Overlapping Coalition Formation via Probabilistic Topic Modeling

Extended Abstract

Michalis Mamakos
Northwestern University
Evanston, IL, USA
mamakos@u.northwestern.edu

Georgios Chalkiadakis
Technical University of Crete
Chania, Greece
gehalk@intelligence.tuc.gr

ABSTRACT

Research in cooperative games often assumes that agents know the coalitional values with certainty, and that they can belong to one coalition only. By contrast, this work assumes that the value of a coalition is based on an underlying collaboration *structure* emerging due to existing but *unknown relations* among the agents; and that agents can form *overlapping coalitions*. Specifically, we first propose *Relational Rules*, a novel representation scheme for cooperative games with overlapping coalitions, which encodes the aforementioned relations, and which extends the well-known MC-nets representation to this setting. We then present a novel decision-making method for decentralized overlapping coalition formation, which exploits *probabilistic topic modeling*—and, in particular, *online Latent Dirichlet Allocation*. By interpreting formed coalitions as documents, agents can effectively learn topics that correspond to profitable collaboration structures.

ACM Reference Format:
Michalis Mamakos and Georgios Chalkiadakis. 2018. Overlapping Coalition Formation via Probabilistic Topic Modeling. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), Stockholm, Sweden, July 10–15, 2018,* IFAAMAS, 3 pages.

1 INTRODUCTION

Cooperative game theory [5] provides a rich framework for the coordination of the actions of self-interested agents. In the related literature, it is usually assumed that an agent can be a member of exactly one coalition. However, in environments where agents hold a divisible resource (e.g., time, money), it is natural for them to split it in order to simultaneously participate in a number of *overlapping coalitions* [4, 6, 11, 12, 14]. Furthermore, it it is more natural than not to assume that agents do not have complete knowledge of the utility that can be yielded by every possible team of agents [3, 9, 10, 13]. We build on the idea of *marginal contribution nets* (MC-nets) [8] and introduce *Relational Rules (RRs)*, a representation scheme for cooperative games with *overlapping* coalitions.

An agent can make an observation of the utility earned by the resource offerings of the members of a coalition, but it is a much more complex task to determine her relations with subsets of agents of that coalition. Towards this end, we propose an agent-learning method that is based on probabilistic topic modeling (PTM) [1]. PTM is a form of *unsupervised learning* which is particularly suitable for unravelling information from massive sets of documents. Probabilistic topic models infer the probability with which each word of a given "vocabulary" is part of a *topic*. Intuitively, the words that have high probability in a topic, are very likely to appear together in a document that refers to this topic with high probability. Hence, a topic, which is a probability distribution of the words of a given vocabulary, reveals the underlying *hidden structure*.

The method we develop employs online Latent Dirichlet Allocation (LDA) [2, 7] to allow agents to learn how well they can cooperate with others. In our setting, agents *repeatedly* form overlapping coalitions, as the game takes place over a number of *iterations*. Thus, we utilize a simple, yet appropriate, protocol, under which in each iteration an agent is (randomly) selected in order to propose (potentially) overlapping coalitions. By interpreting formed coalitions as documents, represented given an appropriate vocabulary, agents are able to use online LDA to update beliefs regarding the hidden collaboration structure—and thus implicitly learn rewarding synergies with others (synergies which are in our experiments described by RRs). We have evaluated our approach against two reinforcement learning (RL) algorithms we developed for this setting, which our algorithm vastly outperforms, implying a high degree of accuracy in agents' beliefs, and a high quality of agent decisions.

To the best of our knowledge, the recent work of [11] is the only one that has so far approached *overlapping* coalition formation *under uncertainty*, but it is concerned with the class of Threshold Task Games [4]. Moreover, ours is the first paper that employs PTM for multiagent learning, introducing thus an entirely novel paradigm for decentralized learning in multiagent settings.

2 RELATIONAL RULES

Agents face uncertainty regarding the value of *synergies* among them, which, in a non-overlapping setting, are concisely described by MC-nets. We now extend MC-nets to overlapping environments by introducing *Relational Rules* (RR), with the following form:

$$A \rightarrow \frac{\sum_{i \in A} \pi_{i,C}}{|A|} \cdot value$$

where $A \subseteq N$ (with $N = \{1, \dots, n\}$ being the set of agents), $value \in \mathbb{R}$; $C \subseteq N$ is a coalition such that $A \subseteq C$; $\pi_{i,C}$ is the portion of her resource that i has invested in coalition C: i.e., $\pi_{i,C} = r_{i,C}/r_i$, where r_i is the total resource quantity (continuous or discrete) that i holds and $r_{i,C}$ is the amount she has invested in C. Therefore, $\pi_{i,C} > 0$, since $i \in C$ ($r_{i,C} = 0$ essentially means that $i \notin C$), and $\pi_{i,C} \leq 1$, since i can offer to C at most r_i.

Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. André, S. Koenig (eds.), July 10–15, 2018, Stockholm, Sweden.

A rule applies to coalition C if and only if $A \subseteq C$, and in that case utility $\frac{\sum_{i \in A} \pi_{i,C}}{|A|} \cdot value$ is added to the coalitional value of C. Note that it is *not* required that an agent's total resource quantity r_i has to be communicated to C's other members, since a rule is applied by the environment. In non-overlapping games, RRs reduce to MC-nets rules without negative literals, as it then holds that $\frac{\sum_{i \in A} \pi_{i,C}}{|A|} = 1$.

3 LEARNING BY INTERPRETING COALITIONS AS DOCUMENTS

Now, we present how agents can employ online LDA in order to effectively learn the underlying collaboration structure. We let each agent maintain and train her own online LDA model. We define the vocabulary of an agent's LDA model to include n words, one for each agent (including herself), indicating their contribution, plus two words for the utility, one representing gain and the other representing loss, since the value earned from a coalition can be either positive or negative. Thus, the vocabulary of an agent consists of $n + 2$ words. Assuming a game that proceeds in rounds, in round t agent $i \in C$ interprets the coalitional configuration regarding C as a document by "writing" in the document $r_{j,C}$ times the word that indicates the contribution of agent $j \in C$—where $r_{j,C} \in \mathbb{N}^+$ is the contribution of j to C. Agent i also "writes" in the document, that corresponds to C, either the word that indicates gain or the one that indicates loss as many times as the absolute value of the utility earned by the coalition is. Since words are discrete data, u_C cannot be real-valued, and so we restrict its value to integers. The number of documents that an agent passes in an iteration to her online LDA is equal to the number of coalitions she is member of.

The intuition behind the notion of a topic is that the words that appear in it with high probability are very likely to appear together in a document that exhibits this topic with high probability. Thus, *the probability with which the word corresponding to an agent's contribution appears in a topic, is correlated with the amount of her contribution.* Hence, the meaning of a topic identified by agent i, is that i has observed in many documents certain agents who contributed a lot, and some that contributed less; and this configuration *results to gain or loss with the corresponding probabilities.* Figure 1 illustrates an example of two typical topics formed by an agent.

4 EXPERIMENTAL EVALUATION

We evaluated our method, coined as OVERPRO, which exploits online LDA, in environments with 50 and 250 agents. OVERPRO requires the parameter tuple $\langle K, \tau_0, \kappa \rangle$ to be passed, where K is the number of topics, and τ_0, κ define the effect that a batch of documents has on the formation on the topics, at a particular round. We compared OVERPRO against two RL algorithms we developed. The first one is termed as Greedy top-k, where k defines the number of the best coalitions maintained, and the second is a Q-learning style algorithm, which requires the learning rate δ_t at round t to be given as a parameter. We have defined *efficiency* as the ratio of social welfare to total resource quantity invested by all agents in every coalition in a round. As depicted in Figure 2, OVERPRO outperforms both RL algorithms. Note that in our experiments the value of a coalition is determined through RRs, but agents do not know the RRs in effect, and hence cannot determine the value of a coalition with certainty. Therefore, agents do not know how well they can do with others, and cannot determine their relations just by an observation of a coalitional value.

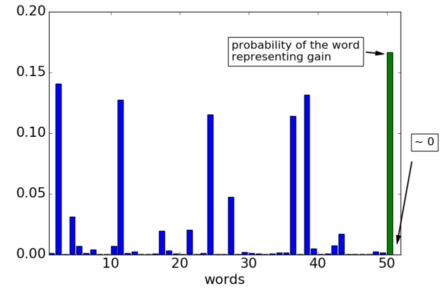

(a) A "profitable" learned topic.

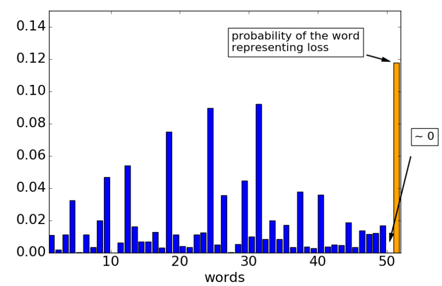

(b) A "non-profitable" learned topic.

Figure 1: Typical topics, as formed by a randomly selected agent at the end of a random iteration in an experiment, where an agent's vocabulary consists of 52 words ($n = 50$). The two last words in a topic indicate the probability of gain and loss respectively, while the rest correspond to agents' contribution.

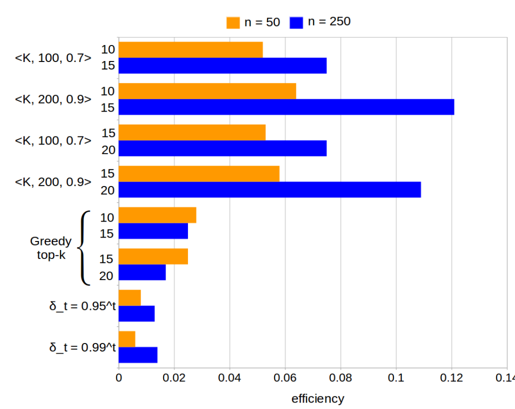

Figure 2: Average efficiency defined as the ratio of social welfare to total resource quantity invested by all agents. For OVERPRO the values of K (number of topics), and respectively for Greedy top-k the values of k, are denoted on the left of each bar.

REFERENCES

[1] David M Blei. 2012. Probabilistic topic models. *Commun. ACM* 55, 4 (2012), 77–84.

[2] David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent dirichlet allocation. *Journal of Machine Learning Research* 3 (2003), 993–1022.

[3] Georgios Chalkiadakis and Craig Boutilier. 2004. Bayesian reinforcement learning for coalition formation under uncertainty. In *Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multiagent Systems*. 1090–1097.

[4] Georgios Chalkiadakis, Edith Elkind, Evangelos Markakis, Maria Polukarov, and Nick R Jennings. 2010. Cooperative games with overlapping coalitions. *JAIR* 39, 1 (2010), 179–216.

[5] Georgios Chalkiadakis, Edith Elkind, and Michael Wooldridge. 2011. Computational aspects of cooperative game theory. *Synthesis Lectures on Artificial Intelligence and Machine Learning* (2011).

[6] Viet Dung Dang, Rajdeep K Dash, Alex Rogers, and Nicholas R Jennings. 2006. Overlapping coalition formation for efficient data fusion in multi-sensor networks. In *AAAI*. 635–640.

[7] Matthew Hoffman, Francis R Bach, and David M Blei. 2010. Online learning for latent dirichlet allocation. In *NIPS*. 856–864.

[8] Samuel Ieong and Yoav Shoham. 2005. Marginal contribution nets: a compact representation scheme for coalitional games. In *Proc. of the 6th Conference on Electronic Commerce*. 193–202.

[9] Samuel Ieong and Yoav Shoham. 2008. Bayesian Coalitional Games.. In *AAAI*. 95–100.

[10] Sarit Kraus, Onn Shehory, and Gilad Taase. 2003. Coalition formation with uncertain heterogeneous information. In *Proceedings of the 2nd International Joint Conference on Autonomous Agents and Multiagent Systems*. 1–8.

[11] Michail Mamakos and Georgios Chalkiadakis. 2017. Probability Bounds for Overlapping Coalition Formation. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*. 331–337.

[12] Onn Shehory and Sarit Kraus. 1998. Methods for task allocation via agent coalition formation. *Artificial Intelligence* 101, 1 (1998), 165–200.

[13] Jeroen Suijs, Peter Borm, Anja De Waegenaere, and Stef Tijs. 1999. Cooperative games with stochastic payoffs. *European Journal of Operational Research* 113, 1 (1999), 193–205.

[14] Yair Zick, Evangelos Markakis, and Edith Elkind. 2014. Arbitration and stability in cooperative games with overlapping coalitions. *Journal of Artificial Intelligence Research* 50 (2014), 847–884.