Adaptive Resonance Theory

ART Networks

Stephen Grossberg &
Gail Carpenter

ART Network Varieties

- Several varieties:
  - ART1: Discrete (e.g. binary) patterns
  - ART2: Continuous patterns …
  - ARTMAP: supervised learning
  - Fuzzy Art: Fuzzy version of ART1

- All work by competitive learning and a sort of **clustering** (even though not identified as such) + nuances

Stability-Plasticity Dilemma

- **stability**: Recognized patterns should be insensitive to noise.
- **plasticity**: System should be capable of learning new patterns.
- The conflict between these is one of the things that ART hopes to resolve.

ART Networks

- Combine supervised and unsupervised (competitive, clustering)
- Dynamically create new categories
- Biologically motivated by an ODE model
- Models short- and long-term memory

ART1 Model

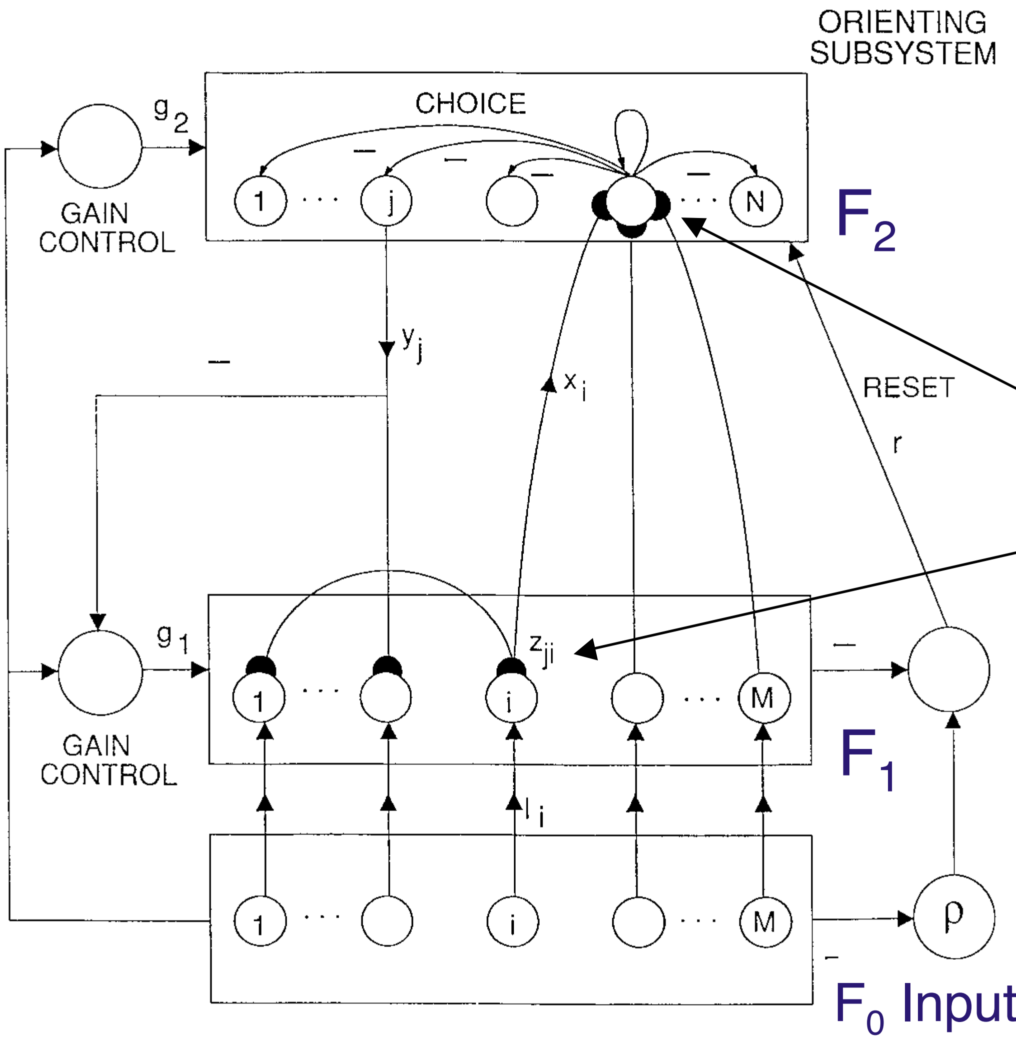

Learning takes place at the synapses denoted by semi-circular endings in the F1 to F2 pathways.

ART General Workings

- F_2 activities can be interpreted as “making a hypothesis” about input at F_1.
- The LTM-gated signals from all the active F_2 nodes are added to generate the total top-down feedback pattern from F_2 to F_1. This pattern plays the role of a **learned expectation**. Activation of this expectation “tests the hypothesis” of the active F_2 category.
- The expected prototype of the category are compared with the bottom-up input pattern (exemplar) to F_1.

Diagram from Carpenter and Grossberg, 1988

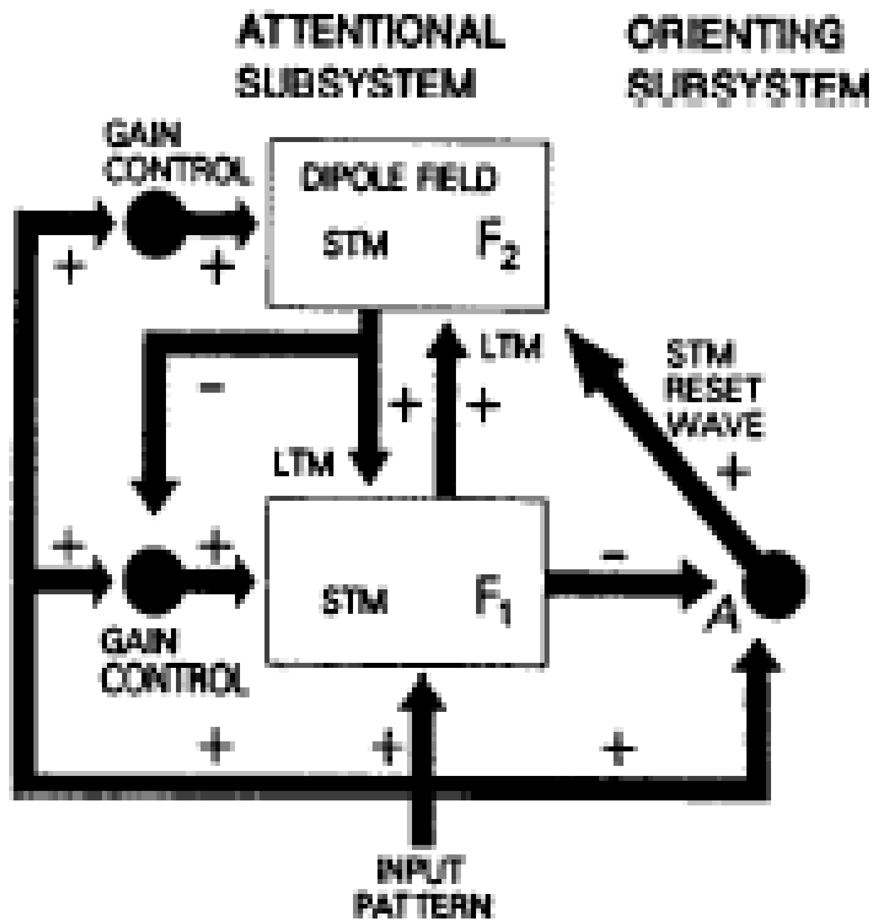

ART search for a recognition code (A of 4)
[Grossberg and Merrill (1996)]

- 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 subsystem A. STM pattern X is represented by the hatched pattern across F_1.
- Pattern X both inhibits A and generates the output pattern S.
- Pattern S is multiplied by long term memory (LTM) traces and added at F_2 nodes to form the input pattern T, which activates the STM pattern Y across the recognition categories coded at level F_2.

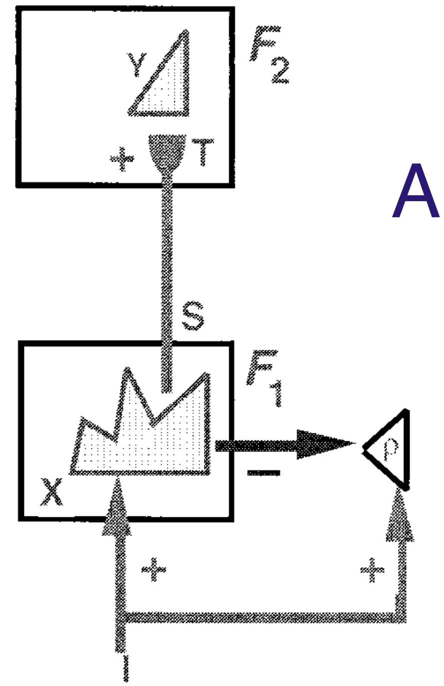

ART search for a recognition code (B of 4)

- 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 the prototype pattern V that encodes the learned expectation of the active F_2 nodes.

- If V mismatches I at F_1, then a new STM activity pattern X* is generated at F_1. X* is represented by the hatched pattern. It includes the features of I that are confirmed by V. Inactivated nodes corresponding to unconfirmed features of X are unhatched.

- 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 A.

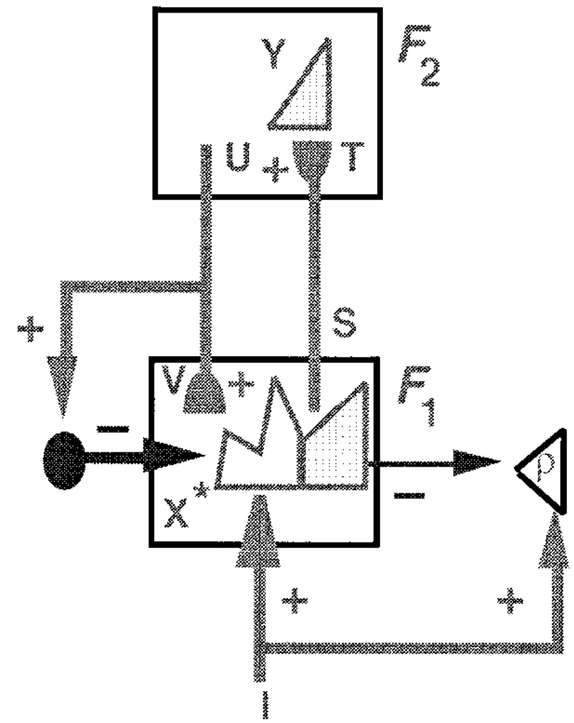

ART search for a recognition code (C of 4)

- If inhibition decreases sufficientl [vigilance condition not met], A releases at a nonspecific arousal wave to F_2, which resets the STM pattern Y at F_2.

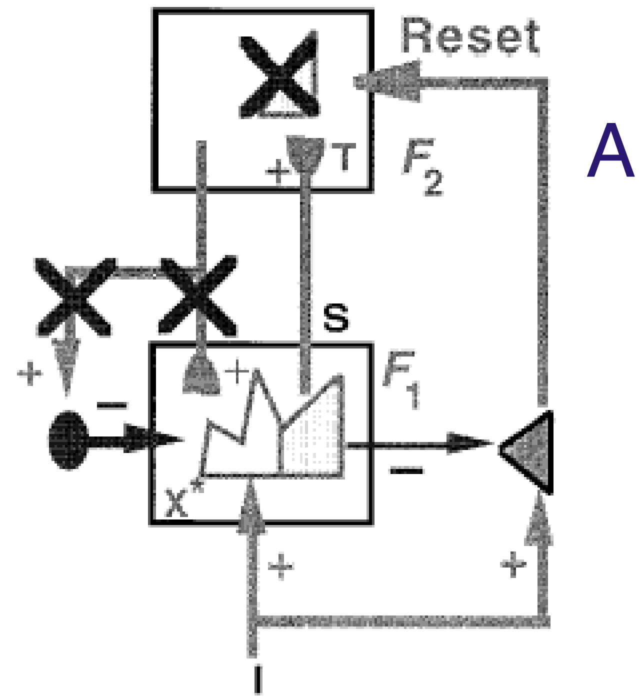

ART search for a recognition code (D of 4)

- After Y is inhibited, its top-down prototype signal is eliminated, and X can be reinstated at F_1.
- Enduring traces of the prior reset lead X to activate a different STM pattern Y* at F_2.
- If the top-down prototype due to Y* 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. Then an attentive resonance develops and learning of the attended data is initiated.

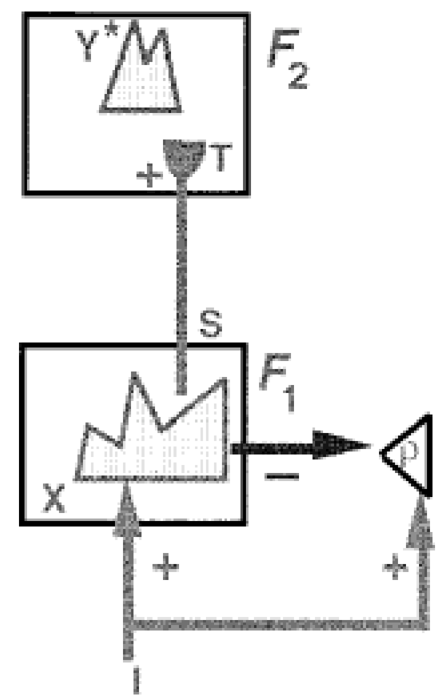

Hypothesis Testing

- In an ART model, learning does not occur as soon as some winning F_2 activities are stored in STM. Instead, F_2 activities can be interpreted as "making a hypothesis" about input at F_1.
- The LTM-gated signals from all the active F_2 nodes are added to generate the total top-down feedback pattern from F_2 to F_1. This pattern plays the role of a learned expectation. Activation of this expectation "tests the hypothesis" of the active F_2 category.
- The expected prototype of the category are compared with the bottom-up input pattern (exemplar) to F_1.

ART Learning

- The resonant state, rather than bottom-up activation, drives the learning process (thus "adaptive resonance" theory):
  - "resonance" = mutual reinforcement between input and storage layers
  - "adaptive" = weights are adjusted when resonance occurs
- ART systems learn prototypes rather than exemplars because the attendant feature vector ***X**** rather than the input exemplar itself is learned.

2 of 3 Rule

- The ART attentive matching process is realized by combining bottom-up inputs and top-down expectations with a nonspecific arousal process that is called attentional gain control – an F_1 node can be fully activated only if **two of the three input sources** that converge on the node send positive signals to the node at a given time – thus the constraint is called the "2/3 rule".

- It is the 2/3 rule that enables an ART network to solve the stability-plasticity dilemma.

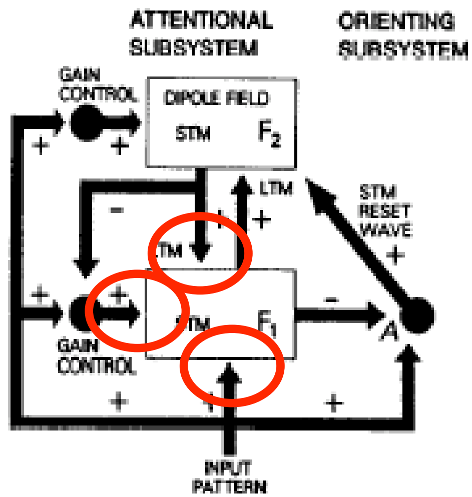

ART1 Viewpoints

- Two kinds of explanations:
  - neural - as just seen
  - algorithmic

- The first is more complicated, since it involves neural explanations for the control aspects of the algorithmic approach.

Basic ART *Algorithm*

- Input pattern presented to input layer.
- Storage layer indicates tentative hypothetical classification.
- Input layer decides if hypothetical is **close enough**; if so, done.
- If not, storage layer indicates alternate hypothesis.
- The above two steps are repeated until the hypothetical classification is accepted.
- *All* hypotheses could be rejected; in this case, a **new** class is created in the storage layer.

ART1 (Discrete Patterns)

- training pattern $x \in \{0, 1\}^n$
- prototype patterns $w_j \in \{0, 1\}^n$
- Storage unit computes
  $y_j = w_j x / || w_j ||^2$ for each prototype
- The winner is the prototype with the largest y_j
- For **acceptance** it is **necessary** that, $y_j > || x ||^2 / n$, where n is the number of dimensions.
- This means that sufficiently-many bits must **match**.

ART1

- Assuming that the acceptance test is passed, it is also required that the **vigilance test** also be:

$$w_j\, x \,/\, ||\, x\, ||^2 > \rho$$

where ρ is an adjustable parameter called **vigilance**.

This means that the input and the pattern share a sufficient fraction of 1's.

ART1

- If the acceptance test is passed, but the vigilance test is not:
  - the prototype in question is temporarily omitted from consideration;
  - a new competition takes place
- until a prototype is found for which both tests are passed.
- If no prototype is found, a new class k is created with

$$w_k = x$$

ART1

- The **higher** the **vigilance**, the more likely a new pattern is to be introduced.

- Lower vigilance will allow one input to pass as another pattern.

ART1 Issues

- Subset-Superset dilemma:
  - If one pattern is **contained in** another, then a given input may have the same inner product with two different prototypes.
  - Can be resolved by allowing weights other than {0, 1} and **normalizing** the prototypes.

ART1 Command-Line Demo

- /cs/cs152/art1/
- art1 < letters.50.in
- Number is % vigilance coded in file.

Matlab ART1 Demo

Increasing **vigilance** causes the network to be more selective, to introduce a new prototype when the fit is not good.

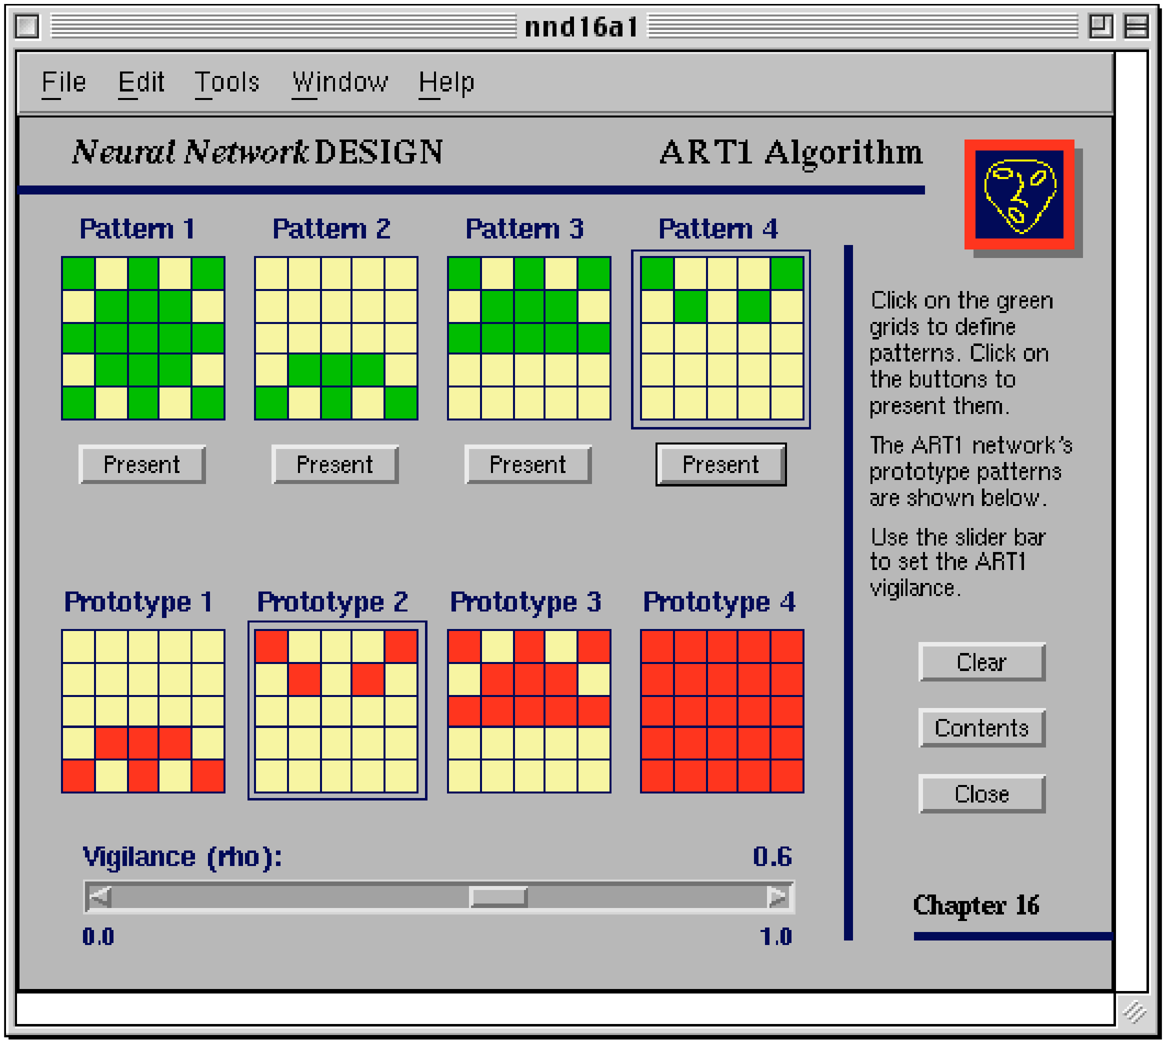

Try different patterns

ART: instars and outstars

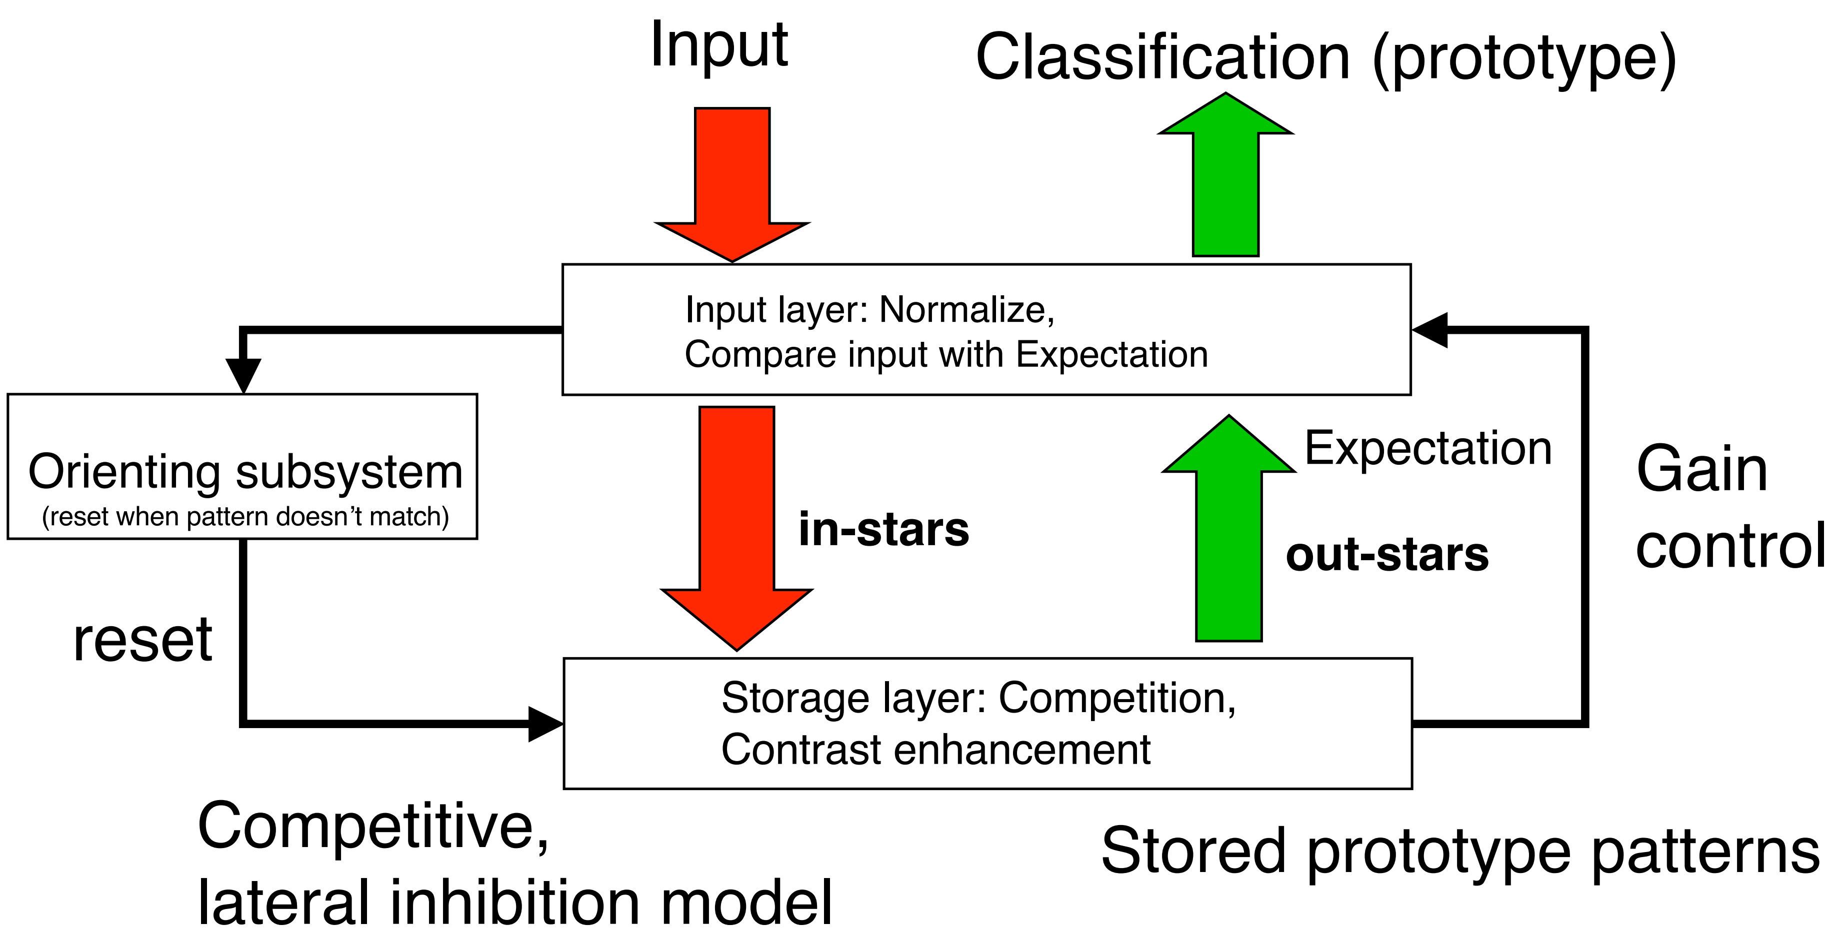

ART2 Clustering (higher vigilance)

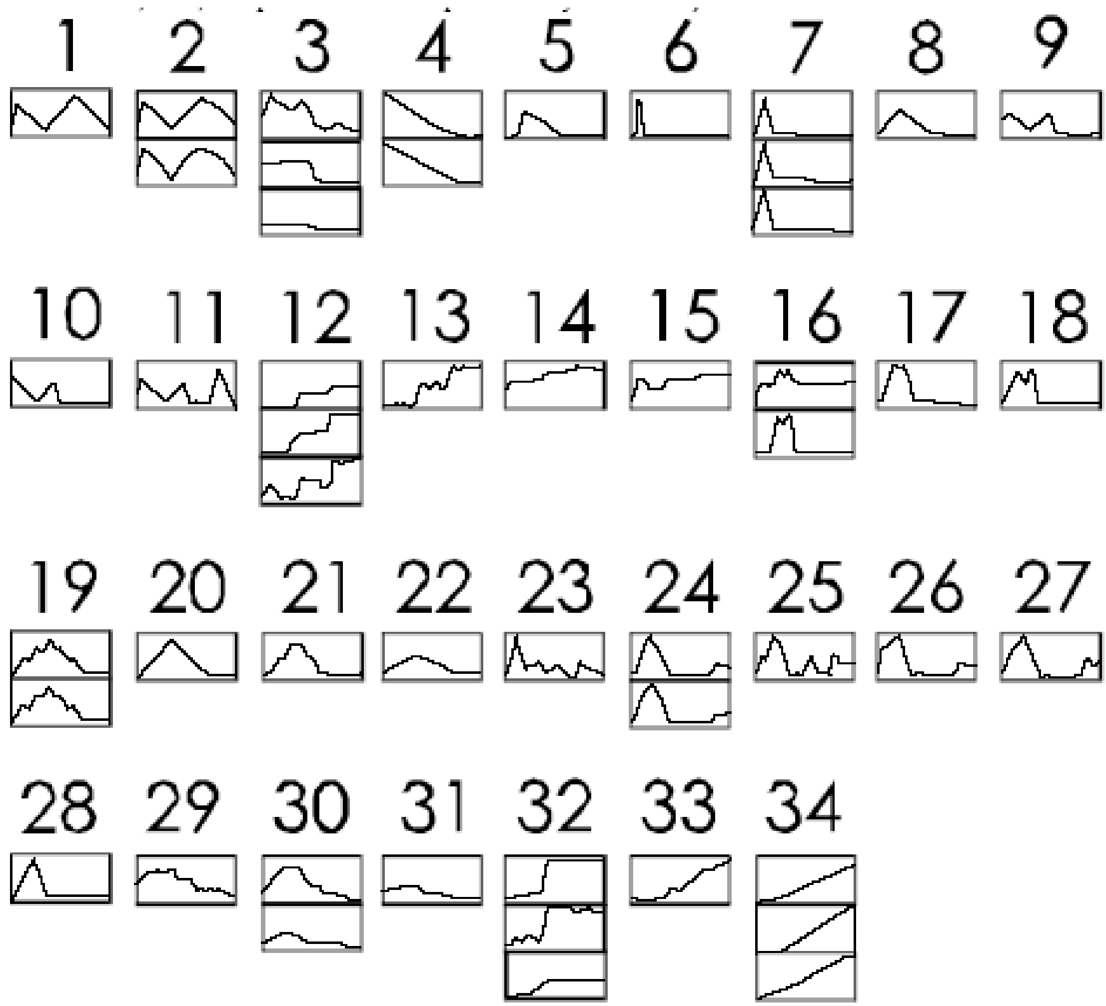

ART2 Clustering (lower vigilance)

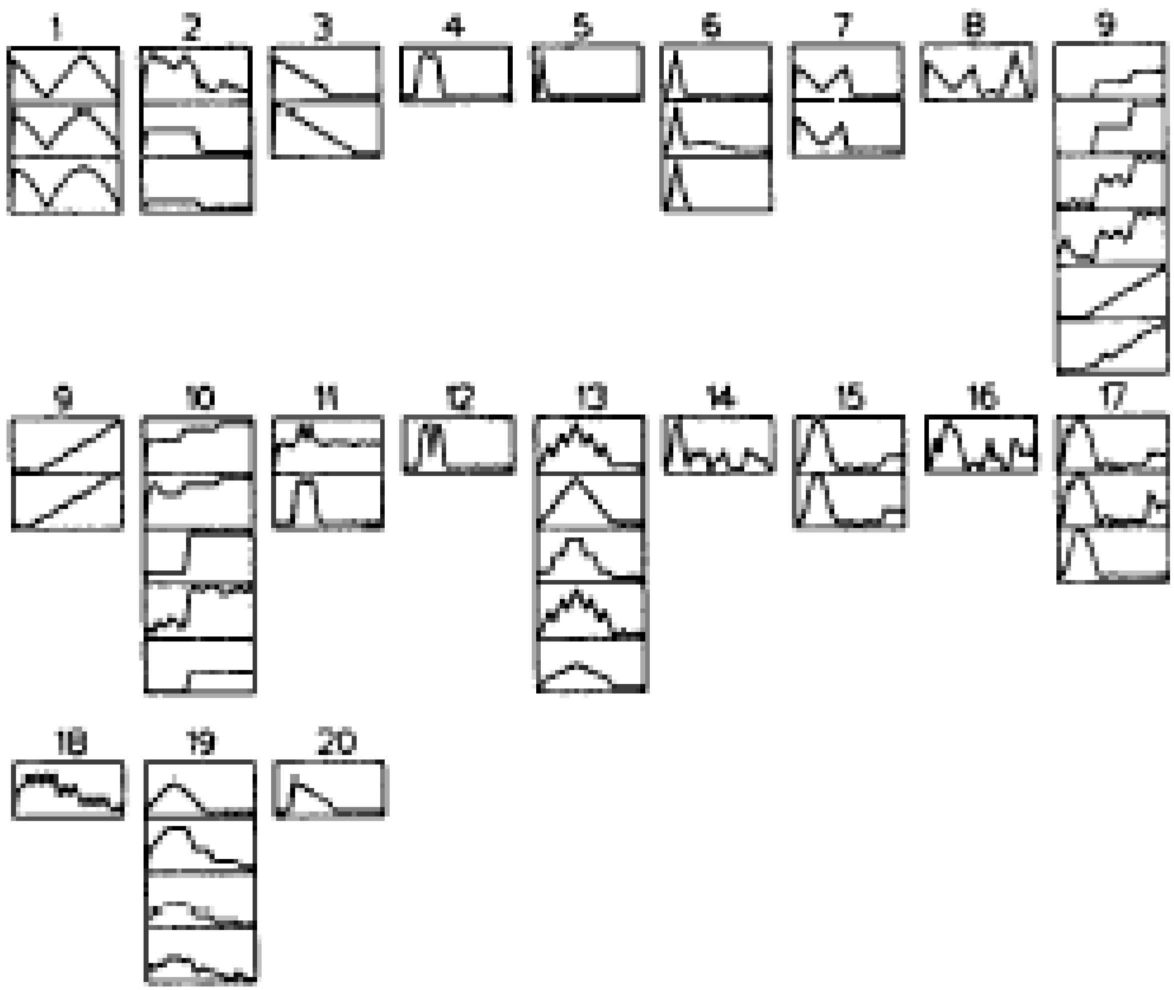

ART Map

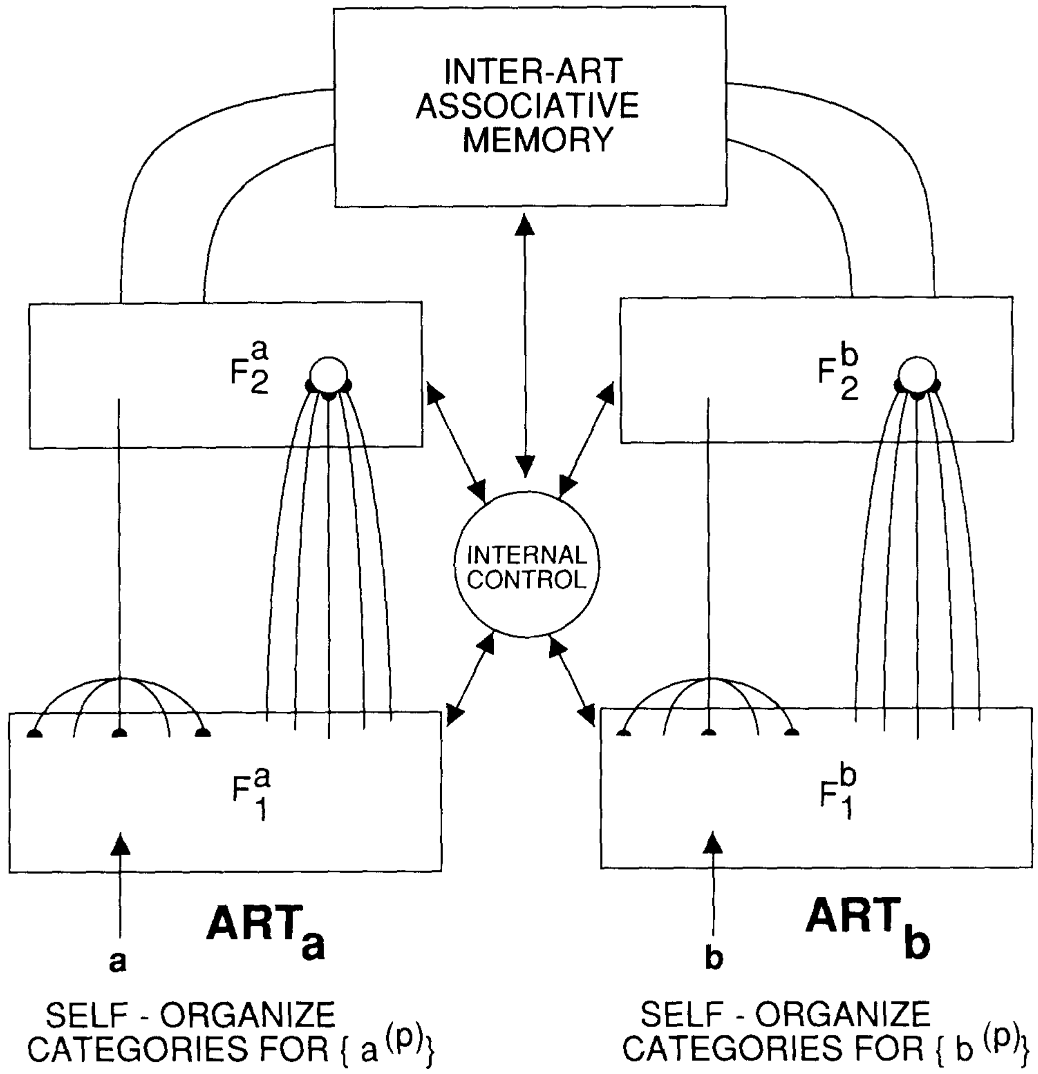

INTER-ART
ASSOCIATIVE
MEMORY
F_2^a
F_2^b
INTERNAL
CONTROL
F_1^a
F_1^b
ART_a
ART_b
a
b
SELF - ORGANIZE
CATEGORIES FOR { a (p)}
SELF - ORGANIZE
CATEGORIES FOR { b (p)}

ART Map vs. BackProp

	Predictive ART	Back Propagation
supervised self-organizing real-time self-stabilizing	yes yes yes yes	yes no no no
learning:	fast or slow match	slow mismatch

Some ART Applications

- Disease identification (HMC Math Clinic)
- Tech support email automation (text similarity) (HMC CS Clinic)
- Satellite data anomaly detection (HMC CS Clinic)
- Music recognition
- Distinguishing poisonous vs. edible mushrooms
- Modeling biological neural processes