

Forming Probably Stable Communities with Limited Interactions

Ayumi Igarashi

Kyushu University
Fukuoka, Japan
igarashi@agent.inf.kyushu-u.ac.jp

Jakub Sliwinski

ETH Zurich
Zurich, Switzerland
jsliwinski@ethz.ch

Yair Zick

National University of Singapore
Singapore
zick@comp.nus.edu.sg

Abstract

A community needs to be partitioned into disjoint groups; each community member has an underlying preference over the groups that they would want to be a member of. We are interested in finding a stable community structure: one where no subset of members S wants to deviate from the current structure. We model this setting as a hedonic game, where players are connected by an underlying interaction network, and can only consider joining groups that are connected subgraphs of the underlying graph. We analyze the relation between network structure, and one’s capability to infer statistically stable (also known as PAC stable) player partitions from data. We show that when the interaction network is a forest, one can efficiently infer PAC stable coalition structures. Furthermore, when the underlying interaction graph is not a forest, efficient PAC stabilizability is no longer achievable. Thus, our results completely characterize when one can leverage the underlying graph structure in order to compute PAC stable outcomes for hedonic games. Finally, given an unknown underlying interaction network, we show that it is NP-hard to decide whether there exists a forest consistent with data samples from the network.

1 Introduction

A professor wants her students to complete a group programming project. In order to do so, students should divide into project groups with a few students in each; naturally, some groups will be objectively better than others. However, students seldom try to find a group that’s objectively optimal for them; they would rather join groups that have at least one or two of their friends. This type of scenario falls into the realm of *constrained coalition formation*; in other words, how should we partition a group of people given that (a) they have preferences over the groups they are assigned to and (b) they have limited interactions with one another? Other scenarios fitting this description include

- (a) Seating arrangements at a wedding (or at conference banquets): some guests should absolutely not be seated together, while others would probably enjoy one another’s company. However, it should always be the case that every guest has at least one acquaintance seated at their table.

- (b) Group formation on social media: given a social media network (e.g. Facebook), people prefer being affiliated with certain groups; however, they are limited to joining groups that already contain their friends.

Constrained coalition formation problems are often modeled as *hedonic games*. Hedonic games formally capture a simple, yet compelling, paradigm: how does one partition players into groups, while factoring individual players’ preferences? The literature on hedonic games is primarily focused on finding “good” *coalition structures* — partitions of players into disjoint groups. A set of coalition structures satisfying certain desiderata is called a *solution concept*. A central hedonic solution concept is *coalitional stability*: given a coalition structure π , we say that a set of players (also known as a *coalition*) S can deviate from π if every $i \in S$ prefers S to its assigned group under π ; a coalition structure π is stable if no coalition S can deviate. In other words, S contains at least one player i who prefers its current coalition (denoted $\pi(i)$) to S . The set of stable coalition structures — also known as the *core* of the hedonic game — may be empty; what’s worse, even when it is known to be non-empty, finding a stable coalition structure may be computationally intractable. Moreover, efficient algorithms for finding stable coalition structures often assume full knowledge of the underlying hedonic game; that is, in order to work, the algorithm needs to have either *oracle access to player preferences* (i.e. queries of the form ‘does player i prefer coalition S to coalition T ?’), or *structural knowledge of the underlying preference structure* (e.g. some concise representation of player preferences that one can leverage in order to obtain a poly-time algorithm).

Neither assumption is realistic in practice: eliciting user preferences is notoriously difficult, especially over combinatorially complex domains such as subsets of players. If one forgoes preference elicitation and opts for mathematically modeling preferences (e.g. assuming that users have additive preferences over coalition members), it is not entirely obvious what mathematical model of user preferences is valid. This leads us to the following natural question: can we find a stable coalition structure when player preferences are unknown? Recent works (Balcan, Procaccia, and Zick 2015; Balkanski, Syed, and Vassilvitskii 2017; Sliwinski and Zick 2017) propose a statistical approach to stability in collaborative environments. In this framework, one assumes the existence of user preference data over some coalitions, which

is then used to construct *probably approximately stable* outcomes (the notion is referred to as PAC stability). In this paper, we explore the relation between structural assumptions on player preferences, and computability of PAC stable outcomes.

Our contribution We assume that there exists some underlying *interaction network* governing player preferences; that is, players are nodes on a graph, and only connected coalitions are feasible. Within this framework, we show that if player preferences are restricted by a forest, one can compute a PAC stable outcome using only a polynomial number of samples. Surprisingly, even if the underlying forest structure is not known to the learner, PAC stabilizability still holds, despite the fact that it may be computationally intractable to find an approximate forest structure that is likely consistent with the true interaction graph. In contrast, we show that it is impossible to find a PAC stable outcome even if the graph contains a single cycle. The latter result is constructive: we show that whenever the underlying interaction graph does contain a cycle, one can construct a sample distribution for which it would be impossible to elicit a PAC stable outcome.

Our positive result for forests is interesting in several respects. First, while one can find PAC stable outcomes in polynomial time, computing stable outcomes for hedonic games on forests is computationally intractable (Igarashi and Elkind 2016); second, unlike (Sliwinski and Zick 2017), we do not require that player preferences are provided in the form of numerical utilities over coalitions. This not only makes our results more general, but also more faithful to the problem we model, which assumes ordinal information about player preferences, rather than cardinal utilities. Finally, in Section 5, we prove a non-trivial technical result on learning forest structures that is of independent interest. Briefly, we study the following problem: we are given samples of subsets of graph vertices, each labeled either ‘connected’ or ‘disconnected’; we need to decide whether there exists some forest T^* that is consistent with the sample — i.e. all connected sets of vertices are connected under T^* and all disconnected sets are not. We show that when all of our vertex samples are connected (i.e. we do not observe any disconnected components), it is possible to efficiently learn an underlying forest structure (if one exists); on the other hand, if one assumes that both connected and disconnected sets are presented to the learner, it is computationally intractable to decide whether there exists a forest, or even a path, that is consistent with the samples.

Related work There exists a rich body of literature studying hedonic games from an economic perspective (e.g. (Banerjee, Konishi, and Sönmez 2001; Bogomolnaia and Jackson 2002)). More recently, the AI community has begun studying both computational and analytical properties of hedonic games (see e.g. (Aziz and Brandl 2012; Deineko and Woeginger 2013; Gairing and Savani 2010; Peters and Elkind 2015), and (Aziz and Savani 2016; Woeginger 2013) for an overview). Interaction networks in cooperative games were first introduced by Myerson (1977). The relation between graph structure and stability in the classic cooperative game setting is also relatively well-understood. Demange (2004) shows that if the un-

derlying interaction network is a forest, then the core is not empty; further studies (Bousquet, Li, and Vetta 2015; Meir et al. 2013) establish relations between approximate stability and the underlying graph structure, while Chalkiadakis, Greco, and Markakis (2016) study the computational complexity of finding core outcomes in graph restricted environments. Igarashi and Elkind (2016) establish both the existence of stable coalition structures in hedonic games over forests, as well as the computational intractability of finding stable coalition structures; Peters (2016) studies the relation between hedonic solution concepts and the treewidth of the underlying interaction graph.

Several works study learning based game-theoretic solution concepts. Sliwinski and Zick (2017) introduce PAC stability in hedonic games, and analyze several common classes of hedonic games. Other works on learning and game theory include learning in cooperative games (Balcan, Procaccia, and Zick 2015; Balkanski, Syed, and Vassilvitskii 2017), rankings (Balcan, Vitercik, and White 2016), auctions (Balcan et al. 2012; Balcan, Sandholm, and Vitercik 2018; Morgenstern and Roughgarden 2016) and noncooperative games (Fearnley et al. 2013; Sinha, Kar, and Tambe 2016).

Full version. An extended version with full proofs is available on arXiv (Igarashi, Sliwinski, and Zick 2018).

2 Preliminaries

Throughout this paper, vectors are denoted by \vec{x} , and sets are denoted by uppercase letters; given a value $s \in \mathbb{N}$, we set $[s] = \{1, \dots, s\}$. A *hedonic game* is given by a pair $\langle N, \succeq \rangle$, where $N = [n]$ is a finite set of *players*, and $\succeq = (\succeq_1, \dots, \succeq_n)$ is a list of preferences players in N have over subsets of N (also referred to as *coalitions*); in more detail, for every $i \in N$, we write $\mathcal{N}_i = \{S \subseteq N \mid i \in S\}$; \succeq_i describes a complete and transitive preference relation over \mathcal{N}_i . For each $i \in N$, let \succ_i denote the strict preference derived from \succeq_i , i.e., $S \succ_i S'$ if $S \succeq_i S'$, but $S' \not\succeq_i S$. An outcome of a hedonic game is a *coalition structure*, i.e., a partition π of N into disjoint coalitions; we denote by $\pi(i)$ the coalition containing $i \in N$. A solution concept is a mapping whose input is a hedonic game $\langle N, \succeq \rangle$, and whose output is a (possibly empty) set of coalition structures. The *core* is the most fundamental solution concept in hedonic games. First, we say that a coalition S *strongly blocks* a coalition structure π if every player $i \in S$ strictly prefers S to its current coalition $\pi(i)$, i.e. $S \succ_i \pi(i)$. A coalition structure π is said to be *core stable* if no coalition $S \subseteq N$ strongly blocks π .

2.1 Interaction Networks

Given an undirected graph $G = \langle N, E \rangle$ whose nodes are the player set, we restrict the space of *feasible coalitions* to be the set of connected subsets of G ; we denote by \mathcal{F}_E the set of feasible coalitions. Intuitively, we restrict our attention to coalition structures where all group members form a social subnetwork of the underlying interaction graph. Note that when G is a clique, all coalitions are feasible, and the result is a standard (unrestricted) hedonic game. From now on, we define a *hedonic graph game* as the tuple $\langle N, \succeq, E \rangle$; here, N

is the set of players, \succeq their preference relations, and E the edges of the underlying interaction network. We focus our attention only on core stable coalition structures that consist of feasible coalitions.

In what follows, it is useful to express player preferences in terms of cardinal utilities. In other words, player i assigns a value $v_i(S) \in \mathbb{R}$ to every coalition $S \in \mathcal{N}_i$; we write a hedonic game as $\langle N, \mathcal{V} \rangle$ where \mathcal{V} is a collection of functions $v_i : \mathcal{N}_i \rightarrow \mathbb{R}$ for each $i \in N$. This representation allows us to seamlessly integrate ideas from PAC learning into the hedonic games model, and is indeed quite common in other works studying hedonic games. However, as we later show, our main result (Theorem 4.1) still holds when we transition from a utility-based cardinal model, to a preference-based ordinal model.

2.2 PAC Learning

We provide a brief introduction to PAC learning¹. The basic idea is as follows: we are given an unknown function $v : 2^N \rightarrow \mathbb{R}$ (a *target concept* in the language of PAC learning) that assigns values to subsets of players. In addition, we are given a set of m samples $((S_1, v(S_1)), \dots, (S_m, v(S_m)))$ where $S_j \subseteq N$ and $v(S_j)$ is the valuation of v over S_j ; we wish to estimate v on subsets we did not observe. We assume that v belongs to a *hypothesis class* \mathcal{H} (say, we know that v is an additive valuation). Our goal is to output a *hypothesis* $v^* \in \mathcal{H}$ (e.g. if v is additive, v^* should be as well) that is likely to match the outputs of v on future observations drawn from some distribution \mathcal{D} . More formally, a hypothesis v^* is ϵ *approximately correct* w.r.t a probability distribution \mathcal{D} over 2^N and an unknown function v if

$$\Pr_{S \sim \mathcal{D}} [v^*(S) \neq v(S)] < \epsilon.$$

A learning algorithm \mathcal{A} takes as input m samples

$$(S_1, v(S_1)), (S_2, v(S_2)), \dots, (S_m, v(S_m))$$

drawn i.i.d. from a distribution \mathcal{D} over 2^N , and two parameters $\epsilon, \delta > 0$.

A class of functions \mathcal{H} is (ϵ, δ) *PAC (probably approximately correctly) learnable* if there exists an algorithm \mathcal{A} that for any $v \in \mathcal{H}$ and probability distribution \mathcal{D} over 2^N , with probability of at least $1 - \delta$, it outputs a hypothesis v^* that is ϵ approximately correct with respect to \mathcal{D} and v . If this holds for any $\epsilon, \delta > 0$, \mathcal{H} is said to be *PAC learnable*; moreover, if the running time of \mathcal{A} , and the number of samples m are polynomial in $\frac{1}{\epsilon}$, $\log \frac{1}{\delta}$ and n , \mathcal{H} is said to be *efficiently PAC learnable*.

The value δ is the *confidence parameter*: intuitively, it is the probability that the random samples drawn from \mathcal{D} do not accurately portray the true sample distribution; for example, if \mathcal{D} is the uniform distribution, then it is possible (though unlikely) that we draw the same subset in every one of our m samples. The value ϵ is called the *error parameter*: it is the likelihood that our hypothesis v^* does not agree

¹What we show here is but one of many variants on the theory of PAC learning. There are many excellent sources on this classic theory; we refer our reader to (Anthony and Bartlett 1999; Kearns and Vazirani 1994; Shashua 2009)

with the target concept v . Not all hypothesis classes are efficiently PAC learnable; learnability is inherently related to the complexity of the hypothesis class. The complexity of real-valued functions is commonly measured using the notion of *pseudo dimension* (see e.g. Chapter 11 of (Anthony and Bartlett 1999)). Given a list of sets $S_1, \dots, S_m \subseteq N$, and corresponding values $r_1, \dots, r_m \in \mathbb{R}$ we say that a class of functions \mathcal{H} can *pseudo-shatter* $(S_j, r_j)_{j=1}^m$ if for any labeling $\ell_1, \dots, \ell_m \in \{0, 1\}$, there is some $v \in \mathcal{H}$ such that $v(S_j) \geq r_j$ iff $\ell_j = 1$. The *pseudo-dimension* of \mathcal{H} , denoted $P_{dim}(\mathcal{H})$ is

$$\max\{m \mid \exists (S_j, r_j)_{j=1}^m \text{ that can be shattered by } \mathcal{H}\}.$$

The following well-known theorem relates the pseudo-dimension and PAC learnability.

Theorem 2.1 ((Anthony and Bartlett 1999)). *A class of functions \mathcal{H} is efficiently (ϵ, δ) PAC learnable using $m = \text{poly}(P_{dim}(\mathcal{H}), \frac{1}{\epsilon}, \log \frac{1}{\delta})$ samples if there exists an algorithm such that given m samples $(S_1, v(S_1)), (S_2, v(S_2)), \dots, (S_m, v(S_m))$ drawn i.i.d. from a distribution \mathcal{D} , it outputs $v^* \in \mathcal{H}$ consistent with the sample, i.e. $v^*(S_j) = v(S_j)$ for all sampled S_j , and runs in time polynomial in $\frac{1}{\epsilon}$, $\log \frac{1}{\delta}$ and m . Furthermore, if $P_{dim}(\mathcal{H})$ is superpolynomial in n , \mathcal{H} is not PAC learnable.*

In other words, in order to establish the PAC learnability of some hypothesis class, it suffices that one shows that its pseudo dimension is low, and that there exists some efficient algorithm that is able to output a hypothesis v^* which matches the outputs of v on all samples. We note that even if an efficient consistent algorithm does not exist (e.g. if the problem of matching a hypothesis to the samples is computationally intractable), a low pseudo dimension is still desirable: it implies that the number of samples needed in order to find a good hypothesis is polynomial.

2.3 PAC Stabilizability

When studying hedonic games, one is not necessarily interested in eliciting approximately accurate user preferences over coalitions using data; in our case, we are interested in identifying core stable coalition structures. Intuitively, it seems that the following idea might work: first, infer player utilities from data and obtain a PAC approximation of the original hedonic game; next, find a coalition structure that stabilizes the approximate hedonic game. This approach, however, may be overcomplicated: first, it may be impossible to PAC learn player preferences from data (this depends on the hypothesis class); moreover, computing a core coalition structure for the learned game may be computationally intractable. (Sliwinski and Zick 2017) propose learning a stable outcome directly from data. They introduce a statistical notion of core stability for hedonic games, which they term *PAC stability* (this term was first used by (Balcan, Procaccia, and Zick 2015) for cooperative transferable utility games).

We say that a partition π is ϵ -PAC stable w.r.t a probability distribution \mathcal{D} over 2^N if

$$\Pr_{S \sim \mathcal{D}} [S \text{ strongly blocks } \pi] < \epsilon.$$

The inputs to our learning algorithms will be samples

$$(S_1, \vec{v}(S_1)), (S_2, \vec{v}(S_2)), \dots, (S_m, \vec{v}(S_m)),$$

where $S_1, \dots, S_m \subseteq N$, and $\vec{v}(S_j)$ is a vector describing players' utilities over S_j ; that is, $\vec{v}(S_j) = (v_i(S_j))_{i \in S_j}$.

Given an unknown hedonic game $\langle N, \mathcal{V} \rangle$ belonging to some hypothesis class \mathcal{H} , a *PAC stabilizing algorithm* \mathcal{A} takes as input m sets S_1, \dots, S_m sampled i.i.d. from a distribution \mathcal{D} , and players' preferences over the sampled sets; in addition, it receives two parameters $\epsilon, \delta > 0$. The algorithm \mathcal{A} *PAC stabilizes* \mathcal{H} , if for any hedonic game $\langle N, \mathcal{V} \rangle \in \mathcal{H}$, distribution \mathcal{D} over 2^N , and parameters $\epsilon, \delta > 0$, with probability $\geq 1 - \delta$, \mathcal{A} outputs an ϵ -PAC stable coalition structure if it exists; again, if the running time of the algorithm \mathcal{A} and the number of samples, m , are bounded by a polynomial in $n, \frac{1}{\epsilon}$ and $\log \frac{1}{\delta}$, then we say that \mathcal{A} *efficiently PAC stabilizes* \mathcal{H} . Similarly, we say that \mathcal{H} is (efficiently) *PAC stabilizable* if there is some algorithm \mathcal{A} that (efficiently) PAC stabilizes \mathcal{H} .

3 Learning Hedonic Graph Games

In what follows we consider the following hypothesis class.

Definition 3.1. For an undirected graph $G = \langle N, E \rangle$, let \mathcal{H}_G be the class of all hedonic games $\langle N, \mathcal{V} \rangle$ where for each player $i \in N$, $v_i(\{i\}) = 0$ and player i strictly prefers its singleton to any disconnected coalition $S \in \mathcal{N}_i \setminus \mathcal{F}_E$, i.e., $v_i(S) < 0$ for all $S \in \mathcal{N}_i \setminus \mathcal{F}_E$.

We first present a baseline negative result: fixing a forest G , the hypothesis class, \mathcal{H}_G is not efficiently PAC learnable. When referring to the PAC learnability of any class of hedonic games, we mean inferring some utility function $v_i^* : 2^N \rightarrow \mathbb{R}$ for all $i \in N$ that PAC approximates the true utilities of players in N . This approximation guarantee can be interpreted in both an ordinal and cardinal manner. If we are given player i 's ordinal preferences, this simply means that v_i^* is consistent with the ordinal preferences; if we are given player i 's cardinal utility function v_i , v_i^* should be a PAC approximation of v_i . As Theorem 3.2 shows, even when we are given additional information about the underlying graph interaction network, players' preferences are not PAC learnable.

Theorem 3.2. *For any graph $G = \langle N, E \rangle$ with exponentially many connected coalitions, the class \mathcal{H}_G is not efficiently PAC learnable.*

Proof. Recall that \mathcal{F}_E is the set of all feasible coalitions over $G = \langle N, E \rangle$; by assumption, $|\mathcal{F}_E|$ is exponential. Let \mathcal{H}_i be the set of all possible utility functions $v_i : \mathcal{N}_i \rightarrow \mathbb{R}$ satisfying $v_i(\{i\}) = 0$ and $v(S) < 0$ for all disconnected coalition $S \in \mathcal{N}_i \setminus \mathcal{F}_E$. The utility player i derives from feasible coalitions in G is unrestricted; in particular, one cannot deduce anything about the utility of some feasible coalition $S \in \mathcal{F}_E$, based on other feasible coalitions' utilities. This immediately implies that the set \mathcal{F}_E can be pseudo-shattered by \mathcal{H}_i . Hence $P_{dim}(\mathcal{H}_i)$ is at least exponential, and by Theorem 2.1, \mathcal{H}_i is not efficiently PAC learnable. \square

As an immediate corollary, forest interaction structures do not admit PAC learnable preference structures in general; this is true even if G is a star graph over n players, since the number of feasible coalitions is exponential in n .

Corollary 3.3. *Let G be a star graph over n players; then \mathcal{H}_G is not PAC learnable.*

Proof. For a star with n nodes, any coalition containing the center of the star is feasible, hence it has 2^{n-1} feasible coalitions. By Theorem 3.2, hedonic games on forests are not PAC learnable. \square

The reason that hedonic games with forest interaction structures are not PAC learnable is that they may have exponentially many feasible coalitions; this is also the reason that finding a core stable coalition structure for hedonic games with forest interaction structures is computationally intractable (Igarashi and Elkind 2016). However, we now show how one can still exploit the structural properties of forest graph structures to efficiently compute PAC stable outcomes.

4 PAC Stabilizability of Hedonic Graph Games

Having established that hedonic games with a forest interaction structure are not, generally speaking, PAC learnable, we turn our attention to their PAC stabilizability. We divide our analysis into two parts. We begin by assuming that the underlying interaction graph structure G is known to us; in other words, we know that our game belongs to the hypothesis class \mathcal{H}_G . In Section 5, we show how one can forgo this assumption.

Theorem 4.1. *If $G = \langle N, E \rangle$ is a forest, \mathcal{H}_G is efficiently PAC stabilizable.*

Proof. We claim that Algorithm 1 PAC stabilizes \mathcal{H}_G . It is related to the algorithm introduced in Demange (2004) used to find core stable outcomes for forest-restricted hedonic games² in the full information setting. Intuitively, instead of identifying the *guaranteed coalition* for each player precisely, Algorithm 1 approximates it. If the input graph $\langle N, E \rangle$ is a forest, we can process each of its connected components separately, so we can assume that $\langle N, E \rangle$ is a tree.

We first provide an informal description of our algorithm, followed by pseudocode. The algorithm first transforms $\langle N, E \rangle$ into a rooted tree with root r by orienting the edges in E towards the leaves. For every player i starting from the bottom to the top, the algorithm identifies B_i - a coalition containing i , the best for i observed in the samples that is entirely contained in i 's subtree, such that others in B_i prefer it to their own best guaranteed coalition; in other words, $B_i \succeq_j B_j$ for all $j \in B_i$. Having identified B_j for every $i \in N$, players are partitioned according to the B_i 's from top-down. The main concern is to ensure that B_i is a good approximation of its

²Demange (2004) presents the algorithm for non-transferable cooperative utility games on trees where each coalition has a choice of action. A hedonic game is a special case of a non-transferable utility game where each coalition has a unique action.

full-information counterpart; this is guaranteed by taking a sufficiently large sample size $m = \lceil \frac{n}{\epsilon} \log \frac{n}{\delta} \rceil$.

In what follows, we assume an orientation of the trees in G , with arbitrary root nodes. Fixing the orientation, we let $\text{desc}(i)$ be the set of descendants of i (we assume that $i \in \text{desc}(i)$). For each coalition $S \subseteq N$, we denote by $\text{child}(S)$ the set of *children* of S , namely,

$$\text{child}(S) = \{i \in N \setminus S \mid i\text{'s parent belongs to } S\}.$$

The *height* of a node $i \in N$ is defined inductively as follows: $\text{height}(i) := 0$ if i is a leaf, i.e., $\text{desc}(i) = \{i\}$, and

$$\text{height}(i) := 1 + \max\{\text{height}(j) \mid j \in \text{desc}(i) \setminus \{i\}\},$$

otherwise.

Algorithm 1 An algorithm finding a PAC stable outcome for forest-restricted games

Input: set \mathcal{S} of $m = \lceil \frac{n}{\epsilon} \log \frac{n}{\delta} \rceil$ samples from \mathcal{D}

- 1: Make a rooted tree with root r by orienting all the edges in E towards the leaves.
 - 2: Initialize $B_i \leftarrow \emptyset$ and $\pi^{(i)} \leftarrow \emptyset$ for each $i \in N$.
 - 3: **for** $t = 0, \dots, \text{height}(r)$ **do**
 - 4: **for** each node $i \in N$ with $\text{height}(i) = t$ **do**
 - 5: set $S^* = \{S \in \mathcal{S} \cap \mathcal{F}_E \mid i \in S \subseteq \text{desc}(i) \wedge v_j(S) \geq v_j(B_j), \text{ for all } j \in S \setminus \{i\} \cup \{i\}\}$
 - 6: choose $B_i \in \text{argmax}\{v_i(S) \mid S \in S^*\}$
 - 7: set $\pi^{(i)} \leftarrow \{B_i\} \cup \bigcup\{\pi^{(j)} \mid j \in \text{child}(B_i)\}$
 - 8: **end for**
 - 9: **end for**
-

Given player i , let \mathcal{B}_i be the collection of coalitions B_j for every descendant $j \neq i$ of i , i.e., $\mathcal{B}_i = \{B_j \mid j \in \text{desc}(i) \setminus \{i\}\}$. For each $i \in N$ and each coalition $X \subseteq N$, we let $P_{\mathcal{B}_i}(X)$ mean that $i \in X \subseteq \text{desc}(i)$, X is connected, and every other player j in $X \setminus \{i\}$ weakly prefers X to B_j . Now, we define a modified preference order for player i , $\succeq_{\mathcal{B}_i}$, that devalues any coalition X for which $P_{\mathcal{B}_i}(X)$ does not hold.

- If $P_{\mathcal{B}_i}(X)$ and $P_{\mathcal{B}_i}(Y)$, then $X \succ_{\mathcal{B}_i} Y \iff X \succ_i Y$
- If $P_{\mathcal{B}_i}(X)$ but $\neg P_{\mathcal{B}_i}(Y)$, then $X \succ_{\mathcal{B}_i} Y$
- If $\neg P_{\mathcal{B}_i}(X)$, then $\forall Y : Y \succeq_{\mathcal{B}_i} X$

Given \mathcal{B}_i and a distribution \mathcal{D} , we say that a coalition X is *top- $\frac{\epsilon}{n}$* for player i , if

$$\Pr_{S \sim \mathcal{D}} [S \succ_{\mathcal{B}_i} X] \leq \frac{\epsilon}{n}.$$

Trivially, for every \mathcal{B}_i the probability of sampling a top- $\frac{\epsilon}{n}$ coalition for player i from \mathcal{D} is at least $\frac{\epsilon}{n}$; moreover, if $\Pr_{S \sim \mathcal{D}} [P_{\mathcal{B}_i}(S)] \leq \frac{\epsilon}{n}$, then any coalition is top- $\frac{\epsilon}{n}$.

Intuitively, \mathcal{B}_i approximates the best coalition i can form with members of the subtree rooted at i . Algorithm 1's objective is to ensure that sampling a coalition S from \mathcal{D} such that $P_{\mathcal{B}_i}(S) \wedge S \succ_i B_i$ is unlikely, namely, the probability of seeing S from \mathcal{D} such that S is better for the highest node i in S than B_i , and every other player in S prefers it to their B_j ,

is smaller than ϵ ; this is done by examining enough coalitions so as to see some top- $\frac{\epsilon}{n}$ coalition for every player.

Examine what happens if \mathcal{B}_i containing B_j 's for i 's descendants is fixed upfront, i.e. not dependent on the sample. Let us bound the probability that for i , none of the coalitions in \mathcal{S} are top- $\frac{\epsilon}{n}$:

$$\begin{aligned} \left(1 - \frac{\epsilon}{n}\right)^m &= \left(1 - \frac{\epsilon}{n}\right)^{\lceil \frac{n}{\epsilon} \log \frac{n}{\delta} \rceil} \\ &\leq \left(\left(1 - \frac{\epsilon}{n}\right)^{\frac{n}{\epsilon}}\right)^{\log \frac{n}{\delta}} < \left(\frac{1}{e}\right)^{\log \frac{n}{\delta}} < \frac{\delta}{n} \end{aligned} \quad (1)$$

Note that Inequality (1) is true irrespective of what \mathcal{B}_i is. Taking a union bound, the probability that there is some player i such that there is no top- $\frac{\epsilon}{n}$ coalition for i in \mathcal{S} is at most δ . Note that S^* can end up not containing any coalition (line 5). But then with high confidence, as a special case of the above consideration, every coalition is top- $\frac{\epsilon}{n}$, and the algorithm can pick $\{i\}$.

Recall that in an actual run of the algorithm the sample \mathcal{S} is drawn, and for every descendant j of i , B_j is computed based on \mathcal{S} , and then B_i is computed based on the same sample. One can ask whether some dependence between the computation of B_i and the B_j 's does not invalidate Inequality (1). This potential problem can be easily solved by taking a larger number of samples: if we take $m = \lceil \frac{n^2}{\epsilon} \log \frac{n}{\delta} \rceil$ samples, we can just use $\frac{n}{\epsilon} \log \frac{n}{\delta}$ samples to compute each B_i and maintain complete independence in the samples.

In order to see the smaller sample size used in Algorithm 1 provides the same guarantee, consider an equivalent reordering of the computation of B_i and B_j 's: first, for every $i \in N$, determine the number k of connected coalitions S in the sample such that i will be the highest node in S . Then, draw the other $m - k$ coalitions and compute B_j 's for every descendant j of i ; finally, based on this, determine the family \mathcal{B}_i . Note that regardless of what \mathcal{B}_i is, each of the undetermined, independently drawn k coalitions has probability of at least $\frac{\epsilon}{n}$ to be top- $\frac{\epsilon}{n}$ for i . Hence, the inequality (1) holds even if \mathcal{B}_i and B_i are computed based on the same sample of coalitions \mathcal{S} .

We are now ready to prove that the coalition structure outputted by Algorithm 1 returns a PAC stable outcome $\pi^{(r)}$. We observe that any coalition included in the returned $\pi^{(r)}$ is a B_i for some i . Note that for every $j \in B_i$, we have that $v_j(B_i) \geq v_j(B_j)$ (line 5). Now, consider any coalition X that strongly blocks $\pi^{(r)}$; let $i = \text{argmax}_{j \in X} \text{height}(j)$. Since X strongly blocks $\pi^{(r)}$,

$$v_j(X) > v_j(\pi^{(r)}(j)) \geq v_j(B_j)$$

for all players $j \in X$. In particular, $v_i(X) > v_i(B_i)$. By construction of B_i and Inequality (1), B_i is top- $\frac{\epsilon}{n}$ for i ; that is,

$$\frac{\epsilon}{n} > \Pr_{S \sim \mathcal{D}} [v_i(S) > v_i(B_i)] \geq \Pr_{S \sim \mathcal{D}} [S = X];$$

thus the probability of drawing a coalition such as X from \mathcal{D} , i.e. strongly blocking $\pi^{(r)}$ and having $i = \text{argmax}_{j \in X} \text{height}(j)$, is less than $\frac{\epsilon}{n}$. Taking a union bound over all players,

$$\Pr_{X \sim \mathcal{D}} [X \text{ strongly blocks } \pi^{(r)}] < \epsilon;$$

this guarantee holds with confidence $1 - \delta$. \square

We conjecture that a similar argument can imply a stronger statement. That is, we can replace ‘strongly block’ in the definition of PAC stabilizability with ‘weakly block’ and still obtain PAC stabilizability on trees. (A coalition S weakly blocks a coalition structure if every player weakly prefers S to their current coalition and at least one player in S has a strict preference) We note that in the full information setting, a strict core outcome does not necessarily exist on trees (Igarashi and Elkind 2016).

Remark 4.2 (From Cardinal to Ordinal Preferences). Note that step 6 of the Algorithm 1 is the *only* step that refers to the numerical representation of agent preferences v_i . The algorithm chooses a coalition with maximal utility value v_i out of some set of possible coalitions; in particular, the only thing required for the successful implementation of Algorithm 1 is players’ ranking of coalitions in the sample. In other words, the particular numerical representation of player preferences plays no role. This is a significant departure from the algorithms devised by (Sliwinski and Zick 2017), where the type of utility representation functions used was crucial for PAC stability.

Next, we show that Theorem 4.1 is ‘tight’ in the sense that if the graph G contains a cycle, \mathcal{H}_G is not PAC stabilizable.

Theorem 4.3. *Given a non-forest graph $G = \langle N, E \rangle$, the class \mathcal{H}_G is not PAC stabilizable.*

Proof Sketch. Since G is not a forest, there is a cycle in G . Without loss of generality, let $C = \{1, 2, \dots, k\}$ be a cycle with $\{i, i+1\} \in E$ for all $i = 1, 2, \dots, k-1$, and $\{k, 1\} \in E$. Let $S_1 = \{1, 2\}$, $S_2 = \{2, 3\}$, $S_3 = \{3, \dots, k, 1\}$. Suppose \mathcal{D} is the uniform distribution on $\{S_1, S_2, S_3\}$ and that the following holds:

$$S_1 \succ_1 S_3, S_2 \succ_2 S_1, S_3 \succ_3 S_2. \quad (2)$$

In this case, nothing beyond (2) can be deduced about the game by examining samples from \mathcal{D} . Consider the following games satisfying (2):

- A game Γ_1 where every player $i \in \{1, 2, 3\}$ strictly prefers $\{i\}$ to any other coalition, and any non-singleton coalition is less preferred than S_i and S_{i-1} , namely, $\{i\} \succ_i S_i \succ_i S_{i-1} \succ_i S'$ for any $S' \in \mathcal{N}_i \setminus \{S_i, S_{i-1}, \{i\}\}$. Here we set $S_0 = S_3$. Every player $j \in S_3$ strictly prefers S_3 to any other coalition.
- A game Γ_2 where every player in C strictly prefers $C = \{1, 2, \dots, k\}$ to any other coalition, and every player $i \in \{1, 2, 3\}$ strictly prefers S_i to any coalition other than C . Every player $j \in S_3$ strictly prefers S_3 to any other coalition other than C .

Suppose towards a contradiction that there is an algorithm \mathcal{A} that returns a $\frac{1}{3}$ -PAC stable partition π . Then, for π to be resistant against deviations supported by \mathcal{D} , π has to include $\{1\}$ or $\{2\}$ or $\{3\}$ for the first game, and C for the second game, which implies that it is impossible to achieve $\varepsilon < \frac{1}{3}$ with any confidence $1 - \delta > 0$. \square

5 Inferring Tree Interaction Networks from Data

Until now, we assume that the underlying interaction network G was given to us as input; this is, naturally, an assumption that we would like to forgo. Suppose the underlying graph is a forest $T = \langle N, E \rangle$, and consider the question of whether it is possible to infer a forest $T^* = \langle N, E^* \rangle$ that agrees with the original graph with high probability. Let \mathcal{T}_n be the set of all possible trees over n vertices, and let \mathcal{F}_n be the set of all possible forests; \mathcal{F}_n is our hypothesis class for guessing an approximate forest. More formally, \mathcal{F}_n consists of functions f_G that given an n vertex forest G , output 1 if a set of vertices is connected, and 0 otherwise. By Cayley’s formula:

$$|\mathcal{T}_n| = n^{n-2} \quad (3)$$

Any forest can be obtained by choosing a tree, and then choosing a subset of its edges, hence:

$$|\mathcal{F}_n| \leq |\mathcal{T}_n| 2^{n-1} = n^{n-2} 2^{n-1} \quad (4)$$

We observe the following variant of Theorem 2.1 for finite hypothesis classes.

Theorem 5.1 (Anthony and Bartlett (1999)). *Let \mathcal{C} be a finite hypothesis class where $\log |\mathcal{C}|$ is polynomial in n . If there exists a polynomial time algorithm that for any $v \in \mathcal{C}$, and samples*

$$\langle S_1, v(S_1) \rangle \dots, \langle S_m, v(S_m) \rangle$$

finds a function $v^ \in \mathcal{C}$ consistent with the samples, i.e., $v^*(S_j) = v(S_j)$ for each $j = 1, 2, \dots, m$, then \mathcal{C} is efficiently PAC learnable.*

Since $\log |\mathcal{F}_n| < 2n \log n$, all we need is to establish the existence of an efficient algorithm to compute a forest consistent with a given sample. More formally, let T be an unknown forest; we are given a set of m subsets of vertices labeled ‘connected’ or ‘disconnected’ according to T , can we find a forest that is consistent with the labeling? First, we consider an easier question and assume all subsets are connected. The answer to this question is affirmative, and appears in Conitzer, Derryberry, and Sandholm (2004).

Theorem 5.2 (Conitzer, Derryberry, and Sandholm (2004)). *Let $T = \langle N, E \rangle$ be a tree. Given a list S_1, \dots, S_m of connected vertices in T , there exists a poly-time algorithm that outputs a tree T^* where every subset S_j is connected in T^* .*

Theorem 5.2 pertains to trees, but immediately generalizes to forests by noting that if T is a forest, any tree whose edgeset is a superset of E is a valid solution as well, hence the same algorithm solves the problem.

In other words, if one only observes subsets of feasible coalitions and players’ preferences over them, it is possible to find a forest structure consistent with the samples.

Corollary 5.3. *If the probability distribution \mathcal{D} supports only connected subgraphs, \mathcal{F}_n is PAC learnable over \mathcal{D} .*

Corollary 5.3 is immediately implied by (4), Theorems 5.1 and 5.2. Theorem 4.1 assumes that the underlying interaction graph is known to us. Leveraging Corollary 5.3, we now show

that this assumption can be forgone; that is, it is possible to PAC stabilize a hedonic game whose underlying interaction graph is a forest, even if the forest structure is unknown to us. Note that we established that the forest structure can be PAC learned efficiently only if the sample contains exclusively connected coalitions, yet we do not have this requirement for PAC stabilizability.

Theorem 5.4. *Let $\mathcal{H}^* = \bigcup\{\mathcal{H}_G \mid G \text{ is a forest}\}$ be the class of all hedonic games whose interaction graph is a forest; then \mathcal{H}^* is efficiently PAC stabilizable.*

Theorem 5.2, while interesting in its own right, provides us with only a partial understanding of the problem: if all one is given is positive examples, it is possible to find a tree structure that is consistent with all connected coalitions. In what follows, we study a more general question of whether we can find a forest consistent with both positive (connected coalitions) and negative (disconnected coalitions) examples. As we show in Theorem 5.5, introducing the possibility of negative examples makes the problem computationally intractable, even if we restrict ourselves to the hypothesis class of paths. Hence, forests cannot be PAC learned efficiently. It is interesting to note that Theorem 5.4 could be achieved despite this negative result.

5.1 The Complexity of Constructing Consistent Trees

We now argue that deciding whether there exists a forest consistent with both positive and negative examples is computationally intractable; in fact, this claim holds even when the desired forest is a path. This result stands in sharp contrast to known computational results in the literature; indeed, there are several efficient algorithms for such restricted networks when only connected coalitions are taken into account³ (Booth and Lueker 1976; Korte and Möhring 1987; Corneil, Olariu, and Stewart 1998; Fulkerson and Gross 1965; Habib et al. 2000; Kratsch et al. 2006; Hsu and Ma 1999).

Specifically, we are given m samples of node subsets S_1, \dots, S_m ; each subset S_j is labeled by a function ℓ_G such that

$$\ell_G(S_j) = \begin{cases} 1 & \text{if } S_j \text{ is connected in } G \\ 0 & \text{otherwise.} \end{cases} \quad (5)$$

We say that a graph $G^* = \langle V, E^* \rangle$ is *consistent with* $G = \langle V, E \rangle$ over the samples $\mathcal{S} = \{S_1, \dots, S_m\} \subseteq V$ if and only if $\ell_{G^*}(S_j) = \ell_G(S_j)$ for all $j \in [m]$. Our objective is to find a forest T^* such that $\ell_{T^*}(S_j) = \ell_G(S_j)$ for all j . Theorem 5.5 states that it is NP-hard to determine whether such a graph exists.

Theorem 5.5. *Given a family of subsets $\mathcal{S} \subseteq 2^N$ such that each set in \mathcal{S} is of size at most 3, and a mapping $\ell : \mathcal{S} \rightarrow \{1, 0\}$, it is NP-hard to decide whether there exists a path $T^* = \langle N, E \rangle$ such that $\ell_{T^*}(S) = \ell(S)$ for each $S \in \mathcal{S}$. The result also holds when T^* is a forest.*

³The problem of deciding the existence of a path consistent with connected coalitions is equivalent to the problem of determining whether the intersection graph of a hypergraph is an interval, which is also closely related to testing the consecutive ones property of a matrix (see, e.g. the survey by Dom (2009) for more details).

To conclude, if one allows observations of both connected and disconnected components, finding a forest consistent with samples is computationally intractable. We note that this does not preclude the existence of efficient heuristics computing consistent forest structures in practice: as previously mentioned, inferring PAC approximations of forest structures has a low communication complexity (Theorem 5.1); thus, given access to strong MILP solvers, we believe that identifying consistent forest structures should be easy in practice.

6 Conclusions and Future Work

This work establishes a strong connection between interaction structure and the ability to guarantee approximate stability in hedonic games; simply put, we show that if one only knows the underlying interaction structure and nothing more, then one can only obtain PAC stable outcomes if the underlying interaction structure is very well-behaved, i.e. a forest. This result seems to imply a natural tradeoff: our work assumes very little knowledge about underlying player preferences, and thus requires a lot of structure; Sliwinski and Zick (2017) make no assumptions on the underlying interaction network, but assume a more restricted player preference model. It would be interesting to explore ‘intermediate’ cases; that is, suppose we make some structural assumptions on the interaction network, what classes of player preferences admit PAC stable outcomes?

We make use of tools from computational learning theory in order to analyze hedonic coalition formation. We believe that as a research paradigm, this is a useful and important methodological approach. Hedonic games (and cooperative games in general) have, by and large, seen sparse application. Other game-theoretic methods have been successfully applied by taking a problem-oriented approach (e.g. stable matching for resident-hospital allocation (Kleinberg and Tardos 2006, Chapter 1.1), or Stackelberg games in the security domain (Tambe 2011)); a concrete problem modeled and solved by a hedonic game framework has not yet been identified, to the best of our knowledge; this is despite the wealth of potential application domains, and rich data environments available nowadays (in particular, social network datasets would be particularly agreeable to the type of analysis presented in this work). Our work makes a fundamental connection between data, community structure, and game-theoretic solution concepts; a connection that we hope will result in a more applicable model of strategic collaborative behavior.

Acknowledgements. We thank reviewers at AAAI-19 for helpful feedback. Jakub Sliwinski was supported by a Singapore MOE Grant #R-252-000-625-133; Jakub Sliwinski and Yair Zick were funded by NRF Fellowship #R-252-000-750-733. Part of this work was conducted while Ayumi Igarashi visited NUS, supported by the NRF Fellowship #R-252-000-750-733.

References

- Anthony, M., and Bartlett, P. 1999. *Neural Network Learning: Theoretical Foundations*. Cambridge University Press.
- Aziz, H., and Brandl, F. 2012. Existence of stability in hedonic

- nic coalition formation games. In *Proc. of the 11th AAMAS*, 763–770.
- Aziz, H., and Savani, R. 2016. Hedonic games. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A., eds., *Handbook of Computational Social Choice*. Cambridge University Press. chapter 15.
- Balcan, M.; Constantin, F.; Iwata, S.; and Wang, L. 2012. Learning valuation functions. In *Proc. of the 25th COLT*, 4.1–4.24.
- Balcan, M.; Procaccia, A.; and Zick, Y. 2015. Learning cooperative games. In *Proc. of the 24th IJCAI*, 475–481.
- Balcan, M. F.; Sandholm, T.; and Vitercik, E. 2018. A general theory of sample complexity for multi-item profit maximization. In *Proc. of the 19th EC*, 173–174.
- Balcan, M.; Vitercik, E.; and White, C. 2016. Learning combinatorial functions from pairwise comparisons. In *Proc. of the 29th COLT*, 1–35.
- Balkanski, E.; Syed, U.; and Vassilvitskii, S. 2017. Statistical cost sharing. In *Proc. of the 30th NIPS*, 6222–6231.
- Banerjee, S.; Konishi, H.; and Sönmez, T. 2001. Core in a simple coalition formation game. *Social Choice and Welfare* 18(1):135–153.
- Bogomolnaia, A., and Jackson, M. 2002. The stability of hedonic coalition structures. *Games and Economic Behavior* 38(2):201–230.
- Booth, K. S., and Lueker, G. S. 1976. Testing for the consecutive ones property, interval graphs, and graph planarity using pq-tree algorithms. *Journal of Computer and System Sciences* 13(3):335 – 379.
- Bousquet, N.; Li, Z.; and Vetta, A. 2015. Coalition games on interaction graphs: A horticultural perspective. In *Proc. of the 16th EC*, 95–112.
- Chalkiadakis, G.; Greco, G.; and Markakis, E. 2016. Characteristic function games with restricted agent interactions: Core-stability and coalition structures. *Artificial Intelligence* 232:76–113.
- Conitzer, V.; Derryberry, J.; and Sandholm, T. 2004. Combinatorial auctions with structured item graphs. In *Proc. of the 19th AAAI*, 212–218.
- Corneil, D. G.; Olariu, S.; and Stewart, L. 1998. The ultimate interval graph recognition algorithm? In *Proc. of the 9th SODA*, 175–180.
- Deincko, V., and Woeginger, G. 2013. Two hardness results for core stability in hedonic coalition formation games. *Discrete Applied Mathematics* 161(13):1837–1842.
- Demange, G. 2004. On group stability in hierarchies and networks. *Journal of Political Economy* 112(4):754–778.
- Dom, M. 2009. Algorithmic aspects of the consecutive-ones property.
- Fearnley, J.; Gairing, M.; Goldberg, P.; and Savani, R. 2013. Learning equilibria of games via payoff queries. In *Proc. of the 14th EC*, 397–414.
- Fulkerson, D. R., and Gross, O. A. 1965. Incidence matrices and interval graphs. *Pacific Journal of Mathematics* 15(3):835–855.
- Gairing, M., and Savani, R. 2010. Computing stable outcomes in hedonic games. In *Proc. of the 3rd SAGT*, 174–185.
- Habib, M.; McConnell, R.; Paul, C.; and Viennot, L. 2000. Lex-bfs and partition refinement, with applications to transitive orientation, interval graph recognition and consecutive ones testing. *Theoretical Computer Science* 234(1):59 – 84.
- Hsu, W.-L., and Ma, T.-H. 1999. Fast and simple algorithms for recognizing chordal comparability graphs and interval graphs. *SIAM Journal on Computing* 28(3):1004–1020.
- Igarashi, A., and Elkind, E. 2016. Hedonic games with graph-restricted communication. In *Proc. of the 15th AAMAS*, 242–250.
- Igarashi, A.; Sliwinski, J.; and Zick, Y. 2018. Forming probably stable communities with limited interactions. *arXiv preprint arXiv:1811.04616*.
- Kearns, M., and Vazirani, U. 1994. *An introduction to computational learning theory*. MIT press.
- Kleinberg, J., and Tardos, E. 2006. *Algorithm Design*. Pearson Education.
- Korte, N., and Möhring, R. H. 1987. *A simple linear-time algorithm to recognize interval graphs*. Berlin, Heidelberg: Springer Berlin Heidelberg. 1–16.
- Kratsch, D.; McConnell, R. M.; Mehlhorn, K.; and Spinrad, J. P. 2006. Certifying algorithms for recognizing interval graphs and permutation graphs. *SIAM Journal on Computing* 36(2):326–353.
- Meir, R.; Zick, Y.; Elkind, E.; and Rosenschein, J. 2013. Bounding the cost of stability in games over interaction networks. In *Proc. of the 27th AAAI*, 690–696.
- Morgenstern, J., and Roughgarden, T. 2016. Learning simple auctions. In *Proc. of the 29th COLT*, 1298–1318.
- Myerson, R. 1977. Graphs and cooperation in games. *Mathematics of Operations Research* 2(3):225–229.
- Peters, D., and Elkind, E. 2015. Simple causes of complexity in hedonic games. In *Proc. of the 24th IJCAI*, 617–623.
- Peters, D. 2016. Graphical hedonic games of bounded treewidth. In *Proc. of the 30th AAAI*, 586–593.
- Shashua, A. 2009. Introduction to machine learning: Class notes 67577. *arXiv preprint arXiv:0904.3664*.
- Sinha, A.; Kar, D.; and Tambe, M. 2016. Learning adversary behavior in security games: A PAC model perspective. In *Proc. of the 15th AAMAS*, 214–222.
- Sliwinski, J., and Zick, Y. 2017. Learning hedonic games. In *Proc. of the 26th IJCAI*, 2730–2736.
- Tambe, M. 2011. *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. Cambridge University Press.
- Woeginger, G. 2013. Core stability in hedonic coalition formation. In *Proc. of the 39th SOFSEM*, 33–50.