Designing adaptive systems that dynamically reprogram themselves based on context

A. Yanguas-Gil (PI, Argonne National Laboratory), **W. M. Severa** (Sandia National Laboratories), **D. Kudithipudi** (Rochester Institute of Technology)

Driving Applications: Lifelong Learning Machines (L2M)

Background

Current Machine Learning (ML) approaches excel in training systems to perform specific tasks. However, they struggle when asked to carry out tasks different than those they were optimized for. Even systems implementing reinforcement learning face the problem of catastrophic forgetting.

These limitations compromise the usefulness of ML for edge processing or on chip processing applications, particularly for noisy and rapidly changing environments.

DAIN in a nutshell

Can we design a system that learns and reconfigures/reprogram itself dynamically based on changing context?

Goal

DAIN will demonstrate architectures capable of the following core L2M capabilities:

1. Continual learning
2. Adaption to new tasks and circumstances
3. Goal driven perception
4. Selective plasticity
5. Monitoring and safety

DAIN's architecture:

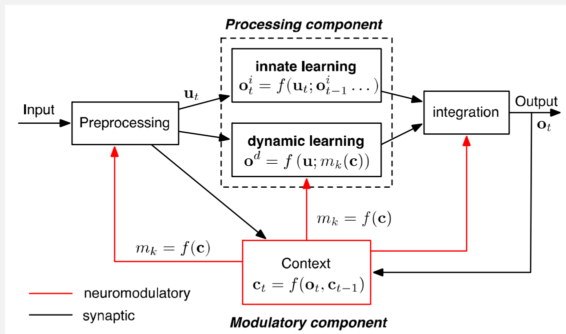

Key innovations

We believe that a key enabling capability of biological systems is the presence of a modulatory system that dynamically controls the way the neural component processes information, stores new knowledge, and carries out action selection.

Our system will:

- Use modulation as a design pattern to build adaptive systems, shifting the focus from training to dynamic reconfiguration/ programming via an external modulatory component.
- Our approach overcomes the timescale/training limitations of RNNs, by offloading slowly changing context information to the modulatory component.

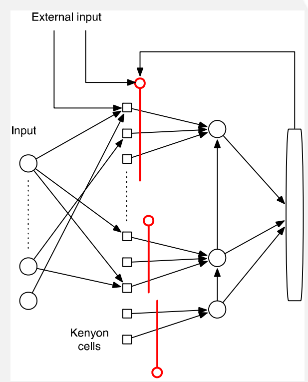

Modulatory neurons (red) are a key component of the insect's mushroom body, providing context-dependent modulation and learning

Impact

The principles developed in this project will help:

- Design broadly trained systems capable of operating in noisy, changing environments and adapting to new tasks and contexts.
- Design hardware optimized for machine learning, yet capable of adapting to changes and reconfiguring itself to carry out multiple tasks.

Contact

ayg@anl.gov

What we are doing

#1 Designing dynamic reconfigurable networks

Core idea: Through modulatory interactions we can switch on/off components and dynamically change how information is processed.

Implementation strategies:

Bilinear operators: context-dependent gains of neuron outputs and dynamic synaptic plasticity.

$$\left.\begin{aligned}\xi_i &= \sum_{j,k} w^{(2)}_{ijk} x_j x_k + \sum_j w^{(1)}_{ij} x_j + b_i \\ \xi_i &= \sum_j W_{ij} x_j + \sum_k M_{ik} m_k\end{aligned}\right\} \quad x_i = f\left(\xi_i, x_{i-1}\right)$$

Context-dependent bias: modulated dynamic thresholds

Modulated plasticity: synaptic plasticity controlled by third neuron

$$w_{AB} = f(C, A, B)$$

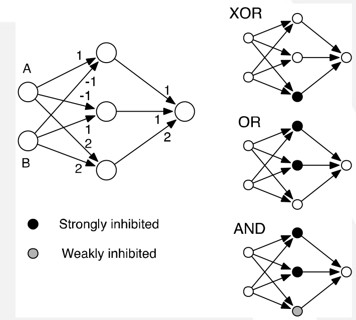

Modulatory interactions can alter the functionality without adjusting synaptic weights

Modulatory component: a small system dynamically reacting to context and external input can drive and dynamically reconfigure potentially much larger processing components for context-dependent processing and learning (including metalearning)

- Development of Turing capable dynamic systems via selective attention
- Exploration of FPGA implementations to augment ML hardware and move to RNNs

#2 Implementing and testing agent-based systems

Core idea: We are implementing an agent based model system using DAIN's architecture and developing tests inspired by bioassays found in literature.
Examples: task/goal dependent navigation, peer teaching, context-dependent learning

Example: goal driven perception enabled by modulatory system leads to complex dynamic behavior under task-independent reward rules

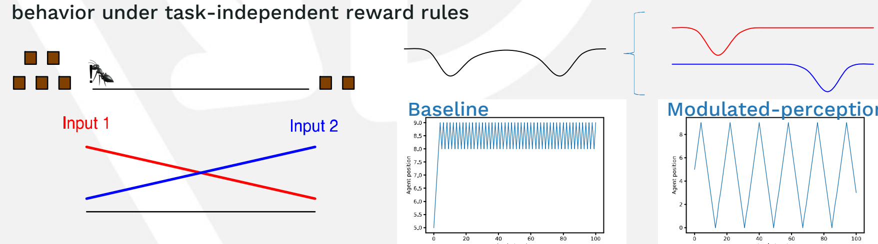

Modulatory systems implementing context dependent perception and failure driven task switching can exhibit complex behaviors even in purely feedforward systems under universal reward rules

This research was developed with funding from the Defense Advanced Research Projects Agency (DARPA)
The views, opinions and/or findings expressed are those of the author and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government

THE ELECTRONICS RESURGENCE INITIATIVE