ART: Self-Organizing Neural Networks for Learning and Memory of Cognitive Recognition Codes

Gail A. Carpenter and Stephen Grossberg
Center for Adaptive Systems, Boston University
111 Cummington Street, Boston, MA 02215

Adaptive resonance (ART) architectures are neural networks that self-organize stable pattern recognition codes in real-time in response to arbitrary sequences of analog or binary input patterns. In ART architectures, top-down learned expectation and matching mechanisms are critical in self-stabilizing the code learning process. A parallel search scheme updates itself adaptively as the learning process unfolds, and realizes a form of real-time hypothesis discovery, testing, learning, and recognition. A parameter called the attentional vigilance parameter determines how fine the categories will be. If vigilance increases (decreases) due to environmental feedback, then the system automatically searches for and learns finer (coarser) recognition categories. Learned representations are encoded in bottom-up and top-down adaptive filters whose long-term memory (LTM) traces vary slowly compared to the rapid short-term memory (STM) information processing.

Adaptive Resonance Theory emerged from an analysis of the instabilities inherent in feedforward adaptive coding structures (Grossberg, 1976a,b). More recent work has led to the development of three classes of ART neural network architectures, specified as systems of differential equations. The first class, ART 1, self-organizes recognition categories for arbitrary sequences of binary input patterns (Carpenter and Grossberg, 1987a). A second class, ART 2, does the same for either binary or analog inputs (Carpenter and Grossberg, 1987b). The talk will include ART 2 simulations that demonstrate how varying vigilance during learning can lead to the stable coexistence of coarse and fine category groupings that depend on the learning history, and show a non-uniform degree of discrimination across the set of inputs.

A third class, ART 3, solves computational problems of ART systems embedded in network hierarchies, where there can, in general, be either fast or slow learning and distributed or compressed code representations (Carpenter and Grossberg, 1990). ART 3 architectures incorporate a third memory, on an intermediate time scale, whose dynamics may be interpreted as chemical transmitter processes. The ART 3 medium-term memory (MTM) equations model the dynamics of production and release of a chemical transmitter substance; the inactivation of transmitter at postsynaptic binding sites; and the modulation of these processes via a nonspecific control signal. The net effect of these transmitter processes is to alter the ionic permeability at the postsynaptic membrane site, thus effecting excitation or inhibition of the postsynaptic cell. Specifically, the presynaptic signal, or action potential, S_i arrives at a synapse whose adaptive weight, or LTM trace, is denoted z_{ij}. The variable z_{ij} is identified with the maximum amount of available transmitter. When the transmitter at this synapse is fully accumulated, the amount of transmitter u_{ij} available for release is equal to z_{ij}. When a signal S_i arrives, transmitter is typically released. The variable v_{ij} denotes the amount of transmitter released into the extracellular space, a fraction of which is assumed to be bound at the postsynaptic cell surface and the remainder rendered ineffective in the extracellular space. Finally, x_j denotes the activity, or membrane potential, of the postsynaptic cell.

Acknowledgements: This research was supported in part by the Air Force Office of Scientific Research (AFOSR F49620-86-C-0037, AFOSR F49620-87-C-0018, and AFOSR 90-0128), the Army Research Office (ARO DAAL-03-88-K-0088), British Petroleum (89-A-1204), DARPA (AFOSR 90-0083), and the National Science Foundation (NSF DMS-90-00530 and IRI-87-16960).

Initially the transmitted signal pattern $\mathbf{S} \cdot \mathbf{u}_j$, as well as the postsynaptic activity x_j, are proportional to the weighted signal pattern $\mathbf{S} \cdot \mathbf{z}_j$ of the linear filter. The activity pattern of the target field is then contrast-enhanced, due to the internal competitive dynamics. The primary ART 3 MTM hypothesis assumes that the transmitter release rate is greatly amplified in proportion to the level of postsynaptic activity. A subsequent reset signal may thus selectively inactivate those pathways that cause an error. Following such a reset wave, the new signal $\mathbf{S} \cdot \mathbf{u}_j$ is no longer proportional to $\mathbf{S} \cdot \mathbf{z}_j$ but is, rather, biased against the previously active representation due to transmitter depletion at those sites. A series of reset events ensue, until an adequate match or a new category is found. Learning occurs on a time scale that is long relative to that of the search process.

The ART 3 MTM serves other functions as well as implementing the ART mismatch-reset-search cycle. In particular it allows the neural network to dispense with special processes to reset STM at onset or offset of an input pattern. The representation of input patterns as a sequence, $\mathbf{I}_1, \mathbf{I}_2, \mathbf{I}_3, \ldots$, corresponds to the assumption that each input is constant for a fixed time interval. In practice, an input vector $\mathbf{I}(t)$ may vary continuously through time. The input need never be constant over an interval, and there may be no temporal marker to signal offset or onset of "an input pattern" per se.

The ART 3 MTM transmitter depletion process can also serve to enhance features that were previously ignored. For example, suppose that the input ($\mathbf{I}$) signal pathways contain an ART 3 MTM process, and that a reset signal is generated, say, by an internal system error or by an external teaching input. Features represented in $\mathbf{I}$ that were not salient in the matched STM pattern $\mathbf{X}$ are enhanced, due to depletion in pathways leading to those features that, for whatever reason, generated an error signal. For instance, the previously ignored color of an object may be brought forth to enhance discrimination between category exemplars.

The mechanisms described thusfar are part of the recognition learning circuit of ART 3. Recognition learning is, however, only one of several processes whereby an intelligent system can learn a correct solution to a problem. We have called Recognition, Reinforcement, and Recall the "3 R's" of neural network learning (Carpenter and Grossberg, 1988). Various types of reaction to reinforcement feedback may be useful in applications. For example, a change in vigilance alters the overall sensitivity of the system to pattern differences; a shift in attention and the reset of active features can help to overcome prior coding biases that may be maladaptive in novel contexts.

In summary, the MTM provides the extra degree of freedom needed to embed ART systems in neural network hierarchies with fast or slow learning and compressed or distributed codes. The ART 3 MTM transmitter processes can also be used in a wide variety of fully or partially connected and adaptive or non-adaptive neural networks. The ART 3 search mechanism serves at least four distinct functions: to correct erroneous category choices; to learn from reinforcement feedback or disconfirmed expectations; to respond to changing input patterns; and, when as error signal occurs, to amplify features that were previously ignored. The talk will illustrate how ART modules embedded in neural network hierarchies carry out various target recognition functions.

References

Carpenter, G.A. and Grossberg, S. (1987a). A massively parallel architecture for a self-organizing neural pattern recognition machine. *Computer Vision, Graphics, and Image Processing*, **37**, 54–115.

Carpenter, G.A. and Grossberg, S. (1987b). ART 2: Self-organization of stable category recognition codes for analog input patterns. *Applied Optics*, **26**, 4919–4930.

Carpenter, G.A. and Grossberg, S. (1988). The ART of adaptive pattern recognition by a self-organizing neural network. *Computer*, **21**, 77–88.

Carpenter, G.A. and Grossberg, S. (1990). ART 3: Hierarchical search using chemical transmitters in self-organizing pattern recognition architectures. *Neural Networks*, in press.

Grossberg, S. (1976a). Adaptive pattern classification and universal recoding, I: Parallel development and coding of neural feature detectors. *Biological Cybernetics*, **23**, 121–134.

Grossberg, S. (1976b). Adaptive pattern classification and universal recoding, II: Feedback, expectation, olfaction, and illusions. *Biological Cybernetics*, **23**, 187–202.