TEAMSTER: Model-Based Reinforcement Learning for Ad Hoc Teamwork (Abstract Reprint)

João G. Ribeiro[1], **Gonçalo Rodrigues**[2], **Alberto Sardinha**[1,3], **Francisco S. Melo**[1]

[1]INESC-ID, IST Taguspark, Av. Prof. Dr. Cavaco Silva, Porto Salvo, 2744-016, Portugal
[2]Google, Google Building 110, Brandschenkestrasse 110, Zürich, 8002, Switzerland
[3]Department of Informatics, Pontifical Catholic University of Rio de Janeiro, Brazil

Abstract Reprint. This is an abstract reprint of a journal article by Ribeiro, Rodrigues, Sardinha, and Melo (2023).

Abstract

This paper investigates the use of model-based reinforcement learning in the context of ad hoc teamwork. We introduce a novel approach, named TEAMSTER, where we propose learning both the environment's model and the model of the teammates' behavior separately. Compared to the state-of-the-art PLASTIC algorithms, our results in four different domains from the multi-agent systems literature show that TEAMSTER is more flexible than the PLASTIC-Model, by learning the environment's model instead of assuming a perfect hand-coded model, and more robust/efficient than PLASTIC-Policy, by being able to continuously adapt to newly encountered teams, without implicitly learning a new environment model from scratch.

References

Ribeiro, J. G.; Rodrigues, G.; Sardinha, A.; and Melo, F. S. 2023. TEAMSTER: Model-based reinforcement learning for ad hoc teamwork. *Artificial Intelligence*, 324: 104013.