resentations become to match world statistics (Carpenter & Grossberg 1991). Other models in which bottom-up and top-down processes are used (e.g., Back Propagation and the Helmholtz Machine) do not yet have these properties.

Edelman criticizes winner-take-all decisions because they violate the "principle of least commitment," but such decisions can quantitatively simulate categorical perception data (e.g., Grossberg et al. 1997a). ART systems such as masking fields (Cohen & Grossberg 1986), ART-EMAP (Carpenter & Ross 1995), Distributed ARTMAP (Carpenter 1996), and Gaussian ARTMAP (Williamson 1996) also show how distributed codes may improve recognition, and how the distribution reflects data uncertainty. Gaussian ARTMAP in particular is a self-organizing RBF (radial basis function) production system.

Self-organizing view-invariant 3-D object categories fuse view-specific categories in ARTMAP systems (e.g., Bradski & Grossberg 1995), as in the IT data reviewed in section 7.2. The 3-D categories occur in the Map Field, wherein outputs from multiple categories, whether of different letter fonts or different object views, are adaptively fused.

Edelman's measurements and dimensionality reduction stages are typically called vision and learned recognition stages. Although ART top-down matching occurs in the vision system, even as peripherally as the LGN (Gove et al. 1995; Grossberg et al. 1997b), vision uses principles and circuits different from those of the recognitions system. Edelman describes measurement as "a convolution with a number of filters, followed by the application of a nonlinearity," including light source compensation and figure-ground separation. Cortical models of visual perception, called FACADE models, suggest additional mechanisms (e.g., Arrington 1994; Chey et al. 1997; Francis & Grossberg 1996; Gove et al. 1995; Grossberg 1994; 1997; Grossberg et al. 1997b; Grossberg & Todorovic 1988). For example, parallel processing streams for boundary representation (interblob stream) and surface representation (blob stream) compute complementary computational properties. Feedback between these streams assures their mutual consistency and initiates figure-ground pop-out. Diffusive filling-in completes surface representations from signals that discount the illuminant.

Edelman summarizes a sensible approach to representation, but one that is limited by its feedforward character. ART models self-organize stable representations that achieve second-order isomorphism to arbitrarily large and changing environments, but only by using learned top-down expectations, attention, and memory search. FACADE models have clarified a lot of data about vision, but only by introducing new concepts about how complementary streams of boundary, surface, and motion processes achieve mutual consistency and coherence using other types of feedback. A major intellectual watershed separates feedforward models from self-organizing feedforward/feedback models. This watershed needs to be crossed for a deeper understanding of how humans autonomously form representations of the real world.

The notion of distal similarity is ill defined

Ulrike Hahn and Nick Chater

Department of Psychology, University of Warwick, Coventry CV7 4AL, United Kingdom. **u.hahn,n.chater@warwick.ac.uk**

Abstract: We argue that the notion of distal similarity on which Edelman's reconstruction of the process of perception and the nature of representation rests is ill defined. As a consequence, the mapping between world and description that is supposedly at stake is, in fact, a mapping between two different descriptions or "representations."

Edelman has shown experimentally that people can extract the underlying parameters used to generate a set of novel stimuli. From the results of multidimensional scaling, he conjectures that the internal space that people recover represents these parameters. This implies that nearby points in the original parameter space are near in the mental space, and it is short step from this to saying that similarity is preserved between the two spaces. Such results are not surprising where the dimensions of variation in the objects are subjectively obvious (e.g., the length and orientation of line segments), and in such cases this correlation between parameter space and mental space is frequently found. But it is impressive with Edelman's stimuli, where the underlying dimensions of variation are far from obvious and interact in a complex way to produce the visual image.

Edelman moves from these results to a general theory of perception founded on similarity. He presents this as an alternative to a "reconstructionist" approach. The goal of perception is assumed to be preserving similarities between things in the environment, rather than building an internal representation of environmental structure. Edelman's target article is important and should act as a valuable stimulus for future research. We believe, however, that there are three difficulties with this viewpoint as a general program in perception.

(1) The notion of "distal" similarity seems ill-defined. Goodman (1972) pointed out that any two objects have infinitely many common and distinctive features, thus "objectively" everything is equally similar to everything else. Watanabe (1985b) illustrates that even choosing for a set of objects only those predicates that are extensionally distinct (which for a finite set of objects is a finite set of predicates) still leaves all between-object similarities equal, unless differential weights for predicates are introduced. This is not just a philosophical nicety. In Edelman's experiments, stimuli are generated artificially by varying a set of parameters; thus nearness in parameter space may be chosen as a reasonable measure of similarity.

But the natural world has not been generated by manipulating a small number of underlying parameters. Variation in natural objects can be considered along a limitless number of dimensions. By choosing (and assigning differential weights to) any subset of these dimensions, all manner of "distal" similarities can be generated. Objects may be compared by overall color, by outline shape using any number of shape representation systems, by nearness to the observer (or to Pluto!), by weight, by perimeter length, and so on, indefinitely. Moreover, any of these measurements can be combined in arbitrary ways (e.g., perimeter length times weight) to produce new measures that can be used to give new dimensions.

Any set of any dimensions seems equally good as a distal measure of similarity. It might be suggested, for example, that physics could supply constraints on what can count as an underlying dimension, but it should be clear that this still leaves an infinite number of possible dimensions along which objects in the environment might be assessed; moreover, it will rule out many psychologically critical dimensions (e.g., the dimensions that define facial structure) since these do not relate to physical quantities. In short, it does not make sense to say that two things are similar without specifying in what way they are similar (Goodman 1972); to specify this, however, requires a cognitive agent to *define* which dimensions of distal variation matter and which do not; then the relation between an "objective" distal similarity structure and the similarity structure in the internal space of an agent breaks down. This means the claim that the perceptual system preserves an objective distal similarity structure loses its sense. Edelman, rather than dealing with objective properties of the world, is dealing with *two different descriptions* or representations – an experimenter-intended one (the underlying parametrization) and one formed by participants (the internal similarity spaces).

The situation seems analogous to the general philosophical difficulty with the correspondence theory of truth: there is no "mind-independent" way to specify which *facts* the world consists of, so the claim that true statements correspond to these facts is circular. In exactly the same way, there is no "mind-independent" way to specify which are the *similarities* in the world, so the claim that similarities in mental space correspond to these external similarities is circular. But if there are no distal similarities, there can be no second-order isomorphism on which to build a theory of representation. The debate about the correspondence theory of truth as stated by us is a philosophical classic. The point we are making – that there is no "picture" relationship between statements and world – is widely accepted (see Strawson, Ayer, Wittgenstein II) even within logical positivism (for example, Neurath).

(2) Perception frequently appears to involve classifying very different patterns as similar. For example, the sequences 101010101010 and 010101010101 appear similar, even though they differ at each spatial location. Similarly, a photograph and its negative will be judged similar, even though they differ in every pixel value. Or again, different pictures of the same face, or different tokens of the same phoneme, will seem very similar, even if, under some obvious physical description, they appear completely different. The point is that the perceptual system identifies the common structure in both stimuli. How does this relate to Edelman's claim that distal structure is preserved in the internal representation of similarity? Using some obvious physical interpretation of the stimulus, the objects are very different, yet they are judged to be very similar, violating Edelman's theory. But using, instead, a perceptually appropriate description for measuring "distal" similarity (e.g., that the stimuli above are both examples of alternating patterns: descriptions in terms of the structure of a face or the identity of a phoneme), the similarities between the distal world and the mind are preserved, but only at the cost of circularity.

(3) Finally, we suggest that the reconstructive approach to perception may not be an *alternative* to Edelman's similarity-based view of perception. Instead, a reconstruction of the perceptual world may be required to explain why the similarities are judged as they are. For example, with Edelman's artificial figures, the parameters of variations may be of interest as part of a specification of the structure of those figures – indeed, only by attempting to reconstruct those figures does it seem possible to realize that there *are* only a small number of underlying parameters of variation (i.e., the recipe for reconstructing each figure is the same, apart from parametric variation). Thus, the parametrization used as a basis for internal similarity judgments may be *based on* the attempt to reconstruct the figure. For example, it is not clear why two pictures of the same face will be judged to be similar unless the same underlying 2/3D structure has been reconstructed (at least partially) for both. Thus, we would argue that the reconstructionist view of perception may be an important component in an account of similarity of relevance to Edelman's empirical results.

Representation of similarities and correspondence structure

Nathan Intrator

School of Mathematical Sciences, Tel-Aviv University, Ramat-Aviv 69978, Israel. **nin@math.tan.ac.il www.math.tan.ac.il/~nin**

Abstract: Apart from the computationally appealing properties of representation by similarities, it is possible to extend this form of representation when needed to include object parts as well as the correspondence between subobject parts.

Edelman provides a solid theory about object representation and its consequences. The idea of representing an object as a vector of distances from several other reference objects is very appealing on computational grounds and demonstrates a simple and probably robust dimensionality reduction. It further suggests a simple algorithm for hierarchical clustering, in which whenever a "sufficiently different" object appears, it may be registered as a new prototype, and when an object that is "not very different" appears and its class label is unexpected, it is again registered as a new prototype.

I would like to elaborate on the issue of "holistic" features versus the feature-representation that correspond to subparts in objects (sect. 6.3). It is a fundamental question in object representation not only whether there is a need to represent objects as wholes or as combinations of features, but also whether the exact topographic relation between subobject parts is essential. There is no doubt, for example, that there is a big difference between a phone that is on or off the hook, although this may be a very small difference in object space. This example demonstrates the need for an explicit feature-based representation with topographic correspondence, but as it would be difficult to argue that there is a prototype for an off-the-hook phone, holistic representations may coexist with more elaborate feature-based representations. If these representations do coexist, then it is likely that those based on prototypes are more specific but computationally simpler and are hence used for very repetitive (everyday) tasks, or tasks that require fast responses. The more elaborate representation is appropriate when the correspondence between object parts is important, for example, to represent walking or running.

One could argue that when a certain part of an object appears to have higher weight for purposes of recognition or discrimination, then that object part can be represented as a prototype or a distance vector from prototypes. The correspondence between object parts carries information that is very important and useful for classification and discrimination (Geman et al., forthcoming). In the case of representation by similarity, the exact relation between subparts and the object (the binding together of object parts) can be encoded via temporal structure such as synfires (Abeles 1981).

The representation of objects as a vector of distances from several prototypes suggests a very simple method for mental object manipulations, in which creating a mental representation of a certain object simply requires stimulating one (or more) of the prototype cells representing an instance of that object.

In summary, it appears that the simple object representation proposed by Edelman is compatible with the need for binding between subparts. Future psychophysics will clarify whether object representation via subparts coexists with holistic representation and whether the binding problem can be addressed by holistic representations and temporal structure.

Representation of similarities – a psychometric but not an explanatory concept for categorization

Martin Jüttner

Institute for Medicinal Psychology, University of Munich, Munich, D-80336, Germany. **martin@imp.med.uni-muenchen.de lrz.muenchen.de/~u7fo1bg/www/**

Abstract: The representation of similarities is a viable concept for a cognitive extension of visual psychophysics to the recognition of shapes, bringing issues such as similarity and categorization back into that field. However, as a framework it appears too general to place constraints on a particular process model for categorization. In particular, a preference for Chorus-like schemes with respect to structure-oriented approaches is unwarranted.

Edelman's conception can be regarded, on a theoretical level, as an extension of classical multidimensional scaling (MDS) to the recognition of shapes. To evaluate the potentials and limitations of such an undertaking, it is useful to recapitulate one of the basic motivations for MDS: it has been observed repeatedly that the probability that a learned response to any stimulus will generalize