

Research Statement

Sam Ganzfried
sam.ganzfried@gmail.com

My main research interests are in game theory and artificial intelligence. Many conceptual and computational challenges arise when developing strong agents for strategic multiagent environments. The first challenge is to determine the desired solution concept. The standard solution concept in game theory is Nash equilibrium. In two-player zero-sum games this is a compelling solution, but for non-zero-sum and multiplayer games there are theoretical limitations; games can have multiple Nash equilibria with different values for the players, and following one of them has no expected payoff guarantee.¹ Even for two-player zero-sum games there can be many Nash equilibria and some may be more desirable than others. The second challenge is to compute the desired solution concept in games of interest. In two-player zero-sum games this can be done in polynomial time, but in non-zero-sum and multiplayer games it is PPAD-hard and widely conjectured that no efficient algorithms exist. Even in the two-player zero-sum case specific games of interest may be so large that standard algorithms are not sufficiently scalable. In addition to having multiple players, games can exhibit many other elements of complexity: very large state and/or action spaces (which may be infinite), imperfect information (where some elements of the game state are not known to all players), and repeated interactions of potentially unknown duration (which can theoretically be infinite). I am interested in developing practical approaches for creating strong agents in games exhibiting some or all of these complexities. Much of standard game-theoretic analysis assumes that all agents are perfectly rational and that no further information is available about the opposing agents. New challenges arise when additional information is available—either in the form of observations from current gameplay or historical data on the specific opponents or a larger population. When such data is available, we would like to devise opponent modeling algorithms that integrate the game-theoretic strategies with machine learning approaches that incorporate the data. A final challenge is that standard game-theoretic algorithms may compute strategies that are unintelligible to humans and not feasible to implement under real-world time constraints.

I have devised efficient algorithms for computing Nash equilibrium strategies in several challenging game classes. I have developed the most efficient exact algorithm for computing Nash equilibrium strategies in games with three or more players [8]. This algorithm is based on a novel mixed-integer programming formulation. This algorithm computes an exact Nash equilibrium quickly for small games, but does not scale to large games. For larger games we must consider iterative algorithms such as fictitious play and counterfactual regret minimization, both of which have no performance guarantee. I have shown that fictitious play leads to better approximation error than counterfactual regret minimization for approximating Nash equilibrium in multiplayer games [9]. I have also developed improved initialization approaches for fictitious play that lead to significantly better performance [13] and showed how this has enabled us to solve a canonical counterexample in which fictitious play does not converge to Nash equilibrium [15]. I have also devised efficient algorithms for computing very close approximations of Nash equilibrium strategies in multiplayer perfect and imperfect-information stochastic games. A stochastic game involves repeated play of stage games until a terminal state is reached, which has unknown (and potentially infinite) duration. These

¹Despite these theoretical limitations, I have shown that exact Nash equilibrium strategies outperformed a variety of agents created for a class project on three-player Kuhn poker, a simplified version of poker that has been used as a testbed in the Annual Computer Poker Competition [21]. So there is still strong reason to expect Nash equilibrium strategies to perform well in practice.

algorithms were applied to a three-player poker tournament [22, 23], and a four-player naval planning scenario [20, 11]. I have also developed improved algorithms using parallelism for more efficiently computing best responses [12], which is useful for determining the approximation error. I have also devised algorithms for computing Nash equilibrium strategies in continuous games (where the pure strategy spaces are infinite). I have applied these to endgames of limit Texas hold 'em [24], and to the continuous Blotto game, a well-studied model of resource allocation with applications to national security, voting, and auctions [10].

Many important games are so large that we cannot hope to directly compute a Nash equilibrium in the full game. For example, no-limit Texas hold 'em has approximately 10^{165} states in its game tree [35]. The standard paradigm for solving such large games had previously been to first apply an abstraction algorithm that computed a smaller strategically similar game, then solve the abstract game, and finally to map the solution of the abstract game back to an approximate solution of the full game. I have developed new improved algorithms for game abstraction that determine which information sets should be grouped together [27, 2]. I have also developed new approaches for reverse-mapping the abstract equilibrium to the full game; I devised a new approach for interpreting opposing actions that have been omitted from the action abstraction [26], and approaches for rounding strategy probabilities from the abstract equilibrium to obtain stronger performance in the full game [30, 2]. Finally, I have devised a new paradigm called endgame solving that has been widely recognized to be the key breakthrough enabling superhuman performance in no-limit Texas hold 'em [28, 6]. Endgame solving involves solving the portion of the game that we have actually reached during game play online in real time to a finer degree of granularity than in the offline computation. While I showed that theoretically this approach may result in a non-equilibrium full-game strategy (even if the trunk strategy were an exact Nash equilibrium), this approach has consistently led to significantly improved performance. This approach was used by the agent Claudico that competed in the inaugural Brains vs. AI competition against the strongest human two-player no-limit Texas hold 'em specialists [7], and subsequently defeated human professional players with the aid of extensive computational resources [36, 3]. This approach was also critical to recent superhuman play in 6-player no-limit Texas hold 'em [4].

In addition to developing algorithms for computing and approximating Nash equilibrium strategies, I have also devised several new game-theoretic solution concepts that may be more appropriate in certain situations. While Nash equilibrium assumes that all players are behaving rationally, if the opponents are irrational (or follow strategies from a different Nash equilibrium) we may obtain an extremely low payoff. I proposed a new solution concept called safe equilibrium that models opponents as behaving rationally with a specified probability and behaving potentially arbitrarily with the remaining probability [18]. Safe equilibrium explicitly balances between rationality and safety, enabling us to construct strategies that are robust to arbitrary degrees of opponents' irrationality. I showed that safe equilibrium has similar existence and computational properties to Nash equilibrium, and developed practical algorithms for its computation. I have also devised a new equilibrium refinement concept for sequential imperfect-information games called observable perfect equilibrium (OPE) [17]. The main equilibrium refinement solution concepts allow the players to make small "trembles" or mistakes arbitrarily on all actions (ensuring that we are playing an equilibrium that is robust to such mistakes), while OPE allows only mistakes from the opponents that are consistent with our observations of the path of play. We show that OPE results in a solution to two-player no-limit poker that differs from the prior refinement concepts. OPE correctly captures the assumption that the opponent is playing as rationally as possible given mistakes that have been observed.

When we are playing against opponents who may be irrational and we have access to historical data and/or observations of their play, we would like to capitalize on this information to obtain higher payoff than following a static equilibrium-based strategy. This is particularly challenging when we are playing against a set of unknown opponents, some of whom may be very strong and potentially deceptive. I developed an algorithm that is successfully able to exploit weaknesses of opponents in extremely large imperfect-information games after only a small number of interactions [25]. It uses an approximate Nash equilibrium strategy as the prior, which is updated based on observations of opponents' play. The algorithm, called

deviation-based best response, creates an opponent model that assumes that the opponent follows the “closest” strategy to the Nash equilibrium approximation subject to being consistent with our observations of their play. This is the first opponent modeling algorithm to incorporate behavioral strategy constraints. This approach led to a large performance improvement against a variety of opponents in two-player limit Texas hold ’em. I have also developed new approaches with theoretical guarantees even against strong deceptive opponents [29]. I have shown that in certain games it is actually possible to deviate from repeatedly playing a one-shot equilibrium strategy in order to exploit perceived weaknesses of an opponent, while still guaranteeing at least the value of the game in expectation against any opponent. Those approaches applied to two-player zero-sum games. Recently I have developed an efficient algorithm for opponent modeling in multiplayer imperfect-information games that significantly outperformed a wide variety of real opponents and several exact Nash equilibrium strategies in a three-player poker variant [31].

In many settings we would like to enable humans to make important decisions under time constraints, not produce massive binary strategy files that are only intelligible to computers. I have designed an algorithm that exploits qualitative information about equilibrium structure to improve the speed of equilibrium finding and produces strategies that are more human understandable [24]. I showed that for the final round of limit Texas hold ’em, equilibrium strategies for any input hand distribution will conform to one of three relatively simple qualitative action structures. I have extended this direction to develop a framework for computing human-understandable strategies from which we have been able to deduce several new fundamental rules of poker strategy that were not previously known for when to make to make extremely large bets or when to not bet at all (unconventional bet sizes were critical for recent poker AI successes) [32]. I have also uncovered a simple poker rule for determining when to call a bet of arbitrary size that significantly outperforms the popular “minimum defense frequency” rule [19]. Recently I have proposed a new framework that enables human decision makers to make fast decisions in parametrized games (games that contain parameters affecting payoffs, action spaces, and/or information states) without the aid of real-time solvers [14]. In this framework we learn a new structure called a parametric decision list, which contains a small set of rules (i.e., a “cheat sheet”) that dictate which strategy should be played for every possible parameter value in a way that can be easily understood and implemented by a human.

I have also done some research that falls outside my primary research interest of computational game theory; I am generally interested in AI, optimization, and Bayesian statistics, and their application to national security, medicine, education, socialization, and weather. I have been working on a new algorithm for educational grading that optimizes over the two main existing approaches (absolute scale and grading curve), creating a new robust, fair, transparent approach that addresses limitations of the existing approaches. I have also studied the effect of modifying question weights to improve the quality of exams [33]. I have studied the problem of determining the optimal number of choices in rating contexts, which can have applications including paper reviewing, exam grading, and online dating [34]. I have supervised several projects on applications of deep learning to hurricane prediction [1, 5]. Recently I have also proposed a novel dynamic model of a relationship between two people where the states depend on the “power” in the relationship and analyzed its stability [16]. While the model is motivated by a social or romantic relationship, it can also be applied to professional or business relationships as well as diplomatic relationships between nations.

To summarize, I have devised novel approaches for computing Nash equilibria in challenging game classes, scalable approaches for extremely large domains, a novel solution concept that balances between rationality and safety, a novel equilibrium refinement concept for sequential imperfect-information games, robust approaches for opponent modeling and exploitation, and approaches for computation of human-understandable strategies. While I have made substantial progress on all these fronts, significant challenges still remain for each of them. In addition to studying foundational questions, I have applied these algorithms to large-scale domains in poker and national security. These approaches are largely domain-independent and would apply to any problem that can be represented within the game classes described. In the future I have particular interest in applying these approaches to disease treatment, national security, and AI safety.

References

- [1] Sheila Alemany, Jonathan Beltran, Adrian Perez, and Sam Ganzfried. Predicting hurricane trajectories using a recurrent neural network. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2019.
- [2] Noam Brown, Sam Ganzfried, and Tuomas Sandholm. Hierarchical abstraction, distributed equilibrium computation, and post-processing, with application to a champion no-limit Texas Hold'em agent. In *Proceedings of the International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2015.
- [3] Noam Brown and Tuomas Sandholm. Superhuman AI for heads-up no-limit poker: Libratus beats top professionals. *Science*, 359:418–424, 2017.
- [4] Noam Brown and Tuomas Sandholm. Superhuman AI for multiplayer poker. *Science*, 365:885–890, 2019.
- [5] Max Chiswick and Sam Ganzfried. Prediction of Bayesian intervals for tropical storms. In *Proceedings of the Florida Artificial Intelligence Research Society Conference (FLAIRS)*, 2020.
- [6] Sam Ganzfried. Endgame solving: The surprising breakthrough that enabled superhuman two-player no-limit Texas hold 'em play. In *International Conference on Game Theory*, 2017.
- [7] Sam Ganzfried. Reflections on the first man vs. machine no-limit Texas hold 'em competition. *AI Magazine*, 38(2), 2017.
- [8] Sam Ganzfried. Fast complete algorithm for multiplayer Nash equilibrium, 2020. arXiv:2002.04734 [cs.GT].
- [9] Sam Ganzfried. Fictitious play outperforms counterfactual regret minimization, 2020. arXiv:2001.11165 [cs.GT].
- [10] Sam Ganzfried. Algorithm for computing approximate Nash equilibrium in continuous games with application to continuous Blotto. *Games*, 12(47), 2021. Special issue on Economics of Conflict and Terrorism.
- [11] Sam Ganzfried. Computing Nash equilibria in multiplayer DAG-structured stochastic games with persistent imperfect information. In *Proceedings of the Conference on Decision and Game Theory for Security (GameSec)*, 2021.
- [12] Sam Ganzfried. Best response computation in multiplayer imperfect-information stochastic games. In *Proceedings of the Florida Artificial Intelligence Research Society Conference (FLAIRS)*, 2022.
- [13] Sam Ganzfried. Fictitious play with maximin initialization. In *IEEE Conference on Decision and Control (CDC)*, 2022.
- [14] Sam Ganzfried. Human strategic decision making in parametrized games. *Mathematics*, 10(1147), 2022. Special issue on Game Theory and Artificial Intelligence.
- [15] Sam Ganzfried. Random initialization solves Shapley's fictitious play counterexample, 2022. arXiv:2209.02154 [cs.GT].
- [16] Sam Ganzfried. Stable relationships, 2022. arXiv:2206.06468 [cs.MA].
- [17] Sam Ganzfried. Observable perfect equilibrium. In *Proceedings of the Conference on Decision and Game Theory for Security (GameSec)*, 2023.
- [18] Sam Ganzfried. Safe equilibrium. In *IEEE Conference on Decision and Control (CDC)*, 2023.
- [19] Sam Ganzfried and Max Chiswick. Most important fundamental rule of poker strategy. In *Proceedings of the Florida Artificial Intelligence Research Society Conference (FLAIRS)*, 2020.
- [20] Sam Ganzfried, Conner Laughlin, and Charles Morefield. Parallel algorithm for Nash equilibrium in multiplayer stochastic games with application to naval strategic planning. In *Proceedings of the International Conference on Distributed Artificial Intelligence (DAI)*, 2020.
- [21] Sam Ganzfried, Austin Nowak, and Joannier Pinales. Successful Nash equilibrium agent for a 3-player imperfect-information game. *Games*, 9(33), 2018.

- [22] Sam Ganzfried and Tuomas Sandholm. Computing an approximate jam/fold equilibrium for 3-player no-limit Texas hold 'em tournaments. In *Proceedings of the International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2008.
- [23] Sam Ganzfried and Tuomas Sandholm. Computing equilibria in multiplayer stochastic games of imperfect information. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*, 2009.
- [24] Sam Ganzfried and Tuomas Sandholm. Computing equilibria by incorporating qualitative models. In *Proceedings of the International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2010.
- [25] Sam Ganzfried and Tuomas Sandholm. Game theory-based opponent modeling in large imperfect-information games. In *Proceedings of the International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2011.
- [26] Sam Ganzfried and Tuomas Sandholm. Action translation in extensive-form games with large action spaces: Axioms, paradoxes, and the pseudo-harmonic mapping. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 2013.
- [27] Sam Ganzfried and Tuomas Sandholm. Potential-aware imperfect-recall abstraction with earth mover's distance in imperfect-information games. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2014.
- [28] Sam Ganzfried and Tuomas Sandholm. Endgame solving in large imperfect-information games. In *Proceedings of the International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2015.
- [29] Sam Ganzfried and Tuomas Sandholm. Safe opponent exploitation. *ACM Transactions on Economics and Computation (TEAC)*, 3(8):1–28, 2015.
- [30] Sam Ganzfried, Tuomas Sandholm, and Kevin Waugh. Strategy purification and thresholding: Effective non-equilibrium approaches for playing large games. In *Proceedings of the International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2012.
- [31] Sam Ganzfried, Kevin A. Wang, and Max Chiswick. Bayesian opponent modeling in multiplayer imperfect-information games, 2022. arXiv:2212.06027 [cs.GT].
- [32] Sam Ganzfried and Farzana Yusuf. Computing human-understandable strategies: Deducing fundamental rules of poker strategy. *Games*, 8(4):1–13, 2017.
- [33] Sam Ganzfried and Farzana Yusuf. Optimal weighting for exam composition. *Education Sciences*, 8(36), 2018. Special issue on Artificial Intelligence and Education.
- [34] Sam Ganzfried and Farzana Yusuf. Optimal number of choices in rating contexts. *Big Data and Cognitive Computing*, 3(48), 2019. Special Issue on Computational Models of Cognition and Learning. Early version presented at 2016 *NIPS* Workshop on Imperfect Decision Makers: Admitting Real-World Rationality.
- [35] Michael Johanson. Measuring the size of large no-limit poker games. Technical report, University of Alberta, 2013.
- [36] Matej Moravčík, Martin Schmid, Neil Burch, Viliam Lisý, Dustin Morrill, Nolan Bard, Trevor Davis, Kevin Waugh, Michael Johanson, and Michael Bowling. Deepstack: Expert-level artificial intelligence in heads-up no-limit poker. *Science*, 356:508–513, 2017.