

REFERENCES

- [1] Martin Arjovsky. 2021. Out of Distribution Generalization in Machine Learning. *arXiv preprint arXiv:2103.02667* (2021).
- [2] Daniel Claes, Frans Oliehoek, Hendrik Baier, Karl Tuyls, et al. 2017. Decentralised online planning for multi-robot warehouse commissioning. In *Proceedings of the 16th international conference on autonomous agents and multiagent systems*. 492–500.
- [3] Elena Congeduti, Alexander Mey, and Frans A. Oliehoek. 2021. Loss Bounds for Approximate Influence-Based Abstraction. In *AAMAS21*.
- [4] Anchal Gupta, Youakim Badr, Ashkan Negahban, and Robin G Qiu. 2021. Energy-efficient heating control for smart buildings with deep reinforcement learning. *Journal of Building Engineering* 34 (2021), 101739.
- [5] Jinke He, Miguel Suau, and Frans Oliehoek. 2020. Influence-Augmented Online Planning for Complex Environments. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 4392–4402.
- [6] Pablo Alvarez Lopez, Michael Behrisch, Laura Bieker-Walz, Jakob Erdmann, Yun-Pang Flötteröd, Robert Hilbrich, Leonhard Lücken, Johannes Rummel, Peter Wagner, and Evamarie Wießner. 2018. Microscopic Traffic Simulation using SUMO, In The 21st IEEE International Conference on Intelligent Transportation Systems. *IEEE Intelligent Transportation Systems Conference (ITSC)*. <https://elib.dlr.de/124092/>
- [7] Elise van der Pol and Frans A. Oliehoek. 2016. Coordinated Deep Reinforcement Learners for Traffic Light Control. Submitted to NIPS'16 Workshop on Learning, Inference and Control of Multi-Agent Systems.
- [8] Frans Oliehoek, Stefan Witwicki, and Leslie Kaelbling. 2021. A sufficient statistic for influence in structured multiagent environments. *Journal of Artificial Intelligence Research* 70 (2021), 789–870.
- [9] Joaquin Quionero-Candela, Masashi Sugiyama, Anton Schwaighofer, and Neil D Lawrence. 2009. *Dataset shift in machine learning*. The MIT Press.
- [10] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347* (2017).
- [11] Miguel Suau, Alexandros Agapitos, David Lynch, Derek Farrell, Mingqi Zhou, and Aleksandar Milenovic. 2021. Offline Contextual Bandits for Wireless Network Optimization. *arXiv preprint arXiv:2111.08587* (2021).
- [12] Miguel Suau, Jinke He, Matthijs T. J. Spaan, and Frans A. Oliehoek. 2022. Influence-Augmented Local Simulators: A Scalable Solution for Fast Deep RL in Large Networked Systems. *Preprint* (2022).
- [13] Cathy Wu, Aboudy Kreidieh, Kanaad Parvate, Eugene Vinitsky, and Alexandre M Bayen. 2017. Flow: A Modular Learning Framework for Autonomy in Traffic. *arXiv preprint arXiv:1710.05465* (2017).