Sim-to-Lab-to-Real: Safe Reinforcement Learning with Shielding and Generalization Guarantees (Abstract Reprint)

Kai-Chieh Hsu[1], Allen Z. Ren[2], Duy P. Nguyen[1], Anirudha Majumdar[2], Jaime F. Fisac[1]

[1]Department of Electrical and Computer Engineering, Princeton University, United States
[2]Department of Mechanical and Aerospace Engineering, Princeton University, United States

Abstract Reprint. This is an abstract reprint of a journal article by Hsu, Ren, Nguyen, Majumdar, and Fisac (2023).

Abstract

Safety is a critical component of autonomous systems and remains a challenge for learning-based policies to be utilized in the real world. In particular, policies learned using reinforcement learning often fail to generalize to novel environments due to unsafe behavior. In this paper, we propose Sim-to-Lab-to-Real to bridge the reality gap with a probabilistically guaranteed safety-aware policy distribution. To improve safety, we apply a dual policy setup where a performance policy is trained using the cumulative task reward and a backup (safety) policy is trained by solving the Safety Bellman Equation based on Hamilton-Jacobi (HJ) reachability analysis. In Sim-to-Lab transfer, we apply a supervisory control scheme to shield unsafe actions during exploration; in Lab-to-Real transfer, we leverage the Probably Approximately Correct (PAC)-Bayes framework to provide lower bounds on the expected performance and safety of policies in unseen environments. Additionally, inheriting from the HJ reachability analysis, the bound accounts for the expectation over the worst-case safety in each environment. We empirically study the proposed framework for ego-vision navigation in two types of indoor environments with varying degrees of photorealism. We also demonstrate strong generalization performance through hardware experiments in real indoor spaces with a quadrupedal robot. See https://sites.google.com/princeton.edu/sim-to-lab-to-real for supplementary material.

References

Hsu, K.-C.; Ren, A. Z.; Nguyen, D. P.; Majumdar, A.; and Fisac, J. F. 2023. Sim-to-Lab-to-Real: Safe reinforcement learning with shielding and generalization guarantees. *Artificial Intelligence*, 314: 103811.