

A Neural Model for Category Learning*

Douglas L. Reilly, Leon N. Cooper, and Charles Elbaum

Center for Neural Science and Department of Physics, Brown University, Providence, RI 02912, USA

Abstract. We present a general neural model for supervised learning of pattern categories which can resolve pattern classes separated by nonlinear, essentially arbitrary boundaries. The concept of a pattern class develops from storing in memory a limited number of class elements (prototypes). Associated with each prototype is a modifiable scalar weighting factor (λ) which effectively defines the threshold for categorization of an input with the class of the given prototype. Learning involves (1) commitment of prototypes to memory and (2) adjustment of the various λ factors to eliminate classification errors. In tests, the model ably defined classification boundaries that largely separated complicated pattern regions. We discuss the role which divisive inhibition might play in a possible implementation of the model by a network of neurons.

I. Introduction

A common concern of neural models has been the problem of relating the function of complex systems of neurons to what is known of individual neurons and their interconnections. In this paper we discuss a neural model that displays a form of learning manifested in human behavior: supervised learning of pattern categories. The terms pattern and event are used here synonymously to refer to a state of the environment that is characterized by a set of measurements. A category of patterns is a set of patterns in the same class. Their members may yield "roughly" the same value for some measurement (or collection of measurements) made on them (e.g. with reference to some feature set). However, one can imagine a category resulting from an association between a collection of

very unlike events and a particular system response (e.g., calling "a" and "A" by the sound of the first letter in the alphabet). In this case, the criterion defining the category is the association itself.

There are several difficulties in the problem of pattern classification that we address here. A given pattern class appears in the primary sensory neurons in a vast variety of manifestations. Consider all of the recognizable distortions of the Arabic numeral "three". All of these must be classified as "three" and at the same time be distinguished from other classes (1, 2, 4, etc.) and all of their distortions. Therefore, the problem of classification involves a separation of "different" classes as well as a grouping together of all distorted members of the same class. Our model is capable of making the separation as well as the grouping with a simple instruction procedure that seems at least roughly comparable to that employed in human learning.

There is a growing body of research dealing with the general problem of learning in an adaptive system composed of neuron like elements. Early work in this field introduced the notion of correlation matrix memories, showing how it was possible for a system to learn associations between pairs of input and output vectors ($\mathbf{x}^i, \mathbf{y}^i$) (Kohonen, 1972). Category learning has frequently been viewed as learning an association between \mathbf{y}^i and a set of noisy versions of \mathbf{x}^i . Models for such concept formation have been proposed which make use of varying amounts of interaction with an external "teacher" (e.g., Amari, 1977; Grossberg, 1978; Barto et al., 1981; Bobrowski, 1982). Among the various approaches in such systems, learning rules incrementally adjust elements of some weight vector \mathbf{w} whose inner product with the input \mathbf{x} is an important contributing factor to the output of the system.

In our approach pattern classification is accomplished through prototype formation. Evidence from psychological experiments suggests that learning of pattern classes might involve abstraction of a pro-

* This work was supported in part by the Alfred P. Sloan Foundation and the Ittleson Foundation, Inc.




Fig. 1. Architecture of the model. Shown are coding neurons (F), prototype cells (G), classification cells (H), mapping (A) from F to G , mapping (B) from G to H , and the external instructor (T). Arrows mark information flow

prototype to represent a category of stimuli (e.g., Posner and Keele, 1968, 1970; Franks and Bransford, 1971). Some knowledge of class variance must also be learned. A closely related argument holds that categories are learned by retaining in memory examples of each class (e.g., Brooks, 1978; Medin and Schaffer, 1978). In pattern recognition theory, the technique of nearest neighbor classification is effectively an exemplar learning scheme (Cover and Hart, 1967; Duda and Hart, 1973). The focus of algorithms for such training has been to find and store the example set of minimal size which can guarantee performance within some acceptable error rate.

Here no distinction is made between the single (prototype) and multiple exemplar theory. Any class member stored in memory will be referred to as a prototype for that class. We will discuss learning in a system of neurons and, in particular, a model for prototype formation and development in a class of distributed memory neural networks.

II. Overview

In the architecture we consider, afferents from coding neurons, F , project onto prototype cells, G , which in turn synapse with classification neurons, H (see Fig. 1). Each class of events will be represented by the activity of a unique H neuron. An input event is coded by a vector of firing rates (\mathbf{f}) in the F bank. If it causes activity in an H cell, it is classified as belonging to the category associated with that cell.

We define four possible network responses to an input pattern. Let $\mathbf{f}(c)$ represent an incoming pattern belonging to the c^{th} class of events, and let \mathbf{h} be the vector of output firing rates of the H neurons. Further, let \mathbf{h}^x be defined as a vector with components

$$(\mathbf{h}^x)_j = \delta_{xj}. \quad (1)$$

Table 1. Classification of system responses for various values of α and Q as defined by (2)

Classification	α	Q	\mathbf{h}
Correct	1	0	\mathbf{h}^c
Unidentified	0	0	0
Incorrect	0	1	$\mathbf{h}^r, r \neq c$
Confusion	0	≥ 2	$\sum_{r \neq c}^Q \mathbf{h}^r$
	1	≥ 1	$\mathbf{h}^c + \sum_{r \neq c}^Q \mathbf{h}^r$

The response \mathbf{h} can be written, in general, as

$$\mathbf{h} = \alpha \mathbf{h}^c + \sum_{r \neq c}^Q \mathbf{h}^r. \quad (2)$$

If $\alpha=1$ and $Q=0$, then the system has correctly classified the input pattern. A response characterized by $\alpha=1$ and $Q \geq 1$, or $\alpha=0$ and $Q \geq 2$, we refer to as confusion, since the system is unable to decide upon any of several pattern classifications. The case where $\alpha=0$ and $Q=1$ is an outright incorrect response. When $\alpha=0$ and $Q=0$, no categorization has been made and the pattern is unidentified. Table 1 summarizes the responses.

The synaptic connections between G and F are represented in the mapping A . In our learning models, a prototype for a class is "imprinted" on the synapses between a G cell and the F set, thus becoming the most effective stimulus for that cell. For any given class, there may be more than one prototype; each will correspond to a different G_i . The mapping B , between cell groups G and H , develops so that the subset of G cells which can cause a given H cell to fire consists of prototypes representing the same class. A sufficient stimulus for an H cell to fire will be supraliminal activity in any member of its corresponding G cell subset.

The H set of neurons (and indirectly, G) has an additional source of input, that diagrammed by the block T in Fig. 1. Through T , an external supervisor can correct the network classification responses. The specific form of the mapping B , along with some aspects of A , will develop as a result of interaction with input patterns and with T . Essentially, T can cause the commitment of a G cell to a prototype and the strengthening of the association between this G cell and the proper classification cell. We assume synaptic modification as the vehicle for these network changes. One can imagine a variety of ways in which synaptic changes can result in cell coupling between the G and H sets. For example, simple Hebbian modification can produce the desired association if the particular H cell was receiving concurrent stimulation from T . The only requirement of this procedure is that cell commitment

never involve a previously committed cell. For simplicity, we further assume that

(1) cell commitment is rapid (i.e., occurring within the duration of event presentation)

(2) only one cell is committed to any one prototype.

In the mapping A , an element A_{ij} represents the logical synapse between G_i and F_j ; i.e., it summarizes the total effectiveness of neuron F_j in firing G_i . In accordance with a distributed memory model studied by Anderson and by Cooper, among others (e.g., Anderson, 1970, 1972; Kohonen, 1972, 1977; Cooper, 1973; Nass and Cooper, 1975; Anderson and Cooper, 1978), we take the firing rate of G_i (call it g_i) to be a weighted sum of the firing rates of the F neurons (f_j), gated by some threshold function

$$g_i = \Theta\left(\sum_j A_{ij}f_j\right), \quad (3)$$

where

$$\begin{aligned} \Theta(x) &= 0 & \text{if } x \leq \theta \\ &= x - b & \text{if } x > \theta. \end{aligned} \quad (4)$$

Given a prototype $\mathbf{P}(c)$ representing a class c of inputs, the equality

$$A_{ij} = P_j(c), \quad \text{all } j \quad (5)$$

establishes a correspondence between the i^{th} G cell and a particular class of patterns c . The synapse vector of G_i takes on the value of the prototype.

Each prototype cell has a "region of influence" in the input space of events. It is defined as the set of input patterns that satisfies the threshold condition for cell firing. For convenience, assume input events to be normalized ($\mathbf{f} \cdot \mathbf{f} = 1$). The region of influence defined by cell G_i with threshold θ is the intersection of the surface of a unit hypersphere with a cone of angular width γ , $\gamma = \cos^{-1} \theta$,

where γ is the angle between $\mathbf{P}(c)$ and an input \mathbf{f} at threshold.

A class of patterns defines a region or set of regions in the pattern space of input events. Class regions corresponding to different pattern categories are assumed to be strictly disjoint. A priori, we choose not to restrict the complexity that the shape of class boundaries may display. To identify the class of an input event, the neural network must characterize and learn the arrangement of class regions. Our model develops by itself a set of prototypes whose influence regions map out the areas belonging to different categories in the pattern space without prior information of what these areas are. One approach to such prototype organization will be discussed. Several others, differing

in their methods of cell modification and in their assumptions about interaction between G cells, or equivalently, between prototypes stored in memory will be discussed elsewhere.

III. Prototype Formation and Development

For the present, we continue the assumption of normalized input patterns ($\mathbf{f} \cdot \mathbf{f} = 1$). Each committed prototype cell has a synapse vector of the form (for the i^{th} cell),

$$\mathbf{A}^i = \lambda_i \mathbf{p}^i, \quad (7)$$

where \mathbf{p}^i is a normalized ($\mathbf{p}^i \cdot \mathbf{p}^i = 1$) prototype vector and $\lambda_i > 1$. The vector \mathbf{p}^i corresponds to some previously seen input pattern whose presentation failed to excite the H cell of the appropriate class. Modification to prototype cell synapses is governed by the following conditions.

1. New Classification

If $\mathbf{f}(c)$ is presented and

$$\mathbf{h} \cdot \mathbf{h}^c = 0 \quad (8)$$

i.e., the H cell for the c^{th} class does not fire, then a new G cell (call it G_k) is committed to $\mathbf{f}(c)$ and the synapse between G_k and H_c is assigned strength 1. The synapses of G_k with F are modified according to

$$A_{kj} \rightarrow P_{kj} = \lambda_0 f_j, \quad (9)$$

where $\lambda_0 > 1$.

2. Confusion

If presentation of $\mathbf{f}(c)$ causes firing rate activity in some H_w where $w \neq c$, then this results in a signal from the T channel to reduce the λ factors of each currently active G cell associated with H_w . The quantity λ is diminished until the response of the cell to $\mathbf{f}(c)$ lies at threshold. If G_r is such a unit, then

$$\lambda_r \rightarrow \lambda'_r$$

such that

$$\lambda'_r \mathbf{p}^r \cdot \mathbf{f}(c) = 1. \quad (10)$$

For convenience, we have taken $\theta = 1$.

These two rules for prototype acquisition and modification will enable the network to learn the geography of the pattern classes.

In an untrained network, all G cells are uncommitted. The strengths of the synapses between G and H are all zero or some arbitrarily small number. When a pattern $\mathbf{f}(c)$ is presented to this system, no H cell

responds above threshold. Information from the T element enters the system, identifying the correct class of the input. A single G cell is committed to $\mathbf{f}(c)$ as a prototype for that class and, simultaneously, the synapse between this G cell and H_c is set equal to 1. Since this input represents the first example of any pattern class, we can let $c = 1$. If the same pattern were to be presented again to the system, the response of the G cell would be

$$\lambda_0 \mathbf{p}^1(c) \cdot \mathbf{f}(c) = \lambda_0 > 1. \quad (11)$$

The output signal, λ_0 , from this G cell would cause H_c to fire.

Suppose a second pattern $\mathbf{f}^2(c')$ is presented to the system. Assume $c' = c$. If

$$\lambda_0 \mathbf{p}^1(c) \cdot \mathbf{f}^2(c) > 1 \quad (12)$$

then H_c will fire and the pattern will be correctly classified. Thus no change occurs. If

$$\lambda_0 \mathbf{p}^1(c) \cdot \mathbf{f}^2(c) < 1 \quad (13)$$

then $\mathbf{f}^2(c)$ will be committed to a new G cell [prototype $\mathbf{P}^2(c)$] and the synapse between this G cell and H_c will be set equal to 1. In this way, a class can be characterized by more than one prototype.

Consider the situation in which $c' \neq c$. Whether or not the existing prototype cell fires past threshold, there will be no active H cells of the class of \mathbf{f}^2 . The subsequent T signal causes a new prototype cell to be committed to \mathbf{f}^2 , along with the setting of the synaptic connection between this G cell and a new H cell. If, in addition,

$$\lambda_0 \mathbf{p}^1(c) \cdot \mathbf{f}^2(c') > 1 \quad (14)$$

then λ_0 is reduced to λ_1 such that

$$\lambda_1 \mathbf{p}^1(c) \cdot \mathbf{f}^2(c') = 1. \quad (15)$$

As the system learns, the λ factors associated with any active incorrect class prototypes will be reduced, leaving only the correct H cell to respond to the pattern.

The strategy of this network learning scheme is made clearer by considering the problem geometrically. The size of the influence region of a prototype cell is directly proportional to the magnitude, λ , of the prototype. Class territories in the space of events are defined by covering them with the overlapping influence fields of a set of prototypes drawn from class samples. Should the influence region of a given prototype extend into the territory of some differing class to the point of incorrectly classifying or confusing a member of that class, the λ factor of the prototype is reduced until its region of influence just excludes the disputed pattern. Prototype modification only decreases λ factors. Influence fields of existing prototype

cells are never enlarged in an effort to include (classify) an event, since for many of these elements, even slightly larger regions of influence have previously resulted in incorrect identifications. Consequently, a pattern that is excluded from the influence regions of all existing prototypes for its class is an occasion for commitment of a new G cell, with the pattern assuming the role of the new prototype.

Note that the prototype cells in memory are completely decoupled in that there are no mutual inhibitory or excitatory interactions among them. In the network's classification response, there is no vote counting among prototypes. The activity of a single prototype cell counts as heavily as the possibly concerted activity of a set of prototype cells, all specific to some other class.

This model was tested in computer simulations using a design set of input patterns. The patterns were vectors randomly generated in a normalized three dimensional pattern space. Samples were constrained to lie on the top half of a unit sphere ($z > 0$) and represented two classes of patterns labelled A and B . In one arrangement the A region was chosen as a spherical cap centered on the z axis and ringed by the B region, a surrounding band on the sphere's surface. The projection of this design is a pair of concentric circles on the $x-y$ plane. A second geometry pictured the A and B regions as separated by a sinusoidal boundary on the sphere's surface.

Patterns arrived in cycles (trials). A trial consisted in presentation of 200 novel A vectors and 200 novel B vectors, randomly distributed with respect to class. After some number of trials, the distribution of prototypes was graphed together with the effective boundaries between the A and B classes. In this space these boundaries are paths along a spherical surface. They are displayed by graphing projections on the $x-y$ plane.

The graphs in Figs. 2-4 illustrate the performance of the model in resolving class boundaries for the two different geometries. In Fig. 2, the class regions were separated by a gap, i.e., an area of pattern space containing no input patterns. When the angular width of this gap is less than $(\lambda_0)^{-1}$, there can develop prototypes for each class which have influence regions extending right up to the boundary with the other class. Consequently the gap is claimed for both pattern categories. Should a pattern from this region be selected as an input, its contested status (response confusion by the model) would cause the influence region of one or the other class to withdraw from a portion of the gap.

Note that in practice, the model need not develop a single decision surface separating pattern classes. In Figs. 3 and 4, there is no gap between the hypothetical

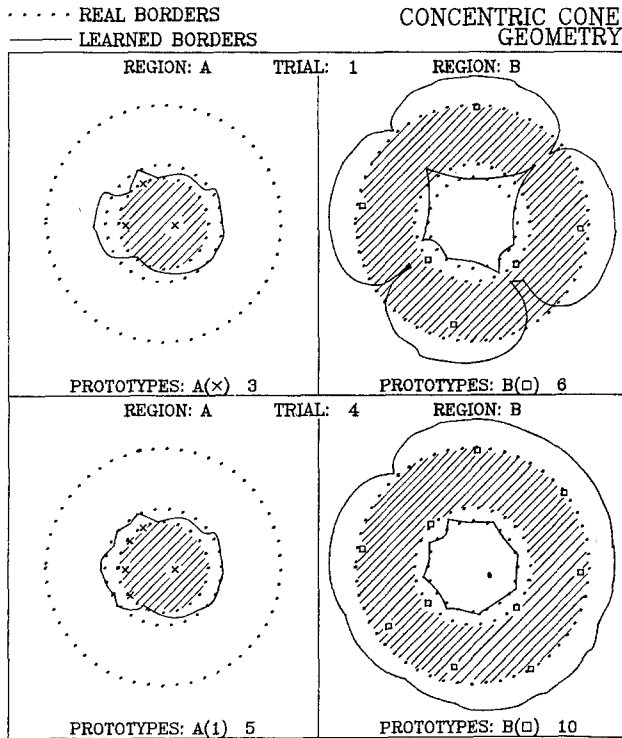


Fig. 2. Prototype regions for the concentric cone geometry with a gap. Region *A*: shaded area within innermost (first) dotted circle. Region *B*: shaded annulus defined by second and third dotted circles. Projections of prototype vectors on sphere's surface are plotted as crosses (*A*) and squares (*B*). Pictured are graphs of prototype boundaries (solid lines) as they appear after the first and fourth trials. Total numbers of prototypes are given below each graph

category regions. A single border separates them, yet in the model, this border is approximated by a double line. If either the prototype or the classification cells were coupled by some mutual interaction (e.g., inhibition), this double border could, in places, be replaced by a single boundary. The nature of such a line would be a function of the specific form of the interaction. Excepting such coupling, it is only in the limit of studying a very large number of design samples that the double line category borders could be expected to merge into a single curve lying along the actual class boundary. The response to any input located in an area where the double lines extend beyond each other will be confusion. Patterns falling in regions from which both prototype generated boundaries have retreated will be identified with neither pattern class.

In the case of prototypes committed to inputs near a class border, the initially large influence regions can result in many incorrect or confused responses until the magnitude of the prototype is appropriately scaled. This creates a somewhat unstable learning process which does not converge smoothly to the final pattern

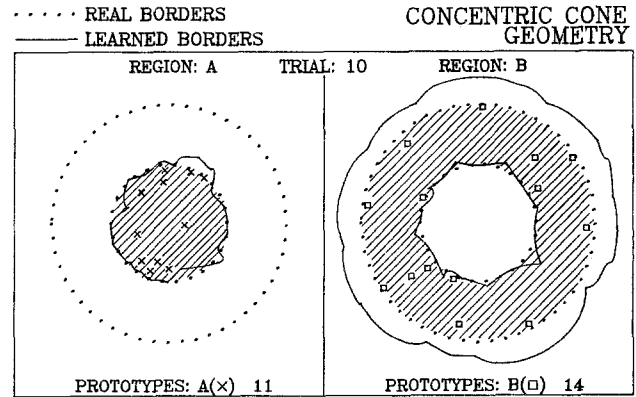


Fig. 3. Prototype regions for concentric cone geometry with no gap. Region *A*: shaded area within innermost (first) dotted circle. Region *B*: shaded annulus bounded by first and second dotted circles. Prototype boundaries (solid lines) pictured after 10 trials




Fig. 4. Prototype regions for sinusoidally separated pattern zones. Region *A*: bottom scalloped semicircle (shaded area within dotted lines). Region *B*: upper scalloped semicircle (shaded area within dotted lines). Prototype boundaries (solid lines) pictured after 50 trials

region mapping¹. Nonetheless, it is clear that this model can resolve pattern classes of arbitrary complexity.

IV. Possible Neural Realization

It is likely that category learning is conducted in different areas of the brain by a variety of cell assemblies. Indeed, one can imagine a number of specific networks of neurons that could implement the important features of our model. We consider a possible

¹ There are a variety of means of improving this. For example, the magnitude of the initial λ_0 may decrease in time so that prototypes committed late in the process leave smaller initial regions of influence. Alternatively, each new prototype may be automatically tested against each existing prototype (treated as an incoming pattern)

neural substrate whose function could relate to one aspect of prototype development in the model.

It has been calculated that under certain conditions, activity in inhibitory fibers whose synapses are located on or near the cell body can have a divisive effect on the somatic membrane potential (Blomfield, 1974). Inhibitory current across these synapses is postulated to increase membrane conductance, thus shunting off a fraction of the summed post-synaptic potential arriving at the cell from its dendrites. The result is to scale the cell output by some multiplicative factor. Inhibitory synapses occurring amidst the excitatory ones further out along the dendritic spines and shafts would have their normal subtractive effect on cell firing rate. Divisive or shunting inhibition has also been considered elsewhere (Poggio, 1981; Kogh et al., 1982).

Cells have been found in different areas of the brain with significant numbers of synapses on or near the perikaryon that are predominantly characterized by flat vesicles and/or symmetric membrane differentiation (e.g., Davis et al., 1979; White et al., 1980). Such morphology is widely considered to be indicative of inhibitory function. By contrast, synapses located on the dendritic shafts and spines of such cells are both excitatory and inhibitory. This anatomy is consistent with that assumed for divisive inhibition. Indeed, other investigators have observed scaling of cell response as a function of inhibitory transmitter released into the soma (Rose, 1977) and under certain conditions of visual stimulus presentation (Dean et al., 1980).

Divisive inhibition is a candidate mechanism for implementing the λ factor scaling of prototype cell response assumed in the prototype learning model. There its principal effect is to provide for a modifiable cell threshold. The distinction which the model makes between prototype commitment and changes in λ is in the same spirit as the functional distinction which Blomfield's model suggests for synapses. The initial commitment of a prototype might involve changes in the spiny synapses and those in general distal to the soma. Such modification could occur according to any of a number of schemes previously suggested (e.g., correlation learning). Cell tuning, on the other hand, would be controlled largely by adjustments to inhibitory synapses proximal to the soma. Long term changes in somatic membrane conductance might even result from very different inhibitory effects (e.g., chemical deposition within the cell body due to active inhibitory afferents).

The processes of modification to sites distal and proximal to the soma might be mutually interactive in a number of ways. For example, one can imagine the somatic membrane conductance of a cell increased to such a point that the cell rarely fires. (In the model,

such was the case for a cell committed to a prototype near a class boundary). Lack of post-synaptic response in conjunction with pre-synaptic activity might cause, as some have suggested (Cooper et al., 1979), the distal synapses of such a cell to lose the information of the stored prototype. This could free the cell to become committed to a new preferred pattern. At the same time, distal modification could be an ongoing process which performs some type of averaging over those inputs able to cause cell firing (the Hebbian requirement). If the environment presented a sequence of smoothly varying events of sufficient duration, the distal modification might cause the cell to "follow" the inputs. In this way, the preferred pattern of the cell could change with only a minimum of change in the degree of cell tuning.

V. Conclusion

Category learning plays an important role in a broad range of mental activity, from learning sequences of task oriented sensori-motor controls to very complex problems in conceptualization. As such it is probably implemented in different ways by different cell assemblies throughout the brain. A successful model for category learning should be consistent with the general features of this host of sub-networks and with their perhaps locally unique architectures. We have presented one such model with properties thought to be characteristic of the neural system as a whole. Among these are: coding of information by neuron firing rates, synaptic transmission of information from cell to cell, excitatory and inhibitory interactions among cells, distributed memory stored over the entire set of synaptic junctions and initially unspecified cell interconnections that are modified by the history of the system's experiences. This model suggests that it is possible to construct plausible neuron networks that incorporate these features and that can display a powerful ability to learn to identify and distinguish categories of events. In a separate publication we will report on the application of this and a related model to a practical problem in categorization (Reilly et al., 1982). The model learning systems were trained to classify examples of unconstrained handwritten numerals. By detecting only very simple information about patterns, the system achieved a high degree of accuracy (approximately 98%) in tests against patterns not viewed during training.

Acknowledgements. We would like to express our appreciation to our colleagues at the Brown University Center for Neural Science for their interest and helpful advice. In particular, we thank Messrs. Paul Munro, Michael Paradiso and Christopher Scofield for several useful discussions.

References

- Amari, S.I.: Neural theory of association and concept-formation. *Biol. Cybern.* **26**, 175–185 (1977)
- Anderson, J.A.: Two models for memory organization using interacting traces. *Math. Biosci.* **8**, 137–160 (1970)
- Anderson, J.A.: A simple neural network generating an interactive memory. *Math. Biosci.* **14**, 197–220 (1972)
- Anderson, J.A., Cooper, L.N.: Les modeles mathematiques de l'organisation biologique de la memoire. *Pluriscience* 168–175, Encyclopaedia Universalis, Paris (1978)
- Barto, A.G., Sutton, R.S., Brouwer, P.S.: Associative search network: a reinforcement learning associative memory. *Biol. Cybern.* **40**, 201–211 (1981)
- Blomfield, S.: Arithmetical operations performed by nerve cells. *Brain Res.* **69**, 115–124 (1974)
- Bobrowski, L.: Rules for forming receptive fields of formal neurons during unsupervised learning processes. *Biol. Cybern.* **43**, 23–28 (1982)
- Brooks, L.: Non-analytical concept formation and memory for instances. In: *Cognition and categorization*, pp. 169–211, Rosch, E., Lloyd, B. (eds.). Hillsdale, N.J.: Lawrence Erlbaum Associates 1978
- Cooper, L.N.: A possible organization of animal memory and learning. In: *Proceedings of the nobel Symposium on collective properties of physical systems*, Vol. 24, pp. 252–264, Lundquist, B., Lundquist, S. (eds.). London, New York: Academic Press 1973
- Cooper, L.N., Liberman, F., Oja, E.: A theory for the acquisition and loss of neuron specificity in visual cortex. *Biol. Cybern.* **33**, 9–28 (1979)
- Cover, T.M., Hart, P.E.: Nearest neighbor pattern classification. *IEEE Trans. Inform. Theor.* **13**, 21–27 (1967)
- Davis, T.L., Sterling, P.: Microcircuitry of cat visual cortex: classification of neurons in layer IV of area 17, and identification of the patterns of lateral geniculate input. *J. Comp. Neur.* **188**, 599–628 (1979)
- Dean, A.F., Hess, R.F., Tolhurst, D.J.: Divisive inhibition involved in directional selectivity. *J. Physiol.* **308**, 84p–85p (1980)
- Duda, R.O., Hart, P.E.: *Pattern classification and scene analysis*. New York: Wiley 1973
- Franks, J.J., Bransford, J.D.: Abstraction of visual patterns. *J. Exp. Psychol.* **90**, 65–74 (1971)
- Grossberg, S.: Adaptive pattern classification and universal recoding. II. Feedback, expectation, olfaction, illusions. *Biol. Cybern.* **23**, 187–202 (1976)
- Kogh, C., Poggio, T., Torre, V.: Retino-ganglion cells: a functional interpretation of dendritic morphology. *Philos. Trans. R. Soc.* (to be published)
- Kohonen, T.: Correlation matrix memories. *IEEE Trans. Comput.* **21**, 353–359 (1972)
- Kohonen, T.: *Associative memory – a system-theoretical approach*. Berlin, Heidelberg, New York: Springer 1977
- Medin, D.L., Schaffer, M.M.: Context theory of classification learning. *Psychol. Rev.* **85**, 207–238 (1978)
- Nass, M.M., Cooper, L.N.: A theory for the development of feature detecting cells in visual cortex. *Biol. Cybern.* **19**, 1–18 (1975)
- Poggio, T.: A theory of synaptic interactions. In: *Theoretical approaches in neurobiology*, pp. 28–38, Reichardt, W., Poggio, T. (eds.). London: MIT Press 1981
- Posner, M.I., Keele, S.W.: On the genesis of abstract ideas. *J. Exp. Psychol.* **77**, 353–363 (1968)
- Posner, M.I., Keele, S.W.: Retention of abstract ideas. *J. Exp. Psychol.* **83**, 304–308 (1970)
- Reilly, D.L., Cooper, L.N., Elbaum, C.: An application of two learning systems to pattern recognition: handwritten characters (to be published)
- Rose, D.: On the arithmetical operation performed by inhibitory synapses onto the neuronal soma. *Exp. Brain Res.* **28**, 221–223 (1977)
- White, E.L., Rock, M.P.: Three-dimensional aspects and synaptic relationships of a Golgi-impregnated spiny stellate cell reconstructed from serial thin sections. *J. Neurocytol.* **9**, 615–636 (1980)

Received: March 12, 1982

Dr. D.L. Reilly
Center for Neural Science and
Department of Physics
Brown University
Providence, RI 02912
USA