Adaptive Resonance Theory (ART)

Davide Bacciu

Dipartimento di Informatica
Università di Pisa
bacciu@di.unipi.it

Applied Brain Science - Computational Neuroscience (CNS)

Incremental Learning in Competitive Networks

- Competitive networks have neuron competing to learn representations of the input stimuli
  - Competition serves to allow neurons to encode different information
  - Two step process with long-range competition and short-range reinforcement (competitive Hebbian learning)
- Incremental learning problem
  - A network should be able to continuously learn from new data
  - Need capabilities to discern what is new (and deserves being learned) from what is old
  - Running out of memory capacity

Stability-Plasticity Dilemma

- Incremental learning requires to address the Stability-Plasticity Dilemma
  - How can a network learn quickly and stably new information without catastrophically forgetting its past knowledge
  - Concept introduced by Stephen Grossberg in 1980
- The Adaptive Resonance Theory (ART, Grossberg 1978)
  - The human brain is very good at solving this dilemma hence we seek inspiration in neurobiology
  - How do the synapses and neurons self-organize to quickly represent information coming as a continuous flow?
  - More of a theory of learning rather than a specific model, based on the concept of resonance

The Adaptive Resonance Theory (ART)

A cognitive and neural theory of how the brain can quickly learn and stably remember and recognize, objects, sounds, events, etc. from a stream of continuous stimuli

- Two key ingredients
  - Category - Abstracted representation of coherent/similar input stimuli encoded in some high level neuron structures
  - Resonance - The synchronous firing activated by hypothesis search when a stimulus matches well an existing category and that enables quick learning
- Originally proposed as a fully unsupervised learning theory
  - Multi-layered competitive learning networks
  - Can incrementally add new neurons when existing ones do not encode sufficiently well the stimulus

ART Incremental Learning Approach

Precondition

- Assume input samples, so far, have been encoded in k categories
- A weight vector $\mathbf{w}_j$ is used to represent the typical stimulus encoded by the category

When a new input vector $\mathbf{x}$ arrives

1. Find the winner j^* among the k category neurons
2. Compare $\mathbf{w}_{j^*}$ with $\mathbf{x}$
   - if they are sufficiently similar ($\mathbf{x}$ resonates with category j^*) then update $\mathbf{w}_{j^*}$ with $\mathbf{x}$
   - else, find/create a free category unit and assign $\mathbf{x}$ as first member

Addressing the Stability-Plasticity Dilemma

- Standard weight decay in competitive learning does not solve the problem
  - Prevents weights to be erased by new incoming data (stability)
  - Freezes learning after a while (no plasticity)
- ART solution
  - A mechanism that checks if current input is a good enough exemplar of winning category
  - Assess both match and mismatch
  - A top-down reviewing of the bottom-up activated response

Basic ART Architecture

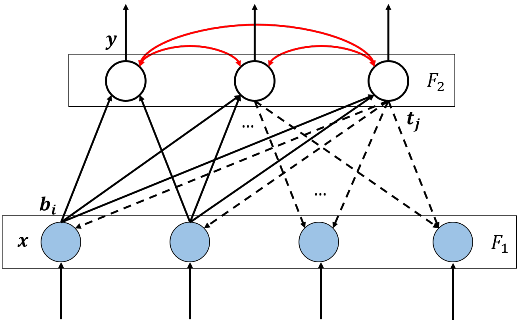

- Inputs **x** and categories **y**
- Bottom-up weights $\mathbf{b}_i$
- Top-down weights $\mathbf{t}_j$
- Lateral competition on categories

- ART-1 - unsupervised with binary neurons
- ART-2 - unsupervised with graded neurons
- ARTMAP - supervised

Bottom-Up Phase (Recognition)

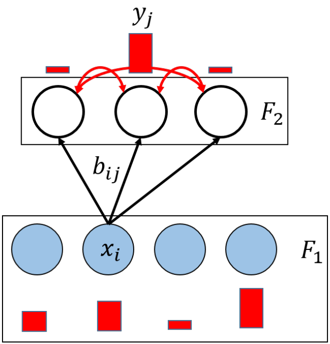

- Recognition phase
  - F_1 units simply forward inputs
  - F_2 units compete to determine winning category
- Typically a winner-takes-all (WTA) F_2 competition
  - Enables quick recognition
- Bottom-up learning rule (instar)

$$\frac{db_{ij}}{d\tau} = \alpha_b x_i y_j - \beta_b y_j b_{ij}$$

Top-down Phase (Comparison)

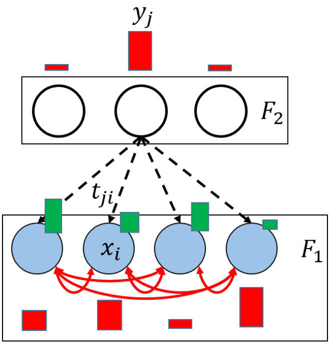

- Top-down learning rule (outstar)

$$\frac{dt_{ji}}{d\tau} = \alpha_t x_i y_j - \beta_t y_j t_{ji}$$

- Comparison phase
  - Winning F_2 unit j creates a reconstruction of the input through the top-down connections

$$\mathbf{s}_j = \mathbf{x}\mathbf{t}_j^T$$

  - F_1 compare the reconstructed vector $\mathbf{s}_j$ with actual activation $\mathbf{x}$
- A soft-competition takes place in F_1 to compare reconstructed vs actual signal

ART-1 - The Vigilance Parameter

A.k.a. soft competition in F_1

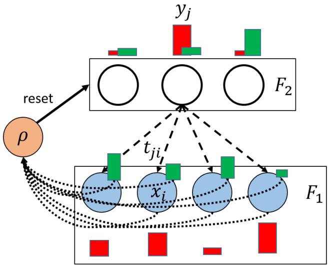

Compute overlap between the expected standard stimulus $\mathbf{s}^{j^*}$ and the actual input $\mathbf{x}$

$$\rho(\mathbf{s}^{j^*}, \mathbf{x}) = \frac{\sum_{i=1}^{N} s_i^{j^*}}{\sum_{i=1}^{N} x_i}$$

Vigilance parameter

- if $\rho(\mathbf{s}^{j^*}, \mathbf{x}) > \rho$ accept categorization and update $\mathbf{b}_{\cdot j}$ and $\mathbf{t}_j$ with current stimulus
- if $\rho(\mathbf{s}^{j^*}, \mathbf{x}) \leq \rho$ test next best category (if available) or recruit a new unit

ART-1 - The Algorithm

- Init $b_{ij} = 1/(N+1)$ $t_{ji} = 1$
- Repeat
  1. Sample a training pattern $\mathbf{x}$, compute $y_j = \mathbf{b}_j^T \mathbf{x}\ \forall j \in F_2$, $A = F_2$
  2. Repeat
     1. Find $j^* = \arg\max_{j \in F_2} y_j$ and compute $\mathbf{s}_{j^*} = \mathbf{x}\mathbf{t}_{j^*}^T$
     2. If $\rho(\mathbf{s}_{j^*}, \mathbf{x}) \leq \rho$ then $A = A/j^*$,
        else assign $\mathbf{x}$ to j^* and update weights
        $$b'_{ij^*} = \frac{s_{j^* i}}{0.5 + \sum_{i=1}^{N} s_{j^* i}} \quad \text{and} \quad t'_{j^* i} = s_{j^* i}$$
     3. Until $A = \emptyset$ or $\mathbf{x}$ is assigned
     4. If $A == \emptyset$ then allocate a new unit with weight vector $\mathbf{x}$
- Until network is stable

ART - The Detailed Picture

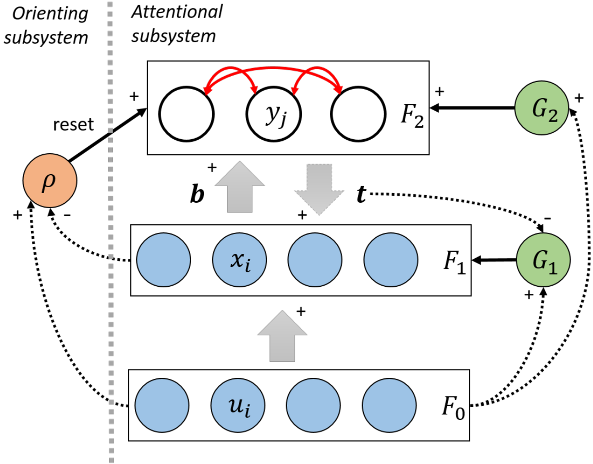

Gain units G_l serve to switch operational phases in the F_l layer

Take Home Messages

- Continuous incremental learning requires maintaining adaptivity without forgetting
  - Stability-plasticity dilemma
- Adaptive Resonance Theory
  - A family of models addressing the dilemma
  - Multi-layer competitive neural networks
  - Double checks the suitability of the encoded memory by measuring how well it can recreate the stimuli (resonance)
- Vigilance parameter determines degree of overlap accepted
  - How do I choose it?
  - What consequences can we expect from having the same vigilance for all neurons?