

REFERENCES

- [1] Martin Allen and Shlomo Zilberstein. 2009. Complexity of decentralized control: Special cases. In *Advances in Neural Information Processing Systems (NIPS)*. 19–27.
- [2] Christopher Amato and Shlomo Zilberstein. 2009. Achieving goals in decentralized POMDPs. In *Autonomous Agents and Multiagent Systems (AAMAS)*. 593–600.
- [3] Mauricio Araya-López, Olivier Buffet, Vincent Thomas, and François Charpillet. 2010. A POMDP Extension with Belief-dependent Rewards. In *Advances in Neural Information Processing Systems (NIPS)*. 64–72.
- [4] Daniel S Bernstein, Robert Givan, Neil Immerman, and Shlomo Zilberstein. 2002. The complexity of decentralized control of Markov decision processes. *Mathematics of operations research* 27, 4 (2002), 819–840.
- [5] Benjamin Charrow, Vijay Kumar, and Nathan Michael. 2014. Approximate representations for multi-robot control policies that maximize mutual information. *Autonomous Robots* 37, 4 (2014), 383–400.
- [6] Morris H DeGroot. 2004. *Optimal Statistical Decisions*. John Wiley & Sons, Inc., Hoboken, NJ. Wiley Classics Library edition.
- [7] Jilles Steeve Dibangoye, Christopher Amato, Olivier Buffet, and François Charpillet. 2016. Optimally solving Dec-POMDPs as continuous-state MDPs. *Journal of Artificial Intelligence Research* 55 (2016), 443–497.
- [8] Mathieu Fehr, Olivier Buffet, Vincent Thomas, and Jilles Dibangoye. 2018. rho-POMDPs have Lipschitz-Continuous epsilon-Optimal Value Functions. In *Advances in Neural Information Processing Systems (NIPS)*. 6933–6943.
- [9] Eric A Hansen, Daniel S Bernstein, and Shlomo Zilberstein. 2004. Dynamic programming for partially observable stochastic games. In *AAAI*. 709–715.
- [10] Mikko Lauri, Eero Heinänen, and Simone Frintrop. 2017. Multi-robot active information gathering with periodic communication. In *IEEE Intl. Conference on Robotics and Automation (ICRA)*. 851–856.
- [11] Liam C MacDermed and Charles L Isbell. 2013. Point Based Value Iteration with Optimal Belief Compression for Dec-POMDPs. In *Advances in Neural Information Processing Systems (NIPS)*. 100–108.
- [12] Kevin Murphy. 2012. *Machine Learning: A probabilistic perspective*. MIT Press.
- [13] Ranjit Nair, Milind Tambe, Makoto Yokoo, David Pynadath, and Stacy Marsella. 2003. Taming decentralized POMDPs: Towards efficient policy computation for multiagent settings. In *Intl. Joint Conference on Artificial Intelligence (IJCAI)*. 705–711.
- [14] Frans A Oliehoek. 2013. Sufficient Plan-Time Statistics for Decentralized POMDPs. In *Intl. Joint Conference on Artificial Intelligence (IJCAI)*. 302–308.
- [15] Frans A Oliehoek and Christopher Amato. 2016. *A Concise Introduction to Decentralized POMDPs*. Springer.
- [16] Frans A Oliehoek, Julian FP Kooij, and Nikos Vlassis. 2008. The cross-entropy method for policy search in decentralized POMDPs. *Informatika* 32, 4 (2008), 341–357.
- [17] Frans A Oliehoek, Matthijs TJ Spaan, Christopher Amato, and Shimon Whiteson. 2013. Incremental clustering and expansion for faster optimal planning in Dec-POMDPs. *Journal of Artificial Intelligence Research* 46 (2013), 449–509.
- [18] Frans A Oliehoek, Matthijs TJ Spaan, Shimon Whiteson, and Nikos Vlassis. 2008. Exploiting locality of interaction in factored Dec-POMDPs. In *Autonomous Agents and Multiagent Systems (AAMAS)*. 517–524.
- [19] Joni K Pajarinen and Jaakko Peltonen. 2011. Periodic Finite State Controllers for Efficient POMDP and DEC-POMDP Planning. In *Advances in Neural Information Processing Systems (NIPS)*. 2636–2644.
- [20] Yash Satsangi, Shimon Whiteson, Frans A Oliehoek, and Matthijs TJ Spaan. 2018. Exploiting submodular value functions for scaling up active perception. *Autonomous Robots* 42, 2 (2018), 209–233.
- [21] Brent Schlotfeldt, Dinesh Thakur, Nikolay Atanasov, Vijay Kumar, and George J Pappas. 2018. Anytime Planning for Decentralized Multirobot Active Information Gathering. *IEEE Robotics and Automation Letters* 3, 2 (2018), 1025–1032.
- [22] Sven Seuken and Shlomo Zilberstein. 2007. Memory-Bounded Dynamic Programming for DEC-POMDPs. In *Intl. Joint Conference on Artificial Intelligence (IJCAI)*. 2009–2015.
- [23] Richard D Smallwood and Edward J Sondik. 1973. The optimal control of partially observable Markov processes over a finite horizon. *Operations research* 21, 5 (1973), 1071–1088.
- [24] Matthijs TJ Spaan, Geoffrey J Gordon, and Nikos Vlassis. 2006. Decentralized planning under uncertainty for teams of communicating agents. In *Autonomous Agents and Multiagent Systems (AAMAS)*. 249–256.
- [25] Matthijs TJ Spaan, Tiago S Veiga, and Pedro U Lima. 2015. Decision-theoretic planning under uncertainty with information rewards for active cooperative perception. *Autonomous Agents and Multi-Agent Systems* 29, 6 (2015), 1157–1185.
- [26] Daniel Szer, François Charpillet, and Shlomo Zilberstein. 2005. MAA*: a heuristic search algorithm for solving decentralized POMDPs. In *Uncertainty in Artificial Intelligence (UAI)*. 576–583.
- [27] Feng Wu, Shlomo Zilberstein, and Xiaoping Chen. 2011. Online planning for multi-agent systems with bounded communication. *Artificial Intelligence* 175, 2 (2011), 487–511.