AI for Power Grids:
Challenges & Opportunities

IRT SystemX, Fraunhofer IEE, University of Amsterdam, INESC TEC

10/11/2026, Online

AI4REALNET has received funding from European Union's Horizon Europe Research and Innovation programme under the Grant Agreement No 101119527, and from the Swiss State Secretariat for Education, Research and Innovation.

Plan

- Introduction
- Context and challenges
- Opportunities
  - AI4REALNET Framework
  - Accelerate Grid simulations
  - AI recommender agent
  - Trustworthiness
  - Joint Human-AI decision-making
- Conclusion & perspectives
- Demonstration
- Q/A

Power Grid context and AI4REALNET framework

Milad Leyli-abadi (IRT SystemX)

AI4REALNET Project – objectives and scope

- Develop next generation of decision-making methods powered by supervised and reinforcement learning for **critical infrastructures**
- Ensure trustworthiness by design and during the interaction between human and AI
- Boost the development and validation of novel AI algorithms via 3 existing open-source AI-friendly digital environments

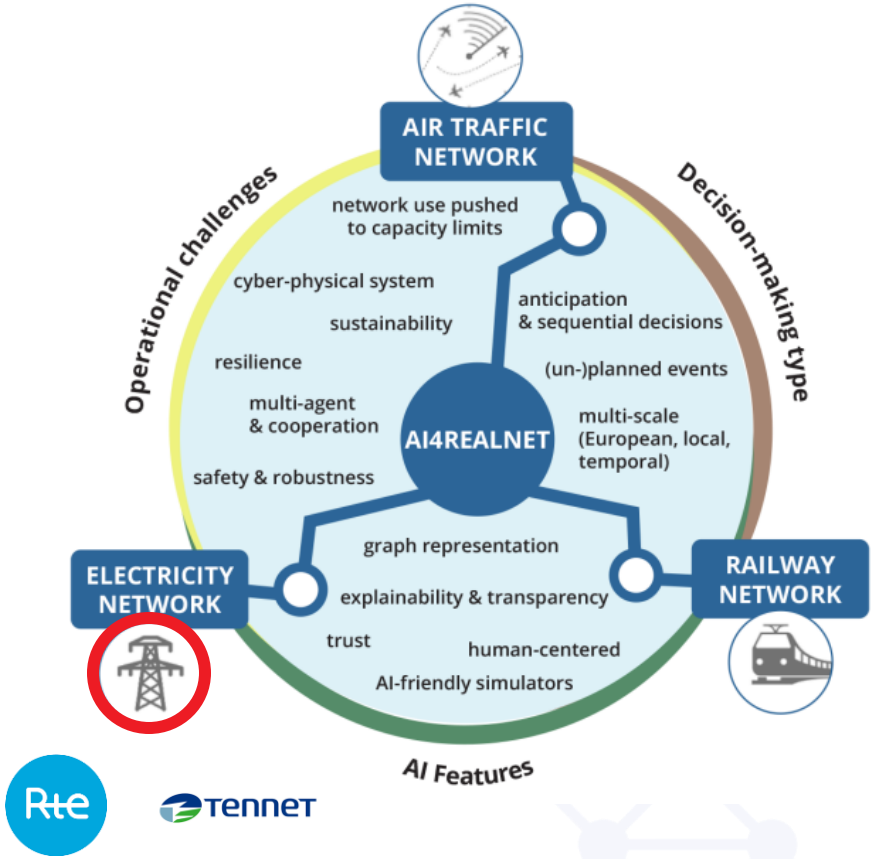

The Use Case: AI-Assisted Congestion Management

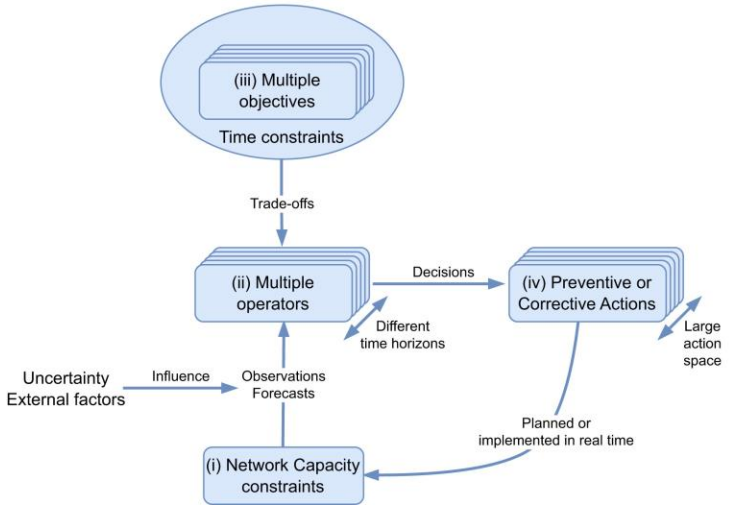

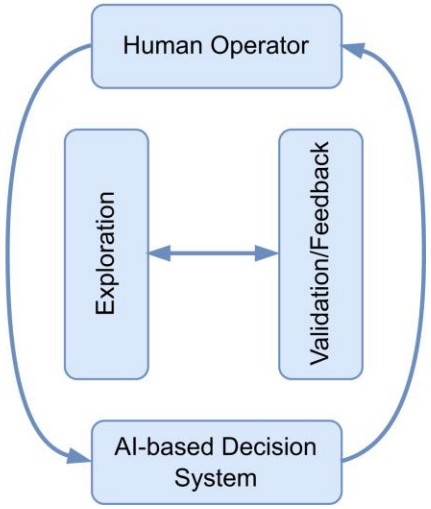





Why it matters:

- **Safety:** avoid overloads and blackouts
- **Sustainability:** maximize renewables, limit emergency redispatch to thermal plants
- **Workload:** ease the operator's load, don't flood them with alerts
- **Trust:** Explainable, and the human keeps the final word (AI-assisted human decision-making)

Context and challenges

A typical task of problem solving and decision-making

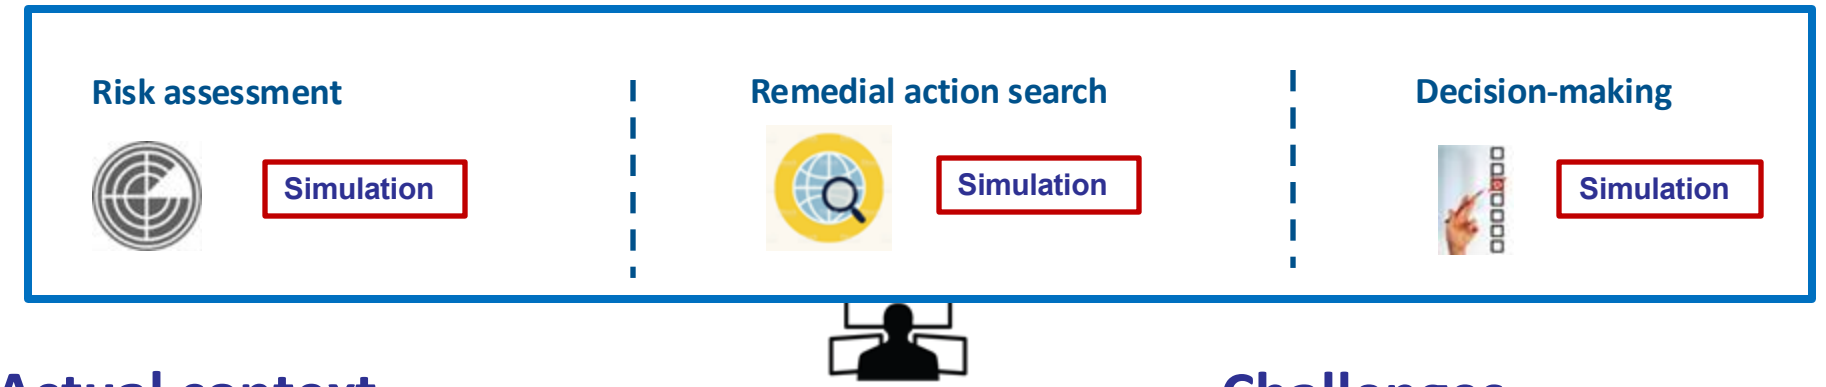



Actual context

- **>25,000 nodes / >10,000 lines** on a real transmission grid (RTE)
- **>4,000-dimensional observation space** and **>65,000 discrete actions -** for a network of just ~100 substations
- Decisions needed **every 5 minutes**, on a horizon from hours-ahead to 30-minutes-ahead
- Rising simulation load as renewables add forecast uncertainty

Challenges

- *Too-slow:* classical simulators can't keep up at this scale → **Accelerated physics solver**
- *Too big:* huge action space, scenarios unseen in training → **AI recommender agent**
- *Too opaque:* Will the operator understand and trust the recommendations → **Joint human-AI decision-making & trustworthiness**

Paper

AI4REALNET Framework

Holistic framework for human-AI collaboration by adopting an interdisciplinary approach

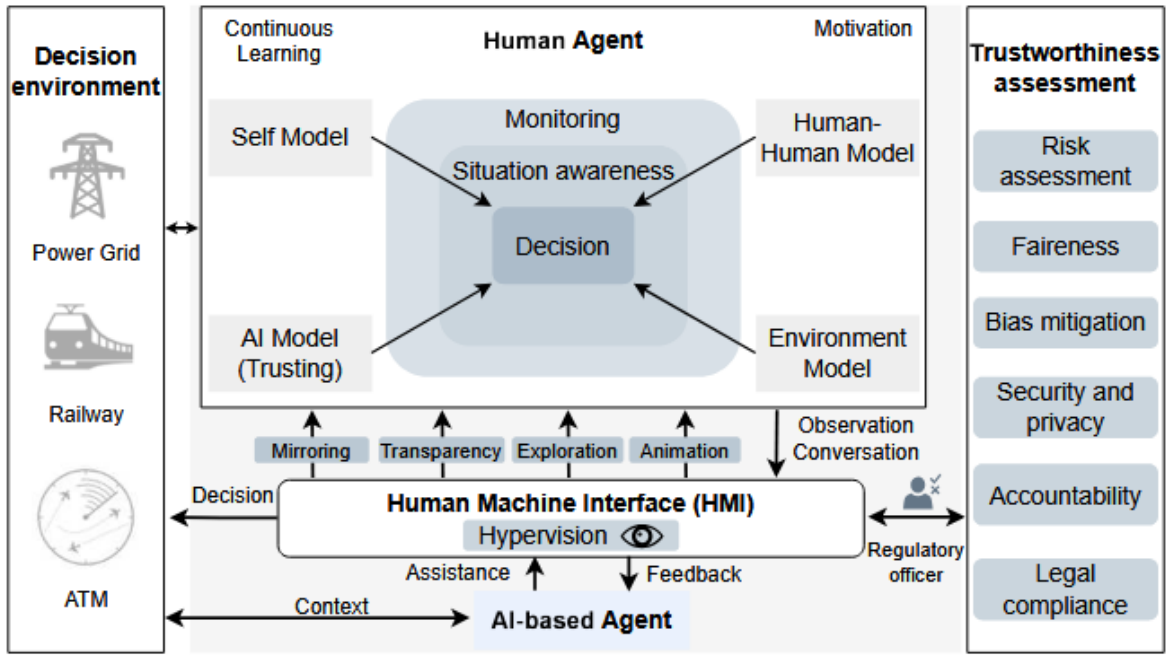

Decision-making from **human agent** point of view

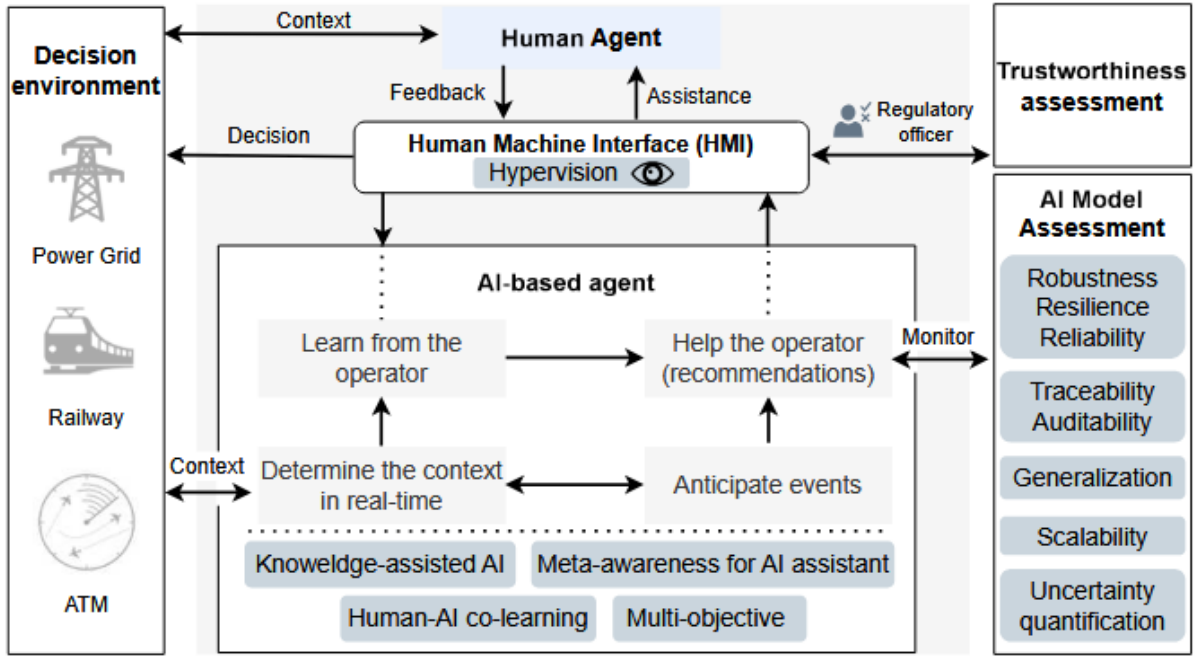

Decision-making from **AI-based agent** point of view

Github

From conceptualisation to real framework

- Digital environment
  - Grid2Op
  - GNN solver ✓
- AI assistants (Agent-as-a-service)
  - CurriculumAgent
  - ExpertAgent
  - Soft GNN ✓
- Trustworthiness
  - Explainability panels
  - Uncertainty Quantification ✓
  - Failure prediction ✓
- Human experience features
  - Co-learning and joint decision ✓
  - Stress & Cognition module
  - Human-related KPIs
  - Information logging

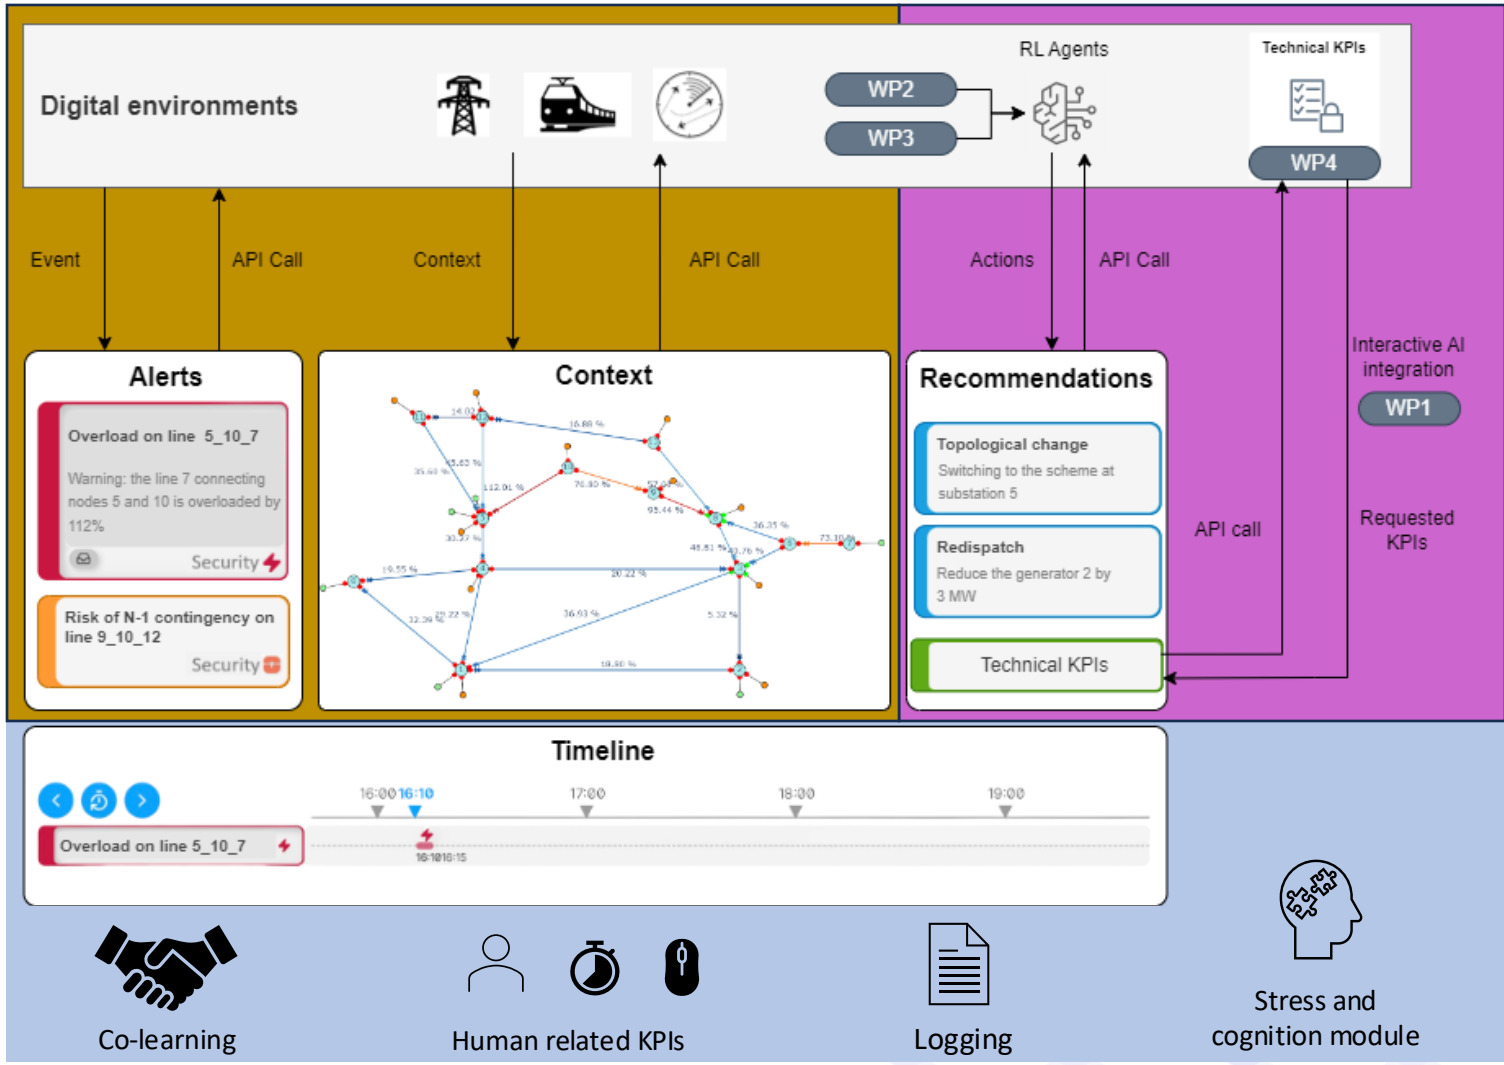

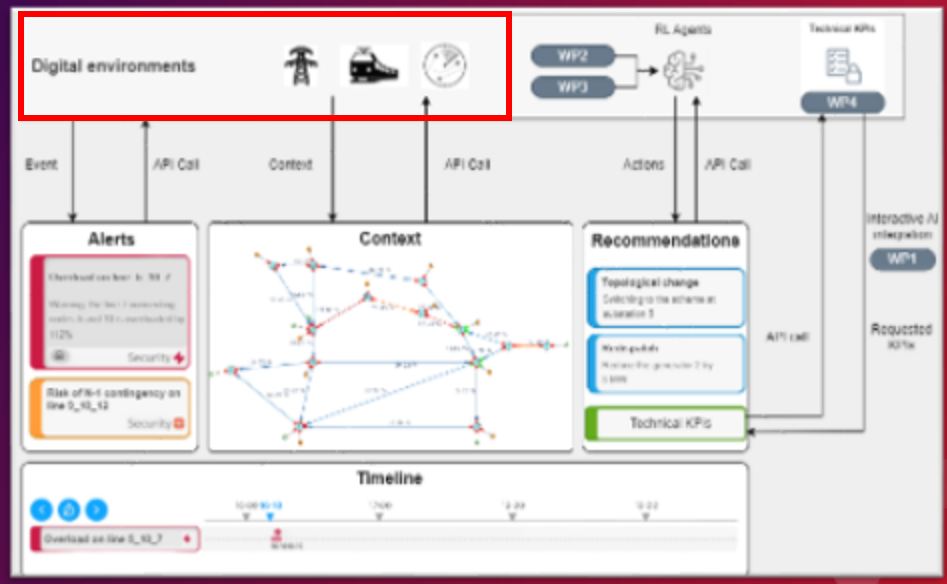

Accelerate grid simulations

Milad Leyli-abadi (IRT SystemX)

Accelerated grid simulation: Problem

- Currently used physical simulators
  - Inputs / Outputs

X = injections
(productions + loads)
T = topology

Numeric Solver
(physics model)

Y = power flows

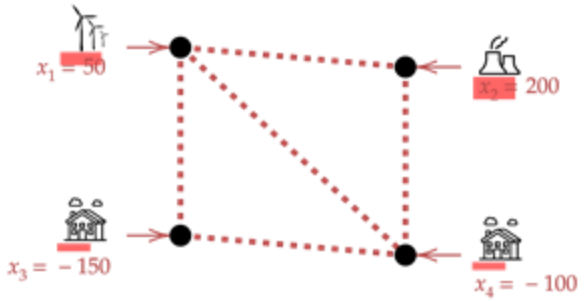

Accelerated grid simulation: Problem

- Currently used physical simulators
  - Inputs / Outputs

X = injections
(productions + loads)
T = topology

Numeric Solver
(physics model)

Y = power flows

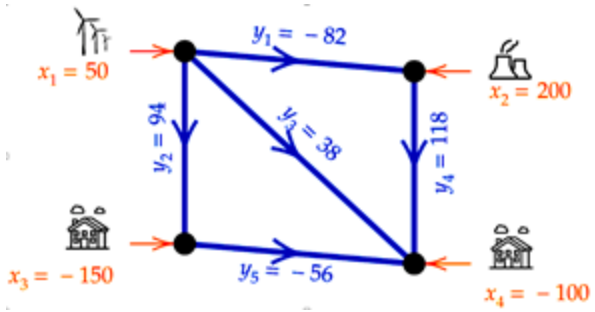

  - Characteristics
    - Relies on **physics equations (Kirchhoff law),** resolved by iterative optimization (Newton-Raphson)
    - Able to predict in a normal condition or different grid conditions

Power Grid equations

$$\begin{cases} 0 = -p_k + \sum_{m=1}^{K} |v_k||v_m|(g_{k,m}\cdot\cos(\theta_k-\theta_m) + b_{k,m}\sin(\theta_k-\theta_m)) & \text{Active power;} \\ 0 = q_k + \sum_{m=1}^{K} |v_k||v_m|(g_{k,s}\cdot\sin(\theta_k-\theta_m) - b_{k,m}\cos(\theta_k-\theta_m)) & \text{Reactive power} \end{cases}$$

Accelerated simulations using surrogates

- Multi-layer perceptron: All the available information are given as a flatten vector in the input space and regularized using Physics Constrained

Regularized multi-layer perceptron

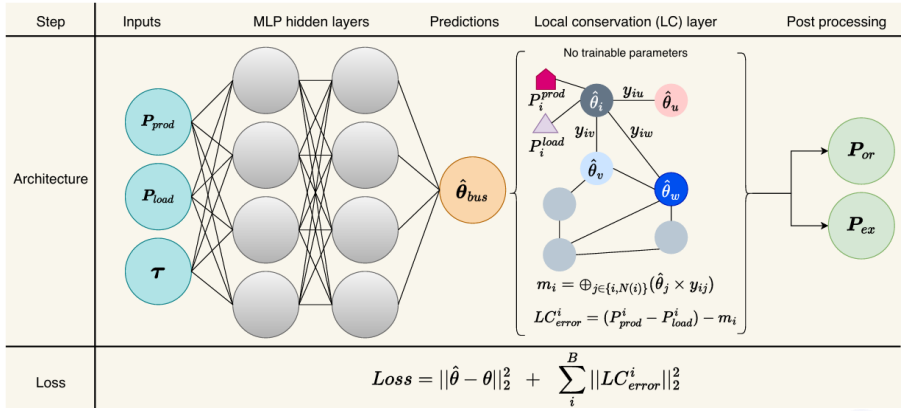

- Grid topology defines the graph structure and messages include a physical sense
  - The initial guess of parameters are learned accelerating their optimization

Physics-informed message-passing

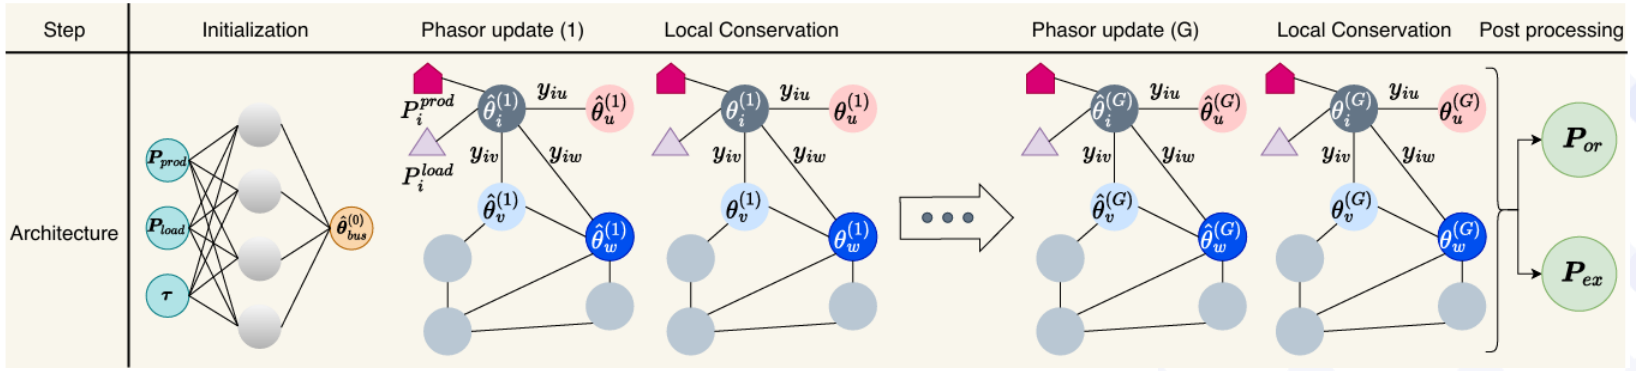

$$p_i^{prod} - p_i^{load} = (\theta_i \times y_{ii}) + \underbrace{(\theta_u \times y_{iu})}_{\text{message from node } u} + \underbrace{(\theta_v \times y_{iv})}_{\text{message from node } v} + \underbrace{(\theta_w \times y_{iw})}_{\text{message from node } w}$$

Datasets and distributions

Training and Test datasets

300 000 observations

100 000 observations

Out-of-distribution Test dataset

200 000 observations

To assess the robustness of models

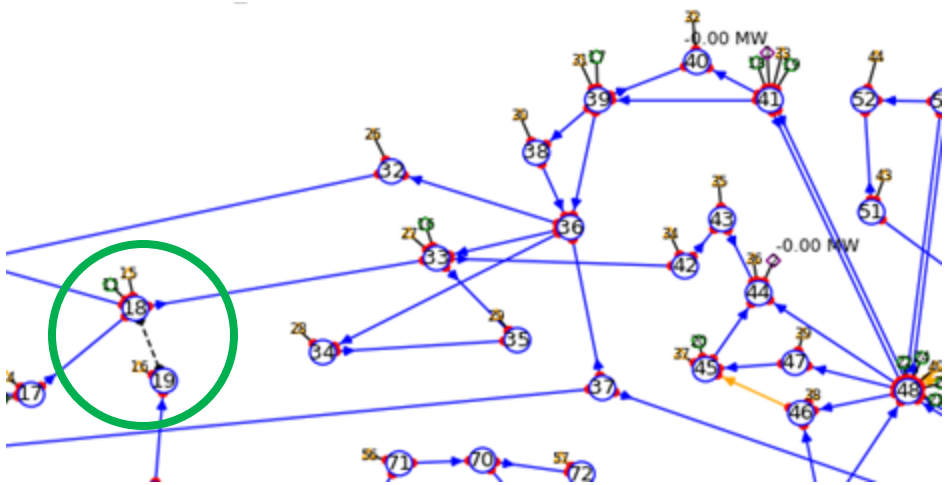

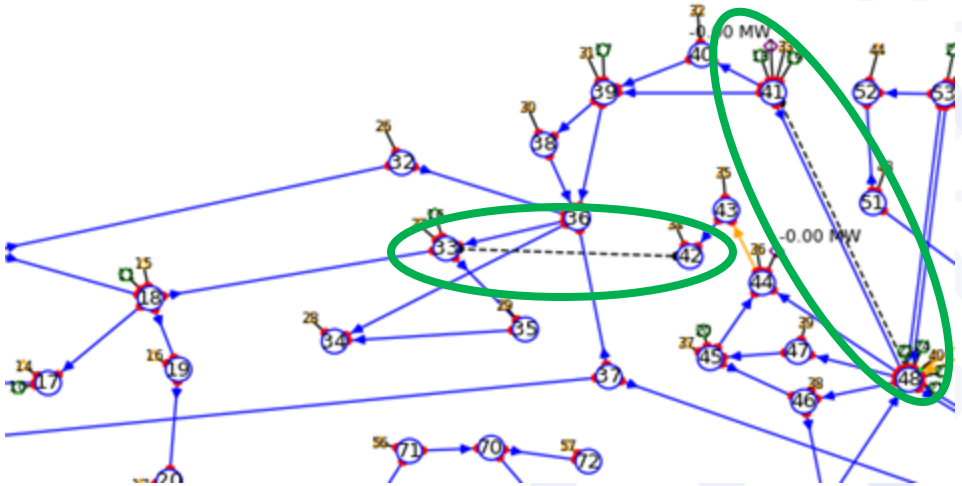

Power Grid case study: Evaluation results

- Results (The values are the violation percentage of the corresponding metric)



Test dataset

Test dataset	Loss target	Output	Disc lines	Loss pos	Energy loss consistency	Global conservation	Local conservation
FC	P	P	0.0	43	0.0	88	91
GNN	Θ	P	**0.0**	**0.0**	**0.0**	**0.0**	**0.0**



Out-of-Distribution Test dataset

Test OOD dataset	Loss target	Output	Disc lines	Loss pos	Energy loss consistency	Global conservation	Local conservation
FC	P	P	0.0	43	0.0	95	93
GNN	Θ	P	**0.0**	**0.0**	**0.0**	**0.0**	**3.08**

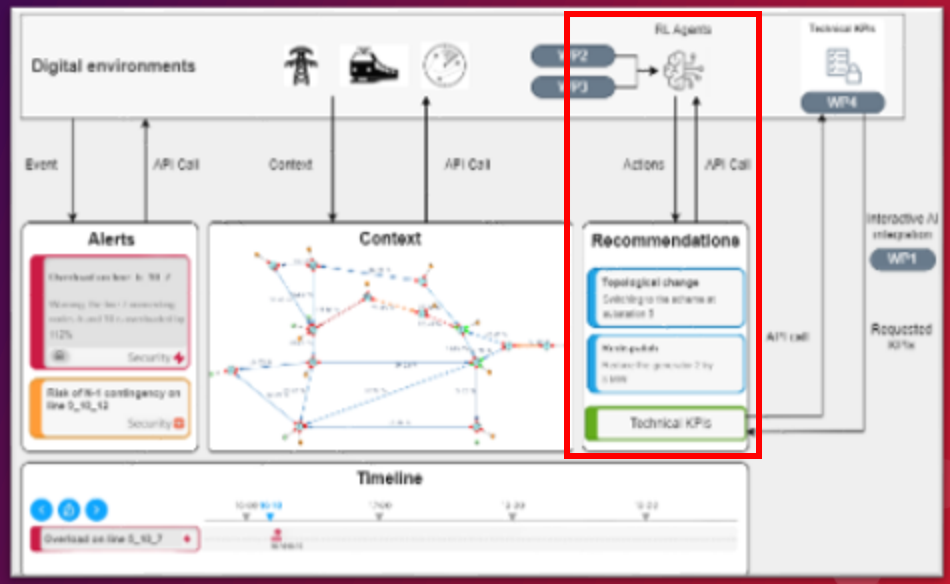

Soft-label imitation learning with Graph Neural Networks for topology control

Mohamed Hassouna (Fraunhofer IEE)

Topology actions to relieve congestion

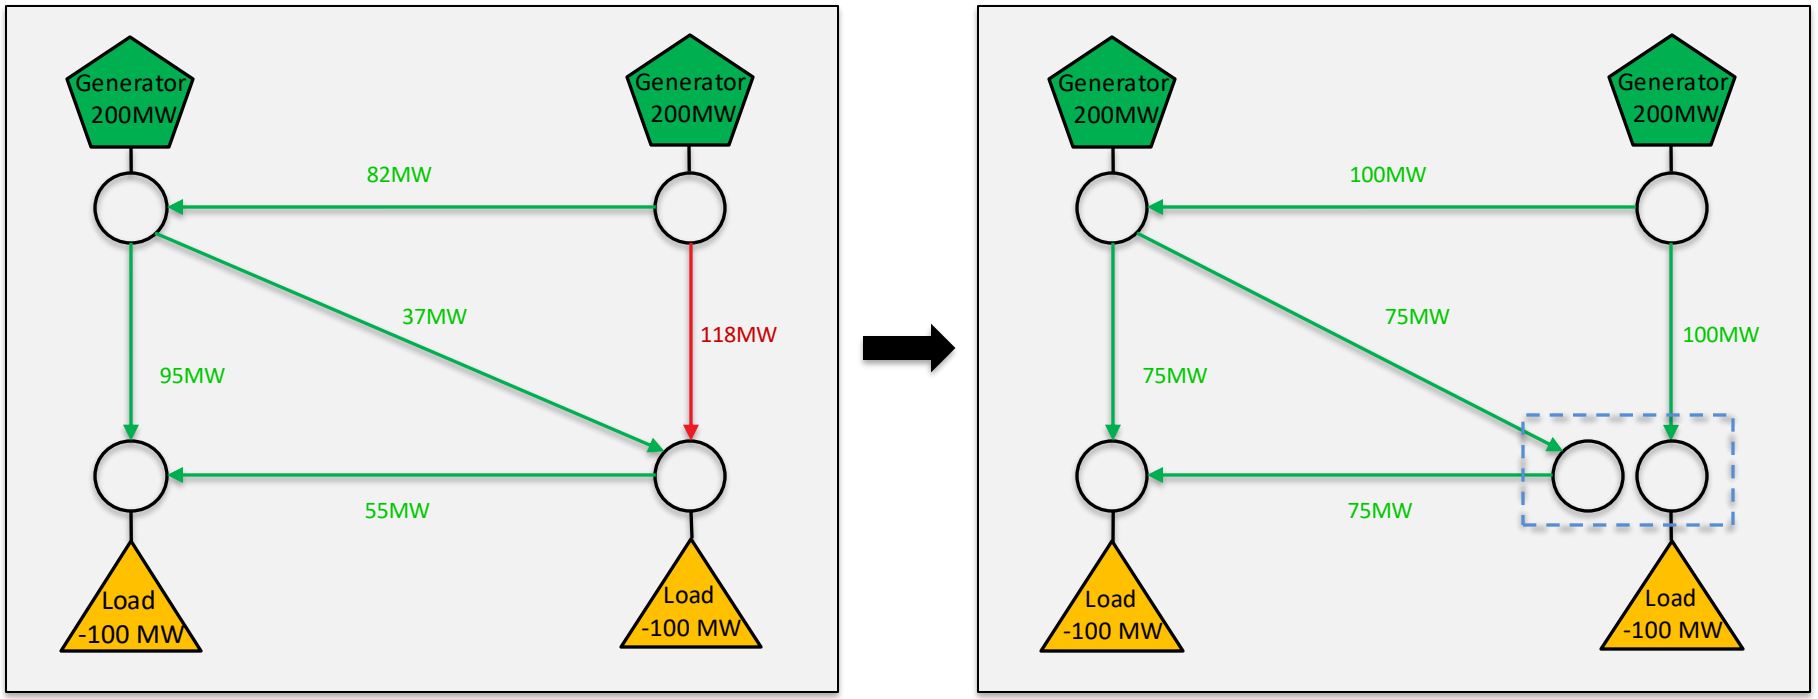

- **Redistribute power flows** by splitting a substation into two busbars – the overload disappears
- **Fast and cost-efficient:** no redispatch of generators, no curtailment of renewables
- **Challenge:** huge combinatorial action space → an AI agent has to recommend good actions within seconds

Starting point: CurriculumAgent

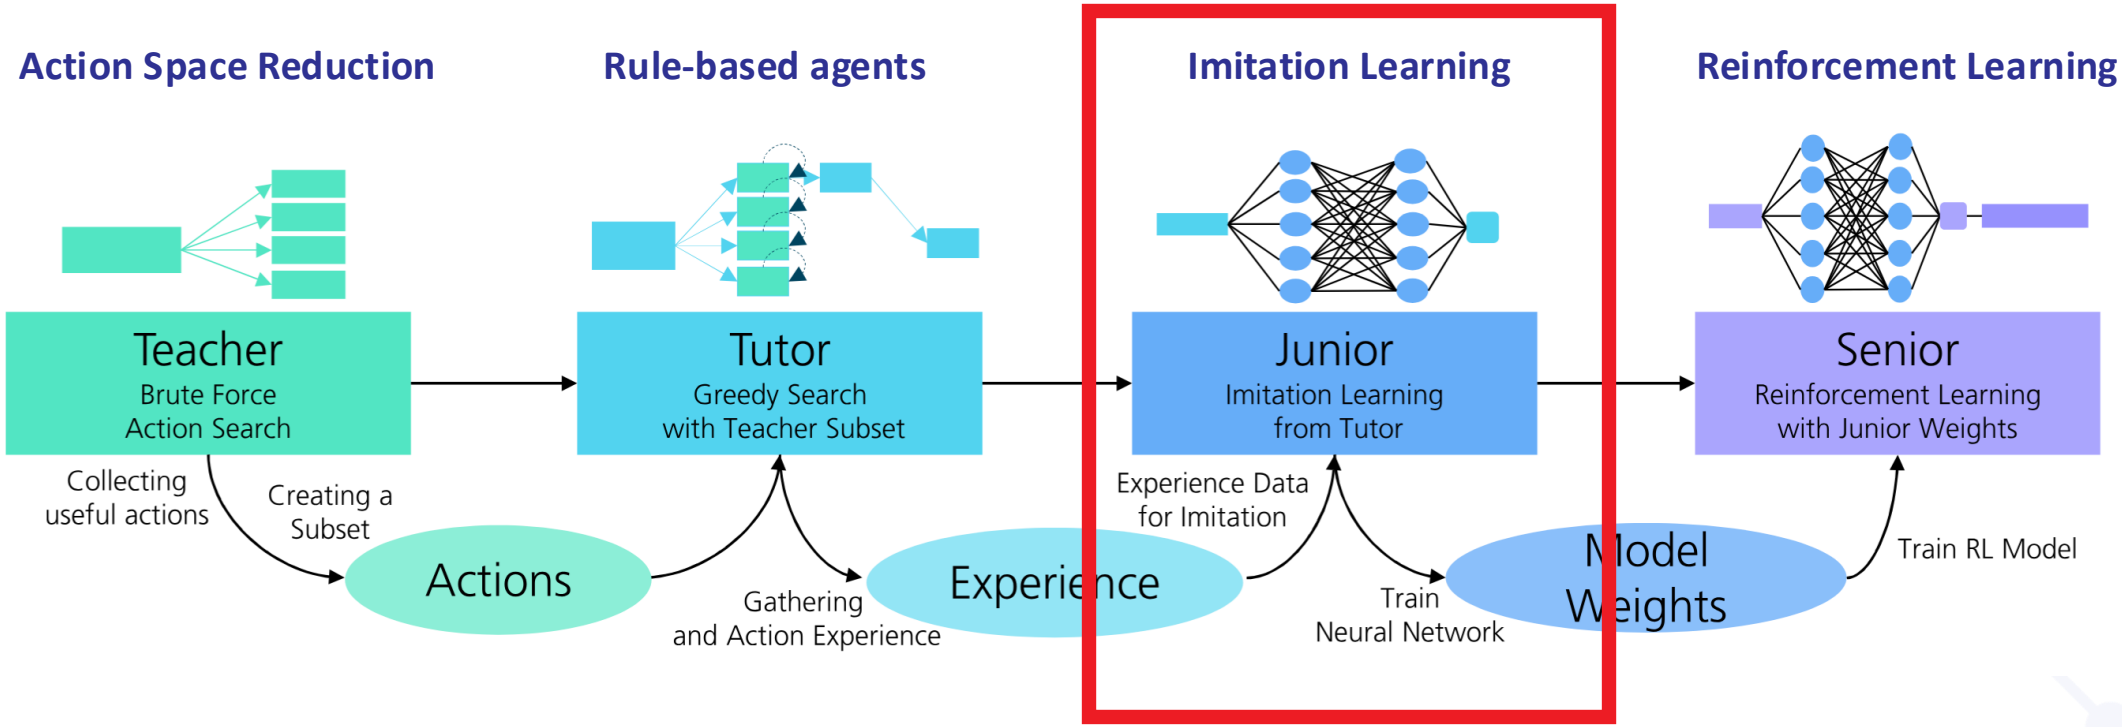

Source: Lehna et al. (2022). Managing power grids through topology actions: A comparative study between advanced rule-based and reinforcement learning agents.

- **Junior** learns by imitation from the Tutor's greedy decisions and is the warm start for the RL-trained **Senior**
- **Imitation learning** yields a strong policy without costly RL exploration – but it is only as good as its labels

Imitation learning: the research gap

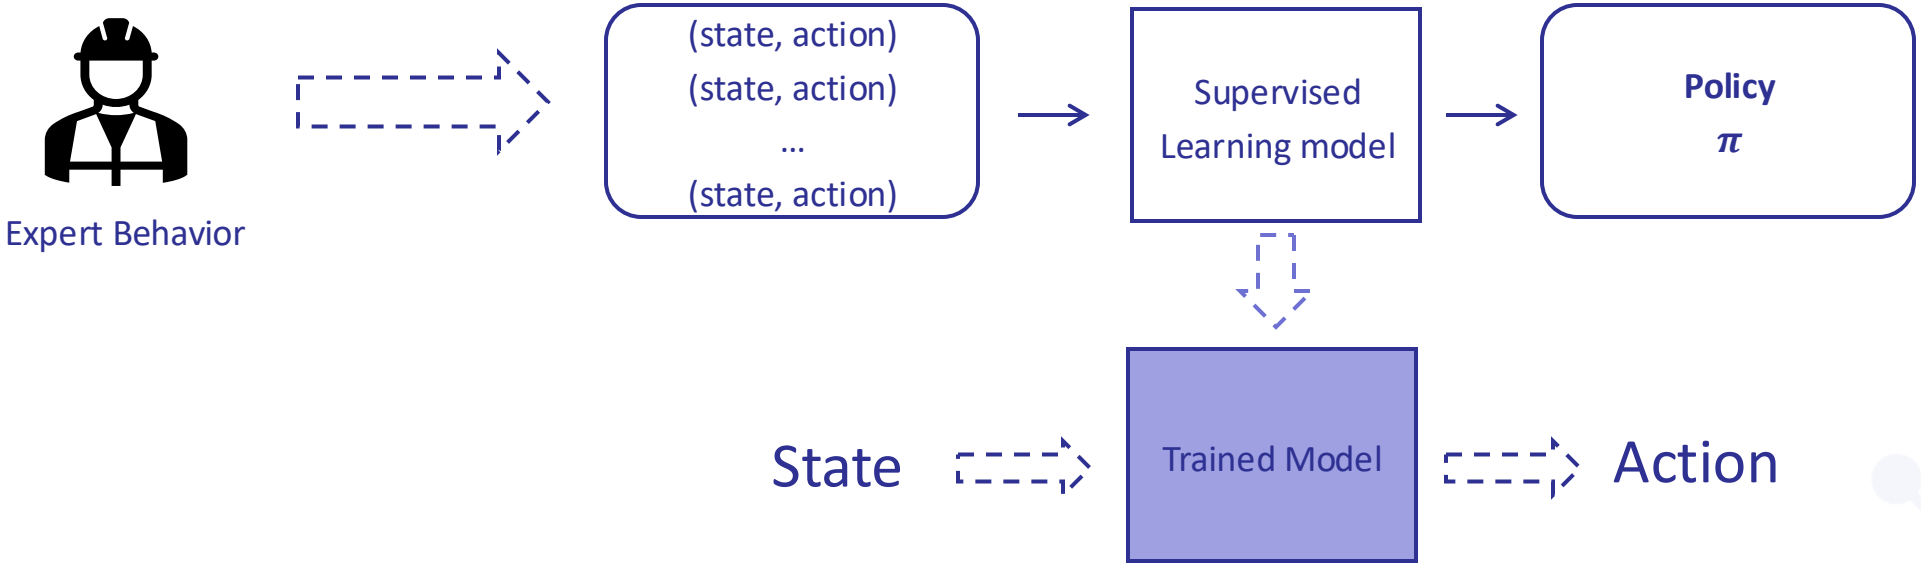

Existing imitation learning

- Assumes **one optimal action per state**
- Rigid policies
- Cannot outperform the expert – inherits its mistakes

Reality in topology control

- Often **many effective actions** for the same congestion
- **Need:** a method that captures the diversity of actions

SoftGNN at a glance

- **Soft-label imitation learning:** captures multiple viable actions
- **Graph Neural Networks (GNNs):** encode grid topology and spatial dependencies
- **Agent framework (SoftGNN):** integrates topology reversion and N-1 security evaluation
- **Significant performance gains** over the expert and state-of-the-art DRL for topology control

Paper

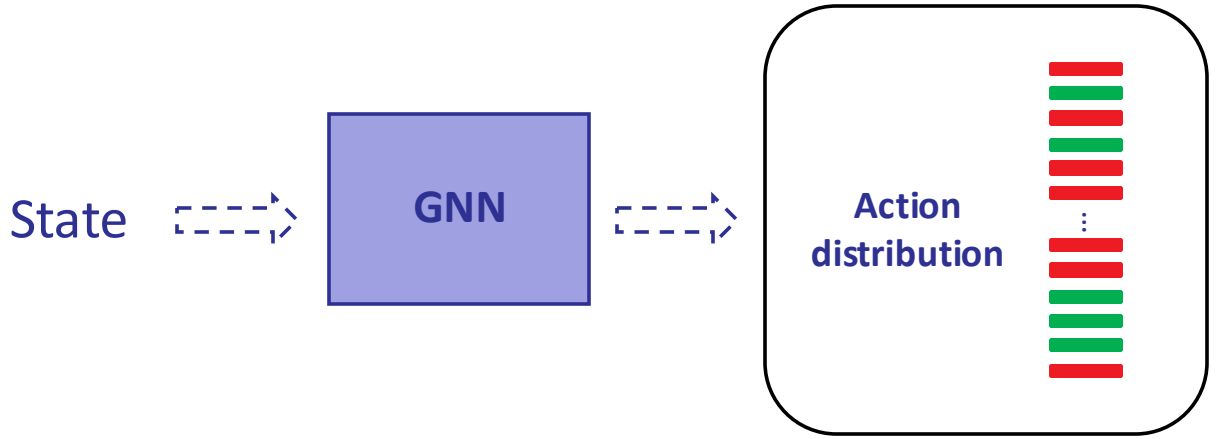

Hassouna, M. et al. (2026). Learning Topology Actions for Power Grid Control: A Graph-Based Soft-Label Imitation Learning Approach. ECML PKDD 2025, Applied Data Science Track, Springer.

Expert agent: greedy heuristic

Greedy expert (Tutor)

- Reactive: activates when the max. line loading $\rho_{max} \geq 90\ \%$
- Simulates **all actions** with a power flow and applies the one with the lowest max. line loading
- **Very costly:** many power-flow simulations per decision
- **Myopic:** only the current time step is considered

Why throw away all simulations of the greedy agent?

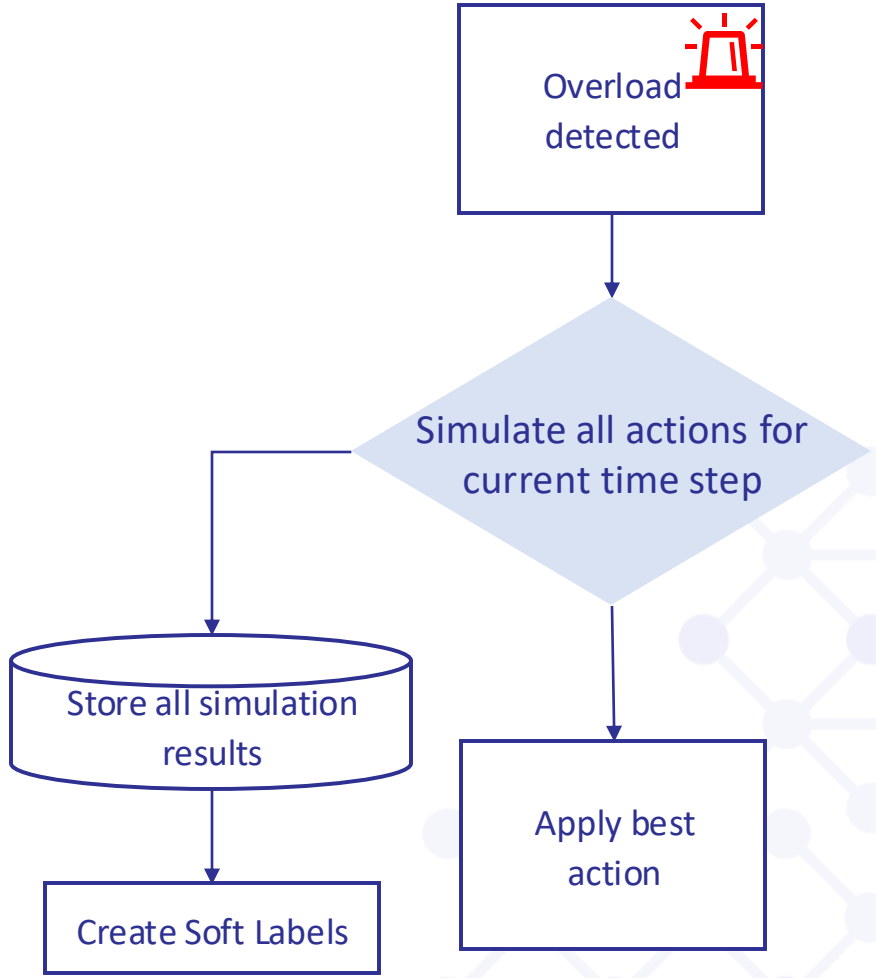

Soft labels from stored simulations

Concept: replace hard labels (only the greedy action) with **soft labels** – a probability distribution over all actions, derived from the stored power-flow results

$\rho_{max}(s, a)$: max. line loading after performing action a in state s τ: softmax temperature

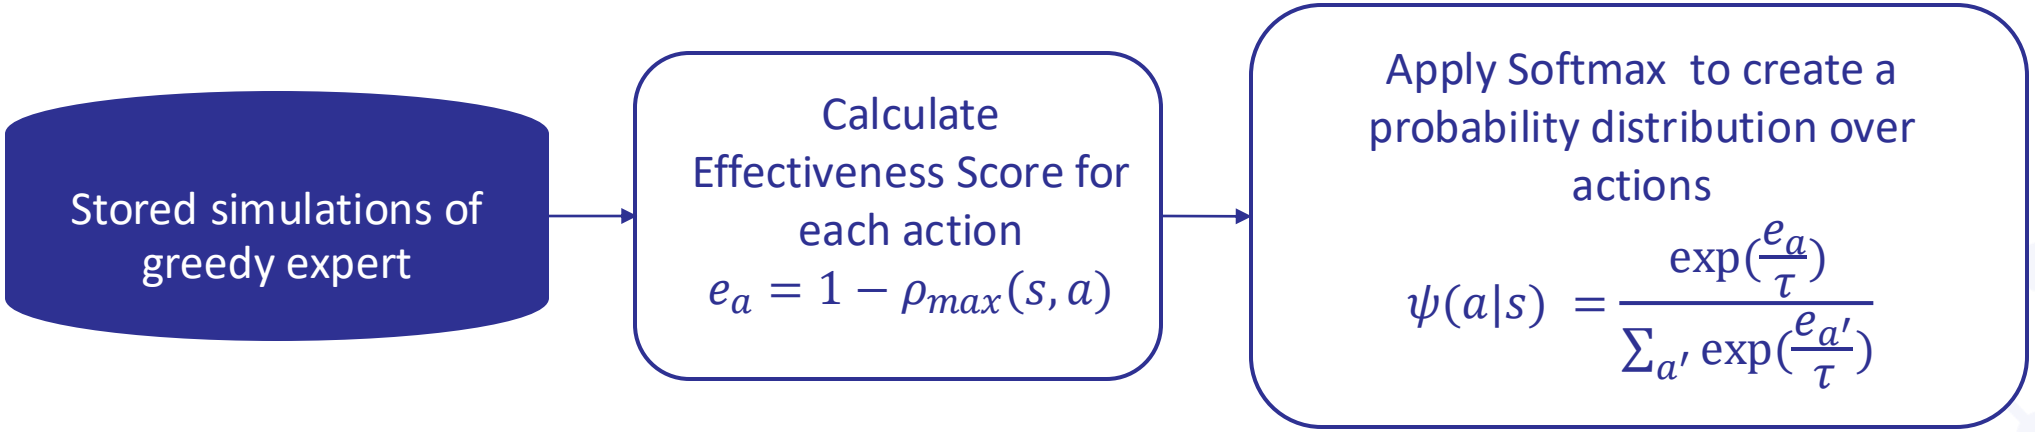

Inspired by knowledge distillation: the agent learns *which actions are good*, not only which one is best → richer labels and more action diversity

Graph representation of the grid

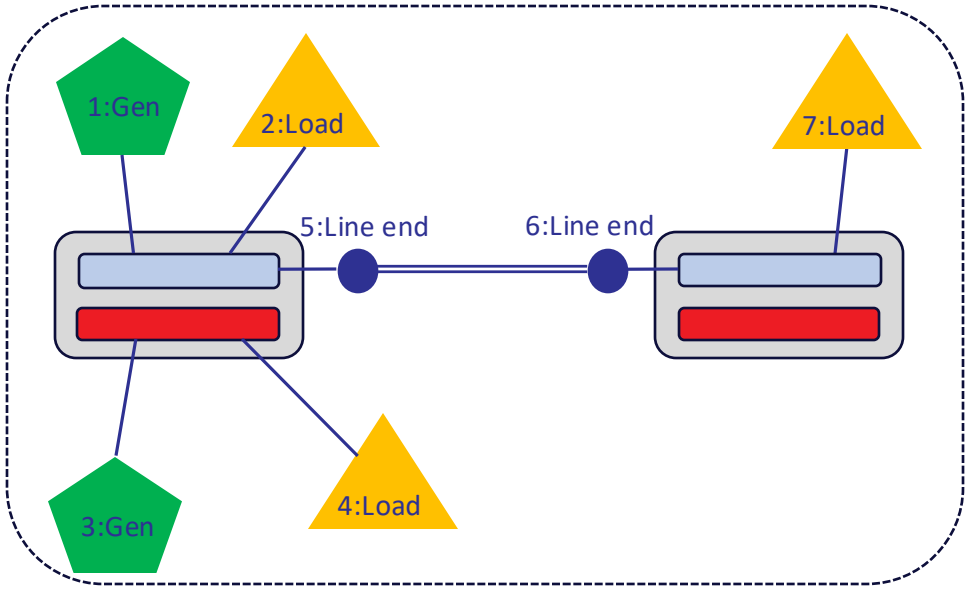

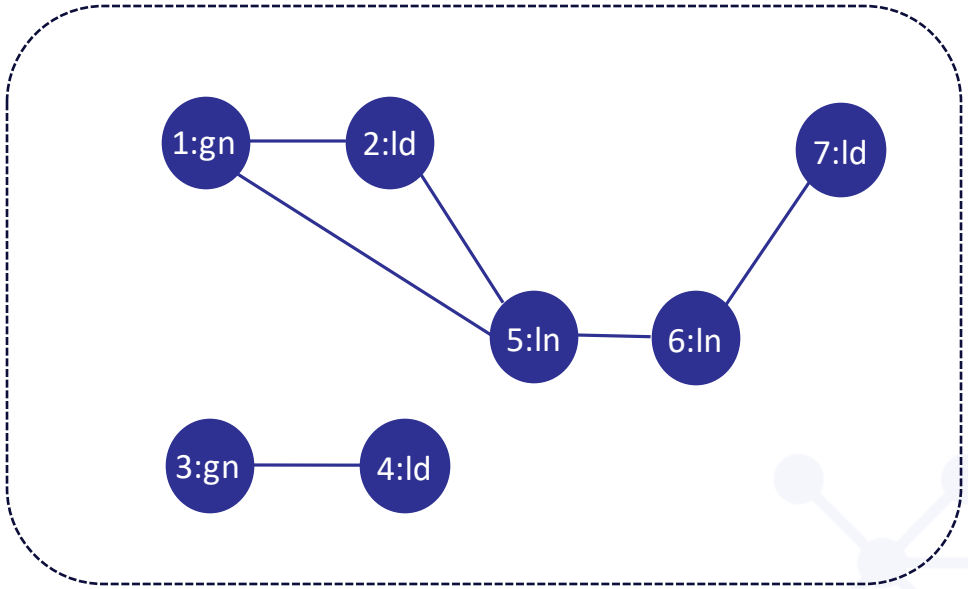

Grid as a graph

- Generators, loads and line ends become nodes
- Edges follow the current busbar assignment and the lines
- A topology action changes the graph itself

Node features

- Voltage, active and reactive power of generators and loads, line loadings
- Current topology assignment of each node
- Operational constraints: cooldowns of lines and substations

Model architecture: SoftGNN agent

Architecture

- Graph → GAT layers → max pooling → FNN → action distribution
- Trained against the soft labels with a KL-divergence loss

SoftGNN agent at runtime

- Activates if the max. line loading $\rho_{max} \geq 90\ \%$
- GNN ranks all actions → top candidates are validated by simulation
- Applies the first feasible action that reduces the overload

Enhancements

- Topology reversion and N-1 security check

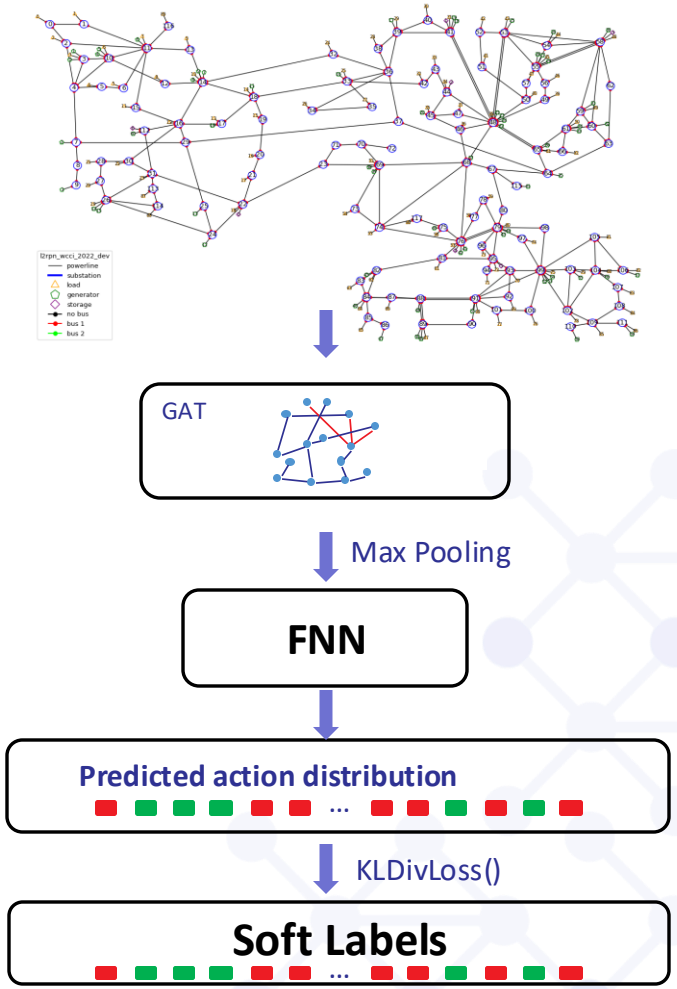

Results on L2RPN WCCI 2022

Agent	Type	L2RPN score	σ	Survival time
DoNothing	Baseline	0.00	0.00	229
$\text{Greedy}_{90\%}$	Expert	37.91	3.89	1014
$\text{Senior}_{95\%}$	DRL	37.13	4.49	988
$\text{TopoAgent}_{85–95\%}$	DRL	41.26	3.01	1232
$\text{HardFNN}_{90\%}$	Hard-label IL	37.54	3.87	1020
$\text{HardGNN}_{90\%}$	Hard-label IL	38.28	3.58	1048
$\text{SoftFNN}_{90\%}$	Soft-label IL	40.73	4.16	1113
$\text{SoftGNN}_{90\%}$	Soft-label IL	43.84	3.60	1293
$\text{SoftGNN}_{90\%}$ + N-1	**Soft-label IL**	**44.43**	**3.27**	**1299**

Key findings

- **Soft > hard labels & expert:** +17 % vs. greedy expert, +16 % vs. HardGNN
- **GNN > FNN** for both label types
- **SoftGNN + N-1 extension:** best overall score
- **Outperforms DRL agents** (Senior, TopoAgent)

L2RPN score: mean and standard deviation σ; survival time: median time steps. Averaged over 20 random seeds on l2rpn_wcci_2022; subscripts give the activation threshold of the max. line loading.
Hassouna, M. et al. (2026). Learning Topology Actions for Power Grid Control: A Graph-Based Soft-Label Imitation Learning Approach. ECML PKDD 2025, Applied Data Science Track, Springer.

➢ Trained now on both ai4realnet_small (36 bus) and ai4realnet_large (updated 118 bus) grids

Multiple recommendations in InteractiveAI

From one action to several options

- Soft labels give a **probability distribution over all actions** → SoftGNN recommends the top-k actions, not only the best one
- Operators choose among alternatives – the human keeps the final decision
- Fallback options if the first action is not acceptable
- Controllable number of recommendations

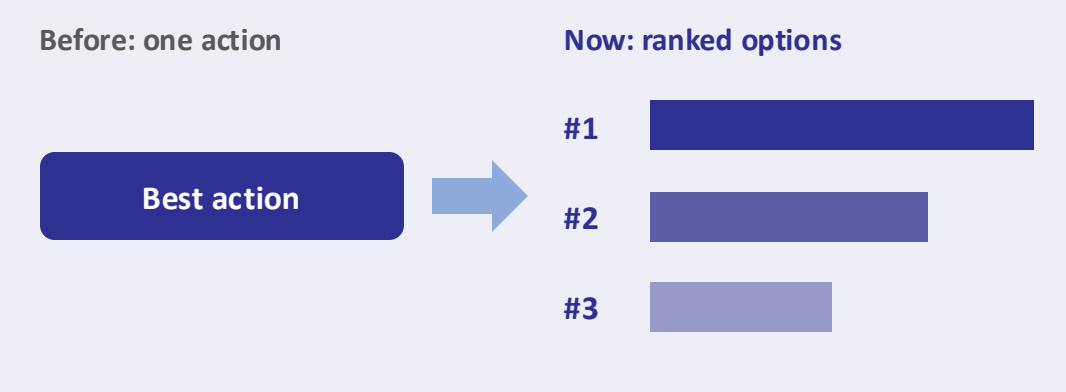

Integrated in InteractiveAI

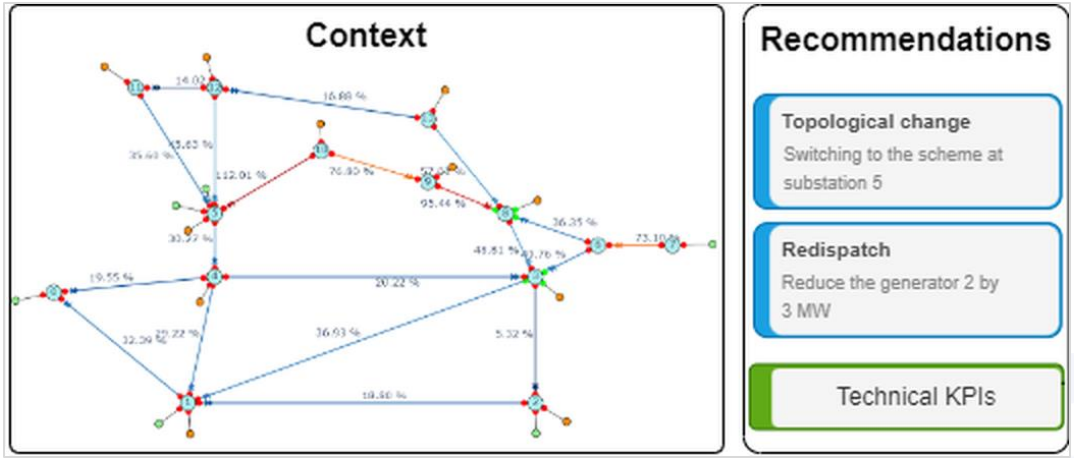

InteractiveAI operator view (schematic)

- SoftGNN runs as **AI assistant** in the AI4REALNET framework
- Ranked recommendations appear next to alerts and grid context in the operator's **Recommendations** panel



Trustworthiness in joint decision-making

Ricardo Jorge Bessa (InescTec)

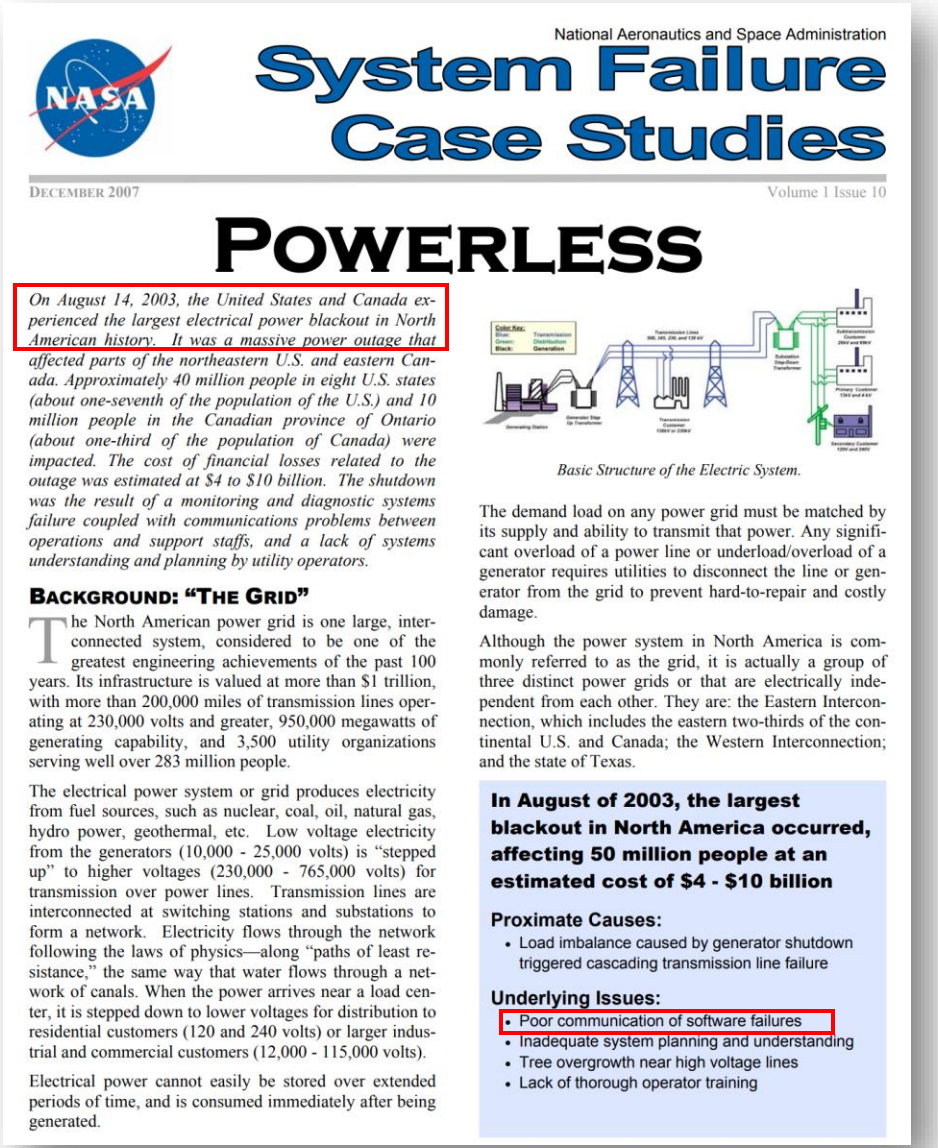

National Aeronautics and Space Administration

System Failure Case Studies

DECEMBER 2007 Volume 1 Issue 10

POWERLESS

On August 14, 2003, the United States and Canada experienced the largest electrical power blackout in North American history. It was a massive power outage that affected parts of the northeastern U.S. and eastern Canada. Approximately 40 million people in eight U.S. states (about one-seventh of the population of the U.S.) and 10 million people in the Canadian province of Ontario (about one-third of the population of Canada) were impacted. The cost of financial losses related to the outage was estimated at $4 to $10 billion. The shutdown was the result of a monitoring and diagnostic systems failure coupled with communications problems between operations and support staffs, and a lack of systems understanding and planning by utility operators.

BACKGROUND: "THE GRID"

The North American power grid is one large, interconnected system, considered to be one of the greatest engineering achievements of the past 100 years. Its infrastructure is valued at more than $1 trillion, with more than 200,000 miles of transmission lines operating at 230,000 volts and greater, 950,000 megawatts of generating capability, and 3,500 utility organizations serving well over 283 million people.

The electrical power system or grid produces electricity from fuel sources, such as nuclear, coal, oil, natural gas, hydro power, geothermal, etc. Low voltage electricity from the generators (10,000 - 25,000 volts) is "stepped up" to higher voltages (230,000 - 765,000 volts) for transmission over power lines. Transmission lines are interconnected at switching stations and substations to form a network. Electricity flows through the network following the laws of physics—along "paths of least resistance," the same way that water flows through a network of canals. When the power arrives near a load center, it is stepped down to lower voltages for distribution to residential customers (120 and 240 volts) or larger industrial and commercial customers (12,000 - 115,000 volts).

Electrical power cannot easily be stored over extended periods of time, and is consumed immediately after being generated.

Basic Structure of the Electric System.

The demand load on any power grid must be matched by its supply and ability to transmit that power. Any significant overload of a power line or underload/overload of a generator requires utilities to disconnect the line or generator from the grid to prevent hard-to-repair and costly damage.

Although the power system in North America is commonly referred to as the grid, it is actually a group of three distinct power grids or that are electrically independent from each other. They are: the Eastern Interconnection, which includes the eastern two-thirds of the continental U.S. and Canada; the Western Interconnection; and the state of Texas.

In August of 2003, the largest blackout in North America occurred, affecting 50 million people at an estimated cost of $4 - $10 billion

Proximate Causes:

- Load imbalance caused by generator shutdown triggered cascading transmission line failure

Underlying Issues:

- Poor communication of software failures
- Inadequate system planning and understanding
- Tree overgrowth near high voltage lines
- Lack of thorough operator training

*"…**software timing error** in FirstEnergy's UNIX-based XA/21 **energy management computer was found to be the primary cause of the grid event alarm failure**. After the alarm system failed silently, the unprocessed events started to queue up and crashed the primary server within 30 minutes.*

*"**The operators only determined they had problems when data from phone calls received from customers**, nearby utilities, and their regional coordinating center calls did not match the information on their screens."*



What if the failure was in an AI assistant?

Self-awareness of the AI assistants

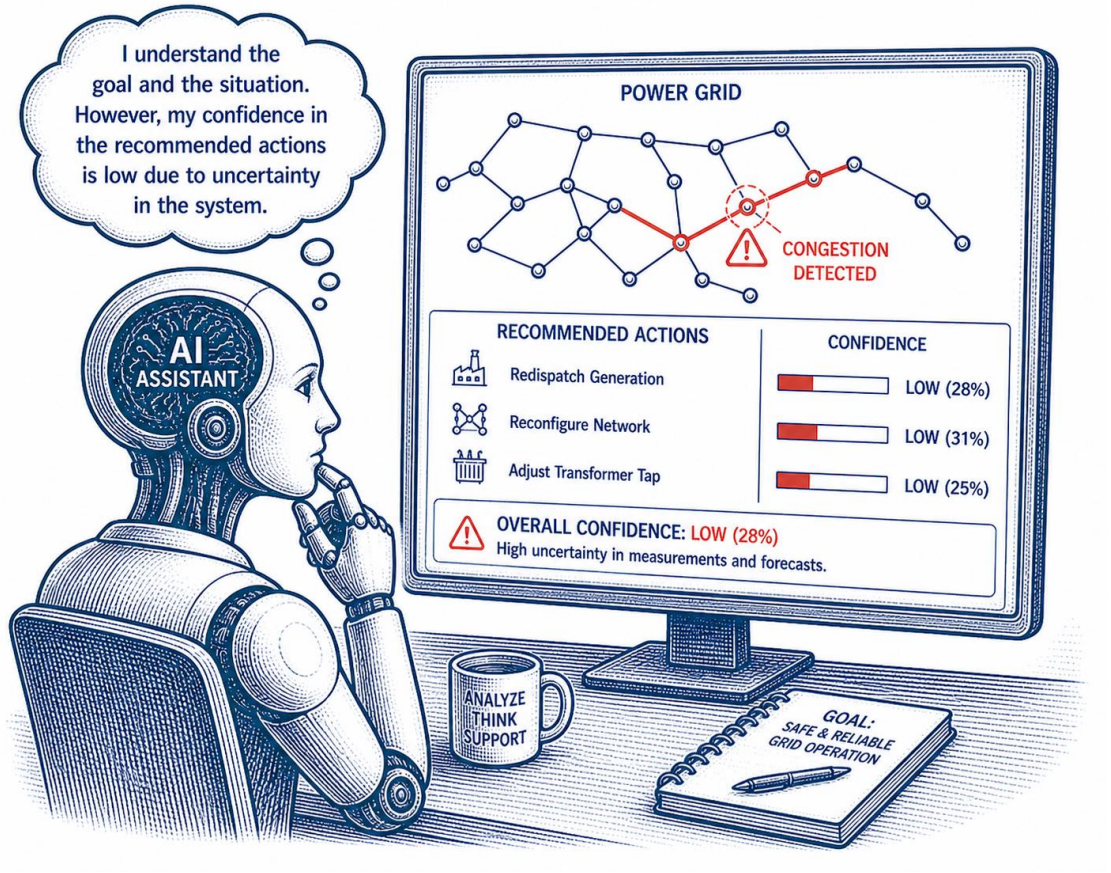

[Generated with ChatGPT based on my design specifications and instructions provided through prompting]



AI **confidence** in its decisions and **failure prediction** can be used to determine when **control** should be **deferred to human** operators

Objective & contributions

- **Goal: predict the reliability of a pre-trained RL agent, and warn the operator early**
  - Estimate the probability that the agent fails to resolve congestion one hour ahead
- **Contribution 1 — AI reliability prediction**
  - Combine epistemic + aleatoric uncertainty with indicators of grid operational complexity
- **Contribution 2 — interpretable models**
  - A novel LLM-based method that distils the predictor into compact if-then rules with natural-language explanations

Prediction target

$$\hat{p}_{\text{fail},\, t+k|t} = P(Y_{t+k} = 1 \mid v_t)$$ *failure of the RL agent at t+k*

Application case: Grid2Op & the RL agent

- **Environment: Grid2Op**
  - l2rpn_icaps_2021_small scenario
  - Based on the IEEE 118-bus system:
    - 36 substations · 59 lines
    - 22 generators · 37 loads
  - 5-minute resolution (one hour ahead = t+12)

- **RL agent: CurriculumAgent**
  - Teacher — exhaustive search over topology actions → reduced action set
  - Tutor — rule-based, generates observation–action trajectories
  - Junior — neural network trained by imitation learning
  - Senior — initialised from Junior, refined with PPO

Two sources of uncertainty

- **Epistemic — model uncertainty**
  - *“Does the agent recognise this state?”*
  - Evidential Neural Network (ENN): parameterises a Dirichlet over the topology actions
  - Trained by behaviour cloning on the agente’s own states
  - Low in familiar states, high out-of-distribution (OoD)
- **Aleatoric — data uncertainty**
  - *“How much irreducible variability is there?”*
  - Two-stage Histogram-based Gradient Boosting (HGB)
    - 1) mean model predicts load & generation
    - 2) second model predicts squared residuals (Poisson) → local variance

Failure-prediction methodology

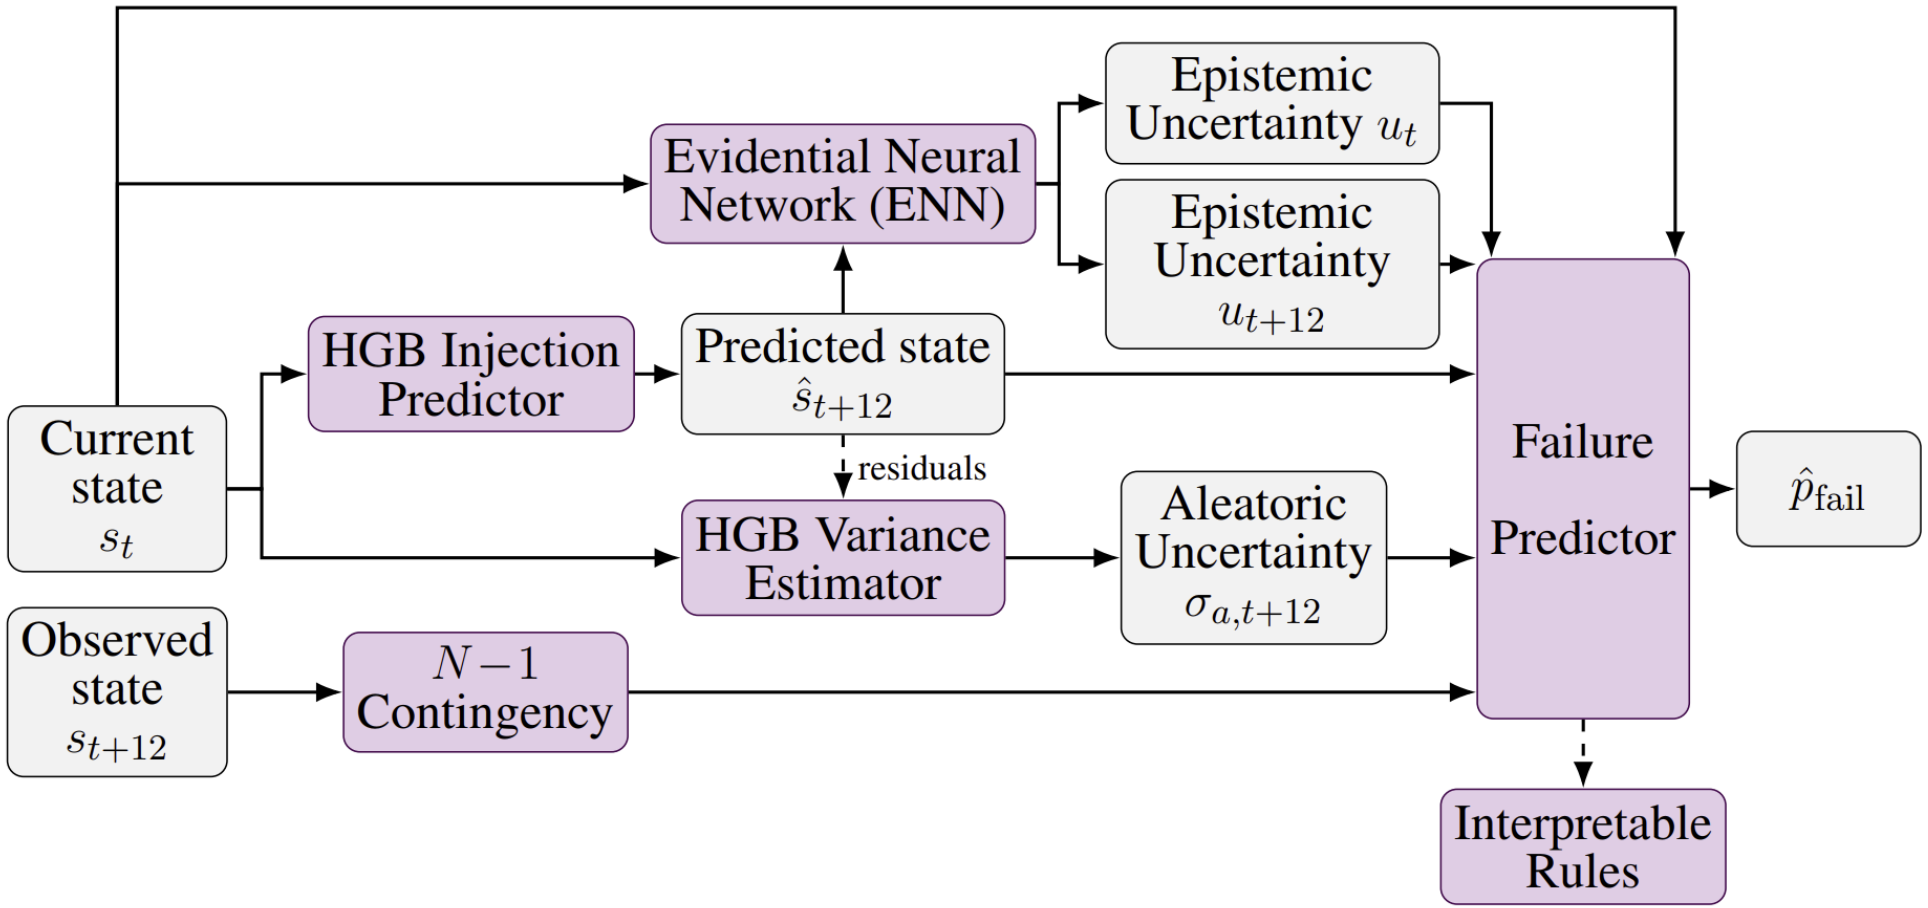

Interpretable rules: a readable output

The LLM distils the HGB predictor into an executable if-then rune – one per contingency, in plain language.

Example – extracted rule for the contingency on line 34_35_110

The RL agent is predicted to fail when the load-to-generation ratio at t is <= 0.986 and either the mean aleatoric reactive power load uncertainty for t+12|t is >= 0.185 or the RL epistemic uncertainty for t + 12 | t is >= 0.803.

HGB (reference)	Decision Tree	LLM Rule
~ 592 trees	737 leaves, depth 14	A few conditions
Most accurate, but opaque	Same data, similar accuracy, but hard to inspect	Readable & translatable to natural language
Same decision, very different transparency		

Results: one-hour-ahead failure prediction

Configuration	F2	Accuracy	Oversight Rate (OVR) ↓	False Alarms ↓
Baseline (grid variables)	70.8	85.7	7.7	14.9
Epistemic only	35.4	42.7	20.1	60.8
Grid + aleatoric	73.6	87.6	7.1	12.9
Grid + aleatoric + epistemic	**76.5**	**89.5**	**6.8**	**10.8**

Aleatoric alone already helps (↑F2, ↓FA, ↓OVR)

Epistemic alone is poor

Combining all three wins: best F2 and accuracy, lowest OVR & FA

Results: compact rules vs the full model

Model	F2	Accuracy	OVR ↓	FA ↓
LLM rule	**73.8**	**89.4**	**13.8**	**11.3**
Decision tree	72.4	87.1	14.4	12.8
HGB (reference)	76.5	89.5	6.8	10.8

Compactness

LLM rule: a few conditions per contingency

Decision tree: 737 leaves

HGB: ~592 trees (opaque)

LLM rules match/beat the tree and approach the HGB — while staying readable



How to predict outcomes under joint human-AI decision making?

Herke van Hoof (Universtiy of Amsterdam)

Explainable AI

- Explain AI models or decisions to improve human understanding of these systems
- Many techniques make more transparent which inputs map to which outputs
  - Mechanistic interpretability (e.g. decision trees)
  - Feature attribution
  - Contrastive explanations
- For decision support, perhaps more relevant: future outcomes of a decision
  - How to do this when actions are finally chosen by (human) decision maker, not AI?

De Peuter, Oliehoek & van Hoof, Performative explanations: Explaining futures under future explanations (submitted)

Outcome predictions

- Outcomes are properties $\psi^{\pi}(s_t, a_t)$ of the future trajectory when decisions are made under some **policy** π *.
- We consider these properties possible *objectives* of the (human) decision maker
- Explanation consists of outcome prediction
- In literature, outcomes predicted under the **AI policy**

* technically, a successor feature: the expected (discounted) cumulative value of some state-action features $\phi(s_t, a_t)$,

De Peuter, Oliehoek & van Hoof, Performative explanations: Explaining futures under future explanations (submitted)

Decision making example



Rawpixel.com (CC0)

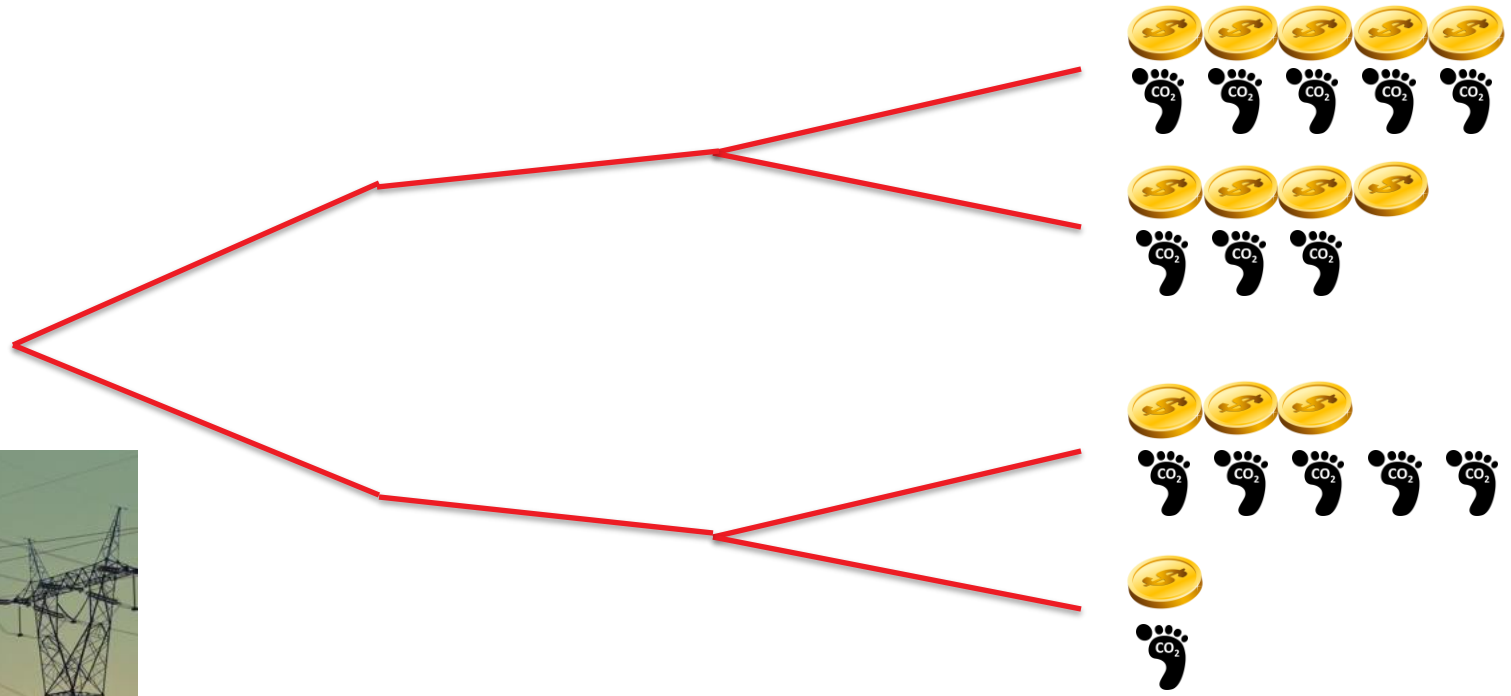

https://publicdomainvectors.org/ (CC0)
Wikimedia / Tommaso.sansone91 (CC0)

De Peuter, Oliehoek & van Hoof, Performative explanations: Explaining futures under future explanations (submitted)

Decision making example

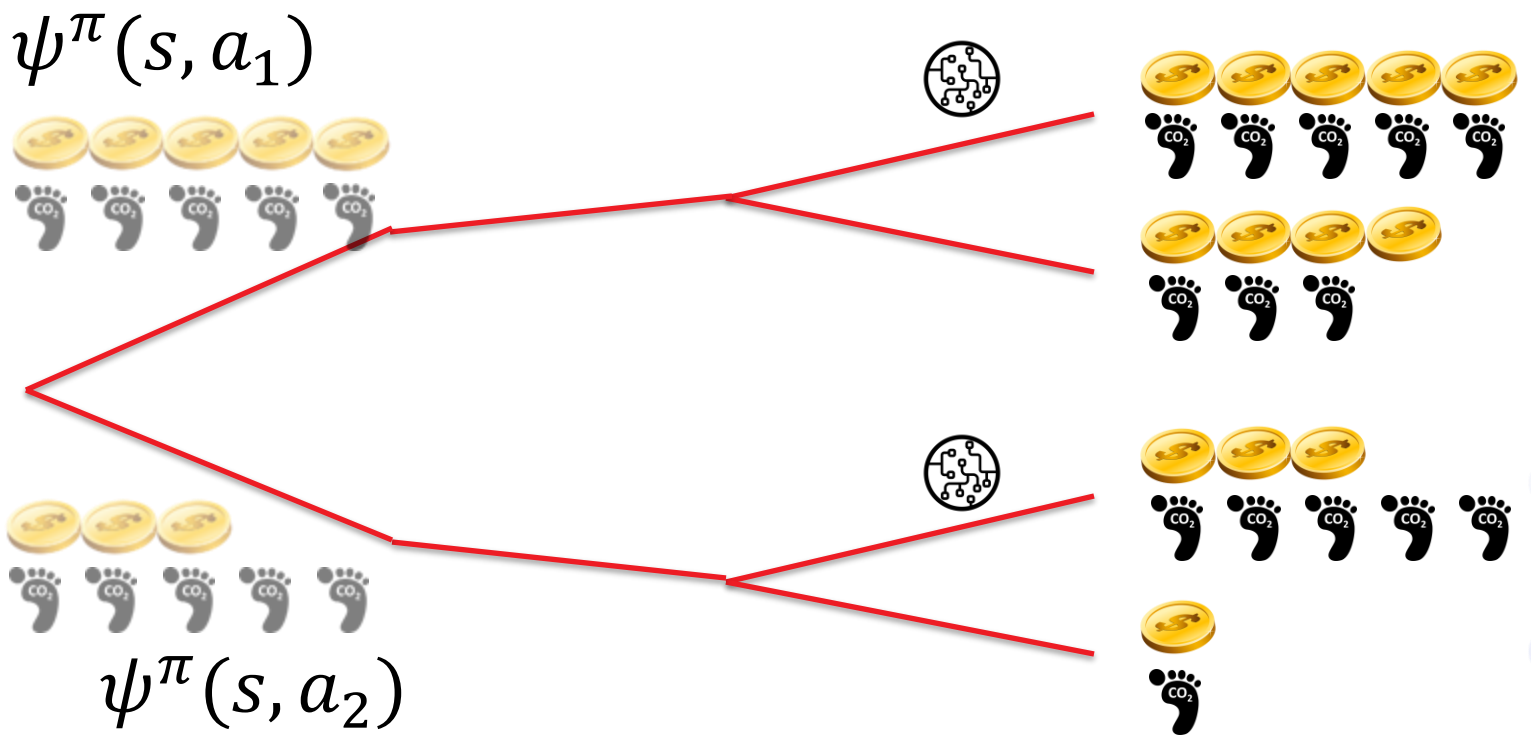

https://publicdomainvectors.org/ (CC0)
Wikimedia / Tommaso.sansone91 (CC0)

- Prediction under AI policy does not recognize that bottom branch is chosen by policies that care about minimizing ⏱

De Peuter, Oliehoek & van Hoof, Performative explanations: Explaining futures under future explanations (submitted)

Explanations interact with decision maker

- We want to explicitly take the (human) decision maker into account when making these predictions
- For each action a_t, explanation policy $\varepsilon(\cdot)$ outputs prediction e_t of outcomes under *decision makers (DM)* future actions
- Explanation policy $\varepsilon(\cdot)$
  - Maximizes return of the DM.
  - Produces accurate explanations

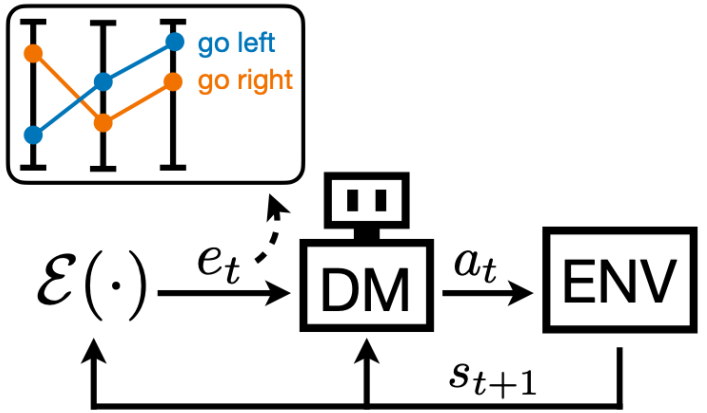

De Peuter, Oliehoek & van Hoof, Performative explanations: Explaining futures under future explanations (submitted)

Why is this difficult?

Although we want to predict the outcomes under DM policy, we don't know their preferences in general

- We assume we do know the distribution of (linear) preferences
- We also assume a known user model: how likely is a user with certain preferences to choose a certain action

De Peuter, Oliehoek & van Hoof, Performative explanations: Explaining futures under future explanations (submitted)

Decision making example

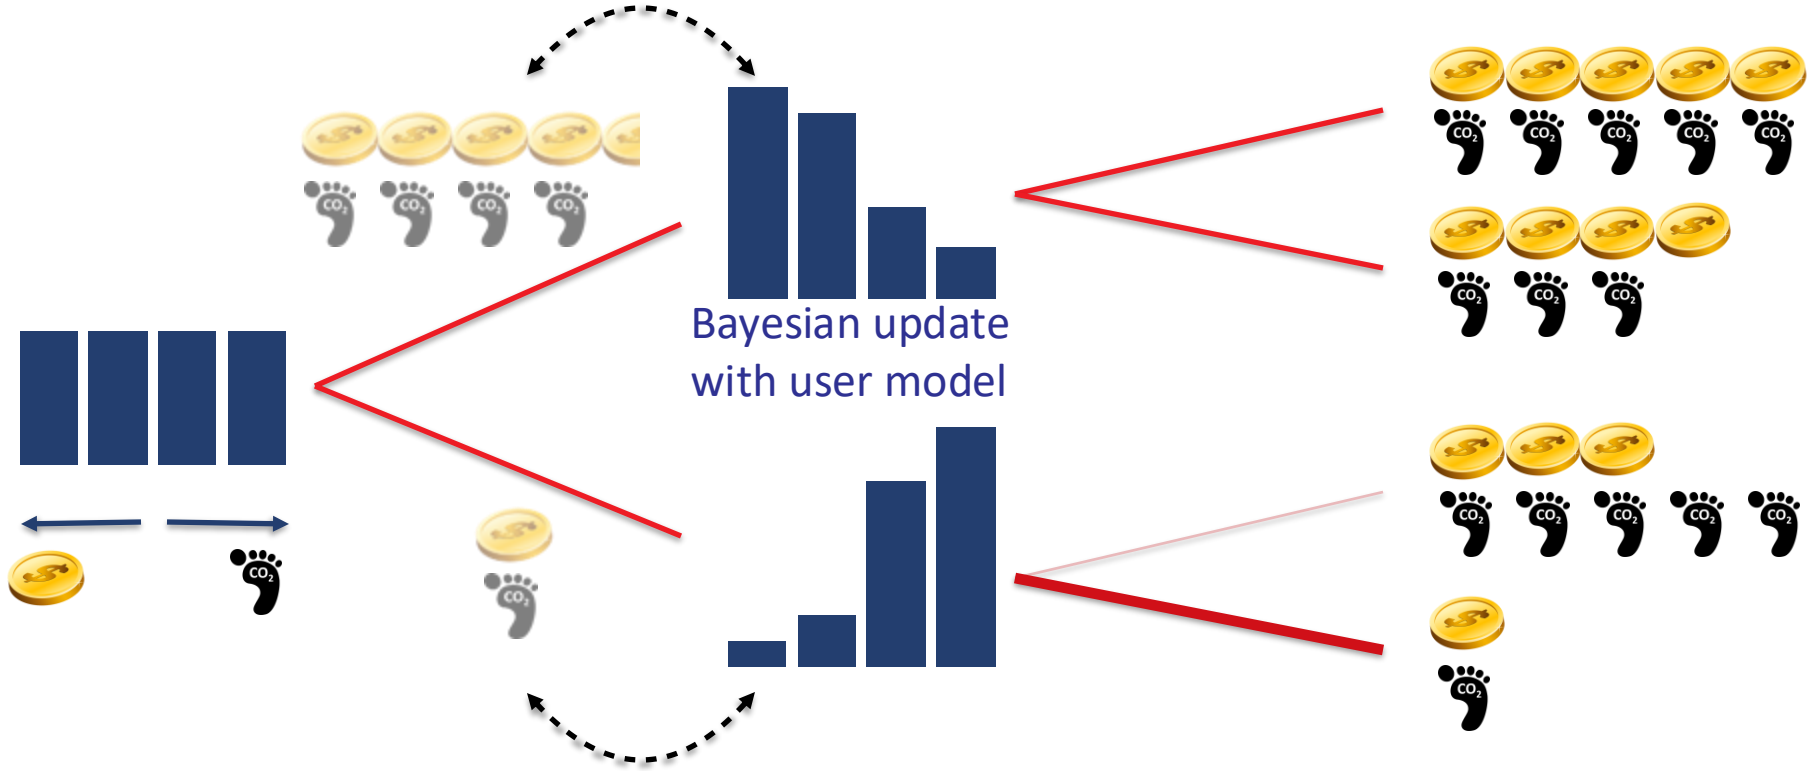

https://publicdomainvectors.org/ (CC0)
Wikimedia / Tommaso.sansone91 (CC0)

- Choices by decision maker reveals something about their preferences
- Predictions and preference distribution influence each other: jointly maintain them

De Peuter, Oliehoek & van Hoof, Performative explanations: Explaining futures under future explanations (submitted)

Experiments

- Baseline: Explanations derived from predictions under neutral policy.
  - Like existing methods, does not take user decisions into account

- Evaluate performance on:
  - **Return** achieved under DM's reward function.
  - Euclidean distance between **realised outcomes** and explanations' predicted outcomes.

De Peuter, Oliehoek & van Hoof, Performative explanations: Explaining futures under future explanations (submitted)

Experiments

- Evaluation on “l2rpn_case14_sandbox" with adversary disconnecting lines
- We can achieve a variety of trade-offs between explanation accuracy and return
- Generally outperform the baseline.

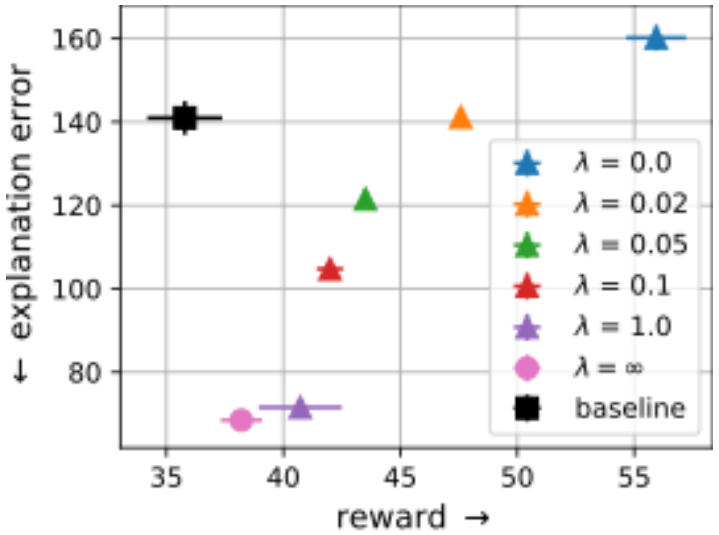

De Peuter, Oliehoek & van Hoof, Performative explanations: Explaining futures under future explanations (submitted)

Takeaways

- Proposition: predictions under an autonomous AI policy are of limited use to (human) decision makers
- Instead, predictions under decision-maker's (likely) actions: requires jointly predicting action outcomes & human preferences
- Our explanations provide various trade-offs between accuracy and return, significantly better on both then baseline explainer

De Peuter, Oliehoek & van Hoof, Performative explanations: Explaining futures under future explanations (submitted)

Conclusions and perspectives

Milad Leyli-abadi (IRT SystemX)

AINETUS Linux Foundation Energy Project

LFENERGY AINETUS

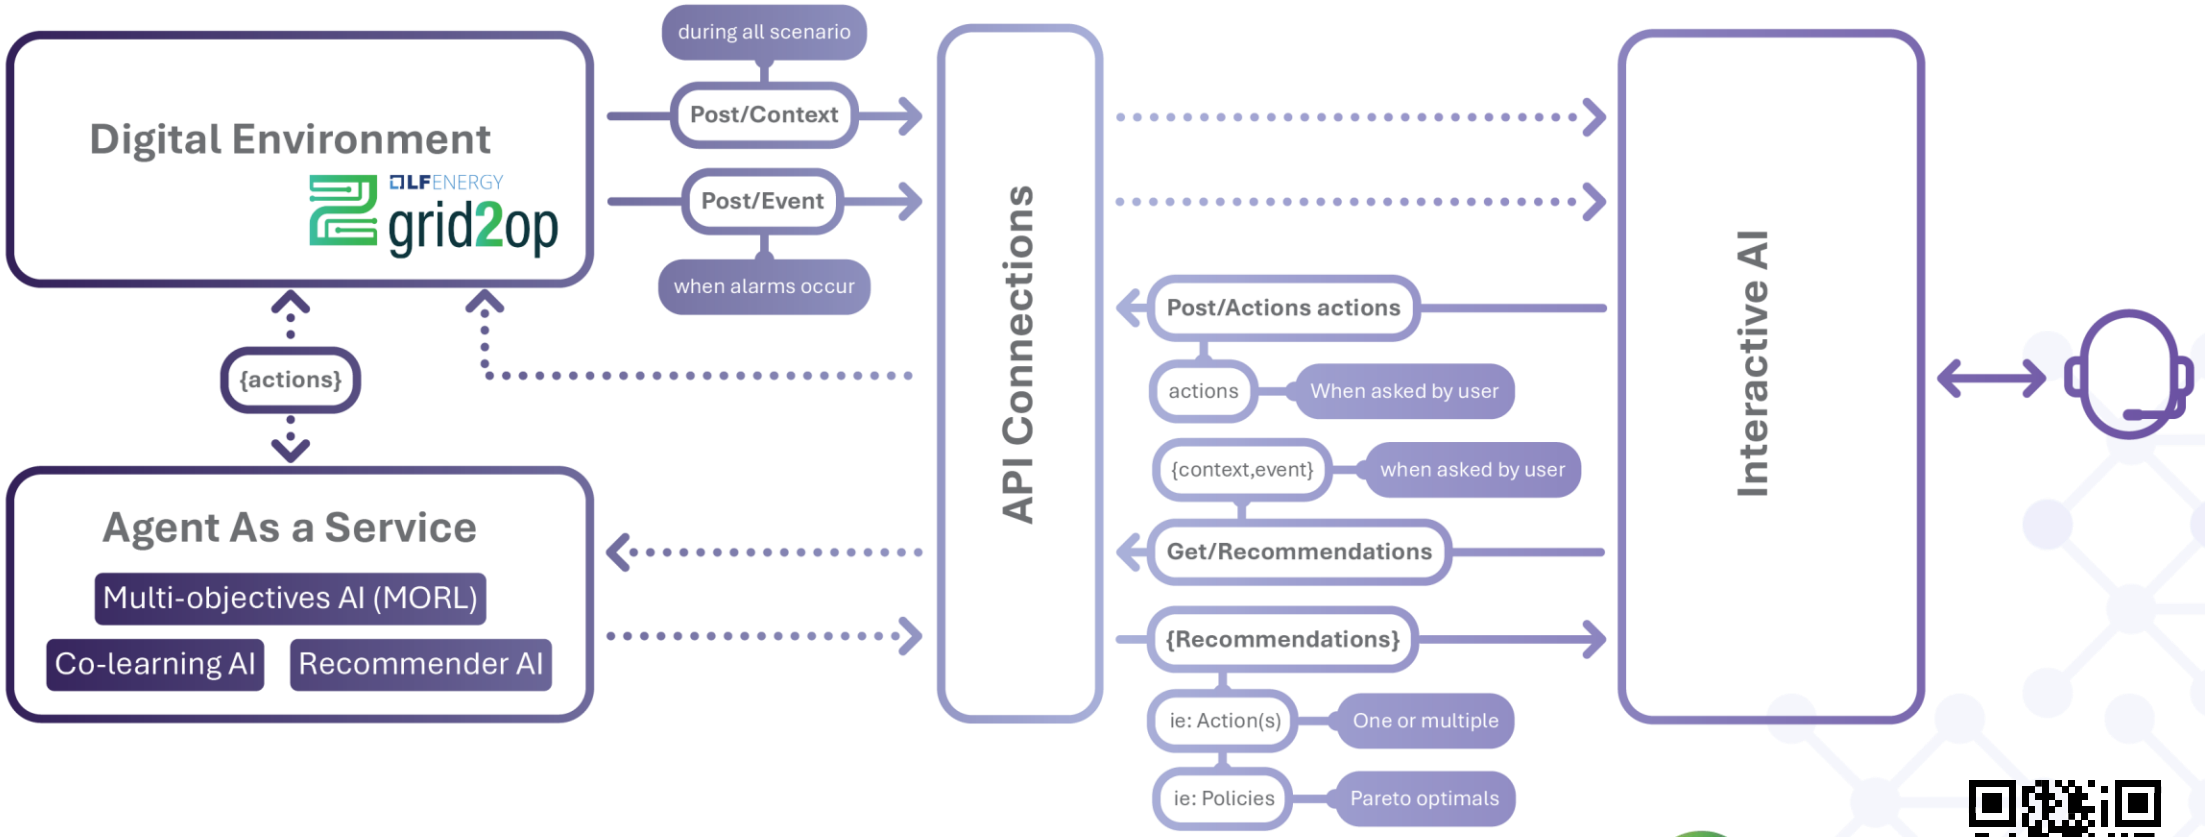

CALL FOR CONTRIBUTORS

Conclusions

- The GNN-based solver replaces costly classical physics simulations
- The suggested AI agents turn a 65K action space into concrete, ranked recommendations
- Failure prediction and interpretable LLM-distilled rules tell the operators when and why to trust, or override, the AI

"An AI assistant that is fast enough, capable enough, and transparent enough to support — not replace — the control-room operator"

Demonstration

Maroua Meddeb (IRT SystemX)

Demonstration

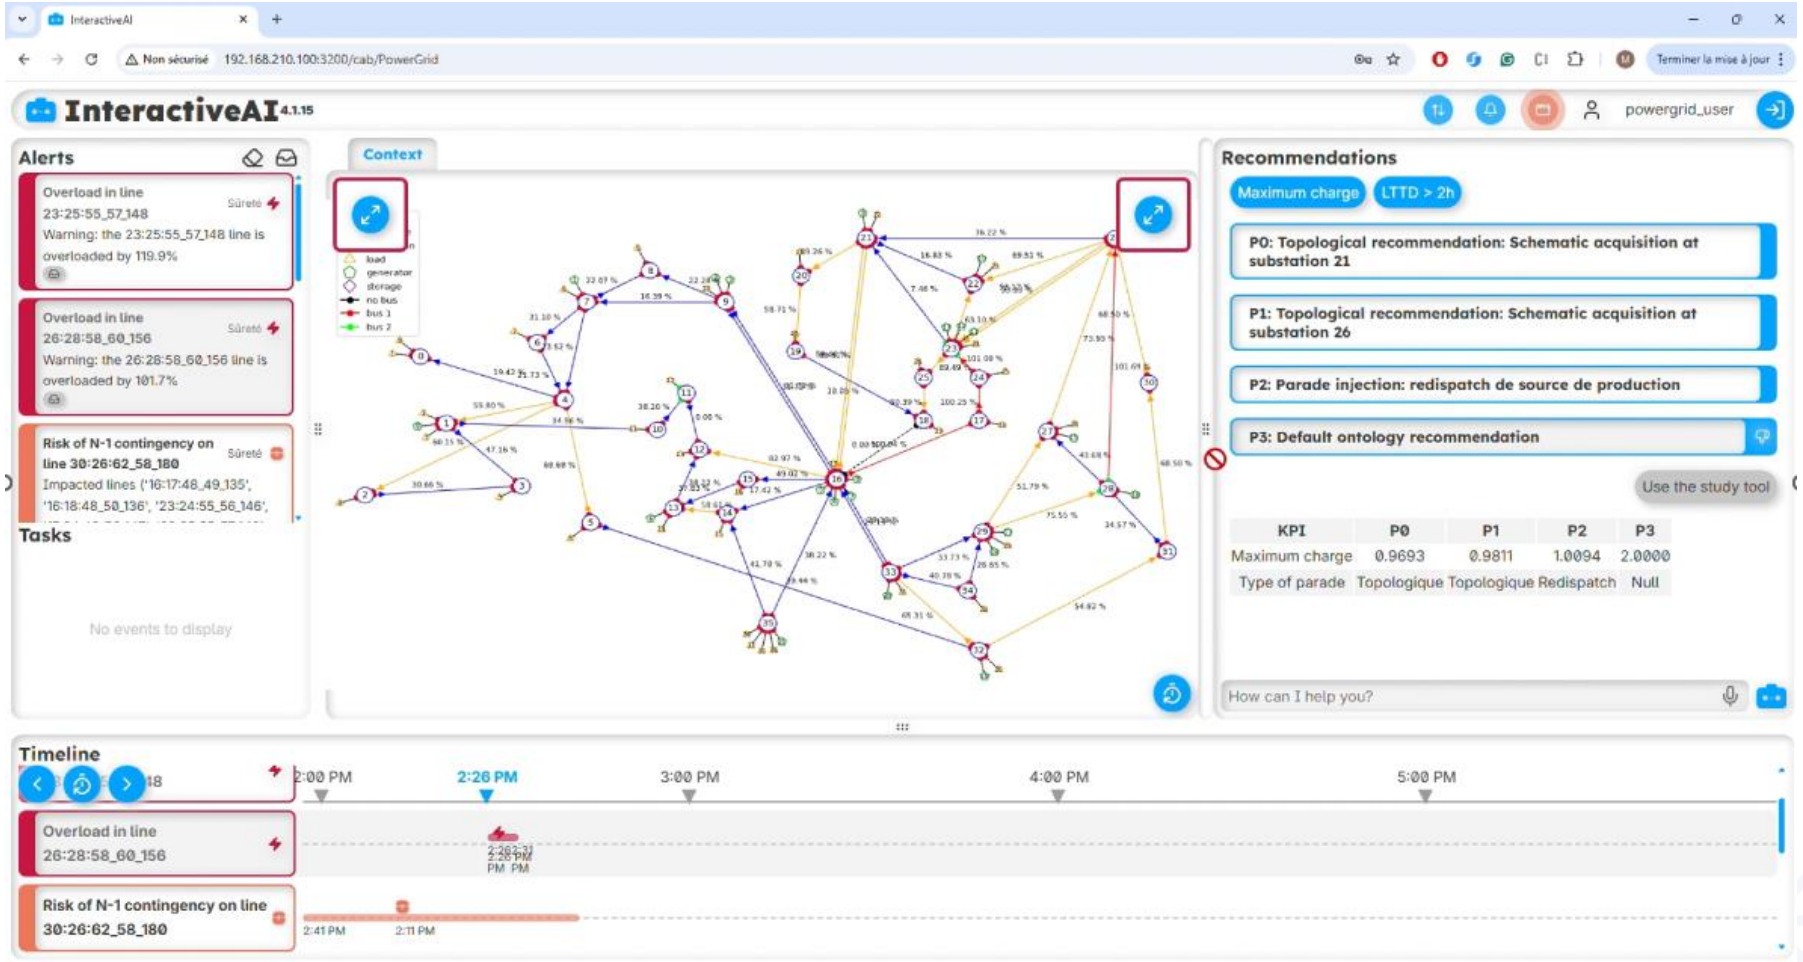

AI4REALNET has received funding from European Union's Horizon Europe Research and Innovation programme under the Grant Agreement No 101119527 and from the Swiss State Secretariat for Education, Research and Innovation

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union. Neither the European Union nor the granting authority can be held responsible for them.