Learning Cooperative Solution Concepts From Voting Behavior: A Case Study on the Israeli Knesset

Extended Abstract

Omer Lev
Ben-Gurion University of the Negev
omerlev@bgu.ac.il

Wei Lu
Massachusetts Institute of Technology
luwei.here@gmail.com

Alan Tsang
Carleton University
akhtsang@gmail.com

Yair Zick
University of Massachusetts, Amherst
yzick@umass.edu

ABSTRACT

Most frameworks for computing solution concepts in hedonic games are theoretical in nature, and require complete knowledge of all agent preferences, an impractical assumption in real-world settings. This paper presents the first application of strategic hedonic game models on real-world data. We show that PAC stable solutions can reflect Members of Knesset' political positions and reveal politicians who are known to deviate from party lines. Moreover, these models compare favorably to machine learning models.

KEYWORDS

Hedonic games, Cooperative games, PAC, PAC stability, Parliamentary politics, Israeli Knesset

ACM Reference Format:
Omer Lev, Wei Lu, Alan Tsang, and Yair Zick. 2021. Learning Cooperative Solution Concepts From Voting Behavior: A Case Study on the Israeli Knesset: Extended Abstract. In *Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), Online, May 3–7, 2021*, IFAAMAS, 3 pages.

1 OUR CONTRIBUTIONS

We use voting records of the Israeli Knesset to construct hedonic game models and corresponding solutions. We compare these against standard clustering and community detection techniques, using the party affiliations as ground truth. Our hedonic game-based partitions reveal Members of Knesset's political positions at large. Moreover, our methodology compares well to standard ML techniques, even identifying 'rogue' MKs who break party lines.

The notion of *PAC stability*, introduced in Sliwinski and Zick [5], allows us to directly find *probably* stable coalition structures. That is, we create a coalition structure that is stable *according to the information we have* (though there is a small probability we have not seen a datapoint that would render our coalition structure unstable). To do so, we design efficient algorithms that compute PAC stable outcomes under various player preference models. We provide an intuitive graphic environment to present our results (https://knesset.s3.amazonaws.com/index.html, which contains additional models and analysis not included in this paper).

Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), U. Endriss, A. Nowé, F. Dignum, A. Lomuscio (eds.), May 3–7, 2021, Online.

2 PRELIMINARIES

We model Members of Knesset (MKs) as players in a *hedonic game*. A hedonic game is given by a set of players $N = \{1, \dots, n\}$. Let $\mathcal{N}_i = \{S \subseteq N : i \in S\}$ be the set of all subsets (known as *coalitions*) containing player i. Each player $i \in N$ has a complete preference order $\succ_i$ over $\mathcal{N}_i$. Most works assume that players' *ordinal* preferences are encoded via *cardinal* utilities; in other words, players have a utility function $v_i : \mathcal{N}_i \to \mathbb{R}$ such that $S \succ_i T$ iff $v_i(S) > v_i(T)$. A coalition structure π is a partition of the player set; π is in the core if for every coalition $S \subseteq N$, at least one member of S weakly prefers their assigned coalition to S; in other words, if $\pi(i)$ is the coalition that i is assigned to under π, then π is in the core if for every $S \subseteq 2^N$ there exists some $i \in S$ such that $\pi(i) \succeq_i S$. Coalition structures in the core are often referred to as *stable*.

2.1 PAC Stability in Hedonic Games

We assume only partial access to preferences — we are given a *dataset* $S_1, \dots, S_m$ of m observations, where each data entry is a coalition $S_j \subseteq N$, and the (cardinal) valuations of players in S_j $(v_i(S_j))_{i \in S}$. We assume that $S_1, \dots, S_m$ are sampled i.i.d. from some distribution $\mathcal{D}$, and as will future coalitions. This is a natural assumption in data analysis, where $S_1, \dots, S_m$ are the *training data* (used to train a model, or in our case, a solution concept), and future samples are taken from the *test data*. Indeed, in our experimental evaluation, we take i.i.d. samples from the Knesset voting data, which forms our training data. Our algorithms offer *probably stable* solutions, as described below.

Hedonic core stability can be considered as capturing *local loss*: given coalition structure π and coalition $S \subseteq N$, loss is $\lambda(\pi, S) = 1$ if π was unable to hedge against S members deviating; 0 otherwise.

Given distribution $\mathcal{D}$, the *expected loss* of π w.r.t. $\mathcal{D}$ is

$$L_{\mathcal{D}}(\pi) = \Pr_{S \sim \mathcal{D}} [\lambda(\pi, S) = 1] \quad (1)$$

This captures a probabilistic core condition: rather than requiring $\lambda(\pi, S) = 0$ for all $S \subseteq N$ (as is the case for the core), we require that it is low w.r.t. $\mathcal{D}$. Thus, our objective is to find coalition structures that incur low expected loss. More formally, a *PAC stabilizing* algorithm takes as input i.i.d. samples $S_1, \dots, S_m \sim \mathcal{D}^m$, and outputs coalition structure π^* (a function of the samples) that guarantees

$$\Pr_{(S_1,\dots,S_m) \sim \mathcal{D}^m} [L_{\mathcal{D}}(\pi^*) \geq \varepsilon] < \delta \quad (2)$$

Intuitively, δ captures the probability that the i.i.d. our observations are 'badly distributed'. In other words, in a vast majority of the

m samples ($\geq 1 - \delta$) the output of our PAC stabilizing algorithm incurs $< \varepsilon$ expected loss. We require that m, the number of samples needed to offer the guarantee in (2), is polynomial in n, $\frac{1}{\varepsilon}$ and $\log \frac{1}{\delta}$. Note that this formulation sidesteps learning player preferences, and directly learns a stable outcome from samples. Indeed, a series of recent works [3–5] present efficient algorithms for computing PAC stable outcomes. Jha and Zick [4] show that only *consistency* with samples is needed to ensure PAC stability, using sample size linear in n: an algorithm is a *consistent solver* if given a set of samples $S_1, \ldots, S_m$ evaluated by a hedonic game (N, v), its output π^* satisfies $\lambda(\pi^*, S_j) = 0$ for all $j \in 1, \ldots, m$. In other words, a coalition structure that is stable w.r.t. to the observed samples is likely to be stable w.r.t. future samples, for a sufficiently large m.

2.2 The Israeli Knesset Data

The Israeli political system consists of multiple parties, partially due to its proportional voting system, and diverse political landscape. The Knesset is the unicameral legislative branch of the national government. We focus on the twentieth Knesset (2015-2019), which included ten parties. However, its political landscape is far more nuanced[1]. Recently Israeli parties generally align along a right-left axis based on their stance on the Israeli-Palestinian conflict. This simplifies our considerations when analyzing and comparing the models. Moreover, due to procedural changes, coalition discipline increased in the past few years (including this dataset).

The Knesset website provides data access through the Open Data Protocol (OData) on past MKs, bills, and member votes on every bill. The Knesset has 120 seats, but the twentieth Knesset has 147 members due to some MKs resigning or joining mid-term. We retrieve data on all 147 MKs' information including name, party affiliation and their votes for all 7515 bills deliberated. A member's vote can take on one of the following values: 0 (vote canceled), 1 (vote for), 2 (vote against), 3 (abstained), and 4 (did not attend).

3 METHODOLOGY

Previous work [1, 6] involved learning the underlying *complete preference profile* before finding a stable partition. This is infeasible because of the preference profiles here are exponentially large representation in the number of players. PAC stability inspires an alternative approach: we directly learn a PAC stable partition from the partial preference relations observed in the Knesset data.

Hedonic game models require a complete ranking of coalitions for each player. We use the voting data to impute a preference relation for each parliament member. This extended abstract focuses on one method: an Appreciation of Friends model using PAC learning.

We sample i.i.d. (with replacement) 3/4 of all bills, repeating 50 times for consistency. The following compares the partitions produced by our models to ground truth party affiliations.

3.1 Appreciation of Friends Model

Players classify others as friends or enemies, and prefer coalitions with more friends and fewer enemies: Formally, let G_i be player i's set of friends, and B_i the set of enemies. $G_i \cup B_i \cup i = N$ and $G_i \cap B_i = \emptyset$. A preference profile P^f is based on *appreciation of friends* if for all players $i \in N, S \succeq_i T$ if and only if $|S \cap G_i| > |T \cap G_i|$, or $|S \cap G_i| = |T \cap G_i|$ and $|S \cap B_i| \leq |T \cap B_i|$. We define a player i's friends as anyone whose votes agreed more often than disagreed. Agreed votes only count if i's vote is either "for" or "against".

Dimitrov et al. [2] proposed an algorithm for finding core stable partitions for these preference profiles. We incorporate Sliwinski and Zick 2017's PAC core finding algorithm to "PAC-ify" their algorithm [2] and use this algorithm to compute a core stable partition.

[1] See http://bit.ly/2sJUZEi-knesset20 for an overview.

[2] We also greatly improve its running time on our data set (see full paper).

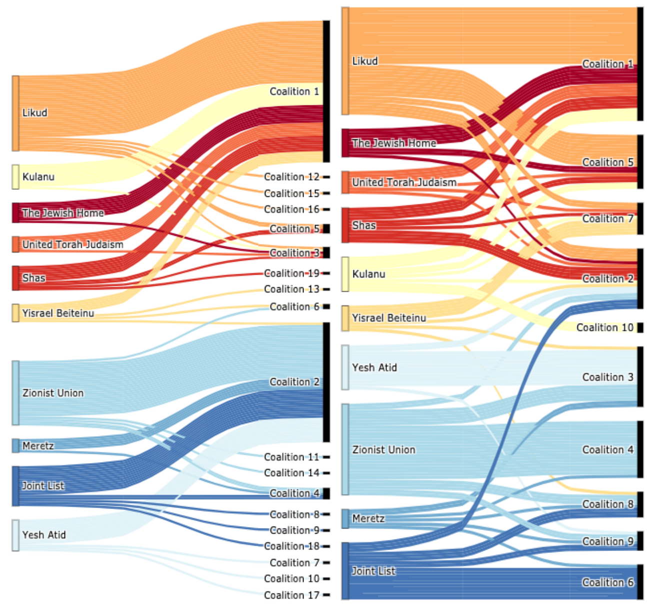

(a) PAC selective friends **(b) k-means ($k = 10$)**

Figure 1: Ground Truth vs. Model Generated Partitions

4 RESULTS

We visualize our results using the *Sankey diagram* with ground truth (party affiliation) on the left and our model partitions on the right. Each link from the left to the right represents a parliament member. Richer, more detailed diagrams can be seen at https://knesset.s3.amazonaws.com/index.html. Right wing parties are colored in reddish hues and left wing parties, in blueish hues.

Fig 1 compares the results of our PAC Friends model against the results of $k - means$ machine learning algorithm ($k = 10$). Our model (Fig 1a) performs well. It is able to effectively separate the government and opposition parties. While on the surface, it makes several "mistakes". On closer inspection, we see the two cross-ideological groups contain MKs that are known to deviate from party lines Coalition 6 is a small coalition combining two low-attendance members, one on the right edge of the left wing and another, who switched between coalition and opposition.

By comparison, the k-means models, however, has trouble identifying the relatively coherent groups of the government and the opposition (Fig 1b), forming an unlikely and sizable cross-ideological coalition (Coalition 2).

REFERENCES

[1] Maria-Florina Balcan, Florin Constantin, Satoru Iwata, and Lei Wang. 2012. Learning Valuation Functions. In *Proceedings of the 25th Conference on Computational Learning Theory (COLT)*. 4.1–4.24.

[2] Dinko Dimitrov, Peter Borm, Ruud Hendrickx, and Shao Chin Sung. 2006. Simple Priorities and Core Stability in Hedonic Games. *Social Choice and Welfare* 26, 2 (2006), 421–433.

[3] Ayumi Igarashi, Jakub Sliwinski, and Yair Zick. 2019. Forming Probably Stable Communities with Limited Interactions. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*. 2053–2060.

[4] Tushant Jha and Yair Zick. 2020. A Learning Framework for Distribution-Based Game-Theoretic Solution Concepts. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*. 355–377.

[5] Jakub Sliwinski and Yair Zick. 2017. Learning Hedonic Games. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*. 2730–2736.

[6] Daniel Vainsencher, Ofer Dekel, and Shie Mannor. 2011. Bundle Selling by Online Estimation of Valuation Functions. In *Proceedings of the 28th International Conference on Machine Learning (ICML)*. 1137–1144.