Explainable machine learning for optimisation of resource use in large-scale heterogenous infrastructures

Prof. Ivana Dusparic

Co-Director SFI Center for Research Training in AI
FI @ CONNECT, Enable
Trinity College Dublin

Pieter Barnard

Doctoral Researcher
CRT Advance, CONNECT
Trinity College Dublin

Resource use optimization in large-scale infrastructures

- Intelligent mobility: urban traffic control, shared on-demand autonomous mobility

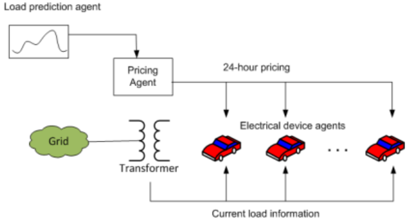

- Smart grid: demand response, renewable energy scheduling

M. Gueriau, F. Cugurullo, R. A. Acheampong, I. Dusparic. Shared Autonomous Mobility-on-Demand: Learning-based approach and its performance in the presence of traffic congestion, IEEE ITS Magazine, 2020.
I. Dusparic, A. Taylor, A. Marinescu, F. Golpayegani, S. Clarke. Residential demand response: Experimental evaluation and comparison of self-organizing techniques. Renewable and Sustainable Energy Reviews, Vol 80, 2017

- 5G+ networks: connectivity to and by UAVs

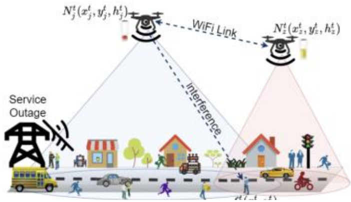

- IoT: sensor scheduling and data processing on energy-constrained devices

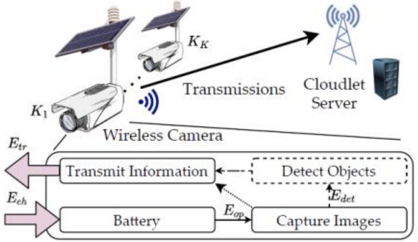

B. Galkin, E. Fonseca, R. Amer, L.A. DaSilva, I. Dusparic. REQIBA: Regression and Deep Q-Learning for Intelligent UAV Cellular User to Base Station Association, IEEE Transactions on Vehicular Technology, 2022.
J. Hribar, R. Shinkuma, G. Iosifidis, I. Dusparic. Analyse or Transmit: Utilising Correlation at the Edge with Deep Reinforcement Learning, IEEE Global Communications Conference (GLOBECOM) 2021

enable

Trinity College Dublin
Coláiste na Tríonóide, Baile Átha Cliath
The University of Dublin

Characteristics of the problems

Large scale, decentralized

Inner city road network

Multiple (and conflicting) goals

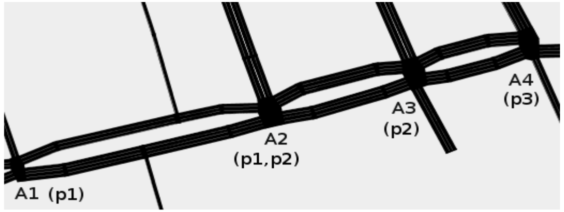

Road network segment: agents A1-A4 meeting goals p1-p3

Dynamically changing operating environment

Non-stationary patterns of renewable energy availability

Why Reinforcement Learning?

- Conditions are dynamic and unpredictable, so are interactions between entities
- Individual behaviours (and their combinations) cannot be predefined at design time – should be learnt!

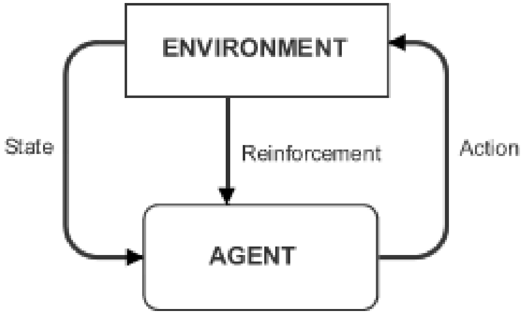

Reinforcement Learning (RL)

1. Through trial and error in the interaction of the environment learns long-term quality of taking a particular action in a particular state
2. Model-free - does not require model of the environment

The not-so-perfect side of Reinforcement Learning

- Single agent using RL, to meet a single goal, in a static environment, is guaranteed to converge to an optimal solution

- However, but what about:
  - Multiple agents?
  - Multiple goals?
  - Dynamic environments?
  - All of these together?
  - Interpretability and explainability of the decisions made?
  - Minimizing “regret” – reducing a period of bad performance while learning?
- Research focus: new multi-agent RL techniques for large-scale infrastructures

RL Algorithms

New RL techniques for large-scale infrastructures

Distributed W-Learning

- W-value – relative importance of action nominated by Q-value
- Learn cooperative coefficient C

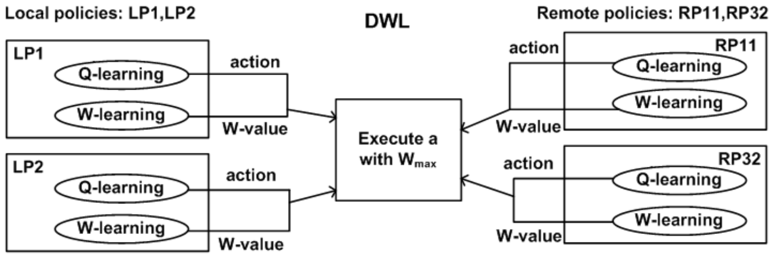

I. Dusparic and V. Cahill. Autonomic Multi-Policy Optimization in Pervasive Systems: Overview and Evaluation. ACM TAAS Vol 7, 1. Apr. 2012

Prediction-Based Multi-Agent RL for Non-stationary Environments

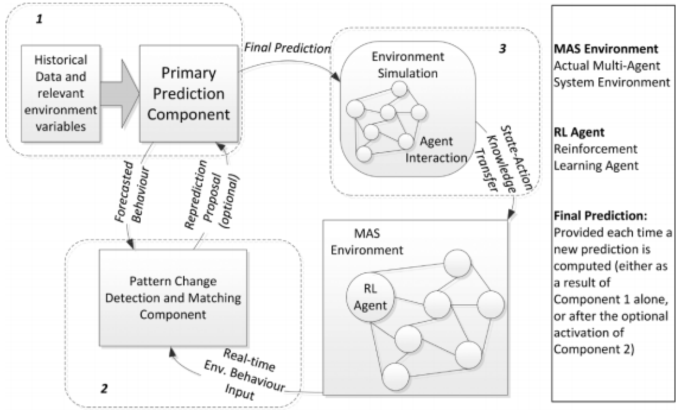

A. Marinescu, I. Dusparic, A. Taylor, V. Cahill, S. Clarke. P-MARL: Prediction-based multi-agent reinforcement learning in inherently non-stationary environments. ACM TAAS, 2017.

New RL techniques for large-scale infrastructures

Parallel Transfer Learning

Runtime knowledge transfers in multi-agent RL

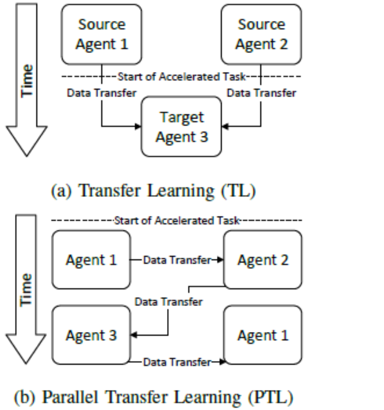

Temporal Relationship of Knowledge Transfers

A. Taylor, I. Dusparic, M. Gueriau, S. Clarke. Parallel Transfer Learning in Multi-Agent Systems: What, when and how to transfer? IJCNN, 2019.

A. Castagna I. Dusparic. Multi-Agent Transfer Learning in Reinforcement Learning-Based Ride-Sharing Systems, ICAART 2022

Online state-space generation

Multi-Layered Growing Neural Gas for learning optimal state space granularity

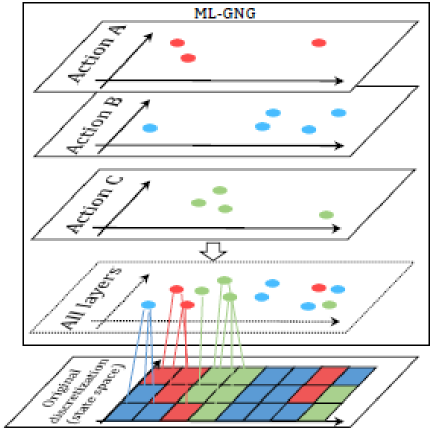

M. Gueriau, N. Cardozo, I. Dusparic. Con-RL: Constructivist Approach to State Space Adaptation in Reinforcement Learning, SASO 2019.

Explainable ML

Interpretability / explainability of ML and RL

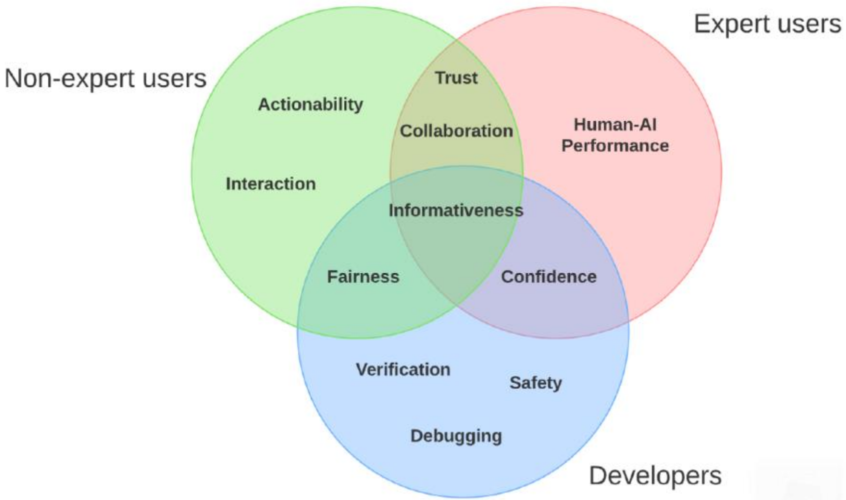

- Machine learning
  - Feature importance

- Reinforcement learning
  - Temporal component/ sequences of actions
  - Policy summarization/extracting rules
  - Contrastive explanations /comparing preferences
  - Discovering causal relationships – generalizability and explainability

- Local vs global (individual decisions vs a model)

Feature Importance

Resource reservation in Sliced Networks

Shared network infrastructure

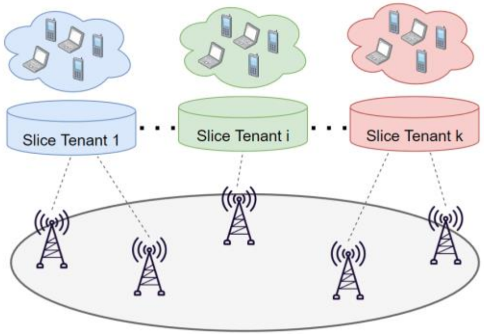

Each tenant wishes to maximize QoS of their end-users, minimize expenditure for resources

AI model used by tenant to autonomously reserve future resources from operator

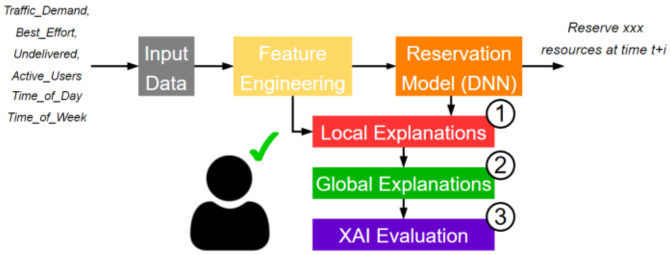

XAI framework added onto conventional AI pipeline

Local explanation to verify real-time operation

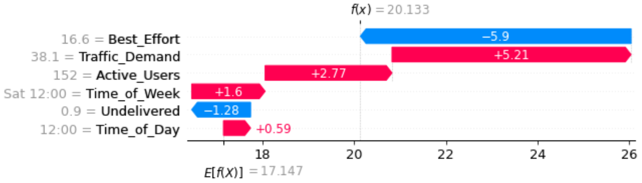

Global explanation to debug harmful behavior

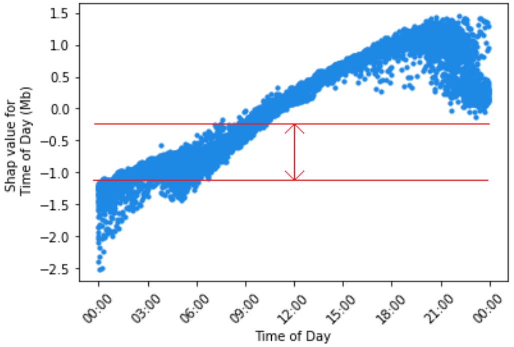

Barnard, Pieter, et al. "Resource reservation in sliced networks: An explainable artificial intelligence (xai) approach." IEEE ICC, 2022.

Network Intrusion Detection

- Current XAI frameworks stop at the human-in-the-loop.
  - To what extent can we use ML to automatically "extract" useful information from explanations?

- In the context of Network Intrusion Detection, where…
  - Model trained to classify normal vs malicious traffic.
  - During testing, new attacks types may appear.

- We show that...
  - Deep Auto-Encoder (DAE) trained specifically to reconstruct explanations taken from training set.
  - During testing, DAE capable of accurately identifying instances where new attacks occur (recall of ~89%).

- PCA plot of "explanation domain" shows greater class separation compared to raw input data.

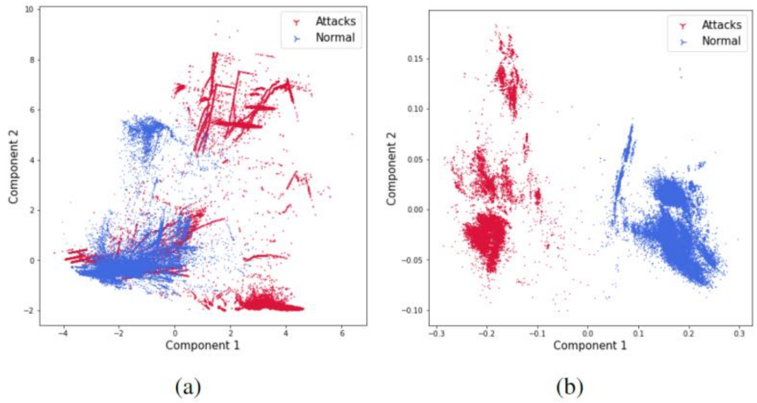

Fig. 4: PCA applied to the training data (a), and the explanations (b).

Barnard, Pieter, et al. "Robust Network Intrusion Detection through Explainable Artificial Intelligence (XAI)." (2022).

Causality in RL

Explainable RL - ReCCoVER

- Recognizing Causal Confusion for Verifiable and Explainable RL
- Untangle causes vs correlations in RL state input features to achieve:
  - Explainability
  - Better generalizability

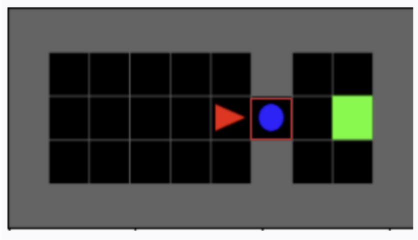

MiniGridworld traffic environment: Agent (red) needs to reach the goal (green), avoid collision with the other vehicle (blue) and obey the traffic light (outlined red).

J. Gajcin and I. Dusparic. ReCCoVER: Detecting Causal Confusion for Explainable Reinforcement Learning, EXTRAAMAS workshop @ AAMAS 2022

Explainable RL - ReCCoVER

1. Identify critical states (eg local maxima)
2. Train feature-parameterized policies
3. Generate alternative worlds (in which certain correlations between feature do not hold) and execute policies
4. Detect causal confusion (failed to perform well in alternative worlds)

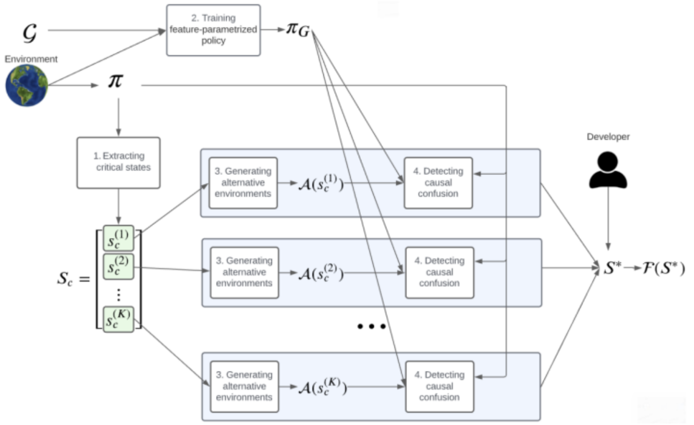

Policy	Number of episodes	Number of collected transitions	Number of critical states	Average number of alternative worlds per critical state	Number of states where causal confusion detected
$\pi_{confused}$	100	1500	2	14	1
$\pi_{correct}$	100	1500	1	15	0

J. Gajcin and I. Dusparic. ReCCoVER: Detecting Causal Confusion for Explainable Reinforcement Learning, EXTRAAMAS workshop @ AAMAS 2022

Wrapping up

Resource use optimization in large-scale infrastructures

- Collaborate with other agents when their goals differ from ours/are potentially conflicting
- Share knowledge with other agents to improve speed of learning
- Detect environment change and adapt to it by re-learning
- Explain the decisions and rationale of the agent in human-understandable terms

What next:

- Learn online ⇨ do more with less data, minimize training time
- Learn from humans when knowledge is available ⇨ human-in-the-loop RL
- Adapt online to previously unseen situations
- Explain all parts of the pipeline and perform (semi)-automated actions ⇨ why did the agent do this? how has the environment changed? what actions should the human take, if any?

enable

Thank You!

ivana.dusparic@tcd.ie, barnardp@tcd.ie

Trinity College Dublin
Coláiste na Tríonóide, Baile Átha Cliath
The University of Dublin

CONNECT

CRT·AI
SFI Centre for Research Training in Artificial Intelligence

ADVANCECRT