

REFERENCES

- [1] George Ainslie. 2001. *Breakdown of Will*. Cambridge University Press.
- [2] Z. Akata, D. Balliet, M. de Rijke, F. Dignum, V. Dignum, G. Eiben, A. Fokkens, D. Grossi, K. Hindriks, H. Hoos, H. Hung, C. Jonker, C. Monz, M. Neerincx, F. Oliehoek, H. Prakken, S. Schlobach, L. van der Gaag, F. van Harmelen, H. van Hoof, B. van Riemsdijk, A. van Wylsberghe, R. Verbrugge, B. Verheij, P. Vossen, and M. Welling. 2020. A Research Agenda for Hybrid Intelligence: Augmenting Human Intellect With Collaborative, Adaptive, Responsible, and Explainable Artificial Intelligence. *Computer* 53, 8 (2020), 18–28.
- [3] Nadav Kiril Altschuler and David Sarne. 2018. Modeling Assistant’s Autonomy Constraints as a Means for Improving Autonomous Assistant-Agent Design. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems* (Stockholm, Sweden) (AAMAS ’18). International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1468–1476.
- [4] Xavier Boyen and Daphne Koller. 1998. Tractable Inference for Complex Stochastic Processes. In *UAI ’98: Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence, University of Wisconsin Business School, Madison, Wisconsin, USA, July 24-26, 1998*, Gregory F. Cooper and Serafin Moral (Eds.). Morgan Kaufmann, 33–42.
- [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 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (Eds.), 5443–5452.
- [6] Prashant Doshi, Xia Qu, Adam Goodie, and Diana L. Young. 2010. Modeling recursive reasoning by humans using empirically informed interactive POMDPs. In *9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2010), Toronto, Canada, May 10-14, 2010, Volume 1-3*, Wiebe van der Hoeek, Gal A. Kaminka, Yves Lesperance, Michael Luck, and Sandip Sen (Eds.). IFAAMAS, 1223–1230.
- [7] Avshalom Elmalech, David Sarne, Avi Rosenfeld, and Eden Shalom Erez. 2015. When Suboptimal Rules. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25-30, 2015, Austin, Texas, USA*, Blai Bonet and Sven Koenig (Eds.). AAAI Press, 1313–1319.
- [8] Owain Evans, Andreas Stuhlmüller, and Noah D. Goodman. 2016. Learning the Preferences of Ignorant, Inconsistent Agents. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA*, Dale Schuurmans and Michael P. Wellman (Eds.). AAAI Press, 323–329.
- [9] Owain Evans, Andreas Stuhlmüller, John Salvatier, and Daniel Filan. 2017. Modeling Agents with Probabilistic Programs. <http://agentmodels.org>. Accessed: 2021-5-17.
- [10] Eyal Even-Dar, Sham M. Kakade, and Yishay Mansour. 2007. The Value of Observation for Monitoring Dynamic Systems. In *IJCAI 2007, Proceedings of the 20th International Joint Conference on Artificial Intelligence, Hyderabad, India, January 6-12, 2007*, Manuela M. Veloso (Ed.), 2474–2479.
- [11] William Fedus, Carles Gelada, Yoshua Bengio, Marc G. Bellemare, and Hugo Larochelle. 2019. Hyperbolic Discounting and Learning over Multiple Horizons. *CoRR* abs/1902.06865 (2019). [arXiv:1902.06865](https://arxiv.org/abs/1902.06865)
- [12] 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.
- [13] Shane Frederick, George Loewenstein, and Ted O’Donoghue. 2002. Time discounting and time preference: A critical review. *Journal of Economic Literature* 40, 2 (2002), 351–401.
- [14] Samuel J Gershman, Eric J Horvitz, and Joshua B Tenenbaum. 2015. Computational rationality: A converging paradigm for intelligence in brains, minds, and machines. *Science* 349, 6245 (2015), 273–278.
- [15] P. Gmytrasiewicz and E. Durfee. 2004. Rational Coordination in Multi-Agent Environments. *Autonomous Agents and Multi-Agent Systems* 3 (2004), 319–350.
- [16] Piotr J. Gmytrasiewicz and Prashant Doshi. 2005. A Framework for Sequential Planning in Multi-Agent Settings. *J. Artif. Intell. Res.* 24 (2005), 49–79.
- [17] Leonard Green, Nathanael Fristoe, and Joel Myerson. 1994. Temporal discounting and preference reversals in choice between delayed outcomes. *Psychonomic Bulletin & Review* 1, 3 (1994), 383–389.
- [18] Nina Grgić-Hlača, Christoph Engel, and Krishna P. Gummadi. 2019. Human Decision Making with Machine Assistance: An Experiment on Bailing and Jailing. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 178 (Nov. 2019), 25 pages.
- [19] A. Heifetz and Y. Spiegel. 2000. On the Evolutionary Emergence of Optimism. *Microeconomic Theory eJournal* (2000).
- [20] Kristina M. Hengen and Georg W. Alpers. 2019. What’s the Risk? Fearful Individuals Generally Overestimate Negative Outcomes and They Dread Outcomes of Specific Events. *Frontiers in Psychology* 10 (2019), 1676.
- [21] Mark K. Ho and Thomas L. Griffiths. 2021. Cognitive science as a source of forward and inverse models of human decisions for robotics and control. *CoRR* abs/2109.00127 (2021). [arXiv:2109.00127](https://arxiv.org/abs/2109.00127)
- [22] Sammie Katt, Frans A. Oliehoek, and Christopher Amato. 2017. Learning in POMDPs with Monte Carlo Tree Search. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017 (Proceedings of Machine Learning Research, Vol. 70)*, Doina Precup and Yee Whye Teh (Eds.). PMLR, 1819–1827.
- [23] Kris N Kirby and Richard J Herrnstein. 1995. Preference reversals due to myopic discounting of delayed reward. *Psychological Science* 6, 2 (1995), 83–89.
- [24] Jon Kleinberg, Sigal Oren, and Manish Raghavan. 2016. Planning Problems for Sophisticated Agents with Present Bias. In *Proceedings of the 2016 ACM Conference on Economics and Computation* (Maastricht, The Netherlands) (EC ’16). Association for Computing Machinery, New York, NY, USA, 343–360.
- [25] Jon M. Kleinberg and Sigal Oren. 2014. Time-inconsistent planning: a computational problem in behavioral economics. In *ACM Conference on Economics and Computation, EC ’14, Stanford, CA, USA, June 8-12, 2014*, Moshe Babaioff, Vincent Conitzer, and David A. Easley (Eds.). ACM, 547–564.
- [26] Abdul Rahman Kreidieh, Cathy Wu, and Alexandre M. Bayen. 2018. Dissipating stop-and-go waves in closed and open networks via deep reinforcement learning. In *21st International Conference on Intelligent Transportation Systems, ITSC 2018, Maui, HI, USA, November 4-7, 2018*, Wei-Bin Zhang, Alexandre M. Bayen, Javier J. Sánchez Medina, and Matthew J. Barth (Eds.). IEEE, 1475–1480.
- [27] Heather C. Lench and Shane W. Bench. 2012. Automatic Optimism: Why People Assume Their Futures will be Bright. *Social and Personality Psychology Compass* 6, 4 (2012), 347–360. [arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1751-9004.2012.00430.x](https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1751-9004.2012.00430.x)
- [28] Richard L. Lewis, Andrew Howes, and Satinder Singh. 2014. Computational Rationality: Linking Mechanism and Behavior Through Bounded Utility Maximization. *Topics in Cognitive Science* 6, 2 (2014), 279–311.
- [29] George Loewenstein and Drazen Prelec. 1992. Anomalies in intertemporal choice: Evidence and an interpretation. *The Quarterly Journal of Economics* 107, 2 (1992), 573–597.
- [30] Harri Oinas-Kukkonen. 2012. A foundation for the study of behavior change support systems. *Personal and Ubiquitous Computing* 17 (2012), 1223–1235.
- [31] Frans A. Oliehoek and Christopher Amato. 2014. Best Response Bayesian Reinforcement Learning for Multiagent Systems with State Uncertainty. In *Proceedings of the Ninth AAMAS Workshop on Multi-Agent Sequential Decision Making in Uncertain Domains (MSDM)*.
- [32] Frans A. Oliehoek and Christopher Amato. 2016. *A Concise Introduction to Decentralized POMDPs*. Springer.
- [33] Raja Parasuraman and Victor Riley. 1997. Humans and automation: Use, misuse, disuse, abuse. *Human factors* 39, 2 (1997), 230–253.
- [34] Todd Rogers, Don A. Moore, and Michael I. Norton. 2017. The Belief in a Favorable Future. *Psychological Science* 28, 9 (2017), 1290–1301. [arXiv:https://doi.org/10.1177/0956797617706706](https://doi.org/10.1177/0956797617706706)
- [35] Stéphane Ross, Brahim Chaib-draa, and Joelle Pineau. 2007. Bayes-Adaptive POMDPs. In *Advances in Neural Information Processing Systems 20, Proceedings of the Twenty-First Annual Conference on Neural Information Processing Systems, Vancouver, British Columbia, Canada, December 3-6, 2007*, John C. Platt, Daphne Koller, Yoram Singer, and Sam T. Roweis (Eds.). Curran Associates, Inc., 1225–1232.
- [36] Dorsa Sadigh, Shankar Sastry, Sanjit A. Seshia, and Anca D. Dragan. 2016. Planning for Autonomous Cars that Leverage Effects on Human Actions. In *Robotics: Science and Systems XII, University of Michigan, Ann Arbor, Michigan, USA, June 18 - June 22, 2016*, David Hsu, Nancy M. Amato, Spring Berman, and Sam Ade Jacobs (Eds.).
- [37] Tali Sharot. 2011. The optimism bias. *Current Biology* 21, 23 (2011), R941–R945.
- [38] Eugene Vinitzky, Aboudy Kreidieh, Luc Le Flem, Nishant Kheterpal, Kathy Jang, Cathy Wu, Fangyu Wu, Richard Liaw, Eric Liang, and Alexandre M. Bayen. 2018. Benchmarks for reinforcement learning in mixed-autonomy traffic. In *Proceedings of The 2nd Conference on Robot Learning (Proceedings of Machine Learning Research, Vol. 87)*, Aude Billard, Anca Dragan, Jan Peters, and Jun Morimoto (Eds.). PMLR, 399–409.
- [39] Eugene Vinitzky, Kanaad Parvate, Aboudy Kreidieh, Cathy Wu, and Alexandre M. Bayen. 2018. Lagrangian Control through Deep-RL: Applications to Bottleneck Decongestion. In *21st International Conference on Intelligent Transportation Systems, ITSC 2018, Maui, HI, USA, November 4-7, 2018*, Wei-Bin Zhang, Alexandre M. Bayen, Javier J. Sánchez Medina, and Matthew J. Barth (Eds.). IEEE, 759–765.
- [40] Louise Whiteley and Maneesh Sahani. 2012. Attention in a Bayesian Framework. *Frontiers in Human Neuroscience* 6 (2012), 100.
- [41] Cathy Wu, Aboudy Kreidieh, Eugene Vinitzky, and Alexandre M. Bayen. 2017. Emergent Behaviors in Mixed-Autonomy Traffic. In *Proceedings of the 1st Annual Conference on Robot Learning (Proceedings of Machine Learning Research, Vol. 78)*, Sergey Levine, Vincent Vanhoucke, and Ken Goldberg (Eds.). PMLR, 398–407.
- [42] Luke Zettlemoyer, Brian Milch, and L. Kaelbling. 2008. Multi-Agent Filtering with Infinitely Nested Beliefs. In *NIPS*.