Adaptive Resonance Theory

Definition

Adaptive Resonance Theory, or ART, is both a cognitive and neural theory of how the brain quickly learns to categorize, recognize, and predict objects and events in a changing world, and a set of algorithms which computationally embody ART principles and are used in large-scale engineering and technological applications where fast, stable, incremental, learning about complex changing environments is needed. ART clarifies the brain processes from which conscious experiences emerge. It predicts a functional link between processes of Consciousness, Learning, Expectation, Attention, Resonance, and Synchrony (CLEARS), including the prediction that "all conscious states are resonant states." This connection clarifies how brain dynamics enable a behaving individual to autonomously adapt in real time to a rapidly changing world. ART predicts how top-down attention works and regulates fast stable learning of recognition categories. In particular, ART articulates a critical role for "resonant" states in driving fast stable learning; hence the name *adaptive resonance*. These resonant states are bound together, using top-down attentive feedback in the form of learned expectations, into coherent representations of the world. ART hereby clarifies one important sense in which the brain carries out predictive computation. ART has explained and successfully predicted a wide range of behavioral and neurobiological data, including data about human cognition and the dynamics of spiking laminar cortical networks. ART algorithms have been used in large-scale applications such as medical data base prediction, remote sensing, airplane design, and the control of autonomous adaptive robots.

Motivation and Background

Many current learning algorithms do not emulate the way in which humans and other animals learn. The power of human and animal learning provides high motivation to discover computational principles whereby machines can learn with similar capabilities. Humans and animals experience the world on the fly, and carry out incremental learning of sequences of episodes in real time. Often such learning is unsupervised, with the world itself as the teacher. Learning can also proceed with an unpredictable mixture of unsupervised and supervised learning trials. Such learning goes on successfully in a world that is non-stationary; that is, whose rules can change unpredictably through time. Moreover, humans and animals can learn quickly and stably through time. A single important experience can be remembered for a long time. ART proposes a solution of this *stability-plasticity dilemma* by showing how brains learn quickly without forcing catastrophic forgetting of already learned, and still successful, memories.

Thus, ART autonomously carries out fast, yet stable, incremental learning under both unsupervised and supervised learning conditions in response to a complex non-stationary world. In contrast, many current learning algorithms are use batch learning in which all the information about the world to be learned is available at a single time. Other algorithms are not defined unless all learning trials are unsupervised. Yet other algorithms become unstable in a non-stationary world, or become unstable if learning is fast; that is, if an event can be fully learned on a single learning trial. ART overcomes these problems.

Some machine learning algorithms are feed-forward clustering algorithms that undergo catastrophic forgetting in a non-stationary world. The ART solution of the stability-plasticity dilemma depends upon feedback, or top-down, expectations that are matched against bottom-up data and thereby focus attention upon critical feature patterns. A good enough match leads to resonance and fast learning. A big enough mismatch leads to hypothesis testing or memory search that discovers and learns a more predictive category. Thus, ART is a self-organizing expert system that avoids the brittleness of traditional expert systems.

The world is filled with uncertainty, so probability concepts seem relevant to understanding how brains learn about uncertain data. This fact has led some machine learning practitioners to assume that brains obeys Bayesian laws. However, the Bayes rule is so general that it can accommodate any system in Nature. Additional computational principles and mechanisms must augment Bayes to distinguish a brain from, say, a hydrogen atom or storm. Moreover, probabilistic models often use non-local computations. ART shows how the brain embodies a novel kind of real-time probability theory, hypothesis testing, prediction, and decision-making whose local computations adapt to a non-stationary world. These ART principles and mechanisms go beyond Bayesian analysis, and are embodied parsimoniously in the laminar circuits of cerebral cortex. Indeed, the cortex embodies a new kind of Laminar Computing that reconciles the best properties of feedforward and feedback processing, digital and analog processing, and data-driven bottom-up processing and hypothesis-driven top-down processing

Structure of Learning System

How CLEARS Mechanisms Interact

Humans are *intentional* beings who learn expectations about the world and make predictions about what is about to happen. Humans are also *attentional* beings who focus processing resources upon a restricted amount of incoming information at any time. Why are we both intentional and attentional beings, and are these two types of processes related? The stability-plasticity dilemma and its solution using resonant states provides a unifying framework for understanding these issues.

To clarify the role of sensory or cognitive expectations, and of how a resonant state is activated, suppose you were asked to "find the yellow ball as quickly as possible, and you will win a $10,000 prize". Activating an expectation of a "yellow ball" enables its more rapid detection, and with a more energetic neural response. Sensory and cognitive top-down expectations hereby lead to *excitatory matching* with consistent bottom-up data. Mismatch between top-down expectations and bottom-up data can suppress the mismatched part of the bottom-up data, to focus attention upon the matched, or expected, part of the bottom-up data.

Excitatory matching and attentional focusing on bottom-up data using top-down expectations generates resonant brain states: When there is a good enough match between bottom-up and top-down signal patterns between two or more levels of processing, their positive feedback signals amplify and prolong their mutual activation, leading to a resonant state. Amplification and prolongation of activity triggers learning in the more slowly varying adaptive weights that control the signal flow along pathways from cell to cell. Resonance hereby provides a global context-sensitive indicator that the system is processing data worthy of learning, hence the name *Adaptive* Resonance Theory.

In summary, ART predicts a link between the mechanisms which enable us to learn quickly and stably about a changing world, and the mechanisms that enable us to learn expectations about such a world, test hypotheses about it, and focus attention upon information that we find interesting. ART clarifies this link by asserting that, in order to solve the stability-plasticity dilemma, only resonant states can drive rapid new learning.

It is just a step from here to propose that those experiences which can attract our attention and guide our future lives by being learned are also among the ones that are conscious. Support for this additional assertion derives from the many modeling studies whose simulations of behavioral and brain data using resonant states map onto properties of conscious experiences in those experiments.

The type of learning within the sensory and cognitive domain that ART mechanizes is *match learning*: Match learning occurs only if a good enough match occurs between bottom-up information and a learned top-down expectation that is read out by an active recognition category, or code. When such an approximate match occurs, previously learned knowledge can be refined. Match learning raises the concern about what happens if a match is not good enough? How does such a model escape perseveration on already learned representations?

If novel information cannot form a good enough match with the expectations that are read-out by previously learned recognition categories, then a memory search, or hypothesis testing, is triggered that leads to selection and learning of a new recognition category, rather than catastrophic forgetting of an old one. Figure 1 illustrates how this happens in an ART model; it will be discussed in greater detail below. In contrast, learning within spatial and motor processes is proposed to be *mismatch learning* that continuously updates sensory-motor maps or the gains of sensory-motor commands. As a result, we can stably learn what is happening in a changing world, thereby solving the stability-plasticity dilemma, while adaptively updating our representations of where objects are and how to act upon them using bodies whose parameters change continuously through time. Brain systems that use inhibitory matching and mismatch learning cannot generate resonances; hence, their representations are not conscious.

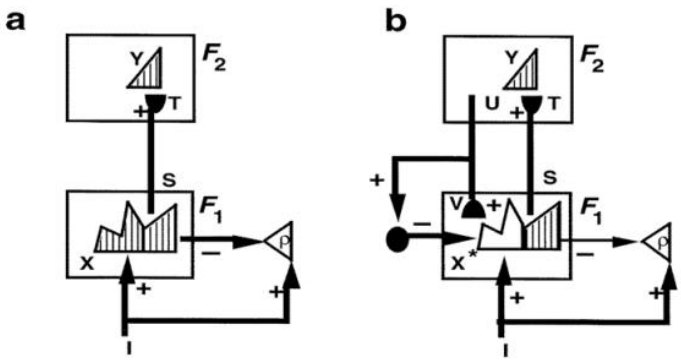

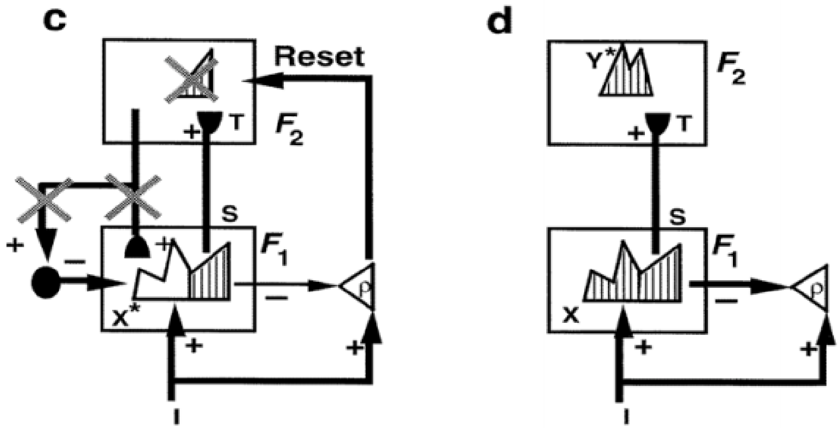

Figure 1. Search for a recognition code within an ART learning circuit: (a) The input pattern I is instated across the feature detectors at level F_1 as a short term memory (STM) activity pattern X. Input I also nonspecifically activates the orienting system with a gain that is called vigilance (ρ); that is, all the input pathways converge with gain ρ onto the orienting system and try to activate it. STM pattern X is represented by the hatched pattern across F_1. Pattern X both inhibits the orienting system and generates the output pattern S. Pattern S is multiplied by learned adaptive weights, also called long term memory (LTM) traces. These LTM-gated signals are added at F_2 cells, or nodes, to form the input pattern T, which activates the STM pattern Y across the recognition categories coded at level F_2. (b) Pattern Y generates the top-down output pattern U which is multiplied by top-down LTM traces and added at F_1 nodes to form a *prototype* pattern V that encodes the learned expectation of the active F_2 nodes. Such a prototype represents the set of commonly shared features in all the input patterns capable of activating Y. If V mismatches I at F_1, then a new STM activity pattern X^* is selected *at* F_1. X^* is represented by the hatched pattern. It consists of the features of I that are confirmed by V. Mismatched features are inhibited. The inactivated nodes corresponding to unconfirmed features of X are unmatched. The reduction in total STM activity which occurs when X is transformed into X^* causes a decrease in the total inhibition from F_1 to the orienting system. (c) If inhibition decreases sufficiently, the orienting system releases a nonspecific arousal wave to F_2; that is, a wave of activation that equally activates all F_2 nodes. This wave instantiates the intuition that "novel events are arousing". This arousal wave resets the STM pattern Y at F_2 by inhibiting Y. (d) After Y is inhibited, its top-down prototype signal is eliminated, and X can be reinstated at F_1. The prior reset event maintains inhibition of Y during the search cycle. As a result, X can activate a different STM pattern Y at F_2. If the top-down prototype due to this new Y pattern also mismatches I at F_1, then the search for an appropriate F_2 code continues until a more appropriate F_2 representation is selected. Such a search cycle represents a type of nonstationary hypothesis testing. When search ends, an attentive resonance develops and learning of the attended data is initiated.

Complementary Computing in the Brain: Resonance and Reset

It has been mathematically proved that match learning within an ART model leads to stable memories in response to arbitrary list of events to be learned . However, match learning also has a serious potential weakness: If you can only learn when there is a good enough match between bottom-up data and learned top-down expectations, then how do you ever learn anything that you do not already know? ART proposes that this problem is solved by the brain by using an interaction between complementary processes of *resonance* and *reset,* that are predicted to control properties of attention and memory search, respectively. These complementary processes help our brains to balance between the complementary demands of processing the familiar and the unfamiliar, the expected and the unexpected.

Organization of the brain into complementary processes is predicted to be a general principle of brain design that is not just found in ART . A complementary process can individually compute some properties well, but cannot, by itself, process other complementary properties. In thinking intuitively about complementary properties, one can imagine puzzle pieces fitting together. Both pieces are needed to finish the puzzle. Complementary brain processes are more dynamic than any such analogy: Pairs of complementary processes interact to form emergent properties which overcome their complementary deficiencies to compute complete information with which to represent or control some aspect of intelligent behavior.

The resonance process in the complementary pair of resonance and reset is predicted to take place in the What cortical stream, notably in the inferotemporal and prefrontal cortex. Here top-down expectations are matched against bottom-up inputs. When a top-down expectation achieves a good enough match with bottom-up data, this match process focuses attention upon those feature clusters in the bottom-up input that are expected. If the expectation is close enough to the input pattern, then a state of resonance develops as the attentional focus takes hold.

Figure 1 illustrates these ART ideas in a simple two-level example. Here, a bottom-up input pattern, or vector, I activates a pattern X of activity across the feature detectors of the first level F_1. For example, a visual scene may be represented by the features comprising its boundary and surface representations. This feature pattern represents the relative importance of different features in the inputs pattern I. In Figure 1a, the pattern peaks represent more activated feature detector cells, the troughs less activated feature detectors. This feature pattern sends signals S through an adaptive filter to the second level F_2 at which a compressed representation Y (also called a recognition category, or a symbol) is activated in response to the distributed input T. Input T is computed by multiplying the signal vector S by a matrix of adaptive weights that can be altered through learning. The representation Y is compressed by competitive interactions across F_2 that allow only a small subset of its most strongly activated cells to remain active in response to T. The pattern Y in the figure indicates that a small number of category cells may be activated to different degrees. These category cells, in turn, send top-down signals U to F_1. The vector U is converted into the top-down expectation V by being multiplied by another matrix of adaptive weights. When V is received by F_1, a matching process takes place between the input vector I and V which selects that subset X* of F_1 features that were "expected" by the active F_2 category Y. The set of these selected features is the emerging "attentional focus".

Binding Distributed Feature Patterns and Symbols during Conscious Resonances

If the top-down expectation is close enough to the bottom-up input pattern, then the pattern X* of attended features reactivates the category Y which, in turn, reactivates X*. The network hereby locks into a resonant state through a positive feedback loop that dynamically links, or binds, the attended features across X* with their category, or symbol, Y.

Resonance itself embodies another type of complementary processing. Indeed, there seem to be complementary processes both within and between cortical processing streams This particular complementary relation occurs between distributed feature patterns and the compressed categories, or symbols, that selectively code them:

Individual features at F_1 have no meaning on their own, just like the pixels in a picture are meaningless one-by-one. The category, or symbol, in F_2 is sensitive to the global patterning of these features, and can selectively fire in response to this pattern. But it cannot represent the "contents" of the experience, including their conscious qualia, due to the very fact that a category is a compressed, or "symbolic" representation. Practitioners of Artificial Intelligence have claimed that neural models can process distributed features, but not symbolic representations. This is not, of course, true in the brain. Nor is it true in ART.

Resonance between these two types of information converts the *pattern* of attended features into a coherent context-sensitive state that is linked to its category through feedback. This coherent state, which binds together distributed features and symbolic categories, can enter consciousness while it binds together spatially distributed features into either a stable equilibrium or a synchronous oscillation. The original ART article predicted the existence of such synchronous oscillations, which were there described in terms of their mathematical properties as "order-preserving limit cycles".

Adaptive Resonance Theory (ART)

- Invented by Grossberg in 1976 and based on unsupervised learning model.
- Resonance means a target vector matches close enough the input vector.
- ART matching leads to resonance and only in resonance state the ART network learns.
- Suitable for problems that uses online dynamic large databases.
- ART 1- classifies binary input vectors, ART 2 – clusters real valued input(continuous valued) vectors.
- Used to solve Plasticity – stability dilemma.

Plasticity –stability dilemma:

How to learn a new pattern without forgetting the old traces (patterns) and how to adapt to the changing environment (i/p). When there is change inthe patterns (plasticity) how to remember previously learned vectors (stability problem) is a problem. ART uses competitive law (self-regulating control) to solve this PLACITICITY – STABILITY Dilemma.

ARCHITCTURE:simplified ART diagram is given below.

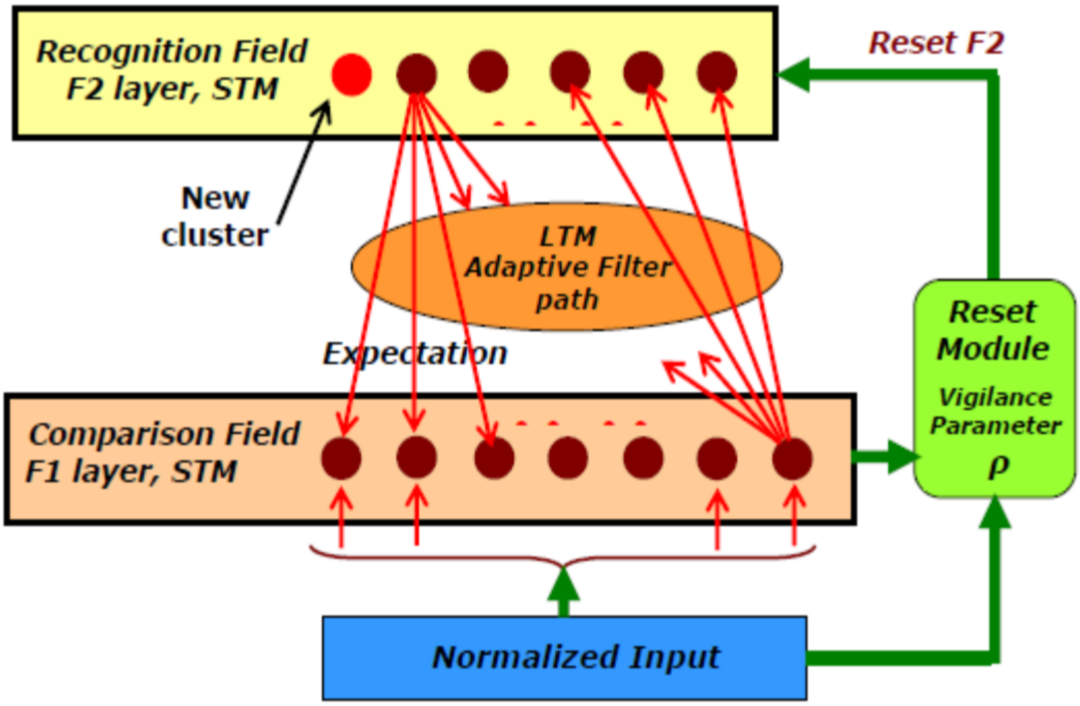

Fig Simplified ART Architecture

ART consists of

(1) F1 Layer: I/P processing unit also called comparison layer.

(2) F2 Layer: clustering or competitive layer.

(3) Reset mechanism.

Comparison Layer: Take 1D i/p vector and transfers it to the best match in recognition field(best match - neuron in recognition unit whose weight closely matches with i/p vector).

Recognition Unit: produces an output proportional to the quality of match. In this way recognition field allows a neuron to represent a category to which the input vectors are classified.

Vigilance parameter: After the i/p vectors are classified the a reset module compares the strength of match to vigilance parameter (defined by the user). Higher vigilance produces fine detailed memories and lower vigilance value gives more general memory.

Reset module: compares the strength of recognition phase. When vigilance threshold is met then training starts otherwise neurons are inhibited until a new i/p is provided.

There are two set of weights (i) bottom up weight - from F1 layer to F2 Layer

(2) Top –Down weight – F2 to F1 Layer

Fast learning: Happens in ART 1 – Weight changes are rapid and takes place during resonance. The network is stabilized when correct match at cluster unit is reached.

Slow Learning: Used in ART 2 .weight change is slow and does not reach equilibrium in each learning iteration.so more memory to store more i/p patterns (to reach stability) is required.

BASIC ART Training steps:

1. Initialize the parameters.
2. If no stop condition do step 3 to 10.
3. For each i/p vector do steps 4 to 9.
4. F1 Layer process starts.
5. If reset condition = true do step 6 to 8.
6. Find f2 unit with largest i/p (to learn current pattern).
7. F1(b) units combine their /p's F1(a) & F2.
8. Test for reset condition. (differs for ART1 & ART2)
9. If reset true candidate is rejected (i.e., neuron is inhibited). Return to step 5.

If reset is false candidate unit is accepted for learning.

10. Learning starts. Weight updation starts as per diff equation.
11. Test for stop condition.

ART1 for binary 0/1 inputs.

ART2 for continuous valued inputs.

1. Stability – plasticity dilemma:

   There exists guarantee of stability of categories formed by competitive learning. If continue presenting same patterns category of a datum may continue changing endlessly.

   To prevent could reduce learning to 0, freezing learned categories, but then net looses plasticity (ability to react to new data).

2. # of output unit categories to be used? (avoid dead units? Or loose plasticity?)
   In ART1 an input pattern when not sufficiently similar (there exists metric for this) to any existing prototype vector causes new category (output unit) to be activated.

   There exists vigilance parameter ρ , $0 < \rho <= 1$ to control similarity.

   If ρ large --> similarity condition very stringent N -> many categories.

   If ρ small -> few categories.

ART1 algorithm:

1. Start with $W_i = \ = (1,1,\ldots,1)$, $\forall$ I (uncommitted states)
2. Present a pattern z.
3. Enable all output units.
4. Find winner i* (exit if there does not exist enabled output unit).
   Among enabled units.

   Winner is one with $\vec{W}_i \cdot z$ largest, where $\vec{W}_i = \overline{W}_i / (\ + \Sigma_j W_{ji})$

5. Test if match between winner $\vec{W_i}$. z is good enough:
   Compute $\gamma = (\vec{W_i} \cdot z) / \Sigma_j \xi_j)$ → fraction of bits in z that are also in W_i^*

   If $\gamma >= \rho$ -> ∃ resonance go to step 6.

   If $\gamma < \rho$ -> reject Wi* ; disable unit i*, go to step 4.

   NOTE: After a finite # of presentation of data patterns the algorithm will stop.

The architecture of an ART1 net consists of two fields of units, the F_1 units and the F_2 (cluster) units, together with a reset unit to control the degree of similarity of patterns placed on the same cluster unit. This main portion of the ART1 architecture is illustrated in Figure 1. The F_1 interface unit X_i is connected to the F_2 cluster unit Y_j by bottom-up weight b_{ij}. Similarly, unit Y_j is connected to unit X_i by top-down weights t_{ji}. The F_1 and F_2 layers are connected by two sets of weighted pathways. In addition, several supplemental units are included in the net to provide for neural control of the learning process. It is designed so that it is not required either that patterns be presented in a fixed order or that the number of patterns to be clustered be known in advance. Updates for both the bottom-up and top-down weights are controlled by differential equations.

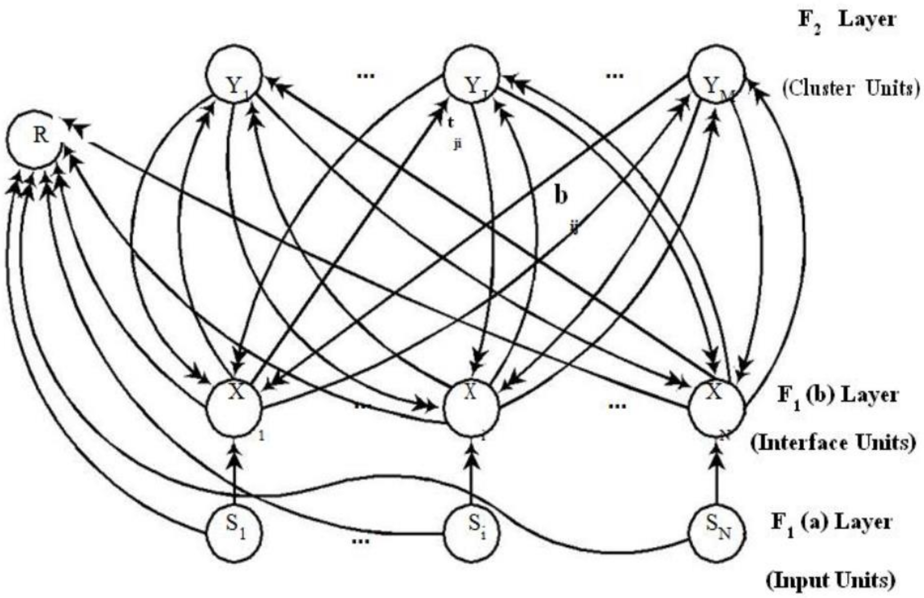

ART 1 Architecture

ART 2 Architecture

The F_1 layer consists of six types of units (W, X, U, V, P, and Q units). There are *n* units of each of these types, where *n* is the dimension of an input vector. The F_2 field contains only one layer, which is denoted by Y and serves as a competitive layer. There are top-down and bottom-up full connections between F_1 and F_2 pattern prototypes are to be preserved on these connections. The input signal is $S = (S_1, ..., S_i, ..., S_n)$ continues to be sent while all of the sections to be described are performed. At the beginning of a learning trail, all activation is set to zero. The computation within the F_1 layer can be thought of as originating with the computation of the activation of unit U_i (the activation of unit V_i normalized to approximately unit length). Next, a signal is sent from each unit U_i to its associated units W_i and P_i. The activation of units Wi and Pi are then computed. Unit W_i sums the signal it received from U_i and the input signal S_i. P_i sum the signal it receives from U_i and the top-down signal it receives if there is an active F_2 unit. The activation of X_i and Q_i are normalized version of the signal at W_i and P_i. An activation function is applied at each of units before the signal is sent to V_i. V_i then sums the signals if receives concurrently from X_i and Q_i. This completes one cycle of updating the F_1 layer.

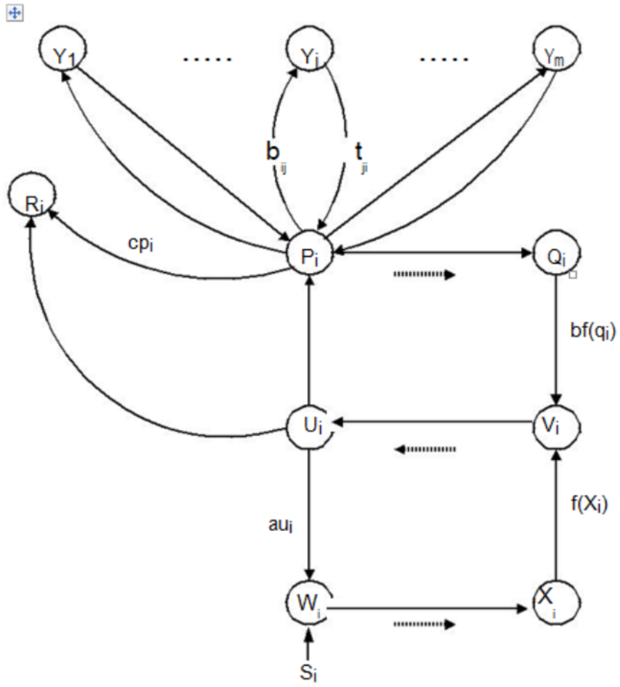

ART 2 Architecture

REFERENCES

• Carpenter, G.A. & Grossberg, S. (2003), Adaptive Resonance Theory Archived 2006-05-19 at the Wayback Machine, In Michael A. Arbib (Ed.), The Handbook of Brain Theory and Neural Networks, Second Edition (pp. 87-90). Cambridge, MA: MIT Press

• Grossberg, S. (1987), Competitive learning: From interactive activation to adaptive resonance Archived 2006-09-07 at the Wayback Machine, Cognitive Science (journal), 11, 23-63

• Carpenter, G.A. & Grossberg, S. (1987), ART 2: Self-organization of stable category recognition codes for analog input patterns Archived 2006-09-04 at the Wayback Machine, Applied Optics, 26(23), 4919-4930

• Carpenter, G.A., Grossberg, S., & Rosen, D.B. (1991a), ART 2-A: An adaptive resonance algorithm for rapid category learning and recognition Archived 2006-05-19 at the Wayback Machine, Neural Networks (Publication), 4, 493-504

• Carpenter, G.A. & Grossberg, S. (1990), ART 3: Hierarchical search using chemical transmitters in self-organizing pattern recognition architectures Archived 2006-09-06 at the Wayback Machine, Neural Networks (Publication), 3, 129-152

• Carpenter, G.A., Grossberg, S., & Reynolds, J.H. (1991), ARTMAP: Supervised real-time learning and classification of nonstationary data by a self-organizing neural network Archived 2006-05-19 at the Wayback Machine, Neural Networks (Publication), 4, 565-588

• Carpenter, G.A., Grossberg, S., & Rosen, D.B. (1991b), Fuzzy ART: Fast stable learning and categorization of analog patterns by an adaptive resonance system Archived 2006-05-19 at the Wayback Machine, Neural Networks (Publication), 4, 759-771

• Carpenter, G.A., Grossberg, S., Markuzon, N., Reynolds, J.H., & Rosen, D.B. (1992), Fuzzy ARTMAP: A neural network architecture for incremental supervised learning of analog multidimensional maps Archived 2006-05-19 at the Wayback Machine, IEEE Transactions on Neural Networks, 3, 698-713

• Mohammad-Taghi Vakil-Baghmisheh and Nikola Pavešić. (2003) A Fast Simplified Fuzzy ARTMAP Network, Neural Processing Letters, 17(3):273–316

• James R. Williamson. (1996), Gaussian ARTMAP: A Neural Network for Fast Incremental Learning of Noisy Multidimensional Maps, Neural Networks, 9(5):881-897

• Y.R. Asfour, G.A. Carpenter, S. Grossberg, and G.W. Lesher. (1993) Fusion ARTMAP: an adaptive fuzzy network for multi-channel classification. In: Proceedings of the Third International Conference on Industrial Fuzzy Control and Intelligent Systems (IFIS).

• Marko Tscherepanow. (2010) TopoART: A Topology Learning Hierarchical ART Network, In: Proceedings of the International Conference on Artificial Neural Networks (ICANN), Part III, LNCS 6354, 157-167

• Georgios C. Anagnostopoulos and Michael Georgiopoulos. (2000), Hypersphere ART and ARTMAP for Unsupervised and Supervised Incremental Learning, In: Proceedings of the International Joint Conference on Neural Networks (IJCNN), vol. 6, 59-64

• Sandia National Laboratories (2017) Lapart-python documentation

• Marko Tscherepanow. (2012) Incremental On-line Clustering with a Topology-Learning Hierarchical ART Neural Network Using Hyperspherical Categories, In: Poster and Industry Proceedings of the Industrial Conference on Data Mining (ICDM), 22–34