Tiered Coalition Formation Game Stability and Simulation

Nathan Arnold
University of Kentucky
Lexington, KY 40506
nathan.arnold@uky.edu

Abstract

Expanding on a 2017 paper by Siler that introduced tiered coalition formation games, I have introduced a variant game and examined the stabilizability of both the original game and its variant. My thesis will contain further theoretical stability findings and the results and interpretation of a simulation based upon real data from video game matchups.

Introduction

In many broad conceptions of competition, it is assumed that competitors are transitive; that is, if X defeats Y and Y defeats Z, then X defeats Z. However, there are a number of scenarios in which transitivity does not appear to hold, such as in hiring problems studied by Tversky (Tversky 1969), the Condorcet paradox, and in pairwise comparison of playable characters in video games such as Pokémon and Super Smash Bros.

In order to advance our understanding of competition in which transitivity does not hold, Siler introduced tiered coalition formation games (Siler 2017). As with other coalition formation games, this setting simulates agents who act according to their own preferences, allows agents to form cooperative groups with each other, and aims to calculate stable outcomes, in which all represented agents are satisfied. An agent in tiered coalition formation games seeks a "tier list" that accurately represents that agent's power relative to other agents. Removing the assumption of transitivity adds considerable complexity to finding such a list, as Siler argued. Siler defined a preference framework and proved that a stable tier list can always be found under that preference framework; however, this stable list is singleton (one agent per tier), enforcing a strict transitive order on an intransitive structure. Siler concluded that this was insufficient for representing, analyzing, and understanding the environments that TCFGs are designed to replicate, and that further research was needed.

For this reason, my research has focused on expanding and advancing undersutanding of TCFGs. I developed a new preference framework that behaves similarly to the original proposed by Siler, but with improved accuracy to the underlying power structure. I also proposed a variant game which restricts the number of tiers, and introduced a new notion of stability for this game that promotes social welfare. Together with my coauthors, we presented these results and created a procedure that finds a stable tier list with high utility and high ranking in fairness criteria (Arnold, Goldsmith, and Snider 2022) (Arnold, Snider, and Goldsmith 2023). Further research will refine our understanding of the stability of TCFGs, and will use simulations with the game Pokémon to compare our results with a real-world tier list. In short, we wish to know which types of stable tier lists are computationally tractable to find, and which among those encourage high utility, high fairness, tiers consisting of multiple agents, and closeness to an externally found result. While our work focuses primarily on those intransitive preferences induced by pairwise competition, we hope it will contribute to the greater understanding and growing literature on intransitive preferences and their functioning.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

TCFG Stability

Most of our work so far, and a significant portion of my upcoming dissertation, is concerned with the stability of tier lists. Our goal was to extend the notion of a satisfactory tier list from that of Siler, in which tiers were singletons, and preferences were limited to deterministic outcomes. I developed a preference framework extending that of Siler into environments in which the outcome of a match is not deterministic, and a given match between two agents has a certain probability of being decided favorably for one agent or the other. I also theorized a variant game, k-tiered coalition formation games, in which the outcome must be a list of exactly k nonempty tiers. Together with Judy Goldsmith, we formalized a new stability concept, named socially conscious stability, such that k-TCFGs are always stabilizable, as well as an algorithm to find an SC stable tier list and notions of optimality. With my other coauthor, Sarah Snider, we implemented the theoretical algorithm and ran experiments to evaluate our results so far. These results appeared in FLAIRS 2022 (Arnold, Goldsmith, and Snider 2022).

A recently published journal article expands upon these results (Arnold, Snider, and Goldsmith 2023). In coalition formation games literature, the two most commonly used notions of stability are Nash stability and core stability, which Siler proved to be equivalent for TCFGs under the deterministic preference framework (Siler 2017). I proved that under the probabilistic preference framework, these two stability notions are equivalent to socially conscious stability for TCFGs, and that the three notions are non-equivalent on k-TCFGs. I also introduced a notion of approximate stability for TCFGs, and with some input from Snider wrote a program to generate approximately stable tier lists, comparing them against each other and the stable tier lists on the same instances. The article includes experimental results expanding on those from FLAIRS, comparing the performance of certain methods of finding stable tier lists on the same instances in terms of utility, fairness, and tier size.

Future Work

In working on my thesis, I hope to develop three more major results on the topic of the stability of TCFGs.

First, it is known that not all k-TCFGs with matchup preferences admit a Nash stable outcome (Arnold, Goldsmith, and Snider 2022). Following from this, the complexity of deciding if a Nash stable outcome is possible in a given instance is unknown. I conjecture that this problem is NP-complete. In order to prove this, I will need to find a reduction from a problem known to be NP-complete. The search for such a reduction is underway.

On the other hand, it remains an open question whether all k-TCFGs admit a core stable outcome (Arnold, Snider, and Goldsmith 2023). I conjecture that the answer is yes. Currently, there is no known polynomially tractable method of verifying that an outcome of a k-TCFG is core stable, which poses a challenge to discovering patterns in the core stable outcomes of different instances. The high-utility and high-fairness socially consciously stable tier list produced by our algorithm (Arnold, Goldsmith, and Snider 2022) is also core stable in a majority of instances, though not all. I am attempting to prove a method exists that will always lead to a core stable tier list in a k-TCFG.

Lastly, Siler proved that under deterministic preferences, a stable tier list of singletons can be found in polynomial time on any instance (Siler 2017). Finding this stable tier list of singletons has become an important preliminary step in our own procedures (Arnold, Goldsmith, and Snider 2022) (Arnold, Snider, and Goldsmith 2023). However, Siler's proof of polynomial time complexity does not hold under probabilistic preferences. I seek a method that will sort agents into a stable singleton tier list in polynomial time under the expanded preference framework.

Because these three questions are theoretical in nature, it is difficult to say with certainty how much progress will be made before February. It is extremely doubtful that all three would be answered by that time.

Simulating Pokemon Battles

Siler's original formulation of TCFGs was based on real-world tier lists, especially those used to sort characters from the Pokémon series of video games (Siler 2017). Namely, a website known as Smogon manages tier lists of Pokémon, and the design and philosophy of Smogon inspired developments on the topic from both Siler (Siler 2017) and myself (Arnold, Goldsmith, and Snider 2022). It felt appropriate, therefore, for the first tests of TCFGs derived from real-world data to be based on simulations of Pokémon.

With major programming contributions from colleagues Sean Roche and BJ Resultay, we have developed a program that accurately simulates all game states and actions from "Pokémon Red and Blue Versions." Using this program, we can further develop a training program so that each agent learns its optimal strategy against each other agent. When pairwise optimal policies are learned, we will execute many repeated simulations for each pair of agents, and use the results to estimate the probabilities of each agent winning in a one-on-one match against each other agent. These data can form the basis for stable tier lists and approximately stable tier lists constructed using the methods we have developed in our prior theoretical work. It is anticipated that this will be complete or nearing completion by February.

Using these data, we wish to compare our results with the real-world tier lists managed by Smogon. I will be making observations on how unstable the Smogon tier list is under our understanding of TCFGs thus far. It is possible, perhaps likely, that patterns will emerge in any inconsistencies between our data and Smogon's tier list. I anticipate that such inconsistencies will be easily related to findings on evolving metagames (such as in team counter-selection games (Crane et al. 2021)), in which the preferability of agents is affected by roles that can only be played by a subset of agents, or by the existence of a dominating agent. The metagame of Pokémon Red and Blue Versions is known to have both factors. Two agents (Golem and Rhydon) are considered more preferable than they would be otherwise because of the role they are able to play in a match. Two other agents (Tauros and Alakazam) are considered dominating agents, and other agents' preferability is altered by those agents' ability to counteract the dominating agents. There will be a plentiful opportunity to develop our understanding of preferences and stability in TCFGs based on these factors, which I expect to be working on at the time of the consortium. There is growing research interest in the concept of evolving metagames (as an extension of evolving preferences), and we are hopeful that our examination of this environment leads to answered questions (and new questions) for other researchers in the field.

References

Arnold, N.; Goldsmith, J.; and Snider, S. 2022. Extensions to Tiered Coalition Formation Games. *The International FLAIRS Conference Proceedings*, 35.

Arnold, N.; Snider, S.; and Goldsmith, J. 2023. Socially conscious stability for tiered coalition formation games. *Annals of Mathematics and Artificial Intelligence*, 1–42.

Crane, D.; Holmes, Z.; Kosiara, T. T.; Nickels, M.; and Spradling, M. 2021. Team counter-selection games. In *2021 IEEE Conference on Games (CoG)*, 1–8. IEEE.

Siler, C. 2017. Tiered Coalition Formation Games. *The International FLAIRS Conference Proceedings*, 30: 210–214.

Tversky, A. 1969. Intransitivity of preferences. *Psychological review*, 76(1): 31–48.