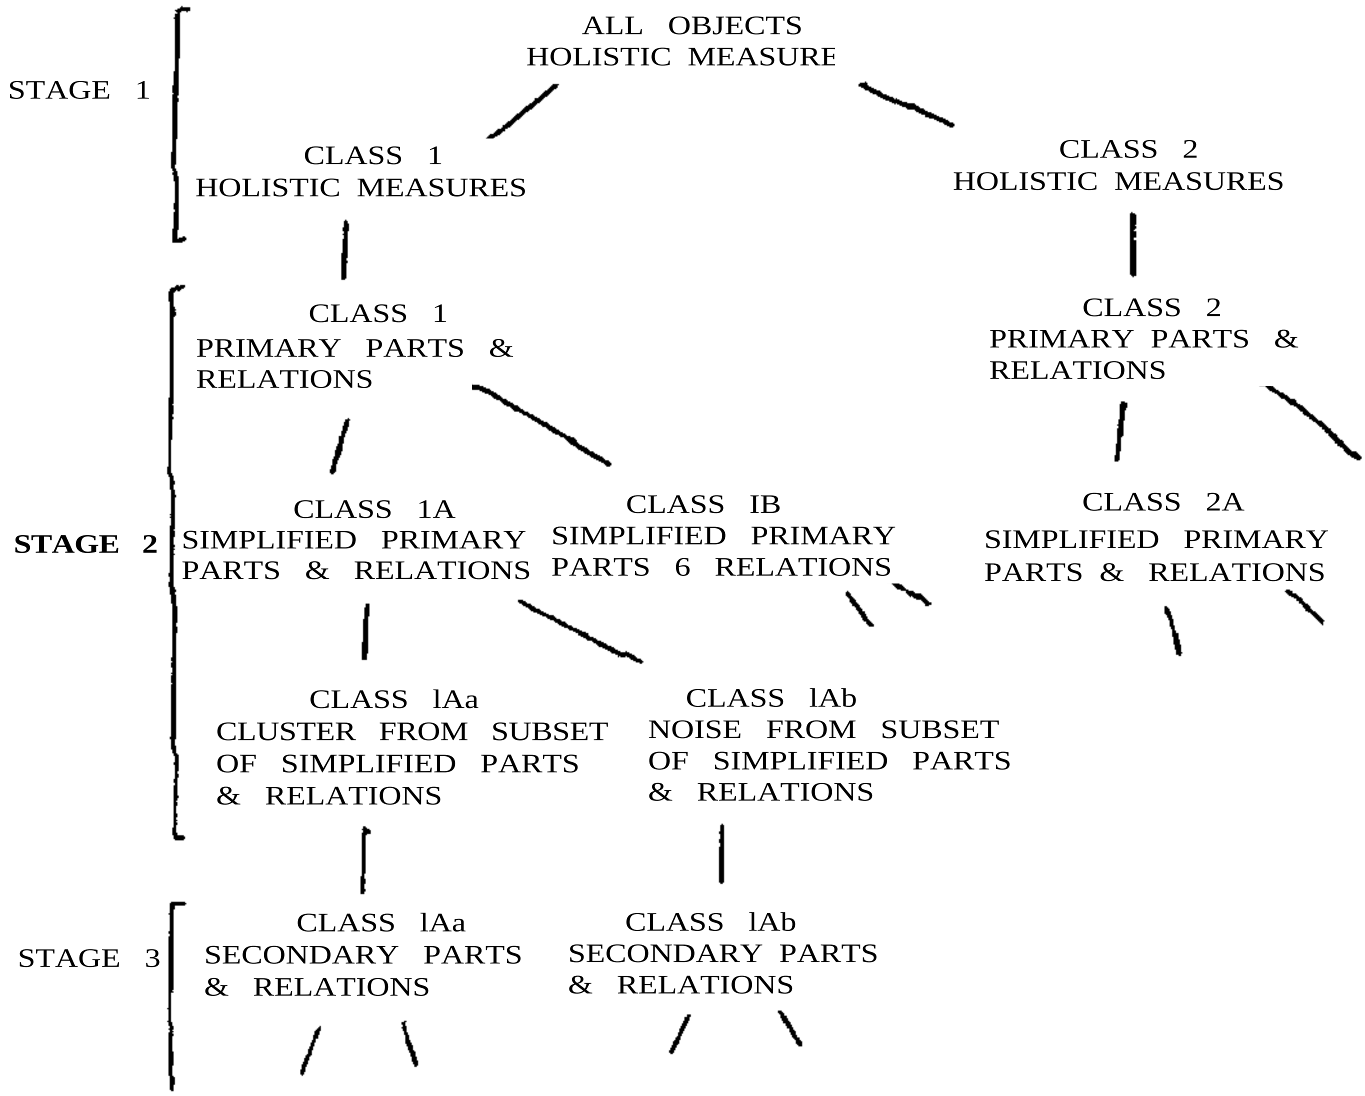
STAGE 1
ALL OBJECTS
HOLISTIC MEASURE
CLASS 1
HOLISTIC MEASURES
CLASS 2
HOLISTIC MEASURES
STAGE 2
CLASS 1
PRIMARY PARTS &
RELATIONS
CLASS 2
PRIMARY PARTS &
RELATIONS
CLASS 1A
SIMPLIFIED PRIMARY
PARTS & RELATIONS
CLASS IB
SIMPLIFIED PRIMARY
PARTS 6 RELATIONS
CLASS 2A
SIMPLIFIED PRIMARY
PARTS & RELATIONS
CLASS lAa
CLUSTER FROM SUBSET
OF SIMPLIFIED PARTS
& RELATIONS
CLASS lAb
NOISE FROM SUBSET
OF SIMPLIFIED PARTS
& RELATIONS
STAGE 3
CLASS lAa
SECONDARY PARTS
& RELATIONS
CLASS lAb
SECONDARY PARTS
& RELATIONS

FIGURE 3.

REFERENCES

1 Basic objects in natural categories. E. Rosch, C. B. Mevis, W. Gray, D. Johnson, Cognit. Psychology 8 pp. 382-439. 1976,

2 Machine Learning. Ed. R. S. Michalski et al. Tiogo Press. 1983.

3 Expertease. ESI. Ltd. Edinburgh. Based on ID3 in |2| .

4 Vision. D. Marr. 1982.

5 Objects, parts and categories. B. Tversky & K. Hemingway. J. Exp. Psychology 113/2 pp. 169-190. 1984.

6 Knowledge Representation for basic visual categories. R. I. Phelps & P. B. Musgrove. Brunei Technical Report. 1984.

7 Brunel Technical Report. In preparation.

8 "Relational Descriptions in Picture Processing'[1]. H. G. Barrow & R, J. Popplestone, Machine Intelligence 6. 1971 Edinburgh University Press, Ed. B. Moltzer & D. Michie.

Horse 1

	Head	Neck	Tail	Body	Front leg 1	Rear leg 1	Rear body/leg Junction	Front leg 2	Rear leg 2	Front body/leg Junction
Compactness	1.13	1.34	0.59	.77	.52	.52	.92	.75	.59	.92
Elongation	.60	.39	.84	.70	.93	.47	1.40	.17	.74	.65
Proportional Area	.11	.09	.07	.48	.04	.04	.06	.04	.04	.02

Horse 2

	Head	Neck	Tail	Body	Front leg 1	Rear leg 1	Rear body/leg junction	Front leg 2	Rear leg 2	Front body/leg junction
Compactness	.78	1.26	.57	.81	.46	.48	.84	.42	.47	.65
Elongation	.49	.43	.89	.74	.92	.92	.57	.91	.92	.89
Proportional Area	.14	.10	.07	.42	.05	.06	.04	.05	.05	.02

Cow

	Body	Neck	Head	Front legs	Back legs(1)	Back legs(2)	Back legs(3)	Front Hooves	Rear Hooves
Compactness	.82	.94	.94	.65	.70	.63	.88	1.28	1.23
Elongation	.66	.65	.39	.76	.93	.88	.77	.62	.37
Proportional Area	.70	.05	.09	.05	.04	.04	.02	.01	.01

Robin

	Body	Head	Tail	Wingtip
Compactness	1.08	.80	.70	.79
Elongation	.50	.62	.90	.71
Proportional Area	.73	.14	.10	.03

TABLE 1