CAREER: Learning for Generalization in Large-Scale Cyber-Physical Systems

PI: Cathy Wu (MIT)

Take a picture to download the full papers

Idea: Address **heterogeneity** in transportation CPS with **learning+models**

BROADER IMPACT

- **Heterogeneity** is fundamental to the design & analysis of **future transportation systems.**

Source of heterogeneity	Examples
Diverse stakeholders	Pedestrian, car, truck; local, state, federal
Geographical contexts	Rural, urban; OECD, non-OECD
Rich objectives	Climate, equity, safety, congestion
Emerging technology	Connectivity, automation, electrification

BACKGROUND

- Under heterogeneity, **existing methods often fail**, either finding good solutions slowly (**inefficient**) or not finding them at all (**non-robust**) [1, 2, 7].
- **Hybridize** to achieve the best of both worlds?

Method class	Robustness	Efficiency
Model-based	✓	X
Model-free	X	✓
Hybrid	✓?	✓?

TECHNICAL APPROACH

- Use a **Contextual Markov Decision Process** (cMDP) to capture **heterogeneity** within a problem class
- cMDP: $\mathcal{M} := \{M_c\}_{c\sim p_C} = \{(\mathcal{S}_c, \mathcal{A}_c, \mathcal{T}_c, r_c, \rho_c)\}_{c\sim p_C}$
- **Aim: Design efficient & robust hybridized methods.**

SCIENTIFIC IMPACT

- Address **increasing system complexity** in CPS.
- Bridge gap between model-based & model-free methods.

Strategy 1: Learning-guided Optimization [3-5]

Problem

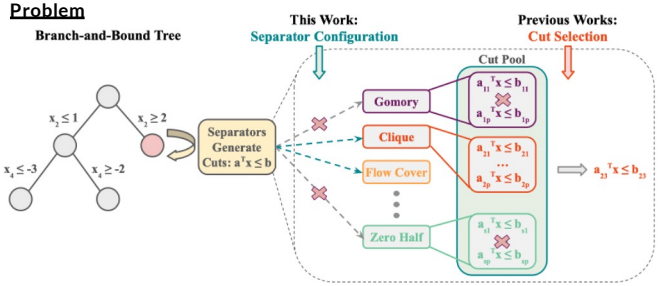

Our work: "Pre-train" classical solvers.
Results: Accelerate SCIP by up to 70% and Gurobi by up to 55%.
Applications: Bin Packing, Max Cut, Facility Location, MIPLIB, etc.

Method

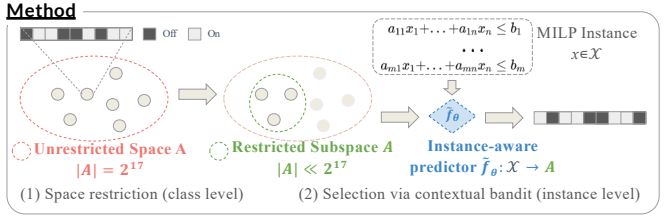

Figure: Learning-to-separate method

Architecture

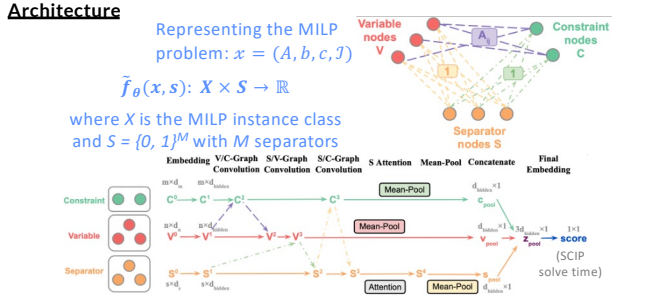

Strategy 2: Model-based Transfer Learning [7]

Problem

Training performance, three trials (rand. seeds), default hyperparameters

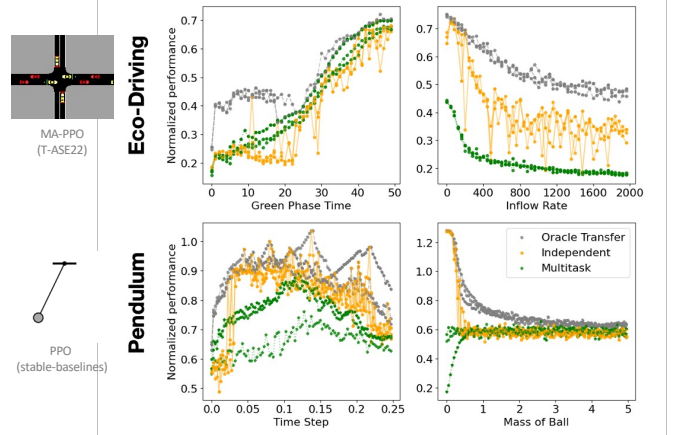

Our work: Approximate Oracle using Bayesian Optimization.
Results: 10-40x more sample efficient.
Applications: Continuous control, Multi-agent traffic control

Method

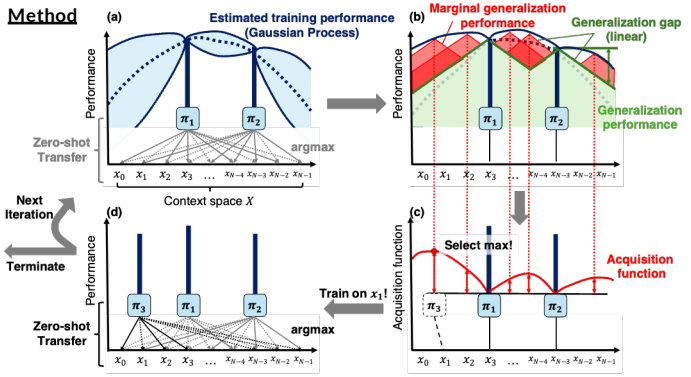

REFERENCES

[1] F. Lamnabhi-Lagarrigue, et al. "Systems & control for the future of humanity, research agenda: Current and future roles, impact and grand challenges," Annual Reviews in Control, 2017.
[2] V. Jayawardana, et al., and **C. Wu**, "The impact of task underspecification in evaluating deep reinforcement learning," in Advances in neural information processing systems (NeurIPS), 2022.
[3] S. Li*, Z. Yan*, and **C. Wu**, "Learning to delegate for large-scale vehicle routing," Advances in Neural Information Processing Systems (NeurIPS), 2021. **Spotlight (<3%)**.
[4] Z. Yan and **C. Wu**, "Neural Neighborhood Search for Multi-agent Path Finding," International Conference on Learning Representations (ICLR), 2024.
[5] S. Li*, W. Ouyang*, M. Paulus, and **C. Wu**, "Learning to Configure Separators in Branch-and-Cut," in Advances in Neural Information Processing Systems (NeurIPS), 2023.
[6] V. Jayawardana, S. Li, **C. Wu**, Y. Farid, and K. Oguchi, "Generalizing Eco-Lagrangian Control via Multi-residual Task Learning," IEEE International Conference on Robotics and Automation (ICRA), 2024.
[7] J.-H. Cho, V. Jayawardana, S. Li, and **C. Wu**, "Model-Based Transfer Learning for Contextual Reinforcement Learning," in Advances in Neural Information Processing Systems (NeurIPS), 2024.

Award ID#: CNS - 2239566