

A Hierarchical ART Network Model for General Pattern Recognition

J. Oliver Ross

Department of Computer Science, University of Arkansas
232 Science-Engineering Building
Fayetteville, AR 72701, USA
joross@csci0.uark.edu

Abstract

A new neural network architecture is introduced which may be used for fault-tolerant general pattern recognition. Images are learned by extracting features at each layer. These same images may later be recognized by extracting features which are then used to constrain a search for additional features to validate one of a set of chosen image representation candidates. Unsupervised learning of feature patterns at each layer is accomplished using adaptive resonance theory (ART1) networks.

recognition phase in which any previously seen pattern may be identified. The paper begins with an overview of the ART1 algorithm along with modifications to make it suitable for a hierarchical structure building block. A description of the model follows in addition to discussion of the learning and recognition algorithms.

Introduction

In 1991, translation-invariant recognition of alphanumeric characters was achieved by Fukushima using a hierarchical neural network architecture, the Neocognitron. However, successful recognition by this model was highly dependent upon a chosen set of local features which the network was trained to extract from images making general pattern recognition unfeasible (Fukushima, 1991).

This paper describes a neural pattern recognition architecture which circumvents this limitation using unsupervised learning to create and store representations of features as it is exposed to images. These feature representations are formed dynamically and no training set is needed. Operation of the model consists of a learning phase in which patterns to be recognized are propagated through the network, and a

ART1 Overview

The ART1 neural network model adaptively clusters input vectors according to their similarity. A vigilance parameter ρ is used to control how similar an input vector must be to a given cluster in order to be grouped with that cluster (Carpenter and Grossberg, 1987). This two layer network consists of bottom-up (from input to output layers) and top-down (from output to input layers) weighted connections. The top-down weights are all initialized to one while the bottom-up weights are initialized to $\frac{1}{n+1}$.

During the propagation of a single image vector through the network, the following algorithm is performed until the vector is classified. An active set of output nodes A initially contains each node in the output layer. For each node in A , we calculate a linear combination

$$y_j = \sum_{i=1}^n b_{j,i} x_i$$

where $b_{j,j}$ is a bottom-up weight connecting the output node to each of the n input nodes and let node $j' \in A$ be the output node with the largest y_j value. Next, the value

$$\frac{\sum_{i=1}^m t_{ij'} x_i}{\sum_{i=1}^m x_i}$$

is computed where $t_{ij'}$ is a top-down weight connecting the output node to an input node. Henceforth, this value will be referred to as the similarity parameter. The similarity parameter is now compared to the vigilance parameter ρ . If the similarity parameter is smaller, the output node j' is removed from the active set A and the algorithm is repeated with the remaining nodes in the set. Otherwise, the input vector is associated with j' and the weights connecting that node to the input layer are updated to increase similarity with the input vector. If the algorithm proceeds without clustering the input vector until the set A is empty, a new node is created with which to associate the input vector (Mehrotra, Mohan and Ranka, 1997).

Description

The hierarchical model is composed of interconnected layers of ART1 networks. Each network in layer k has a receptive field or subset of the networks in layer $k-1$ from which it receives its input. As an image vector is propagated through the model, the output of each network in the receptive field is a cluster vector which is transferred as input to the recipient network in the next layer. This network in turn produces as its output a cluster vector which will be transferred to the next layer. At this stage, a new mapping of weighted connections is

formed from each activated node from the networks in the receptive field to the activated node of the recipient network. These new connections are made between every ART1 network and its receptive field in all layers above the first. These extra connections are used to transfer signals in both top-down and bottom-up directions. In the recognition phase, signals flow from output nodes along the corresponding connections from the receptive field to the recipient ART network activating one or more of its cluster nodes. One cluster node from this subset must be chosen to be the feature representation extracted from that network. Every activated node sends signals back down to each of the ART networks in its receptive field activating a unique cluster node in every ART output layer. These activated nodes correspond to the rest of the lower level features which must be present in order for the presence of the higher level feature to be verified. The input vectors to these ART networks are then forced to be clustered as the activated output nodes. The similarity parameter from this forced clustering is used as a confidence value representing the probability of the low-level feature being present.

Learning Dynamics

In the training phase, new connections are formed between output vectors of ART networks in adjacent layers. Let $P = \{\vec{p}_1, \vec{p}_2, \dots, \vec{p}_n\}$ be the set of local representation cluster vectors of ART network layer k in the receptive field of ART network α_j in layer $k+1$. Every representation cluster vectors of ART set P is compressed from a local to a distributed representation P' . Let $\vec{q}_j = \alpha_j(P')$

represent the output of α_j where and $\bar{q}_j = (q_{j1}, q_{j2}, \dots, q_{jn})$ and $q_{j*} \in \bar{q}_j$ is the cluster node into which the ART input vector is classified. For each $\bar{p}_i \in P$, let p_{i*} be the cluster node into which the corresponding ART input vector is classified. Connection mappings $\phi: \{p_{i*}\} \mapsto q_{j*}$ are produced from each p_{i*} to be used as pathways to route signals inbetween adjacent hierarchical layers.

Figure 1 shows a k th layer ART network (right) and its receptive field (left) of ART output layers and their connections to distributed representation neurons. The distributed representation signals are transmitted over direct connections to the input layer of the ART network on the right. Here and in figure 2, we see two sets of dotted lines representing two dynamically formed function mappings of routing connections. These routing connections connect the ART network to its receptive field output vectors.

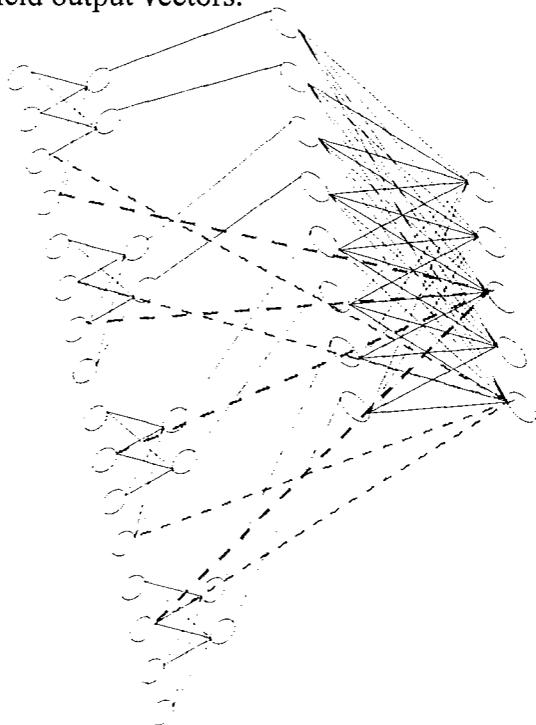


Figure 1: an arbitrary network and connections to its receptive field

Recognition Dynamics

In the recognition phase, once pattern vectors are clustered by ART networks in the first layer, all nodes p_{i*} in the receptive field of α_j in layer two form a subset $A \subset P$. A function $\lambda: A \mapsto B$ is synthesized out of any connections mapping the elements of A onto a elements in \bar{q}_j . Here, B is the subset contained in \bar{q}_j representing the possible clusters into which the receptive field vector may be categorized. In this way subsets B_{kj} are formed for j receptive fields in each layer k .

The recognition objective is to choose an element from each B_{kj} as the proper categorization corresponding to P . Weights are assigned to each connection in each ϕ_{kj} . At any given layer k , only one connection mapping ϕ'_{kj} may be active at a time although signals may still be transferred across inactive connections. This active mapping is represented by assigning connection weight values of one while all other connection weights are zero. Additionally, an n -dimensional activation vector corresponds to each element in all \bar{q}_j where n is the dimension of P in the receptive field of α_j . First layer activation vectors contain similarity values generated by the first layer ART networks propagated across the connection mappings ϕ'_{ij} . No activation vector is entirely filled unless all p_{i*} map to the same q_{j*} . Activation vectors are filled by a recurrent algorithm applicable at any layer. If the activation vector of a preimage element $\phi_{kj}^{-1}(B_{kj})$ is complete, the element to be filled is set equal to the average of these vector elements. Otherwise, the algorithm is used to find the activation vector of $\phi_{kj}^{-1}(B_{kj})$. If this

strategy proceeds without completion to the first layer, the pattern vector is forced to be categorized as cluster node q_{j^*} , and its similarity value is transferred to complete the corresponding activation vector. Competitions for active connections between elements in B_{kj} are resolved by choosing the element with highest activation vector value. When an element B_{kj} wins a competition, the connections formed by the corresponding function ϕ_{kj} are set to one. When this occurs and the element B_{kj} is in a high layer of the model hierarchy, weight values of the tree of nodes which map to B_{kj} are also switched to one symbolizing the certainty of the feature categorization.

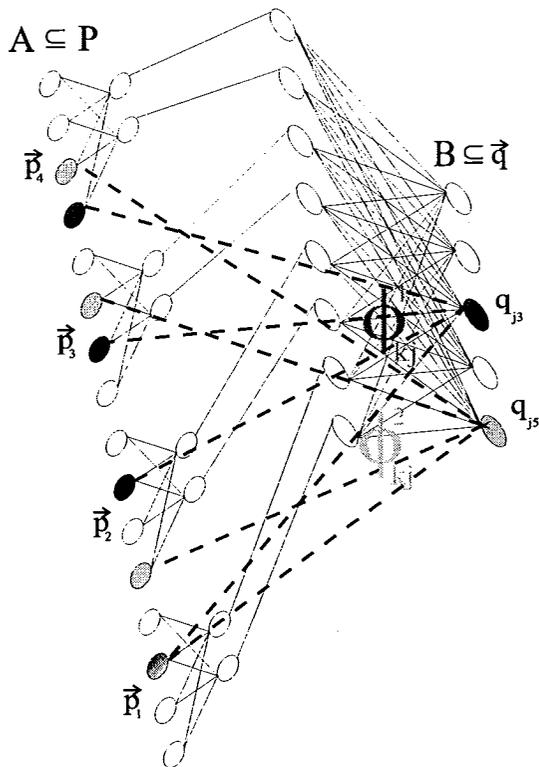


Figure 2: function mappings between a network and its receptive field

Figure 2 shows the two function mappings from the receptive field output

layer nodes to the output nodes of the recipient ART network. From this diagram it is clear that the map formed by the collection of these connections is not necessarily one-to-one. Thus, an activated node in a receptive field may activate not only one, but a subset of the output layer nodes of the recipient ART network.

Experimental Results

Figure 3 shows the results of a simple experiment performed using the network model as proof of concept. The two letter images on the left were learned by the network. The row vector above these two letters shows how they were clustered by the model.

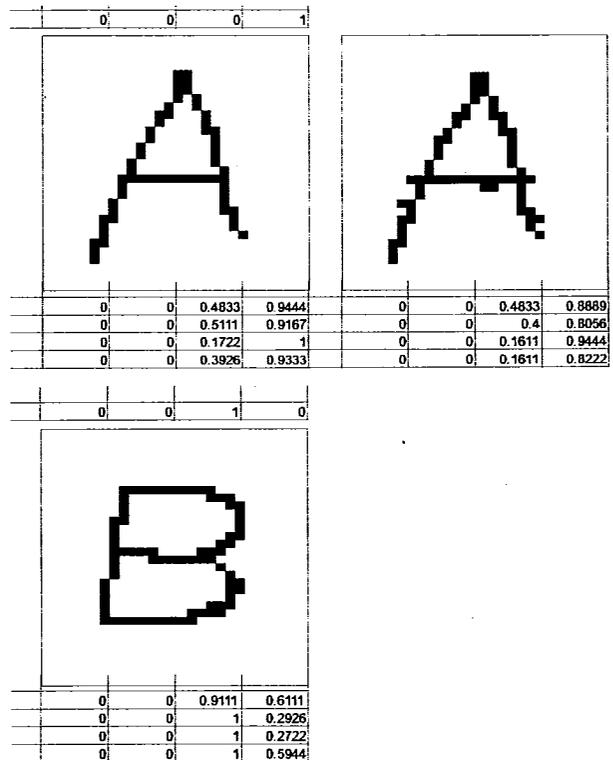


Figure 3: Learning and Recognition of Character Images

These images were correctly recognized by the model in addition to correct recognition of a slightly altered version of the first image (shown at the right). The four column vectors under each letter are activation vectors of the four cluster nodes in the last layer of the model.

Recognition by the Neocognitron. *IEEE Transactions on Neural Networks* 2(3):355-365.

Mehrotra, K., Mohan, C. K. and Ranka, S. (1997). *Elements of Artificial Neural Networks*, Cambridge, MA: MIT Press.

Conclusion

This hierarchical adaptive resonance theory architecture introduces a new level of flexibility in visual pattern recognition through its use of unsupervised learning. When traditional bottom-up feature extraction techniques are used for pattern recognition and parts of the image which deviate from a training image are misclassified early in the feature extraction hierarchy, the classification error is compounded as the erroneous signal is propagated through the additional hierarchical layers. The proposed model is fault-tolerant in that recall of a pattern is the result of many interactions across layers of the model between signals originating at the input layer. However, it is clear that further analysis of the model's performance is in order to be able to more accurately characterize its behavioral properties in certain conditions.

References

- Carpenter, G. A., and Grossberg, S. (1987). A Massively Parallel Architecture for a Self-Organizing Neural Pattern Recognition Machine. *Computer Vision, Graphics, Image Processing* 37:54-115.
- Fukushima, K., and Wake, N. (1991). Handwritten Alphanumeric Character