Learning Hedonic Games via Probabilistic Topic Modeling

Athina Georgara, Thalia Ntiniakou, and Georgios Chalkiadakis

School of Electrical and Computer Engineering, Technical University of Crete, Greece
{ageorgara, tntiniakou, gehalk} @intelligence.tuc.gr

Abstract. A usual assumption in the hedonic games literature is that of complete information; however, in the real world this is almost never the case. As such, in this work we assume that the players' preference relations are hidden: players interact within an unknown hedonic game, of which they can observe a small number of game instances. We adopt probabilistic topic modeling as a learning tool to extract valuable information from the sampled game instances. Specifically, we employ the online Latent Dirichlet Allocation (LDA) algorithm in order to learn the latent preference relations in Hedonic Games with Dichotomous preferences. Our simulation results confirm the effectiveness of our approach.

Keywords: Adaptation and Learning, Cooperative Game Theory

1 Introduction

Hedonic games constitute a class of cooperative games that intuitively attempts to capture the interpersonal relations amongst the players and the social bonds of the formed coalitions. That is, such games model settings where each player defines an ordinal preference relation over all the possible groups she can participate in, depending *exclusively* on the identity of her co-partners. Though most previous works on hedonic games considers complete information over the game, in a more realistic framework that would not be a plausible assumption. In real-life settings that can be modelled as hedonic games, we face the problem of uncertainty, i.e. the agents have little or no information about the overall game. For instance, the players find themselves in a new, unknown environment, and need *to discover their own preferences* through the interaction with others. For this reason, it is essential for an agent to be able to *learn* the underlying hidden game. This was in fact the motivation for the authors in [18] to explore the *Probably Approximately Correct (PAC) learnability* of several classes of hedonic games: it studied how good probabilistic "hedonic" utility function approximations can be derived from a (polynomial) number of samples, and proposed algorithms to do so for specific problem instances in the process.

In this work, we use *probabilistic topic modeling* in order to extract information about each agent's preferences, and thus learn the underlying game. Probabilistic topic models (PTMs) is a statistical approach used in analyzing words of documents that was originally used in data mining to discover a distribution over topics related to a given text document. Here, instead, we are inspired by recent work of [15], and employ a widely used PTM algorithm, *online Latent Dirichlet Allocation (LDA)*, to operate on instances of formed coalitions, in order to discover the ordinal preferences relations of each agent.

As such, our contributions in this paper are as follows. We propose a novel way to learn the player utility functions, capturing their ordinal preferences. In order for our method to work, we need to prescribe a way to represent coalitions and preferences orderings into text documents that will be the input of the algorithm. For this reason, we present in Section 3 a novel procedure to interpret a pair of coalition-preference order into a 'bag-of-words', which can then be channelled into the PTM algorithm. We conducted a preliminary yet systematic evaluation of our approach. Our results show that given a small number–compared to the coalitional space of a game–of observations, we can form beliefs that to a great extent reflect the agents' real preference relations. Note that a concrete benefit derived from our work here, is that once the validity of a topic has been established, an agent can use this information in the future to propose coalitions during some coalition formation protocol. Moreover, our approach can be used by, e.g., recommender systems, to promote bundles of goods.

2 Background and Related Work

In this section we discuss the fundamental notions of our work, hedonic games and probabilistic topic modeling, along with previous related works.

2.1 Hedonic Games

Hedonic Games were initially introduced by [9] to describe economic situations where individuals act in collaboration, and have personal preferences for belonging in a specific coalition. In [4], the authors studied concepts of stability on hedonic coalitions, while [2] provides a thorough and extensive study of hedonic games. To begin, a hedonic game G is given by a tuple $\langle N; \succsim \rangle$. $N = \{1, 2, \cdots, n\}$ is a finite set of agents, and $\succsim = (\succsim_1, \cdots, \succsim_n)$ is a preference relation, where each $\succsim_i \subseteq N_i \times N_i$ is a complete, reflexive and transitive preference relation over the coalitions i can possibly belong to. Given a hedonic game $G = \langle N, \succsim \rangle$, agent $i \in N$ *strongly* prefers coalition $S \in N_i$ over coalition $T \in N_i$ if and only if $S \succ_i T$, and agent i is indifferent between S and T *if and only if* $S \sim_i T$; thus, agent i *weakly* prefers (or just prefers) coalition S over T *if and only if* $S \succsim_i T$.

The outcome of a hedonic game is a coalition structure, i.e. a partition of N. Unlike most class of cooperative games, in hedonic games there is no payoff (reward) assigned to the participants. However, since hedonic games is considered to be a subclass of *non-transferable utility (NTU) games*, one can naturally assume that in such games each agent earns some individual 'moral satisfaction' by simply being part of a certain coalition; that is, each agent is rewarded with personal satisfaction by collaborating with specific other agents. In this light, we can intuitively claim, as Sliwinski and Zick do in their work in [18], that each agent forms a personal utility function v_i which mirrors her personal relation $\succsim_i$; in other words, we translate the abstract notion of preferences over coalitions into a mathematical representation. Formally, given a game $G = \langle N, \succsim \rangle$, for each agent $i \in N$ there is a utility function $v_i : N_i \to \mathbb{R}$ such that for all coalitions $S, T \in N_i$ holds $v_i(S) \geq v_i(T)$ if and only if $S \succsim_i T$.

A particularly appealing class of hedonic games is when agent preferences are *dichotomous*. The notion of dichotomous preferences was initially studied within economic settings, as in [5] to describe situations where outcomes can deterministically be distinguished into *good* and *bad*. In the context of game theory, the concept of dichotomous preferences within hedonic games was introduced in [1] through Boolean Hedonic Games, and several aspects regarding their complexity were thoroughly studied in [17]. In hedonic games with dichotomous preferences, or in boolean hedonic games (BHGs), each agent partitions the coalitional space into two disjoint sets. Formally, let $N_i = \{S \subseteq N : i \in S\}$ be the set of all coalitions that contains agent i; thus, N_i can be partitioned into N_i^+, and N_i^- such that $N_i^+ \cup N_i^- = N_i$ and $N_i^- \cap N_i^+ = \emptyset$. Intuitively, N_i^+ is the set of *desirable* coalitions corresponding to good outcomes and N_i^- is the set of *non-desirable* coalitions. That is, agent i strictly prefers all coalition in N_i^+ to those in N_i^-, and she is indifferent otherwise; i.e. $S \succ_i T$ if and only if $S \in N_i^+$ and $T \in N_i^-$, and $S \sim_i T$ if and only if $S, T \in N_i^+$ or $S, T \in N_i^-$. We refer to coalitions in N_i^+ as *satisfactory*, and to coalitions in N_i^- as *dissatisfactory*.

2.2 Probabilistic Topic Modeling

Probabilistic topic models (PTMs), are statistical methods, introduced in the linguistic scenario of uncovering underlying (latent) topics in a collection of documents. Topic modeling algorithms have been also adapted to other scenarios as well, for example in genetic data, images and social networks. A widely known and successful PTM is that of Latent Dirichlet Allocation (LDA)[3].

Latent Dirichlet Allocation We first describe the basic terms behind latent Dirichlet allocation (LDA) following [3] and [15]:

- ▷ A *word* is the basic unit of discrete data. A vocabulary consists of words and is indexed by $\{1, 2, ..., V\}$. The vocabulary is fixed and is fed as input to the LDA model.
- ▷ A *document* is a series of L words, $(w_1, w_2, \ldots, w_L)$.
- ▷ A *corpus* is a collection of D documents.
- ▷ A *topic* is a distribution over a vocabulary.

LDA is a Bayesian probabilistic topic model, in which each document can be described by a mixture of topics. A generative process for each document in the collection is assumed, where a random distribution over topics is chosen, and for each word in a document a topic is chosen from the distribution; finally, a word is drawn from the chosen topic. The same set of topics is shared by all documents in corpus, but each document exhibits topics in different proportions.

Latent variables, describing the hidden structure LDA indents to uncover, are assumed to be included to the generative process. The topics are $\beta_{1:K}$, where K is the dimensionality of the topic variable, which is known and fixed. Each topic β_k, is a distribution over the vocabulary of the corpus, where $k \in \{1, \dots, K\}$; and β_{kw} is the probability of word w in topic k. For the d^{th} document, θ_d is the distribution over topics; and θ_{dk} is the topic proportion of topic k in d. The topic assignments for the d^{th} document are indicated by z_d, and the topic assignment for the l^{th} word of the d^{th} document is denoted by z_{dl}. Consequently, w is the only observed variable of the model and w_{dl} represents the l^{th} word seen in the d^{th} document, while β, θ and z are latent variables. The posterior of the topic structure given the documents is:

$$p(\beta_{1:K}, \theta_{1:D}, z_{1:D} \mid w_{1:D}) = \frac{\beta_{1:K}\theta_{1:D}z_{1:D}w_{1:D}}{p(w_{1:D})}$$

where D is the number of documents, and the computation of the denominator, i.e. the probability of seeing the given document under any topic structure, is intractable [3]. Moreover, LDA includes priors, so that β_k is drawn from a Dirichlet distribution with parameter η and θ_d is drawn respectively from a Dirchlet also, with parameter α.

As mentioned, the posterior cannot be computed. To approximate it, the two most prominent approaches are (a) variational inference introduced in [13] and (b) Markov Chain Monte Carlo sampling methods proposed in [12]. In variational inference, the true posterior is approximated by a simpler distribution q, which depends on parameters $\phi_{1:D}$,$\gamma_{1:D}$ and $\lambda_{1:K}$ defined as:

$$\phi_{dwk} \propto \exp\{E_q[\log \theta_{dk}] + E[\log \beta_{kw}]\},$$

$$\gamma_{dk} = \alpha + \sum_{w} n_{dw}\phi_{dwk}\lambda_{kw} = \eta + \sum_{d} n_{dw}\phi_{dwk}$$

The probability that topic assignment of word w in d is k, under distribution q, is denoted by ϕ_{dwk}. Variable n_{dw} represents how many times the word w has been seen in document d. The variational parameters $\gamma_{1:D}$ and $\lambda_{1:K}$ are associated with variable n_{dw}. The variational inference algorithm's intuition is to minimize the *Kullback-Leibler(KL) divergence* between the variation distribution and the true (intractable) posterior. This is accomplished by iterating between assigning values to document-level variables, and updating topic-level variables.

Online Latent Dirichlet Allocation In the online version of LDA topic model [11], the documents are received on streams rather than a single batch of the original LDA algorithm. In this approach, the exact number of documents is not required to be known, though an estimation is at least required. As a result, online LDA can adapt to very large corpora. The value of the variational parameter $\lambda_{1:K}$ is updated every time a new batch arrives, while the rate at which the documents of batch t actually affects the value of $\lambda_{1:k}$ is controlled by $\rho_t = (\tau_0 + t)^{-k}$. The variational inference for the online version of LDA is shown in Algorithm 1, where it eventuates that α and η are assigned to a value once, and remain fixed. Ultimately, the estimated probability of the term w in topic k is $\beta_{kw} = \frac{\lambda_{kw}}{\sum \lambda_k}$.

Initialize λ randomly
for $t = 1$ **to** ∞ **do**
  $\rho_t = (\tau_0 + t)^{-k}$
  Expectation step:
  Initialize γ_{tk} randomly
  repeat
    set $\phi_{twk} \propto \exp\{E_q[log\theta_{tk}] + E_q[log\beta_{kw}]\}$
    set $\gamma_{tk} = \alpha + \sum_w n_{tw}\phi_{twk}$
  until $(\frac{1}{k}\sum|$ *change in* $\gamma_{tk}| \leq \epsilon)$;
  Maximization step:
  compute $\tilde{\lambda}_{kw} = \eta + Dn_{nt}\phi_{twk}$
  set $\lambda = (1 - \rho_t)\lambda + \rho_t\tilde{\lambda}$
end

Algorithm 1: Online Variational Inference for LDA [11]

2.3 Uncertainty

In most real-life settings, a common phenomenon is the absence of partial or complete information. That is, in situations modeled by a game, the participants (players) usually lack knowledge regarding other players (especially if the number of agents is very large), or are uncertain about even their own contribution in the overall game. For this reason, it is essential for the agents to be able to learn the unknown aspects of the game. That is, the agents should be capable of correctly estimating the outcome (e.g. the potential gain or loss) of a given synergy.

A recent work that studies the uncertainty problem is that of [15]. In this study, the authors exploit online LDA model, a probabilistic topic model, to discover coalitional values during an iterative overlapping coalition formation process. The coalitional values are based on an underlying collaboration structure, emerging due to existing but unknown relations among the agents. All agents extract, from a number of sample text documents, information regarding the formed coalitions' utilities. Specifically, they proposed the OVERPRO method which allows each agent to refine significant agents, and therefore characterize coalitions as profitable and unprofitable. In the model of [15], each agent trains her own online LDA probabilistic topic model, during the iterative process, and discover coalitions significant to her. This is the only work we are aware of that explicitly tries to learn the underlying inter-agent synergies or the underlying cooperation structure, and actually does so using probabilistic topic modeling.

Of course, there are many other papers that attempt to deal with various aspects of uncertainty in cooperative game settings. For instance, Kraus et al. in [14] study a series of strategies for revenue distribution in environments with incomplete information. In [7], the authors propose a Bayesian reinforcement learning model, in order to reduce uncertainty regarding coalitional values.

Sliwinski and Zick in [18] tackle the uncertainty problem in hedonic games, by adopting *probably approximately correct* algorithms in order to approximate agents' preference relations and discover core stable partitions. The authors work on discovering the underlying hedonic game, by estimating each agent's personal utility; and to that extend, find stable coalitions, i.e. find coalitions where agents are less (stable) or more (unstable) probable to diverse. [18] exploits the characteristics a hedonic game should have in order to be PAC learnable and/or PAC stabilizable. In particular, the authors show that certain classes of hedonic games, such as Additevelt Separable Hedonic Games and $\mathcal{W}-$Hedonic Games, are PAC learnable; and therefore focus on PAC stabilizability of hedonic games.

In our work, we adopt the probabilistic topic modeling framework of [15] and apply it to hedonic games. In particular, we lay our interest on a special class of hedonic games, that of *hedonic games with dichotomous preferences* (also known as boolean hedonic games, discussed in Section 2.1). Of course, for our approach to work, we need first to devise a way for coalitional instances to be transformed into input "documents" the LDA can operate on. In Section 3 we describe in detail how this is achieved.

3 Game Interpretation as Documents

Let $G = \langle N, \boldsymbol{\phi}_1, \cdots, \boldsymbol{\phi}_n \rangle$ be a hedonic game with dichotomous preferences, where N is a set of players and $n = |N|$. $\boldsymbol{\phi}_i$ represents a logic formula correlated to agent $i \in N$, which allows agent i to 'approve' or 'disapprove' a given coalition. The formula $\boldsymbol{\phi}_i$ consists a concise representation of the preference relation $\succsim_i$ (the preference relation $\succsim_i$ may be of size exponential in n, as opposed to formula $\boldsymbol{\phi}_i$ which may be significantly shorter). In hedonic games with dichotomous preferences, each agent classifies the coalitions related to her into satisfactory coalitions and dissatisfactory ones. Intuitively, formula $\boldsymbol{\phi}_i$ expresses agent i's goal, and agent i is satisfied if her goal is achieved, or dissatisfied otherwise.

We define an instance κ of game G as a tuple $\langle \pi^{\kappa}, satisfied_1, \cdots, satisfied_n \rangle$, where π^{κ} is a partition of N, and $satisfied_i$ is a auxiliary boolean variable that indicates whether agents i's goal is achieved or not. We let each instance κ produce n documents, one document per agent. Since preferences are assumed to be exclusively personal, it is natural to consider n formulae to be independent. Under this assumption, it is natural to train and maintain n different LDA models each of which learns a single formula. Therefore, we assign a single model to each agent, which corresponds to the formula related to that agent; i.e. agent i is responsible for the i^{th} model, which is used to discover $\boldsymbol{\phi}_i$.

We sample m instances of the game G. Every instance produces n documents, where each document refers to exactly one's agent formula. Thus, in total we have $m \cdot n$ documents to train n different models. That is, each agent i uses in her own probabilistic topic model exactly m documents, which correspond to her own formula $\boldsymbol{\phi}_i$. Thus, the corpus of each model is of size $(1/n)\%$ of the total number of produced documents.

This approach is similar to the one used in [15]. However, in [15] agents belonging in the same coalition process identical documents describing this coalition; while in our case, agents belonging to the same coalition process *different* documents describing the same coalition. This divergences from [15] makes our representation distinct, and is due to the nature of the games: the work of [15] concentrates in transferable utility games, where a coalition is related to a single utility; by contrast, we work with hedonic games with dichotomous preferences, where each agent participating in a coalition S characterizes S as *satisfactory* or *disatisfactory*, regardless of the corresponding characterization by her partners.

A document $\mathbf{d}_{i,\kappa}$ is related to agent i in instance κ, and contains the following: an indicative word for each agent in coalition $\pi^\kappa(i)$, the indicative word 'gain' if the coalition $\pi^\kappa(i)$ is satisfactory (i.e. if $\pi^\kappa(i) \in N_i^+$), or the indicative word 'loss' if $\pi^\kappa(i)$ is disatisfactory ($\pi^\kappa(i) \in N_i^-$). For example, with $N = \{1, 2, 3, 4, 5\}$, and an instance $S^\kappa = \langle\{1, 2\}, \{3, 5\}, \{4\}, True, False, True, False, True\rangle$, the produced documents are of the form:

$$\begin{array}{ll} \mathbf{d}_{1,\kappa} = \big(ag_1, ag_2, gain\big) & \mathbf{d}_{2,\kappa} = \big(ag_1, ag_2, loss\big) \\ \mathbf{d}_{3,\kappa} = \big(ag_3, ag_5, gain\big) & \mathbf{d}_{4,\kappa} = \big(ag_4, loss\big) \\ \mathbf{d}_{5,\kappa} = \big(ag_3, ag_5, gain\big) & \end{array}$$

As we can see, agents belonging in the same coalition may have identical documents, $d_{3,\kappa} \equiv d_{5,\kappa}$; while other agents may have different documents, $d_{1,\kappa} \not\equiv d_{2,\kappa}$.

The approach can easily represent instances of quite complex preferences, where, e.g., agent 1 realizes she is happy to work with agent 2 or 3 alone, however, together they turn out to be insupportable, and thus she dislikes being with both of them. Of course, using a carefully crafted for a domain of interest translation method, and a possibly different unsupervised learning method, one can potentially extract even more subtle preferences (relating, e.g., to the degree by which an agent prefers a coalition or even a specific partner to others).

4 Experimental Evaluation

In order to evaluate the performance of probabilistic topic modeling in discovering good and bad outcomes in hedonic games, we work in environments of hedonic games with dichotomous preferences with 50 agents ($n = 50$). The coding for the simulation framework was in Python 3.5; all the experiments ran on a PC with an i5@3.2GHz processor and 4GB of RAM.

4.1 Dataset and Setting Escalation

For the simulations we created several game instances of hedonic games with dichotomous preferences, according to the following procedure:

1. for each game G define the preference relations through $\boldsymbol{\phi}$ formulae
2. generate randomly partitions of N, π
3. for each agent i decide whether $\pi(i)$ is satisfactory regarding formula $\boldsymbol{\phi}_i$
4. interpret and log the instance information $\langle \pi(i), satisfied_i \rangle$ into documents

Formulae Construction A formula $\boldsymbol{\phi}_i$ expresses agent i's goal in a concise and short representation. That is, each $\boldsymbol{\phi}_i$ consists of two subsets of agents: (a) the "must be included" (denoted below as included) agents and (b) the "must be excluded" (denoted below as excluded) agents. In its simplest form, agent i is satisfied within a coalition where all agents in the included set are members of this coalition, and no agent in the excluded set is participating. However, an agent i may have several pairs of $\langle included, excluded \rangle$ subsets of agents, and be satisfied with a coalition if this coalition is consistent with at least one such pair. Therefore, in the general form we have that: $\boldsymbol{\phi}_i = \bigvee_{l=1:L} \phi_{i,l} = \phi_{i,1} \vee \phi_{i,2} \vee \cdots \vee \phi_{i,L}$ where $\phi_{i,l} = included_{i,l} \wedge excluded_{i,l}$. The complexity of a formula $\boldsymbol{\phi}_i$ depends on (a) the number of $\phi_{i,l}$ and (b) the number of agents each pair $\langle included, excluded \rangle$ contains. We ran our experiments on settings with escalated complexity of the structure of the formulae $\boldsymbol{\phi}$.

- ▷ **Low Complexity:** $\boldsymbol{\phi}_i$s with low complexity consist of a single $\langle included,$ $excluded \rangle$ pair, and the total number of agents appearing in this pair of subsets is fixed to $n/10$. That is, each agent has a must include or exclude demand on 10% of the agents.
- ▷ **High Complexity:**[1] we escalate the complexity by increasing the number of $\langle included,$ $excluded \rangle$ pairs to 3. The number of agents appearing in each pair uniformly ranges in $[n/10, n/5]$.[2]

Instance Generation & Information Logging For a given hedonic game G, i.e. for a given set of formulae $\boldsymbol{\phi}$, we generate a number of game instances. In each instance we randomly partition agents into coalitions. Each coalition in the formed partition is characterized as satisfactory or dissatisfactory according to formula $\boldsymbol{\phi}_i$, for every agent i within the coalition. After characterizing the coalition of a certain agent, we log the contained information into a text document using the interpretation described in Section 3.[3] The total number of documents produced per game varies depending on the game's complexity, and ranges between $[500K, 2.5M]$. However, the size of the corpus each online LDA

[1] We have preliminary results showing our approach can be quite effective in even more complex settings.

[2] The number of agents participating in each $\boldsymbol{\phi}_i$, along with the total number of formulae per agent within each level of complexity environment were chosen so that the required dataset could be generated within a reasonable time frame; these numbers do not impose any burden on the LDA algorithm itself.

[3] For practical reasons, the logged information is repeated more than once within a document. That is, we boost the term frequency of the agents' indicative words, and the characterization's 'gain'/'loss', to avoid misleading words with low frequencies.

model is fed with, corresponds only to the 2% of the total number of produced documents. That is, we use $10,000$ and $50,000$ game instances for the low and the high complexity environments respectively.

Note that the game instances do not have stochasticity. That is, for each generated sample (document) of a game instance, it is deterministically guaranteed that if the sample contains the word 'gain'('loss') then the respective coalition is satisfactory (dissatisfactory). In a real life situation of a boolean hedonic game formation, we would sample instances of the game by letting participants form groups, and then receive by each one a feedback on whether or not they were satisfied in their coalition.

LDA Model For the implementation, we used the *scikit-learn* Python 3.5 library [16]. As mentioned, the online version of LDA was used for a range of topics, iterations and batch sizes related to the formulae complexity. In Table 1 we present the different parameters used in the simulations.

Complexity	$\|corpus\|$	$\|\phi_i\|$	number of Topics	batch size	iterations
Low	10000	1	5-12	5-25 (step 5)	50
High	50000	3	5-19(step 2)	5-25 (step 5)	50

Table 1. Simulation parameters

The number of topics K, is a parameter that the (online) LDA model needs to be provided with. In situations such as the ones we are studying, the exact number of topics cannot be known a priori; however, it can systematically be chosen depending on the problem at hand. Moreover, there exist other LDA variations, such as HDP [19], that are non-parametric on the number of topics.

4.2 Significant Agents and Valid Topics

In order to evaluate our experimental results, we first need to determine the coalition that is primarily described by each topic. For this reason, we define *significant agents* within each topic. Given a topic k, the agents that appear with probability greater than a small number ϵ are considered to be significant: formally, agent i is significant with respect to topic k if $Pr(i|topic = k) \geq \epsilon$. Therefore, the coalition related to topic k is $S = \{j \in N : Pr(j|topic = k) \geq \epsilon\}$.

After establishing the coalition described in each topic, we move to assessing the *validity* of the topic. Intuitively, the validity of a topic signifies whether the topic reflects a sub-formula describing the agent's hedonic preferences. Thus, given the significant agents within a topic, we characterize it as valid or invalid for agent i. A topic k is valid when:

$\triangleright$ $Pr(gain|topic = k) \geq Pr(loss|topic = k)$ and
$S = \{j : Pr(j|topic = k) \geq \epsilon\} \in N_i^+$, **or**

$$\triangleright \; Pr(loss|topic = k) \geq Pr(gain|topic = k) \text{ and}$$
$$S = \{j : Pr(j|topic = k) \geq \epsilon\} \in N_i^-$$

otherwise the topic is invalid. In words, the meaning of the latter characterization is essentially the actual cross-validation of the topics result with the corresponding formula $\boldsymbol{\phi}_i$.

Given these definitions, we can then adopt as an evaluation metric the percentages of valid and invalid topics found by the algorithm. Intuitively, we would like the algorithm to discover *valid* topics; that is, they reflect preferences subformulae that correspond to satisfactory/ dissatisfactory collaboration patterns.

4.3 Results

Before presenting our results, in Figure 1 we show two examples of topics (distributions over the "words"), in the low complexity environment. That is, each bar depicts the probability of a specific word belonging to the given topic. To clarify, the x-axis of the graphs depicts the vocabulary of our corpus, i.e. the bars $0 - 49$ correspond to the agent names; and the last two bars $50, 51$ represent the words 'gain' and 'loss', respectively. In the left graph, the distribution over the vocabulary is exhibited for a valid topic, which infers a satisfactory coalition–due to the relatively high probability of the word 'gain' (the 50^{th} word in the axis shown with green bar). Similarly, in the plot we see the graph for a dissatisfactory coalition scenario (the red bar corresponds to the word 'loss').

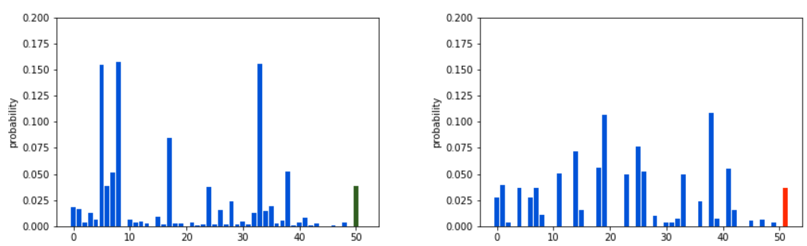

Fig. 1. Environment Complexity: *low*. Example of a topic describing 'satisfactory' coalition (left), and a topic corresponding to a 'dissatisfactory' coalition (right).

We now present our actual results. First, we conducted a series of game simulations for the low complexity environment. Specifically, we constructed 4 different games (denoted $G_{l,1}, G_{l,2}, G_{l,3}$ and $G_{l,4}$ in the Appendix A) that are subject to the characteristics of the low complexity environment. For each one of these games we ran the learning process 5 times, using batches of different size (of the same documents) channeled to the LDA model.[4] The resulting topics were evaluated using the metrics described above, and we computed the average percentage of *valid* and *invalid* topics per game per number of topics.

Number of Topics	(%) **valid**	(%) **invalid**
5	89.83%	10.17%
6	91.81%	8.16%
7	96.79%	3.21%
8	92.08%	7.92%
9	93.98%	6.02%
10	92.67%	7.33%
11	95.53%	4.47%
12	88.89%	11.11%

Table 2. Environment Complexity: *low*. Percentage of *valid* and *invalid* topics over 4 different hedonic games with dichotomous preferences.

Table 2 shows our average results in the low complexity setting. As we can see, generally, the model learns correctly at least 88% of the sub-formulae $\hat{\phi}_i$, for different numbers of topics.[5] At the same time, the percentage of incorrectly learned sub-formulae does not exceed 11.5%. In Table 6 (in Appendix A), the reader can see the detailed results per game in the low complexity environment. There, we can observe that in some games percentage of the correctly learned topics may even reach 100%, and it is for all consistently greater than 80%.

A similar evaluation process was followed for the high complexity environment. Again, we constructed 4 different games (denoted $G_{h,1}, G_{h,2}, G_{h,3}$ and $G_{h,4}$ in the Appendix A) that are subject to the characteristics of the high complexity environment. So, Table 3 shows the respective average percentage of average percentages per game per number of topics for *valid* and *invalid* topics, while Table 7 in Appendix A) shows the detailed average percentages per game. Here, we see that the average percentage of valid topics learned is greater than 78% for various number of topics. The average percentage of incorrectly learned topics does not exceed 22%. As the environment complexity increases along with the number of topics we intend to discover, the accuracy of learned collaboration patterns drops. Looking at the detailed results in Table 7, the correctly learned topics may even reach 100%, and it is for all always greater than 69%.

Topic Significance Constraint In a more realistic scenario, knowing that the probability of 'gain' is greater than the one of 'loss', and vice versa, within a

[4] As we have already mentioned, there is no stochasticity during the dataset creation. However, by employing this repetition of the learning procedure per game, we ensure the robustness of our results.

[5] A larger number of topics allows for more preferences sub-formulae to be learned, but it naturally increases complexity.

Number of Topics	(%) valid	(%) invalid
5	84%	16%
7	87%	13%
9	90%	10%
11	90%	10%
13	86.92%	13.08%
15	84.64%	15.36%
17	80.59%	19.41%
19	78.16%	21.84%

Table 3. Environment Complexity: *high*. Percentage of *valid*, *invalid* and *insignificant* topics over 4 different hedonic games with dichotomous preferences.

topic, may not be enough. Intuitively, we would like to be *confident* whether the coalition described within a topic is satisfactory or dissatisfactory. For this reason, we introduce a *topic significance constraint*, according to which a topic is labeled as significant if the absolute difference of the probability of term 'gain' and the probability of term 'loss', exceeds some small number δ, i.e. $|Pr(gain|topic = k) - Pr(loss|topic = k)| \geq \delta$. Thus, each topic is assessed as 'significant' with confidence level δ; and if a topic is 'significant' it can therefore assessed as 'valid' or 'invalid'. Now the definition of a valid topic k becomes:

$$\triangleright\ Pr(gain|topic = k) \geq Pr(loss|topic = k) + \delta \text{ and } S = \{j : Pr(j|topic = k) \geq \epsilon\} \in N_i^+, \textbf{ or}$$
$$\triangleright\ Pr(loss|topic = k) \geq Pr(gain|topic = k) + \delta \text{ and } S = \{j : Pr(j|topic = k) \geq \epsilon\} \in N_i^-$$

By taking into account the topic significance constraint, we re-evaluated the topics arisen from the LDA model for the low and the high complexity environment. The results depicted in Table 4 show an expected drop in the average percentage of valid topics, since we discard a portion of the topics by assessing them as 'insignificant'. It is worth noting that the difference between the average percentage in the unconstrained and constrained cases reaches up to 15.27 percentage points for the assessment of valid topics.

"Anytime" Behaviour Last but not least, we conducted a simulation experiment to examine how our model behaves during an ongoing learning process. That is, assume that agents at a certain time t_1 have access to a part of the corpus. The agents train their models with the sub-corpus that is available to them at the time. After this first phase training, agents have some beliefs over satisfactory and dissatisfactory coalitions, that they could use in a decision-making process. At time t_2 a second part of the corpus is revealed to the agents, thus the agents update their already partially trained models with the new documents; and so on and so forth. In each of the later phases, the values of parameters of the model regarding the number of batches and iterations are maintained. Equivalently, these later phases correspond to LDA processes with prior distributions over topics and documents. Intuitively, using priors leads to faster convergence.

Average (%) valid topics							
Topics	**low complexity**			**Topics**	**high complexity**		
	(%) *unconstr*	(%) *constr*	diff		(%) *unconst*	(%) *constr*	diff
5	89.83%	88%	−1.83	5	84%	82%	−2
6	91.81%	91.11%	−0.7	7	87%	82.86%	−4.14
7	96.79%	94.88%	−1.91	9	90%	85.56%	−4.34
8	92.08%	86.77%	−5.31	11	90%	82.27%	−7.73
9	93.98%	86.76%	−7.22	13	86.92%	73.08%	−13.84
10	92.67%	83.83%	−8.84	15	84.67%	71%	−13.67
11	95.53%	83.86%	−11.67	17	80.59%	67.35%	−13.24
12	88.89%	75.35%	−13.54	19	78.16%	62.89%	−15.27

Table 4. Environment Complexity: *low* / *high*. Average percentage of *valid* topics over 4 different games **with** and **without** the significance constraint.

In Table 5 we show the results of this procedure, for games in the high complexity environment and for varying numbers of topics. We let the agents train in 3 phases and recorded the time needed and the average percentage of valid topics for each phase. 2500 documents were fed to the model in each phase. As we can see, **Phase 0** requires more time, while the average percentages of valid topics are not particularly encouraging; **Phase 1** requires approximately $75 - 80\%$ of time required by **Phase 0**, and the percentages rise by up to 22 percentage points; similarly, **Phase 2** requires approximately 40% of the time required by **Phase 0**, and the average percentage of valid topics consistently gets close to or even reaches 100%.

Topics	**Phase 0**		**Phase 1**		**Phase 2**	
	time (sec)	valid (%)	time (sec)	valid (%)	time (sec)	valid (%)
5	$43.97s$	80%	$35.66s$	100%	$24.0s$	100%
7	$43.57s$	77.14%	$33.04s$	97.14%	$25.62s$	100%
9	$45.63s$	71.11%	$35.45s$	93.33%	$27.64s$	97.78%
11	$43.91s$	67.27%	$35.48s$	90.18%	$28.31s$	90.91%

Table 5. Environment Complexity: *high*. Anytime behaviour of the training model.

5 Conclusion and Future Work

In this work, we presented a probabilistic topic modeling approach for learning boolean hedonic games. To this end, we presented a novel method for translating game samples into "documents" fed to our algorithm as input. We conducted a systematic evaluation of our approach: first, we studied the performance of online LDA discovering collaboration patterns within environments of different preference relation complexity; and then, we examined the "anytime" performance of our method, by progressively revealing to the model parts of the total corpus. Our results verified the effectiveness of our approach. This work could inspire recommendation methods to be used, e.g., by online advertisers to suggest bundles of goods, or even by travel agencies to promote vacation packages.

As future work, we intend to apply our approach to different classes of hedonic games or non-transferable utility games, in general; and also test it in partition function game environments [8]. We actually already have preliminary results for additively separable hedonic games. Moreover, we intend to compare LDA with different machine learning methods (e.g. neural networks) applied on this domain. It is worth noting that this approach can potentially be used for learning (dichotomous or other) preferences in more generic domains, and not just hedonic games [6, 10].

Acknowledgements *We thank Michalis Mamakos for code sharing.*

References

1. Aziz, H., Harrenstein, P., Lang, J., Wooldridge, M.: Boolean hedonic games. In: 15th International Conference on Principles of Knowledge Representation and Reasoning (KR). pp. 166–175 (2016)
2. Aziz, H., Savani, R., Moulin, H.: Hedonic Games, pp. 356–376. Handbook of Computational Social Choice, Cambridge University Press (2016)
3. Blei, D.M., Ng, A.Y., Jordan, M.I.: Latent dirichlet allocation. Journal of machine Learning research **3**, 993–1022 (Mar 2003)
4. Bogomolnaia, A., Jackson, M.O.: The stability of hedonic coalition structures. Games and Economic Behavior **38**(2), 201 – 230 (2002)
5. Bogomolnaia, A., Moulin, H., Stong, R.: Collective choice under dichotomous preferences. Journal of Economic Theory **122**(2), 165 – 184 (2005)
6. Brandt, F., Conitzer, V., Endriss, U., Lang, J., Procaccia, A.D.: Handbook of Computational Social Choice. Cambridge University Press, USA (2016)
7. Chalkiadakis, G., Boutilier, C.: Bayesian reinforcement learning for coalition formation under uncertainty. In: Proceedings of the 3rd AAMAS - Volume 3. pp. 1090–1097. IEEE Computer Society, Washington, DC, USA (2004)
8. Chalkiadakis, G., Elkind, E., Wooldridge, M.: Computational Aspects of Cooperative Game Theory (Synthesis Lectures on Artificial Inetlligence and Machine Learning). Morgan & Claypool Publishers, 1st edn. (2011)
9. Dréze, J.H., Greenberg, J.: Hedonic coalitions: Optimality and stability. Econometrica **48**(4), 987–1003 (1980)
10. Frnkranz, J., Hllermeier, E.: Preference Learning. Springer-Verlag, Berlin, Heidelberg, 1st edn. (2010)
11. Hoffman, M.D., Blei, D.M., Bach, F.: Online learning for latent dirichlet allocation. In: Proceedings of NIPS. pp. 856–864. USA (2010)
12. Jordan, M.I. (ed.): Learning in Graphical Models. MIT Press, Cambridge, MA, USA (1999)
13. Jordan, M.I., Ghahramani, Z., Jaakkola, T.S., Saul, L.K.: An introduction to variational methods for graphical models. Machine Learning **37**(2), 183–233 (1999)
14. Kraus, S., Shehory, O., Taase, G.: The advantages of compromising in coalition formation with incomplete information. In: Proceedings of the 3rd AAMAS. pp. 588–595 (2004)

15. Mamakos, M., Chalkiadakis, G.: Overlapping coalition formation via probabilistic topic modeling. In: Proceedings of the 17th AAMAS. pp. 2010–2012 (2018), Full version: http://arxiv.org/abs/1804.05235
16. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: Machine learning in Python. Journal of Machine Learning Research **12**, 2825–2830 (2011)
17. Peters, D.: Complexity of hedonic games with dichotomous preferences. In: AAAI (2016)
18. Sliwinski, J., Zick, Y.: Learning hedonic games. In: Proceedings of the 26th IJCAI-17. pp. 2730–2736 (2017)
19. Teh, Y.W., Jordan, M.I., Beal, M.J., Blei, D.M.: Hierarchical dirichlet processes. Journal of the American Statistical Association **101** (2004)

A Appendix: Detailed results

Topics	**valid**(%)				**invalid**(%)			
	$G_{l,1}$	$G_{l,2}$	$G_{l,3}$	$G_{l,4}$	$G_{l,1}$	$G_{l,2}$	$G_{l,3}$	$G_{l,4}$
5	83.33%	84.00%	92.00%	100.00%	16.67%	16.00%	8.00%	0.00%
6	80.56%	96.67%	93.33%	96.67%	19.44%	3.33%	6.67%	3.33%
7	92.86%	100.00%	97.14%	97.14%	7.14%	0.00%	2.86%	2.86%
8	83.33%	97.50%	95.00%	92.50%	16.67%	2.50%	5.00%	7.50%
9	87.04%	100.00%	93.33%	95.56%	12.96%	0.00%	6.67%	4.44%
10	86.67%	100.00%	94.00%	90.00%	13.33%	0.00%	6.00%	10.00%
11	89.39%	100.00%	98.18%	94.55%	10.61%	0.00%	1.82%	5.45%
12	80.56%	90.00%	96.67%	88.33%	19.44%	10.00%	3.33%	11.67%

Table 6. Environment complexity: *low*. Detailed results of average percentage per different game for valid, invalid and insignificant topics for varying number of topics.

Topics	**valid**(%)				**invalid**(%)			
	$G_{h,1}$	$G_{h,2}$	$G_{h,3}$	$G_{h,4}$	$G_{h,1}$	$G_{h,2}$	$G_{h,3}$	$G_{h,4}$
5	96.00%	84.00%	84.00%	72.00%	4.00%	16.00%	16.00%	28.00%
7	94.29%	85.71%	94.29%	74.29%	5.71%	14.29%	5.71%	25.71%
9	97.78%	82.22%	95.56%	84.44%	2.22%	17.78%	4.44%	15.56%
11	94.55%	78.18%	100.00%	87.27%	5.45%	21.82%	0.00%	12.73%
13	89.23%	75.38%	100.00%	83.08%	10.77%	24.62%	0.00%	16.92%
15	82.67%	84.00%	97.33%	74.67%	17.33%	16.00%	2.67%	25.33%
17	76.47%	82.35%	87.06%	76.47%	23.53%	17.65%	12.94%	23.53%
19	69.47%	69.47%	92.63%	81.05%	30.53%	30.53%	7.37%	18.95%

Table 7. Environment complexity: *high*. Detailed results of average percentage per different game for valid, invalid and insignificant topics for varying number of topics.