Enhancing the Misreport Network for Optimal Auction Design

Haiying Wu HAIYINGWU@TJU.EDU.CN
Shuyuan You TJUYSY@TJU.EDU.CN
Zhiqiang Zhuang ZHUANG@TJU.EDU.CN
Tianjin University, Tianjin, China

Kewen Wang K.WANG@GRIFFITH.EDU.AU
Zhe Wang ZHE.WANG@GRIFFITH.EDU.AU
Griffith University, Brisbane, Australia

Abstract

Optimal auction mechanism design has long been a focus in computer science and economics. While substantial progress has been made in single-item auctions, optimal design for multi-item auctions has yet to be derived. Recent years have seen a surge in deriving near-optimal auctions through deep learning. As one of the approaches, ALGNet models the bidding process as a two-player game. The ALGNet model however adopted a rather simple design for generating optimal misreports to derive the regret of the trained auction mechanisms. We show that this design can be improved both in network structure and the testing method. Specifically, we train a misreport network tailored for each individual bidder which leads to better misreports. This approach is especially effective when the auctions are asymmetric. By studying misreport, we can get a more accurate estimate of the regret in the auction mechanism thus enhancing its robustness. Experimental results demonstrate that our approach can detect misreport more effectively than previous methods resulting in an increase in regret values as large as 70%. The new misreport network can also be applied to train auction mechanisms, allowing for a better description of the auction process.

1. Introduction

Optimal auction mechanism design is an important and active research area at the intersection of computer science, game theory, and economics [2–4, 10, 11, 14, 20, 21, 23, 24]. The focus is on designing an incentive-compatible auction mechanism that maximizes the seller's revenue. One prominent usage of such auctions is spectrum auctions [1, 22].

Myerson [15] solved the problem of optimal auction mechanism design for the single-item case. However, in the context of multi-item auctions, even in the simplest scenario of two bidders and two items, the theoretically optimal auction remains unknown. Due to this theoretical hurdle, subsequent research adopts an approach that treats auction design as a learning problem. To meet specific requirements, such as incentive compatibility (i.e. zero regret for bidders) and individual rationality (i.e. non-negative revenue for the auctioneer), researchers introduce various constraints in the learning process. Among the learning techniques, neural networks dominate auction mechanism training for their superior generalization and flexibility. RegretNet proposed by Duetting et al. [8] is a pioneering approach in the field. This type of learning framework has enabled researchers to approximate the optimal design, although it may not be the theoretically optimal one.

To improve the training process of RegretNet, ALGNet [18] models the auction process as a two-player game by using a neural network to approximate the misreport (i.e. fake bids made by bidders to improve their utilities). However, one issue with ALGNet is that although a neural network is utilized for misreport optimization in the training phase, there does not exist an equally powerful neural network for misreport optimization in the testing phase. As a result, inferior misreports are generated while testing the mechanism which leads to a lower regret value than the actual one.

To address this shortcoming, we propose to also train a misreport neural network for testing the mechanism. Experimental results confirm that the misreport network consistently generates better misreport than that used in the testing phase of ALGNet. Another issue with ALGNet is that it utilizes a single misreport network for all bidders. This design is optimal only in symmetric auctions in which the bidders draw their valuations from the same distribution. The design is not very effective in handling asymmetric auctions in which the bidders draw their valuations from different distributions. We therefore extend the ALGNet approach by training a distinct misreport network for each bidder. We refer to this approach as Multi-MisreportNet. Results show that the method captures each bidder's behavior more accurately.

In summary, we confirm and improve upon the shortcomings of ALGNet by using a separately trained misreport network in testing and an individualized misreport network for each bidder.

2. Preliminaries

Consider n bidders represented by $N = \{1, \ldots, n\}$ and m items available for auction represented by $M = \{1, \ldots, m\}$. Each bidder i has a valuation function $v_i : 2^M \to \mathbb{R}_{\geq 0}$. The value that bidder i assigns to a subset of items $S \subseteq M$ is denoted as $v_i(S) = \sum_{j \in S} v_i(\{j\})$, where v_i is independently drawn from a distribution F_i over possible valuation functions V_i. We denote the set of all possible valuation profiles as $V = \prod_{i=1}^{n} V_i$, and $v = (v_1, \ldots, v_n) \in V$ a the valuation profile.

Although the auctioneer knows the distributions $F = (F_1, \ldots, F_n)$, the actual valuations v of the bidders are unknown. Bidders report their bids and an auction mechanism determines the allocation of items to the bidders and the corresponding payments. An auction mechanism (g, p) is a combination of allocation rules $g_i : V \to 2^M$ and payment rules $p_i : V \to \mathbb{R}_{\geq 0}$. Given a bid profile $b = (b_1, \ldots, b_n) \in V$, the auction computes an allocation $g(b)$ and payments $p(b)$. We denote the valuation profile v without bidder i as v_{-i}, and similarly for bid profile b_{-i}. The set of possible valuation profiles for all bidders except bidder i is denoted as $V_{-i} = \prod_{j \neq i} V_j$.

A bidder's utility, denoted as $u_i(v_i, b) = v_i(g_i(b)) - p_i(b)$, is determined by their valuation v_i and the bid profile b. An auction is dominant strategy incentive compatible (DSIC) if truthful reporting is the optimal strategy for every bidder, regardless of the other bidders' reports. In formal terms, this means that $u_i(v_i, (v_i, b_{-i})) \geq u_i(v_i, (b_i, b_{-i}))$ holds true for all bidders i, all valuations $v_i \in V_i$, all bids $b_i \in V_i$, and all bids $b_{-i} \in V_{-i}$ from others. On the other hand, an auction is said to be (ex-post) individually rational (IR) if every bidder obtains a non-zero utility, meaning that $u_i(v_i, (v_i, b_{-i})) \geq 0$. This ensures that every bidder receives some benefit from participating in the auction and does not suffer a loss.

A DSIC auction is designed so that each bidder has a dominant strategy to report the true valuation, and the revenue on the valuation profile v is $\sum_i p_i(v)$. The goal is to identify an auction mechanism that satisfies DSIC and IR as much as possible and minimize the negated expected revenue $-\mathbb{E}_F\left[\sum_{i=1}^{n} p_i(v)\right]$. DSIC is measured the regret of the mechanism defined as follows:

$$rgt_i = \mathbb{E}\left[\max_{v_i' \in V_i} u_i\left(v_i, \left(v_i', v_{-i}\right)\right) - u_i\left(v_i, \left(v_i, v_{-i}\right)\right)\right] \tag{1}$$

Intuitively, a smaller regret means a mechanism is close to being DSIC. RegretNet [8] is the first to use neural networks to find near-optimal auction mechanisms, which use a neural network ω to approximate (g, p) to derive u and compute rgt. ALGNet [18] is a refinement of RegretNet that enhances the training process by using another neural network to approximate the optimal misreport $v_i' = argmax_{v_i' \in V_i} u_i\left(v_i, (v_i', v_{-i})\right)$. For the details of ALGNet, please refer to Appendix A.

3. MisreportNet

As we explained in Section 1, during the training process of ALGNet, there is no equally powerful network available for validation during testing. This limitation can lead to the discovery of slightly suboptimal misreports during testing, resulting in a lower regret value than the actual one.

To address the above issue, this section describes the design of MisreportNet during testing. Firstly, we present the utilization of Single-MisreportNet as an evaluation method for auction mechanisms. Secondly, we expand upon this model to develop Multi-MisreportNet, with the aim of identifying optimal misreports for different bidders.

3.1. Single-MisreportNet

The primary aim of our study is to design a suitable neural network denoted M^{φ} that accepts a valuation profile V as input and generates a matrix $M^{\varphi}(V) \in \mathbb{R}^{n \times m}$ as misreports to test the auction mechanism, as shown in Equation (2). Each row of this matrix corresponds to the optimal misreport value $\boldsymbol{v}_i'$ for bidder i.

$$M^{\varphi}(V)_{ij} = \operatorname{argmax}_{\boldsymbol{v}_i' \in F} u_i\left(\boldsymbol{v}_i, V_{-i}, \boldsymbol{v}_i'\right) \tag{2}$$

φ represents the parameters of neural network M. Given a valuation vector $\boldsymbol{v}_i$ for bidder i and the valuation vectors V_{-i} of other bidders, the neural network M^{φ} outputs the optimal misreport value $\boldsymbol{v}_i'$ for bidder i, such that bidder i obtains the maximum utility u_i.

$$\mathcal{L} = -\mathbb{E}_{V \in F}\left[\sum_{i=1}^{n}\left[u_i^w\left(\left[M^{\varphi}(V)\right]_i, V_{-i}\right) - u_i^w\left(\boldsymbol{v}_i, V_{-i}\right)\right]\right] \tag{3}$$

Given the parameters φ of the MisreportNet and the parameters w of the auction mechanism, the MisreportNet's performance is evaluated using a utility loss function $\mathcal{L}$, as shown in Equation (3), where the differences between the utility obtained by each bidder under the true valuation $\boldsymbol{v}_i$ and the utility obtained under the misreported valuation $\left[M^{\varphi}(V)\right]_i$ are computed. The results for all bidders are summed to obtain the total utility difference. A larger value of the utility difference ($-\mathcal{L}$) indicates a better performance of the MisreportNet, meaning greater accuracy in misreport and resulting in higher regret.

Table 1: Results of Single-MisreportNet (Single) under ALGNet Testing

3×10	**Regret per bidder**	**Bias**	**5×10**	**Regret per bidder**	**Bias**
Single	0.00263 ± 0.0002	70.78%	Single	0.00448 ± 0.0003	71.65%
ALGNet	0.00154 ± 0.0002		ALGNet	0.00261 ± 0.0004	

We present the Single-MisreportNet method tests performed in the ALGNet environment in Table 1. The mechanism testing using the neural network-generated misreports yields higher regret values compared to the ALGNet of direct bidding optimization. This indicates that the testing method of ALGNet is not as effective in identifying the optimal misreports for bidders, resulting in slightly underestimated regret values. Therefore, by employing the neural network to search for the optimal misreports, we can obtain more realistic misreport values.

3.2. Multi-MisreportNet

Through our validation in Section 3.1, the MisreportNet testing method's effectiveness was determined. However, this approach of using a single neural network to find multiple bidders for misreports is inherently inconsistent with the realities of auctions. Therefore, we propose a multi-MisreportNet with a 'one-bidder-one-misreport' structure. This design allows us to capture the misreport strategy of each bidder individually, leading to more realistic regret outcomes.

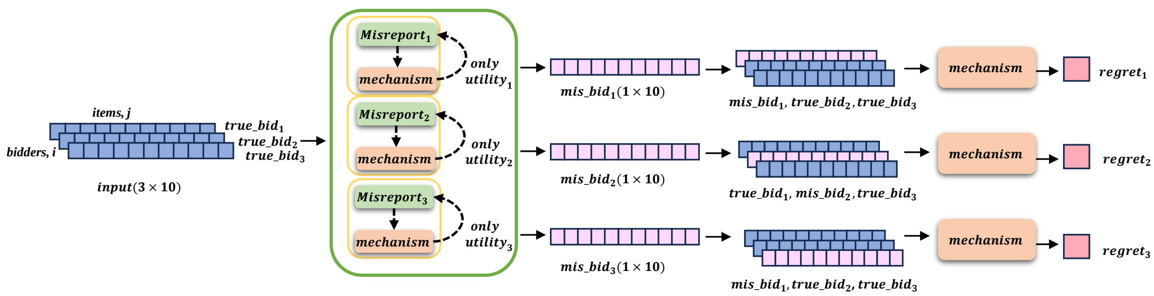

Figure 1: Multi-MisreportNet Architecture.

We introduce an intuitive framework, as illustrated in Figure 1, that can be easily utilized in various settings. We proceed by initializing $n-$misreport networks, each associated with a bidder's misreport strategy. The input consists of an $n \times m$ bidding matrix, while the output is a $1 \times m$ misreport vector specific to each bidder's misreport network. Note that the mechanism used to train the misreport in the figure is fixed. Our objective is to determine the optimal misreport for each individual. During the training phase of each misreport network, our focus is solely on optimizing the corresponding bidder's bid to achieve their individual optimal bid, denoted as bidder$_i$. To create a new bidding matrix, we combine the true bids of the remaining bidders with the output of the optimal misreport. This matrix is subsequently injected into the auction mechanism used for testing, enabling us to perform allocation and payment calculations.

Finally, the individual regrets are calculated and subsequently presented, serving as a quantification of the post-regret induced by the optimal misreport for each bidder. These regrets elucidate the relative disparity between the bidder's optimal misreport and their true bid, thereby offering a relatively more accurate metric to gauge their retrospective remorse. To gain a deeper understanding of how multiple networks operate, we show the pseudo-code of the training part in the Appendix B.

4. Experiments

To evaluate the robustness of the auction mechanism and the efficiency of our proposed multi-MisreportNet, we conduct experiments in different settings within the ALGNet framework. We denote the auction setting as $n \times m$, with n and m representing the number of bidders and items, correspondingly. See Appendix C for specific environment setup details. We demonstrate in Appendix D that the auction mechanism of ALGNet performs adequately for simple auction environments on a smaller scale. Therefore, we focus our experiments on more complex auction settings, such as 3×10 and 5×10. In this section, we offer a detailed analysis of the experimental outcomes, which demonstrate that the MisreportNet can detect misreport more effectively.

In Table 2, we present the outcomes of the regret examination for the multi-MisreportNet with configurations 3×10 and 5×10. Additionally, we contrast the outcomes of the single-Misreport and ALGNet (where ALGNet is considered as the baseline).

Table 2: Results of Multi-MisreportNet (Multi) under ALGNet Testing

Setting	**Bidder**	**Multi regret**	**Single regret**	**ALGNet regret**
3×10	bidder$_1$	0.00389 ± 0.0003	0.00263 ± 0.0002	0.00154 ± 0.0002
	bidder$_2$	0.00398 ± 0.0004		
	bidder$_3$	0.00439 ± 0.0004		
5×10	bidder$_1$	0.00478 ± 0.0002	0.00448 ± 0.0003	0.00261 ± 0.0004
	bidder$_2$	0.00483 ± 0.0003		
	bidder$_3$	0.00476 ± 0.0003		
	bidder$_4$	0.00489 ± 0.0003		
	bidder$_5$	0.00484 ± 0.0003		

Compared to the single-misreport method, it appears that designing a network for each person leads to a larger difference in regret and is therefore more aligned with the actual auction conditions. While the disparity between single and multi of 5×10 is relatively minor than that of 3×10, it may be explained by the intricacies involved with auction models on a larger scale. Consequently, the ALGNet mechanism shows less efficacy than what is observed in the 3×10 setting.

We also show the outcomes in the asymmetric auction setting in Table 3. Our results show that Multi-MisreportNer enhances the identification of disparities in the distribution of individuals, leading to more precise findings.

Table 3: Comparison between Multi-MisreportNet and ALGNet in 3×10 asymmetric auction setting

Setting	**Bidder**	**Distribution**	**Multi regret**	**ALGNet regret**
3×10	bidder$_1$	$v_1 \sim U[3, 6]$	0.09260 ± 0.0186	0.03611 ± 0.0069
	bidder$_2$	$v_2 \sim U[4, 7]$	0.01411 ± 0.0094	0.00081 ± 0.0015
	bidder$_3$	$v_3 \sim U[5, 8]$	0.12206 ± 0.0375	0.03768 ± 0.0190

5. Conclusion and future work

This paper proposes to introduce a separately trained misreport network for mechanism testing to detect better misreports so as to obtain a more realistic regret value. We also employed a multi-MisreportNet approach to handle bidder heterogeneity, which resulted in increased and more accurate regret value. In future work, we aim to integrate multi-MisreportNet into mechanism training to address optimal asymmetric auction mechanism design and evaluate its effectiveness in asymmetric auction design.

References

[1] Gianluca Brero, Benjamin Lubin, and Sven Seuken. Machine learning-powered iterative combinatorial auctions. *arXiv preprint arXiv:1911.08042*, 2019.

[2] Yang Cai, Constantinos Daskalakis, and S Matthew Weinberg. An algorithmic characterization of multi-dimensional mechanisms. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 459–478, 2012.

[3] Yang Cai, Constantinos Daskalakis, and S Matthew Weinberg. Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization. In *2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*, pages 130–139. IEEE, 2012.

[4] Vincent Conitzer and Tuomas Sandholm. Complexity of mechanism design. *arXiv preprint cs/0205075*, 2002.

[5] Michael Curry, Ping-Yeh Chiang, Tom Goldstein, and John Dickerson. Certifying strategyproof auction networks. *Advances in Neural Information Processing Systems*, 33:4987–4998, 2020.

[6] Michael Curry, Tuomas Sandholm, and John Dickerson. Differentiable economics for randomized affine maximizer auctions. *arXiv preprint arXiv:2202.02872*, 2022.

[7] Zhijian Duan, Jingwu Tang, Yutong Yin, Zhe Feng, Xiang Yan, Manzil Zaheer, and Xiaotie Deng. A context-integrated transformer-based neural network for auction design. In *International Conference on Machine Learning*, pages 5609–5626. PMLR, 2022.

[8] Paul Dütting, Zhe Feng, Harikrishna Narasimhan, David Parkes, and Sai Srivatsa Ravindranath. Optimal auctions through deep learning. In *International Conference on Machine Learning*, pages 1706–1715. PMLR, 2019.

[9] Zhe Feng, Harikrishna Narasimhan, and David C Parkes. Deep learning for revenue-optimal auctions with budgets. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems*, pages 354–362, 2018.

[10] Sergiu Hart and Noam Nisan. Approximate revenue maximization with multiple items. *Journal of Economic Theory*, 172:313–347, 2017.

[11] Jason D Hartline and Tim Roughgarden. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM conference on Electronic commerce*, pages 225–234, 2009.

[12] Dmitry Ivanov, Iskander Safiulin, Igor Filippov, and Ksenia Balabaeva. Optimal-er auctions through attention. *Advances in Neural Information Processing Systems*, 35:34734–34747, 2022.

[13] Kevin Kuo, Anthony Ostuni, Elizabeth Horishny, Michael J Curry, Samuel Dooley, Ping-yeh Chiang, Tom Goldstein, and John P Dickerson. Proportionnet: Balancing fairness and revenue for auction design with deep learning. *arXiv preprint arXiv:2010.06398*, 2020.

[14] Xinye Li and Andrew Chi-Chih Yao. On revenue maximization for selling multiple independently distributed items. *Proceedings of the National Academy of Sciences*, 110(28): 11232–11237, 2013.

[15] Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.

[16] Neehar Peri, Michael Curry, Samuel Dooley, and John Dickerson. Preferencenet: Encoding human preferences in auction design with deep learning. *Advances in Neural Information Processing Systems*, 34:17532–17542, 2021.

[17] Tian Qin, Fengxiang He, Dingfeng Shi, Wenbing Huang, and Dacheng Tao. Benefits of permutation-equivariance in auction mechanisms. *Advances in Neural Information Processing Systems*, 35:18131–18142, 2022.

[18] Jad Rahme, Samy Jelassi, and S Matthew Weinberg. Auction learning as a two-player game. *arXiv preprint arXiv:2006.05684*, 2020.

[19] Jad Rahme, Samy Jelassi, Joan Bruna, and S Matthew Weinberg. A permutation-equivariant neural network architecture for auction design. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pages 5664–5672, 2021.

[20] Tuomas Sandholm and Anton Likhodedov. Automated design of revenue-maximizing combinatorial auctions. *Operations Research*, 63(5):1000–1025, 2015.

[21] Pingzhong Tang and Zihe Wang. Optimal auctions for negatively correlated items. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 103–120, 2016.

[22] Jakob Weissteiner and Sven Seuken. Deep learning—powered iterative combinatorial auctions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 2284–2293, 2020.

[23] Andrew Chi-Chih Yao. An n-to-1 bidder reduction for multi-item auctions and its applications. In *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*, pages 92–109. SIAM, 2014.

[24] Andrew Chi-Chih Yao. On solutions for the maximum revenue multi-item auction under dominant-strategy and bayesian implementations. *arXiv preprint arXiv:1607.03685*, 2016.

Appendix A. Related Work

Designing the optimal mechanism under varying scenarios is a fundamental research area in economics. Initial approaches conceptualized the problem as a linear program — a strategy encumbered with scalability issues due to the substantial data prerequisites and time consumption. An alternate solution was to explore within a specific family of DSIC (Dominant Strategy Incentive Compatible) auctions, such as the Virtual Valuation Combinatorial Auctions (VVCAs) family[20]. However, this approach also comes with certain constraints. The restriction to a specific family of auctions subtly introduces the threat that the optimal mechanism may not be encapsulated within this family, thereby deeming the pursuit of achieving the optimal mechanism as unattainable.

The design of auctions gradually began to incorporate an assortment of traditional machine-learning algorithms. However, the continual evolution and refinement of deep learning brought about a realization: these machine learning algorithms were seen as less adaptive, less diverse, and less proficient than the solutions facilitated by deep learning. As a result, RegretNet capitalized on neural networks and pursued near-optimal auction mechanisms, which marked an innovative milestone in the field. This methodology fused the auction's assignment and payment rules into unique allocation networks and payment networks, respectively. The results derived from RegretNet's model effectively offered an approximation that was in close proximity to the performance of all previously deemed optimal auctions. Such a groundbreaking approach wielded substantial influence, acting as a cornerstone to guide the design blueprint for numerous ensuing studies.

ALGNet, an advanced development over the pre-existing RegretNet, introduces a pioneering evaluation criterion in auction mechanisms- the square root disparity between anticipated revenue and expected regret ($\sqrt{P} - \sqrt{R}$). This novel concept mitigates the use of time-bound and hyperparameter-sensitive loss functions, thereby alleviating the typically resource-intensive quest for the perfect hyperparameters, an issue commonly observed with RegretNet.In parallel with this, ALGNet implements a misreporting network, analogizing the training process to a strategic game between two players. The aim is to optimize the misreport throughout the training process in such a manner that the aggregate benefit is maximized. The significant limitation of this approach, however, is the potential for the identified misreports to plateau at local optima.

Recognizing this limitation, our proposed research centers around undertaking a deeper dive into the misreporting network. Specifically, we seek to better understand its structure and function, with the ultimate goal of enhancing its performance. This, we believe, will help unlock the full potential of ALGNet in yielding a more robust, reliable, and advantageous auction mechanism.

Following the development of RegretNet, numerous subsequent works have made valuable contributions. One notable example is RegretFormer [12], in which the authors introduced a novel network structure encompassing attention layers and a new loss function capable of predefining a regret budget. Additionally, You and colleagues have incorporated two performance-enhancing generic methods into the deep learning auction mechanism, particularly the integration with RegretFormer, which resulted in outcomes reaching the state of the art.

Beyond the aforementioned studies, various unique angles have been explored within the field of auction mechanism design. Kuo et al.[13] proposed ProportionNet, introducing the concept of fairness as a solution to maintain high revenue and ensure strong incentives. Similarly, Peri et al.[16] presented PreferenceNet, which incorporated human preferences into auction design, creating a relatable example through expected allocation encoding. Duan et al.[7] transcended the narrow confines of earlier studies by integrating common contextual information about bidders and projects within the auction learning framework. Feng et al.[9] expanded upon the concept of RegretNet to accommodate private budget constraints and facilitate Bayesian incentive compatibility, yielding auction designs closely aligned with incentive compatibility principles. Curry et al.[6] revisited the learning challenge of affine maximization auctions, introducing an architectural system that supports multiple bidders with perfect honesty. Rahme et al.[19] analyzed the auction design challenge from the perspective of permutation equivariant symmetry, devising a neural network capable of flawlessly retrieving the optimal permutation equivariant mechanism and demonstrating its superior generalization abilities. In a further study, Curry et al.[5] tested the efficacy of application-specific methods using neural networks to explicitly verify the strategyproofness within specific valuation profiles. Their modifications to the RegretNet architecture paved the way for more precise regret values. Lastly, Qin et al.[17] expanded upon the benefits of permutation equivariant neural networks within additive valuation and symmetric valuation settings, explained through rigorous proofs.

Appendix B. Training Algorithm

Algorithm 1: Multi-Misreport Net training

Input: n - number of agents, m - number of objects, B - batch size
Output: $V'[n, m]$ - Best misreport bid
Parameter: R - number of iterations for misreport network training, α - learning rate
Initialize: $M[1 \text{ to } n]$ - different misreport nets for each bidder, A - auction mechanism

1. **for** $k = 1, ..., R$ **do**
2. Sample valuation batch S of size B
3. **for** $i = 1, ..., n$ **do**
4. $v_i' = M_i(S)$
5. $u_i =$ Utility(A, V_i, V_i')
6. $\varphi_i^{k+1} \leftarrow \varphi_i^k - \alpha \nabla_{\varphi_i} \mathcal{L}\left(\varphi_i^k\right)(u_i)$
7. **end for**
8. **end for**
9. **return** $V' = \{v_1', v_2', ...v_n'\}$

The method of Multi-MisreportNet is summarized in Algorithm 1. We can employ this algorithm to evaluate any existing auction mechanism A, which in this study refers to the ALGNet model. The algorithm takes as input the number of agents n, the number of objects m, and the batch size B. The output is the combination of the optimal misreport for each bidder, i.e. $V'[n, m]$. The algorithm iterates over the training process R times (line 1). At each iteration, the algorithm samples a valuation batch S of size B (line 2). The algorithm then trains each misreportNet by computing the utility of the auction mechanism with the true value estimates and the misreported value estimates (lines 3-5). The gradient of the loss function concerning the weights of each misreport net is computed using backpropagation. Gradient descent is then used to update the weights of each misreport net (line 6). The algorithm iterates over the training process R times to improve the quality of the misreport nets. The proposed MMN algorithm has the potential to improve the efficiency and fairness of auction mechanisms and can be applied to a variety of auction settings.

Appendix C. Setup

We focused our experimental efforts on the 3×10 and 5×10 auction settings. To implement the misreport network framework, we utilize the PyTorch deep learning tool, which is appropriate for complex and computationally intensive tasks.

We conducted our experiments on a server equipped with an Intel Xeon processor and an NVIDIA GPU that boasts 32GB of memory. The misreportNet we employed comprises seven layers, with each layer containing 100 hidden nodes. We utilized the tanh activation function at the hidden nodes. During the training phase, we employed a batch size of 200 and conducted 300 iterations, while testing the network on 20,000 samples. In each batch of data, we optimized the parameters using the Adam optimizer with a learning rate of 0.001. The bidders' valuations were sampled from a uniform distribution between 0 and 1, i.e., $v_{ij} \in [0, 1]$.

Appendix D. Other results

Table 4: Single-MisreportNet Results in 1×2 and 1×10 Auction Environments

Setting	**Single regret**	**baseline**
1×2	$0.00006 \pm 1.7080e - 06$	$0.00012 \pm 5.6765e - 06$
1×10	$0.00257 \pm 2.7801e - 05$	$0.00186 \pm 7.3553e - 05$

Through experiments conducted in multiple environments, we have observed a more pronounced evaluation of regret by the neural network when optimizing for misreports in scenarios characterized by a larger number of items. This observation can be attributed to the augmented dimensionality of features stemming from an augmented quantity of items, thereby endowing the neural network with an enhanced and more elaborate information repository. Consequently, the neural network can effectively capture and leverage varied feature combinations and correlations, enhancing its understanding and predictive capabilities regarding participants' behavior.

Moreover, in scenarios with a larger number of items, participants exhibit a wider array of bidding strategies, necessitating the neural network to better comprehend and predict the diverse strategies and behaviors employed by different participants. This heightened complexity of decision-making contributes to improved performance in the neural network's decision-making process.

Table 5: Single-MisreportNet and Multi-MisreportNet Results in 2×2 Auction Environments

Bidder	**Multi regret**	**Single regret**	**baseline**
bidder$_1$	$0.00075 \pm 1.0045e - 05$	$0.00066 \pm 6.8621e - 06$	$0.00064 \pm 2.6137e - 05$
bidder$_2$	$0.00061 \pm 5.1336e - 06$		

Additionally, we speculate that in the 1×2 and 2×2 environments, the mechanism itself may have already been trained to a near-optimal state, thereby resulting in less discernible effects in the misreport testing.