The 7th Annual Workshop on Object Perception and Memory

Proceedings

November 18,1999 http://vision.psych.umn.edu/www/opam99/

Organized by:
Brian J. Stankiewicz
University of Minnesota
Thomas Sanocki
University of South Florida

MODELING THE DISCOVERY OF ABSTRACT INVARIANTS

Philip J. Kellman, Timothy Burke and John E. Hummel

UCLA

We present the beginnings of a model of the human capacity to learn abstract invariants, such as square. The model is founded on three primary assumptions, which we believe to be neurally plausible and generic: Metric space, Topology, Comparison operations (subtraction, greater-than/less-than), and Extraction of vertices. The model successfully learns to discriminate simple planar quadrilaterals, and generalizes that learning across variations in viewpoint and modest variations in shape.

A key feature of human information processing is the ability to detect abstract invariants. Many important outputs of visual perception are descriptions of shapes and spatial relations. As the Gestalt psychologists pointed out early in this century, these are relational notions. A square, or a melody, is not definable in terms of any particular elements, but in terms of invariant relationships. We not only encode characteristic relations, but we discover new ones with experience. This ability is crucial in the growth of expertise in many domains.

Perceptual learning of abstract relations remains deeply mysterious. How does the visual system discover abstract invariants, such as squareness, roundness, or parallelism? By abstract invariant we mean a property that, while computable from, is not definable in the vocabulary of primitive features from which it is derived. For example, no logical concatenation of neural responses in primary visual defines the invariant squareness. Squareness is both more and less than any finite set of such features. It is more because some new activation pattern might also form a square, and it is less because many of the attributes of any given activation pattern (e.g., its precise location, size and orientation) have nothing to do with their squareness.

We assume that an invariant such as "squareness" is not detected simply by constructing a large number of detectors for specific squares and then summing their outputs. Such an approach cannot account for the human ability to generalize to new examples. Although perceptual learning of abstract relationships is well documented (e.g., Chase & Simon, 1973; Gibson, 1969; Goldstone, 1998), little modeling has addressed the learning of abstract invariants. Our aim is to model invariant detection and learning in shape perception.

Assumptions

The openendedness of the human ability to learn abstract invariants calls for a system formally like a grammar with recursive rules. Perceptual learning of abstract invariants may depend on some basic set of relations and some processes for concatenating them. A key challenge is to postulate only those steps for recoding or finding relations that can be justified on independent grounds. Success will not consist of finding a special-purpose processor for squareness, but allowing squareness to emerge naturally from basic operations that we would expect, on independent grounds, to be part of visual processing. Using this approach, we can develop the fundamental grammar of invariant processing and ultimately test the scope and limits of what is learnable.

The vocabulary of a "visual grammar" consists of a finite set of basic operations corresponding to (presumably innate) assumptions on the part of the visual system about the nature of the visual world. We assume that the visual system comes equipped with (at least) the following knowledge/capacities:

I. Metric Space: Neurons in early visual processing (e.g., retina, LGN, V1) "know" (perhaps implicitly via their interconnections) about metric space: They know that their receptive fields correspond to finite regions of larger metric space, and they know (approximately) where in that space their receptive fields are located. This knowledge manifests itself in a neuron's (or hypercolumn's) ability to signal its location independently of any of its other properties, namely by activating other neurons representing location (e.g., in Euclidean coordinates) independently of the nature of whatever visual features happen to reside there. A corollary is that neurons in early visual processing "know about" their adjacency relations (and possibly other topological relations). This kind of knowledge is manifest in the local lateral connectivity among, for example, neurons in visual area V1.

II. Difference Operations: The nervous system is equipped with routines for performing difference operations, including subtraction, and evaluating greater-than and less-than relations.

III. Vertex Finding: Early visual routines find vertices and other local changes in contour.

These three assumptions form the foundation of the current effort. Our goal is to demonstrate that these assumptions are sufficient to bootstrap the learning of abstract invariants.

The Model

The assumptions are embodied in a four-layer "neural"-style network. Units in the first layer represent retinal coordinates of the vertices in an image of a quadrilateral. We assume that the vertices are detected and their spatial coordinates registered by an early preprocessing stage (Assumptions III and I, above). Units in the second layer compute the pair-wise Euclidean distances between the coordinates coded in the first layer (Assumptions I and II). In the current implementation, each unit represents one distance (e.g., there is one unit for the distance between vertex 1 and vertex 2, another for vertices 2 and 3, etc.), and distance is represented as activation. That is, in Layer 2 layer, distance is rate-coded. Coordinates are "read into" the model in a fixed order, starting from some corner on the stimulus, and proceeding around the figure clockwise. This convention is a simplified implementation of our more general assumption that the system knows which vertices are connected to which by virtue of an intervening contour (Assumption I, topological relations). (In the current model, implicit knowledge of sequential order of vertices in a connected figure is important, although the particular starting point in the sequence is not.)

Units in the third layer take their inputs from pairs of distance units in the second layer, and compare them for equality (Assumption II). These comparisons serve as the input to the fourth layer of units, which learn to classify their inputs as four-sided geometrical figures (squares, parallelograms, trapezoids, etc.).

Simulations

We trained the model on seven figures: a square, a rhombus, a rectangle, a parallelogram, two trapezoids and an irregular quadrilateral. After training, we tested the model for its ability to recognize these figure at new sizes, orientations and locations in the visual field. We also tested the tolerance of the modelís classification performance to varying amounts of shape distortion.

The model successfully recognized the figures at new orientations, sizes and locations, and, similar to human observers, it tolerated modest but not extreme distortions in assigning shape labels.

Discussion

Our prototype shape network can discover, from plausible building blocks, the abstract invariants that determine a simple shape classification. Building on this foundation, we hope to discover the visual grammar and computational processes that make possible and constrain human perceptual learning.

References

Chase, W. & Simon, W. (1973). Perception in chess. *Cognitive Psychology*, *4*, 55-81.

Gibson, E.J. (1969). *Principles of perceptual learning and development*. New York: Prentice-Hall.

Goldstone, R.L. (1997). Perceptual learning. *Annual Review of Psychology*, *49*, 585-612.