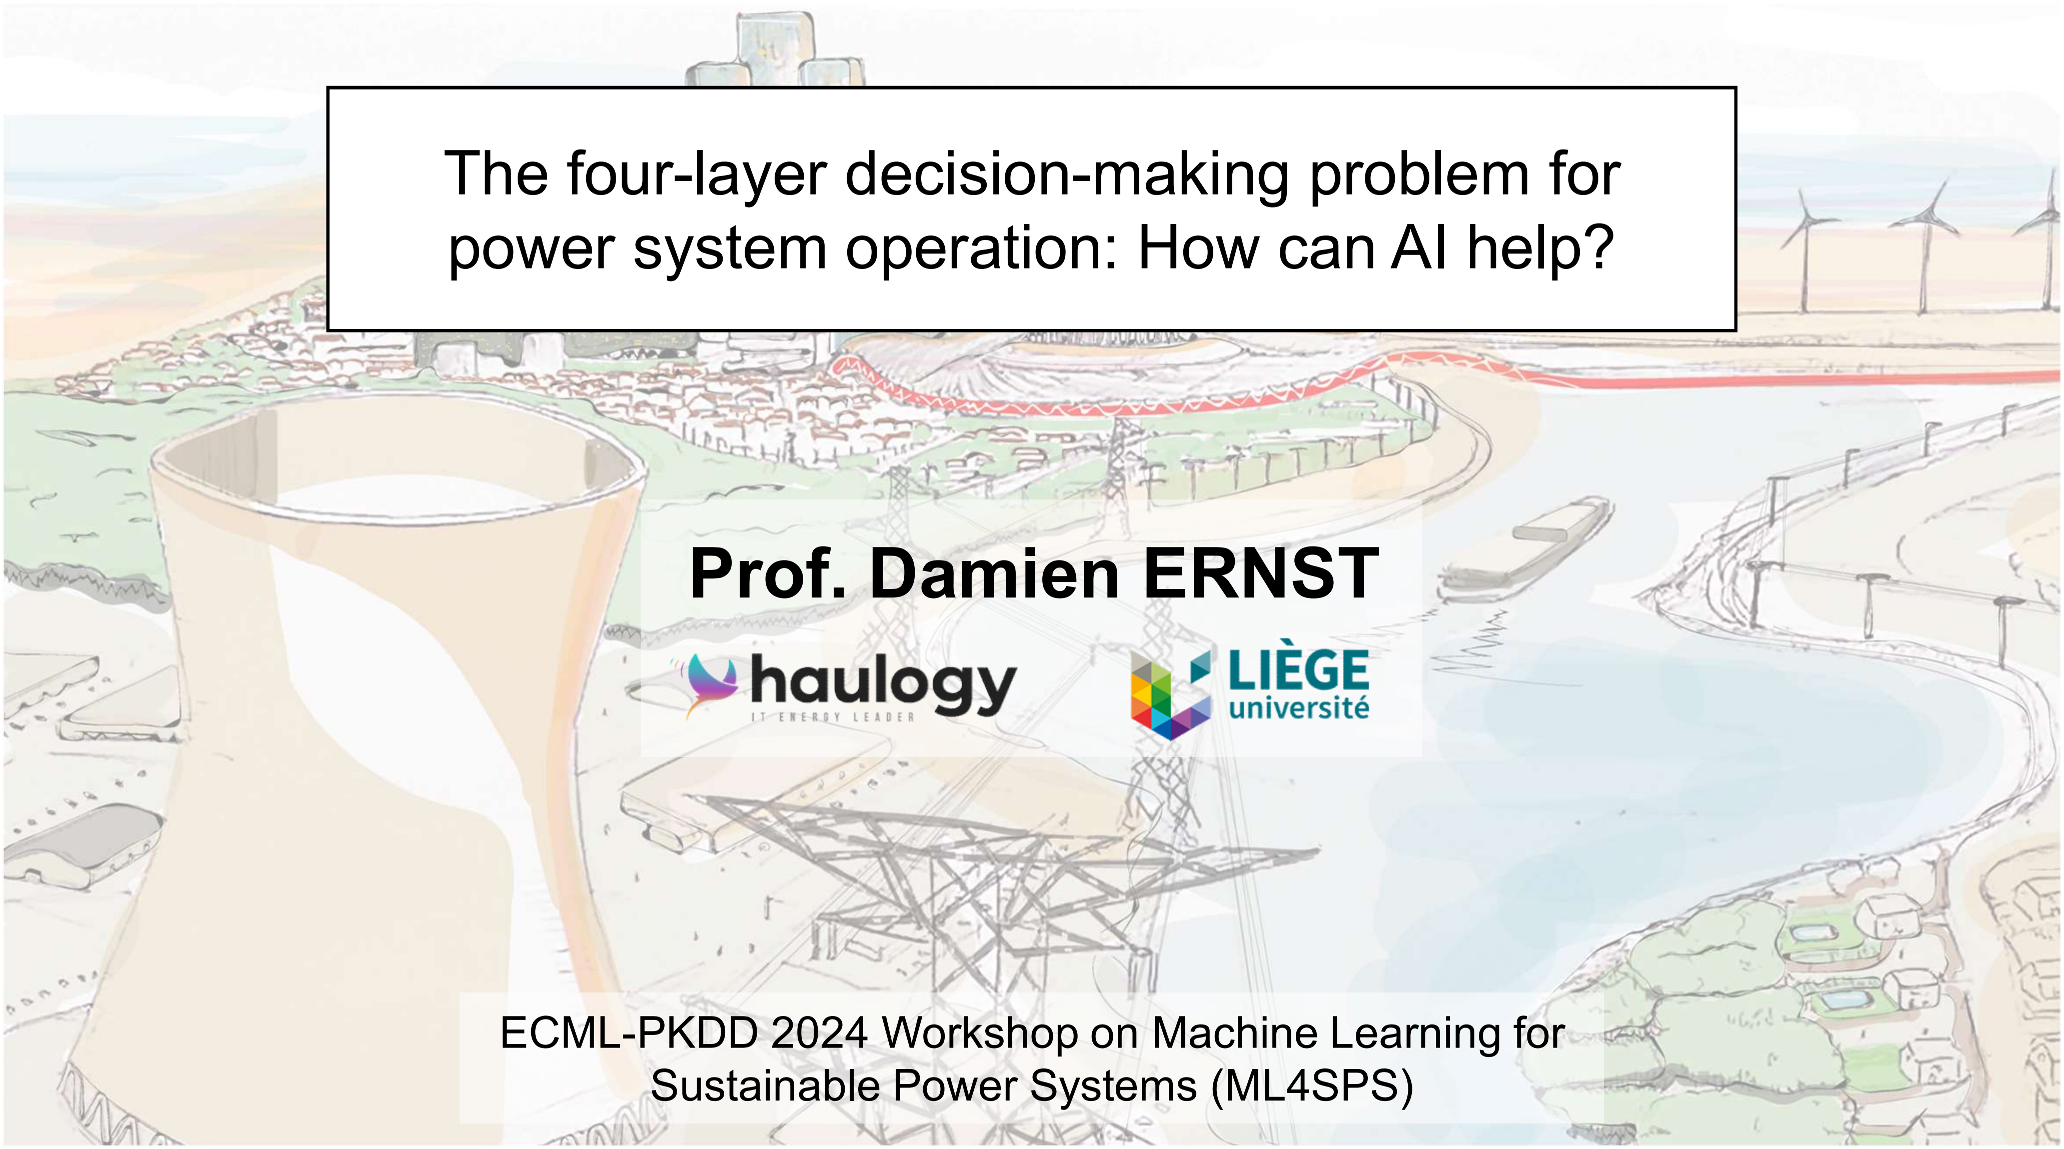

The four-layer decision-making problem for power system operation: How can AI help?

Prof. Damien ERNST

haulogy

LIÈGE université

ECML-PKDD 2024 Workshop on Machine Learning for Sustainable Power Systems (ML4SPS)

Power systems are one of the most complex engineering structures ever built by mankind.

This talk takes you on a journey through the different layers of power systems, with the aim of showing to machine learning (ML) researchers the numerous decision-making problems each layer presents.

The solutions to these problems will help considerably to accelerate the decarbonization of our societies.

The first layer: Microgrids

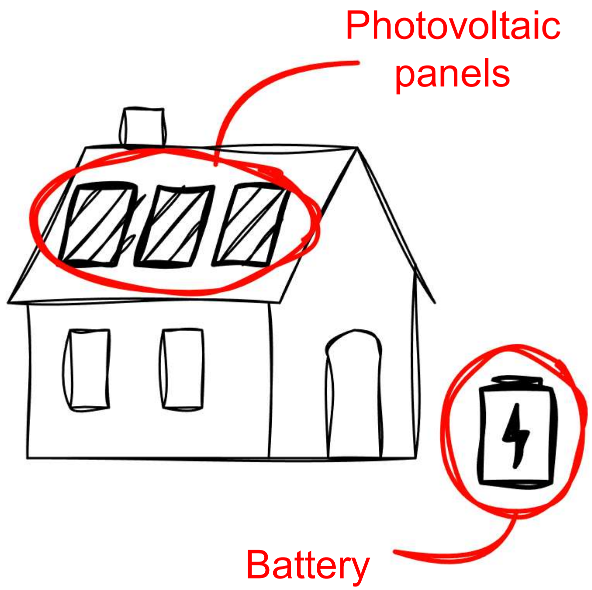

The illustration depicts a simple and standard version of a microgrid, featuring a single load and two distributed energy resources (a photovoltaic (PV) installation and a battery).

(Near) optimal control policies for microgrids

A microgrid is a system that includes one or more loads and one or more distributed energy resources (DERs), which can operate in parallel with the broader electric network.

Let us consider the challenges related to the computation of (near) optimal policies **for controlling microgrids* to minimise energy costs**.

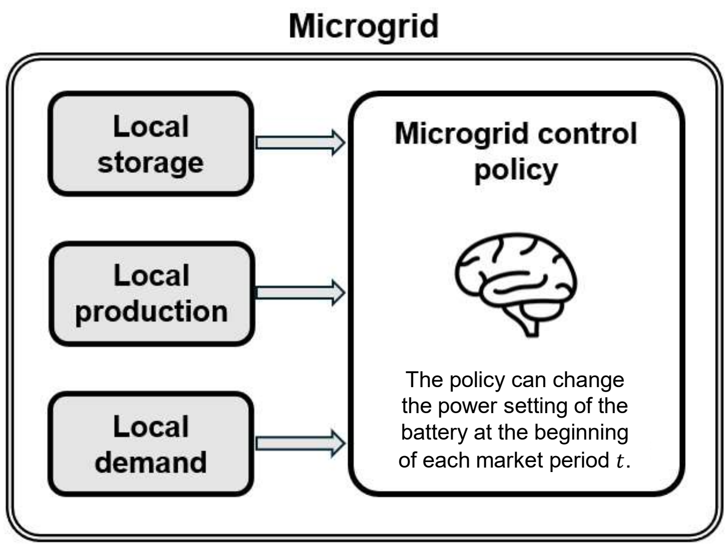

*We focus in the following on microgrids which are located in the European Union with its specific electricity market rules.

Understanding electricity trading in the European Union

Before proceeding further, it is important to understand how electricity is traded in the European Union.

(i) Electricity **is treated as a commodity, bought and sold on a quarter-hourly basis**. This means that the electricity commodity traded always refers to quantities of electrical energy associated with market periods of 15 minutes;

(ii) Electricity **can be bought and sold in various markets**, such as the day-ahead market or the intraday market;

(iii) The retailer **can act as an interface between markets and consumers**. The retail products are defined by, for example, a feed-in tariff for electricity injection (selling) and a purchase tariff for electricity consumption (buying).

Case 1 – Perfect forecast

Context: The computation of a (near) optimal microgrid control policy used to dictate to what level the battery is charged and discharged (**power setting**) for every market period t with knowledge of all relevant time-series data (e.g., weather, load, feed-in and purchase tariffs).

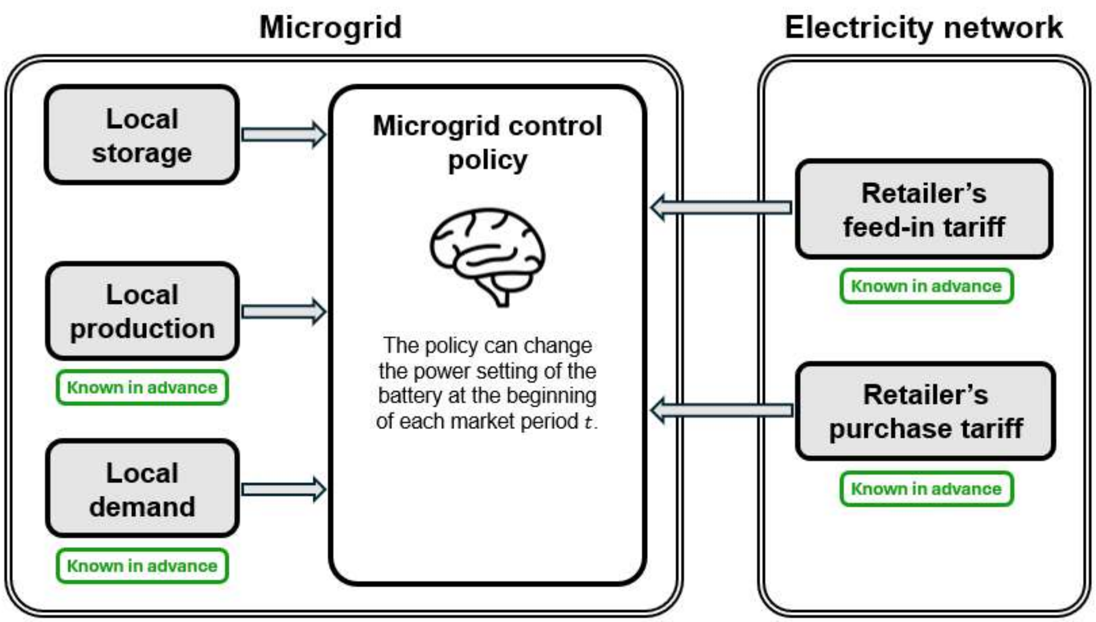

How to compute an optimal microgrid control policy (1/2)

A policy is an object that, at the beginning of each **market period** $\boldsymbol{t} \in \{\mathbf{0}, \mathbf{1}, \ldots, \boldsymbol{T}-\mathbf{1}\}$ within the optimisation horizon $\boldsymbol{T}$, selects an **action** $\boldsymbol{u_t}$ from the **action space** $\boldsymbol{\mathcal{U}}$, based on **a piece of information** $\boldsymbol{i_t}$ from the set of **available information** $\boldsymbol{\mathcal{I}}$.

Representation of t within the time horizon T.

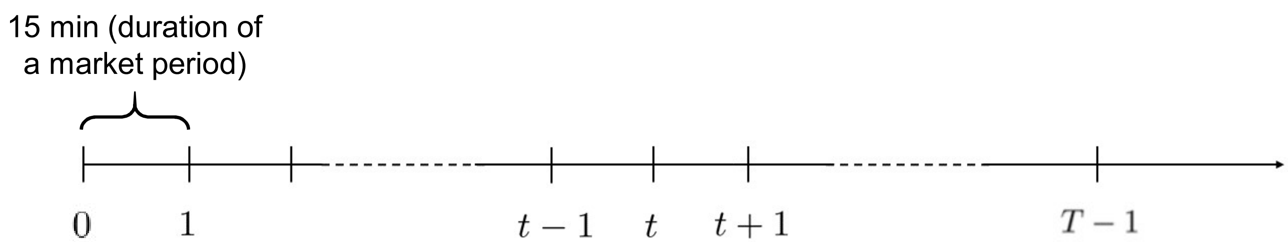

How to compute an optimal microgrid control policy (2/2)

In our context, **assuming linear behaviour for the battery dynamics**, the energy cost minimisation problem for the microgrid can be approached as a linear programming problem. This problem involves determining an action $u_t \in \mathcal{U}$ at each time step $t \in \{0, 1, \dots, T-1\}$ that sets the electrical power to charge (discharge) into (from) the battery.

$$\left(u_0^*, \dots, u_{T-1}^*\right) \in \underset{(u_0,\dots,u_{T-1}) \in \mathcal{U}^T}{\arg\min} \; Cost(u_0, \dots, u_{T-1})$$
$$\text{where } \mathcal{U} = [Power_{\min}, Power_{\max}]$$

However, **if we want to be closer to reality, we need to model the non-linear short-term and long-term dynamics of the battery**. This includes factors such as losses that do not vary linearly with the power setting and battery aging.

Why should you be excited as ML researchers

- A (near) optimal control problem of a non-linear system that can be formalised as an **MDP($\boldsymbol{\varepsilon}$)**, which refers to a Markov Decision Process that includes **a time-series variable $\boldsymbol{\varepsilon}$**. A long time horizon T needs to be considered. Very little work has been done so far in the fields of MDP(ε) with a long time horizon. [Case 1]

Case 2 – Load and weather to be forecasted

Context: The computation of a (near) optimal microgrid control policy for a battery power setting with time-series of both feed-in and purchase tariffs as input data.

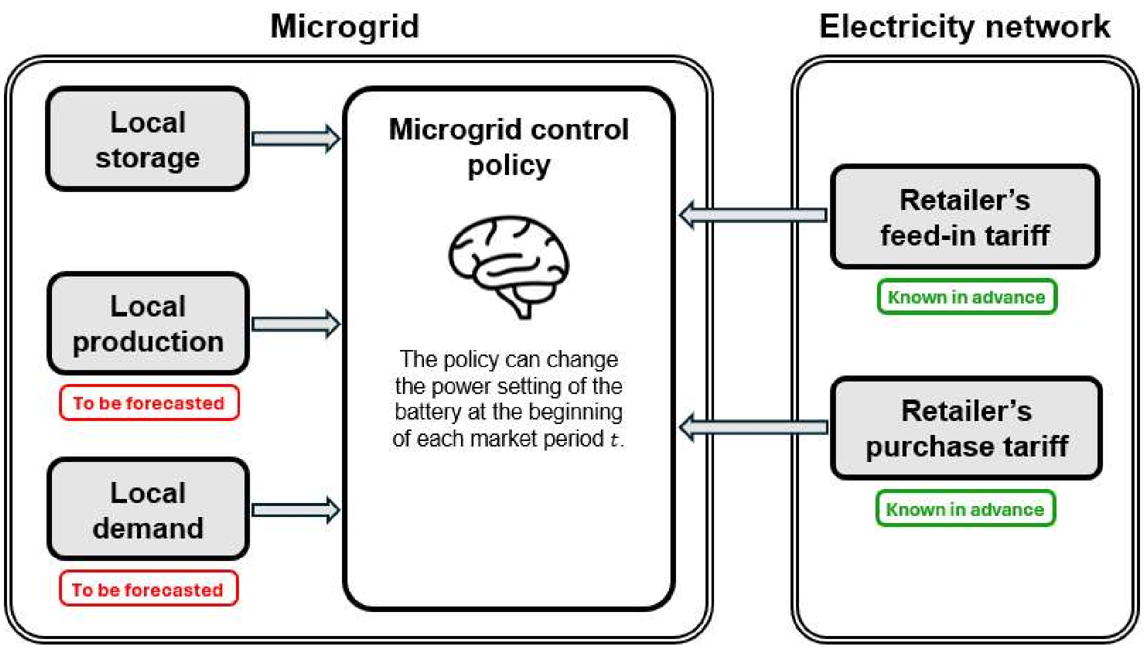

This second case adds two challenges

Policies take values in the same action space $\mathcal{U}$ as in the first case. However, the key difference from Case 1 is the set $\mathcal{I}$ of available information used to compute optimal actions, as the weather and load time-series data are no longer available, resulting in local production and demand being unknown.

This leads to the two following challenges:

1. Forecasting weather and load time-series at a very local level;

2. Even with accurate probabilistic models for generating forecasts, optimal control requires solving a sequential and stochastic problem across numerous time steps. We believe that **policy gradient techniques in RL (direct policy search methods) could be effective for addressing this challenge**.

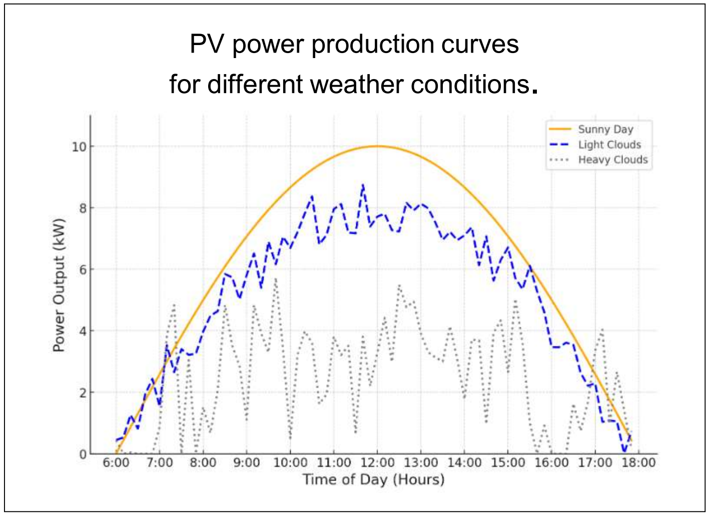

Why should you be excited as ML researchers

- A (near) optimal control problem of a non-linear system that can be formalised as an **MDP($\boldsymbol{\varepsilon}$)**, which refers to a Markov Decision Process that includes **a time-series variable $\boldsymbol{\varepsilon}$**. A long time horizon needs to be considered. Very little work has been done so far in the fields of MDP(ε) with a long time horizon. [Cases 1 and 2]
- The forecasting of very local weather and load time-series requires one to rely on probabilistic forecasts. This problem remains largely unsolved to date. [Case 2]

Case 3 – Market enters the game

Context: The computation of a (near) optimal microgrid control policy for the battery power setting without any time-series as input data (e.g., weather, load, feed-in and purchase tariffs) as input, while also enabling bidding on the day-ahead market and accounting for the imbalance tariff imposed by the Transmission System Operator (TSO).

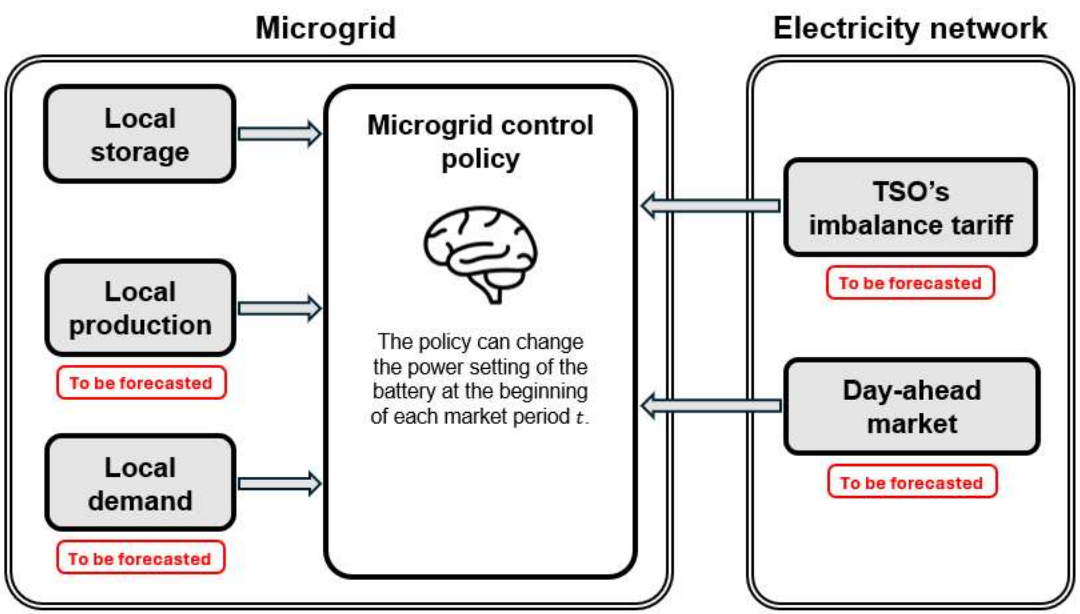

Complexity of day-ahead market interactions for the policy (1/2)

In Case 3, no longer are there feed-in and purchase tariffs; the microgrid interacts directly with the day-ahead market. This requires submitting at each $t = 47 + k \times 96, k \in \{0,1,2,\ldots\}$ a vector $v_t \in \mathcal{V} \subset \mathbb{R}^{24}$ of energy quantities to be traded on the day-ahead market.

The energy quantity $v_{47+k\times 96}[i]$ where $i \in 1,2,\ldots,24$ represents four times the quantity of energy that is bought for every quarter of an hour i of the next day $k+1$. We therefore have to compute at every $t = 47 + k \times 96, k \in \{0,1,2,\ldots\}$ an additional action $v_t \in \mathcal{V}$.

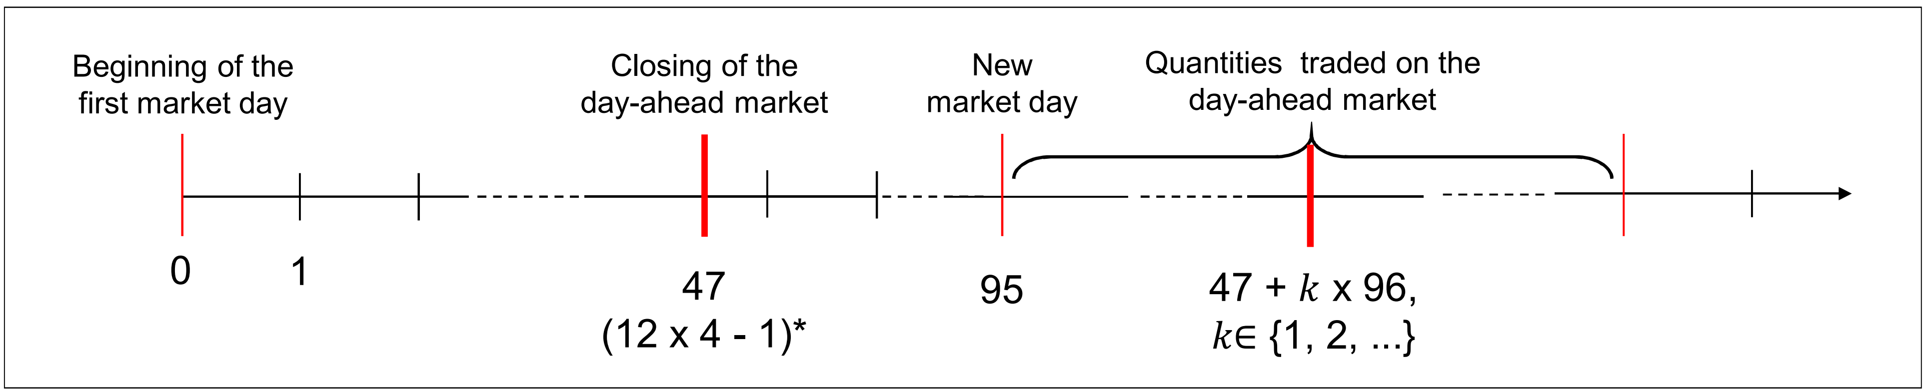

*Bidding closes at noon.

Why should you be excited as ML researchers?

- A (near) optimal control problem of a non-linear system that can be formalized as an **MDP(ε)**, which refers to a Markov Decision Process that includes **a time-series variable ε**. A long time horizon needs to be considered. Very little work has been done so far in the fields of the MDP(ε) with long time horizon. [Cases 1 to 3]
- The forecasting of very local weather and load time-series requires one to rely on probabilistic forecasts. This problem remains largely unsolved to date. [Cases 2 and 3]
- Dealing with very large action spaces, where some actions can only be taken at specific time steps (day-ahead market). [Case 3]

Case 4 – forecasting with all relevant markets

Context: Computation of a (near) optimal microgrid control policy for the battery power setting without any time-series data as input, while also considering as additional market the intraday market.

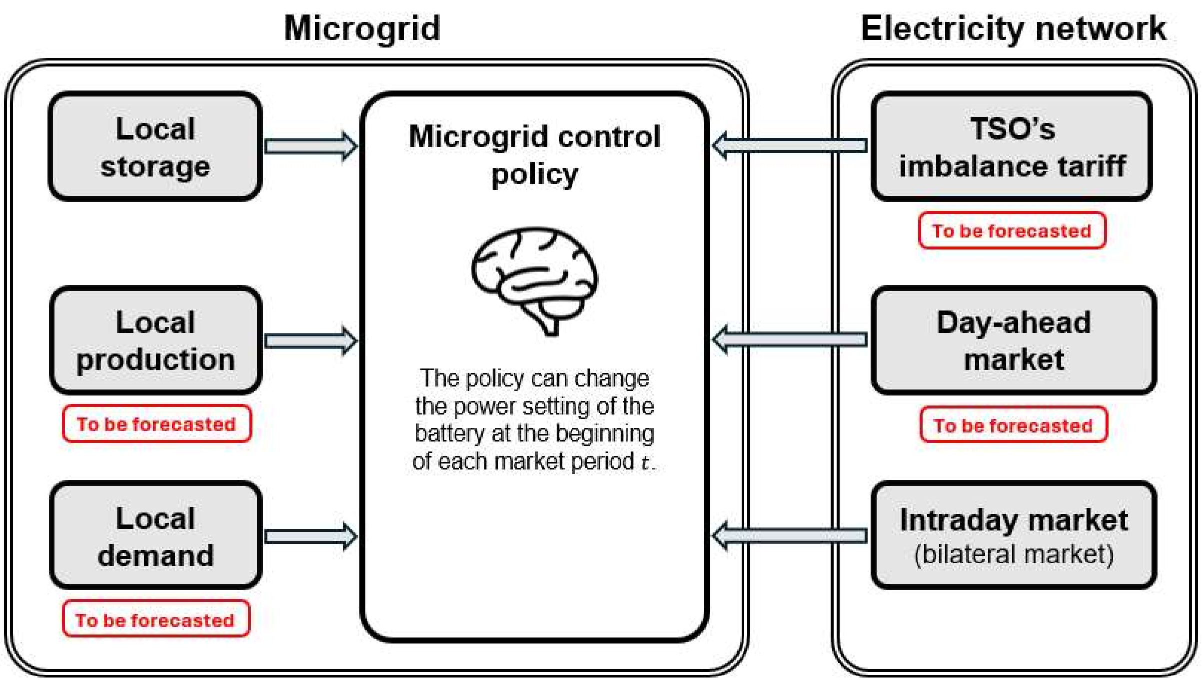

Action space and bid acceptance for intraday market decisions

In Case 4, a specific action space $\mathcal{W}$ related to the intraday market must be considered in addition to the previous action spaces $\mathcal{U}$ and $\mathcal{V}$. We assume that at the beginning of each market period t, a Boolean vector w_t must be computed to decide whether to accept or reject each available bid in the order book of the intraday market.*

If there are n_t available bids, then the action space $\mathcal{W}$ at t is:

$$
\boxed{\begin{gathered} w_t \in \mathcal{W}, \text{ where } \mathcal{W} = \{0,1\}^{n_{\max}}, \\ n_{\max} = \max\{n_t | t \in \{0, \ldots, T-1\}\} \end{gathered}}
$$

w_t

Trading (continuous and discrete) and delivery timelines for products $Q_1$to Q_4.

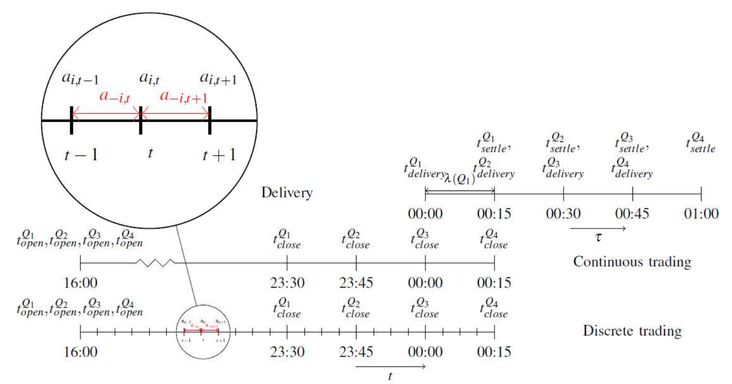

Note that there may be thousands and thousands of bids ($n_t >> 1$)!

*The microgrid could also generate its own bids for the intraday market and access existing bids during a market period.

Why should you be excited as ML researchers

- A (near) optimal control problem of a non-linear system that can be formalized as an **MDP(ε)**, which refers to a Markov Decision Process that includes **a time-series variable ε**. A long time horizon needs to be considered. Very little work has been done so far in the fields of MDP(ε) with long time horizon. [Cases 1 to 4]
- The forecasting of very local weather and load time-series requires one to rely on probabilistic forecasts. This problem remains largely unsolved to date. [Cases 2 to 4]
- Dealing with very large action spaces, where some actions can only be taken at specific time steps (day-ahead market). [Cases 3 and 4]
- The intraday market integration adds additional forecasting challenges that depend on multitude of factors (e.g., weather, geopolitical, behaviour of traders, etc). The resulting MDP(ε) shows immensely large action spaces, with various action types that can be taken. [Case 4]

Case 5 – Operation and sizing

Context: Computation of a (near) optimal microgrid control policy for the battery power setting while addressing all the previously mentioned challenges, and a (near) optimal sizing of both the PV installation capacity and the battery storage capacity.

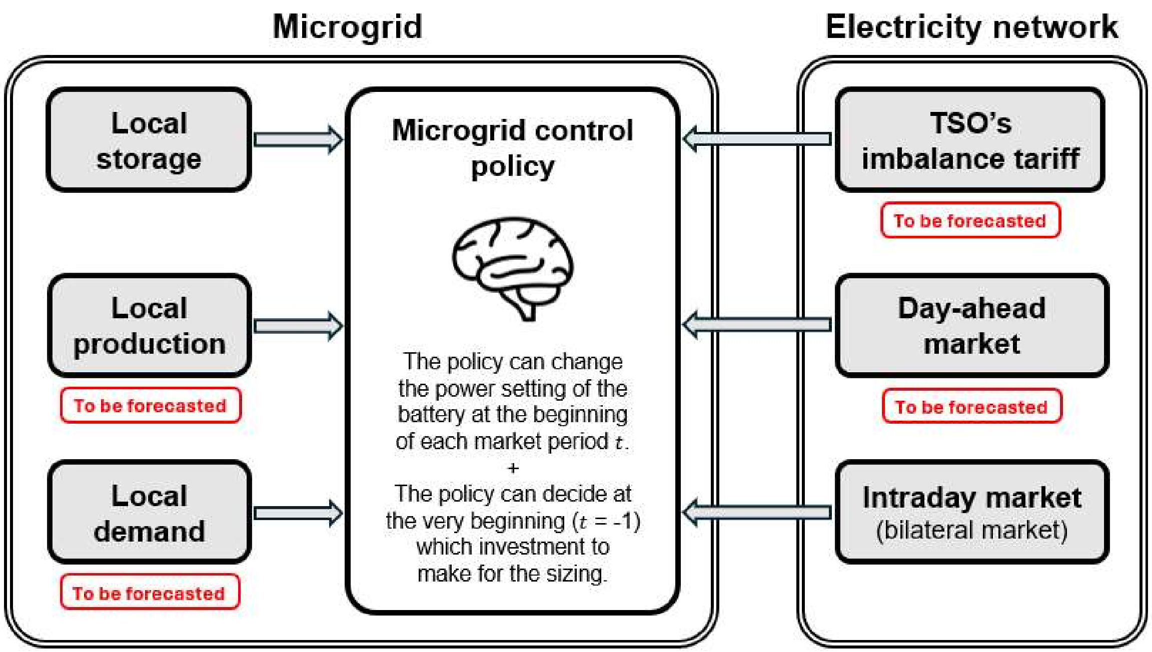

Sizing problem: an additional one-time decision

The parameters related to the sizing of the capacities belong to a new action space $\mathcal{Z}$ and need to be optimized.

More specifically, we assume that an additional one-time decision must be made at time -1, corresponding to the choice of the size of the PV installation capacity and the battery storage capacity:

$$\boxed{z_{-1} \in \mathcal{Z} \text{ where } \mathcal{Z} \subset \left(\mathbb{R}^+\right)^p, p \in \mathbb{N}}$$

A reinforcement learning point of view on the problem

The problem of sizing and operation can be seen from RL researchers as a problem of optimizing both a policy and an environment. Here, the environment is parameterised by z_{-1} (which includes PV installation capacity and battery storage capacity) and the policy by θ.

$\pi_\theta : \mathcal{I} \times \{0, \ldots, T-1\} \rightarrow \mathcal{U} \times \mathcal{V} \times \mathcal{W}$: policy-taking decision based on accessible information $i_t \in \mathcal{I}$

$\theta \in \Theta$: policy parameterisation

$$(z^*_{-1}, \theta^*) \in \underset{(z_{-1},\theta)}{\arg\max}\ V(z_{-1}, \theta) \text{ where } V(z_{-1}, \theta) = \underset{\substack{s_0 \sim P_0(\cdot) \\ a_t = \pi_\theta(i_t, t) \\ \xi_t \sim P_\xi(\cdot|t) \\ r_t = \rho_{z_{-1}}(s_t, a_t, \xi_t)\ :\ \text{reward function} \\ s_{t+1} = f_{z_{-1}}(s_t, a_t, \xi_t)\ :\ \text{dynamics}}}{\mathbb{E}} \sum_{t=0}^{T-1} r_t$$

More information in: Bolland, A., Boukas, I., Berger, M., & Ernst, D. (January 2022). **Jointly learning environments and control policies with projected stochastic gradient ascent**. *Journal of Artificial Intelligence Research*, 73, 117-171. https://hdl.handle.net/2268/247965

Timeline of actions across the different action spaces

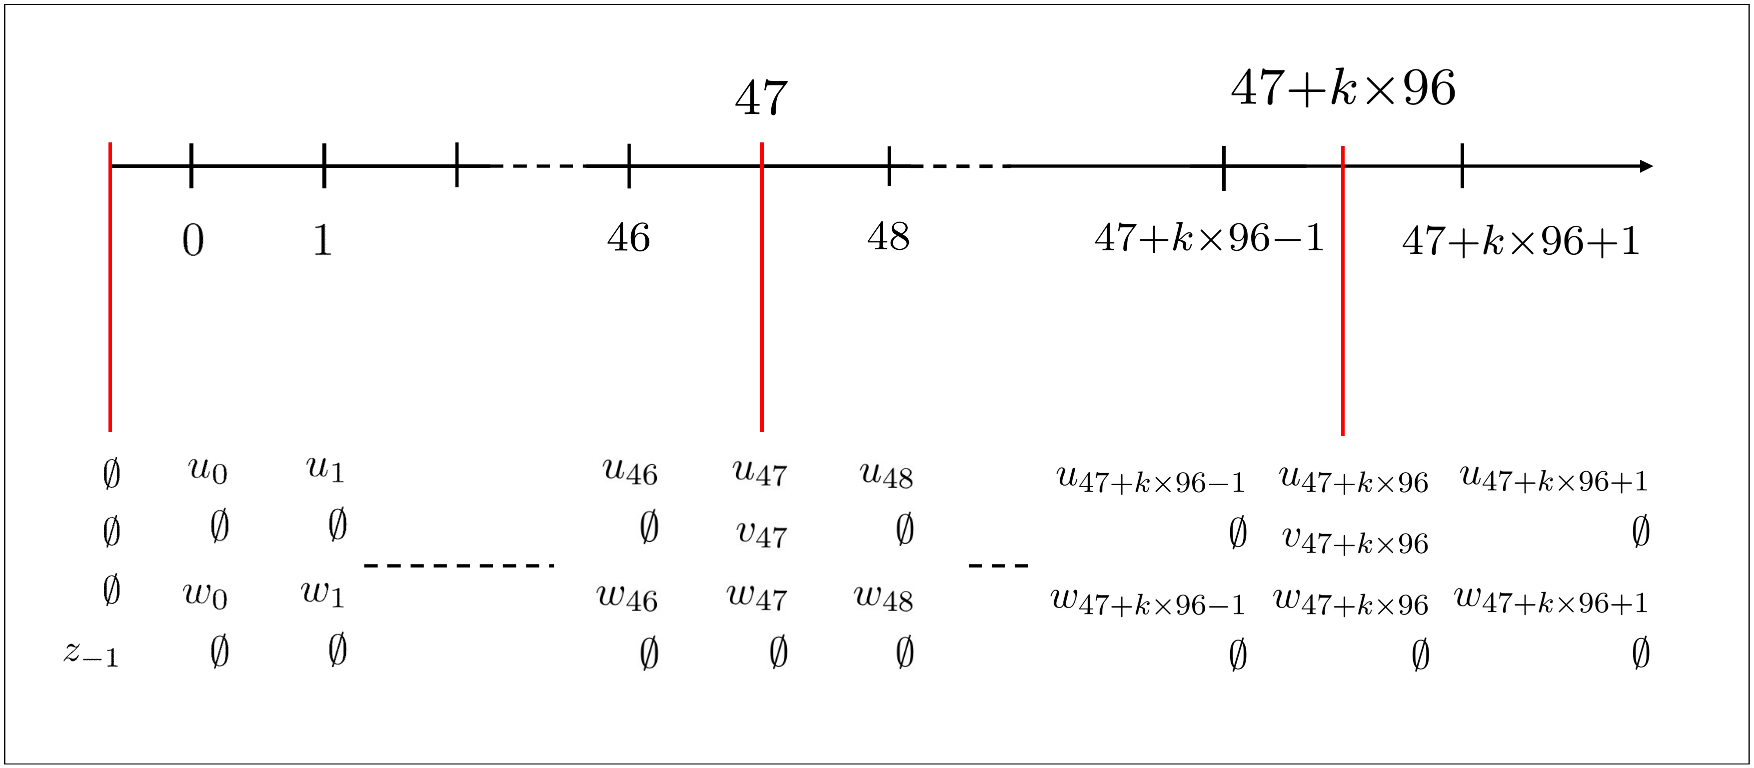

Why should you be excited as ML researchers

- A (near) optimal control problem of a non-linear system that can be formalized as an **MDP(ε)**, which refers to a Markov Decision Process that includes **a time-series variable ε**. A long time horizon needs to be considered. Very little work has been done so far in the fields of MDP(ε) with long time horizon. [Cases 1 to 5]
- The forecasting of very local weather and load time-series requires one to rely on probabilistic forecasts. This problem remains largely unsolved to date. [Cases 2 to 5]
- Dealing with very large action spaces, where some actions can only be taken at specific time steps (day-ahead market). [Cases 3 to 5]
- The intraday market integration adds additional forecasting challenges that depend on multitude of factors (e.g., weather, geopolitical, behaviour of traders, etc). The resulting MDP(ε) shows immensely large action spaces, with various action types that can be taken. [Cases 4 and 5]
- Determining the optimal sizing of the microgrid to minimize overall costs, including both investment and operational expenses. This can be formalised as a joint optimization of the control policy and the environment. [Case 5]

The second layer: Distribution networks

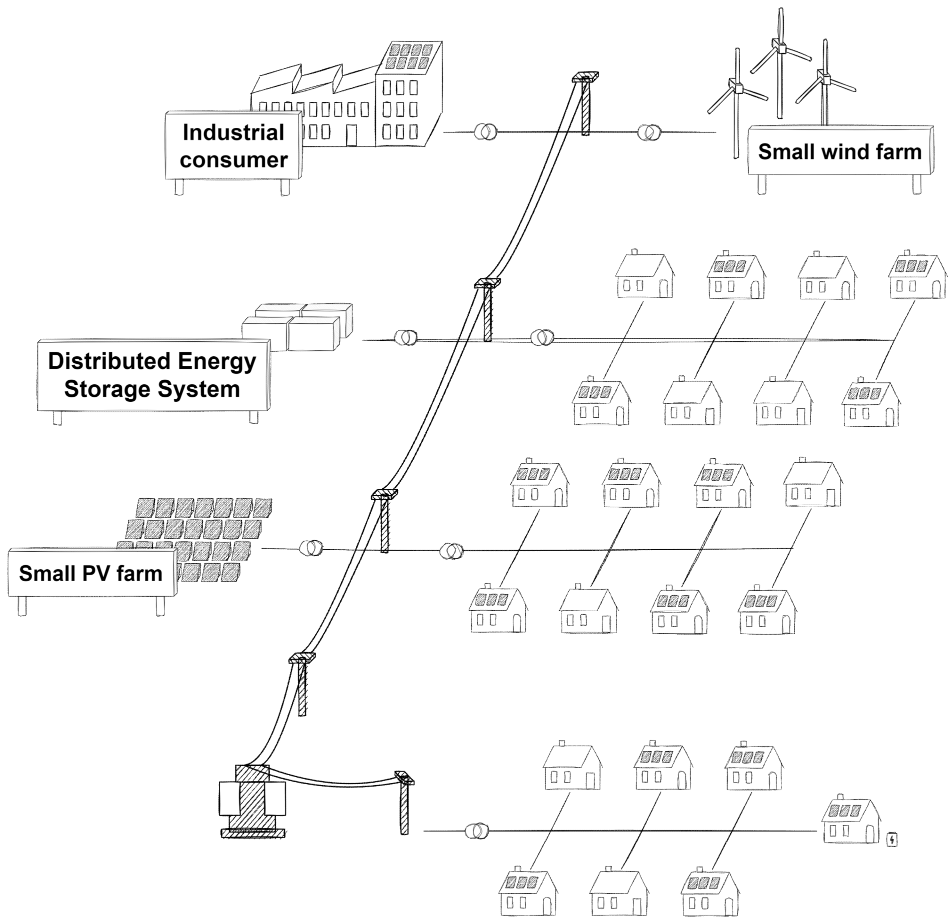

The “fit and forget” doctrine for managing distribution networks has reached its limitations

Ten years ago, issues in distribution networks were almost non-existent. These networks operated under the “fit and forget” doctrine, meaning they were built robustly enough so that any grid user could be connected without facing power quality issues (e.g., interruption of supply, overvoltages). Investment decisions were based on established “good practice” rules.

Today, the situation has changed drastically. **The rise of Distributed Energy Resources (DERs)** like PV installations and of new loads like heat pumps and Electric Vehicles (EVs), is creating serious power quality issues. These include **congestions,** which refer to situations where excessive current flows through the network elements, leading to overloads, as well as **voltage issues** characterised by voltages often going above or below admissible values.

The pace at which DERs and new loads **are being integrated into the network is exceeding the financial capabilities of Distribution System Operators (DSOs)** to implement their "good practice" investment rules to maintain the fit and forget doctrine.

An example of a DER causing a distribution network issue: disconnection of domestic PV installations.

Domestic PV installations tend to turn consumers into producers during sunny hours as these consumers are unable to fully consume the electricity generated, causing voltage to rise as excess electricity is pushed into the low-voltage distribution network. When the voltage reaches a certain level (e.g., nominal voltage of 230V + 10%), the PV panels are disconnected.

To mitigate such issues, it is essential for DSOs to **establish (near) optimal investment policies taking into account their limited investment capabilities.**

Production of a PV installation affected by voltage issues.

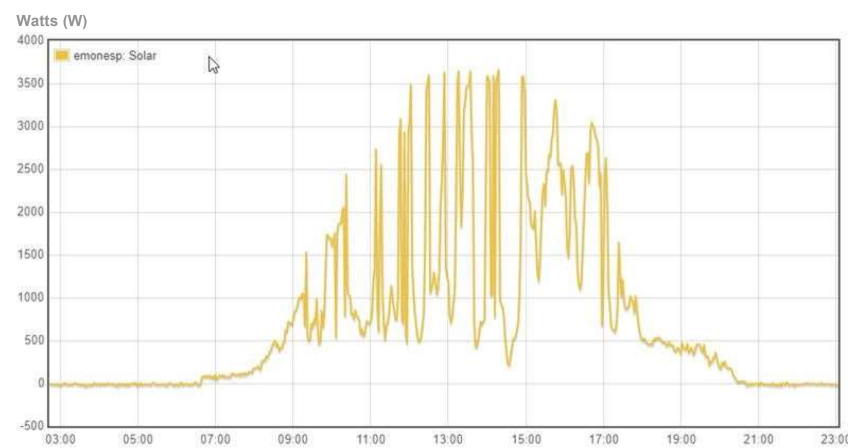

Computing optimal investment decision-making strategies (1/3)

Each year, the DSO has a specific investment budget to maintain and/or expand its network.

The challenge is to **allocate this investment budget in a way that minimises network issues over time**.

This is an extremely difficult optimisation problem. Why?

Difficulty 1:

This challenge has an **inherently sequential nature**. Investment budget decisions made each year affect a sequence of interdependent decisions, requiring consideration of both immediate impacts and long-term consequences.

Computing optimal investment decision-making strategies (2/3)

Difficulty 2:

The decision-making process is also challenging due to the **vast number of potential investment options**. Even with perfect foresight – turning it into a deterministic problem – the complexity remains significant.

For a given network, each year, a set of **components** $\mathcal{C}$ can either be upgraded or not. In the simplest case, the **action space** $\mathcal{U}$ is defined as:

$$\boxed{\mathcal{U} = \{0,1\}^{|\mathcal{C}|}}$$

where 1 represents an upgrade and 0 represents no upgrade.

A more realistic (and complex) approach to the problem would involve determining the type of upgrade (e.g., capacity increase or technology enhancements) and the addition of new components (e.g., transformers).

An illustration of an investment decision-making process

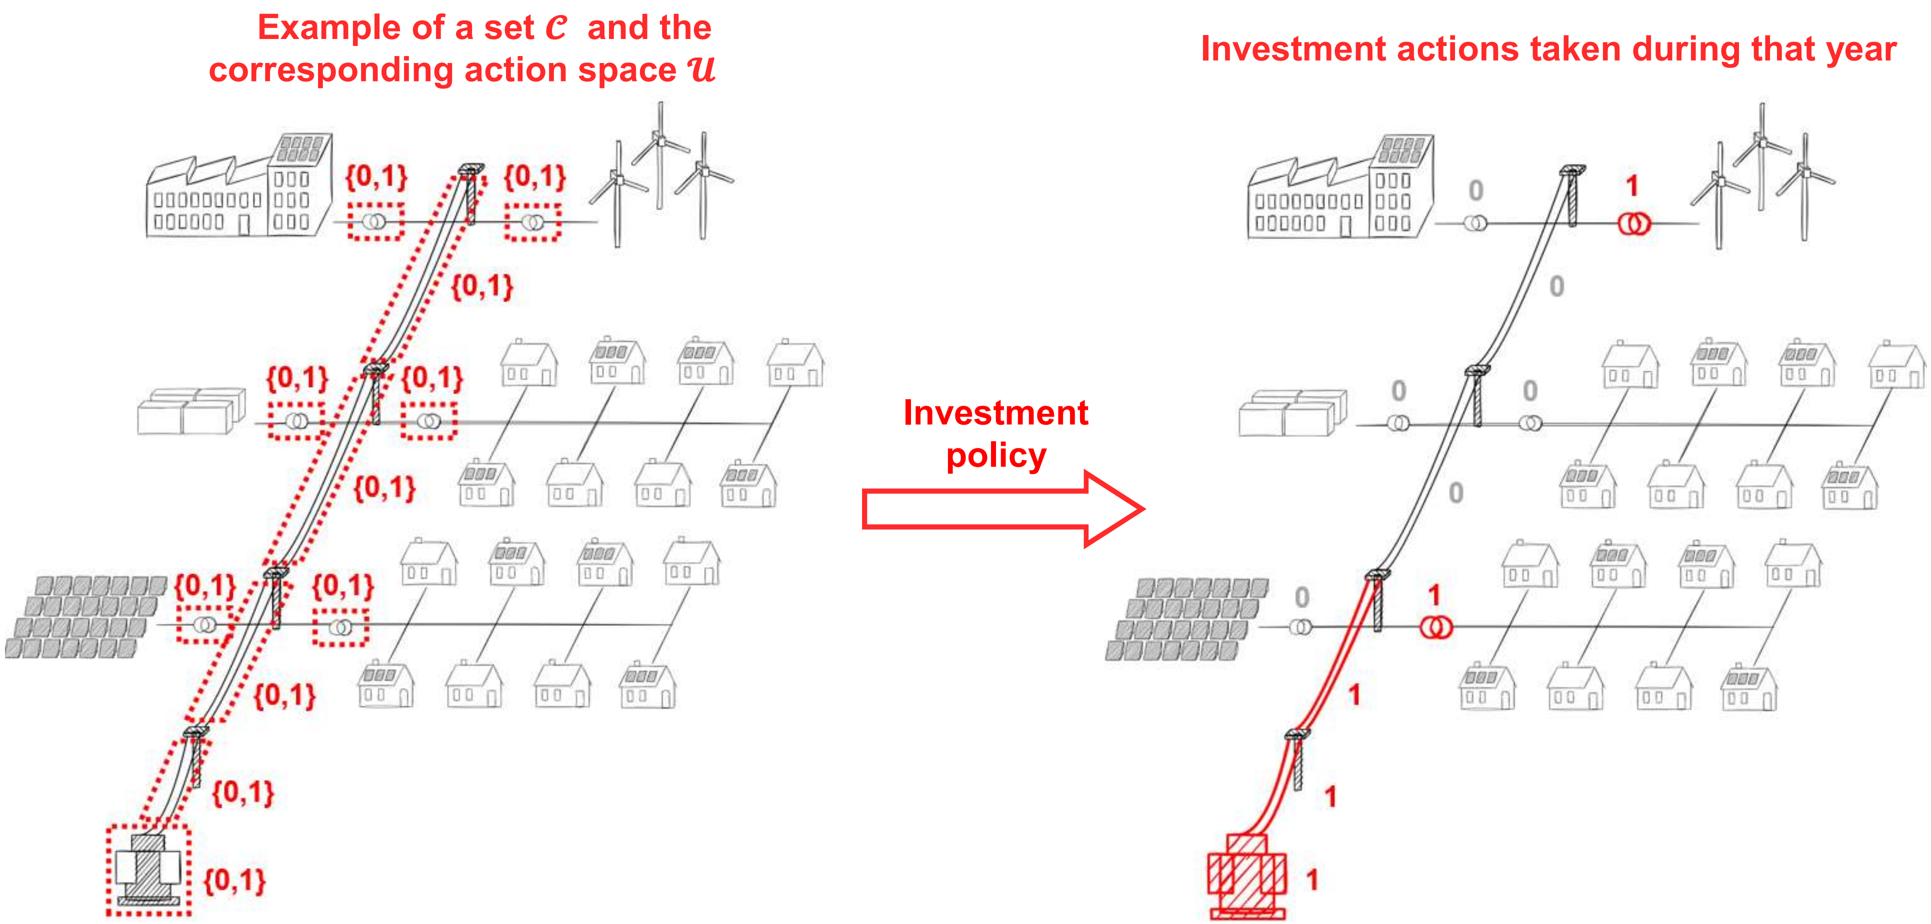

Computing optimal investment decision-making strategies (3/3)

Difficulty 3:

The problem is **highly stochastic** and uncertain, involving weather forecasts, customer behaviour and technology prices. These uncertainties can be represented as probability distributions, which can be used to generate future scenarios.

AI has shown promising results for weather forecasting, with numerous works demonstrating its potential. However, predicting the evolution of customer behaviours, technological advancements, and their associated costs remains a significant challenge, and to date, **AI has not provided reliable solutions**.

Developing ML techniques to address these prediction problems would be a great challenge for AI researchers!

A smart investment policy alone will not be enough: Active Network Management (ANM) is also necessary

ANM schemes use a policy **to modulate power generation sources, loads and batteries** to avoid congestions and voltage issues on the distribution networks.

ANM policies can help reduce investment costs!

The problem of computing ANM policies can be modelled as a stochastic sequential optimal control problem. The sequential nature is amplified due to time-dependent constraints introduced by batteries. **This makes it a fascinating challenge for RL researchers to explore and develop effective solutions.**

GYM-ANM: A user-friendly RL framework for network management

Are you an RL researcher looking to develop an RL algorithm for solving ANM problems without having to know too much about distribution networks?

The **GYM-ANM** environment is for you!

GYM-ANM is a framework for designing RL environments that model ANM tasks in distribution networks. These environments provide new playgrounds for RL research in the management of electricity networks that do not require an extensive knowledge of the underlying dynamics of such systems.

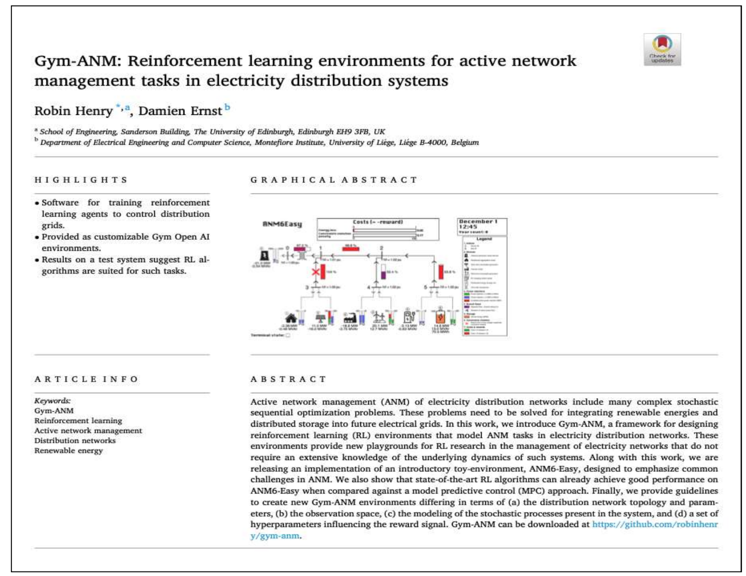
Gym-ANM: Reinforcement learning environments for active network management tasks in electricity distribution systems

Robin Henry [*,a], Damien Ernst [b]

[a] School of Engineering, Sanderson Building, The University of Edinburgh, Edinburgh EH9 3FB, UK
[b] Department of Electrical Engineering and Computer Science, Montefiore Institute, University of Liège, Liège B-4000, Belgium

H I G H L I G H T S

- Software for training reinforcement learning agents to control distribution grids.
- Provided as customizable Gym Open AI environments.
- Results on a test system suggest RL algorithms are suited for such tasks.

G R A P H I C A L A B S T R A C T

A R T I C L E I N F O

Keywords:
Gym-ANM
Reinforcement learning
Active network management
Distribution networks
Renewable energy

A B S T R A C T

Active network management (ANM) of electricity distribution networks include many complex stochastic sequential optimization problems. These problems need to be solved for integrating renewable energies and distributed storage into future electrical grids. In this work, we introduce Gym-ANM, a framework for designing reinforcement learning (RL) environments that model ANM tasks in electricity distribution networks. These environments provide new playgrounds for RL research in the management of electricity networks that do not require an extensive knowledge of the underlying dynamics of such systems. Along with this work, we are releasing an implementation of an introductory toy-environment, ANM6-Easy, designed to emphasize common challenges in ANM. We also show that state-of-the-art RL algorithms can already achieve good performance on ANM6-Easy when compared against a model predictive control (MPC) approach. Finally, we provide guidelines to create new Gym-ANM environments differing in terms of (a) the distribution network topology and parameters, (b) the observation space, (c) the modeling of the stochastic processes present in the system, and (d) a set of hyperparameters influencing the reward signal. Gym-ANM can be downloaded at https://github.com/robinhenry/gym-anm.

A graphical example of an ANM policy modulating a six-bus distribution network

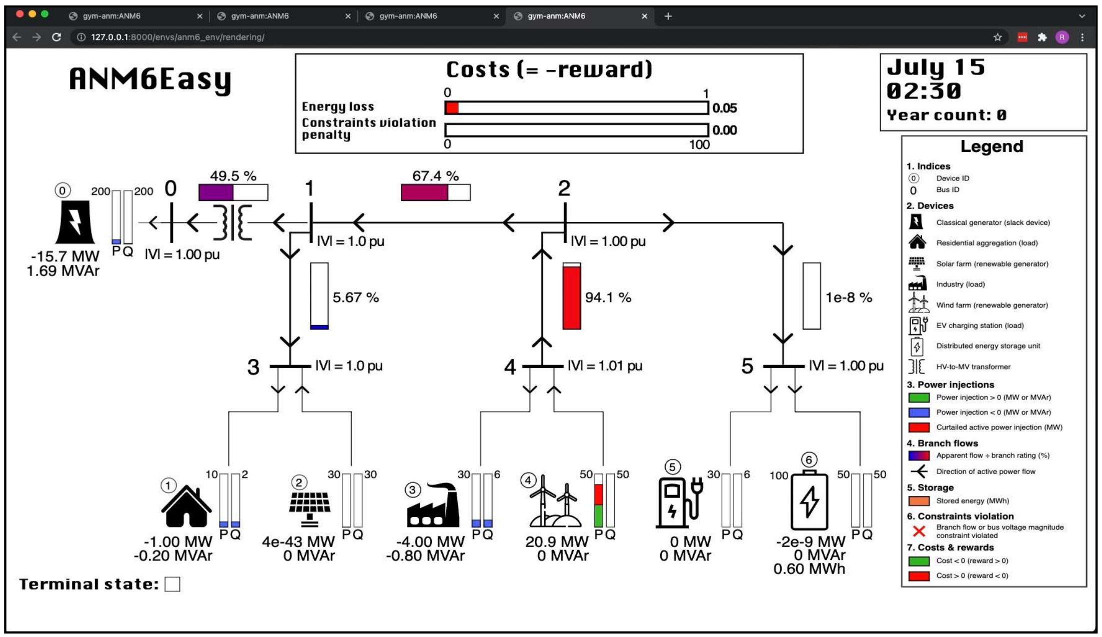

https://youtu.be/D8kGH94kavY

What you observe in the video

The distribution network consists of six buses, one high to medium-voltage transformer, three aggregated passive loads, two renewable energy generators, one battery, and one fossil fuel generator.

This video showcases two situations:

1. A windy night with low consumption: During this period, PV production is zero and wind production is nearly at its maximum. Due to the low demand, the ANM policy curtails wind production to prevent overheating of the transmission lines [2:30]. Towards the end of the night, the policy sets the batteries to charge [4:45] in preparation for the morning peak of EVs [08:00].

2. A sunny, windy weekday: During the day, residential PV and wind production exceed the demand [11:30]. The ANM policy curtails both sources [12:00 for PV, 13:00 for Wind] of energy while charging the batteries [14:00] to store extra energy to prepare for the late afternoon EV charging period [16:15].

Coupling investment policies with the computation of ANM policies

To achieve optimality, we need to identify the **best pair (investment policy, ANM policy)**.

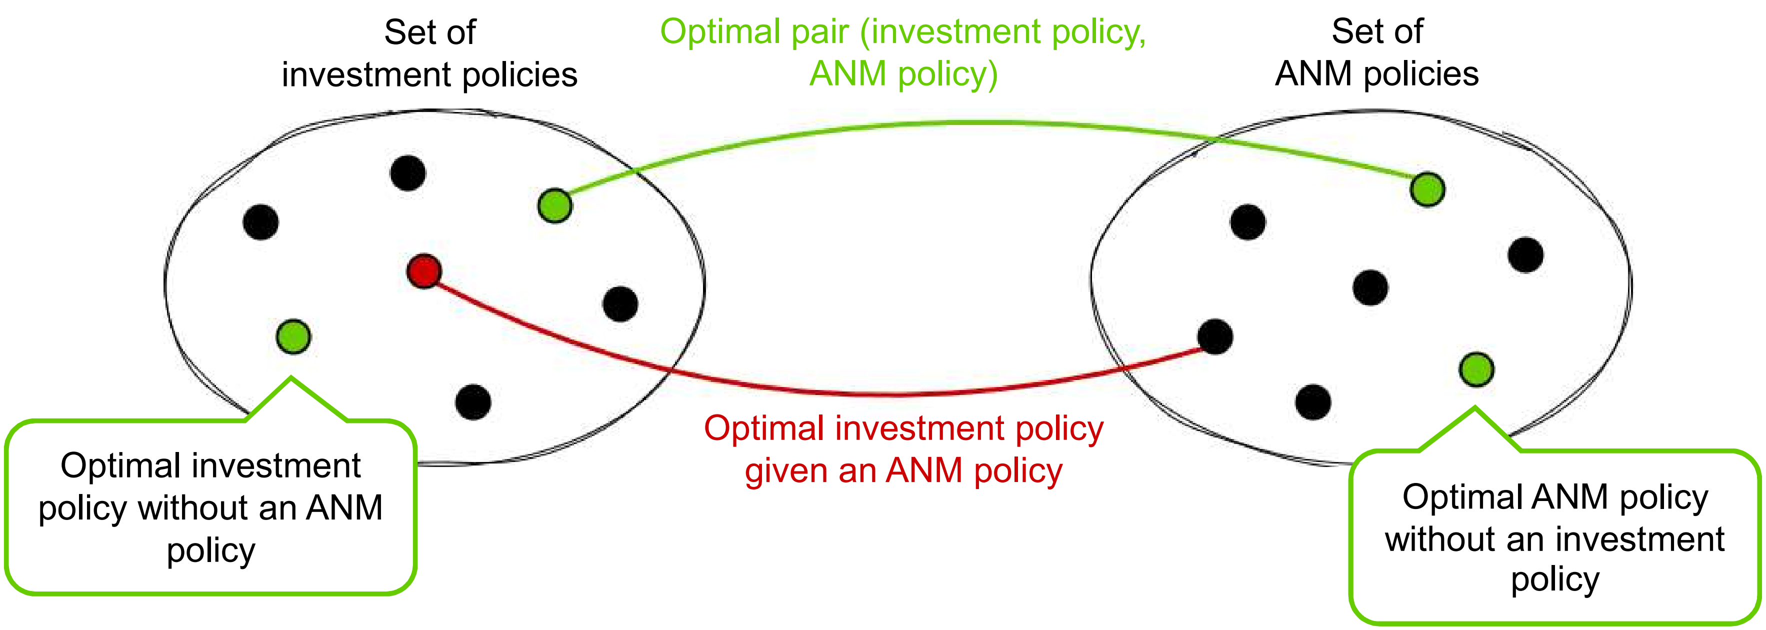

The third layer: Transmission networks

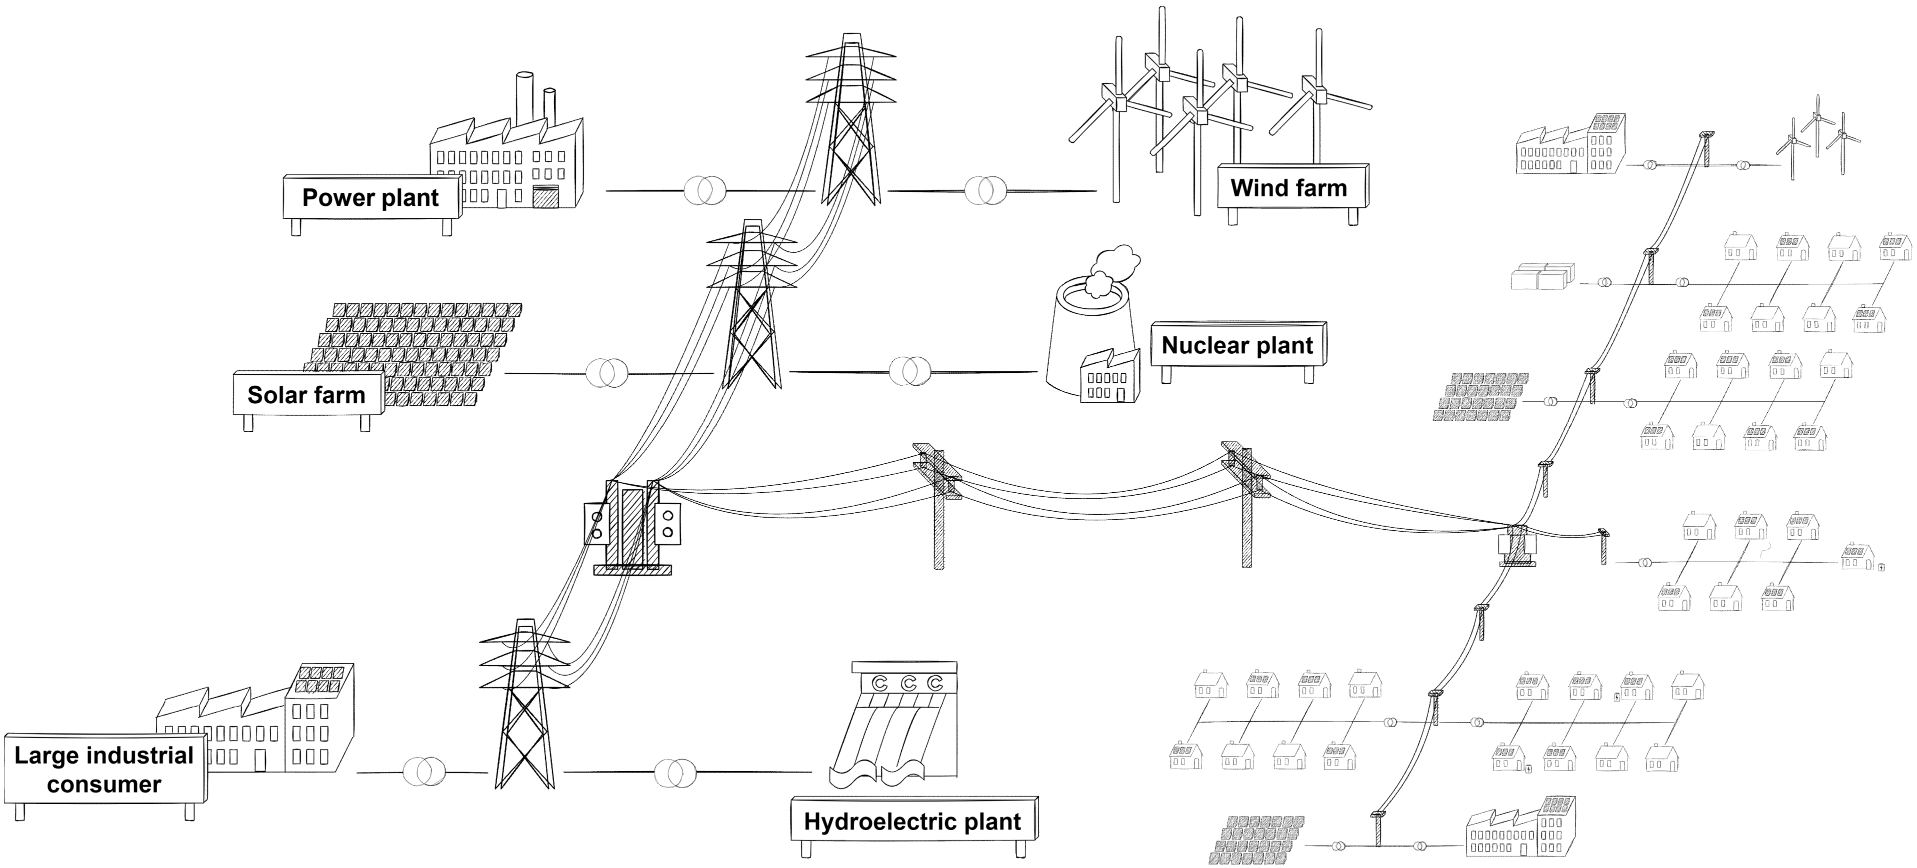

Decision-making strategies in transmission networks

The transmission network covers various decision-making strategies occurring at different time scales, ranging from a few decades (investment policies) to a few milliseconds (e.g., protecting devices against short circuits).

Power system control centres are where critical decisions to **ensure the safe operation of the power system** are made. Operators in these centres play a crucial role, especially in high-pressure situations.

Poor decisions in control centres can be fatal, potentially leading to blackouts. Such failures may result in significant economic losses and can endanger lives!

A picture of the national control centre of RTE, the French Transmission System Operator (TSO).

Machine learning in power system control centres

Resolution schemes for many decision-making strategies in control centres could involve ML techniques. A recurring issue with existing decision-making tools that rely on optimisation techniques is that **their computational complexity is often too high to generate the right decisions fast enough** to handle critical situations (e.g., the trip of transmission lines, the disconnection of a large-scale wind farm).

Researchers are exploring the use of **Deep Neural Networks (DNNs)** as alternatives to these optimisation programmes, aiming for quicker decision-making based on current conditions.

Deep Neural Networks for voltage control

In transmission networks, the static voltage control problem seeks to minimise voltage constraint violations by determining optimal controls (e.g., generator voltage setpoints, transformer ratios, shunt/self activations, line openings) for a given power network context.

The main approach involves:

(i) Using a DNN to process the power network context as input and generate the control variables as output;

(ii) Training the DNN with the (RL) REINFORCE algorithm to avoid the need to solve numerous optimisation problems, which would be required in a supervised learning setting;

(iii) Implementing a Graph Neural Network (GNN) to effectively exploit the topology of the power network.

A little bit of mathematics behind the approach (1/2)

The static voltage control problem in transmission networks can be expressed as the following parametric optimisation problem:

$$\boxed{\min_{y \in \mathcal{Y}(x)} f(y, x)}$$

where the context x and the controls y are structured as a graph.

The function f is considered as a **black-box function** and may **not** be **differentiable** due to discontinuities in the underlying power network simulator. It also encodes or penalises the constraints.

The actions y can be **high-dimensional**, and the available controls $\mathcal{Y}$ depend on the context x.

A little bit of mathematics behind the approach (2/2)

The approach learns a neural network that maps the context of a power network to control variables $\boldsymbol{y}_{\boldsymbol{\theta}}: \boldsymbol{x} \rightarrow \boldsymbol{y}$.

(i) It relies on the computation of $\nabla_{\theta} f(y_{\theta}(x), x)$ using the REINFORCE gradient estimation with a Gaussian policy $\Pi_{\theta}: \mathcal{N}(y_{\theta}(x), \sigma 1)$.

(ii) It requires a neural network that can adapt to varying topologies and the number of controls in power networks. This is accomplished using a general GNN architecture that can accommodate for **hyper, heterogenous and multi graphs (H2MG)**.

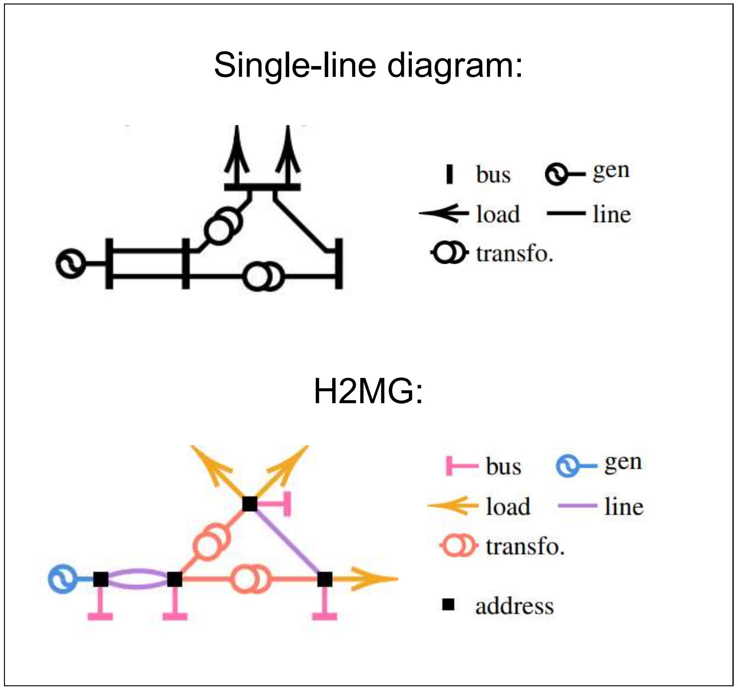

Power system operators can make mistakes

On November 4, 2006, a major blackout occurred in Europe, leaving more than 15 million customers without electricity for several hours.

The blackout was caused by human error, specifically the failure to adhere to the N-1 security doctrine.

Human Error To Blame For Europe Blackout

Nov. 15, 2006

The switching off of a line coupled with outage of a second line triggered blackout.

Agence France-Presse

Germany's biggest power supplier said on Nov. 15 that human error was to blame for the electricity cut that plunged parts of Western Europe into darkness on November 4. E.ON said the switching-off of an electricity line over the Ems River in western Germany to allow a cruise ship to pass through, coupled with the outage of a second transmission line, "set off the domino effect which led to the temporary disconnection of the European inter-connected power grid."

Could AI replace control centre operators to ensure better security and efficiency in power systems?

The Grid2Op environment for training your power system operator in the GYM framework

RTE has developed a GYM environment for training your power system operator agent.

One of the proposed objectives is to control the power system at minimal cost to avoid congestions. There are two types of actions:

Costly actions: Modifying generation or flexible loads (including batteries).

Non-costly actions: Actions related to elements owned by the TSO (e.g., network topology changes, control of tap changers).

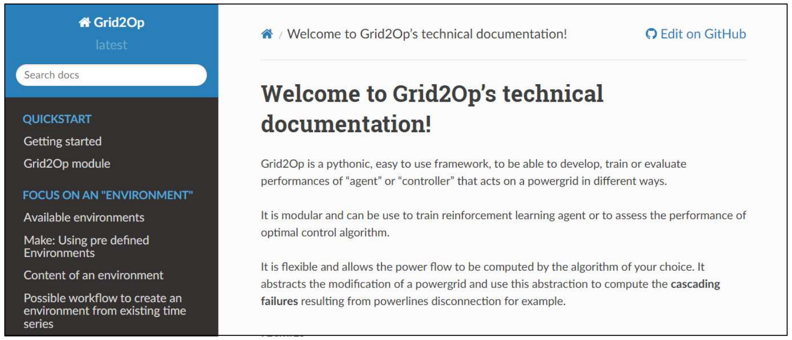

Does RL work on voltage control?

The winning approach of the 2022 L2RPN challenge utilised Monte Carlo Tree Search (MCTS) techniques. Heuristics were employed to reduce the action space, which is immense.

RL algorithms based on function approximators seem to struggle with changes in system topologies. It is as if they can only successfully generalise around a **reference topology**.

Indeed, so far, they seem to offer good results only when combined with a heuristic that keeps them within the vicinity of the reference topology.

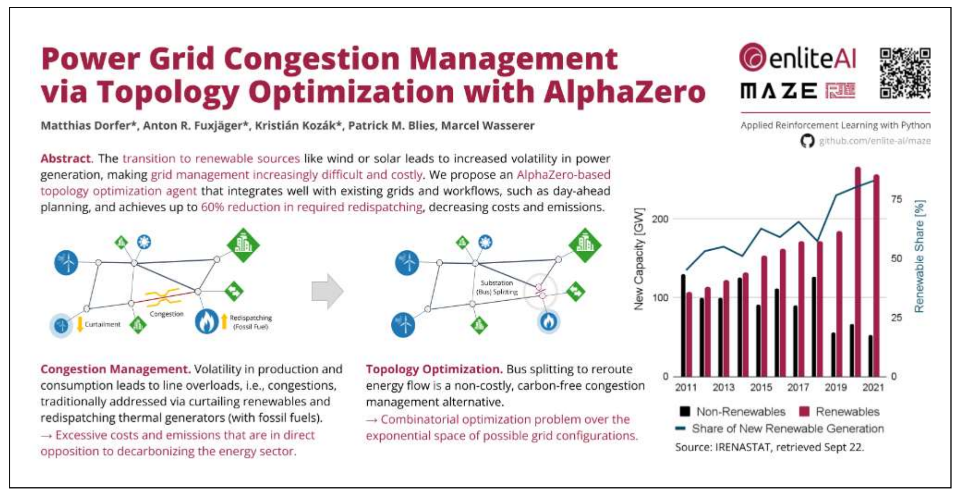

Fundamental RL research driven by Grid2Op

This problem is linked to very fundamental RL research questions, such as:

(i) Defining metrics related to the difficulty of generalisation in RL. The size of the state space is not the only thing that matters;

(ii) Developing specific techniques for generalisation in RL when parts of the state space can be represented by a graph, with a structure that can be determined by a history of actions (e.g., opening/closing lines, topological changes in substations, etc.).

It is likely that if more RL researchers seriously engaged with this environment, we would see significant advances in RL.

The fourth layer: the supergrid and the global grid

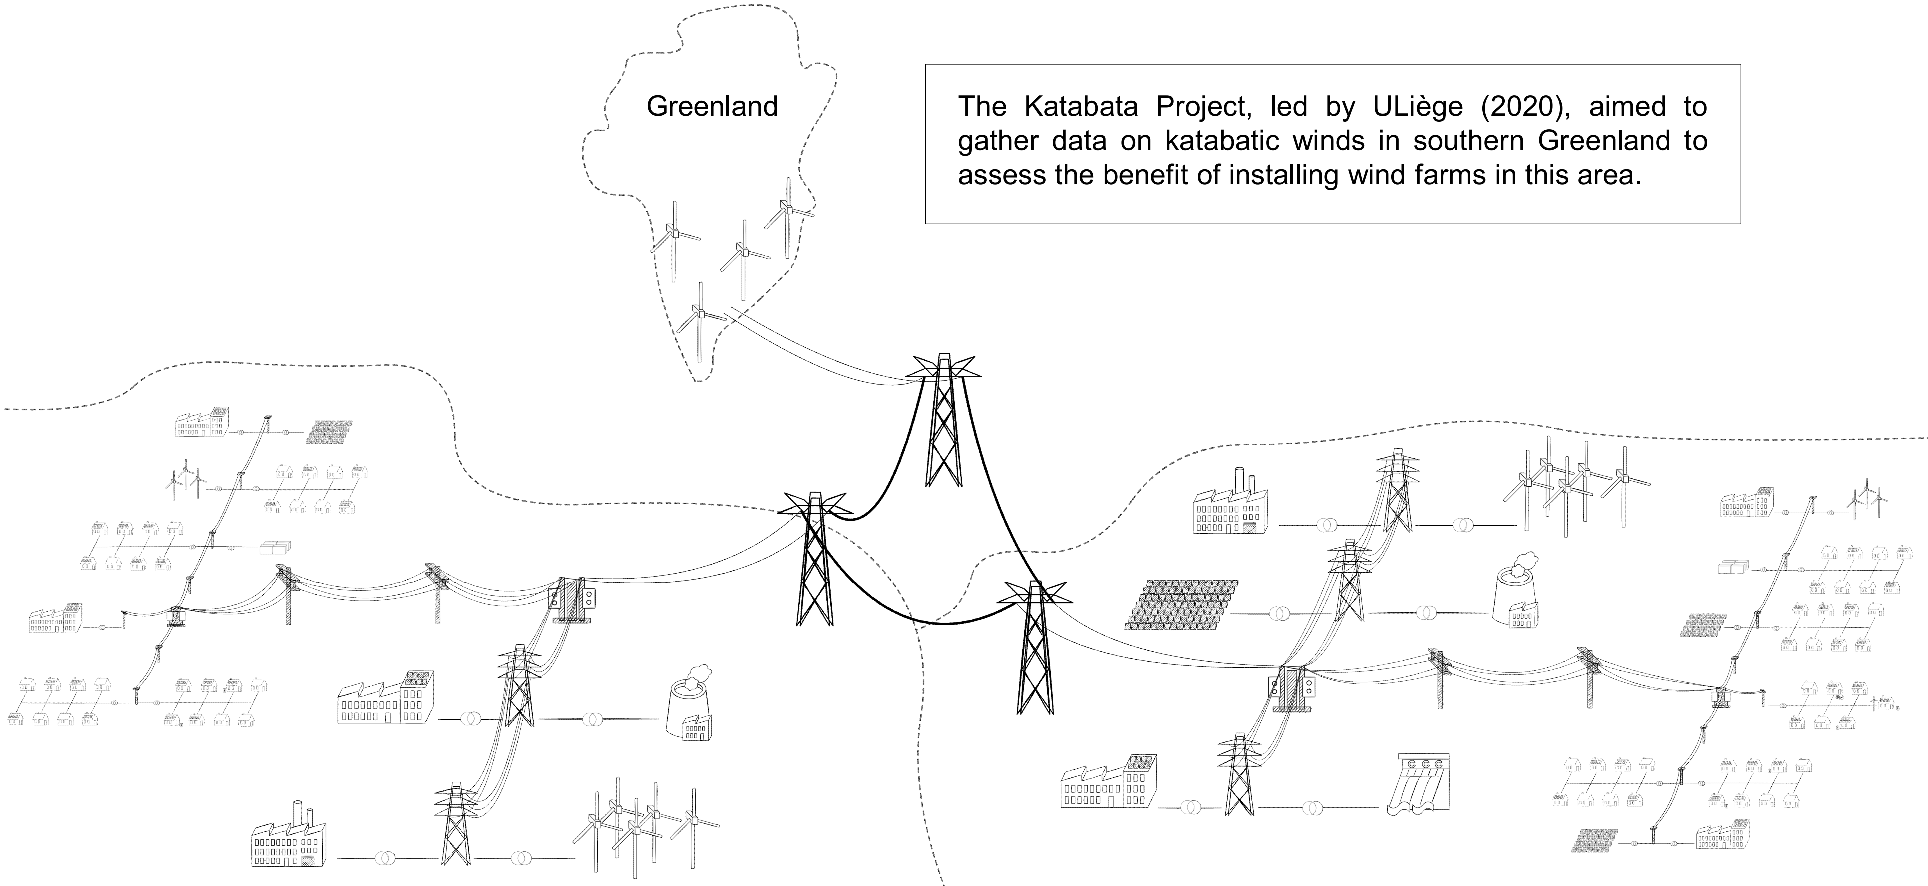

The Katabata Project, led by ULiège (2020), aimed to gather data on katabatic winds in southern Greenland to assess the benefit of installing wind farms in this area.

More about Katabata project in: Ernst, D., Fettweis, X., Fonder, M., & Louis, J. (2020). **Extreme engineering for fighting climate change and the Katabata project**. https://hdl.handle.net/2268/251827

Decision-making strategies and supergrid

The supergrid is a Direct Current (DC) network that overlays existing Alternating Current (AC) networks. It is still in its early stages. Currently, TSOs often regard them simply as (negative) loads in their standard decision-making tools. This approach works adequately with a few point-to-point DC links.

However, this will change as renewable energy resources become more prevalent, creating strong business cases for extensive DC links. These long links can help smooth out the fluctuations of renewable energy production and enable the harvesting of renewable resources in remote areas with abundant sun and wind.

As a result, decision-making challenges will emerge, which are crucial for the energy transition. This is especially true if the world comes together to pursue the ambitious goal of building a **global grid**, the ultimate step towards decarbonising our societies.

The definition of the global grid

A global grid is an electrical network that spans the entire planet, connecting the world's electricity consumers and producers. Its backbone would be composed of very long High Voltage Direct Current (HVDC) links. **This extensive network could drive renewable electricity costs significantly down, potentially putting fossil fuels out of business.**

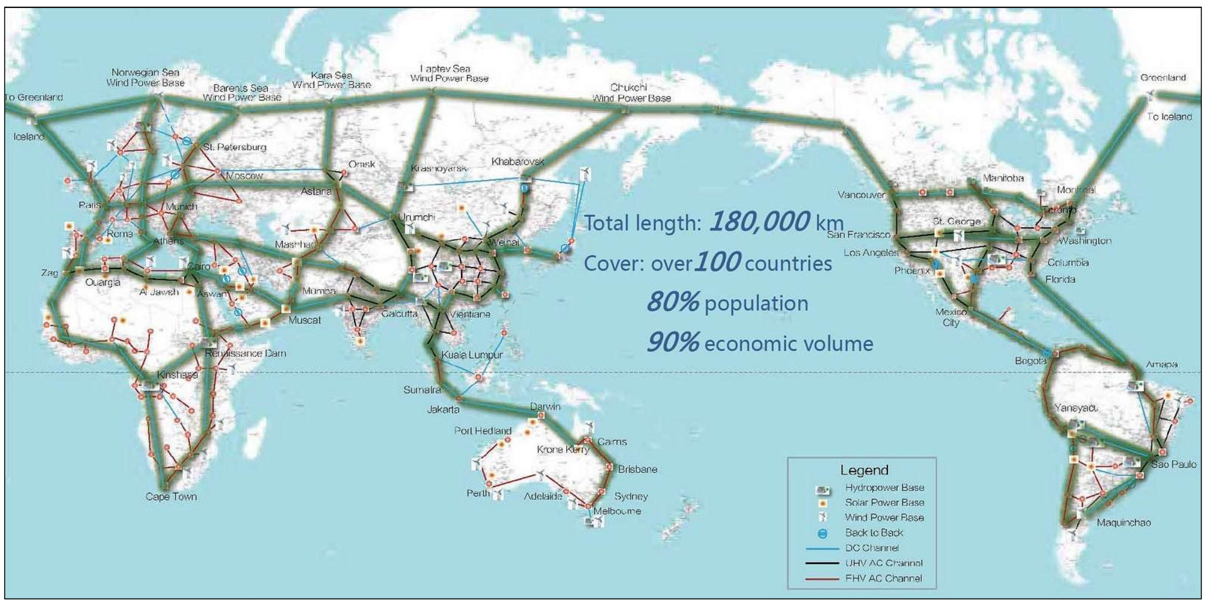

A mapped prototype of the Global Energy Interconnection Backbone Grid.

The planification problem behind the global grid

The primary reason why the different countries of the world have not yet come together for building the global grid is the lack of an acceptable solution to the **“right planning problem”** for its construction.

As a preliminary outline of the main features of this problem, the solution space is a set of **sequences of investments in transmission lines** such that:

(i) each new investment should ensure that all directly involved parties benefit from it;
(ii) no party should have the incentive to unilaterally exploit the global grid infrastructure for its own gain at the expenses of others.

Addressing this problem and finding its solution present fascinating challenges for machine learning researchers!

More information in: Muñoz, J. C., Sauma, E., Muñoz, F. D., & Moreno, R. (2023). **Analysis of generation investments under price controls in cross-border trade of electricity**. *Energy Economics*, 123, 106722. https://doi.org/10.1016/j.eneco.2023.106722

Coordination between the four layers

The physical coupling between different layers of the network system is becoming very significant. There is **a need for decision-making strategies that operate in a coordinated manner**.

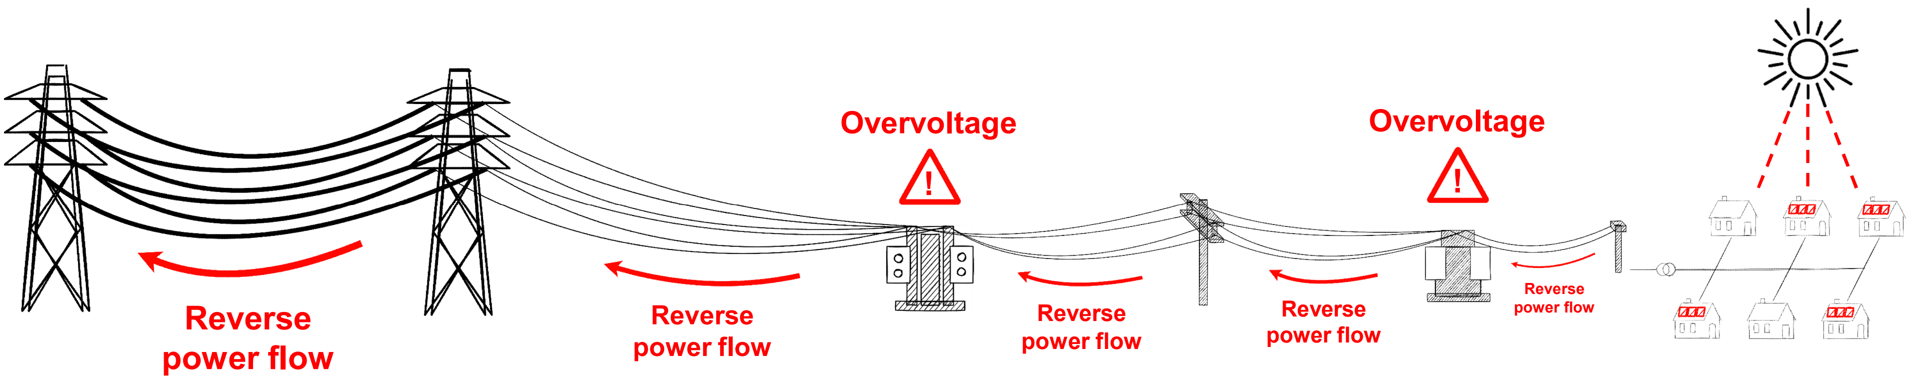

Example: When there is abundant solar energy, the surplus solar energy is fed into the electrical grid, potentially causing reverse power flow across several or even all layers of the network. This can lead to voltage control problems in several layers of the power system.

Are power systems really limited to four physical layers?

Yes, if you limit your power system to the physical network through which electrical energy is transmitted.

No, if you take a broader perspective that includes other physical layers transmitting energy in the form of chemical energy as part of the power system.

In this broader view, known as **sector coupling**, energy-rich molecules (e.g., CH_4, H_2, etc.) are increasingly produced from decarbonised electricity. This approach is gaining in popularity as it integrates different energy forms into a cohesive system.

An example of sector coupling

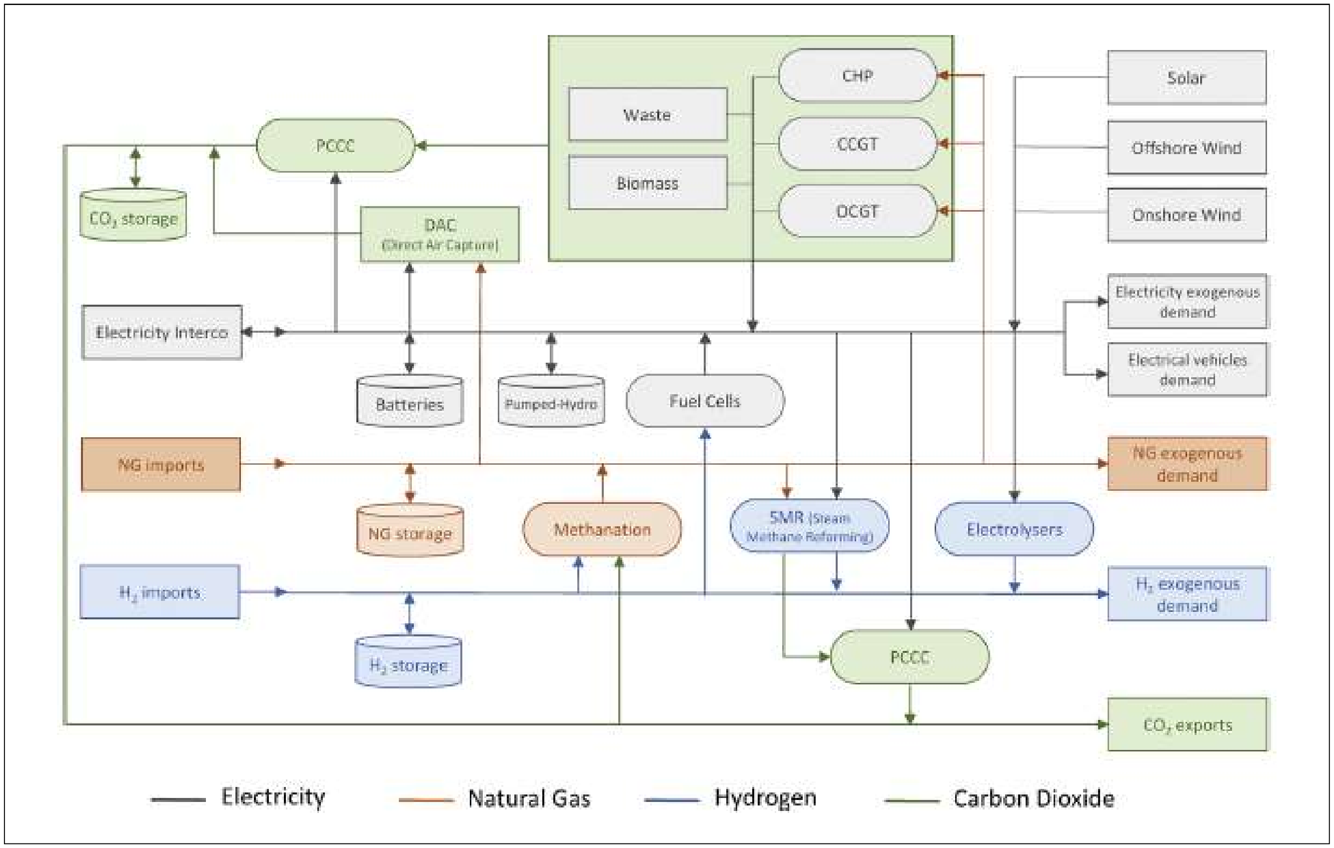

Picture taken from: Berger, M., Radu, D.-C., Fonteneau, R., Deschuyteneer, T., Detienne, G., & Ernst, D. (2020). **The role of power-to-gas and carbon capture technologies in cross-sector decarbonisation strategies**. *Electric Power Systems Research, 180*. https://hdl.handle.net/2268/235110

Two relevant challenges for ML in sector coupling (1/2)

Challenge 1:

The emergence of new devices for generating and storing molecules (e.g., CO_2, CH_4, H_2, etc.) with sector coupling will lead to complex sequential decision-making strategies with **long-time horizons**.

Additionally, some of these molecules are used to smooth out intraseasonal fluctuations in renewable energy.

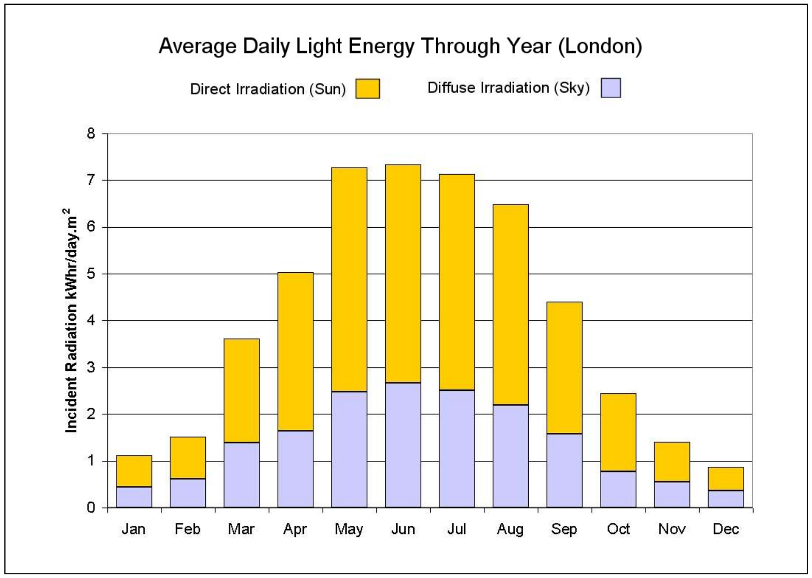

Two relevant challenges for ML in sector coupling (2/2)

Challenge 2:

Significant investments in infrastructure, such as CO_2 and H_2 transport networks, are required to benefit from sector coupling. Current investment planning techniques often solve large optimisation problems under the assumption of perfect foresight. However, this approach falls short for policy makers for two main reasons:

(i) There are too many uncertainties in the energy sector, making the assumption of perfect foresight unrealistic.

(ii) Implementing solutions is often complicated by issues such as Not In My BackYard (NIMBY) concerns and geopolitical factors. Policy makers need flexibility in these proposed solutions.

Addressing these challenges requires moving beyond "classical" deterministic or stochastic optimisation problems. Two research avenues for ML researchers include developing decision-making techniques related to **no-regret decisions** and identifying necessary conditions for **ϵ-optimality**.

The emerging European hydrogen backbone

Can you imagine the ML challenges related to the computation of the optimal European hydrogen backbone?

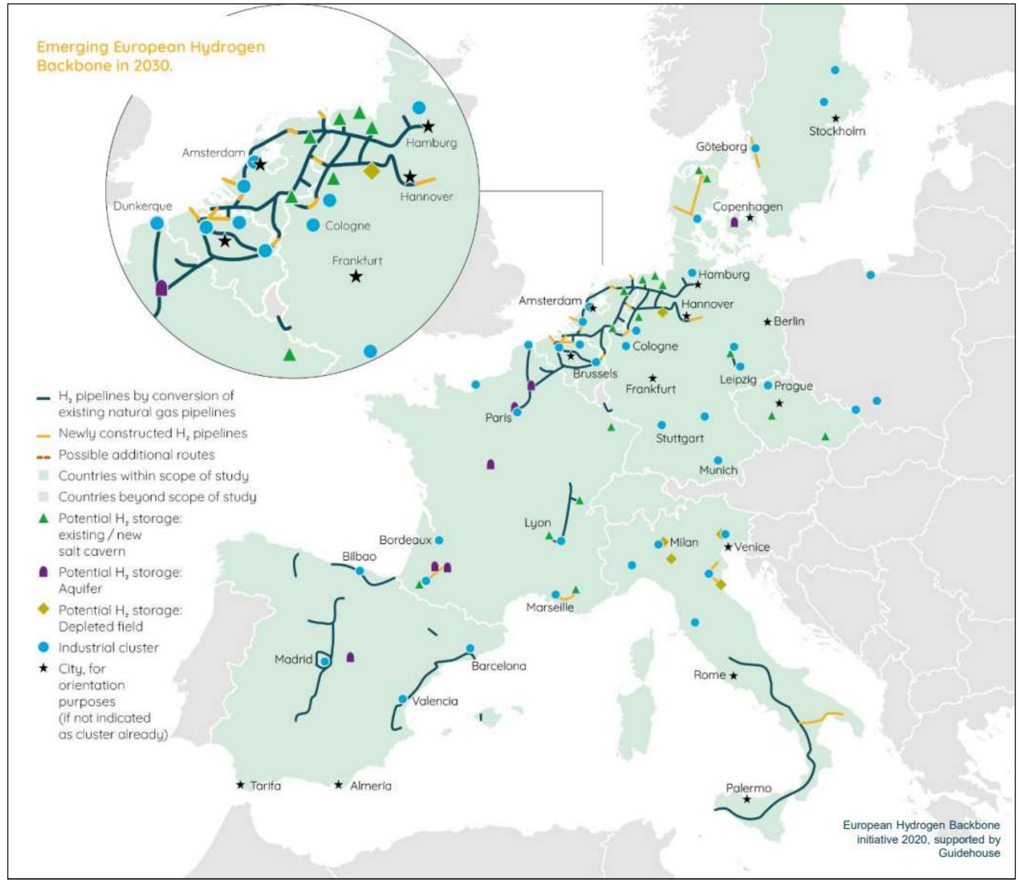

No-regret decision

A no-regret decision is one that **will not lead to regrets in the future.**

<u>Example:</u>

Consider you have three non-zero probability scenarios (A, B, and C) representing the uncertainty in parameters influencing the planning of your energy system.

Required installed capacities of electrolysers for each scenario:

Scenario A: 10 GW,
Scenario B: 8 GW,
Scenario C: 12 GW.

If you choose to install 8 GW, this is a no-regret decision because it is optimal for Scenario B and is also below the optimal capacities computed for Scenarios A and C.

Necessary conditions for ϵ-suboptimality

Necessary conditions for ϵ-optimality are conditions that any ϵ-optimal solution must satisfy. These can be valuable for guiding (near) optimal investment decisions.

Example:

(i) Suppose your optimisation problem aims to compute the optimal installed capacity of electrolysers (in GW) and H_2 storage capacity (in GWh).

(ii) Assume that your 0.05-optimal space includes three elements: (10 GW, 1200 GWh), (8 GW, 600 GWh), and (12 GW, 2400 GWh).

Therefore, installing more than 8 GW of electrolysers or more than 600 GWh of H_2 storage are two necessary conditions for 0.05-optimality.

Three-dimensional representation of an ϵ- optimal space.

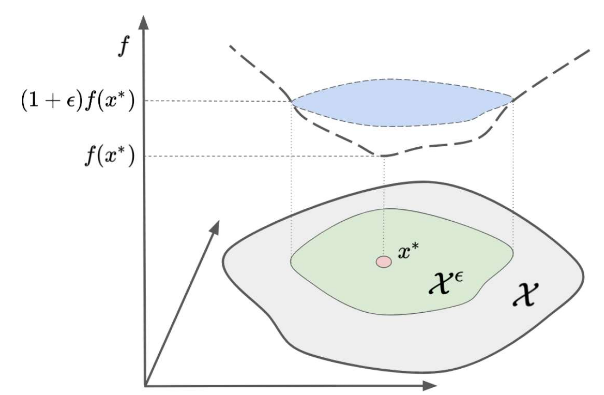

Picture taken from: Dubois, A., & Ernst, D. (2022). **Computing necessary conditions for near-optimality in capacity expansion planning problems**. *Electric Power Systems Research*, Volume 211. https://hdl.handle.net/2268/263897

What a typical ML researcher may say at the end of this talk

All these problems are fascinating, but I am an ML scientist working on drones. Are there any challenges in power systems where I could apply my skills more directly?

Of course! Let me tell you a story…

In August 2003, a major blackout occurred in North America, initially leaving 50 million people without electricity. The event resulted in 61,800 MW of load being cut in the USA and Canada, with an estimated cost in the USA ranging from $4 to $10 billion. The restoration time varied from a few hours to up to 4 days.

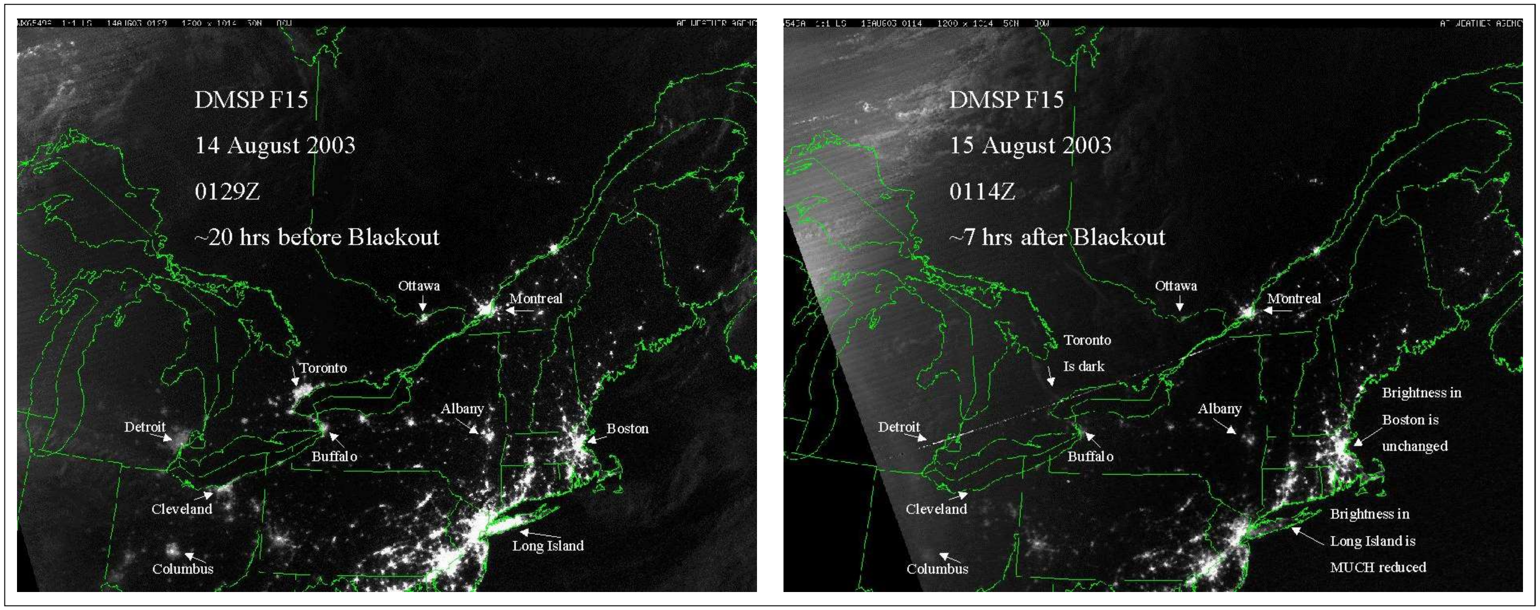

Causes of the 2003 North America blackout

The official report on the blackout indicates that a generating plant in Eastlake, Ohio, a suburb northeast of Cleveland, went offline due to high demand. This put a strain on high-voltage power lines in Walton Hills, Ohio, a southeast suburb of Cleveland, which later went out of service after coming into contact with **“overgrown trees”**.

This trip caused the load to shift to other transmission lines, which couldn’t handle the increased demand, leading to their breakers tripping. As multiple trips occurred, many generators suddenly lost parts of their loads, causing them to accelerate out of phase with the grid at different rates. To prevent damage, the generators tripped offline. This cascading effect ultimately forced the shutdown of more than 100 power plants.

Building autonomous drones for the monitoring of tree growth

Overgrown trees pose a significant threat to power system security. You could focus on designing autonomous drones to identify areas where vegetation management is needed. Even better, you could develop drones that autonomously trim the problematic trees.

Take home messages

Power systems present numerous complex ML challenges.

With the increasing use of batteries and other energy storage devices, the temporal coupling of decisions in power system operations is becoming more significant. This makes power systems an excellent playground for RL researchers!

There is a wide range of decision-making problems, some of which are specific to local aspects of the power system. Because the machine learning community often focuses on more mainstream problems, such as autonomous driving, researchers in this field may find themselves more isolated and receive fewer citations. However, addressing decision-making problems for power systems plays a vital role in accelerating the decarbonization of our societies, and is thus a **good opportunity to contribute to a more sustainable environment for current and future generations**.

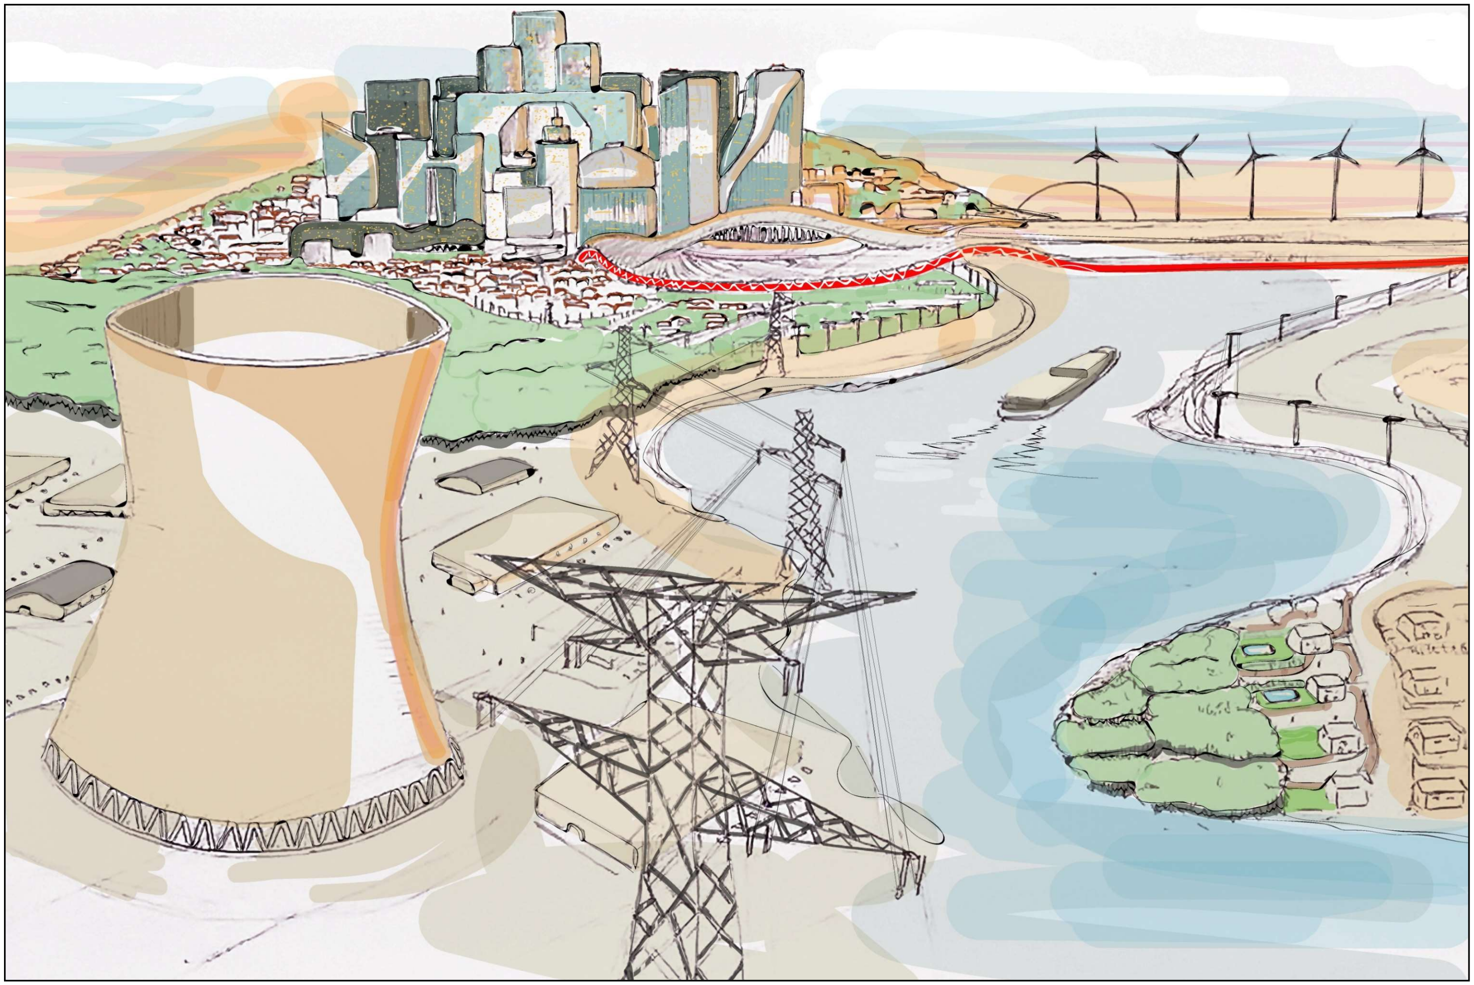

Illustration by Roberto Perry.

Acknowledgments

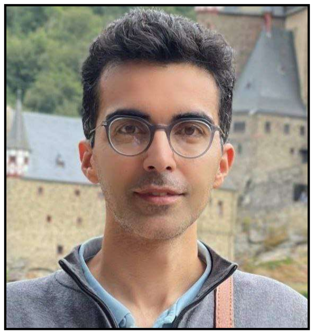

Alireza Bahmanyar
Haulogy
Senior Electric Power
Systems Expert

Amina Benzerga
ULiège
PhD Student

François Cubelier
ULiège
PhD Student

Benjamin Donnot
RTE
Project/Research Officer

Balthazar Donon
ULiège & RTE
Postdoctoral Researcher

Laurine Duchesne
Haulogy
Head of Consultancy Services

Raphaël Fonteneau
ULiège
Senior Researcher

Thibaut Techy
ULiège
Energy Expert

Research in the lab – the first layer

Aittahar, S., Manuel de Villena Millan, M., Derval, G., Castronovo, M., Boukas, I., Gemine, Q., & Ernst, D. (2023). **Optimal control of renewable energy communities with controllable Assets**. *Frontiers in Energy Research*. https://hdl.handle.net/2268/264828

Manuel De Villena, M., Aittahar, S., Mathieu, S., Boukas, I., Vermeulen, E., & Ernst, D. (2022). **Financial optimization of renewable energy communities through optimal allocation of locally generated electricity**. *IEEE Access, 10*, 77571 - 77586. https://hdl.handle.net/2268/296721

Bolland, A., Boukas, I., Berger, M., & Ernst, D. (January 2022). **Jointly learning environments and control policies with projected stochastic gradient ascent**. *Journal of Artificial Intelligence Research*, 73, 117-171. https://hdl.handle.net/2268/247965

François-Lavet, V., Taralla, D., Ernst, D. & Fonteneau, R. (2016). **Deep reinforcement learning solutions for energy microgrids management**. *European Workshop on Reinforcement Learning (EWRL 2016)*. https://hdl.handle.net/2268/203831

François-Lavet, V., Gemine, Q., Ernst, D., & Fonteneau, R. (2016). **Towards the minimization of the levelized energy costs of microgrids using both long-term and short-term storage devices**. *Smart Grid: Networking, Data Management, and Business Models* (pp. 295-319). CRC Press. https://hdl.handle.net/2268/195799

Research in the lab – the second layer

Vassallo, M, Bahmanyar, A., Duchesne, L., Leerschool, A., Gerard, S., Wehenkel, T. & Ernst D. (2024). **A systematic procedure for topological path identification with raw data transformation in electrical distribution networks**. *7th International Conference on Energy, Electrical and Power Engineering (CEEPE 2024)*. https://hdl.handle.net/2268/311815

Benzerga, A., Bahmanyar, A., Mathieu, S. & Ernst D. (2023). **Smart meter data analytics case study: Identification of LV distribution network topology to design optimal planning solutions**. *Applications of Big Data and Artificial Intelligence in Smart Energy Systems,* p. 161-192, River Publishers. https://hdl.handle.net/2268/304363

Benzerga, A., Mathieu, S., Bahmanyar, A. & Ernst D (2021). **Probabilistic capacity assessment for three-phase low-voltage distribution networks.** *2021 IEEE 15th International Conference on Compatibility, Power Electronics and Power Engineering (CPE-POWERENG).* https://hdl.handle.net/2268/262549

Henry, R., & Ernst, D. (2021). **GYM-ANM: Reinforcement learning environments for active network management tasks in electricity distribution systems**. *Energy and AI, 5*. https://hdl.handle.net/2268/257972

Mathieu, S., Ernst, D., & Gemine, Q. (2020). **Short-term active distribution network operation under uncertainty**. *Proceedings of the16th International Conference on Probabilistic Methods Applied to Power Systems (PMAPS 2020).* https://hdl.handle.net/2268/247385

Gemine, Q., Ernst, D., & Cornélusse, B. (2016). **Active network management for electrical distribution systems: Problem formulation, benchmark, and approximate solution**. *Optimization and Engineering, 18* (3), 587-629. https://hdl.handle.net/2268/173942

Research in the lab – the third layer

Donon, B., Cubelier, F., Karangelos, E., Wehenkel, L., Crochepierre, L., Pache, C., Saludjian, L., & Panciatici, P. (2024). **Topology-aware reinforcement learning for tertiary voltage control**. *Proceedings of the XXIII Power Systems Computation Conference*, Paris, France. https://hdl.handle.net/2268/315490

Bahmanyar, A., Duchesne, L., Dekeyne, G., Goubet, C., Gemine, Q., Panciatici, P., & Ernst, D. (2024). **Identification of the critical cluster of generators by during fault angle trajectory estimation for transient stability analysis**. *Electric Power Systems Research, 234*. https://hdl.handle.net/2268/319512

Radu, D., Berger, M., Dubois, A., Fonteneau, R., Pandžić, H., Dvorkin, Y., Louveaux, Q. & Ernst, D. (2022). **Assessing the impact of offshore wind siting strategies on the design of the European power system**. *Applied energy*, Volume 305, Elsevier. https://hdl.handle.net/2268/251079

Bahmanyar, A., Ernst, D., Vanaubel, Y., Gemine, Q., Pache, C., & Panciatici, P. (2021). **Extended equal area criterion revisited: A direct method for fast transient stability analysis**. *Energies, 14* (21). https://hdl.handle.net/2268/263320

Duchesne, L., Karangelos, E., & Wehenkel, L. (2020). **Recent developments in machine learning for energy systems reliability management**. *Proceedings of the IEEE*. https://hdl.handle.net/2268/246570

Research in the lab – the fourth layer

Dubois, A., & Ernst, D. (2022). **Computing necessary conditions for near-optimality in capacity expansion planning problems**. *Electric Power Systems Research*, Volume 211, Elsevier. https://hdl.handle.net/2268/263897

Radu, D.-C., Berger, M., Fonteneau, R., Hardy, S., Fettweis, X., Le Du, M., Panciatici, P., Balea, L., & Ernst, D. (2019). **Complementarity assessment of South Greenland katabatic flows and West Europe wind regimes**. *Energy*, 175, 393-401. https://hdl.handle.net/2268/230016

Yu, J., Bakic, K., Kumar, A., Iliceto, A., Beleke Tabu, L., Ruaud, J. L., Fan, J., Cova, B., Li, H., Ernst, D., Fonteneau, R., Theku, M., Sanchis, G., Chamollet, M., Le Du, M., Zhang, Y., Chatzivasileiadis, S., Radu, D.-C., Berger, M., ... Ranjbar, M. (2019). **Global electricity network - Feasibility study**. https://hdl.handle.net/2268/239969

Chatzivasileiadis, S., Ernst, D., & Andersson, G. (2013). **The global grid**. *Renewable Energy, 57*, 372-383. https://hdl.handle.net/2268/144423

Sarlette, A., Dai, J., Phulpin, Y., & Ernst, D. (2012). **Cooperative frequency control with a multi-terminal high-voltage DC network**. *Automatica, 48* (12), 3128–3134. https://hdl.handle.net/2268/134354

Research in the lab – sector coupling

Mbenoun, J., Benzerga, A., Miftari, B., Detienne, G., Deschuyteneer, T., Vazquez, J., Derval, G., & Ernst, D. (2024). **Integration of offshore energy into national energy system: A case study on Belgium**. *ORBi-University of Liège*. https://hdl.handle.net/2268/314586

Dachet, V., Lokotar, I., & Ernst, D. (2024). **Remote renewable energy hubs: A leadership opportunity that Europe must seize**. *Confrontations Europe*. https://hdl.handle.net/2268/318392

Dachet, V., Benzerga, A., Coppitters, D., Contino, F., Fonteneau, R., & Ernst, D. (2024). **Towards CO2 valorization in a multi remote renewable energy hub framework with uncertainty quantification**. *Journal of Environmental Management,* 363, 121262. https://hdl.handle.net/2268/309679

Miftari, B., Berger, M., Derval, G., Louveaux, Q., & Ernst, D. (2023). **GBOML: A structure-exploiting optimization modelling language in python**. *Optimization Methods and Software*. https://hdl.handle.net/2268/296930

Berger, M., Radu, D.-C., Fonteneau, R., Deschuyteneer, T., Detienne, G., & Ernst, D. (2020). **The role of power-to-gas and carbon capture technologies in cross-sector decarbonisation strategies**. *Electric Power Systems Research*, 180. https://hdl.handle.net/2268/235110