

# Deep Learning-Powered Iterative Combinatorial Auctions with Active Learning

Extended Abstract

Benjamin Estermann  
ETH Zurich  
Zurich, Switzerland  
estermann@ethz.ch

Roger Wattenhofer  
ETH Zurich  
Zurich, Switzerland  
wattenhofer@ethz.ch

Stefan Kramer  
ETH Zurich  
Zurich, Switzerland  
stefan.kramer@inf.ethz.ch

Ye Wang  
University of Macau  
Macau, China  
wangye@um.edu.mo

## ABSTRACT

Deep learning-powered iterative combinatorial auctions (DL-ICA) are auctions that utilize machine learning techniques. Unlike traditional auctions, bidders in DL-ICA do not need to report the valuations for all bundles upfront. Instead, they report their value for certain bundles iteratively, and the allocation of the items is determined by solving a winner determination problem. During this process, the bidder profiles are modeled with neural networks. However, DL-ICA may not always achieve the optimal winner allocation due to the relatively low number of reported bundles, resulting in reduced economic efficiency. This paper proposes an algorithm that uses active learning for initial sampling strategies to improve the resulting economic efficiency (social welfare). The proposed algorithm outperforms previous studies in real-world combinatorial auction models across various domains while using fewer samples on average.

## KEYWORDS

Active Learning for Regression; Deep Learning; Combinatorial Auctions

### ACM Reference Format:

Benjamin Estermann, Stefan Kramer, Roger Wattenhofer, and Ye Wang. 2023. Deep Learning-Powered Iterative Combinatorial Auctions with Active Learning: Extended Abstract. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 3 pages.

## 1 INTRODUCTION

Traditional auctions only allow bidders to bid on individual items, which can result in the exposure problem. For example, consider an auction for advertisement slots on a TV channel where a bidder needs three consecutive slots to broadcast a 30-second commercial. If the first two slots have unexpectedly high prices due to intense competition, the bidder may not have enough funds to acquire the third slot, rendering the first two slots useless and decreasing the social welfare of the auction. Combinatorial auctions (CA) allow

bidders to bid on bundles of items to avoid the exposure problem [1, 3, 6], but the bundle space grows exponentially with the number of items, making it impossible for bidders to report their full value function. To address this issue, iterative algorithms have been developed [5, 8], which interact with bidders and ask for a limited number of bundles in each round. However, these algorithms may not always lead to the optimal allocation, resulting in reduced economic efficiency.

Our work builds on previous research in deep learning-powered iterative combinatorial auctions (DL-ICA) [12] and proposes a modification of the machine learning-based elicitation algorithm by selecting a set of initial bundles more efficiently in order to improve the efficiency of the final allocation. Previous studies have used uniform random sampling for the initial request [2, 4, 12]. Uniform sampling is a simple and straightforward method for selecting initial bundles in combinatorial auctions, but it has a fundamental limitation: it does not take into account the complexity of the bidder’s valuation of different bundles. Since the bundle space can be exponentially large, a random sample of bundles is unlikely to explore the entire space and may not include bundles that are highly valued by bidders. This can lead to a poor quality of elicited valuations and, consequently, a suboptimal allocation.

Active learning is a solution that can significantly reduce the amount of labeled data needed to train a model [10, 13]. In active learning, the machine learning algorithm selects the most informative data points from a pool of unlabeled data points, and asks an annotator to label them [11]. In DL-ICA, active learning allows the algorithm to select the most informative bundles to ask the bidders about. By doing so, the algorithm can gain a better understanding of the bidder’s valuation function with fewer queries compared to uniform sampling. This can lead to a more accurate and complete estimation of the bidder’s valuation function, resulting in a higher probability of finding an optimal allocation. Therefore, active learning is a promising approach to improve the efficiency of the elicitation process in combinatorial auctions.

Specifically, we propose a new algorithm, Greedy Active Learning on Input Values (GALI), which uses a greedy approach [14] to select the most informative initial bundles in DL-ICA. Our contribution is significant because it addresses an important challenge in combinatorial auctions: improving economic efficiency while reducing the number of bundles required from bidders. By using

*Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom. © 2023 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.*

**Table 1: Comparison of Uniform Sampling (UF) and Greedy Active Learning on Input Values (GALI) in the LSVM and GSVM models. The results “Average #Queries” and “Max #Queries” are measured per bidder. The value in parantheses in the Efficiency column is the standard deviation of the efficiency.**

Model	Sampling Technique	Average #Queries	Max. #Queries	Average #Iterations	Max. #Iterations	Efficiency in %
GSVM	UF	39.7	49.6	4.6	10.0	<b>97.95</b> (0.32)
	GALI	37.2	45.5	3.9	8.4	<b>99.18</b> (0.20)
LSVM	UF	50.9	57.2	5.0	10.3	<b>96.80</b> (0.41)
	GALI	47.5	52.3	4.0	8.8	<b>97.55</b> (0.32)

machine learning combined with active learning techniques, our proposed algorithm achieves better results than previous studies.

## 2 GREEDY ACTIVE LEARNING ON INPUT VALUES (GALI)

The goal of GALI is to ensure maximum diversity among the bundles initially queried. To determine which bundle to query, the active learning algorithm proposed by [13] iterates over all unlabeled bundles in the pool and computes their distance to the closest labeled bundle. The next bundle to be queried is then the bundle with the greatest distance to the closest labeled bundle. Since the size of our bundle space grows exponentially with the number of items (for each item, you can either take it or not take it), finding the next bundle to sample would lead to an exponential number of computations. We propose to frame this problem as a set of integer linear programs (ILPs), which can then be computed with an efficient ILP-solver. Let  $m$  be the number of items,  $\mathcal{X} = \{0, 1\}^m$  the bundle space and  $S \subset \mathcal{X}$  the set of already sampled bundles. The main idea is that for every sampled bundle  $s \in S$ , we want to find some  $x \in \mathcal{X} \setminus S$  that is furthest away from  $s$  but still not closer to any other bundle in  $S \setminus s$ . We need to find this bundle without iterating over all the bundles in  $\mathcal{X}$ . The following ILP (1) accomplishes this by exploiting the structure of  $\mathcal{X}$  and using a linearized notion of the distance norm.

$$\begin{aligned}
 & \arg \max_{x \in \mathcal{X}} \sum_{j=1}^m x_j + s_j - 2x_j s_j \\
 & \text{s.t.} \\
 & \sum_{j=1}^m x_j + s_j - 2x_j s_j \leq \sum_{j=1}^m x_j + s'_j - 2x_j s'_j \quad \forall s' \in S, s' \neq s \\
 & x, s \in \{0, 1\}^m \\
 & s' \in \{0, 1\}^m \quad \forall s' \in S, s' \neq s
 \end{aligned} \tag{1}$$

To determine all the bundles that should be queried from the bidders, GALI starts with a random bundle. It then iteratively solves the ILP (1) for each labeled bundle in the bundle space, where the next bundle to be queried is the one with the largest distance to its nearest already labeled bundle. Since the total number of bundles to be queried from the bidders is mostly constrained by the finite time the bidders have to properly evaluate them, this algorithm is able to terminate with a relatively small overhead compared to the uniform sampling approach. In particular, the bidder’s bundle bids

are not used to decide the next sample, so the bidders could submit all of their bids at the end instead of being queried in each iteration.

## 3 EXPERIMENTAL EVALUATION

We evaluate our approach on two well-known auction models, Global Synergy Value Model (GSVM) [7] and Local Synergy Value Model (LSVM) [9], both of which include bidder profiles with regional and national bidders. GSVM models an auction with 18 items and 7 bidders, 6 regional and one national bidder. The items are arranged in two circles, a national circle with 12 items and a regional circle with 6 items. The national bidder is interested in all of the items in the national circle, while the regional bidders are interested in 2 items in the national circle and 4 items in the regional circle. LSVM consists of 18 items and 6 bidders. As in GSVM, one of them is of national type, while the others are regional bidders. For each model, we run 51 different instances. All hyperparameters were kept consistent across both sampling strategies and were set according to the optimal values found by [12]. The results are summarized in Table 1. GALI is able to consistently outperform the UF baseline in both auction models in terms of efficiency achieved. It even does so while using on average fewer samples and fewer iterations.

## 4 CONCLUSION

This paper presents a novel approach to the initial sampling strategy of the machine learning-based elicitation algorithm of Brero et al. [5], with the aim of improving economic efficiency. The approach involves the use of active learning to acquire initial bundle-value pairs, specifically through the method of Greedy Active Learning on Input Values (GALI). The experiments show that the use of GALI can lead to higher efficiency in the allocation of goods for the GSVM and LSVM auction models while requiring fewer bundles to be queried from the bidders. In future work, it may be useful to explore other active learning methods and evaluate their potential for improving the efficiency of DL-ICA. In addition, further research could be conducted to determine the optimal number of initial bundle-value pairs required to achieve the desired level of efficiency. It may also be worthwhile to consider how the proposed modification to the iterative phase of the elicitation algorithm could be further refined or combined with other techniques to improve performance.

## ACKNOWLEDGMENTS

We thank the AAMAS 2023 reviewers for their feedback. This work was supported in part by University of Macau grants SRG2022-00032-FST and MYRG-CRG2022-00013-IOTSC.

## REFERENCES

- [1] Jawad Abrache, Teodor Crainic, Michel Gendreau, and Monia Rekik. 2007. Combinatorial auctions. *Annals of Operations Research* 153, 1 (2007), 131–164. <https://EconPapers.repec.org/RePEc:spr:annopr:v:153:y:2007:i:1:p:131-164>. [10.1007/s10479-007-0179-z](https://doi.org/10.1007/s10479-007-0179-z)
- [2] Manuel Beyeler, Gianluca Brero, Benjamin Lubin, and Sven Seuken. 2021. iMLCA: Machine Learning-powered Iterative Combinatorial Auctions with Interval Bidding. In *Proceedings of the 22nd ACM Conference on Economics and Computation*. 136–136.
- [3] Vitor Bosshard, Ye Wang, and Sven Seuken. 2018. Non-decreasing Payment Rules for Combinatorial Auctions. In *IJCAI*. 105–113.
- [4] Gianluca Brero, Sébastien Lahaie, and Sven Seuken. 2019. Fast iterative combinatorial auctions via bayesian learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 1820–1828.
- [5] Gianluca Brero, Benjamin Lubin, and Sven Seuken. 2018. Combinatorial Auctions via Machine Learning-based Preference Elicitation. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*. International Joint Conferences on Artificial Intelligence Organization, 128–136. <https://doi.org/10.24963/ijcai.2018/18>
- [6] Robin Fritsch, Younjo Lee, Adrian Meier, Ye Wang, and Roger Wattenhofer. 2022. Understanding the Relationship Between Core Constraints and Core-Selecting Payment Rules in Combinatorial Auctions. *arXiv preprint arXiv:2204.11708* (2022).
- [7] Jacob K Goeree and Yuanchuan Lien. 2016. On the impossibility of core-selecting auctions. *Theoretical economics* 11, 1 (2016), 41–52.
- [8] Jonathan Levin and Andrzej Skrzypacz. 2016. Properties of the combinatorial clock auction. *American Economic Review* 106, 9 (2016), 2528–2551.
- [9] Tobias Scheffel, Georg Ziegler, and Martin Bichler. 2012. On the impact of package selection in combinatorial auctions: an experimental study in the context of spectrum auction design. *Experimental Economics* 15, 4 (December 2012), 667–692. <https://doi.org/10.1007/s10683-012-9321-0>
- [10] Burr Settles. 2009. Active learning literature survey. (2009).
- [11] Masashi Sugiyama and Shinichi Nakajima. 2009. Pool-based active learning in approximate linear regression. *Machine Learning* 75, 3 (2009), 249–274.
- [12] J. Weissteiner and S. Seuken. 2020. Deep Learning-powered Iterative Combinatorial Auctions. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI-20)*. New York City, United States of America, 2284–2293. <https://doi.org/10.1609/aaai.v34i02.5606>
- [13] Dongrui Wu, Chin-Teng Lin, and Jian Huang. 2018. Active Learning for Regression Using Greedy Sampling. <https://doi.org/10.48550/ARXIV.1808.04245>
- [14] Hwanjo Yu and Sungchul Kim. 2010. Passive sampling for regression. In *2010 IEEE International Conference on Data Mining*. IEEE, 1151–1156.