SCIENCES & BIOENGINEERING SCIENCES

The Research Group

Artificial Intelligence Lab

has the honor to invite you to the public defence of the PhD thesis of

Florent Delgrange

to obtain the degree of Doctor of Sciences

Joint PhD with Universiteit Antwerpen

Title of the PhD thesis:

Activating Formal Verification of Deep Reinforcement Learning Policies by Model Checking Bisimilar Latent Space Models

Promotors:
Prof. dr. Ann Nowé (VUB)
Prof. dr. Guillermo A. Pérez (UAntwerpen)

The defence will take place on
Tuesday, August 20, 2024 at 4 p.m. in Building D, Room D2.01 (promotiezaal)

The defence can also be followed through a live stream: https://tinyurl.com/yenmtswb

Members of the jury

Prof. dr. Bernard Manderick (VUB, chair)
Dr. Roxana Rădulescu (VUB/Utrecht University NL, secretary)
Prof. dr. Bart Goethals (UAntwerpen)
Prof. dr. Nils Jansen (Radboud University Nijmegen, NL/Ruhr-Universität Bochum, DE)
Prof. dr. Pablo Samuel Castro (Google/Université de Montréal, CA)

Curriculum vitae

Florent Delgrange graduated with an MSc in Computer Science from the University of Mons in 2018. In 2020, he joined the Artificial Intelligence Lab at the Vrije Universiteit Brussel (VUB) and the Computer Science Department of the University of Antwerp (UAntwerpen). Florent's research focuses on the intersection of formal verification and reinforcement learning. He has published seven peer-reviewed articles in top-tier international conferences and journals, as well as three peer-reviewed articles in workshops.
During his PhD, Florent was also a teaching assistant for the Theory of Computation course at VUB.

Abstract of the PhD research

Intelligent agents are computational entities that autonomously interact with an environment to achieve their design objectives. On the one hand, reinforcement learning (RL) encompasses machine learning techniques that allow agents to learn by trial and error a control policy, prescribing how to behave in the environment. Although RL is proven to converge to an optimal policy under some assumptions, the guarantees vanish with the introduction of advanced techniques, such as deep RL, to deal with high-dimensional state and action spaces. This prevents them from being widely adopted in real-world safety-critical scenarios.

On the other hand, formal methods are mathematical techniques that guarantee systems' correctness. In particular, model checking allows formally verifying the agent's behaviors in the environment. However, this typically relies on a formal description of the interaction and an exhaustive exploration of the state space. This poses significant challenges because the environment is seldom explicitly accessible. Even when it is, model checking suffers from the curse of dimensionality and struggles to scale to high-dimensional state and action spaces, which are common in deep RL.

In this thesis, we leverage the strengths of deep RL to handle realistic scenarios while integrating formal methods to provide guarantees on the agent's behaviors. Specifically, we activate formal verification of deep RL policies by learning a latent model of the environment, over which we distill the deep RL policy. The outcome is amenable for model checking and is endowed with bisimulation guarantees, which allows to lift the verification results to the original environment.
Beyond distillation, we show that our method is also useful for learning representation in the context of deep RL, facilitating the learning of the policy in complex environments. We present a framework for partially observable environments.

We finally show how our method can be leveraged in synthesis, i.e., the automatic generation of controllers from logical specifications with formal guarantees. Precisely, we present how deep RL components learned via our latent space models facilitate synthesis in typically intractable environments.