ERI ELECTRONICS RESURGENCE INITIATIVE SUMMIT & MTO Symposium

WORKSHOP:
New Opportunities for Lifelong Learning Machines

PROGRAM MANAGER: Ted Senator	
DATE: Wednesday, October 20, 2021	**TIME:** 3:15pm – 5:30pm
ROOM NAME: Ben Epstein	

DESCRIPTION

DARPA's Lifelong Learning Machines (L2M) Program is now in Phase 2. This phase focuses on the development of practical systems that address complete lifelong learning solutions covering continuous learning, transferability and mission life sustainability as applied to various domains. The workshop will be an opportunity for the L2M Phase 2 Systems Group teams to describe their L2M systems and targeted applications, and enable attendees to better understand how to make use of outcomes from the program, transitioning L2M program results to deployed self-learning systems.

AGENDA

3:15 to 3:20pm	***Overview of the L2M Program*** Ted Senator, DARPA, L2M Program Manager
3:20 to 3:30pm	***Autonomous Navigation and Classification*** Mario Aguilar-Simon, Teledyne Scientific, Fellow
3:30 to 3:40pm	***Autonomous Navigation and Game Play Domains*** Praveen Pilly, HRL, Senior Research Scientist
3:40 to 3:50pm	***Lifelong Learning at the Edge*** Angel Yanguas-Gil, Argonne National Labs, Principal Materials Scientist
3:50 to 4:00pm	***The Eigentask framework*** Aswin Nadamuni Raghavan, SRI International, Senior Computer Scientist
4:00 to 4:15pm	**Q&A**
Afternoon Break: 4:15pm-4:30pm	
4:30 to 4:40pm	***Overcoming Catastrophic Forgetting with Meta-Learned Neuromodulation*** Nick Cheney, Univ. of Vermont, Assistant Professor
4:40 to 4:50pm	***A NeuRoBot That Learns Locomotion Online*** Francisco Valero-Cuevas, Univ. of Southern California, Professor
4:50 to 5:00pm	***L2M Hardware Implementations*** Dhireesha Kudithipudi, Univ of Texas, San Antonio, Professor
5:00 to 5:10pm	***Lifelong Learning through Random Forests*** Josh Vogelstein, Johns Hopkins University, Professor
5:10 to 5:30pm	***Discussion and Q&A***
Workshops Conclude at 5:30pm	

QUESTIONS

Please contact Ben Epstein (benjamin.epstein.ctr@darpa.mil) or ERI-Summit@darpa.mil for more information on this workshop.

New Opportunities for Lifelong Learning Machines

Mario Aguilar-Simon

Intelligent Systems Laboratory

mario.aguilar-simon@teledyne.com

919.323.3485

ERI Summit

October 20, 2021

Uncertainty Modulated Learning (UML)

- Current machine learning systems are brittle or scope-limited because they
  - → degrade if conditions are different from those they were trained on or cannot learn new tasks
  - → cannot reliably learn from data encountered during execution time
  - → cannot adapt on-line to changes they encounter in real environments
- Our algorithm, UML, is inspired by the brain's mechanism of neuromodulation, chemical signaling that continuously modulates neural activity and plasticity to regulate changes and adaptation

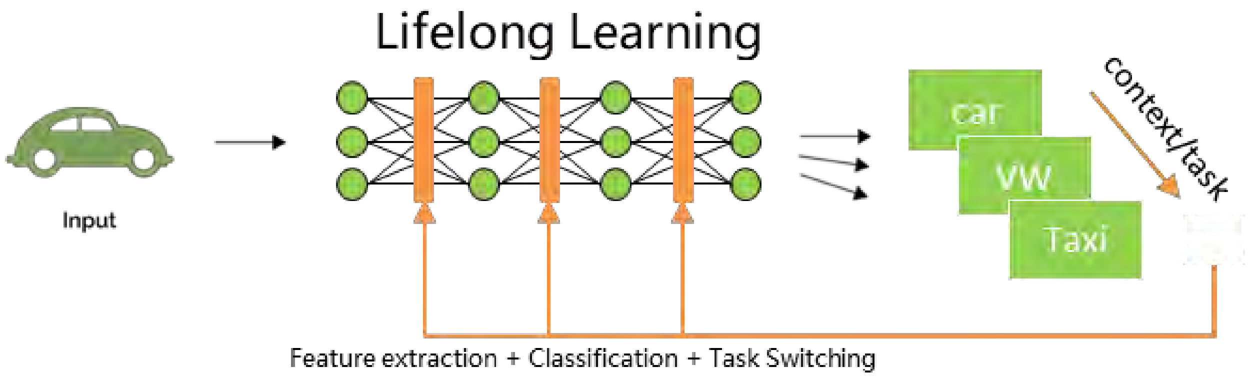

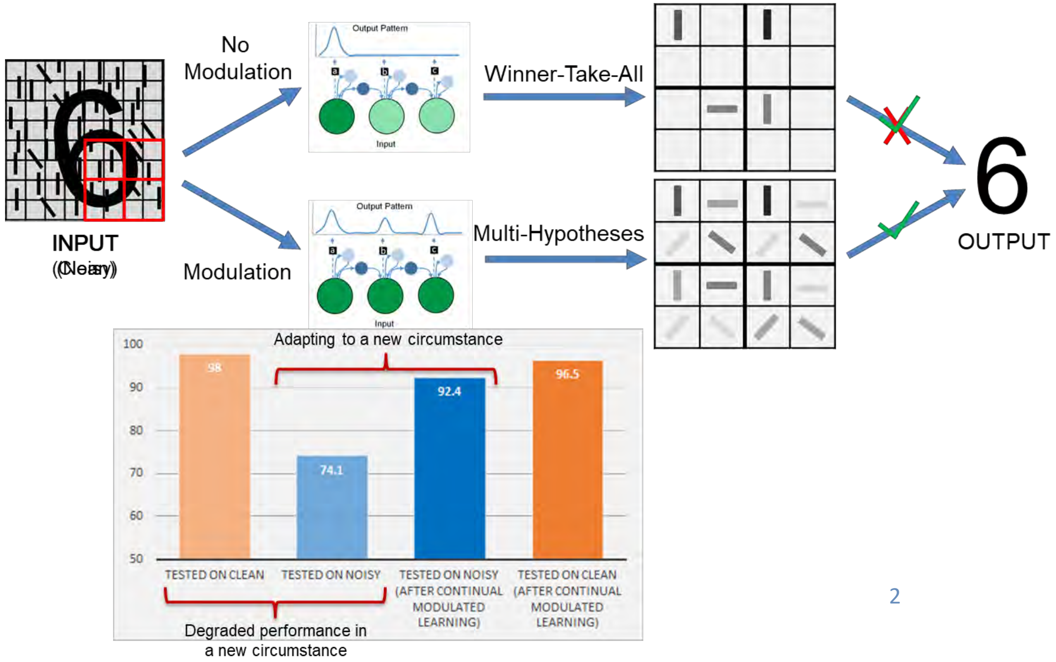

L2M Project Overview

- Algorithm Approach: From Biology to Novel Machine Learning
  1. Improved representations to support robust performance
  2. Converging Bottom-Up and Top-Down signals to evaluate incoming signals
  3. Discriminative and Generative modeling to quantify decision hypotheses
  4. **Uncertainty tracking to regulate processing across all layers**

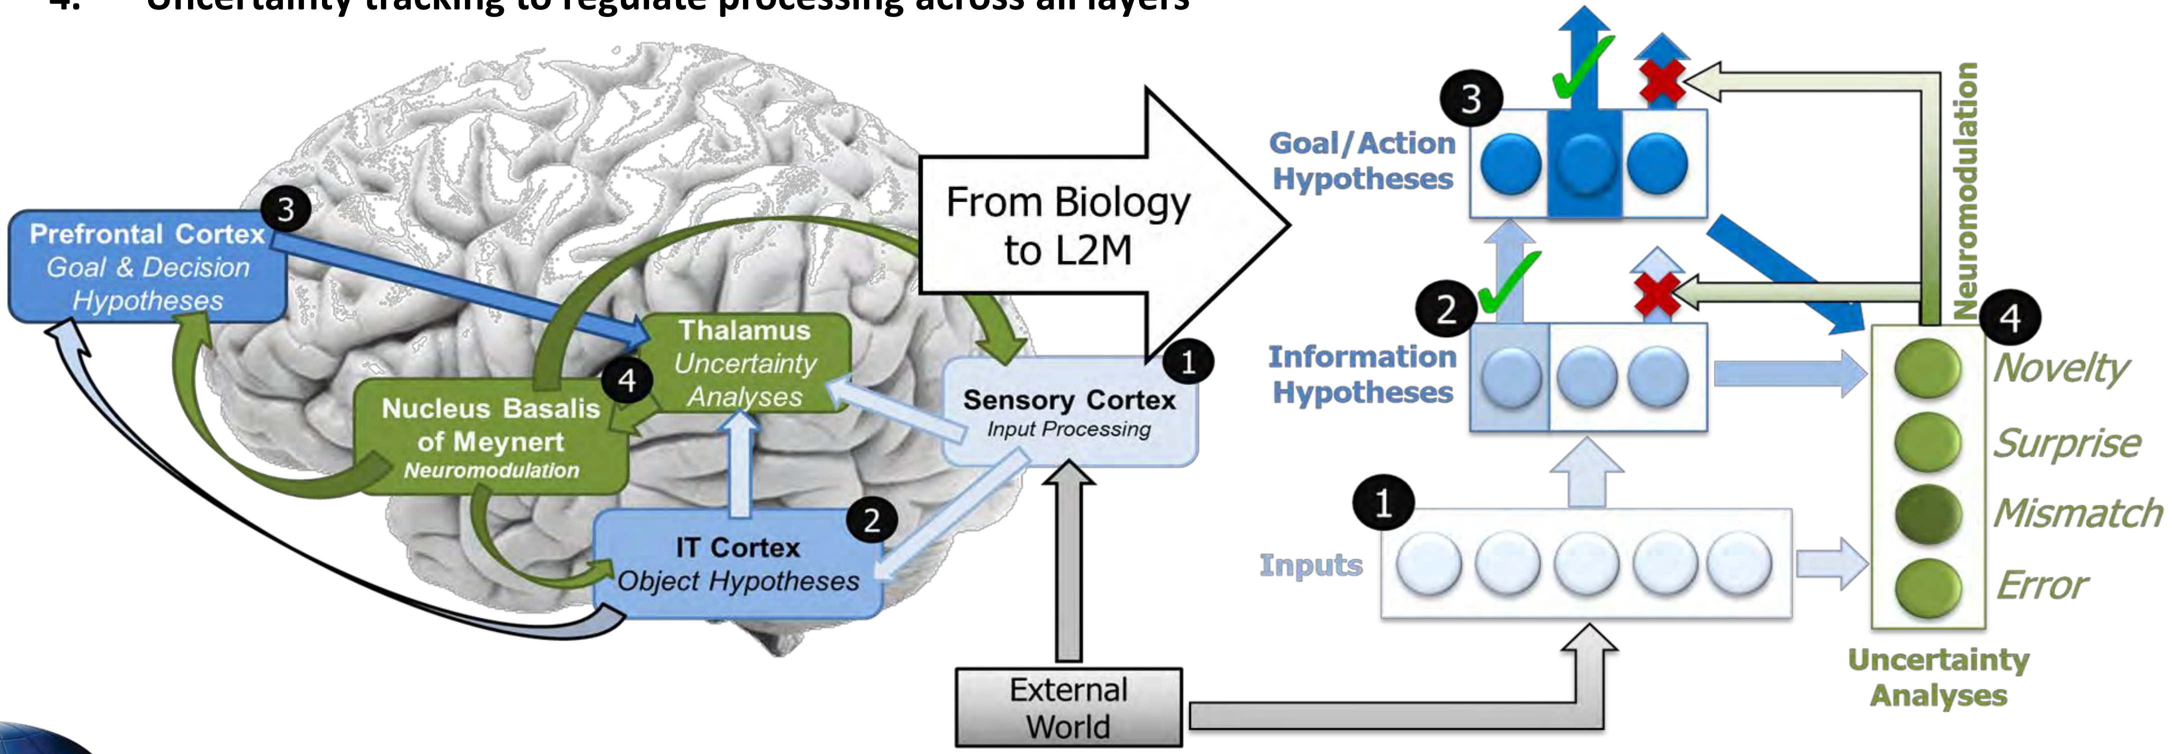

Self-Supervised Learning on a Drone "Driven" by Uncertainty

TELEDYNE SCIENTIFIC COMPANY
Intelligent Systems Laboratory

We implemented the UML algorithm on-board a simulated drone to support real-time self-supervised adaptation to new circumstances and continual learning

Learning to recognize known objects (sedans and vans) from novel views

1

2

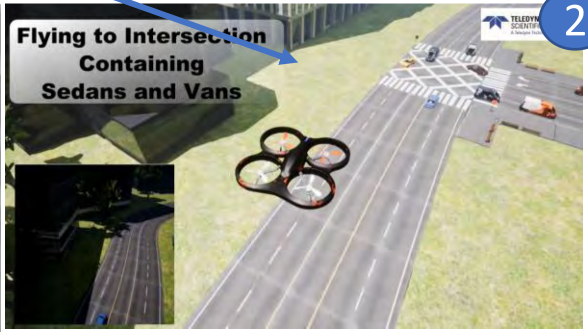

4

Learning to avoid novel obstacles

5

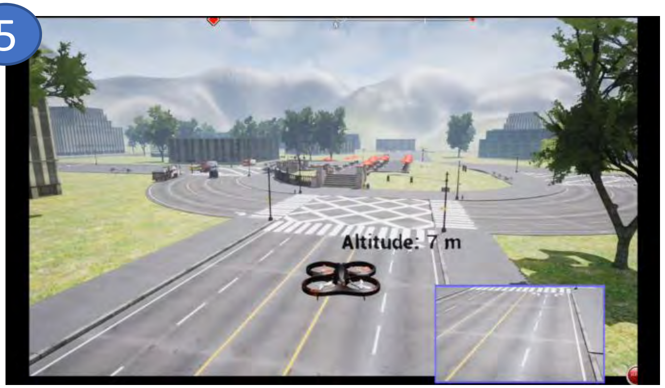

Learning novel objects
without supervision

3

Sedans

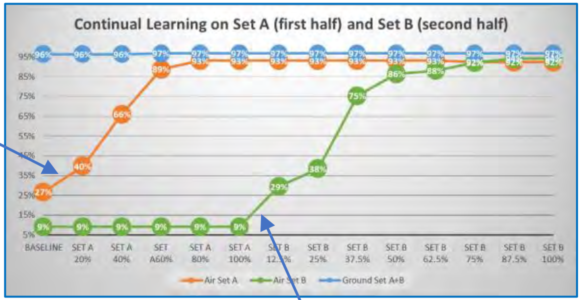

Vans

DARPA

Demonstrated ability to improve performance and learn new tasks without supervision and without forgetting previous knowledge.

https://www.youtube.com/watch?v=_asgnjHTqpc

UML for Classification and Performance Prediction

TELEDYNE SCIENTIFIC COMPANY
Intelligent Systems Laboratory

- Improved ATR performance in novel conditions through uncertainty tracking
  - A critical risk for DNNs is performance in the presence of novel targets or clutter

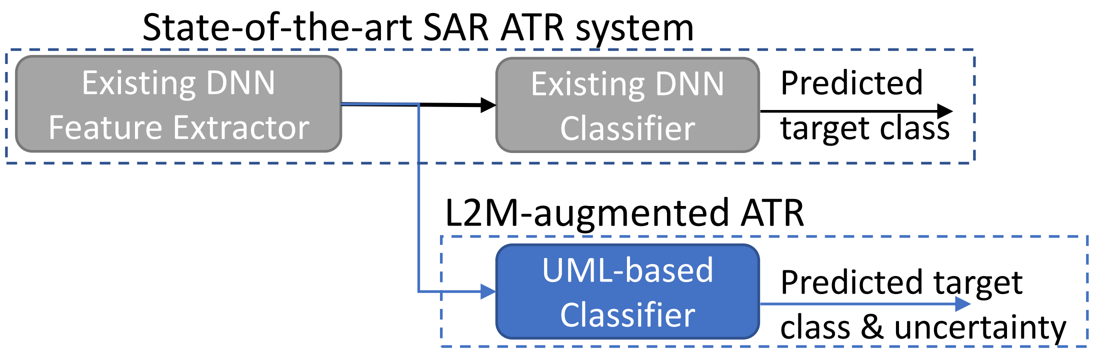

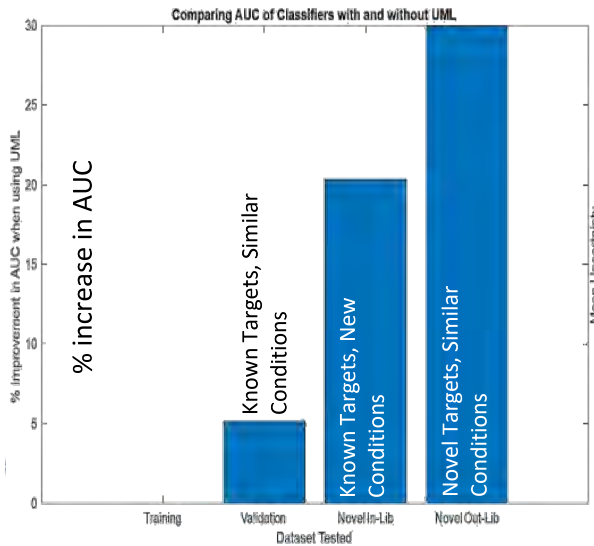

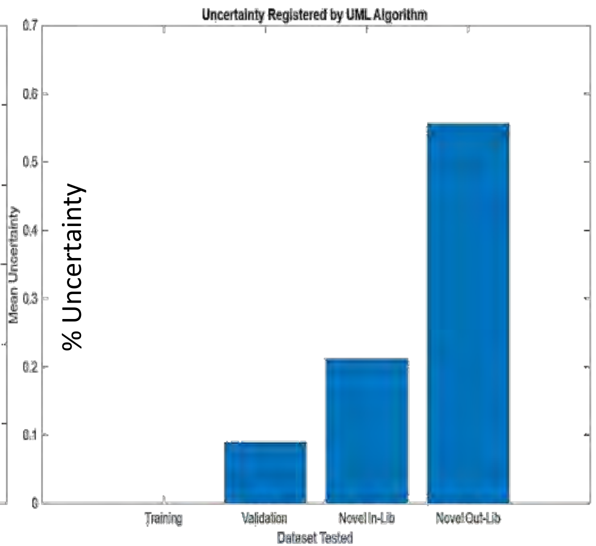

- Under DARPA CAML, demonstrated the ability to predict the performance of a Machine Learning system (MLS) on novel datasets

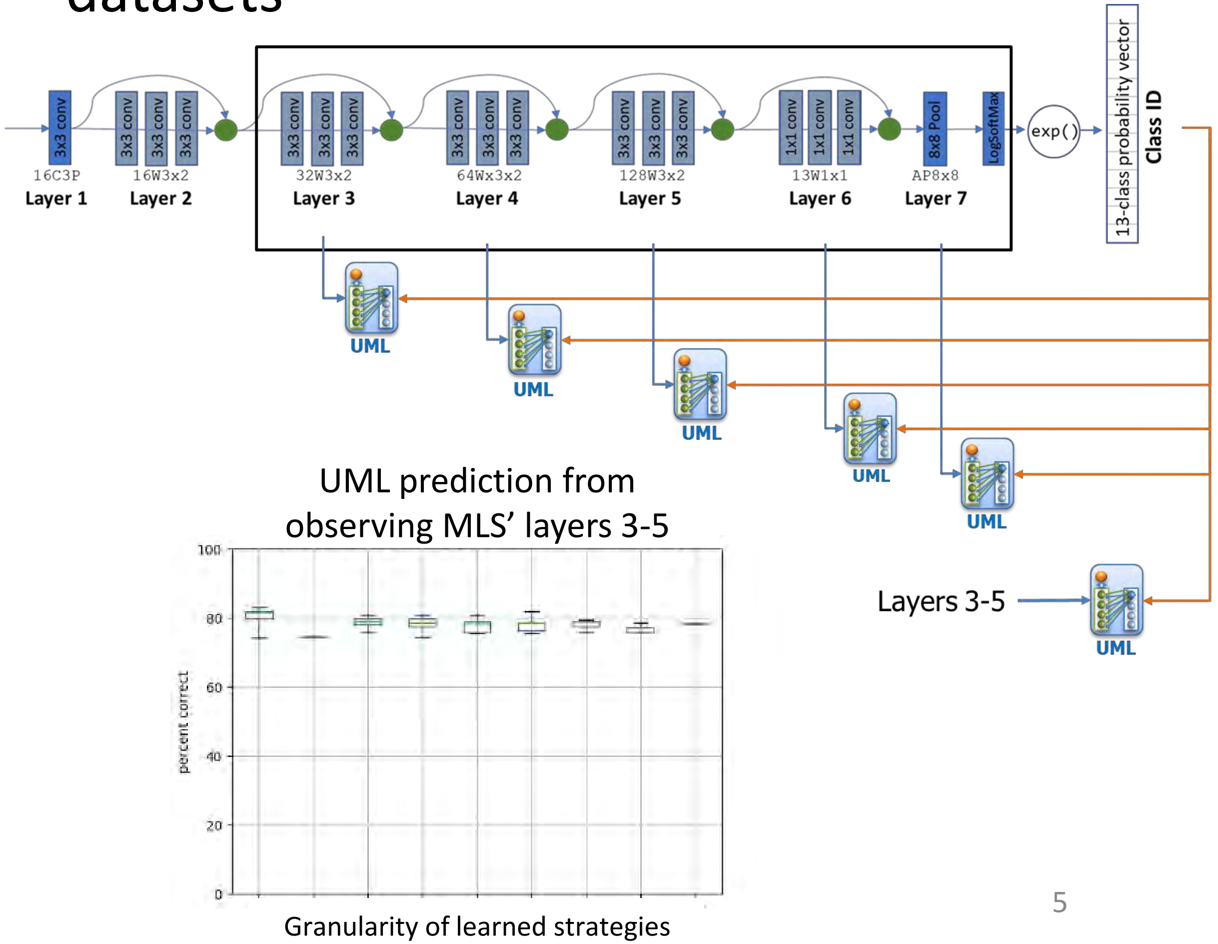

UML for Object Classification: Recyclables Sorting Use Case

The EPA estimated that 8% of plastic in the US is recycled, partially due to contamination in the sorting streams

- Enabling better sorting through AI will have **a significant environmental impact**

UML was compared to the best performing classifier, DenseNet, selected based on its performance on pre-trained classes (~92% accuracy on 10 classes)

- After training both on the full set of 27 classes, DenseNet's performance degraded substantially

Current Approach

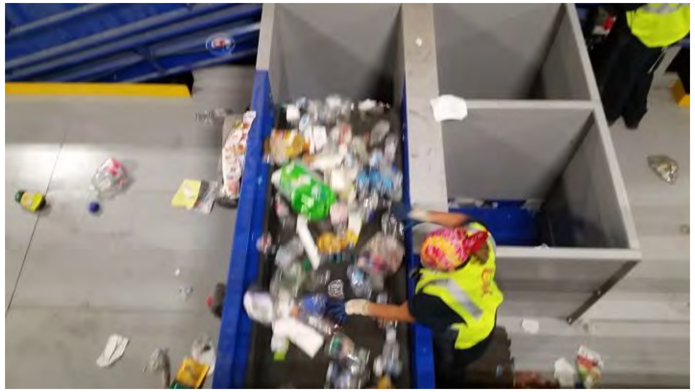

Lifelong Learning-Enabled Approach

	Accuracy (%) on 27 Classes
YOLO + DenseNet	64.0
YOLO + UML	88.1

Did not scale

Minimal impact

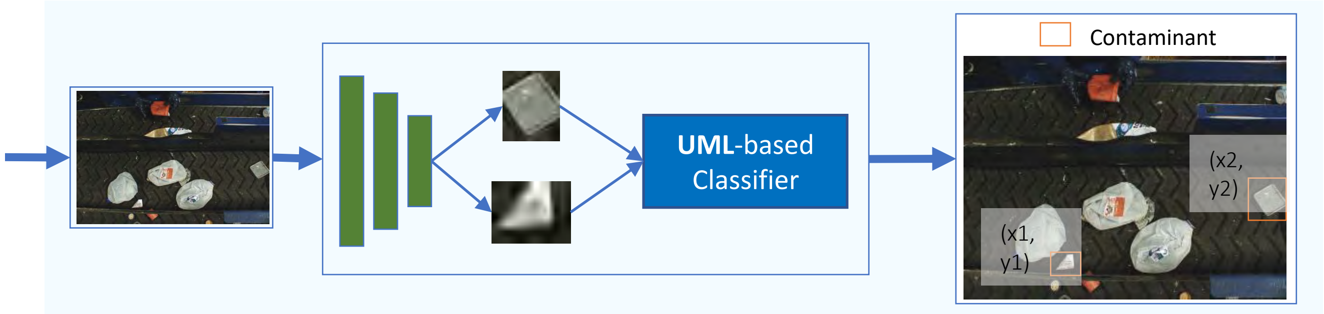

DARPA

UML for Out-of-Nominal Detection and Correction in RL

TELEDYNE SCIENTIFIC COMPANY
Intelligent Systems Laboratory

- Applied UML to improve robustness in a RL agent under changing conditions:
  - Train this model to evaluate the predicted state against learned templates of experienced states, detect when there is high uncertainty (out-of-nominal conditions), and override initial action response

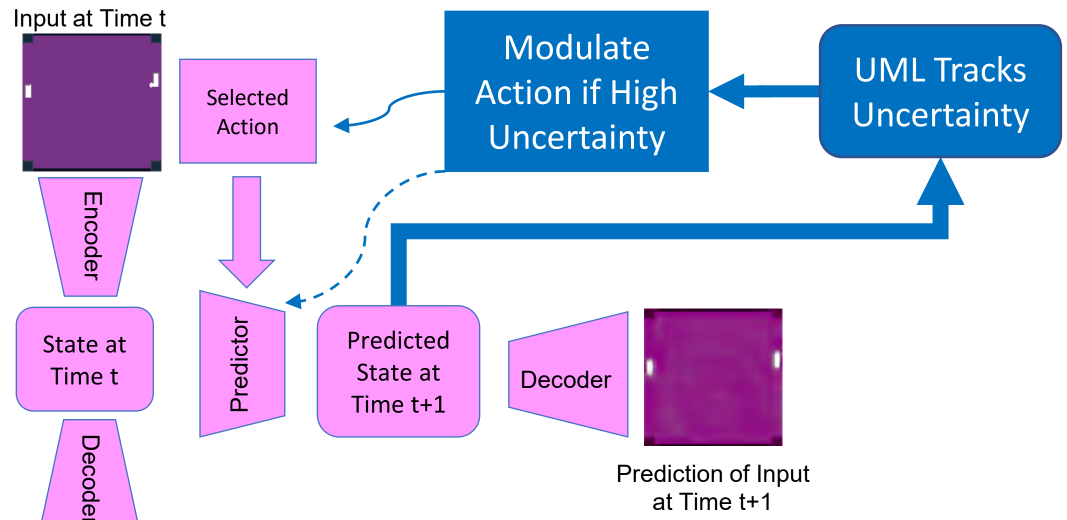

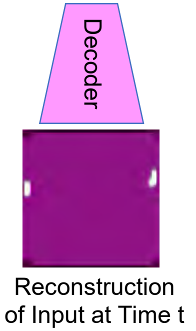

We implemented this system and conducted experiments for the Atari game Pong in OpenAI Gym

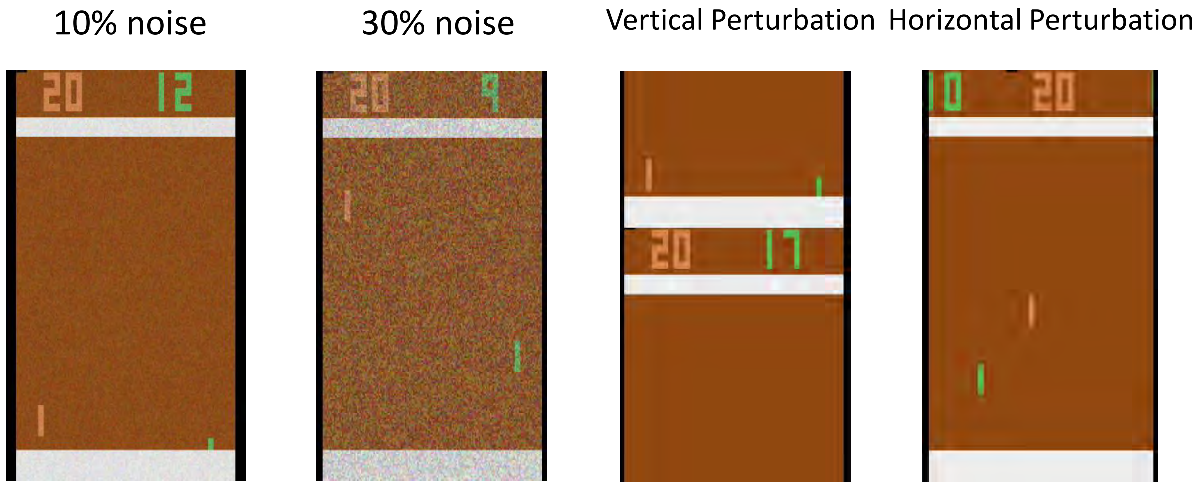

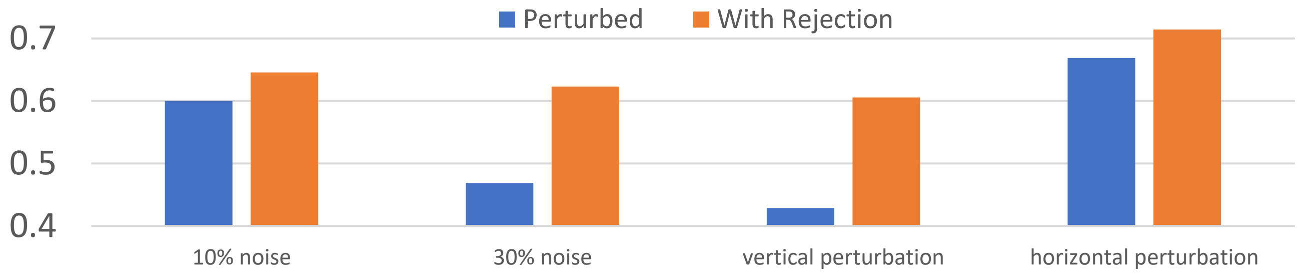

Portion of points scored before and after UML modulation in the presence of out-of-nominal conditions

Piergiovanni, A. J., Wu, A., & Ryoo, M. S. (2018). Learning real-world robot policies by dreaming. *arXiv preprint arXiv:1805.07813*.

Transition Opportunities

TELEDYNE SCIENTIFIC COMPANY
Intelligent Systems Laboratory

- Our UML algorithm (C4) can be used as a drop-in replacement in any system/program where either supervised, unsupervised or semi-supervised algorithms are being used, which will equip the underlying system with a continual and self-supervised learning capability. TRL = 4-5
- Our uncertainty tracking algorithm (C2) can be used to inform the performance of any machine learning algorithm or system that relies on machine learning algorithms. Furthermore, such systems can be modified to use information about uncertainties to adapt their execution. TRL = 3-4
- Both of the components above, when integrated in a system for object recognition, can support capabilities for rapid adaptation to new environments or the addition of new object types or variations. TRL=5-6 (under testing in a commercial environment)
  - Ongoing efforts include testing this system on a Teledyne FLIR SkyRaider
  - The algorithms can run in real-time on a Nvidia TX2

R80D SkyRaider VTOL aircraft

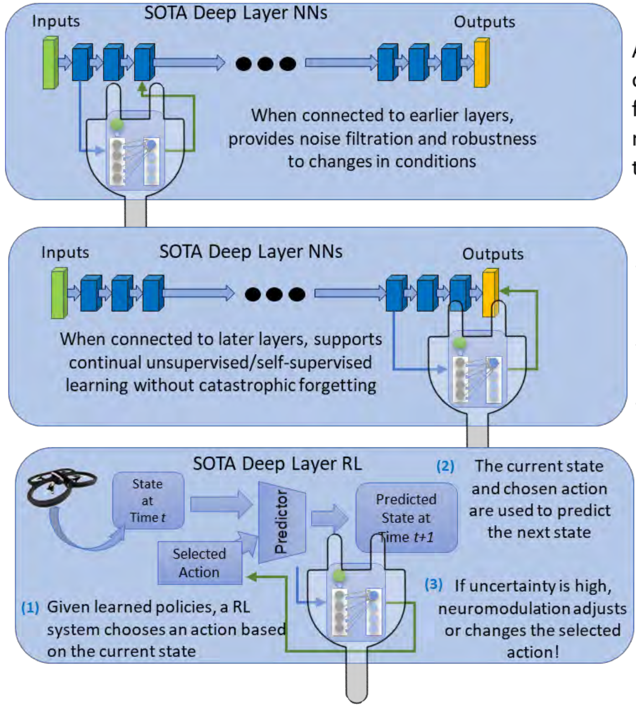

Able to support adaptive data enhancement / feature extraction by monitoring uncertainty in the signals

Able to support adaptive decision making and learning in real-time and from few samples by monitoring uncertainty in the decision

Able to support adaptive policy selection and updating by monitoring uncertainty in the inference process

DARPA

Thanks

Questions?

Mario Aguilar-Simon
Intelligent Systems Laboratory
mario.aguilar-simon@teledyne.com
919.323.3485

Lifelong Learning for Autonomous Driving

Praveen K. Pilly
HRL Laboratories

Lifelong Learning (L2) for Autonomous Driving

pink: lane marking
red: another car

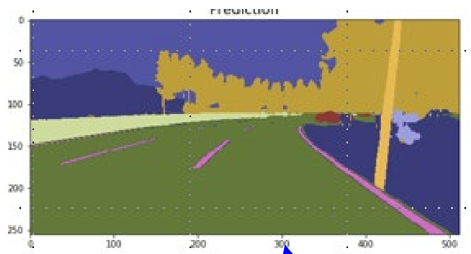

Changed domain creates misperceptions & wrong behaviors

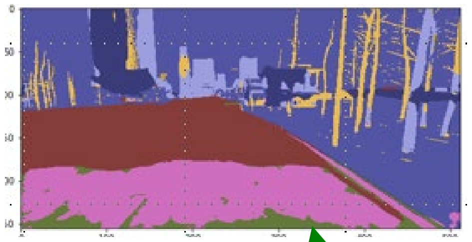

Training

Fielded

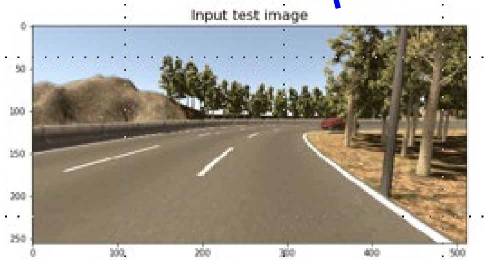

Deployed

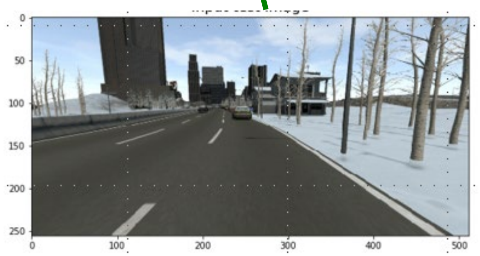

Current CONOPS (unacceptable)

Cloud → Re-train/update base model → Test and verify all core capabilities → Software update

Log of recent experiences leading up to “surprise events”

several weeks later

Challenges

- Need to learn new conditions (surprises), but without forgetting old skills
- Minimize storage of data for all previous conditions
- Minimize re-training base model with all stored data

Lifelong Learning for Autonomous Systems

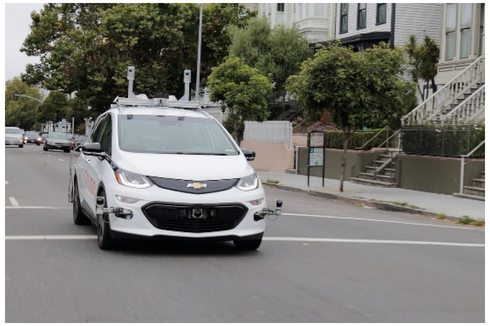

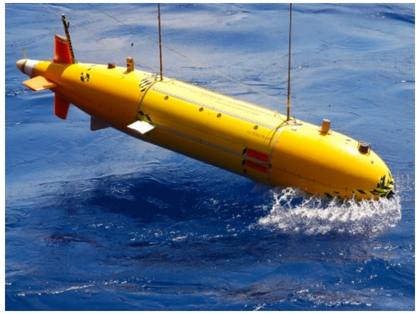

theverge.com
polaris.com
eos.org

L2 Capability	L2 Metric
Generalize to new tasks as well as improve on old tasks	Forward Transfer Backward Transfer
Rapidly recover performance for an old task	Performance Recovery
Rapidly adapt to new tasks	Performance Relative to a Single Task Expert Sample Efficiency
Safely adapt to new tasks	Learn Burn
Avoid catastrophic forgetting of old tasks	Performance Maintenance
Scale up learning of new tasks	Cumulative Gain
Minimize storage of data for old tasks and re-training	SWaP

HRL System Group (SG): L2 Components

Primary Responsibility	Component	Component Description
HRL	Sliced Cramer Preservation	To avoid catastrophic forgetting
HRL	World Model	For backward transfer to variants of old tasks
HRL	Neurogenesis	To scale up the learning of new tasks
UCI	Neuromodulated Attention	To rapidly re-adapt to old tasks
UTA	Context-Skill Model	For forward transfer to variants of old tasks
LU	Modulated Hebbian Network	To rapidly re-adapt to old tasks
LU	Plastic Neuromodulated Network	To rapidly adapt to new tasks
INRIA	Reflexive Adaptation	To safely adapt to new tasks
ITU	Meta-Learned Instinct Network	To safely adapt to new tasks
Baylor	Metaplasticity Kernel Model	To avoid catastrophic forgetting
JHU	Matrix Sketching (SCP++)	To avoid catastrophic forgetting

HRL SG System Architecture

SCP	Sliced Cramer Preservation
NG	Neurogenesis
MP	Metaplasticity

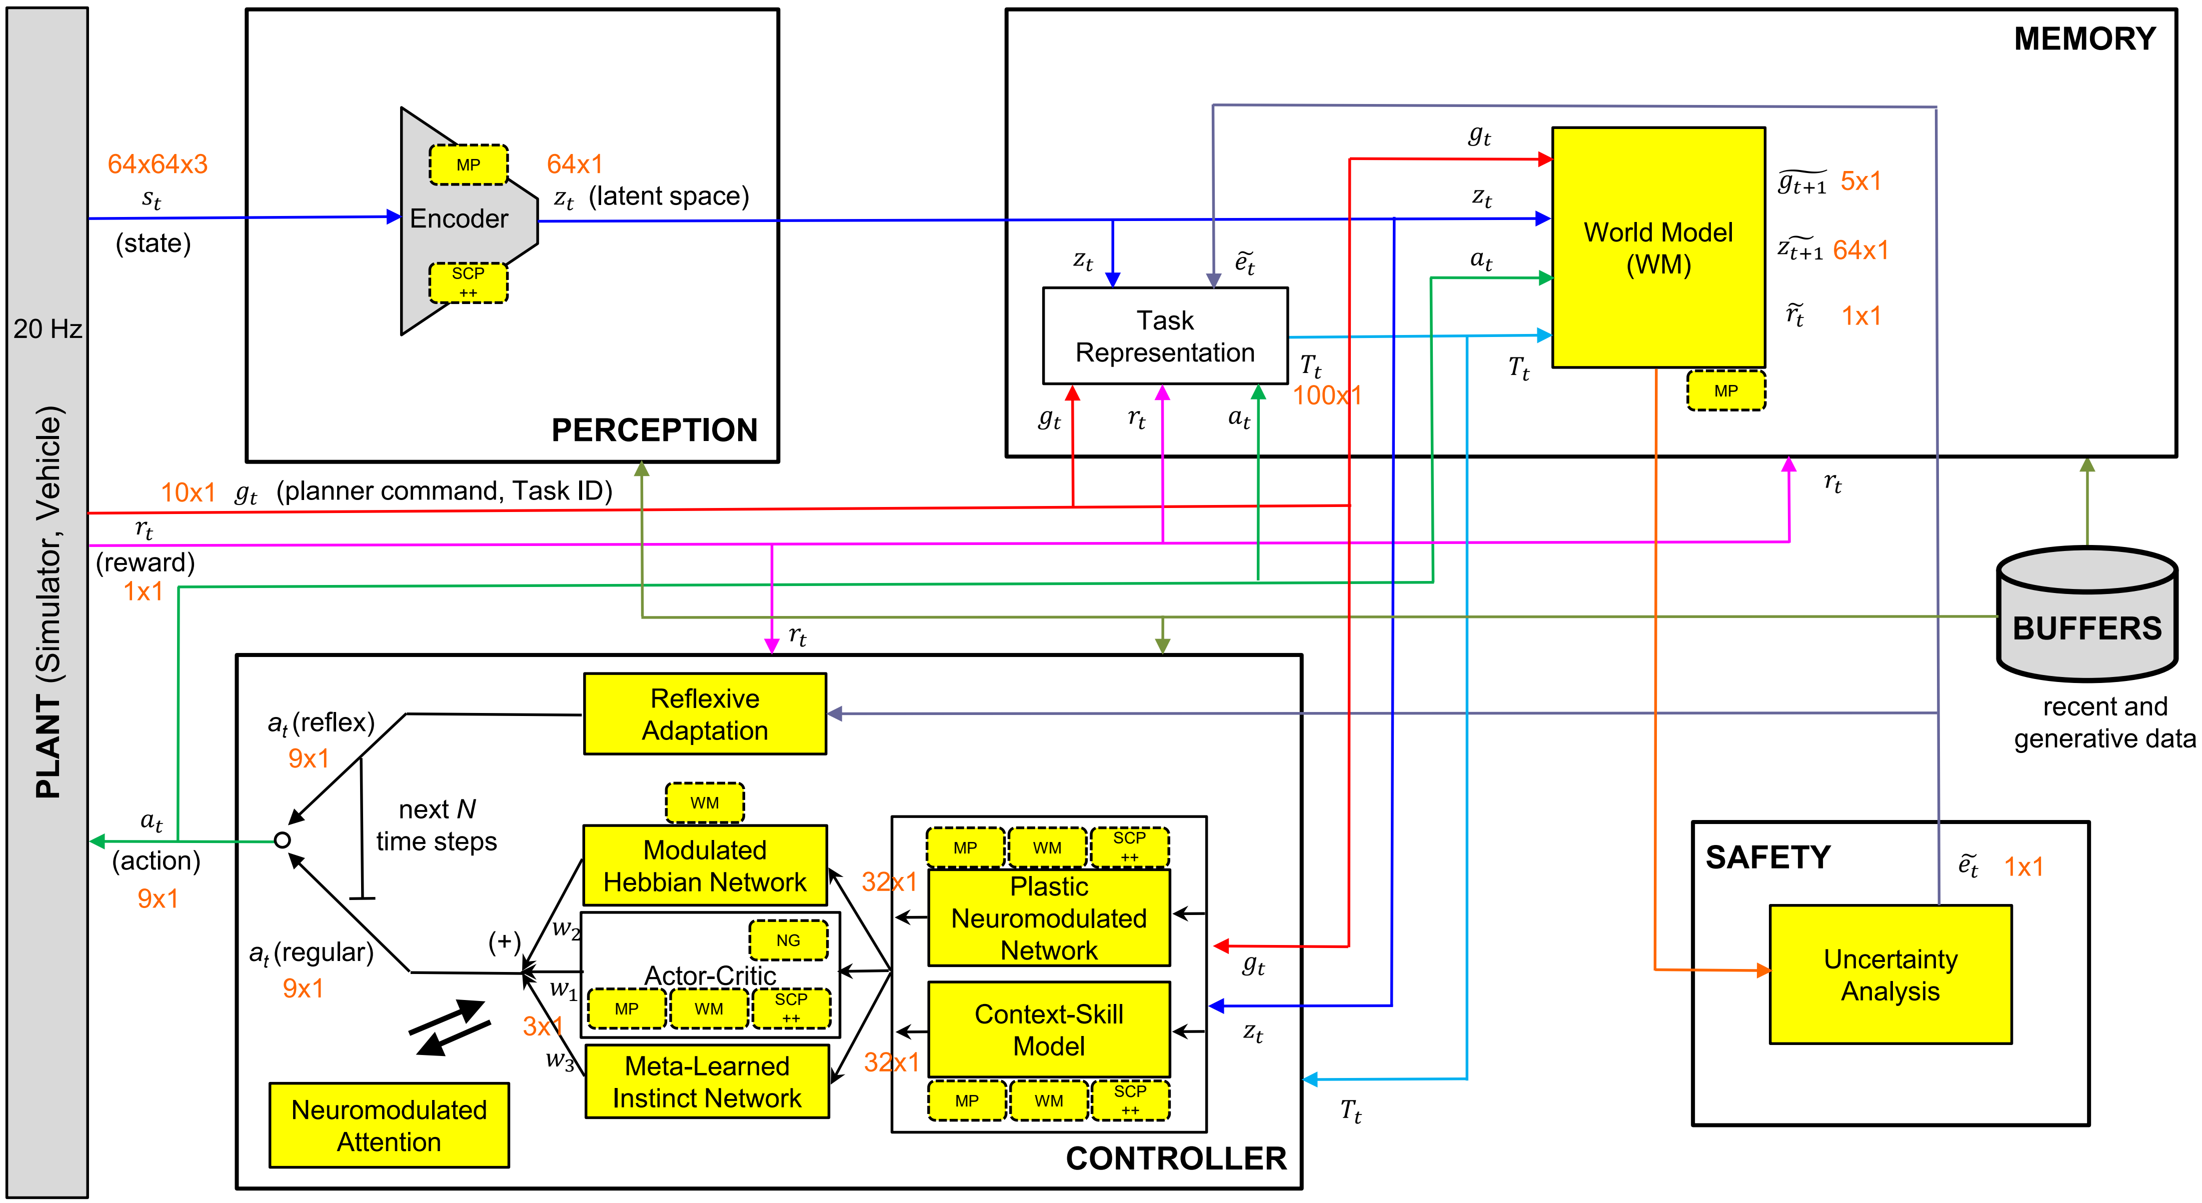

How do various components facilitate L2?

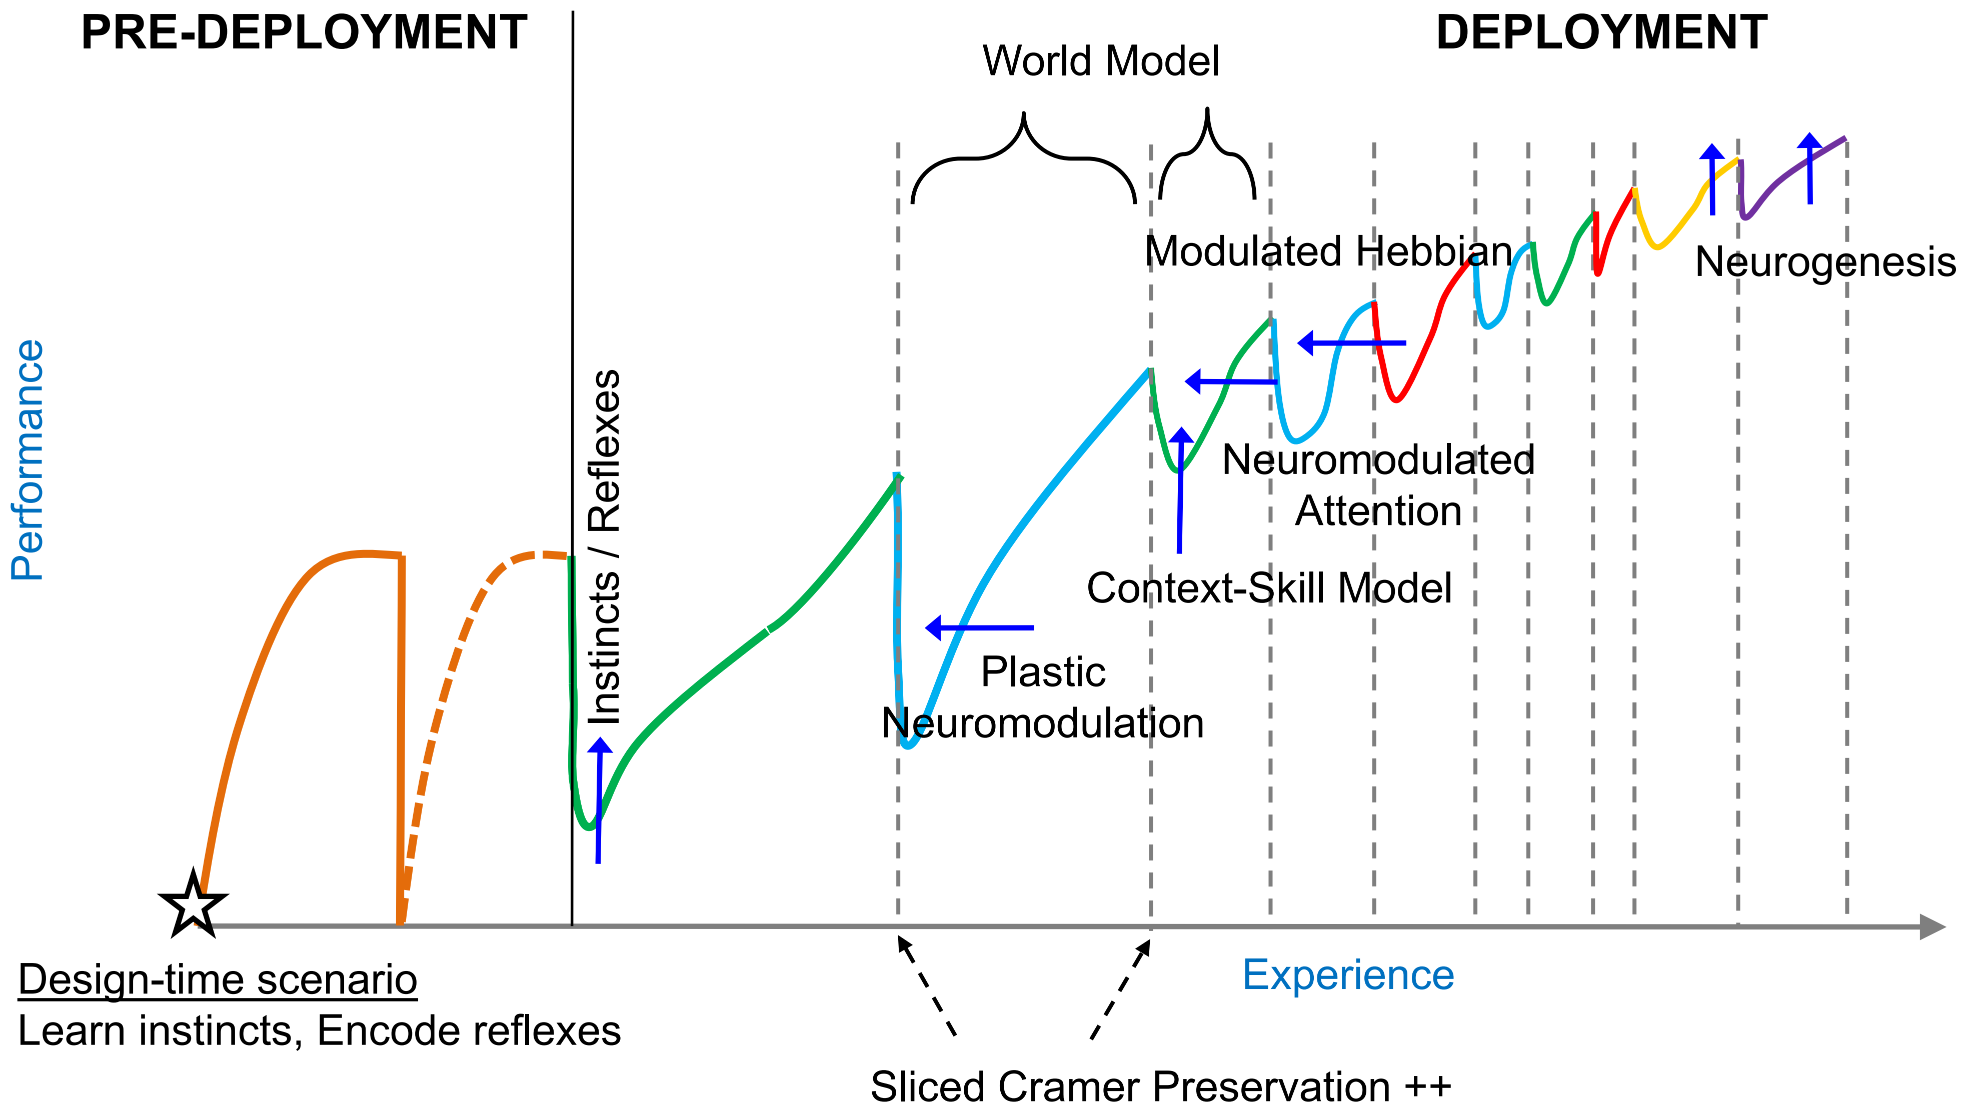

Phase 2 Evaluation Tasks

Task	Vehicle Model	Driving Lane	Precipitation	Sun Altitude
T2_1	Car	Correct	0% (Clear)	15^{o} (Dusk)
T2_2	Car	Correct	100% (Rainy)	15^{o} (Dusk)
T3_1	Motorcycle	Correct	0% (Clear)	-60^{o} (Night)
T3_2	Motorcycle	Correct	100% (Rainy)	-60^{o} (Night)
T1_1	Motorcycle	Opposite	0% (Clear)	90^{o} (Midday)
T1_2	Motorcycle	Opposite	100% (Rainy)	90^{o} (Midday)

Audi TT

Kawasaki Ninja

T1_1

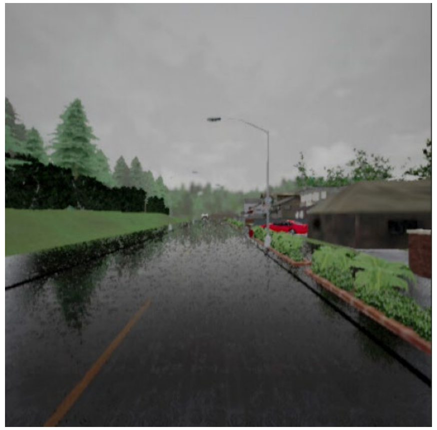

T2_2

Differences in physical parameters (body: mass, drag coefficient; wheels: friction, damping rate, max. steering angle, radius, etc.)

Phase 2 Evaluation Results

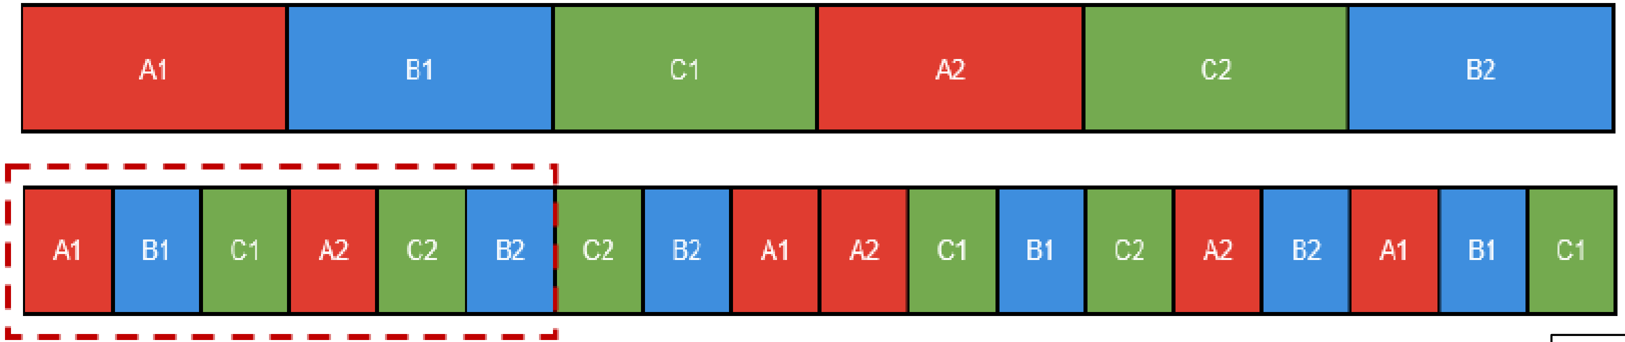

Condensed scenario

Dispersed scenario

- Cost of switching
- Greater interference from other tasks

L2 Metric	Condensed (n=30)	Dispersed (n=28)
L2 Components	11	11
Performance Maintenance (>0)	-0.24	-2.09
Forward Transfer (>1)	12.08	10.74
Backward Transfer (>1)	1.19	1.09
Performance Relative to a STE (>1)	2.51	1.84
Sample Efficiency (>1)	11.52	6.13

Personal Autonomous Driving

- When does the user trust the car enough to give it full autonomy?
- How can user feedback and imitation help the car to improve human trust?

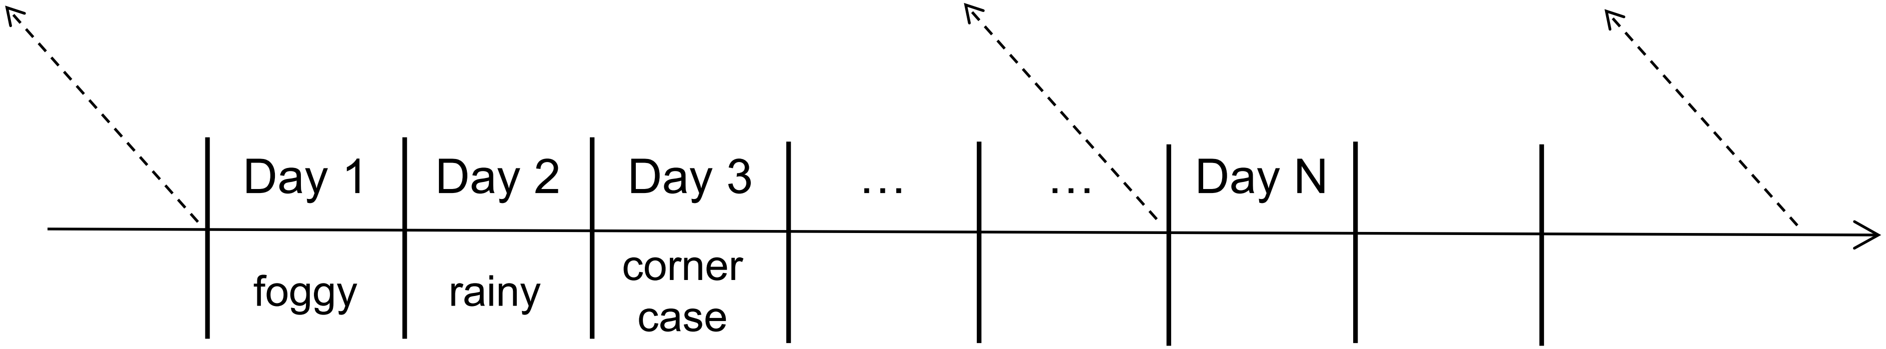

Continual Imitation Learning

DARPA LIFELONG LEARNING MACHINES

LIFELONG LEARNING AT THE EDGE

ANGEL YANGUAS-GIL
Applied Materials Division
Email: ayg@anl.gov

Argonne National Laboratory is a U.S. Department of Energy laboratory managed by UChicago Argonne, LLC.

10/20/2021

WHAT PROBLEM ARE WE TRYING TO SOLVE?

Many applications would benefit from AI capable of learning after deployment

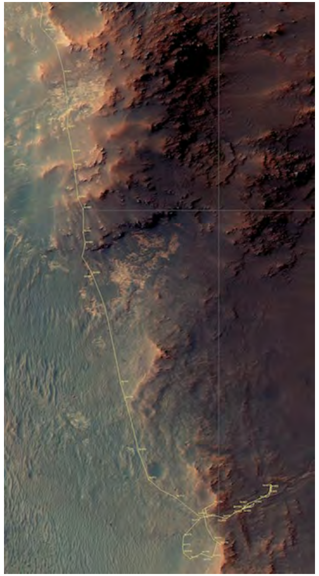

Source: NASA (public domain)

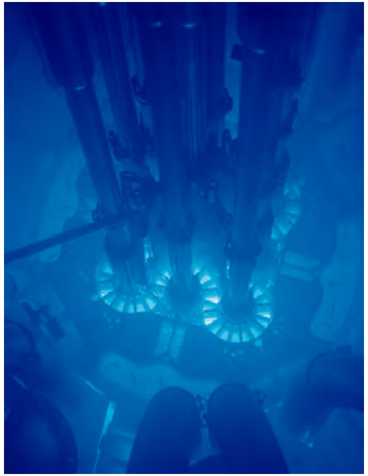

Extreme environments: nuclear reactors

Edge processing applications

Autonomous vehicles

Smart sensing

Manufacturing and control

Systems embedded in extreme environments

… and more

INSECT BRAIN AS MODEL SYSTEM FOR THE EDGE

Insects are smart, autonomous sensors in a compact package

~ 1 mW

Bee: 1,000,000 neurons
Fruit fly: 100,000 neurons

Insects embody many key requirements for having smart sensors capable of continuously learning at the edge

(bees can learn to navigate to multiple locations miles away from the hive depending of the time of day)

Insect brains are within reach of existing, globally available semiconductor fabrication capabilities

Our current gap is knowledge: we don’t know how to put the pieces together

INSECT BRAIN AS MODEL SYSTEM FOR THE EDGE

Our approach in this program

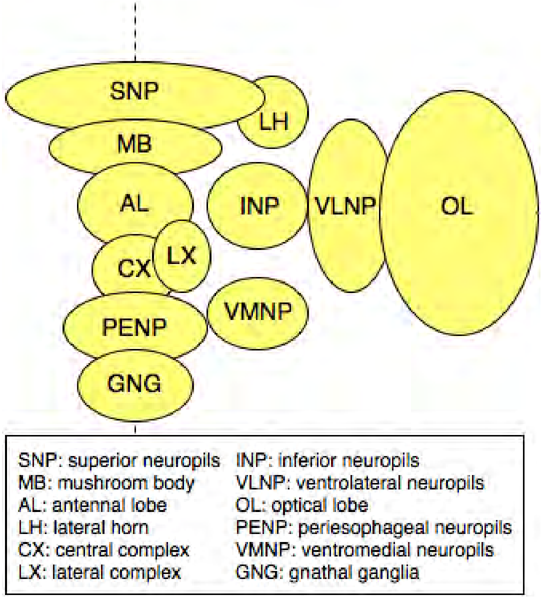

(From neuroscience to math)

→ Algorithms →

Online learning of multiple tasks without forgetting

Ability to operate in noisy, degraded environments

Introspection: use past experience and context to understand when and how to learn

Understand how to optimize architectures for lifelong learning at the edge

INSECT BRAIN AS MODEL SYSTEM FOR THE EDGE

Our approach in this program

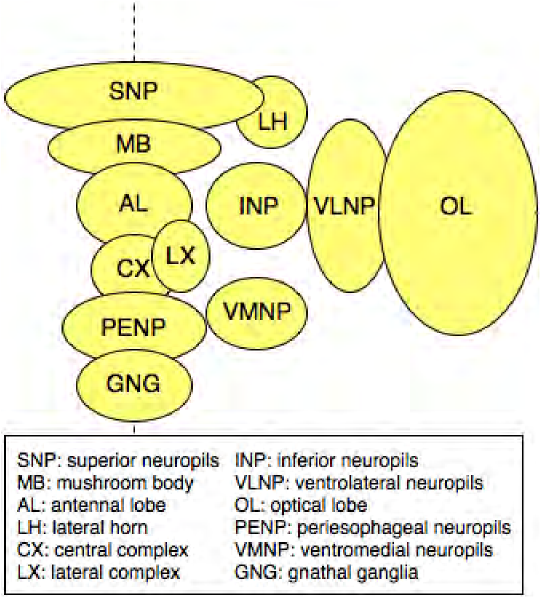

(From neuroscience to math)

→ Algorithms →

Online learning of multiple tasks without forgetting

Ability to operate in noisy, degraded environments

Introspection: use past experience and context to understand when and how to learn

Understand how to optimize architectures for lifelong learning at the edge

Key characteristics: learning based on local learning rules, sparse representation, neurochemistry codifying internal state, persistent and broad representations

INSECT BRAIN AS MODEL SYSTEM FOR THE EDGE

Our approach in this program

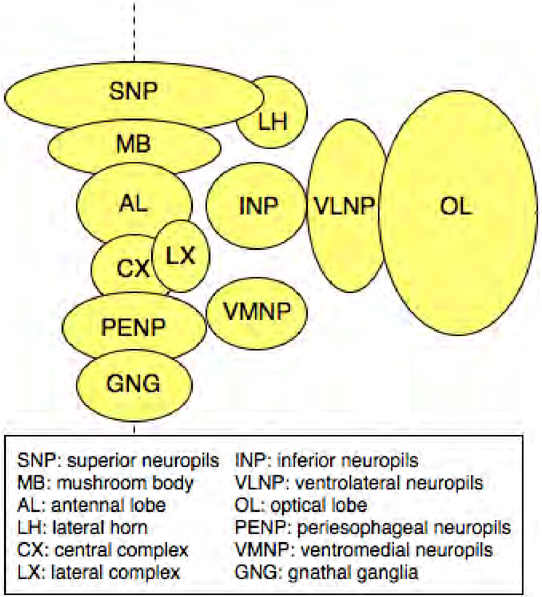

→ Algorithms →

Applications

Signal processing - i.e. RF

Reinforcement learning / interactive systems

New benchmarks for online learning at the edge (learning sparse inputs)

INSECT BRAIN AS MODEL SYSTEM FOR THE EDGE

Our approach in this program

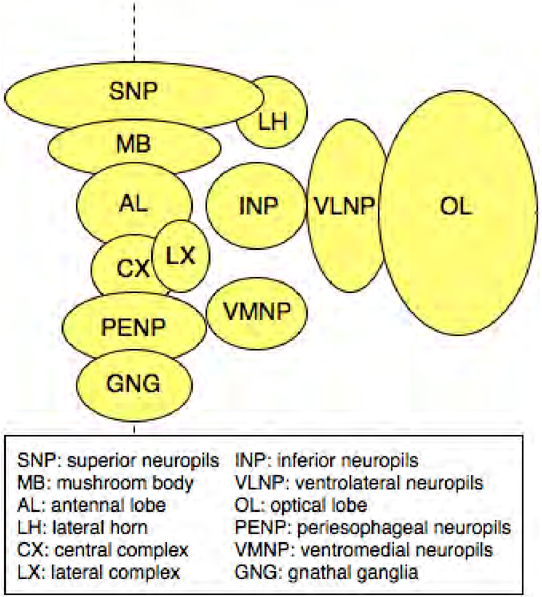

→ Algorithms →

Hardware

FPGA implementations

Neuromorphic chips (Loihi)

Emergent materials

Applications to extreme environments

INSECT ARCHITECTURES: HOW GOOD THEY ARE?

In some cases they can outperform standard ML algorithms

		Task-Incremental Learning		Class-Incremental Learning	
	Method	Split-MNIST	Split-CIFAR-10	Split-MNIST	Split-CIFAR-10
Baseline	iid-offline	99.50	95.43	95.93	80.64
	Fine-Tune	97.78	70.28	19.65	17.34
Continual Learning Memory-free	Online EWC	97.94	60.64	19.66	17.41
	SI	97.84	60.20	19.81	17.67
	LwF	99.16	59.95	21.37	18.68
Continual Learning Memory-based	A-GEM	99.31	68.82	50.36	17.94
	iCaRL	98.50	82.44	72.49	38.92
	GSS	98.46	86.22	53.69	48.37
	RPSNet	–	67.0	–	–
	InstAParam	–	83.8	–	–
	ER-MIR	–	–	87.6	40.0
	CN-DPM	–	–	**93.81**	47.11
	DER++	99.36	87.11	92.34	54.08
	Ours (transfer w/ INEL)	65.25	81.20	21.56	22.33
	Ours (transfer w/ MSE)	**99.60**	**94.65**	21.84	22.85
	Ours (Opt)	–	–	78.76	**55.74**

Evolved architectures perform surprisingly well despite not storing past information

EXAMPLE: LIFELONG LEARNING OF RF SIGNALS

Key contributions

Architecture based on the mushroom body of the insect brain with sparse representation, local plasticity rules

Ternary synapses:

Modulatory

Presynaptic 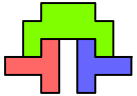 Postsynaptic

8PSK 16QAM

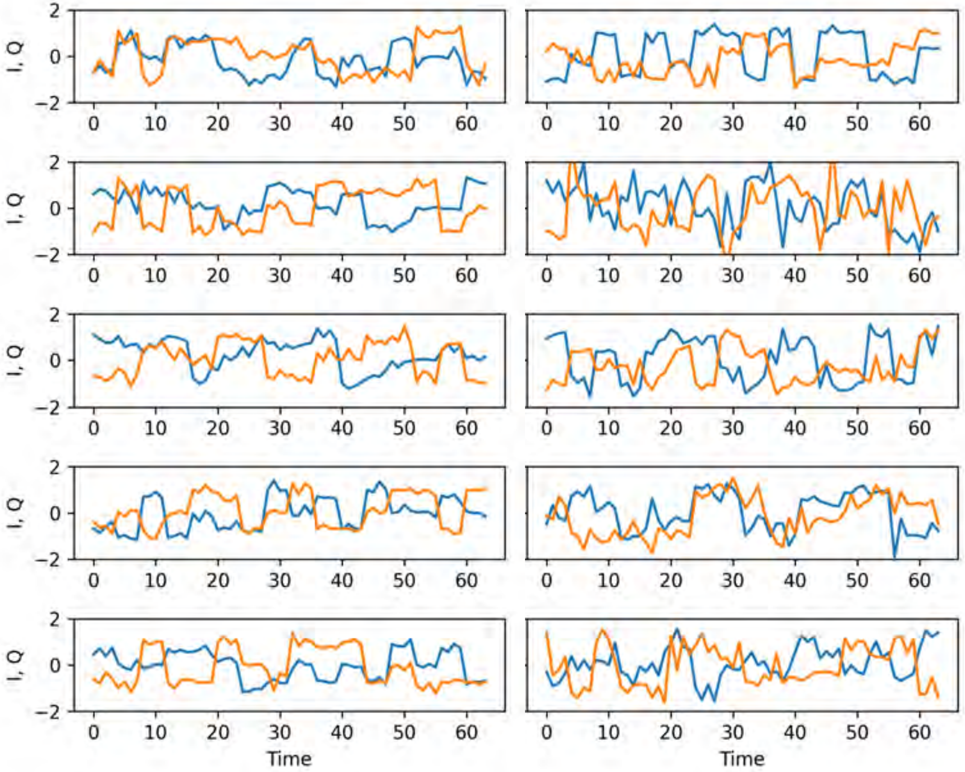

EXAMPLE: LIFELONG LEARNING OF RF SIGNALS

Continual and online learning results

Learns using a few examples

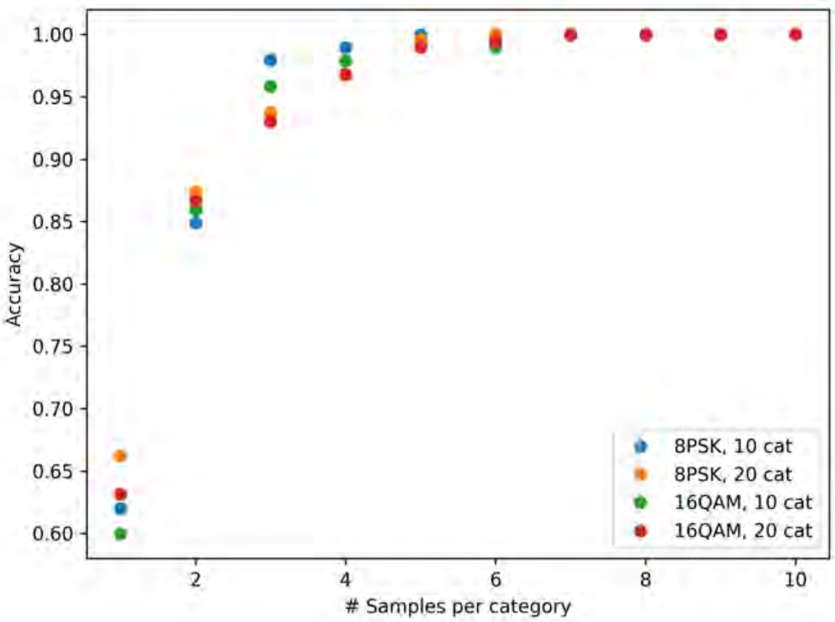

Learns multiple tasks without forgetting

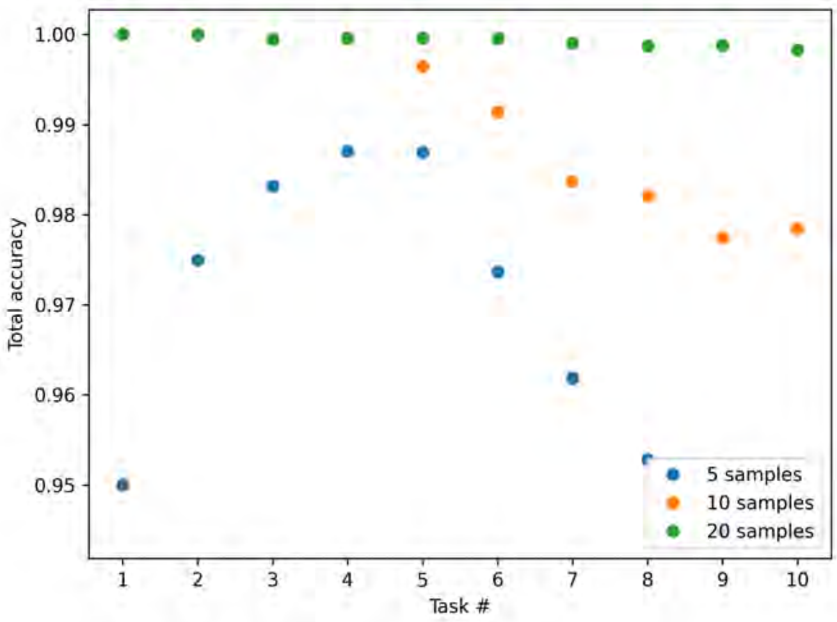

Processing component can be built using ~2,000-5,000 transistors

CONCLUSIONS

Insects provide a good model system for lifelong learning at the edge

Small sizes and lightweight nature are a good model system for novel architectures optimized to perform under power and resource constrained scenarios

We have explored their capabilities both using ML benchmarks and application-inspired tasks

Optimization of the architectures and learning rules lead to configurations that can compete with traditional machine learning algorithms

Same co-design approach has helped us identify optimal emerging materials for on chip, online learning architectures

They can help extend the range of computing towards extreme environments

THANKS

Email: ayg@anl.gov

Yanguas-Gil, IEEE Space Computing Conference, 60-66 (2019) [best paper award]
Daram, 20th International Symposium on Quality Electronic Design, 191 (2019)
Yanguas-Gil, APL Materials 7 (9), 091102 (2019)
Daram, Front. Neurosci. (2020)
Madireddy, arXiv:2007.08159 (2020)

U.S. DEPARTMENT OF ENERGY
Argonne National Laboratory is a U.S. Department of Energy laboratory managed by UChicago Argonne, LLC.

The Eigentask Framework

Aswin Raghavan* (SRI International)

DARPA ERI Summit 2021

Workshop on New Opportunities for Lifelong Learning Machines

October 20, 2021

SRI International

* aswin.raghavan@sri.com

Limitations of train-once test-anywhere ML

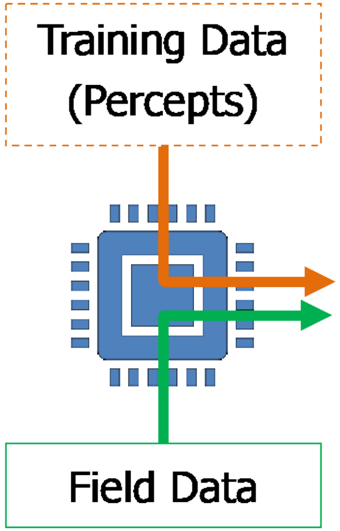

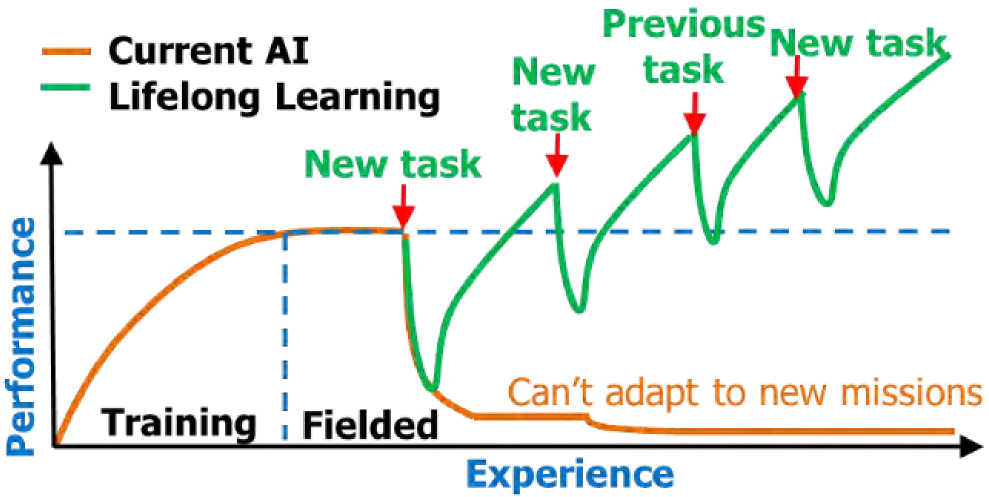

- Stagnant performance despite additional data collected in the field
- Stops working unexpectedly: data drift, sensor failure, biases in data
- Unusable for new "tasks" without cloud-side data collection and training, degradation of performance on old "tasks"

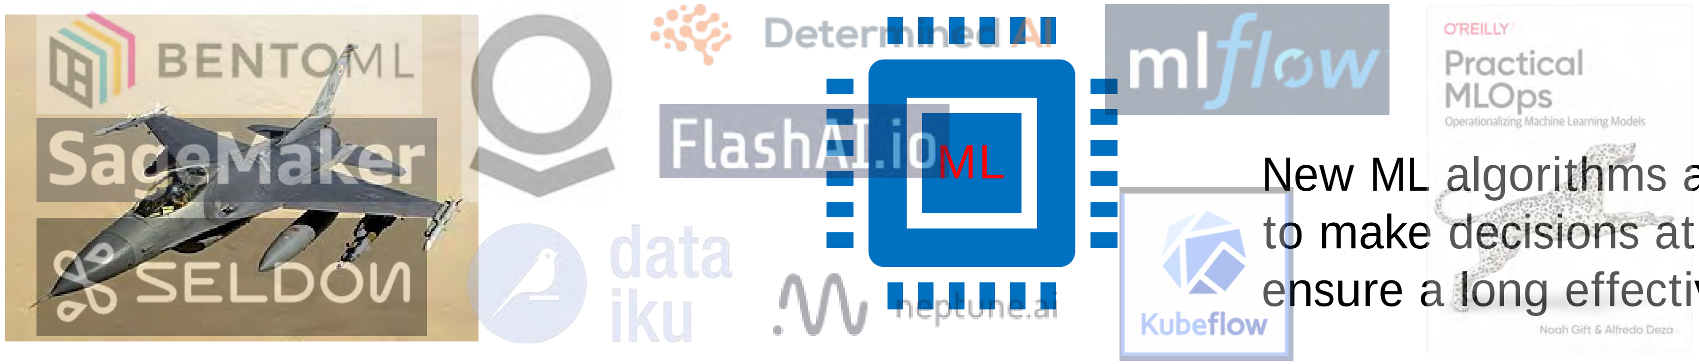

New ML algorithms are needed for AI to make decisions at the edge and ensure a long effective fielded lifetime

F-16 airframe life: 8,000 hours.
$7,000-$24,000 per hour.

Lifetime: ??? unpredictable
Very High $??? per hour.

Sculley, D. et. al. "Machine Learning: The High Interest Credit Card of Technical Debt." In SE4ML: Software Engineering for Machine Learning (NIPS 2014 Workshop), 2014.
Isbell et. al. "You Can't Escape Hyperparameters and Latent Variables: Machine Learning as a Software Engineering Enterprise", NeurIPS 2020 Invited Talk
Rosser et. al. "A Systems Perspective on Technical Debt." In 2021 IEEE Aerospace Conference (50100), 1–10, 2021. https://doi.org/10.1109/AERO50100.2021.9438359.
Rosser et. al. "Technical Debt in Hardware Systems and Elements." In 2021 IEEE Aerospace Conference (50100), 1–10, 2021. https://doi.org/10.1109/AERO50100.2021.9438332.
Tang et. al. "An Empirical Study of Refactorings and Technical Debt in Machine Learning Systems." In 2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE)

Our Focus Areas: Edge AI for C2 and ISR

- On-board learning over cloud-based training. Data collection and storage is hard. (Stealth)
- Large space of tasks dictated by adversary: active evasion and adaptation by adversary
- Evolving capabilities, platforms and armaments

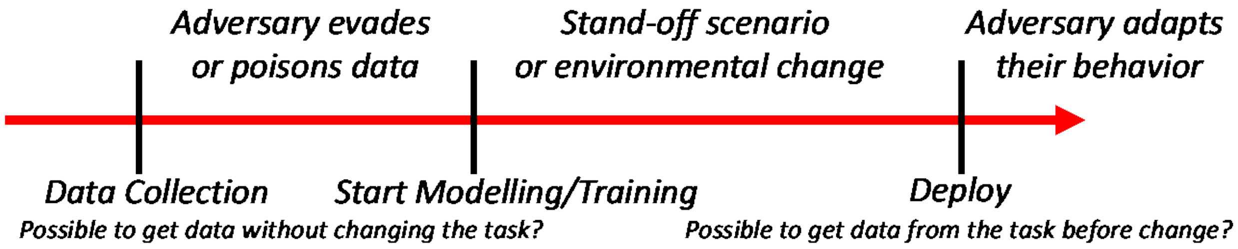

Image: OneSAF simulator

Our DARPA L2M research uses Starcraft-II for simulation of a lifetime of Reinforcement Learning (RL) tasks

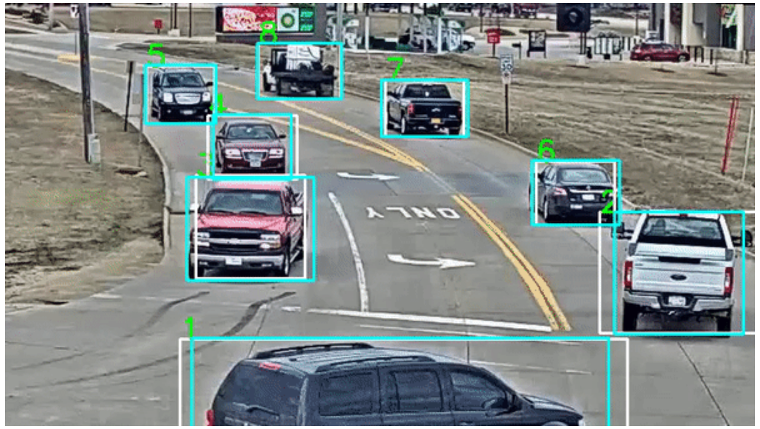

The Eigentask framework is applicable to ISR also e.g., Detect-Track-ID at the edge

SRI International

Eigentask Framework

Eigentask framework: Continuously learn a set of simple and composable skills in order to

- Wake phase: Express new tasks as a combination of eigentasks. For example:

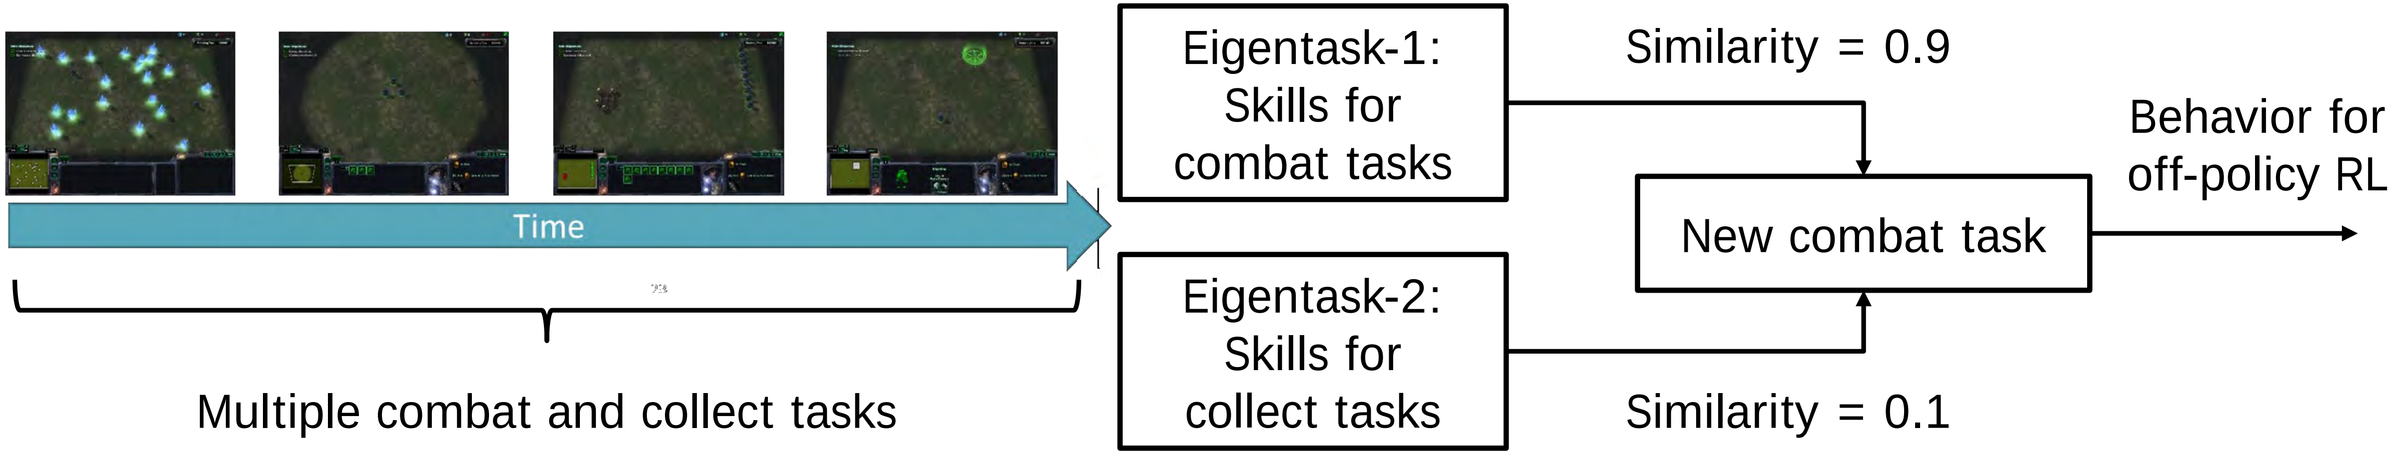

- Sleep phase: Consolidate skills by minimizing forgetting. For example:

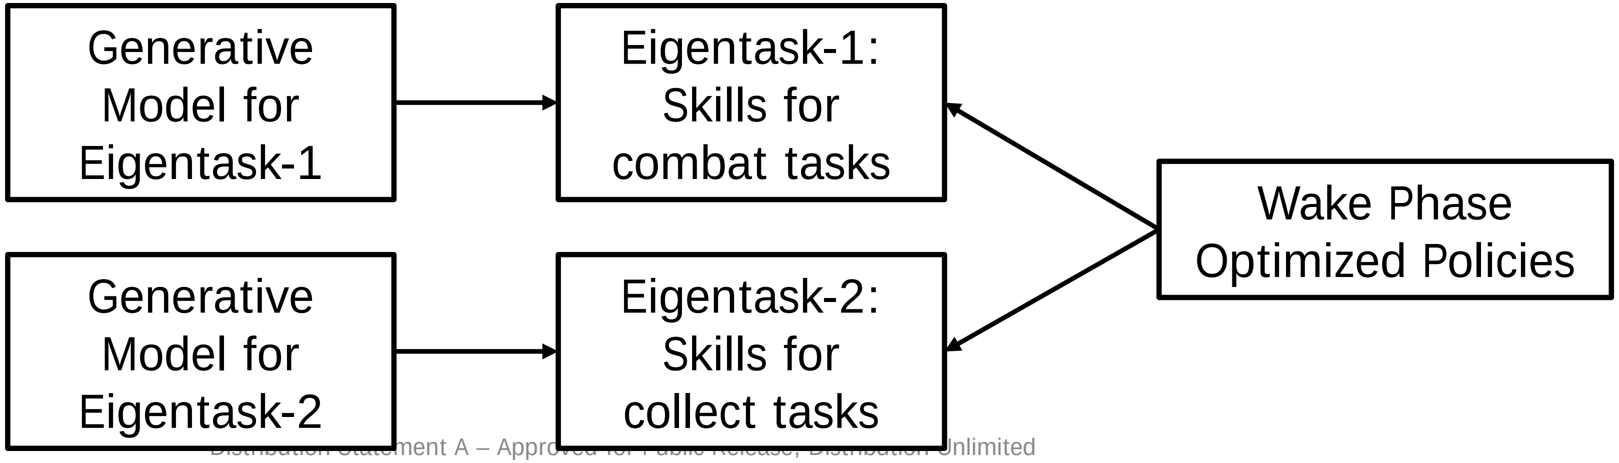

Eigentask Framework

- An Eigentask is a triplet (g, τ, d)
  - $g: \epsilon \rightarrow X$ generative model of inputs
  - $\tau: X \rightarrow [0,1]$ inputs where skill is applicable
  - $d: X \rightarrow Y$ policy/skill to execute

$$E_{x,y}[E_{\tau(x)}[L_G(x, g(\epsilon)) + L_D(y, d(y|g(\epsilon), x))]]$$

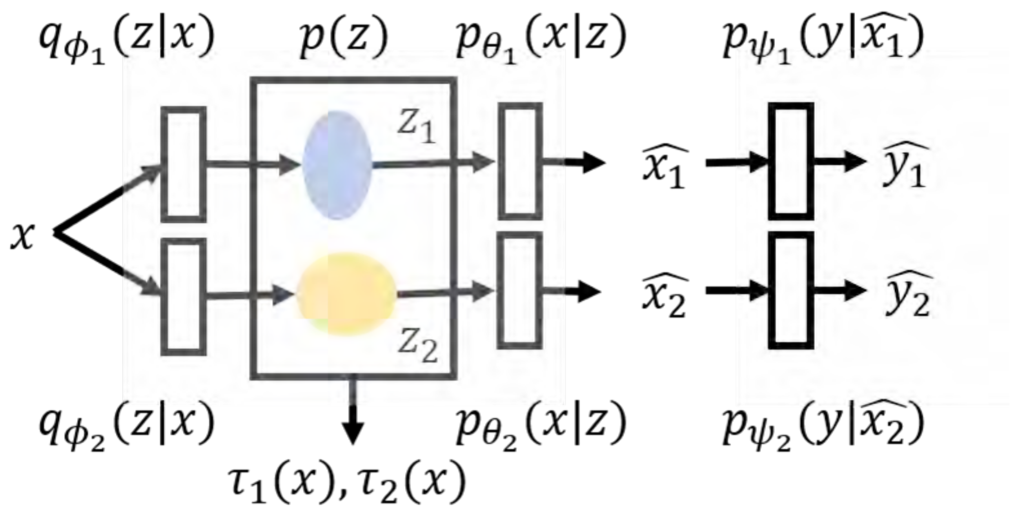

Likelihood Ratio for OOD:

$$\hat{\tau} = softmax_T(LR(x_i = x)) = softmax_T\left(\frac{\Phi(z = z_1)}{\max_{z_i} \Phi(z = z_i)}\right)$$

$$\min_{\theta,\phi,\psi} \boldsymbol{E_{x,y}}[\boldsymbol{E_\tau}[E_{q_\phi(z|x)}[\log p_\theta(\boldsymbol{x}|z)] - D_{KL}[q_\phi(z|x), p(z)] + \log p_\psi(y|z)]$$

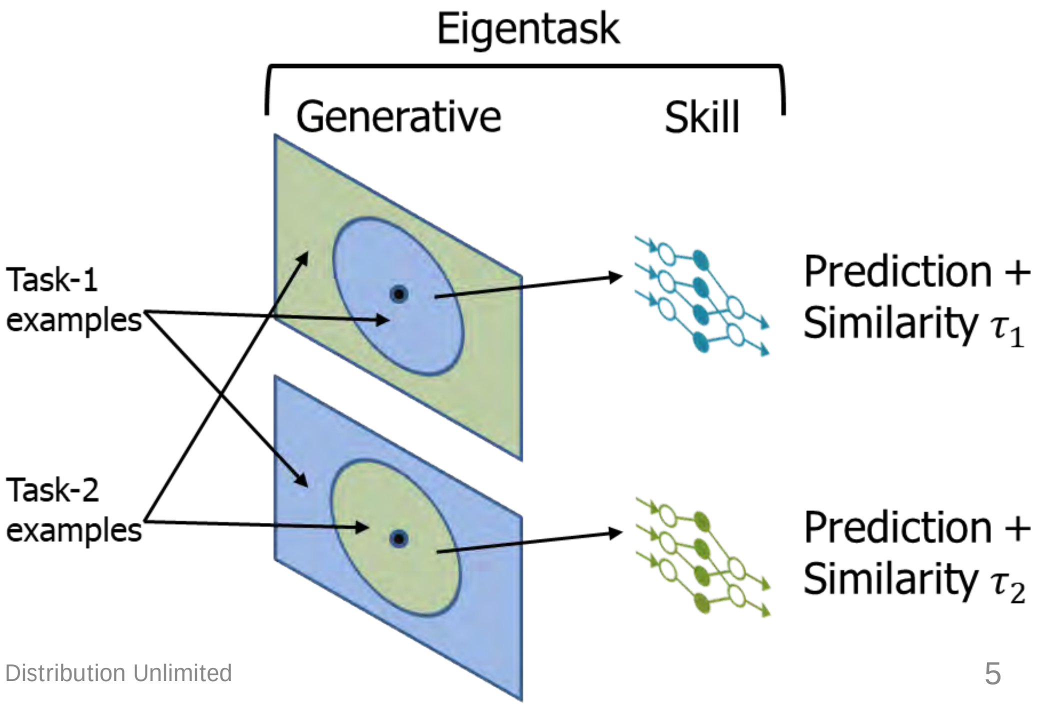

Evaluating Lifelong Learning

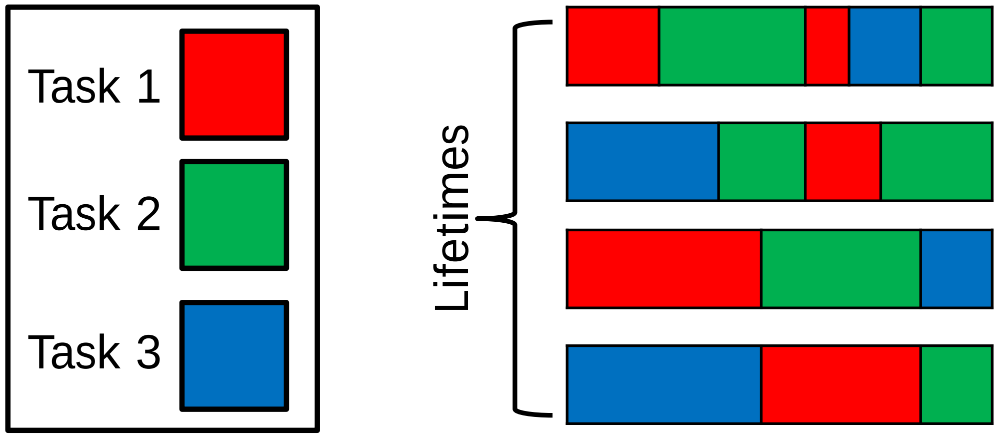

- No task ID given to learner at each time step
- Learned is not given the task boundaries
- Each task is experienced for a fixed amount of time
- Tasks share same observation and action space
- But different rewards and dynamics
- Performance is measured on ALL tasks at every task change

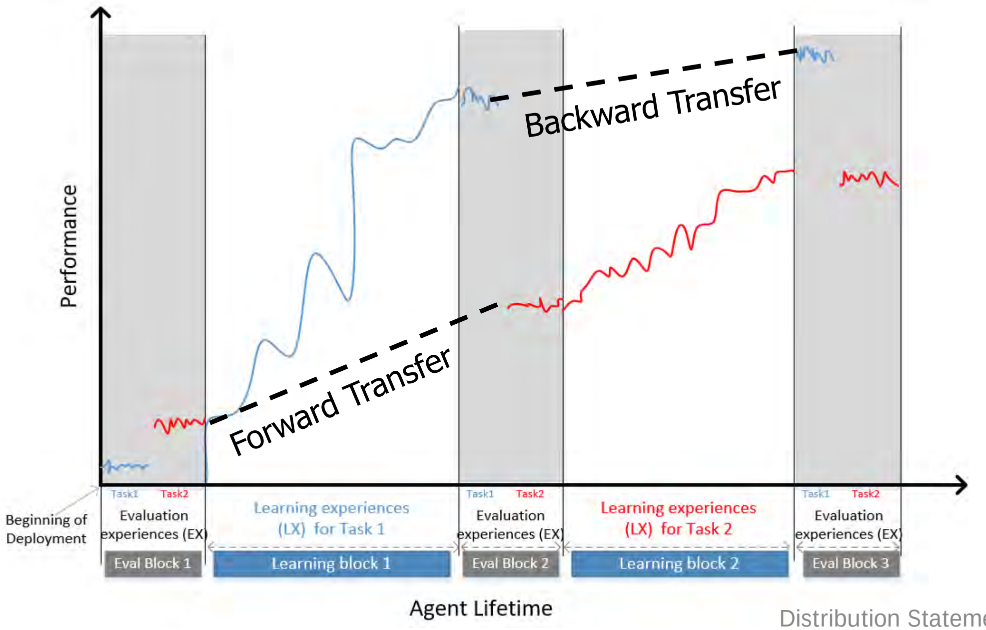

- Forward Transfer (FT): Does an L2 agent improve learning on a new task by leveraging data from previous tasks?
  - FT > 0: (GOOD) Indicates jumpstart
  - FT < 0: (BAD) Indicates interference
- Backward Transfer (BT): Does an L2 agent improve performance on a previously learned task by leveraging data from new tasks?
  - BT < 0: (BAD) Indicates forgetting
- Other metrics include performance relative to single-task expert

Benefits of Eigentasks

- Positive forward and backward transfer

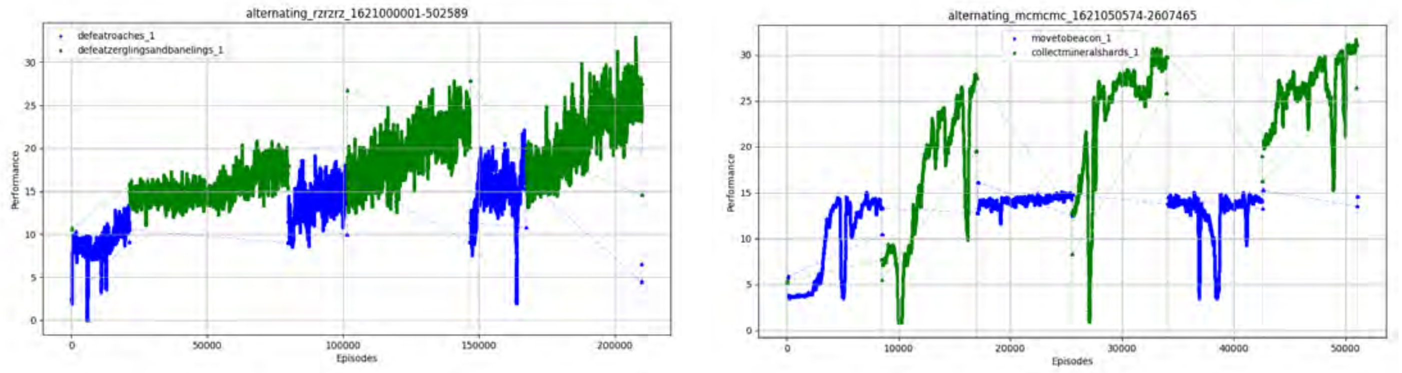

Metrics averaged over condensed scenarios			
FT	BT	Relative perf. to expert	Perf. Maint.
1.42 (0.1)	1.00 (0.03)	1.17 (0.1)	-3.05 (1.8)

- Improvement over sample complexity of RL

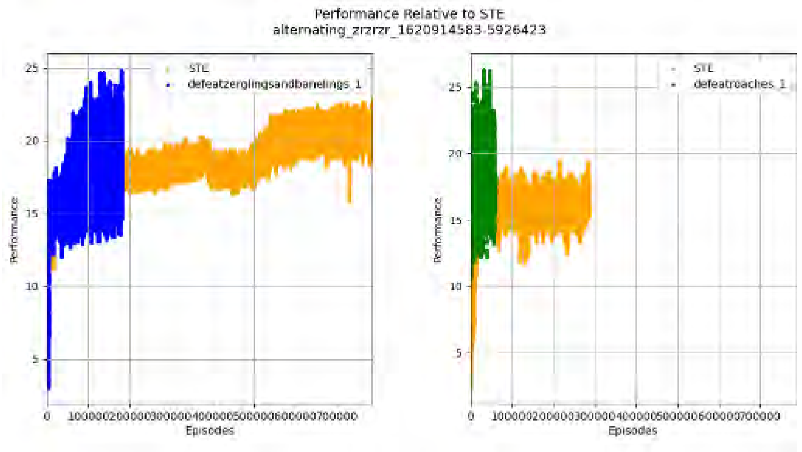

- Improvement over asymptotic single-task RL

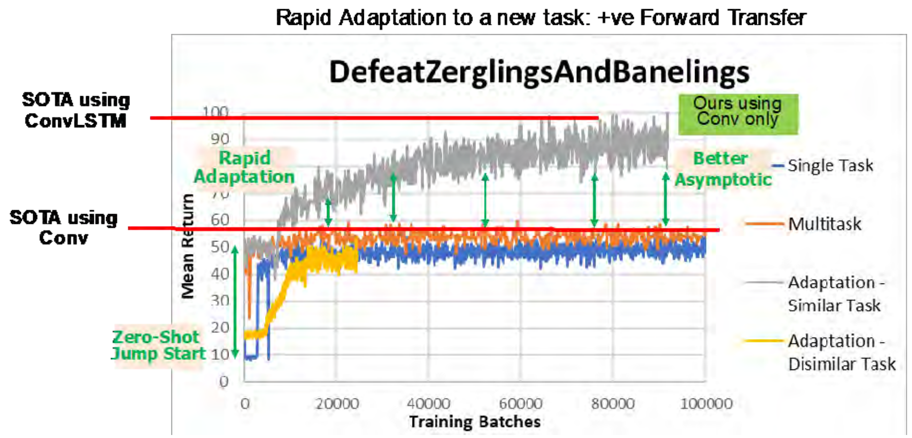

*Our L2M beats DeepMind's agent by 1.5x with >10x reduction in training, superior policy than previously known**

System Integration for Real-Time Games

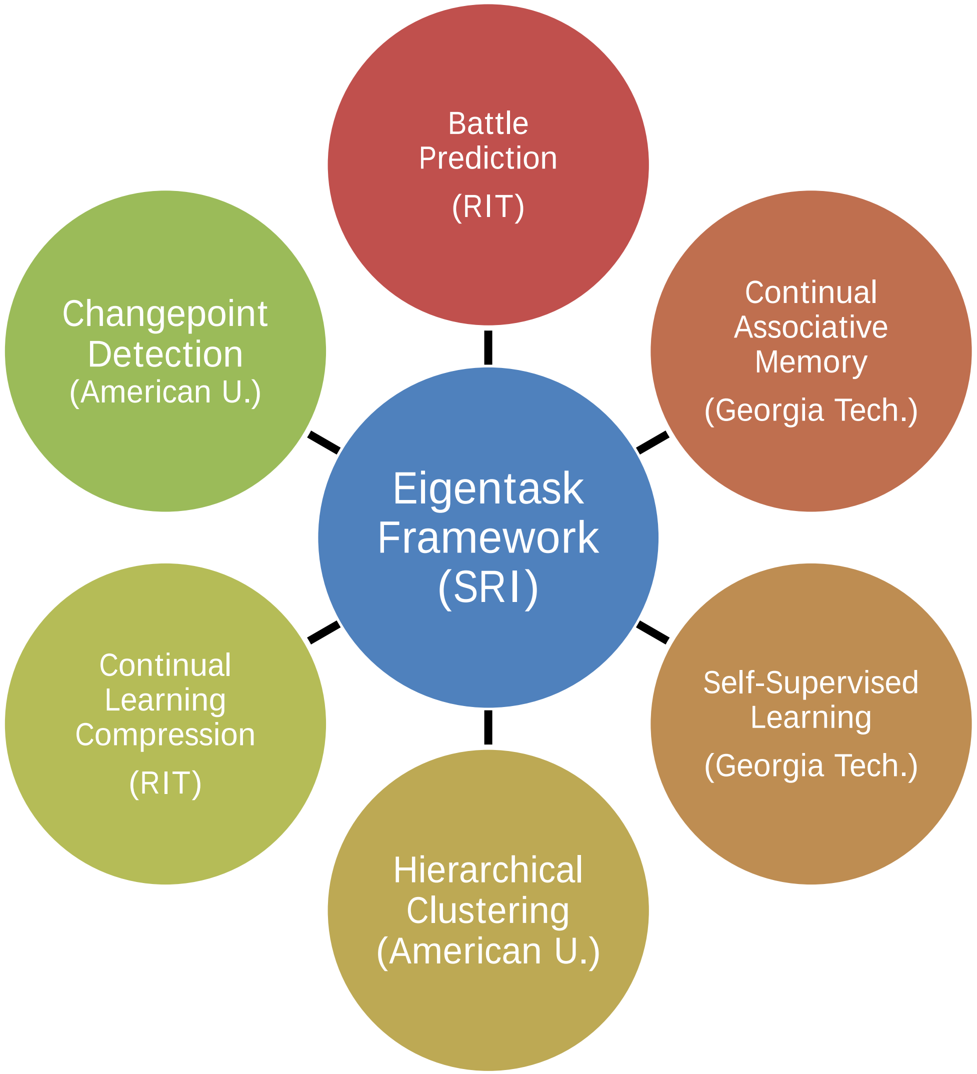

- Ongoing work under DARPA L2M Phase 2
- Initial experiments are positive in configurations with two components
- Expected to increase TRL by the end of Q2 2022
  - Expand task set to more complex tasks
  - Expand to other real-time games
  - Expand to extend to C2 simulators like OneSAF

SRI International

Application to Computer Vision for ISR

- We tested Eigentasks on “mixed datasets”
- Eigentasks naturally learns separate skills corresponding to the different datasets
- Image classification accuracy is higher than SOTA incremental class learning

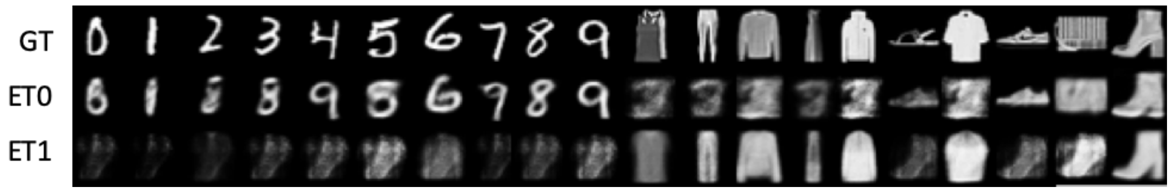

Model reconstructions showing Eigentasks have specialized to different datasets

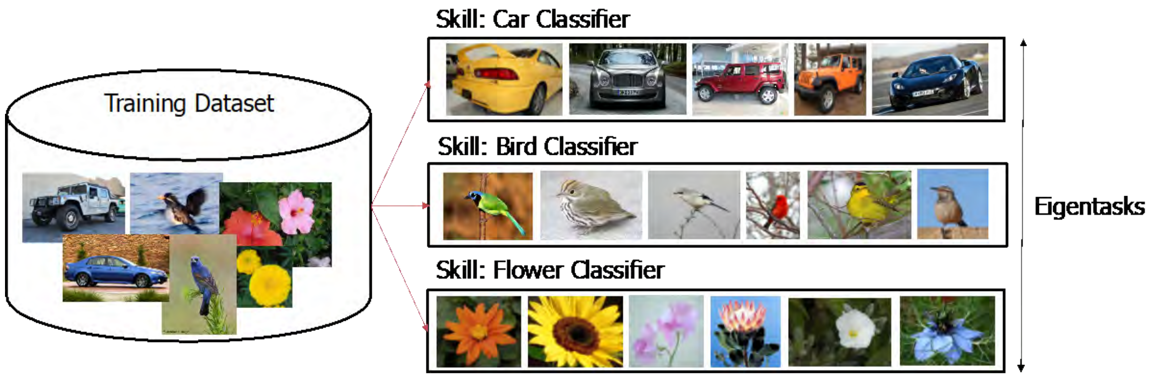

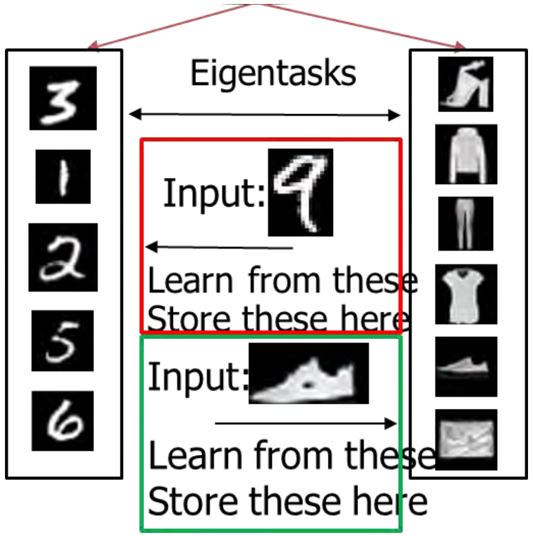

Approach	Method	D1	D2
Baselines	None - lower bound	19.90	10.22
	Offline - upper bound	97.94	90.89
Regularization	EWC	20.01	10.00
	Online EWC	19.96	10.00
	SI	19.99	10.00
Replay	LwF	23.85	10.07
	DGR	90.79	73.36
	DGR x2	91.83	65.82
	DGR+distill	91.79	72.40
	DGR+distill x2	94.01	67.37
	RtF	92.56	61.15
	RtF x2	92.86	61.41
Replay+Exemplars	iCaRL	94.57	82.69
Replay+Eigentask	ET1-BaseAug	87.68	69.29
	ET1-BAug	90.99	74.11
	ET1-VAug	87.33	63.34
	ET1-VBAug	90.69	77.43
	ET2-BaseAug	88.93	57.91
	ET2-BAug	**91.27**	69.95
	ET2-VAug	82.08	69.55
	ET2-VBAug	90.25	**76.81**

Table 1. Average test accuracy over all tasks on splitMNIST (D1) and split(MNIST+FashionMNIST) (D2) benchmarks. ET1 and ET2 denote the number of eigentasks in an OWVAE model. Meth-

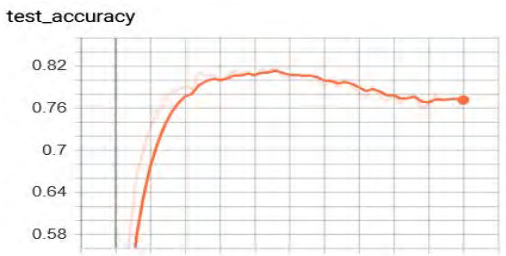

Mixed-Dataset Natural Image Classification With Eigentasks:
Task Separation w/out task label
Performance
comparable to SOA

SRI International

Summary

- The Eigentask framework is a powerful and promising framework for building lifelong learning systems
  - Learns fragments of reusable knowledge and composable skills
- Allows RL policies, classifiers, object detectors etc. to coexist and be learned jointly
  - Better curriculum learning for complex problems (e.g., to play full game of SC2 quicker and better)
- Biologically inspired by mixture-of-experts and sleep in mammalian brain
- Enables long-lived AI at the edge: without any data storage, transmission and cloud-side training
- Actively seeking transition partners in ISR and C2 applications

Overcoming Catastrophic Forgetting with Meta-Learned Neuromodulation

University of Wyoming, University of Vermont

PIs: Nick Cheney, Jeff Clune

Problem: Neural Networks Catastrophically Forget by Overwriting

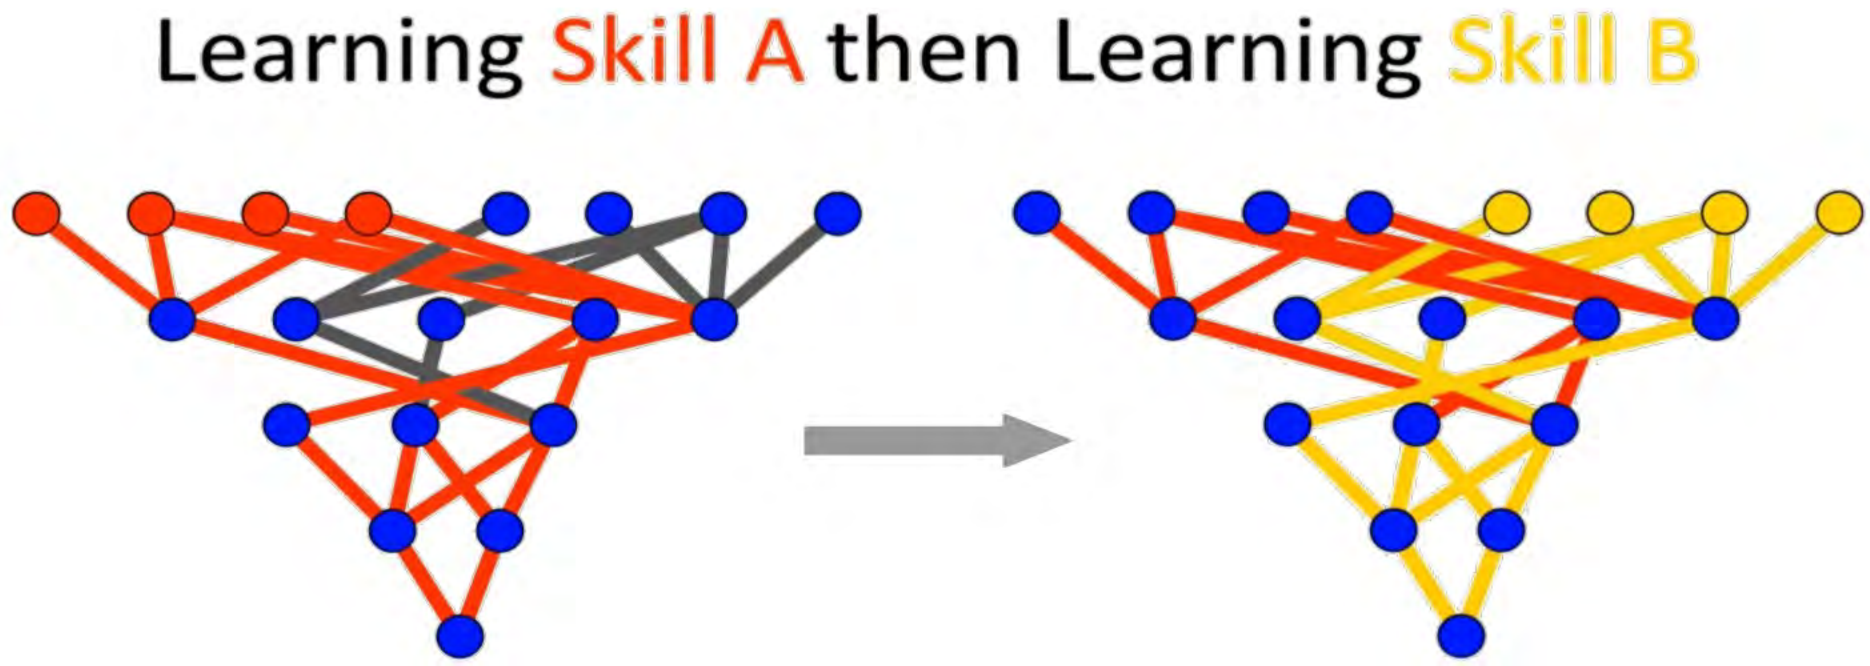

(Ellefsen et al., 2015)

ML Inspiration: Trend from hand-designed to learned solutions

  → 

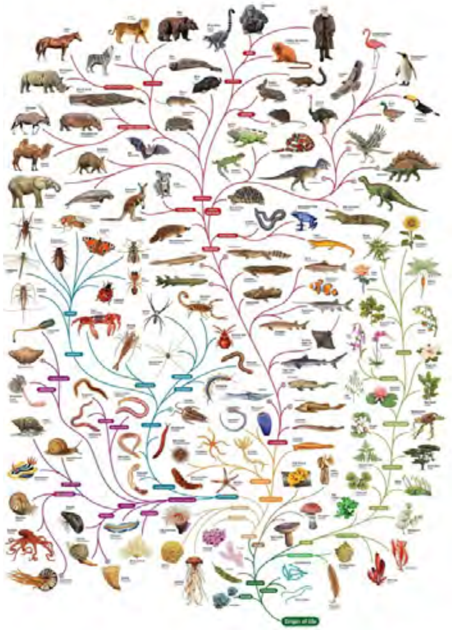

Features:	HOG/SIFT/SURF	→	Deep Learning
Architectures:	Hand Designed	→	Architecture Search
Hyperparams:	Manually Tuned	→	Adaptive/Learned
(RL) Algorithms:	Hand Designed	→	Meta-learned

(Clune, 2019)

Approach: A Neuromodulated Meta-Learning (ANML) algorithm

1) Create a Prediction Learning Network (PLN), which is a deep neural network to solve a particular task (here: classification of a handwritten character or photo)

2) Create a Neuromodulatory Network (NM), which takes in the same input as the PLN, and decides which nodes (or synapses) of the PLN are to be used for this prediction (i.e. context-dependent gating)

3) After making a sequences of predictions (and updates to) the PLN, update the strategy of the NM such that prior learned tasks are not degraded when learning new ones (i.e. meta-learning via backpropogation through time)*

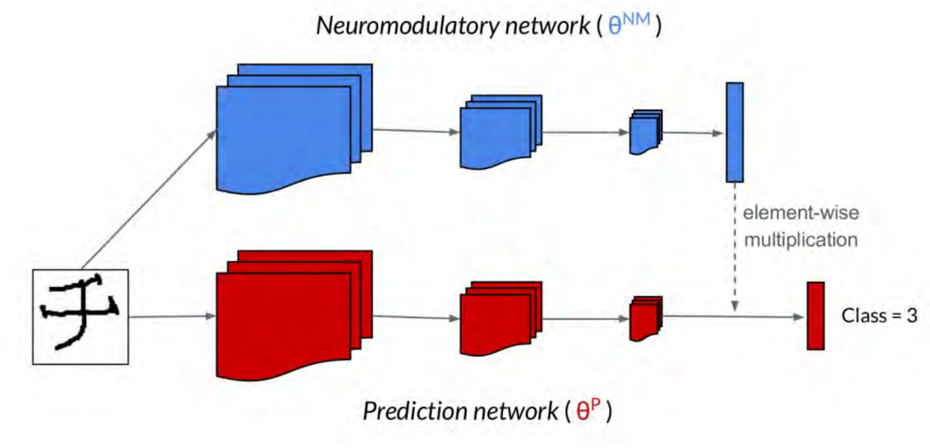

* Note: this is setting the loss/fitness/reward function of the entire system to be our ultimate goal of reducing catastrophic forgetting (or for some other goal in future work)!

(Beaulieu et al., 2020)

Results: Avoiding forgetting on long (held-out) sequences

ANML achieves 90% of max (non-sequential) performance despite sequential learning

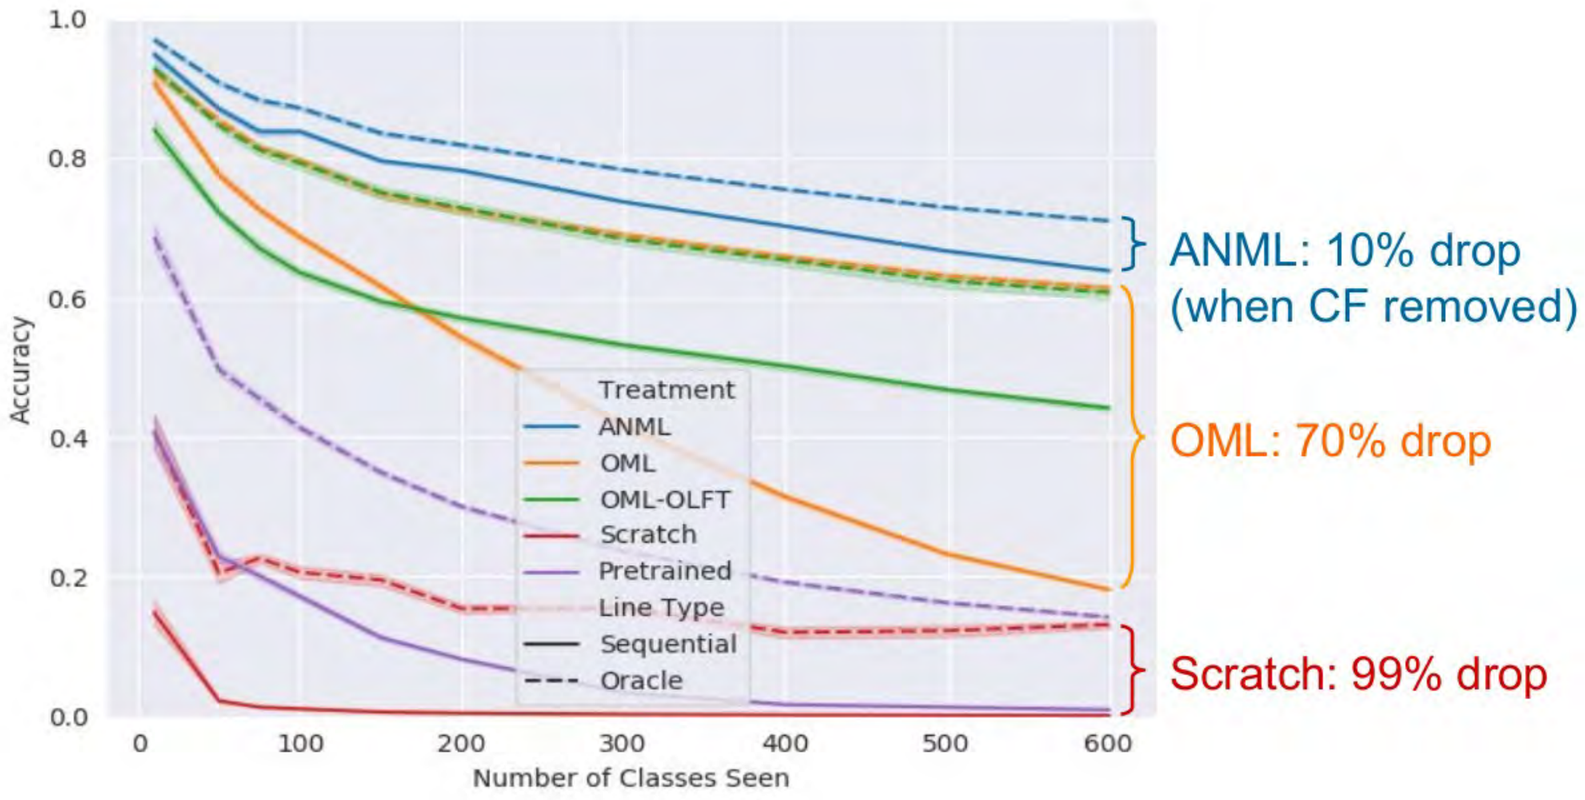

(Beaulieu et al., 2020)

Questions, Limitations, and/or Opportunities:

Out-of-distribution data (domain shift)

(Omniglot meta-training → Imagenet meta-testing)

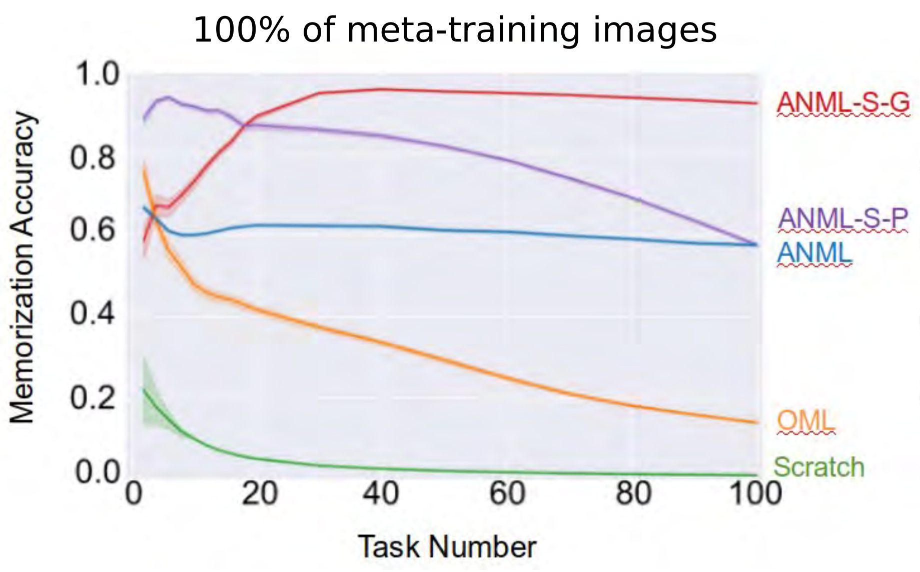

Resilience to limited training data

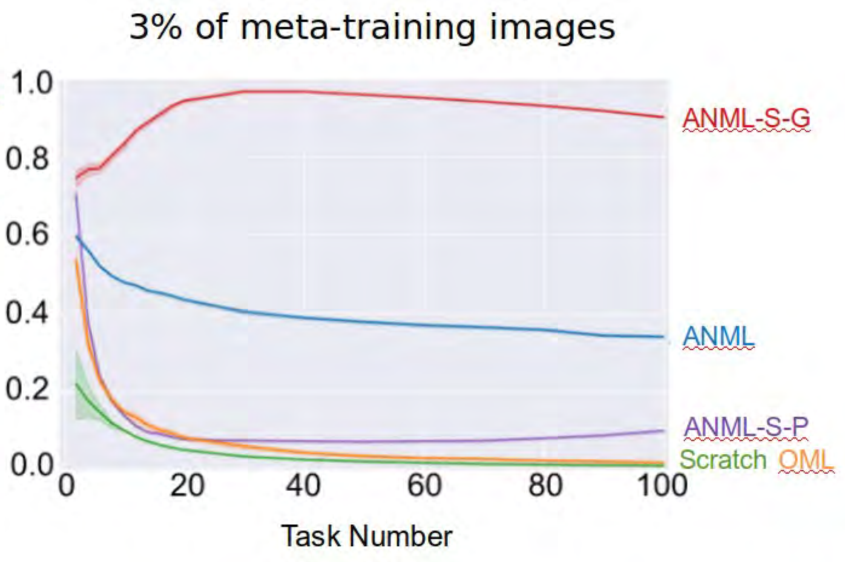

Improvement in non-continual learning

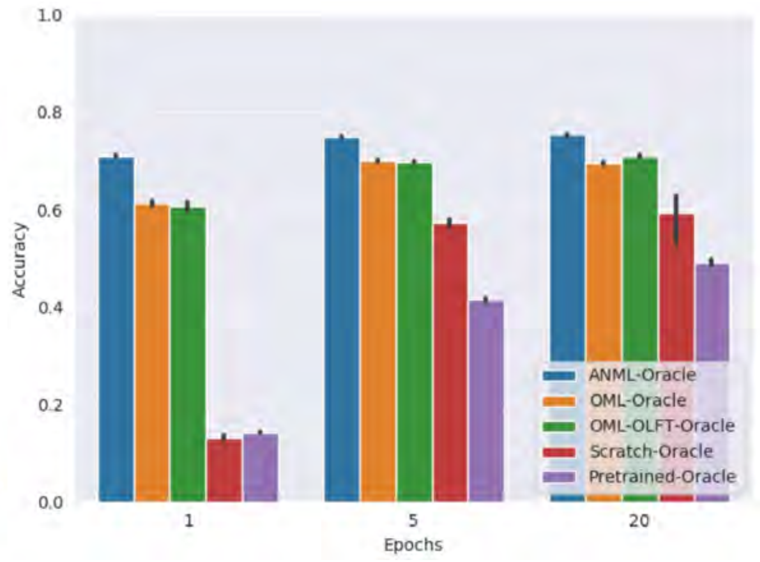

(Beaulieu et al., 2021; Frati et al., in prep; Beaulieu et al., 2020)

Potential DoD/Industry Transition Partners: Actively Underway

U.S. ARMY ENGINEER RESEARCH AND DEVELOPMENT CENTER
ERDC

CRREL
COLD REGIONS RESEARCH AND ENGINEERING LABORATORY

OpenAI

Thank you!

reach us at: ncheney@uvm.edu

Thanks to our team/collaborators/partners:

Jeff Clune
Shawn Beaulieu
Lapo Frati
Ken Stanley
Joel Lehman
Thomas Miconi

OpenAI

Uber AI Labs

University of
British Columbia

“A NeuRoBot That Learns Locomotion Online”

L2M PI:
Francisco J. Valero-Cuevas, PhD
L2M CoPI:
Alice Parker, PhD

Team:
University of Southern California (USC)

USC Viterbi School of Engineering

University of Southern California

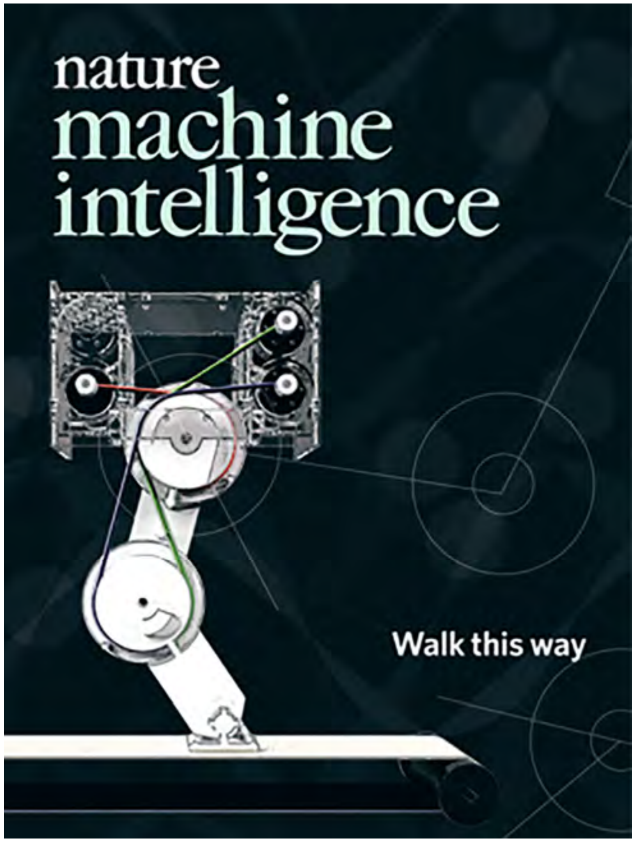

Marjaninejad et. al.
Nature Machine Intelligence (2019)

The G2P Algorithm

Model-agnostic autonomous control of tendon-driven systems

…Capable of using “motor babbling” to learn locomotor actions based on limited experience

Implementation of G2P on a quadruped system (in simulation)

Expand the algorithm to control 4 legs

Test the on-the-go adaptability for different tasks (in air, on ground, with payload…)

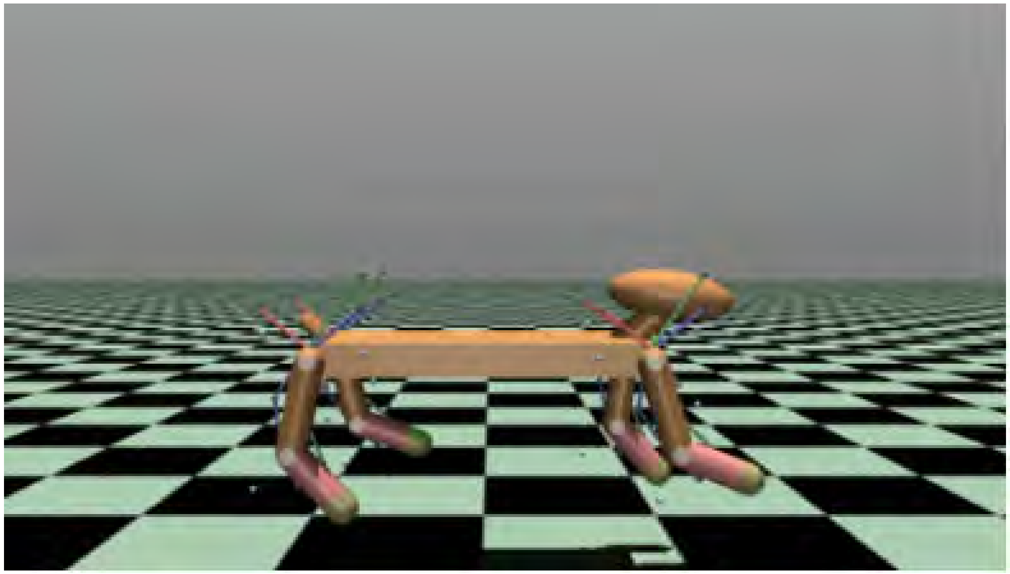

The quadruped system

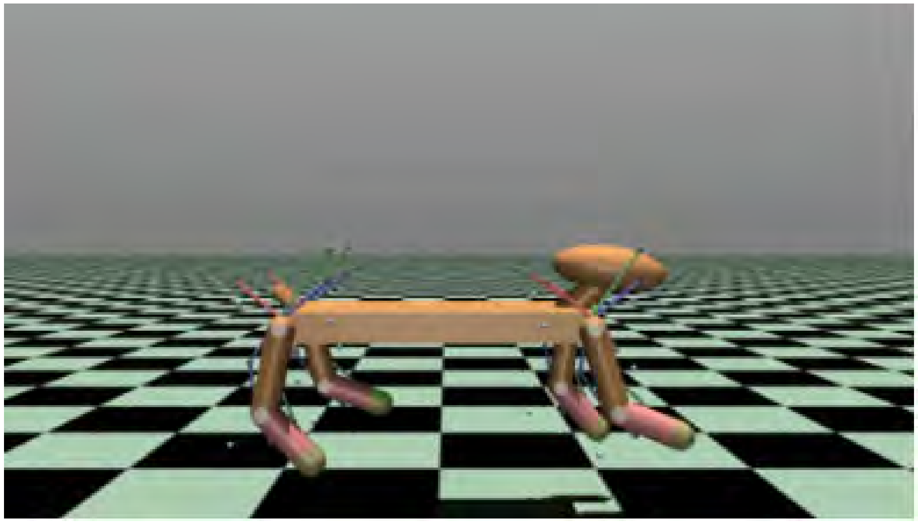

- Four limbs (N+1 tendon design)
  - 3 tendons
  - 2 DoFs
- Tactile force sensors
  - 1 contact force sensor per "paw"
- 0.5 meter length
  - Body, ρ = 1000Kg/m^3

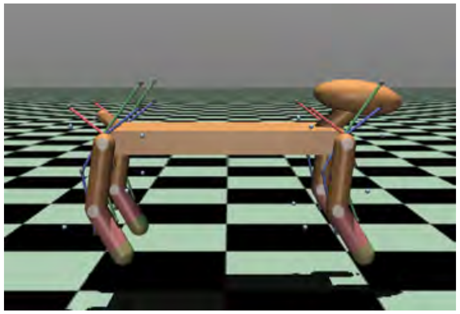

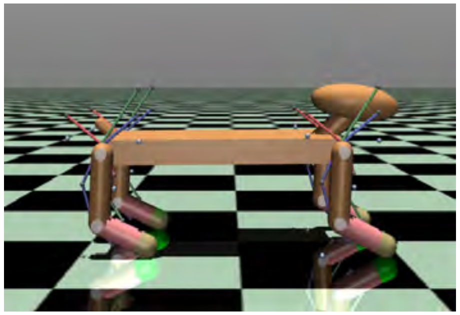

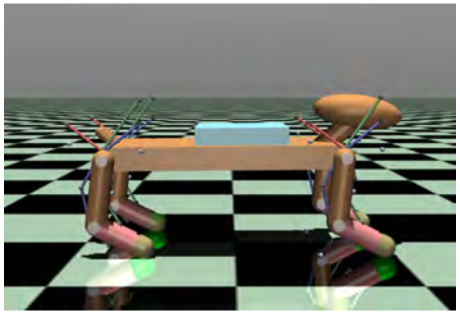

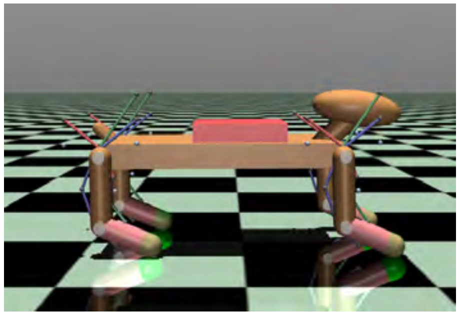

The learning curriculum

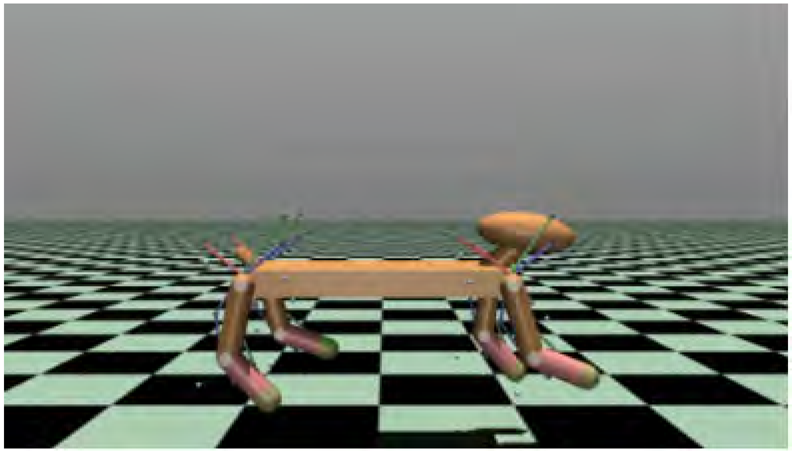

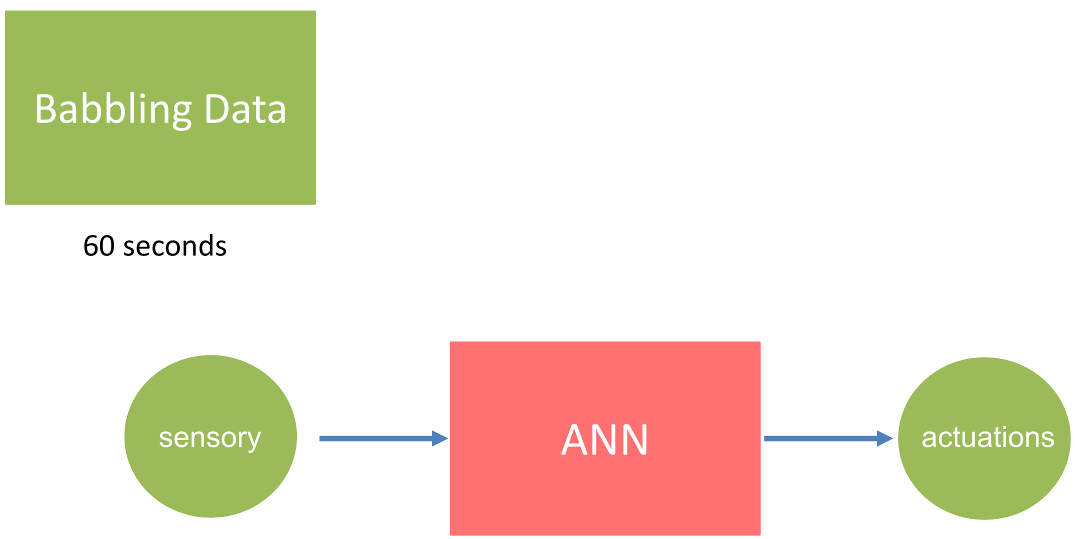

The learning curriculum

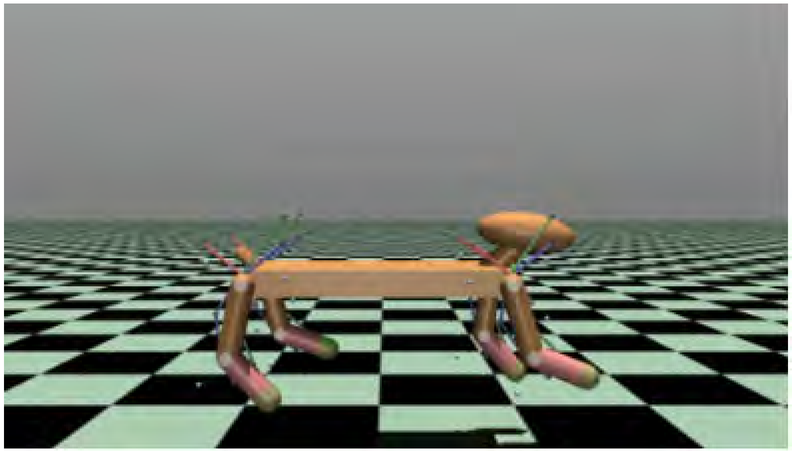

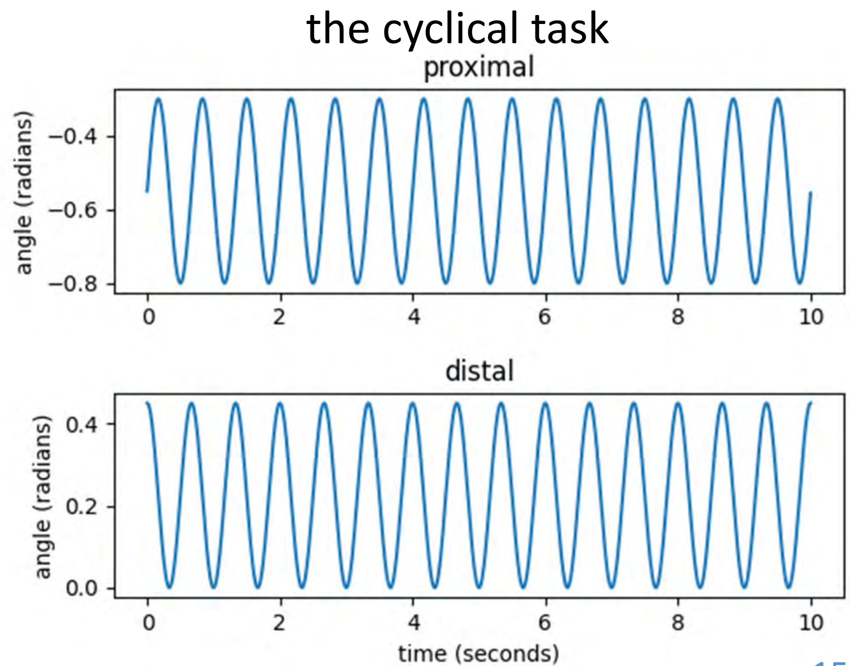

ω : 15
ϕ : $\pi/2$

Refinements

The learning curriculum

Babbling

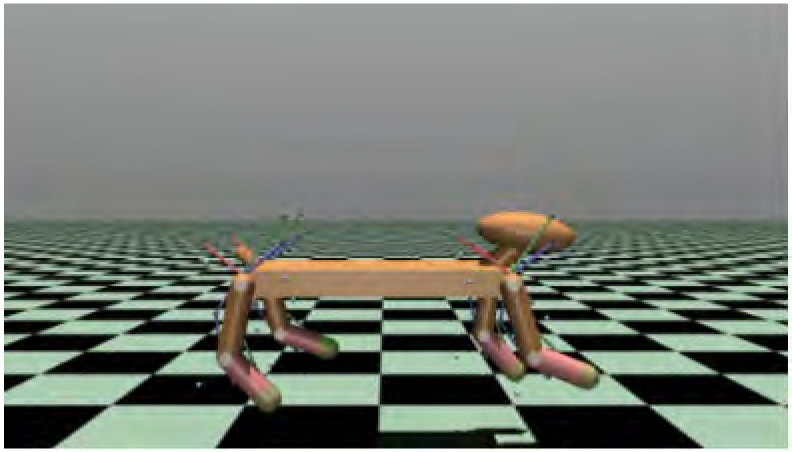

Babbling Data

60 seconds

Task specific refinement

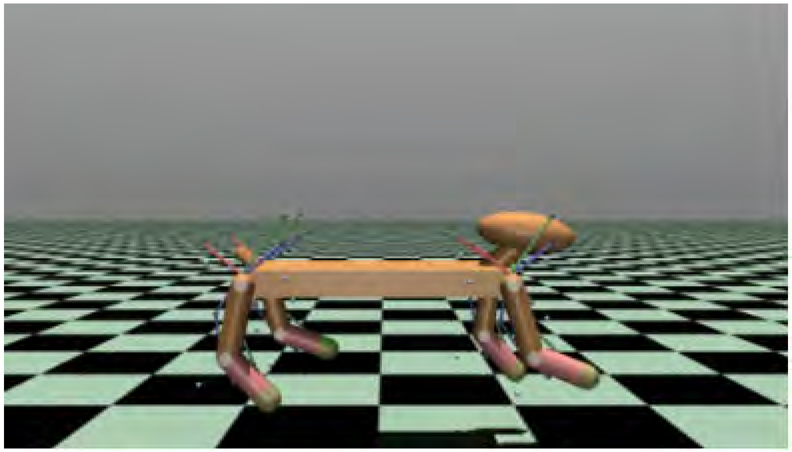

10 seconds

sensory → ANN → actuations

The learning curriculum with 2 mins of experience

Babbling

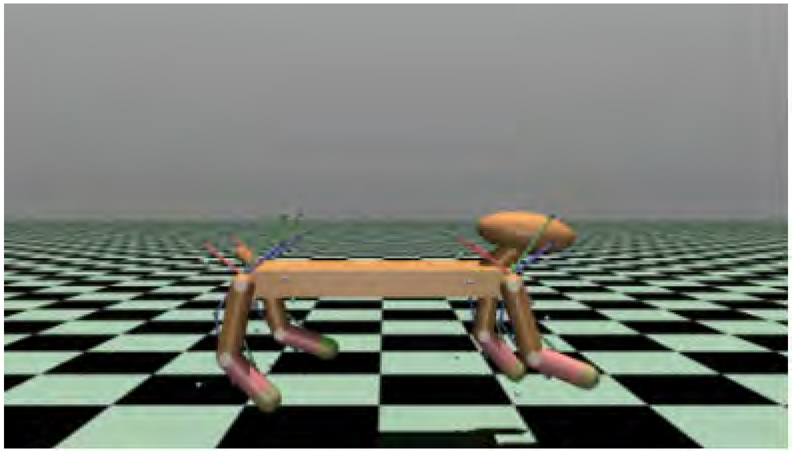

Babbling Data

60 seconds

Task specific refinement

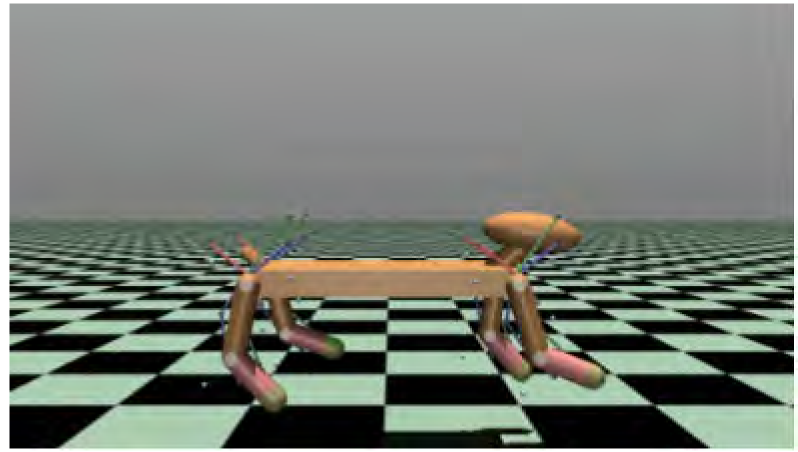

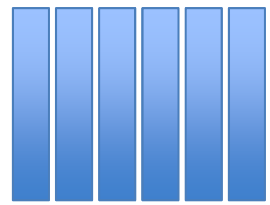

60 seconds

sensory → ANN → actuations

The learning curriculum

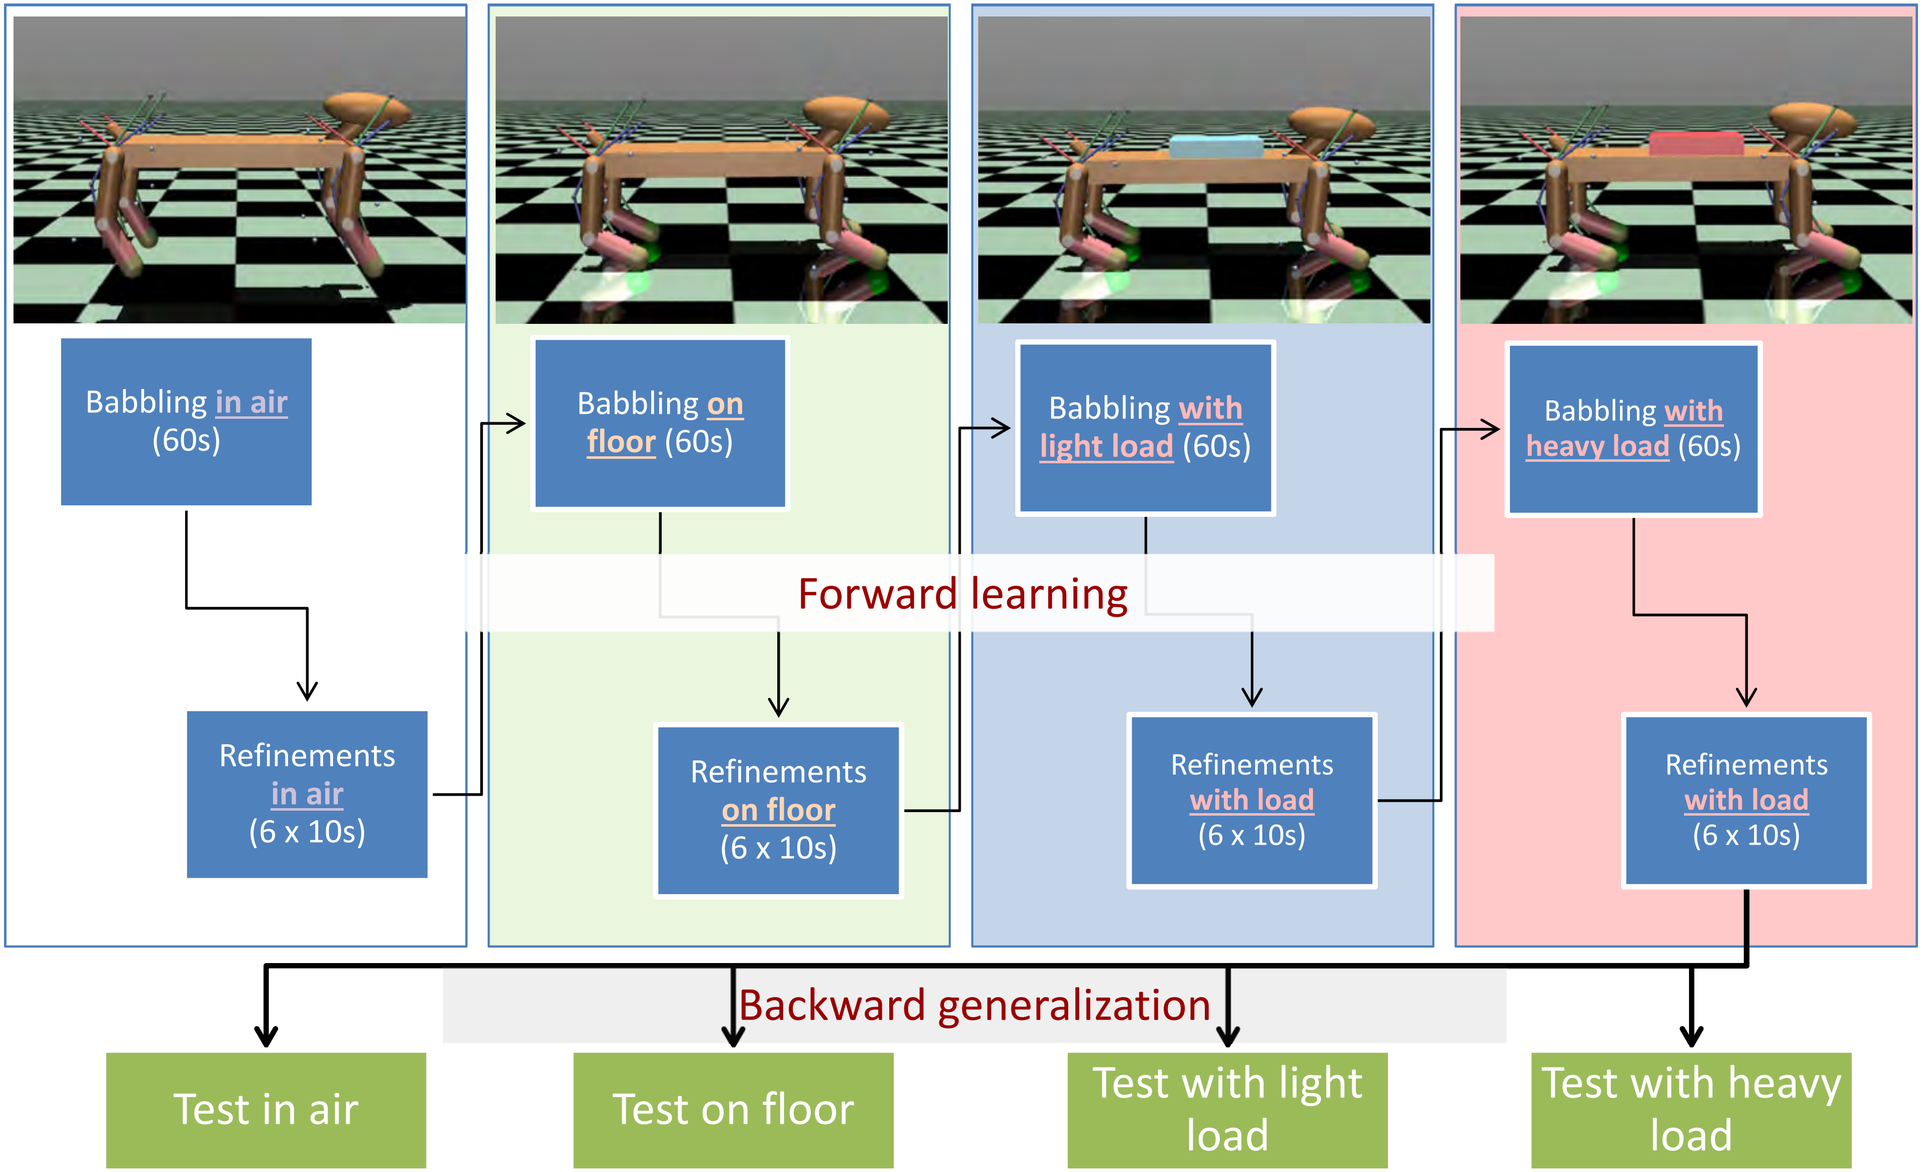

Extension 1: Studied the effects of

1. ANN structure: One for the quadruped vs. one per leg

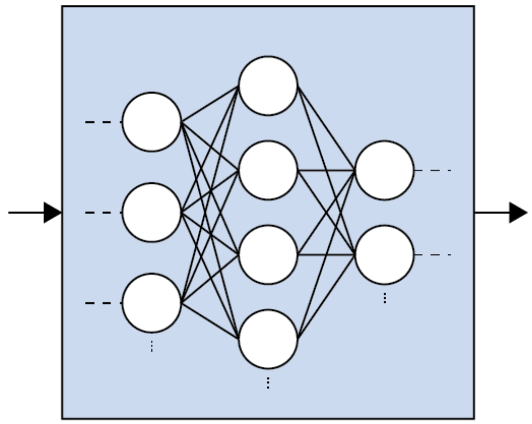

VS

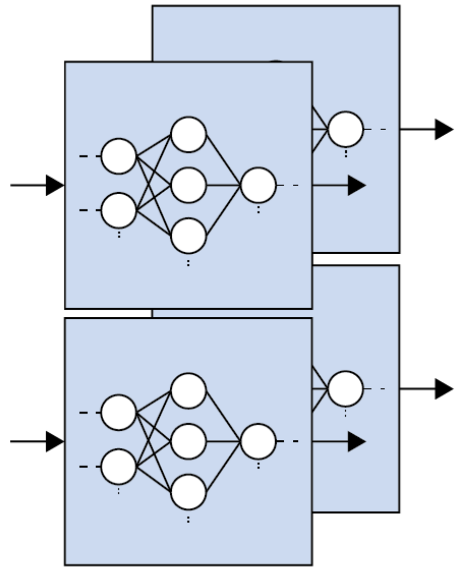

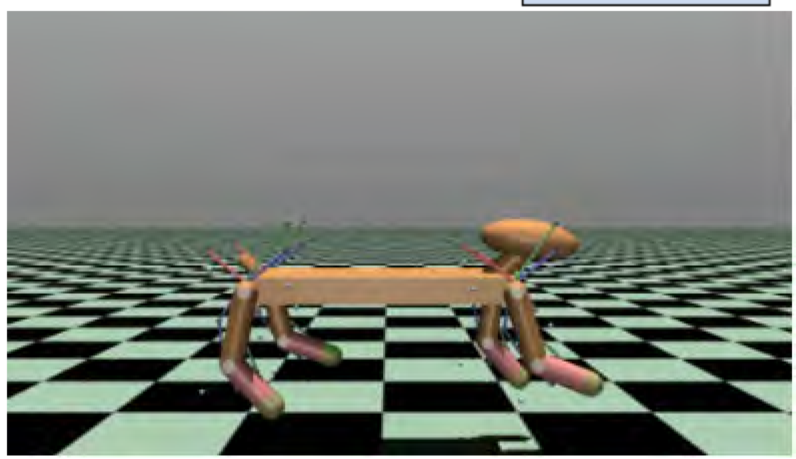

Studied the effects of:

1. ANN structure

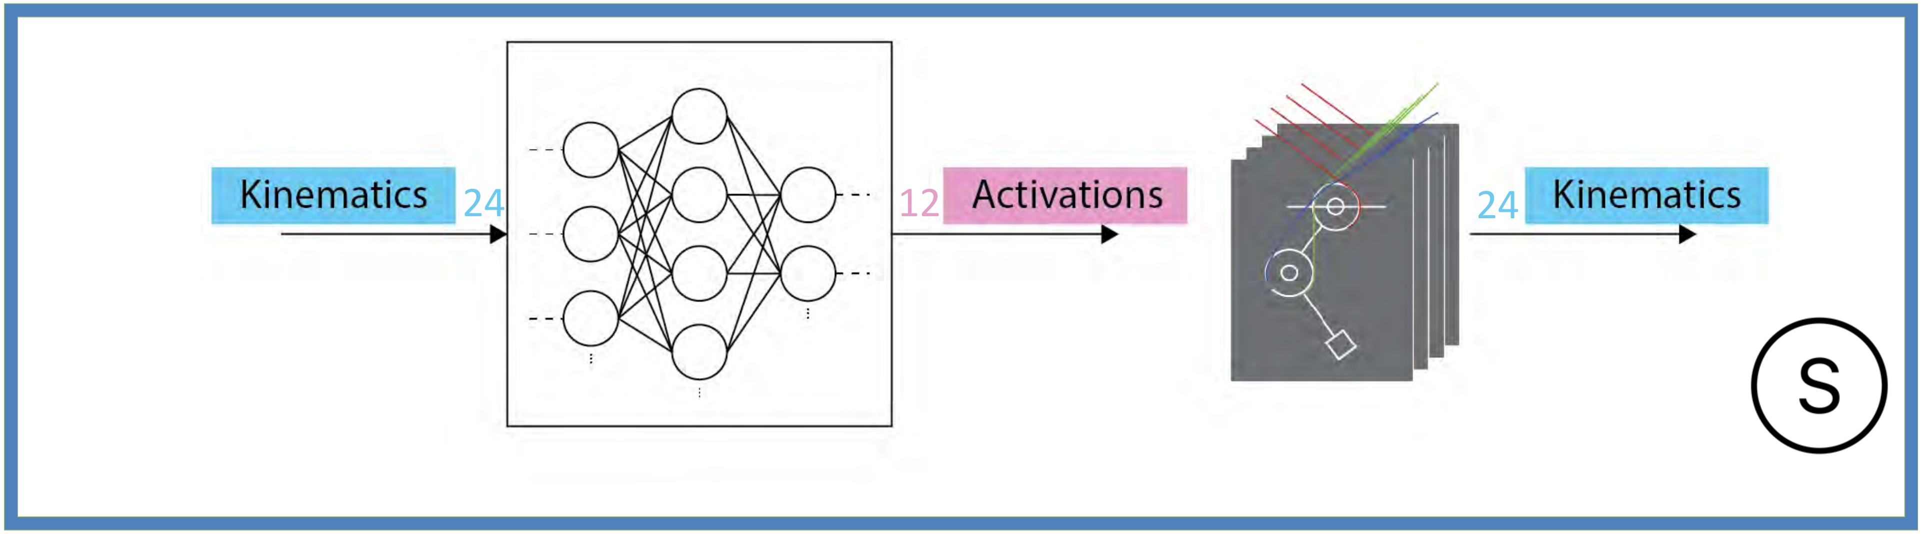

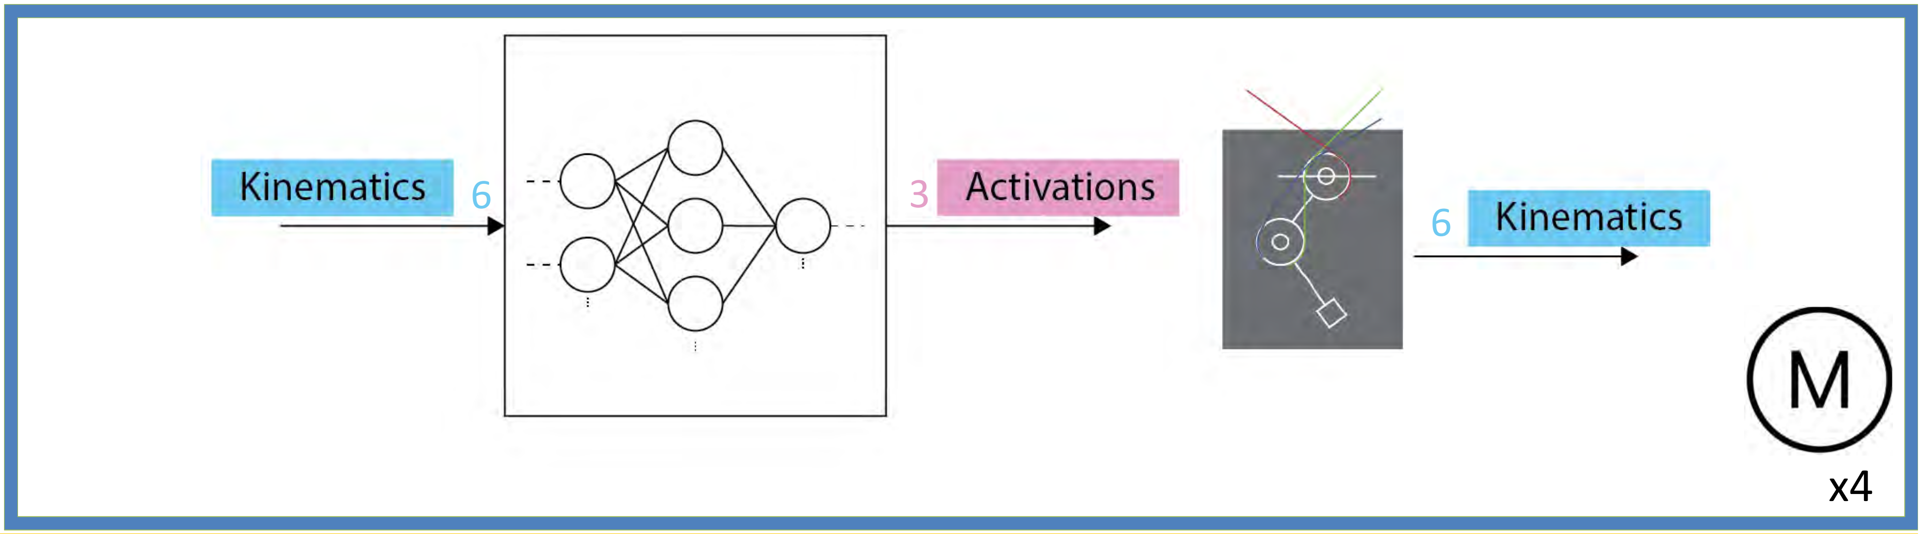

Extension 2: Studied the effects of

1. ANN structure
2. tactile sensing of normal contact force

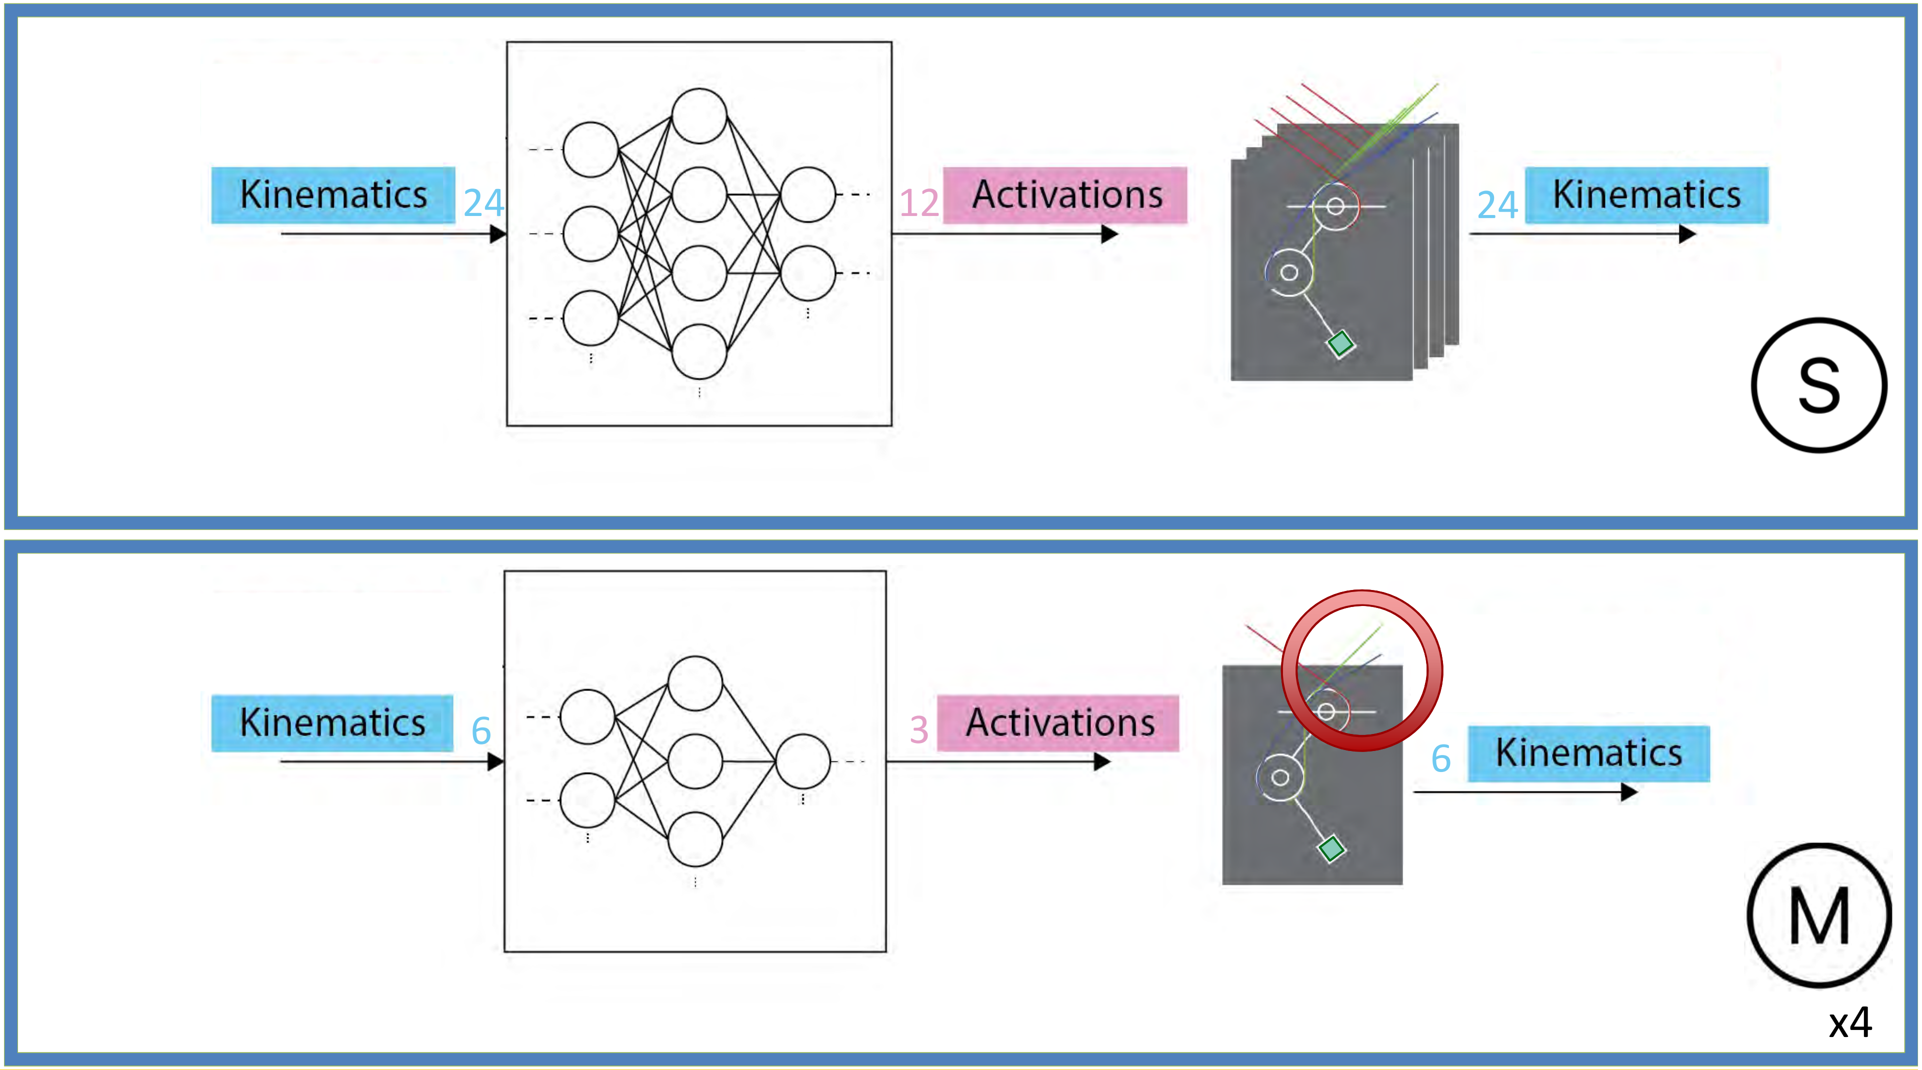

Extension 3: Studied the effects of

1. ANN structure
2. tactile force
3. kinematic feedback

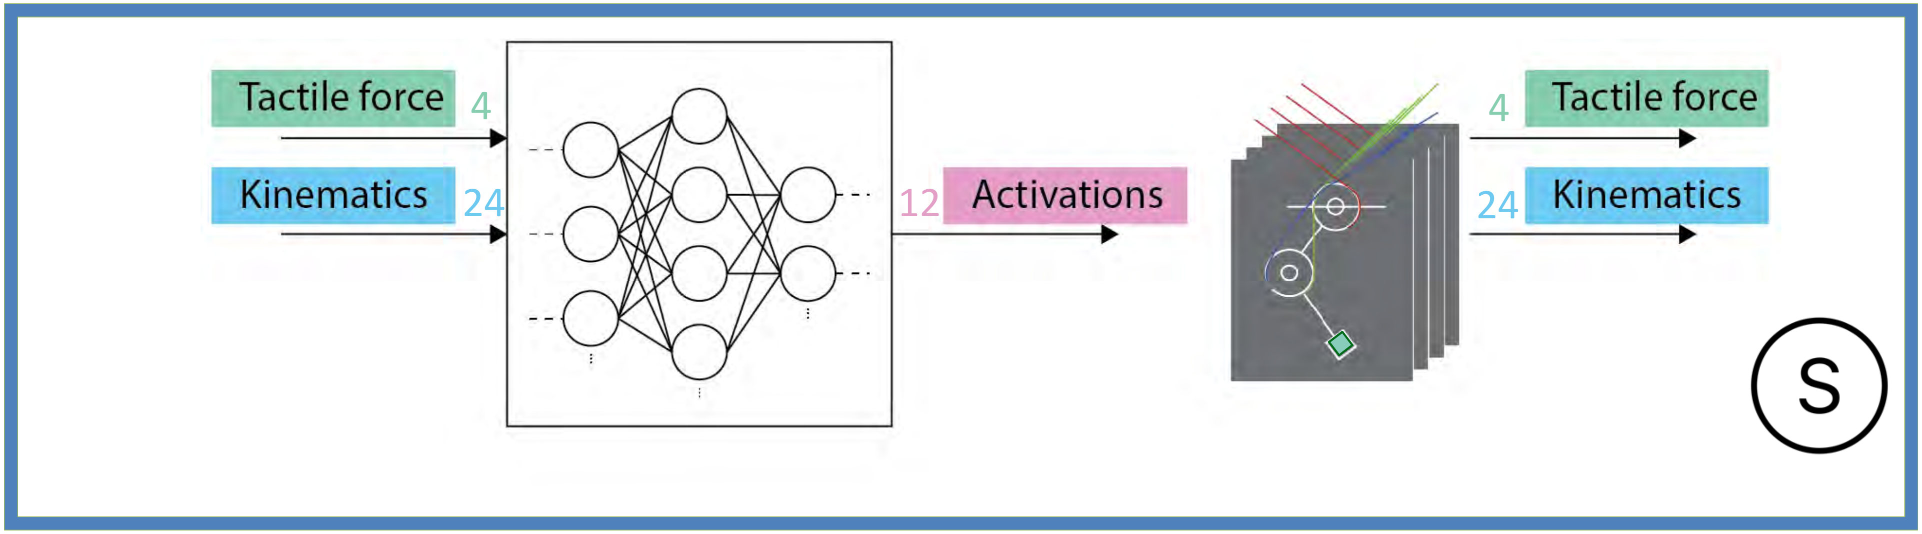

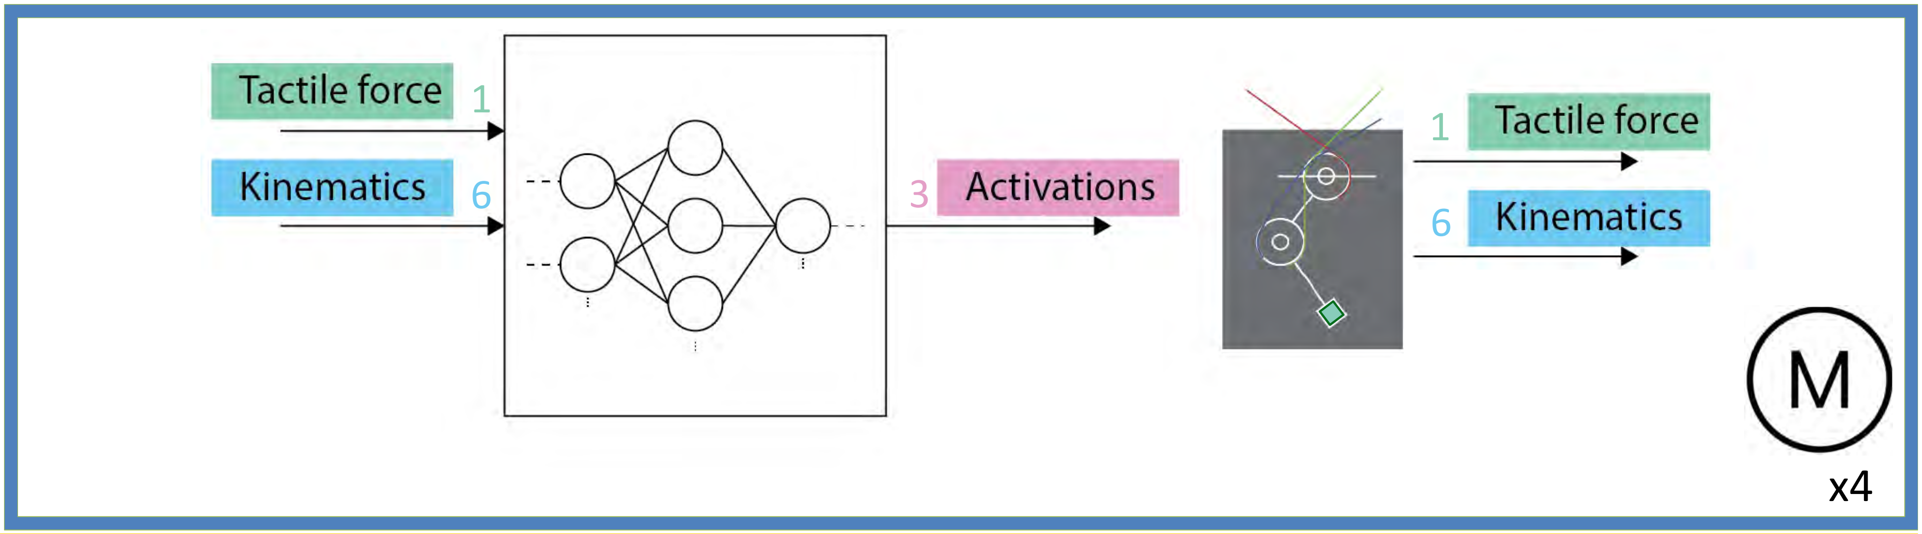

Extension 3: Studied the effects of

1. ANN structure
2. tactile force
3. kinematic feedback

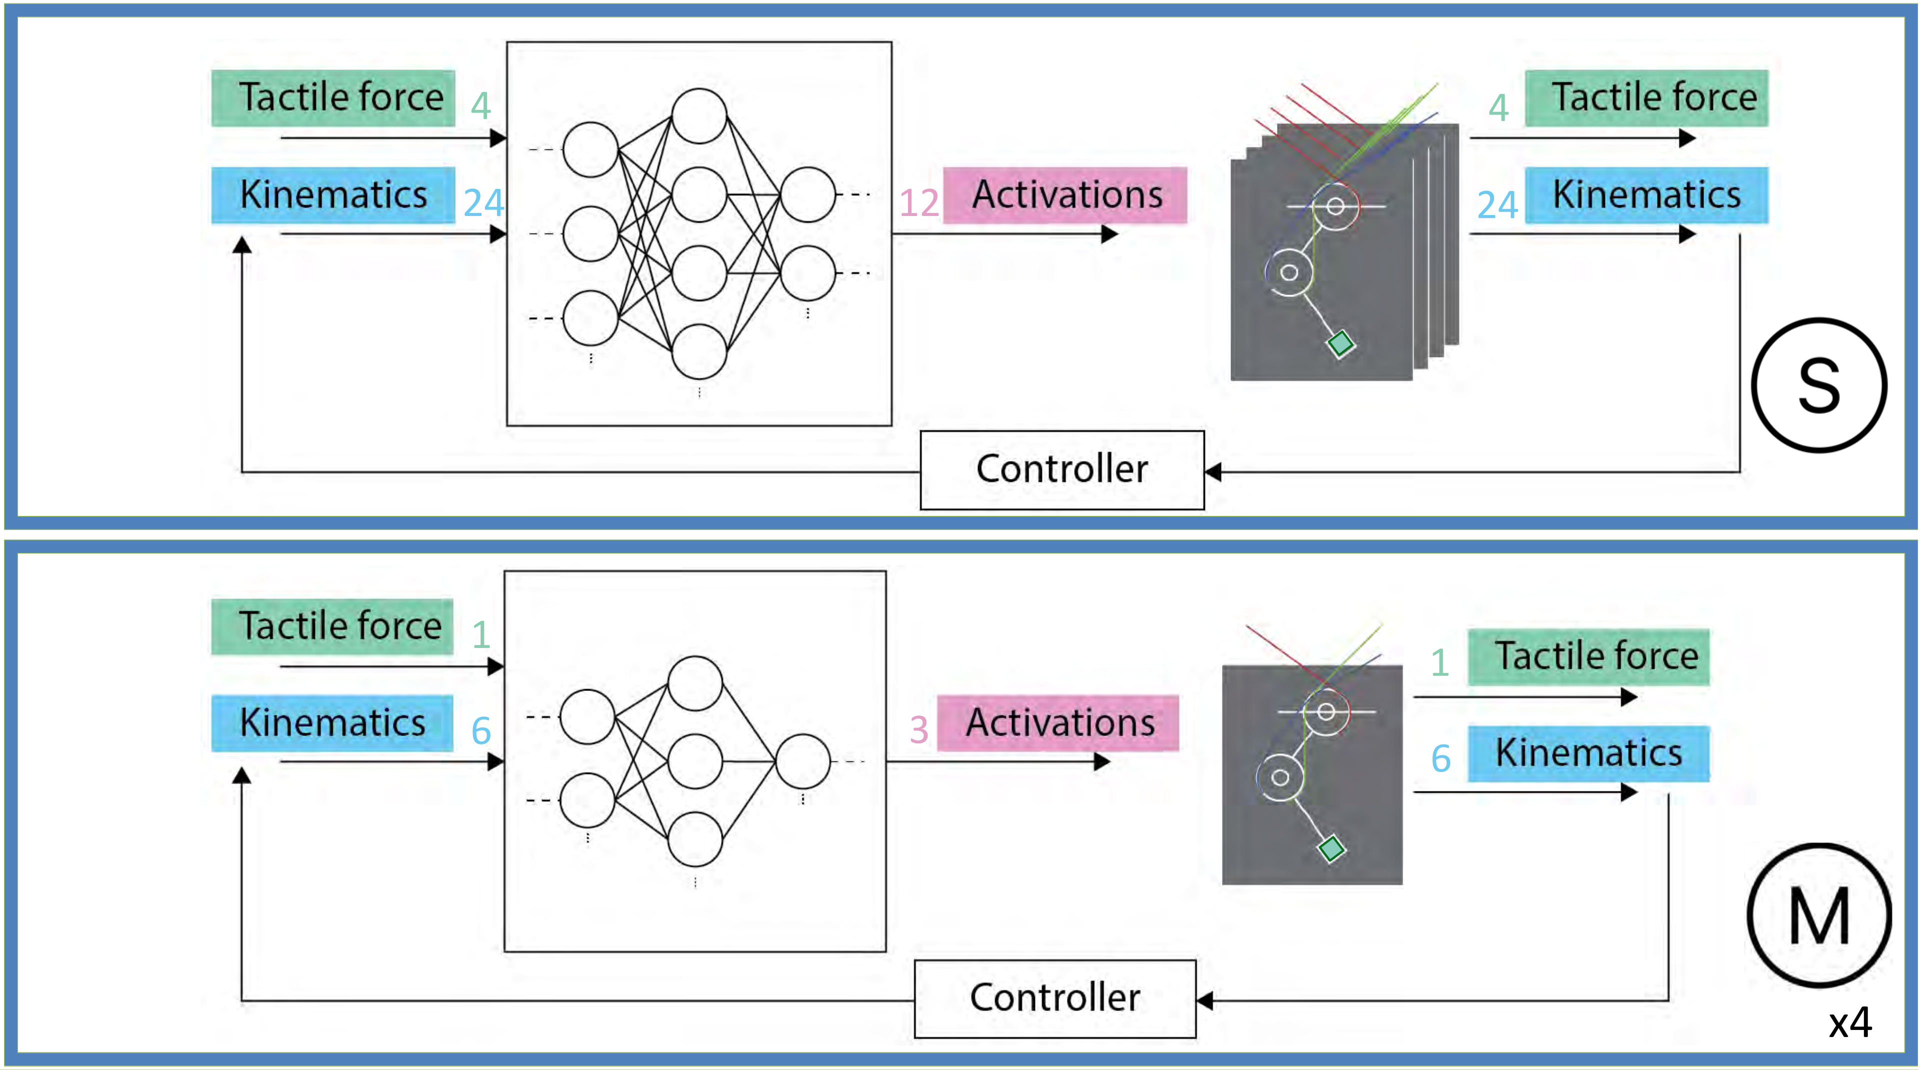

Stacking all the best configurations

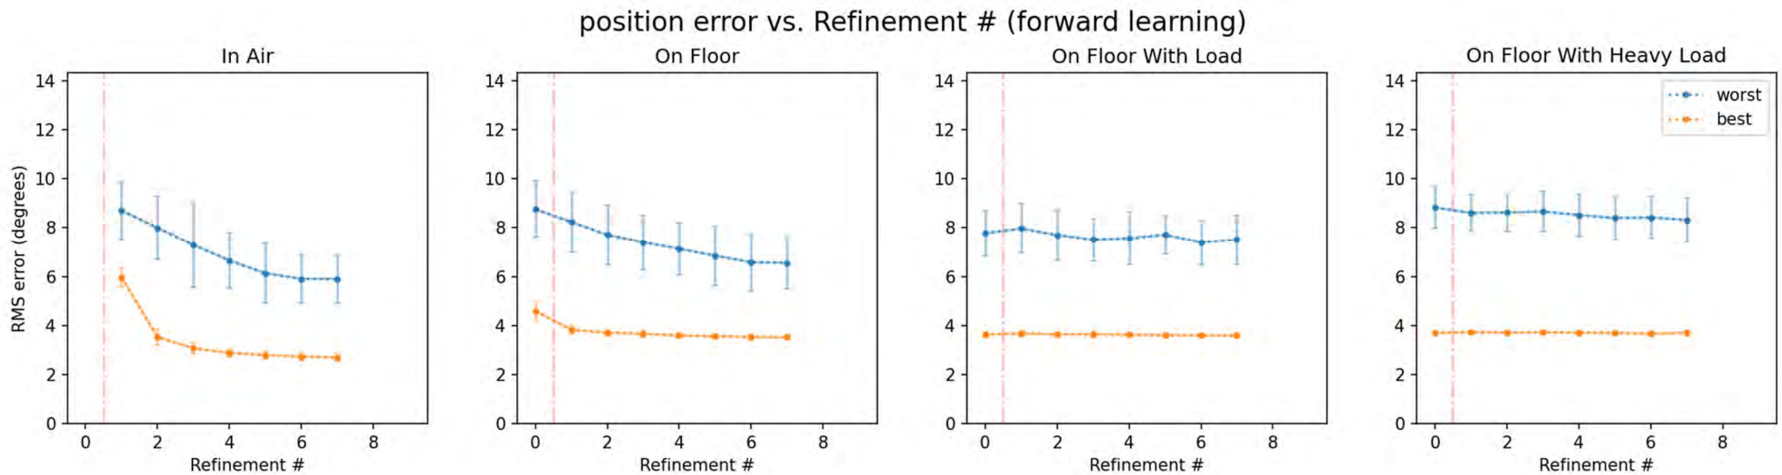

Configurations

Feedback	Tactile	ANN arch
✗	✗	S

original (worst)

VS

Configurations

Feedback	Tactile	ANN arch
✓	✓	M

With extensions (best)

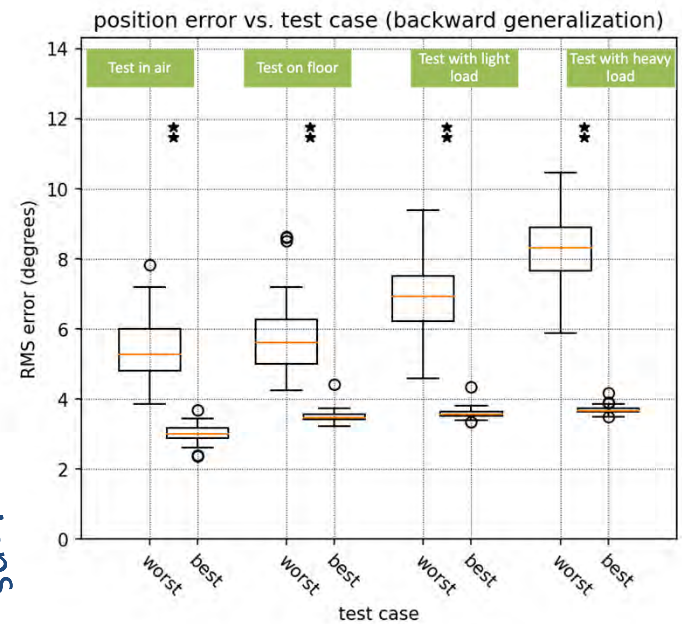

Stacking all of the best configurations

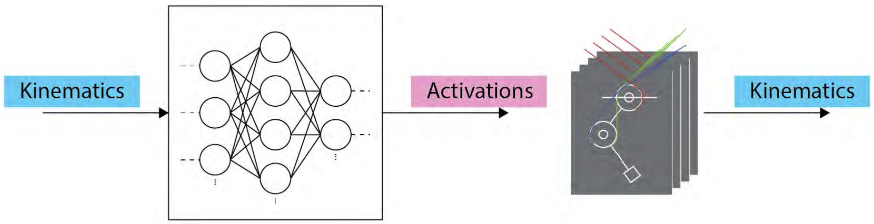

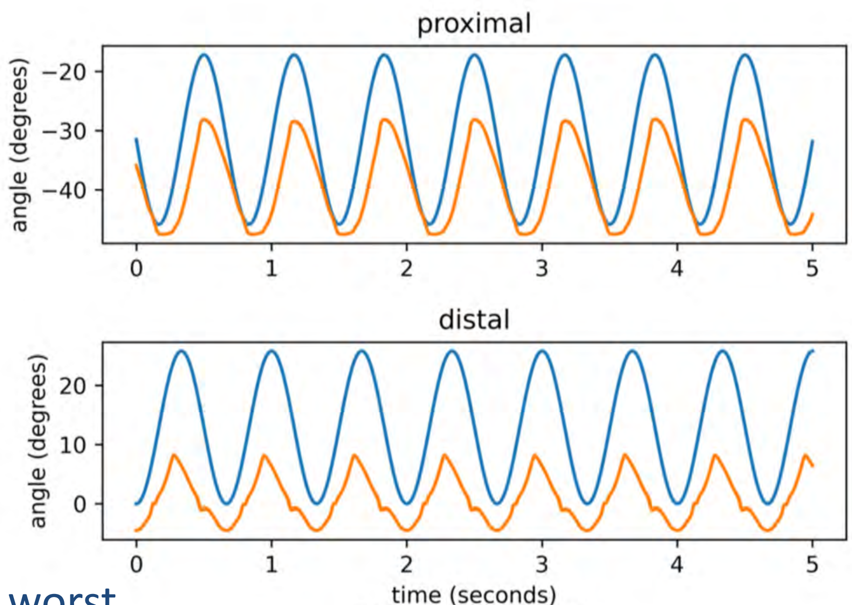

RMSE (degrees): 9.13

worst

Configurations	Feedback	Tactile	ANN arch
	✗	✗	S

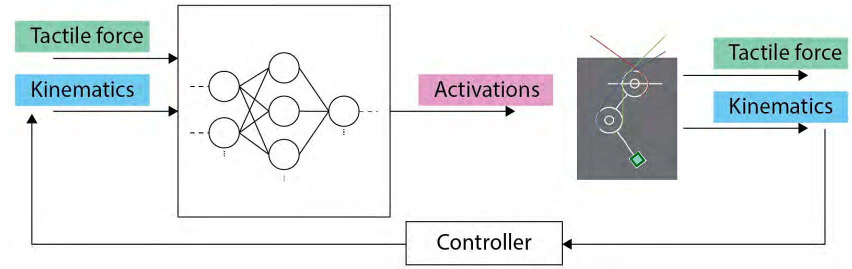

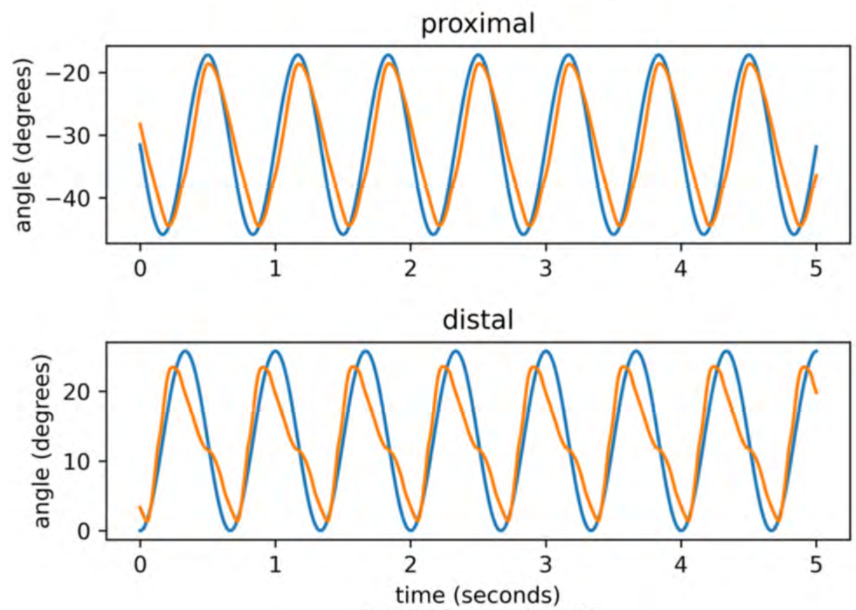

RMSE (degrees): 3.76

Feedback	Tactile	ANN arch
✓	✓	M

best

Stacking all of the best configurations

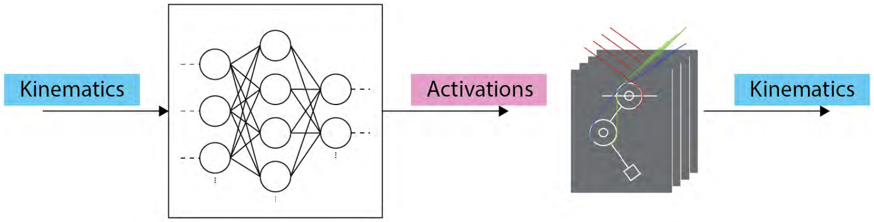

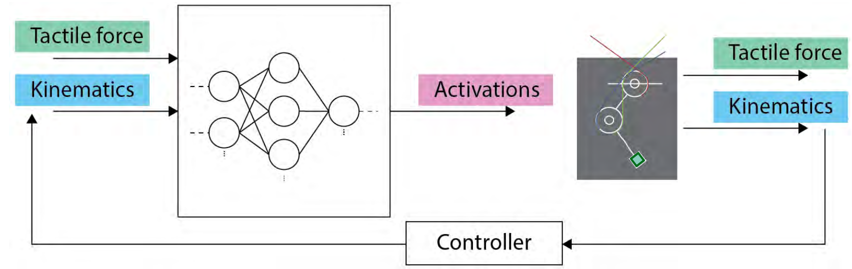

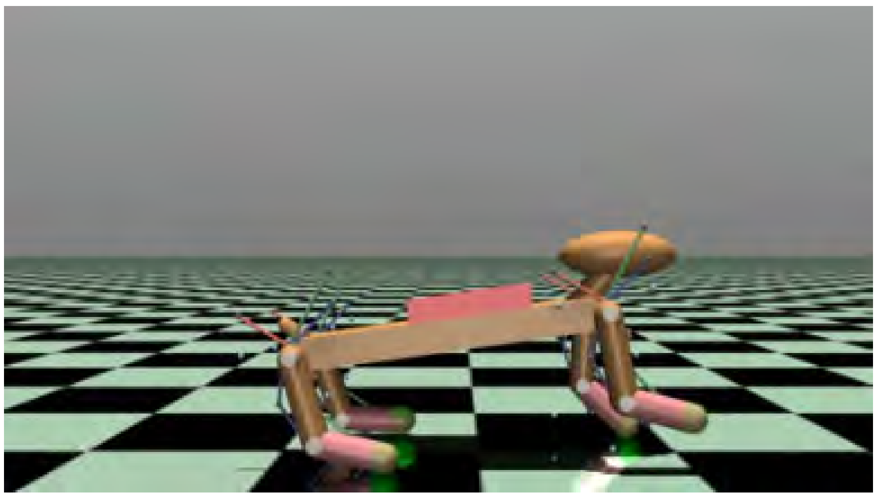

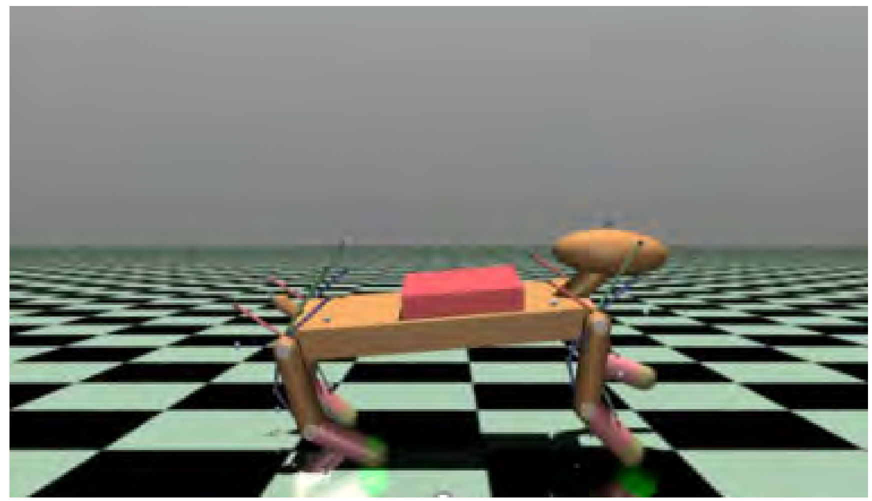

Configurations	Feedback	Tactile	ANN arch	Feedback	Tactile	ANN arch	
	✗	✗	S	✓	✓	M	best

Conclusion

- We were able to demonstrate:
  - o Fast, data-efficient learning of kinematic control in a tendon-driven quadruped **with minimal experience, and limited information** (kinematics of the leg and contact sensors at the paw)
  - o Robust performance across testcases (in air, on floor, with loads) with median RMS error < 4 degrees
  - o Studied the effects of different ANN configurations and sensory conditions
    - Multiple (leg specific) ANNs arch. significantly outperforms the all-to-all approach
    - Adding kinematic feedback significantly improves performance
    - Tactile can contribute to improved generalizability; but further experiments are needed to find the optimal inter-network connectivity

Future directions for this research:

- Hardware implementation, and extensions to manipulation
  - Testing in real-world scenarios
  - Model-agnostic: makes it robust to variabilities in hardware design
- Developing hybrid ANN structures
  - Add communication <u>across</u> multiple ANNs (as in biology) to study hierarchical and distributed control and robustness

Our Team

Acknowledgements

- USC Provost fellowship
- USC Graduate School's Research Enhancement fellowship

USC

NIH

National Institutes of Health

National Institute of Neurological Disorders and Stroke

Congressionally Directed Medical Research Programs

CDMRP

Department of Defense

DARPA

NSF

Collaborative Research in Computational Neuroscience

Thank you!

L2M Hardware Implementations

Dhireesha Kudithipudi, PhD

NEUROMORPHIC
ARTIFICIAL
INTELLIGENCE LAB

RESEARCH CONTRIBUTORS : Tej Pandit, Anurag Daram, Nicholas Soures and Peter Helfer

UTSA The University of Texas at San Antonio™

Argonne NATIONAL LABORATORY

Sandia National Laboratories

DARPA

ERI Summit, Workshop on New Opportunities for **Lifelong Learning Machines** 20th October 2021

Lifelong Learning at the Edge

LIFELONG LEARNING EXAMPLE

L2 agents need to be able to learn tasks continually on the edge, while adapting to dynamic environment conditions

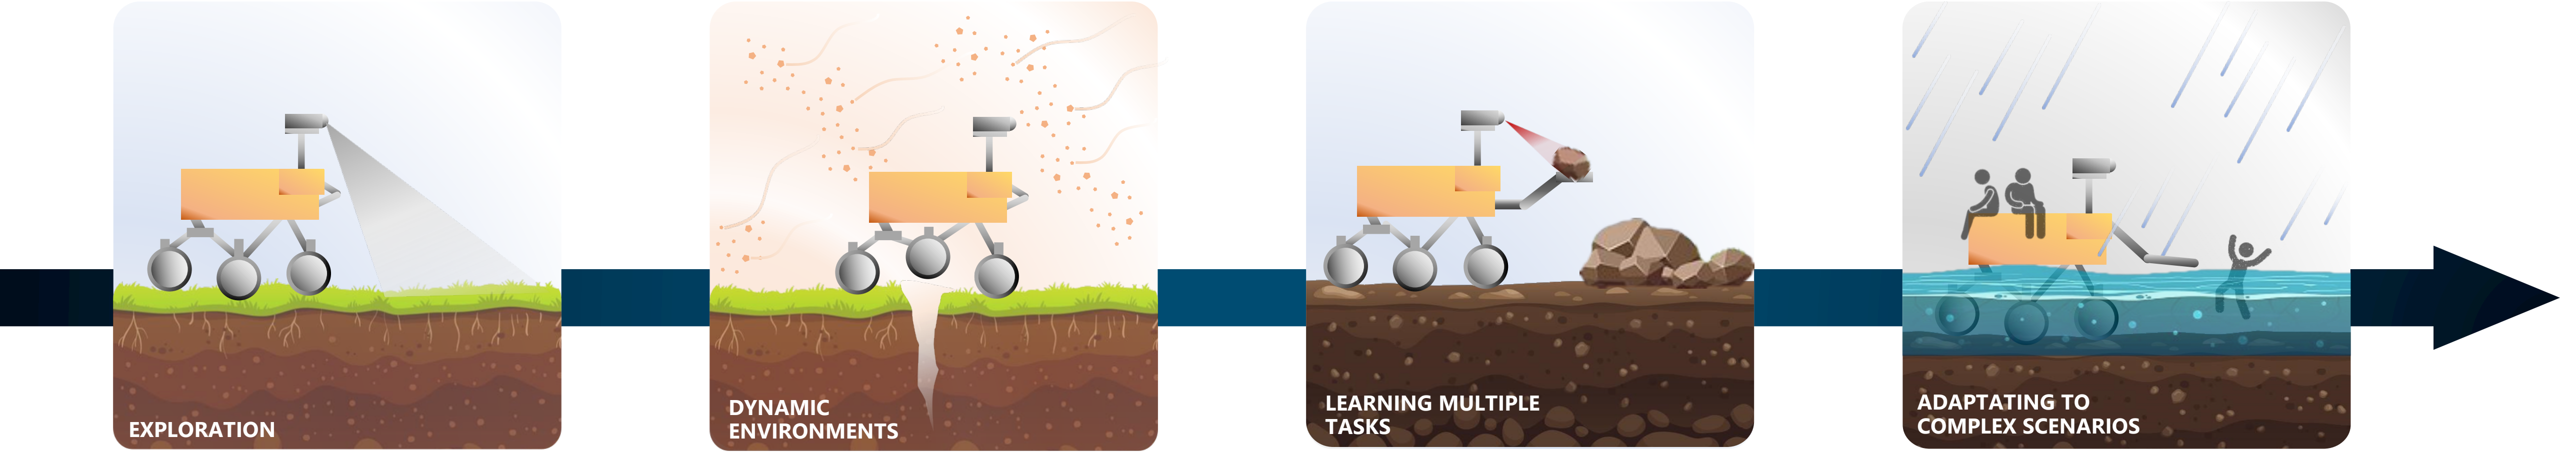

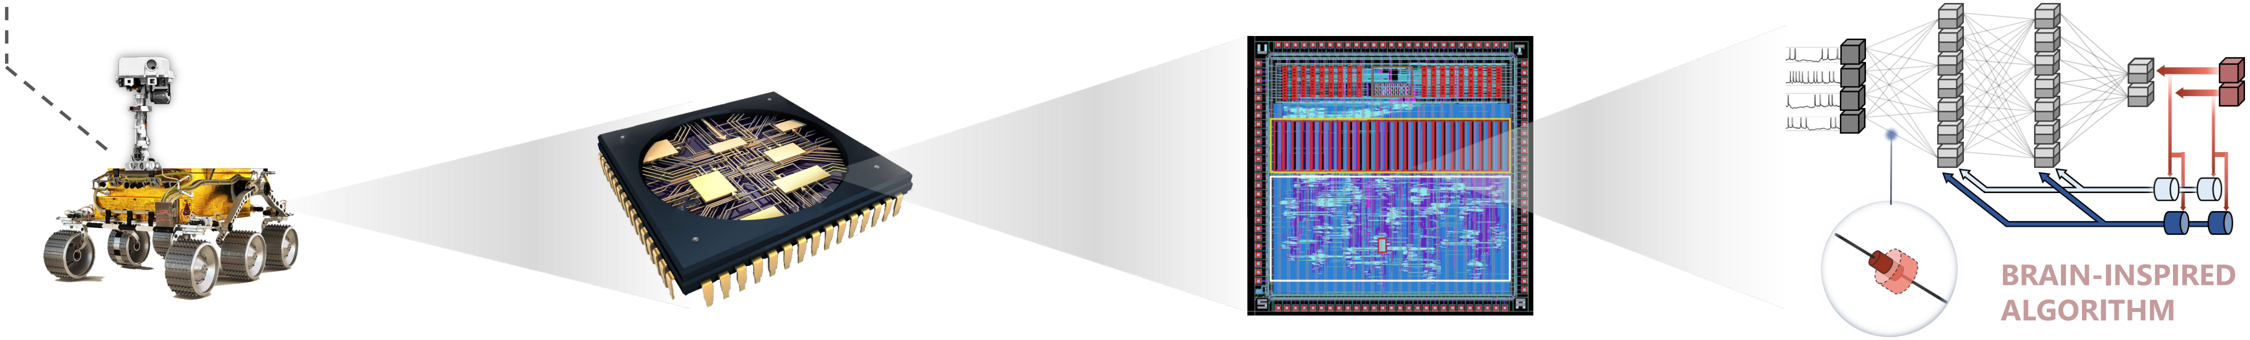

NEUROMORPHIC HARDWARE

L2 agents operating at the edge require specialized hardware

Design Considerations for Lifelong Learning Hardware

SUPPORT STRUCTURAL PLASTICITY

EVER-CHANGING LATENCY REQUIREMENTS

HANDLING MULTI-MODAL NON-IID DATA DISTRIBUTION

ROBUST ON-DEVICE LEARNING

Multiscale Co-design Approach Spanning Algorithms, Architectures And Devices

DYNAMIC/UNBOUNDED MEMORY

LIMITED ENERGY BUDGET

VARIABLE DIMENSION DATA MOVEMENT

Insect-brain Inspired Architectures for Rapid Learning

Concept : Neuromodulatory interactions demonstrate quick and robust learning of short-term associations

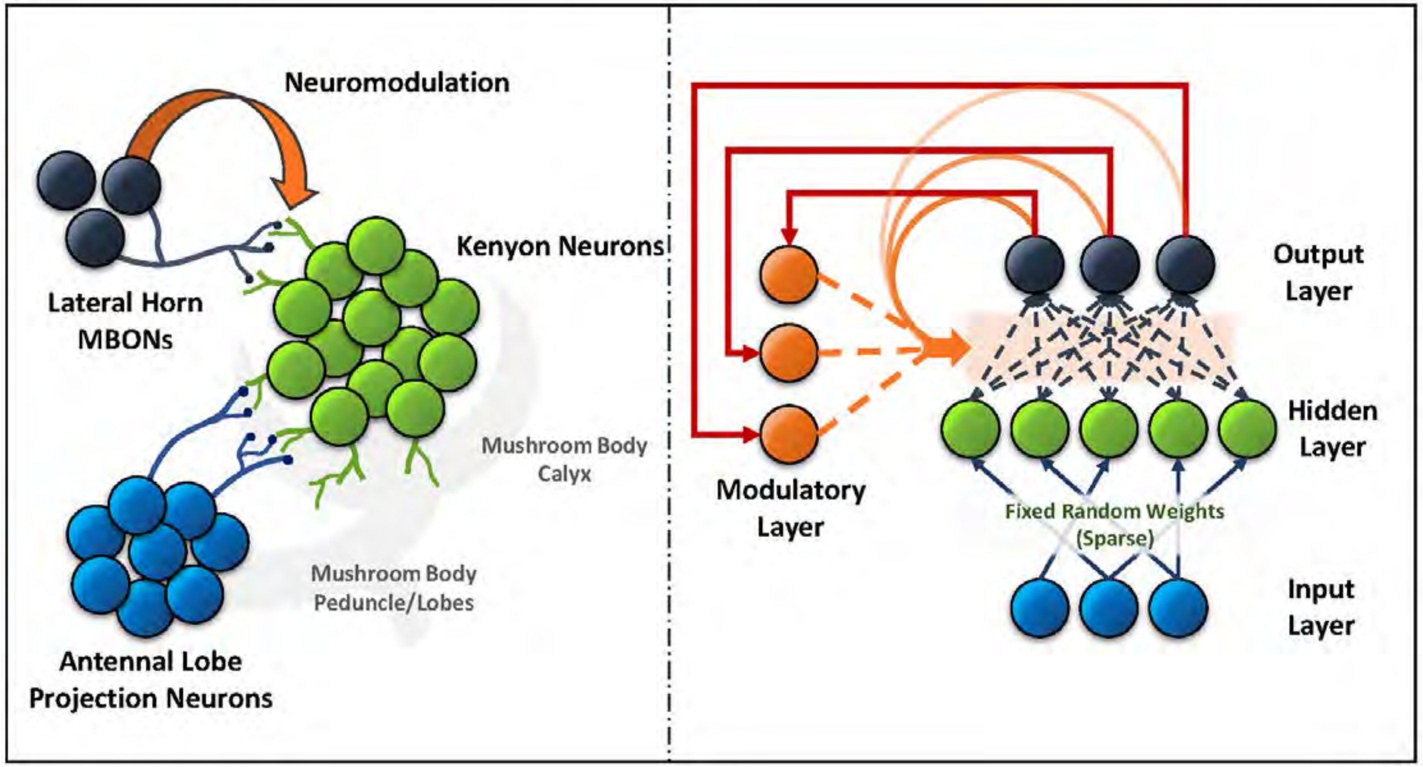

Available at : https://github.com/Nu-AI/Neuromodulatory_OneShotLearning

FEATURES

Heterogenous local plasticity

- Learn tasks quickly
- Real time learning and execution

Compartmentalization

- Learning with minimal forgetting
- Adaptation to new tasks

Insect-brain Inspired Architectures for Rapid Learning

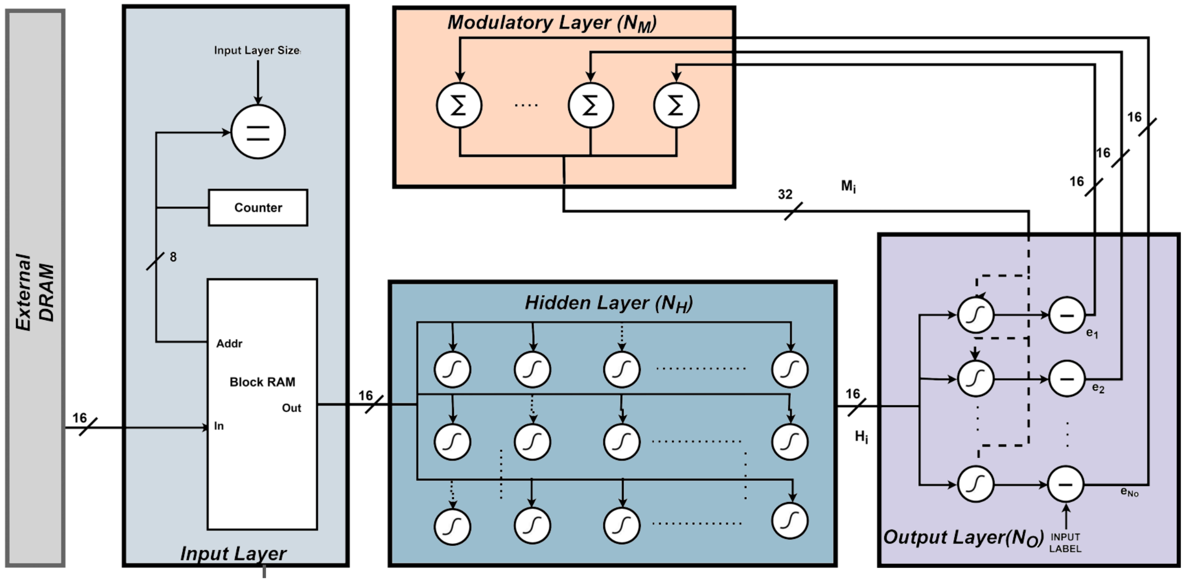

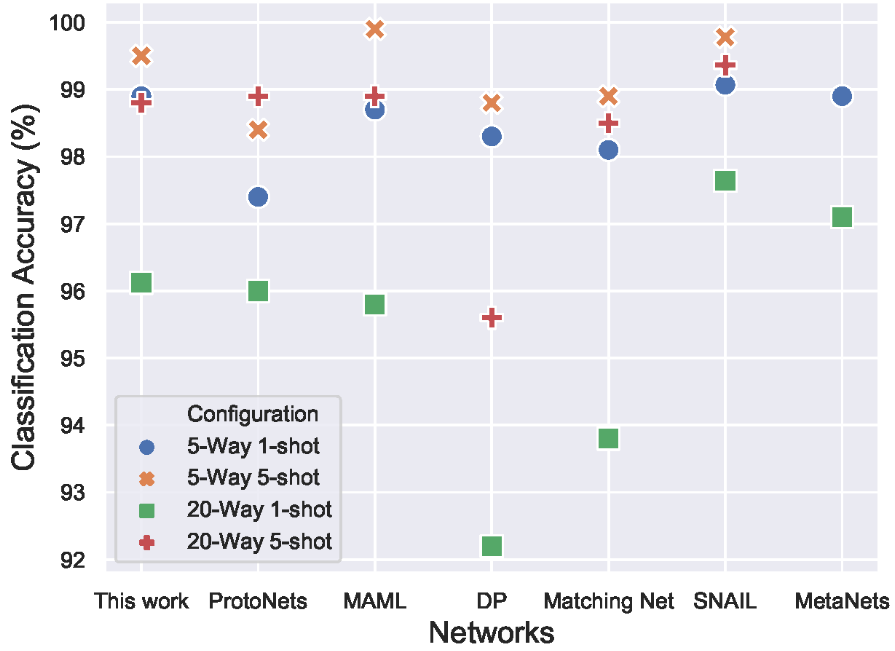

FPGA Implementation

The proposed architecture can process **~3000** images per second (training) under **<1W**

The proposed architecture learns rapidly in **0.15 epochs**

SOTA accuracy with <**~20x** training parameters and episodes

Few-shot image recognition on Omniglot dataset

TASK AGNOSTIC CONTINUAL LEARNING IN SPIKING NEURAL NETWORKS

Presented at Theory and Foundations of Continual Learning workshop ICML 2021

Heterogeneous neural-plasticity mechanisms can aid in designing continual learning systems in an energy efficient way

Local Plasticity

Does not rely on task knowledge, Sparse updates, Low communication overhead

Complex Synapses

Exploit heterosynaptic decay and metaplasticity to mitigate catastrophic interference

Model size does not grow as a function of number of tasks

Spiking Neural Networks

Energy efficiency in communicating with single spike [~20pJ per spike]

TACOS

TASK AGNOSTIC CONTINUAL LEARNING IN SPIKING NEURAL NETWORKS

SPONSORED BY:

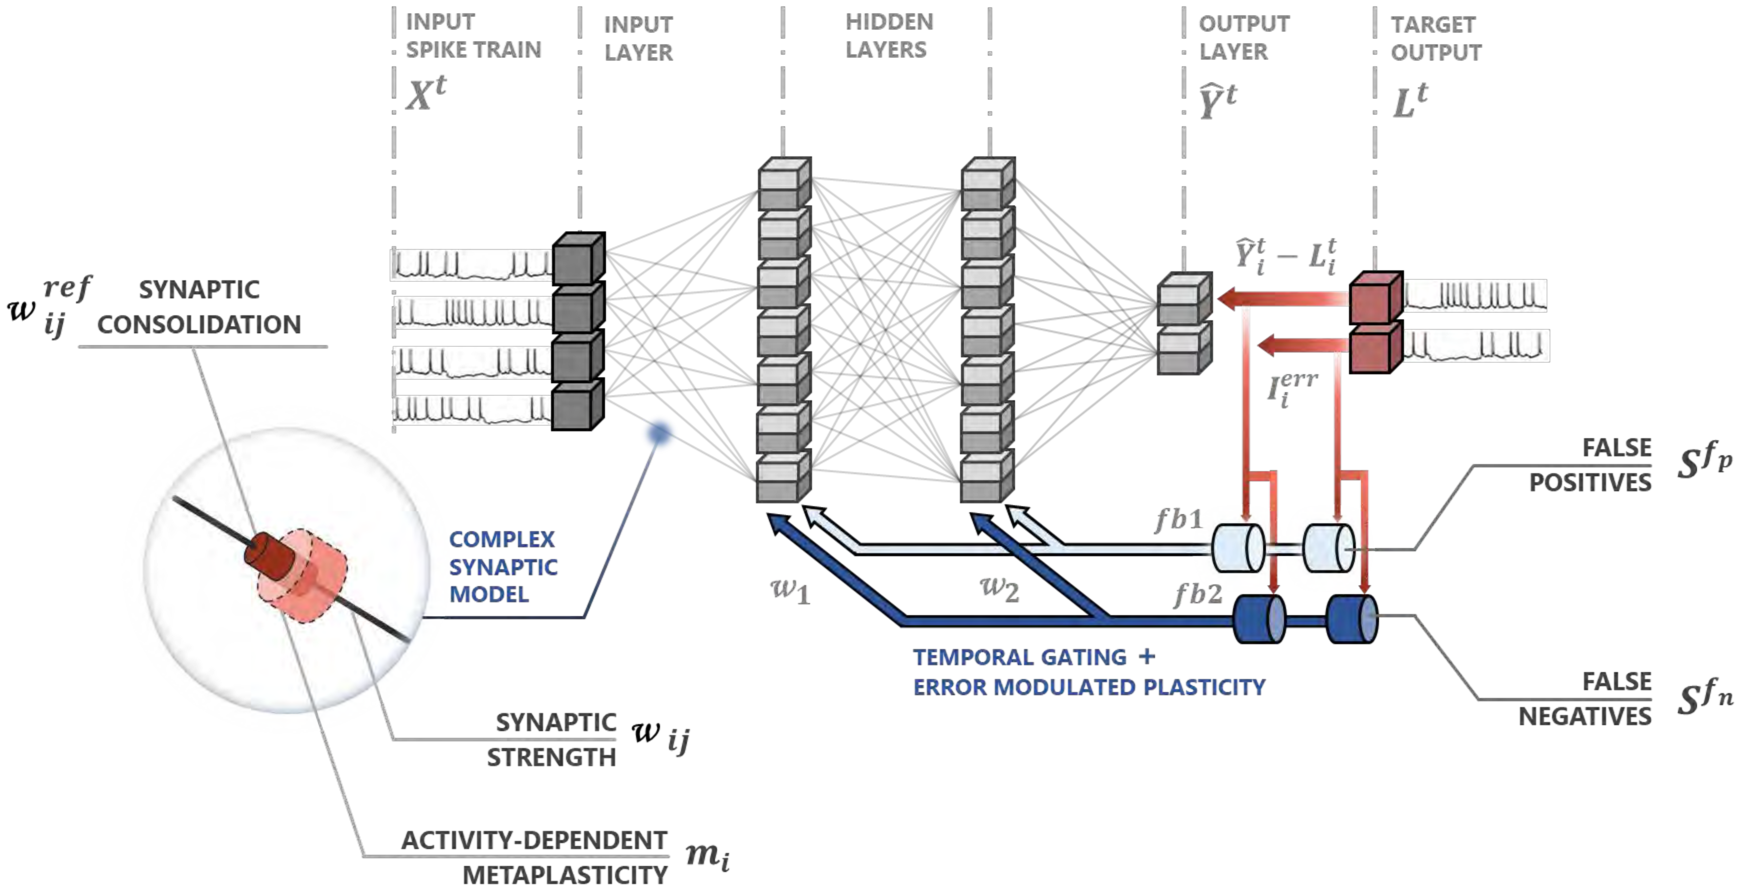

TACOS

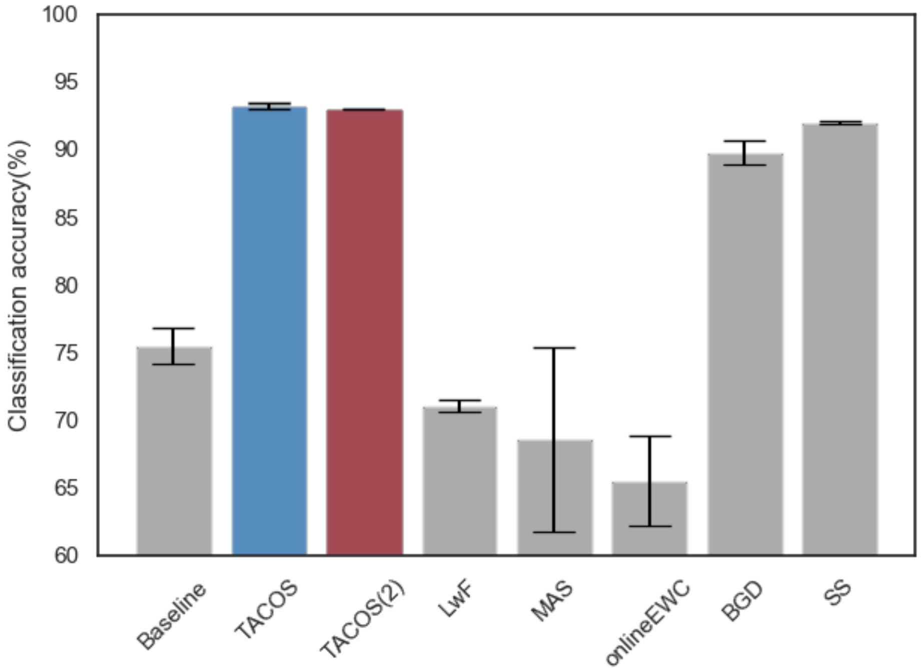

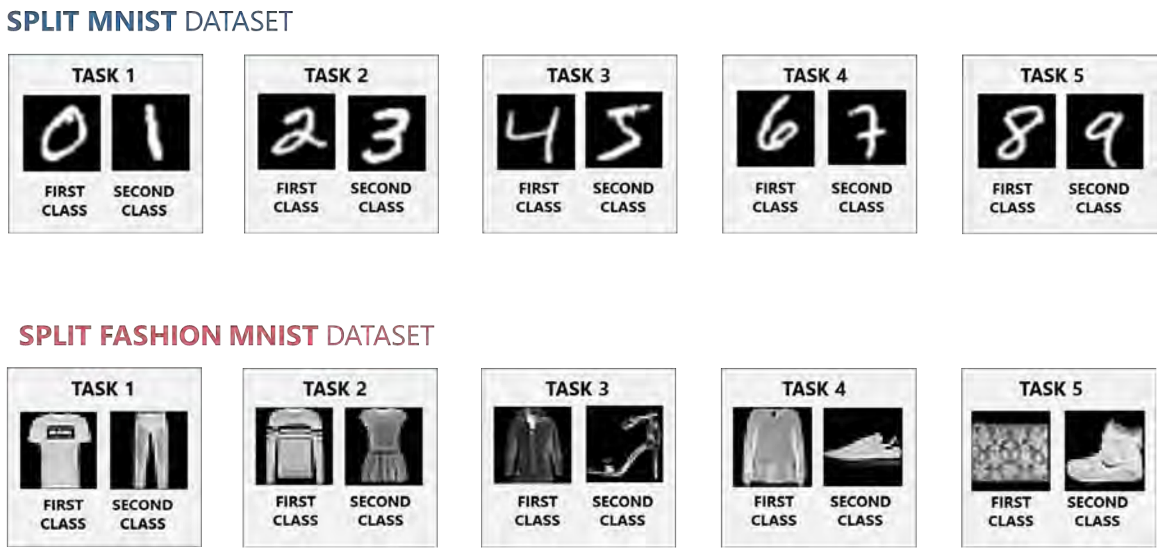

First demonstration of full-scale spiking neural network with continual learning capability

TACOS mean performance across all tasks is better than **state-of-the-art regularization models** in domain-IL scenario

Achieved through local plasticity functions **which do not require task information** or grow with the number of tasks

Trained only on a **single epoch**

GENESIS: First Lifelong Learning ASIC

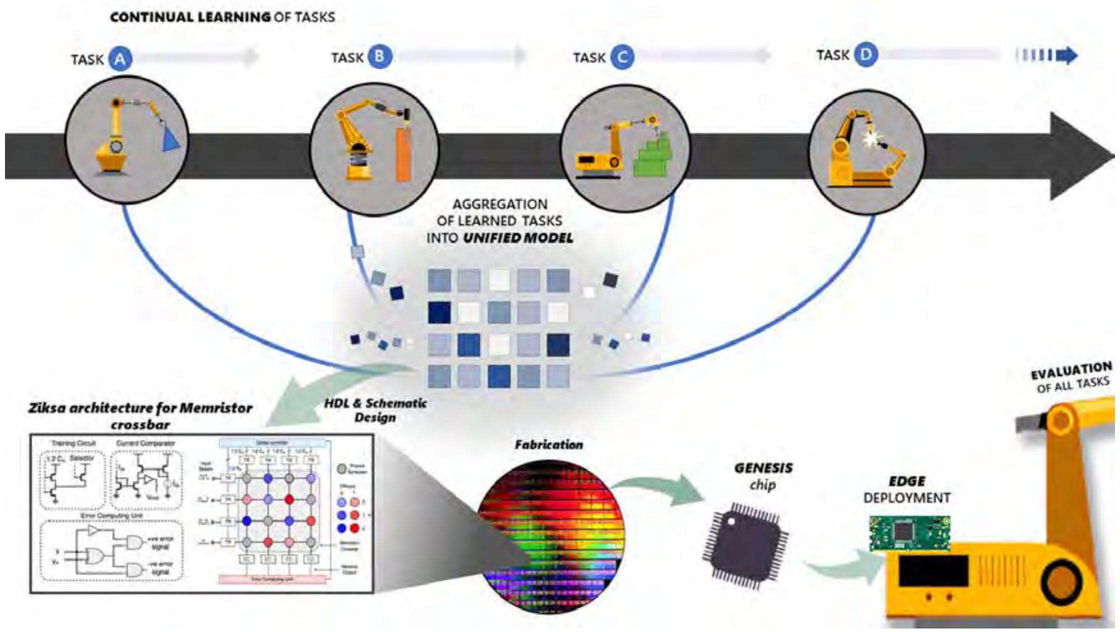

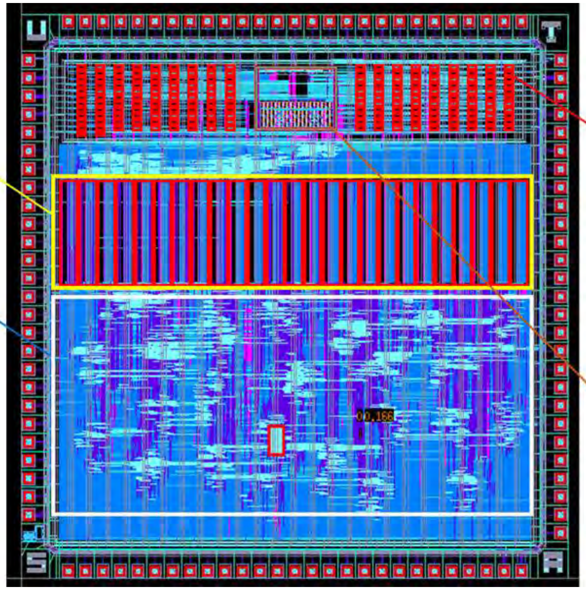

First mixed-signal AI accelerator for lifelong learning (Tape out: IBM 65nm)

Architecture for spiking continual learning (TACOS)

On-device learning with CMOS/memristor crossbars

161x power savings and **2x** speedup compared to SOTA accelerators

SPONSORED BY:

Publications: IEEE TC'20, SOCC'15, IJCNN'19, ISCAS'17 '18 '20, ACM JETC '18'19

DoD Applications

SEARCH AND RESCUE OPERATIONS
Simulations of Natural Disasters and how AI Agents can help identify and rescue people in dangerous situations.

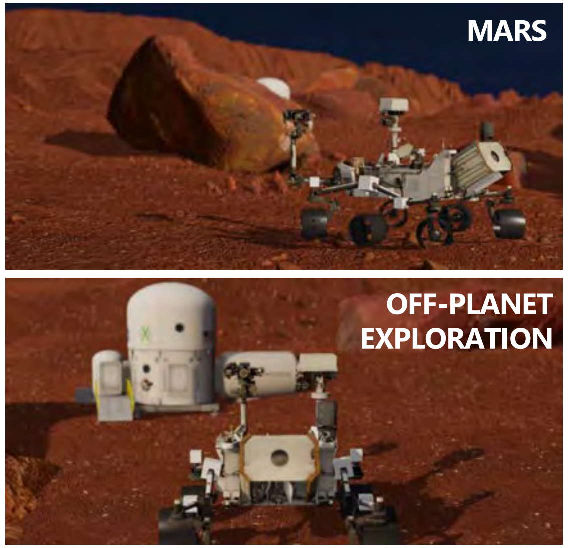

ENVIRONMENTAL EXPLORATION
Agents exploring procedurally generated environments can quickly adapt to new/unseen environments and carry out exploratory tasks.

Thank You

Contact

dk@utsa.edu

https://www.nuailab.com/

Omnidirectional Lifelong Learning via Representer Ensembling

JHU: Jayanta Dey | Ali Geisa | Hayden Helm | Ronak Mehta | Will LeVine | Carey E. Priebe | Joshua T. Vogelstein
Microsoft Research: Weiwei Yang | Jonathan Larson | Bryan Tower | Chris White

Lifelong Learning

- Goal : Improve on
  a. past tasks
  b. current tasks
  c. future or yet unseen tasks

Motivation

- Challenges with the Current Approaches:

  a. Can improve performance on the future tasks
  b. Struggles to maintain performance on the past tasks
  c. Suffers from catastrophic forgetting
  d. Has worse space, capacity and time complexity

- Our Claim :

  a. Only algorithm to show transfer in both directions- past to future and future to past tasks
  b. Has quasilinear space and time complexity
  c. Shows synergistic learning on CIFAR 100 dataset

Learning Schemas

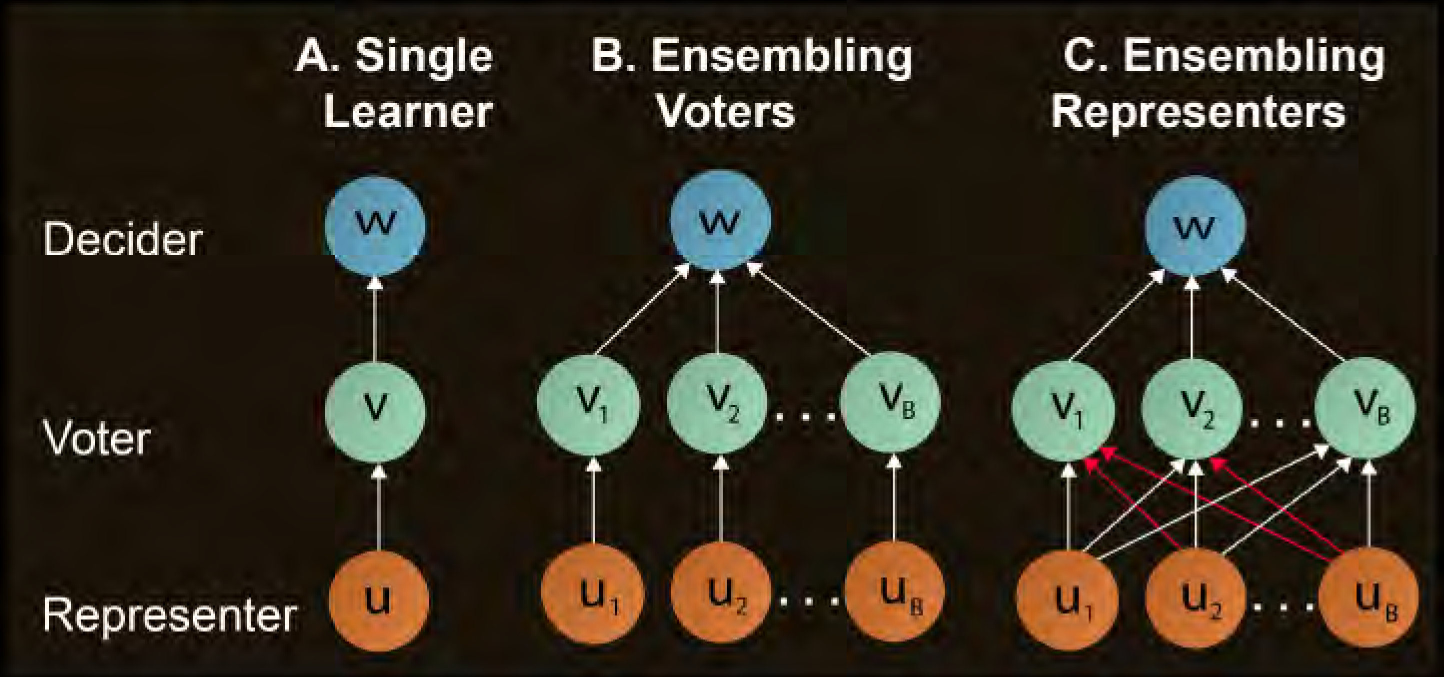

Omnidirectional Algorithms can Transfer Between XOR and XNOR

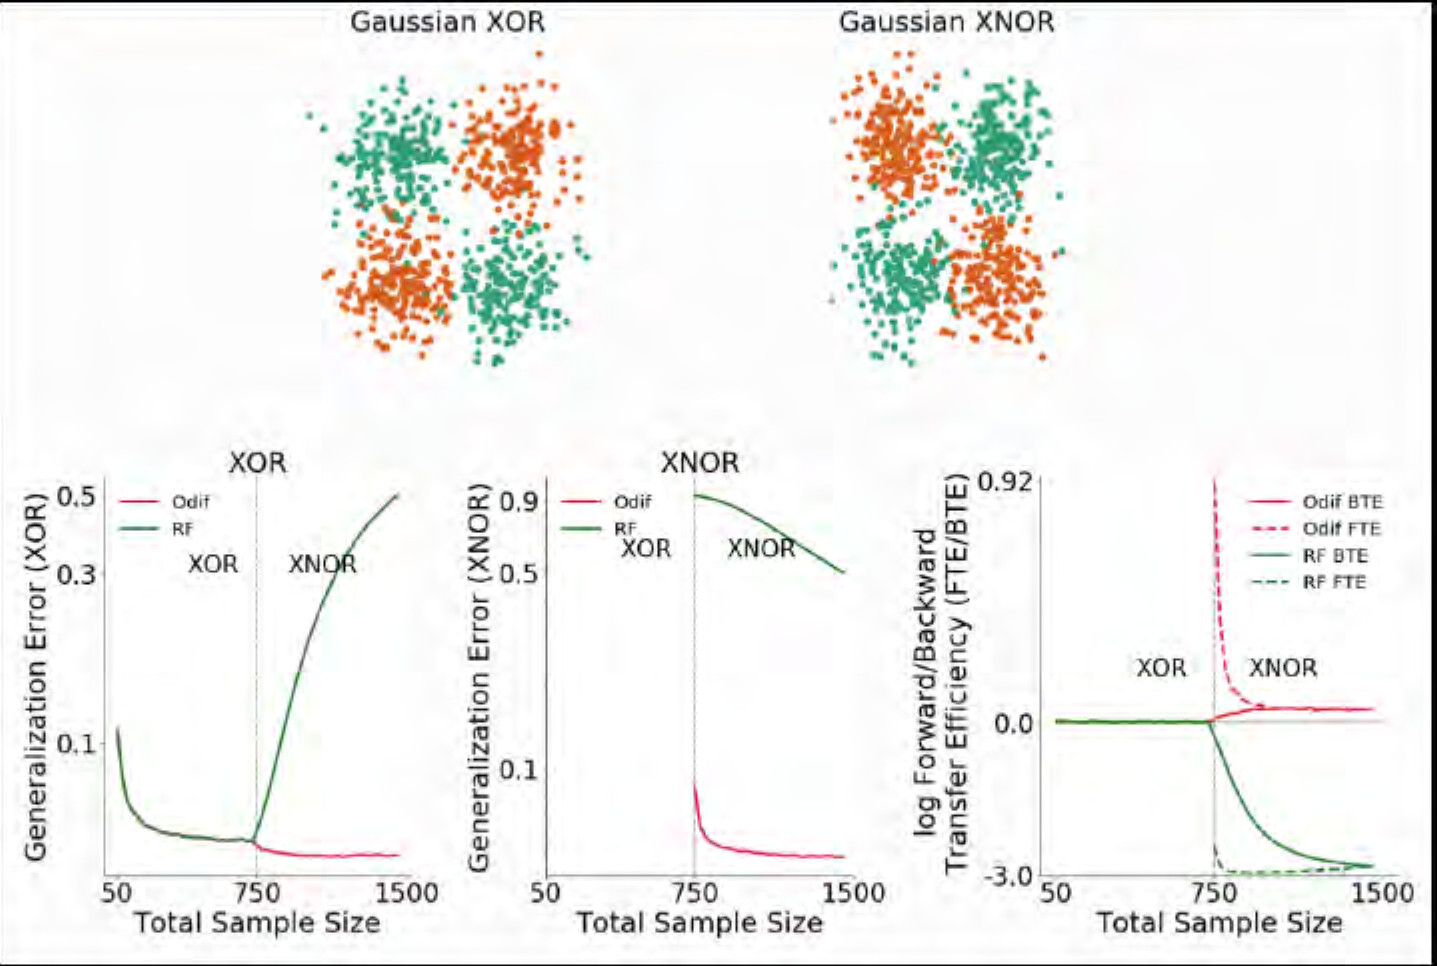

CIFAR 10x10

- *CIFAR 100* is a popular image classification dataset with 100 classes of images.
- 500 training images and 100 testing images per class.
- All images are 32x32 color images.
- CIFAR 10x10 breaks the 100-class task problem into 10 tasks, each with 10-class.

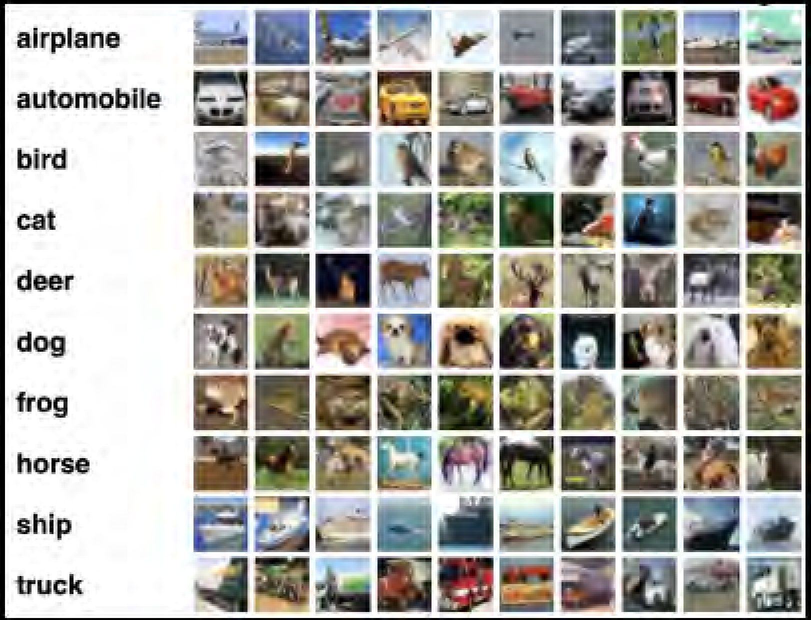

Omnidirectional Algorithms Show Forward Transfer for the CIFAR 10x10 Tasks

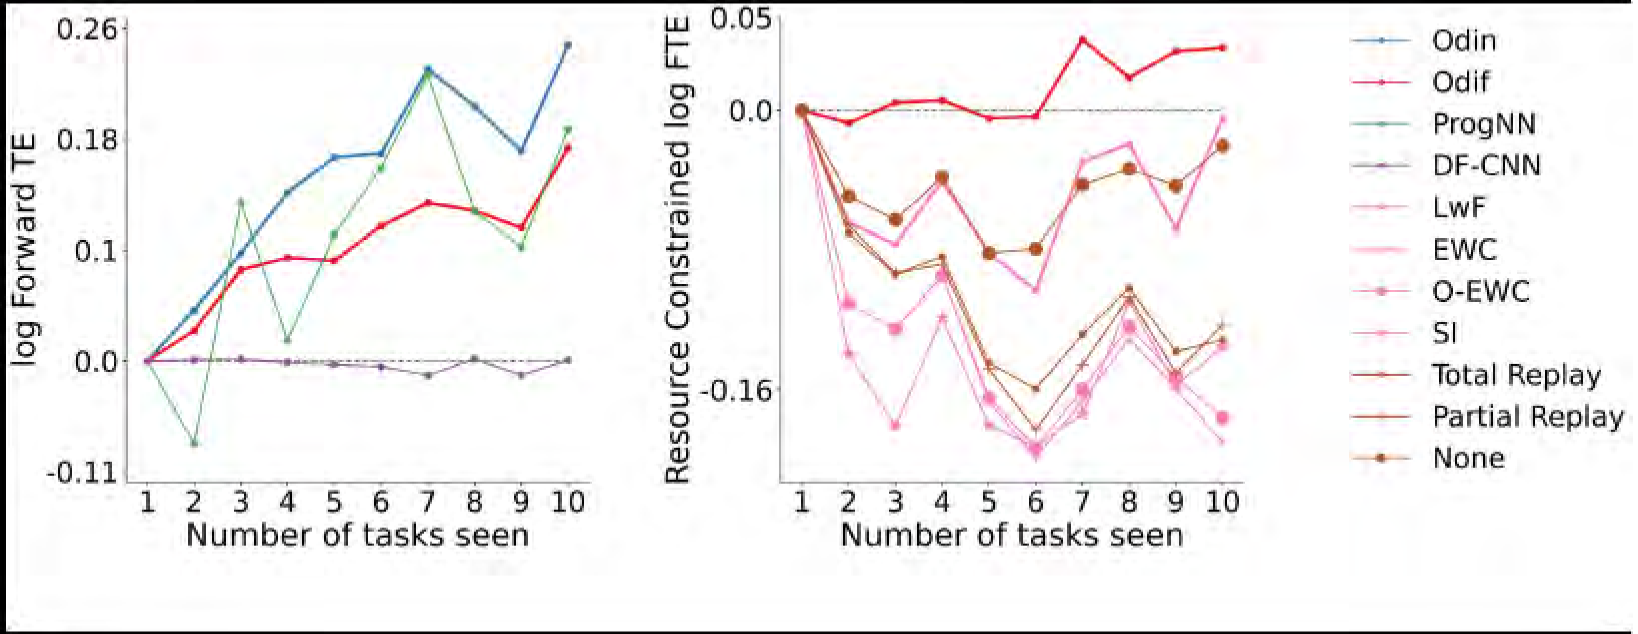

Omnidirectional Algorithms Uniquely Show Backward Transfer for Each CIFAR 10x10 Task

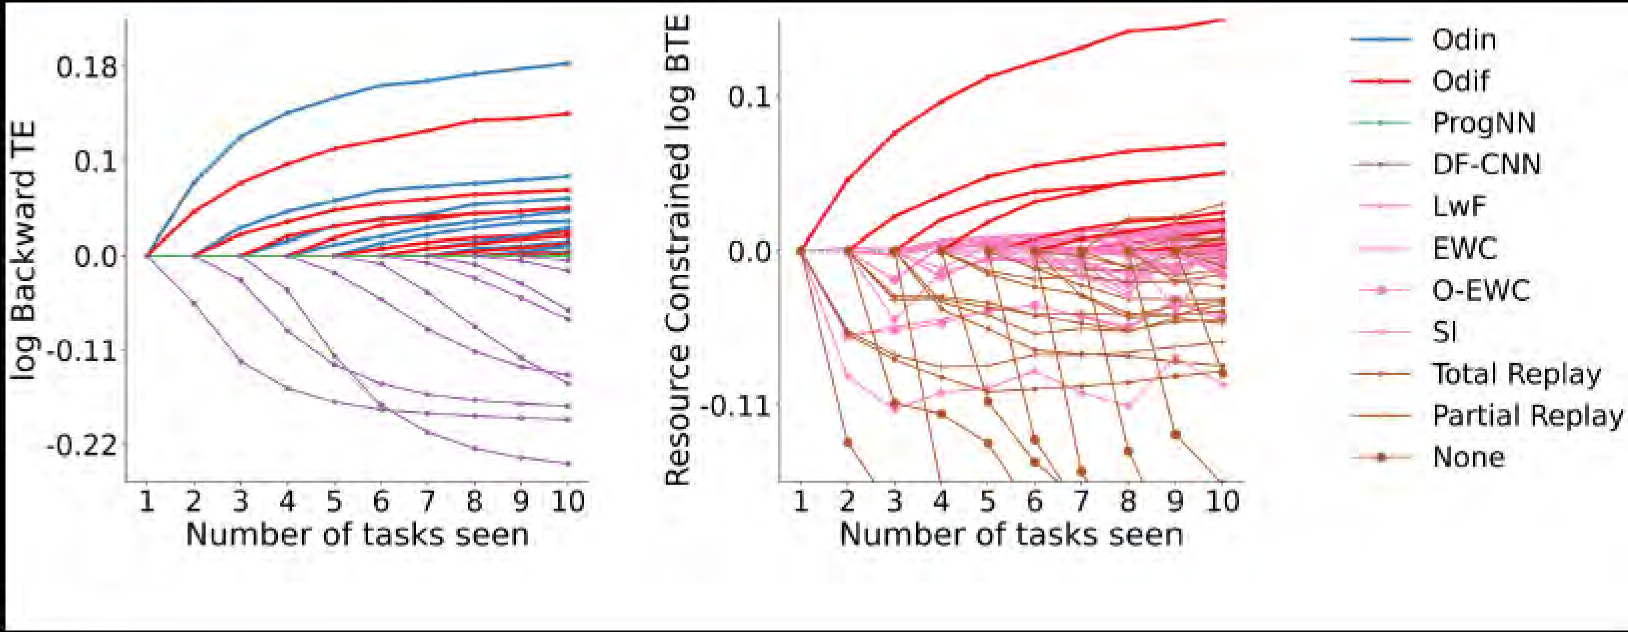

Acknowledgements

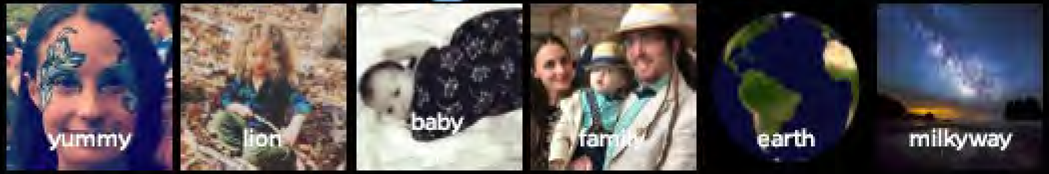

JHU

Microsoft Research

DARPA L2M

{BME,CIS, ICM, KNDI}@JHU | neurodata
jovo@jhu.edu | http://neurodata.io/talks | @neuro_data