




PRAVEEN PILLY

HRL LABORATORIES
DARPA L2M PROGRAM




LIFELONG LEARNING MACHINES (L2M)

This material is based upon work supported by the United States Air Force and DARPA under Contract No. FA8750-18-C-0103. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Air Force and DARPA.

THE STATE OF AI IS CONFUSING

Credit: DARPA PM Siegelmann


Beyond human capabilities



i2.kknews.cc/SIG=29vnh65/2175/3455714929.jpg



© DeepMind Technologies



But lacking basic capabilities

Original Inputs



Modified Inputs



Wrong ML Behavior



Not robust to open, novel situations

SOME KEY CHALLENGES FOR AUTONOMOUS CARS

Surprising events



<https://abcosafety.com/c-63-traffic-work-zone.aspx>

Sensor failure



<https://www.weather.gov/safety/fog-driving>



Lidar



GPS




Camera



RADAR

Training data
distribution

Unforeseen data
distribution



- Continual learning
- Robust generalization
- Rapid adaptation

CURRENT CONOPS

**Not sustainable without
automated continual learning**

Training




**Fleet
deployed**



Detect lane markings
Detect objects
Detect drivable surface

Handoff to human operator
Human intervention
Crash

Car N Car 1



Log of recent
experiences
leading up to
"surprise event"

**Re-train/update
base model**

**Test and verify
all core
capabilities**



**several
weeks
later**

**Software
update**

commons.wikimedia.org/wiki/File:Tesla_Model_S_%26_X_side_by_side_at_the_Gilroy_Supercharger.jpg

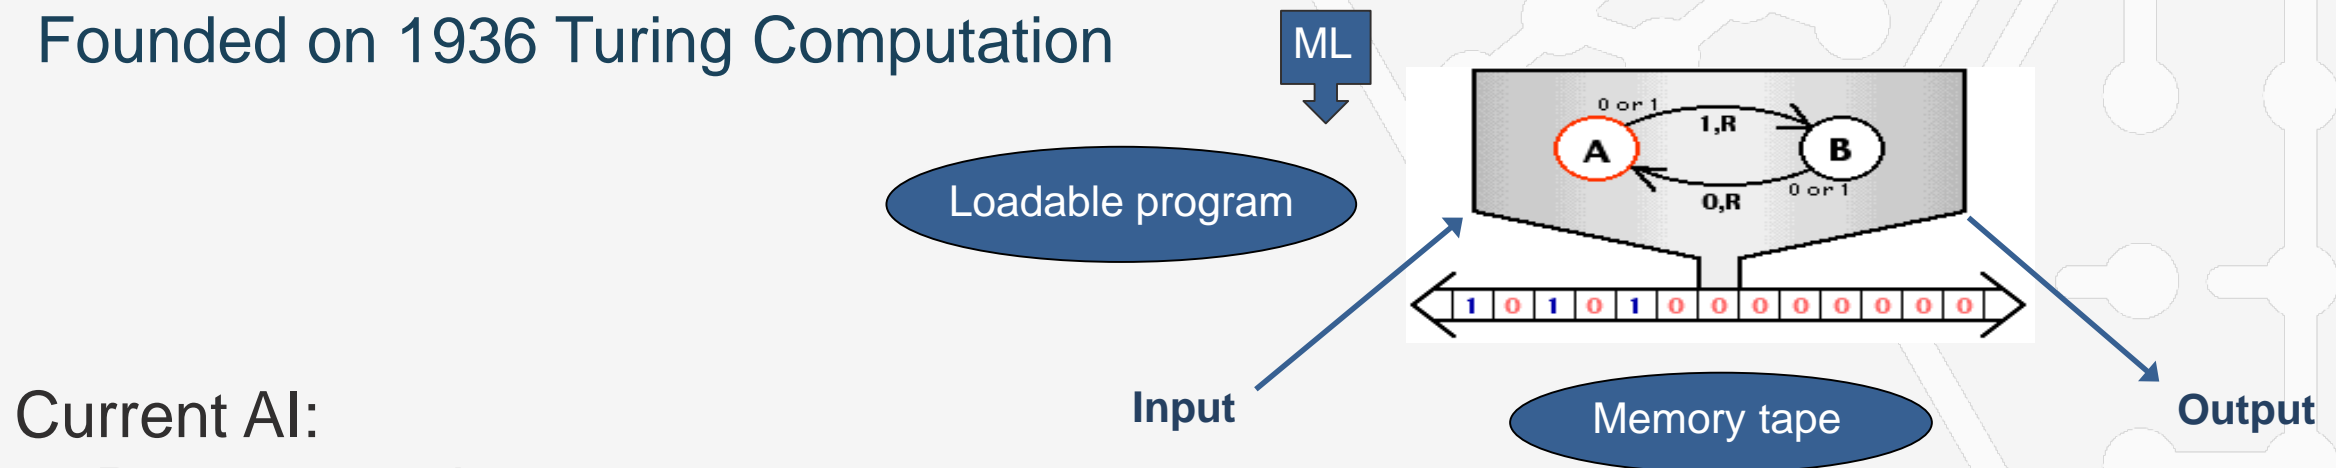
IDENTIFYING THE KEY LIMITATION

Credit: DARPA PM Siegelmann

Current AI systems only compute with what they've been programmed or trained for in advance

- No way to prepare a training set for all possible futures
- Malfunctions in unseen circumstances

Founded on 1936 Turing Computation




Current AI:

- Program and rules
- Machine learning: fitting parameters: training data & training method

NEW FOUNDATION: SUPER-TURING COMPUTATION

Credit: DARPA PM Siegelmann



Stronger computation: Continuum of computational hierarchy
From Turing Machines (fixed programs) to Super-Turing
Computation (modifiable programs)

Larger: Exponentially more functions

$$\aleph_0 \longrightarrow 2^{\aleph_0}$$


Instead of Turing Machines, employ Analog and Recurrent Neural Networks:
Super-Turing is the foundation of Lifelong Learning Machines

TOWARDS BETTER MACHINE LEARNING

Credit: DARPA PM Siegelmann

From frozen not-so-intelligent systems to systems that learn from experience and improve performance with time

We combine fixed programs and lifelong learning just like the brain combines Turing and Super-Turing



PROGRAM RESULT 1: GHOST ARM FOR INTERNAL "SELF MODELING"

Credit: DARPA PM Siegelmann

Columbia University

Self-modeling for speedy adaptation to new conditions
(change in task or robot)

Kwiatkowski & Lipson (2019)

Sensor failure



<https://www.weather.gov/safety/fog-driving>




Lidar



GPS



Camera



RADAR



PROGRAM RESULT 2: UNSUPERVISED ASSOCIATIONS TO INTELLIGENT SEARCH


Credit: DARPA PM Siegelmann

NYU, Toyota-TIC, HRL

Self-play kick starts learning in the absence of explicit tasks / labels

Jiang et al. (2018)

Depth
from 2D



Time →

Use consistency in spatial and temporal predictions as a proxy supervision signal

U Mass

Use self-learned associations for fast search


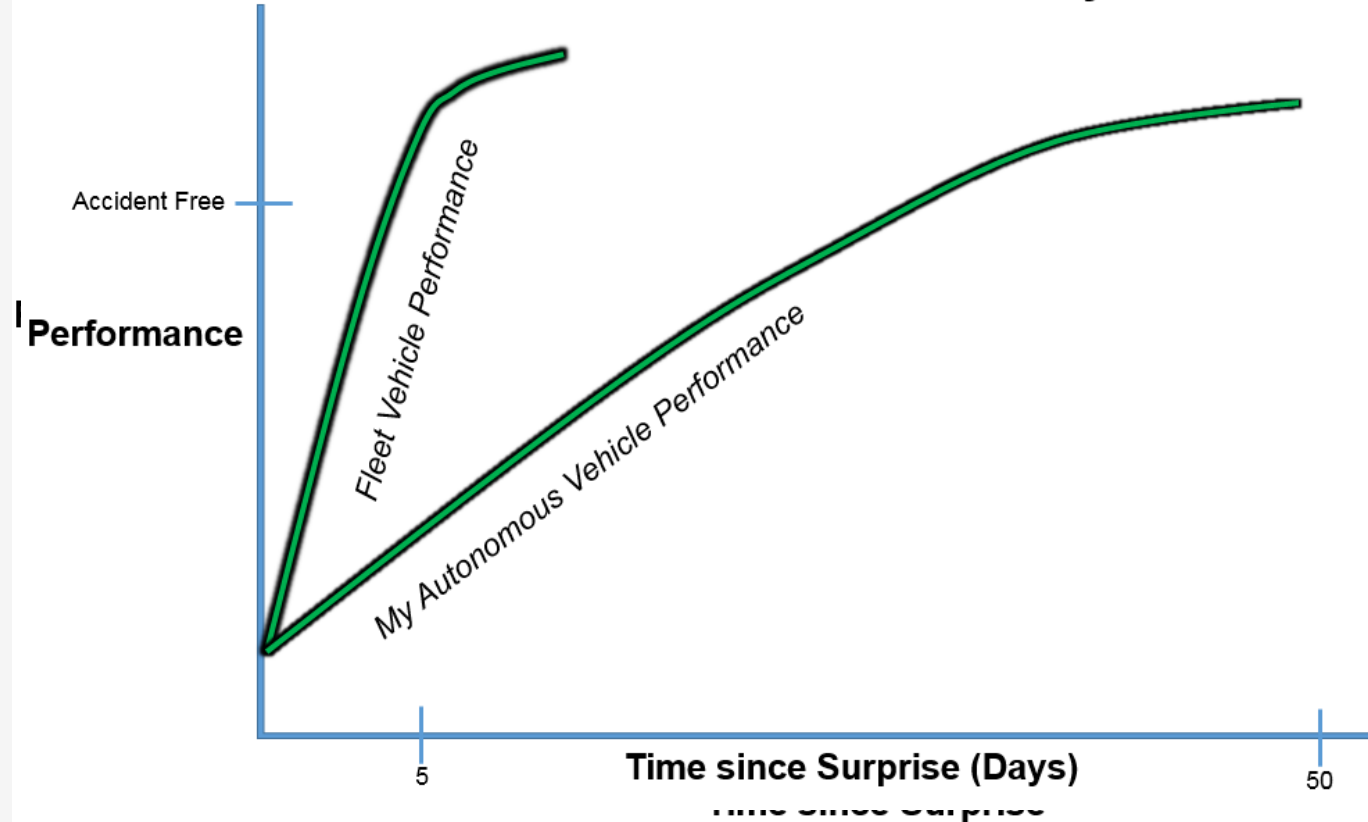


<https://www.hennepin.us/residents/transportation/traffic-signal-report>



HRL WORK: SUPER-TURING EVOLVING LIFELONG LEARNING ARCHITECTURE (STELLAR)



Autonomous Vehicle Fleet on My Commute




TOP-DOWN ATTENTION DRIVES SELECTIVE PLASTICITY





Bottom-up
stimulus-driven
pathway

<https://www.youtube.com/watch?v=eoaK0tBB3Yw>


● Vote for ● Vote against



Top-down signal



Top-down and goal-driven
identification of task-relevant neurons
enable us to selectively reduce the
plasticity of synapses between
important neurons



TOP-DOWN ATTENTION FACILITATES SEQUENTIAL TASK LEARNING

Training



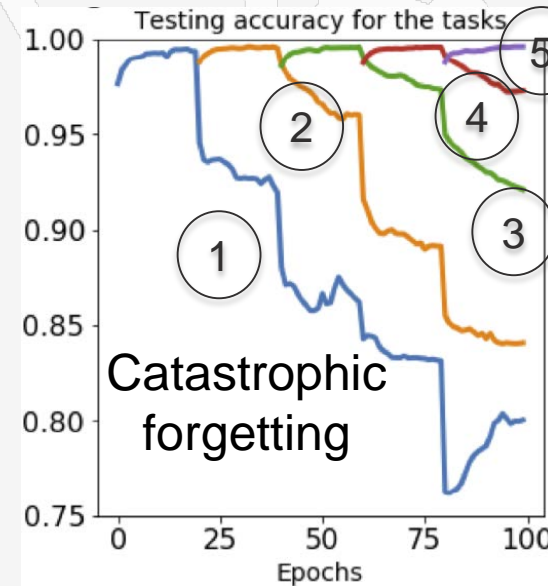
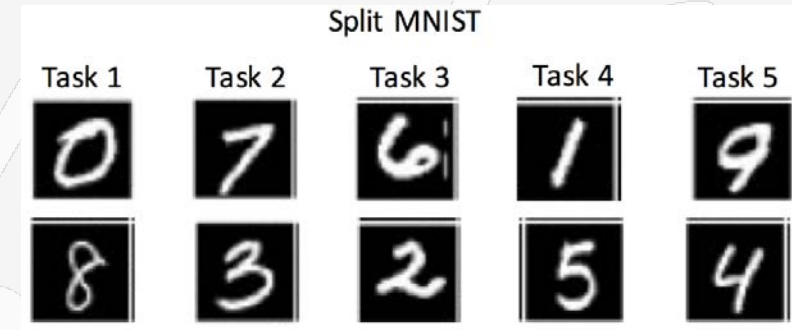
Deployment




<https://navoshta.com/detecting-road-features/>

<https://abcosafety.com/c-63-traffic-work-zone.aspx>

Neural network parameters are overwritten in response to new experiences



Vanilla



Our method

SELF-SUPERVISED LEARNING: CONSOLIDATION OF OLD EXPERIENCES


Perception

Memory
(Hippocampus)

Memory
(Cortex)

Learn to
predict the
future

- Complementary learning systems (hippocampus and cortex)
- Memory system mimics hippocampal circuits to model world dynamics
- Offline cortical replays of prior experiences interleaved with hippocampal replays of recent experiences can solve catastrophic forgetting



<https://navoshta.com/detecting-road-features/>

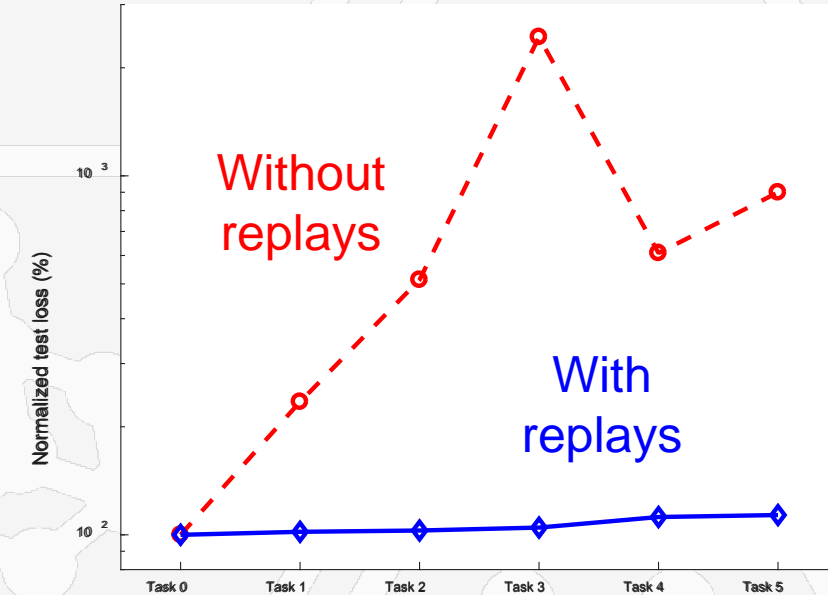




GENERATIVE SEQUENTIAL REPLAYS: LIFELONG LEARNING MEMORY

- 6 Atari games are sequentially experienced
- Memory network is trained to predict the next frame for each game in a self-supervised manner
- Accuracy of next-frame predictions for old games is preserved using interleaved replays

134% increase in loss
(without replays)

<2% increase in loss
(with replays)



Reconstruction of
Task 0
from memory

LIFELONG LEARNING IS CRITICAL FOR AUTONOMY

Attentive car when driving




<https://www.cbc.ca/news/technology/google-autopilot-waymo-1.4379755>

Dreaming car when parked



commons.wikimedia.org/wiki/File:Tesla_Model_S_%26_X_side_by_side_at_the_Gilroy_Supercharger.jpg

Rapidly adapt to and learn from novel, surprising situations
without catastrophic forgetting



ERI **ELECTRONICS RESURGENCE INITIATIVE**

S U M M I T

2019 | Detroit, MI | **July 15 - 17**