Institute for the Wireless Internet of Things at Northeastern University

institute IMDEA networks

EXPLORA: AI/ML *EXPL*ainability for the Open *RAN*

Claudio Fiandrino[1], Leonardo Bonati[2], Salvatore d'Oro[2], Michele Polese[2], Tommaso Melodia[2], Joerg Widmer[1]

[1]IMDEA Networks Institute

[2]Institute for the Wireless Internet of Things, Northeastern University

CoNEXT 2023

Artifacts Available V1.1 · Artifacts Evaluated Functional V1.1 · Results Reproduced V1.1

[Developing the Science of Networks]

A Primer on AI Explainability

- DARPA introduced XAI to make accessible an interpretable the internal logic and the outcome of AI models

Today

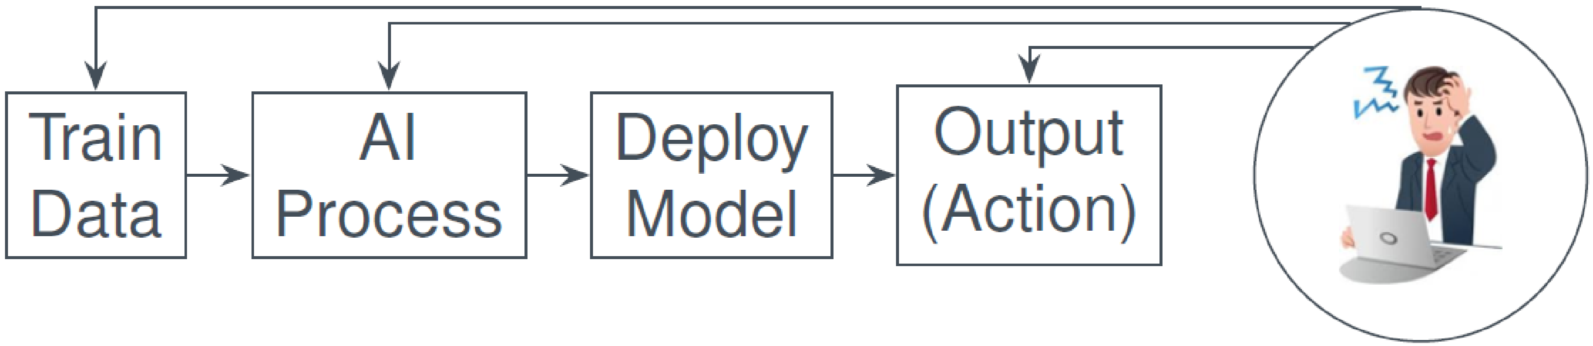

- Why did you do that?
- When do you succeed/fail?
- When can I trust you?
- How do I correct errors?

With XAI

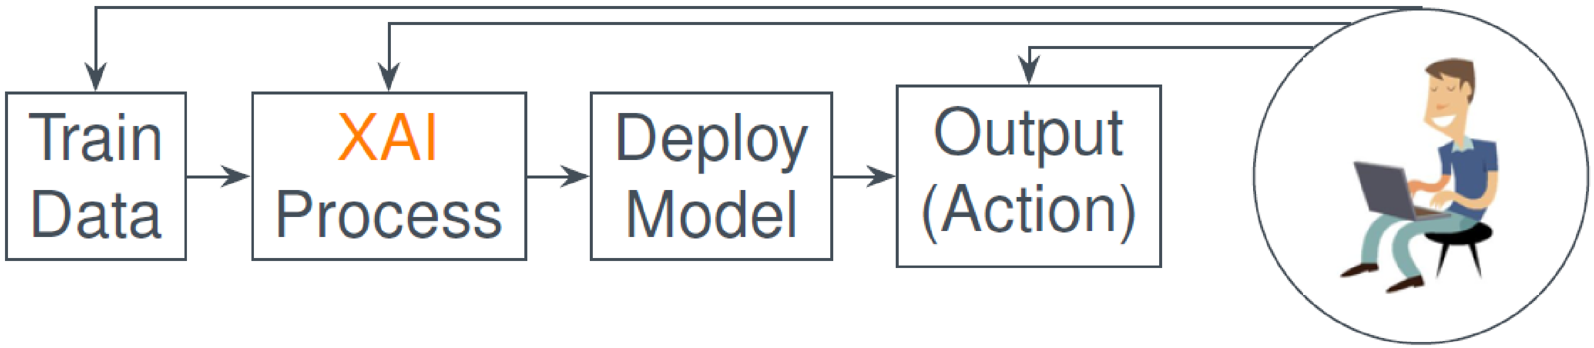

- I understand why
- I know why you succeed/fail
- I know when to trust
- I know when you err

Resource Allocation in Network Slicing

- The agent decides per slice Physical Resource Block allocation (PRB) and Scheduling Policy (e.g., proportional fair)
- Slices: eMBB, eMTC, URLLC; Agents: high throughput (HT), low latency (LL)
- Users report several KPIs (e.g., buffer status, transmitted bit rate/packets, etc..)

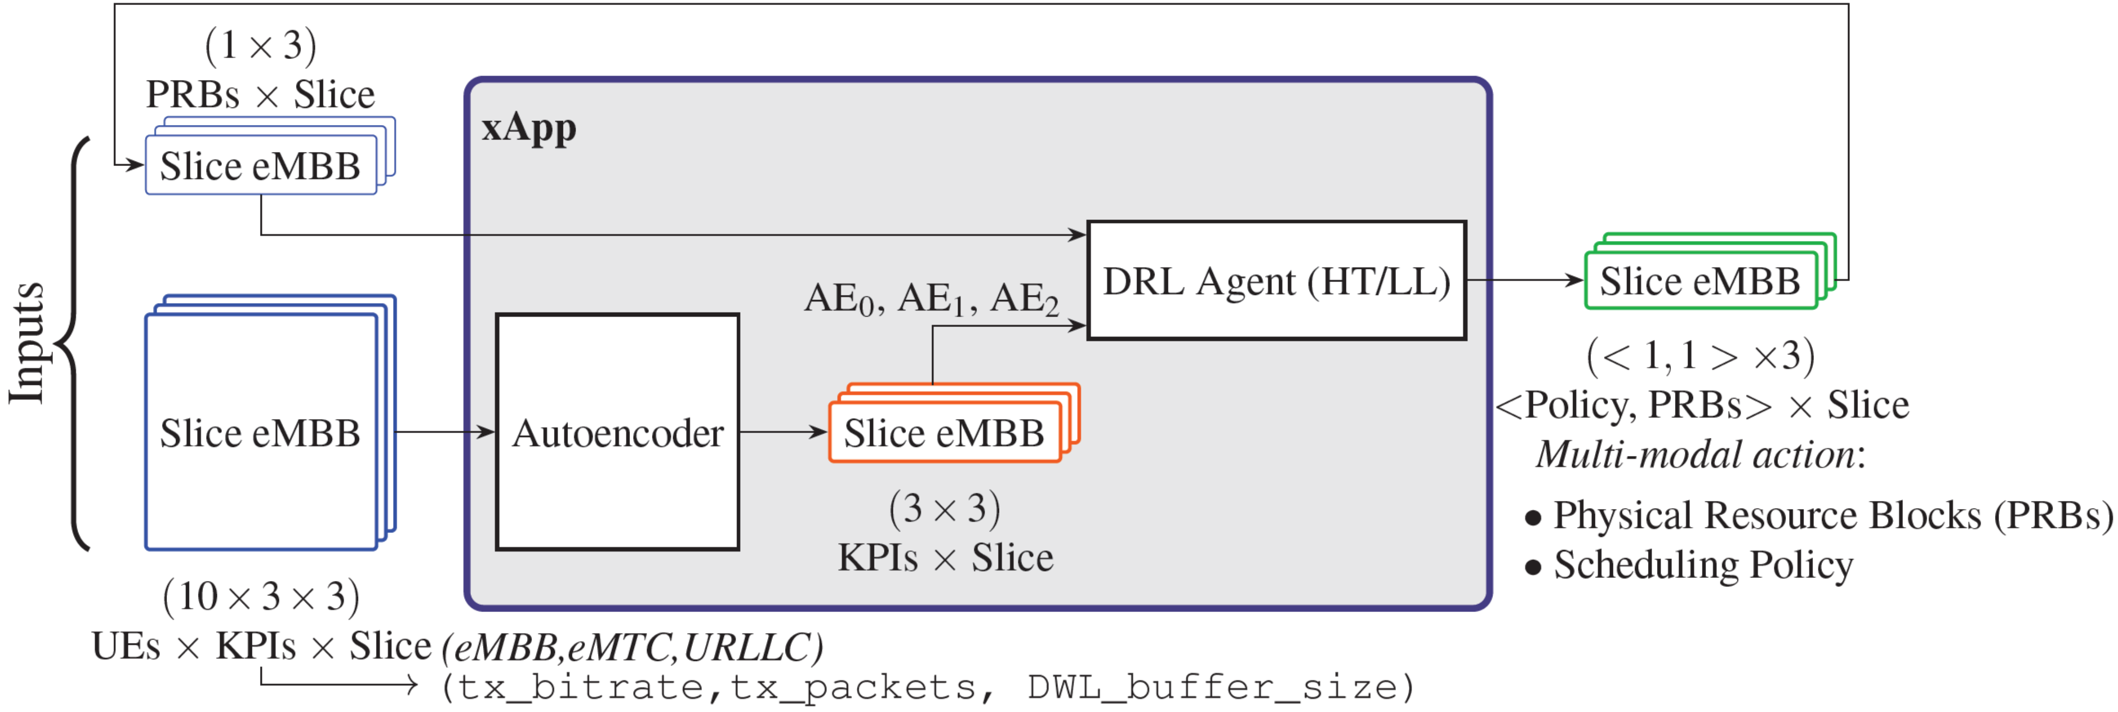

Explaining Deep Reinforcement Learning Agents

- Unique characteristics of Deep Reinforcement Learning (DRL):
  - Learning actions (finite, possibly large, set)
  - Inhomogenous data (inputs and outputs are not necessarily the same piece of information)
- We asked: "would surrogate self-interpretable models suffice to explain a complex agent?"
  - For this, we train decision trees onto tabular data (one per slice, one on all slices)
- How does this method compare to the case of applying known XAI tools directly to the agent?
  - We look for a semantic equivalence between these two types of explanations

Limitations of the State-of-the-Art

- We analyze known-techniques when applied to the model
  - SHAP
  - Decision trees

Both apply to the agent, hence do not provide feature relevance of the input, but on the outputs of the autoencoder

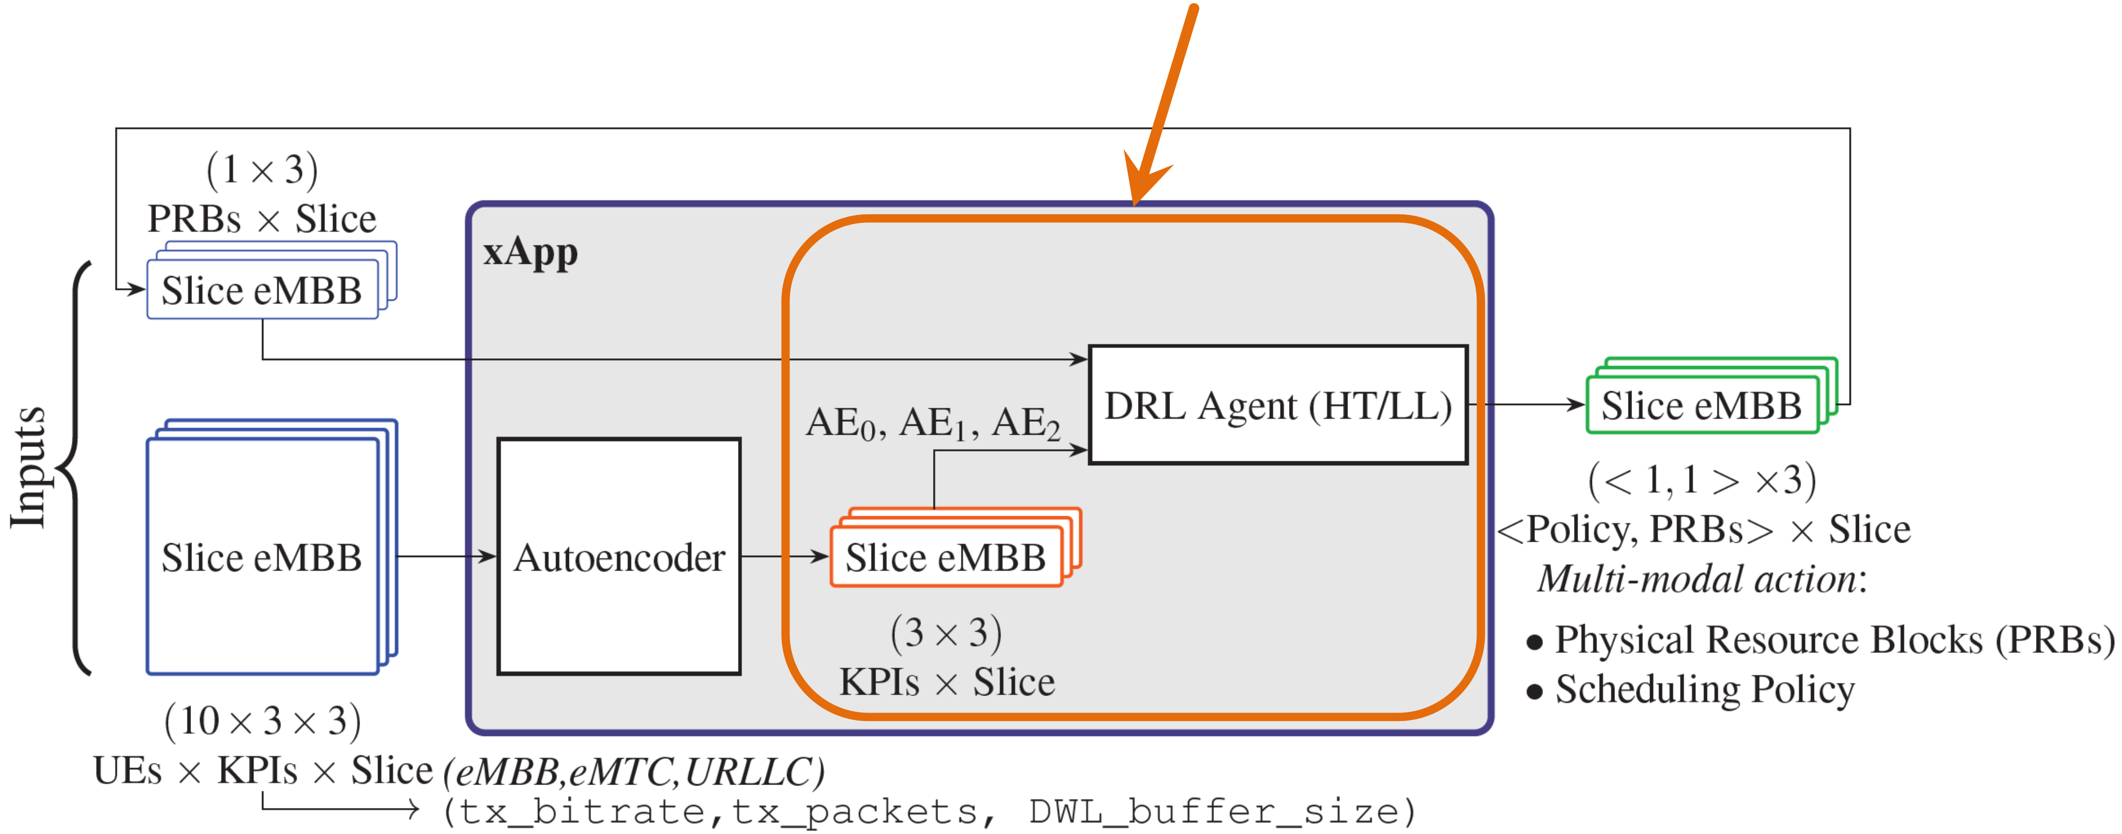

Limitations of the State-of-the-Art: SHAP

- SHAP: a game theory-based approach for ranking feature relevance
- Explanations: "*low importance prior change of PRB allocation*"

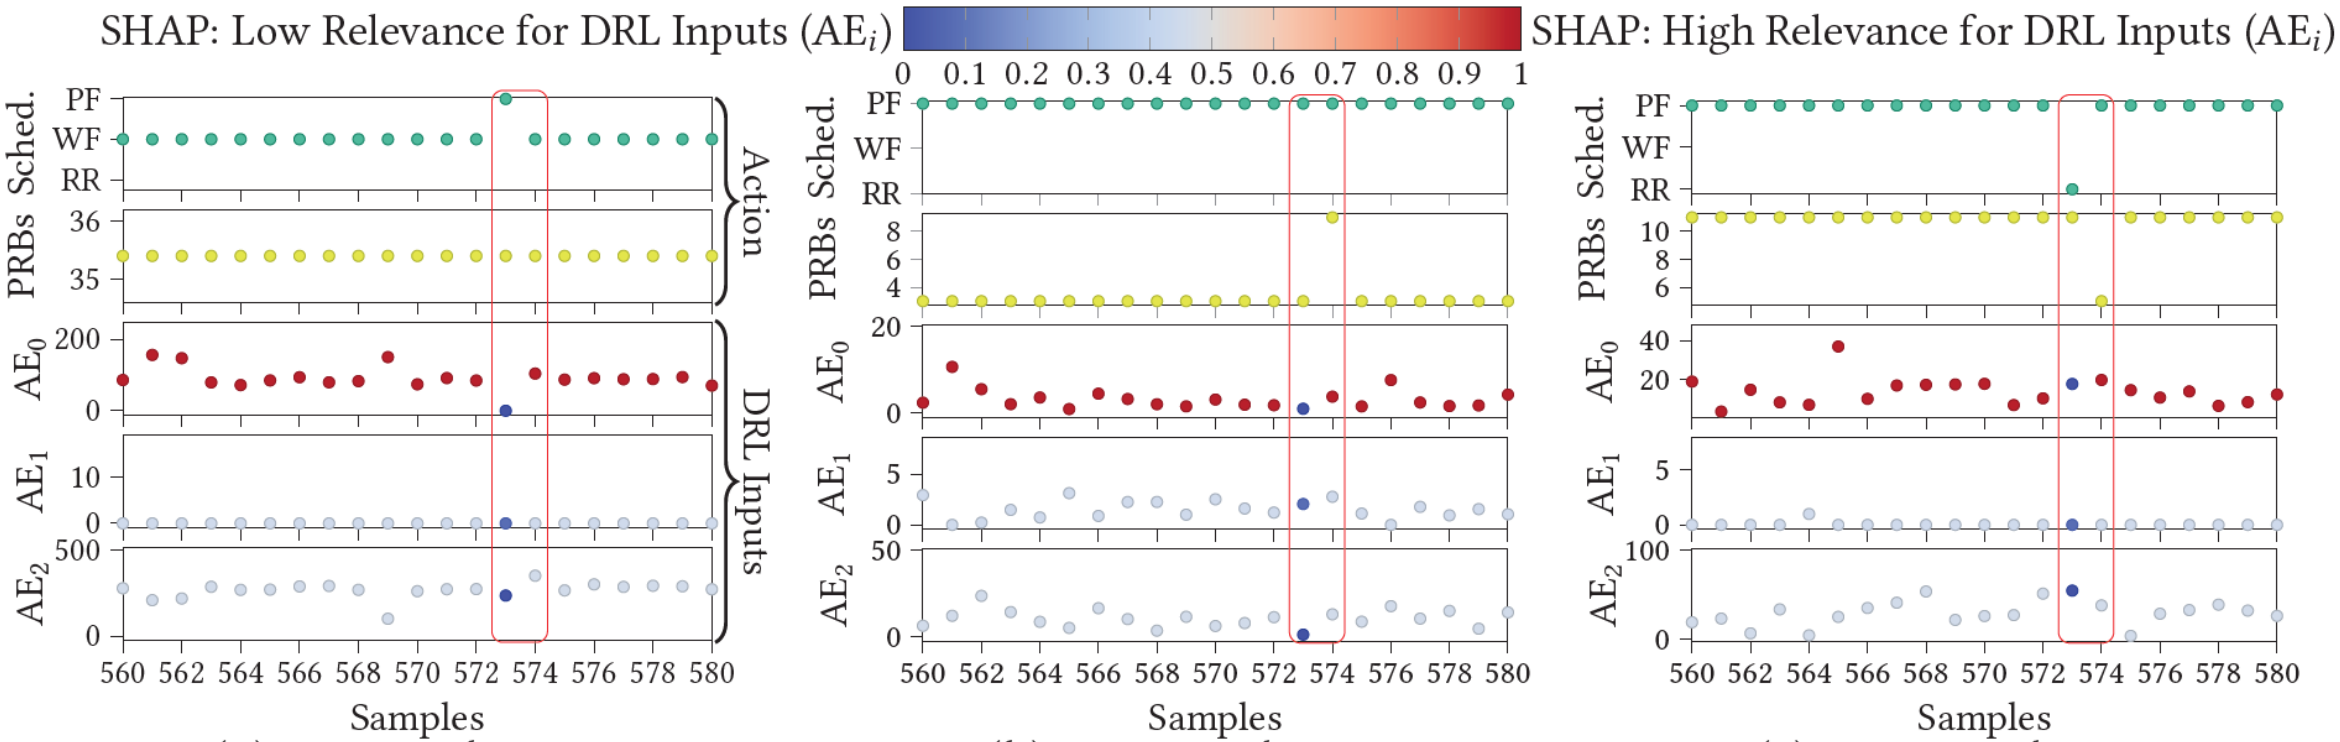

(a) eMBB slice (b) mMTC slice (c) URLLC slice

Limitations of the State-of-the-Art: SHAP (II)

- SHAP is extremely costly from a computational perspective when applied to DRL agents

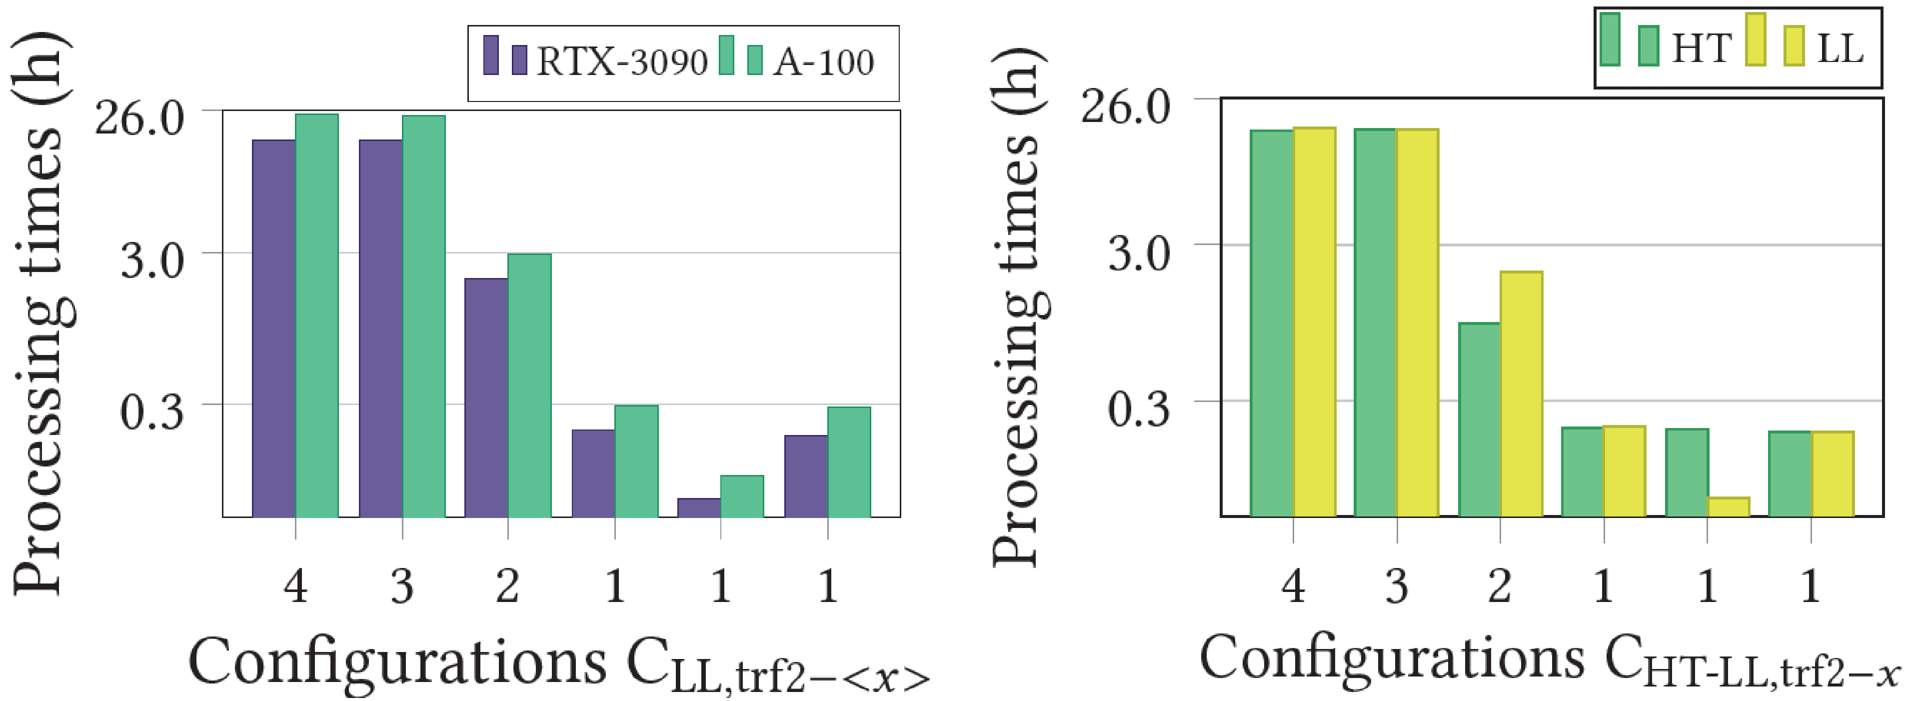

(a) On different machines (b) On different agents

The Challenges that We Solve

- CH#1: Autoencoders break direct *input-output* connection
  - Known techniques like SHAP or LRP only reveal importance of the latent representation, *not the actual input*
- CH#2: If actions depend on either past actions or states of the environment like past PRB allocation, the process has memory
  - Primary cause and primary effect of an action unclear: can not use casual learning models
- CH#3: Actions may be multi-modal, i.e., one decision controls several parameters
  - Hard to understand the effect of each mode

Our Approach for Explainability: EXPLORA

- EXPLORA is a framework: ❶ explains DRL agent behavior, ❷ can refine action decision based on intents
- We use *attributed graphs* as data structure for the explanations:
  - Nodes are actions and attributes are the monitored effect of the action
  - We study the transitions to synthesize the explanations

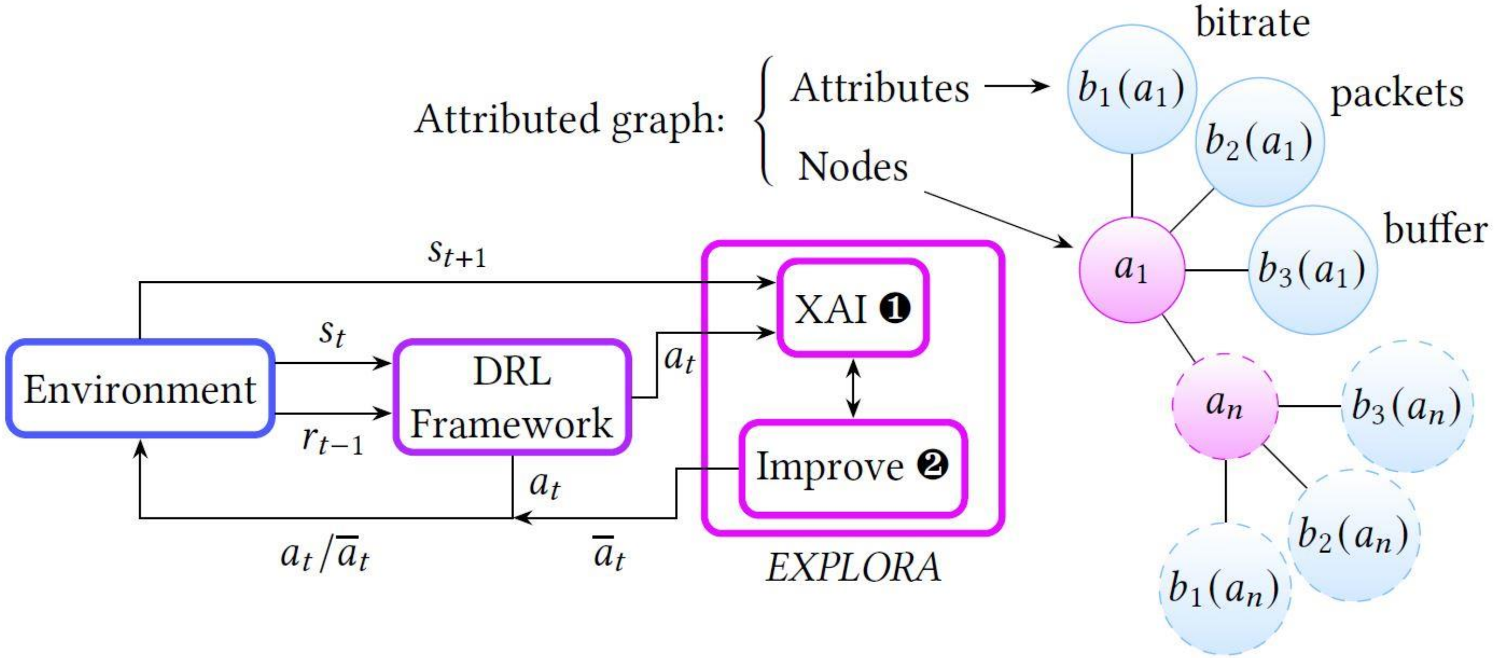

Building the Attributed Graph

- The workflow for three time-steps (t_0, t_1, t_2) and 2 users
  - [36,3,11] => PRBs per slice (sum to 50)
  - [1,2,2] => scheduling policy

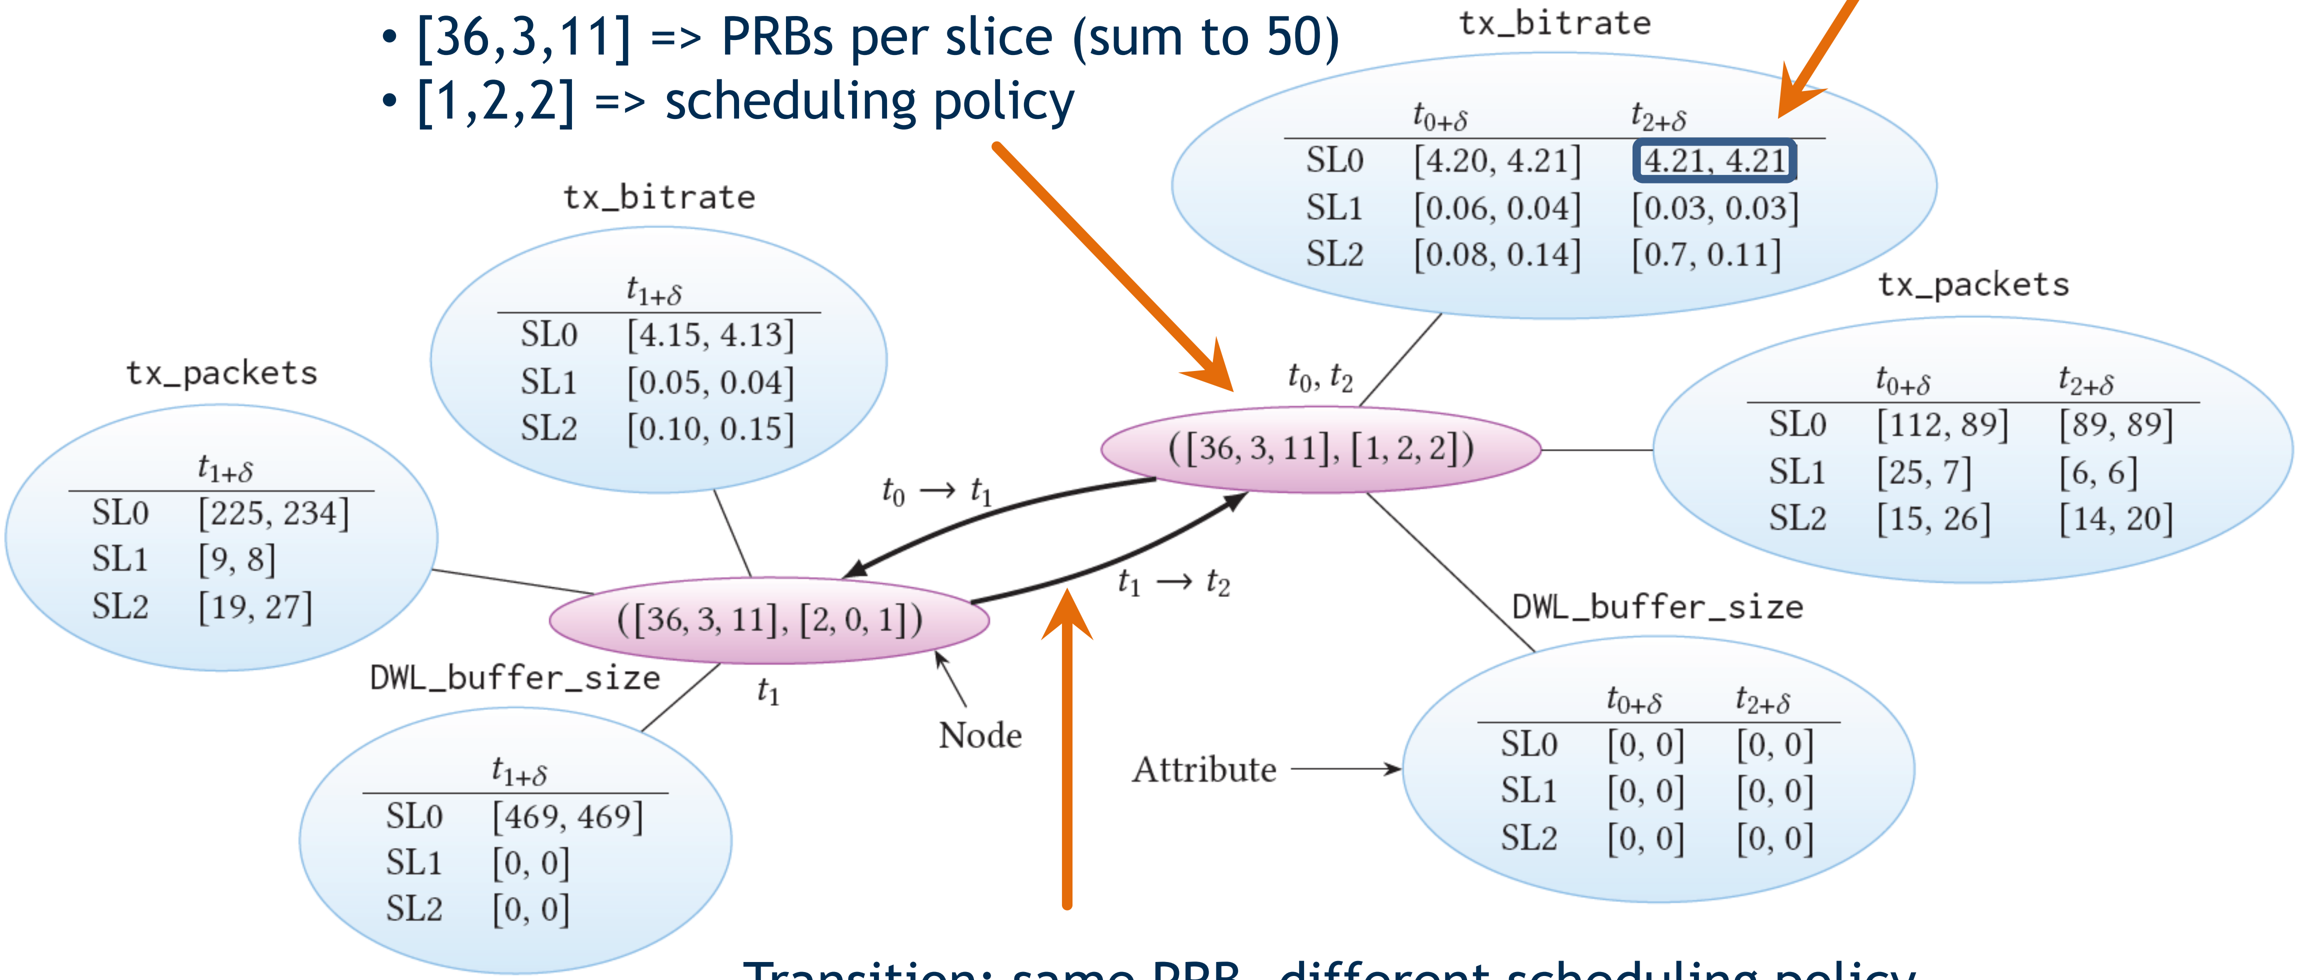

Transition: same PRB, different scheduling policy

An Example of a Graph

- The multi-modal action gives rise to classes of transitions

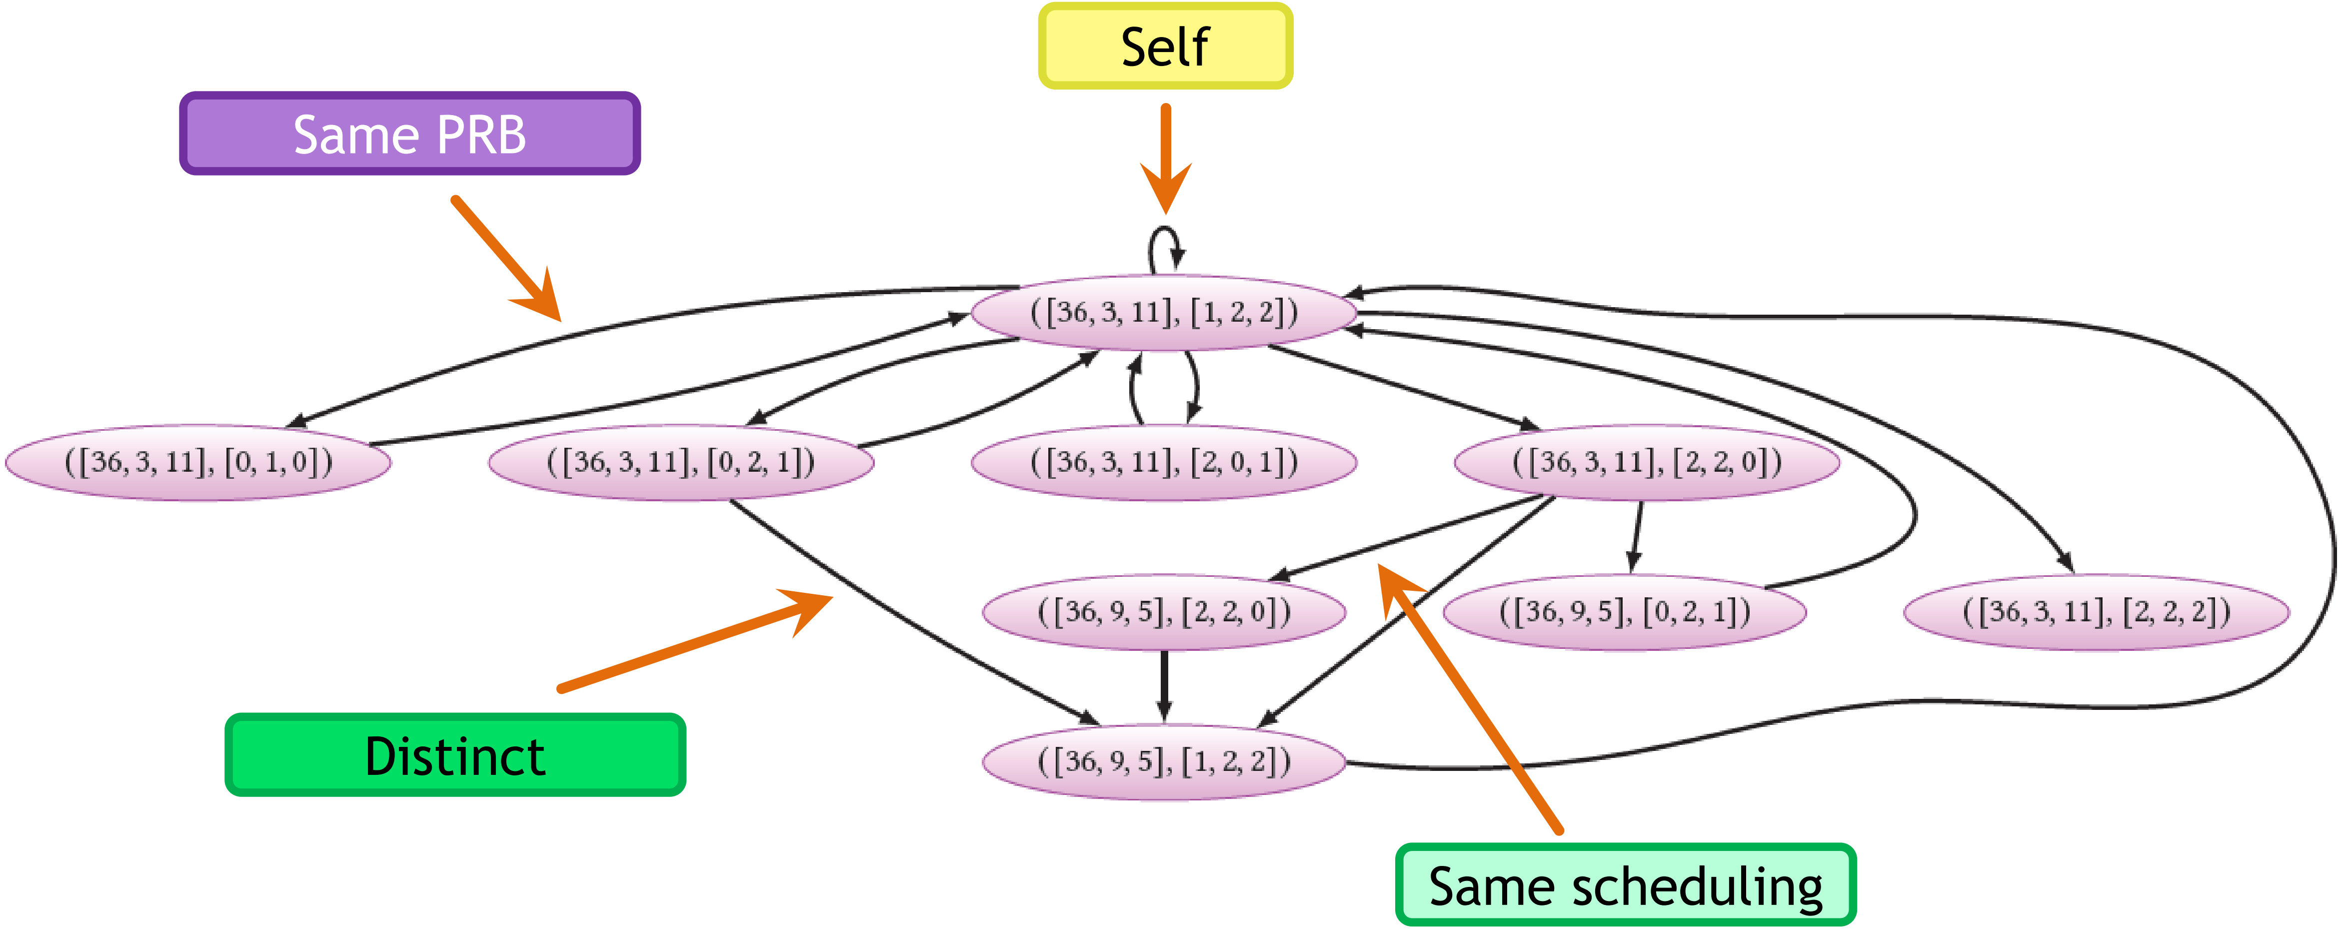

Implementing EXPLORA

- We designed a *O-RAN compatible solution*: DRL agent and EXPLORA are xApps running on the near RealTime Radio Intelligent Controller
- Distribute Units (O-DU) monitor and stream KPIs to the near RT-RIC through E2 interface. DRL agent takes actions, sends to the RIC Message Router (RMR) who streams to E2 such control enforced at O-DU

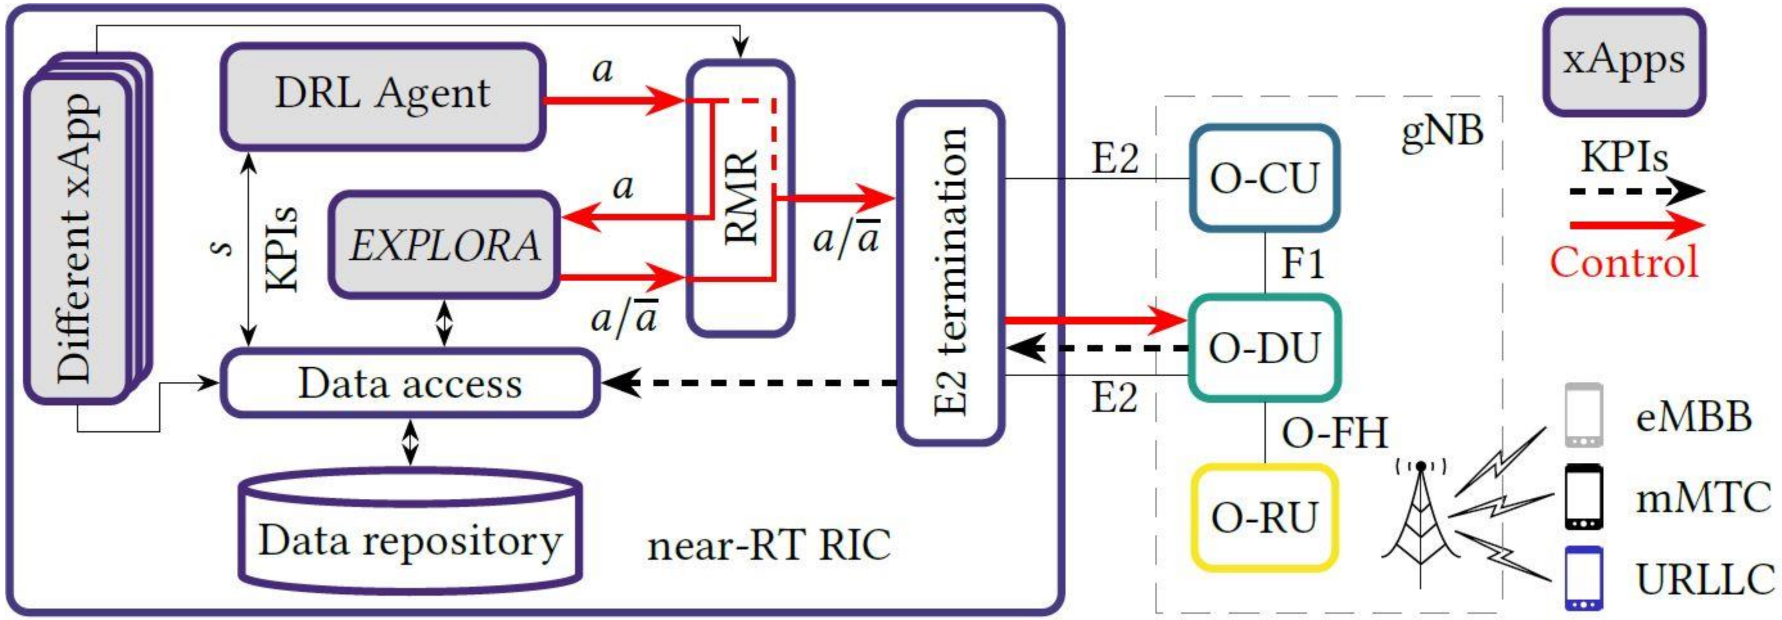

Understanding Agents

- We study the transitions of the attributed graph:
  - Comparing the distributions of classes of actions per slice
  - Building a decision tree

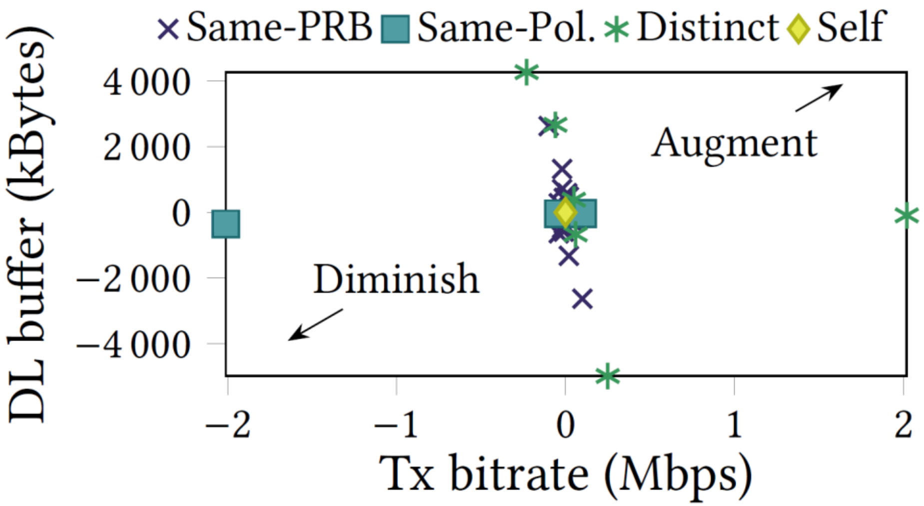

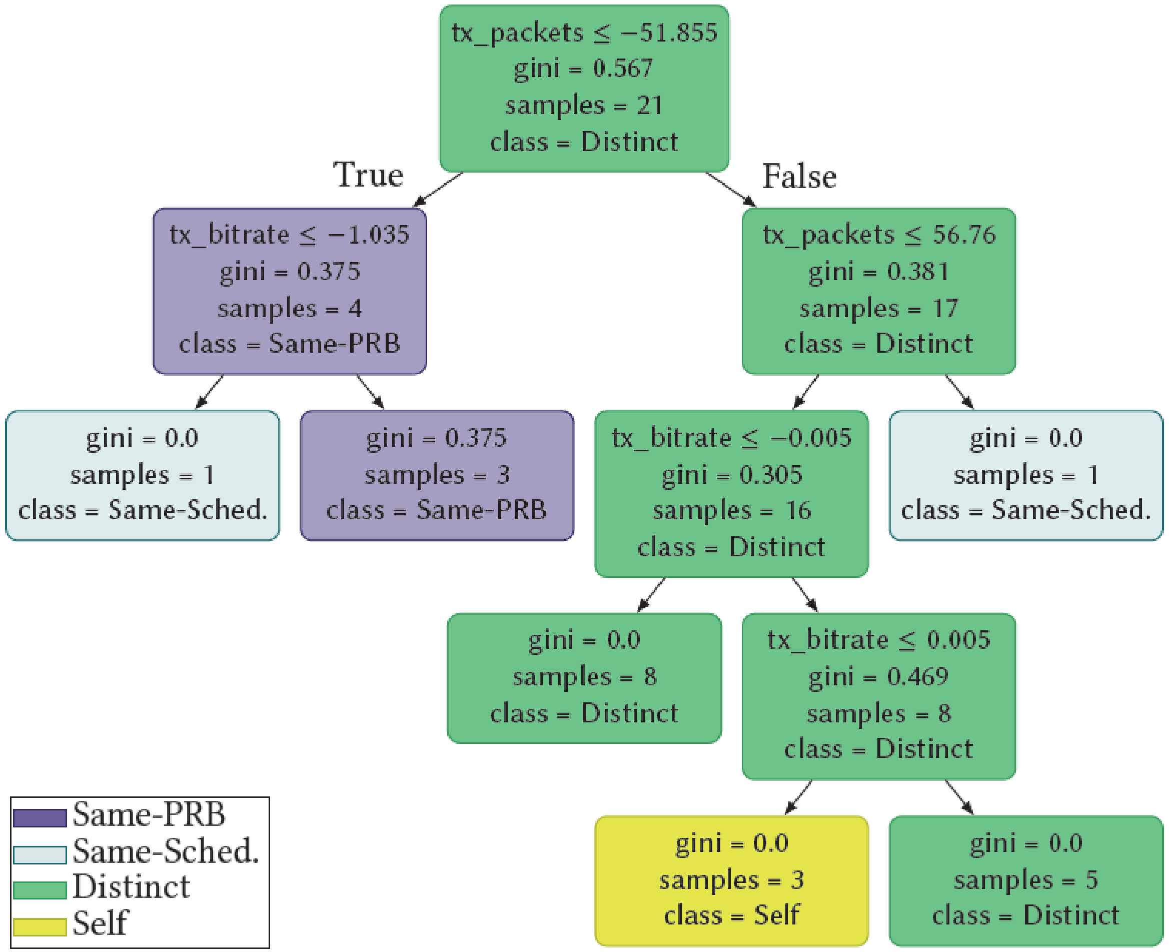

Agent HT, slice eMBB: same scheduling reduces throughput

Intent-based Action Steering

- Specific policies based on intents can improve throughput
  - Before taking action *a*, explore the graph for better alternatives

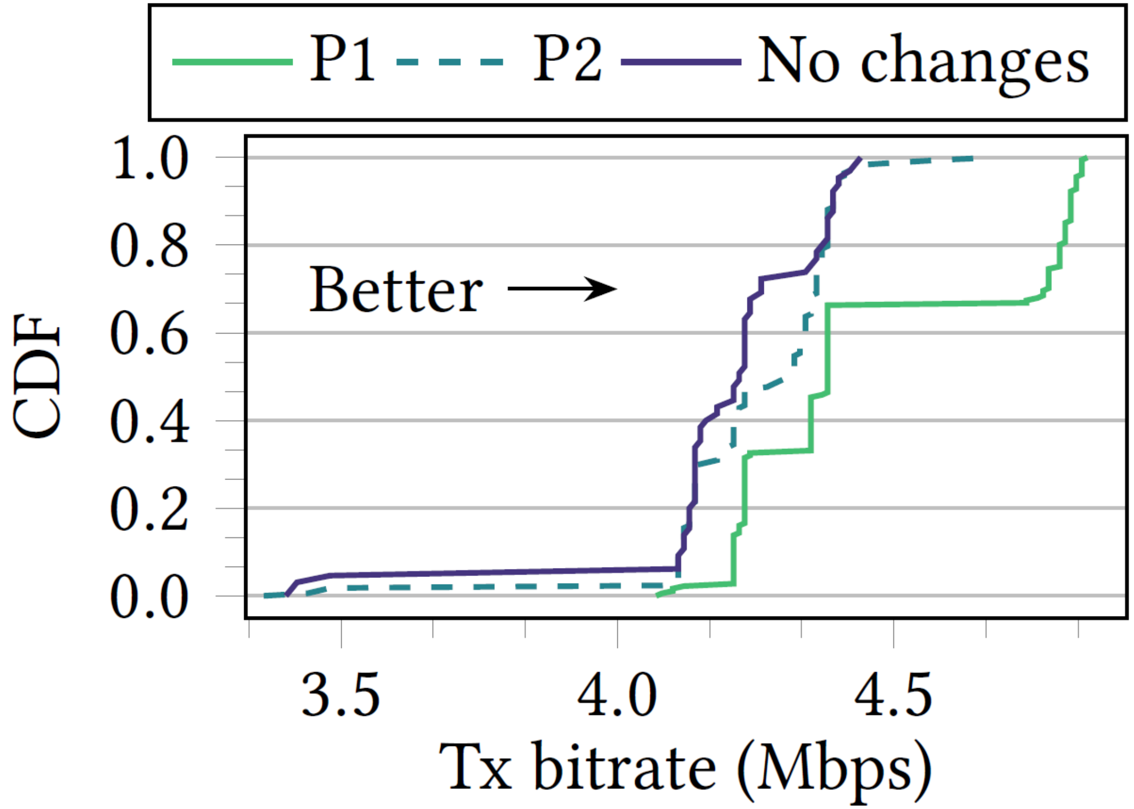

P1 looks for actions that in the past delivered higher throughput than the current one

P2 looks for actions that in the past obtained higher reward than the current one

Take-Home Messages

- EXPLORA generates network-aware explanations on agents' behavior
- We show that agents' decisions could be revisited and improved programmatically with intents
- We make public our code for reproducibility

https://github.com/wineslab/explora

Thank You!

<claudio.fiandrino@imdea.org>