Scalable Reinforcement Learning for Large-Scale Coordination of Electric Vehicles

Stavros Orfanoudakis, Pedro Vergara

Delft, 7-11-2024

TU Delft

IEPG Intelligent Electrical Power Grids

Projected EV Number

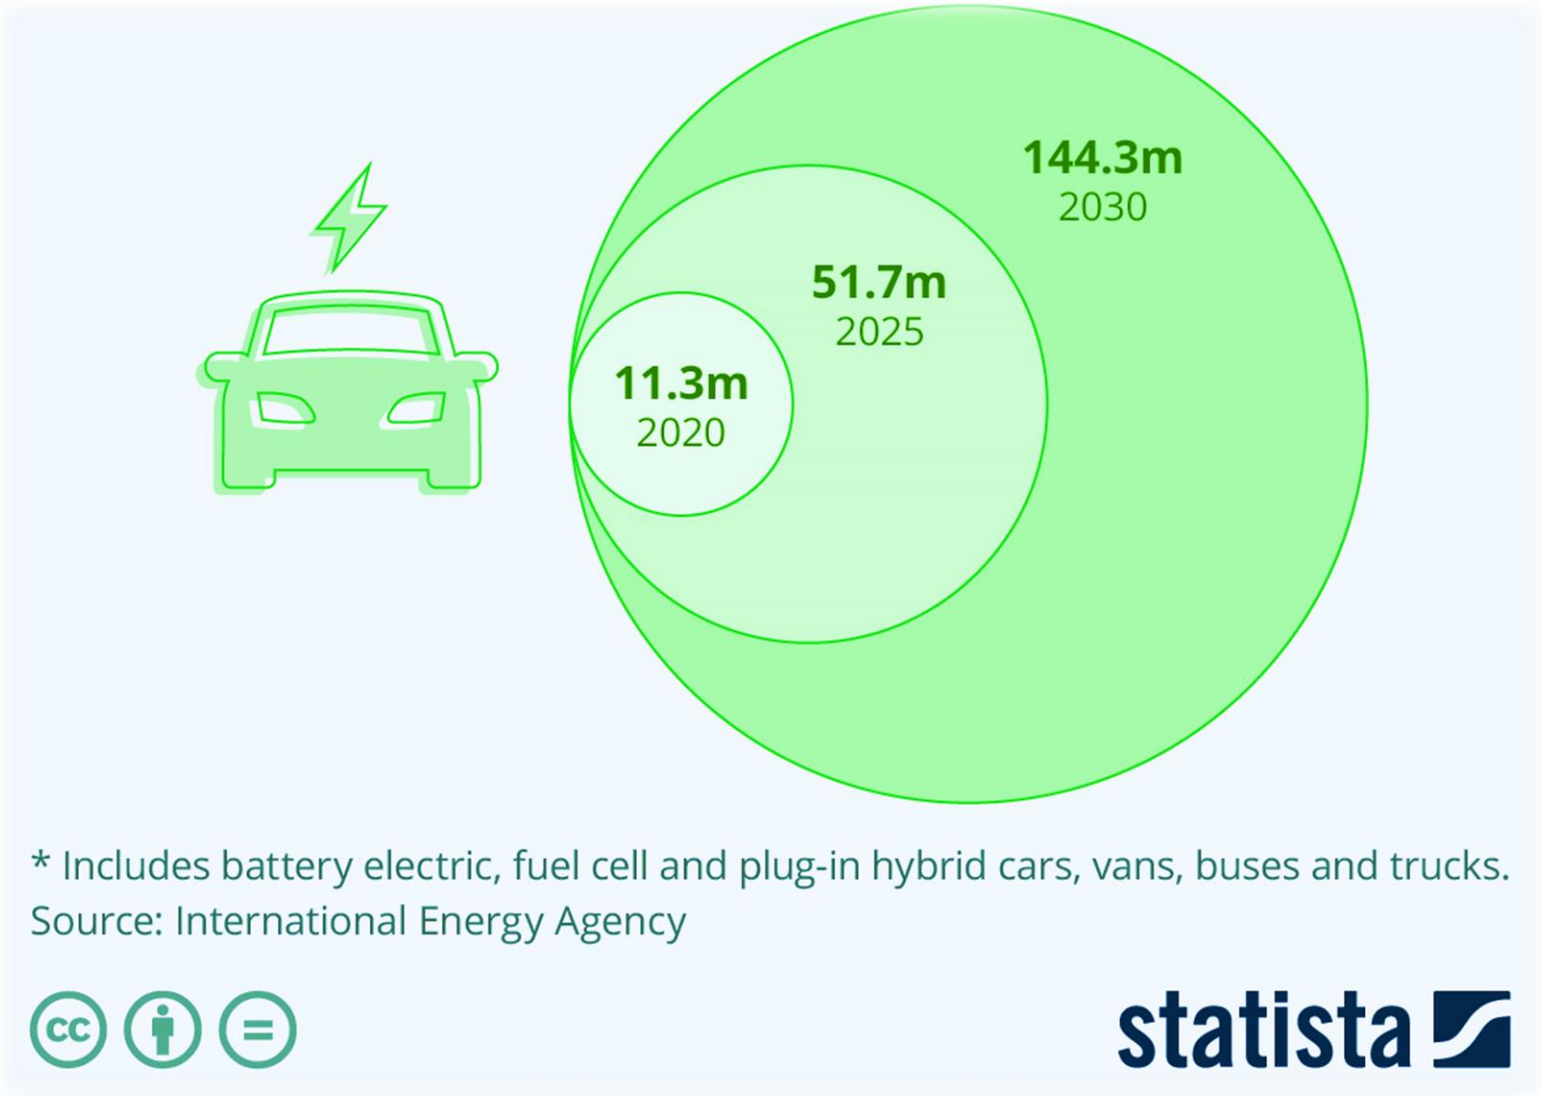

What happens to the electricity grid with many thousand EVs?

TU Delft

Huge Energy Demand Peaks!

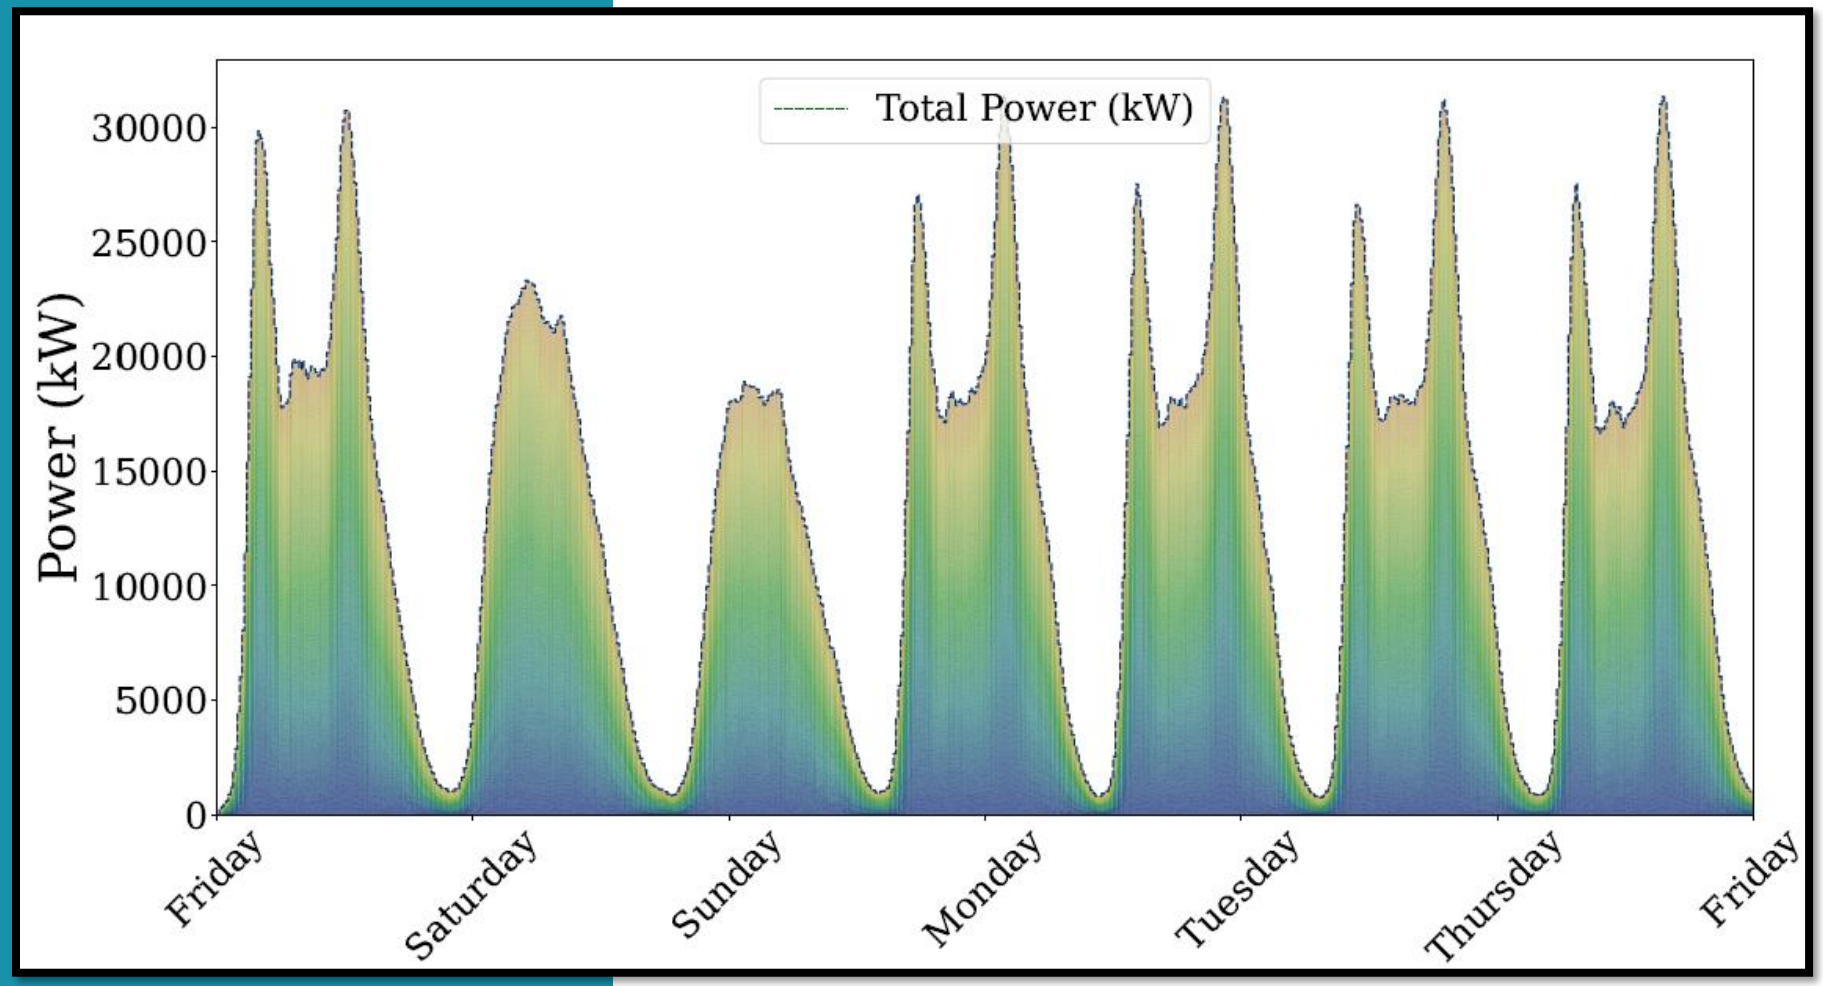

[2] Elaadnl open datasets for electric mobility research | update april 2020, https://platform.elaad.io/analyses/ElaadNL_ opendata.php, Accessed: [23/11/2023].

The Solution

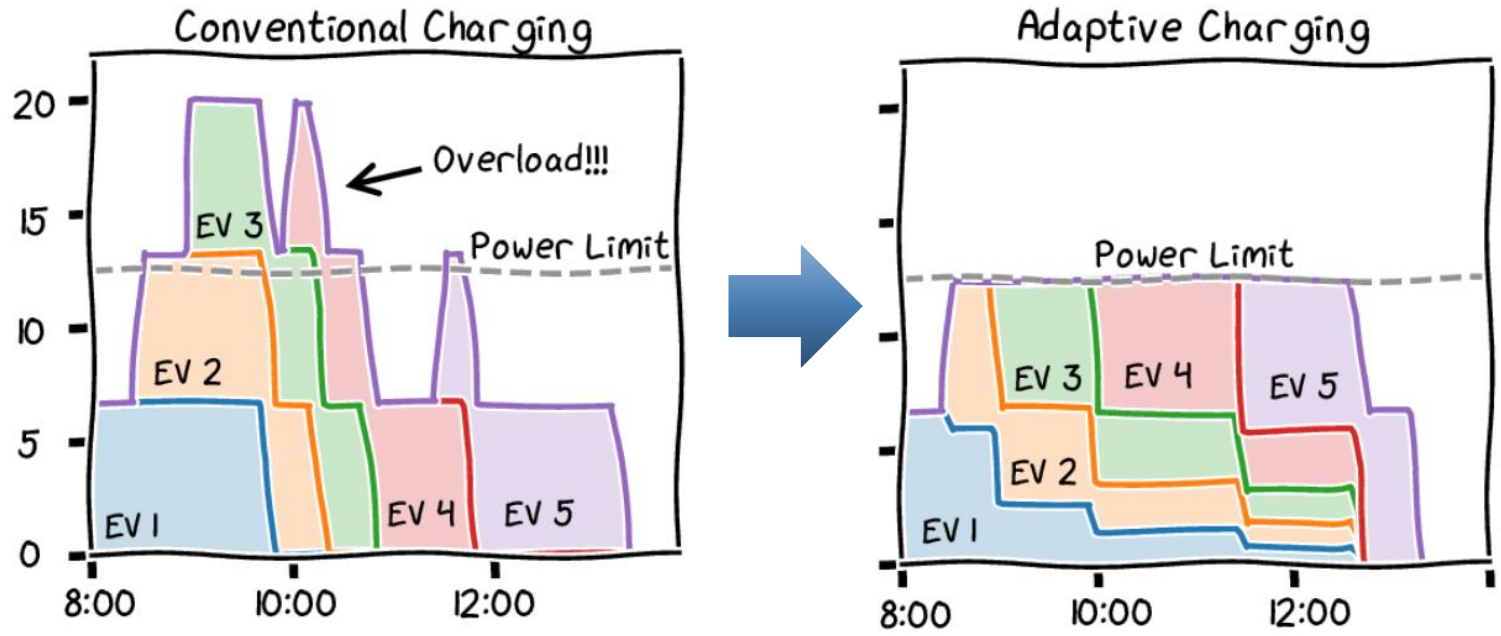

Source of Uncertainty

How long is the EV staying connected?

Communication with DSO?

What is the type of EV and the state of charge?

[1] https://ev.caltech.edu/info

Simulation

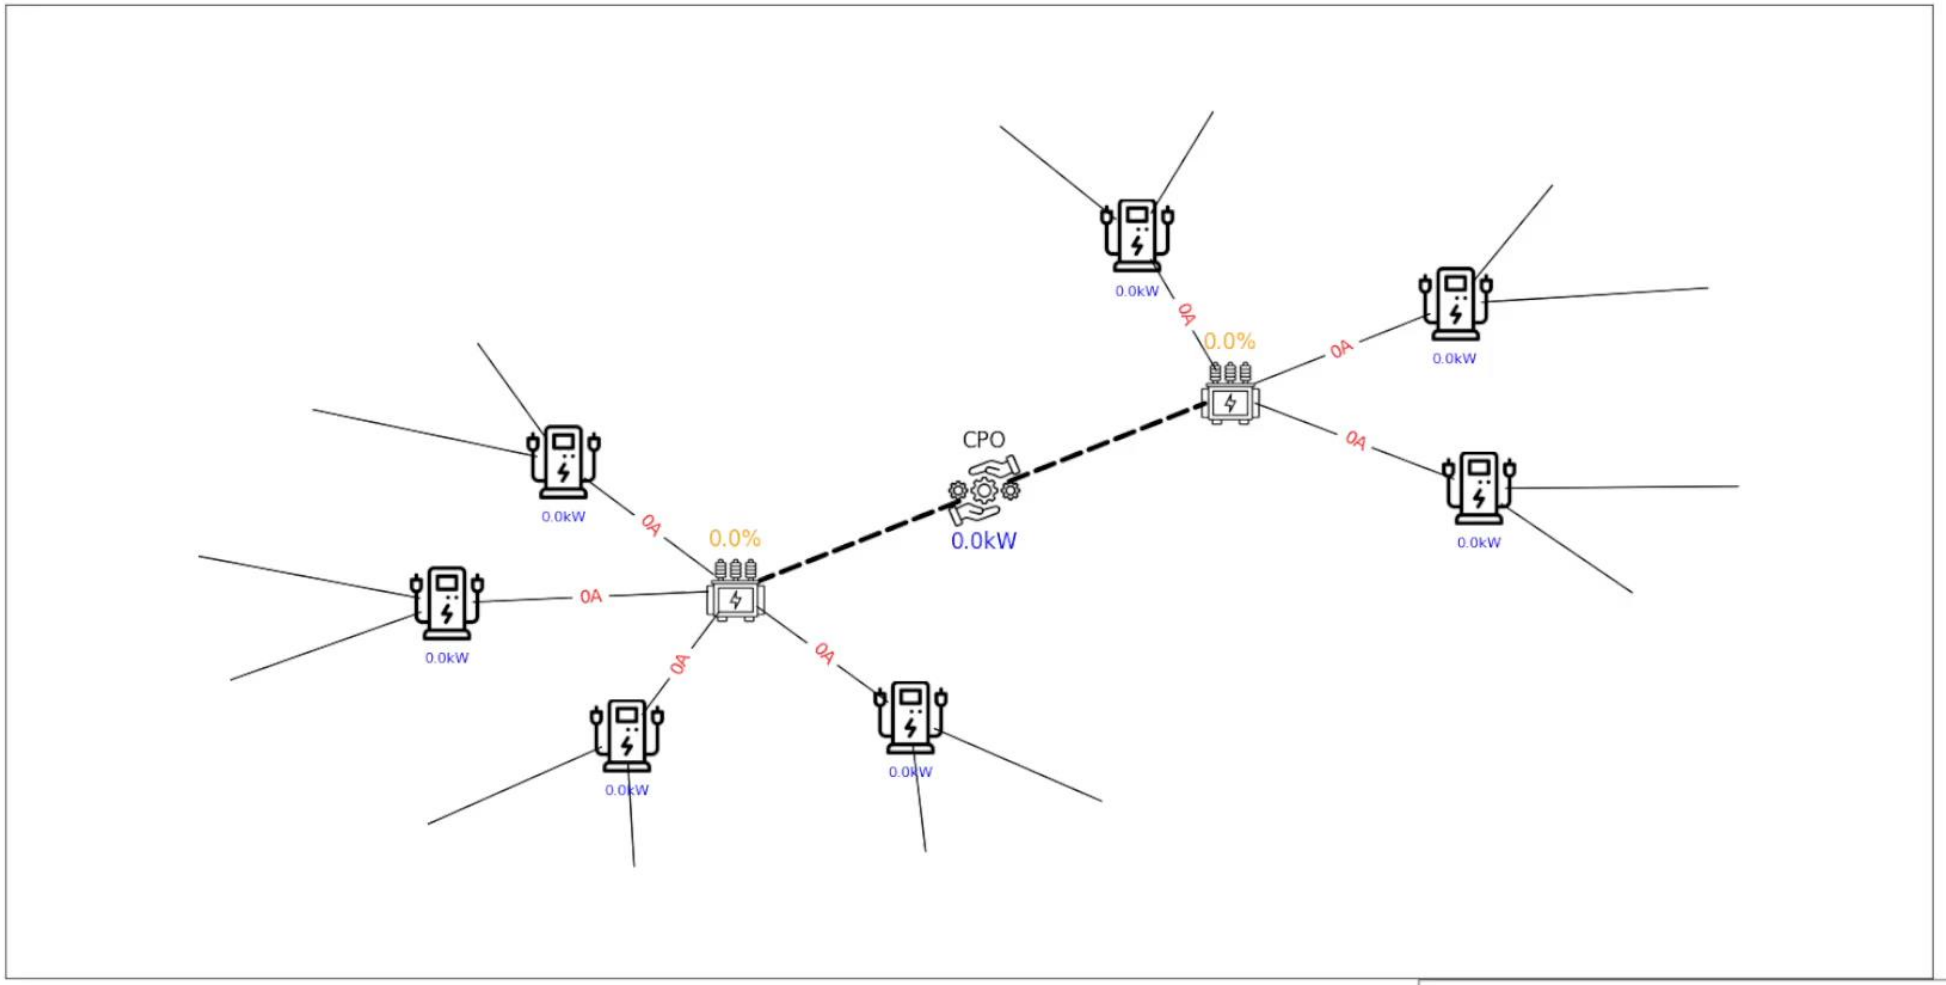

EV2GYM
A REALISTIC V2G GYM SIMULATOR

https://github.com/StavrosOrf/EVsSimulator

Time step: 5 / 96
2022-10-13 06:15:00
Scenario: public_PowerSetpointTracking
Simulation Name: ev_city_96_2024-01-09_21-28-08-110206_replay

Modeling the problem

A vector with many variables

$$S_t = [t, P_t^{set}, P_{t-1}^{tot}] \cup [d_{j,i}, E_{j,i,t} - E_{j,i}^{arr}, t - t_{j,i}^{arr}] \quad \forall j, \forall i$$

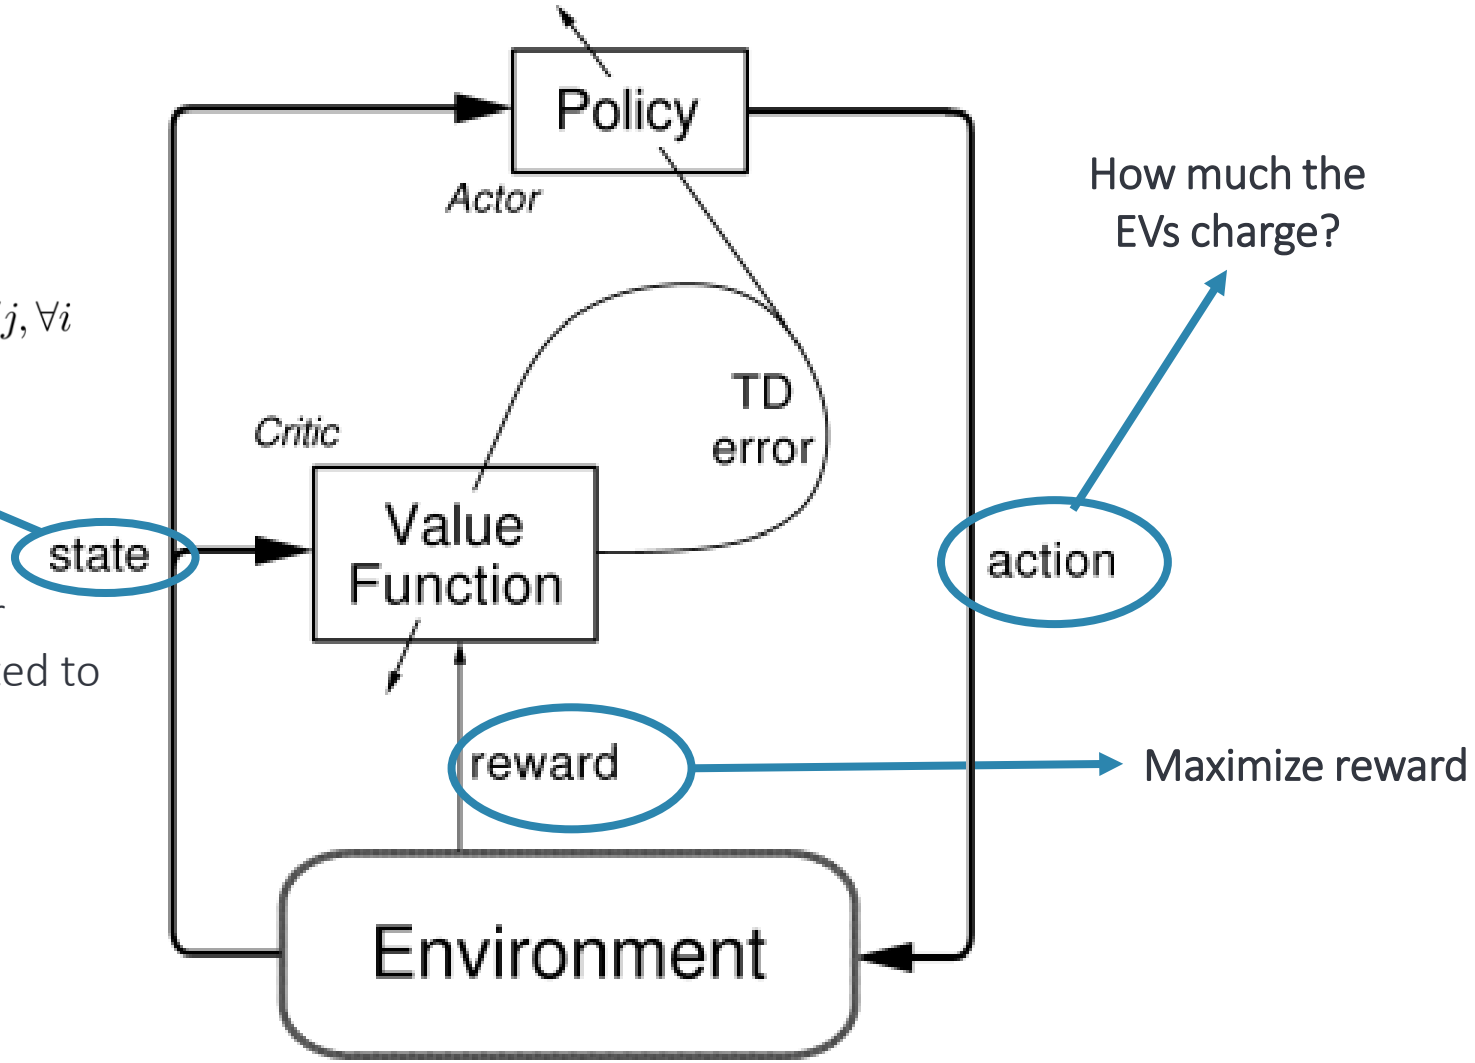

Limitations

- **Scales linearly** with the EV charge number
- Uses zeros when there are no EV connected to a charger **(dynamic size state)**
- Produce actions even if there is **no EV connected**

Reinterpreting EV charging

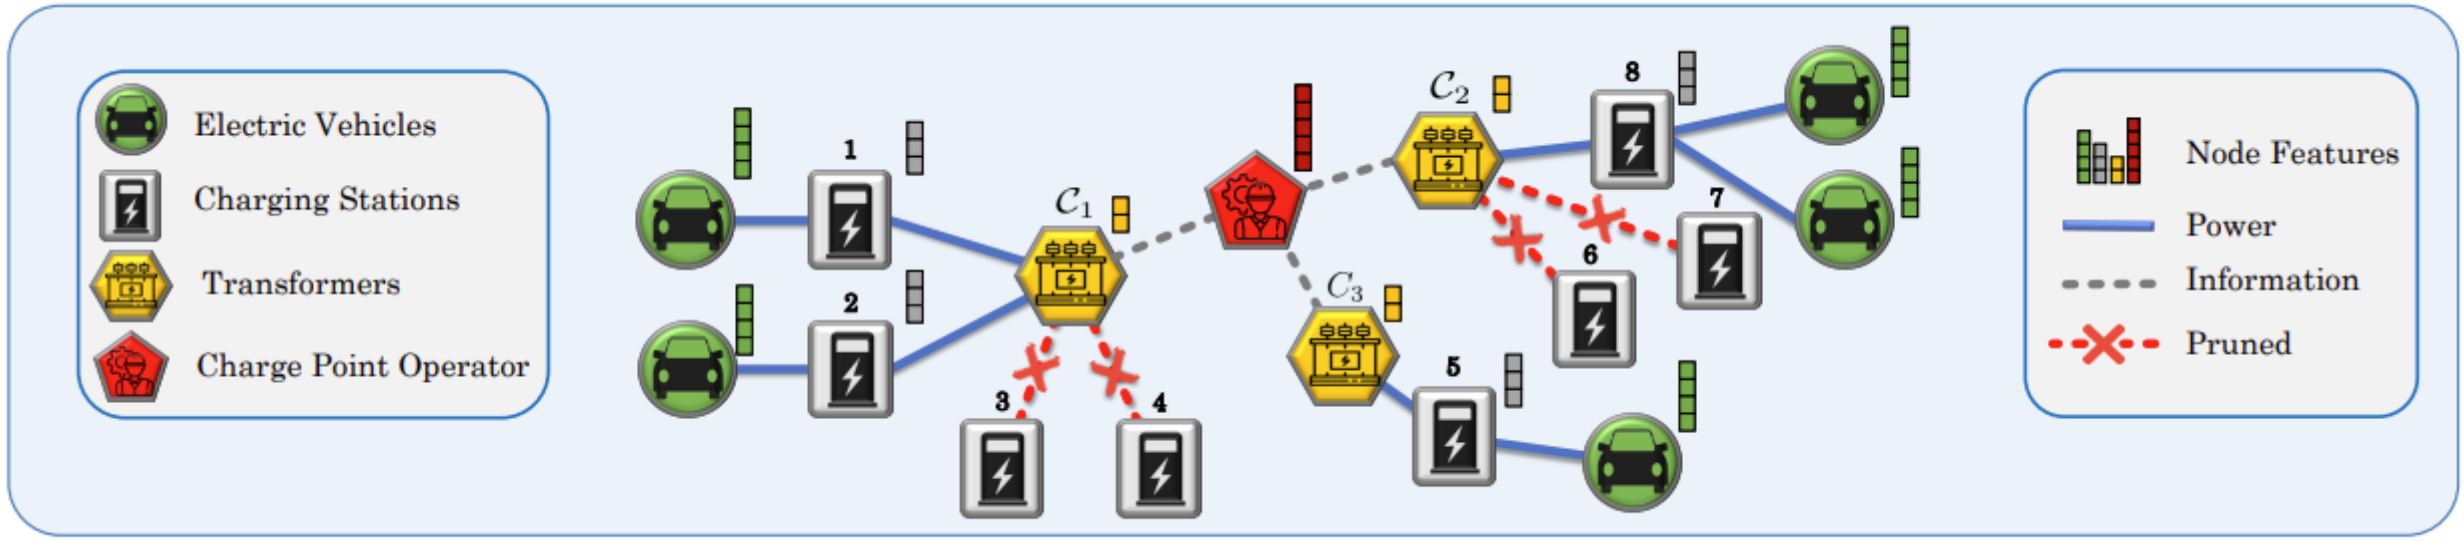

a. Graph Problem Structure

EV-GNN Architecture

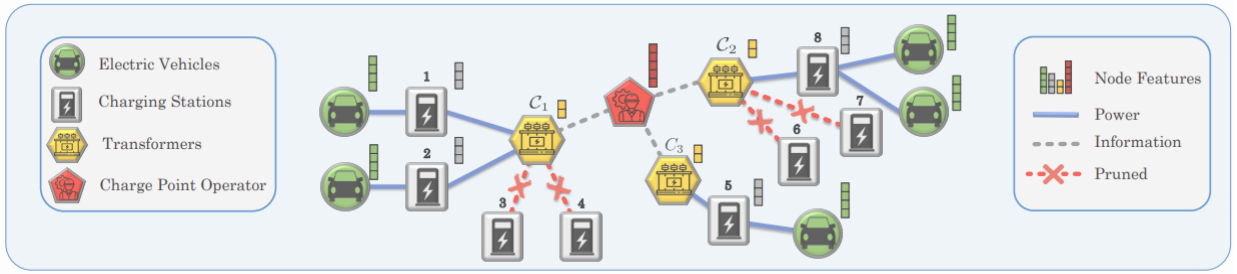

a. Graph Problem Structure

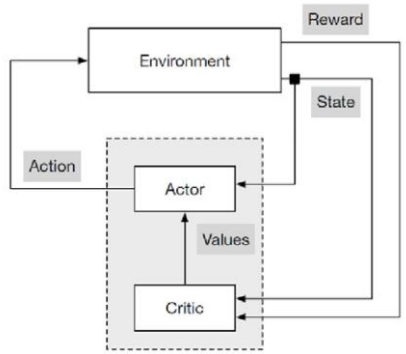

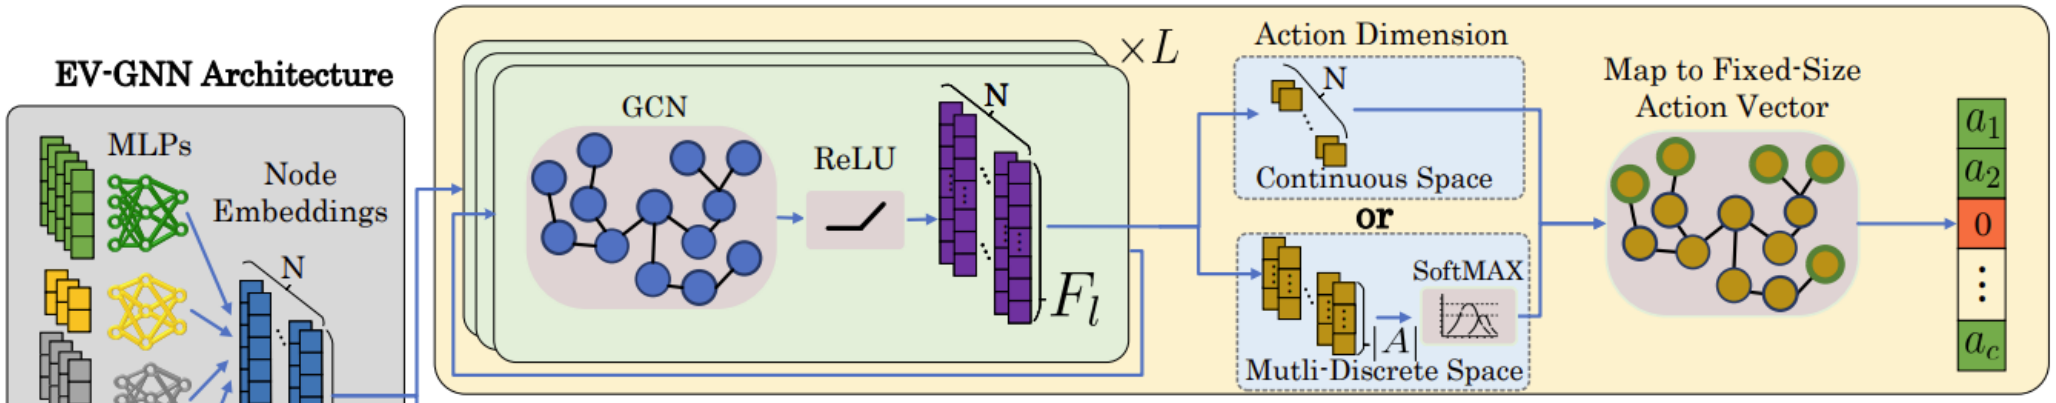

b. Feature Extractor

c. Actor NN

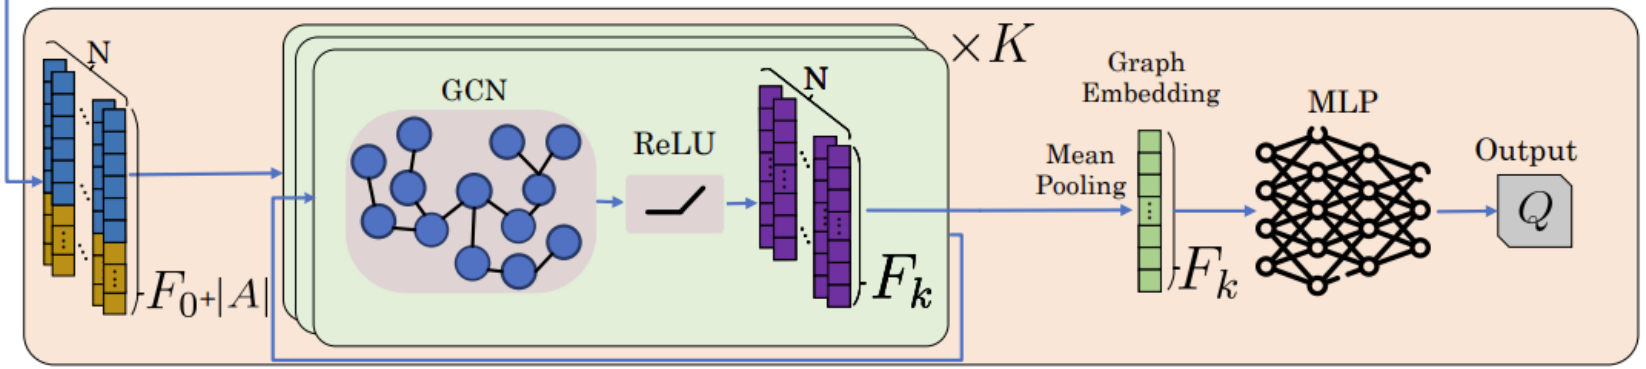

d. Critic NN

Optimality Gap as a function of RL algorithm and experiment scale

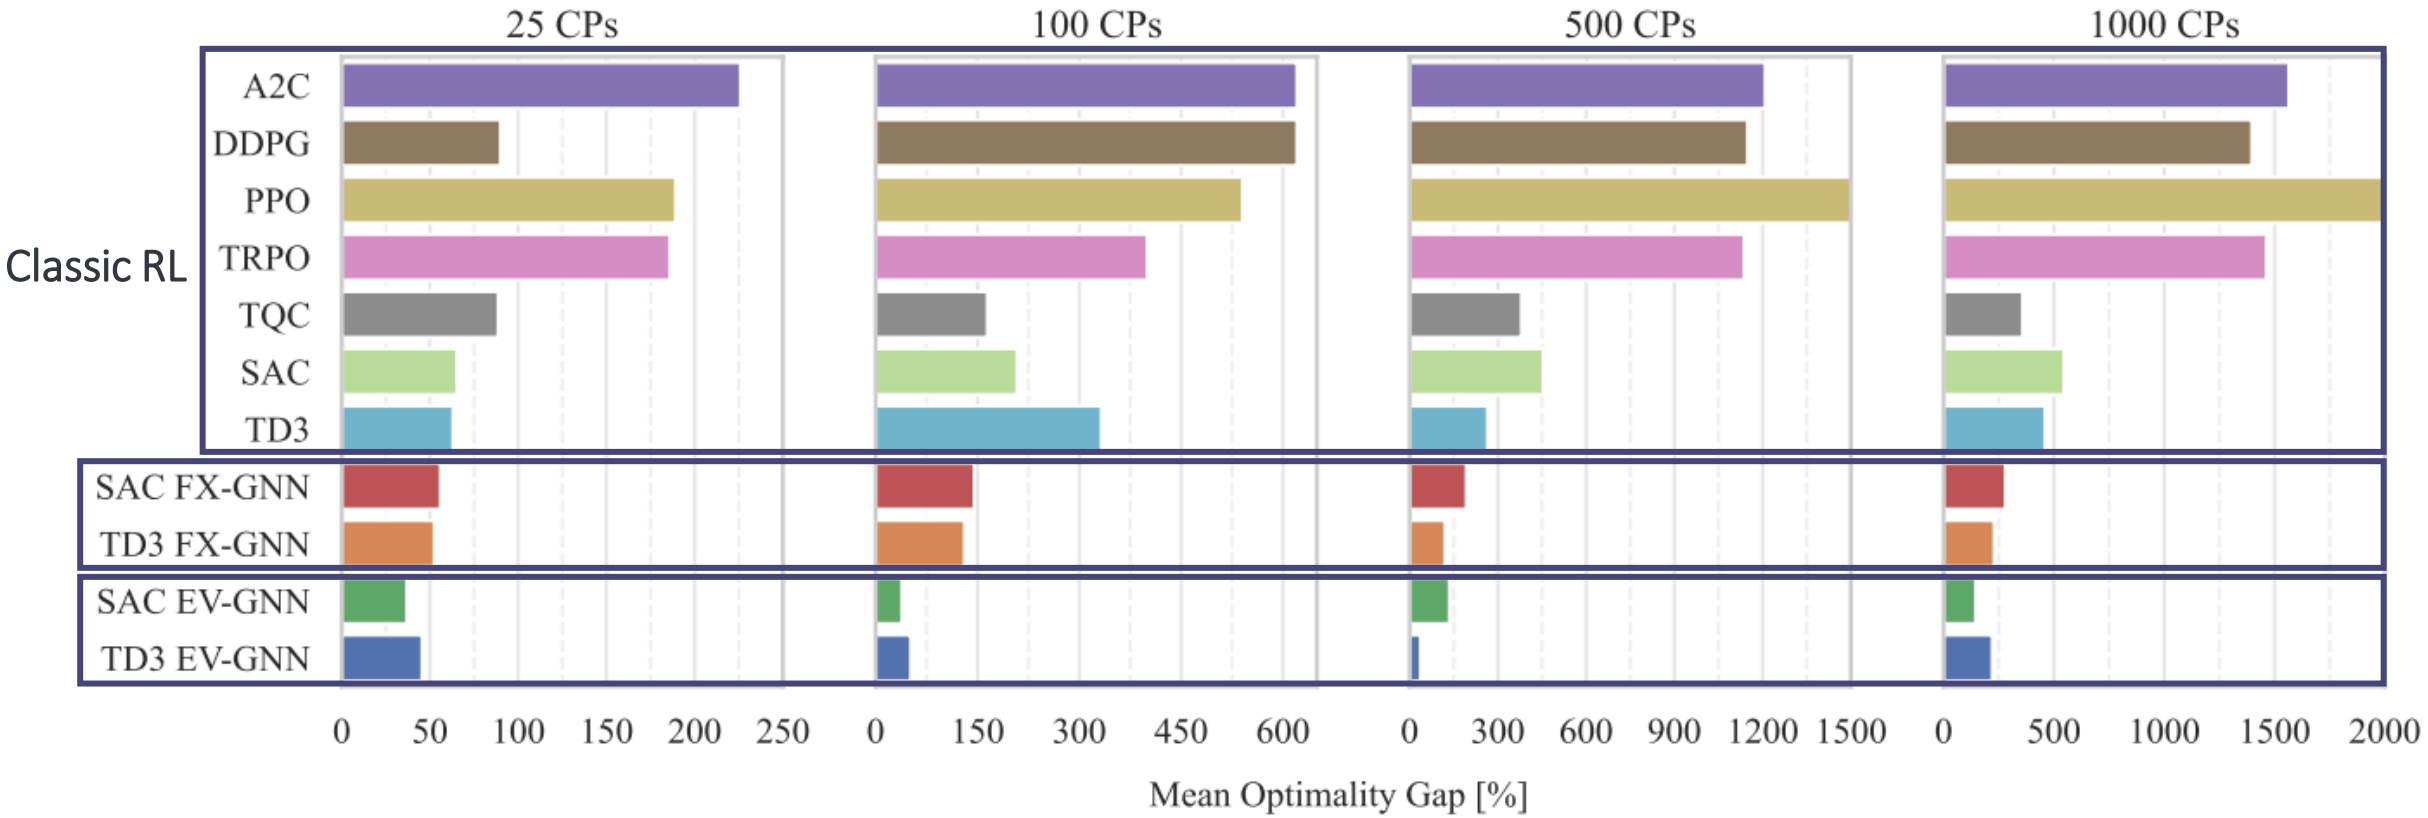

Oracle reward

$$G = \frac{|r - r^*|}{|r^*|} \cdot 100\%,$$

Sample Efficiency

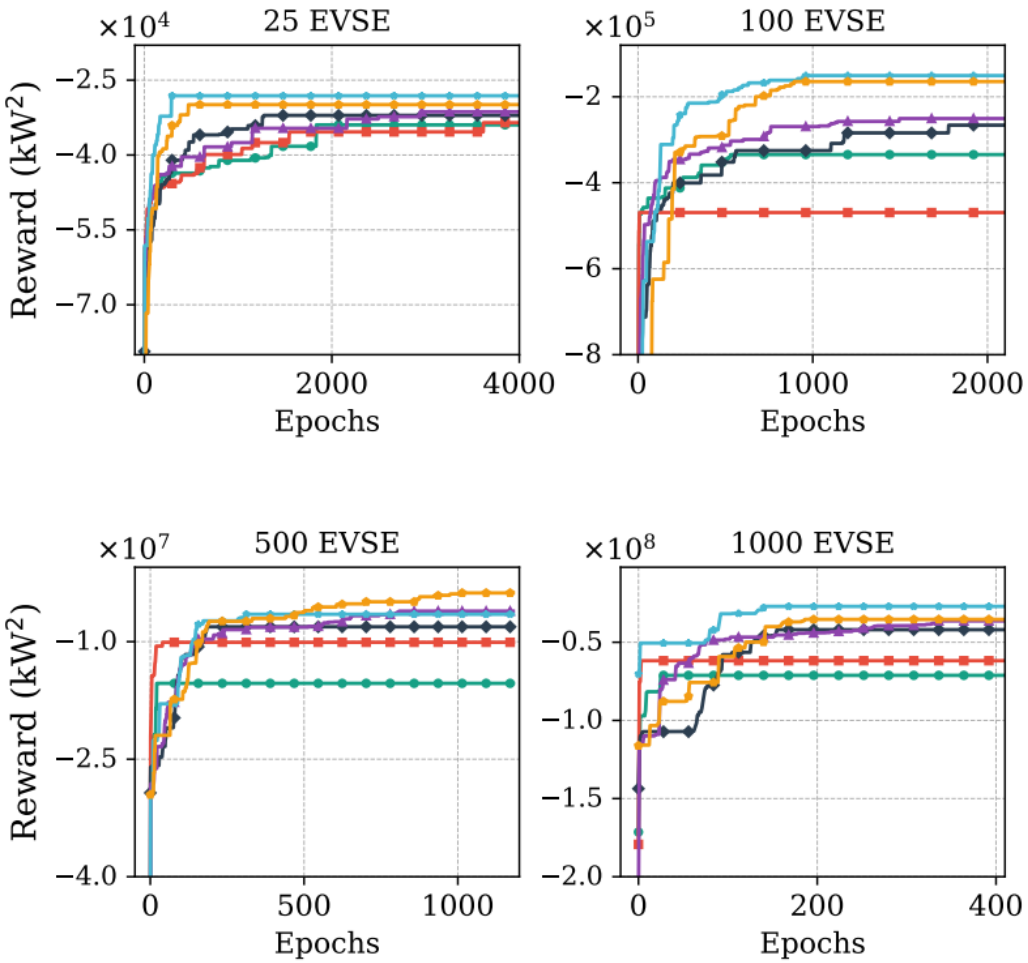

Invalid actions

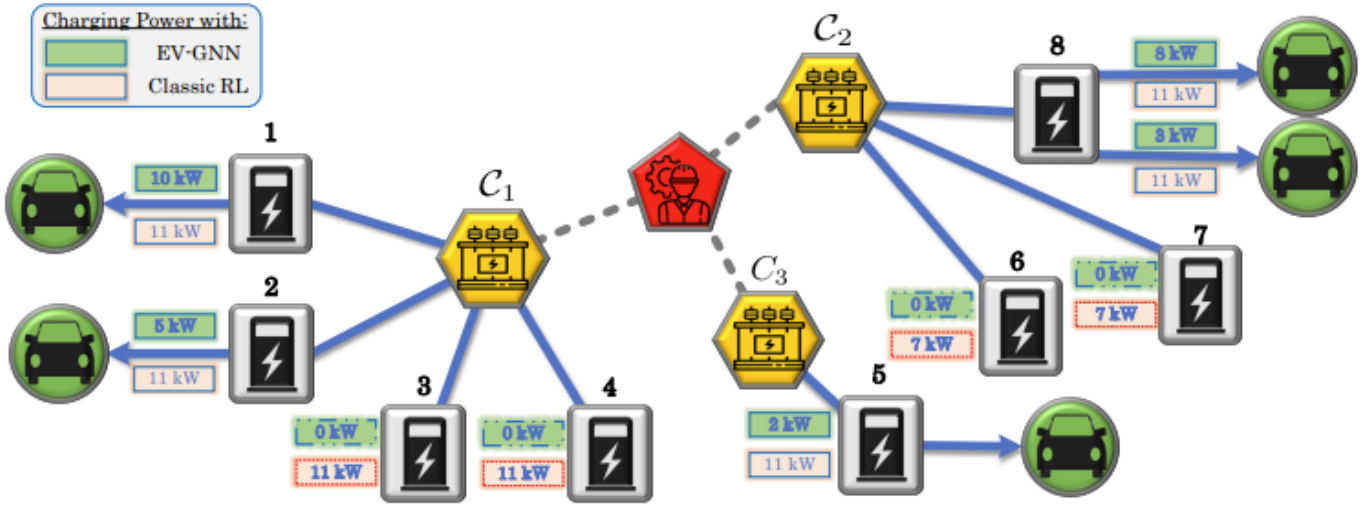

Generalization

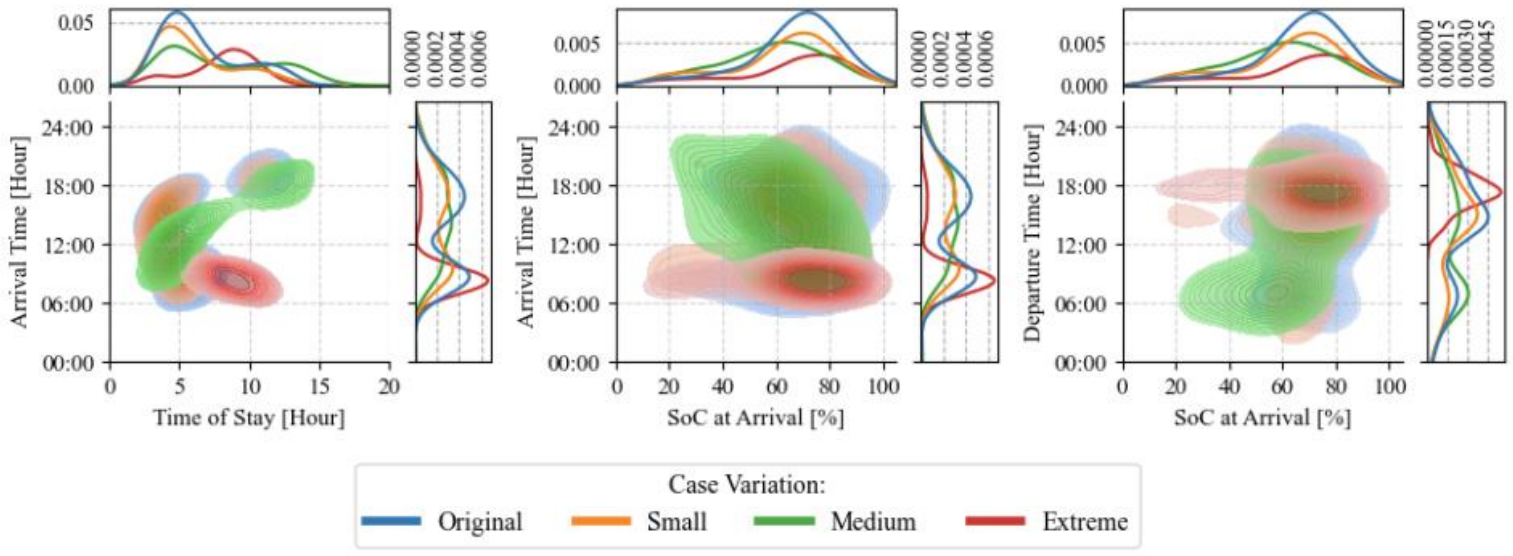

(a)

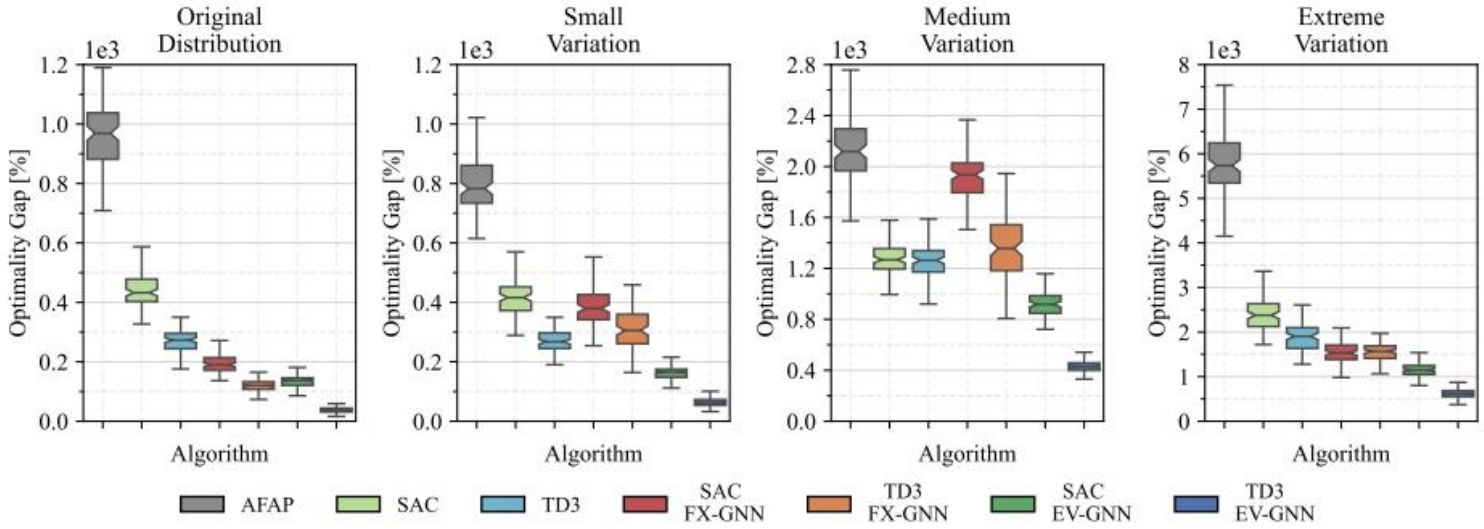

Wide applicability

Multi-Discrete Action Spaces

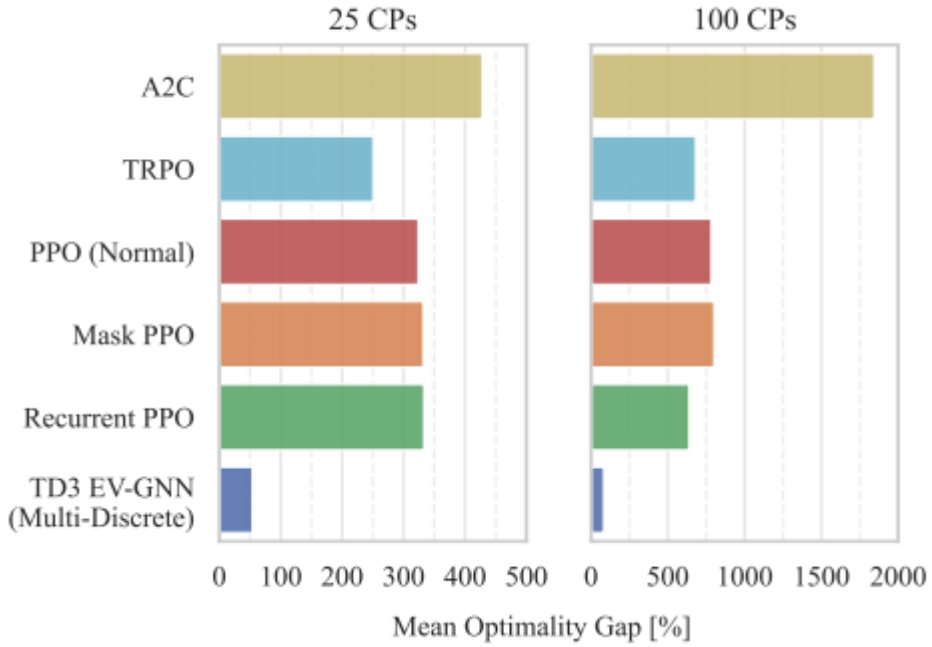

Multi-Objective EV charging problems

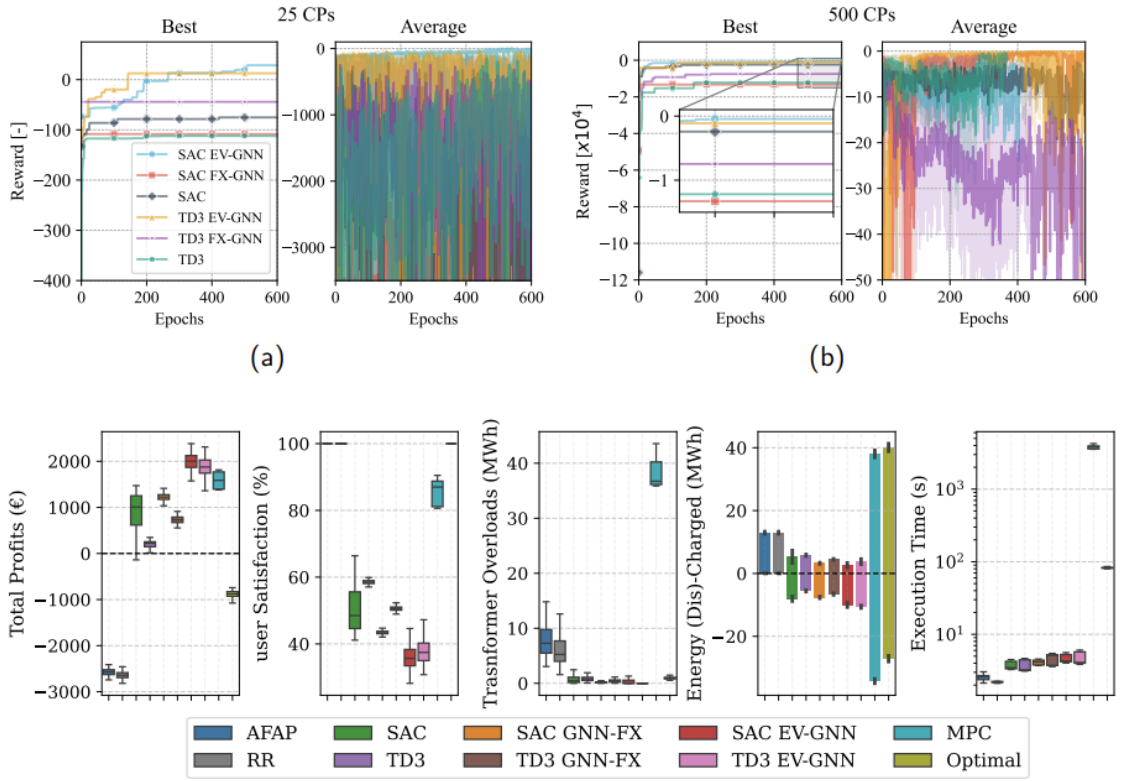

Conclusions

- Scalable Solution for Large-Scale Charging
  - Addresses traditional methods' limitations
  - Efficient, graph-based architecture for large, complex systems
  - Practical for CPOs managing thousands of EVs daily
- Improved Scalability & Generalization
  - Outperforms traditional RL in adapting to new environments
  - Filters irrelevant data, leveraging graph symmetries
- Versatile Across Control Domains
  - Adapts to continuous and discrete RL (e.g., TD3)
  - Robust in V2G profit maximization scenarios
- Promising Future Directions
  - Apply to dynamic tasks (e.g., vehicle routing, portfolio optimization)
  - Potential for Safe RL integration to enhance constraint satisfaction

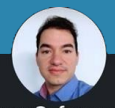

Stavros Orfanoudakis

AI Researcher | PhD @ EEMCS | IEPG | TU Delft 🔥

Thank you!

s.orfanoudakis@tudelft.nl