

CLAN: Exploring Continuous Learning on Commodity Edge Devices using Asynchronous Distributed Neuroevolution

Parth Mannan, Ananda Samajdar and Tushar Krishna

Georgia Institute of Technology

<http://synergy.ece.gatech.edu>

Email: pmannan@nvidia.com

August 24, 2020

The Future

Training AI at the Edge

A Revolution in Privacy Constraints for the World of AI



Edgify.ai

Follow

Dec 21, 2018 • 9 min read

Training Up to Train Developers

Deloitte
Insights



Jun Wu Contributor

COGNITIVE WORLD Contributor Group ©

AI

FIGURE 3

The edge AI c

Edge AI chips by

■ Smartphone

2020

2024

0 20

Sources: Marketsar
smart speakers, wea
end user industry, ai

Edge AI Is The Next Wave of AI

Why do you need to know about Edge AI? How do you get into the wave?



Jun Wu

Follow

Apr 19 • 4 min read ★

Training at the Network Edge with AI: 3 Key Benefits

By Yasser Khan | May 11, 2020

Tweet

Save



Using AI to train at the edge is critical for IoT implementations, especially as they scale. But the sheer horsepower of the cloud means the two should go hand in hand.

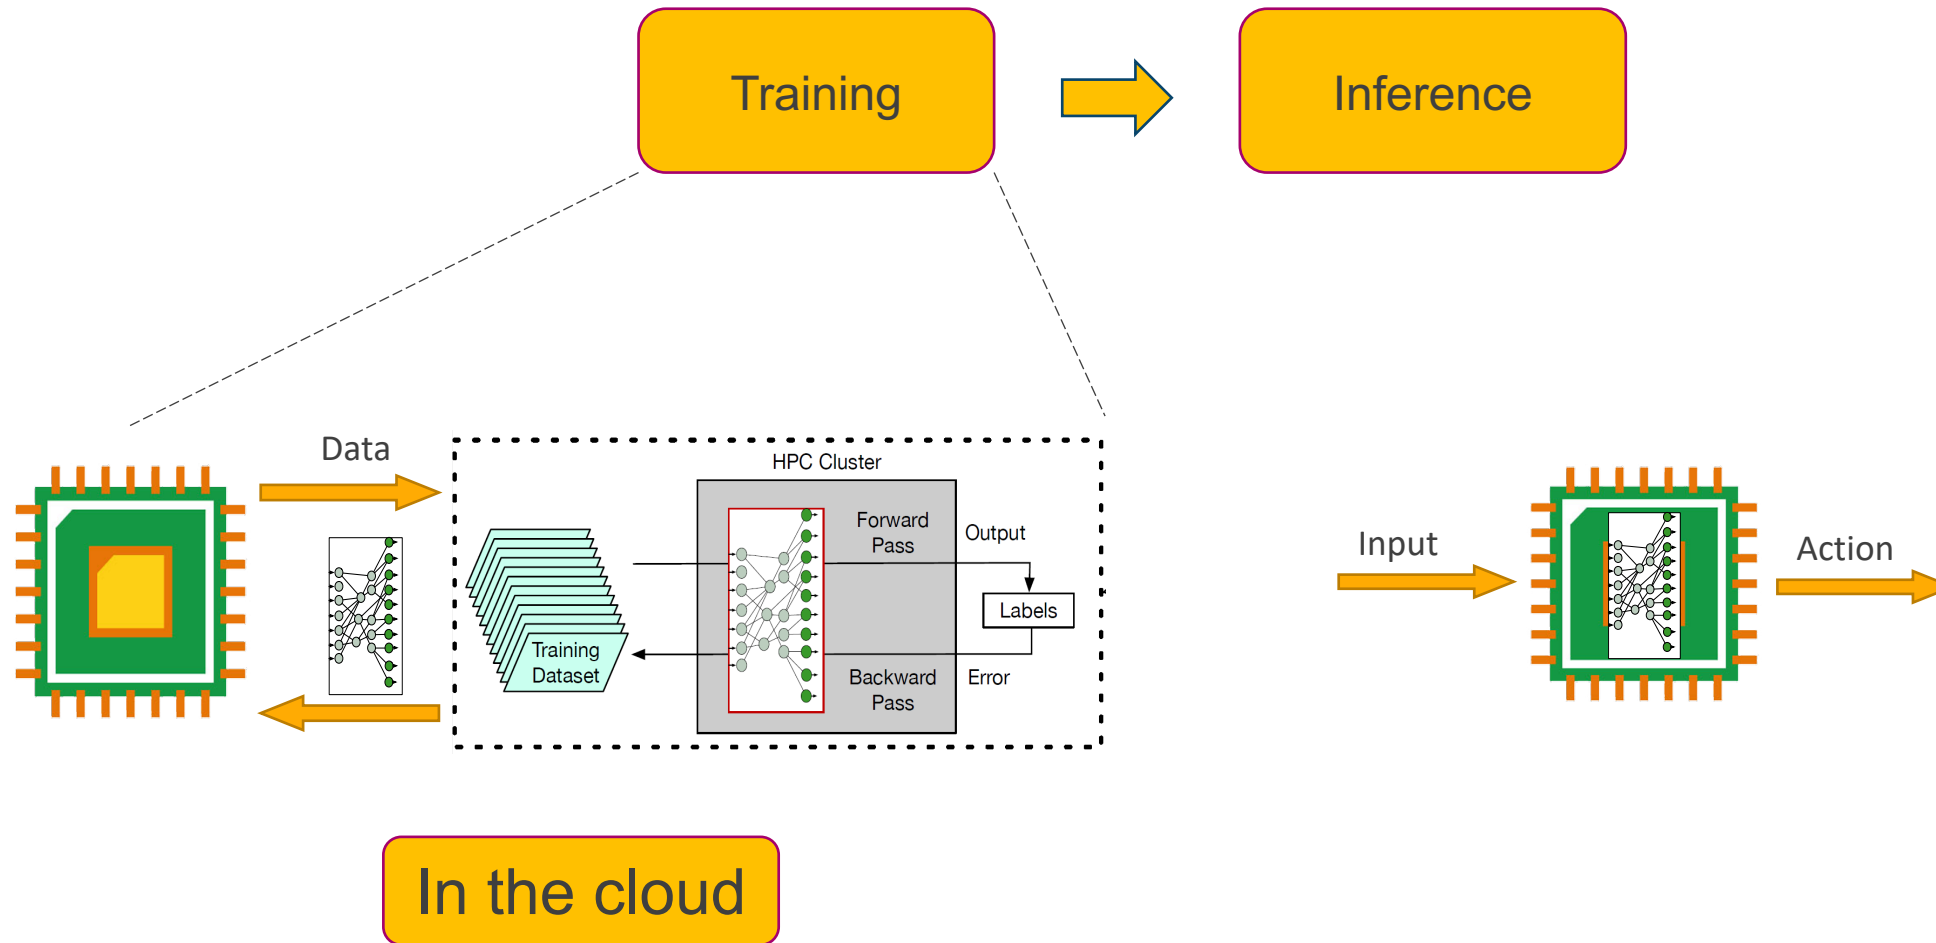
It's no surprise that the Internet of Things is growing rapidly. Each time an enterprise undertakes an IoT implementation, hundreds or perhaps thousands of new devices are added to the tally. And each time an implementation is successful, exponential growth is often seen as companies grasp the value



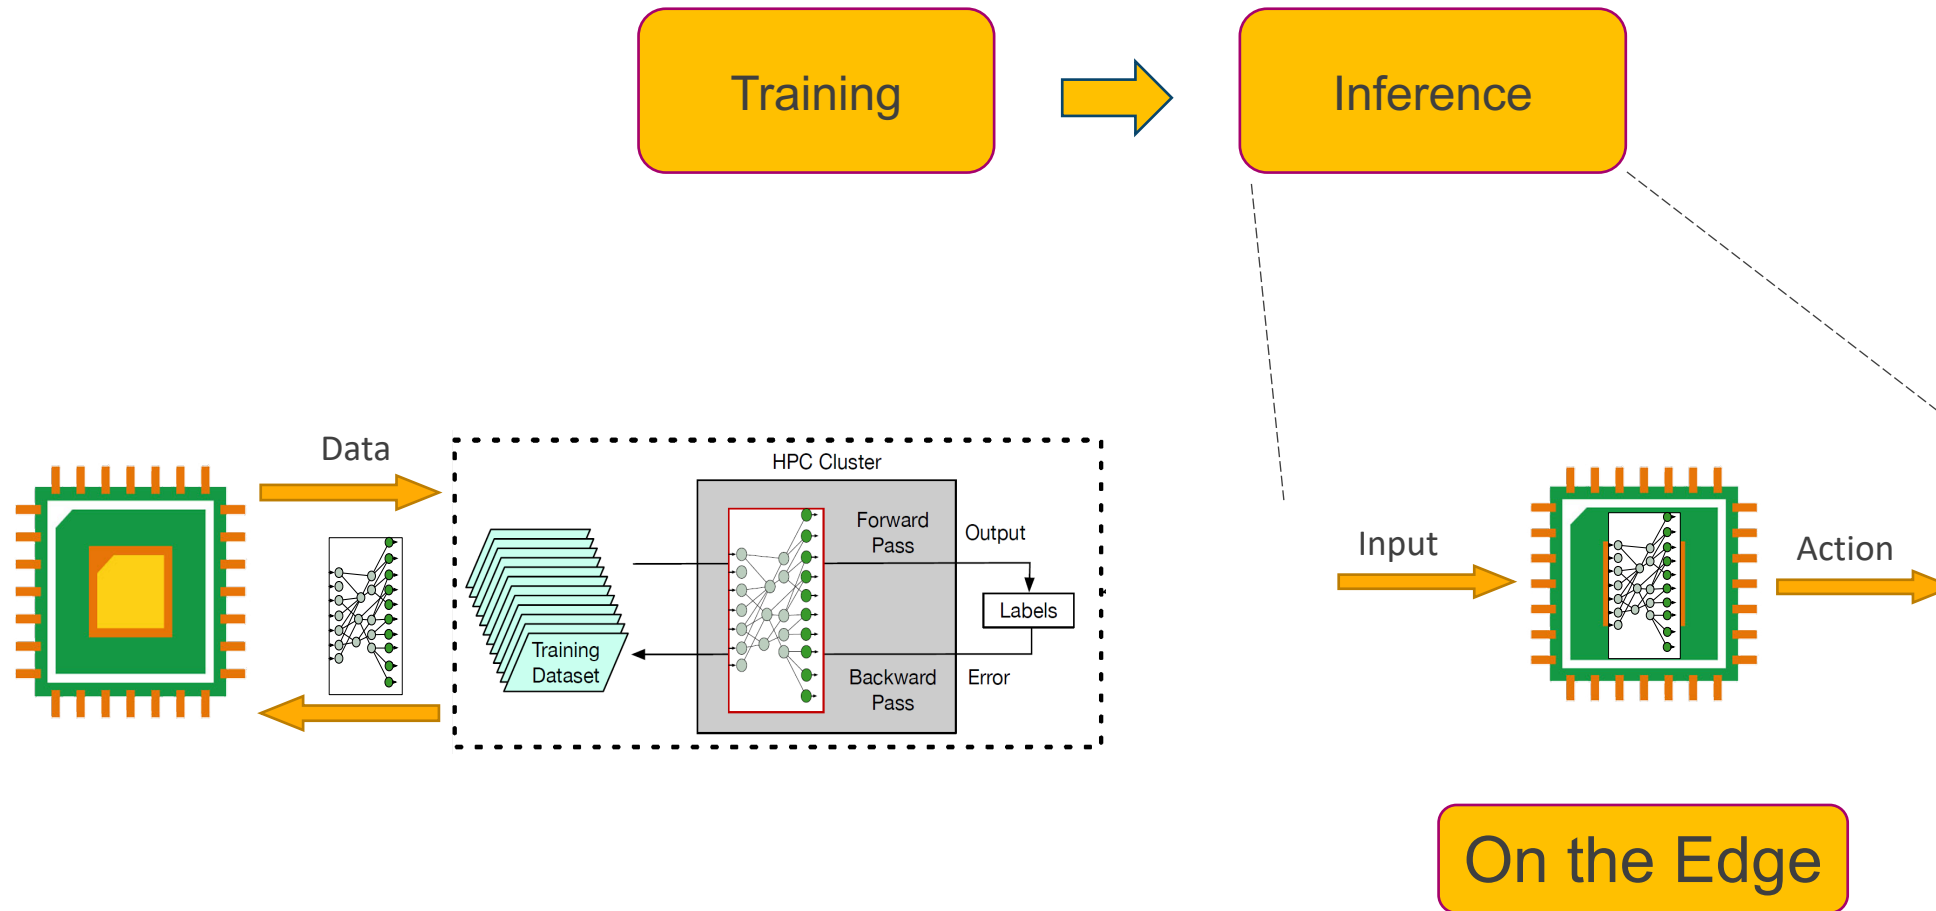
Outline

- A.I. at edge landscape
- Challenges with the current landscape
 - Learning at the Edge
- Neuro-Evolutionary algorithms
 - Introduction to NEAT
- Scaling NEAT at the Edge
 - Setup
 - Design choices and Evaluations
- Conclusion

The A.I. at edge landscape



The A.I. at edge landscape

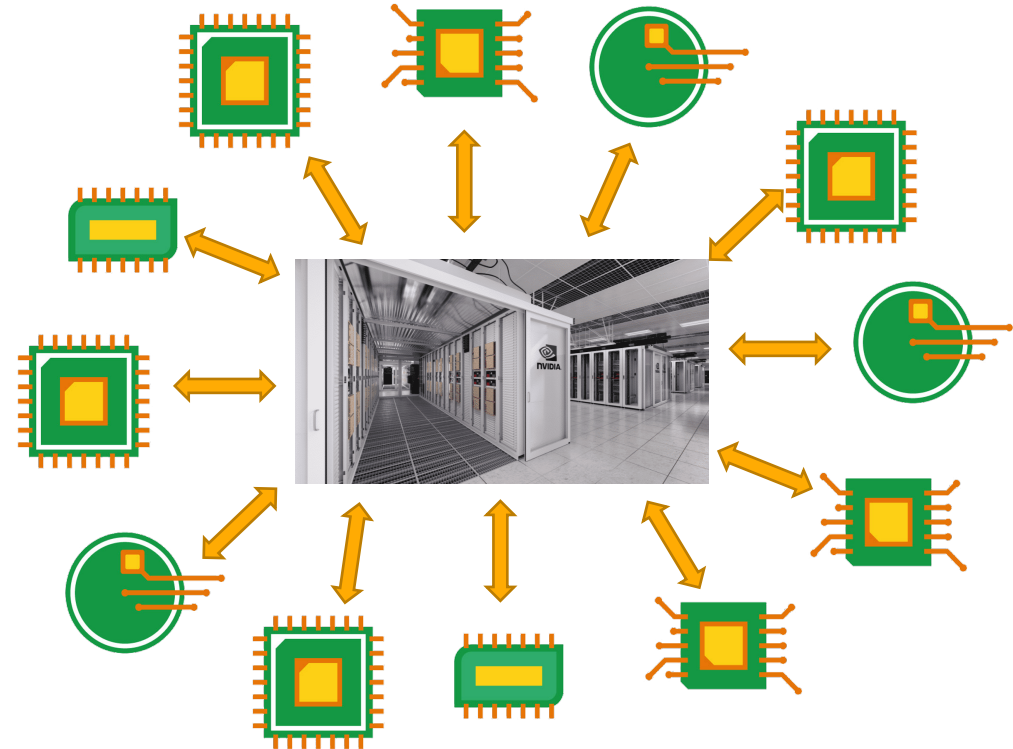


Outline

- A.I. at edge landscape
- Challenges with the current landscape
 - Learning at the Edge
- Neuro-Evolutionary algorithms
 - Introduction to NEAT
- Scaling NEAT at the Edge
 - Setup
 - Design choices and Evaluations
- Conclusion

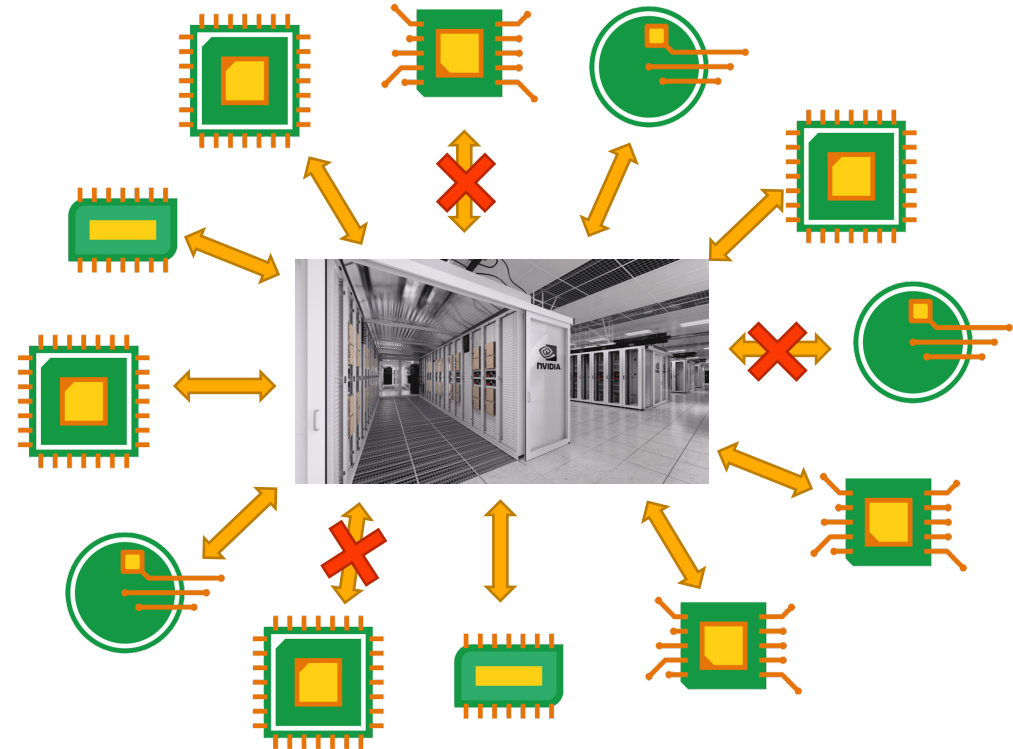
Challenges with current landscape

- Limited uplink and downlink bandwidth to the cloud



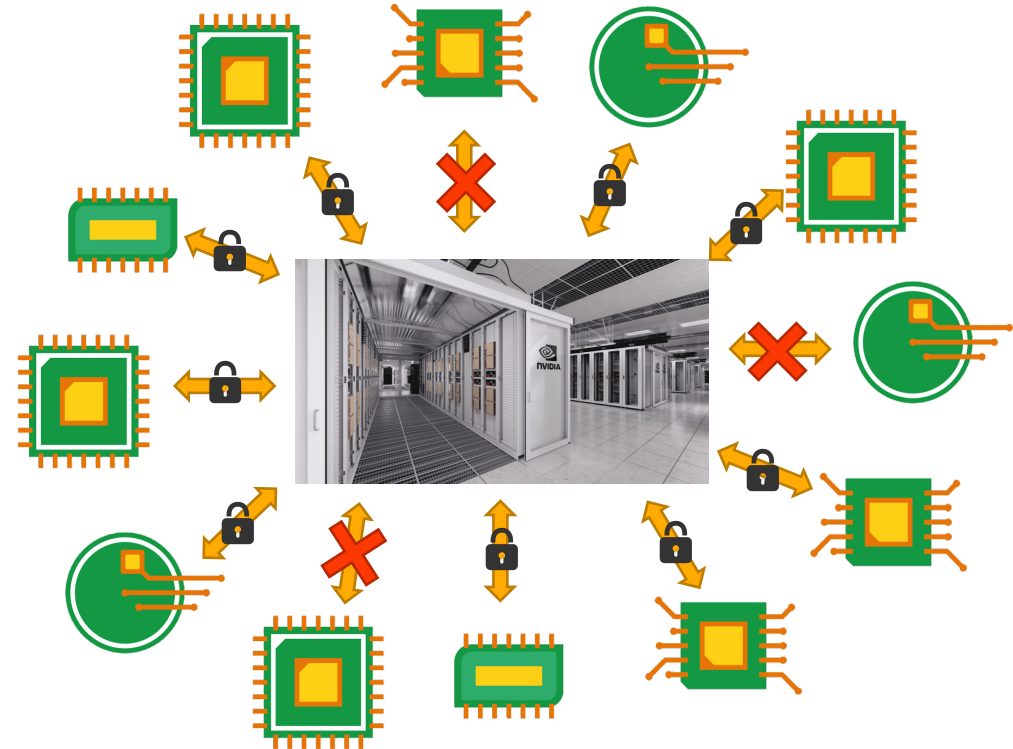
Challenges with current landscape

- Limited uplink and downlink bandwidth to the cloud
- Constant connectivity to the cloud



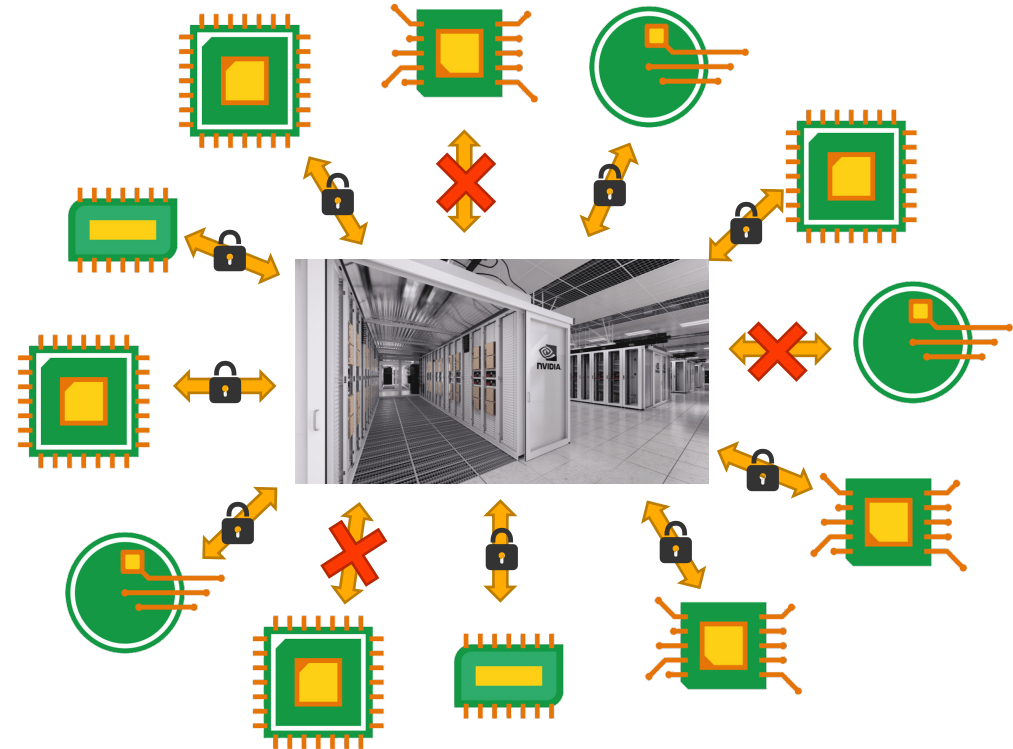
Challenges with current landscape

- Limited uplink and downlink bandwidth to the cloud
- Constant connectivity to the cloud
- Privacy and Security



Challenges with current landscape

- Limited uplink and downlink bandwidth to the cloud
- Constant connectivity to the cloud
- Privacy and Security



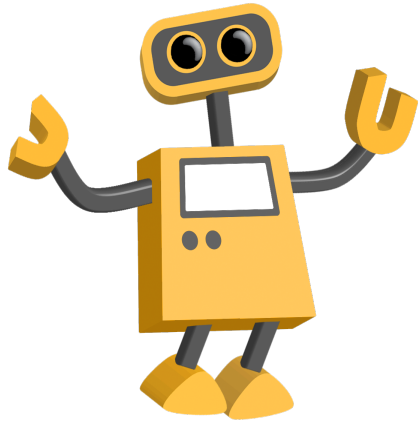
Solution: Train at the edge

Outline

- A.I. at edge landscape
- Challenges with the current landscape
 - Learning at the Edge
- Neuro-Evolutionary algorithms
 - Introduction to NEAT
- Scaling NEAT at the Edge
 - Setup
 - Design choices and Evaluations
- Conclusion

Learning at the Edge

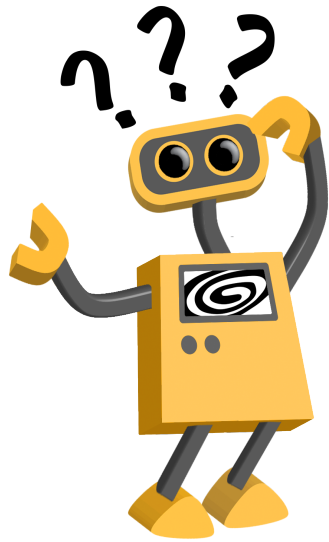
- The environment at the edge can change rapidly



Training

Learning at the Edge

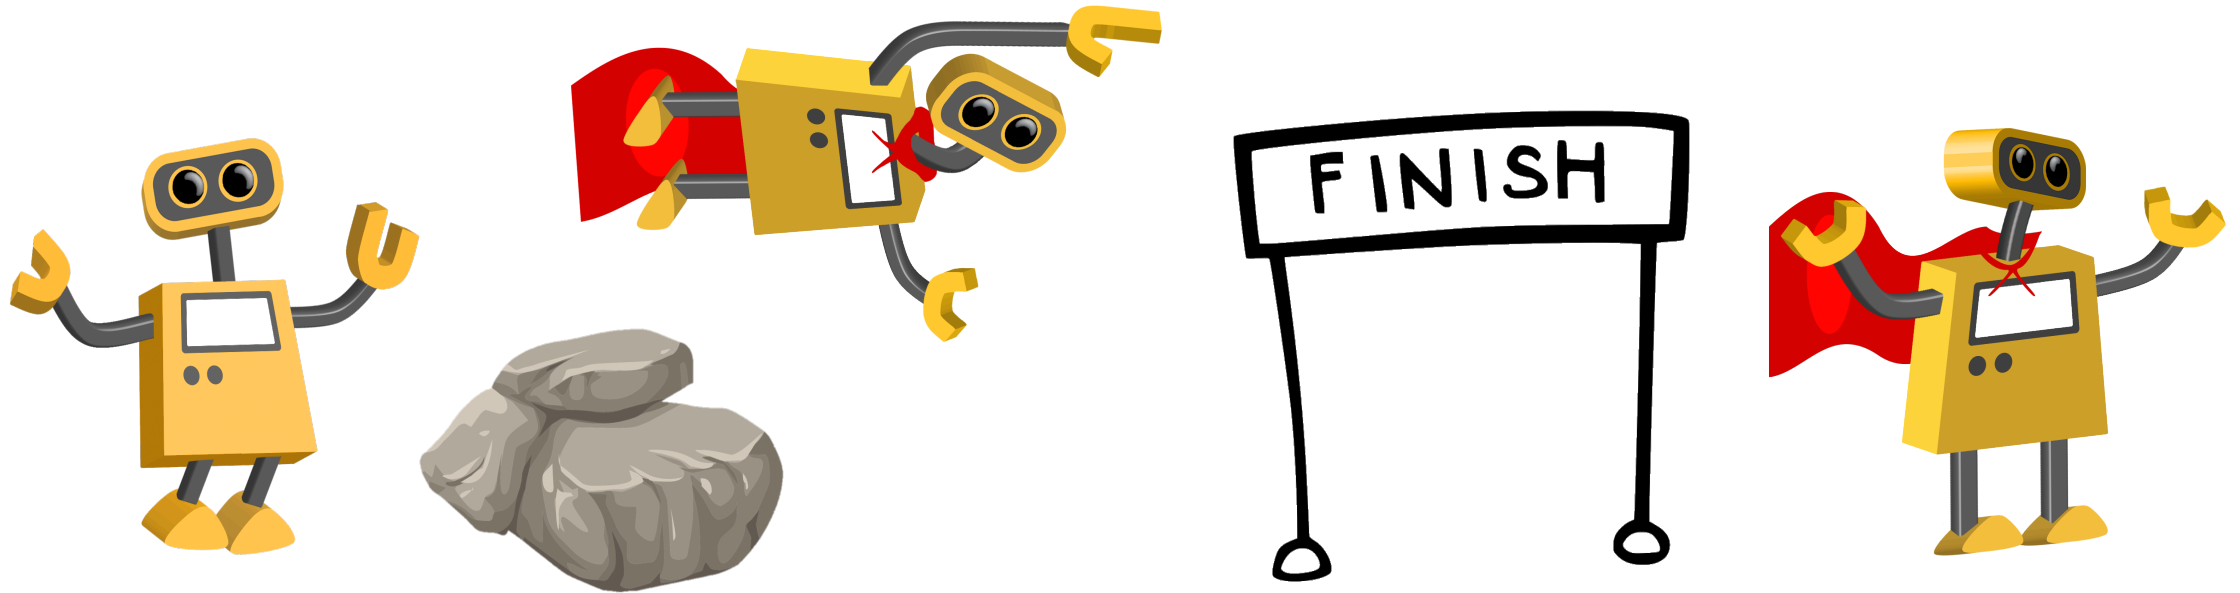
- The environment at the edge can change rapidly



Real Environment

Learning at the Edge

- The environment at the edge can change rapidly



Learn new skill – Continuous learning

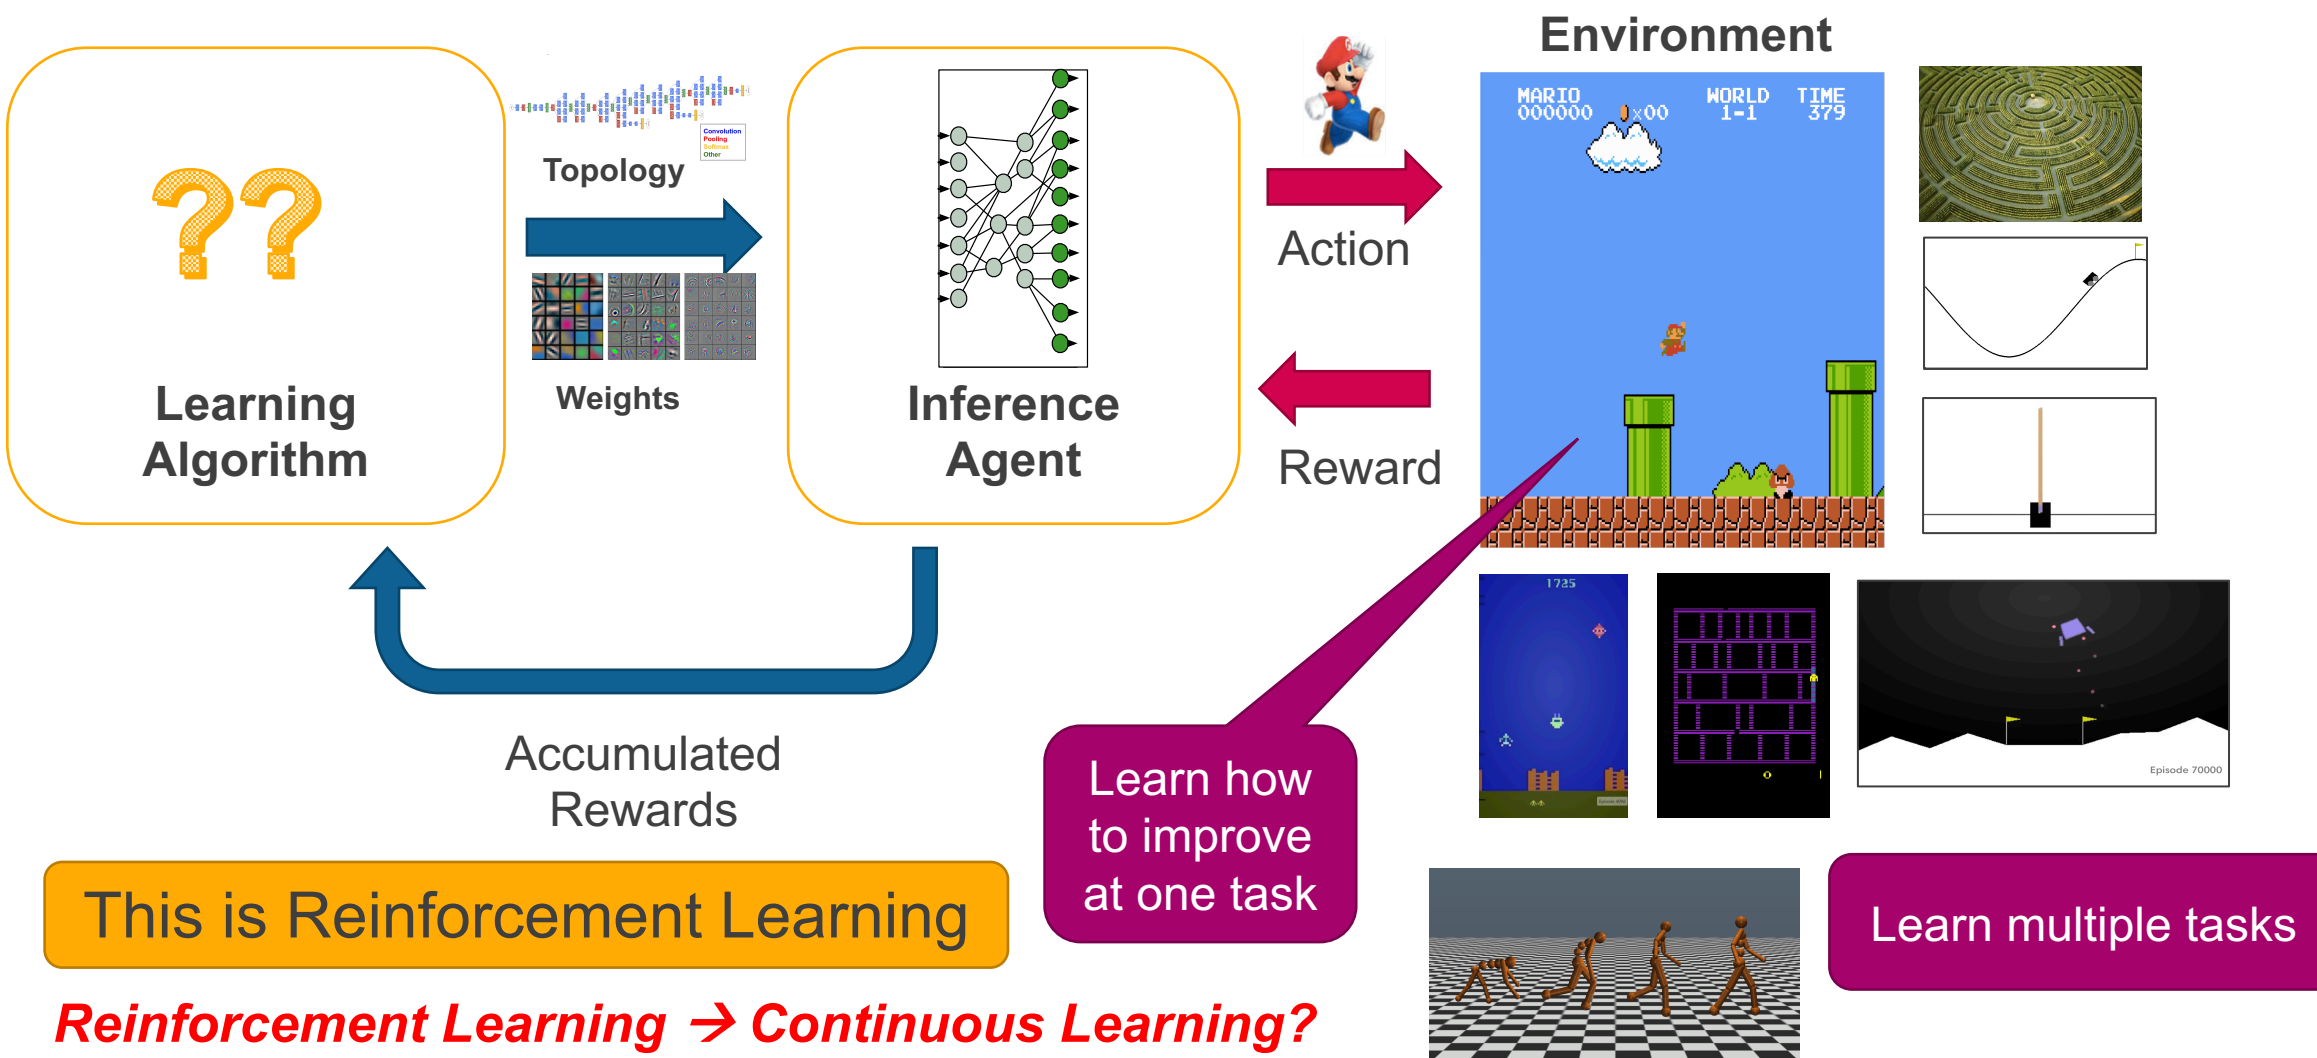
Learning at the edge

- The environment at the edge can change rapidly
 - Needs continuous learning
- Limited Compute and memory at the edge
 - Needs to fit in the hardware budget at the edge
- No mechanism to label collected data
 - Needs ability to process unlabeled and unstructured data

Algorithms for Learning at the edge

	Data	Hyper-param Tuning	DNN Plasticity	Compute	Memory
Supervised DL	Labeled X	Manual X	Designed for one problem X	Backprop (gradients) X	Backprop (gradients) X

Template for Continuous Learning



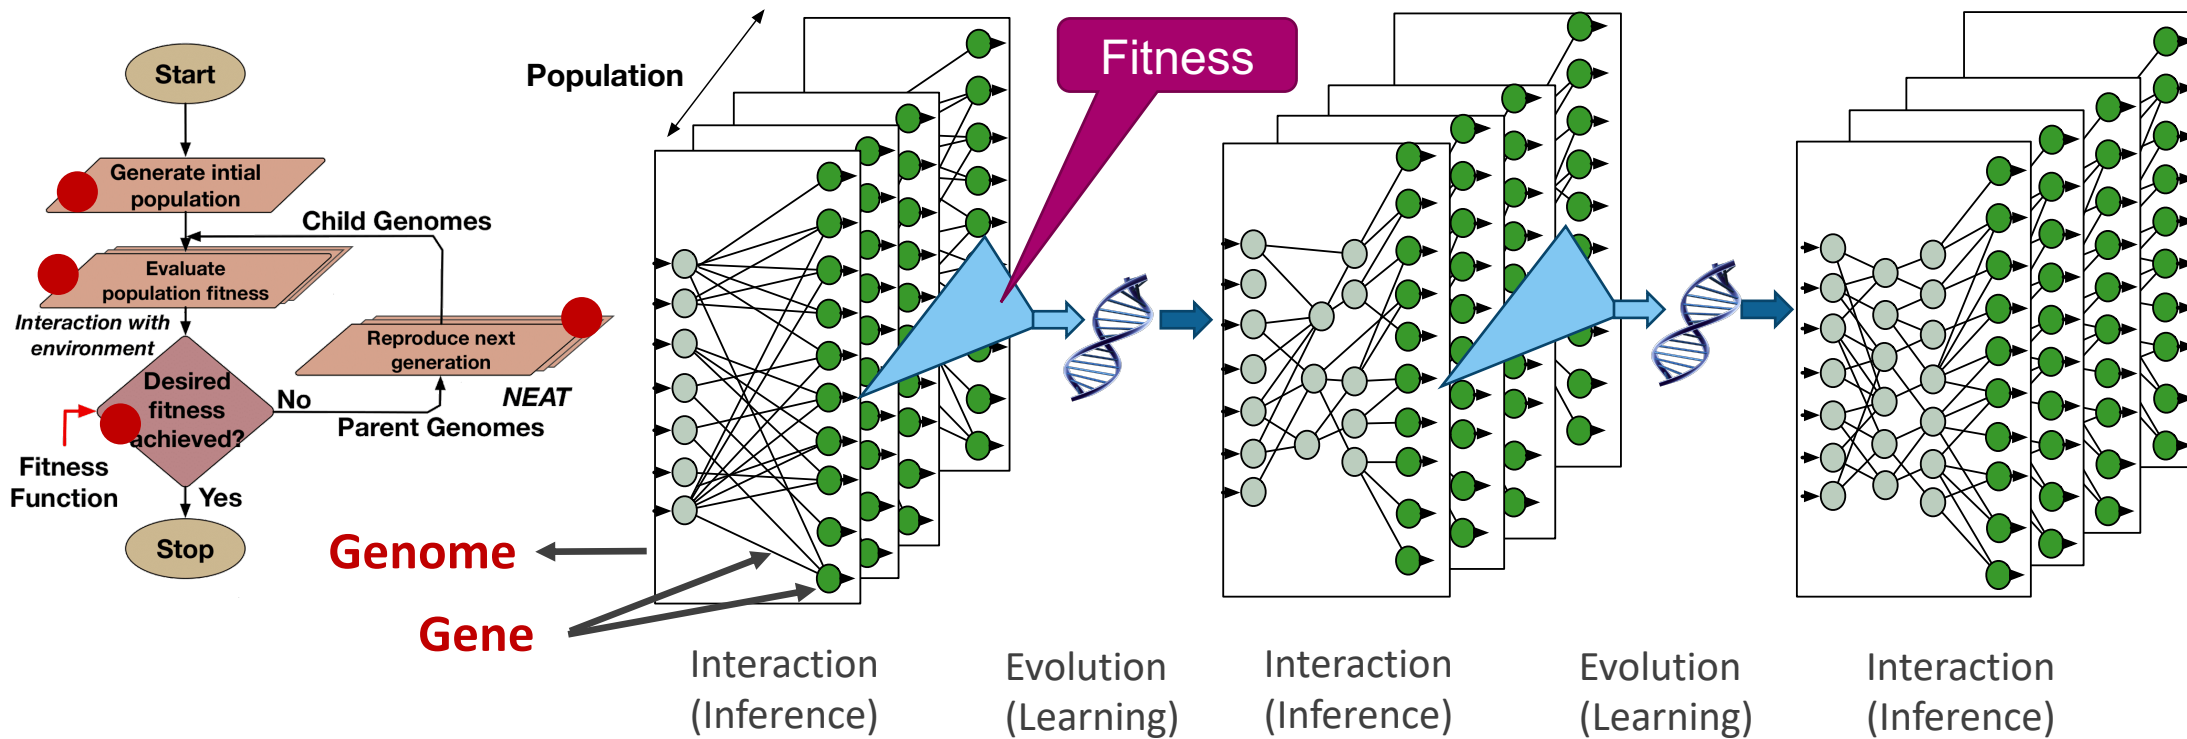
Candidates for Continuous Learning

	Data	Hyper-param Tuning	DNN Plasticity	Compute	Memory
Supervised DL	Labeled X	Manual X	Designed for one problem X	Backprop (gradients) X	Backprop (gradients) X
Reinforcement Learning	Unlabeled ✓	Manual X	Reward Function ✓	Backprop (gradients) X	Backprop (gradients) X
	RL is not viable for continuous learning on the edge				

Outline

- A.I. at edge landscape
- Challenges with the current landscape
 - Learning at the Edge
- Neuro-Evolutionary algorithms
 - Introduction to NEAT
- Scaling NEAT at the Edge
 - Setup
 - Design choices and Evaluations
- Conclusion

Neuro-Evolutionary (NE) Algorithm



Neural Network (NN) expressed as a graph

Gene: Vertex or Edge in the graph

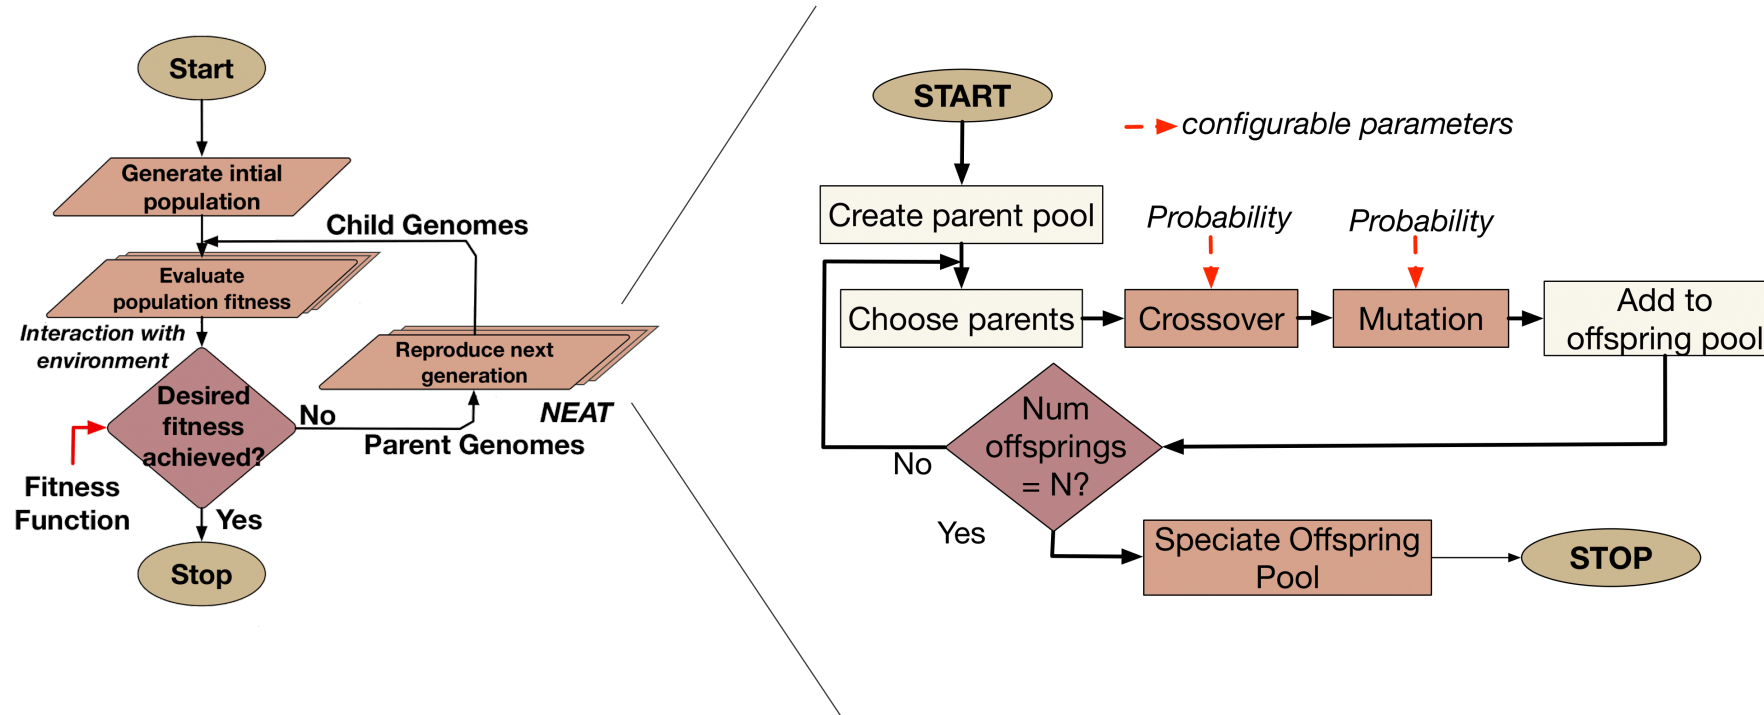
Genome: Collection of all genes (i.e., a NN)

NeuroEvolution of Augmented Topologies (NEAT) [1]

[1] Stanley, K. O., & Miikkulainen, R. (2002). Evolving neural networks through augmenting topologies. *Evolutionary computation*, 10(2), 99-127.

This slide adapted from – Samajdar et. al. GeneSys: Enabling Continuous Learning through Neural Network Evolution in Hardware.

Neuro-Evolutionary (NE) Algorithm



NeuroEvolution of Augmented Topologies (NEAT) [1]

[1] Stanley, K. O., & Miikkulainen, R. (2002). Evolving neural networks through augmenting topologies. *Evolutionary computation*, 10(2), 99-127.

Neuro-evolutionary Algo → Continuous Learning?

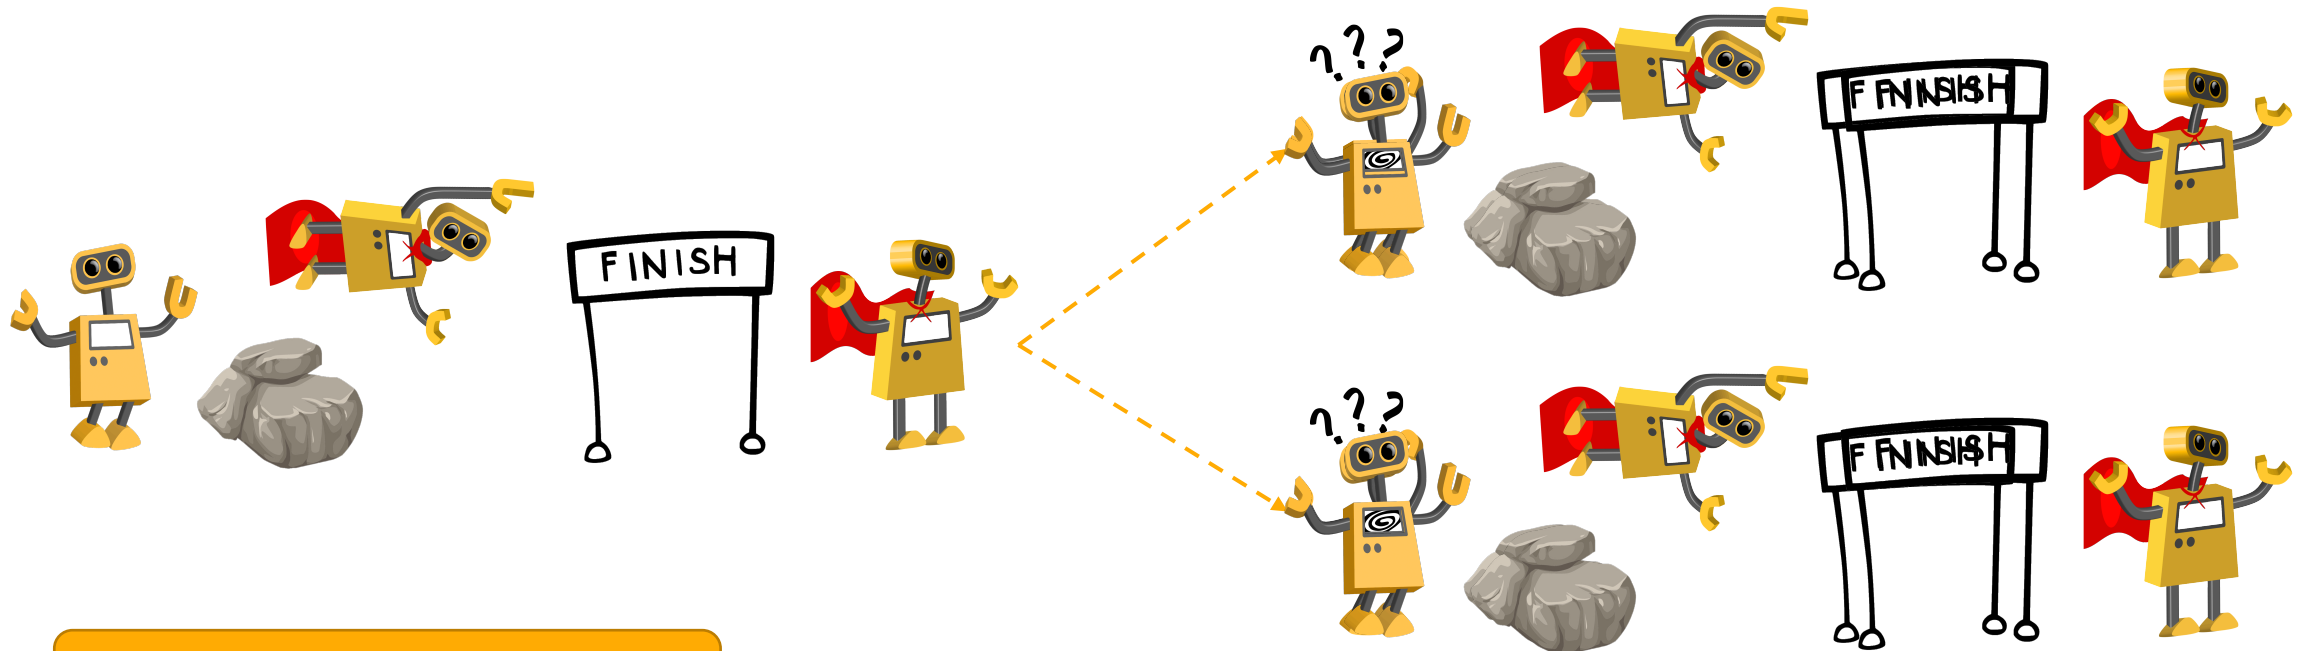
This slide adapted from – Samajdar et. al. GeneSys: Enabling Continuous Learning through Neural Network Evolution in Hardware.

Learning Algorithms at Edge

	Data	Hyper-param Tuning	DNN Plasticity	Compute	Memory
Supervised DL	Labeled X	Manual X	Designed for one problem X	Backprop (gradients) X	Backprop (gradients) X
Reinforceme nt Learning	Unlabeled ✓	Manual X	Reward Function ✓	Backprop (gradients) X	Backprop (gradients) X
Evolutionary	Unlabeled ✓	Automated ✓	Reward Function ✓	Massive Parallelism ✓	Only store model ✓

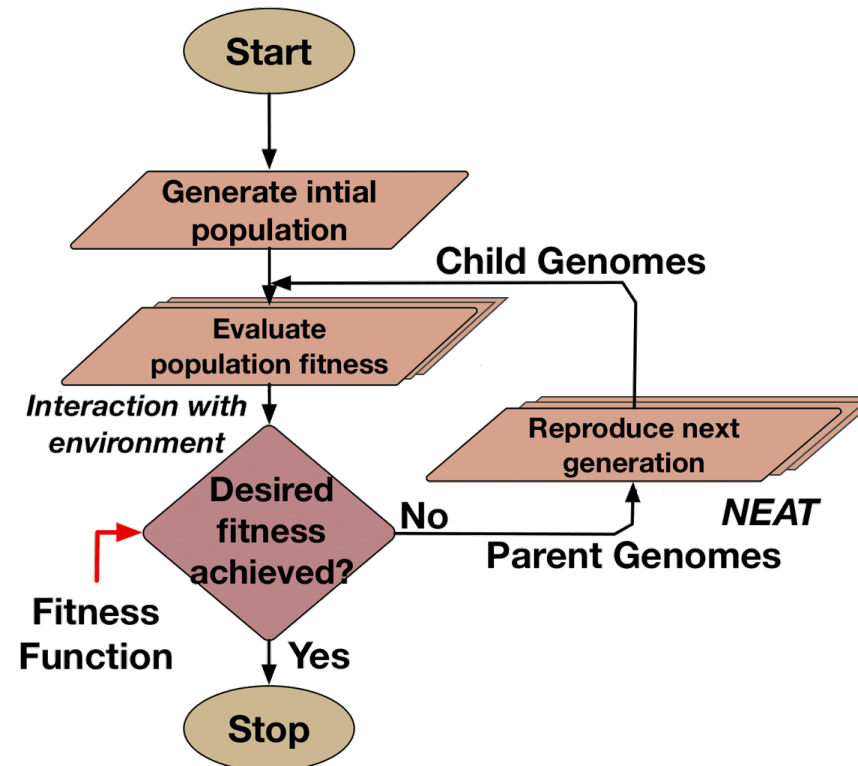
Learning at the Edge

- Is continuous learning enough?
- Edge devices often deployed in groups
- Can knowledge gained by one help the entire system?



Distributed learning

Neuro-Evolutionary (NE) Algorithm



Could we scale NEAT at the edge?

NeuroEvolution of Augmented Topologies (NEAT) [1]

[1] Stanley, K. O., & Miikkulainen, R. (2002). Evolving neural networks through augmenting topologies. *Evolutionary computation*, 10(2), 99-127.

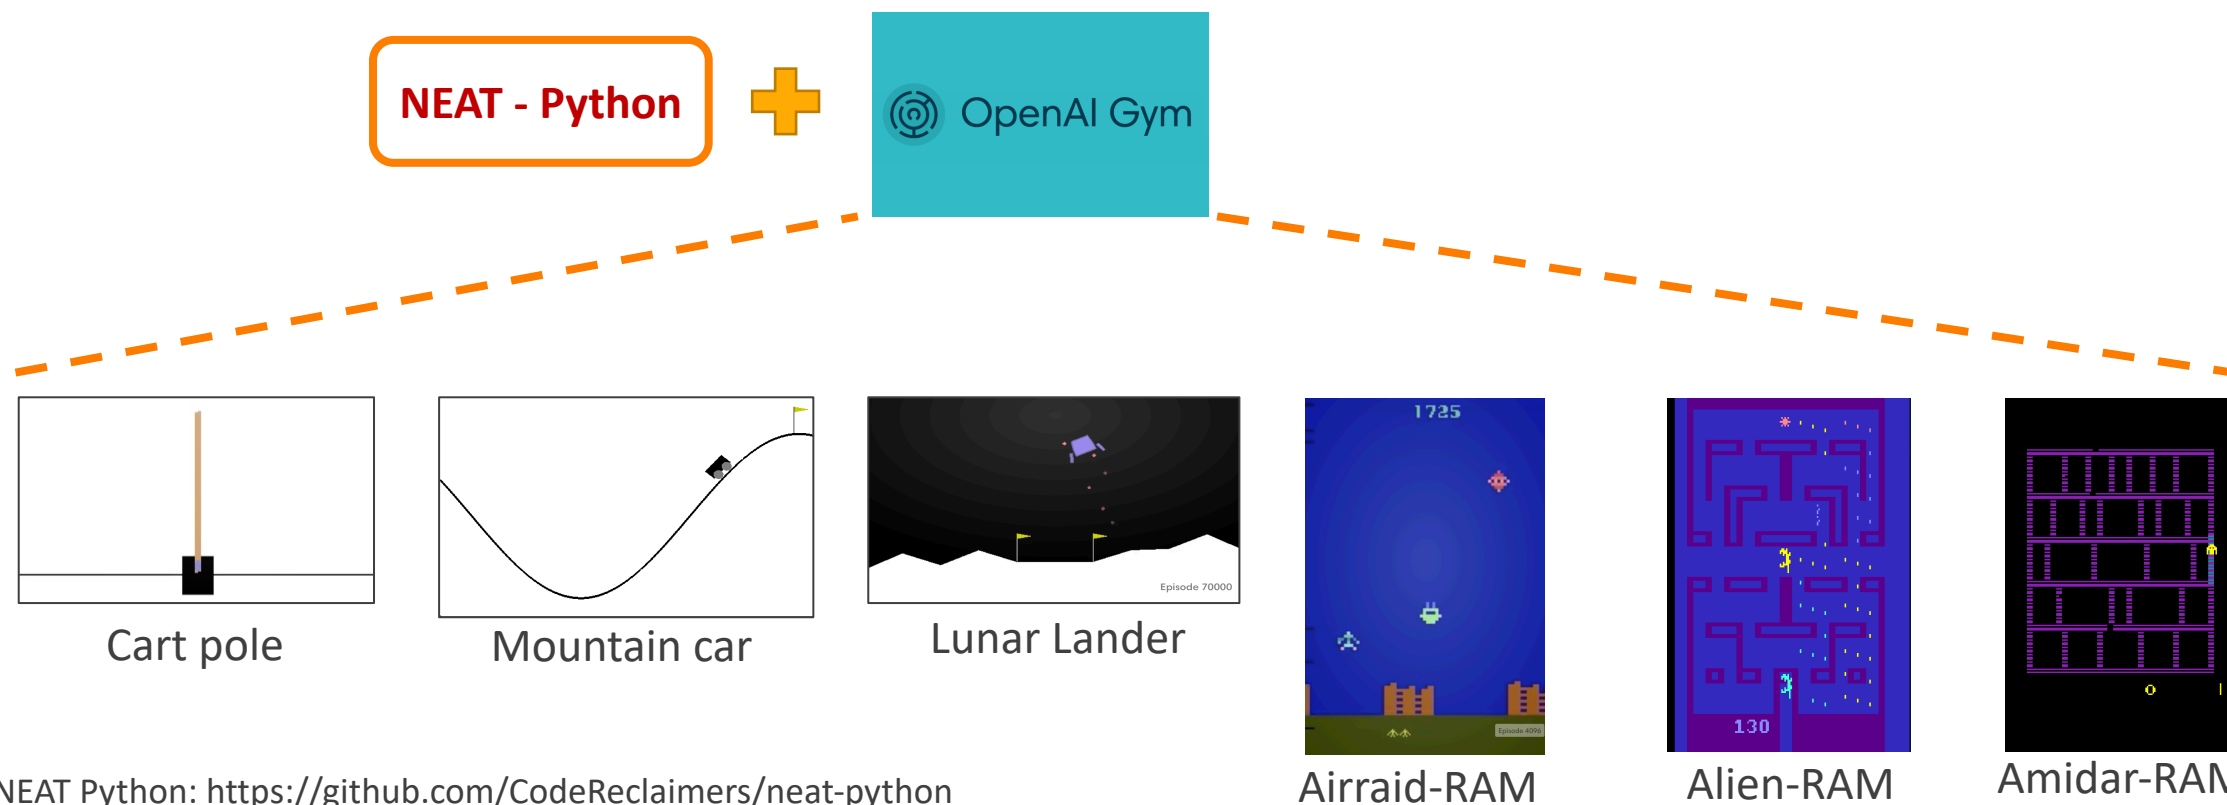
This slide adapted from – Samajdar et. al. GeneSys: Enabling Continuous Learning through Neural Network Evolution in Hardware.

Outline

- A.I. at edge landscape
- Challenges with the current landscape
 - Learning at the Edge
- Neuro-Evolutionary algorithms
 - Introduction to NEAT
- Scaling NEAT at the Edge
 - Setup
 - Design choices and Evaluations
- Conclusion

Scaling NEAT at the Edge

- Test bed with 15 Raspberry Pi
- Connected with 62.24Mbps local WiFi network



NEAT Python: <https://github.com/CodeReclaimers/neat-python>

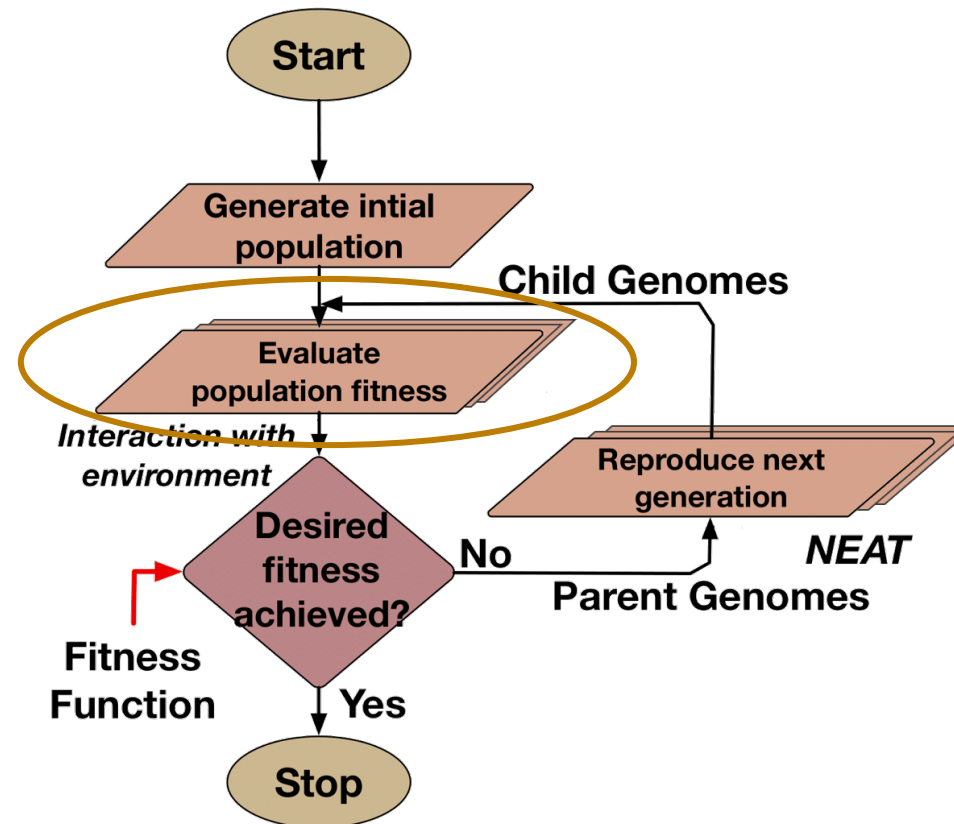
Scaling NEAT at the Edge

- Test bed with 15 Raspberry Pi
- Connected with 62.24Mbps local WiFi network
- Only fitness function is changed between workloads
- All workloads are run till convergence (or failure) multiple times
 - Multi-step inference – Rewards accumulated over multiple time steps between each generation
 - Single-step inference – Rewards from each action leads to a new generation

Outline

- A.I. at edge landscape
- Challenges with the current landscape
 - Learning at the Edge
- Neuro-Evolutionary algorithms
 - Introduction to NEAT
- Scaling NEAT at the Edge
 - Setup
 - Design choices and Evaluations
- Conclusion

Neuro-Evolutionary (NE) Algorithm

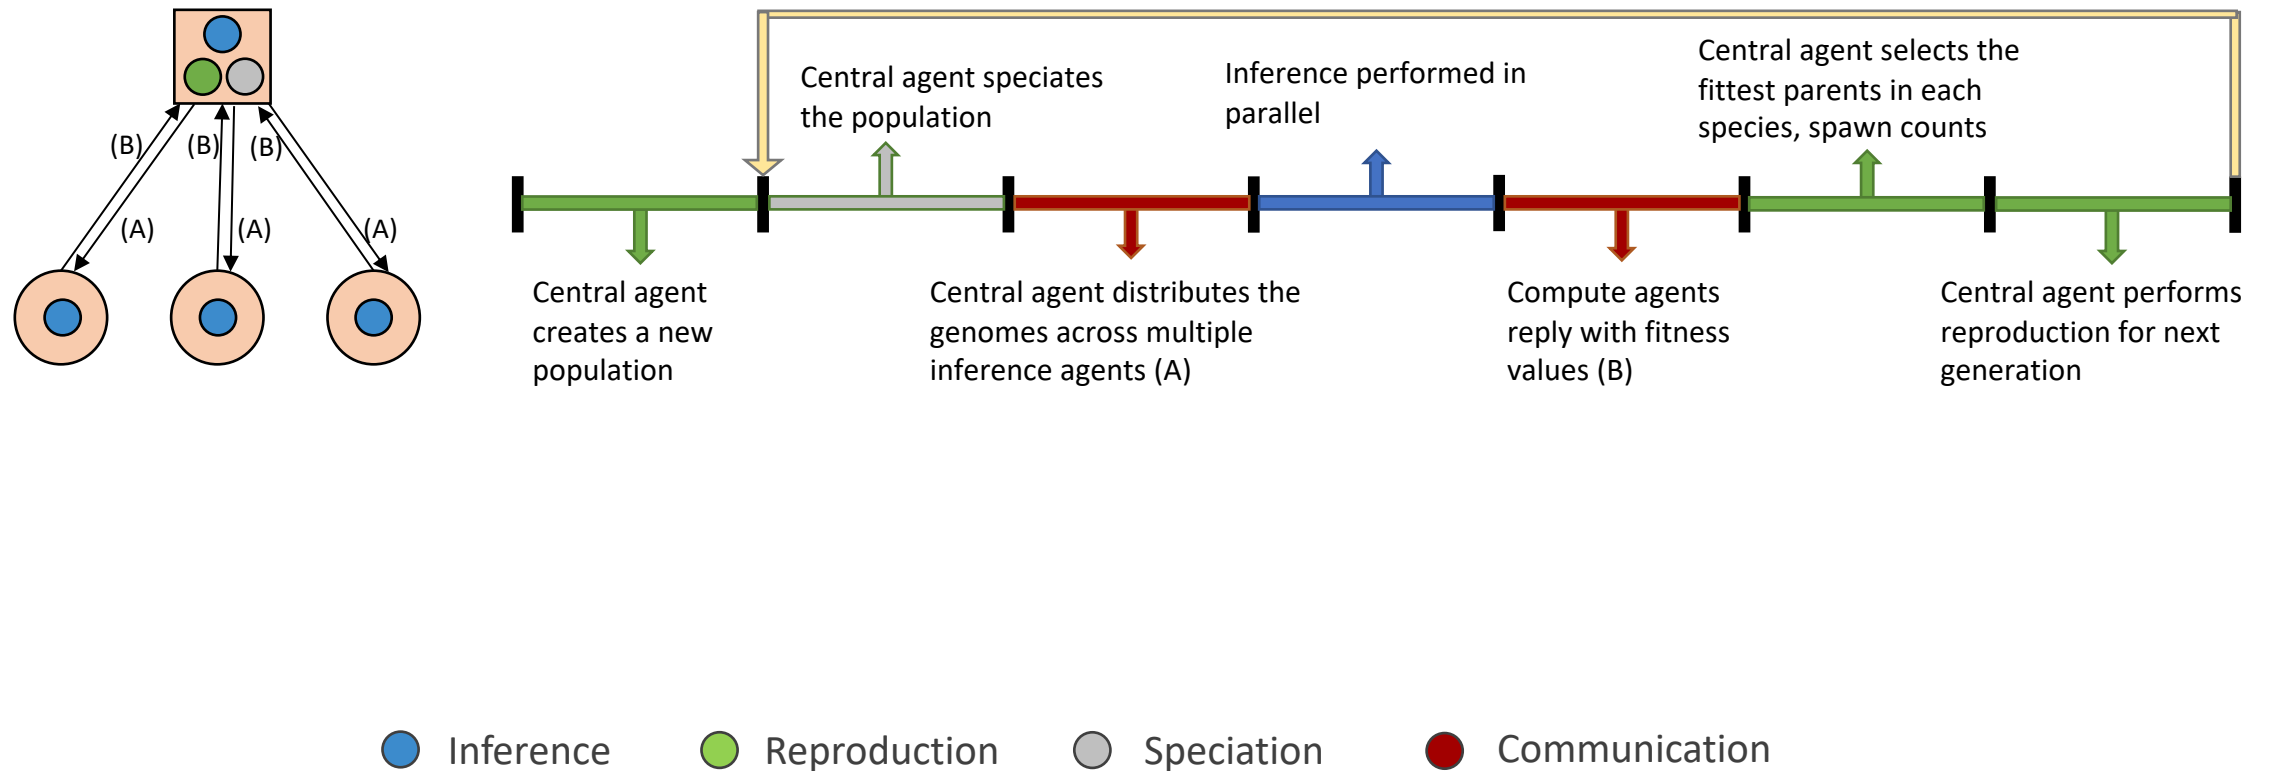


NeuroEvolution of Augmented Topologies (NEAT) [1]

[1] Stanley, K. O., & Miikkulainen, R. (2002). Evolving neural networks through augmenting topologies. *Evolutionary computation*, 10(2), 99-127.

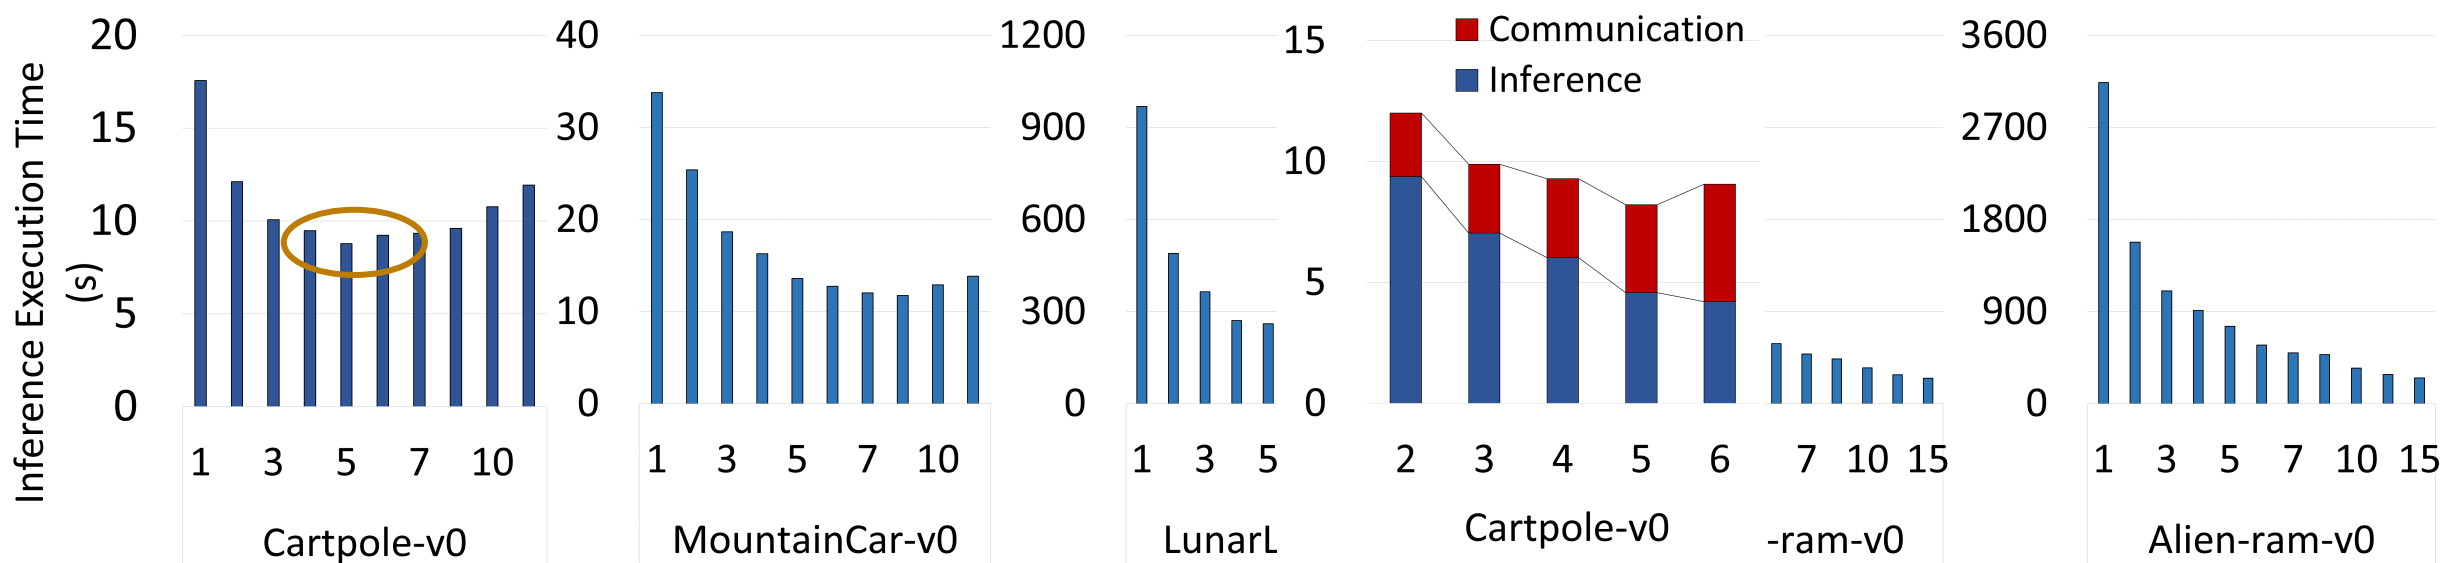
Scaling NEAT at the Edge

- Scaling Inference – Fitness Calculation

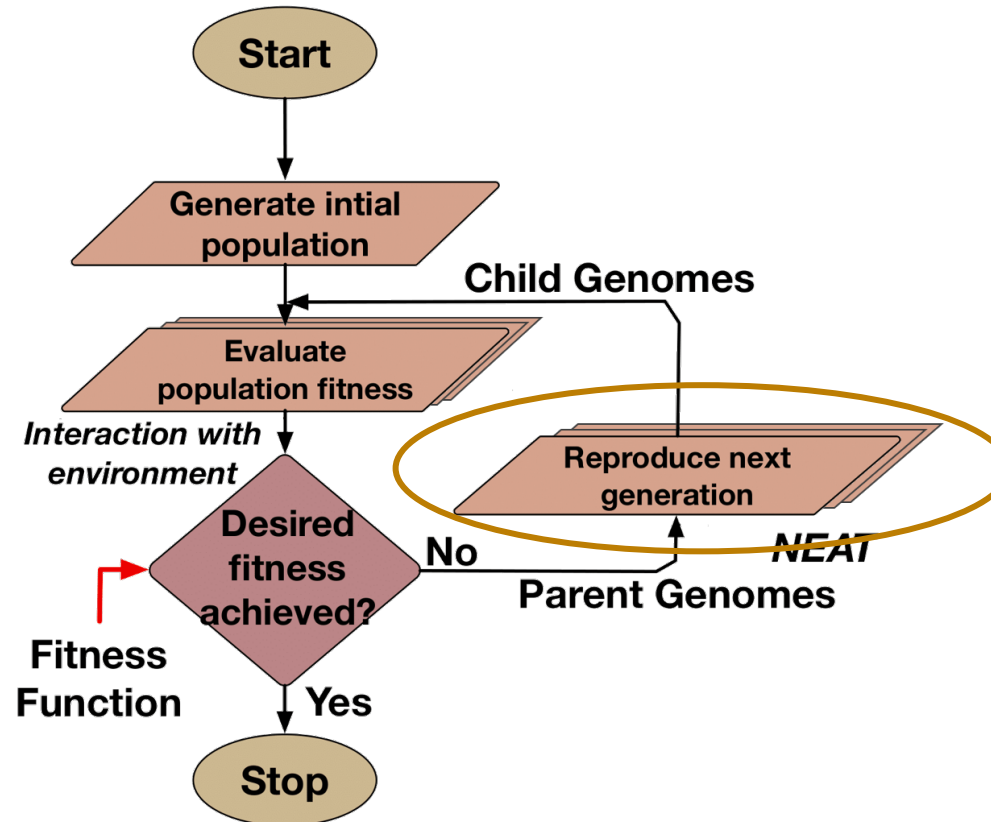


Scaling NEAT at the Edge

- Scaling Inference – Fitness Calculation



Neuro-Evolutionary (NE) Algorithm

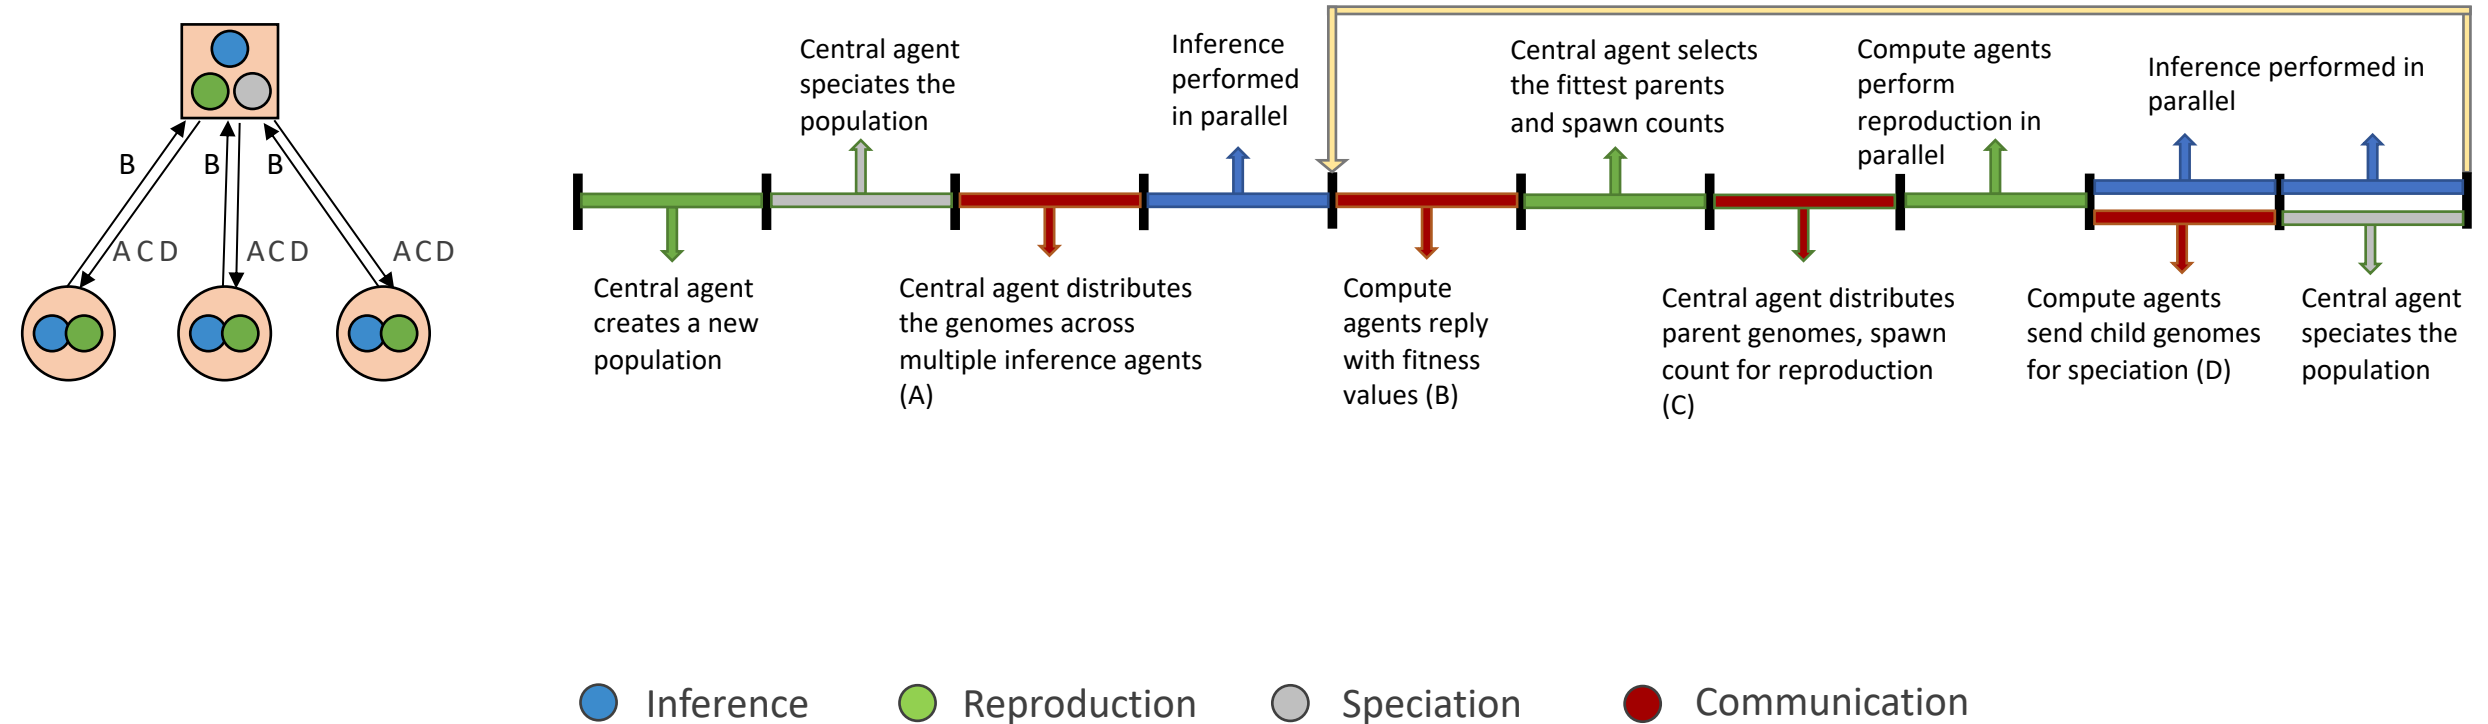


NeuroEvolution of Augmented Topologies (NEAT) [1]

[1] Stanley, K. O., & Miikkulainen, R. (2002). Evolving neural networks through augmenting topologies. *Evolutionary computation*, 10(2), 99-127.

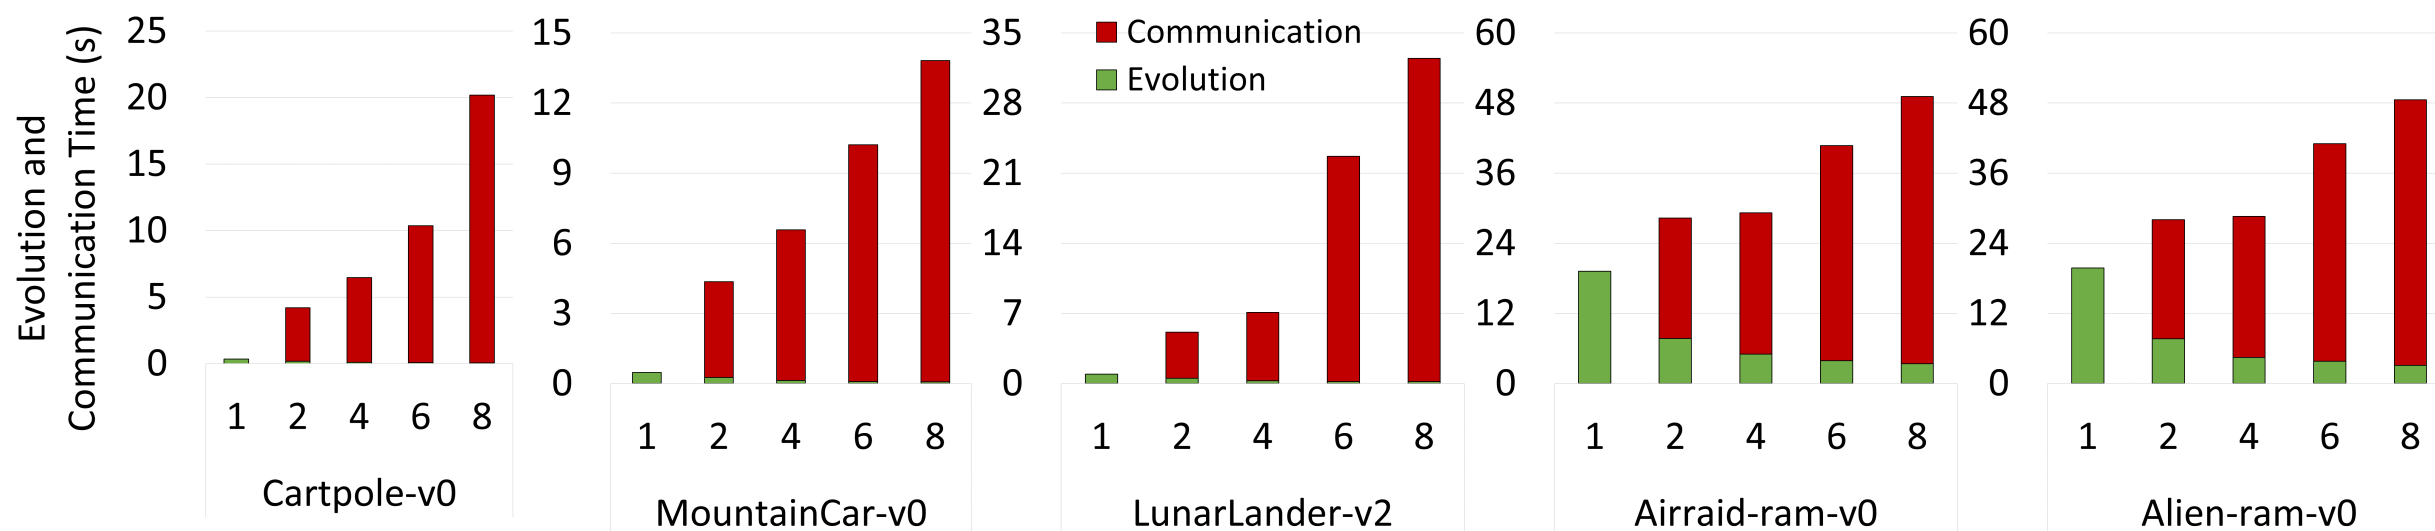
Scaling NEAT at the Edge

- Scaling Learning – Reproduction

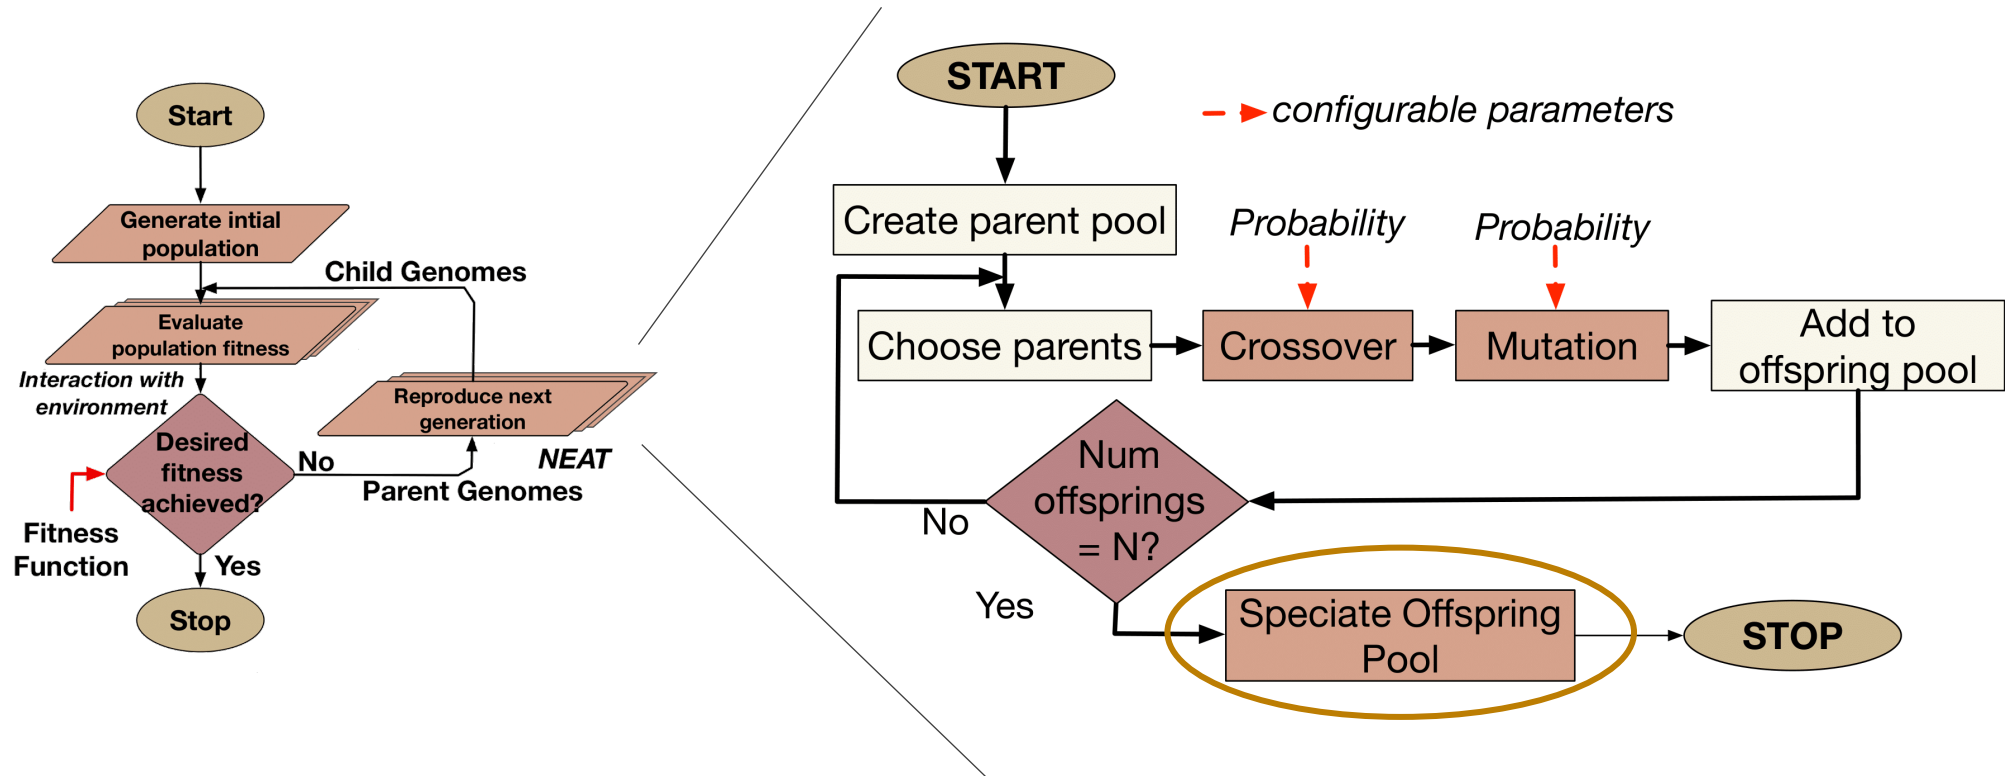


Scaling NEAT at the Edge

- Scaling Learning – Reproduction



Neuro-Evolutionary (NE) Algorithm

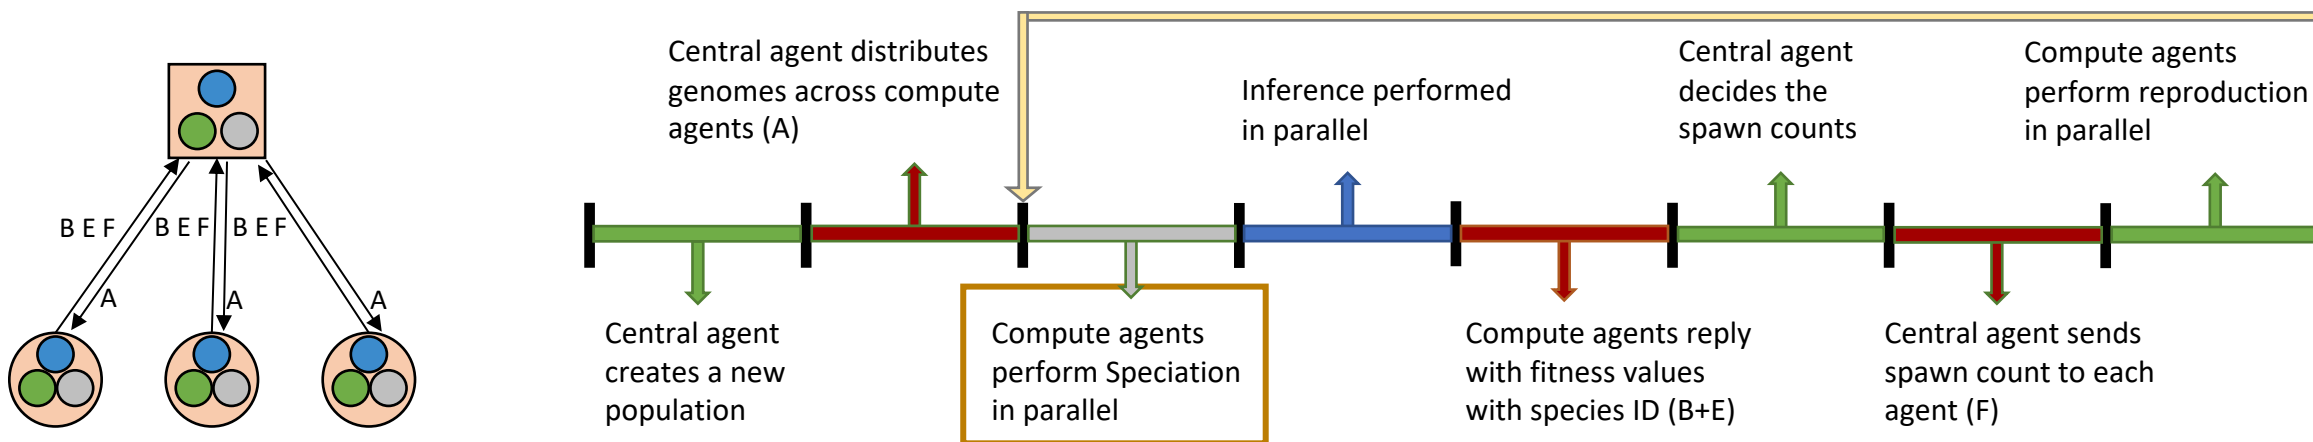


NeuroEvolution of Augmented Topologies (NEAT) [1]

[1] Stanley, K. O., & Miikkulainen, R. (2002). Evolving neural networks through augmenting topologies. *Evolutionary computation*, 10(2), 99-127.

Scaling NEAT at the Edge

- Scaling Selection – Speciation and selection

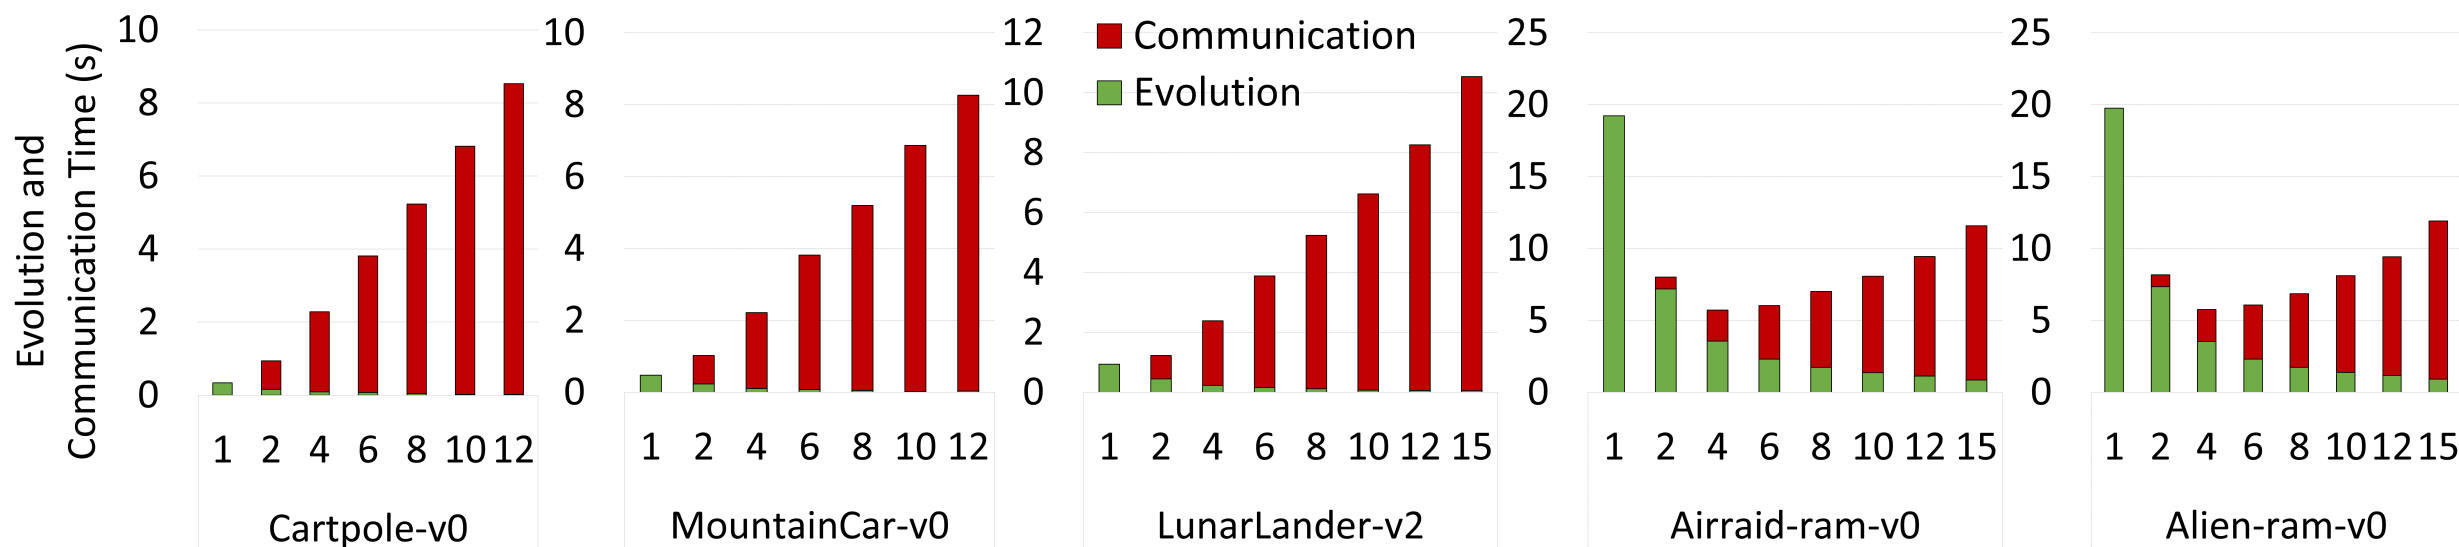


Asynchronous NeuroEvolution

● Inference
 ● Reproduction
 ● Speciation
 ● Communication

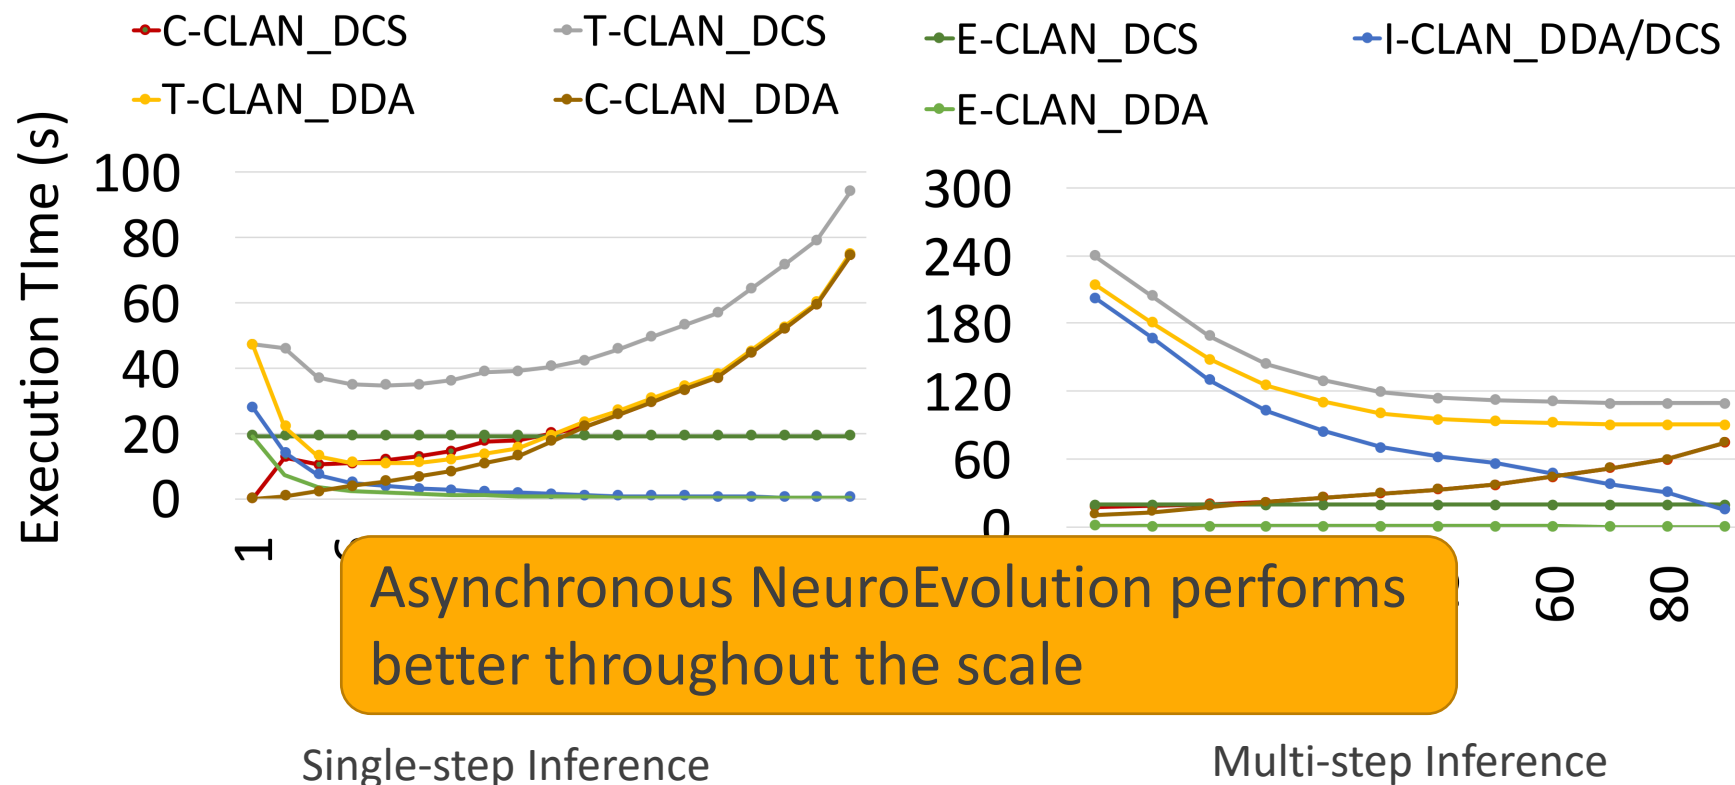
Scaling NEAT at the Edge

- Scaling Selection – Speciation and selection



Scaling NEAT at the Edge

- Extrapolating Scalability



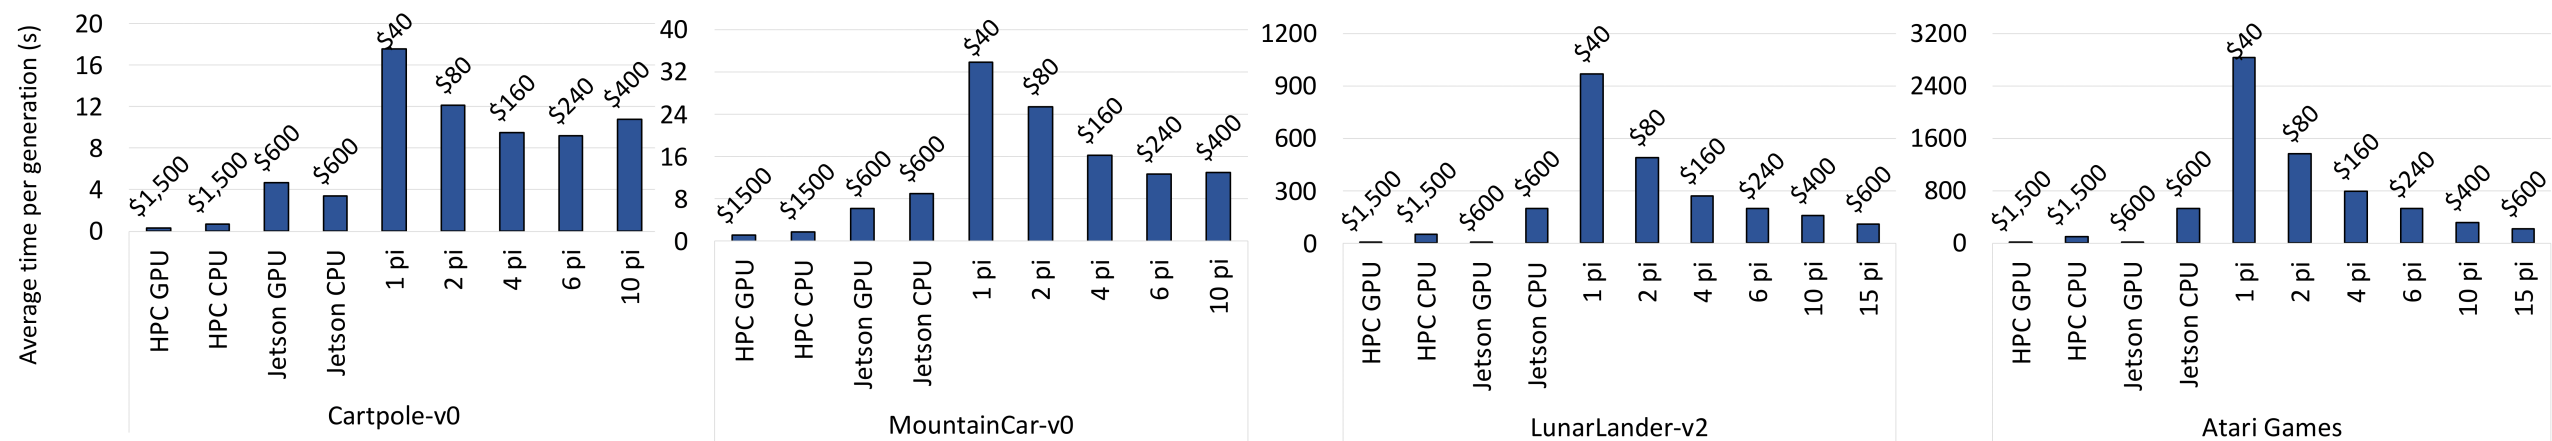
CLAN_DCS – Distributed Inference

CLAN_DDA – Distributed Asynchronous Speciation

Scaling NEAT at the Edge

- Performance per dollar

Platform	Processor	Cost
HPC CPU	6th gen i7	\$1500
HPC GPU	Nvidia GTX 1080	\$1500
Jetson Tx2 CPU	CPU ARM Cortex A57	\$600
Jetson Tx2 GPU	Pascal	\$600
Raspberry Pi CPU	ARM Cortex A53	\$40



Outline

- A.I. at edge landscape
- Challenges with the current landscape
 - Learning at the Edge
- Neuro-Evolutionary algorithms
 - Introduction to NEAT
- Scaling NEAT at the Edge
 - Setup
 - Design choices and Evaluations
- Conclusion

Conclusion

- We demonstrate a system of agents running on Raspberry Pis learning collaboratively using neuro-evolutionary algorithms
- Proposed modifications to the algorithm allow
 - scaling to continue up to 65 Raspberry Pi nodes
 - showing a 2x performance improvement over naive scaling
 - reducing communication overhead by over 3.6x
- The proposed system using cheap Raspberry Pi hardware can outperform higher end computing platforms showcasing a Price-Performance Product benefit of 2.5x.