

REFERENCES

- [1] Ruzena Bajcsy, Yiannis Aloimonos, and John K Tsotsos. 2018. Revisiting active perception. *Autonomous Robots* 42, 2 (2018), 177–196.
- [2] Ricson Cheng, Arpit Agarwal, and Katerina Fragkiadaki. 2018. Reinforcement Learning of Active Vision for Manipulating Objects under Occlusions. In *Conference on Robot Learning*. 422–431.
- [3] Maxime Chevalier-Boisvert. 2018. gym-miniworld environment for OpenAI Gym. <https://github.com/maximecb/gym-miniworld>.
- [4] Gokhan Uzunbas M. Nam Lim S. Devrim Kaba, M. 2017. A reinforcement learning approach to the view planning problem. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 6933–6941.
- [5] Wei Jing, Chun Fan Goh, Mabarani Rajaraman, Fei Gao, Sooho Park, Yong Liu, and Kenji Shimada. 2018. A computational framework for automatic online path generation of robotic inspection tasks via coverage planning and reinforcement learning. *IEEE Access* 6 (2018), 54854–54864.
- [6] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. 2016. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*. 1928–1937.
- [7] Steven D. Whitehead and Dana H. Ballard. 1990. Active Perception and Reinforcement Learning. In *Machine Learning Proceedings 1990*, Bruce Porter and Raymond Mooney (Eds.). Morgan Kaufmann, San Francisco (CA), 179 – 188. <https://doi.org/10.1016/B978-1-55860-141-3.50025-0>
- [8] Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning* 8, 3-4 (1992), 229–256.