NEUROMORPHIC **ARTIFICIAL**
INTELLIGENCE LAB

RELATIONAL NEUROGENESIS
For LIFELONG LEARNING AGENTS

by **Tej Pandit** advisor **Dhireesha Kudithipudi**

NICE-2021

Tej Pandit and Dhireesha Kudithipudi. 2020. Relational Neurogenesis for Lifelong Learning Agents. In Proceedings of the Neuro-inspired Computational Elements Workshop (NICE '20)

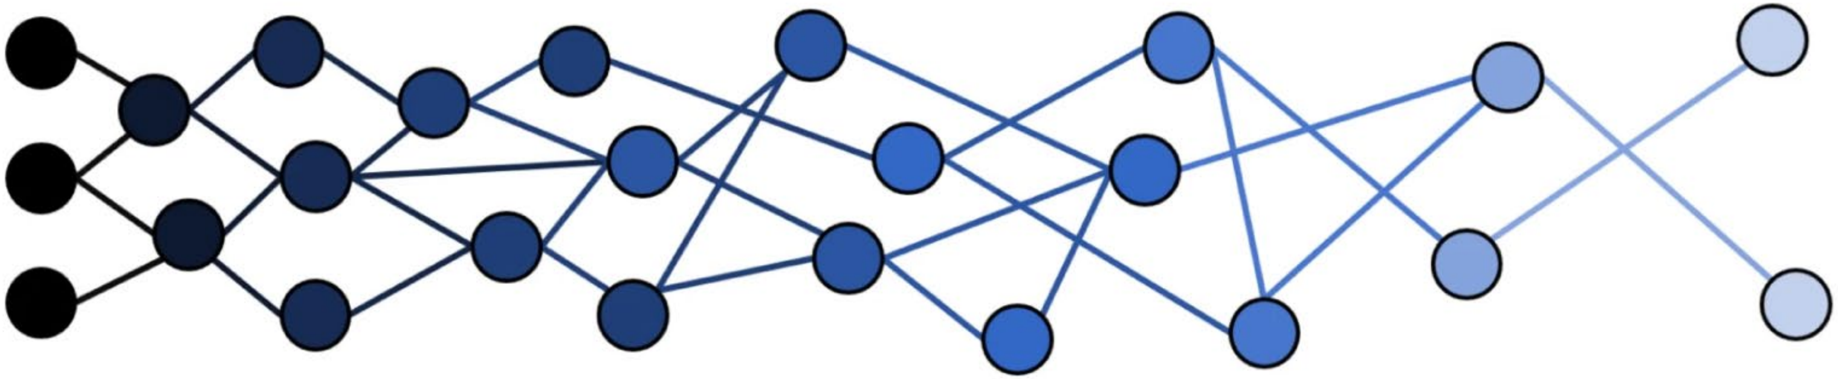

UTSA The University of Texas at San Antonio™

19th March - 2021

Lifelong-learning & Its Challenges

WHAT IS LIFELONG-LEARNING ?

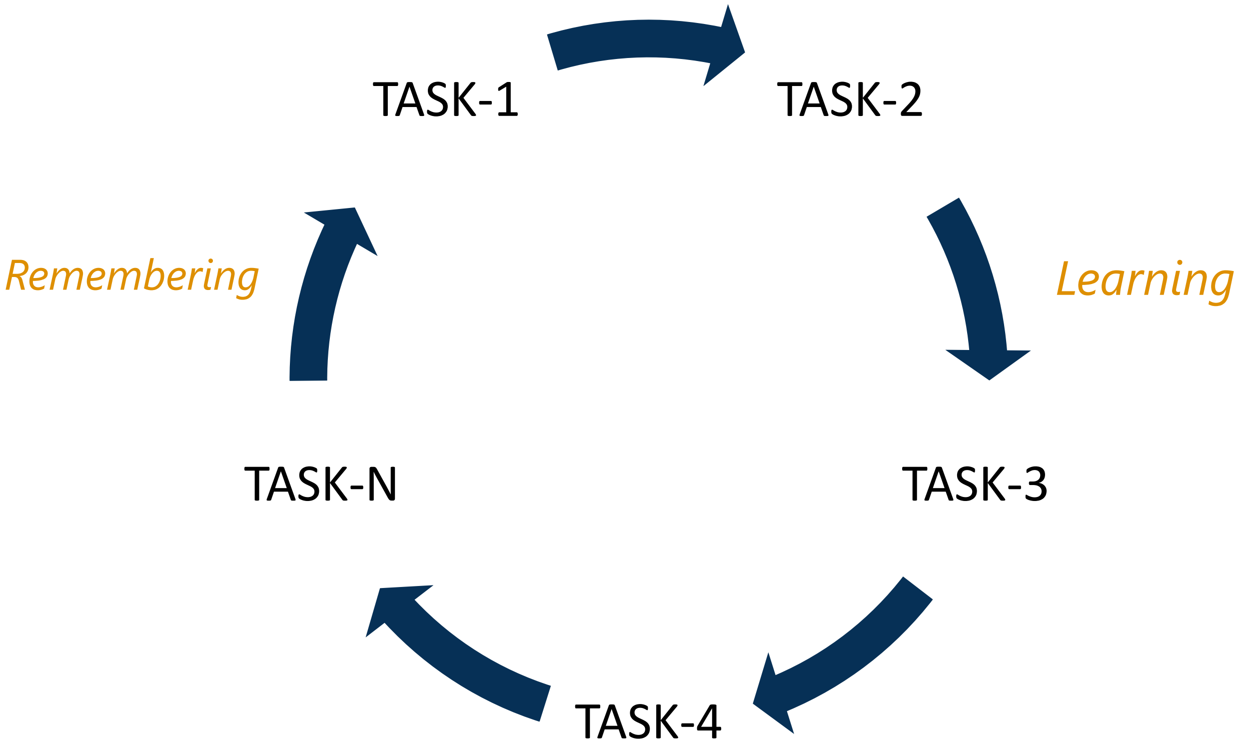

- Tasks are learned sequentially
- Ability to recollect previously learned tasks and continually learn new tasks is considered **lifelong learning**

- Inability to remember previously learned tasks is called **catastrophic forgetting**
- Learning mechanism can be supervised, unsupervised or reinforcement based

Neuroevolution

Neuroevolution is the process of evolving or modifying the architecture of a neural network

- Neurogenesis is the generation of new Neurons
  → Addition of NODES
- Synaptogenesis is the generation of new Synapses
  → Addition of EDGES
- Neuronal Death/Termination is the removal of Neurons
  → Removal of NODES
- Synapse Termination is the removal of Synapses
  → Removal of EDGES

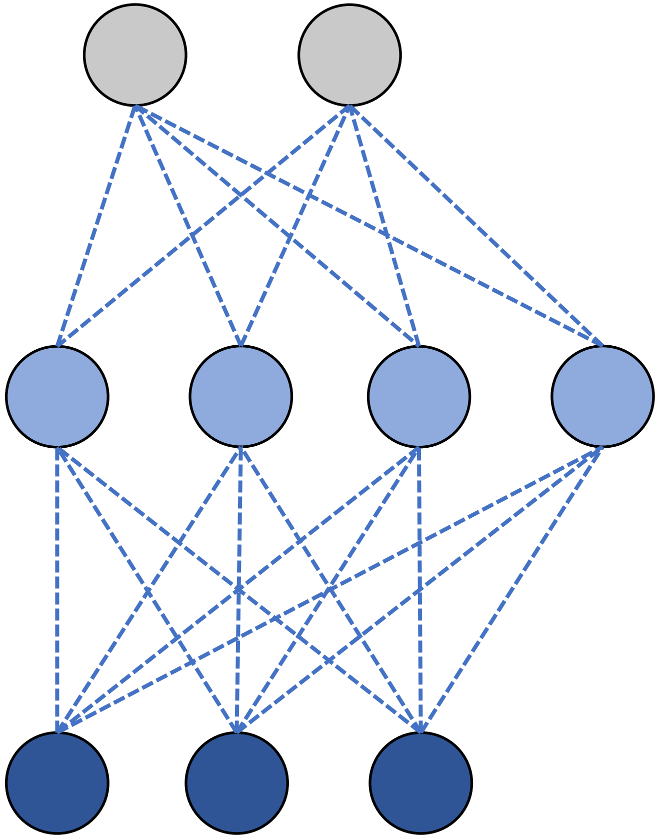

Research Questions & Contributions

I How can neuroevolution assist lifelong learning?

II In the absence of a supervisor providing context, how does evolution occur?

III Which mechanisms can aid information preservation?

I New rules for neuroevolution through neurogenesis and synaptogenesis

II Mechanism for preserving information through activity tracking

III Simulated environments for evaluating lifelong learning

IV Library for bridging reinforcement learning algorithms and simulations

Reinforcement Learning

CONTEXT

The problem is defined through a
RL context

Reinforcement Learning Example

Spider
AGENT

Fire
NEGATIVE REWARD

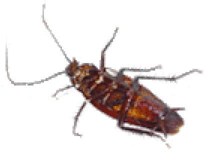

Food
POSITIVE REWARD

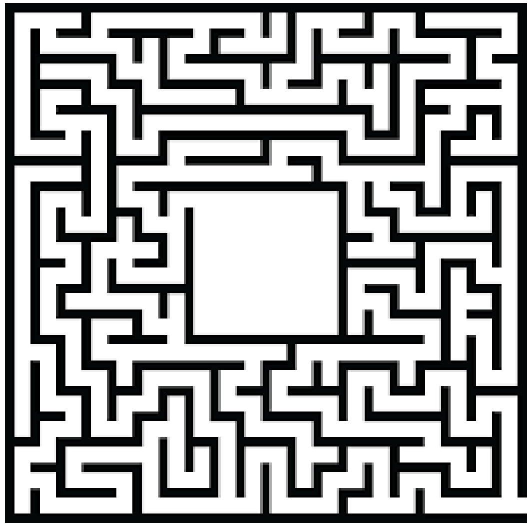

Maze
ENVIRONMENT

LIFELONG LEARNING SETUP

FOOD Env-1	→	FIRE Env-2	→	FOE Env-3	→	FRIEND Env-4
Task-1		Task-2		Task-3		Task-4

Example Environment

Auditory, Olfactory and Vibratory Sense Zones

based on radial distance from the agent.

Line-Of-Sight vision-based sense vectors

receiving distance from 5 points

inferring object type from the color of received point of contact

Reinforcement Learning Paradigm & Algorithm

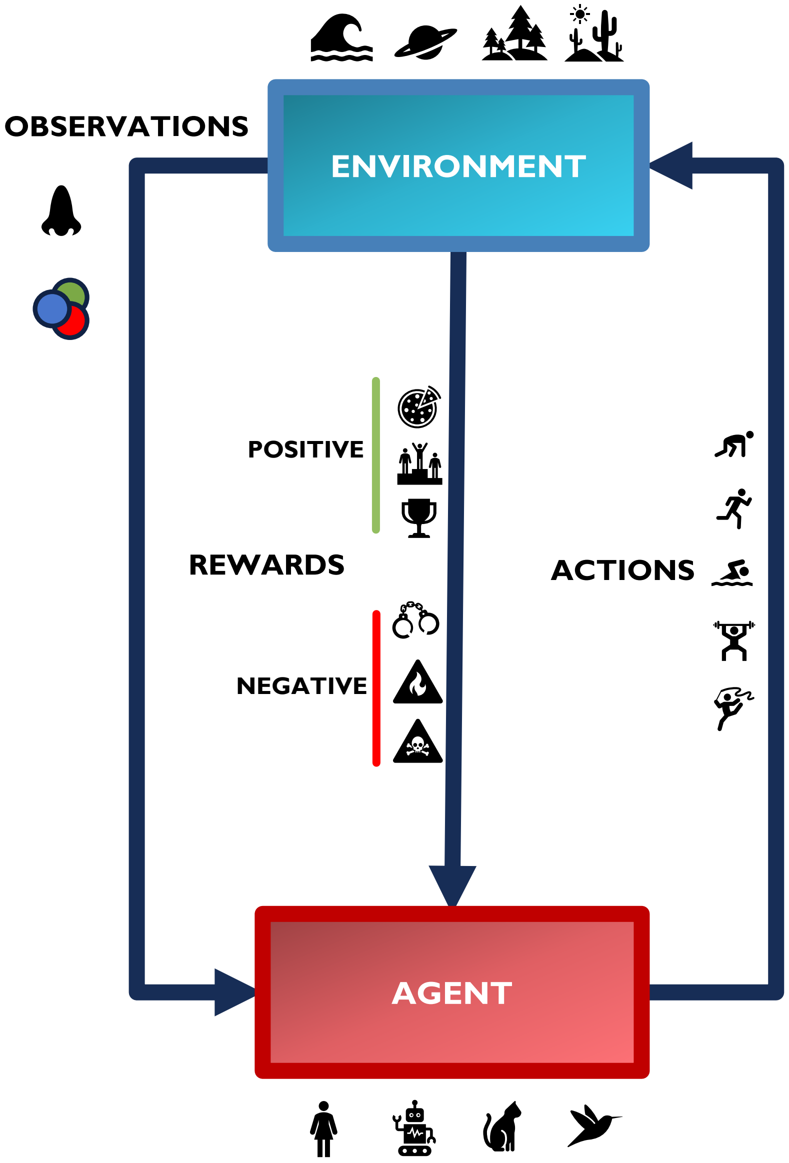

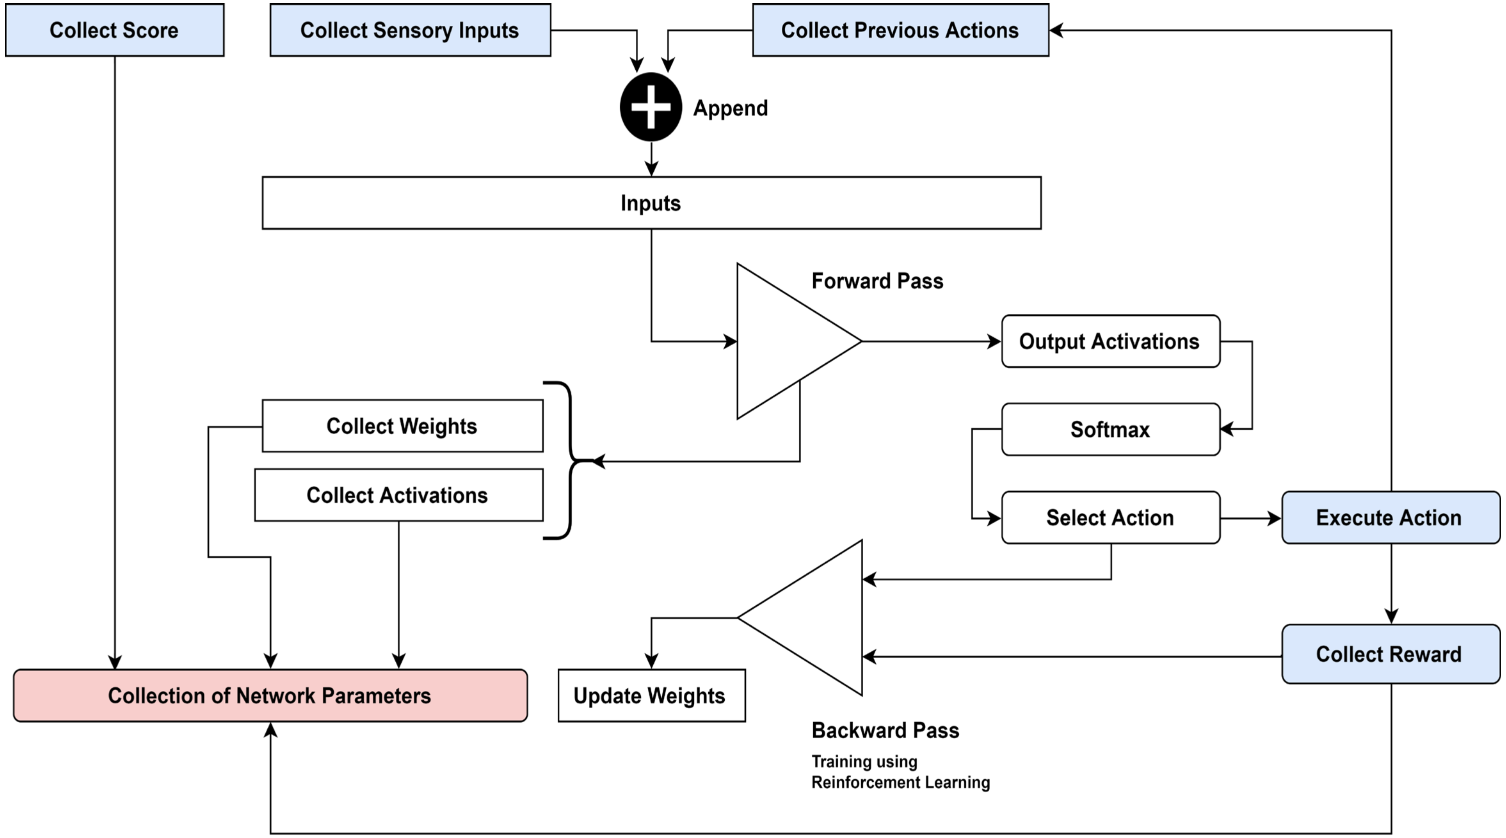

Relational Neurogenesis

ALGORITHM & ARCHITECTURE

Structure of the RN algorithm
Overview of the architecture

DRL Network

- Forward Pass

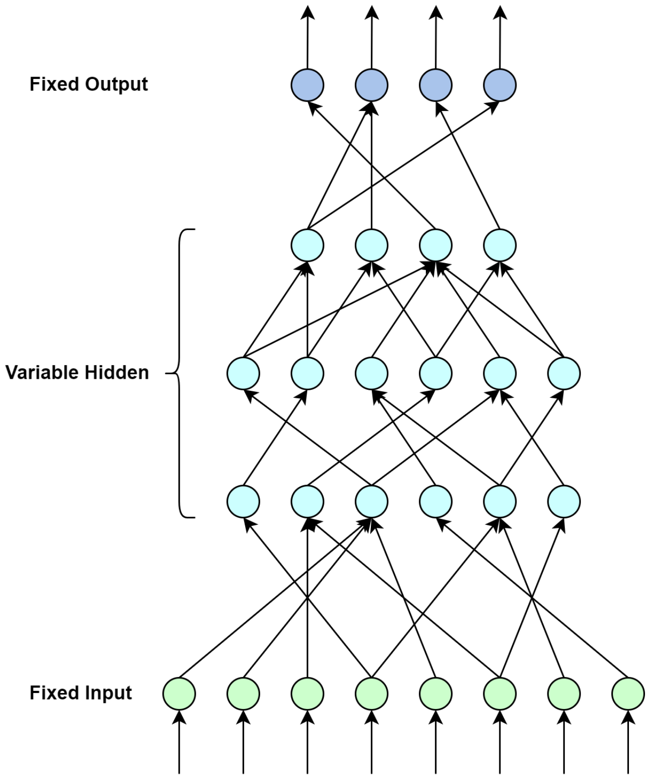

- Activation Measurement

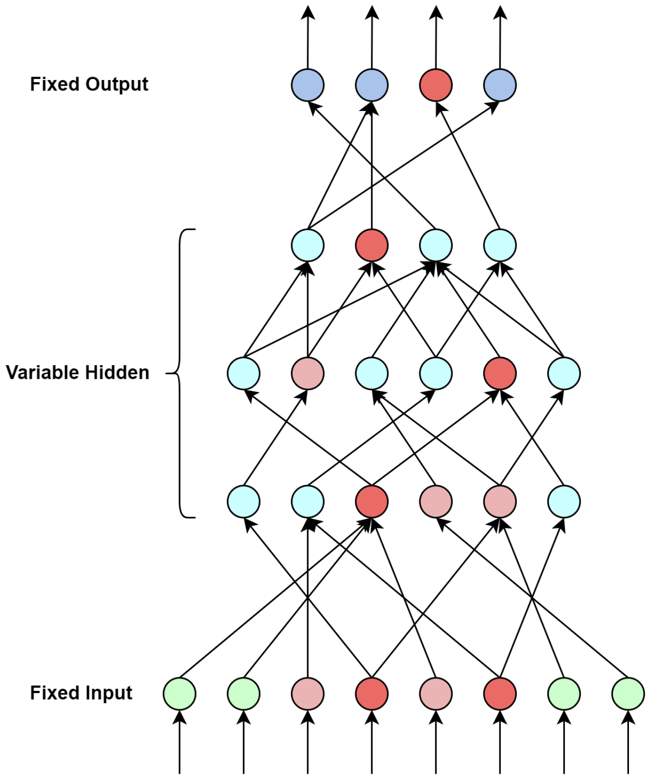

- Input Activation Matrix $[A_I]$

-	-	0.1	0.2	0.6	0.1	-	-
-	-	0.1	1	0.2	0.2	-	-
-	0.3	0.5	0.3	0.2	1	0.1	-
-	0.1	0.2	1	0.5	0.5	0.1	-
0	0	0.5	1	0.5	1	0	0

Relational Neurogenesis Framework

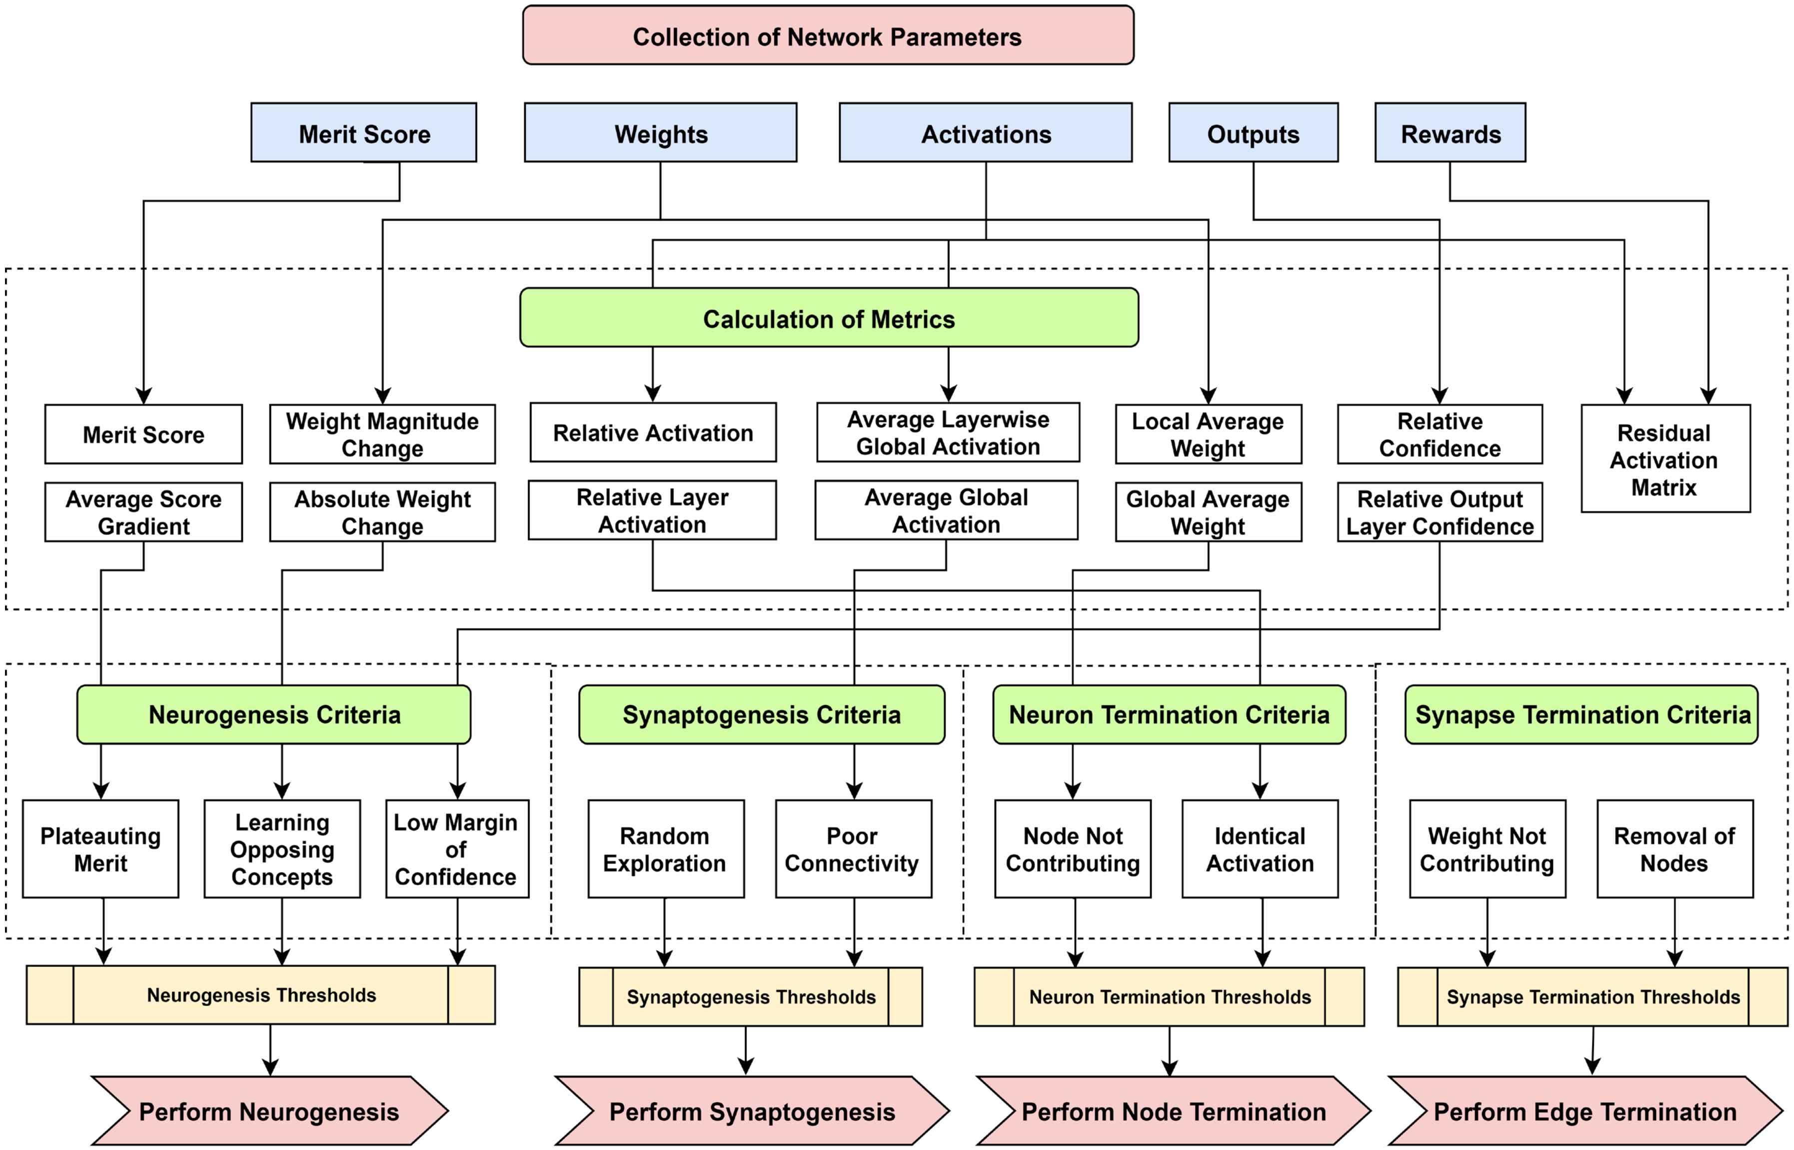

Deep Reinforcement Learning Network	
	Weights
	Activations
	Outputs
	Score

Calculate Intermediate Metrics	
	Relative Layer Activation
	Weight Change Magnitude
	Average Global Activation
	Relative Output Confidence
	Score Gradient

Compare Threshold Criteria	
	Neurogenesis
	Synaptogenesis
	Neuron Pruning
	Synapse Pruning

Relational Neurogenesis

MECHANISMS

Mechanisms and Methods developed to support Neuroevolution

Neurogenesis Mechanisms

Plateauing Merit

The Merit Score is an evaluation metric of the agent's performance.

The Merit Score should keep rising as the agent learns.

A plateauing or descending merit score curve is undesirable as it shows lack or loss of learning respectively

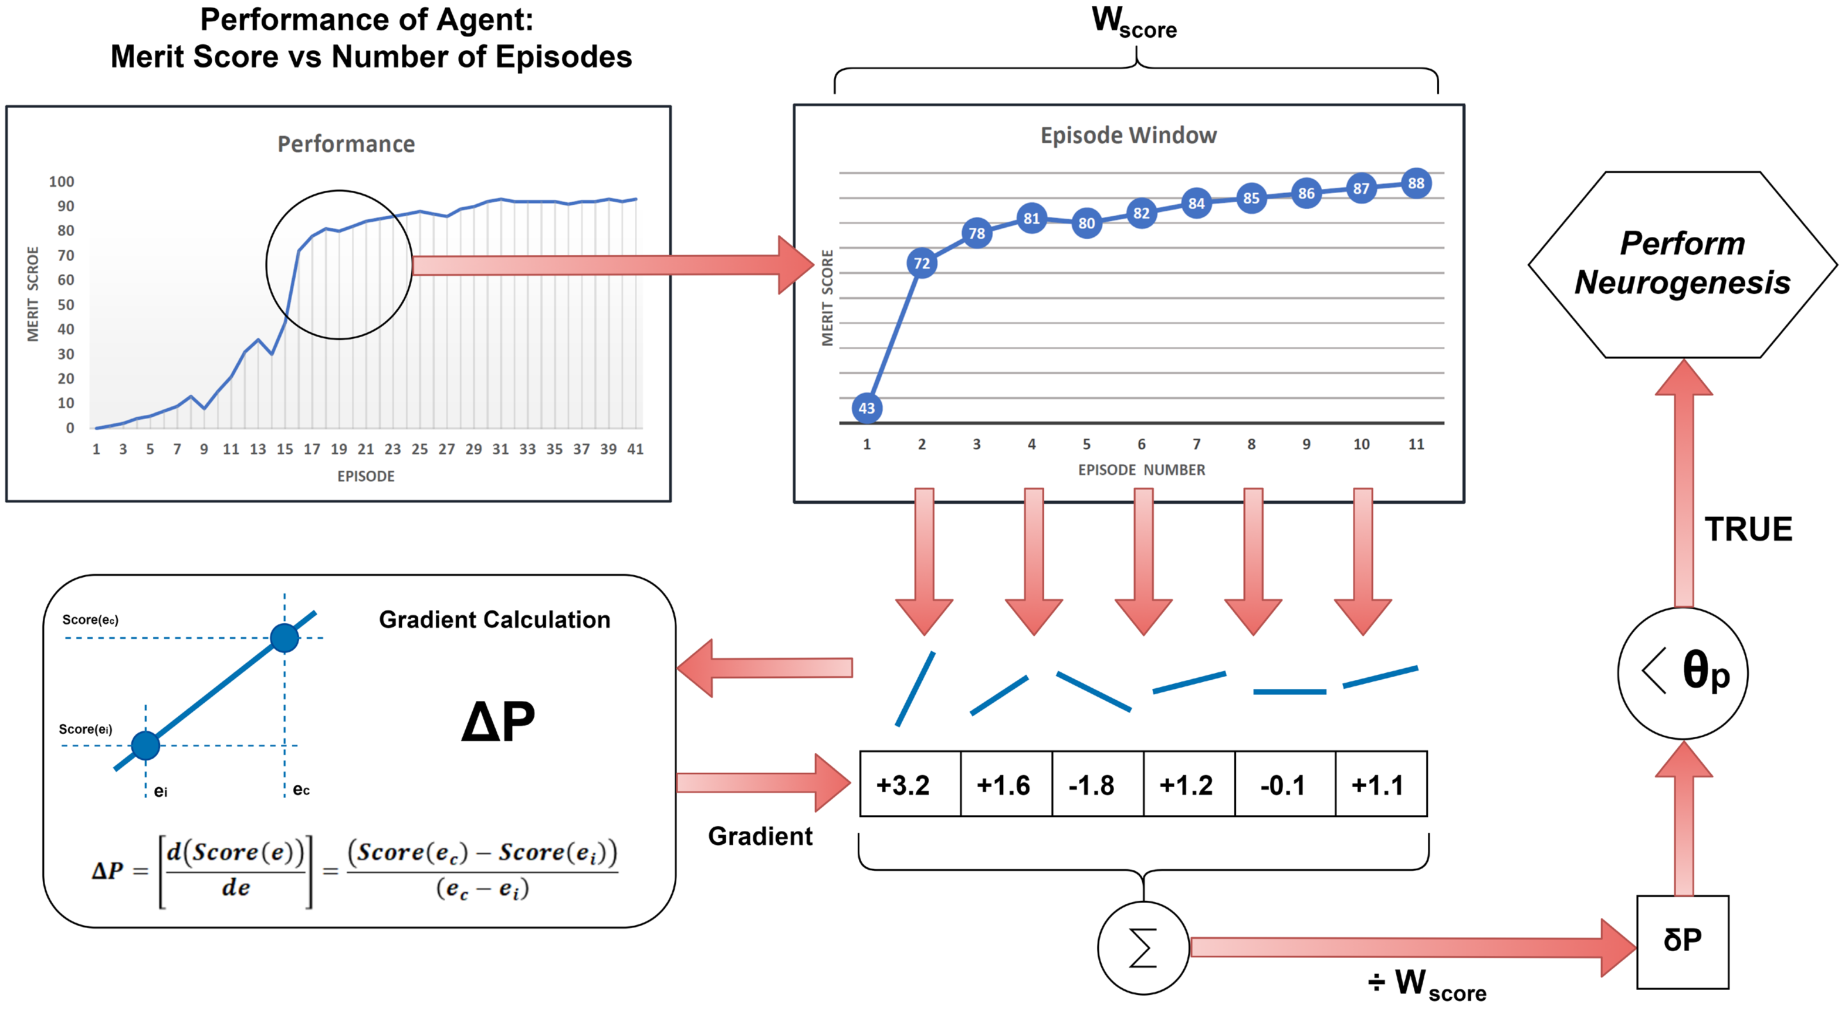

Neurogenesis Mechanisms

Learning Opposing Concepts

Learning opposing ideas or concepts causes the network weights to oscillate

The oscillation of weights results in poor representation

The oscillation of weights is caused by pulling of nodes in opposing directions

Learning Opposing Concepts

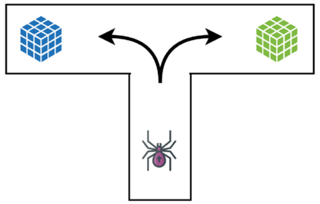

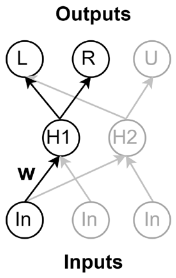

$$\Delta w(f) = w(f) - w(f-1)$$

Δw = change in weight

$|\Delta w|$ = change in magnitude

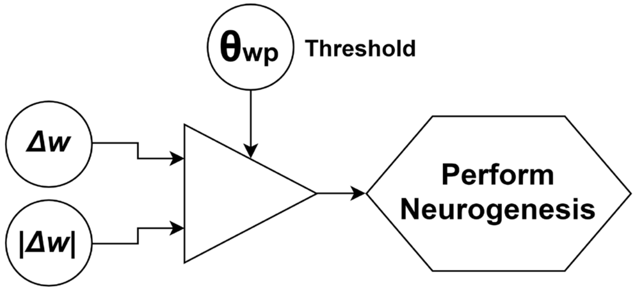

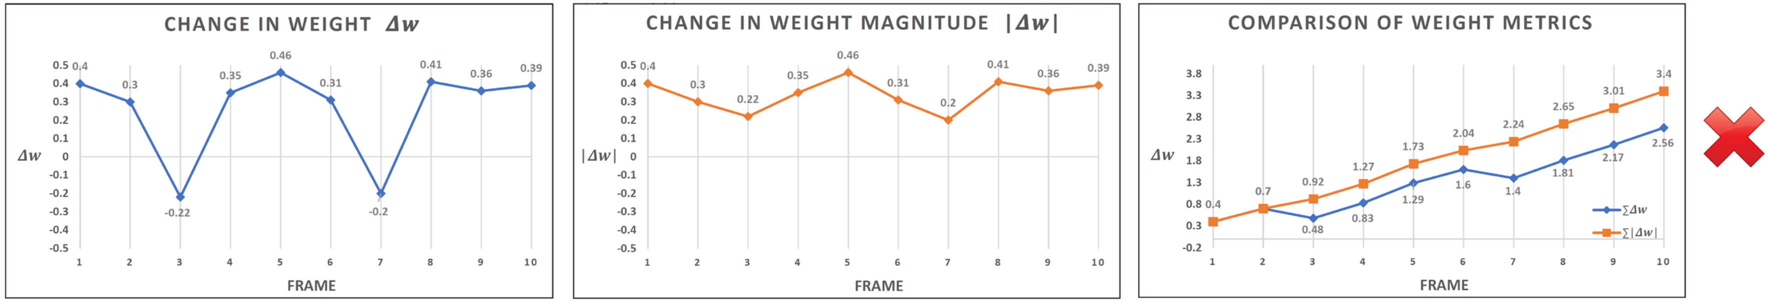

Minimal Weight Oscillation

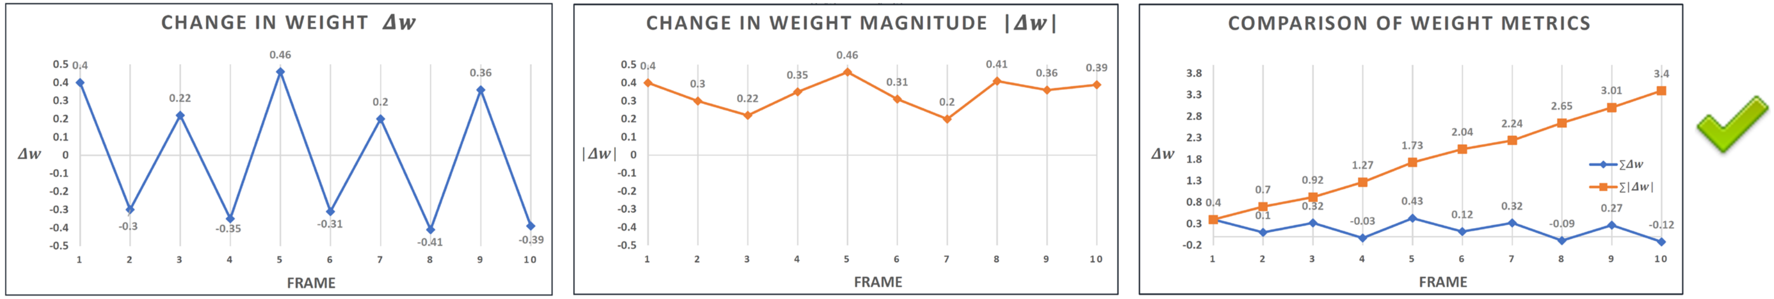

Large Weight Oscillation

Neurogenesis Mechanisms

Low Margin of Confidence

The output activations are passed through a softmax

The resultant can be viewed as a confidence metric of each output

If the confidence is low or similar between outputs, it shows poor class separability and additional nodes are needed to separate the classes

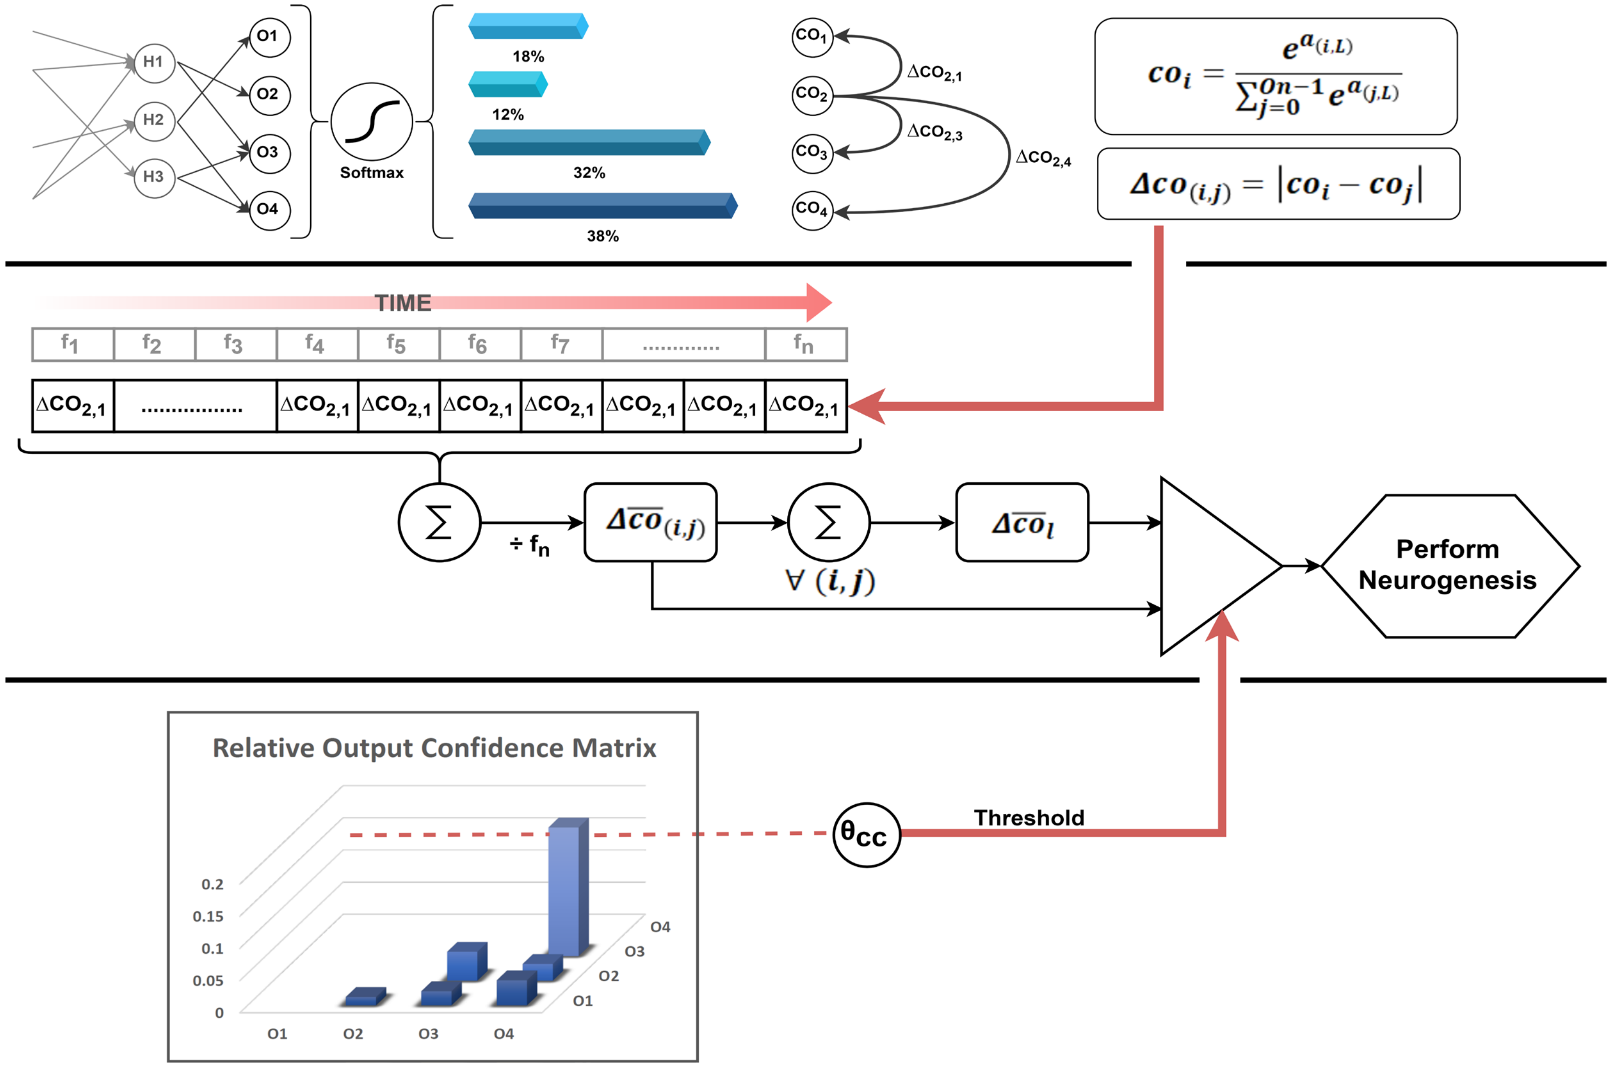

$$co_i = \frac{e^{a_{(i,L)}}}{\sum_{j=0}^{On-1} e^{a_{(j,L)}}}$$

$$\Delta co_{(i,j)} = |co_i - co_j|$$

Simulated Environments

EVALUATION

Virtual Environments developed in
Unity Engine for evaluating
Lifelong Learning

Environment Set – A *Spider Survival*

Before Training

Episode 1

Environment – A1

FIRE

After Training

Episode 500

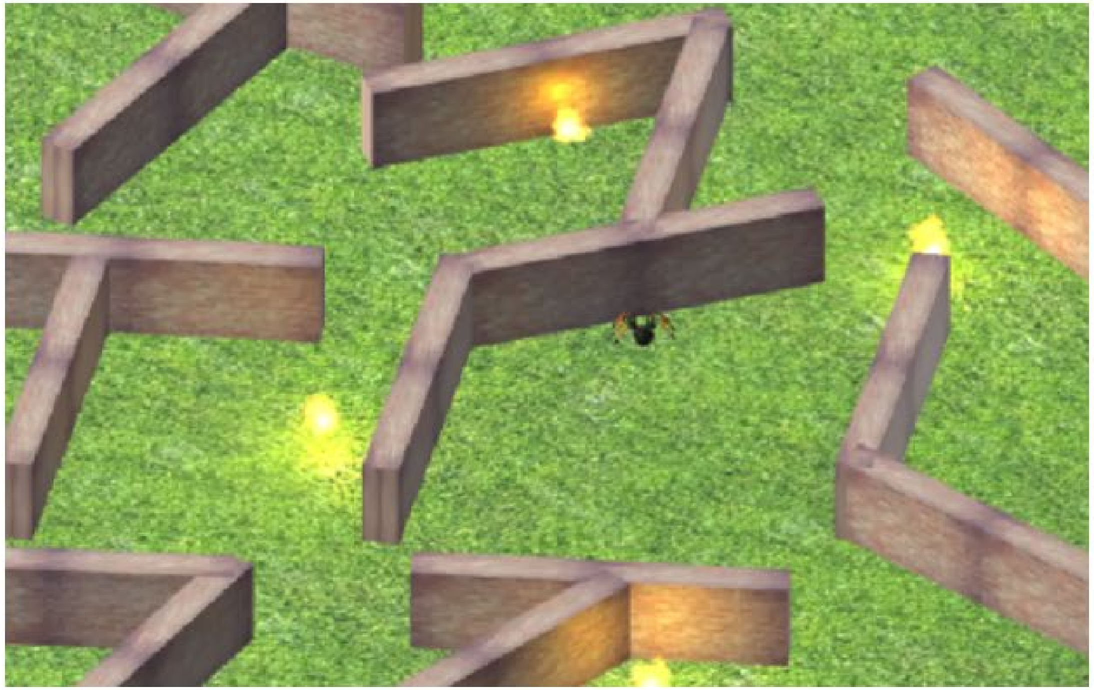

Training Time

60 minutes

Environment Set – A *Spider Survival*

Environment - A2

FOOD

Environment - A4

FRIEND

Coexisting

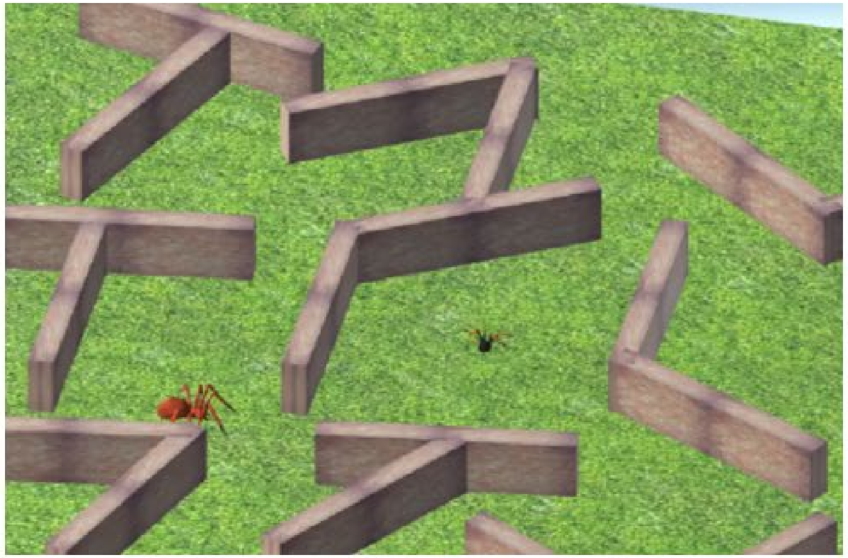

Environment - A3

FOE

Environment - A4

FRIEND

Competing

Environment Set – B *Forest Fire*

AGENT

Agent Drone

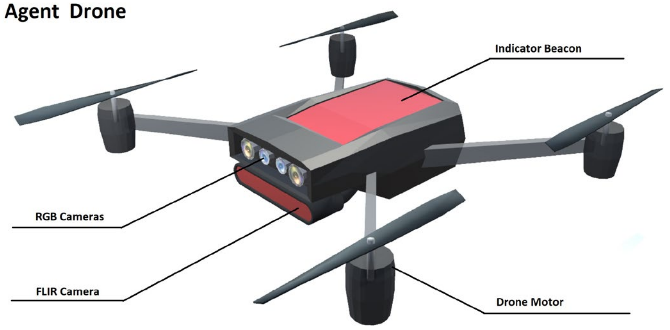

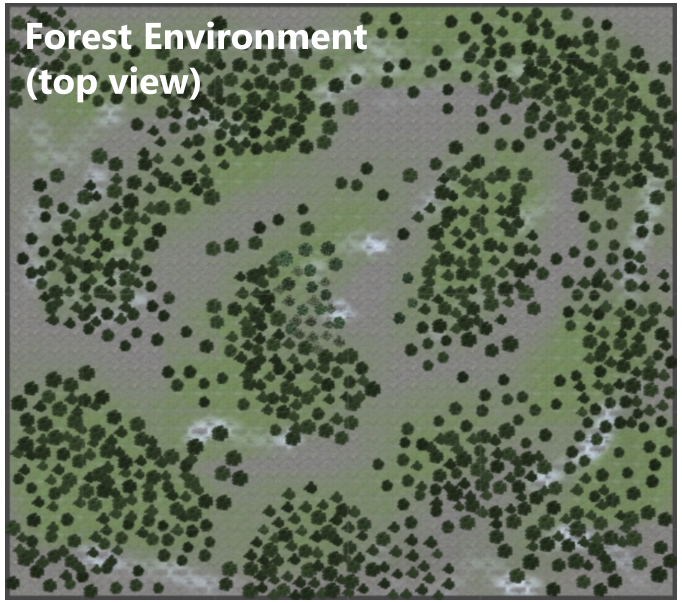

- Task I

  Navigating Forest Fires

- Task II

  Locating Trapped Civilians

- Task III

  Rescuing Civilians

- Task IV

  Multi-agent cooperation

Environment Set – B *Forest Fire*

Trained on B1

Successfully Navigating Forest Fires

Trained on B2

Successfully Locating Trapped Civilians

RGB Camera (HD)

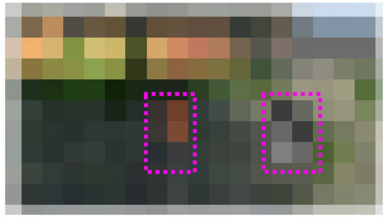

Processed Image (16x9)

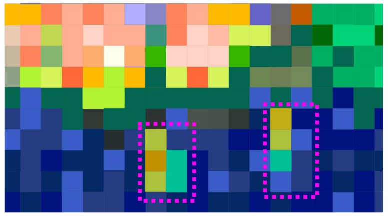

FLIR Camera (16x9)

Object Detection from Visible and IR Spectra

Results

EVALUATION

Results obtained in single task as well as continual learning scenarios

Results: Single Task Performance

RN is compared against RL and Continual Learning Algorithms:

I. Observed to learn at an **accelerated** pace

II. **Nearly matches** SOTA performance in individual task

III. Heavy **computational** overheads

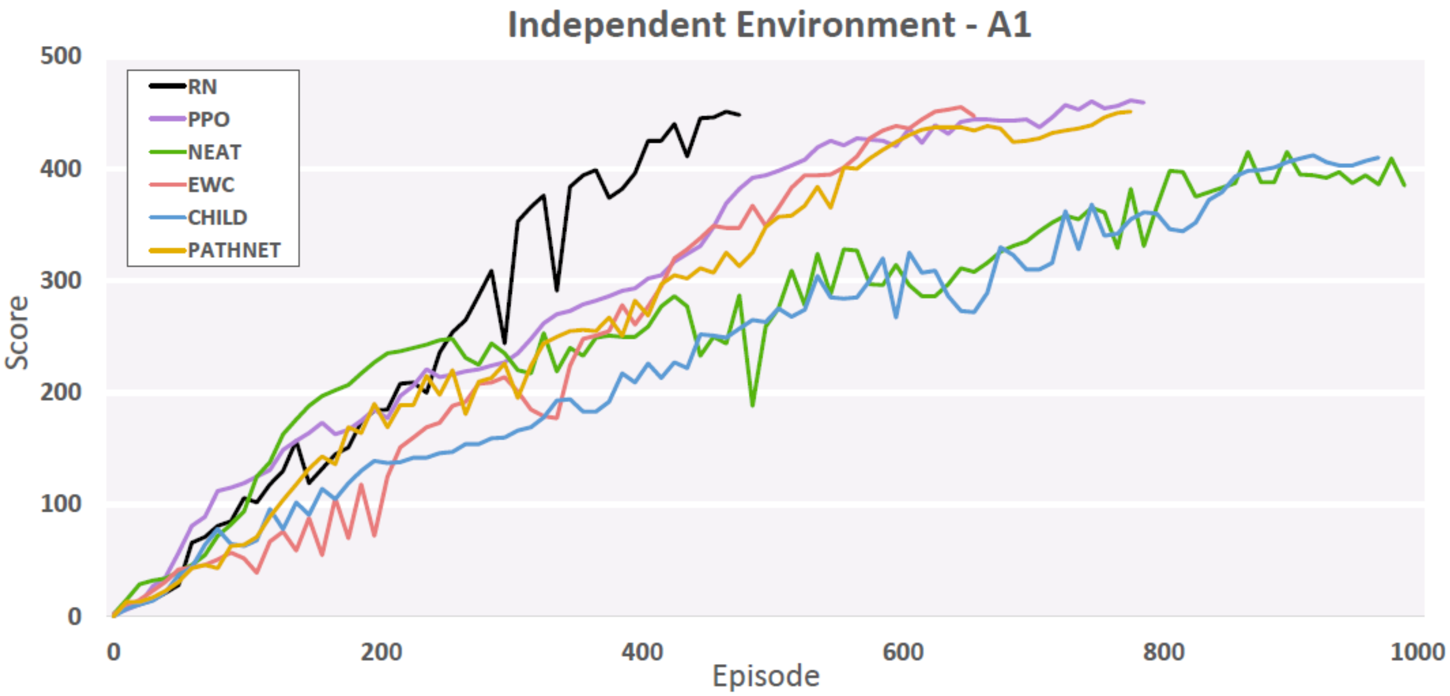

Network Growth

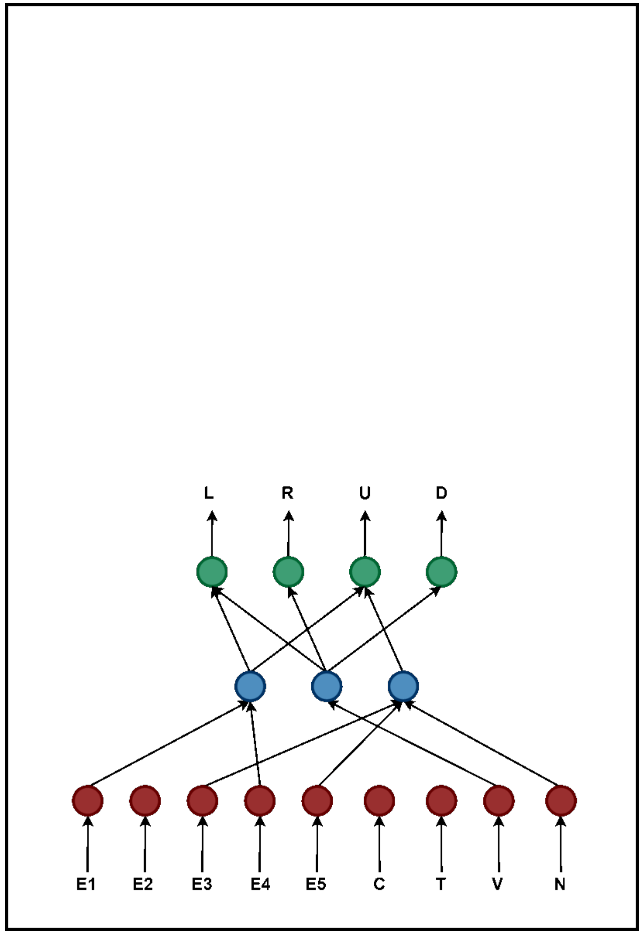

Results: Continual Learning Performance

Continual Learning Performance

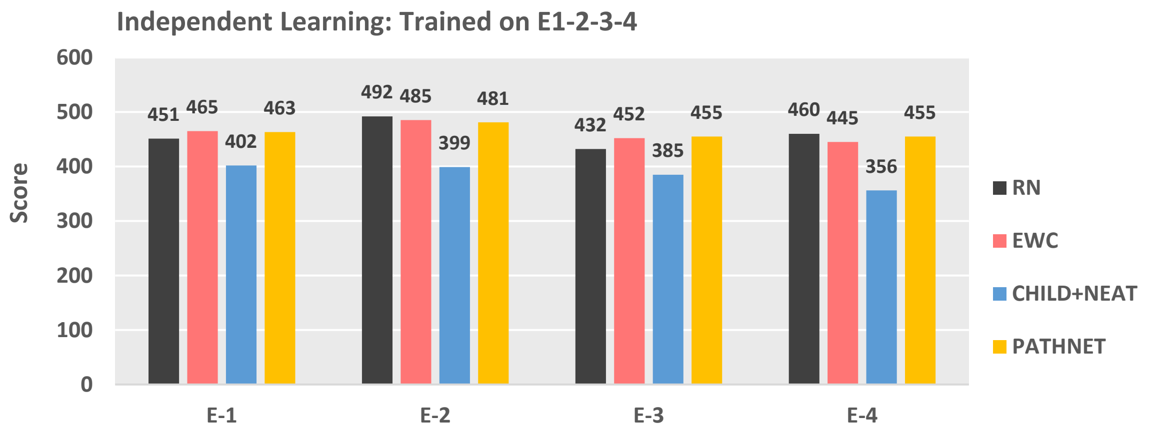

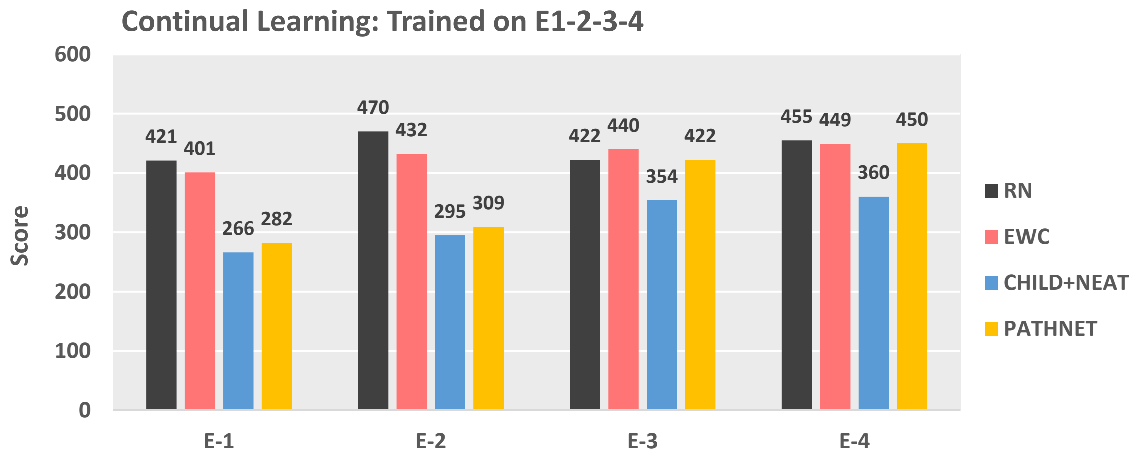

Network Analysis

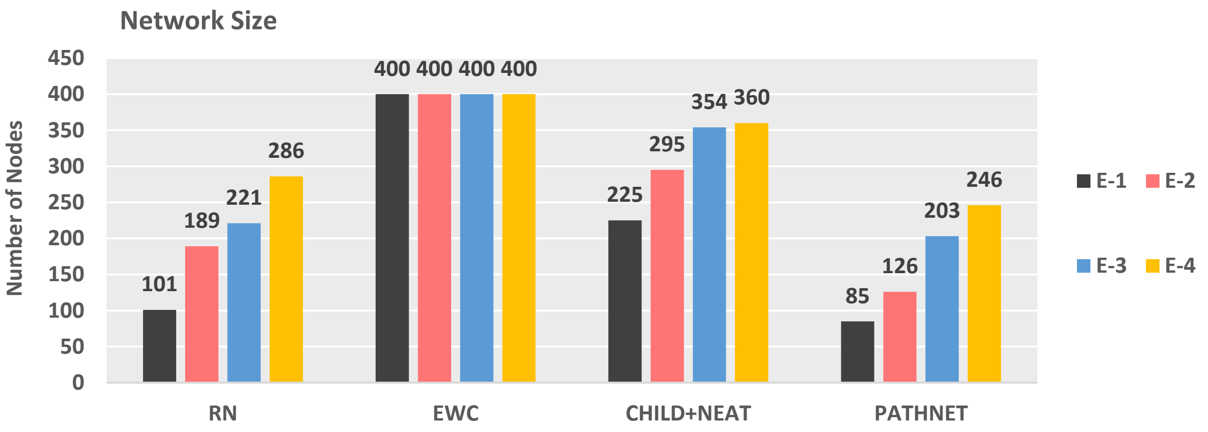

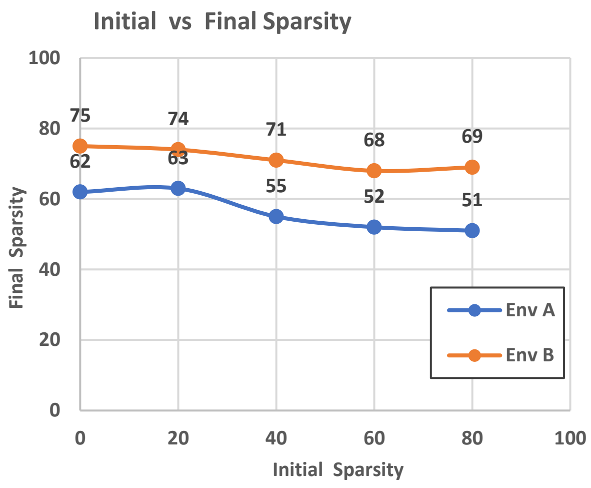

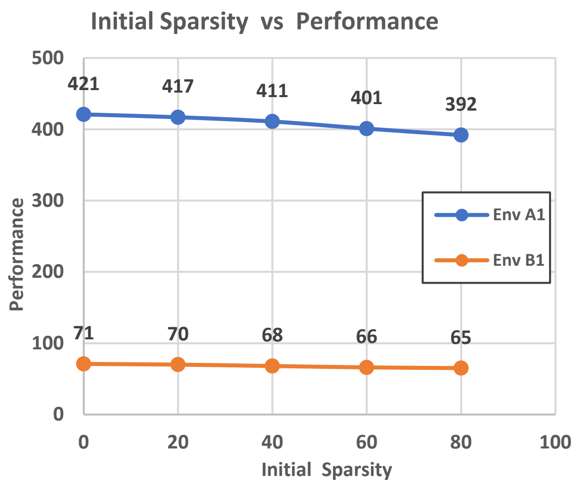

What's Next ?

CONCLUSION

FUTURE SCOPE

More Refined Mechanisms

Transfer Learning

Improved Environments

Conclusion & Future Scope

Takeaways

- ***Relational Neurogenesis*** is a combination of **evolutionary algorithms** and **deep reinforcement learning**
- It can **learn continually** with minimal **catastrophic forgetting**
- It minimizes and optimizes **network growth**
- It converges (episodically) much quicker than other algorithms
- **No supervisor** needed for task-switching
- But Relational Neurogenesis is ***computationally expensive***

Future Work

- Optimize and **unify** diverse neuroevolutionary mechanisms
- Reduce computational overheads
- Transfer learn between virtual and real-world scenarios
- Explore **extent** of lifelong learning supported by expandable networks

Thank you for attending the talk 😇