

REFERENCES

- [1] J. Audibert, S. Bubeck, and R. Munos. 2010. Best Arm Identification in Multi-Armed Bandits. In *Proceedings of the Conference On Learning Theory (COLT)*. 41–53.
- [2] L. Bisi, G. De Nittis, F. Trovò, M. Restelli, and N. Gatti. 2017. Regret Minimization Algorithms for the Followers Behaviour Identification in Leadership Games. In *Proceedings of the Conference on Uncertainty in Artificial (UAI)*. 1–10.
- [3] Mario Bravo, David Leslie, and Panayotis Mertikopoulos. 2018. Bandit learning in concave N-person games. In *Proceeding of the conference on Neural Information Processing Systems (NIPS)*. 5661–5671.
- [4] N. Brown and T. Sandholm. 2017. Safe and nested subgame solving for imperfect-information games. In *Proceeding of the conference on Neural Information Processing Systems (NIPS)*. 689–699.
- [5] Noam Brown and Tuomas Sandholm. 2018. Superhuman AI for heads-up no-limit poker: Libratus beats top professionals. *Science* 359, 6374 (2018), 418–424.
- [6] Noam Brown and Tuomas Sandholm. 2019. Superhuman AI for multiplayer poker. *Science* 365, 6456 (2019), 885–890.
- [7] S.R. Chowdhury and A. Gopalan. 2017. On kernelized multi-armed bandits. In *Proceedings of the International Conference on Machine Learning (ICML)*. 844–853.
- [8] Abraham D Flaxman, Adam Tauman Kalai, Adam Tauman Kalai, and H Brendan McMahan. 2005. Online convex optimization in the bandit setting: gradient descent without a gradient. In *Proceedings of the ACM-SIAM Symposium On Discrete Algorithms (SODA)*. 385–394.
- [9] A. Garivier, E. Kaufmann, and W.M. Koolen. 2016. Maximin action identification: A new bandit framework for games. In *Proceedings of the Conference On Learning Theory (COLT)*. 1028–1050.
- [10] Nicola Gatti, Alessandro Lazaric, Marco Rocco, and Francesco Trovò. 2015. Truthful learning mechanisms for multi-slot sponsored search auctions with externalities. *ARTIF INTELL* 227 (2015), 93–139.
- [11] Nicola Gatti and Marcello Restelli. 2011. Equilibrium approximation in simulation-based extensive-form games. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 199–206.
- [12] E. Kaufmann, O. Cappé, and A. Garivier. 2016. On the complexity of best-arm identification in multi-armed bandit models. *J MACH LEARN RES* 17, 1 (2016), 1–42.
- [13] Robert D Kleinberg. 2005. Nearly tight bounds for the continuum-armed bandit problem. In *Advances in Neural Information Processing Systems*. 697–704.
- [14] C.K. Ling, F. Fang, and J.Z. Kolter. 2018. What game are we playing? end-to-end learning in normal and extensive form games. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*. 396–402.
- [15] C.K. Ling, F. Fang, and J.Z. Kolter. 2019. Large Scale Learning of Agent Rationality in Two-Player Zero-Sum Games. In *Proceedings of the Conference on Artificial Intelligence (AAAI)*. 6104–6111.
- [16] Alberto Marchesi, Francesco Trovò, and Nicola Gatti. 2019. Learning Probably Approximately Correct Maximin Strategies in Simulation-Based Games with Infinite Strategy Spaces. (2019). arXiv:cs.GT/1911.07755
- [17] J. Nocedal and S. Wright. 2006. *Numerical optimization*. Springer Science & Business Media.
- [18] K.B. Petersen and M.S. Pedersen. 2008. The matrix cookbook. *Technical University of Denmark* 7, 15 (2008), 1–25.
- [19] J. Rong, T. Qin, and B. An. 2019. Competitive Bridge Bidding with Deep Neural Networks. In *Proceedings of the International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. 16–24.
- [20] Mark Rowland, Shayegan Omidshafiei, Karl Tuyls, Julien Perolat, Michal Valko, Georgios Piliouras, and Remi Munos. 2019. Multiagent Evaluation under Incomplete Information. In *Advances in Neural Information Processing Systems*. 12270–12282.
- [21] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. 2016. Mastering the game of Go with deep neural networks and tree search. *Nature* 529, 7587 (2016), 484–489.
- [22] M. Sion. 1958. On general minimax theorems. *PAC J MATH* 8, 1 (1958), 171–176.
- [23] S. Sokota, C. Ho, and B. Wiedenbeck. 2019. Learning Deviation Payoffs in Simulation-Based Games. In *AAAI Technical Track: Game Theory and Economic Paradigms*. 1–8.
- [24] N. Srinivas, A. Krause, S.M. Kakade, and M. Seeger. 2010. Gaussian process optimization in the bandit setting: No regret and experimental design. In *Proceedings of the International Conference on Machine Learning (ICML)*. 1015–1022.
- [25] M. Tambe. 2011. *Security and game theory: algorithms, deployed systems, lessons learned*. Cambridge university press.
- [26] K. Tuyls, J. Perolat, M. Lanctot, J.Z. Leibo, and T. Graepel. 2018. A generalised method for empirical game theoretic analysis. In *Proceedings of the International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. 77–85.
- [27] E.A. Viqueira, C. Cousins, Y. Mohammad, and Greenwald. A. 2019. Empirical Mechanism Design: Designing Mechanisms from Data. In *Proceedings of the Conference on Uncertainty in Artificial (UAI)*. 1–11.
- [28] E.A. Viqueira, A. Greenwald, C. Cousins, and E. Upfal. 2019. Learning Simulation-Based Games from Data. In *Proceedings of the International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. 1778–1780.
- [29] H. von Stackelberg. 1934. *Marktform und Gleichgewicht*. Springer, Vienna.
- [30] Y. Vorobeychik and M.P. Wellman. 2009. Strategic analysis with simulation-based games. In *Proceedings of the IEEE Winter Simulation Conference (WSC)*. 359–372.
- [31] Y. Vorobeychik, M.P. Wellman, and S. Singh. 2007. Learning payoff functions in infinite games. *MACH LEARN* 67, 1-2 (2007), 145–168.
- [32] B. Wiedenbeck, F. Yang, and M.P. Wellman. 2018. A regression approach for modeling games with many symmetric players. In *Proceedings of the Conference on Artificial Intelligence (AAAI)*. 1266–1273.
- [33] B. Wilder, B. Dilkina, and M. Tambe. 2019. Melding the data-decisions pipeline: Decision-focused learning for combinatorial optimization. In *Proceedings of the Conference on Artificial Intelligence (AAAI)*, Vol. 33. 1658–1665.
- [34] C.K.I. Williams and C.E. Rasmussen. 2006. *Gaussian processes for machine learning*. Vol. 2. MIT press Cambridge, MA.
- [35] M. Wright and M.P. Wellman. 2019. Probably Almost Stable Strategy Profiles in Simulation-Based Games. In *Proceedings of the Games, Agents, and Incentives Workshops (GAIW)*. 1–9.