Knowledge assisted AI for Real-World Network Infrastructure

Webinar (University of Amsterdam, IRT SystemX & TU Delft)

November 29th 2024

Herke van Hoof, Milad Leyli-Abadi, Joost Ellerbroek

AI4REALNET has received funding from European Union's Horizon Europe Research and Innovation programme under the Grant Agreement No 101119527

Introduction

November 29th, Herke van Hoof

AI4REALNET has received funding from European Union's Horizon Europe Research and Innovation programme under the Grant Agreement No 101119527

Welcome!

- Knowledge assisted AI for Real-World Network Infrastructure
- Presenters:
  - Milad Leyli-Abadi
  - Joost Ellerbroek
  - Herke van Hoof
- Organization:
  - Bianca Silva
  - Milad Leyli-Abadi
  - Herke van Hoof

Welcome!

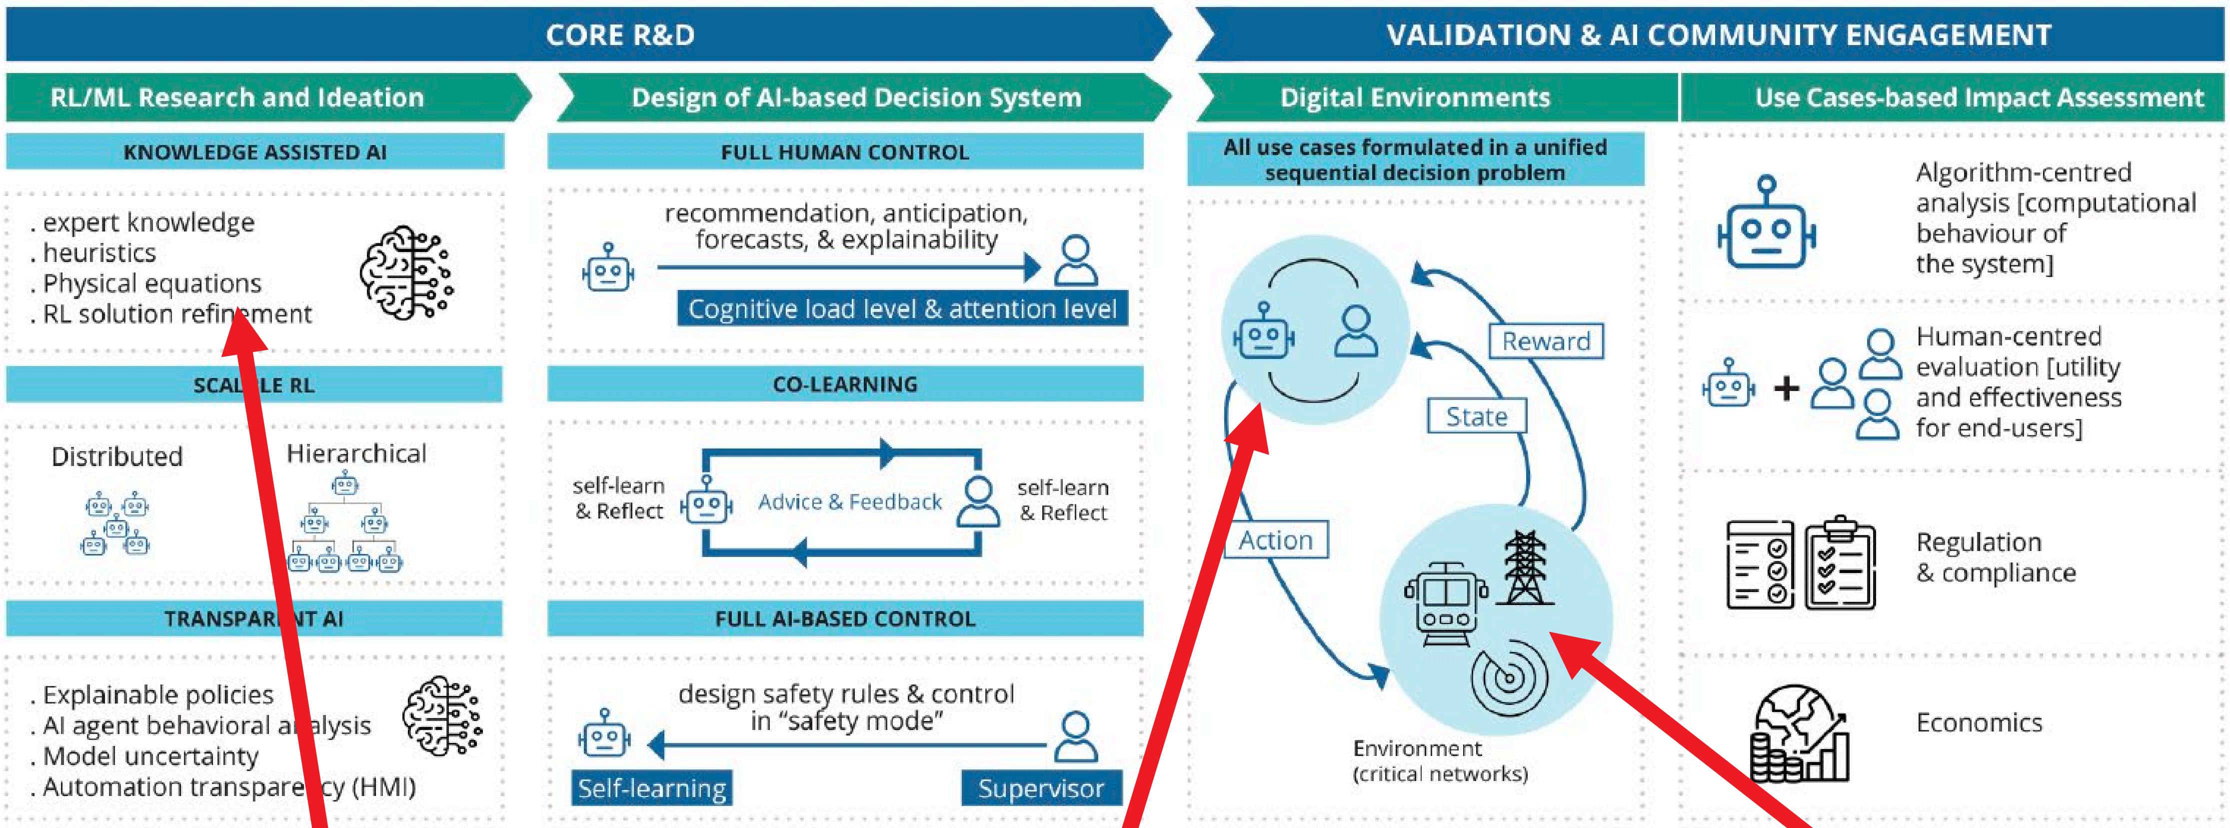

Knowledge-assisted AI

Decision making

Real-world infrastructure networks

Welcome!

- Today's schedule:
  - Introduction
  - Decision making in AI4REALNET
  - Knowledge-assisted AI: Definition, overview, and state-of-the-art
  - Case study: Air traffic control
  - Case study: Power network control
  - Questions & Discussion

Knowledge assisted AI in AI4REALNET

Context and usecases

November 29th, Milad Leyli-Abadi

AI4REALNET has received funding from European Union's Horizon Europe Research and Innovation programme under the Grant Agreement No 101119527

AI4REALNET project objectives and scope

- Develop next generation of **decision-making methods** powered by supervised and reinforcement learning for **critical infrastructures**
- Ensure **trustworthiness in**
  - AI-assisted human control
  - Human-AI co-learning
  - Autonomous AI
- **Boost the development and validation of novel AI algorithms** via 3 existing open-source AI-friendly digital environments

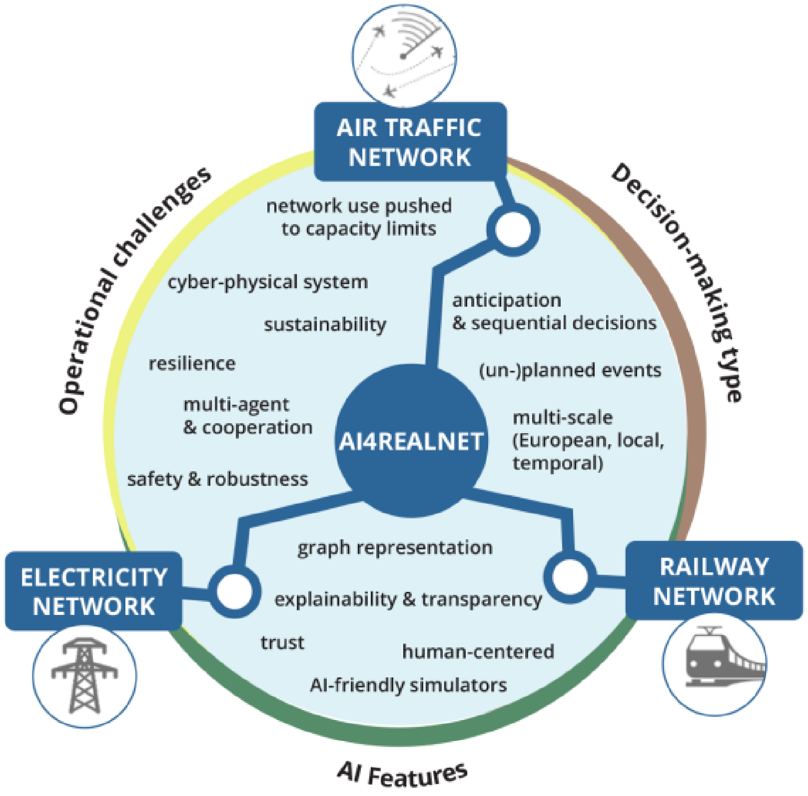

Project use cases: focus on critical infrastructures

Rte

TenneT

DB

SBB CFF FFS

NAV PORTUGAL

TU Delft

AI4 REALNET — UC1 POWER GRID

AI assistant supporting human operators' decision-making in managing power grid congestion

AI ROLE

Provide a human operator with remedial action recommendations aimed at safely managing overloads on the electrical lines and easing the workload of the human operator.

7 AFFORDABLE AND CLEAN ENERGY
13 CLIMATE ACTION

FULL HUMAN CONTROL

recommendation, anticipation, forecasts, & explainability

Cognitive load level & attention level

AI4 REALNET — UC2 POWER GRID

Sim2Real, transfer AI-assistant from simulation to real-world operation

AI ROLE

Provide a human operator with remedial action recommendations, considering a transfer from training (digital) to real-world environments.

7 AFFORDABLE AND CLEAN ENERGY
13 CLIMATE ACTION

FULL HUMAN CONTROL

recommendation, anticipation, forecasts, & explainability

Cognitive load level & attention level

AI4 REALNET — UC1 RAILWAY

Automated re-scheduling in railway operations

AI ROLE

The re-scheduling task is performed in a highly automated manner by an AI-based re-scheduling system. It observes the real-time state of all the trains and tracks in the control area of interest and automatically detects the need to intervene, decides on an intervention, and executes this intervention.

9 INDUSTRY, INNOVATION AND INFRASTRUCTURE
11 SUSTAINABLE CITIES AND COMMUNITIES
13 CLIMATE ACTION

FULL AI-BASED CONTROL

design safety rules & control in "safety mode"

Self-learning

Supervisor

AI4 REALNET — UC2 RAILWAY

AI-assisted human re-scheduling in railway operations

AI ROLE

Assist the human dispatcher in railway operations in re-scheduling train runs to fulfil all offered services and minimize delays for the customer.

9 INDUSTRY, INNOVATION AND INFRASTRUCTURE
11 SUSTAINABLE CITIES AND COMMUNITIES
13 CLIMATE ACTION

JOINT DECISION MAKING

self-learn & Reflect

Advice & Feedback

self-learn & Reflect

AI4 REALNET — UC1 ATM

Airspace sectorization assistant

AI ROLE

Partially and fully automate the sectorization process to assist or replace the staff manager in deciding when and how to split and merge sectors to balance the workload of tactical ATCOs.

9 INDUSTRY, INNOVATION AND INFRASTRUCTURE
11 SUSTAINABLE CITIES AND COMMUNITIES
13 CLIMATE ACTION

FULL HUMAN CONTROL

recommendation, anticipation, forecasts, & explainability

Cognitive load level & attention level

AI4 REALNET — UC2 ATM

Flow and airspace management assistant

AI ROLE

Provide advice to air traffic controller about deviations with better sector capacity adherence and performance measured by an indicator of environmental area. Also consider the need to review the sectorization plan due to the activation of military areas and required trajectory efficient deviations.

9 INDUSTRY, INNOVATION AND INFRASTRUCTURE
11 SUSTAINABLE CITIES AND COMMUNITIES
13 CLIMATE ACTION

JOINT DECISION MAKING

self-learn & Reflect

Advice & Feedback

self-learn & Reflect

AI4REALNET conceptual framework

- The AI4REALNET conceptual framework defined based on an interdisciplinary approach by integrating diverse fields, such as **psychology** and **cognitive engineering**, with **AI**

- Four layers addressing
  - Decision environment and context
  - Human agent decision making
  - AI agent decision making
  - Trustworthiness assessment

Example: Human-AI co-learning

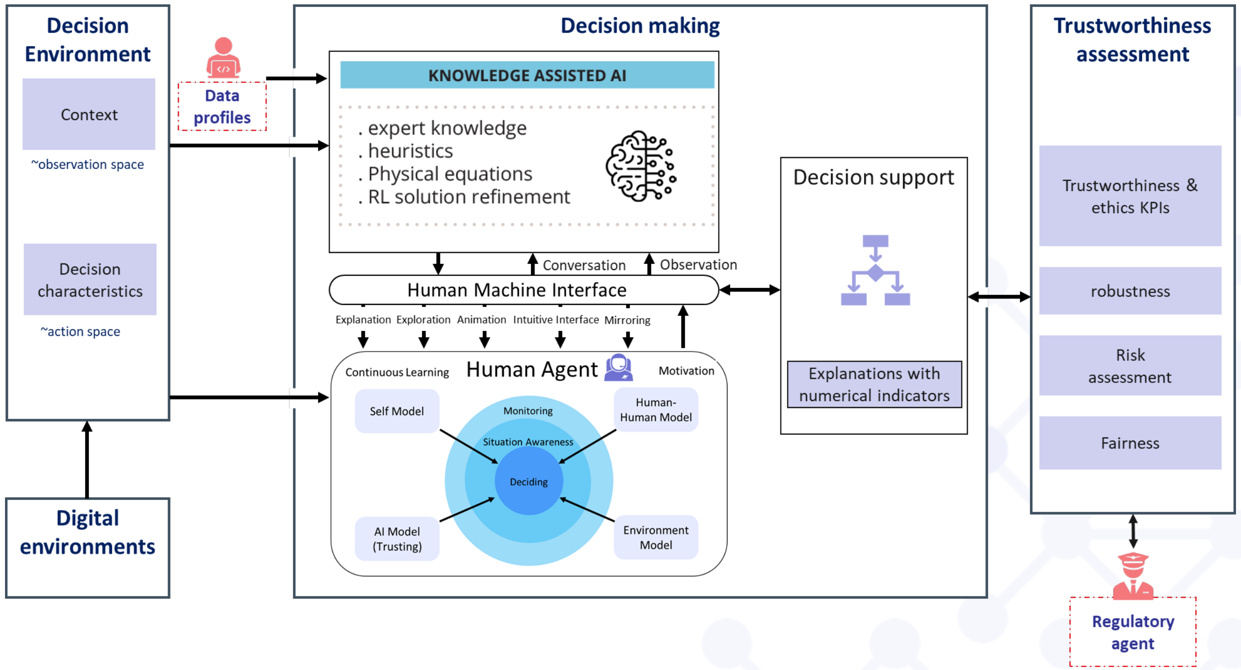

Characteristics of critical infrastructure domains that make KAI important

- **Complexity:**
  - Real-world dynamics often involve non-linear behaviors and rare events

  → Data alone might not capture these dynamics and prior knowledge consideration is crucial
- **Data scarcity**:
  - History data may be incomplete and noisy and might not fully represent future scenarios

  → It makes domain knowledge consideration crucial
- **Regulatory and safety requirements:**
  - Regulations and safety standards demand systems to align with predefined rules and guidelines

  → Knowledge-assisted AI can ensure compliance and simplify audits
- **Need for Interpretability:**
  - Operators and stakeholders must trust AI decisions, which necessitates explainability

  → Explicit knowledge makes it easier to trace and justify AI actions

Definition, overview, and SOTA

November 29th, Herke van Hoof

AI4REALNET has received funding from European Union's Horizon Europe Research and Innovation programme under the Grant Agreement No 101119527

Definition

- Knowledge assisted AI in AI4REALNET

  “[d]evelop AI technologies that can leverage the strength of both classical planning or optimisation heuristics, as well as ML techniques.”

- Informed Machine Learning [von Rueden et al., 2021]

  "learning from […] data and prior knowledge. The prior knowledge comes from an independent source, is given by formal representations, and is explicitly integrated into the machine learning pipeline"

- Neural-symbolic or hybrid systems [van Harmelen & ten Teije, 2019; Yu et al., 2023; Sarker et al., 2021]

  Various architectures for combining learning and reasoning or symbolic systems, including deliberative components inside a learning system

Overview

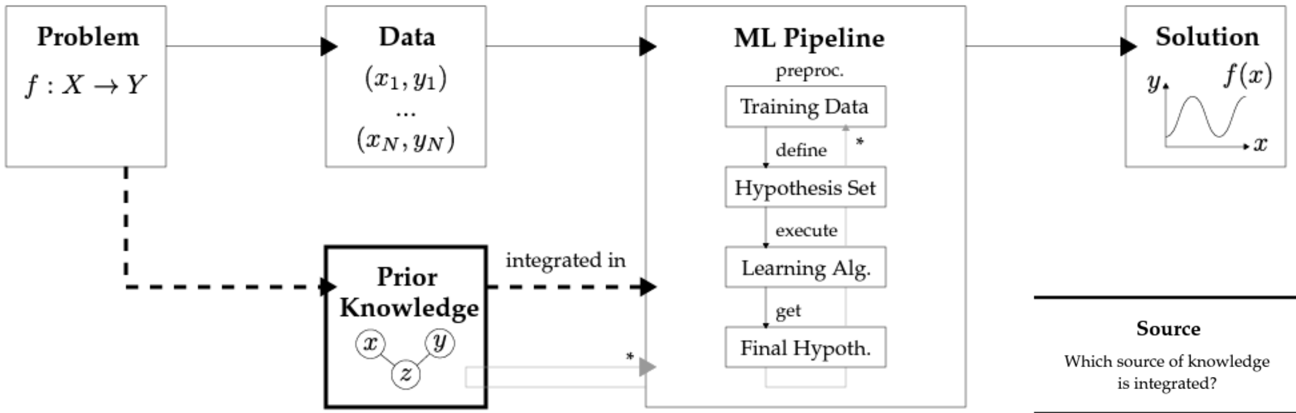

Von Rueden et al., 2021

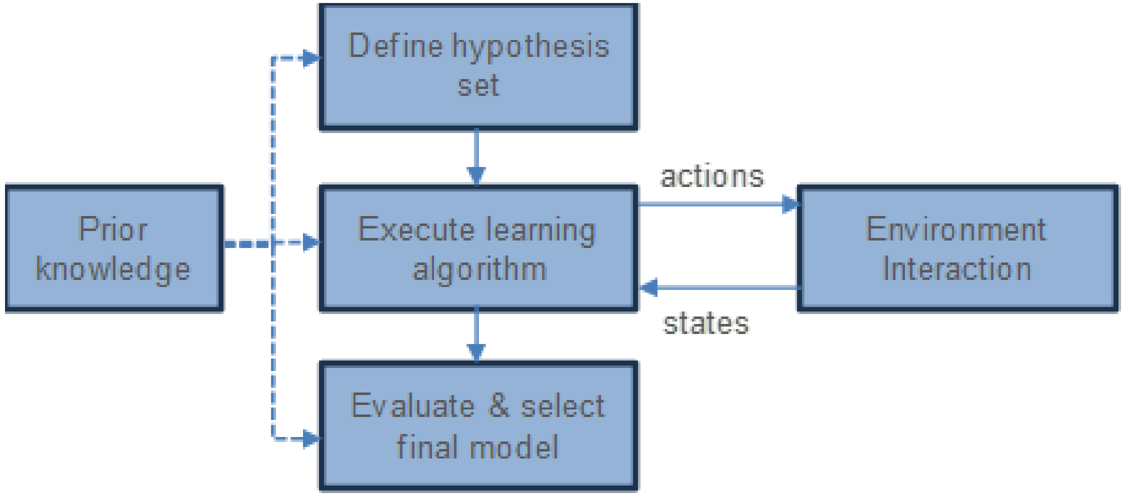

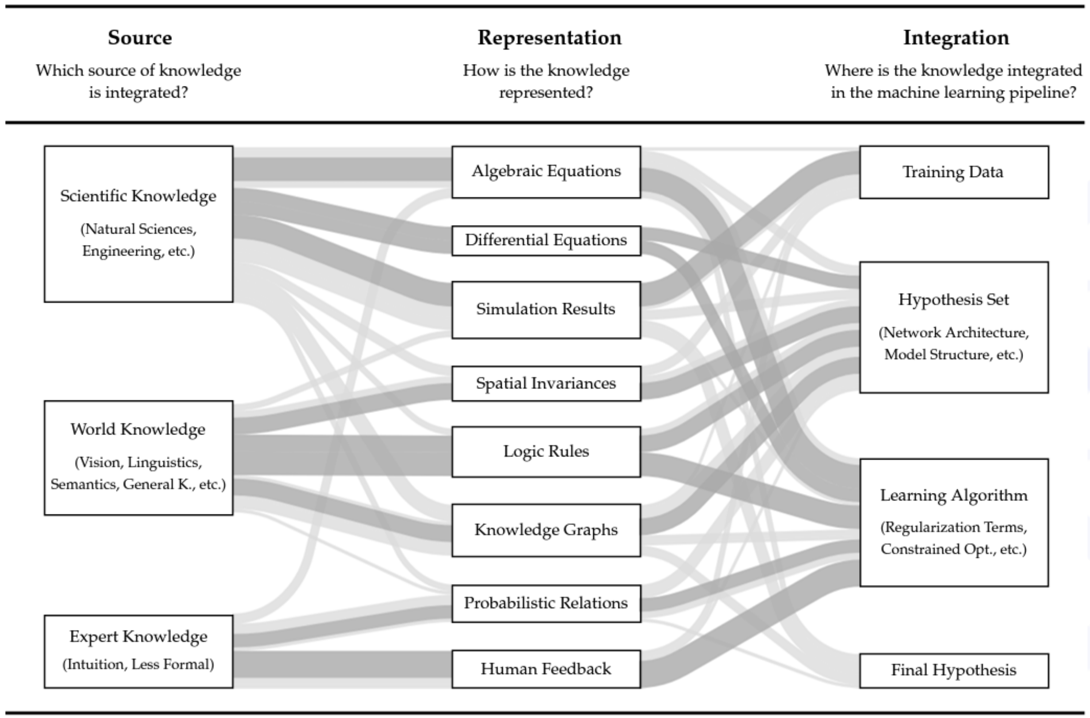

State of the art

- Increasing body of work in informed machine learning and neural-symbolic AI
- Most of this work in supervised learning (regression, classification)
- Infrastructure control requires *decision making, e.g.: reinforcement learning*
- Less work available – topic of today

State of the art

- Knowledge assisted methods for reinforcement learning fall predominantly in four categories:
    - Prior information about the desired system behavior
    - Prior information desired system states
    - Symbolic components within (neural network) models
    - High-level symbolic planning with low-level (neuronal) learning
- Alternative: Decision making based on informed prediction methods

State of the art: Information about policy

- Policy is a rule or (learned) function that decides which actions to take in which situation
- Learned policy bad at first: also generating data using a prior policy, speeds up learning (Zhao et al., 2020, 2022).
- Alternatively, encourage learned policy to be close to a known guiding policy (Dai et al., 2022)
- Use knowledge to exclude actions known to be bad or dangerous: shielding (e.g.: Al-Shiekh et al., 2018).
- Use information about the functional form of the policy, such as (geometric) invariances (van der Pol et al., 2020).

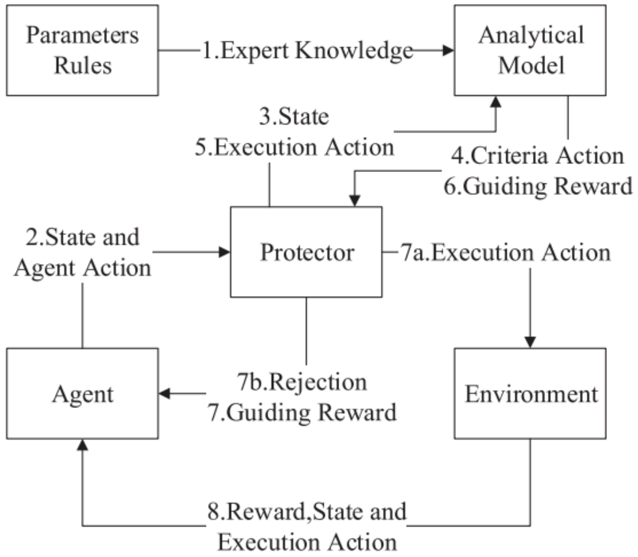

Figure: Zhao et al., 2020

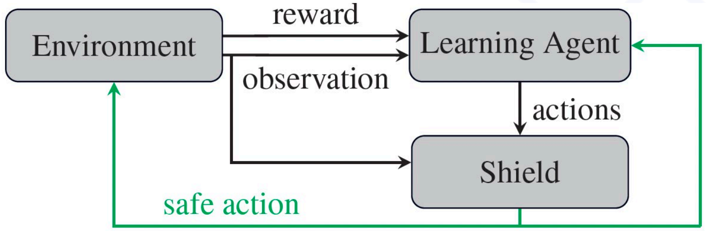

Figure: Al-Shiekh et al., 2018

State of the art: Information about value

- Value function represent whether system state is good or bad, helpful to evaluate actions
- Initially values bad when learning from scratch, instead refine a coarse value function obtained from optimizing an approximation (Wöhlke et al., 2022).
- Alternatively, shape reward function using knowledge or assumptions about the problem (Xie et al., 2024)
- Specify rewards using human feedback or preferences (e.g. Christiano et al., 2017)

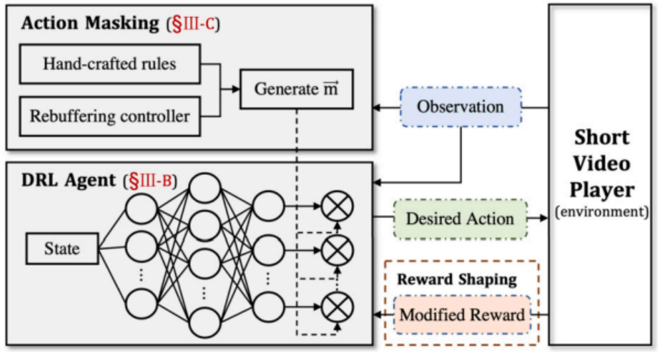

Figure: Xie et al., 2024

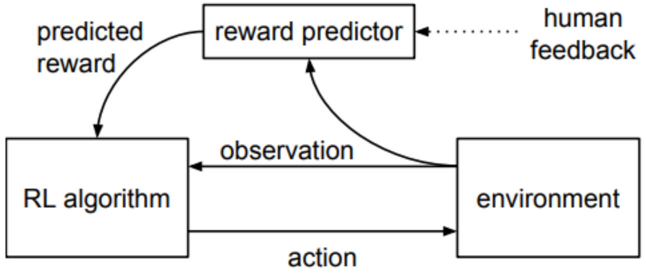

Figure: Christiano et al., 2017

State of the art: Symbolic components

- Often, general-purpose architecture (e.g.: neural network) trained as value function or policy
- Instead, include symbolic components (e.g., reasoning engine, optimizer, planner) into such architectures
- E.g. learn to extract symbolic representation, further processed by symbolic system (Garnelo et al., 2016, Garcez et al., 2018).
- Or generalize using known symbolic relations between inputs (Höpner et al., 2022).

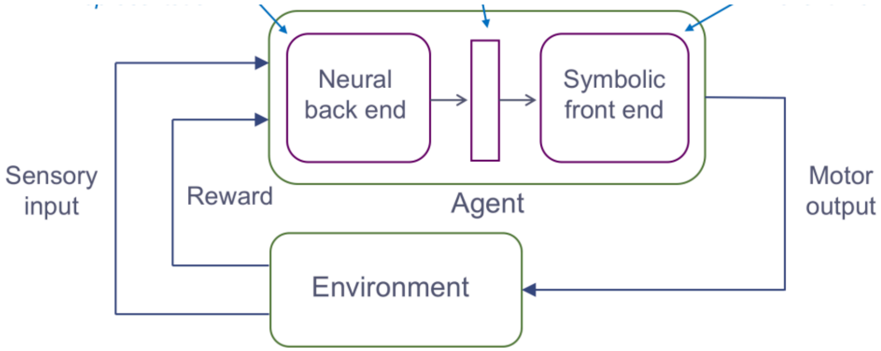

Figure: Garnelo et al., 2016

State of the art: Hierarchical

- Data-driven methods excel at processing large-volume sensory data, while planning or reasoning usually takes place at a more abstract level
- Use planning vs. learning components in different layers of *decision hierarchy* (e.g. Araki et al., 2021)
- Feed high-level task description into decision making architecture (e.g. Vaezipoor et al., 2021)
- Learn symbolic policies on top of pre-trained sensory-motor skills (e.g., Mitchener et al., 2022)

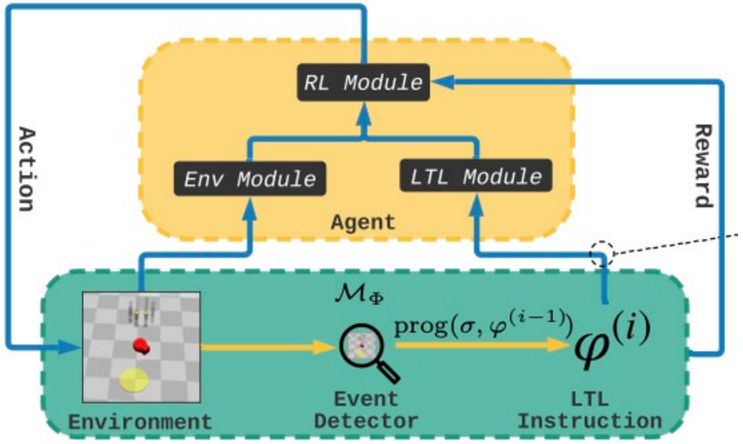

Figure: Vaezipoor et al., 2021

Types of knowledge in assisted RL methods

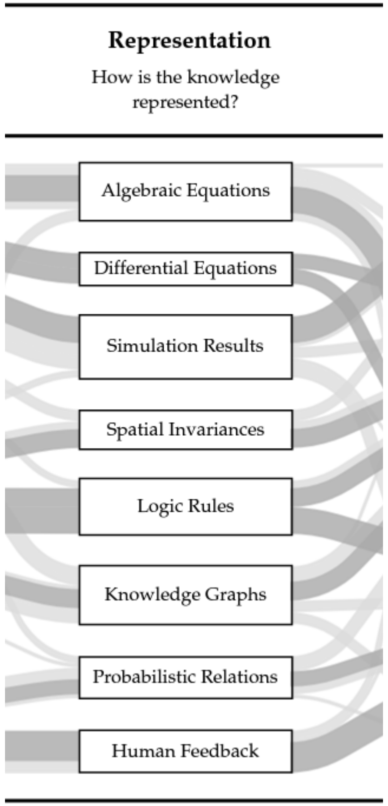

Critical in model based reinforcement learning (Moerland et al., 2023)

Critical in model based reinforcement learning (Moerland et al., 2023)

Critical in model based reinforcement learning (Moerland et al., 2023)

e.g. Araki et al., 2021, Vaezipoor et al., 2021.

e.g. Van der Pol et al., 2020

e.g. Höpner et al., 2022

Critical in model based reinforcement learning (Moerland et al., 2023)

e.g. Christiano et al., 2017

Von Rueden et al., 2021

Knowledge integration in assisted RL methods

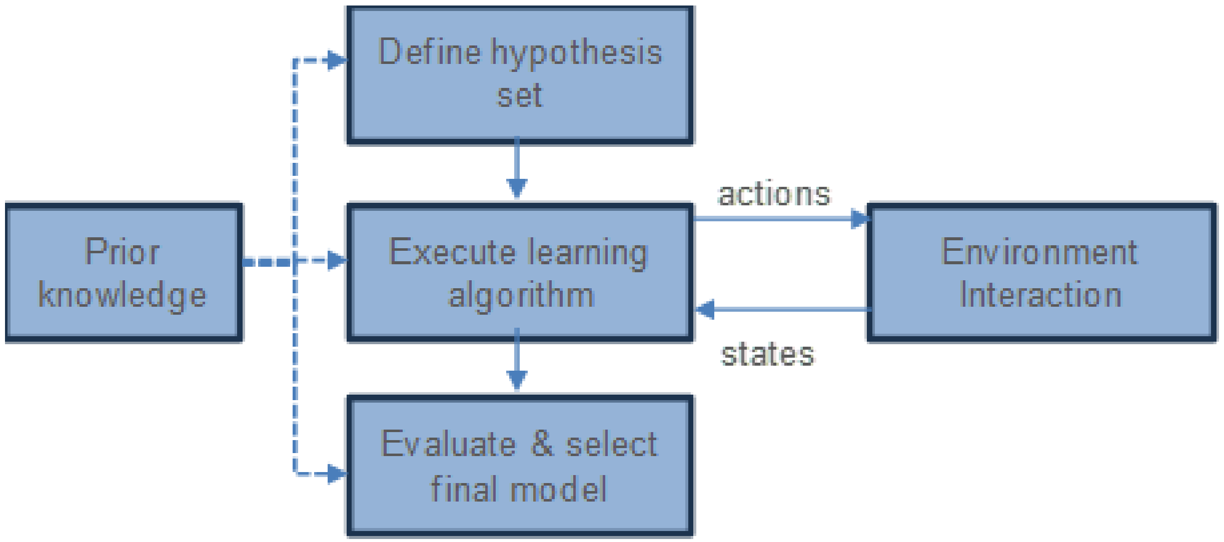

- By far most discussed methods constrain the hypothesis set
- Some exceptions:
  - Mitchener et al. (2022), describe a system that tunes symbolic system
  - Christiano et al. (2017), describe a system learning from human feedback
  - Zhao et al. (2020, 2022) use a predefined policy to generate data to train system

Challenges

- Applying these techniques in network topologies with critical safety constraints.
- Can algebraic or differential equations directly be used in model-free reinforcement learning methods?
- Further study of reasoning or other deliberative components inside a neural network
- How to integrate knowledge in selection and evaluation of the final model?

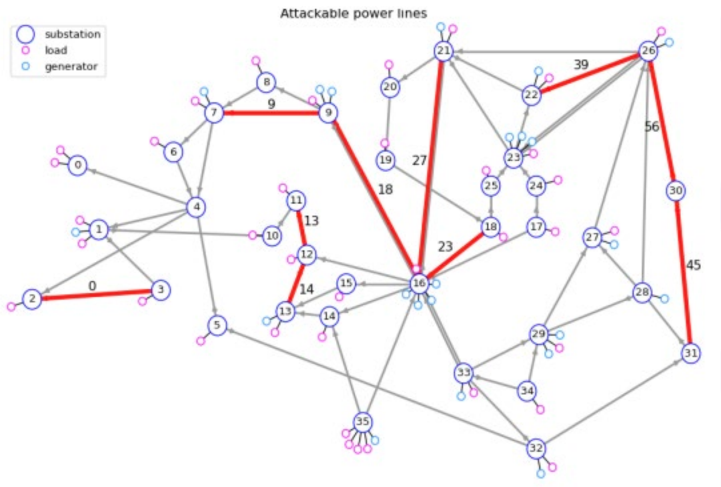

Figure: Marot et al., 2021

Case study – Air traffic control

November 29th, Joost Ellerbroek

AI4REALNET has received funding from European Union's Horizon Europe Research and Innovation programme under the Grant Agreement No 101119527

What ATC is all about

*"To ensure a **safe**, **orderly**, and **expeditious** flow of traffic"*

A safe flow of traffic

- Radar separation:
  - En-route: mostly 5 nautical miles (sometimes 10 nm)
  - TMA: 3 nm

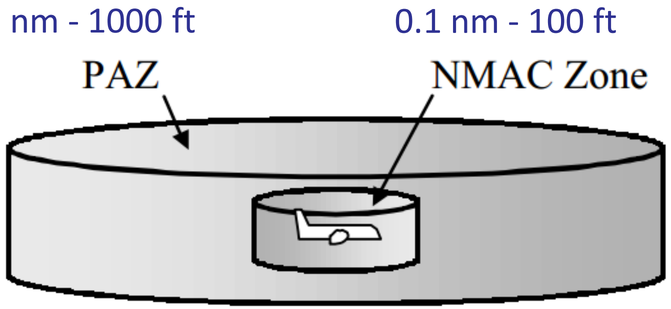

- When aircraft within distance less than 5 nautical miles and less than 1000 ft altitude difference, this is called a **loss-of-separation**

- A **Conflict** is a **predicted** loss of separation, uses protected aircraft zone **(PAZ or PZ)**

- Near miss/Near Mid-Air Collision (**NMAC)** (US) /Airprox (UK CAA reports)

An orderly flow of traffic

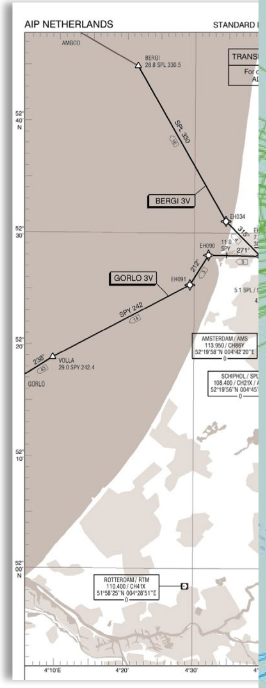

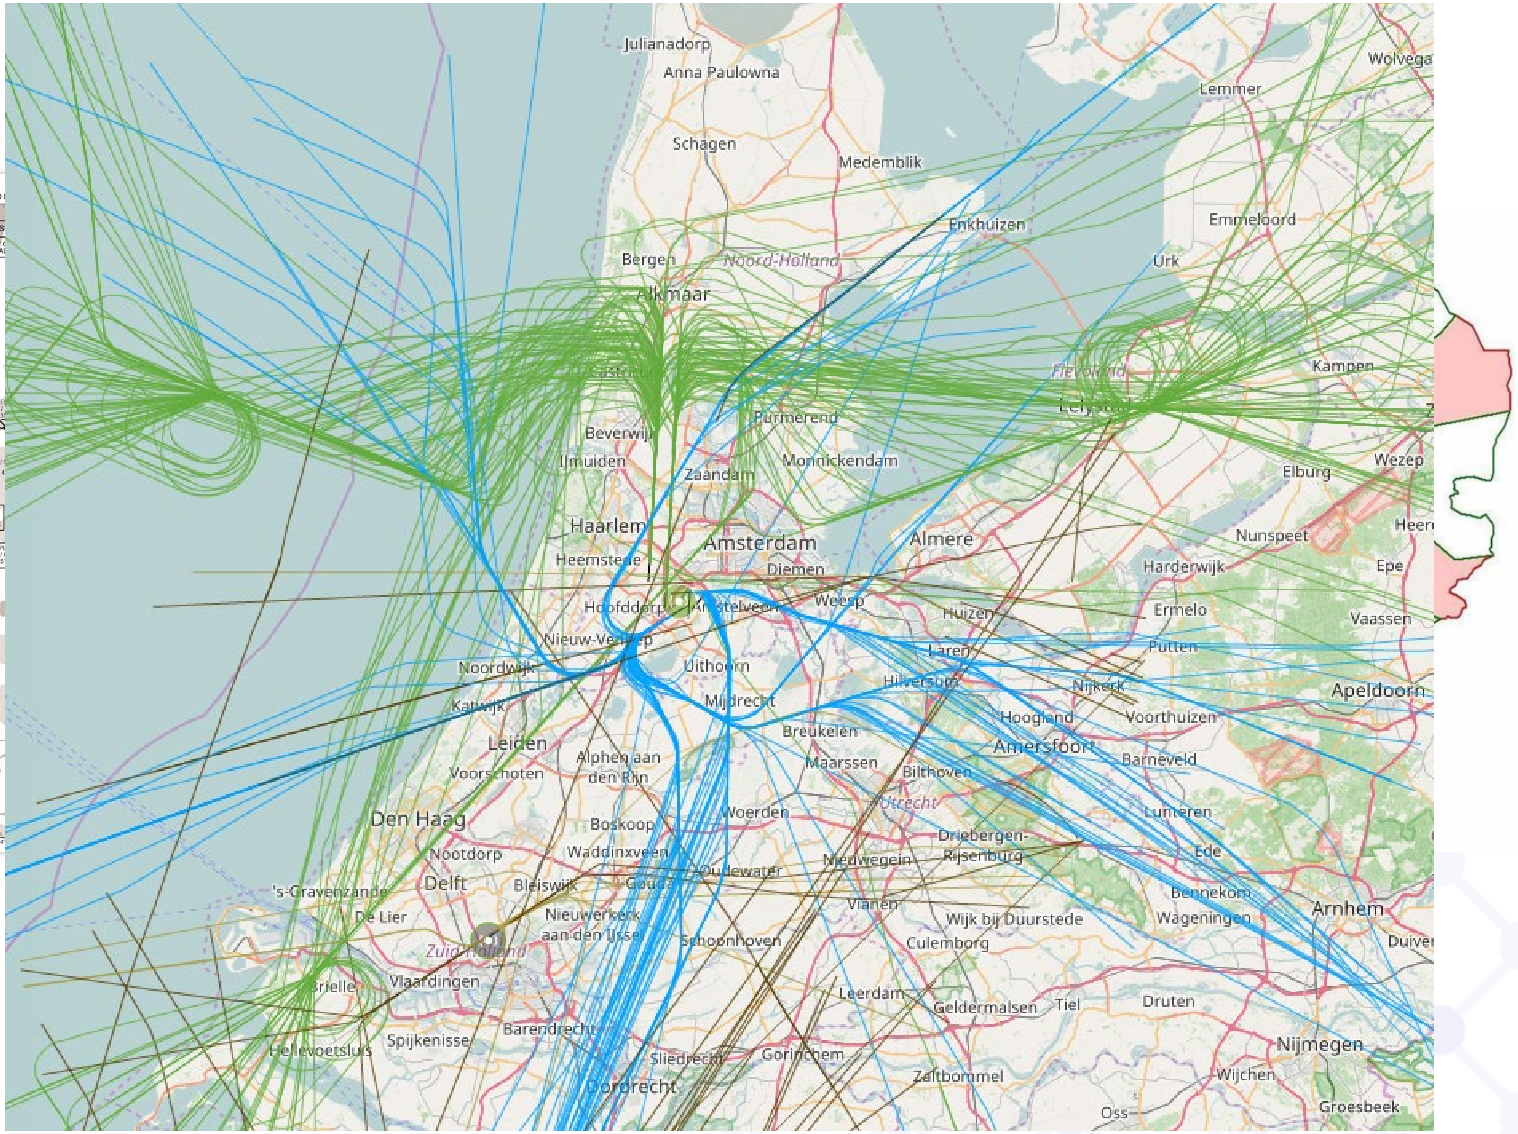

Efficient flight?

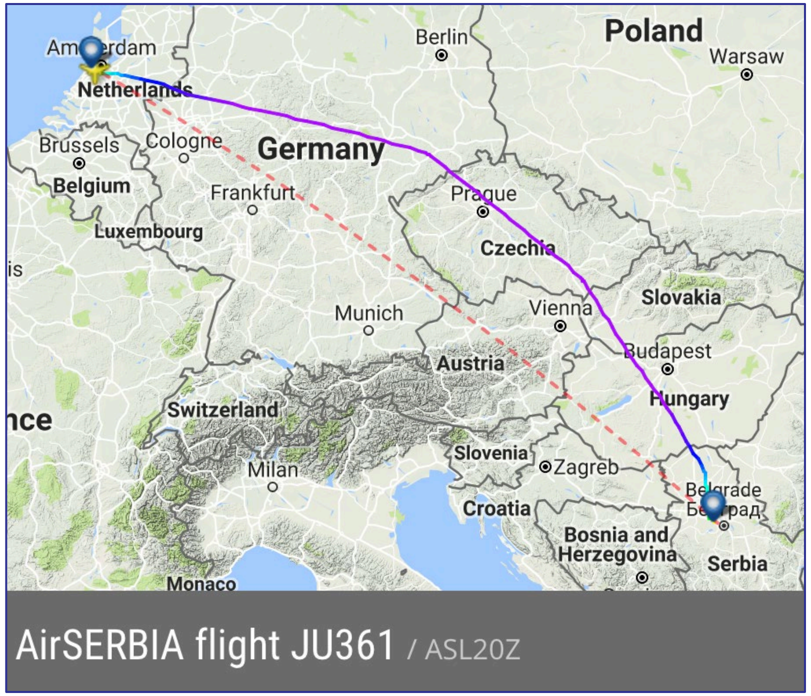

Great-circle distance	1400 km
Flown distance	1500 km

Improving efficiency by delegating separation task

- Enable user-preferred (direct) routing
- En-route separation performed on flight deck
- Development of **geometric** and classical optimisation methods since ‘90s

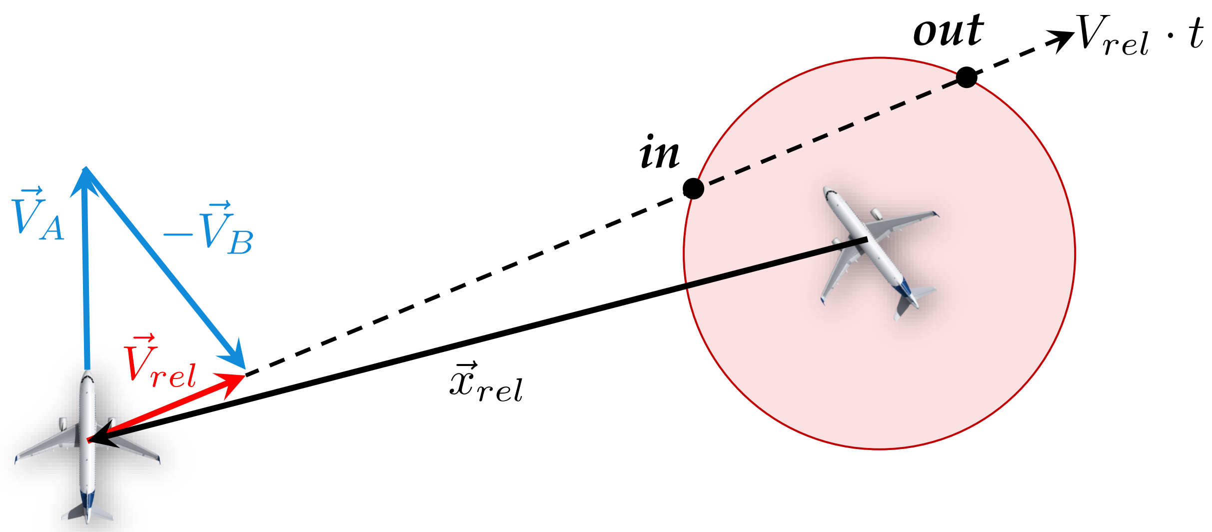

Improving efficiency by delegating separation task

- Enable user-preferred (direct) routing
- En-route separation performed on flight deck
- Development of **geometric** and classical optimisation methods since ‘90s

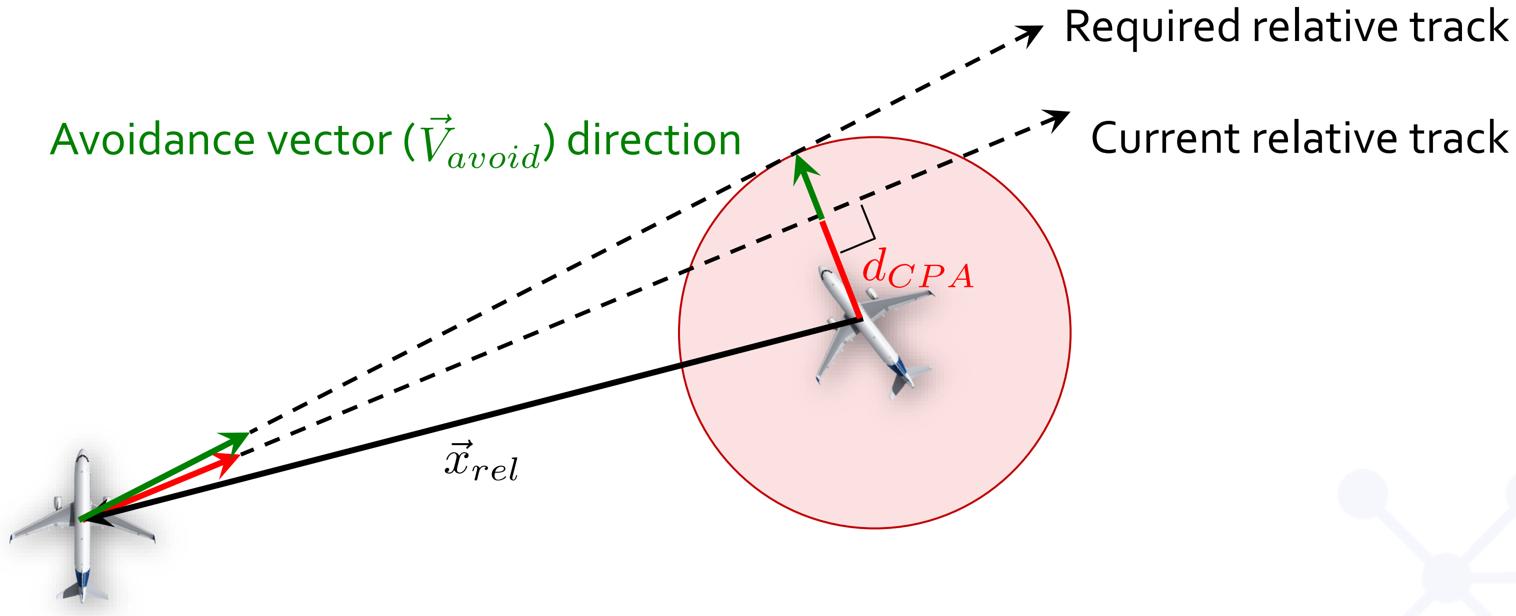

Implicit coordination in geometric methods

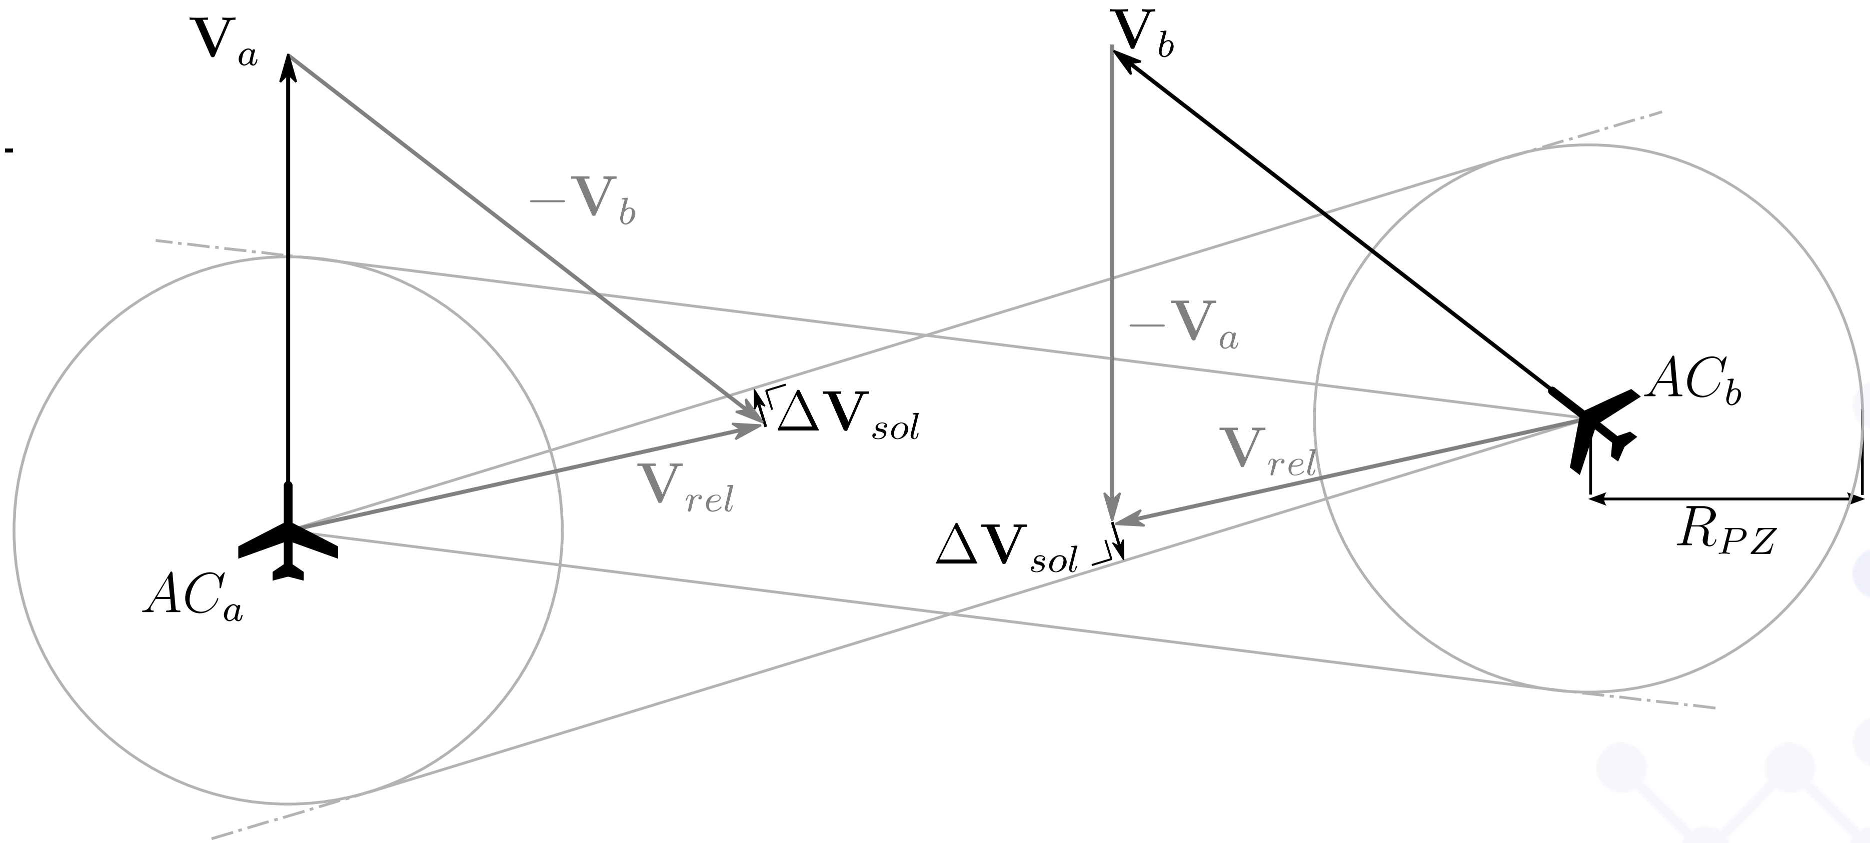

Emergence: the main concern of distributed separation

- For distributed systems, behaviour on the global scale cannot be predicted from local rules and behaviour

- This is the case for even the simplest example: Conway's game of Life

- **Micro-level: simple rule, If sum cells around cell**

0,1	= cell 'dies'	3	= 'birth'
2	= cell 'survives'	4-8	= cell 'dies'

- **Macro-level: complex patterns**

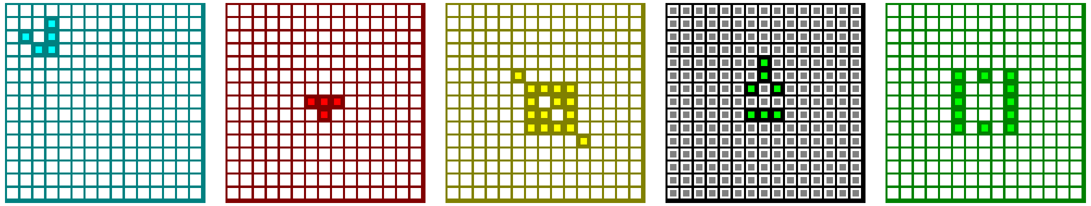

Hybrid application of RL

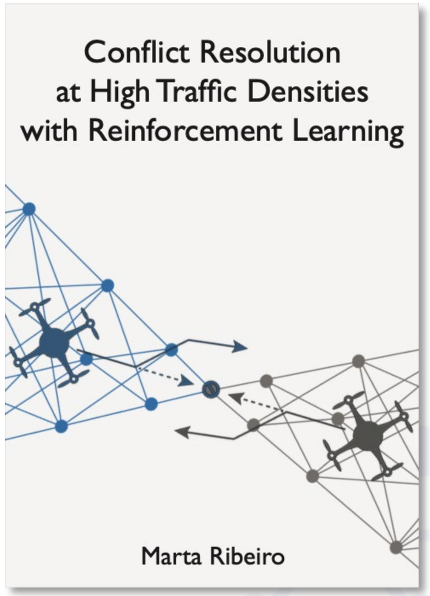

- **Hypothesis**: RL techniques are good at pattern recognition: potential of learning emerging patterns
- Hybrid geometric+RL approach, **RL model** will define:
  - **the look-ahead time,** and
  - **how many degrees of freedom to employ**
    (i.e., heading, speed, or altitude variation)
- **Geometric algorithm** performs resolution actions based on these parameters

Ribeiro, M., Ellerbroek, J., & Hoekstra, J. (2022). Improving Algorithm Conflict Resolution Manoeuvres with Reinforcement Learning. *Aerospace*, *9*(12), 847.

Hybrid geometric/RL conflict resolution model

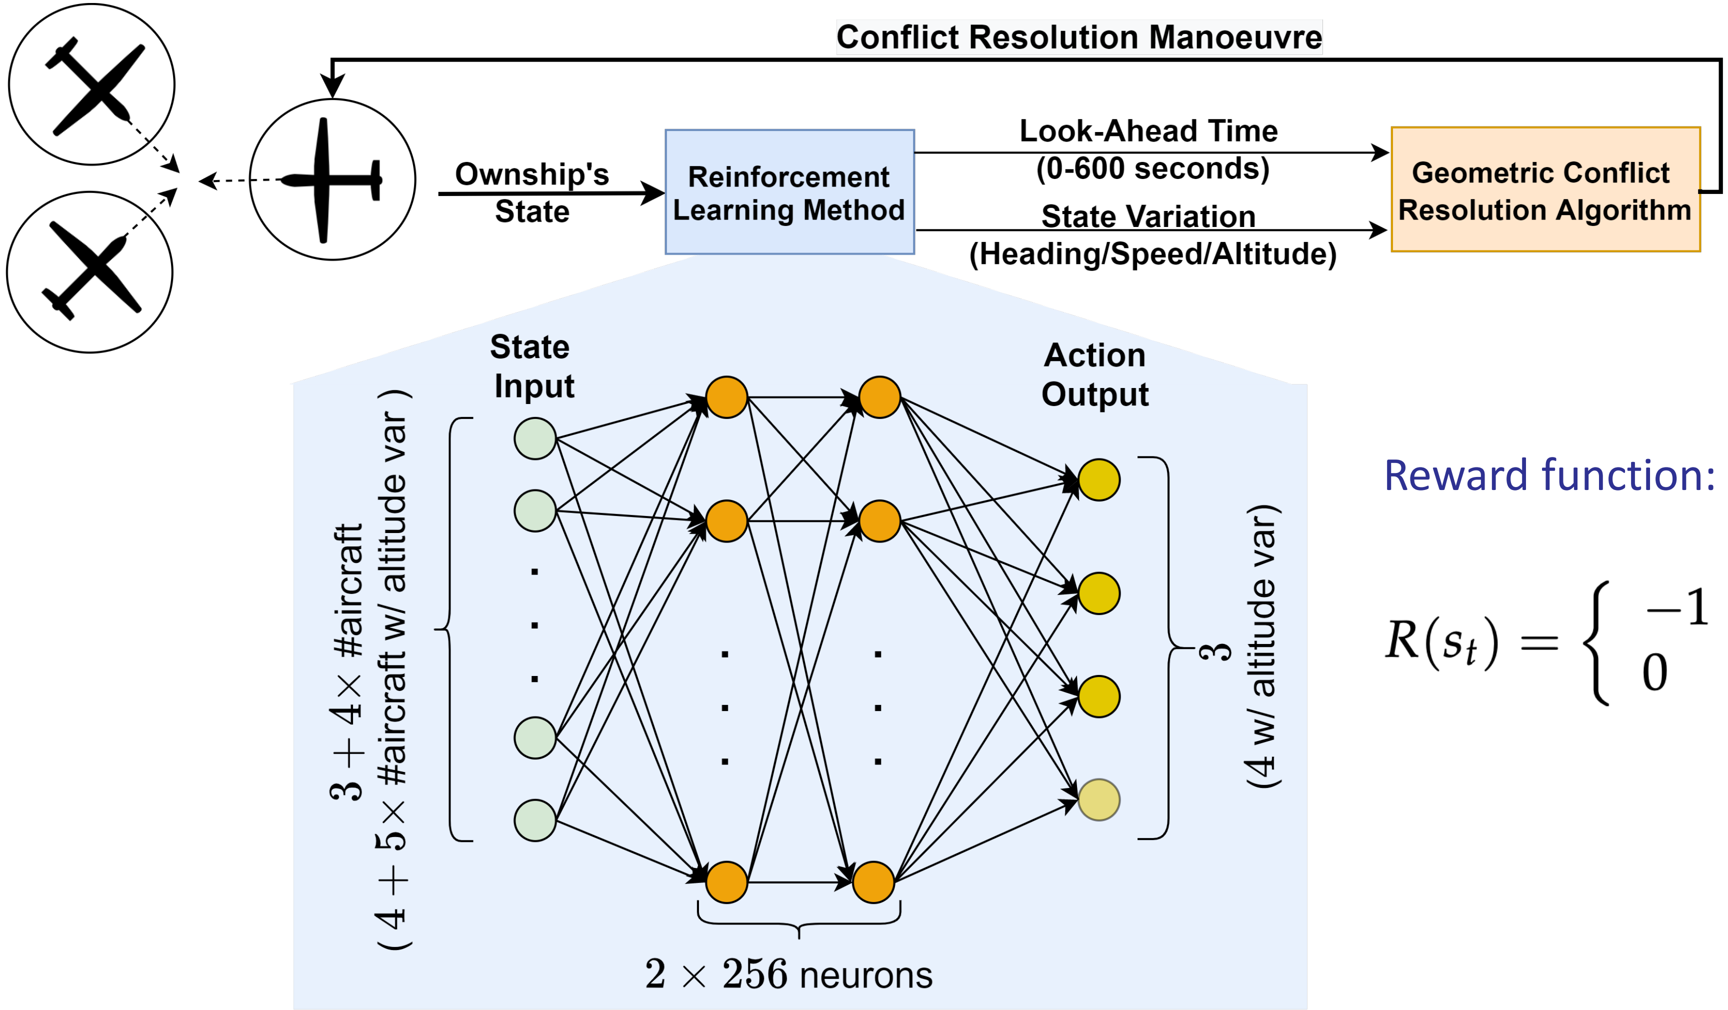

Reward function:

$$R(s_t) = \begin{cases} -1 & \text{Loss of separation} \\ 0 & \text{otherwise} \end{cases}$$

State Space

Dimension	Element	Limits
1	Current heading	−180 ◦ to 180 ◦
1	Relative bearing to next waypoint	−180 ◦ to 180 ◦
1	Current speed	m/s to 18 m/s
#Surrounding aircraft	Current distance to #surrounding aircraft	0 m to 3000 m
#Surrounding aircraft	Distance at CPA with #surrounding aircraft	0 m to 3000 m
#Surrounding aircraft	Time to CPA with #surrounding aircraft	0 s to 600 s
#Surrounding aircraft	Relative heading to #surrounding aircraft	180 ◦ to 180 ◦
Only when the geometric CR method can also perform altitude variation:		
1	Current altitude	0 ft to 100 ft
#Surrounding aircraft	Relative altitude to #surrounding aircraft	0 ft to 100 ft

- Efficiency
- Safety

Action Space

Dimension	Action	Limits	Units
1	Look-ahead time (for CR only)	[−1, + 1] transforms to [0, 600]	Seconds
1	Heading variation	Yes if ≥ 0, no otherwise	Yes/no
1	Speed variation	Yes if ≥ 0, no otherwise	Yes/no
Only when the geometric CR method can also perform altitude variation:			
1	Vertical speed variation	Yes if ≥ 0, no otherwise	

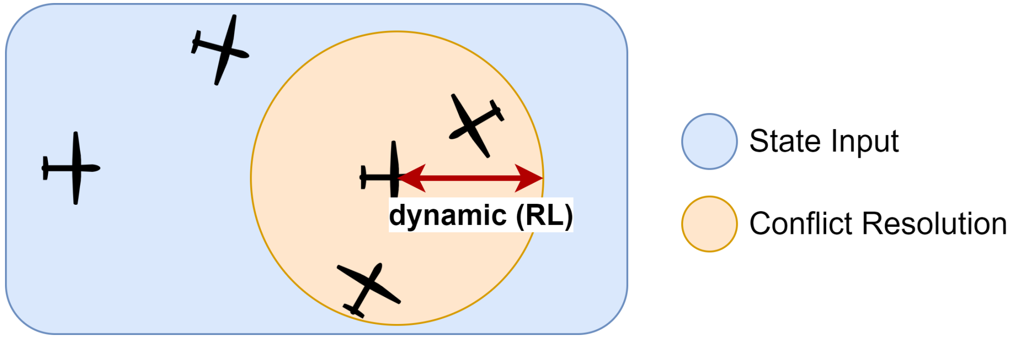

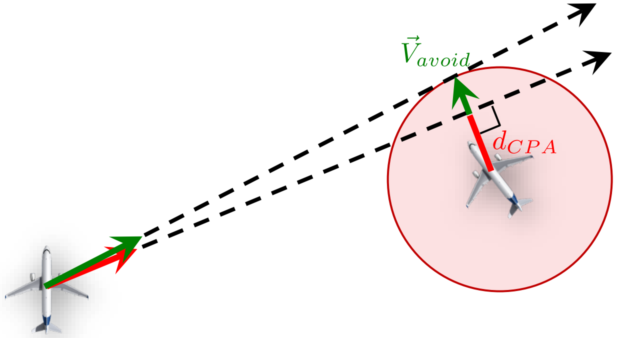

Results

- **Reduced number of LoSs** on all traffic densities, even at a higher traffic density than the RL method was trained on

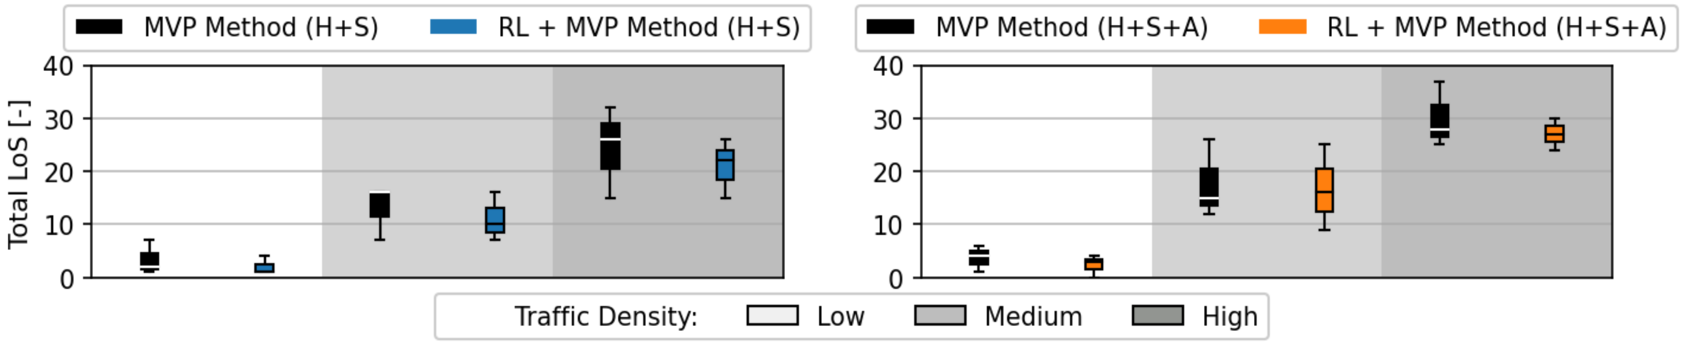

- **Increased number of conflicts**

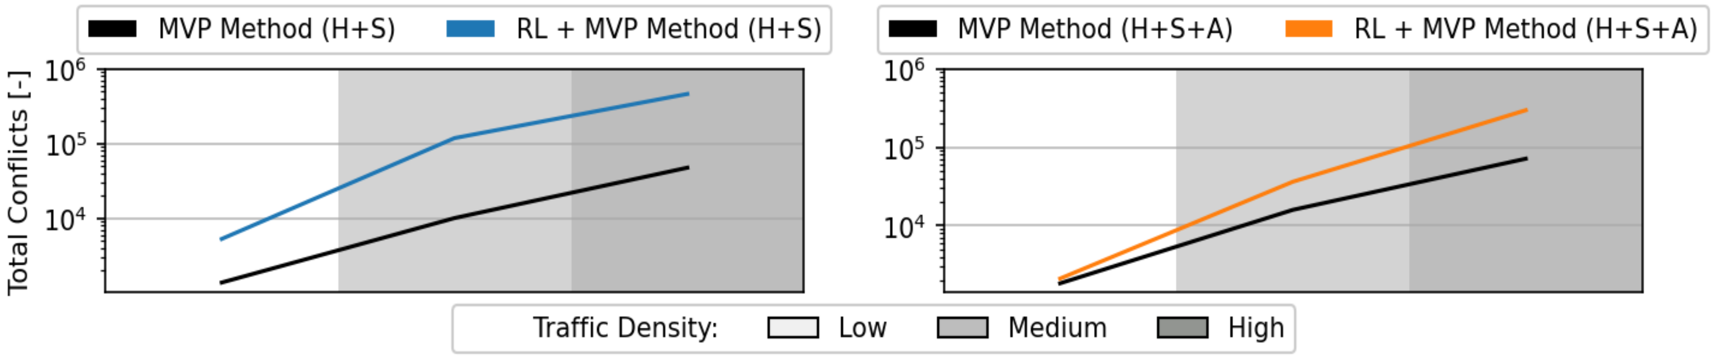

Takeaways and follow-on studies

- The hybrid model generated **fewer losses of minimum separation** than the geometric baseline CR method

- This was caused by two mechanisms:

  1. The **prioritisation** of conflicts depending on the degrees of freedom
  2. The **heterogeneity** of deconflicting directions between aircraft in a conflict situation

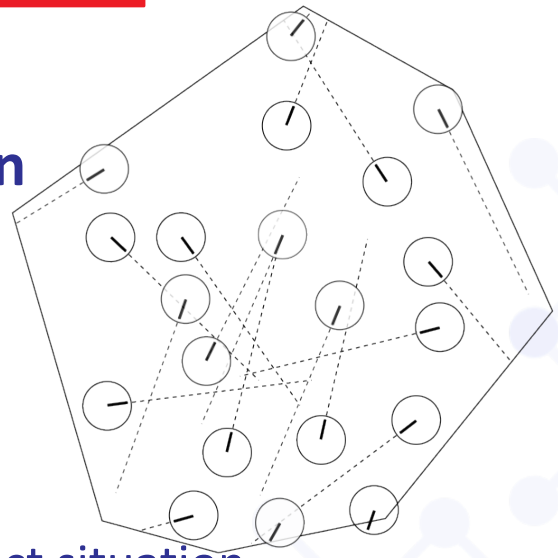

- However, this is still tied reactively to detected conflicts; follow-on studies looked at different structures

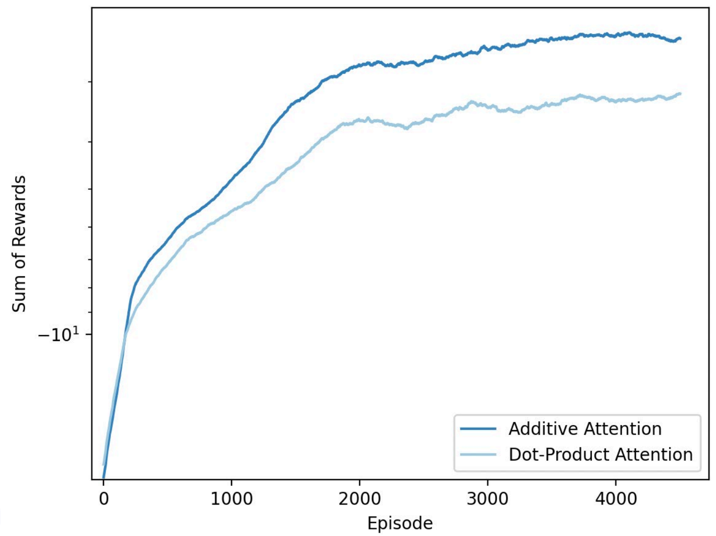

Case study – Power Grid control

November 29th, Milad Leyli-Abadi

AI4REALNET has received funding from European Union's Horizon Europe Research and Innovation programme under the Grant Agreement No 101119527

Power Grid case study: context

- Context - Increase of required number of simulations:
  - Emergence of renewable energy source: less predictable, hard to control
  - Globalization of energy market / exchanges with neighboring countries
  - A wider range of uncertainties to take into account and assess on power-flows

- Physical simulators - limitations
  - However, computation time of a physical simulation on real-grids: 100ms

- Solution: Hybridizing physical models with machine learning
  - Expectation of performance improvement using a ML model: x100 minimum

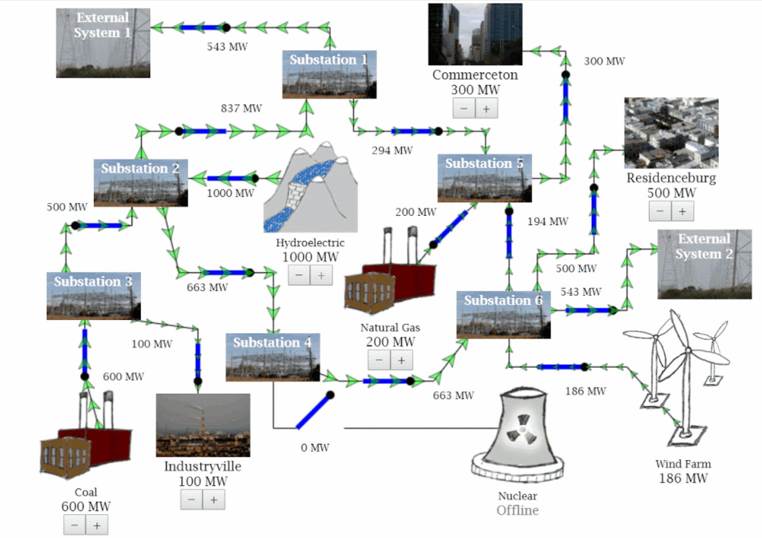

Power Grid case study: problem

- Currently used physical simulators
  - Inputs / Outputs

X = injections
(productions + loads)
T = topology

Numeric Solver
(physics model)

Y = power flows

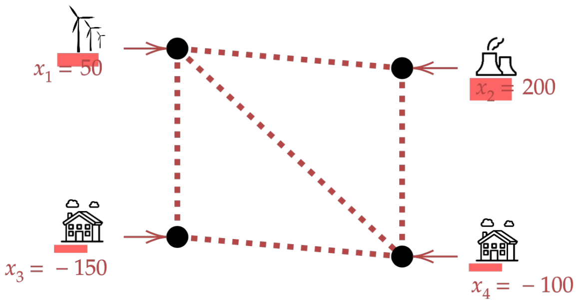

Power Grid case study: problem

- Currently used physical simulators
  - Inputs / Outputs

X = injections
(productions + loads)
T = topology

Numeric Solver (physics model)

Y = power flows

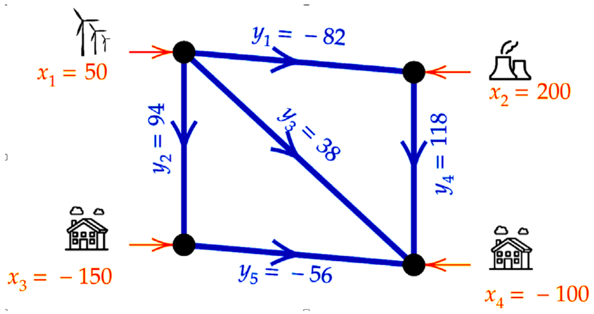

  - Characteristics
    - Relies on **physics equations (Kirchhoff law),** resolved by iterative optimization (Newton-Raphson)
    - Able to predict in a normal condition or different grid conditions

Power Grid equations

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

Power Grid case study: grid representation

- The power grid could be represented naturally as graph

Power Grid
Substations and lines

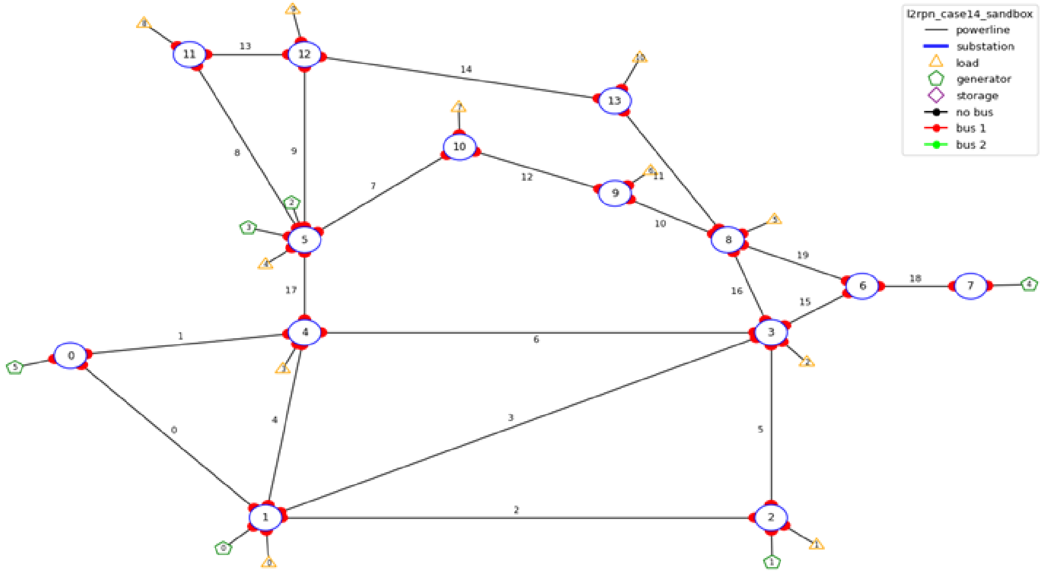

Graph
Nodes and edges

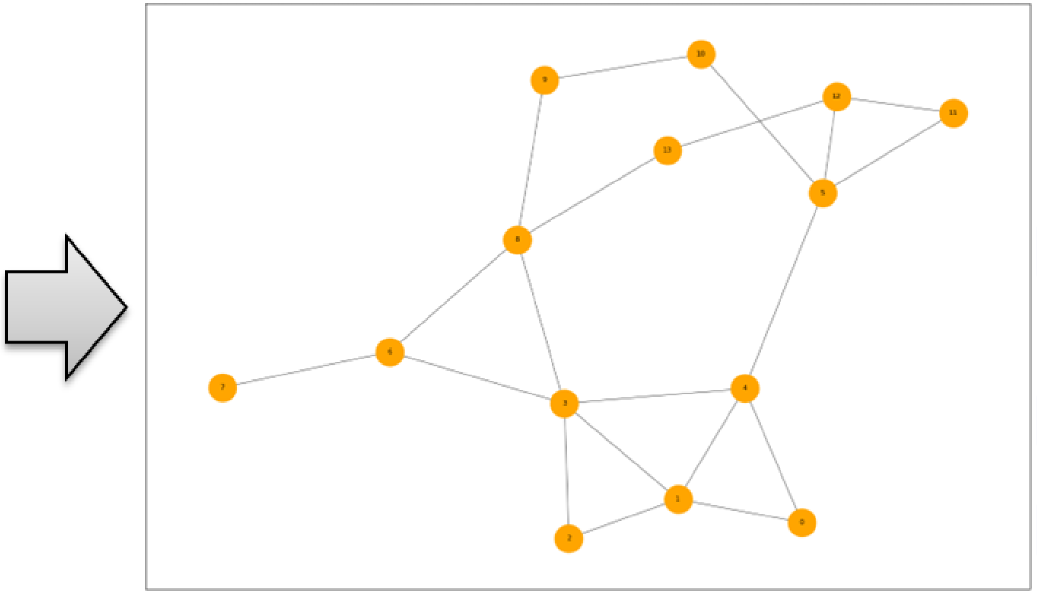

Power Grid case study: Physics equations

- Local conservation law for a given substation i from power flow p of connected lines ℓ

$$p_i^{prod} - p_i^{load} = \sum_{\ell \in N(i)} p_i^{\ell}$$

- Equivalently, we can write the active powers p, in terms of voltage angles θ and admittances y of neighboring nodes j

$$p_i^{prod} - p_i^{load} = \sum_{j \in \{i, N(i)\}} \theta_j \times y_{ij}$$

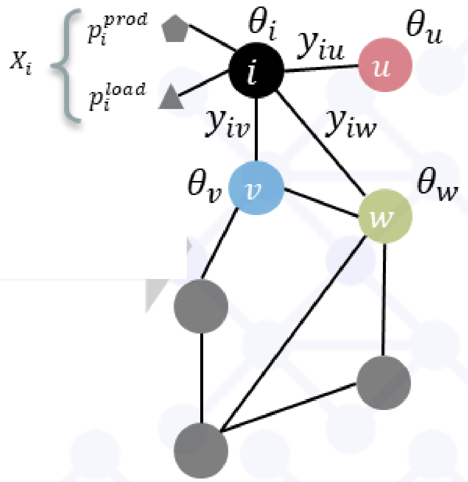

- Considering the neighborhood of node i as $N(i) = \{u, v, w\}$, this becomes

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

- The new value of θ for layer k could be computed as following

$$\theta_i^{(k)} = \frac{(p_i^{prod} - p_i^{load}) - [(\theta_u^{(k-1)} \times y_{iu}) + (\theta_v^{(k-1)} \times y_{iv}) + (\theta_w^{(k-1)} \times y_{iw})]}{y_{ii}}$$

Power Grid case study: Physics Informed GNN

- GNN layers followed by local conservation (LC) layers to compute the error

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

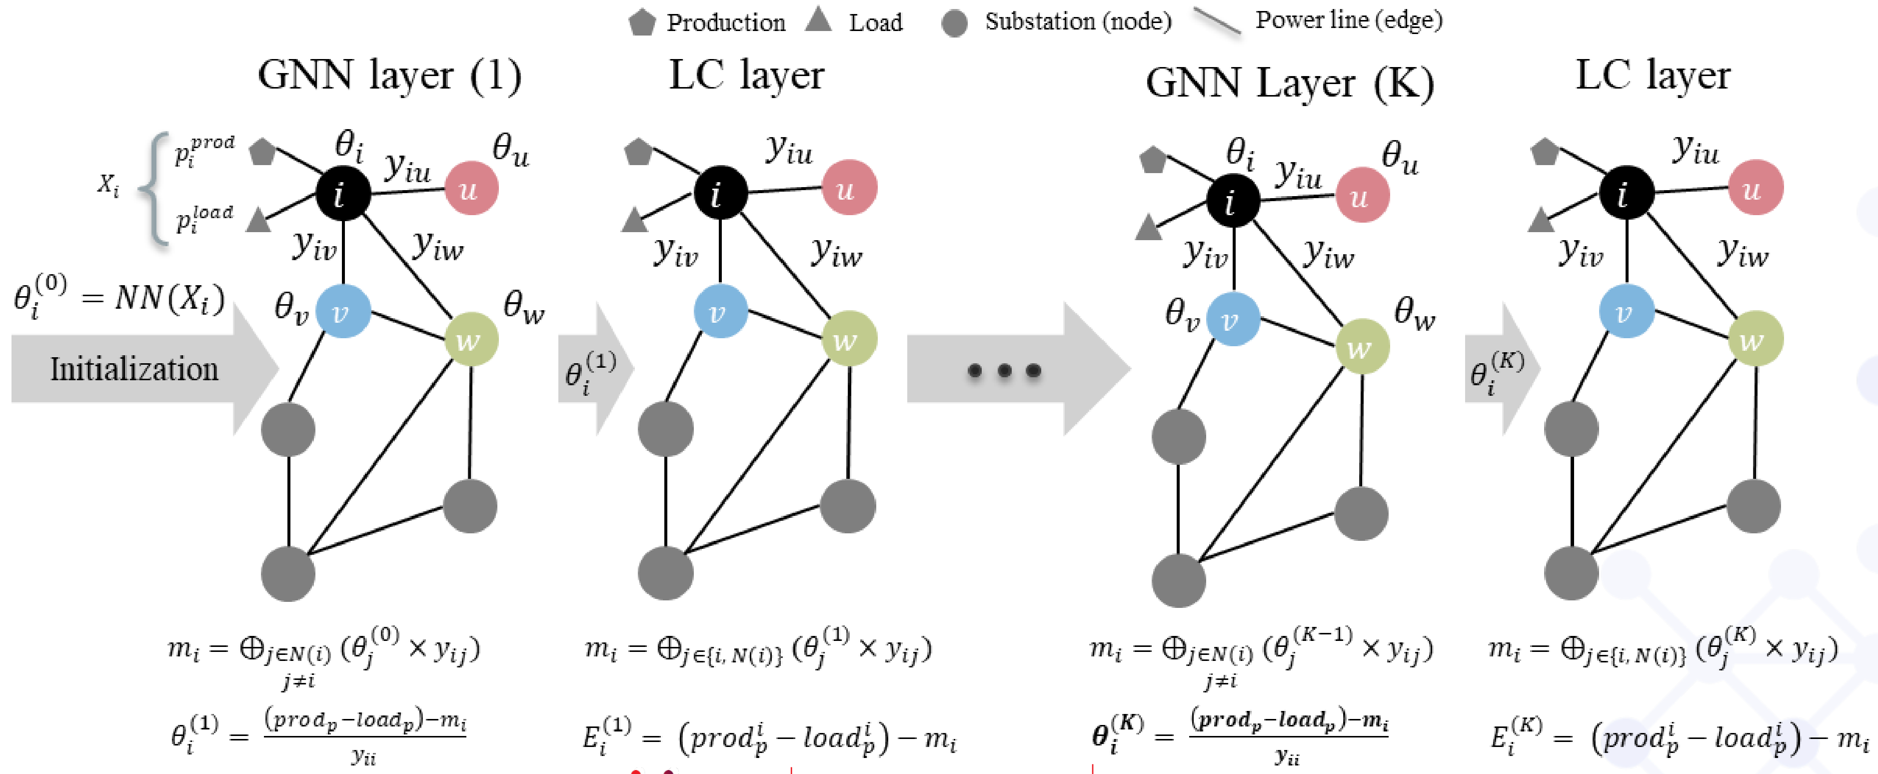

Datasets and distributions

Training and Test datasets

300 000 observations

100 000 observations

Out-of-distribution Test dataset

200 000 observations

⇨ To assess the robustness of models

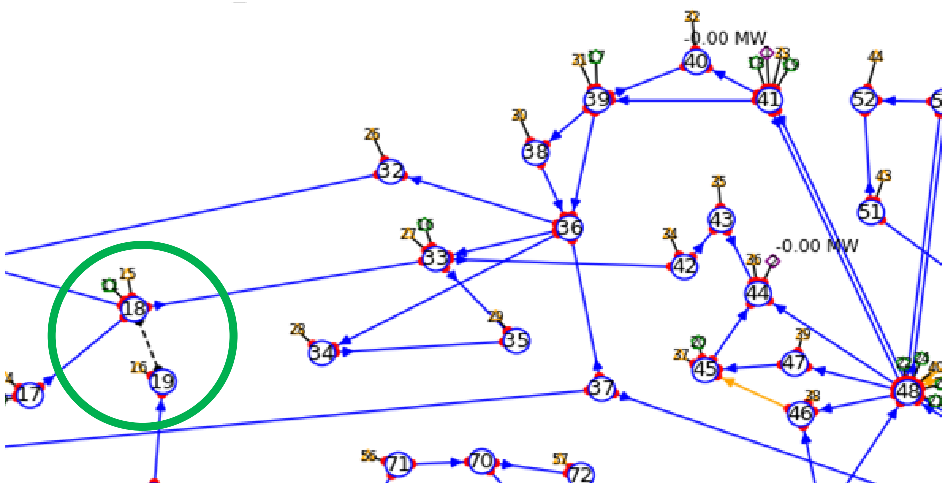

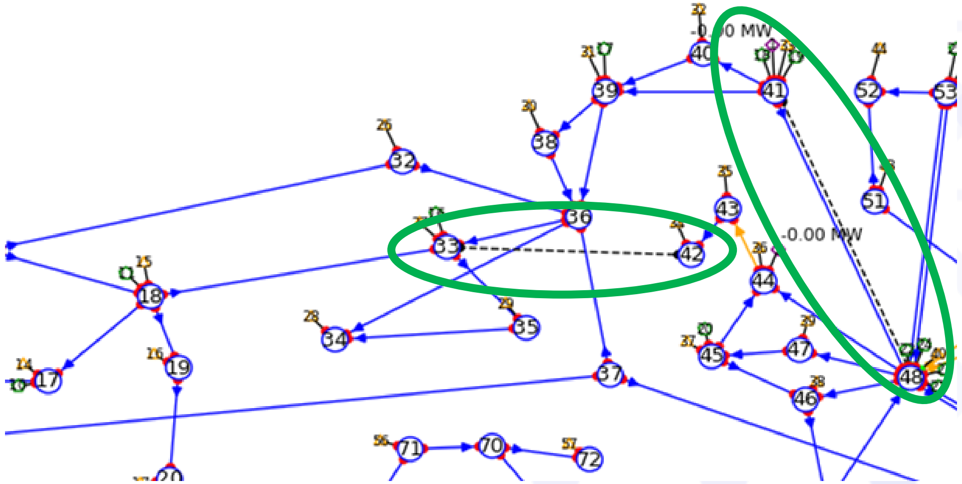

Power Grid case study: Evaluation results

- 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**

Power Grid case study: Multi-criteria evaluation

Evaluation based in multiple categories of metrics

ML-related

MAPE 90
currents

MAPE 10
powers

MAE
voltages

Physics compliance

ID	Type	Measure
		Basic
P1	Current positivity	$\frac{1}{L}\Sigma_{\ell}^{L}\mathbb{1}_{(\hat{a}_{or,ex}^{\ell}<0.)}$
P2	Voltage positivity	$\frac{1}{L}\Sigma_{\ell}^{L}\mathbb{1}_{(\hat{v}_{or,ex}^{\ell}<0.)}$
P3	Losses positivity	$\frac{1}{L}\Sigma_{\ell}^{L}\mathbb{1}_{(\hat{p}_{ex}^{\ell}+\hat{p}_{or}^{\ell}<0.)}$
P4	Disconnected Line	$\frac{1}{L_{disc}}\Sigma_{\ell_{disc}}^{L_{disc}}\mathbb{1}_{(\lvert\hat{x}_{ex}^{\ell}\rvert+\lvert\hat{x}_{or}^{\ell}\rvert>0.)}$
P5	Energy Losses	$\frac{\Sigma_{\ell=1}^{L}(\hat{p}_{ex}^{(\ell)}+\hat{p}_{or}^{(\ell)})}{Gen} \in [0.005, 0.04]$
		Uni-dimension law
P6	Global Conservation	$\mathrm{MAPE}((Prod - Load) - (\Sigma_{\ell=1}^{L}(\hat{p}_{ex}^{\ell}+\hat{p}_{or}^{\ell})))$
P7	Local Conservation	$\mathrm{MAPE}((p_k^{prod} - p_k^{load}) - (\Sigma_{l\in neig(k)}\hat{p}_k^{\ell}))$
P8	Voltage equality	$\sum_{\substack{i,j \\ i,j\in k \\ i\neq j}}\mathbb{1}_{(\lvert v_i - v_j\rvert>0)}$

Generalization

Model capacity to generalize on unobserved out-of-distribution (OOD) dataset

Industrial readiness

$$Speed\ Up = \frac{Time_{ClassicalSolver}}{Time_{Inference}}$$

Evaluation pipeline: LIPS framework

- LIPS: Learning Industrial Physical Simulation
- A modular framework to evaluate hybrid models
- Open-source framework based on various categories of evaluation criteria
- Multiple competitions are organized on the basis of LIPS framework

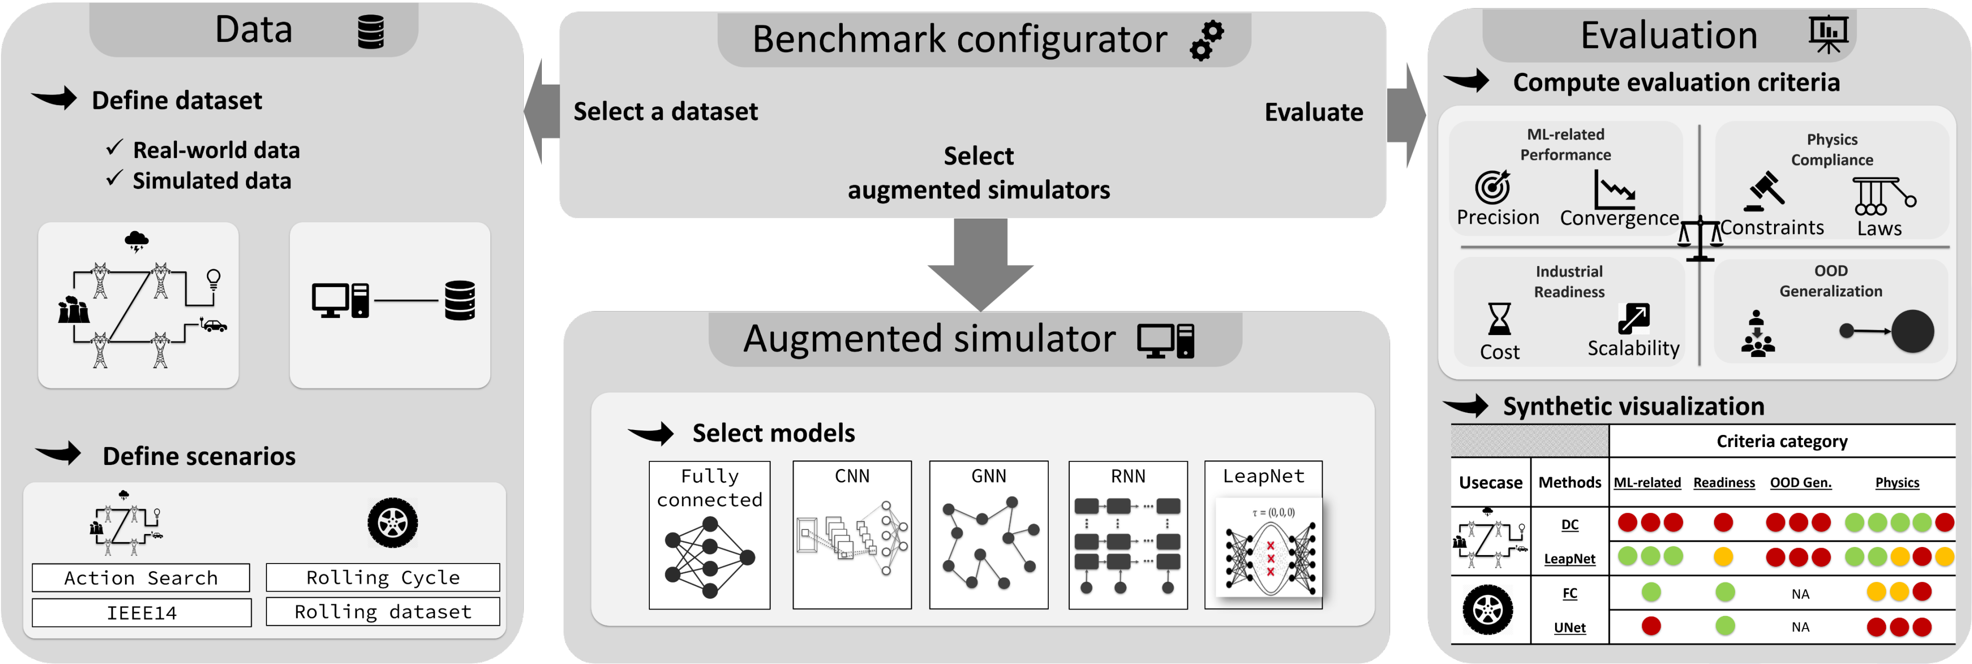

Github

Power Grid case study: Physics criteria to respect

- The AI-based solutions should be conform wrt. various physics criteria/law

ID	Type	Measure	Description
		Basic	
P1	Current positivity	$\frac{1}{L}\sum_{\ell}^{L}\mathbb{1}_{(\hat{a}^{\ell}_{or,ex}<0.)}$	Proportion of negative current
P2	Voltage positivity	$\frac{1}{L}\sum_{\ell}^{L}\mathbb{1}_{(\hat{v}^{\ell}_{or,ex}<0.)}$	Proportion of negative voltages
P3	Losses positivity	$\frac{1}{L}\sum_{\ell}^{L}\mathbb{1}_{(\hat{p}^{\ell}_{ex}+\hat{p}^{\ell}_{or}<0.)}$	Proportion of negative energy losses
P4	Disconnected Line	$\frac{1}{L_{disc}}\sum_{\ell_{disc}}^{L_{disc}}\mathbb{1}_{(\lvert\hat{x}^{\ell}_{ex}\rvert+\lvert\hat{x}^{\ell}_{or}\rvert>0.)}$	Proportion of non-null a, p or q values
P5	Energy Losses	$\frac{\sum_{\ell=1}^{L}(\hat{p}^{(\ell)}_{ex}+\hat{p}^{(\ell)}_{or})}{Gen} \in [0.005, 0.04]$	energy losses range consistency
		Uni-dimension law	
P6	Global Conservation	$\mathrm{MAPE}((Prod-Load)-(\sum_{\ell=1}^{L}(\hat{p}^{\ell}_{ex}+\hat{p}^{\ell}_{or})))$	Mean energy losses residual
P7	Local Conservation	$\mathrm{MAPE}((p^{prod}_{k}-p^{load}_{k})-(\sum_{l\in neig(k)}\hat{p}^{\ell}_{k}))$	Mean active power residual at nodes
P8	Voltage equality	$\sum_{\substack{i,j\\ i,j\in k\\ i\neq j}}\mathbb{1}_{(\lvert v_i-v_j\rvert>0)}$	Proportion of not equal voltages at nodes

Competition on a real-world application

- Fast contingency screening
  - Penetration of renewable energy (30% wind + solar here)
  - Changing topologies at substations
  - Trust from the operators with acceptable compliance to physical laws

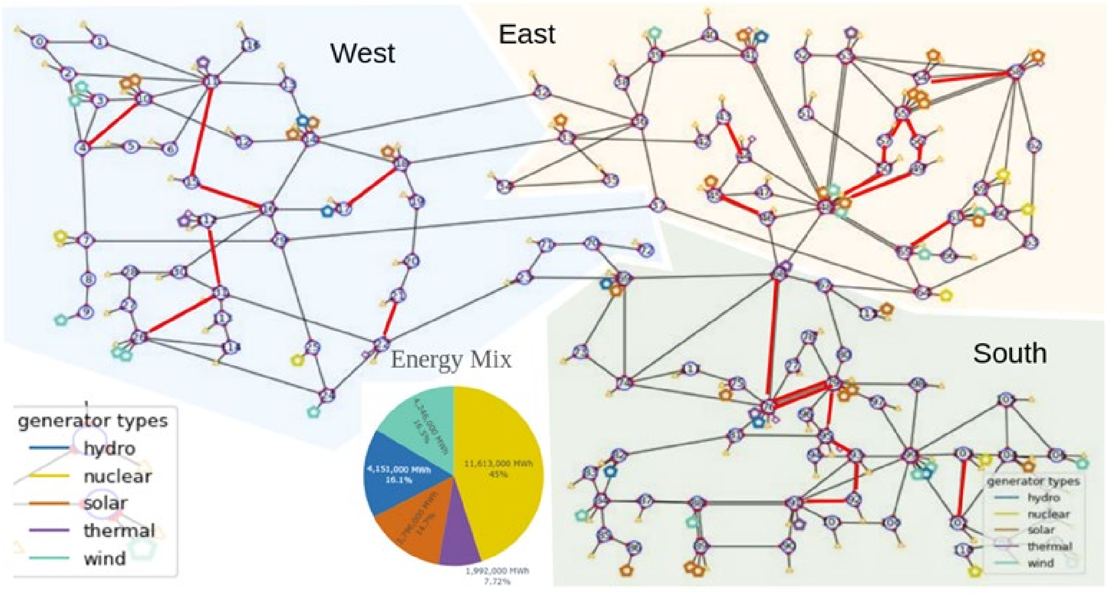

Method	Test dataset (30%) ML-related(66%) a_{or}	a_{ex}	p_{or}	p_{ex}	v_{or}	v_{ex}	Test dataset (30%) Physics(34%) p_1	p_2	p_3	p_4	p_5	p_6	p_7	p_8	Criteria category Speed-up (40%) Speed-up	OOD generalization (30%) ML-related(66%) a_{or}	a_{ex}	p_{or}	p_{ex}	v_{or}	v_{ex}	OOD generalization (30%) Physics(34%) p_1	p_2	p_3	p_4	p_5	p_6	p_7	p_8	Global Score (%)
Powerflow (LightSim2grid)	🟢	🟢	🟢	🟢	🟢	🟢	🟢	🟢	🟢	🟢	🟢	🟢	🟢	🟢	1	🟢	🟢	🟢	🟢	🟢	🟢	🟢	🟢	🟢	🟢	🟢	🟢	🟢	🟢	60.2
Competition baselines																														
Fully Connected	🟠	🟠	🔴	🔴	🟠	🟠	🟢	🟢	🔴	🔴	🟠	🔴	🔴	🔴	15.45	🔴	🔴	🔴	🔴	🔴	🔴	🟢	🟢	🔴	🔴	🟠	🔴	🔴	🔴	33.5
LeapNet	🟢	🟢	🟠	🟠	🔴	🔴	🟢	🟢	🔴	🔴	🟠	🔴	🔴	🔴	11.9	🟠	🟠	🟠	🟠	🔴	🔴	🟢	🟢	🔴	🔴	🟠	🔴	🔴	🔴	37.6
Competition Ranking																														
1-Arizona State University	🟢	🟢	🟢	🟢	🟢	🟢	🟢	🟢	🟢	🟢	🟠	🟢	🔴	🟢	7.87	🟢	🟢	🟢	🟢	🟢	🟢	🟢	🟢	🟢	🟢	🟠	🟢	🔴	🟢	64.2 ± .62
2-XI'AN JIAOTONG University	🟢	🟢	🟠	🟠	🟢	🟢	🟢	🟢	🔴	🔴	🟠	🟠	🟢	🔴	9.69	🟢	🟢	🟠	🟠	🟢	🟢	🟢	🟢	🔴	🔴	🟠	🟠	🟢	🔴	57.89 ± 1.42
3-Xlerator ThinAir	🟢	🟢	🟢	🟢	🔴	🔴	🟢	🟢	🔴	🔴	🟠	🔴	🔴	🔴	12.42	🟠	🟠	🟠	🟠	🔴	🔴	🟢	🟢	🔴	🔴	🟠	🔴	🔴	🔴	41.15 ± 1.27

Concluding remarks

November 29th, Milad Leyli-Abadi and Herke van Hoof

AI4REALNET has received funding from European Union's Horizon Europe Research and Innovation programme under the Grant Agreement No 101119527

Concluding remarks

- Knowledge assisted AI promising for control of real-world networks
- Various techniques for using knowledge-assisted AI in decision making
- Demonstration of applications in power networks and air traffic management

Future directions

- **Deep integration** of optimization, network structure, and constraints
- Facilitate **Human-AI collaboration** for better decision-making
- **Autonomous adaptation** of AI systems in response to changing environments

References

- M. Alshiekh, R. Bloem, R. Ehlers, B. Könighofer, S. Niekum, and U. Topcu. Safe reinforcement learning via shielding. In Proc. AAAI conference on artificial intelligence, volume 32, 2018.
- B. Araki, X. Li, K. Vodrahalli, J. DeCastro, M. Fry, and D. Rus. The logical options framework. In Int. Conf. Machine Learning, pages 307–317. 2021.
- P.F. Christiano, J. Leike, T. Brown, M. Martic, S. Legg, and D. Amodei. Deep reinforcement learning from human preferences. Advances in neural information processing systems, 30, 2017.
- H. Dai, P. Chen, and H. Yang. Driving torque distribution strategy of skid-steering vehicles with knowledge-assisted reinforcement learning. Applied Sciences, 12(10):5171, 2022.
- A. d'Avila Garcez, A.R.R. Dutra, and E. Alonso. Towards symbolic reinforcement learning with common sense. arXiv preprint arXiv:1804.08597, 2018.
- M. Garnelo, K. Arulkumaran, and M. Shanahan. Towards deep symbolic reinforcement learning. arXiv preprint arXiv:1609.05518, 2016.
- N. Höpner, I. Tiddi, and H. van Hoof. Leveraging class abstraction for commonsense reinforcement learning via residual policy gradient methods. In Int. Joint Conf. Artificial Intelligence, IJCAI 2022, pages 3050–3056.
- A. Marot, B. Donnot, G. Dulac-Arnold, A. Kelly, A. O'Sullivan, J. Viebahn, M. Awad, I. Guyon, P. Panciatici, and C. Romero. "Learning to run a power network challenge: a retrospective analysis." In NeurIPS 2020 Competition and Demonstration Track, pp. 112-132. PMLR, 2021.
- L. Mitchener, D. Tuckey, M. Crosby, and A. Russo. Detect, understand, act: A neuro-symbolic hierarchical reinforcement learning framework. Machine Learning, 111(4):1523– 1549, 2022.
- T.M. Moerland, J. Broekens, A. Plaat, C.M. Jonker, et al. Model-based reinforcement learning: A survey. Foundations and Trends in Machine Learning, 16(1):1–118, 2023.
- Md K. Sarker, L. Zhou, A. Eberhart, and P. Hitzler. Neuro-symbolic artificial intelligence. AI Communications, 34(3):197–209, 2021.
- P. Vaezipoor, A.C. Li, R.A. Toro Icarte, and S.A. Mcilraith. Ltl2action: Generalizing ltl instructions for multi-task rl. In Int. Conf. Machine Learning, 2021.
- E. Van der Pol, D. Worrall, H. van Hoof, F. Oliehoek, and M. Welling. MDP homomorphic networks: Group symmetries in reinforcement learning. Advances in Neural Information Processing Systems, 33:4199–4210, 2020.
- F. Van Harmelen and A. Ten Teije. A boxology of design patterns for hybrid learning and reasoning systems. J. Web Engineering, 18(1-3):97–123, 2019.
- L. Von Rueden, S. Mayer, K. Beckh, B. Georgiev, S. Giesselbach, R. Heese, B. Kirsch, J. Pfrommer, A. Pick, R. Ramamurthy, et al. Informed machine learning–a taxonomy and survey of integrating prior knowledge into learning systems. IEEE Trans. Knowledge and Data Engineering, 35(1):614–633, 2021.
- Jan Wöhlke, Felix Schmitt, and Herke Hoof. Value refinement network (VRN). In Int. Joint Conf. Artificial Intelligence, 2022.
- Y. Xie, Y. Zhang, T. Lin, Z. Pan, S.-Z. Qian, B. Jiang, and J. Yan. Short video preloading via domain knowledge assisted deep reinforcement learning. Digital Communications and Networks, 2024.
- D. Yu, B. Yang, D. Liu, H. Wang, and S. Pan. A survey on neural-symbolic learning systems. Neural Networks, 2023.
- H. Zhao, J. Zhao, J. Qiu, G. Liang, and Z. Y. Dong. Cooperative wind farm control with deep reinforcement learning and knowledge-assisted learning. IEEE Trans. Industrial Informatics, 16(11):6912–6921, 2020.
- H. Zhao, Z. Liu, X. Mai, J. Zhao, J. Qiu, J. Liu, Z.Y. Dong, and A.MYM Ghias. Mobile battery energy storage system control with knowledge-assisted deep reinforcement learning. Energy Conversion and Economics, 3(6):381–391, 2022.

Thanks for your attention!

Questions?

AI4REALNET has received funding from European Union's Horizon Europe Research and Innovation programme under the Grant Agreement No 101119527

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.

Project use cases: focus on critical infrastructures

1. Human operators are aided in their decision-making by an AI-assistant to congestion problems
2. Transfer from simulation to real-world (*Sim2Real*)

1. AI assistant exploring different modes of co-learning for train re-dispatching
2. AI-based system that makes re-dispatching decisions in a fully automated way

1. Airspace sectorization assistant
2. Flow & airspace management assistant