Bridging the Gap Between Estimated and True Regret in Deep Learning–Based Auction Mechanisms

Extended Abstract

Shuyuan You
Griffith University
Brisbane, Australia
shuyuan.you@griffithuni.edu.au

Zhiqiang Zhuang
Qiannan Normal University for Nationalities
Duyun, China
zoldlady@gmail.com

Haiying Wu
Tianjin University
Tianjin, China
haiyingwu@tju.edu.cn

Kewen Wang
Griffith University
Brisbane, Australia
k.wang@griffith.edu.au

Zhe Wang
Griffith University
Brisbane, Australia
zhe.wang@griffith.edu.au

ABSTRACT

Recent advances, such as RegretNet, ALGnet, RegretFormer, and CITransNet, use deep learning to approximate optimal multi-item auctions and measure their violation of incentive compatibility via ex-post regret. However, the accuracy of these regret estimates remains unclear. We show that existing methods systematically underestimate actual regret due to hyperparameter sensitivity and the non-convexity of the optimization landscape. In some models, the true regret is found to be orders of magnitude larger than reported, leading to overstated claims of IC and revenue. To address this, we derive a theoretical lower bound on regret and introduce an efficient item-wise regret approximation. Building on this, we propose Item-wise Guided Gradient Refinement that substantially improves regret estimation accuracy while reducing computational cost. Our method provides a more reliable foundation for evaluating incentive compatibility in deep learning-based auction mechanisms and highlights the need to reassess prior performance claims in this area.

KEYWORDS

Mechanism Design; Deep Learning; Incentive Compatibility

ACM Reference Format:
Shuyuan You, Zhiqiang Zhuang, Haiying Wu, Kewen Wang, and Zhe Wang. 2026. Bridging the Gap Between Estimated and True Regret in Deep Learning–Based Auction Mechanisms: Extended Abstract. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 3 pages. https://doi.org/10.65109/WTCL6718

1 INTRODUCTION

Optimal auction design is a central problem in algorithmic game theory with substantial practical relevance[4, 10]. While Myerson [11] characterized the optimal auction for the single-item setting, the problem becomes substantially more complex for multiple distinct items. Despite extensive research [1–3, 8, 13, 19, 20], the exact optimal mechanism remains unknown even for relatively simple settings. Recent work has turned to deep learning-based approaches, such as RegretNet [7], which parameterize auction mechanisms using neural networks. These models relax the IC constraint and quantify violations via ex-post regret, aiming to achieve higher revenue than traditional methods [12] while keeping regret low. And lots of work [5–7, 9, 14–18, 22] follows RegretNet.

However, a critical yet underexplored question remains regarding whether these models accurately estimate incentive-compatibility violations. From a theoretical perspective, computing a bidder's optimal ex-post regret requires searching over the entire combinatorial valuation space, a task that suffers from the curse of dimensionality. Existing models relying on gradient-based optimization are highly sensitive to hyperparameters. Our empirical analysis reveals that suboptimal configurations often lead to systematic underestimation of regret.

To address the critical challenge of regret underestimation, we develop a theoretically grounded and computationally efficient evaluation framework. We first derive a mathematical lower bound on the optimal regret and introduce Item-wise Regret as a novel proxy metric. Building on this, we introduce an Item-wise Guided Gradient Refinement strategy. Our method uses discrete item-wise solutions to construct a structured set of initial points for gradient-based optimization. This design bridges discrete grid search and continuous optimization, substantially improving the optimizer's ability to converge to solutions closer to the global optimum. Our framework not only improves regret estimation accuracy but also achieves orders-of-magnitude efficiency gains compared to exhaustive random search methods. The full paper can be seen in [21].

2 REGRET UNDERESTIMATION

We scrutinized the evaluation protocols of leading models, revealing a significant lack of consensus (e.g., initialization counts L range from 1 to 1000). To test reliability, we cross-evaluated RegretNet, ALGnet, and RegretFormer under swapped hyperparameter settings.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). https://doi.org/10.65109/WTCL6718

As shown in Table 1, default protocols systematically underestimate violations. Most notably, RegretFormer exhibits extreme fragility: under ALGnet's settings, its detected regret explodes by 28-fold ($5.17 \times 10^{-3} \rightarrow 144.02 \times 10^{-3}$), suggesting its reported performance relies on undetected violations. Furthermore, accurate convergence is prohibitively expensive. In 5×10 settings, finding the true global maximum requires 216.9 hours, while standard fast checks (3 hours) underestimate violations by 35 times.

Table 1: Cross-evaluation of Regret ($\times 10^{-3}$) on 3×10

	Hyperparameter Source		
Model	RegretNet	ALGnet	RegretFormer
RegretNet	**1.62**	0.40	1.55
ALGnet	**4.78**	1.56	3.89
RegretFormer	10.59	**144.02**	5.17

3 METHODOLOGY

To resolve the trade-off between accuracy and efficiency, we propose a framework combining theoretical bounds with guided optimization.

3.1 Lower Bound and Item-wise Regret

Exact regret computation involves searching the joint valuation space with complexity $O(n \cdot Q^m)$. To bypass this, we decompose it. We define $Rgt_{i,j}$ as the regret of bidder i restricted to item j.

Theorem 3.1. *For any bidder i, the ex-post regret is lower-bounded by the maximum single-item regret:* $Rgt_i \geq \max_{j \in M} Rgt_{i,j}$.

We further define **Item-wise Regret** as $\sum_{j \in M} Rgt_{i,j}$. This proxy reduces computational complexity to $O(n \cdot Q \cdot m)$ (linear in items). While not a strict upper bound in pathological cases, it serves as an efficient, high-quality heuristic for the true regret landscape.

3.2 Item-wise Guided Gradient Refinement

Standard gradient methods fail because random initializations struggle to navigate the non-convex loss landscape. We propose a hybrid strategy that uses item-wise discrete solutions to construct a structured initialization set $\mathcal{B}_{init}$ for the gradient optimizer. $\mathcal{B}_{init}$ comprises five strategic components:

(1) **Combinatorial Candidate (1):** Aggregating optimal item-wise misreports to capture basic multi-item synergies.
(2) **Single-Item Candidates (m):** Vectors optimizing for a single item j while keeping other dimensions truthful, ensuring local optima are covered.
(3) **Perturbed Combinatorial (k):** Adding Gaussian noise (σ) to the combinatorial candidate to explore its neighborhood.
(4) **Perturbed Truthful (k):** Adding Gaussian noise (σ) to the truthful profile to detect violations near truth-telling.
(5) **Global Random (k):** A safety buffer of uniform random points.

This structured initialization bridges the gap between discrete grid search and continuous optimization, guiding the optimizer toward the global maximum with significantly fewer iterations than random restarts.

4 EXPERIMENTS

We evaluate our framework on ALGnet and re-examine SOTA models (RegretNet, ALGnet, RegretFormer) in n-bidder m-item additive settings ($v_i \sim U[0,1]$). We used strict noise settings for RegretFormer ($k = 80, \sigma = 0.6$) to navigate its complex landscape, while others utilized deterministic priors ($k = 0$).

4.1 Validating on ALGnet

Table 2 compares our method against baselines. In small settings (2×2) where exact optima are computable, our method matches the ground truth (0.71 vs 0.72×10^{-3}) with negligible cost. In large settings (5×10), converging to true regret via random restarts requires **216.9 hours**. Our method achieves comparable fidelity (5.25×10^{-3}) in just **1.19 hours**, delivering a $\sim 200\times$ speedup.

Table 2: Validation on ALGnet: Accuracy & Runtime

Setting	Metric	Regret (10^{-3})	Time (h)
2×2	Optimal (Ground Truth)	0.72	25.5
	Original Report	0.71	0.85
	Ours (Refined)	**0.71**	**0.025**
5×10	Near Optimal (Intensive)	5.24	216.9
	Original Report	2.48	9.2
	Ours (Refined)	**5.25**	**1.19**

4.2 Re-evaluating SOTA Models

We applied our framework to re-evaluate leading models (Table 3). **RegretNet** demonstrates high robustness, with our estimates matching reported values. **ALGnet** shows moderate underestimation ($\sim 3\times$ in 3×10). Most critically, **RegretFormer** exhibits catastrophic hidden violations. In the 3×10 setting, we detect a regret of 365.80×10^{-3}, which is over **70 times higher** than the reported 5.17. This suggests its high revenue claims are likely artifacts of failing to satisfy incentive compatibility.

Table 3: Re-evaluation of Regret ($\times 10^{-3}$)

Setting	Model	Reported	Ours	Revenue
3×10	RegretNet	1.62	**1.62**	5.54
	ALGnet	1.56	**4.73**	5.55
	RegretFormer	5.17	**365.80**	6.18
5×10	RegretNet	4.72	**5.25**	6.76
	ALGnet	2.48	**5.19**	6.74

5 CONCLUSION

We demonstrated that standard evaluation protocols for deep auction mechanisms systematically underestimate regret, masking potential incentive compatibility violations. We introduced an efficient Item-wise Guided Gradient Refinement framework that provides accurate regret estimates at a fraction of the computational cost of random restarts. Our findings highlight the urgent need to reassess performance claims in the field and offer a rigorous tool for future mechanism design.

REFERENCES

[1] Saeed Alaei, Hu Fu, Nima Haghpanah, Jason D. Hartline, and Azarakhsh Malekian. 2012. Bayesian optimal auctions via multi- to single-agent reduction. In *Proceedings of the 13th ACM Conference on Electronic Commerce, EC 2012, Valencia, Spain, June 4-8, 2012*. 17.

[2] Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S Matthew Weinberg. 2020. A simple and approximately optimal mechanism for an additive buyer. *Journal of the ACM (JACM)* 67, 4 (2020), 1–40.

[3] Yang Cai and Mingfei Zhao. 2017. Simple mechanisms for subadditive buyers via duality. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*. 170–183.

[4] Peter Cramton. 2013. Spectrum auction design. *Review of industrial organization* 42 (2013), 161–190.

[5] Michael Curry, Tuomas Sandholm, and John Dickerson. 2022. Differentiable Economics for Randomized Affine Maximizer Auctions. *arXiv preprint arXiv:2202.02872* (2022).

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

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

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

[9] Da Huo, Zhenzhe Zheng, and Fan Wu. 2025. Learning Optimal Auctions with Correlated Value Distributions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 39. 13944–13952.

[10] Xiangyu Liu, Chuan Yu, Zhilin Zhang, Zhenzhe Zheng, Yu Rong, Hongtao Lv, Da Huo, Yiqing Wang, Dagui Chen, Jian Xu, et al. 2021. Neural auction: End-to-end learning of auction mechanisms for e-commerce advertising. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 3354–3364.

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

[12] Thomas R Palfrey. 1983. Bundling decisions by a multiproduct monopolist with incomplete information. *Econometrica: Journal of the Econometric Society* (1983), 463–483.

[13] Gregory Pavlov. 2011. Optimal mechanism for selling two goods. *The BE Journal of Theoretical Economics* 11, 1 (2011).

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

[15] Jad Rahme, Samy Jelassi, Joan Bruna, and S Matthew Weinberg. 2021. A permutation-equivariant neural network architecture for auction design. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 5664–5672.

[16] Jad Rahme, Samy Jelassi, and S. Matthew Weinberg. 2021. Auction Learning as a Two-Player Game. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*.

[17] Weiran Shen, Pingzhong Tang, and Song Zuo. 2019. Automated Mechanism Design via Neural Networks. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*. 215–223.

[18] Tonghan Wang, Yanchen Jiang, and David C Parkes. 2024. GemNet: Menu-Based, strategy-proof multi-bidder auctions through deep learning. *arXiv preprint arXiv:2406.07428* (2024).

[19] Zihe Wang and Pingzhong Tang. 2014. Optimal mechanisms with simple menus. In *Proceedings of the fifteenth ACM conference on Economics and computation*. 227–240.

[20] Andrew Chi-Chih Yao. 2017. Dominant-strategy versus bayesian multi-item auctions: Maximum revenue determination and comparison. In *Proceedings of the 2017 ACM Conference on Economics and Computation*. 3–20.

[21] Shuyuan You, Zhiqiang Zhuang, Kewen Wang, and Zhe Wang. 2026. Bridging the gap between estimated and true regret towards reliable regret estimation in deep learning based mechanism design. *arXiv preprint arXiv:2601.13489* (2026).

[22] Shuyuan You, Zhiqiang Zhuang, Haiying Wu, Kewen Wang, and Zhe Wang. 2023. Improving Deep Learning Powered Auction Design. In *International Conference on Neural Information Processing*. Springer, 556–569.