



















- International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob)*. IEEE, 228–233.
- [12] Matthew Grounds and Daniel Kudenko. 2005. Combining reinforcement learning with symbolic planning. In *Adaptive Agents and Multi-Agent Systems III. Adaptation and Multi-Agent Learning*. Springer, 75–86.
- [13] Marek Grzes and Daniel Kudenko. 2008. Plan-based reward shaping for reinforcement learning. In *2008 4th International IEEE Conference Intelligent Systems*, Vol. 2. IEEE, 10–22.
- [14] Chad Hogg, Ugur Kuter, and Hector Munoz-Avila. 2010. Learning methods to generate good plans: Integrating htn learning and reinforcement learning. In *Twenty-Fourth AAAI Conference on Artificial Intelligence*. Citeseer.
- [15] León Illanes, Xi Yan, Rodrigo Toro Icarte, and Sheila A McClraith. 2020. Symbolic Plans as High-Level Instructions for Reinforcement Learning. In *Proceedings of the International Conference on Automated Planning and Scheduling*, Vol. 30. 540–550.
- [16] Sergio Jiménez, Tomás De La Rosa, Susana Fernández, Fernando Fernández, and Daniel Borrajo. 2012. A review of machine learning for automated planning. *The Knowledge Engineering Review* 27, 4 (2012), 433–467.
- [17] Saket Joshi, Paul Schermerhorn, Roni Khardon, and Matthias Scheutz. 2012. Abstract planning for reactive robots. In *2012 IEEE International Conference on Robotics and Automation*. IEEE, 4379–4384.
- [18] George Konidaris, Leslie Pack Kaelbling, and Tomas Lozano-Perez. 2018. From skills to symbols: Learning symbolic representations for abstract high-level planning. *Journal of Artificial Intelligence Research* 61 (2018), 215–289.
- [19] Matteo Leonetti, Luca Iocchi, and Peter Stone. 2016. A synthesis of automated planning and reinforcement learning for efficient, robust decision-making. *Artificial Intelligence* 241 (2016), 103–130.
- [20] Daoming Lyu, Fangkai Yang, Bo Liu, and Steven Gustafson. 2019. SDRL: interpretable and data-efficient deep reinforcement learning leveraging symbolic planning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 2970–2977.
- [21] Tom M Mitchell, Richard M Keller, and Smadar T Kedar-Cabelli. 1986. Explanation-based generalization: A unifying view. *Machine learning* 1, 1 (1986), 47–80.
- [22] Sanmit Narvekar, Bei Peng, Matteo Leonetti, Jivko Sinapov, Matthew E. Taylor, and Peter Stone. 2020. Curriculum Learning for Reinforcement Learning Domains: A Framework and Survey. *Journal of Machine Learning Research* 21, 181 (2020), 1–50. <http://jmlr.org/papers/v21/20-212.html>
- [23] Jun Hao Alvin Ng and Ronald PA Petrick. 2019. Incremental Learning of Planning Actions in Model-Based Reinforcement Learning. In *IJCAI*. 3195–3201.
- [24] Vasanth Sarathy and Matthias Scheutz. 2018. MacGyver problems: Ai challenges for testing resourcefulness and creativity. *Advances in Cognitive Systems* 6 (2018), 31–44.
- [25] Wei-Min Shen. 1989. *Learning from the environment based on percepts and actions*. Ph.D. Dissertation. Carnegie Mellon University.
- [26] Wei-Min Shen and Herbert A Simon. 1989. Rule Creation and Rule Learning Through Environmental Exploration. In *IJCAI*. Citeseer, 675–680.
- [27] Armando Solar-Lezama. 2009. The sketching approach to program synthesis. In *Asian Symposium on Programming Languages and Systems*. Springer, 4–13.
- [28] Richard S Sutton and Andrew G Barto. 2018. *Reinforcement learning: An introduction*. MIT press.
- [29] Richard S Sutton, Doina Precup, and Satinder Singh. 1999. Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial intelligence* 112, 1-2 (1999), 181–211.
- [30] Kartik Talamadupula, J Benton, Subbarao Kambhampati, Paul Schermerhorn, and Matthias Scheutz. 2010. Planning for human-robot teaming in open worlds. *ACM Transactions on Intelligent Systems and Technology (TIIST)* 1, 2 (2010), 1–24.
- [31] Christopher JCH Watkins and Peter Dayan. 1992. Q-learning. *Machine learning* 8, 3-4 (1992), 279–292.
- [32] John Winder, Stephanie Milani, Matthew Landen, Erebus Oh, Shane Parr, Shawn Squire, Marie desJardins, and Cynthia Matuszek. 2019. Planning with Abstract Learned Models While Learning Transferable Subtasks. *arXiv preprint arXiv:1912.07544* (2019).
- [33] Zhanhao Xiao, Hai Wan, Hankui Hankz Zhuo, Jinxia Lin, and Yanan Liu. 2019. Representation Learning for Classical Planning from Partially Observed Traces. *arXiv preprint arXiv:1907.08352* (2019).
- [34] Fangkai Yang, Daoming Lyu, Bo Liu, and Steven Gustafson. 2018. Peorl: Integrating symbolic planning and hierarchical reinforcement learning for robust decision-making. *arXiv preprint arXiv:1804.07779* (2018).
- [35] Qiang Yang, Kangheng Wu, and Yunfei Jiang. 2007. Learning action models from plan examples using weighted MAX-SAT. *Artificial Intelligence* 171, 2-3 (2007), 107–143.
- [36] Terry Zimmerman and Subbarao Kambhampati. 2003. Learning-assisted automated planning: looking back, taking stock, going forward. *AI Magazine* 24, 2 (2003), 73–73.