Antoine Marot – Lead AI Scientist @ RTE

Friday 29th of April 2022

Trustworthy, robust and adaptable AI assistant for grid operations

Control room operators today

Quite a crowded work environment !

Need for change in operations given the Energy Transition

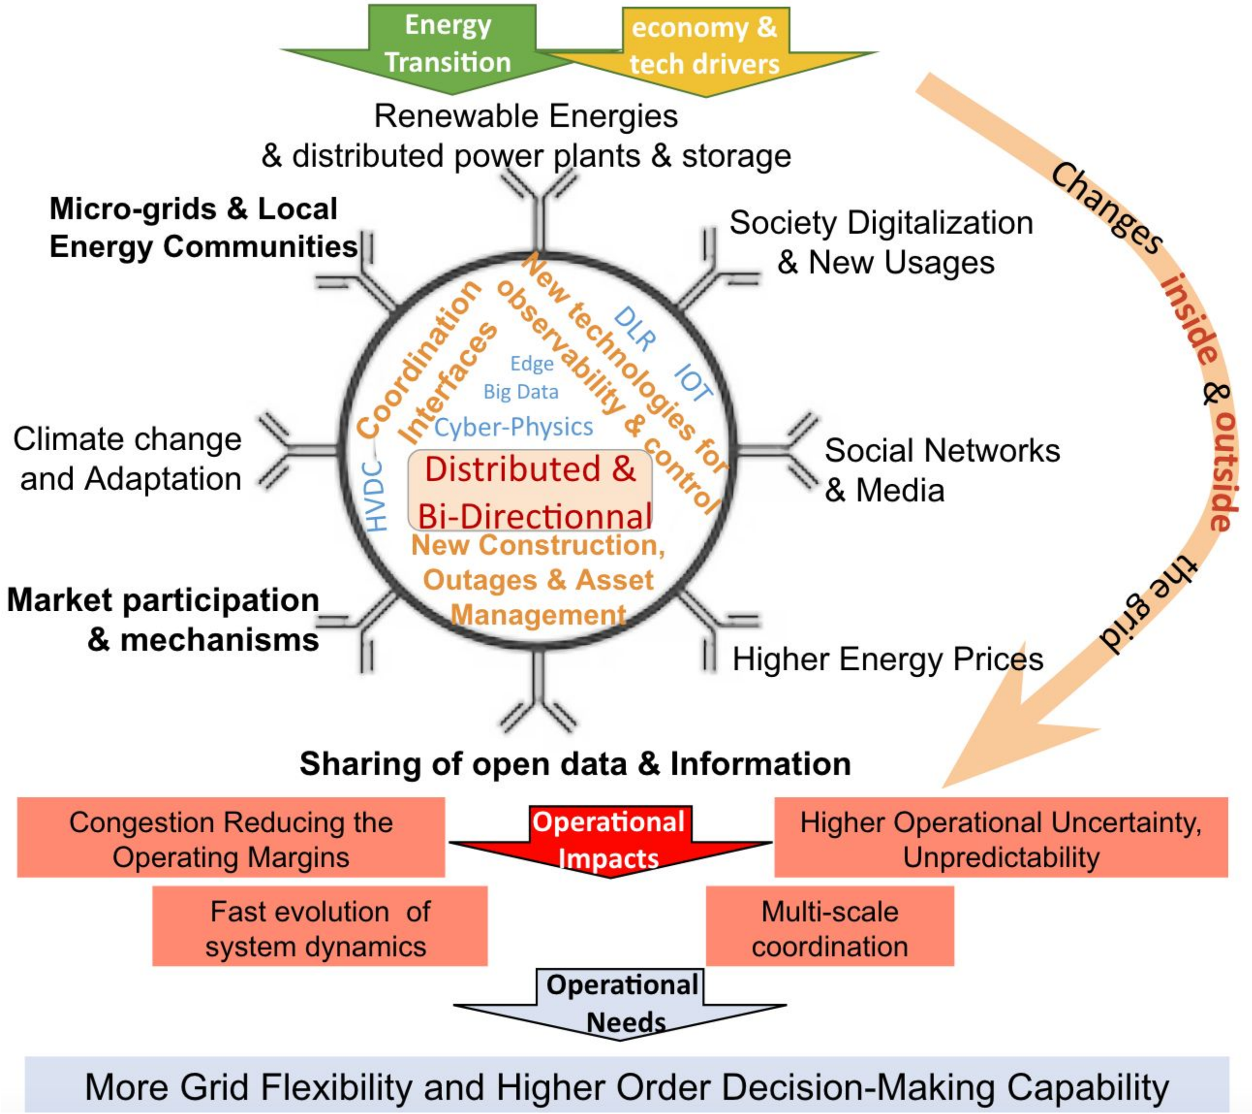

Perspectives for Future Power System Control Centers for The Energy Transition*, Journal of Modern Power System and Clean Energy, 2022, A. Marot, A. Kelly, J. Cremer et al.*

Evolution of control rooms past, present and future

Towards operation planners under a single unified and shared interface

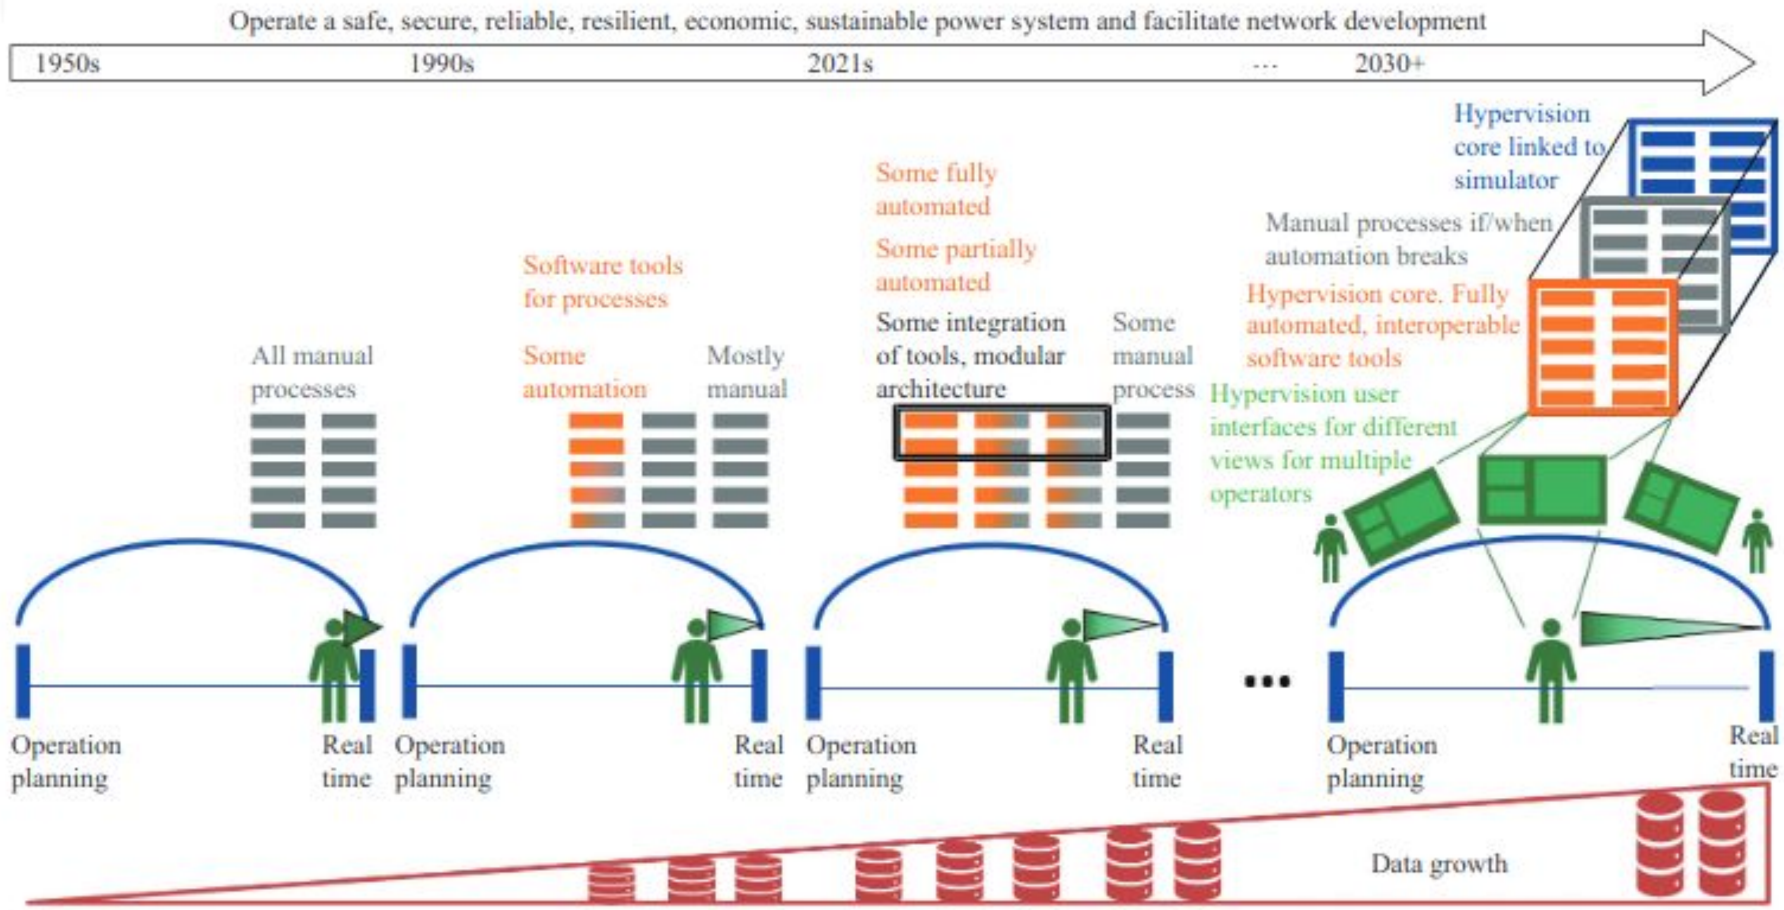

Fig. 2. Evolution of operator's decision-making environment over decades with increasing number of tasks.

Designing Assistants for Operators

Desired Features

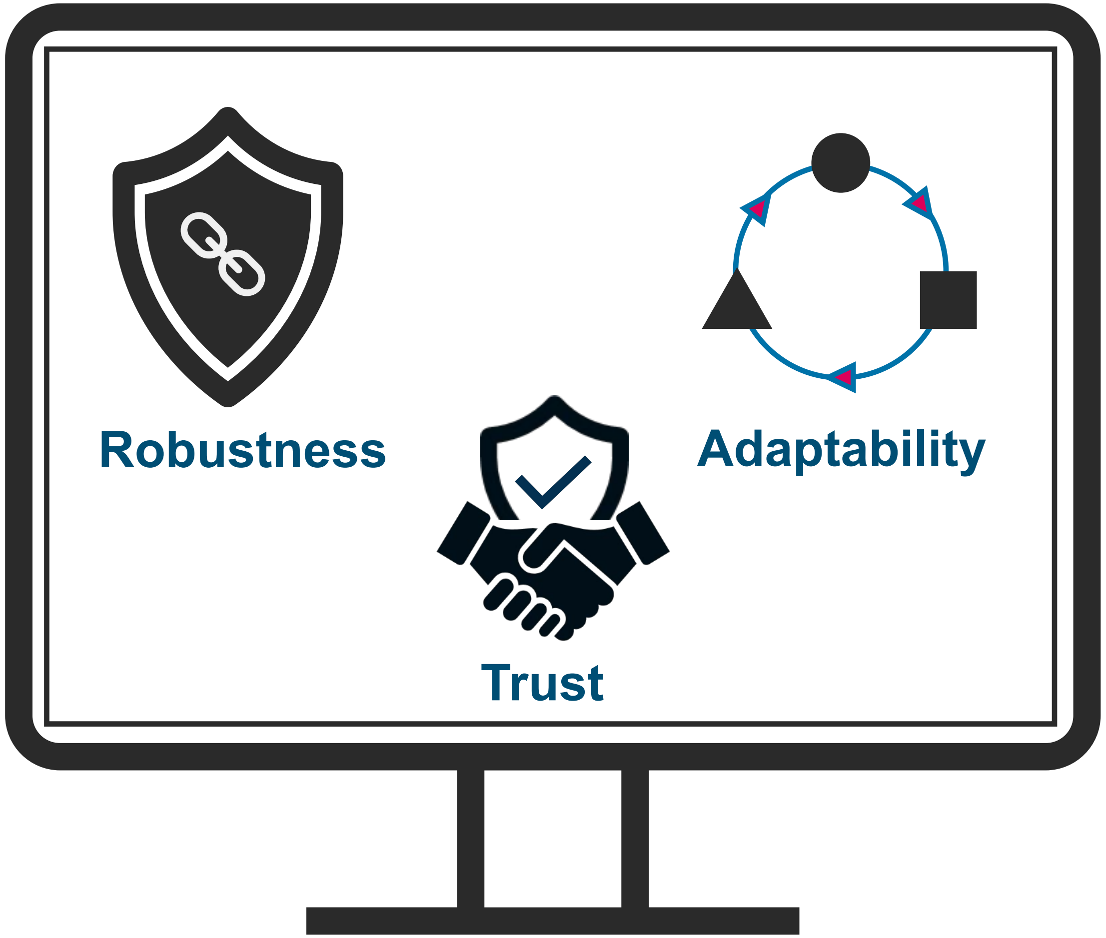

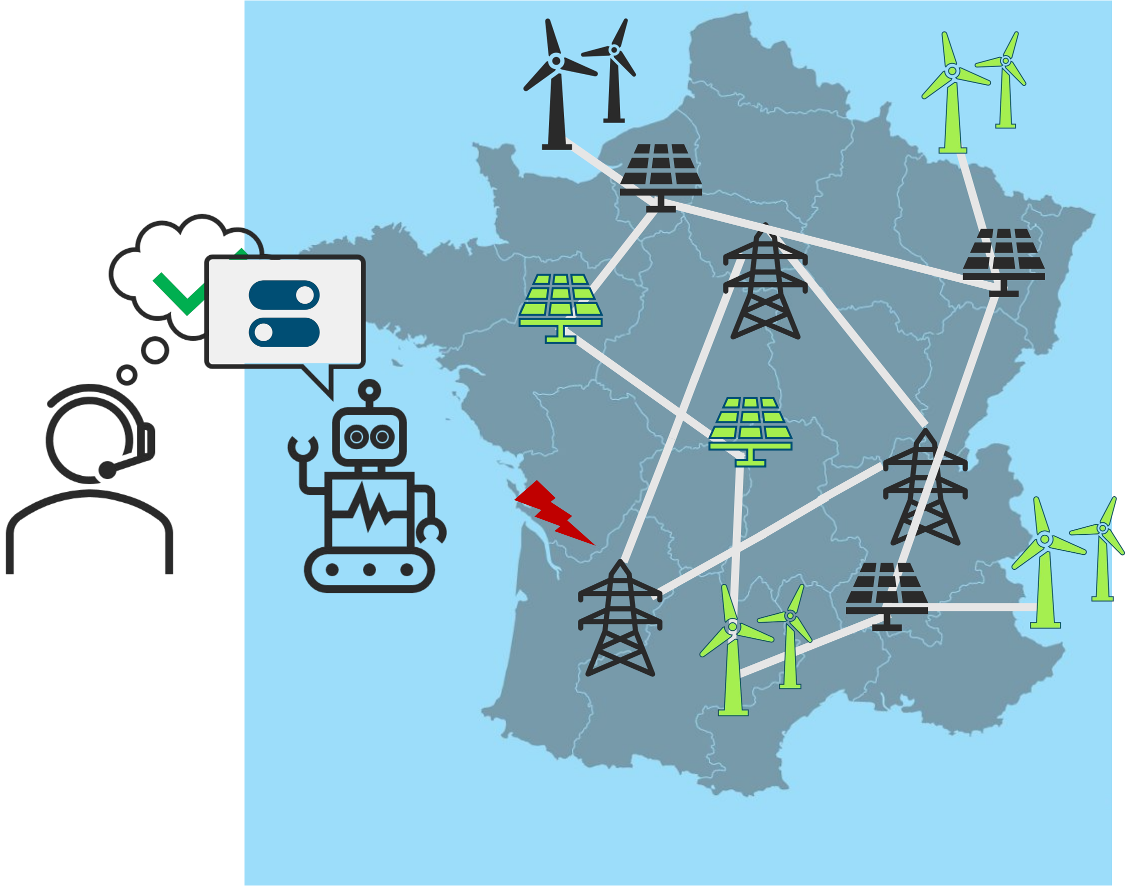

Learning to run a Power Network Challenge: a Retrospective Analysis, PMLR Journal, 2021, A. Marot, A. Kelly et al.

01

Problem design

« Learning To run a power network »
L2RPN Challenge

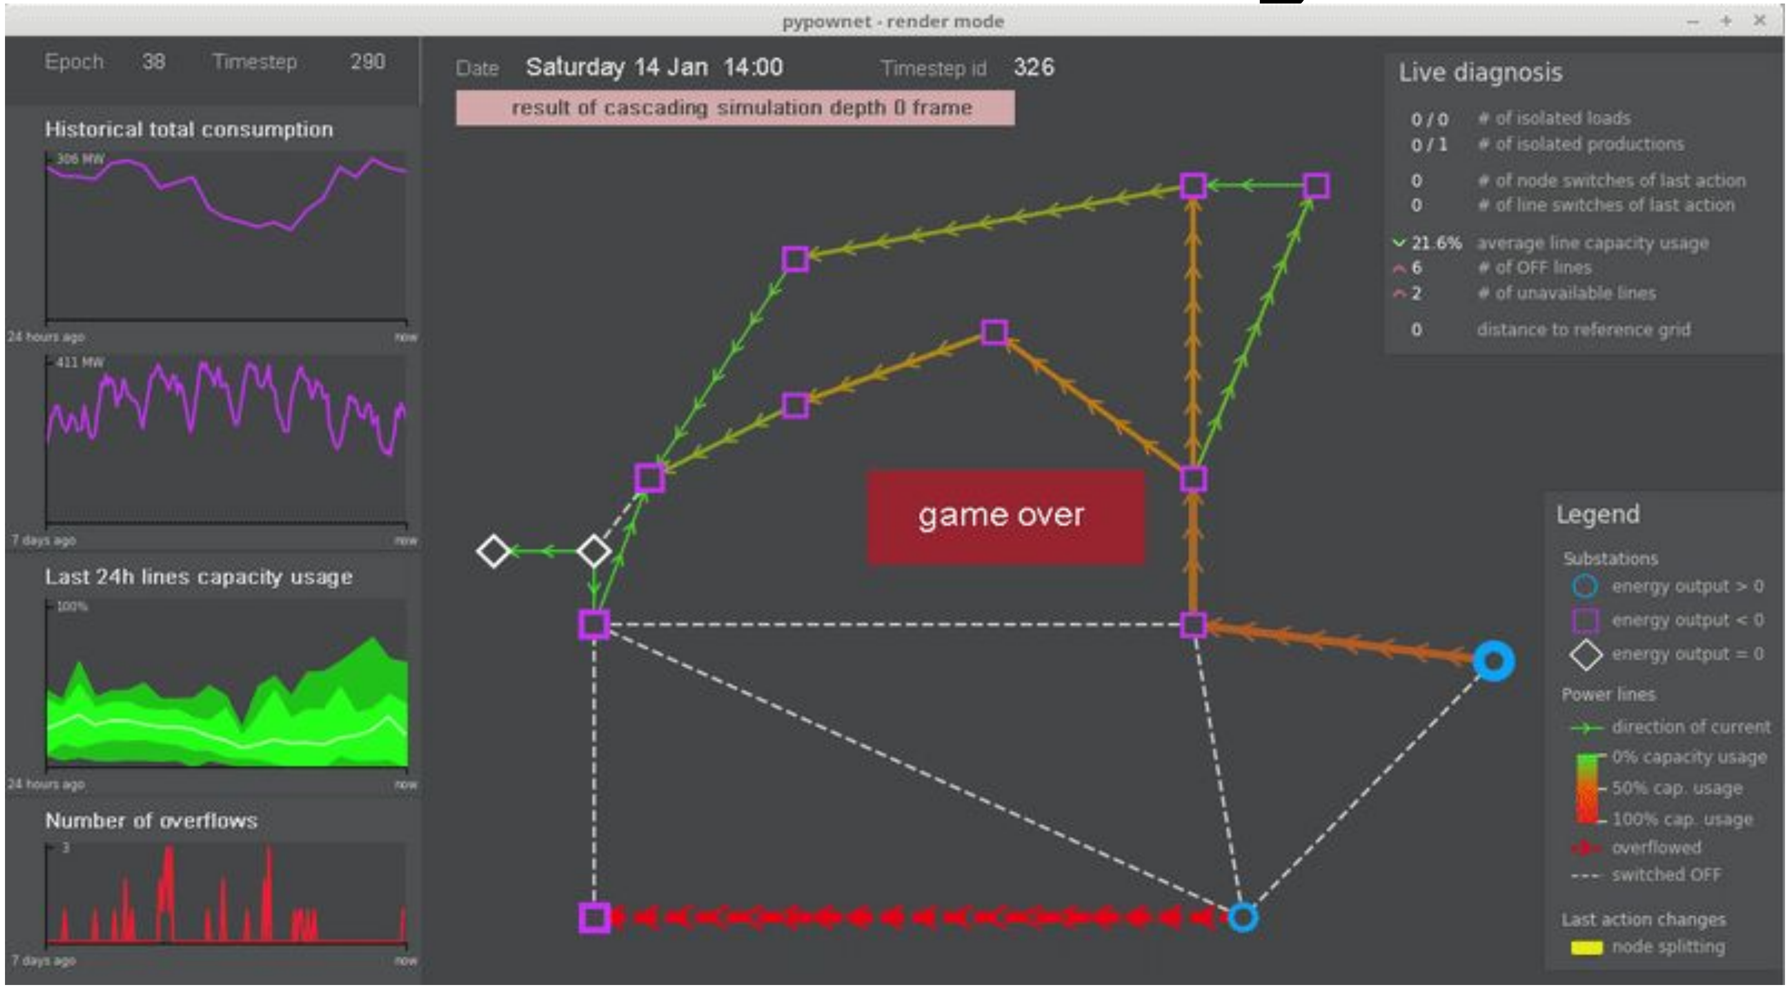

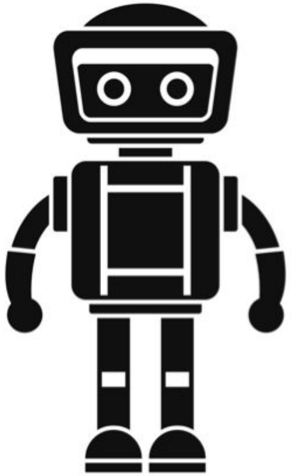

Test the potential of AI to robustly operate a power grid in real-time given operational constraints.

Modeling of real-time operation decision making

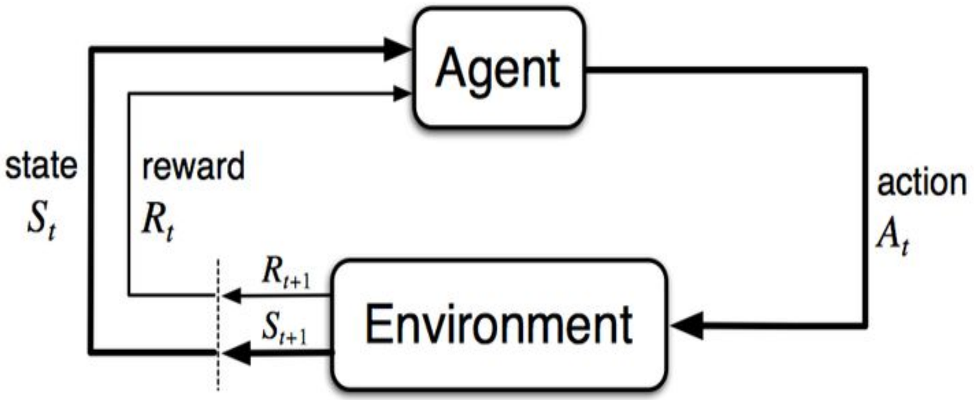

Fig. 1 - *Reinforcement Learning like decision-making framework*

State: flows, productions, consumptions, power grid topology, month, day, hour, etc

Action: connect/disconnect **one** line **or** change electrical configuration at a substation **or** change productions

Reward/cost function: number of overflowed lines, cost of operations, … **Participant design choice**

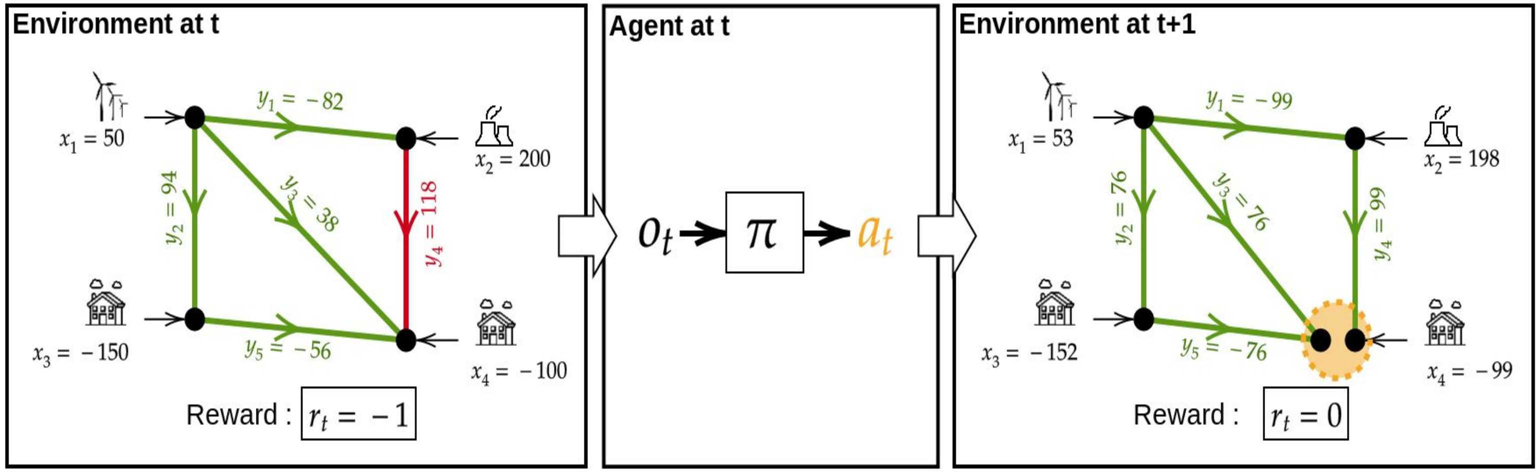

Fig. 2 - *Step-by-step evolution of the RL environment*

Time resolution considered: 5 minutes (human operators work with snapshots every 5 minutes)

Operational Constraints

Line capacities have been calibrated to have **congestions** to be managed appearing few % of the time.

The game should represent additional operational constraints:

1) A limited time to react to a congestion (2 ts)

1) A limited number of simultaneous action (1/ts)

1) A cooldown time before reusing an asset (3 ts)

> There is hence a **budget** associated to the actions you take: they should be picked up carefully !

PS: Problem formulation to be related to stochastic Multi-period AC SCOPF as in [Alizadeh & Capitanescu 2021]

GridAlive - full ecosystem to design & play with

https://github.com/rte-france/gridAlive

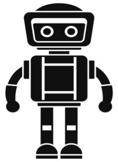

L2RPN-Baseline

Repo to share open-source models & winning approaches

AGENT

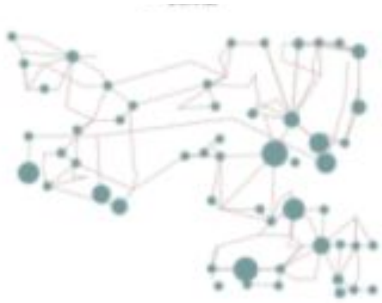

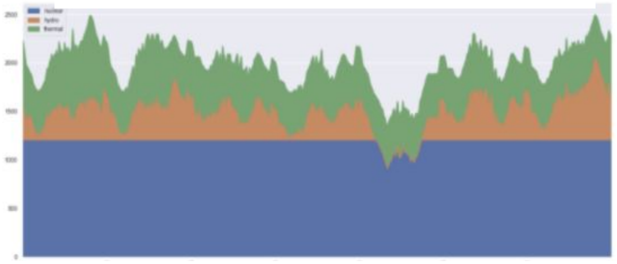

ChroniX2Grid

Power Grid time serie generation

Data

Grid2Op: **testbed platform** to model real-time operations, run & benchmark control algorithms

Power Grid Simulator

Analysis

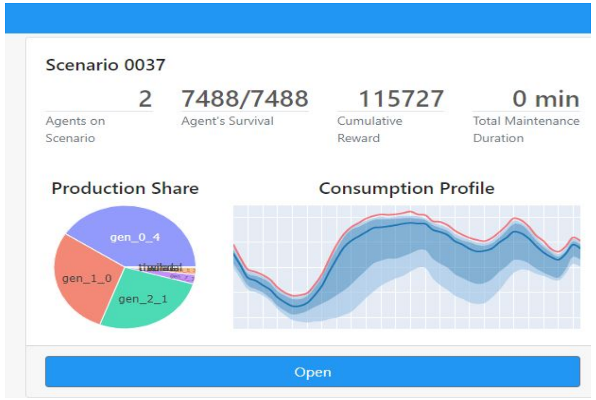

Grid2Viz

Scenario & agents study

02

L2RPN - Adaptability and Robustness

L2RPN serie competitions

IJCNN Feasibility challenges

NeurIPS Sustainable World challenge

2019

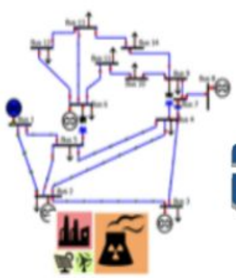

Small Grid, no events,
Winter month,
only topology

Spring 2020

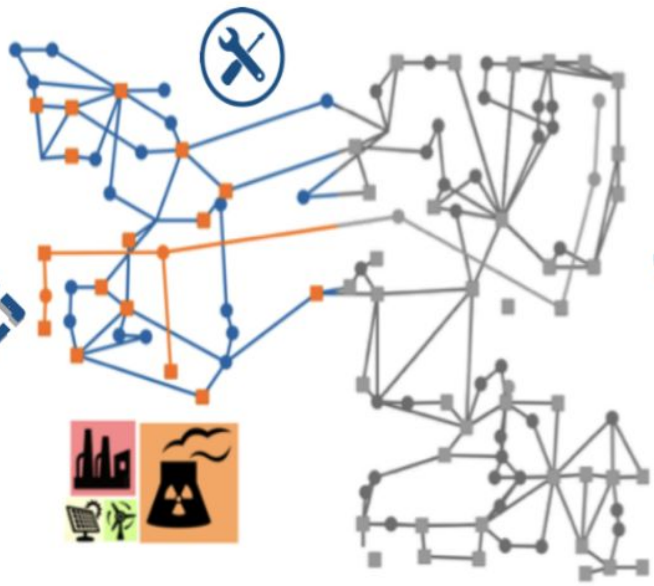

***Medium** Grid, **maintenance**,*
all year long, only topology

Robustness

Summer 2020

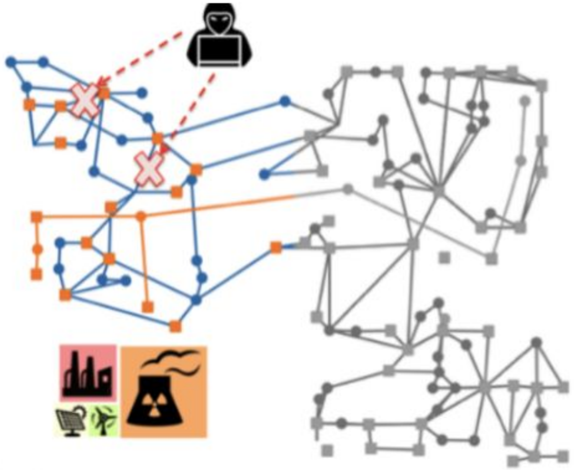

*Medium Grid**, adverserial attacks**,*
*Topology & **redispatching** actions*

Adaptability

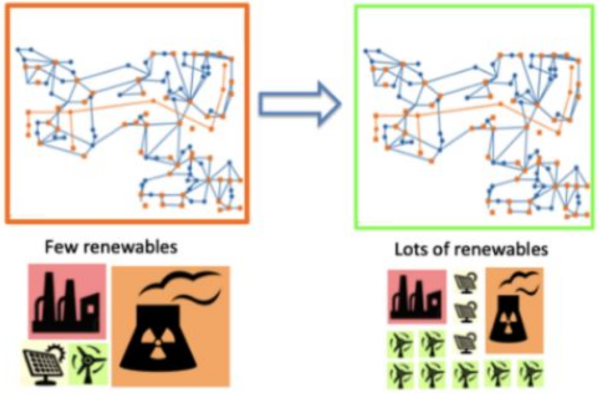

***Large** Grid, **2 environments**,*
*Topology & **redispatching** actions*

Participate from:
https://l2rpn.chalearn.org/

2 sub problems

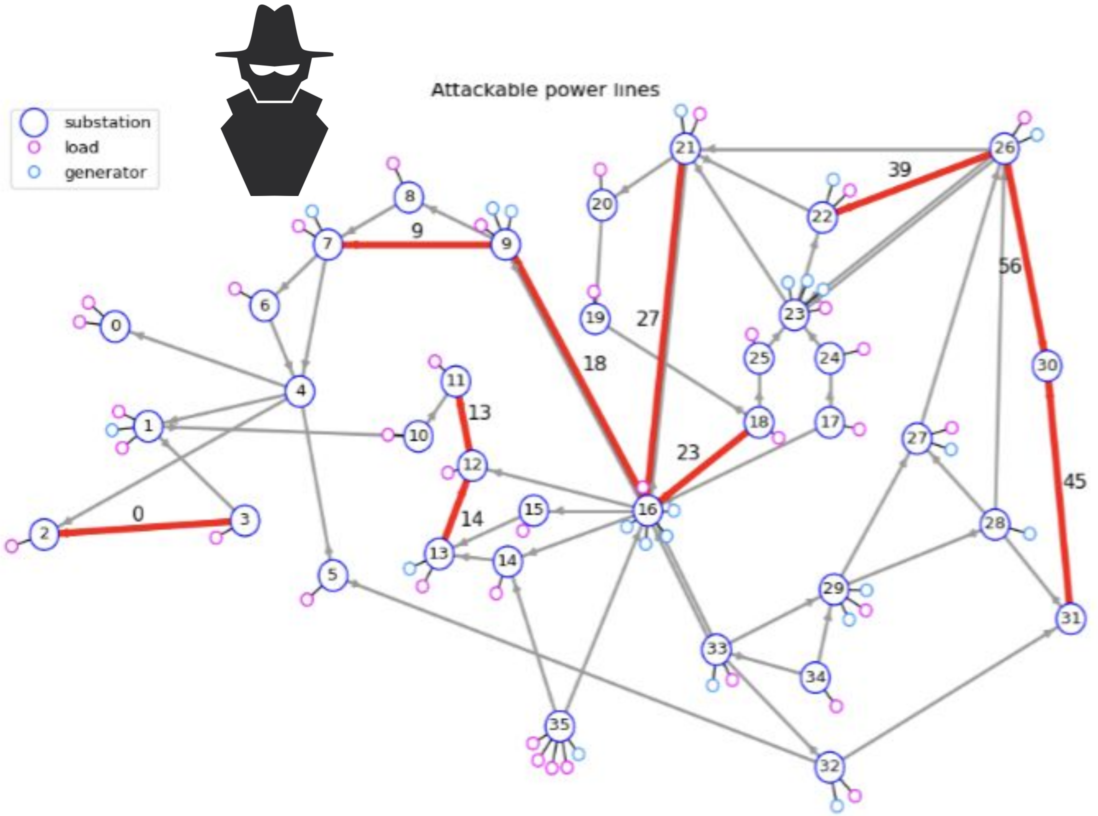

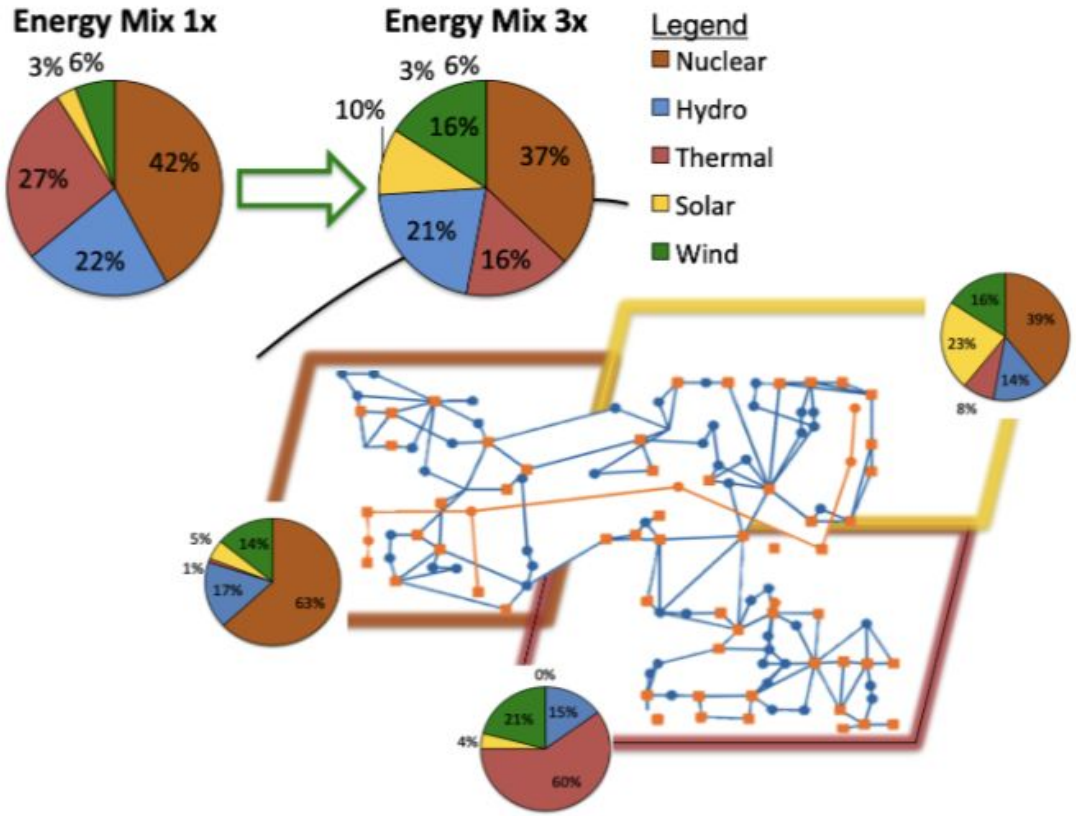

Robustness Track: 10 attackable line by an adverserial oponent (testing for N-1 robustness)

PowerTech 2021 paper: *Adversarial training for continuous robustness control in power*

Adaptability Track: 5 energy mix (1x - 1.5x - 2x - 2.5x - 3x) for training
Tesing over other mix (1.7x - 2.2x - 2.7x - 3.2x)

Agent evaluation – score function

Control the power flows to **optimize the cost of operations** on the power grid while being **robust** to blackouts.

We can hence define our overall operational cost $\boldsymbol{c}_{\text{operations}}(\boldsymbol{t})$:

$$c_{\text{operations}}(t) = c_{\text{loss}}(t) + c_{\text{redispatching}}(t) \quad \& \quad c_{\text{blackout}}(t) = \text{Load}(t) * \beta * p(t), \ \beta \geqslant 1$$

Now we can define our cost $\boldsymbol{c}$ for an episode:

$$c(\boldsymbol{e}) = \sum_{t=1}^{t_{\text{end}}} c_{\text{operations}}(t) + \sum_{t=t_{\text{end}}}^{T_e} c_{\text{blackout}}(t)$$

Under N episodes, the final score to minimize is:

$$Score = \sum_{i=1}^{N} c(e_i)$$

test scenarios = 24 weekly scenarios at 5 - min resolution
Rescaled between [-100,100] for better interpretability

Final Leaderboards

L2RPN NEURIPS 2020 - Robustness Track

Organized by BDonnot

Train controllers to conduct a power grid for as long as possible while avoiding incidents.

Jul 08, 2020-Nov 02, 2020

197 participants

USD $15,000 reward

L2RPN NEURIPS 2020 - Adaptability Track

Organized by BDonnot

Train controllers to conduct a power grid with various energy mixes.

Jul 08, 2020-Nov 02, 2020

144 participants

USD $15,000 reward

Score					
#	User	Entries	Date of Last Entry	score ▲	Computation time
1	rl_agnet	44	10/31/20	59.26 (1)	483.22 (41)
2	binbinchen	55	10/30/20	46.89 (2)	437.85 (40)
3	lujixiang	116	10/28/20	44.62 (3)	778.02 (44)

Score					
#	User	Entries	Date of Last Entry	Score ▲	Computation time
1	rl_agnet	12	11/01/20	25.53 (1)	474.35 (27)
2	KunjieTang	10	11/02/20	24.66 (2)	414.09 (26)
3	lujixiang	17	11/02/20	24.63 (3)	518.29 (29)

PS: multiply by 2 the scores to compare on the same scale with Robustness

- Top Ranking Teams are the same on both tracks. They all come from China (Baidu, Huawei, State Grid of China) !
    - Ranking is stable accross validation and test scenarios
- In Robustnees track, the wining team won with a good margin.
- Quite tight however on Adaptability Track

Video avalaible at
https://l2rpn.chalearn.org/competitions

Quick overview of best submission

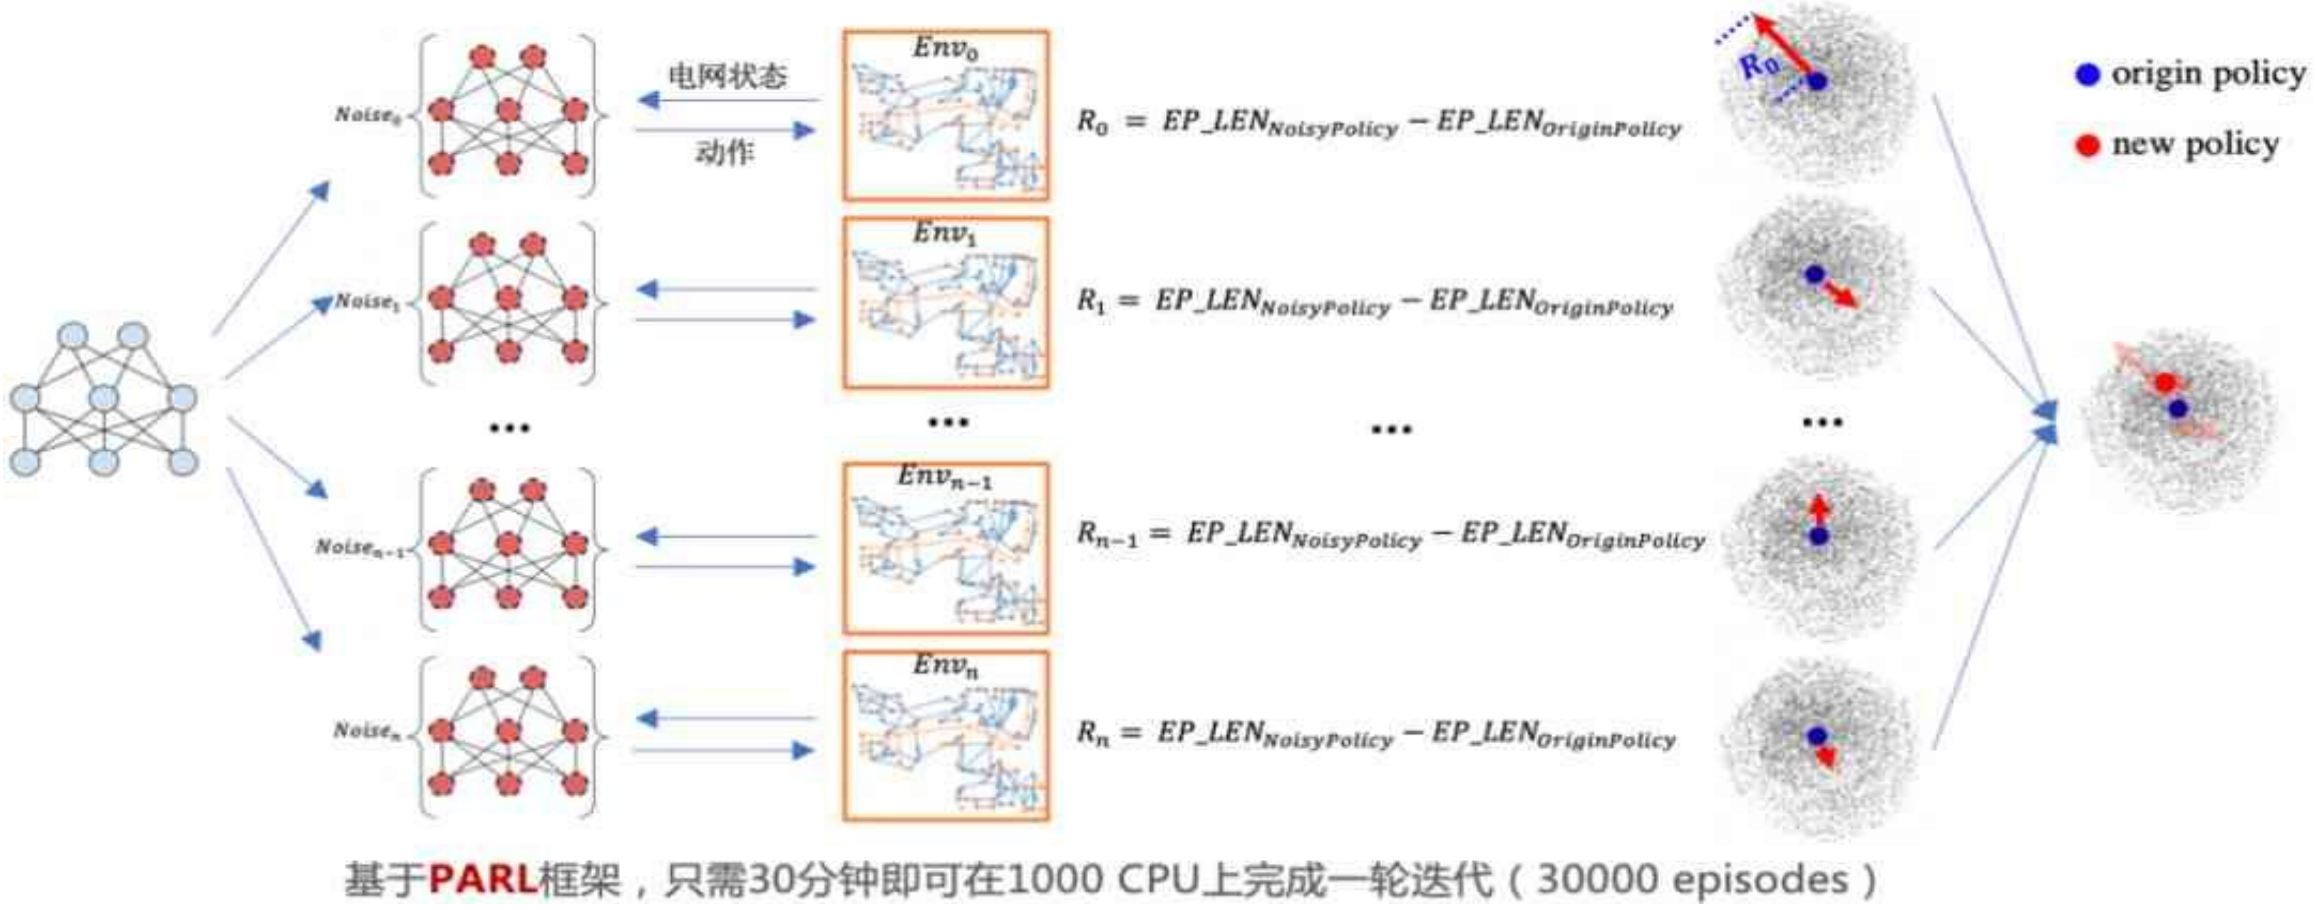

Learning with DeepRL combined with evolutionnary approach **a**:

- Perturbations ("**mutations**") of current model parameters (generation n)
- Reward evaluation in // of survival time (simpler possible reward)
- **Best mutation selection -** new model generation n+1

03

Trustworthyness

Behavior Analysis

Hard and failing january scenario

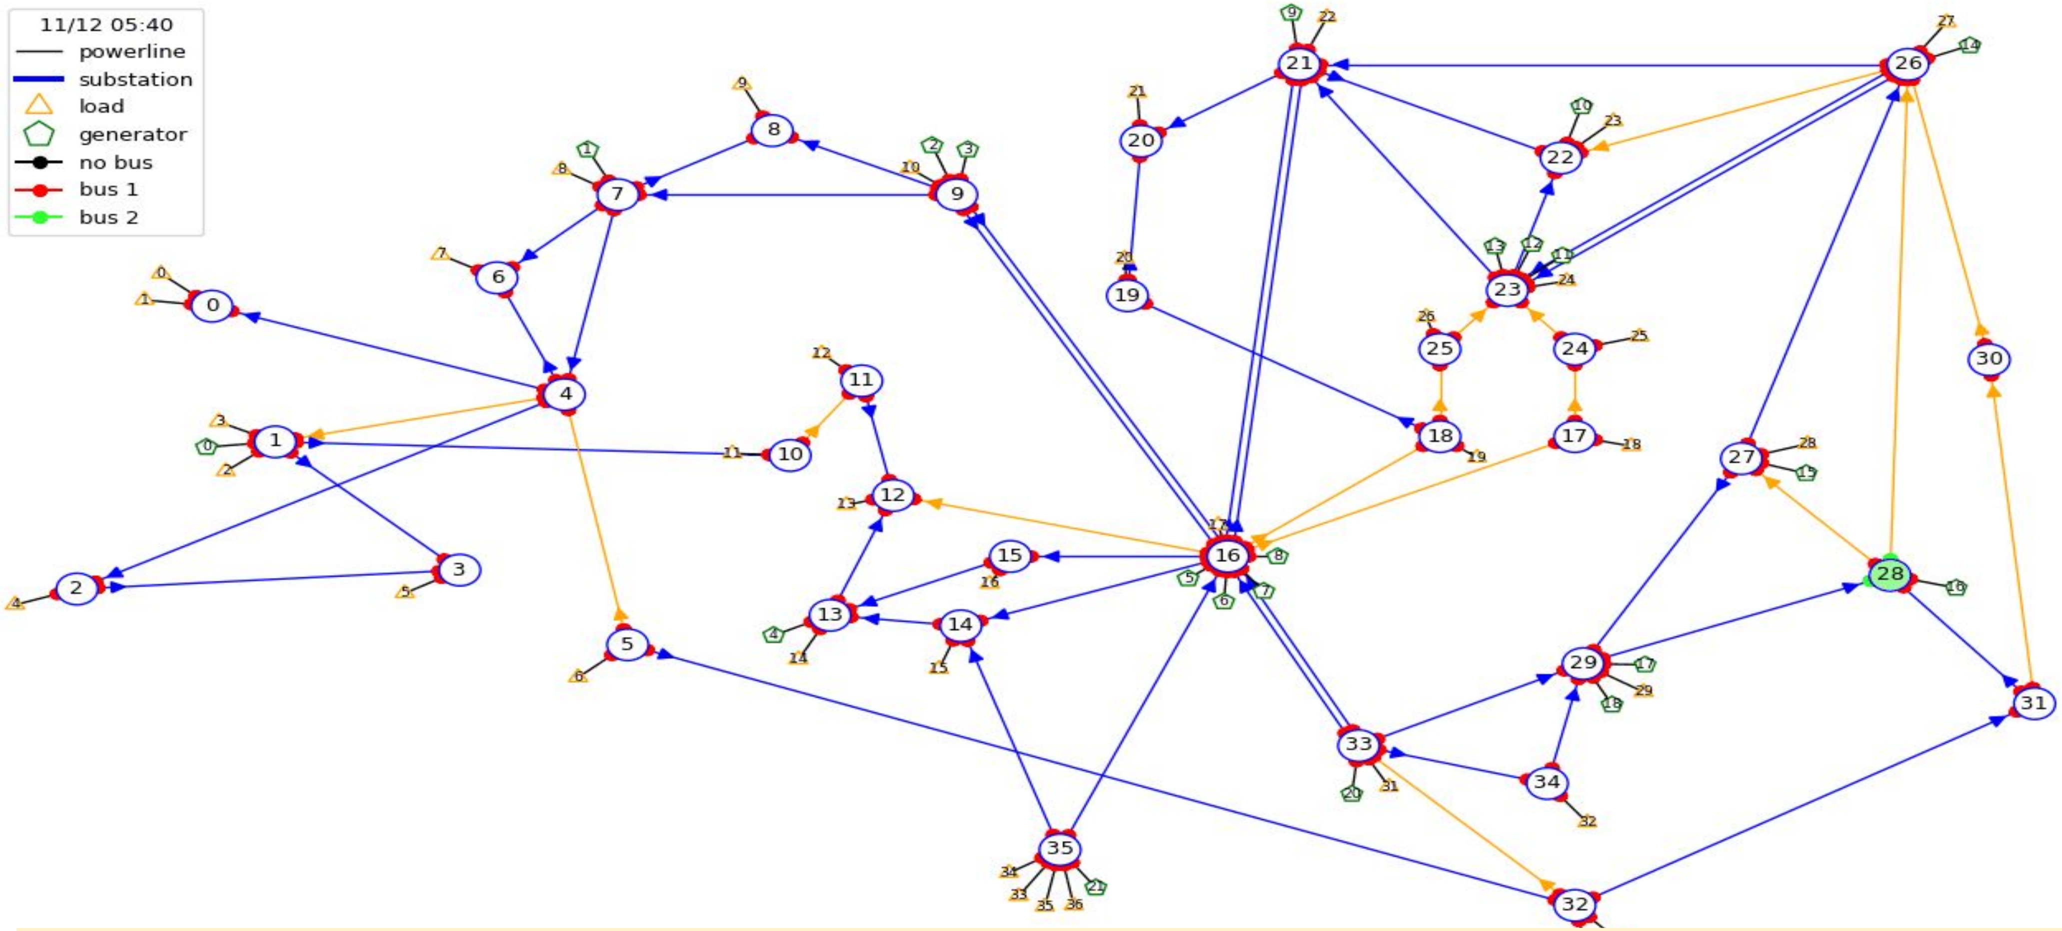

Check out our “Advanced Behavior Analysis of best AI agents from L2RPN NeurIPS competition” video: https://www.youtube.com/watch?v=xlqS-CzvMwk

Autonomous-only agent not Trusted

From previous competitions

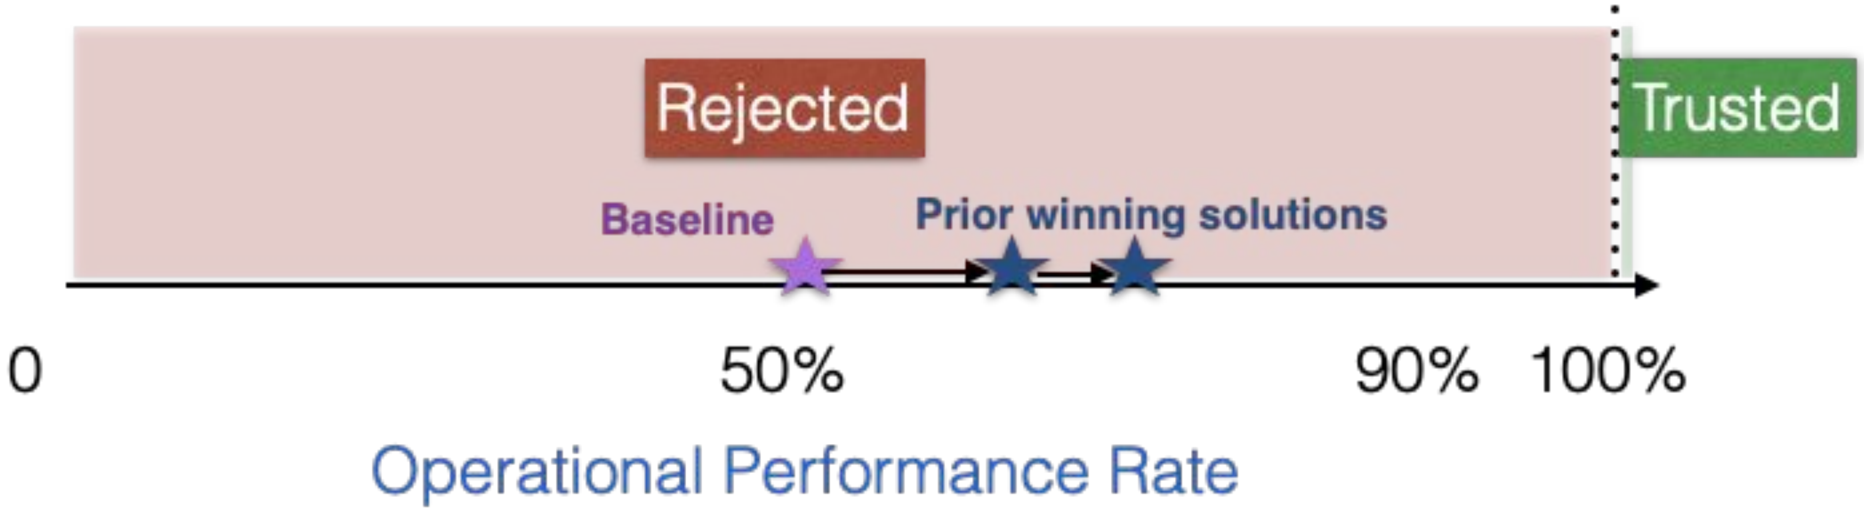

- Prior competitions have improved operational performance, succeeding at more scenarios
- But still far from a near perfect autonomous-only agent

➔ Such agent will not be deployed autonomously, neither they will be trusted by operators

We aim rather at an **assistant** that can communicate with operators and give them control

Trust when predicting failure

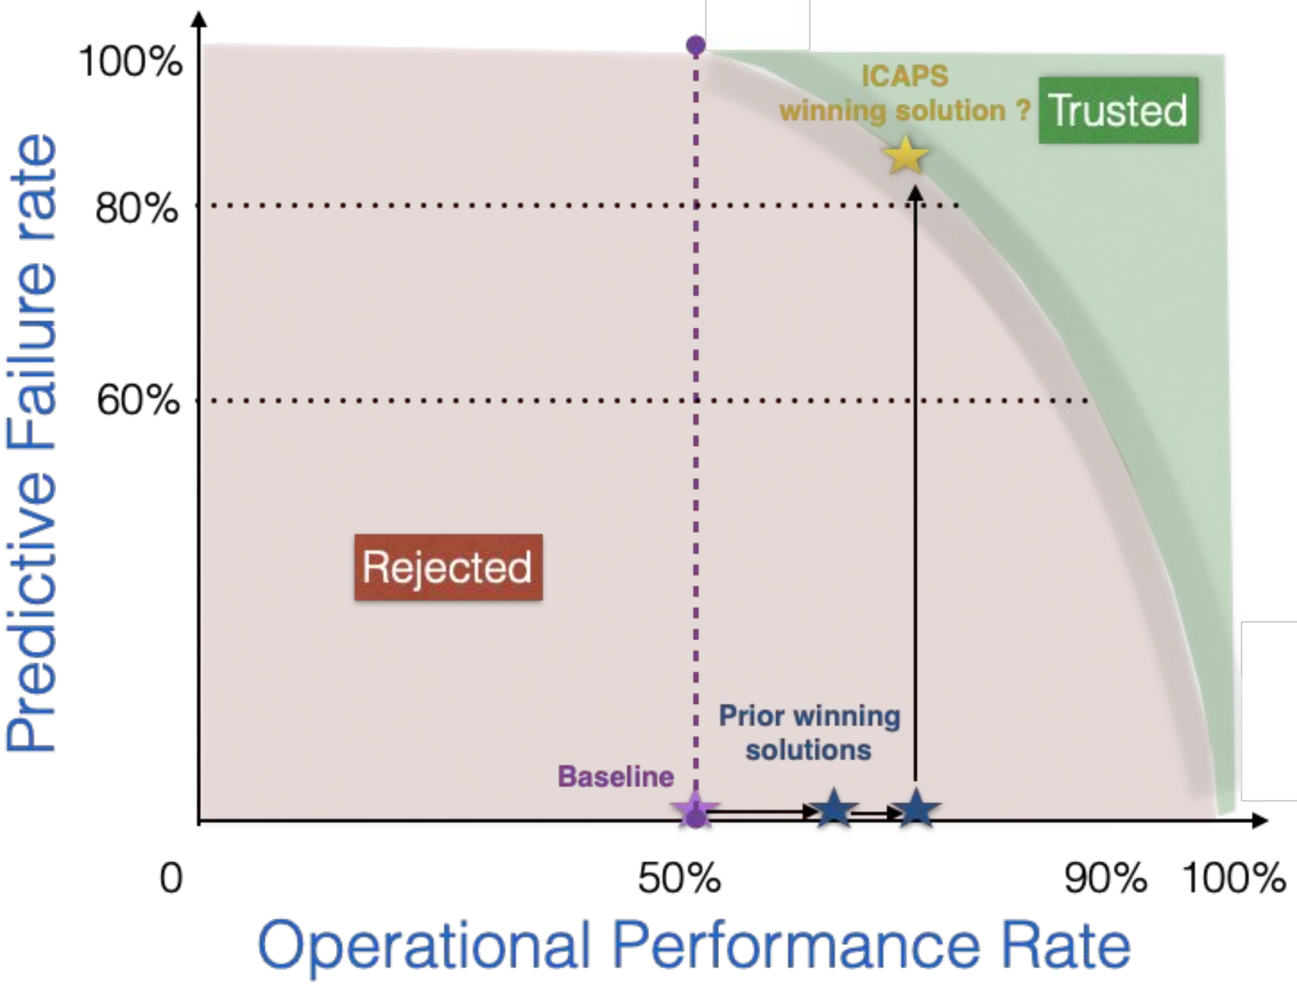

One can build trust for a given level of good enough performance by alerting when risk of failure

ICAPS 2021 Competition: L2RPN with Trust

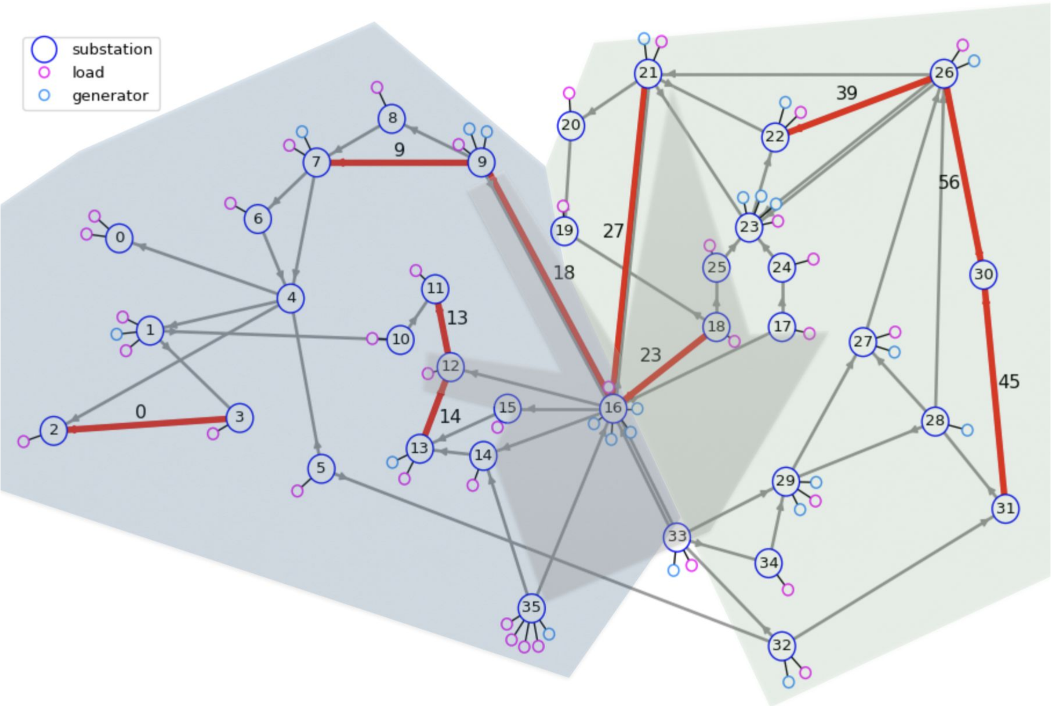

3 possible areas that an agent can alert upon (& 10 attackable lines)

In this competition, participants, while operating the grid under a higher penetration of renewable energy, were asked in addition **to design trustworthy agents that are able to communicate when they are in trouble**, especially when they might fail.

Score = 0.3 ScoreAlarm + 0.7 ScoreOperationCost

[-200,100] [-100,100]

Alarms sent over scenario jan28_1

6 substations changed

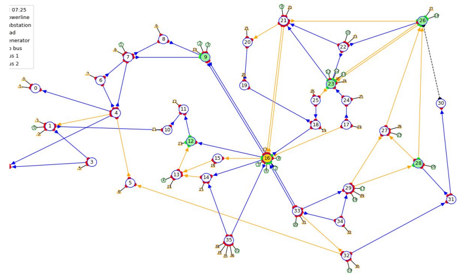

reference topology

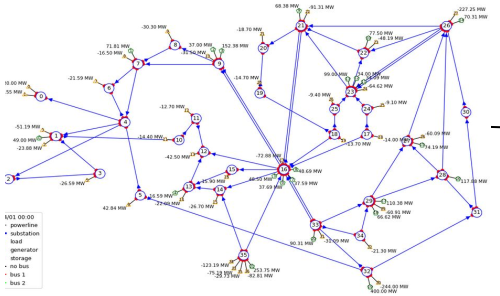

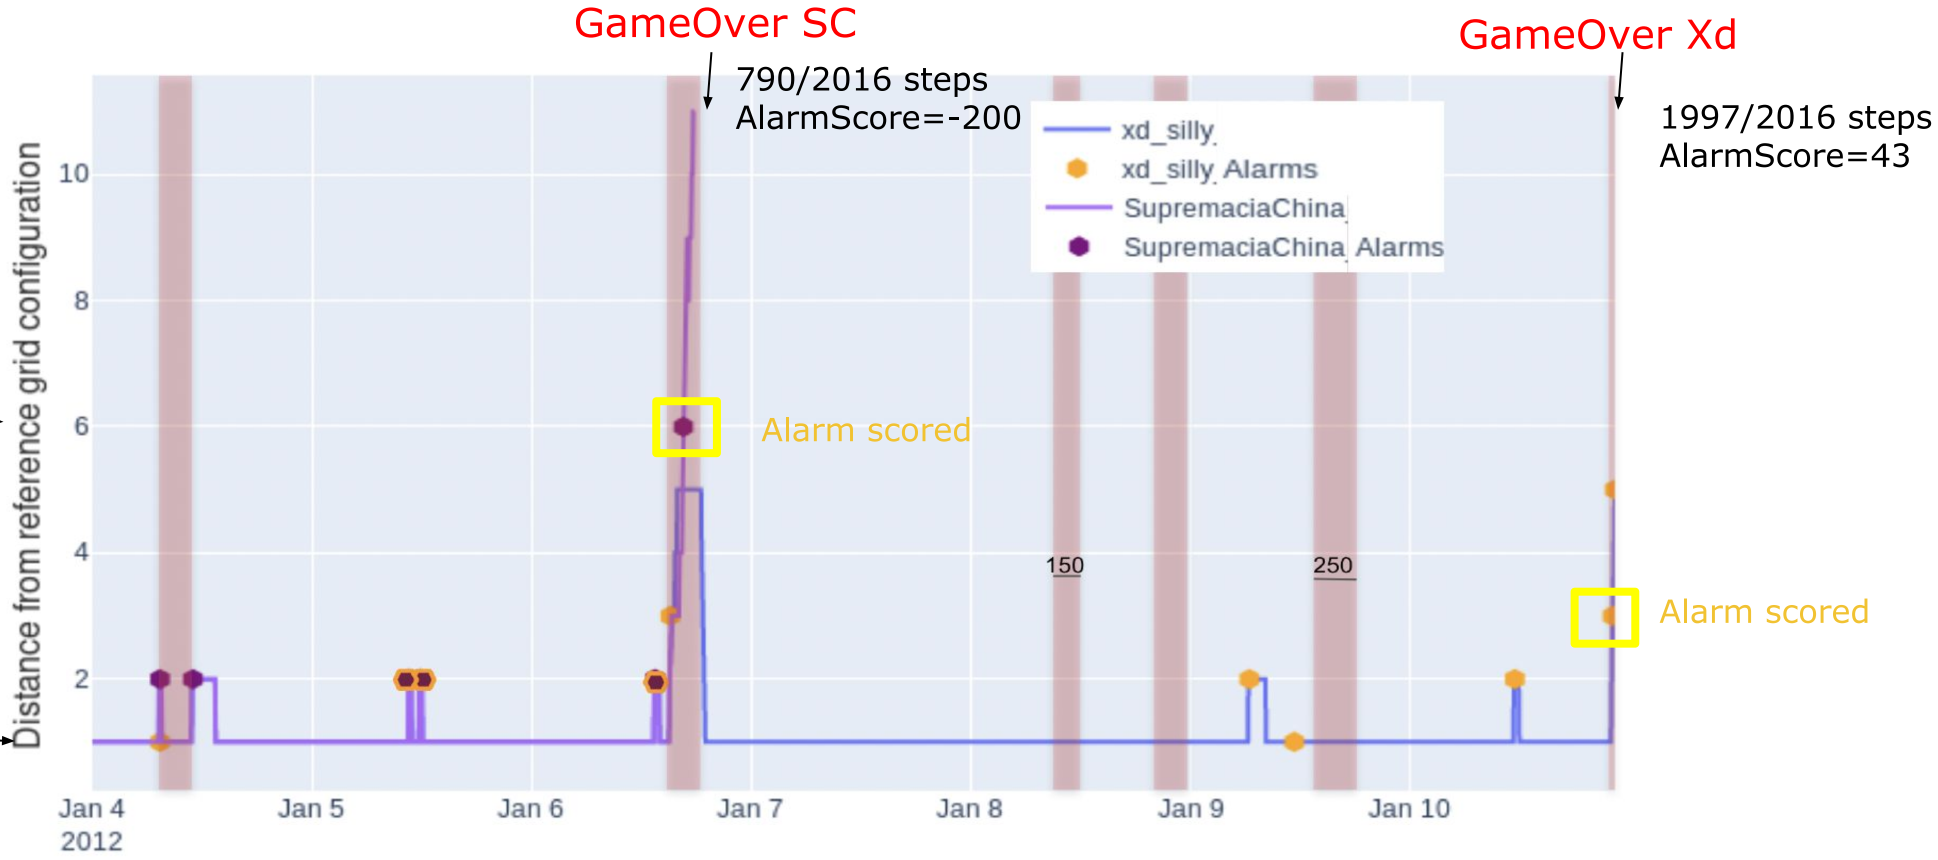

Fig. 6: 2 best agents behaviour over time and before respective failures in scenario jan28_1. It shows times of actions (as the topology distance varies) and alarms, and periods of attacks)

Statistiques budget and alarm for Xd_silly (vs SupremaciaChina)

- 0.63 alarm per day on average (respectively 0.78)
- keeps an average budget of 2.5 (respectively 2.2)
- only spend 1.5% of the time with attention budget < 1 (respectively 10% of the time)

Trust framework

Before a human can trust an agent, high levels of (i) credibility, (ii) reliability and (iii) intimacy are required according to the **Trust Equation** (by Charles Green in *The Trusted Advisor book*):

- **Credibility** = increases when the agent is transparent and explains the proposed action
    - Modeled by requesting time and localisation of alert
- **Reliability** = agent act consistently for similar situations and "knows" the limits of its capabilities
    - Modeled by operational performance and alerting capability evaluation
- **Intimacy** = the relationship quality overtime with the right level of relevant interactions and as little misunderstandings as possible
    - Modeled by the attention budget of the operator

04

Conclusions

Conclusions

1. Machine Learning approaches are promising for control problems when fast computation is required and here showed to **generalize** better than Expert Systems

2. It is now possible to **experiment with adaptable, robust and trustworthy AI assistant**, opening a new field for research and smart grid flexibilities.

3. Challenge **helps develop benchmarks** and enforce reproducibility to make faster and stronger progress as a community

4. We should keep working on **attracting AI researchers** and collaborating with them on power system related problems

References
"L2RPN challenge: a Retrospective Analysis", PMLR 2021
"L2RPN with Trust", PSCC 2022
"Towards an AI assistant for the human grid operator", HHAI 2022

Sponsors & Collaborators

Principal coordinators:

- Antoine Marot (RTE, France)
- Isabelle Guyon (U. Paris-Saclay; UPSud/INRIA, France and ChaLearn, USA)

Protocol and task design:

- Gabriel Dulac-Arnold (Google Research, France)
- Olivier Pietquin (Google Research, France)
- Isabelle Guyon (U. Paris-Saclay; UPSud/INRIA, France and ChaLearn, USA)
- Patrick Panciatici (RTE, France)
- Antoine Marot (RTE, France)
- Benjamin Donnot (RTE, France)
- Camilo Romero (RTE, France)
- Jan Viebahn (TenneT, Netherlands)
- Adrian Kelly (EPRI, Ireland)
- Mariette Awad (American University of Beirut, Lebanon)
- Yang Weng (Arizo State Univ., USA)

Data format, software interfaces, and metrics:

- Benjamin Donnot (RTE, France)
- Mario Jothy (Artelys, France)
- Gabriel Dulac-Arnold (Google Research, France)
- Aidan O'Sullivan (UCL/Turing Institute, UK)
- Zigfried Hampel-Arias (Lab 41, USA)
- Jean Grizet (EPITECH & RTE, France)

Environment preparation and formatting:

- Carlo Brancucci (Encoord, USA)
- Vincent Renault (Artelys, France)
- Camilo Romero (RTE, France)
- Bri-Mathias Hodge (NREL, USA)
- Florian Schäfer (Univ. Kassel/pandapower, Germany)
- Antoine Marot (RTE, France)
- Benjamin Donnot (RTE, France)

Baseline methods and beta-testing:

- Kishan Prudhvi Guddanti (Arizo State Univ., USA)
- Loïc Omnes (ENSAE & RTE, France)
- Jan Viebahn (TenneT, Netherlands)
- Medha Subramanian (TenneT & TU Delft, Netherlands)
- Benjamin Donnot (RTE, France)
- Jean Grizet (EPITECH & RTE, France)
- Patrick de Mars (UCL, UK)
- Lucas Tindall (Lab 41 & UCSD, USA)

L2RPN keeps running in 2022

Energies of the future and carbon neutrality

Participate and join us! https://l2rpn.chalearn.org/

legacy phases:

- Robustness https://competitions.codalab.org/competitions/25426
- Adaptability https://competitions.codalab.org/competitions/25427
- Trust https://competitions.codalab.org/competitions/33121