IFIP Networking 2024

TX4Nets 2024

CHALMERS
UNIVERSITY OF TECHNOLOGY

# Explainable Reinforcement Learning: Towards Trustworthy Autonomous Network Operations

Carlos Natalino

Researcher, Optical Networks Unit

Department of Electrical Engineering

Chalmers University of Technology, Gothenburg, Sweden

# Outline

- Network automation & programmability
- Trustworthy AI
- Understanding reinforcement learning
- Representative use case
- Preliminary investigation
- Challenges and opportunities

# Network automation & programmability

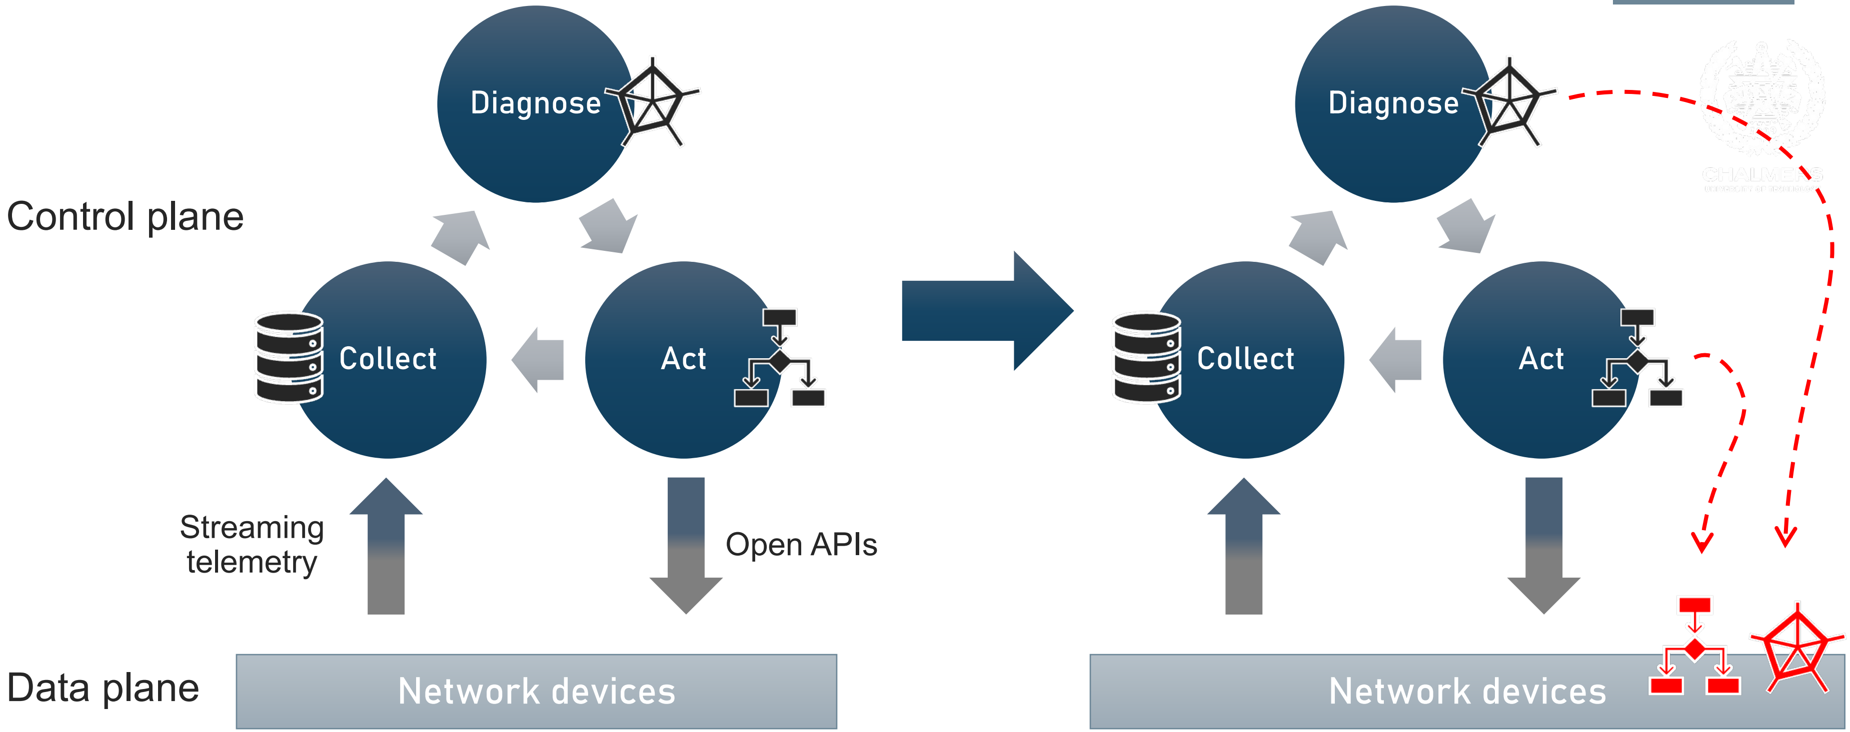


*Achim Autenrieth, "Carrier Grade AI/ML for Network Automation", invited talk, OFC 2022, 9 March 2022

# Fixed networks

## Network automation and programmability

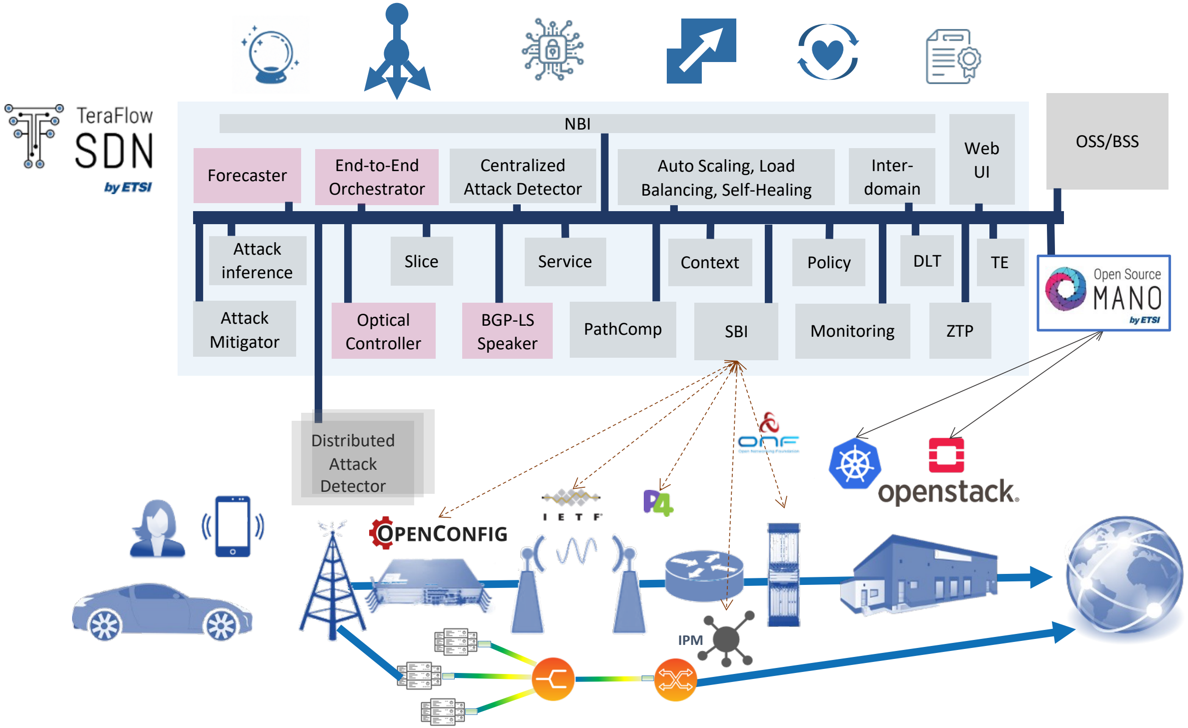


Source: ETSI SDG TFS

# Wireless networks

## Network automation and programmability

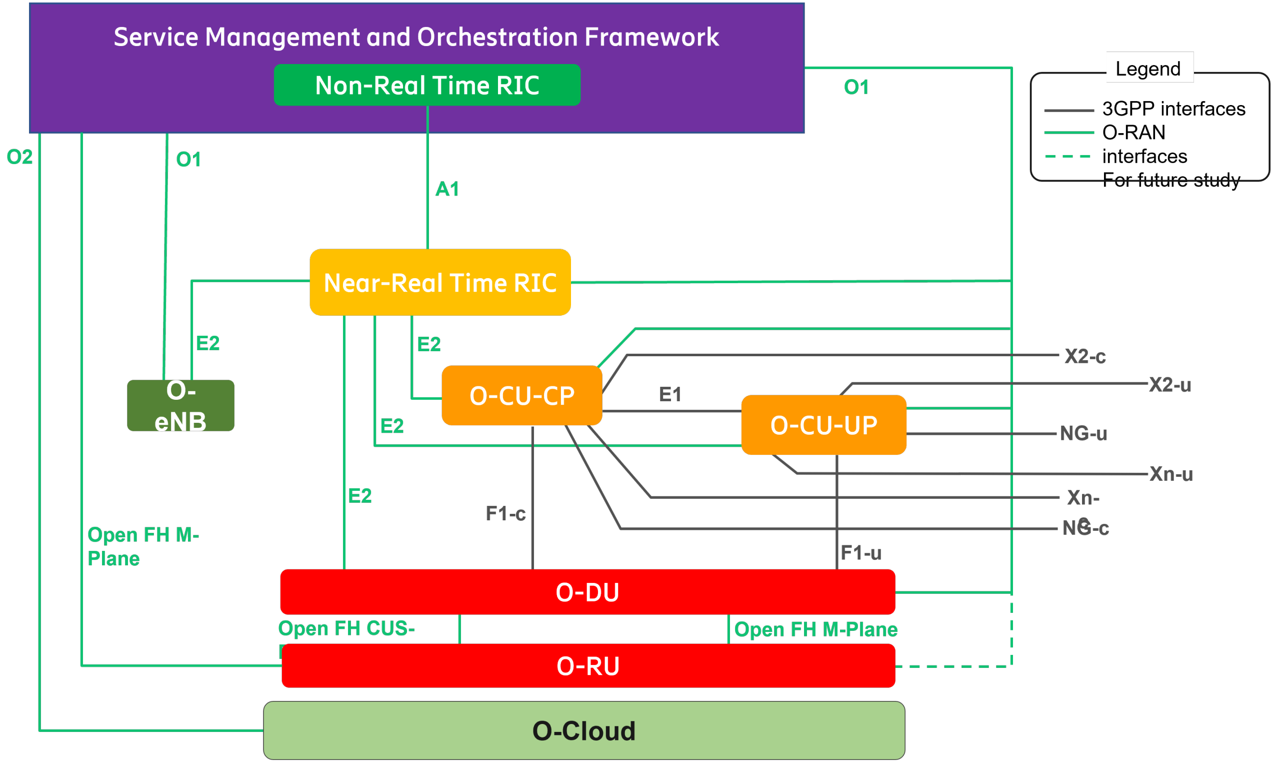


Source: O-RAN Alliance

# Current level of network autonomy

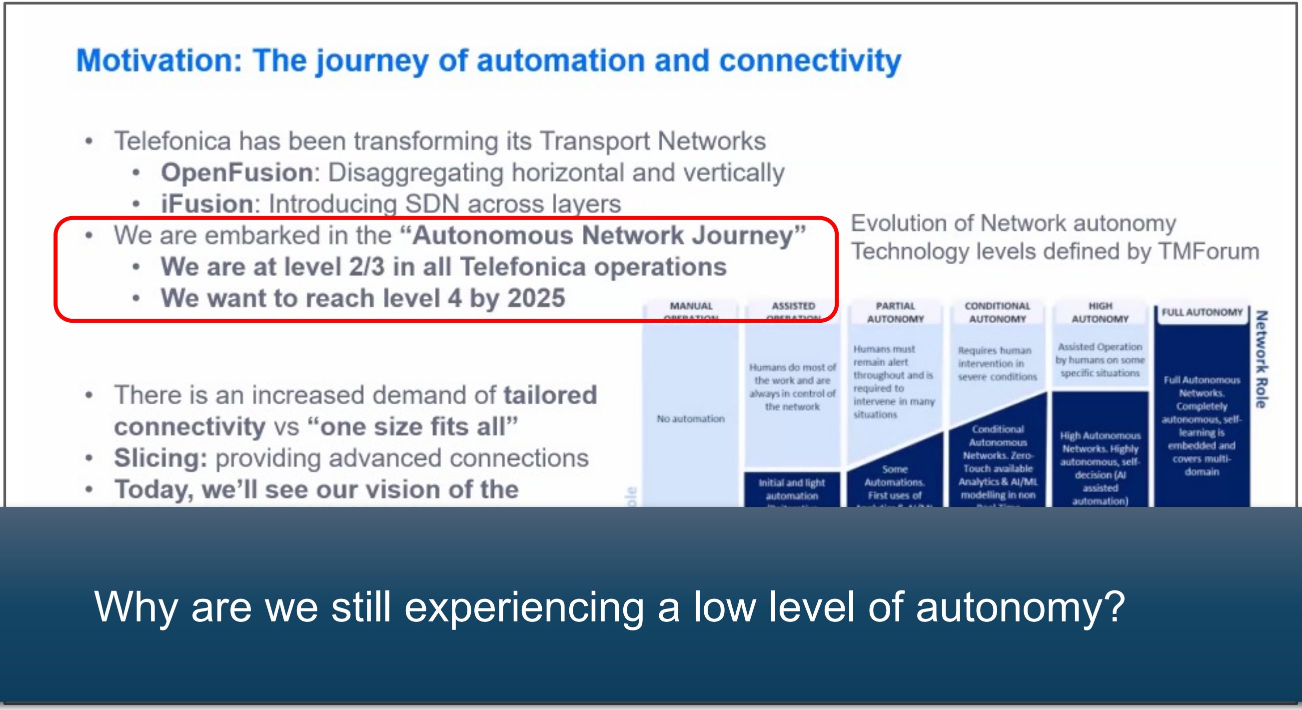


Why are we still experiencing a low level of autonomy?

* Oscar Gonzáles de Dios, "AI-based automation of multi-layer multi-domain transport networks," OFC, pp. W4I.2, March 2024.

# Outline

- Network automation & programmability
- **Trustworthy AI**
- Understanding reinforcement learning
- Representative use case
- Preliminary investigation
- Challenges and opportunities

# Trustworthy AI

1 Human agency and oversight
- Including fundamental rights, human agency and human oversight

2 Technical robustness and safety
- Including resilience to attack and security, fall back plan and general safety, accuracy, reliability and reproducibility

3 Privacy and data governance
- Including respect for privacy, quality and integrity of data, and access to data

4 Transparency
- Including traceability, explainability and communication

5 Diversity, non-discrimination and fairness
- Including the avoidance of unfair bias, accessibility and universal design, and stakeholder participation

6 Societal and environmental wellbeing
- Including sustainability and environmental friendliness, social impact, society and democracy

7 Accountability
- Including auditability, minimisation and reporting of negative impact, trade-offs and redress.

Ethics guidelines for trustworthy AI, EU, April 2019.

# Outline

- Network automation & programmability
- Trustworthy AI
- **Understanding reinforcement learning**
- Representative use case
- Preliminary investigation
- Challenges and opportunities

# The reinforcement learning loop

**RL framework**

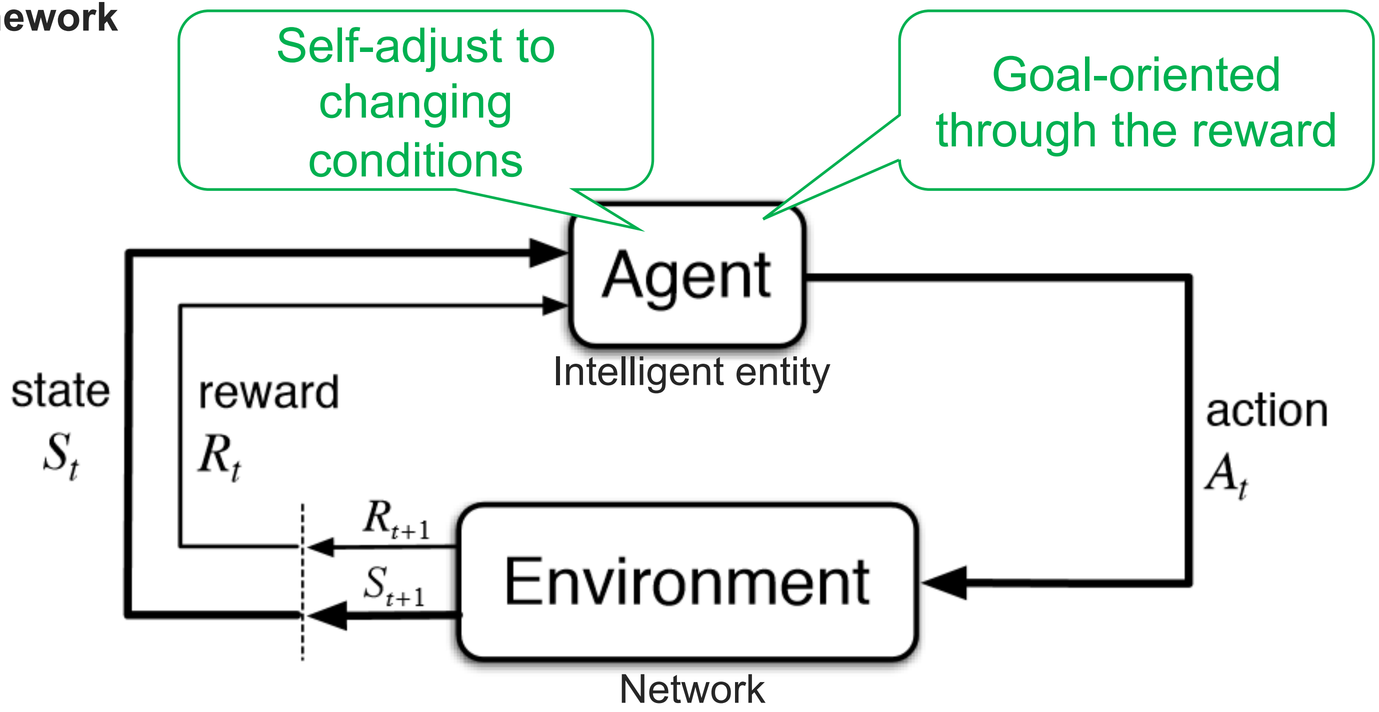

# Discounter reward

## RL framework

- Long-term value of each state/action pair
- Term $\gamma$ regulates the value of future rewards to the current action
- Number of interactions depends on the *episode length*

$$R_t = \sum_{k=0}^{\infty} \gamma^k \times R_{t+k+1}$$

$$R_t = R_{t+1} + \gamma \times R_{t+2} + \gamma^2 \times R_{t+3} + \cdots$$

# RL vs. DRL

## RL framework

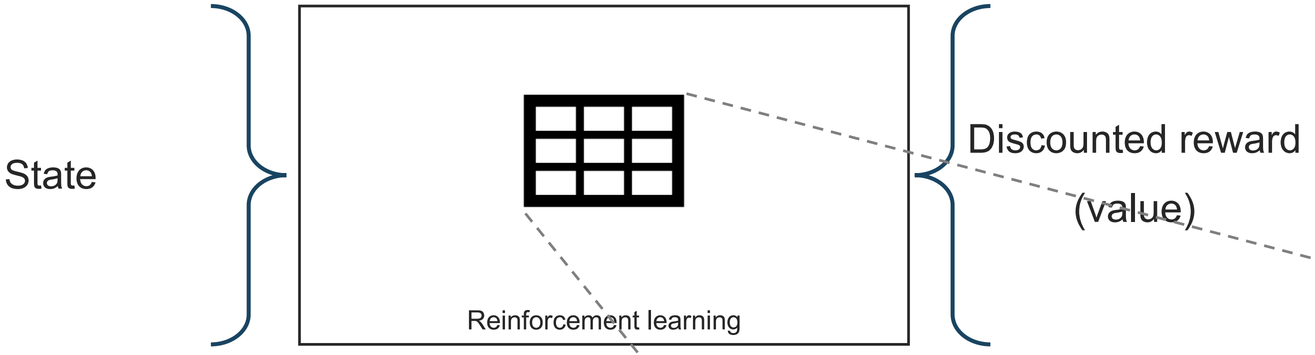


| State | Action 1 | Action 2 | … | Action $\|A\|$ |
|---|---|---|---|---|
| **(0, 0, …, 0)** | 10 | 5 | … | 1 |
| **(0, 0, …, 1)** | 3 | 8 | … | 4 |
| **…** | … | … | … | … |
| **(1, 1, …, 1)** | 6 | 2 | … | 11 |

# RL vs. DRL

## RL framework

State

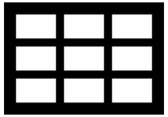

Reinforcement learning

Discounted reward

(value)

State

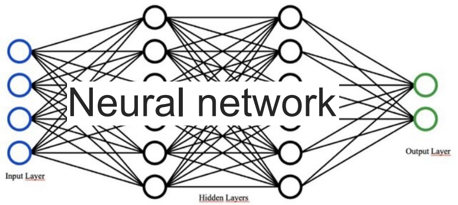


Deep reinforcement learning

Discounted reward (value)

or

Action probability (policy)

# Inference

## RL framework

**State**

- Source
- Destination
- Bit rate
- Block spectrum information
- Path spectrum information
- Etc.

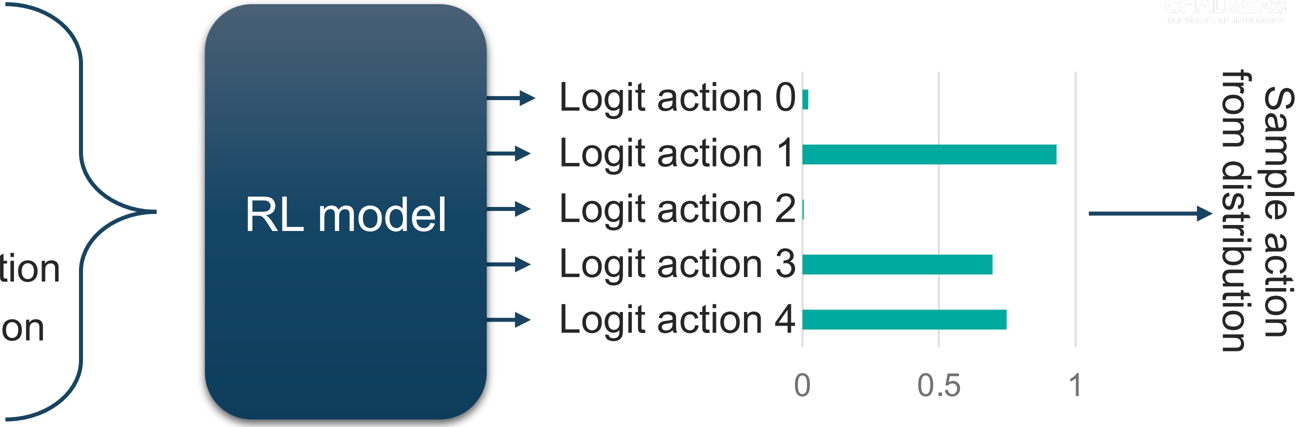

# How does RL learn?

## RL framework

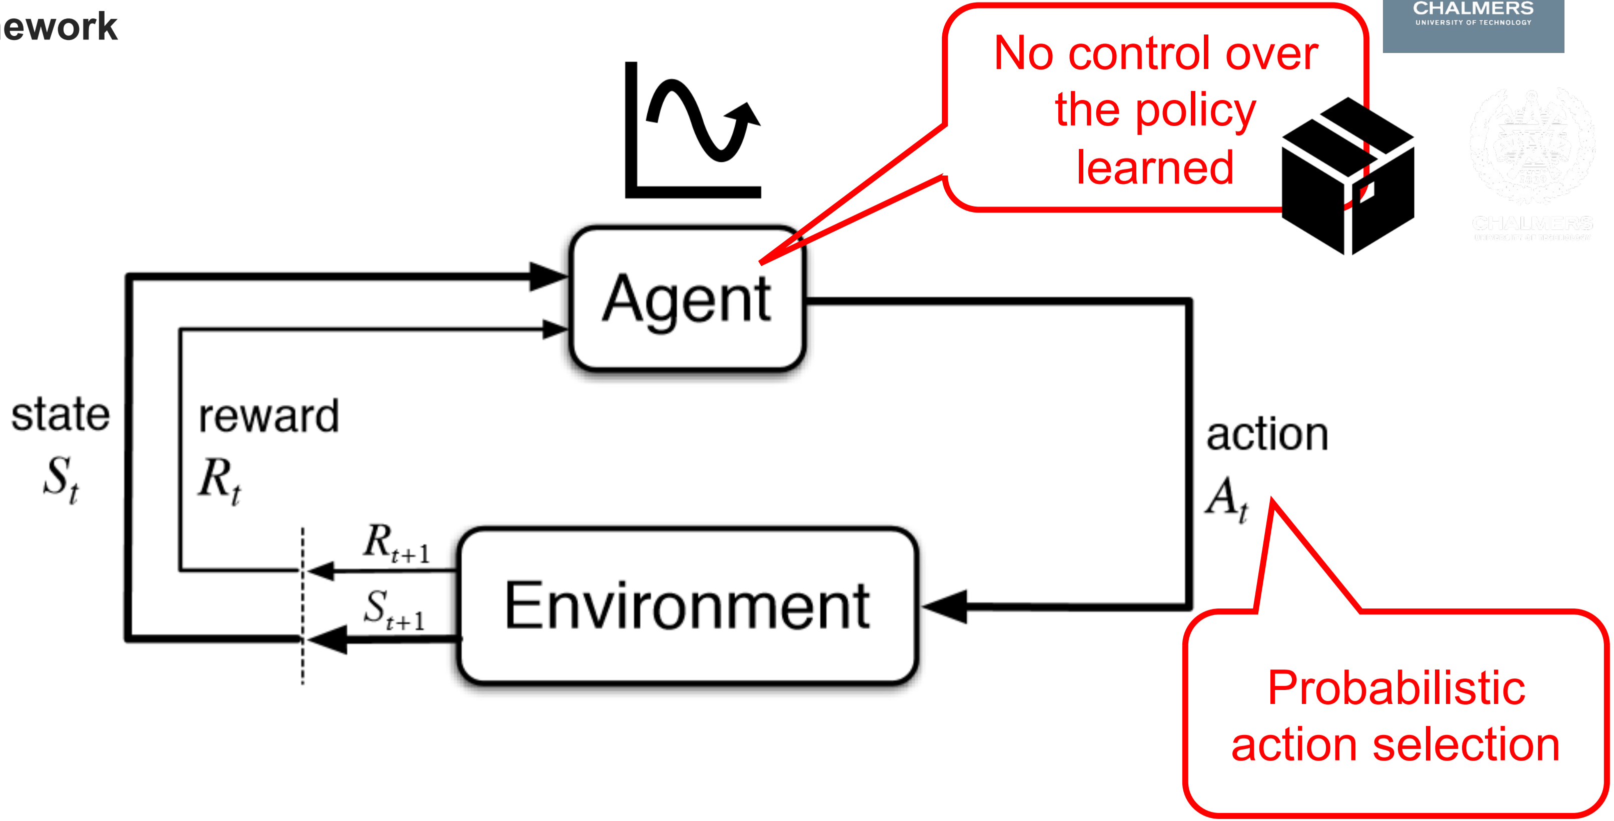

# Outline

- Network automation & programmability
- Trustworthy AI
- Understanding reinforcement learning
- **Representative use case**
- Preliminary investigation
- Challenges and opportunities

# Use cases for reinforcement learning

Suitable use cases

- Complex problems
  - Scenario-specific behavior
  - Need for context-specific thresholds
- Unknown future
  - Random arrivals
  - External events

Unsuitable use cases

- Simple problems
  - Heuristics are sufficient
- Known sequence of steps
  - Concurrent optimization (e.g., linear programming) is a better fit

# Problem statement

## Routing, modulation, core, and spectrum allocation

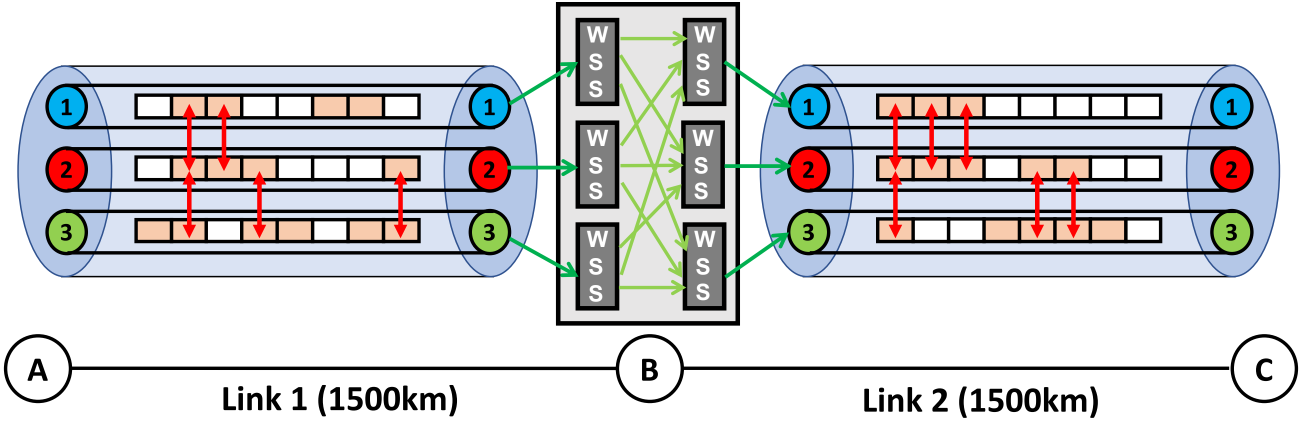


*Rt (A, C, 35 gbit/s)*

QPSK *s* = 3 slots → Available Core Path

Route: *A-B-C*

CP: [ $\Phi_1^2, \Phi_2^1$ ]

*Core2* on *Link1*

*Core1* on *Link2*

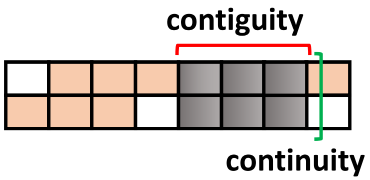


free slot

occupied slot

assigned slot

core switching

crosstalk

Teng et al., JOCN, 2024

# RL modeling

## Routing, modulation, core, and spectrum allocation

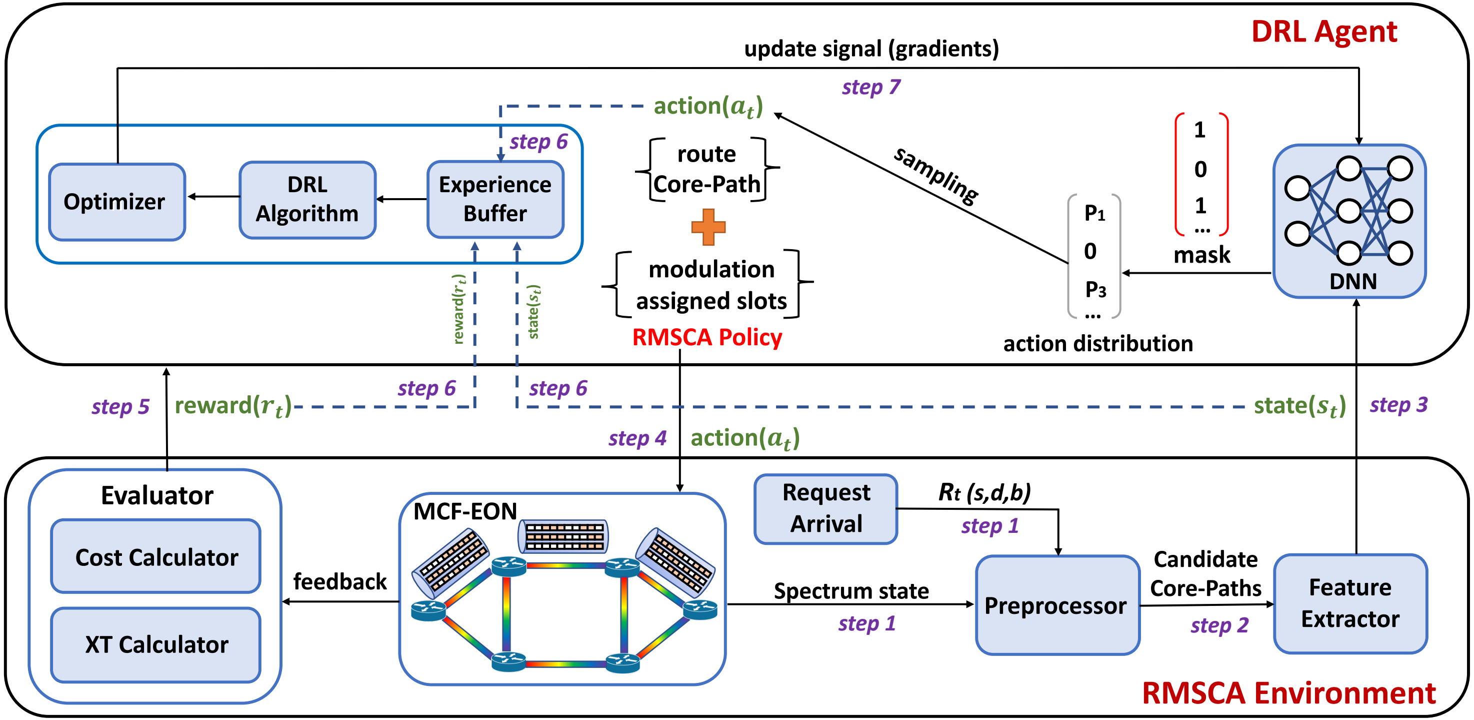


Teng et al., JOCN, 2024

# Representative results

## Routing, modulation, core, and spectrum allocation

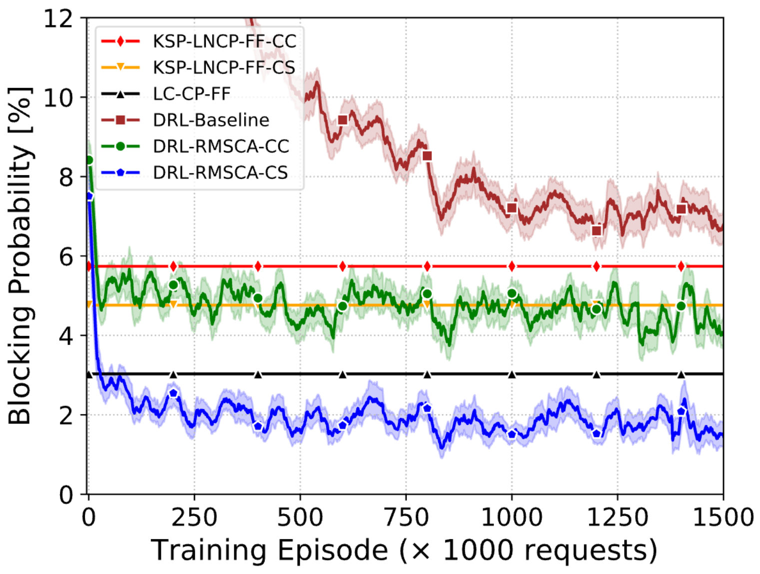


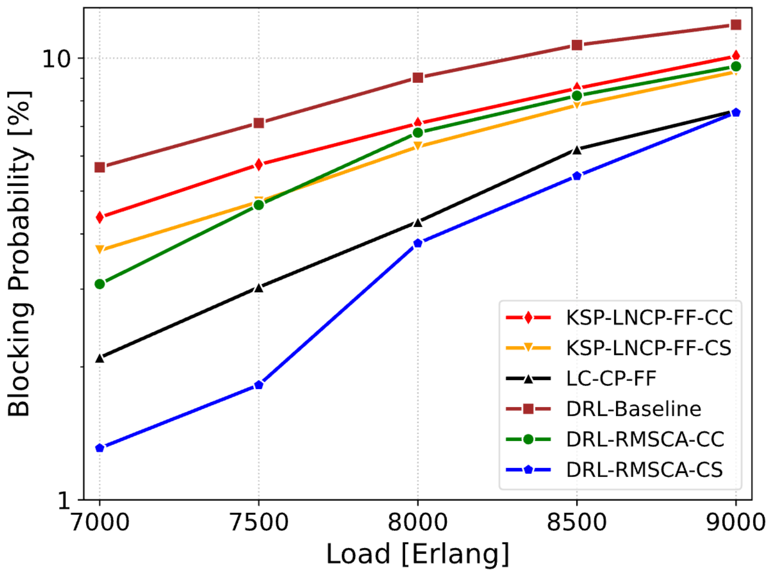


Teng et al., JOCN, 2024

# Outline

- Network automation & programmability
- Trustworthy AI
- Understanding reinforcement learning
- Representative use case
- **Preliminary investigation**
- Challenges and opportunities

# How does RL learn?

## RL framework

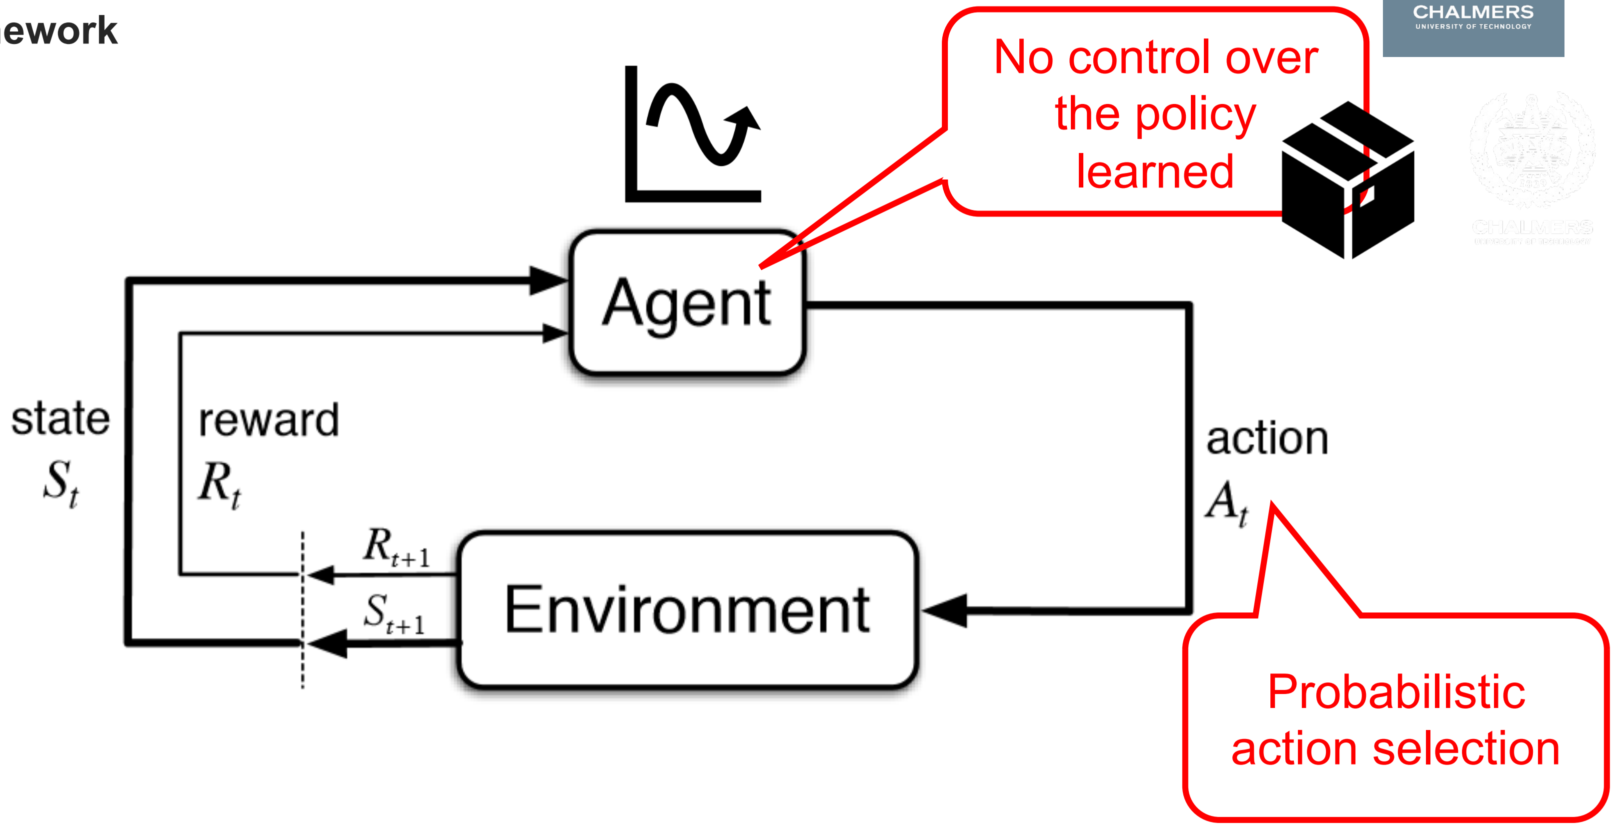

# Research questions

## RL framework

**State**

- Source
- Destination
- Bit rate
- Block spectrum information
- Path spectrum information
- Etc.

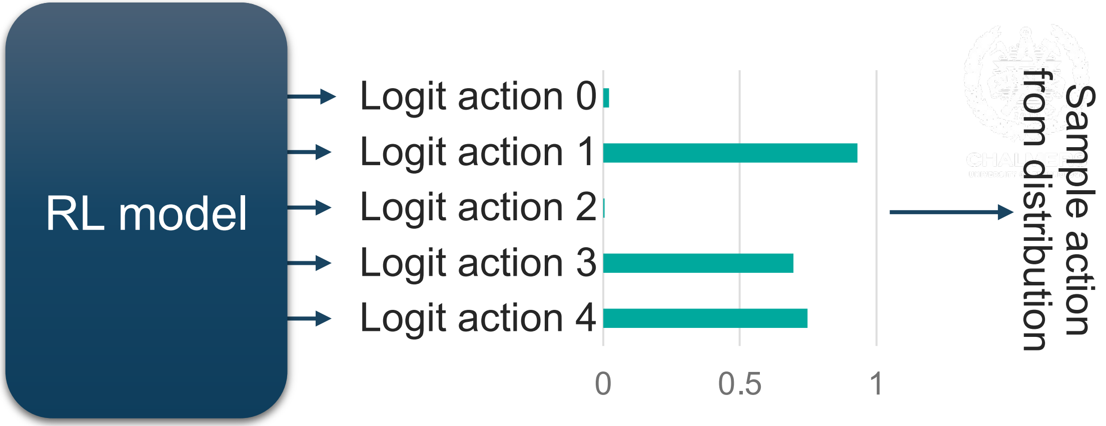


Q1: How do features impact the RL decision?
Q2: How does the impact of features change for different RL scenarios?
Q3: Can we identify unwanted behavior?

Ayoub et al., OFC’24, p. W4I.6

# Proposed approach

Train RL → Learn classifier → Assess feature impact

Ayoub et al., OFC’24, p. W4I.6

# Proposed approach

Train RL → Learn classifier → Assess feature impact

**Train RL**

- DeepRMSA*
- Maximize acceptance / minimize blocking

**Learn classifier**

- ~ to imitation learning
- Choose the most accurate model
- Accuracy assessed as classification

**Assess feature impact**

- Use Shapley Additive exPlanations (SHAP)

SHAP

* Chen et al., JLT 37 (16), 2019.
Ayoub et al., OFC'24, p. W4I.6

# RF classifier performance

## Results

| Model | Infeasible allowed | Infeasible penalized | Infeasible masked |
|---|---|---|---|
| **F1-score** | 0.75-0.85 | 0.74-0.84 | 0.77-0.84 |
| **Accuracy** | 0.84-0.94 | 0.77-0.91 | 0.84-0.94 |
| **Precision (class 5)** | 0.75-1.00 | 0.67-0.90 | 0.75-1.00 |

Fairly uniform accuracy across scenarios

Ayoub et al., OFC’24, p. W4I.6

# Feature importance

## Results

**Features:**

**Request:**

0. Bit rate
1. Source node
2. Destination node

**Path/block option:**

**Block:**

3. Initial index
4. Number of free slots

**Path:**

5. Number of requested slots
6. Total free slots in the path
7. Average free block length

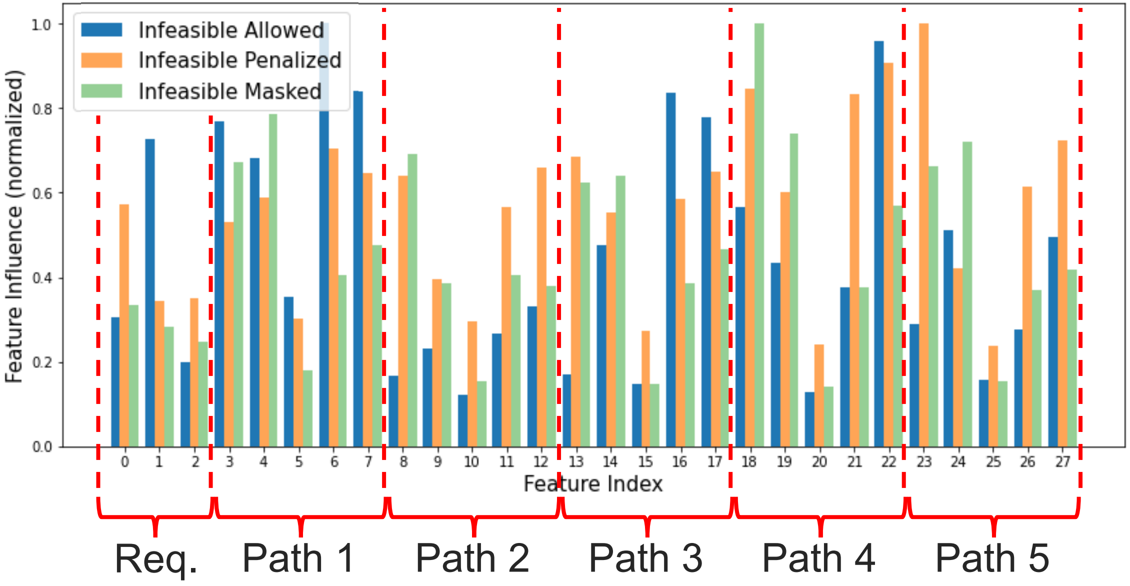


Ayoub et al., OFC'24, p. W4I.6

# Feature importance

## Results

**Features:**

**Request:**
0. Bit rate
1. Source node
2. Destination node

**Path/block option:**

**Block:**

3. Initial index
4. Number of free slots

**Path:**

5. Number of requested slots
6. Total free slots in the path
7. Average free block length

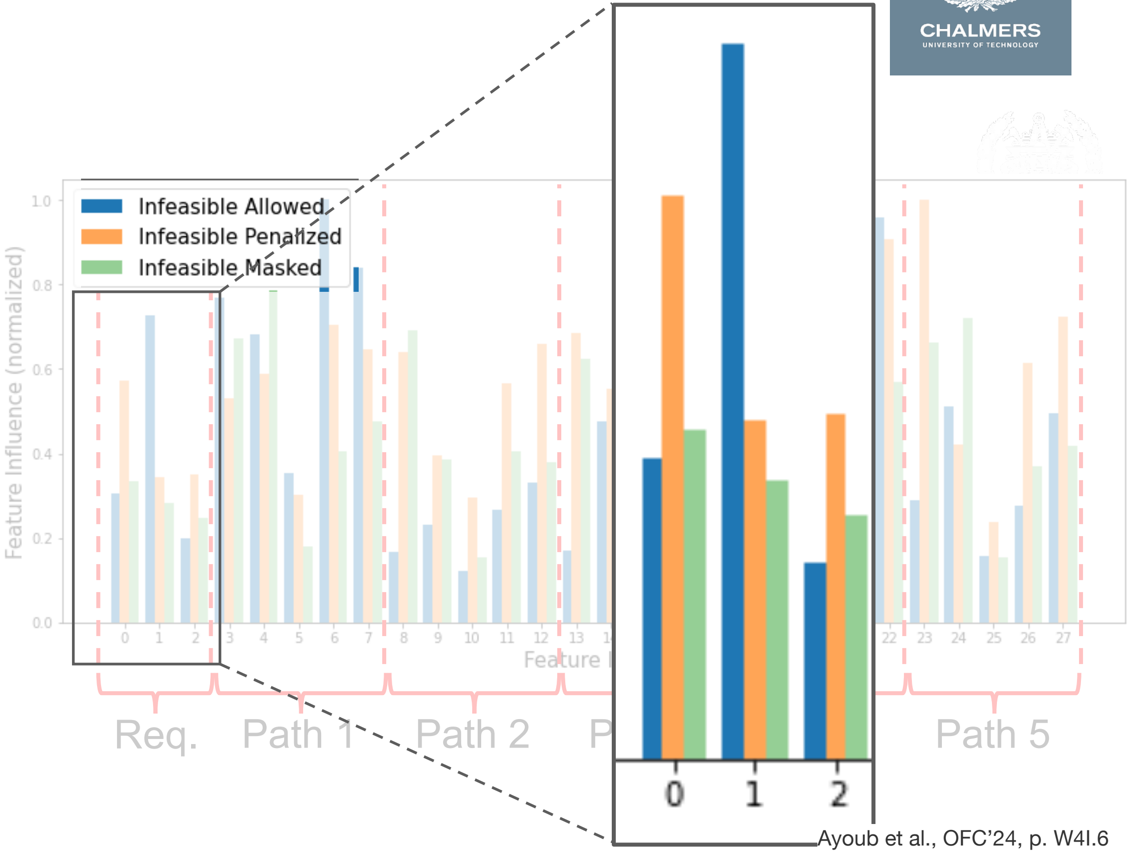


Ayoub et al., OFC'24, p. W4I.6

# Feature influence/impact on rejection

## Results

Low values in path 0 influence rejection

High values in paths 2 and 3 also influence rejection

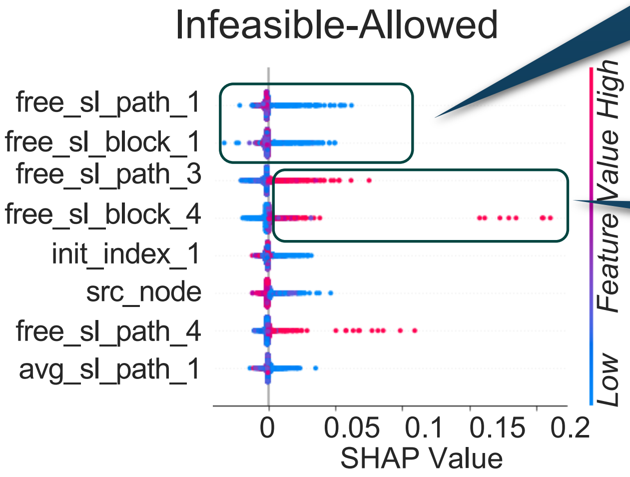


Ayoub et al., OFC’24, p. W4I.6

# Feature influence/impact on rejection

**Results**

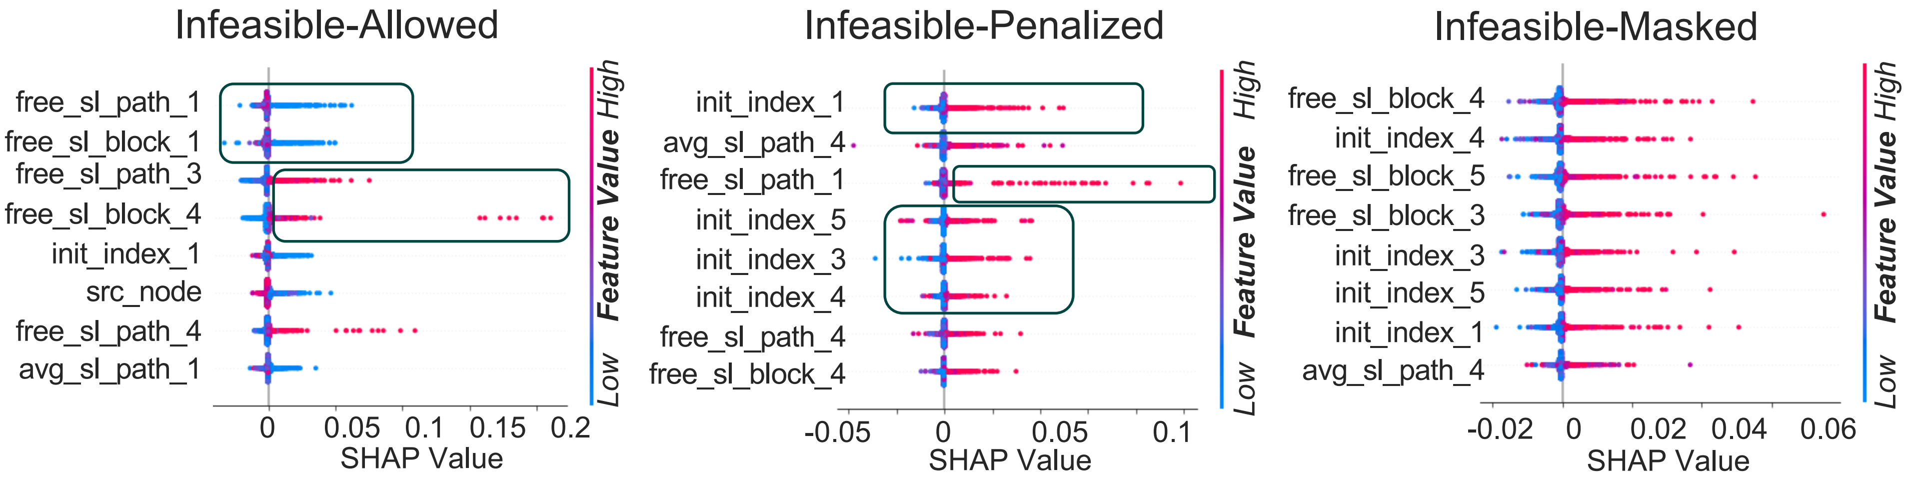


Ayoub et al., OFC’24, p. W4I.6

# Outline

- Network automation & programmability
- Trustworthy AI
- Understanding reinforcement learning
- Representative use case
- Preliminary investigation
- **Challenges and opportunities**

# Challenges and opportunities

1 Human agency and oversight

- Including fundamental rights, human agency and *human oversight*

2 Technical robustness and safety

- Including resilience to attack and security, fall back plan and general safety, *accuracy*, *reliability* and *reproducibility*

3 Privacy and data governance

- Including respect for *privacy*, *quality and integrity of data*, and *access to data*

4 Transparency

- Including *traceability*, *explainability* and communication

5 Diversity, non-discrimination and fairness

- Including the avoidance of unfair *bias*, accessibility and universal design, and stakeholder participation

6 Societal and environmental wellbeing

- Including *sustainability* and environmental friendliness, social impact, society and democracy

7 Accountability

- Including *auditability*, minimisation and reporting of negative impact, trade-offs and redress.

Ethics guidelines for trustworthy AI, EU, April 2019.

# Acknowledgements

- Paolo Monti
- Yiran Teng
- Shuangyi Yan
- Anders Lindgren
- Stefan Melin
- Achim Autenrieth
- Wolfgang John
- Ali Balador

- Celtic-Next project AI-NET-PROTECT
- ECO-eNET

# References and further reading


- Y. Teng *et al.*, "Deep-reinforcement-learning-based RMSCA for space division multiplexing networks with multi-core fibers," in *Journal of Optical Communications and Networking*, vol. 16, no. 7, pp. C76-C87, July 2024, doi: 10.1364/JOCN.518685.
- O. Ayoub, C. Natalino and P. Monti, "Towards Explainable Reinforcement Learning in Optical Networks: The RMSA Use Case," *Optical Fiber Communications Conference and Exhibition (OFC)*, San Diego, CA, USA, 2024, pp. W4I.6.
- E. Etezadi et al., “Deep reinforcement learning for proactive spectrum defragmentation in elastic optical networks,” in Journal of Optical Communications and Networking, vol. 15, no. 10, pp. E86-E96, October 2023. DOI: 10.1364/JOCN.489577.
- Wolfgang John, “The journey towards 6G: Going beyond connectivity services,” IEEE NetSoft, Madrid, Spain, June 2023.
- Achim Autenrieth, “Carrier Grade AI/ML for Network Automation”, invited talk, OFC 2022, 9 March 2022.
- Whitepaper, “European Vision for the 6G Network Ecosystem,” 5GPPP, 2021. DOI: 10.5281/zenodo.5007671.
- Whitepaper, “Ethics guidelines for trustworthy AI,” EU, April 2019. URL: https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai


*In chronological order

# Thank you! ☺

Chalmers profile

GitHub page

IFIP Networking 2024

TX4Nets 2024

CHALMERS
UNIVERSITY OF TECHNOLOGY

# Explainable Reinforcement Learning: Towards Trustworthy Autonomous Network Operations

Carlos Natalino
Researcher, Optical Networks Unit
Department of Electrical Engineering
Chalmers University of Technology, Gothenburg, Sweden

CHALMERS

UNIVERSITY OF TECHNOLOGY