

- 21 Treisman, A.M. and Gelade, G. (1980) A feature-integration theory of attention *Cognit. Psychol.* 12, 97-136
- 22 Moore, C.M. et al. (1996) Are attentional dwell times inconsistent with serial visual search? *Psychonomic Bull. Rev.* 3, 360-365
- 23 Chun, M.M. and Potter, M.C. (1995) A two-stage model for multiple target detection in rapid serial visual presentation. *J. Exp. Psychol. Hum. Percept. Perform.* 21, 109-127
- 24 Seiffert, A.E. and Di Lollo, V. (1997) Low-level masking in the attentional blink. *J. Exp. Psychol. Hum. Percept. Perform.* 23, 1061-1073
- 25 Enns, J.T. and Di Lollo, V. (1997) Object substitution: A new form of masking in unattended visual locations. *Psychol. Sci.* 8, 135-139
- 26 Rensink, R.A., O'Regan, J.K. and Clark, J.J. (1997) To see or not to see: The need for attention to perceive changes in scenes. *Psychol. Sci.* 8, 368-373
- 27 Pashler, H. (1994) Dual-task interference in simple tasks: Data and theory. *Psychol. Bull.* 2, 220-244
- 28 Welford, A.T. (1952) The 'psychological refractory period' and the timing of high-speed performance - A review and theory. *Br. J. Psychol.* 43, 2-19
- 29 Spence, C. and Driver, J. (1997) Audiovisual links in exogenous covert spatial orienting. *Percept. Psychophys.* 59, 1-22
- 30 Posner, M.I. and Boies, S.J. (1971) Components of attention. *Psychol. Rev.* 78, 391-408
- 31 Potter, M.C. and Chun, M.M. Two attentional deficits in serial target search: The visual attentional blink and an amodal task-switch deficit. *J. Exp. Psychol. Learn. Mem. Cog.* (in press)
- 32 Duncan, J., Martens, S. and Ward, R. (1997) Restricted attentional capacity within but not between sensory modalities. *Nature* 369, 313-315
- 33 Husain, M. et al. (1997) Temporal dynamics of visual attention reveal a non-spatial abnormality in spatial neglect. *Nature* 385, 154-156
- 34 Posner, M.I. et al. (1987) How do the parietal lobes direct covert attention? *Neuropsychologia* 25, 135-145
- 35 Joseph, J.S., Chun, M.M. and Nakayama, K. (1997) Attentional requirements in a 'preattentive' feature search task. *Nature* 387, 805-808

# Computational theories of object recognition

Shimon Edelman

**This paper examines four current theoretical approaches to the representation and recognition of visual objects: structural descriptions, geometric constraints, multidimensional feature spaces and shape-space approximation. The strengths and weaknesses of the four theories are considered, with a special focus on their approach to categorization - a computationally challenging task which is not widely addressed in computer vision, where the stress is rather on the generalization of recognition across changes of viewpoint.**

The study of visual object recognition has seen such rapid development recently that its comprehensive survey would not fit within the confines of a journal paper. In this short review I concentrate on some aspects of what Marr<sup>1</sup> termed the 'computational theory' of object representation. Recognition algorithms stemming from the different computational formulations of the problem of representation are also mentioned. Very little space is devoted to implementational issues, and none at all to the evaluation of various theories as models of human performance or as explanations of the functional neurobiology of object recognition in primates (see Refs 2-5 and the forthcoming special issues of *Vision Research* and *Cognition*). Traditionally, in cognitive science, debates concerning theories of object representation centre on computational problems stemming from the effect of viewpoint on the appearance of objects<sup>6-12</sup>. The emergence of powerful formal methods for overcoming the effect of viewpoint<sup>7,13,14</sup> and the recent successes of surprisingly simple empirical approaches to recognition<sup>15,16</sup> are likely to shift the focus of theoretical discussion to other topics. Indeed, my chief aim here is to bring to the foreground a class of computational problems that differ from those related to viewpoint dependence, yet confront any

recognition system. These problems arise from the need to 'categorize', or make sense of, novel objects.

## Perception, recognition and categorization

The spectrum of problems arising in connection with object recognition is best understood in terms of two basic distinctions. The first of these has to do with the perception of the shape of an observed object on the one hand, and the recognition of objects on the basis of their shapes on the other hand. A classical observation of this distinction was made by Wittgenstein<sup>17</sup>, who discussed at length the difference between seeing a shape and seeing it 'as something'. Unlike merely perceiving a shape (a problem not addressed in the present review), recognizing it as something involves memory, that is, representations of shapes seen in the past. The form of these representations is constrained by the various factors, such as orientation and illumination, that affect the appearance of objects. Because of the effects of orientation, for instance, simply storing a particular snapshot of an object for future reference would not do: another view of the same object may turn out to be less similar to the stored view than to a view of a different object, leading to an erroneous recognition. As noted above, contemporary theoretical treatments

S. Edelman is at the School of Cognitive and Computing Sciences, University of Sussex, Falmer, Brighton, UK BN1 9QJ.

tel: +44 1273 678659  
fax: +44 1273 671320  
e-mail: shimon@ccs.sussex.ac.uk

## Glossary

**Basis pursuit:** In signal processing, choosing the optimal 'alphabet' of features for the description of a data set (for example, basis functions in function approximation), and, simultaneously, estimating the optimal contribution (for example, weight) of each feature to the observed data.

**Correspondence:** A mapping that assigns the same label to different projections of the same feature (for example, a point on an object's surface), as seen in two images taken from different vantage points.

**Curse of dimensionality:** The exponential dependence of the number of examples required for learning a task on the number of dimensions of the representation space<sup>41</sup>. Suppose that filling a region in a one-dimensional feature space with representative examples requires ten data points; a comparable coverage of a three-dimensional feature space would then require 1000 data points<sup>53</sup>.

**Fiducial geometric features:** Features that are associated with fixed locations on the object's surface and therefore can be trusted to convey information about its geometry and orientation. A surface marking or a corner formed by two surfaces meeting at an angle are good geometric features; a smooth bend in a surface is not.

**Instability of structural interpretation:** The possibility of assigning multiple structural interpretations to a given image, typically exacerbated by a sensitivity of the interpretation process to fine details of the data (which are most prone to corruption by noise).

**Labelled graph matching:** The establishment of a mapping between the vertices and the edges of two graphs in such a manner that the labels carried by the elements of one graph match those of the other one.

**Latent variable analysis:** In statistics, explaining the variation observed in a data set in terms of a few postulated 'hidden variables', which give rise to the data through the process of observation or measurement. The aim of the analysis is to specify both the variables and the observation process, to obtain an optimal fit to the data.

**Manifold:** Intuitively, a smooth curve or surface embedded in a higher-dimensional space. Performing, for instance, a thousand simultaneous measurements on the image of an object effectively

maps it to a point in a 1000-dimensional abstract space. If the object is then made to rotate around a fixed axis, the curve ascribed by that point as the values of the various measurements change with rotation is a one-dimensional manifold embedded in the 1000 dimensions.

**Non-accidental properties:** Two-dimensional image features that can be used to make inferences about three-dimensional object structure because of the low likelihood that they arise by chance<sup>21</sup>. A representative example of such a feature is a pair of parallel lines; because a chance image alignment of two segments that are in fact not parallel in 3D is unlikely, two parallel lines in the image are a good indicator of the presence of a three-dimensional shape such as a cylinder 'out there' in the scene. For a Bayesian treatment of related issues, see Ref. 54.

**Regularization:** A common mathematical technique applied to problems that are formally ill-posed. By extending the definition borrowed from the theory of differential equations, a problem is considered to be ill-posed if its solution does not depend continuously on the data, or if more than one solution exists, as in the case of structural interpretation. Regularization attempts to reduce the solution space, by imposing additional constraints, over and above those contained in the data. References to the mathematical literature on regularization and a discussion of its relevance to low-level visual tasks can be found in Ref. 55.

**Shape space:** An abstraction, introduced by Kendall<sup>33,55</sup> to allow a rigorous statistical treatment of problems having to do with variation of shape in a sample of objects. Objects are treated as points in a metric space, where similar shapes correspond to nearby locations. Categorization thus becomes a matter of determining the location of the stimulus relative to other points in the shape space.

**Viewpoint consistency constraint:** The mathematical relationship between the projected (2D) coordinates of the fiducial features, imposed by their arrangement in the 3D space. For a rigid object, the projected location of some initially detected features constrains the possible range of its orientation and predicts the location of other features. If this prediction is verified (as in the second stage of recognition by alignment<sup>7</sup>), the initial hypothesis of the presence of the object in the image is confirmed.

of recognition concentrate precisely on this problem; state of the art algorithms<sup>18</sup> are capable of overcoming it, provided that several views per object are available, and that certain auxiliary problems such as **correspondence** (see Glossary) are solved.

These algorithms, however, do not necessarily observe another distinction: that between recognition and categorization. In most everyday recognition tasks, memory is involved in a peculiar manner: whereas recognition is (literally) remembering a thing one saw before, making sense of a novel object (as when a child realizes that a giraffe, glimpsed at a zoo for the first time, is a quadruped animal) is more like remembering a thing never seen before. Obviously, to be capable of categorization, either the system's memory traces, or the process whereby these are compared to the stimulus, or both, must be structured appropriately; no algorithm designed merely to counter the effects of viewpoint would do. One may note that the distinction between familiar and novel shapes need not be of

any consequence at the 'front end' of a visual system (the processing stage whose task is to form a description of the current stimulus). To wit, completely unfamiliar objects can be, in principle, described in terms of their constituent edges, surfaces and other spatial features, without recourse to memory. However, even a complete specification of the shape of an object does not qualify as its categorization: the latter has to do with other objects that the stimulus resembles, rather than with the details of the shape of the stimulus itself. In other words, faithful geometrical description is no substitute for recognition or categorization; computing such a description constitutes a separate problem, which may have little to do with the problem of representing objects in a form suitable for recognition or categorization.

Among the various computational approaches to representation proposed in the past, the one intuitively most suitable for categorization is the variant of structural decomposition usually attributed to Marr and Nishihara<sup>19</sup> and, subsequently, popularized as a psychological theory by

### Box 1. Structural descriptions

*Summary:* an object is represented as a collection of parts (chosen from a small alphabet common to all objects), along with their spatial relationships.

*Main mathematical tools used in representation and recognition:* combinatorial optimization theory.

*Examples:* recognition by components<sup>a</sup> and minimum description length interpretation<sup>b</sup>.

#### Strengths

- (1) Conceptual parsimony – a handful of primitives can allow a very large number of object classes to be represented concisely.
- (2) Invariance to changes in viewing conditions, as long as the parts and their relationships can be identified.
- (3) Good support for categorization, stemming from the possibility of repressing novel objects in terms of the same primitives as the familiar ones.

#### Weaknesses

- (1) No reliable method for the extraction of structural primitives from raw images exists at present.
- (2) Choosing the right structural interpretation among the many that are possible in a given situation is likely to be problematical.

#### References

- a Hummel, J.E. and Biederman, I. (1992) Dynamic binding in a neural network for shape recognition *Psychol. Rev.* 99, 480–517
- b Bienenstock, E., Geman, S. and Potter, D. (1997) Compositionality, priors, and object recognition, in *Neural Information Processing Systems* (Vol. 9) (Mozer, M.C., Jordan, M.I. and Petsche, T., eds), MIT Press

Biederman<sup>8</sup>. Until very recently, it seemed that this structural theory was the only one offering a principled treatment of novel objects. For this reason, it is the first theory I discuss; the emerging alternatives are presented in later sections.

#### Structural decomposition

In any structural decomposition model, the shape of an object is described in terms of relatively few generic components, joined by spatial relationships that are chosen from an equally small fixed set (see Box 1). The standardization of the primitives (the components and their relationships) is crucial in that it allows representation of novel objects. Mathematically, comparison between objects then amounts to **labelled graph matching** (a difficult combinatorial problem)<sup>20</sup> and categorization – to the attribution of an object to an equivalence class of graphs corresponding to shapes that are structurally identical, yet may be geometrically distinct.

#### Recognition By Components

A typical structural theory, Biederman's Recognition By Components (RBC)<sup>8</sup>, postulates a set of 30 or so primitive shapes (geons), claimed to be detected easily in images as a result of their '**non-accidental**' properties. The latter are three-dimensional features that are almost always (that

is, barring an accident of viewpoint) preserved by the imaging (projection) process<sup>21</sup>.

The use of non-accidental features to infer the presence of geons, and the distributed computation of the graph structure of the input object, are the cornerstones of the implementation of RBC described by Hummel and Biederman<sup>22</sup>. This work demonstrated the ability of a carefully engineered multilayer neural network to derive structural representations from labelled line drawings. In many respects, however, it also served to highlight the shortcomings of RBC, three of which are discussed below (see Fig. 1).

#### Computational problems

*The need for metric information.* Although, in principle, RBC is capable of representing novel shapes (via their structural decomposition), this ability comes at the expense of ignoring fine (metric, or quantitative, as opposed to structural, or qualitative) distinctions among shapes. This shortcoming was recognized and amended by Stankiewicz and Hummel<sup>23</sup>, who augmented RBC by quantitative variables, encoding, for instance, the lengths of the various parts of an object in addition to their qualitative characteristics such as convexity or cross-section shape.

*Difficulties with the recovery of parts.* A more severe problem faced by the structural approaches is the need for reliable detection of parts in images. One aspect of the detection problem – inferring three-dimensional structure from a two-dimensional projection – is essentially solved by the use of non-accidental features. Nevertheless, the difficulty of finding in the image lines and junctions whose non-accidental relationships can be used to infer the presence of geons has so far precluded RBC from being applied to the recognition of objects in grey-level images. The model described by Hummel and Biederman<sup>22</sup> worked for hand-labelled line drawings; invariably, attempts to apply RBC to images of real objects involve highly simplified shapes consisting of two or three clearly distinguishable parts, and are likely to use range data instead of photometric images.

*Instability of description in terms of parts.* Even if the input to a structural decomposition system is given in the form of a collection of labelled lines, its interpretation in all but the most artificial examples is problematical, because of an inherent **instability of structural interpretation** that affects all structural approaches. The instability stems from the possibility to decompose any shape in a number of ways, depending on the primitives that are assumed to exist. For example, a handprinted letter 'A' can be decomposed into either three or five approximately straight lines, depending on whether the sides of the A are represented as single long lines or concatenations of two shorter segments each. It should be realized that the same problem arises in any combinatorially structured domain. For example, the problem of **basis pursuit** in signal processing<sup>24</sup> consists of choosing a subset of basis functions, whose weighted sum best approximates a given signal. The difficulty here stems from the possibility of decomposing the signal in many different ways, depending on the choice of basis functions and on the optimization criterion. Likewise, in **latent-variable analysis** in statistics<sup>25</sup> one is faced with the problem of selecting (or rather, postulating) the set of

variables, and of estimating their contributions to the observables.

It has been suggested that the instability problem may be alleviated by imposing a prior expectation (in the Bayesian sense) on each possible solution, at a number of levels of a structural hierarchy, as in pixel - edge element - curve<sup>26</sup>. This 'compositional' approach attempts to combat instability by 'regularization' of the solution, and by using top-down expectations descending from the higher levels of the hierarchy to help disambiguate the interpretations at the lower levels.

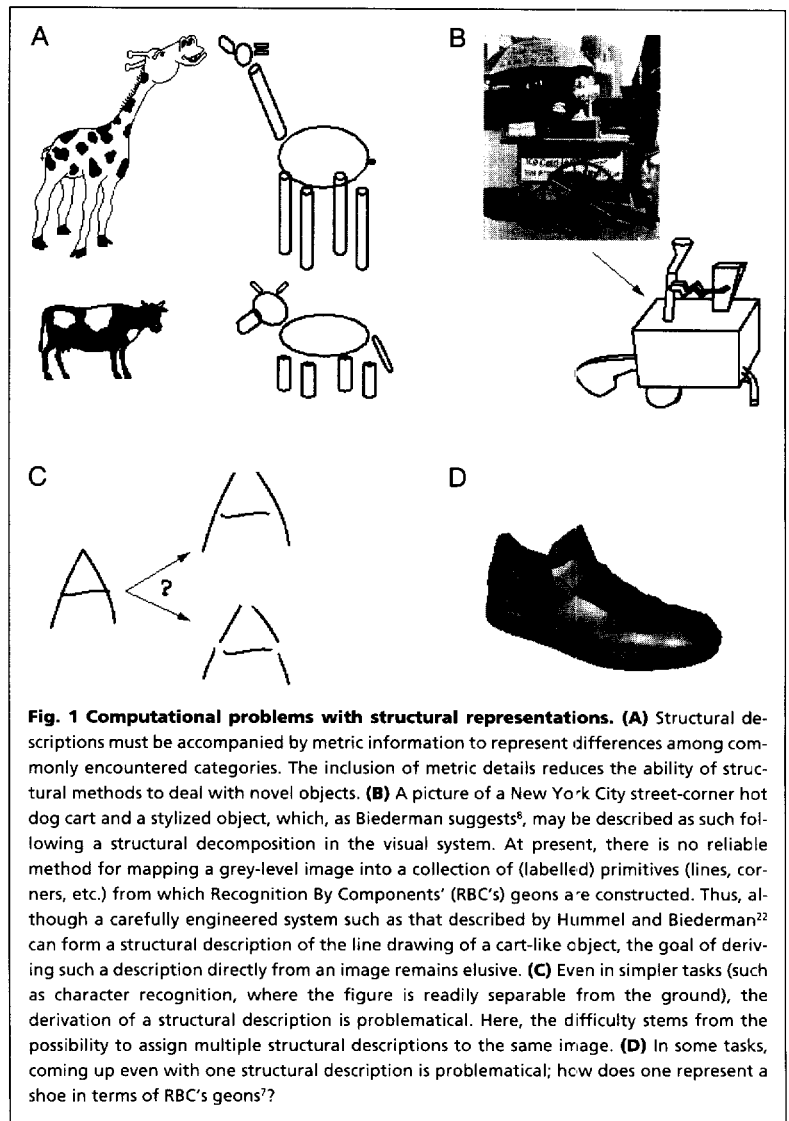
Because structural interpretation guided by these principles has not been attempted for unconstrained object classes or for 'raw' grey-level images, it is difficult to estimate the effectiveness of the compositional approach in overcoming the instability problem, or the sheer combinatorics of representing moderately complex objects as structural hierarchies<sup>27,28</sup>. Experience with large-scale projects that adopted this approach has not been encouraging. A vintage example, the MIT Vision Machine program<sup>29</sup>, which explored bottom-up structured solutions to low-level visual tasks such as edge detection, confronted computational difficulties and ran into a representational dead end. On the one hand, its recovery of shape from various low-level cues proved to be unreliable. On the other hand, nobody seemed to have any use for the recovered shape even when the low-level computations did work. The abandonment of this and similar projects in the early 1990s suggests that the current attempts at reviving the structural interpretation methods<sup>26,30</sup> are about to face severe difficulties, both computational and conceptual.

### Geometric constraints

Whereas structural methods ignore much of the quantitative information inherent in the image locations of object features, geometric methods such as alignment<sup>7</sup> use this information to identify the object and to compute its pose with respect to the observer (Box 2). These methods rely on the following 'viewpoint consistency constraint'<sup>31</sup>: the establishment of correspondence between localized features in the object and in the image constrains the relative placement of the object features and, therefore, the object's geometry.

### Varieties of alignment

Given a library of object models, each accompanied by a set of **fiducial geometric features** and a corresponding set of features in the images, one can compute the hypothetical viewing position of each candidate object, and verify the hypothesis of its presence in the image by transforming the model (aligning it to the image) and evaluating the goodness of the resulting match. Ullman<sup>7</sup> proved that the locations of as few as three features in the image suffice, under certain conditions, for unique alignment (the robustness of the method can be improved by using more features, or anchor points, than is strictly necessary). Soon afterwards, Ullman and Basri<sup>14</sup> realized that storing a few views per object (with full correspondence) obviates the need for maintaining three-dimensional models of objects. This work prompted the development of a variety of algebraic methods for view-based recognition, all based on the observation



**Fig. 1 Computational problems with structural representations.** (A) Structural descriptions must be accompanied by metric information to represent differences among commonly encountered categories. The inclusion of metric details reduces the ability of structural methods to deal with novel objects. (B) A picture of a New York City street-corner hot dog cart and a stylized object, which, as Biederman suggests<sup>8</sup>, may be described as such following a structural decomposition in the visual system. At present, there is no reliable method for mapping a grey-level image into a collection of (labelled) primitives (lines, corners, etc.) from which Recognition By Components' (RBC's) geons are constructed. Thus, although a carefully engineered system such as that described by Hummel and Biederman<sup>22</sup> can form a structural description of the line drawing of a cart-like object, the goal of deriving such a description directly from an image remains elusive. (C) Even in simpler tasks (such as character recognition, where the figure is readily separable from the ground), the derivation of a structural description is problematical. Here, the difficulty stems from the possibility to assign multiple structural descriptions to the same image. (D) In some tasks, coming up even with one structural description is problematical; how does one represent a shoe in terms of RBC's geons<sup>7</sup>?

that views (that is, vectors of image coordinates of a set of fiducial points) of a rigid object reside in a low-dimensional (linear, if the projection is orthographic) subspace of views of all possible objects.

### Computational problems

*Need for feature correspondence.* Because the establishment of correspondence is an absolute prerequisite for all the above methods, poor performance of the feature extraction and correspondence stage can disrupt subsequent recognition completely. Given the difficulty of detecting features (either points<sup>8,13</sup> or regions<sup>32</sup>) reliably in a bottom-up fashion, it seems that alignment will remain practical only in the context of industrial object identification.

*Lack of abstraction of category information.* A more serious problem with alignment-like methods is their too literal treatment of object geometry. Alignment attempts to account for the observed location of every feature of the object; in comparison, categorization of novel objects requires abstraction of geometrical detail. This seems to call for a conceptual framework that is inherently statistical<sup>33,34</sup>. Mere tolerance to certain variation in model parameters does not seem to suffice, as indicated by the relatively disappointing performance of a version of alignment that adopted this approach<sup>35</sup>.

## Box 2. Geometric constraints

*Summary:* an object is represented by the relative coordinates of a small set of its features.

*Main mathematical tools:* algebra and tensor calculus.

*Examples:* interpretation trees<sup>a</sup>, varieties of alignment<sup>b,c</sup>, linear combination of views<sup>d</sup> and trilinear tensor<sup>e</sup>.

### Strengths

- (1) Amenability to rigorous mathematical (algebraic) treatment.
- (2) Invariance to changes in viewing conditions, as long as corresponding features in both the input and the stored representation can be identified.
- (3) Demonstrable effectiveness in practical situations involving industrial objects.

### Weaknesses

- (1) Feature detection is unreliable for natural objects.
- (2) At present, there is no clear extension from identification of individual shapes to categorization of shape classes.
- (3) Feature correspondence must be established prior to recognition.

### References

- a Grimson, W.E.L. and Lozano-Perez, T. (1987) Localizing overlapping parts by searching the interpretation tree *IEEE Trans. Patt. Anal. Mach. Intell.* 9, 469-482
- b Lowe, D.G. (1987) Three-dimensional object recognition from single two-dimensional images *Artif. Intell.* 31, 355-395
- c Ullman, S. (1989) Aligning pictorial descriptions: an approach to object recognition *Cognition* 32, 193-254
- d Ullman, S. and Basri, R. (1991) Recognition by linear combinations of models *IEEE Trans Patt. Anal. Mach. Intell.* 13, 992-1005
- e Shashua, A. (1995) Algebraic functions for recognition *IEEE Trans. Patt. Anal. Mach. Intell.* 17, 779-789

*Lack of an explicit representation of object statistics.* From a statistical standpoint, alignment is deficient because it is geared to treating two objects at a time, instead of capturing several dimensions of variation within an ensemble. A step towards combining alignment with statistics has been made by Basri<sup>36</sup>, who developed an algorithm for representing object classes by their prototypes (defined as statistical averages of exemplars). However, because this method assigns the input to the closest known category (an approach known in pattern recognition as the nearest-neighbour decision), essentially, it is limited to the processing of objects that resemble one of the prototypes much more closely than any of the others.

### Multidimensional feature spaces

The limitation imposed by the use of the nearest-neighbour decision procedure can be removed by reformulating the problems of recognition and categorization as clustering in multidimensional feature spaces<sup>37</sup> (see Box 3). Importantly, this framework facilitates the representation of a novel object by its membership in a number of clusters (categories) simultaneously. The idea of using multiple reference classes may be illustrated by thinking of a giraffe: one may imagine that its ancient Roman name, *camelopardalis*, reflects the observation of its similarity to a camel (in its shape), and to a leopard (in its visual texture) (see Fig. 2).

## Box 3. Multidimensional feature spaces

*Summary:* an object is represented by a vector of feature values; the features can be geometric, photometric or any others.

*Main mathematical tools:* multivariate statistics.

*Examples:* eigenfaces and related methods<sup>a-c</sup> and histograms of local measurements<sup>d,e</sup>.

### Strengths

- (1) The features tend to be very easy to detect.
- (2) The statistical approach provides a common framework for recognition and categorization.

### Weaknesses

- (1) Structure is represented implicitly rather than explicitly.
- (2) Decision spaces tend to be high-dimensional, with the associated computational difficulties.

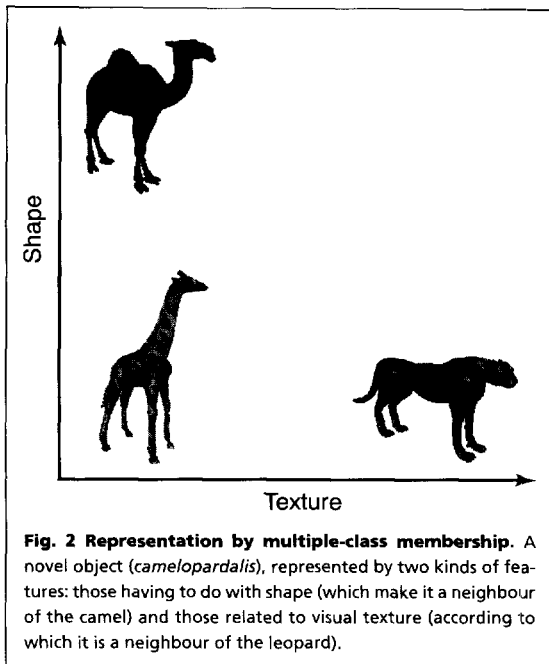
### References

- a Kirby, M. and Sirovich, L. (1990) Application of the Karhunen-Loeve procedure for characterization of human faces *IEEE Trans. Patt. Anal. Mach. Intell.* 12, 103-108
- b Turk, M. and Pentland, A. (1991) Eigenfaces for recognition *J. Cogn. Neurosci.* 3, 71-86
- c Murase, H. and Nayar, S. (1995) Visual learning and recognition of 3D objects from appearance *Int. J. Comput. Vis.* 14, 5-24
- d Mel, B. (1997) Combining color, shape, and texture histogramming in a neurally-inspired approach to visual object recognition *Neural Comput.* 9, 777-804
- e Schiele, B. and Crowley, J.L. (1996) Object recognition using multidimensional receptive field histograms, in *Lecture Notes in Computer Science* (Vol. 1) (Buxton, B. and Cipolla, R., eds), pp. 610-619, Springer

### Multidimensional histograms

A judicious choice of features in this classical pattern-recognition approach to vision abolishes the need for precise correspondence, and can lead to considerable invariance with respect to object transformations. This has been demonstrated by the success of systems that represent objects by histograms (computed over the entire input image) of multidimensional vectors of measurements related to colour<sup>16,38</sup> and to local distributions of intensity<sup>15,16</sup>.

Proponents of theories of representation based on encoding qualitative structure or metric details of objects may find it difficult to accept the possibility of representing shapes without specifying their geometry explicitly (for example, by tallying the frequency of various measurements). It seems to me that this difficulty stems from a confusion between the phenomena of perception (seeing a shape) and of representation (seeing it as something; cf. Wittgenstein's point raised in the introduction). Whereas the structure delivered by a perceptual system should be geometrically faithful to its object, there is no *a priori* computational reason to assume that this structure is to be retained in the memory trace laid down by the representational system into which perception feeds. In fact, the computational approaches discussed here and below postulate that representations are not geometrically or structurally analogous to percepts; whether or not the human visual system maintains structurally or



geometrically faithful representations of objects is an empirical question that is not addressed here.

#### Computational problems

*Combining diagnosticity with invariance.* The main problem of feature-space methods is finding features that afford reliable discrimination among similar objects, along with invariance across object transformations (this corresponds to the issue of stability vs sensitivity of features, mentioned by Marr)<sup>1</sup>. Simple geometrical arguments<sup>37,39</sup> can be used to show that these are conflicting requirements that can be met jointly only as a result of a compromise, or following special training<sup>40</sup>. That is, unless the features are both absolutely diagnostic and inherently invariant to the transformation in question, as in the case of the bar codes used to label goods in stores.

*Difficulty of learning from examples in multidimensional spaces.* Because statistical representations such as those involving feature histograms must be learned, the issue of dimensionality assumes a central role in determining the viability of any given scheme. Learning from examples in a high-dimensional space is computationally problematical. The problem, known as the '**curse of dimensionality**'<sup>41</sup>, lies in the exponential dependence of the required number of examples on the number of dimensions of the representation space. Thus, dimensionality reduction becomes of primary importance<sup>42</sup>. The challenge, then, is to reduce dimensionality while preserving the ability of the representational system to deal with novel objects, without having to come up with novel features.

#### Approximation in feature space

An effective way to increase the likelihood that most possible differences between two objects would be captured by (that is, would have a non-vanishing projection on to) at least some of the dimensions of the feature space is, somewhat paradoxically, to increase the number of features (that is, distinct measurements performed by the system on the

### Box 4. Approximation in feature spaces

*Summary:* an object is represented as a low-dimensional 'surface' (manifold) defined by prescribed points in a feature space.

*Main mathematical tools:* approximation theory and morphometrics (theory of shape spaces).

*Examples:* view space approximation<sup>a</sup> and representation by similarities to prototypes<sup>b,c</sup>.

#### Strengths

- (1) The features tend to be easy to detect.
- (2) The statistical approach provides a common framework for recognition and categorization.
- (3) Approximation by smooth functions alleviates the problems stemming from dimensionality.

#### Weakness

- (1) Structure is represented implicitly.

#### References

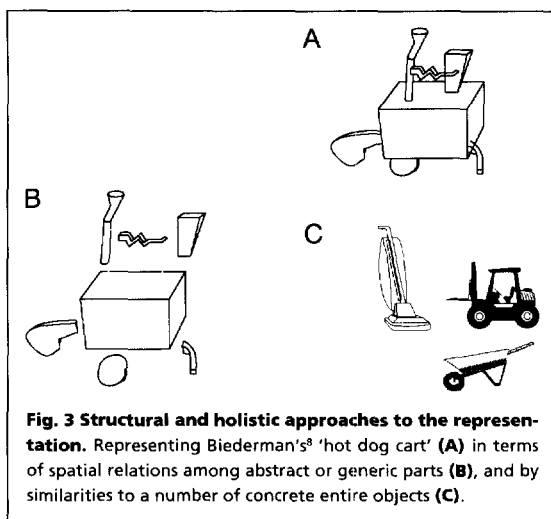
- a Poggio, T. and Edelman, S. (1990) A network that learns to recognize three-dimensional objects *Nature* 343, 263-266
- b Edelman, S. (1995) Representation, similarity, and the chorus of prototypes *Minds Mach.* 5, 45-68
- c Edelman, S. and Duvdevani-Bar, S. (1997) Similarity, connectionism, and the problem of representation in vision *Neural Comput.* 9, 701-720

image). The dimensionality of the resulting powerful yet unwieldy representation space must be reduced. The first stage in this process can be based on the observation that an object undergoing a transformation (such as rotation) is mapped into a low-dimensional '**manifold**' in the feature space<sup>43,44</sup>, provided that the measured features are smoothly related to the transformation parameters (Box 4).

Given a few points known to belong to the proper manifold (that is, a few views of the object in question), it is possible to recognize any other view of the object by interpolation. Note that the nominal dimensionality of this representation is equal to the number of views used to interpolate its view space; the true dimensionality can be much lower, and is determined by the number of degrees of freedom of the object (equal to three for objects that are only allowed to rotate in space). Although this method seems to be irrelevant to the representation of novel objects (for which no example views are available in advance, by definition), it can in fact be extended to encompass both recognition and categorization.

#### Representation by similarities to prototypes

The main requirement for such an extension is that the mechanism used for interpolating the view-space manifolds of familiar objects (prototypes) respond to unfamiliar objects as well<sup>45</sup>. Under this scheme, novel objects are represented by their similarities to the prototypical or reference shapes (Fig. 3), which, in turn, are represented by stored chosen views. The nominal dimensionality of the resulting **shape space** is determined by the number of reference shapes. An implementation of this scheme described by Edelman and



**Fig. 3 Structural and holistic approaches to the representation.** Representing Biederman's<sup>8</sup> 'hot dog cart' (A) in terms of spatial relations among abstract or generic parts (B), and by similarities to a number of concrete entire objects (C).

Duvdevani<sup>46</sup> contains ten reference-shape detection modules, each of which computes the similarity of its preferred shape to the input. The vector of similarities is then subjected to further processing (for example, compared with stored vectors to determine category membership). This system was tested on tens of grey-level shaded images of each of 50 novel objects (man-made and natural), achieving satisfactory results in both recognition and categorization.

The most prominent feature of this approach (and the source of a potentially serious shortcoming, as discussed below) is its focus on similarities among shapes. Under certain conditions, its representation of similarities is formally veridical – a property that holds a great philosophical appeal (for some of the mathematics behind this see Refs 47, 48). However, the geometry of individual shapes is not made explicit; the present method shares this characteristic with all other approaches based on abstract feature spaces.

#### *Computational problems*

*Potential proliferation of prototypes.* So far, representation by similarities to prototypes has only been implemented on a rather limited scale. It is not clear, therefore, how well this approach will scale up with the number of possible input objects. Although it is easy to program a system to acquire new prototypes at need, the rate of such acquisition must decrease to zero, and the number of prototypes must asymptote at some small fraction of the total number of objects, if the system is to be viable. Another aspect of this problem is related to the instability of interpretation that occurs in the structural methods: if too many prototypes are found to be nearly equally similar to the input, it may be difficult to choose the best subset while attempting to lower the dimensionality of the representation.

#### **Outstanding questions**

- How can the alignment approach be made to represent qualitative structure and not only quantitative (metric) details?
- How can the structural approach be made to work in practice, on real images?
- How can the feature-space approaches be made to represent structure explicitly?

*Lack of representation of structure.* A further challenge for this approach is posed by the need to make explicit the dimensions of similarity. Returning to the example in Fig. 2, it seems reasonable to require that a representational system should not only note the resemblance between the giraffe, the camel and the leopard, but that it should also realize that the resemblance of the giraffe to the camel involves certain dimensions of the giraffe's shape whereas its resemblance to the leopard involves colour and visual texture. The structural underpinnings of similarity also need to be represented explicitly. Consider two objects: a sphere attached to the top of a cube, and a cube on top of a sphere. By all accounts, both of these objects are equally similar to a sphere; one expects that their structural difference should be represented as well. It has been claimed that this can be achieved by associating similarity with certain locations in the image, and by maintaining pointers to the relevant locations along with the similarity<sup>49</sup>; this approach has yet to be tested in practice.

#### **Conclusions**

This paper presented a critical review of four theoretical approaches to object representation: structural descriptions, geometrical models, high-dimensional feature spaces and a low-dimensional representation based on similarity of the object to several prototypes. When judged by the criteria of computational plausibility and functional adequateness for the purposes of recognition and categorization, all the approaches were found to be deficient, although the nature and the severity of the problems that were identified varied from one theory to another.

On purely computational grounds, the best choice available to a designer of a visual system is probably a library of geometric models of objects accompanied by an alignment mechanism, but only if the range of tasks is restricted to the identification of one of the stored models at a time. If the system is to carry out categorization in addition to identification, a structural approach may be resorted to. Unfortunately, the extraction of fiducial features and the establishment of feature correspondence required for alignment are not always sufficiently reliable, while the structural decomposition and the matching algorithms needed for recognition by components tend to suffer from instability and from combinatorial problems. These difficulties, which are a matter of computational principle rather than implementational detail, limit the appeal of these two theories (especially as far as categorization is concerned). This, in turn, casts a certain doubt (which may or may not prove to be well-founded by an empirical investigation) on their validity as models of object representation in biological systems.

The other two theories discussed above take a completely different route to representation, by adopting the feature-space paradigm. The early work in pattern recognition abstracted away the issues surrounding the choice of features and their behaviour under object transformations. In contrast, contemporary approaches rely on feature spaces derived from studies of low-level vision, and use mathematical concepts and techniques that only recently became known outside their narrow fields of study (such as shape space approximation and dimensionality reduction). These

foundations may be partly responsible for the impressive performance of some recently implemented feature-based systems in tasks that have confounded computer vision researchers for decades.

The most intriguing outstanding question in this context is whether the feature-based representations that hold promise for identification and categorization can also be made to support tasks that require an explicit manipulation of object structure or refer to the object's parts. A potentially fruitful approach here may be to label features by their approximate origin in the image<sup>49</sup>. For example, two instances of an eye in conjunction with a nose and a mouth (all properly positioned in the image) qualify as a representation of a face<sup>50</sup>. There are similarities here to the part-based structural approach, but there are also important differences. Firstly, the features can be concrete shapes, as in the representation by similarities to prototypes. This allows for efficient recognition of structures, circumventing the problematical need for composition from generic primitives. Secondly, the features may span several levels of a hierarchy of parts and wholes, facilitating concise representations optimized for individual tasks. Thirdly, the required structural relationships hold in the two-dimensional image, not in the three-dimensional object-centred space, and can be determined by an easily implementable attention-like mechanism. It is interesting to note that these modifications to feature-based representations correspond closely to Grenander's<sup>51</sup> distinction between pattern recognition (a process whereby patterns as complete entities are attributed to various classes) and pattern theory – a computational framework that aims to account for the structure of each pattern and for the processes that may have generated it, and not merely classify it<sup>52</sup>. Further developments in this direction may lead to a more comprehensive, yet computationally plausible, computational theory of recognition and categorization.

#### Acknowledgements

I thank Nathan Intrator and Peter Dayan for useful discussions and Sharon Duvevani-Bar, Kalanit Grill-Spector and two anonymous reviewers for comments on earlier versions of the manuscript.

#### References

- Marr, D. (1982) *Vision*, W.H. Freeman
- Jolicoeur, P. and Humphrey, G.K. Perception of rotated two-dimensional and three-dimensional objects and visual shapes, in *Perceptual Constancies* (Walsh, V. and Kulikowski, J., eds), Cambridge University Press (in press)
- Logothetis, N.K. and Scheinberg, D.L. (1996) Visual object recognition *Annu. Rev. Neurosci.* 19, 577–621
- Tanaka, K. (1996) Inferotemporal cortex and object vision *Annu. Rev. Neurosci.* 19, 109–139
- Rolls, E.T. (1996) Visual processing in the temporal lobe for invariant object recognition, in *Neurobiology* (Torre, V. and Conti, T., eds), pp. 325–353, Plenum Press
- Pinker, S. (1985) *Visual Cognition* (special issue of *Cognition*), MIT Press
- Ullman, S. (1989) Aligning pictorial descriptions: an approach to object recognition *Cognition* 32, 193–254
- Biederman, I. (1987) Recognition by components: a theory of human image understanding *Psychol. Rev.* 94, 115–147
- Tarr, M.J. and Bulthoff, H.H. (1995) Is human object recognition better described by geon-structural-descriptions or by multiple-views? *J. Exp. Psychol. Hum. Percept. Perform.* 21, 1494–1505
- Biederman, I. and Gerhardstein, P.C. (1995) Viewpoint-dependent mechanisms in visual object recognition *J. Exp. Psychol. Hum. Percept. Perform.* 21, 1506–1514
- Bulthoff, H.H., Edelman, S. and Tarr, M.J. (1995) How are three-dimensional objects represented in the brain? *Cereb. Cortex* 5, 247–260
- Ullman, S. (1995) Sequence-seeking and counter-streams: a model for information flow in the cortex *Cereb. Cortex* 5, 1–11
- Lowe, D.G. (1987) Three-dimensional object recognition from single two-dimensional images *Artif. Intell.* 31, 355–395
- Ullman, S. and Basri, R. (1991) Recognition by linear combinations of models *IEEE Trans. Patt. Anal. Mach. Intell.* 13, 992–1005
- Schiele, B. and Crowley, J.L. (1996) Object recognition using multidimensional receptive field histograms, in *Lecture Notes in Computer Science* (Vol. 1) (Buxton, B. and Cipolla, R., eds), pp. 610–619, Springer
- Mel, B. (1997) Combining color, shape, and texture histogramming in a neurally-inspired approach to visual object recognition *Neural Comput.* 9, 777–804
- Wittgenstein, L. (1973) *Philosophical Investigations*, Blackwell
- Ullman, S. (1996) *High Level Vision*, MIT Press
- Marr, D. and Nishihara, H.K. (1978) Representation and recognition of the spatial organization of three dimensional structure *Proc. R. Soc. London Ser. B* 200, 269–294
- Garey, M.R. and Johnson, D.S. (1979) *Computers and Intractability: a Guide to the Theory of Completeness*, W.H. Freeman
- Lowe, D.G. and Binford, T.O. (1985) The recovery of three-dimensional structure from image curves *IEEE Trans. Patt. Anal. Mach. Intell.* 7, 320–326
- Hummel, J.E. and Biederman, I. (1992) Dynamic binding in a neural network for shape recognition *Psychol. Rev.* 99, 480–517
- Stankiewicz, B. and Hummel, J. (1996) MetriCat: a representation for basic and subordinate-level classification, in *Proceedings of 18th Annual Conference of the Cognitive Science Society* (Cottrell, G.W., ed.), pp. 254–259, San Diego, CA
- Chen, S. and Donoho, D.L. (1994) Basis pursuit, in *Proceedings of 1994 28th Conference on Signals, Systems and Computers* (Vol. 1), pp. 41–44, IEEE Computer Society Press
- Joreskog, K. and Wold, H. (1982) *Systems Under Indirect Observation: Causality, Structure, Prediction*, North-Holland
- Bienenstock, E., Geman, S. and Potter, D. (1997) Compositionality, priors, and object recognition, in *Neural Information Processing Systems* (Vol. 9) (Mozer, M.C., Jordan, M.I. and Petsche, T., eds), MIT Press
- Brooks, R. (1983) Model-based three-dimensional interpretations of two-dimensional images *IEEE Trans. Patt. Anal. Mach. Intell.* 5, 140–149
- Connell, J.H. (1985) Learning shape descriptions: generating and generalizing models of visual objects. *A.I. TR No. 853*, Artificial Intelligence Laboratory, Massachusetts Institute of Technology
- Little, J.J., Poggio, T. and Gamble, E.B., Jr (1988) Seeing in parallel: The vision machine *Int. J. Supercomp. Appl.* 2, 13–28
- Geman, S. Minimum description length priors for object recognition, in *Challenging the Frontiers of Knowledge Using Statistical Science* (in press)
- Lowe, D.G. (1986) *Perceptual Organization and Visual Recognition*, Kluwer Academic Publishers
- Basri, R. and Jacobs, D.W. Recognition using region correspondences *Int. J. Comput. Vis.* (in press)
- Kendall, D.G. (1984) Shape manifolds, metrics and complex projective spaces *Bull. London Math. Soc.* 16, 81–121
- Bookstein, F.L. (1996) Biometrics, biomathematics and the morphometric synthesis *Bull. Math. Biol.* 58, 313–365
- Shapira, Y. and Ullman, S. (1991) A pictorial approach to object classification, in *Proceedings IJCAI*, pp. 1257–1263, American Association for Artificial Intelligence
- Basri, R. (1996) Recognition by prototypes *Int. J. Comput. Vis.* 19, 147–168
- Duda, R.O. and Hart, P.E. (1973) *Pattern Classification and Scene Analysis*, Wiley
- Swain, M.J. and Ballard, D.H. (1991) Color indexing *Int. J. Comput. Vis.* 7, 11–32
- Green, D.M. and Swets, J.A. (1996) *Signal Detection Theory and Psychophysics*, Wiley
- Intrator, N. and Edelman, S. (1997) Learning low dimensional

- representations of visual objects with extensive use of prior knowledge *Network* 8, 259-281
- 41 Bellman, R.E. (1961) *Adaptive Control Processes*, Princeton University Press
- 42 Edelman, S. and Intrator, N. (1997) Learning as extraction of low-dimensional representations, in *Mechanisms of Perceptual Learning* (Medin, D., Goldstone, R. and Schyns, P., eds), pp. 353-380, Academic Press
- 43 Poggio, T. and Edelman, S. (1990) A network that learns to recognize three-dimensional objects *Nature* 343, 263-266
- 44 Jacobs, D.W. (1996) The space requirements of indexing under perspective projections *IEEE Trans. Patt. Anal. Mach. Intell.* 18, 330-333
- 45 Edelman, S. (1995) Representation, similarity, and the chorus of prototypes *Minds Mach.* 5, 45-68
- 46 Edelman, S. and Duvdevani-Bar, S. (1997) A model of visual recognition and categorization *Philos. Trans. R. Soc. London Ser. B* 352 (1358), 1191-1202
- 47 Edelman, S. Representation is representation of similarity *Behav. Brain Sci.* (in press)
- 48 Edelman, S. and Duvdevani-Bar, S. (1997) Similarity, connectionism, and the problem of representation in vision *Neural Comput.* 9, 701-720
- 49 Edelman, S. (1994) Biological constraints and the representation of structure in vision and language *Psychology* Vol. 5
- 50 Riesenhuber, M. and Dayan, P. (1997) Neural models for the part-whole hierarchies, in *Advances in Neural Information Processing* (Vol. 9) (Jordan, M., ed.), pp. 17-23, MIT Press
- 51 Grenander, U. (1993) *General Pattern Theory*, Oxford University Press
- 52 Mumford, D. (1994) Neuronal architectures for pattern-theoretic problems, in *Large-scale Neuronal Theories of the Brain* (Koch, C. and Davis, J.L., eds), pp. 125-152, MIT Press
- 53 Stone, C.J. (1982) Optimal global rates of convergence for non-parametric regression *Ann. Stat.* 10, 1040-1053
- 54 Richards, W., Jepson, A. and Feldman, J. (1996) Priors, preferences and categorical percepts, in *Perception as Bayesian Inference* (Knill, D. and Richards, W., eds), pp. 93-122, Cambridge University Press
- 55 Poggio, T., Torre, V. and Koch, C. (1985) Computational vision and regularization theory *Nature* 317, 314-319
- 56 Kendall, D.G. (1989) A survey of the statistical theory of shape *Stat. Sci.* 4, 87-120

# Risk-sensitivity: crossroads for theories of decision-making

Alex Kacelnik and Melissa Bateson

**Most actions result in one of a set of possible outcomes. To understand how this uncertainty, or risk, affects animals' decision-making some researchers take a normative approach, asking how an animal should respond to risk if it is maximizing its fitness. Others focus on predicting responses to risk by generalizing from regularities in behavioural data, without reference to cognitive processes. Yet others infer cognitive processes from observed behaviour and ask what actions are predicted when these processes interact with risk. The normative approach (Risk-sensitivity Theory; RST) is unique in predicting a shift in a subject's response to risk as a function of its resource budget, but the predictions of this theory are not yet widely confirmed. In fact, evidence suggests a strong bias towards risk-proneness when delay to reward is risky and risk-aversion when amount of reward is risky, a pattern not readily explained by RST. Extensions of learning theory and of Scalar Expectancy Theory provide process-based explanations for these findings but do not handle preference shifts or provide evolutionary justification for the processes assumed. In this review we defend the view that risk-sensitivity must be studied with theoretical plurality.**

Under natural circumstances, most actions have a set of possible consequences, rather than a single well-defined outcome. A choice is described as risky if the probabilities of the different possible outcomes are known, but the precise outcome is not known. The best that an animal can do when faced with a risky choice is to use information about the probabilities of different outcomes to decide which action to perform. For evolutionary biologists committed to the study of decision-making in animals, this raises an in-

triguing and challenging question: has natural selection produced sensitivity to risk, and if such an adaptation exists, what form does it take? Posed in this form, the question of risk-sensitivity belongs in the normative field of behavioural ecology. However, recently it has become clear that a number of different fields of behavioural research have converged on the problem of how animals respond to risk, each bringing a different perspective on how the behavioural phenomena should be explained. Here, we review some of

Alex Kacelnik is at the Department of Zoology, University of Oxford, South Parks Road, Oxford, UK OX1 3PS and Melissa Bateson is at the Department of Psychology, Experimental Duke University, Durham, NC 27706, USA.

tel: +44 1865 273164  
fax: +44 1865 271120  
e-mail: alex.kacelnik@zool.oxford.ac.uk