

REFERENCES

- [1] Pieter Abbeel and Andrew Y Ng. 2004. Apprenticeship learning via inverse reinforcement learning. In *International Conference on Machine Learning (ICML)*.
- [2] Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N. Bennett, Kori Inkpen, Jaime Teevan, Ruth Kikin-Gil, and Eric Horvitz. 2019. Guidelines for human-AI interaction. In *CHI*. ACM, Article 3, 13 pages.
- [3] Ofra Amir, Ece Kamar, Andrey Kolobov, and Barbara Grosz. 2016. Interactive teaching strategies for agent training. In *International Joint Conferences on Artificial Intelligence (IJCAI)*.
- [4] Felipe Leno da Silva and Anna Helena Reali Costa. 2019. A Survey on Transfer Learning for Multiagent Reinforcement Learning Systems. *Journal of Artificial Intelligence Research* 64 (2019), 645–703.
- [5] Christos Dimitrakakis, David C Parkes, Goran Radanovic, and Paul Tylkin. 2017. Multi-view Decision Processes: The helper-AI problem. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- [6] Eyal Even-Dar and Yishay Mansour. 2003. Approximate Equivalence of Markov Decision Processes. In *Learning Theory and Kernel Machines*, Bernhard Schölkopf and Manfred K. Warmuth (Eds.). Springer, Berlin, Heidelberg, 581–594.
- [7] Richard Everett and Stephen J. Roberts. 2018. Learning Against Non-Stationary Agents with Opponent Modelling and Deep Reinforcement Learning. In *AAAI Spring Symposia 2018*.
- [8] Alan Fern, Sriraam Natarajan, Kshitij Judah, and Prasad Tadepalli. 2014. A Decision-Theoretic Model of Assistance. *Journal of Artificial Intelligence Research* 50 (2014), 71–104.
- [9] Jakob N. Foerster, Richard Y. Chen, Maruan Al-Shedivat, Shimon Whiteson, Pieter Abbeel, and Igor Mordatch. 2018. Learning with Opponent-Learning Awareness. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 122–130.
- [10] Aditya Grover, Maruan Al-Shedivat, Jayesh K. Gupta, Yuri Burda, and Harrison Edwards. 2018. Learning Policy Representations in Multiagent Systems. In *International Conference on Machine Learning (ICML)*. 1797–1806.
- [11] Dylan Hadfield-Menell, Stuart J. Russell, Pieter Abbeel, and Anca D. Dragan. 2016. Cooperative Inverse Reinforcement Learning. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- [12] Luis Haug, Sebastian Tschiatschek, and Adish Singla. 2018. Teaching inverse reinforcement learners via features and demonstrations. In *Advances in Neural Information Processing Systems (NeurIPS)*. 8464–8473.
- [13] Matteo Hessel, Hubert Soyer, Lasse Espeholt, Wojciech Czarnecki, Simon Schmitt, and Hado van Hasselt. 2019. Multi-Task Deep Reinforcement Learning with PopArt. In *AAAI*. 3796–3803.
- [14] Yeping Hu, Alireza Nakhaei, Masayoshi Tomizuka, and Kikuo Fujimura. 2019. Interaction-aware Decision Making with Adaptive Strategies under Merging Scenarios. *arXiv preprint arXiv:1904.06025* (2019).
- [15] Stephen James, Paul Wohlhart, Mrinal Kalakrishnan, Dmitry Kalashnikov, Alex Irgan, Julian Ibarz, Sergey Levine, Raia Hadsell, and Konstantinos Bousmalis. 2019. Sim-to-real via sim-to-sim: Data-efficient robotic grasping via randomized-to-canonical adaptation networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 12627–12637.
- [16] Shervin Javdani, Henny Admoni, Stefania Pellegrinelli, Siddhartha S. Srinivasa, and J. Andrew Bagnell. 2018. Shared autonomy via hindsight optimization for teleoperation and teaming. *The International Journal of Robotics Research* 37, 7 (2018), 717–742.
- [17] Parameswaran Kamalaruban, Rati Devidze, Volkan Cevher, and Adish Singla. 2019. Interactive Teaching Algorithms for Inverse Reinforcement Learning. In *International Joint Conference on Artificial Intelligence (IJCAI)*.
- [18] Joel Z. Leibo, Vinicius Flores Zambaldi, Marc Lanctot, Janusz Marecki, and Thore Graepel. 2017. Multi-agent Reinforcement Learning in Sequential Social Dilemmas. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 464–473.
- [19] Owen Macindoe, Leslie Pack Kaelbling, and Tomás Lozano-Pérez. 2012. Pomcpoc: Belief space planning for sidekicks in cooperative games. In *AIIDE*.
- [20] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin A. Riedmiller, Andreas Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharmarajan Kumar, Daan Wierstra, Shane Legg, and Demis Hassabis. 2015. Human-level control through deep reinforcement learning. *Nature* 518, 7540 (2015), 529–533.
- [21] Stefanos Nikolaidis, Swaprava Nath, Ariel D Procaccia, and Siddhartha Srinivasa. 2017. Game-theoretic modeling of human adaptation in human-robot collaboration. In *International Conference on Human-Robot Interaction (HRI)*. 323–331.
- [22] Stefanos Nikolaidis, Ramya Ramakrishnan, Keren Gu, and Julie A. Shah. 2015. Efficient Model Learning from Joint-Action Demonstrations for Human-Robot Collaborative Tasks. In *International Conference on Human-Robot Interaction (HRI)*. 189–196.
- [23] Lerrel Pinto, James Davidson, Rahul Sukthankar, and Abhinav Gupta. 2017. Robust Adversarial Reinforcement Learning. In *International Conference on Machine Learning (ICML)*. 2817–2826.
- [24] Neil C. Rabinowitz, Frank Perbet, H. Francis Song, Chiyuan Zhang, S. M. Ali Eslami, and Matthew Botvinick. 2018. Machine Theory of Mind. In *International Conference on Machine Learning (ICML)*. 4215–4224.
- [25] Goran Radanovic, Rati Devidze, David Parkes, and Adish Singla. 2019. Learning to Collaborate in Markov Decision Processes. In *International Conference on Machine Learning (ICML)*. 5261–5270.
- [26] Md Shihanur Rahman and AMT Oo. 2017. Distributed multi-agent based coordinated power management and control strategy for microgrids with distributed energy resources. *Energy conversion and management* 139 (2017), 20–32.
- [27] Roberta Raileanu, Emily Denton, Arthur Szlam, and Rob Fergus. 2018. Modeling Others using Oneself in Multi-Agent Reinforcement Learning. In *International Conference on Machine Learning (ICML)*. 4254–4263.
- [28] Tabish Rashid, Mikayel Samvelyan, Christian Schroeder Witt, Gregory Farquhar, Jakob Foerster, and Shimon Whiteson. 2018. QMIX: Monotonic Value Function Factorisation for Deep Multi-Agent Reinforcement Learning. In *International Conference on Machine Learning*. 4292–4301.
- [29] Aurko Roy, Huan Xu, and Sebastian Pokutta. 2017. Reinforcement Learning under Model Mismatch. In *Advances in Neural Information Processing Systems (NeurIPS)*. 3043–3052.
- [30] Steindór Sæmundsson, Katja Hofmann, and Marc Peter Deisenroth. 2018. Meta Reinforcement Learning with Latent Variable Gaussian Processes. In *UAI*.
- [31] Shai Shalev-Shwartz, Shaked Shammah, and Amnon Shashua. 2016. Safe, multi-agent, reinforcement learning for autonomous driving. *arXiv preprint arXiv:1610.03295* (2016).
- [32] Guni Sharon, Roni Stern, Ariel Felner, and Nathan R Sturtevant. 2015. Conflict-based search for optimal multi-agent pathfinding. *Artificial Intelligence* 219 (2015), 40–66.
- [33] David Silver. 2005. Cooperative Pathfinding. *AIIDE* 1 (2005), 117–122.
- [34] Adish Singla, Seyed Hamed Hassani, and Andreas Krause. 2018. Learning to Interact With Learning Agents. In *AAAI*. 4083–4090.
- [35] Adish Singla, Ryen W White, Ahmed Hassan, and Eric Horvitz. 2014. Enhancing personalization via search activity attribution. In *SIGIR*. 1063–1066.
- [36] Richard S Sutton and Andrew G Barto. 2018. *Reinforcement learning: An introduction*. MIT press.
- [37] Josh Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. 2017. Domain randomization for transferring deep neural networks from simulation to the real world. In *International Conference on Intelligent Robots and Systems (IROS)*. 23–30.
- [38] Sebastian Tschiatschek, Ahana Ghosh, Luis Haug, Rati Devidze, and Adish Singla. 2019. Learner-aware Teaching: Inverse Reinforcement Learning with Preferences and Constraints. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- [39] Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. 2019. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature* (2019), 1–5.
- [40] Ryen W White, Wei Chu, Ahmed Hassan, Xiaodong He, Yang Song, and Hongning Wang. 2013. Enhancing personalized search by mining and modeling task behavior. In *International Conference on World Wide Web (WWW)*. 1411–1420.
- [41] Ryen W White, Ahmed Hassan, Adish Singla, and Eric Horvitz. 2014. From devices to people: Attribution of search activity in multi-user settings. In *International Conference on World Wide Web (WWW)*. 431–442.
- [42] H James Wilson and Paul R Daugherty. 2018. Collaborative intelligence: Humans and AI are joining forces. *Harvard Business Review* 96, 4 (2018), 114–123.
- [43] Jiachen Yang, Alireza Nakhaei, David Isele, Hongyuan Zha, and Kikuo Fujimura. 2018. CM3: Cooperative Multi-goal Multi-stage Multi-agent Reinforcement Learning. *arXiv preprint arXiv:1809.05188* (2018).
- [44] Haoqi Zhang, David C Parkes, and Yiling Chen. 2009. Policy teaching through reward function learning. In *Conference on Electronic Commerce (EC)*. 295–304.
- [45] Yan Zheng, Zhaopeng Meng, Jianye Hao, Zongzhang Zhang, Tianpei Yang, and Changjie Fan. 2018. A Deep Bayesian Policy Reuse Approach Against Non-Stationary Agents. In *Advances in Neural Information Processing Systems (NeurIPS)*. 962–972.
- [46] Brian D. Ziebart. 2010. *Modeling Purposeful Adaptive Behavior with the Principle of Maximum Causal Entropy*. Ph.D. Dissertation. Advisor(s) Bagnell, J. Andrew.