[21] Le Cong Dinh, Stephen Marcus McAleer, Zheng Tian, Nicolas Perez-Nieves, Oliver Slumbers, David Henry Mguni, Jun Wang, Haitham Bou Ammar, and Yaodong Yang. 2022. Online Double Oracle. *Transactions on Machine Learning Research* (2022). https://openreview.net/forum?id=rrMK6hYNSx

[22] John Fearnley, Martin Gairing, Paul W Goldberg, and Rahul Savani. 2015. Learning equilibria of games via payoff queries. *J. Mach. Learn. Res.* 16 (2015), 1305–1344.

[23] John Fearnley, Tobenna Peter Igwe, and Rahul Savani. 2015. An empirical study of finding approximate equilibria in bimatrix games. In *International Symposium on Experimental Algorithms*. Springer, 339–351.

[24] John Fearnley and Rahul Savani. 2016. Finding approximate Nash equilibria of bimatrix games via payoff queries. *ACM Transactions on Economics and Computation (TEAC)* 4, 4 (2016), 1–19.

[25] Filiberto Fele and Kostas Margellos. 2020. Probably approximately correct Nash equilibrium learning. *IEEE Trans. Automat. Control* (2020), 4238–4245.

[26] Xidong Feng, Oliver Slumbers, Ziyu Wan, Bo Liu, Stephen McAleer, Ying Wen, Jun Wang, and Yaodong Yang. 2021. Neural auto-curricula in two-player zero-sum games. *Advances in Neural Information Processing Systems* 34 (2021), 3504–3517.

[27] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. *Deep learning*. MIT press.

[28] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. *Advances in neural information processing systems* 27 (2014).

[29] Keegan Harris, Ioannis Anagnostides, Gabriele Farina, Mikhail Khodak, Steven Wu, and Tuomas Sandholm. 2023. Meta-Learning in Games. In *International Conference on Learning Representations*. https://openreview.net/forum?id=uHaWaNhCvZD

[30] Sergiu Hart and Andreu Mas-Colell. 2000. A simple adaptive procedure leading to correlated equilibrium. *Econometrica* 68, 5 (2000), 1127–1150.

[31] David Haussler. 1990. *Probably approximately correct learning*. University of California, Santa Cruz, Computer Research Laboratory.

[32] Howard Heaton, Daniel McKenzie, Qiuwei Li, Samy Wu Fung, Stanley Osher, and Wotao Yin. 2021. Learn to Predict Equilibria via Fixed Point Networks. *arXiv preprint arXiv:2106.00906* (2021).

[33] Kurt Hornik, Maxwell B. Stinchcombe, and Halbert L. White. 1989. Multilayer feedforward networks are universal approximators. *Neural Networks* 2 (1989), 359–366.

[34] Junling Hu and Michael P Wellman. 2003. Nash Q-learning for general-sum stochastic games. *Journal of machine learning research* 4, Nov (2003), 1039–1069.

[35] Chi Jin, Qinghua Liu, and Sobhan Miryoosefi. 2021. Bellman eluder dimension: New rich classes of rl problems, and sample-efficient algorithms. *Advances in neural information processing systems* 34 (2021), 13406–13418.

[36] Chi Jin, Qinghua Liu, Yuanhao Wang, and Tiancheng Yu. 2022. V-Learning – A Simple, Efficient, Decentralized Algorithm for Multiagent RL. In *ICLR 2022 Workshop on Gamification and Multiagent Solutions*.

[37] Spyros C Kontogiannis, Panagiota N Panagopoulou, and Paul G Spirakis. 2006. Polynomial algorithms for approximating Nash equilibria of bimatrix games. In *International Workshop on Internet and Network Economics*. Springer, 286–296.

[38] Spyros C Kontogiannis and Paul G Spirakis. 2007. Efficient algorithms for constant well supported approximate equilibria in bimatrix games. In *International Colloquium on Automata, Languages, and Programming*. Springer, 595–606.

[39] Akshay Krishnamurthy, Alekh Agarwal, and John Langford. 2016. PAC reinforcement learning with rich observations. *Advances in Neural Information Processing Systems* 29 (2016).

[40] Marc Lanctot, Vinicius Zambaldi, Audrunas Gruslys, Angeliki Lazaridou, Karl Tuyls, Julien Pérolat, David Silver, and Thore Graepel. 2017. A unified game-theoretic approach to multiagent reinforcement learning. *Advances in neural information processing systems* 30 (2017).

[41] Jiayang Li, Jing Yu, Yu Nie, and Zhaoran Wang. 2020. End-to-end learning and intervention in games. *Advances in Neural Information Processing Systems* 33 (2020).

[42] Chun Kai Ling, Fei Fang, and J. Zico Kolter. 2018. What Game Are We Playing? End-to-end Learning in Normal and Extensive Form Games. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*. International Joint Conferences on Artificial Intelligence Organization, 396–402. https://doi.org/10.24963/ijcai.2018/55

[43] Chun Kai Ling, Fei Fang, and J Zico Kolter. 2019. Large scale learning of agent rationality in two-player zero-sum games. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 6104–6111.

[44] Edward Lockhart, Marc Lanctot, Julien Pérolat, Jean-Baptiste Lespiau, Dustin Morrill, Finbarr Timbers, and Karl Tuyls. 2019. Computing Approximate Equilibria in Sequential Adversarial Games by Exploitability Descent.. In *IJCAI*, Sarit Kraus (Ed.). ijcai.org, 464–470.

[45] Alberto Marchesi, Francesco Trovò, and Nicola Gatti. 2020. Learning Probably Approximately Correct Maximin Strategies in Simulation-Based Games with Infinite Strategy Spaces. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*. 834–842.

[46] Luke Marris, Ian Gemp, Thomas Anthony, Andrea Tacchetti, Siqi Liu, and Karl Tuyls. 2022. Turbocharging Solution Concepts: Solving NEs, CEs and CCEs with Neural Equilibrium Solvers. In *Advances in Neural Information Processing Systems*, Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (Eds.). https://openreview.net/forum?id=RczPtvlaXPH

[47] H Brendan McMahan, Geoffrey J Gordon, and Avrim Blum. 2003. Planning in the presence of cost functions controlled by an adversary. In *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*. 536–543.

[48] Dov Monderer and Lloyd S Shapley. 1996. Fictitious play property for games with identical interests. *Journal of economic theory* 68, 1 (1996), 258–265.

[49] John F Nash et al. 1950. Equilibrium points in n-person games. *Proceedings of the national academy of sciences* 36, 1 (1950), 48–49.

[50] Eugene Nudelman, Jennifer Wortman, Yoav Shoham, and Kevin Leyton-Brown. 2004. Run the GAMUT: A comprehensive approach to evaluating game-theoretic algorithms. In *AAMAS*, Vol. 4. 880–887.

[51] Nicolas Perez-Nieves, Yaodong Yang, Oliver Slumbers, David H Mguni, Ying Wen, and Jun Wang. 2021. Modelling Behavioural Diversity for Learning in Open-Ended Games. In *International Conference on Machine Learning*. PMLR, 8514–8524.

[52] Kevin Scaman and Aladin Virmaux. 2018. Lipschitz regularity of deep neural networks: analysis and efficient estimation. In *NeurIPS*. 3839–3848.

[53] Peter Schuster and Karl Sigmund. 1983. Replicator dynamics. *Journal of theoretical biology* 100, 3 (1983), 533–538.

[54] Pier Giuseppe Sessa, Ilija Bogunovic, Andreas Krause, and Maryam Kamgarpour. 2020. Contextual games: Multi-agent learning with side information. *Advances in Neural Information Processing Systems* 33 (2020), 21912–21922.

[55] Shai Shalev-Shwartz and Shai Ben-David. 2014. *Understanding machine learning: From theory to algorithms*. Cambridge university press.

[56] Yoav Shoham and Kevin Leyton-Brown. 2008. *Multiagent systems: Algorithmic, game-theoretic, and logical foundations*. Cambridge University Press.

[57] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus. 2014. Intriguing properties of neural networks. In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, Yoshua Bengio and Yann LeCun (Eds.).

[58] Haralampos Tsaknakis and Paul G Spirakis. 2007. An optimization approach for approximate Nash equilibria. In *International Workshop on Web and Internet Economics*. Springer, 42–56.

[59] Leslie G Valiant. 1984. A theory of the learnable. *Commun. ACM* 27, 11 (1984), 1134–1142.

[60] Enrique Areyan Viqueira, Cyrus Cousins, Eli Upfal, and Amy Greenwald. 2019. Learning equilibria of simulation-based games. *arXiv preprint arXiv:1905.13379* (2019).

[61] Yaodong Yang and Jun Wang. 2020. An Overview of Multi-Agent Reinforcement Learning from Game Theoretical Perspective. *arXiv preprint arXiv:2011.00583* (2020).

[62] Jing Zhang and Ioannis Ch Paschalidis. 2017. Data-driven estimation of travel latency cost functions via inverse optimization in multi-class transportation networks. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. IEEE, 6295–6300.

[63] Tong Zhang. 2002. Covering number bounds of certain regularized linear function classes. *Journal of Machine Learning Research* 2, Mar (2002), 527–550.

[64] Ding-Xuan Zhou. 2002. The covering number in learning theory. *Journal of Complexity* 18, 3 (2002), 739–767.