CAREER: Towards an Intermittent Learning Framework for Smart and Efficient Cyber-Physical Autonomy

Prof. Kyriakos G. Vamvoudakis (kyriakos@gatech.edu)
Daniel Guggenheim School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA, USA

http://kyriakos.ae.gatech.edu/NSFCareer.html

Intermittent Learning Through Operant Conditioning:

- Convince CPS to learn like humans
- Extending optimal behavior by reducing frequency and predictability of positive reward.
- Ideas from behavioral psychology: B.F. Skinner's Operant

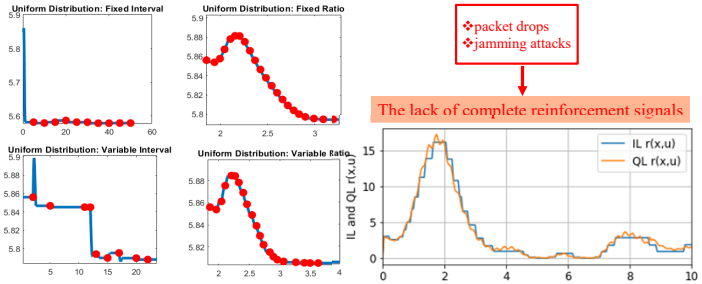

❖packet drops
❖jamming attacks

The lack of complete reinforcement signals

Drawbacks of existing learning mechanisms:

- ➢Continuous data sharing
- ➢Over utilization of communication channels
- ➢Vulnerability to attacks

Intermittent Data Sharing Schemes

- Central nervous system
- Q-learning + Intermittent Learning (IL) for critic + actor
- Error based triggers in update

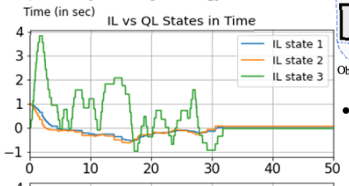

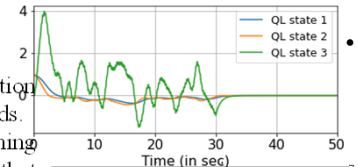

Optimal control with intermittent sensor signals:

- Alleviation of communication and computation burdens via intermittent policies and sparse rewards.
- Combining Operant conditioning with Q-learning yields the Intermittent Learning (IL) framework that incorporate error-based triggers to update built in neural networks.

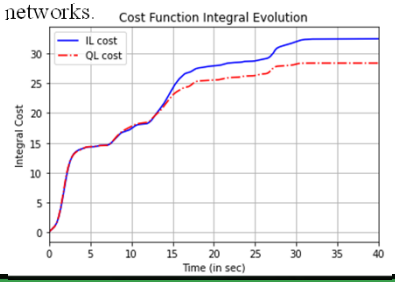

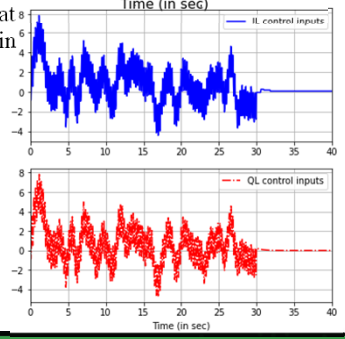

Intermittent Autonomous Path Planning:

- A Q-learning controller for waypoint navigation
- Unknown system dynamics, external disturbances, and by means of intermittent communication between subsystems
- Intermittent transmission of control inputs to minimize communication overhead.
- A relaxed persistence of excitation technique is employed to improve the convergence speed of the Q-learning controller

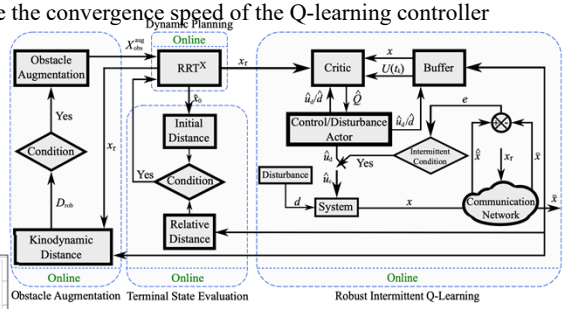

- By limiting the squared norm error with a calculated threshold, the intermittent control policy guarantees stability of the equilibrium point and conserves computational and communication efforts.
- Since neither robust intermittent Q-learning nor RRT-X requires offline computation a priori, this is a fully online framework without any offline computation.

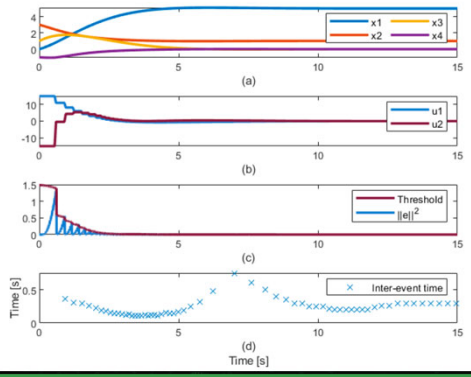

Challenge:

- Revolutionize Reinforcement Learning for CPS to make it less computationally and communicationally expensive
- Augment CPS with the ability to co-exist and interact with humans-in-the-loop

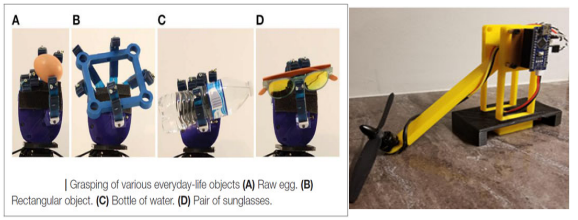

Grasping of various everyday-life objects (A) Raw egg, (B) Rectangular object, (C) Bottle of water, (D) Pair of sunglasses.

Scientific Impact:

- Can we inform autonomous agents to operate in human-centric environments and increase their learning capabilities? Can we be model agnostic when we do so?
- Can we mimic learning in human beings for machine applications?
- Can we provide assured security for IoT like the human brain conserves resources?
- Can we generalize transferrable learning mechanisms between different components of CPS? Can we learn from other agents, human or otherwise, by observation?
- Can we leverage safe policies that are deployed intermittently to enforce secure learning in the presence of adversaries?

Broader Impact:

- Function in the interconnection of learning theory, information theory, control theory, and game theory to introduce a flavor of behavioral perspective to engineering problems in CPS
- Transcend policies developed for CPS to general engineering problems by leveraging model agnostic capabilities.
- Incorporation of a simple, compact and affordable personal control system in undergraduate class in collaboration with Eric Feron
- Design and development of affordable open-source robotic hands as an experimental platform

Solution:

- Intermittent human-like learning in CPS through novel RL adaptations inspired by behaviorism
- Non-equilibrium game theory will be used to model bounded rationality decision making in the context of intermittency

Scholar work and outcomes for year 2021-2022:
See kyriakos.ae.gatech.edu

Award ID#: 1851588