Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI-16)

Evaluating the Robustness of Game Theoretic Solutions When Using Abstraction

Oscar Veliz
University of Texas at El Paso
500 W. University Ave, El Paso, TX, 79968, USA
(915) 747-6373 osveliz@miners.utep.edu

Abstract

Game theory is a tool for modeling multi-agent decision problems and has been used to analyze strategies in domains such as poker, security, and trading agents. One method for solving very large games is to use abstraction techniques to shrink the game by removing detail, solve the reduced game, and then translate the solution back to the original game. We present a methodology for evaluating the robustness of different game-theoretic solution concepts to the errors introduced by the abstraction process. We present an initial empirical study of the robustness of several solution methods when using abstracted games.

Introduction

Games that model real world interactions are often complex, with huge numbers of possible strategies and information states. We are interested in better understanding the effect of abstraction in game-theoretic analysis. In particular, we focus on the *strategy selection problem*: how should an agent choose a strategy to play in a game, based on an abstracted game model? This problem has three interacting components: (1) the method for abstracting the game, (2) the method for selecting a strategy based on the abstraction, and (3) the method for mapping this strategy back to the original game. This approach has been studied extensively for poker, which is a 2-player, zero-sum game. However, much less is known about how abstraction interacts with strategy selection in more general games.

The main contributions of our work are as follows. First, we specify a model of the strategy selection problem when players use asymmetric abstractions as a *meta-game*. In this model players can use different methods for abstracting the game, solving the game, and reverse-mapping the solution. We introduce a collection of specific methods for abstracting and solving games intended to be representative of popular methods used in the literature. Finally, we present the results of extensive tournament simulations of abstraction and solution methods which show show differences in robustness.

Abstraction Meta-Games

We first introduce a formal model that can be used to study the situation where players select strategies based on abstracted game models. Our model is based on the meta-game framework introduced by Kiekintveld et al. (Kiekintveld and Wellman 2008), which focused on situations where players received noisy observations of the same underlying game and had to select strategies based on these observations. The situation where players use abstractions is similar in that players make strategy choices based on imperfect game abstractions. Players may use different abstractions which may cause problems for solutions that use coordination.

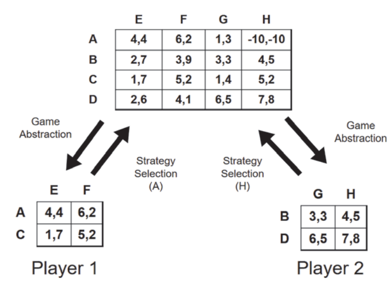

Figure 1: 2-players asymmetric abstractions

An example of an *abstraction meta-game* is shown in Figure 1. In this example, we have two players who are playing the one-shot normal form game shown at the top of the figure; this is the *base game*. They each perform their own (unspecified) abstraction to reduce the game. Then they both solve their abstracted games using their own solution methods. If both players use Nash equilibrium their actions when reversed mapped are A and H leading to *-10,-10*!

Abstraction and Solution Methods

We define an abstraction method as a function that maps one normal-form game into a second (smaller) normal-form game. We identify two broad categories of abstractions that are common in the literature: *strategy elimination* and *strategy combination*. Strategy elimination abstractions remove some strategies completely to reduce the size of the game. Strategy combination simplify games by merging multiple strategies into a single representative strategy. Our goal in this paper is not to develop novel abstraction methods nor to exhaustively evaluate the many existing techniques. The first abstraction method we consider is *TopN*, which is representative of strategy elimination. It creates a smaller game by selecting a subset of size N based on the overall payoff sum for each action to form the abstracted game. Reverse mapping strategies is trivial since a strategy in the abstracted game is also a strategy in the original game.

The second abstraction method we use is *KMeans*, which is representative of strategy combination. This method uses k-means clustering to group strategies into clusters based on the similarity of their payoffs in Euclidian distance. The abstracted game payoffs are computed by taking the average for all payoff outcomes in the cluster. The reverse mapping assumes that all actions in a cluster have equal probability.

We consider several candidate solution methods for selecting strategies in abstracted games. All are based on known solution concepts or simple heuristics for playing games, and are intended to provide a diverse pool of plausible strategies. We use baseline agents such as Uniform Random (UR) and common strategies like Nash Equilibrium (MSNE). It was also important to consider pessimistic strategies like MaxMin which could be considered as robust. We also include some sophisticated strategies like Quantal Response Equilibria (QRE) (McKelvey and Palfrey 1995) and Cognitive Hierarchies (CH) (Camerer, Ho, and Chong 2004) which originates in behavioral game theory.

Results

We conducted many experiments to identify the best parameter settings for QRE, CH, and QLK agents and then ran round robin tournaments on over 300 games, 18 unique agents, and 4 levels of abstractions (including no abstraction). The full data set is difficult to visualize, so we present selected subsets results to illustrate key points.

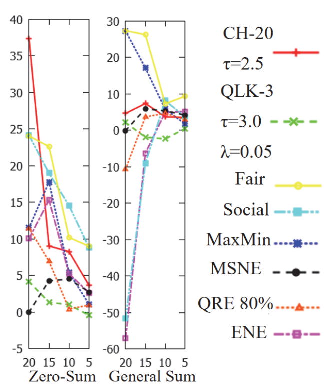

Figure 2: Stability of agents using KMeans

For example, Figure 2 shows the stability of different strategies which is the agent's incentive to change to a different solution concept (lower is better). The domain shows the level of abstraction which increases as it goes to the right starting from games with 20 actions to games with 3 actions. Notice that the Nash equilibrium strategy is incredibly stable for Zero Sum games but is mediocre for General Sum.

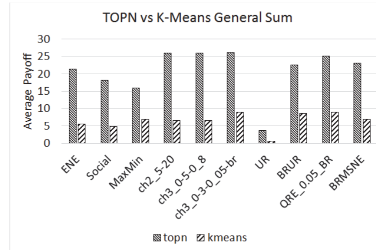

Figure 3: TopN vs Kmeans in General Sum

Figure 3 shows the average payoffs for the different solution methods in a tournament that combines agents using both abstractions. We use only the best parameter settings for the QRE, CH, and QLK agents. This tournament was run on general-sum games, using a high level of abstraction (from 20 actions down to 5 actions). The results are quite striking as TopN outperforms KMeans in *every* solution method.[1]

Conclusion

Our results demonstrate that using abstraction to solve games is a complex endeavor, and the type of abstraction, the solution methods used to analyze the abstracted games, and the class of games all have a strong influence on the results. Many of the strongest results using abstraction to analyze large games (e.g., Poker) have focused on zero-sum games. One of the most interesting observations from our results is that abstraction often works very differently in zero-sum games than it does general-sum games. In particular, solution methods based on finding Nash equilibrium seem to work much better in zero-sum games than they do in the other classes of games in our experiments. Another important observation is that Nash equilibrium often does not perform well in cases where abstraction is used as part of the solution process. It is still effective when the games are zero-sum, but in the other cases it was not robust to the introduction of error based on game abstraction.

Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant No. IIS-1253950.

References

Camerer, C. F.; Ho, T.-H.; and Chong, J.-K. 2004. A cognitive hierarchy model of games. *The Quarterly Journal of Economics* 119(3):861–898.

Kiekintveld, C., and Wellman, M. 2008. Selecting strategies using empirical game models: an experimental analysis of meta-strategies. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems, AAMAS '08*, volume 2, 1095–1101.

McKelvey, R., and Palfrey, T. 1995. Quantal response equilibria for normal form games. *Games and economic behavior* 10(1):6–38.

[1] http://www.cs.utep.edu/kiekintveld/students/veliz/index.html