AI for Energy: Reinforcement Learning for data-efficient and explainable control algorithms to exploit energy flexibility

Chris Develder, et al.
AI4E Team, IDLab, Ghent University – imec

GHENT UNIVERSITY

IDLab
INTERNET & DATA LAB

imec

Energy Transition – Supply

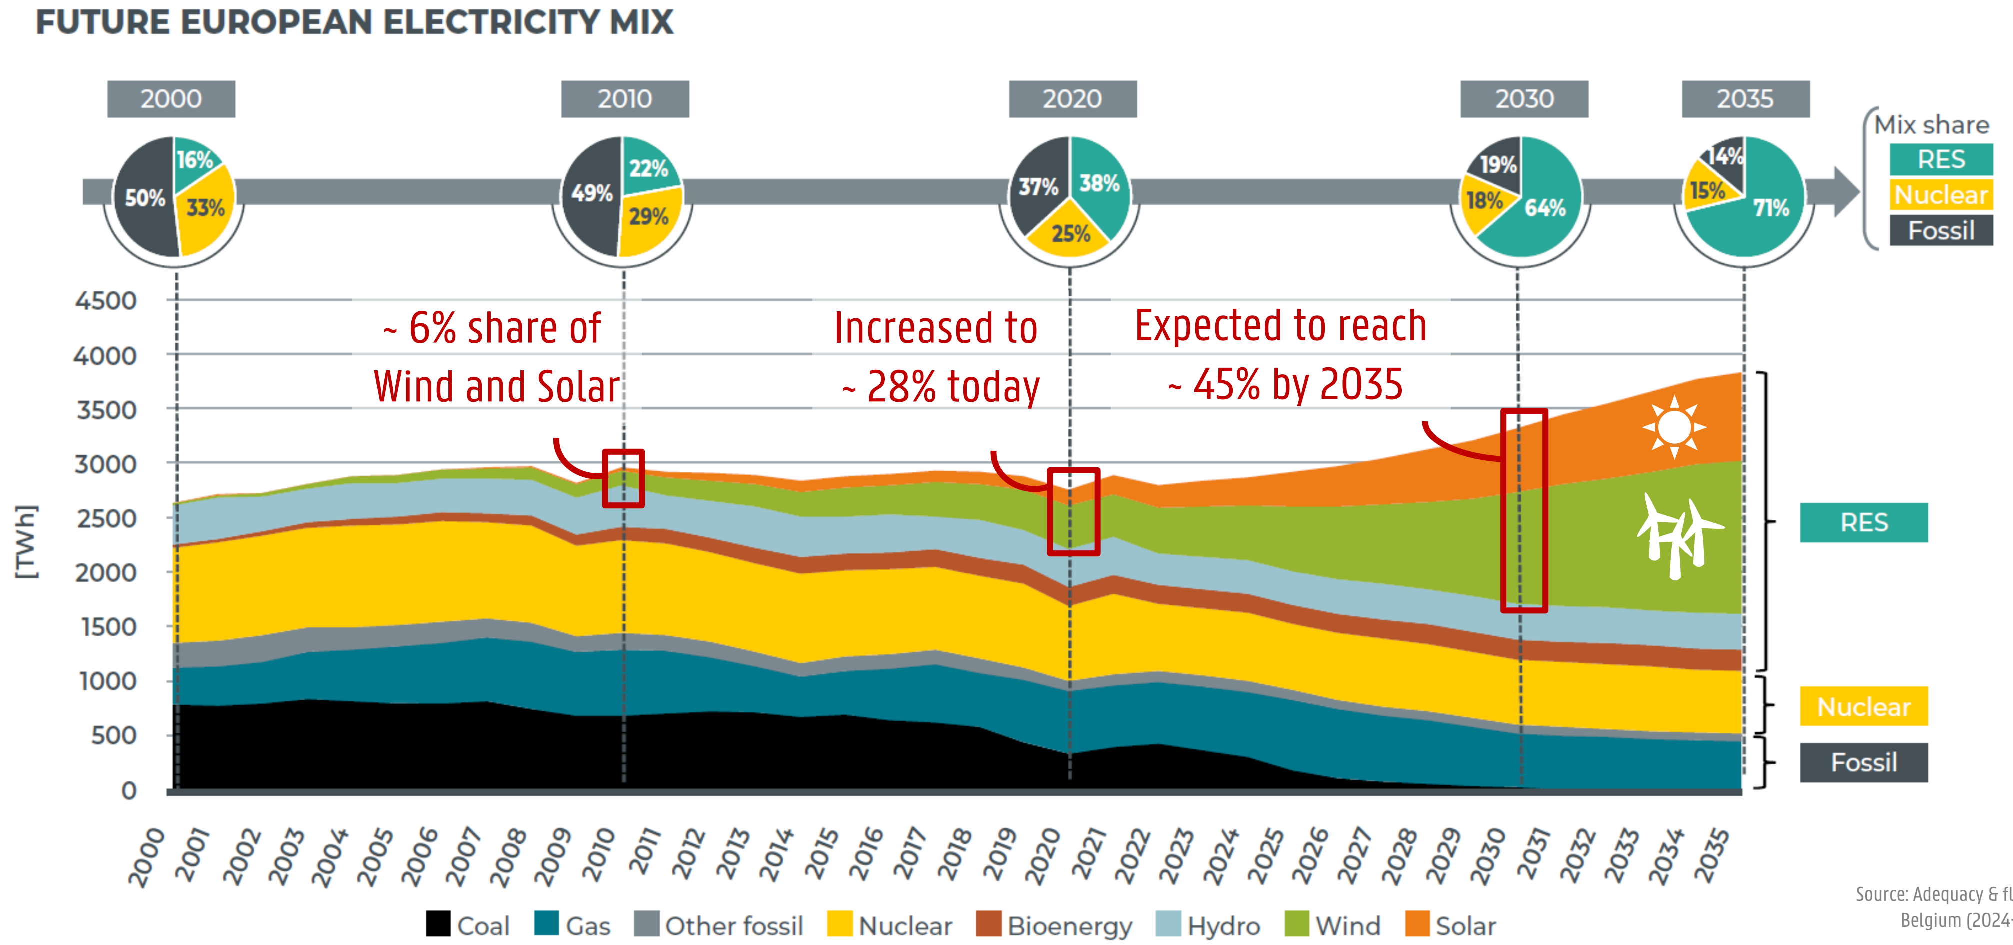

Source: Adequacy & flexibility study for Belgium (2024-2034), ELIA

GHENT UNIVERSITY

Energy Transition – Demand

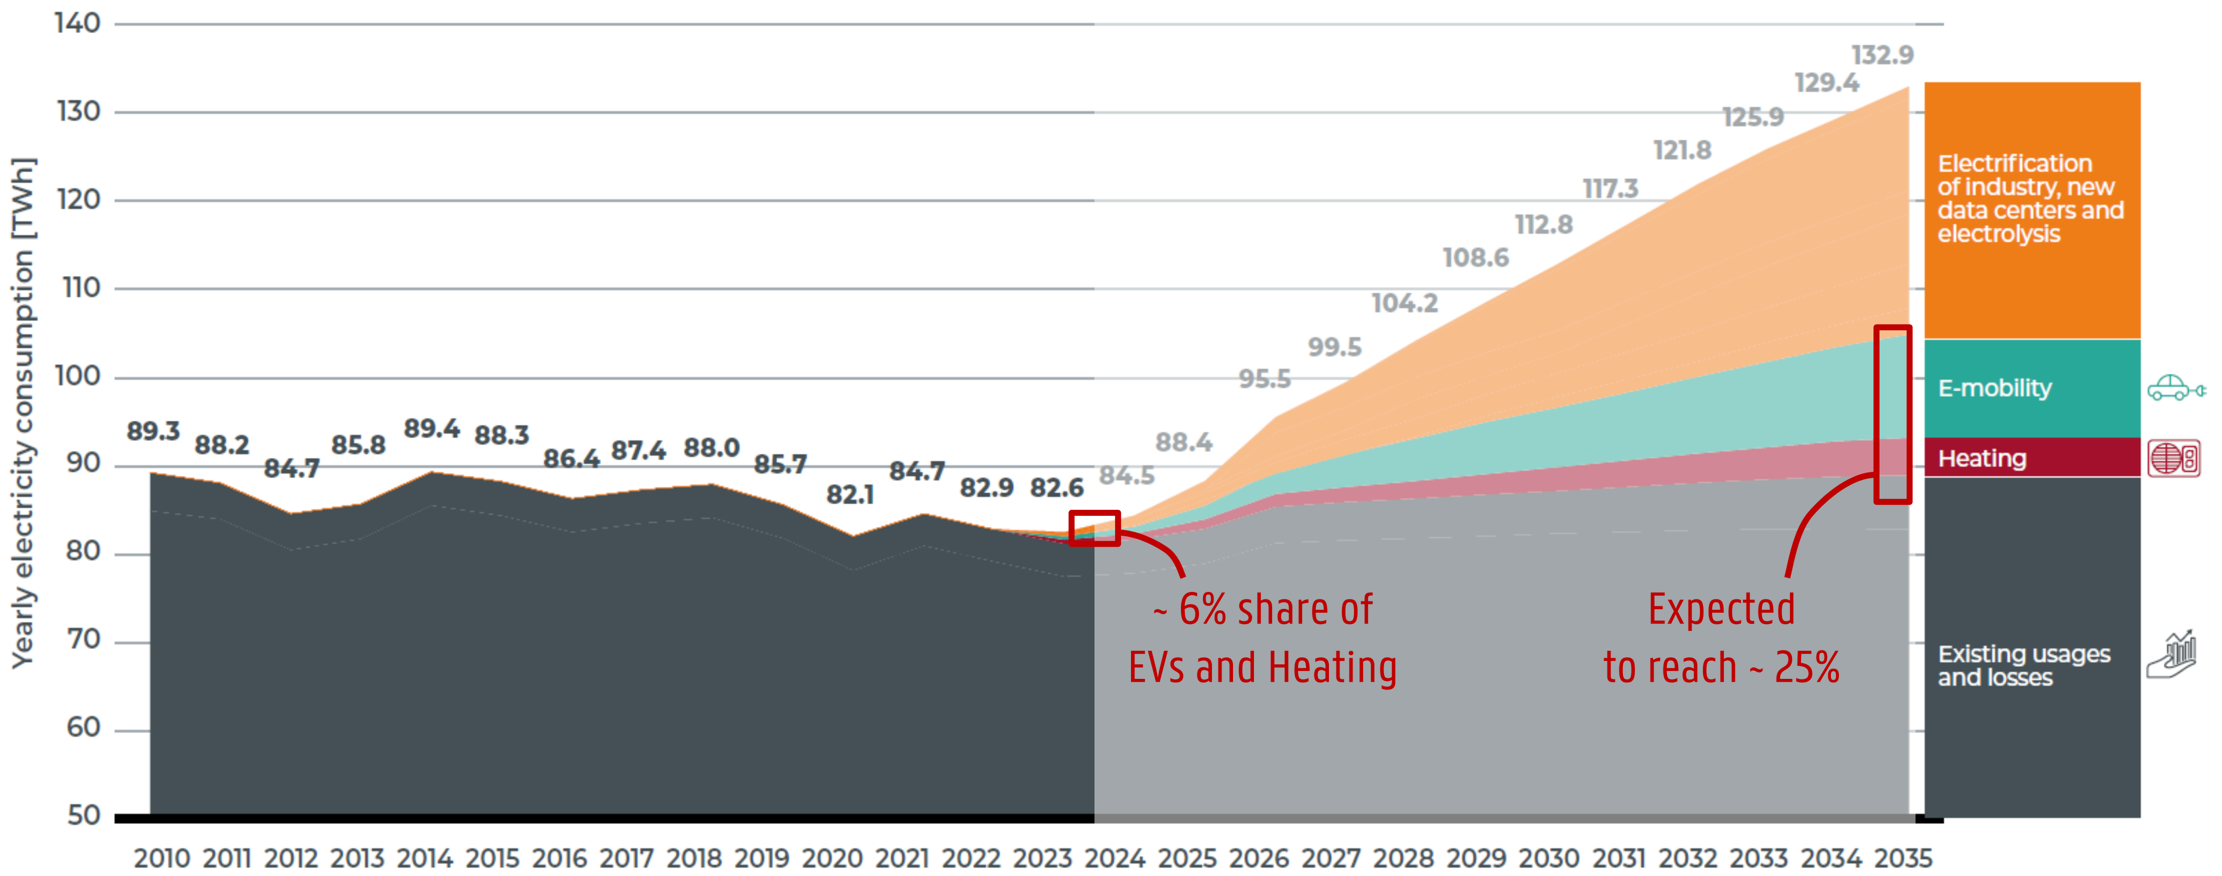

Source: Adequacy & flexibility study for Belgium (2024-2034), ELIA

SMART GRID

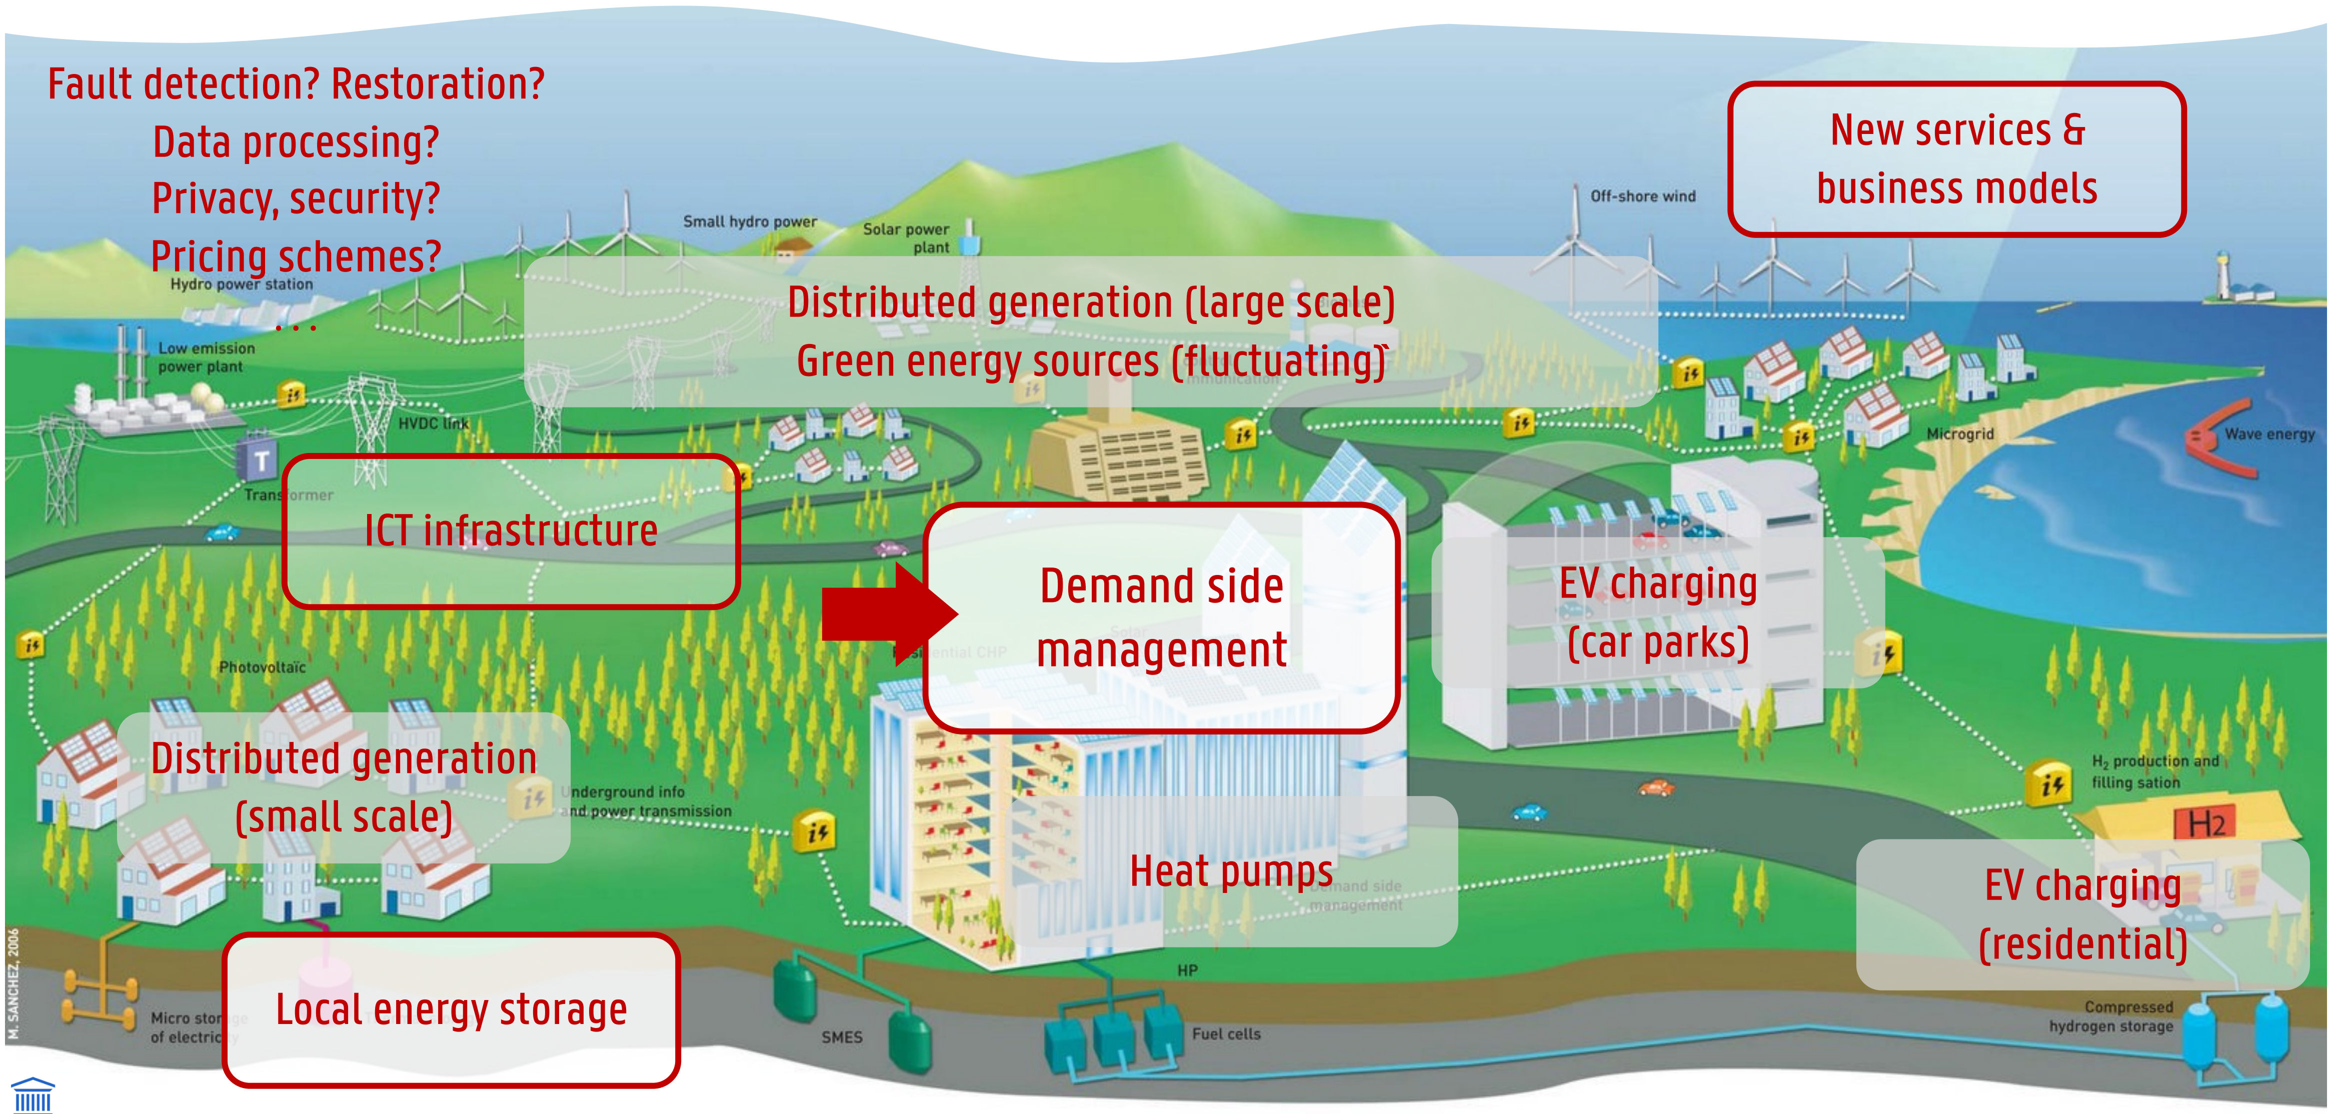

GHENT UNIVERSITY

C. Develder, et al., "Home Energy Management in the context of consumer-centric power systems", Powering the future, VAIA - Flanders AI Research, 26 Nov. 2024

Outline

PART I: Physics-informed reinforcement learning

- Challenge: Learn suitable control policy from limited data
- Solution: Infuse a priori physics knowledge into neural network model

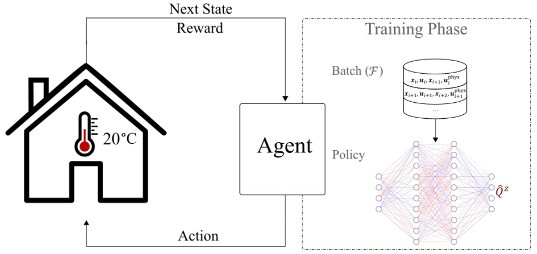

PART II: Interpretable decision trees

- Challenge: Policies from RL are black-box and non-trivial to deploy
- Solution: Distill a learned policy into decision tree

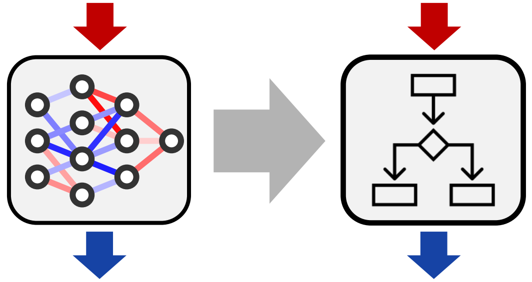

PART I:
Physics-informed reinforcement learning for residential heat pump control

G. Gokhale, B. Claessens, and C. Develder, **"Physics-informed neural networks for control-oriented thermal modeling of buildings"**, Appl. Energy, Vol. 314, 15 May 2022, pp. 1-10. doi:10.1016/j.apenergy.2022.118852

G. Gokhale, B. Claessens, and C. Develder, **"PhysQ: A physics-informed reinforcement learning framework for building controls"**, Arxiv preprint, 21 Nov. 2022, arXiv:2211.11830v1

Thanks Gargya!

GHENT UNIVERSITY

Residential heat pump control: Problem statement & proposed solution

- Objective:

  Minimize cost of energy consumed by a household while maintaining thermal comfort

- Controller should:
  - Use data efficiently
  - Scale across different households
  - Yield high-quality control policies

- Proposed Solution:

  Off-policy reinforcement learning with physics informed neural networks

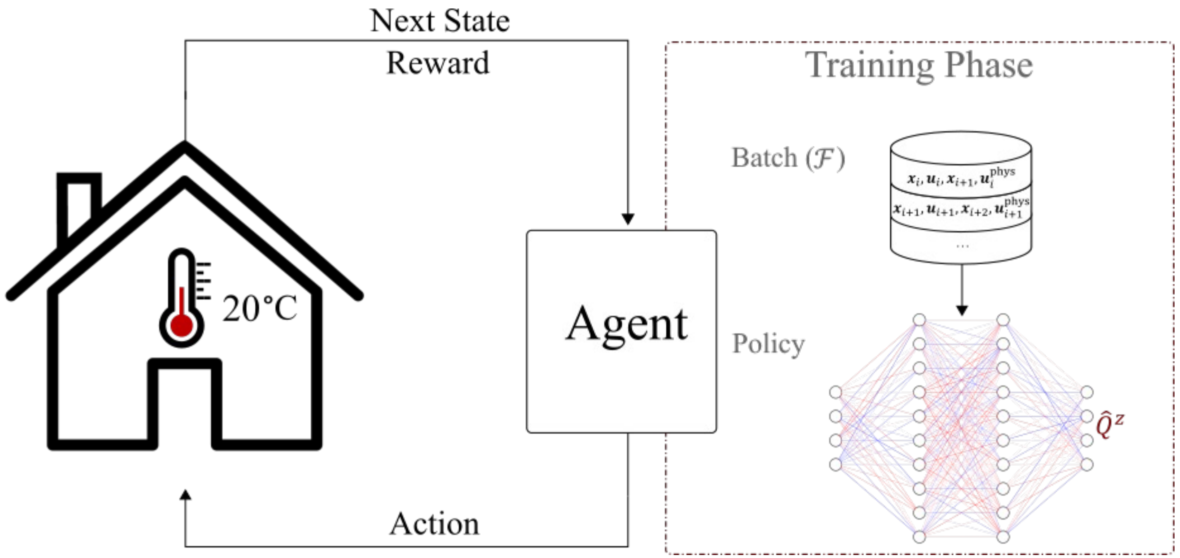

Reinforcement Learning: Agent acts in an environment

- At each timestep t the agent
  - Gets **state** observation s_t
  - Gets scalar **reward** r_t (depending on current state and thus previous action(s))
  - Decides to take **action** a_t
- The environment
  - Receives action a_t
  - Emits **state** s_{t+1} (which will depend on action)
  - Emits scalar **reward** r_{t+1}

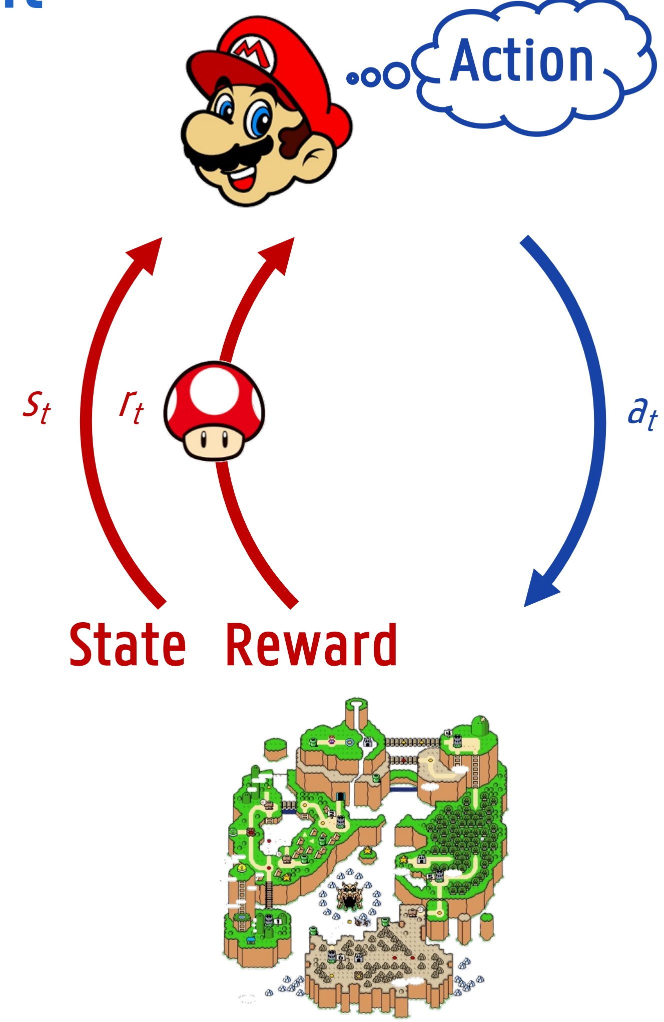

RL idea: learn a policy that maps state s_t to action a_t, such that long-term reward $\sum r_t$ is maximized

Reinforcement Learning: Agent acts in an environment

- At each timestep t the agent
  - Gets **state** observation s_t
  - Gets scalar **reward** r_t (depending on current state and thus previous action(s))
  - Decides to take **action** a_t
- The environment
  - Receives action a_t
  - Emits **state** s_{t+1} (which will depend on action)
  - Emits scalar **reward** r_{t+1}

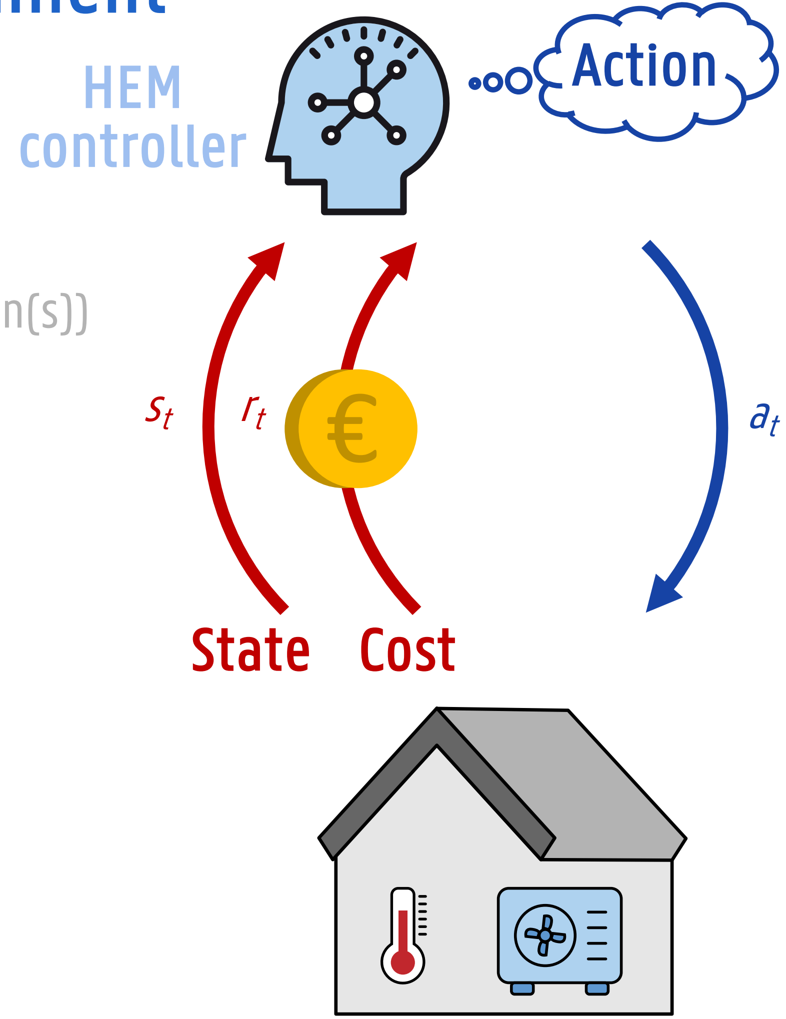

RL idea: learn a policy that maps state s_t to action a_t, such that long-term reward $\sum r_t$ is maximized

C. Develder, et al., "Home Energy Management in the context of consumer-centric power systems", Powering the future, VAIA - Flanders AI Research, 26 Nov. 2024

Reinforcement Learning: Heat pump control

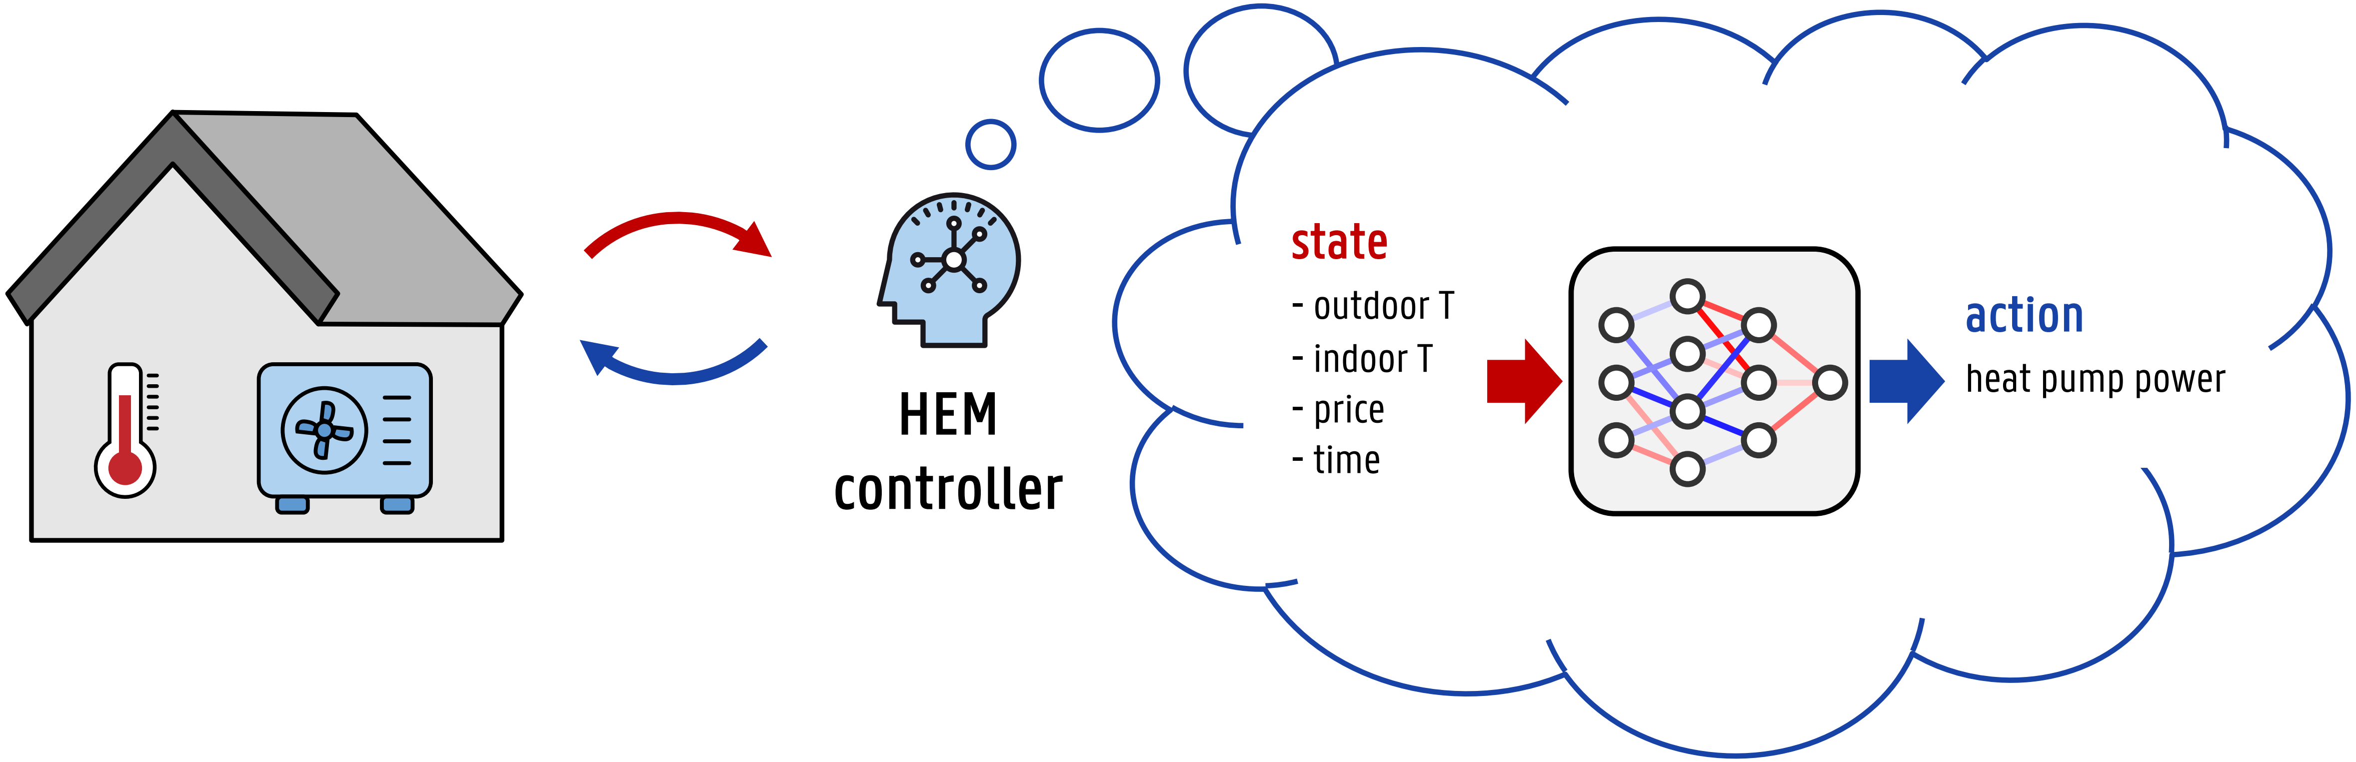

C. Develder, et al., "Home Energy Management in the context of consumer-centric power systems", Powering the future, VAIA - Flanders AI Research, 26 Nov. 2024

PhysQ: A physics-informed reinforcement learning framework

- **PhysQ**:
  - Physics informed RL framework
  - Leverages prior physics to learn physically relevant latent representations
  - Learns a control policy using this latent representation
- Two-step Algorithm:
  - **Step 1:** Learn physically relevant representation (bulk temperature T_m)
  - **Step 2:** Learn a low-dimensional Q-function based on this representation

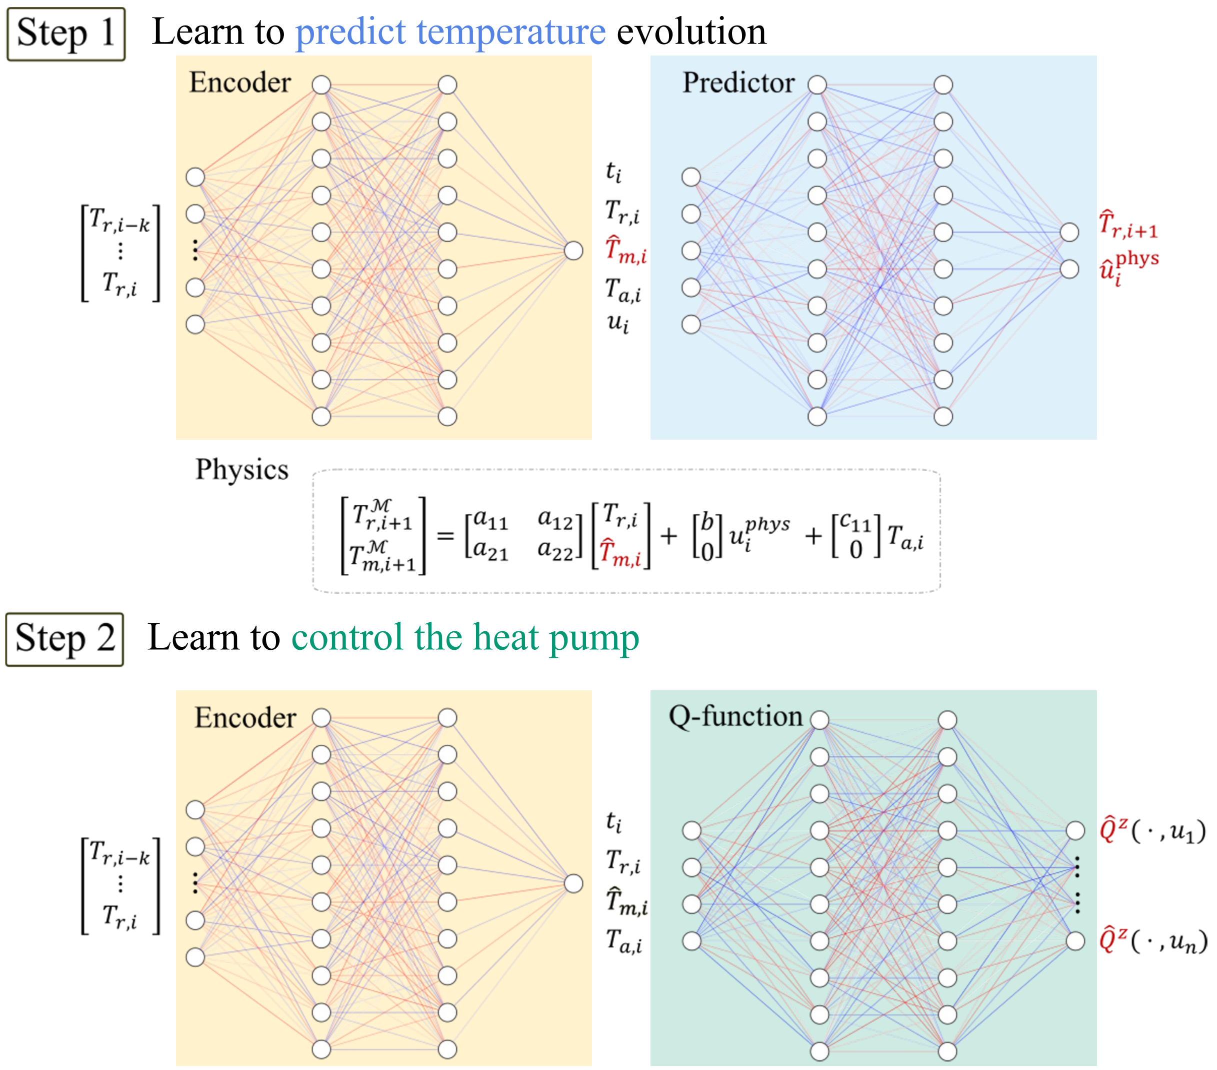

C. Develder, et al., "Home Energy Management in the context of consumer-centric power systems", Powering the future, VAIA - Flanders AI Research, 26 Nov. 2024

Simulation results

Results for BELPEX prices using a building simulator

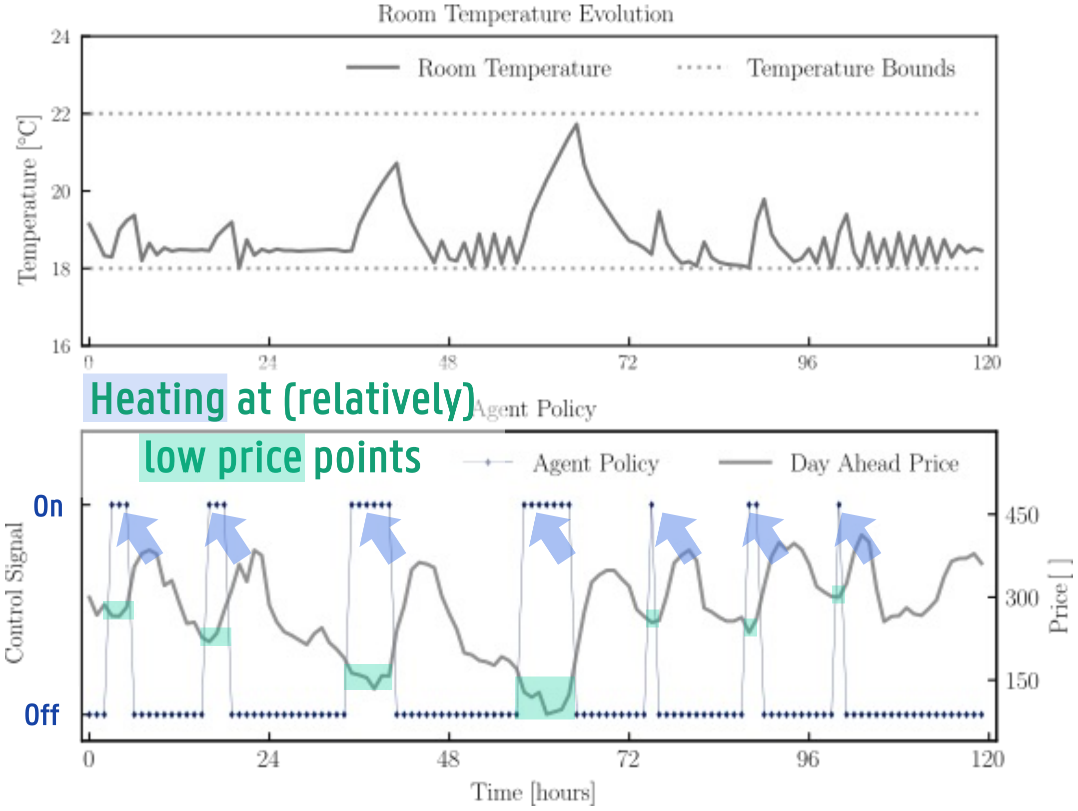

GHENT UNIVERSITY

Simulation results

Results for BELPEX prices using a building simulator

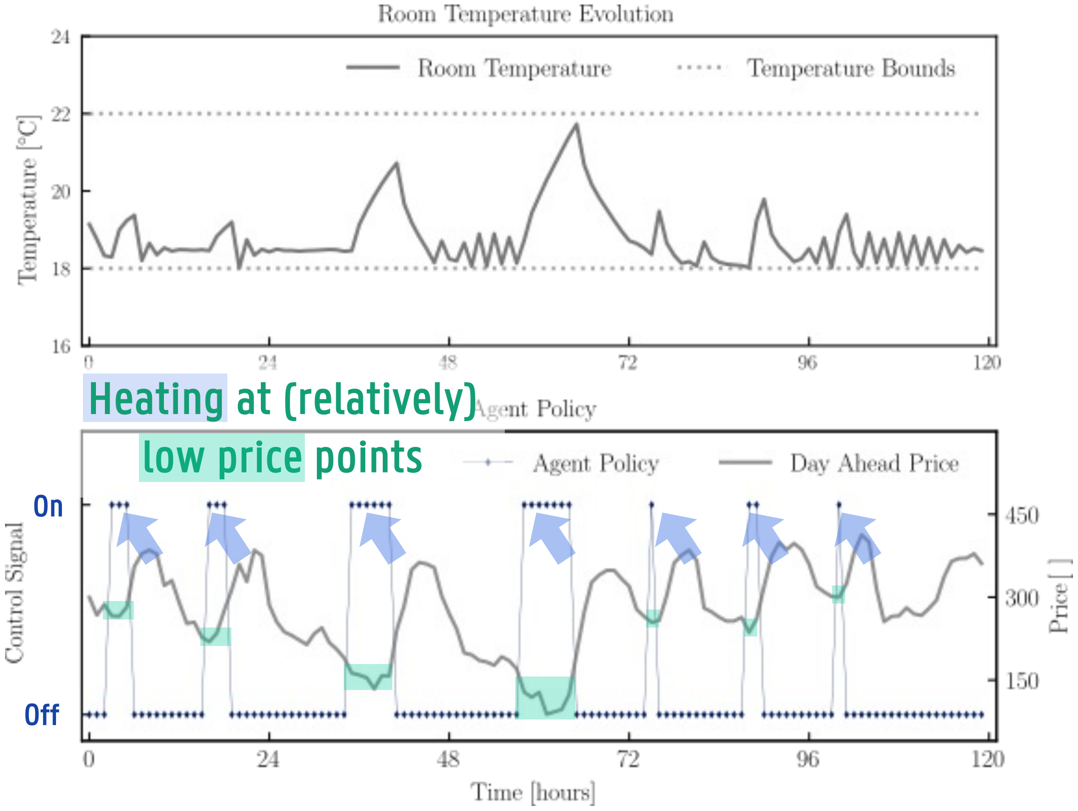

Comparison over different training data sizes and agent types

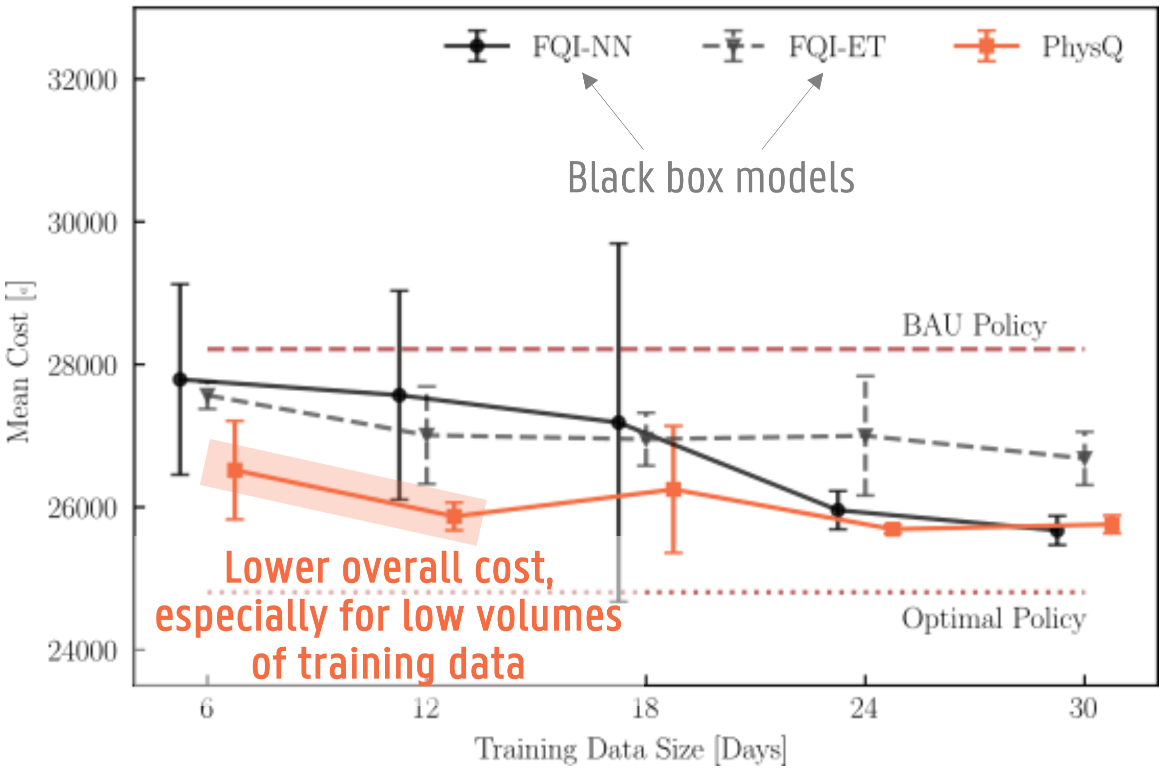

GHENT UNIVERSITY

C. Develder, et al., "Home Energy Management in the context of consumer-centric power systems", Powering the future, VAIA - Flanders AI Research, 26 Nov. 2024

PART II:
Interpretable decision trees distilled from black-box RL policies

G. Gokhale, S.S. Karimi Madahi, B. Claessens and C. Develder, **"Distill2Explain: Differentiable decision trees for explainable reinforcement learning in energy application controllers"**, in Proc. 15th ACM Int. Conf. Future Energy Sys. (e-Energy 2024), Singapore, 4-7 Jun. 2024, pp. 1-8. [doi:10.1145/3632775.3661937](https://doi.org/10.1145/3632775.3661937)

Thanks Gargya!

Black-box Reinforcement Learning

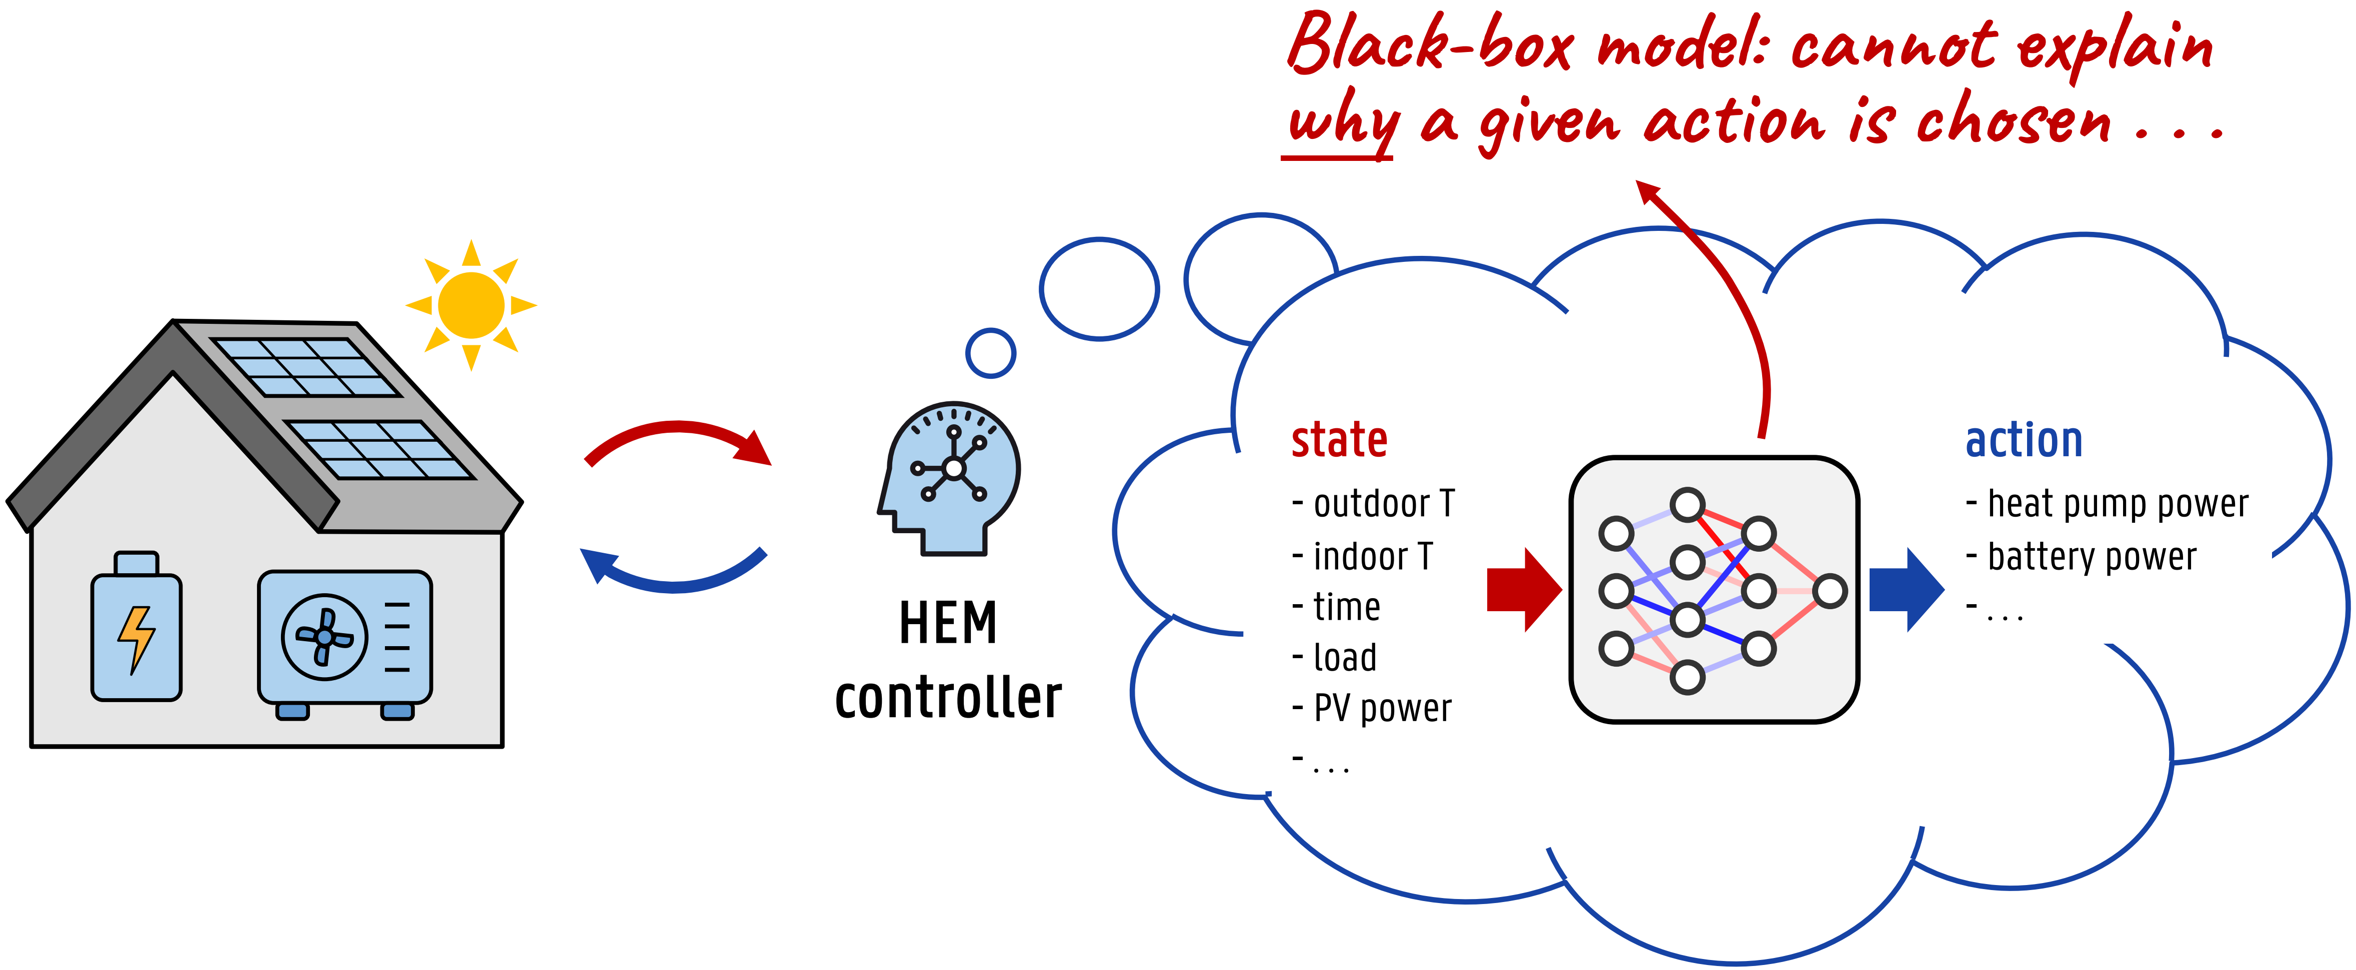

Black-box Reinforcement Learning

Idea: replace neural networks by something more explainable!

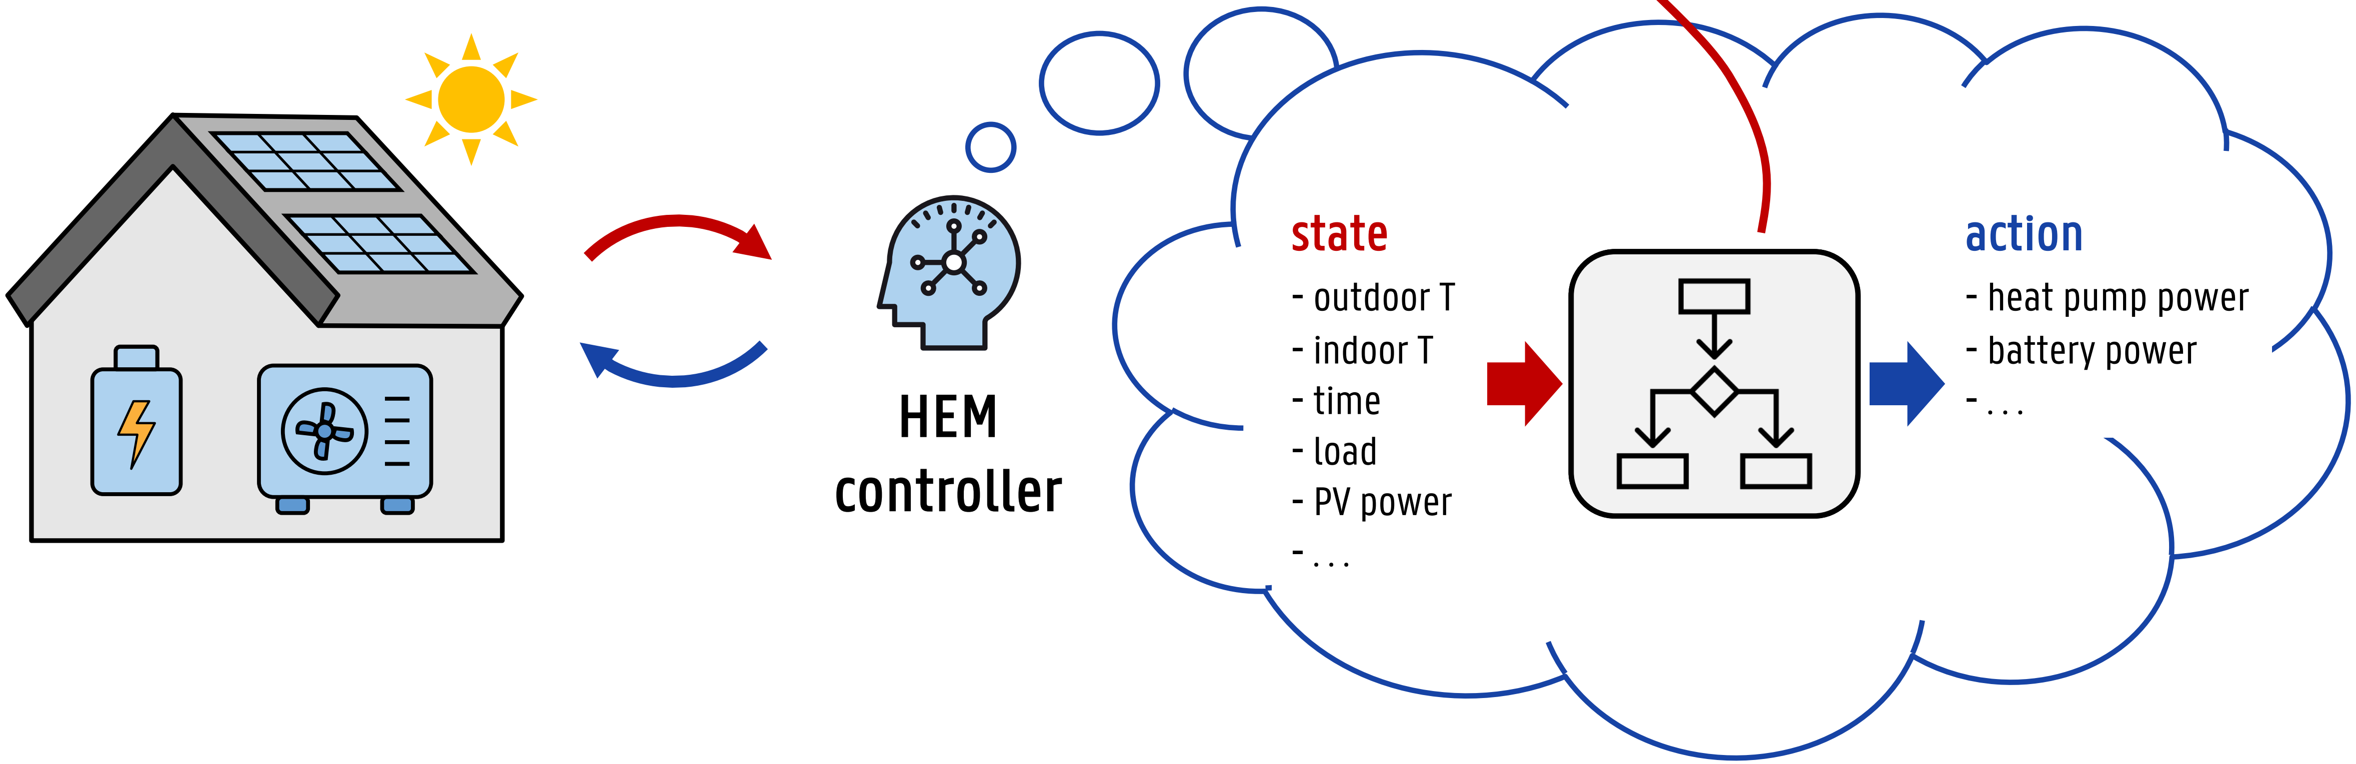

C. Develder, et al., "Home Energy Management in the context of consumer-centric power systems", Powering the future, VAIA - Flanders AI Research, 26 Nov. 2024

Enhancing explainability: Differentiable Decision Trees

"If it is sunny and household load is low, then charge battery"

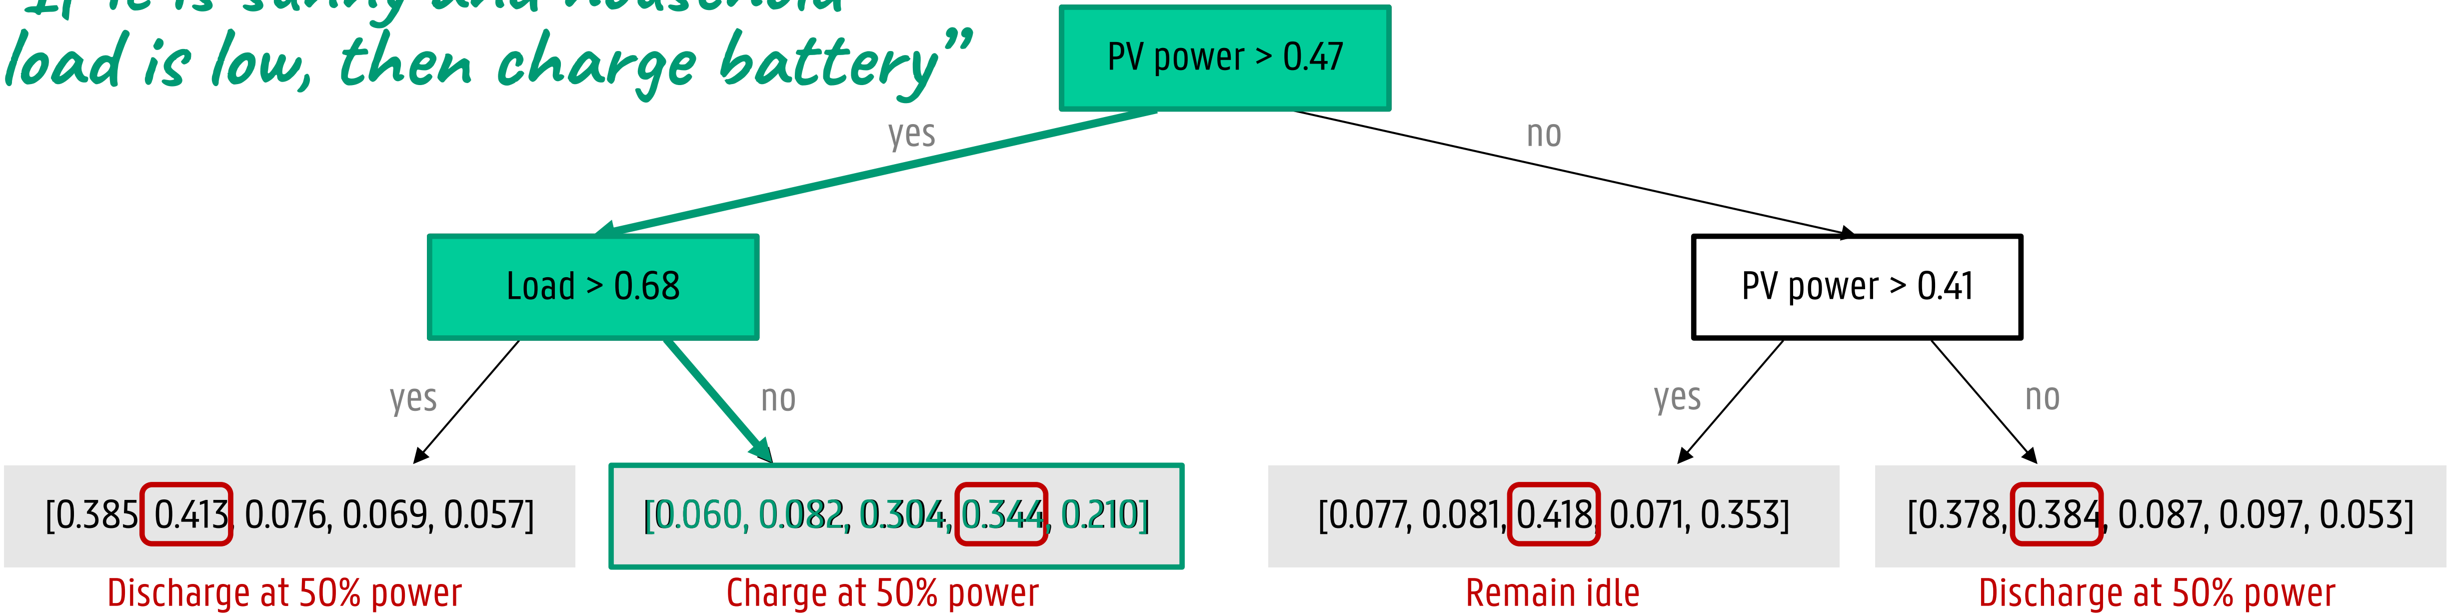

Distill2Explain: Learns the tree from trained RL policy

Step 1: train an RL "teacher" model

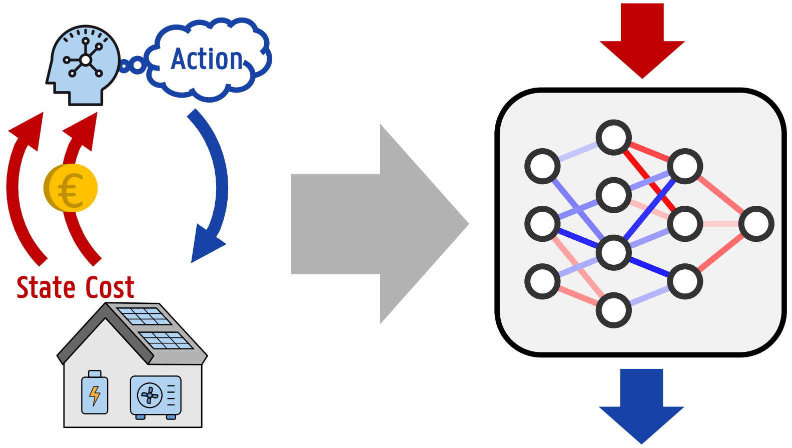

Step 2: distill a decision tree "student" model

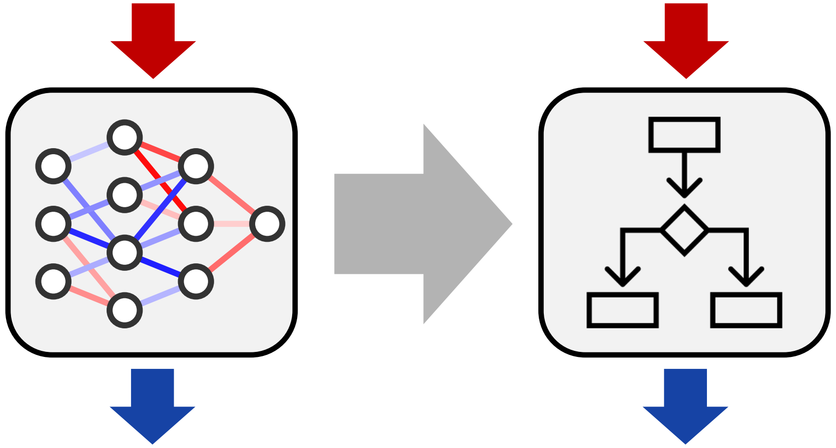

Distill2Explain: Learns the tree from trained RL policy

Step 1: train an RL "teacher" model

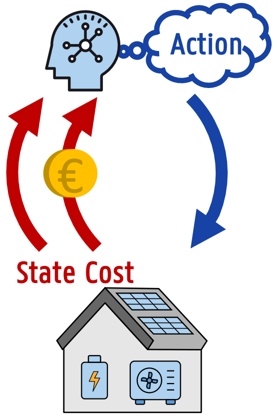

Step 2: distill a decision tree "student" model

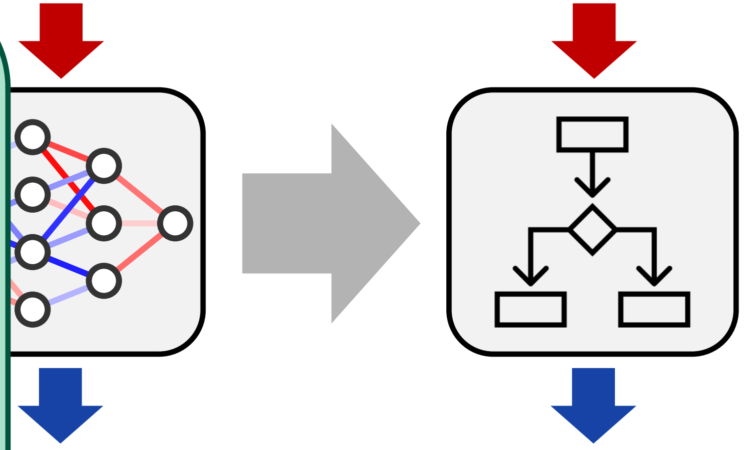

How? E.g., see [1] on how to distill **soft** decision trees. We learn **deterministic** trees

[1] Y. Coppens, K. Efthymiadis, T. Lenaerts, A. Nowé, T. Miller, R. Weber, and D. Magazzeni, "Distilling deep reinforcement learning policies in soft decision trees", In Proc. IJCAI 2019 Workshop on Explainable AI, pp. 1–6.

Sample result: Home battery control under time-varying pricing

- Household with rooftop PV and battery
- Day-ahead pricing (qh based) and capacity tariff (= based on max. peak load)

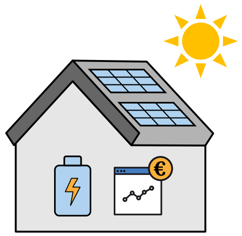

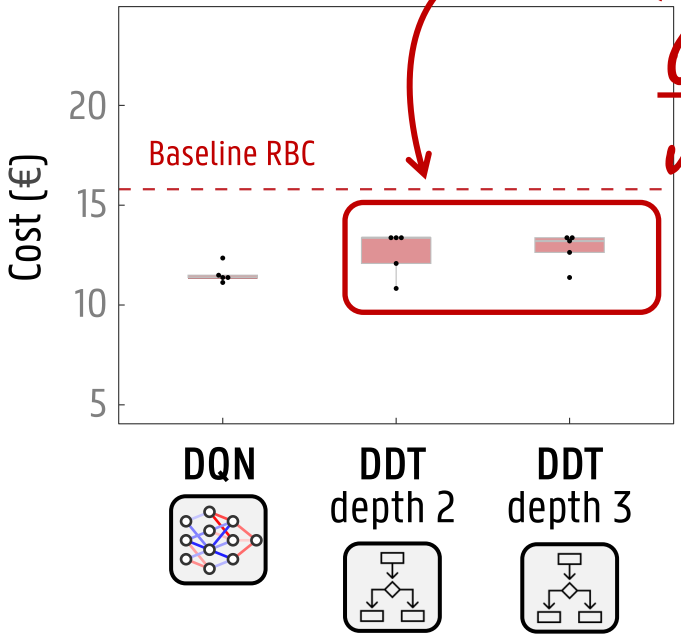

Comparable cost . . .
with simpler policies

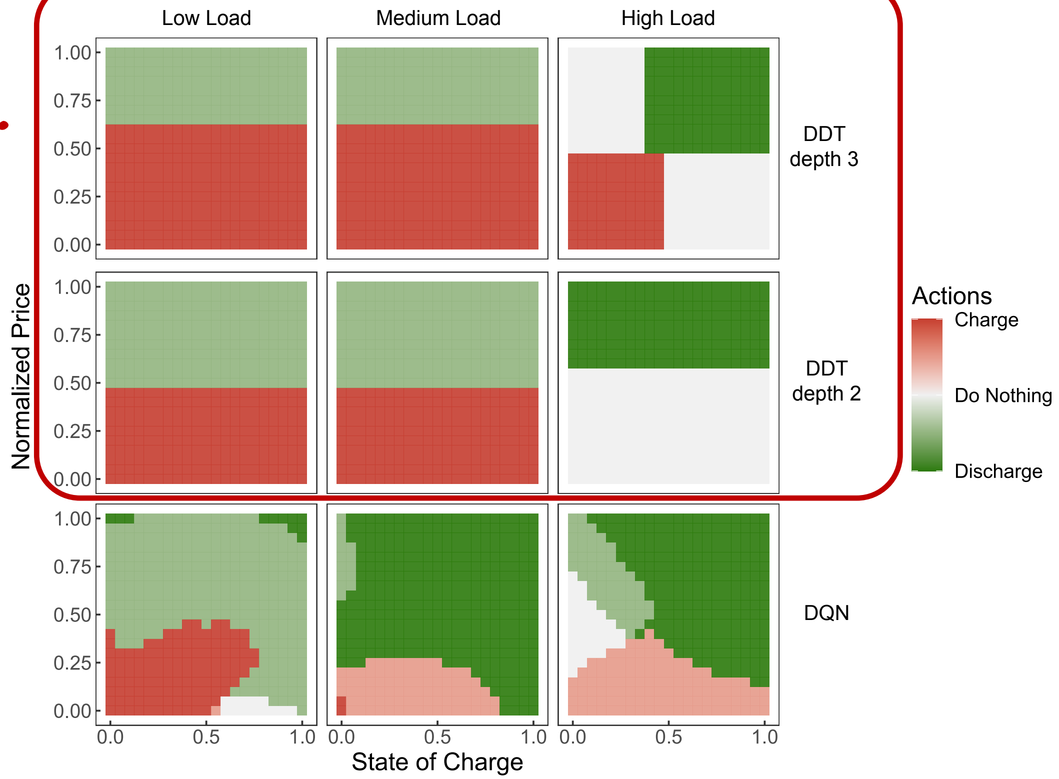

GHENT
UNIVERSITY

C. Develder, et al., "Home Energy Management in the context of consumer-centric power systems", Powering the future, VAIA - Flanders AI Research, 26 Nov. 2024

Thank you!
Any questions?

chris.develder@ugent.be

http://users.ugent.be/~cdvelder

It's not easy
being green ...

GHENT
UNIVERSITY

The AI4E team at IDLab, Ghent University

https://ugentai4e.github.io/