



















## REFERENCES

- [1] Chris A. B. Baker, Sarvapali Ramchurn, W. T. Luke Teacy, and Nicholas R. Jennings. 2016. Planning Search and Rescue Missions for UAV Teams. In *Proceedings of the 22nd European Conference on Artificial Intelligence (ECAI-2016)*, G. A. Kaminka, M. Fox, P. Bouquet, Hullermeijer E., Dignum F., Dignum V., and van Harmalen F. (Eds.). ECCAI, IOS Press, The Hague, The Netherlands, 1777–1782.
- [2] Rafael H. Bordini, Ana L. C. Bazzan, Rafael de O. Jannone, Daniel M. Basso, Rosa M. Vicari, and Victor R. Lesser. 2002. AgentSpeak(XL): Efficient Intention Selection in BDI Agents via Decision-Theoretic Task Scheduling. In *Proceedings of the First International Conference on Autonomous Agents and Multiagent Systems (AAMAS'02)*. 1294–1302.
- [3] Rafael H Bordini, Jomi Fred Hübner, and Michael Wooldridge. 2007. *Programming Multi-Agent Systems in AgentSpeak Using Jason*. Vol. 8. John Wiley & Sons.
- [4] Cameron Browne, Edward J. Powley, Daniel Whitehouse, Simon M. Lucas, Peter I. Cowling, Philipp Rohlfshagen, Stephen Tavener, Diego Perez, Spyridon Samothrakis, and Simon Colton. 2012. A Survey of Monte Carlo Tree Search Methods. *IEEE Trans. Comput. Intellig. and AI in Games* 4, 1 (March 2012), 1–43.
- [5] L. Busoniu, R. Babuska, and B. De Schutter. 2008. A Comprehensive Survey of Multiagent Reinforcement Learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* 38, 2 (2008), 156–172.
- [6] Guillaume Chaslot, Sander Bakkes, Istvan Szita, and Pieter Spronck. 2008. Monte-Carlo Tree Search: A New Framework for Game AI. In *Proceedings of the Fourth Artificial Intelligence and Interactive Digital Entertainment Conference*. Stanford, California, USA.
- [7] Bradley J. Clement and Edmund H. Durfee. 1999. Theory for Coordinating Concurrent Hierarchical Planning Agents Using Summary Information. In *Proceedings of the Sixteenth National Conference on Artificial Intelligence and Eleventh Conference on Innovative Applications of Artificial Intelligence*, Jim Hendler and Devika Subramanian (Eds.). AAAI Press / The MIT Press, Orlando, Florida, USA, 495–502.
- [8] Bradley J. Clement and Edmund H. Durfee. 2000. Performance of Coordinating Concurrent Hierarchical Planning Agents Using Summary Information. In *4th International Conference on Multi-Agent Systems*. IEEE Computer Society, Boston, MA, USA, 373–374.
- [9] Bradley J. Clement, Edmund H. Durfee, and Anthony C. Barrett. 2007. Abstract Reasoning for Planning and Coordination. *J. Artif. Intell. Res. (JAIR)* 28 (2007), 453–515.
- [10] K. S. Decker and V. R. Lesser. 1993. Quantitative Modeling of Complex Environments. *International Journal of Intelligent Systems in Accounting, Finance and Management* 2 (1993), 215–234.
- [11] Eithan Ephrati and Jeffrey S. Rosenschein. 1993. A Framework for the Interleaving of Execution and Planning for Dynamic Tasks by Multiple Agents. In *From Reaction to Cognition, 5th European Workshop on Modelling Autonomous Agents, MAAMAW '93, Neuchatel, Switzerland, August 25-27, 1993, Selected Papers*, Cristiano Castelfranchi and Jean-Pierre Müller (Eds.). Springer, Neuchatel, Switzerland, 139–153.
- [12] Jakob Foerster, Ioannis Alexandros Assael, Nando de Freitas, and Shimon Whiteson. 2016. Learning to Communicate with Deep Multi-Agent Reinforcement Learning. In *Advances in Neural Information Processing Systems 29*, D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett (Eds.). Curran Associates, Inc., 2137–2145.
- [13] Jakob Foerster, Nantas Nardelli, Gregory Farquhar, Triantafyllos Afouras, Philip H. S. Torr, Pushmeet Kohli, and Shimon Whiteson. 2017. Stabilising Experience Replay for Deep Multi-Agent Reinforcement Learning. In *Proceedings of the 34th International Conference on Machine Learning*. 1146–1155.
- [14] Levente Kocsis and Csaba Szepesvári. 2006. Bandit Based Monte-Carlo Planning. In *17th European Conference on Machine Learning*, Johannes Fürnkranz, Tobias Scheffer, and Myra Spiliopoulou (Eds.). Springer, Berlin, Germany, 282–293.
- [15] Brian Logan, John Thangarajah, and Neil Yorke-Smith. 2017. Progressing Intention Progression: A Call for a Goal-Plan Tree Contest. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2017, São Paulo, Brazil, May 8-12, 2017*. ACM, 768–772.
- [16] Anuj Mahajan, Tabish Rashid, Mikayel Samvelyan, and Shimon Whiteson. 2019. MAVEN: Multi-Agent Variational Exploration. In *Advances in Neural Information Processing Systems 32*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Eds.). Curran Associates, Inc., 7613–7624.
- [17] Alexander Pokahr, Lars Braubach, and Winfried Lamersdorf. 2005. *Jadex: A BDI Reasoning Engine*. Springer US, Boston, MA, 149–174.
- [18] Anand S. Rao and Michael P. Georgeff. 1991. Modeling Rational Agents within a BDI-Architecture. In *Proc. of KR'91*. Cambridge, MA, 473–484.
- [19] Sebastian Sardiña, Lavindra de Silva, and Lin Padgham. 2006. Hierarchical Planning in BDI Agent Programming Languages: A Formal Approach. In *Proceedings of the 5th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2006)*, Hideyuki Nakashima, Michael P. Wellman, Gerhard Weiss, and Peter Stone (Eds.). ACM, Hakodate, Japan, 1001–1008.
- [20] Maarten P. D. Schadd, Mark H. M. Winands, Mandy J. W. Tak, and Jos W. H. M. Uiterwijk. 2012. Single-Player Monte-Carlo Tree Search for SameGame. *Knowl.-Based Syst.* 34 (2012), 3–11.
- [21] Patricia H. Shaw and Rafael H. Bordini. 2007. Towards Alternative Approaches to Reasoning About Goals. In *Declarative Agent Languages and Technologies V, 5th International Workshop*, Matteo Baldoni, Tran Cao Son, M. Birna van Riemsdijk, and Michael Winikoff (Eds.), Vol. 4897. Springer, Honolulu, HI, USA, 104–121.
- [22] Patricia H. Shaw and Rafael H. Bordini. 2010. An Alternative Approach for Reasoning about the Goal-Plan Tree Problem. In *Proceedings of the 19th European Conference on Artificial Intelligence (ECAI 2010)*, Helder Coelho, Rudi Studer, and Michael Wooldridge (Eds.), Vol. 215. IOS Press, Lisbon, Portugal, 1035–1036.
- [23] David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dhruv Kumar, Thore Graepel, et al. 2017. Mastering Chess and Shogi by Self-Play with a General Reinforcement Learning Algorithm. *arXiv preprint arXiv:1712.01815* (2017).
- [24] Sainbayar Sukhbaatar, arthur szlam, and Rob Fergus. 2016. Learning Multi-agent Communication with Backpropagation. In *Advances in Neural Information Processing Systems 29*, D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett (Eds.). Curran Associates, Inc., 2244–2252.
- [25] John Thangarajah and Lin Padgham. 2011. Computationally Effective Reasoning About Goal Interactions. *Journal of Automated Reasoning* 47, 1 (2011), 17–56.
- [26] John Thangarajah, Lin Padgham, and Michael Winikoff. 2003. Detecting & Avoiding Interference Between Goals in Intelligent Agents. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI-03)*, Georg Gottlob and Toby Walsh (Eds.). Morgan Kaufmann, Acapulco, Mexico, 721–726.
- [27] John Thangarajah, Lin Padgham, and Michael Winikoff. 2003. Detecting & Exploiting Positive Goal Interaction in Intelligent Agents. In *The Second International Joint Conference on Autonomous Agents & Multiagent Systems, AAMAS 2003*. ACM, Melbourne, Victoria, Australia, 401–408.
- [28] John Thangarajah, Sebastian Sardiña, and Lin Padgham. 2012. Measuring Plan Coverage and Overlap for Agent Reasoning. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 2*. International Foundation for Autonomous Agents and Multiagent Systems, 1049–1056.
- [29] Régis Vincent, Bryan Horling, Victor Lesser, and Thomas Wagner. 2001. Implementing Soft Real-Time Agent Control. In *Proceedings of the Fifth International Conference on Autonomous Agents (AGENTS'01)*. ACM Press, New York, NY, USA, 355–362. <https://doi.org/10.1145/375735.376329>
- [30] T. Wagner, A. Garvey, and V. Lesser. 1998. Criteria-Directed Heuristic Task Scheduling. *International Journal of Approximate Reasoning* 19 (1998), 91–118.
- [31] Tonghan Wang, Jianhao Wang, Yi Wu, and Chongjie Zhang. 2020. Influence-Based Multi-Agent Exploration. In *Proceedings of the 8th International Conference on Learning Representations (ICLR)*. <https://arxiv.org/pdf/1910.05512.pdf>
- [32] Max Waters, Lin Padgham, and Sebastian Sardiña. 2014. Evaluating Coverage Based Intention Selection. In *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014)*, Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (Eds.), 957–964.
- [33] Max Waters, Lin Padgham, and Sebastian Sardiña. 2015. Improving domain-independent intention selection in BDI systems. *Autonomous Agents and Multi-Agent Systems* 29, 4 (2015), 683–717. <https://doi.org/10.1007/s10458-015-9293-5>
- [34] David E Wilkins, Karen L Myers, John D Lowrance, and Leonard P Wesley. 1995. Planning and reacting in uncertain and dynamic environments. *Journal of Experimental & Theoretical Artificial Intelligence* 7, 1 (1995), 121–152.
- [35] Michael Winikoff. 2005. JACK Intelligent Agents: An Industrial Strength Platform. In *Multi-Agent Programming*. Springer, New York, NY, 175–193.
- [36] Yuan Yao, Lavindra de Silva, and Brian Logan. 2016. Reasoning About the Executability of Goal-Plan Trees. In *Proceedings of the 4th International Workshop on Engineering Multi-Agent Systems (EMAS 2016)*. Singapore, 181–196.
- [37] Yuan Yao and Brian Logan. 2016. Action-Level Intention Selection for BDI Agents. In *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2016)*, J. Thangarajah, K. Tuyls, C. Jonker, and S. Marsella (Eds.). IFAAMAS, Singapore, 1227–1235.
- [38] Yuan Yao, Brian Logan, and John Thangarajah. 2014. SP-MCTS-based Intention Scheduling for BDI Agents. In *Proceedings of the 21st European Conference on Artificial Intelligence (ECAI-2014)*, Torsten Schaub, Gerhard Friedrich, and Barry O'Sullivan (Eds.). ECCAI, IOS Press, Prague, Czech Republic, 1133–1134.
- [39] Yuan Yao, Brian Logan, and John Thangarajah. 2016. Intention Selection with Deadlines. In *Proceedings of the 22nd European Conference on Artificial Intelligence (ECAI-2016)*, G. A. Kaminka, M. Fox, P. Bouquet, Hullermeijer E., Dignum F., Dignum V., and van Harmalen F. (Eds.). ECCAI, IOS Press, The Hague, The Netherlands, 1700–1701. <https://doi.org/10.3233/978-1-61499-672-9-1700>
- [40] Yuan Yao, Brian Logan, and John Thangarajah. 2016. Robust Execution of BDI Agent Programs by Exploiting Synergies Between Intentions. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI-16)*, Dale Schuurmans and Michael P. Wellman (Eds.). AAAI Press, Phoenix, USA, 2558–2564. <http://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/12148>