Guided A* Search for Scheduling Power Generation Under Uncertainty

Patrick de Mars & Aidan O’Sullivan
UCL Energy Institute

The Unit Commitment (UC) Problem

- **Fundamental task in power systems operation**: determining on/off schedules of power generators for future period (e.g. day ahead)

- Objective: **minimise expected operating costs** over uncertain demand, wind and other stochastic processes

- Typically solved by mixed-integer linear programming (MILP) using a deterministic reserve constraint (e.g. proportion of demand or 'N-1' criterion) to manage uncertainty

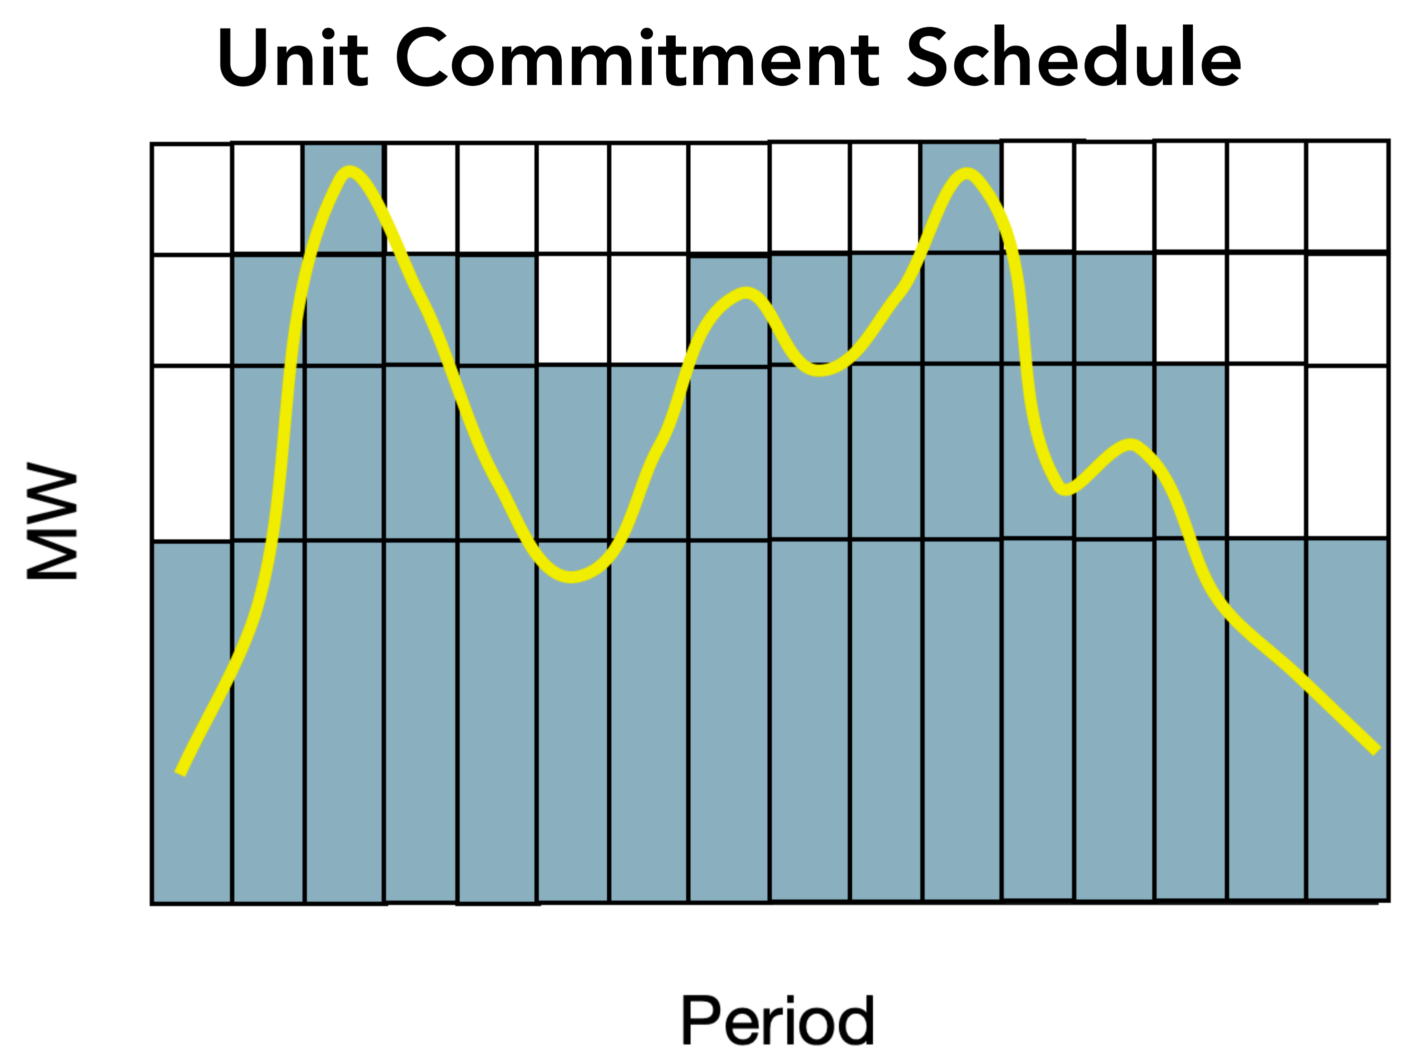

Generators (blue) are scheduled based on a demand forecast (yellow)

Motivation

- Uncertainty increasing due to: renewables penetration, behind-the-meter generation, ‘prosumers’, electrification of end-use sectors etc.

- Deterministic approaches are sub-optimal in high uncertainty power systems [1]

- Scenario-based stochastic optimisation approaches are computationally expensive [2]

- Large and growing size of power systems means **small efficiency improvements of existing assets can result in large absolute CO_2 emissions reductions**

Applying RL to the UC Problem

- RL is an attractive framework:
  - Suited to stochastic sequential decision making problems
  - Most of the computation (i.e. training) conducted in advance
  - Reward can be shaped to reflect societal values (**energy trilemma**)
- Challenges:
  - Large discrete (combinatorial) action space (up to 2^N actions)
  - Extreme penalties for lost load (blackouts), requiring **safe operation**
  - **Long time dependencies** (generators cannot be switched on/off frequently)
- **Existing research has only considered small power systems** (up to 12 generators) and hasn't considered generalisability to unseen problems (training and testing on same profiles)

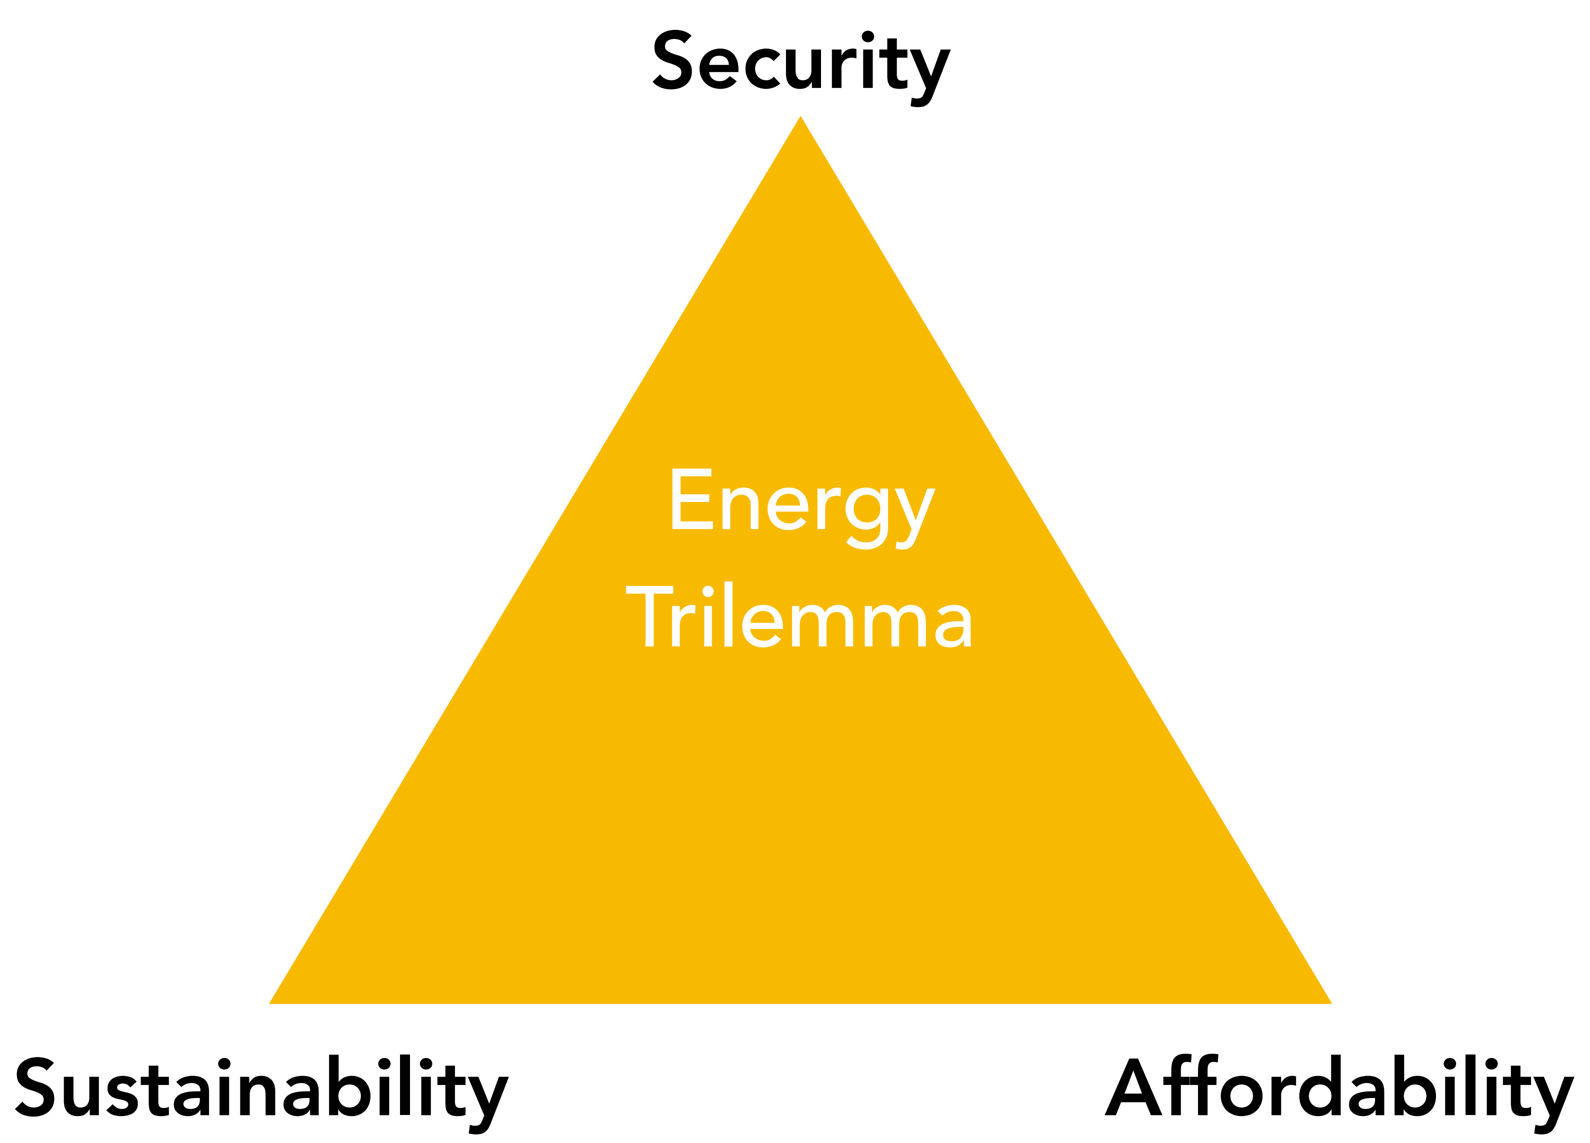

UC as a Markov Decision Process

- We formulate the UC problem as an episodic MDP with T decision periods and N generators
- Agent observes forecasts and current generator up/down times; actions are **combinatorial commitment decisions**
- **Stochastic demand and wind** modelled as auto-regressive moving average (ARMA) processes
- Reward reflects operating cost comprised of: fuel cost, carbon cost, startup cost, lost load cost (penalty for blackouts)
- Search tree representation: replace edge costs with expected cost using Monte Carlo approach
    - Solve the UC problem by finding lowest cost path
    - Note: 2^N branches for N generators!

States	$\boldsymbol{u}_t$: generator up/down times $\in \mathbb{Z}^N$ $\boldsymbol{d}$: demand forecast $\in \mathbb{R}^T$ $\boldsymbol{w}$: wind forecast $\in \mathbb{R}^T$ x_t: demand forecast error $\in \mathbb{R}$ y_t: wind forecast error $\in \mathbb{R}$ t: timestep $0 \leq t \leq T \in \mathbb{Z}$
Observations	$\{\boldsymbol{u}_t, \boldsymbol{d}, \boldsymbol{w}, t\}$
Actions	a_t: commitment decisions $\{0,1\}^N$
Rewards	r_t: negative operating cost $\in \mathbb{Z}$
Transitions	$\mathrm{u}_{i,t+1} = \begin{cases} u_{i,t}+1, & \text{if } a_{i,t}=1 \text{ and } u_{i,t}>0 \\ 1, & \text{if } a_{i,t}=1 \text{ and } u_{i,t}<0 \\ -1, & \text{if } a_{i,t}=0 \text{ and } u_{i,t}>0 \\ u_{i,t}-1, & \text{if } a_{i,t}=0t \text{ and } u_{i,t}<0 \end{cases}$ $x_t \sim X_t$: sample demand forecast error (from ARMA) $y_t \sim Y_t$: sample wind forecast error (from ARMA)

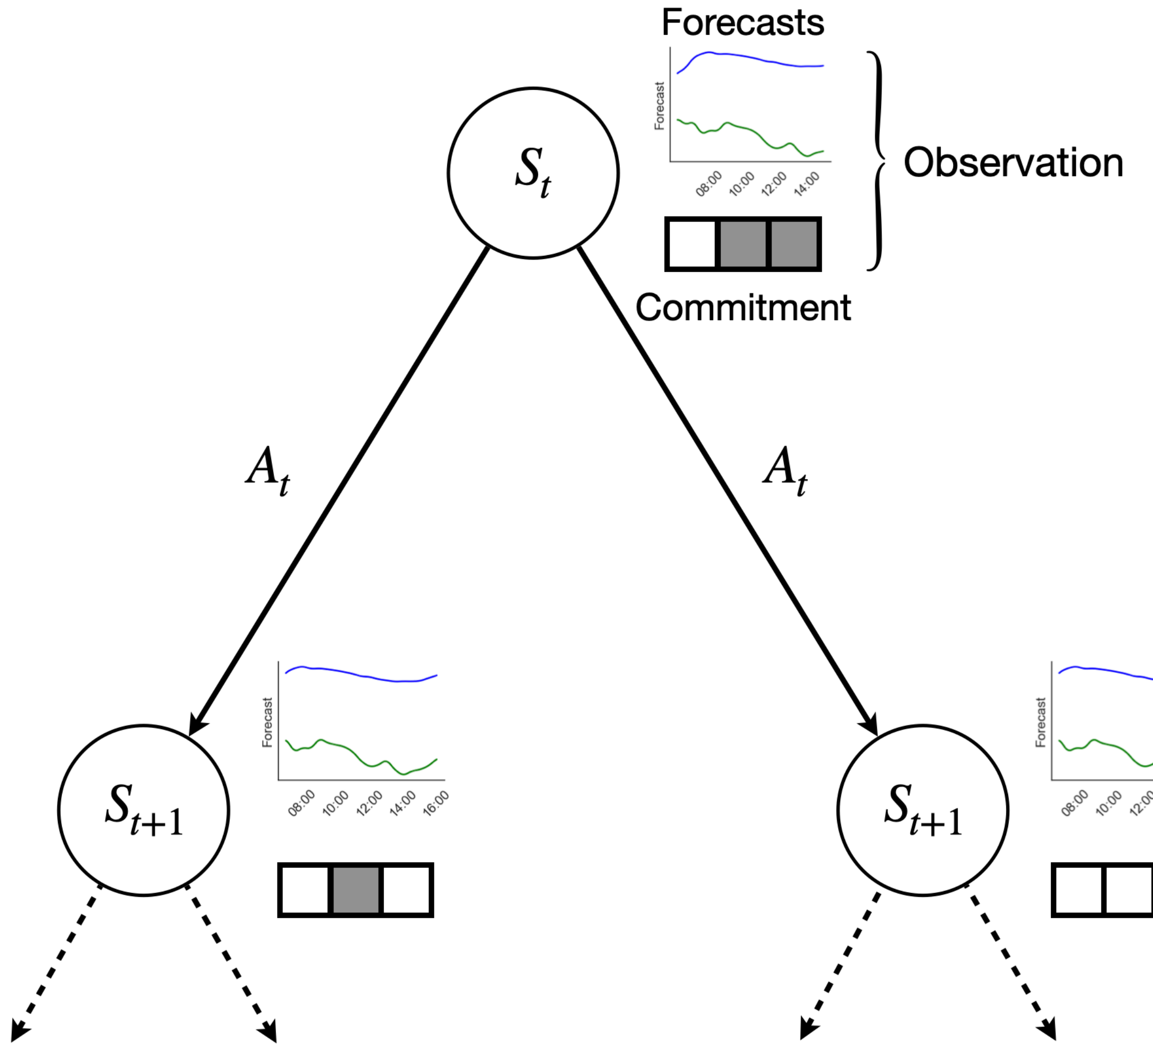

Search Tree Representation

Solution Method: Guided A*

- Train a policy $\pi(a|s)$ using model-free RL (PPO)

- **Guided expansion** used to reduce search breadth, pruning low probability branches:

$$A_{\pi}(s) = \{a \in A(s) | \pi(a|s) \geq \rho\}$$

$\pi(a|s)$:= expansion policy

ρ := branching threshold

- Use A* search [3] with a **priority list heuristic** to find lowest cost path through tree to fixed depth H

- In practice the UC problem is time constrained. We used iterative-deepening A* (IDA*) [4] as an anytime algorithm: incrementally increase H, terminate when time budget is spent

Priority List Heuristic

- A* can exploit a problem-specific heuristic to improve search efficiency
- PL heuristic: commit generators in order of cost; ignore most constraints to improve speed

Experimental Setup

- Experiments conducted on power systems of 10, 20 and 30 generators considered, based on data from [5] (widely used UC benchmark)
- Demand and wind forecasts based on GB power system data (4 years of training data with 20 held out days for testing)
- MDP represented in a Gym-style environment (https://github.com/pwdemars/rl4uc)
- Two experiments conducted:
    - Comparison with MILP with no carbon price
    - Impact of carbon price of $50 per tCO_2

Generator cost curves

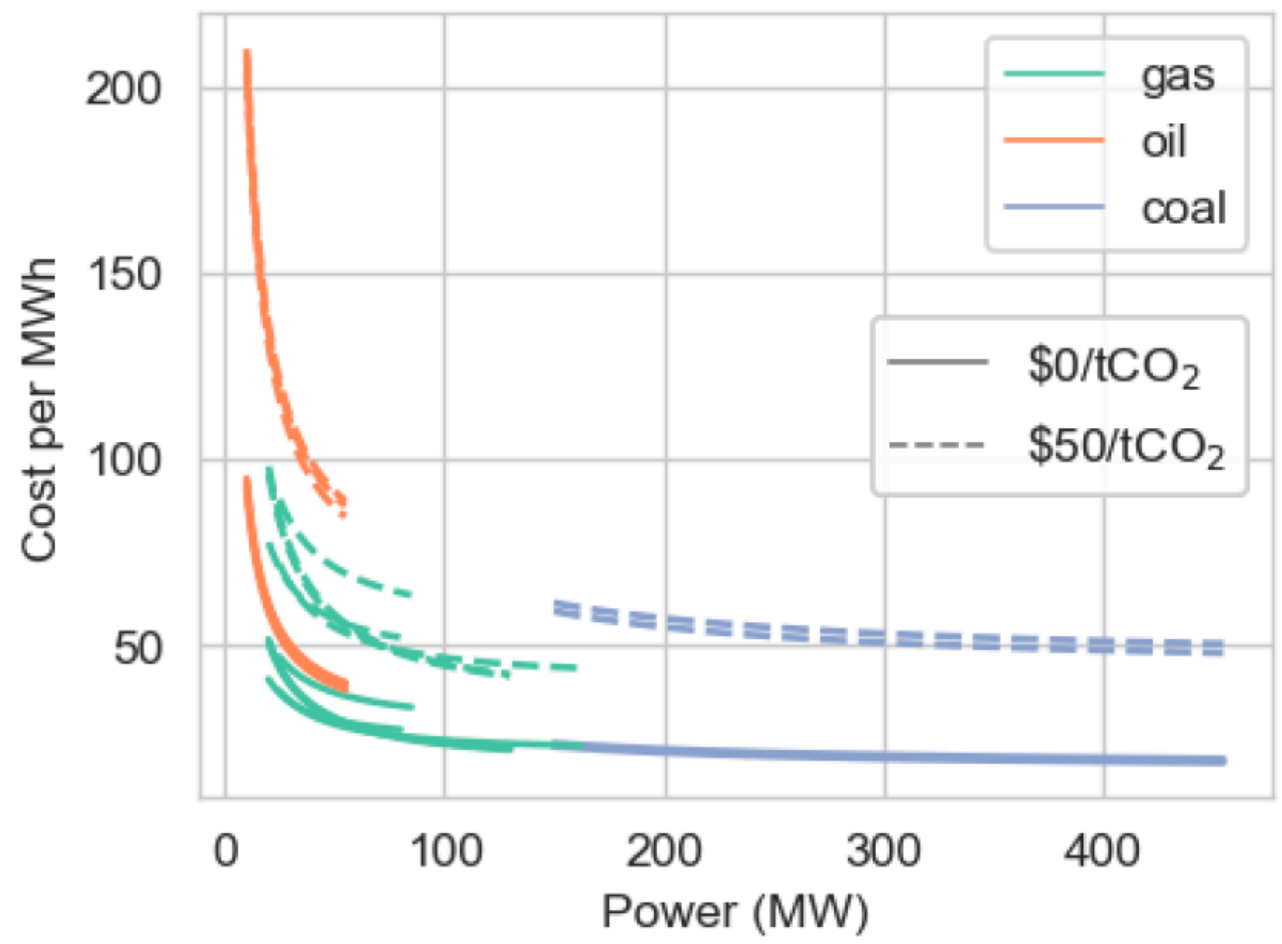

Policy training (PPO)

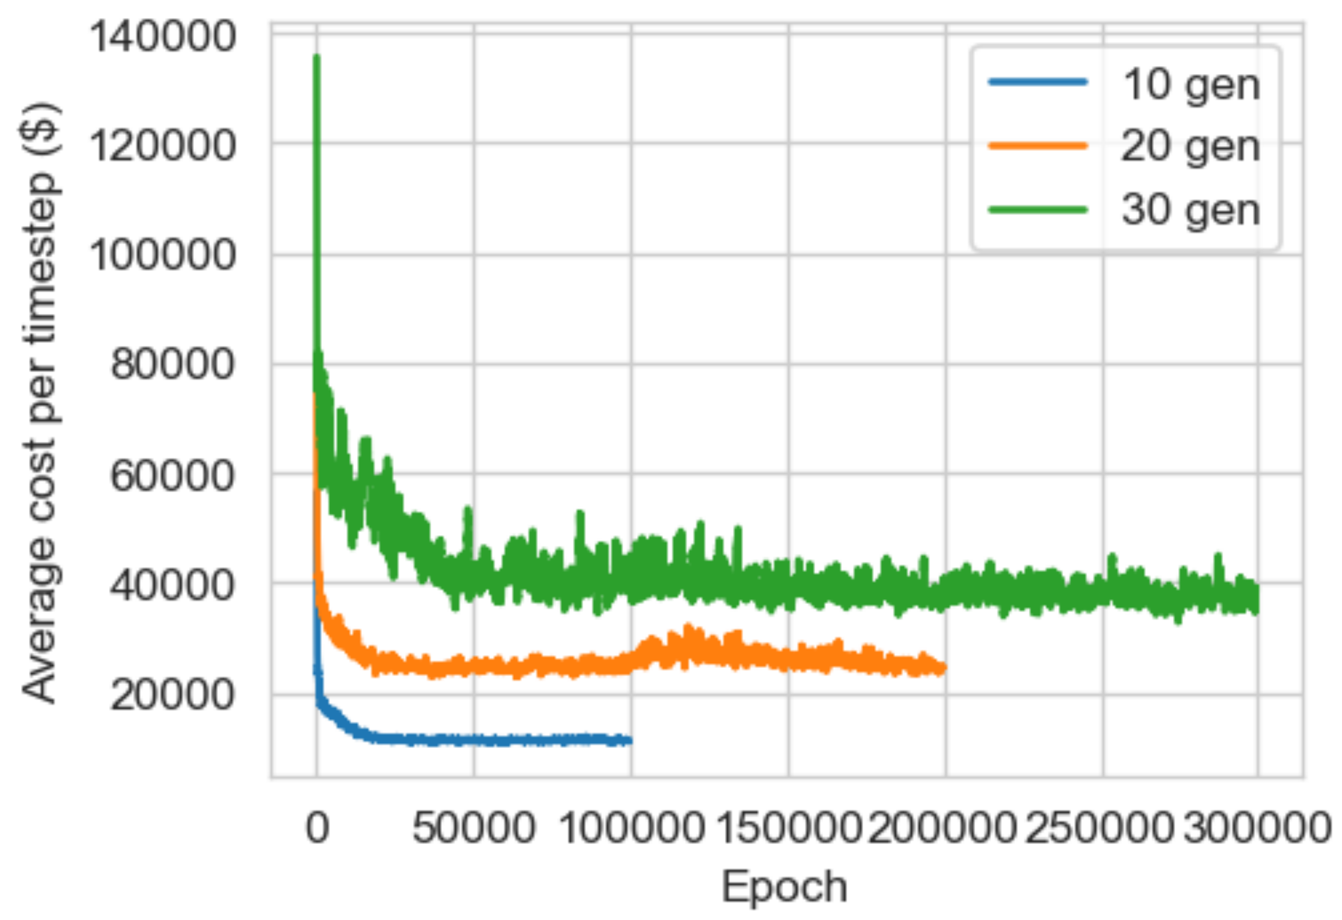

Experiment 1: Guided A* vs. MILP (no carbon price)

- Guided A* schedules were **0.8—1.2% cheaper** than MILP with a deterministic reserve constraint
    - Comparable to improvements of stochastic over deterministic MILP methods
- More secure operation: **loss of load probability roughly 50% lower** for guided A* compared with MILP

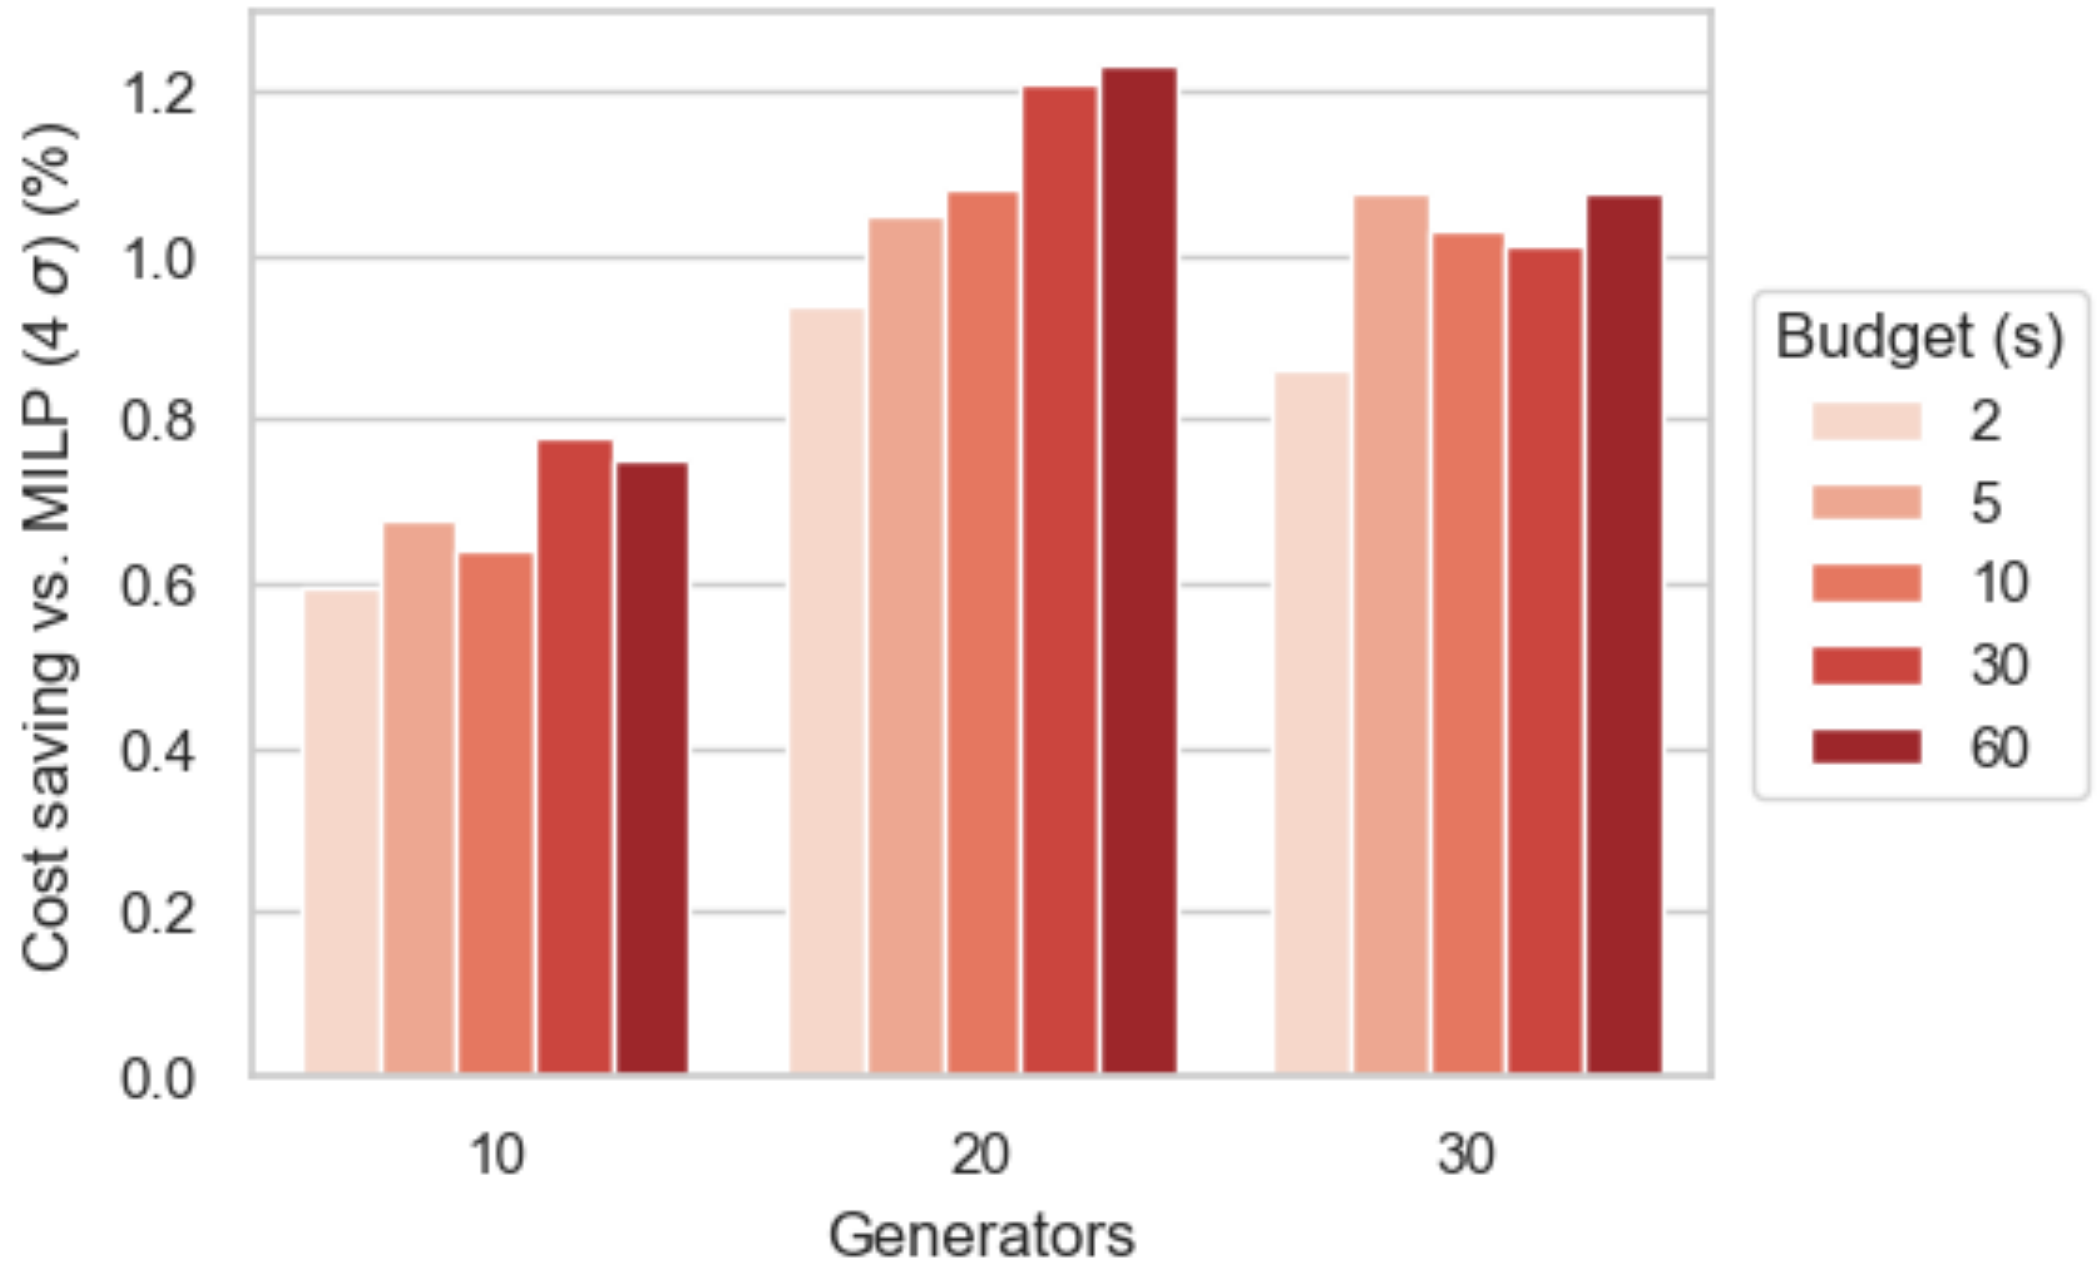

Experiment 2: Guided A* with Carbon Price

# Gens	$/tCO$_2$	LOLP (%)	ktCO$_2$	Coal (%)	Gas (%)	Oil (%)	Startups
10	0	0.12	264.03	99.64	41.88	6.28	141
10	50	0.12	245.89	91.30	61.37	13.19	114
20	0	0.11	527.56	99.09	40.74	8.21	235
20	50	0.09	476.62	86.38	66.24	5.38	164
30	0	0.16	780.43	99.10	40.89	5.69	346
30	50	0.17	724.81	88.59	67.86	12.67	215

- Including a carbon price of $50 per tCO$_2$ **reduces total carbon emissions by between 7—10%**
- Usage of generators (% periods online) shifts from coal towards lower carbon intensity generation (gas)
- Fewer startups, smaller reserve margins with carbon price

Generator cost curves

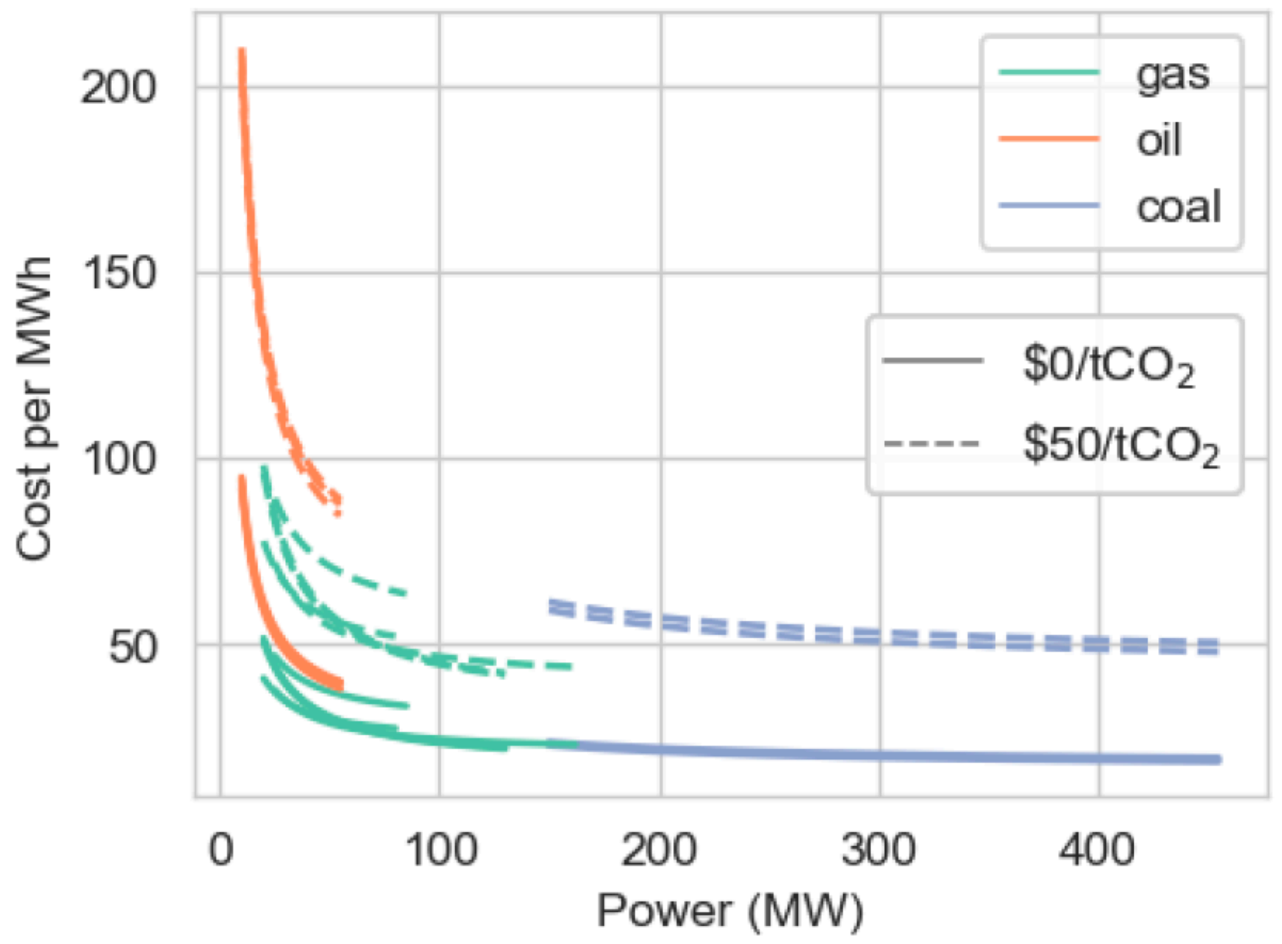

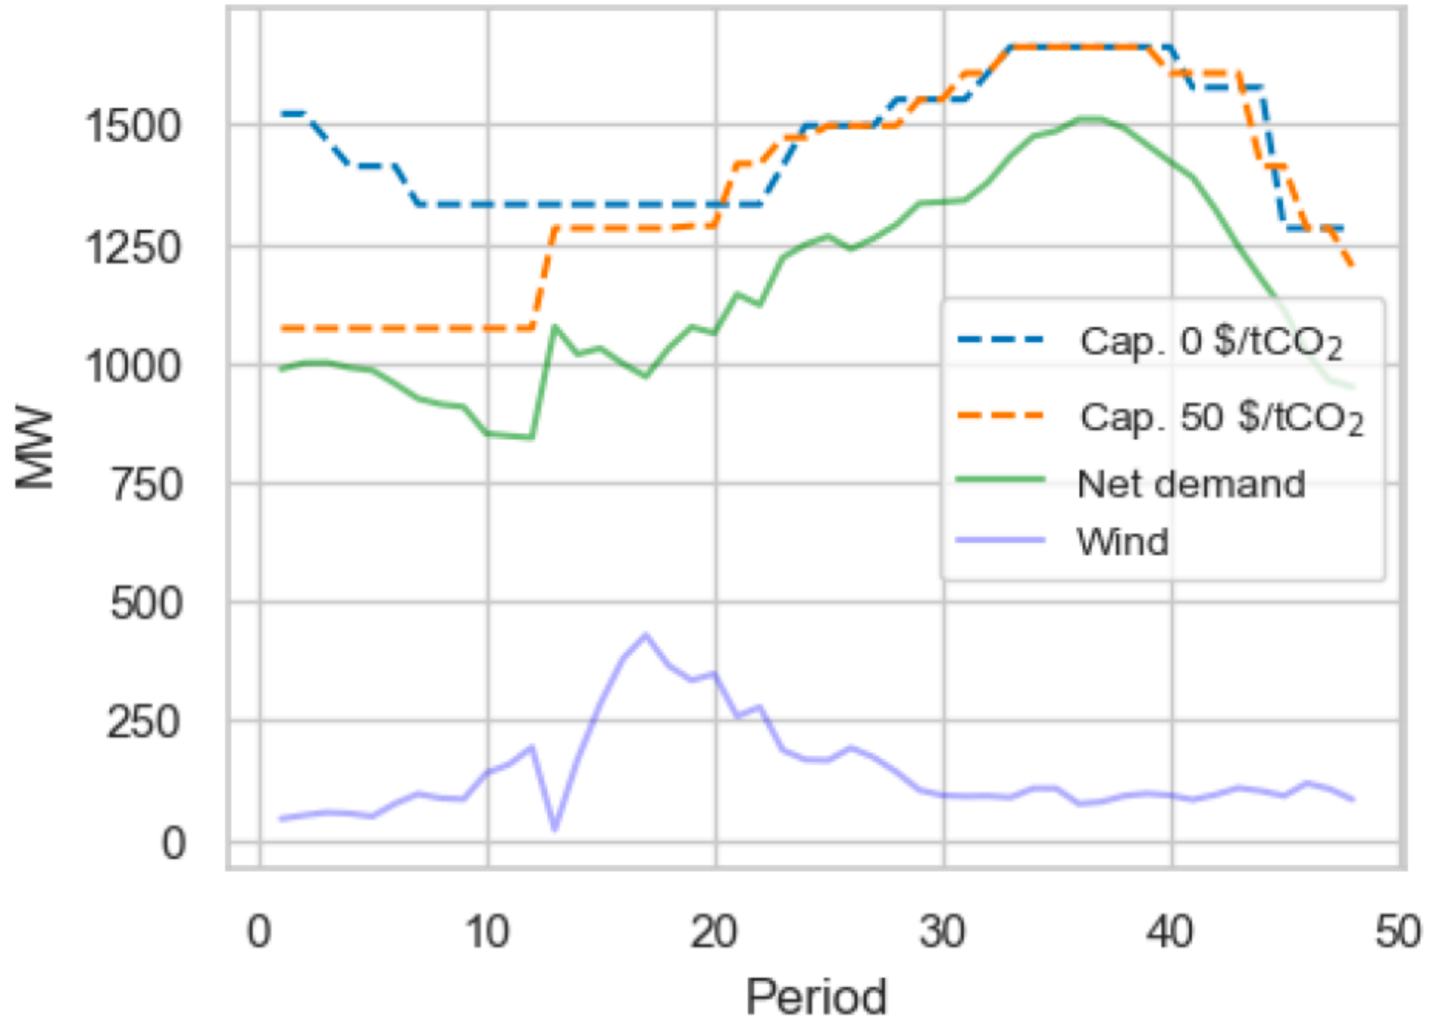

Conclusions

- RL can be successfully applied to the UC problem when combined with planning methods
- Reward shaping significantly alters behavioural strategies
- RL for power systems requires domain expertise: methods can't be applied out-of-the-box!

Thank you for listening, please get in touch if you have any questions!

patrick.demars.14@ucl.ac.uk

References

- [1] Ruiz, P. A., Philbrick, C. R., Zak, E., Cheung, K. W., and Sauer, P. W. Uncertainty management in the unit commitment problem. IEEE Transactions on Power Systems, 24 (2):642–651, 2009.

- [2] Bertsimas, D., Litvinov, E., Sun, X. A., Zhao, J., and Zheng,T. Adaptive robust optimization for the security constrainedunit commitment problem. IEEE Transactions on Power Systems, 28(1):52–63, 2012.

- [3] Russell, S. and Norvig, P. Artificial Intelligence: A ModernApproach. Prentice Hall Press, USA, 3rd edition, 2009. ISBN 0136042597.

- [4] Korf, R. E. Real-time heuristic search. Artificial Intelligence, 42(2-3):189–211, 1990.

- [5] Kazarlis, S. A., Bakirtzis, A., and Petridis, V. A genetic algorithm solution to the unit commitment problem. IEEE Transactions on Power Systems , 11(1):83–92, 1996.