Machine Learning for Robust, Safe and Adaptive Cyber-Physical Energy System

Dileep Kalathil

Department of Electrical and Computer Engineering
Texas A&M University

NSF Workshop on Cyber-enabled Infrastructure to Support Carbon-neutral Electricity and Mobility

A Short Self-Introduction

Dileep Kalathil

dileep.kalathil [at] tamu.edu

I am an assistant professor in the Department of Electrical and Computer Engineering at Texas A&M University. My research is in the areas of Reinforcement Learning and Control Theory, with applications in large scale engineering systems such as power systems, communication networks and mobile robotics. Before joining TAMU, I was a postdoctoral researcher in the EECS department at University of California, Berkeley, working with Prof. Pravin Varaiya and Prof. Kameshwar Poolla. I received my PhD from University of Southern California (USC) in 2014, working with Prof. Rahul Jain.

I co-direct the Learning and Emerging Networked Systems (LENS) Laboratory at TAMU.

Research Area:

- Theory: **reinforcement learning**, control theory, game theory
- Applications: power systems, communication networks, mobile robotics

My Experience in ML + Power Systems

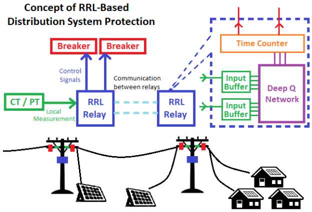

RL for Power System Protection
[Wu et al, 2022], IEEE Open Access Journal of Power and Energy

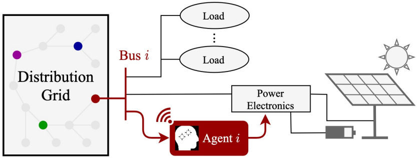

RL for Distribution Systems Voltage Regulation
[El Helou, 2022], IEEE Open Access Journal of Power and Energy

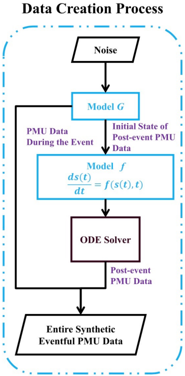

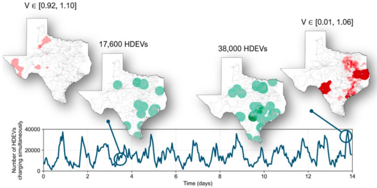

Case study of Heavy-duty Vehicle Electrification
[El Helou et al, 2022], Advances in Applied Energy

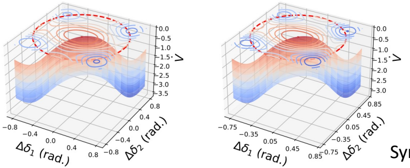

Distributed Learning of Lyapunov Function for Microgrids [Jena et al, 2022], arXiv

Synthetic PMU Data Generation [Zheng et al , 2021], IEEE Open Access Journal of Power and Energy

My Experience in Power Systems

- Sharing economy for energy systems [Kalathil et al, 2019, TSG], [Henriquez-Auba et al, 2021, Applied Energy]
- Mechanism design for demand response [Muthirayan et al, 2020, TSG] [Muthirayan et al, 2021 TSG]
- Learning for demand response [Kalathil et al, Allerton, 2017]

ML in Power Systems

- There exists a number of works on using ML for carbon-neutral energy systems

In collaboration with
BloombergNEF and
Deutsche Energie-Agentur (dena)

WORLD ECONOMIC FORUM

Harnessing Artificial Intelligence to Accelerate the Energy Transition

Machine Learning for Sustainable Energy Systems

Priya L. Donti[1,2] and J. Zico Kolter[1,3]

Energy system digitization in the era of AI: A three-layered approach toward carbon neutrality

Le Xie,[1,*] Tong Huang,[2] Xiangtian Zheng,[1] Yan Liu,[3] Mengdi Wang,[4,5,8] Vijay Vittal,[6] P.R. Kumar,[1] Srinivas Shakkottai,[1] and Yi Cui[7]

ML in Power Systems

- There exists a number of works on using ML for carbon-neutral energy systems

In collaboration with
BloombergNEF and
Deutsche Energie-Agentur (dena)

WORLD ECONOMIC FORUM

Harnessing Artificial Intelligence to Accelerate the Energy Transition

Machine Learning for Sustainable Energy Systems

Priya L. Donti[1,2] and J. Zico Kolter[1,3]

Energy system digitization in the era of AI: A three-layered approach toward carbon neutrality

Le Xie,[1,*] Tong Huang,[2] Xiangtian Zheng,[1] Yan Liu,[3] Mengdi Wang,[4,5,8] Vijay Vittal,[6] P.R. Kumar,[1] Srinivas Shakkottai,[1] and Yi Cui[7]

- Can we use off-the-shelf ML algorithms for cyber-physical energy systems?

ML in Power Systems

- There exists a number of works on using ML for carbon-neutral energy systems

In collaboration with
BloombergNEF and
Deutsche Energie-Agentur (dena)

WORLD ECONOMIC FORUM

Harnessing Artificial Intelligence to Accelerate the Energy Transition

Machine Learning for Sustainable Energy Systems

Priya L. Donti[1,2] and J. Zico Kolter[1,3]

Energy system digitization in the era of AI: A three-layered approach toward carbon neutrality

Le Xie,[1,*] Tong Huang,[2] Xiangtian Zheng,[1] Yan Liu,[3] Mengdi Wang,[4,5,8] Vijay Vittal,[6] P.R. Kumar,[1] Srinivas Shakkottai,[1] and Yi Cui[7]

- Can we use off-the-shelf ML algorithms for cyber-physical energy systems?
- My perspective as an RL researcher:

Reinforcement Learning

- Reinforcement Learning (RL): how to learn the optimal sequence of actions in an unknown and evolving environment to maximize the cumulative long-term reward

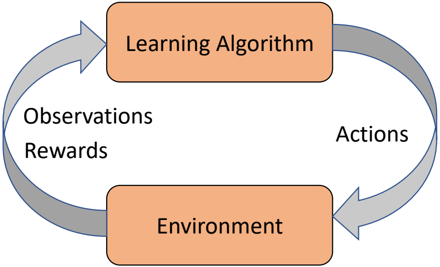

Reinforcement Learning

- Reinforcement Learning (RL): how to learn the optimal sequence of actions in an unknown and evolving environment to maximize the cumulative long-term reward

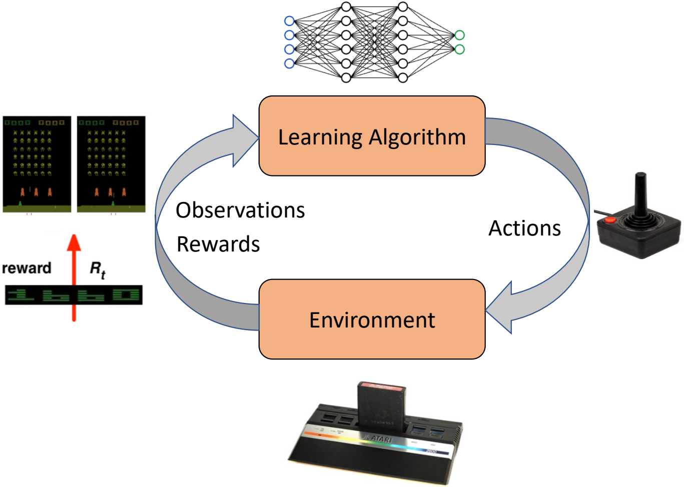

Reinforcement Learning

- Reinforcement Learning (RL): how to learn the optimal sequence of actions in an unknown and evolving environment to maximize the cumulative long-term reward

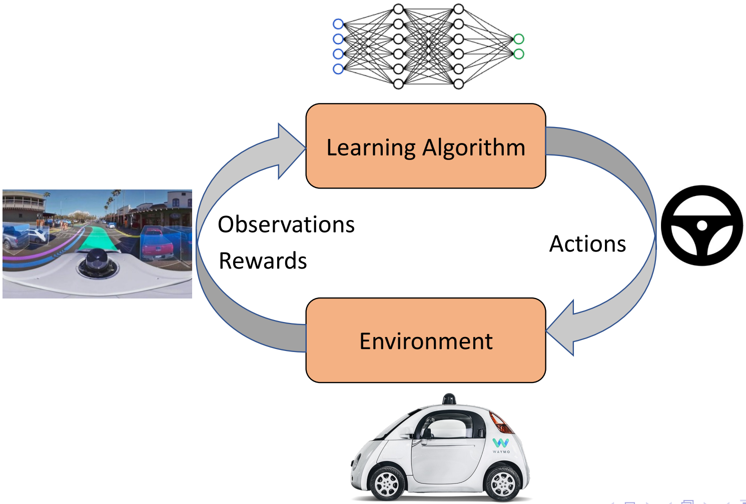

Reinforcement Learning: Shining Successes

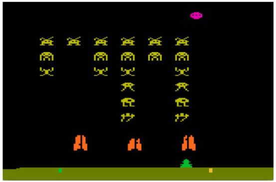

DQN for playing Atari game,
DeepMind (2015)

AlphaGo for playing Go/Chess/Shogi,
DeepMind (2017)

Sensorimotor robotics,
UC Berkeley (2015)

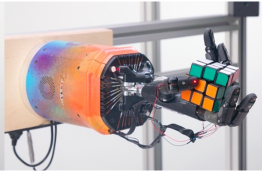

Robotic hand solving Rubik's cube
OpenAI (2019)

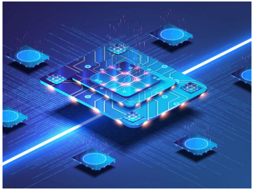

Chip placement design
Google (2020)

Recommendation Systems
Netflix (2022);
Microsoft Vowpal Wabbit (2022)

Reinforcement Learning in the Real-World

- However, most RL success stories are limited to very structured or simulated environments
  - Games, simple robotic settings
- The success stories of RL in real-world engineering systems are rare/limited

Reinforcement Learning in the Real-World

- However, most RL success stories are limited to very structured or simulated environments
  - Games, simple robotic settings
- The success stories of RL in real-world engineering systems are rare/limited
- What is holding up RL from emerging as the go-to solution for the control of real-world engineering systems?

Reinforcement Learning in the Real-World

- However, most RL success stories are limited to very structured or simulated environments
  - Games, simple robotic settings
- The success stories of RL in real-world engineering systems are rare/limited
- What is holding up RL from emerging as the go-to solution for the control of real-world engineering systems?
- RL algorithms are not resilient (lacks robustness, safety and adaptability guarantees!)

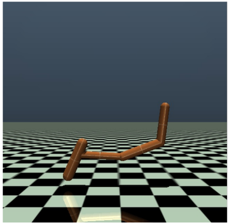

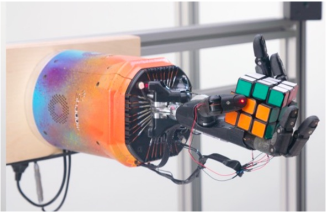

Robotic hand solving Rubik's cube, OpenAI (2019)
Only 32% success!

Reinforcement Learning in the Real-World

- However, most RL success stories are limited to very structured or simulated environments
  - Games, simple robotic settings
- The success stories of RL in real-world engineering systems are rare/limited
- What is holding up RL from emerging as the go-to solution for the control of real-world engineering systems?
- RL algorithms are not resilient (lacks robustness, safety and adaptability guarantees!)
- Naively using ML/RL algorithms can lead to catastrophic failures!

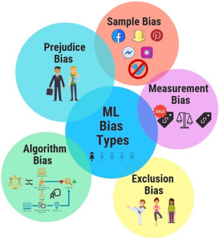

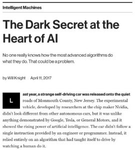
Intelligent Machines

The Dark Secret at the Heart of AI

No one really knows how the most advanced algorithms do what they do. That could be a problem.

by Will Knight April 11, 2017

Last year, a strange self-driving car was released onto the quiet roads of Monmouth County, New Jersey. The experimental vehicle, developed by researchers at the chip maker Nvidia, didn't look different from other autonomous cars, but it was unlike anything demonstrated by Google, Tesla, or General Motors, and it showed the rising power of artificial intelligence. The car didn't follow a single instruction provided by an engineer or programmer. Instead, it relied entirely on an algorithm that had taught itself to drive by watching a human do it.

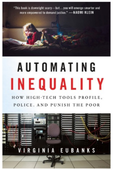

Why Do We Need Resilient RL?

simulation-to-reality (sim-to-real) gap

Why Do We Need Resilient RL?

simulation-to-reality (sim-to-real) gap

- RL algorithms typically learn a policy by training on a simulator
- However, the real-world systems parameters can be different from that of the simulator
  - Due to the approximation errors incurred while modeling, changes in the real-world parameters over time, adversarial disturbances in the real-world

Why Do We Need Resilient RL?

simulation-to-reality (sim-to-real) gap

- RL algorithms typically learn a policy by training on a simulator
- However, the real-world systems parameters can be different from that of the simulator
  - Due to the approximation errors incurred while modeling, changes in the real-world parameters over time, adversarial disturbances in the real-world
  - For example, in robotics, mass, friction, sensor noise, action delays etc. can be different between the real-world system and the simulator (which represents the real-world system)

Why Do We Need Resilient RL?

simulation-to-reality (sim-to-real) gap

- RL algorithms typically learn a policy by training on a simulator
- However, the real-world systems parameters can be different from that of the simulator
  - Due to the approximation errors incurred while modeling, changes in the real-world parameters over time, adversarial disturbances in the real-world
  - For example, in robotics, mass, friction, sensor noise, action delays etc. can be different between the real-world system and the simulator (which represents the real-world system)
- RL algorithms trained on a simulator may perform poorly in the real-world systems due to sim-to-real gap [Tobin, et al, 2017] (Panaganti and **K.**, *AISTATS, 2022)*

Why Do We Need Resilient RL?

simulation-to-reality (sim-to-real) gap

- RL algorithms typically learn a policy by training on a simulator
- However, the real-world systems parameters can be different from that of the simulator
  - Due to the approximation errors incurred while modeling, changes in the real-world parameters over time, adversarial disturbances in the real-world
  - For example, in robotics, mass, friction, sensor noise, action delays etc. can be different between the real-world system and the simulator (which represents the real-world system)
- RL algorithms trained on a simulator may perform poorly in the real-world systems due to sim-to-real gap [Tobin, et al, 2017] (Panaganti and K., *AISTATS, 2022)*

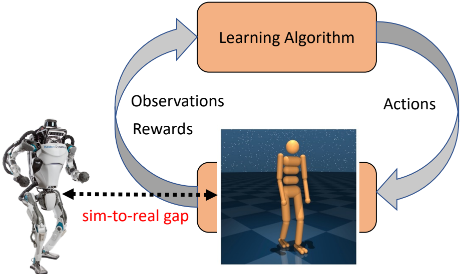

Three RL Challenges for Cyber-Physical Energy Systems

- Robustness: Algorithm must be robust against: the parameter mismatches between the simulator model and real-world system system, adversarial disturbances in real-world system, noisy or partial observation, adversarial attacks, ...

Three RL Challenges for Cyber-Physical Energy Systems

- Robustness: Algorithm must be robust against: the parameter mismatches between the simulator model and real-world system system, adversarial disturbances in real-world system, noisy or partial observation, adversarial attacks, . . .
  - Preliminary studies from my group: [Xu et al, 2023, AISTATS] [Panaganti et al, 2022, NeurIPS], [Panaganti et al, 2021, ICML]

Three RL Challenges for Cyber-Physical Energy Systems

- Robustness: Algorithm must be robust against: the parameter mismatches between the simulator model and real-world system system, adversarial disturbances in real-world system, noisy or partial observation, adversarial attacks, . . .
  - Preliminary studies from my group: [Xu et al, 2023, AISTATS] [Panaganti et al, 2022, NeurIPS], [Panaganti et al, 2021, ICML]
- Safety: Algorithm should maintain the necessary safety constraints during learning and deployment; algorithm should satisfy the stability criteria, . . .

Three RL Challenges for Cyber-Physical Energy Systems

- Robustness: Algorithm must be robust against: the parameter mismatches between the simulator model and real-world system system, adversarial disturbances in real-world system, noisy or partial observation, adversarial attacks, ...
  - Preliminary studies from my group: [Xu et al, 2023, AISTATS] [Panaganti et al, 2022, NeurIPS], [Panaganti et al, 2021, ICML]
- Safety: Algorithm should maintain the necessary safety constraints during learning and deployment; algorithm should satisfy the stability criteria, ...
  - Preliminary studies from my group: [Bura et al, 2022, NeurIPS], [Liu et al, 2021, NeurIPS]

Three RL Challenges for Cyber-Physical Energy Systems

- Robustness: Algorithm must be robust against: the parameter mismatches between the simulator model and real-world system system, adversarial disturbances in real-world system, noisy or partial observation, adversarial attacks, ...
  - Preliminary studies from my group: [Xu et al, 2023, AISTATS] [Panaganti et al, 2022, NeurIPS], [Panaganti et al, 2021, ICML]
- Safety: Algorithm should maintain the necessary safety constraints during learning and deployment; algorithm should satisfy the stability criteria, ...
  - Preliminary studies from my group: [Bura et al, 2022, NeurIPS], [Liu et al, 2021, NeurIPS]
- Adaptability and Scalability: Algorithm should be able to rapidly adapt to new/changing environments; should be able to scale to large-scale systems comprised of many interacting subsystems; should be able to exploit the historical operational data and known models

Three RL Challenges for Cyber-Physical Energy Systems

- Robustness: Algorithm must be robust against: the parameter mismatches between the simulator model and real-world system system, adversarial disturbances in real-world system, noisy or partial observation, adversarial attacks, ...
  - Preliminary studies from my group: [Xu et al, 2023, AISTATS] [Panaganti et al, 2022, NeurIPS], [Panaganti et al, 2021, ICML]
- Safety: Algorithm should maintain the necessary safety constraints during learning and deployment; algorithm should satisfy the stability criteria, ...
  - Preliminary studies from my group: [Bura et al, 2022, NeurIPS], [Liu et al, 2021, NeurIPS]
- Adaptability and Scalability: Algorithm should be able to rapidly adapt to new/changing environments; should be able to scale to large-scale systems comprised of many interacting subsystems; should be able to exploit the historical operational data and known models
  - Preliminary studies from my group: [Rengarajan et al, 2022, NeurIPS] [Rengarajan et al, 2022, ICLR], [Jena et al, 2022, arXiv]

Three RL Challenges for Cyber-Physical Energy Systems

- Robustness: Algorithm must be robust against: the parameter mismatches between the simulator model and real-world system system, adversarial disturbances in real-world system, noisy or partial observation, adversarial attacks, . . .
  - Preliminary studies from my group: [Xu et al, 2023, AISTATS] [Panaganti et al, 2022, NeurIPS], [Panaganti et al, 2021, ICML]
- Safety: Algorithm should maintain the necessary safety constraints during learning and deployment; algorithm should satisfy the stability criteria, . . .
  - Preliminary studies from my group: [Bura et al, 2022, NeurIPS], [Liu et al, 2021, NeurIPS]
- Adaptability and Scalability: Algorithm should be able to rapidly adapt to new/changing environments; should be able to scale to large-scale systems comprised of many interacting subsystems; should be able to exploit the historical operational data and known models
  - Preliminary studies from my group: [Rengarajan et al, 2022, NeurIPS] [Rengarajan et al, 2022, ICLR], [Jena et al, 2022, arXiv]

Three RL Challenges for Cyber-Physical Energy Systems

- Robustness: Algorithm must be robust against: the parameter mismatches between the simulator model and real-world system system, adversarial disturbances in real-world system, noisy or partial observation, adversarial attacks, ...
  - Preliminary studies from my group: [Xu et al, 2023, AISTATS] [Panaganti et al, 2022, NeurIPS], [Panaganti et al, 2021, ICML]
- Safety: Algorithm should maintain the necessary safety constraints during learning and deployment; algorithm should satisfy the stability criteria, ...
  - Preliminary studies from my group: [Bura et al, 2022, NeurIPS], [Liu et al, 2021, NeurIPS]
- Adaptability and Scalability: Algorithm should be able to rapidly adapt to new/changing environments; should be able to scale to large-scale systems comprised of many interacting subsystems; should be able to exploit the historical operational data and known models
  - Preliminary studies from my group: [Rengarajan et al, 2022, NeurIPS] [Rengarajan et al, 2022, ICLR], [Jena et al, 2022, arXiv]

How do we develop scalable RL algorithms that are provably robust, safe and adaptive for real-world cyber-physical engineering systems?