Machine Learning for Cyber Physical Systems: the Good and Bad Uses

Xing Liu, Hansong Xu, Fan Liang, William Grant Hatcher, Weixian Liao and Wei Yu
Cyber Physical Networked System and Security Research Laboratory
Department of Computer and Information Sciences, Towson University
Web: http://wp.towson.edu/wyu Email: wyu@towson.edu
Acknowledgement: NSF CAREER Award - CNS-1350145

Overview

- Cyber-physical Systems (CPS) integrate computing, networking, and control to facilitate smart-world systems.
- Machine Learning schemes, which have proven effective in numerous fields (robotics, automation, prediction, etc.) can be leveraged as intelligent solutions for problems in the complex and dynamic CPS.
- Reinforcement Learning can make precise decisions automatically to maximize cumulative reward via systematic trial-and-error interactions in an unknown environment.

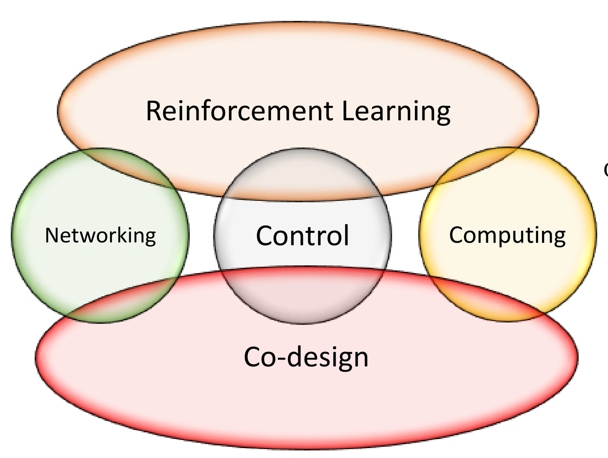

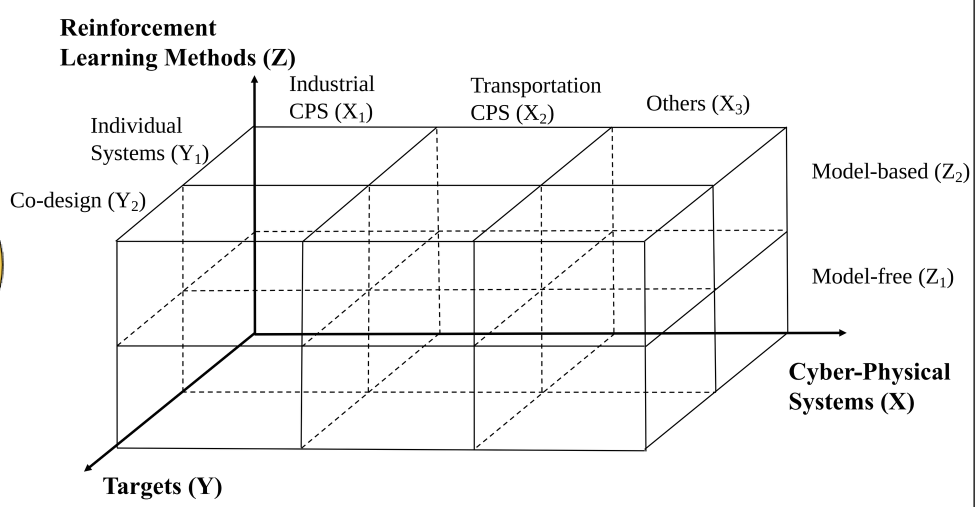

Research Focus

- Investigate existing research works that consider research problems in CPS and apply/adapt reinforcement learning algorithms as solutions.
- Use reinforcement learning algorithms (e.g., Q-learning) to improve the performance of Transportation CPS and Industrial CPS.
- Outline several promising future research directions for reinforcement learning in CPS, as well as machine learning in both good and bad uses.

Our Contributions

- Propose a three-dimensional framework to investigate existing research works in terms of CPS domains, targets, and reinforcement learning methods.
- Conduct two case studies leveraging reinforcement learning to: (1) solve routing efficiency problems in Transportation CPS, and (2) improve control performance in Industrial CPS.
- Outline and recommend research directions toward leveraging reinforcement learning for CPS.
- Investigate good (benign) and bad (risky) uses of machine learning in CPS.

Investigating Reinforcement Learning in CPS

- Reinforcement learning in control subsystems:
  - Controllers continuously supervise industrial systems for process automation, such as the flotation process control in mineral processing.
- Reinforcement learning for networking subsystems:
  - Reinforcement learning has been adopted to address critical problems in networking subsystems, including access management, routing, and resource allocation, among others.
- Reinforcement learning for computing subsystems:
  - Reinforcement learning has been used to solve storage and computing resource allocation problems in dynamic edge computing.
- Reinforcement learning for co-design (i.e., networking-control, networking-computing, control-computing, and networking-control-computing):
  - Reinforcement learning can be leveraged to conduct control and networking co-design to find optimal control and network action pairs while communication and control systems are tightly interacting.

Two Case Studies

- Applying Q-learning to solve the routing efficiency problems in vehicular ad hoc networks in Transportation CPS
  - Q-learning setup:
    - System states contain three parameters (i.e., vehicle distance, velocity difference, and channel bandwidth).
    - Actions are changing the modulation types (e.g., BPSK, QPSK, 16QAM, and 64QAM) for all vehicles in the communication distance.
    - Reward function is defined by the number of hops to the destination and the delivery rate.
  - Simulation setup:
    - 200 vehicles, random data transmission rate, and vehicle velocity between 0-60 (Km/H) in three traffic conditions: *one-way road* (all vehicles move in one direction), *two-way road* (all vehicles move in two (opposite) directions in two lanes), and *cross-road* (vehicle junction area).
- Applying Q-learning to improve control system performance in Industrial CPS
  - Q-learning setup:
    - System states are temperature and trend of temperature change.
    - Actions are changes to the rate of temperature increase/decrease.
    - Reward function is defined by the stability of the physical plant.
  - Physical system (i.e., continuous stirred-tank reactor (CSTR)):
    - A feeder supplies raw material to the reactor.
    - The flow rate of the steam pipe is controlled by the controller to heat the reactor to maintain a target reaction temperature.

Experimental Results

- **Q-learning for Networking**
  - In the top figure, we compare the Packet Delivery Rate of the reinforcement learning-based ad hoc on-demand distance vector (RLAODV), ad hoc on-demand distance vector (AODV), and ad hoc on-demand distance vector and link quality (AODVL) under One-way Road, Two-way Road, and Cross-Road Scenarios.
  - In the bottom figure, we compare End-to-End Delay of RLAODV, AODV, and AODVL under One-way Road, Two-way Road, and Cross-Road Scenarios.

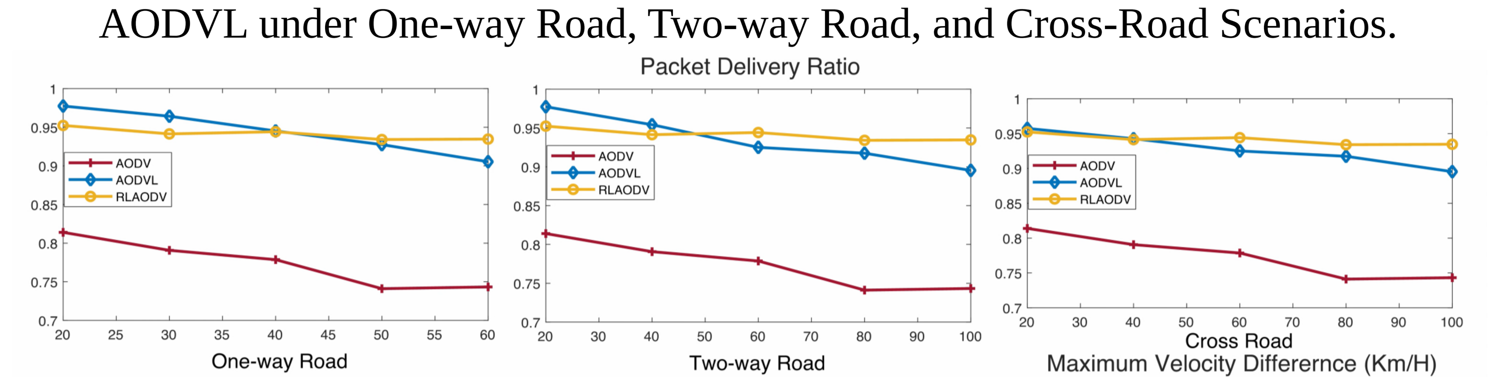

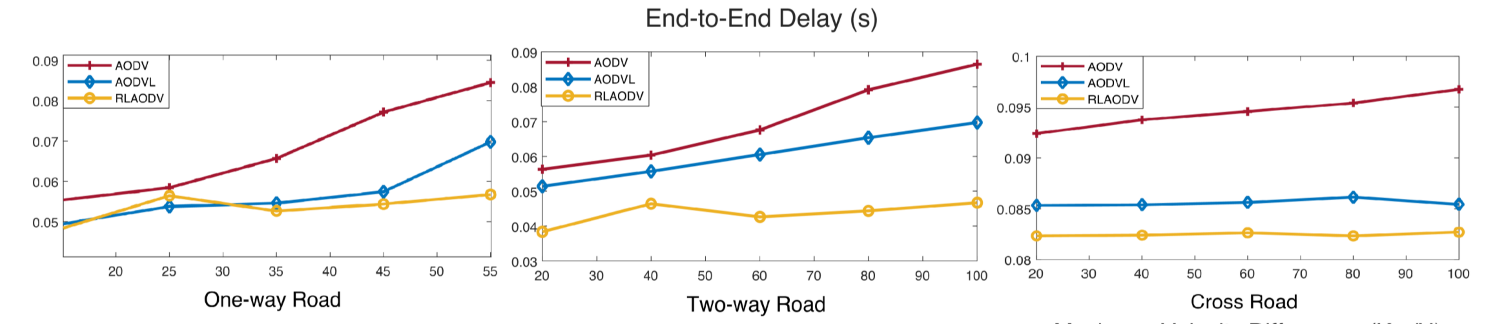

- **Q-learning for Control**

In the figure on the right, we compare control performance of the Feedback Control System, Feedforward Control System, Joint Feedback and Feedforward Control System, and Reinforcement Learning Control System in the subfigures from top to bottom, left to right.

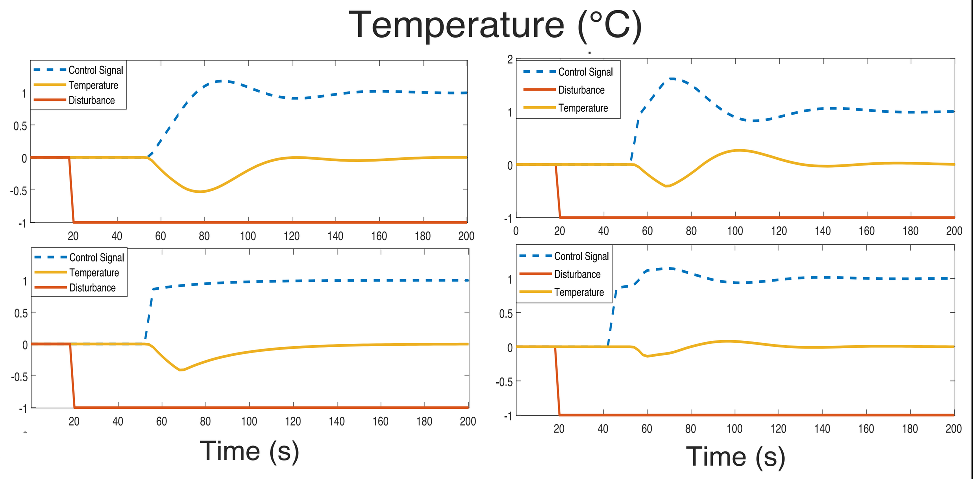

- We implement the reinforcement learning controller for the CSTR as the physical plant in our simulation using MATLAB/Simulink.

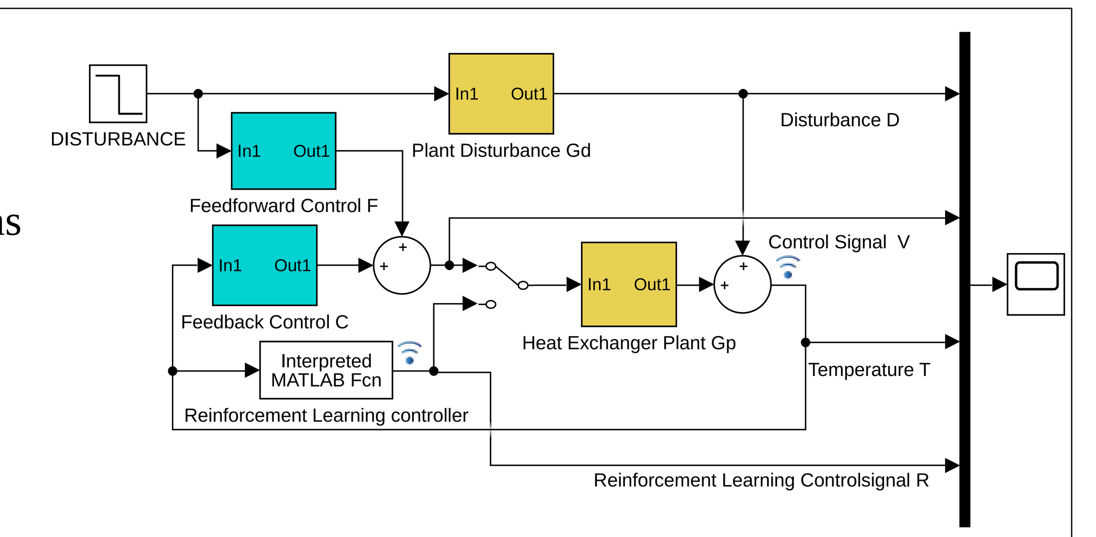

Ongoing Research

- Roadmap of Reinforcement Learning for CPS.
- Good and Bad uses of Machine Learning in CPS
  - Good: using machine learning to improve the performance of CPS.
  - Bad: potential risks for leveraging machine learning in CPS and using machine learning-based techniques to attack CPS.

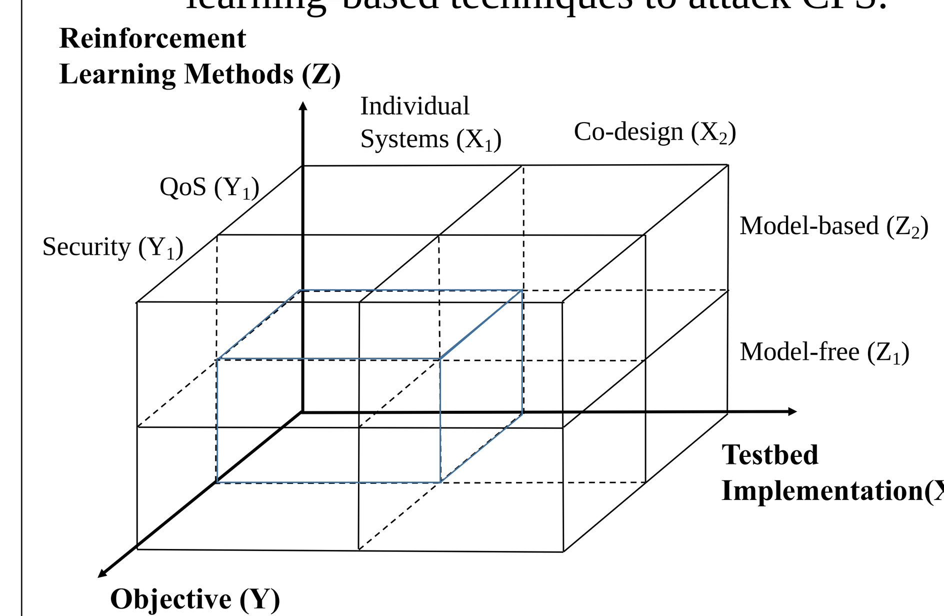

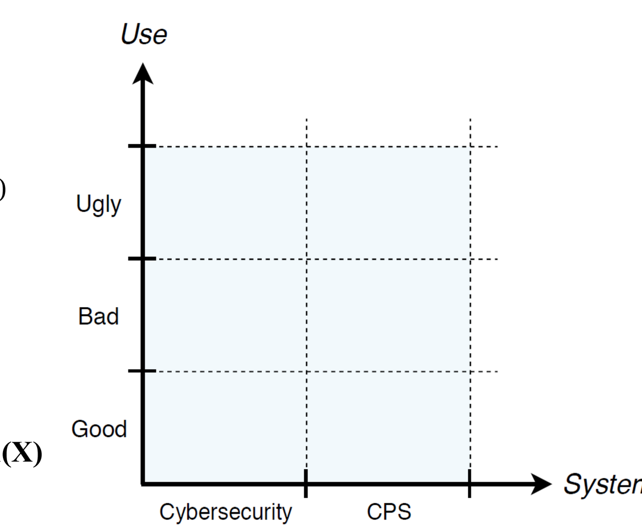

Good and Bad Uses of Machine Learning in CPS

- **Good uses of Machine Learning in CPS**
  - Addressing security issues:
    - Machine learning can assist CPS to address security issues from network, control, and computing perspectives.
  - Challenges of leveraging machine learning in CPS:
    - Extend the applicability of machine learning models to fit different applications.
    - Training process is a black box process and computationally expensive.
  - Improving performance of CPS:
    - Machine learning has the potential to be utilized in a variety of CPS.
    - Machine learning can be leveraged for control, network, and computing subsystems to promote system automation and intelligence.
- **Bad uses of Machine Learning in CPS**
  - Data collection phase: injecting false data to compromise machine learning results and affect the decision-making process of CPS.
  - Data training phase: adversaries can leverage weaknesses in machine learning and tamper with learning subsystems to induce erroneous results.
  - Data testing phase: adversaries can inject false testing datasets to disturb the testing results and obtain fallacious results.

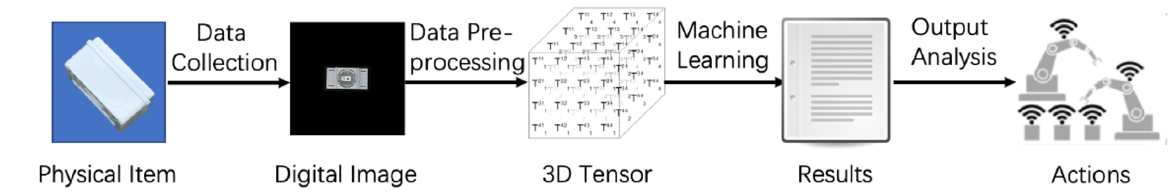

  - Malicious machine learning in CPS:
    - Network: malicious machine learning can attack the network resource scheduling of CPS to affect network performance.
    - Control: adversaries can use malicious machine learning approaches to analyze the state and status of CPS, and control CPS to a weakened state.
    - Computing: adversaries can leverage malicious machine learning to inject additional computing tasks to deteriorate computing performance.