Monitoring Signal Temporal Logic Specifications for Safe Reinforcement Learning

Technical University of Munich

School of Computation, Information, and Technology

Chair of Robotics, Artificial Intelligence and Real-time Systems

Background

Recent research in reinforcement learning (RL) has shown promising results in training agents to satisfy complex specifications expressed in Linear Temporal Logic (LTL). LTL is a formal language that allows for the specification of temporal properties, making it suitable for describing complex behaviors and safety requirements in RL tasks. As an example, LTL can be used to specify that an agent must eventually reach a goal while avoiding certain unsafe states, or that it must repeatedly perform a certain action. However, LTL specifications cannot express timing constraints, which are often crucial in real-world applications. Signal Temporal Logic (STL) extends LTL by allowing for the specification of timing constraints.

Description

This thesis aims to investigate the use of Signal Temporal Logic (STL) for specifying and monitoring complex behaviors in multi-task RL. Current approaches typically translate the specifications into automata [1, 2]. For STL specifications, this translation can lead to a state-space explosion, making it often computationally infeasible to apply in practice. To address this issue, this work aims to investigate better monitoring techniques for STL specifications that can be applied in the context of multi-task RL.

Tasks

- Literature review on temporal logic in RL and monitoring techniques for STL
- Design of a monitoring framework for STL specifications in multi-task RL
- Implementation of the proposed methods in a suitable RL framework
- Evaluation of the methodology on benchmark control tasks
- Documentation and thesis writing

References

[1] Z. Guo, İ. Işık, H. M. S. Ahmad, and W. Li, "One subgoal at a time: Zero-shot generalization to arbitrary linear temporal logic requirements in multi-task reinforcement learning," in *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. [Online]. Available: https://openreview.net/forum?id=NGgLhJKttI

[2] M. Jackermeier and A. Abate, "DeepLTL: Learning to efficiently satisfy complex LTL specifications for multi-task RL," in *The Thirteenth International Conference on Learning Representations*, 2025. [Online]. Available: https://openreview.net/forum?id=9pW2J49flQ

Supervisor:
Prof. Dr.-Ing. Matthias Althoff

Advisor:
Marlon Müller, M.Sc.

Research project:
—

Type:
Master's thesis

Research area:
Machine Learning, Formal Methods

Required skills:
Python, Git

Beneficial skills:
Reinforcement Learning, Temporal Logic, Control Systems

Language:
English

Date of submission:
31. März 2026

For more information please contact us:

Phone: —

E-Mail: marlon.mueller@tum.de

Website: www.ce.cit.tum.de/cps