Analyzing Games with a Variable Number of Players

Madelyn Gatchel
Davidson College
magatchel@davidson.edu

Abstract

We introduce a novel technique that uses a multi-headed neural network to analyze symmetric games with a variable number of players, where the number of participants falls in a specified range. We hypothesize that the payoffs in a game with x players are similar or related to the same game with $x \pm 1$ players, given a large value of x. With this hypothesis, we generalize prior work to analyze games with a large, variable number of players.

Game theory is the branch of economics that aims to understand self-interested agents interacting and making decisions. The normal-form representation of simultaneous move games is a mathematical model of incentives. This model includes a fixed set of players, their corresponding strategy sets, and a payoff matrix that describes each player's payoff function. Using this model we can approximate Nash equilibria in the game, which are predictions about behavior in the real-world interaction. Because the interaction we are studying might have an uncertain number of participants, a game with a fixed number of players might be an insufficient model. Thus we focus on analyzing games with a variable number of players, where the number of players falls in a specified range.

In real-world settings it is also likely that the number of participants is large. Sokota et al. (2019) apply machine learning techniques to analyze games with a large number of players. We hypothesize that the payoffs in a game with x players are similar or related to the payoffs in the same game with $x \pm 1$ players, given a large value of x. With this hypothesis, we generalize Sokota, Ho, and Wiedenbeck's results to analyze games with a large, variable number of players. The new goal is to develop analyses that accommodate this uncertainty in the number of players, such as finding equilibria that are robust within the range of possible player values.

Our techniques are particularly relevant for analyzing simulation-based games, where the payoff matrix is not known in advance but can be filled through a series of multi-agent simulations. Because mathematical models often require greater precision or a deeper understanding of the interaction than is feasible, simulation-based game theory is a realistic alternative that utilizes agent-based modeling to construct a game-theoretic model.

For example, Wah and Wellman (2016) construct a simulation-based game to analyze the effects of latency arbitrage in financial markets. They analyze the same game with a variable number of background traders (24, 58, 238), but analyze each instance of the game separately. For each of the three instances, they find symmetric equilibria in the game and then evaluate background-trader surplus and latency arbitrageur profit (if applicable) to conclude that latency arbitrage reduces surplus overall. Using our techniques, Wah and Wellman could analyze the entire game at once, with the number of players ranging from 24 to 238, instead of analyzing each instance separately. The hope is that having equilibria that are robust within the range of player values will enable more meaningful analysis of and predictions about real-world games.

Related Work

The normal-form payoff matrix stores each player's payoff for every possible combination of actions the players can simultaneously play. As the number of players increases, the size of the payoff matrix increases exponentially. While there are many compact representations of games, they generally require some knowledge about the underlying structure of the game, which is not known in simulation-based settings. Several papers have addressed this issue by learning game models from data (Areyan Viqueira et al., 2020; Wiedenbeck et al., 2018; Sokota et al., 2019).

In particular, Sokota et al. (2019) use a neural network to learn the deviation payoff function. A deviation payoff is the expected payoff a player would receive by deviating or changing strategies, given the mixed strategies everyone else is playing. This learned deviation payoff function is used in Nash-finding algorithms to find approximate equilibria in simulation-based games with a large number of players, without constructing an explicit payoff table.

While several papers consider the same game with a different number of players (Tuffin and Maillé 2006; Honorio and Ortiz 2015; Petruzzi, Pitt, and Busquets 2017), the authors do so to validate the scalability of their model and not as the focus of their analysis. We believe that this demonstrates a need for analysis that spans different numbers of players and that our techniques would be applicable.

Learning Model

Sokota et al. (2019) use a multi-headed neural network to learn the deviation payoff function for games with a fixed number of players. Games with a variable number of players essentially have a different, but related, payoff function for each instance of the game within the player value range. As a result, the learning problem is much more complex.

Our preliminary results indicate that it is possible to learn the deviation payoff function for the entire game with an added input dimension specifying the number of players. We have evaluated network performance on 50 randomly generated symmetric games (from game classes used in prior literature) with 3 and 5 strategies and a player range of 50 to 100. We achieve decent deviation payoff approximations by training on about half as much data as we would need if we trained each instance separately. However, there is room for improvement by continuing to optimize the hyperparameters and determining how and when to refit the network after running the Nash-finding algorithm.

Equilibrium Robustness

Typical game-theoretic analysis involves finding approximate Nash equilibria in a game with a fixed number of players. However, in games with a variable number of players, finding approximate Nash equilibria for the entire game is not as straightforward. For example, an approximate equilibrium in one instance of the game might not be an approximate equilibrium in any other instances and therefore it is not a good prediction of a stable state of the game. Thus we seek to find a metric to evaluate the robustness of a candidate equilibrium, or the stability of an equilibrium across all instances of the game.

We have considered two metrics to evaluate whether a profile is an equilibrium across the range of player values: average regret and equilibrium frequency. Regret is the maximum payoff amount any player can gain by deviating to any other strategy. The regret of a candidate Nash equilibrium is computed for each instance of the game and then averaged. If the average regret falls under some specified regret threshold, we say the candidate equilibrium is a robust equilibrium in the game. While in many cases this metric properly filters out mixed strategies that are not robust across all instances, it is possible that this metric filters out a candidate equilibrium that is indeed robust. For example, a candidate Nash equilibrium might be an approximate Nash equilibrium for many instances of the game but have a high enough regret for one or a few instances and thus not be classified as a robust approximate equilibrium based on average regret.

The equilibrium frequency metric involves counting the number of instances in which the candidate Nash equilibrium is an approximate equilibrium, and if the count is higher than some threshold, then the candidate equilibrium is a robust equilibrium in the game. This metric does not overpenalize a candidate equilibrium for having a high regret for a few instances but being a good approximate Nash equilibrium overall. Figure 1 shows how many times each mixture in the simplex was an approximate equilibrium (for a fixed ϵ) in a randomly generated 3-strategy game with

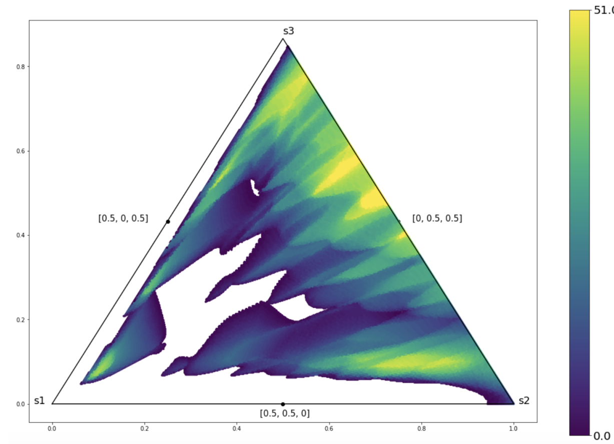

Figure 1: Approximate equilibrium frequency simplex

50 to 100 players. Each point in the simplex corresponds to a symmetric mixed strategy, and the color shows how many times that profile was an approximate equilibrium. The white points in the simplex correspond to profiles that were never approximate Nash equilibria, and the brighter points correspond to mixed strategies that were approximate Nash equilibria in many instances of the game.

Ongoing Work

We are currently exploring how to extend our method to generalize over other continuous parameters of the simulation environment. We are also exploring other robustness metrics that are specific to the context of a given real-world game.

References

Areyan Viqueira, E.; Cousins, C.; and Greenwald, A. 2020. Improved Algorithms for Learning Equilibria in Simulation-Based Games. In *AAMAS*, 79–87.

Honorio, J.; and Ortiz, L. 2015. Learning the structure and parameters of large-population graphical games from behavioral data. *Journal of Machine Learning Research* 16(1): 1157–1210.

Petruzzi, P. E.; Pitt, J.; and Busquets, D. 2017. Electronic Social Capital for Self-Organising Multi-Agent Systems. *ACM Trans. Auton. Adapt. Syst.* 12(3).

Sokota, S.; Ho, C.; and Wiedenbeck, B. 2019. Learning deviation payoffs in simulation-based games. *AAAI* 33(1): 2173–2180.

Tuffin, B.; and Maillé, P. 2006. How many parallel TCP sessions to open: A pricing perspective. In *Performability Has its Price*, 2–12. Springer Berlin Heidelberg.

Wah, E.; and Wellman, M. P. 2016. Latency arbitrage in fragmented markets: A strategic agent-based analysis. *Algorithmic Finance* 5(3-4): 69–93.

Wiedenbeck, B.; Yang, F.; and Wellman, M. P. 2018. A regression approach for modeling games with many symmetric players. In *AAAI*, 1266–1273.