
A Context-Integrated Transformer-Based Neural Network for Auction Design

Zhijian Duan¹ Jingwu Tang¹ Yutong Yin¹ Zhe Feng² Xiang Yan³ Manzil Zaheer⁴ Xiaotie Deng¹

Abstract

One of the central problems in auction design is developing an incentive-compatible mechanism that maximizes the auctioneer’s expected revenue. While theoretical approaches have encountered bottlenecks in multi-item auctions, recently, there has been much progress on finding the optimal mechanism through deep learning. However, these works either focus on a fixed set of bidders and items, or restrict the auction to be symmetric. In this work, we overcome such limitations by factoring *public* contextual information of bidders and items into the auction learning framework. We propose CITransNet, a context-integrated transformer-based neural network for optimal auction design, which maintains permutation-equivariance over bids and contexts while being able to find asymmetric solutions. We show by extensive experiments that CITransNet can recover the known optimal solutions in single-item settings, outperform strong baselines in multi-item auctions, and generalize well to cases other than those in training.

1. Introduction

Auction design is a classical problem in computational economics, with many applications on sponsored search (Jansen & Mullen, 2008), resource allocation (Huang et al., 2008) and blockchain (Galal & Youssef, 2018). Designing an incentive-compatible mechanism that maximizes the auctioneer’s expected revenue is one of the central topics in auction design. The seminal work by Myerson (1981) provides an optimal auction design for the single-item setting; however, designing a revenue-optimal auction is still not fully understood even for two bidders and two items setting

¹Peking University, Beijing, China ²Google Research, Mountain View, US ³Shanghai Jiao Tong University, Shanghai, China ⁴Google DeepMind, Mountain View, US. Correspondence to: Zhijian Duan <zjduan@pku.edu.cn>, Xiaotie Deng <xiaotie@pku.edu.cn>.

after four decades (Dütting et al., 2019).

Recently, pioneered by Dütting et al. (2019), there is rapid progress on finding (approximate) optimal auction through deep learning, e.g., (Shen et al., 2019; Luong et al., 2018; Tacchetti et al., 2019; Nedelec et al., 2021; Shen et al., 2020; Brero et al., 2021; Liu et al., 2021). Typically, we can formulate auction design as a constrained optimization problem and find near-optimal solutions using standard machine learning pipelines. However, existing methods only consider simple settings: they either focus on a fixed set of bidders and items, e.g. (Dütting et al., 2019; Rahme et al., 2021b) or ignore the identity of bidders and items so that the auction is restricted to be symmetric (Rahme et al., 2021a). As a comparison, in practice, auctions are much more complex beyond the aforementioned simple settings. For instance, in e-commerce advertising, there are a large number of bidders and items (i.e., ad slots) with various features (Liu et al., 2021), and each auction involves a different number of bidders and items. To handle such a practical problem, we need a new architecture that can incorporate public features and take a different number of bidders and items as inputs.

Main Contributions. In this paper, we consider contextual auction design, in which each bidder or item is equipped with context. In contextual auctions, the bidder-contexts and item-contexts can characterize various bidders and items to some extent, making the auctions close to those in practice. We formulate the contextual auction design as a learning problem and extend the learning framework proposed in Dütting et al. (2019) to our setting. Furthermore, we present a sample complexity result to bound the generalization error of the learned mechanism.

To overcome the aforementioned limitations of the previous works, we propose CITransNet: a Context-Integrated Transformer-based neural Network architecture as the parameterized mechanism to be optimized. CITransNet incorporates the bidding profile along with the bidder-contexts and item-contexts to develop an auction mechanism. It is built upon the transformer architecture (Vaswani et al., 2017), which can capture the complex mutual influence among different bidders and items in an auction. As a result, CITransNet is permutation-equivariant (Rahme et al., 2021a) over bids and contexts, i.e., any permutation of bidders (or items) in the bidding profile and bidder-contexts (or

item-contexts) would cause the same permutation of auction result (We will provide a formal definition in Remark 3.1). Moreover, in CITransNet, the number of parameters does not depend on the auction scale (i.e., the number of bidders and items), which brings CITransNet the potential of generalizing to auctions with various bidders or items, which we denote as *out-of-setting generalization*.

We show by extensive experiments that CITransNet can almost reach the same result as Myerson (1981) in single-item auctions and can obtain better performance in complex multi-item auctions compared to those strong baseline algorithms we use. Additionally, we also justify its out-of-setting generalization ability. Experimental results demonstrate that, under the same contextual setting, CITransNet can still perform well in auctions with a different number of bidders or items than those in training.

Further Related Work. As discussed before, it is an intricate task to design optimal auctions for multiple bidders and multiple items. Many previous works focus on special cases (to name a few, Manelli & Vincent (2006); Pavlov (2011); Giannakopoulos & Koutsoupias (2014); Yao (2017); Daskalakis et al. (2017); Haghpanah & Hartline (2021)) and the algorithmic characterization of optimal auction (e.g., Chawla et al. (2010); Cai et al. (2012); Babaioff et al. (2014); Yao (2014); Cai & Zhao (2017); Hart & Nisan (2017)). In addition, machine learning has also been applied to find approximate solutions for multiple items settings (Balcan et al., 2008; Lahaie, 2011; Dütting et al., 2015), and there are also many works analyzing the sample complexity of designing optimal auctions (Cole & Roughgarden, 2014; Devanur et al., 2016; Balcan et al., 2016; Guo et al., 2019; Gonczarowski & Weinberg, 2021). In our paper, we follow the paradigm of *automated mechanism design* (Conitzer & Sandholm, 2002; 2004; Sandholm & Likhodedov, 2015).

Dütting et al. (2019) propose the first neural network framework, RegretNet, to automatically design optimal auctions for general multiple bidders and multiple items settings by modeling an auction as a multi-layer neural network and using standard machine learning pipelines. Feng et al. (2018) and Golowich et al. (2018) modify RegretNet to handle different constraints and objectives. Curry et al. (2020) extend RegretNet to be able to verify strategyproofness of the auction mechanism learned by neural network. ALGNet (Rahme et al., 2021b) models the auction design problem as a two-player game through parameterizing the misreporter as well. PreferenceNet (Peri et al., 2021) encodes human preference (e.g. fairness) into RegretNet. Rahme et al. (2021a) propose a permutation-equivariant architecture called EquivariantNet to design *symmetric* auctions, a special case that is anonymous (bidder-symmetric) and item-symmetric. In contrast, we study optimal contextual auction design, and our proposed CITransNet is permutation-

equivariant while not restricted to symmetric auctions.

Existing literatures of contextual auction mainly discuss the online setting of some *known* contextual repeated auctions, e.g., *posted-price auctions* (Amin et al., 2014; Mao et al., 2018; Drutsa, 2020; Zhiyanov & Drutsa, 2020), in which at every round the item is priced by the seller to sell to a strategic buyer, and *second price auctions* (Golrezaei et al., 2021). As a comparison, we consider the offline setting of contextual *sealed-bid auction*. We learn the mechanism from historical data and optimize the expected revenue for the auctioneer. Besides, we do not assume the conditional distribution of the bidder’s valuation when given both the bidder-context and item-context.

Organization. This paper is organized as follows: In Section 2 we introduce contextual auction design, model the problem as a learning problem and derive a sample complexity for it; In Section 3 we present the structure of CITransNet, along with the training and optimization procedure; We conduct experiments in Section 4 and draw the conclusion in Section 5.

2. Contextual Auction Design

In this section, we set up the problem of contextual auction design. Then, we extend the learning framework proposed by Dütting et al. (2019) to our contextual setting.

2.1. Contextual Auction

We consider a contextual auction with n bidders $N = \{1, 2, \dots, n\}$ and m items $M = \{1, 2, \dots, m\}$. Each bidder $i \in N$ is equipped with bidder-context $x_i \in \mathcal{X} \subset \mathbb{R}^{d_x}$ and each item $j \in M$ is equipped with item-context $y_j \in \mathcal{Y} \subset \mathbb{R}^{d_y}$, in which d_x and d_y are the dimensions of bidder-context variables and item-context variables, respectively. Denote $x = (x_1, x_2, \dots, x_n)$ as the bidder-contexts and $y = (y_1, y_2, \dots, y_m)$ as the item-contexts. x and y are sampled from underlying joint probability distribution $\mathcal{D}_{x,y}$. Let v_{ij} be the valuation of bidder i for item j . Conditioned on bidder-context x_i and item-context y_j , v_{ij} is sampled from a distribution $\mathcal{D}_{v_{ij}|x_i,y_j}$, i.e., the distribution of v_{ij} depends on both x_i and y_j .

The valuation profile $v = (v_{ij})_{i \in N, j \in M} \in \mathbb{R}^{n \times m}$ is unknown to the auctioneer, however, she knows the sampled bidder-contexts x and item-contexts y . In this paper, we only focus on *additive* valuation setting, i.e., the valuation of each bidder i for a set of items $S \subseteq M$ is the sum of valuation for each item $j \in S$: $v_{iS} = \sum_{j \in S} v_{ij}$. At an auction round, each bidder bids for each item. Given the bidding profile (or bids) $b = (b_{ij})_{i \in N, j \in M}$, the contextual auction mechanism is defined as follows:

Definition 2.1 (Contextual Auction Mechanism). A contextual auction mechanism (g, p) consists of an allocation rule

g and a payment rule p :

- The allocation rule $g = (g_{ij})_{i \in N, j \in M}$, in which $g_{ij} : \mathbb{R}^{n \times m} \times \mathcal{X}^n \times \mathcal{Y}^m \rightarrow [0, 1]$ computes the probability that item j is allocated to bidder i , given the bidding profile $b \in \mathbb{R}^{n \times m}$, bidder-contexts $x \in \mathcal{X}^n$ and item-contexts $y \in \mathcal{Y}^m$. For all b, x, y , and $j \in M$, we have $\sum_{i=1}^n g_{ij}(b, x, y) \leq 1$ to guarantee no item is allocated more than once.
- The payment rule $p = (p_1, p_2, \dots, p_n)$, in which $p_i : \mathbb{R}^{n \times m} \times \mathcal{X}^n \times \mathcal{Y}^m \rightarrow \mathbb{R}_{\geq 0}$ computes the price bidder i need to pay, given the bidding profile $b \in \mathbb{R}^{n \times m}$, bidder-contexts $x \in \mathcal{X}^n$ and item-contexts $y \in \mathcal{Y}^m$.

Define $\mathcal{V} = \mathcal{V}_1 \times \mathcal{V}_2 \times \dots \times \mathcal{V}_n$ be the joint valuation profile domain set, in which \mathcal{V}_i is the domain set of all the possible valuation profiles $v_i = (v_{i1}, v_{i2}, \dots, v_{im})$ of bidder i . Let $\mathcal{V}_{-i} = (\mathcal{V}_1, \dots, \mathcal{V}_{i-1}, \mathcal{V}_{i+1}, \dots, \mathcal{V}_n)$ be the joint valuation profile domain set except \mathcal{V}_i . Similarly, we denote $v_{-i} = (v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_n)$ and $b_{-i} = (b_1, \dots, b_{i-1}, b_{i+1}, \dots, b_n)$. Without loss of generality, we assume $b_i \in \mathcal{V}_i$ for all $i \in N$. Each bidder $i \in N$ aims to maximize her utility, defined as follows,

Definition 2.2 (Quasilinear utility). In an additive valuation auction setting, the utility of bidder i under mechanism (g, p) is defined by

$$u_i(v_i, b, x, y) = \sum_{j=1}^m g_{ij}(b, x, y) v_{ij} - p_i(b, x, y) \quad (1)$$

for all $v_i \in \mathcal{V}_i, b \in \mathcal{V}, x \in \mathcal{X}^n, y \in \mathcal{Y}^m$.

In this work, we want the auction mechanism to be *dominant strategy incentive compatible* (DSIC)¹, defined as below,

Definition 2.3 (DSIC). An auction (g, p) is *dominant strategy incentive compatible* (DSIC) if for each bidder, the optimal strategy is to report her true valuation no matter how others report. Formally, for each bidder $i \in N$, for all $x \in \mathcal{X}^n, y \in \mathcal{Y}^m$ and for arbitrary $b_{-i} \in \mathcal{V}_{-i}$, we have

$$u_i(v_i, (v_i, b_{-i}), x, y) \geq u_i(v_i, (b_i, b_{-i}), x, y), \quad (2)$$

for all $b_i \in \mathcal{V}_i$.

Besides, the auction mechanism needs to be *individually rational* (IR), defined as follows,

Definition 2.4 (IR). An auction (g, p) is *individually rational* (IR) if for each bidder, truthful bidding will receive a non-negative utility. Formally, for each bidder $i \in N$, for all $x \in \mathcal{X}^n, y \in \mathcal{Y}^m$ and for arbitrary $v_i \in \mathcal{V}_i, b_{-i} \in \mathcal{V}_{-i}$, we have

$$u_i(v_i, (v_i, b_{-i}), x, y) \geq 0. \quad (\text{IR})$$

¹There is another weaker notion of incentive compatibility, Bayesian incentive compatibility (BIC), in the literature. In practice, DSIC is more desirable than BIC. It doesn't require prior knowledge of the other bidders and is more robust. In this work, we only focus on DSIC, similar to Dütting et al. (2019).

In a DSIC and IR auction, rational bidders would truthfully report their valuations. Therefore, let $\mathcal{D}_{v,x,y}$ be the joint distribution of v, x and y , the expected revenue is:

$$rev := \mathbb{E}_{(v,x,y) \sim \mathcal{D}_{v,x,y}} \left[\sum_{i=1}^n p_i(v, x, y) \right]. \quad (3)$$

Optimal contextual auction design aims to find an auction mechanism that maximizes the expected revenue while satisfying the DSIC and IR conditions.

2.2. Contextual Auction Design as a Learning Problem

Similar to Dütting et al. (2019), we formalize the problem of optimal auction design as a learning problem. First, we define *ex-post regret*:

Definition 2.5 ((Ex-post) Regret). The ex-post regret for a bidder i under mechanism (g, p) is the maximum utility gain she can achieve by misreporting when the bids of others are fixed, i.e.,

$$rgt_i(v, x, y) := \max_{b_i \in \mathcal{V}_i} u_i(v_i, (b_i, v_{-i}), x, y) - u_i(v_i, v, x, y). \quad (4)$$

In particular, similar to Dütting et al. (2019), the DSIC condition is equivalent to $rgt_i(v, x, y) = 0, \forall i \in N, v \in \mathcal{V}, x \in \mathcal{X}^n, y \in \mathcal{Y}^m$. By assuming that $\mathcal{D}_{v,x,y}$ has full support on the space of (v, x, y) and recognizing that the regret is non-negative, an auction satisfies DSIC (except for measure zero events) if

$$\mathbb{E}_{(v,x,y) \sim \mathcal{D}_{v,x,y}} \left[\sum_{i=1}^n rgt_i(v, x, y) \right] = 0. \quad (\text{DSIC})$$

Let \mathcal{M} be the set of all the auction mechanisms that satisfy Equation (IR). By setting Equation (DSIC) as a constraint, we can formalize the problem of finding an optimal contextual auction as a constraint optimization:

$$\begin{aligned} \min_{(g,p) \in \mathcal{M}} & - \mathbb{E}_{(v,x,y) \sim \mathcal{D}_{v,x,y}} \left[\sum_{i=1}^n p_i(v, x, y) \right] \\ \text{s.t.} & \mathbb{E}_{(v,x,y) \sim \mathcal{D}_{v,x,y}} \left[\sum_{i=1}^n rgt_i(v, x, y) \right] = 0, \end{aligned} \quad (\text{I})$$

This optimization problem is generally intractable due to the intricate constraints². To handle such a problem, we parameterize the auction mechanism as (g^w, p^w) , where $w \in \mathbb{R}^{d_w}$ are the parameters (with dimension d_w) to be optimized.

²In the automated mechanism design literature (Conitzer & Sandholm, 2002; 2004), Equation (I) can be formulated as a linear programming. However, this LP is hard to solve in practice because of the exponential number of constraints, even for discrete value distribution settings.

All the expectation terms are computed empirically by L samples of (v, x, y) independently drawn from $\mathcal{D}_{v,x,y}$. The empirical ex-post regret for bidder i under parameters w is defined as

$$\widehat{rgt}_i(w) := \frac{1}{L} \sum_{\ell=1}^L rgt_i^w(v^{(\ell)}, x^{(\ell)}, y^{(\ell)}), \quad (5)$$

where $rgt_i^w(v, x, y)$ is computed based on the parameterized mechanism (g^w, p^w) . On top of that, the learning formulation of Equation (I) is

$$\begin{aligned} \min_{w \in \mathbb{R}^{d_w}} & - \frac{1}{L} \sum_{\ell=1}^L \sum_{i=1}^n p_i^w(v^{(\ell)}, x^{(\ell)}, y^{(\ell)}) \\ \text{s.t.} & \widehat{rgt}_i(w) = 0, \forall i \in N \end{aligned} \quad (II)$$

Equation (IR) can be satisfied through the architecture design. See Section 3.4 for the discussion.

2.3. Sample Complexity

We provide a sample complexity to bound the two gaps at the same time: the gap between empirical revenue and expected revenue, and the gap between empirical regret and expected regret. Such result justifies the feasibility to approximately solve Equation (I) by Equation (II).

For contextual auction mechanism class \mathcal{M} , similar to Dütting et al. (2019), we measure the capacity of \mathcal{M} via *covering numbers* (Shalev-Shwartz & Ben-David, 2014). We define the $\ell_{\infty,1}$ -distance between two auction mechanisms $(g, p), (g', p') \in \mathcal{M}$ as $\max_{v,x,y} \sum_{i \in N, j \in M} |g_{ij}(v, x, y) - g'_{ij}(v, x, y)| + \sum_{i \in N} |p_i(v, x, y) - p'_i(v, x, y)|$. For all $r > 0$, let $\mathcal{N}_{\infty,1}(\mathcal{M}, r)$ be the minimum number of balls with radius r that cover all the mechanisms in \mathcal{M} under $\ell_{\infty,1}$ -distance (called the r -covering number of \mathcal{M}). We have the following result:

Theorem 2.6. *For each bidder i , assume w.l.o.g. that the valuation function v_i satisfies $v_i(S) \leq 1, \forall S \subseteq M$. Fix $\delta, \epsilon \in (0, 1)$, for any $(g^w, p^w) \in \mathcal{M}$, when*

$$L \geq \frac{9n^2}{2\epsilon^2} \left(\ln \frac{4}{\delta} + \ln \mathcal{N}_{\infty,1}(\mathcal{M}, \frac{\epsilon}{6n}) \right), \quad (6)$$

with probability at least $1 - \delta$ over draw of training set S of L samples from $\mathcal{D}_{v,x,y}$, we have both

$$\left| \sum_{i=1}^n \left(\mathbb{E}_{(v,x,y)} p_i^w(v, x, y) - \sum_{\ell=1}^L \frac{p_i^w(v^{(\ell)}, x^{(\ell)}, y^{(\ell)})}{L} \right) \right| \leq \epsilon, \quad (7)$$

and

$$\left| \mathbb{E}_{(v,x,y) \sim \mathcal{D}_{v,x,y}} \left[\sum_{i=1}^n rgt_i^w(v, x, y) \right] - \sum_{i=1}^n \widehat{rgt}_i(w) \right| \leq \epsilon. \quad (8)$$

See Appendix E for detailed proofs.

3. Model Architecture

In this section, we describe CITransNet, the proposed context-integrated transformer-based neural network for computing allocation and payment in Equation (II).

3.1. Overview of CITransNet

As shown in Figure 1, CITransNet takes the bidding profile $b \in \mathbb{R}^{n \times m}$, bidder-contexts x and item-contexts y as inputs. An input layer is used first to compute a d -dimensional feature vector for each bidder-item pair. Afterward, the features of all the bidder-item pairs, i.e., $I \in \mathbb{R}^{n \times m \times d}$, are fed into one or multiple interaction layers. Such transformer-based interaction layers model the interactions between bidders and items. The global feature maps $F \in \mathbb{R}^{n \times m \times 3}$ are obtained through the last interaction layer. Finally, we compute the allocation result $g^w(b, x, y)$ and payment result $p^w(b, x, y)$ through the final output layer.

3.2. Input Layer

First, we apply a pre-processing to obtain a representation $e_{x_i} \in \mathbb{R}^{d'_x}$ for each bidder context x_i and $f_{y_j} \in \mathbb{R}^{d'_y}$ for each item context y_j :

- If x_i (or y_j) is drawn from a continuous space, simply set $e_{x_i} = x_i$ (or $f_{y_j} = y_j$).
- If x_i (or y_j) is only drawn from some finite types, embed it into a continuous space, similarly as the common procedure in word embedding (Mikolov et al., 2013). The corresponding embedding is e_{x_i} (or f_{y_j}).

We construct the initial representation for each bidder-item pair: $E = (E_{i,j})_{i \in N, j \in M}$, in which

$$E_{i,j} = [b_{ij}; e_{x_i}; f_{y_j}] \in \mathbb{R}^{1+d'_x+d'_y}, \quad (9)$$

Afterwards, two 1×1 convolutions with a ReLU activation are applied to E and reduce the third-dimension of E from $1 + d'_x + d'_y$ to $d - 1$. Formally,

$$E' = \text{Conv}_2(\text{ReLU}(\text{Conv}_1(E))) \in \mathbb{R}^{n \times m \times (d-1)}, \quad (10)$$

where both Conv_1 and Conv_2 are 1×1 convolutions, and $\text{ReLU}(x) := \max(x, 0)$. By concatenating E' and the bids b , we get $I \in \mathbb{R}^{n \times m \times d}$, the output of our input layer:

$$I = [b; E'] \in \mathbb{R}^{n \times m \times d}, \quad (11)$$

where feature $I_{ij} \in \mathbb{R}^d$ in I captures the bidding and context information of the corresponding bidder-item pair.

3.3. Interaction Layer

Given the representation for all bidder-item pairs $I \in \mathbb{R}^{n \times m \times d}$, we move on to model the interactions between

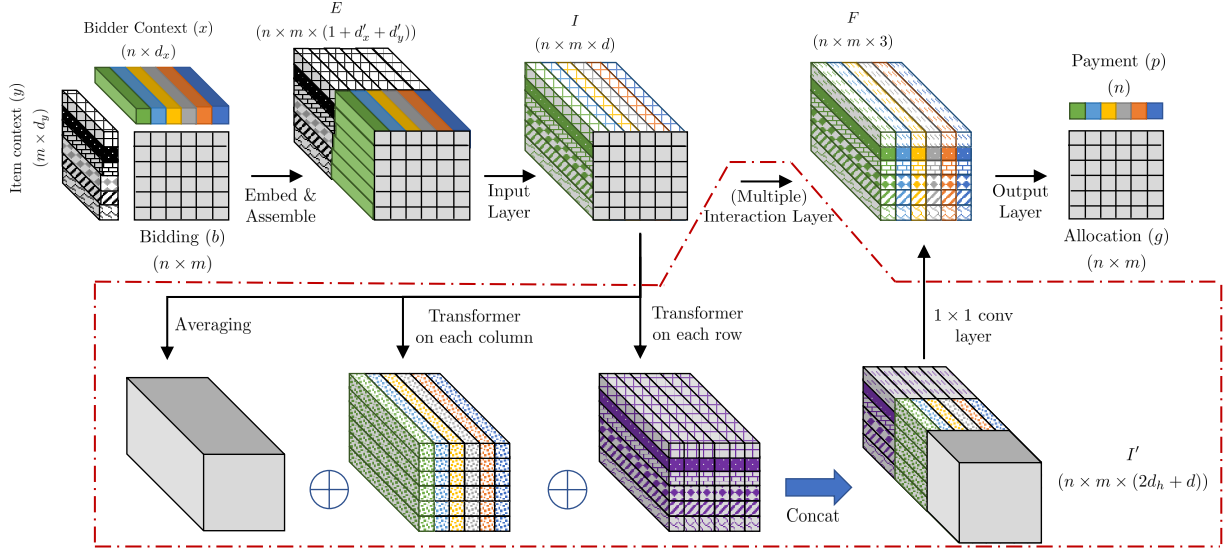


Figure 1: A schematic view of CITransNet, which takes the bidding profile $b \in \mathbb{R}^{n \times m}$, bidder-contexts $x \in \mathcal{X}^n$ and item-contexts $y \in \mathcal{Y}^m$ as inputs. We first embeds x and y into $e_x \in \mathbb{R}^{d'_x}$ and $f_y \in \mathbb{R}^{d'_y}$, and then assemble e_x , f_y and b into $E \in \mathbb{R}^{n \times m \times (1+d'_x+d'_y)}$, the initial representation for each bidder-item pair. The remaining part of our input layer along with one or more transformer-based interaction layers are adopted to model the mutual interactions among different bidders and items. Based on the output $F \in \mathbb{R}^{n \times m \times 3}$ of the last interaction layer, we compute the allocation and payment result via the final output layer.

different bidders and items, which is illustrated in the lower part of Figure 1. The interaction layer is built upon transformer model (Vaswani et al., 2017), which can be used to capture the high-order feature interactions of input through the multi-head self-attention module (Song et al., 2019). See Appendix A for a description of transformer.

Precisely, for each bidder i , we model its interactions with all the m items through transformer on the i -th row of I (denoted as $I_{i,\cdot} \in \mathbb{R}^{m \times d_h}$):

$$I_{i,\cdot}^{\text{row}} = \text{transformer}(I_{i,\cdot}) \in \mathbb{R}^{m \times d_h}, \forall i \in N, \quad (12)$$

where d_h is the size of the hidden nodes in the MLP part of the transformer. Symmetrically, for each item j , we model its interactions with all the n bidders through another transformer on the j -th column of I (called $I_{\cdot,j} \in \mathbb{R}^{n \times d_h}$):

$$I_{\cdot,j}^{\text{column}} = \text{transformer}(I_{\cdot,j}) \in \mathbb{R}^{n \times d_h}, \forall j \in M. \quad (13)$$

Afterwards, the global representation for all the bidder-item pairs is obtained by the average of all the features

$$e^{\text{global}} = \frac{1}{nm} \sum_{i=1}^n \sum_{j=1}^m I_{ij} \in \mathbb{R}^d. \quad (14)$$

Combining $I_{i,\cdot}^{\text{row}}$, $I_{\cdot,j}^{\text{column}}$ and e^{global} together, we get new features I'_{ij} for each bidder-item pair

$$I'_{ij} := [I_{i,\cdot}^{\text{row}}; I_{\cdot,j}^{\text{column}}; e^{\text{global}}] \in \mathbb{R}^{2d_h+d} \quad (15)$$

Finally, as what we did in input layer, two 1×1 convolutions with a ReLU activation are applied on I' in order to reduce the third dimension of I' from $2d_h + d$ to d_{out} . Formally,

$$F = \text{Conv}_4(\text{ReLU}(\text{Conv}_3(I'))) \in \mathbb{R}^{n \times m \times d_{\text{out}}}, \quad (16)$$

where both Conv_4 and Conv_3 are 1×1 convolutions, and F is the output of the interaction layer. By stacking multiple interaction layers, we can model higher-order interactions among all the bidders and items.

3.4. Output Layer

In the last interaction layer, we set $d_{\text{out}} = 3$ and get the global feature maps $F = (F^h, F^q, F^p) \in \mathbb{R}^{n \times m \times 3}$, which will be used to compute the final allocation and payment in the output layer.

The first feature map $F^h \in \mathbb{R}^{n \times m}$ is used to compute the original allocation probability $h^w(b, x, y) \in [0, 1]^{n \times m}$ by softmax activation function on each column of F^h , i.e.,

$$h_{i,j}^w = \text{Softmax}(F_{\cdot,j}^h), \forall j \in M. \quad (17)$$

Here $h_{i,j}^w$ is the probability that item j is allocated to bidder i and we have $\sum_{i=1}^n h_{i,j}^w = 1$ for each item $j \in M$.

Since some item j may not be allocated to any bidder, we use the second feature map F^q to adjust h_w . The weight $q^w(b, x, y) \in (0, 1)^{n \times m}$ of each probability is computed

through sigmoid activation on F^q :

$$q_{i,j}^w = \text{Sigmoid}(F_{i,j}^q), \forall i \in N, \forall j \in M, \quad (18)$$

where $\text{Sigmoid}(x) := \frac{1}{1+e^{-x}} \in (0, 1)$.

The allocation result g^w is then obtained by combining h^w and q^w together:

$$g_{ij}^w(b, x, y) = q_{ij}^w(b, x, y)h_{ij}^w(b, x, y). \quad (19)$$

As a result, we have $0 < \sum_{i=1}^n g_{i,j}^w(b, x, y) < 1$ for each item $j \in M$.

For payment, we compute payment fraction $\tilde{p}^w(b, x, y) \in (0, 1)^n$ via the third feature map F^p :

$$\tilde{p}_i^w = \text{Sigmoid}\left(\frac{1}{m} \sum_{j=1}^m F_{ij}^p\right), \forall i \in N, \quad (20)$$

where \tilde{p}_i^w is the fraction of bidder i 's utility that she has to pay to the auctioneer. Given the allocation g^w and payment fraction \tilde{p}^w , the payment for bidder i is

$$p_i^w(b, x, y) = \tilde{p}_i^w(b, x, y) \sum_{j=1}^m g_{ij}^w(b, x, y)b_{ij}. \quad (21)$$

By doing so, Equation (IR) is satisfied.

Remark 3.1 (Permutation-equivariant). Similar to the definition in Rahme et al. (2021a), we say an auction mechanism (g^w, p^w) is permutation-equivariant if for any two permutation matrices $\Pi_n \in \{0, 1\}^{n \times n}$ and $\Pi_m \in \{0, 1\}^{m \times m}$, and any input (including bids $b \in \mathbb{R}^{n \times m}$, bidder-contexts $x \in \mathbb{R}^{n \times d_x}$ and item-contexts $y \in \mathbb{R}^{m \times d_y}$), we have $g^w(\Pi_n b \Pi_m, \Pi_n x, \Pi_m^T y) = \Pi_n g^w(b, x, y) \Pi_m$ and $p^w(\Pi_n b \Pi_m, \Pi_n x, \Pi_m^T y) = \Pi_n p^w(b, x, y)$. Transformer is known to be permutation-equivariant, since it maps each embedding in input to a new embedding that incorporates the information of the *set* of all the input embeddings. Moreover, the 1×1 convolutions we use in CITransNet are all per bidder-item wise, i.e., acting on each bidder-item pair. As a result, CITransNet maintains permutation-equivariant.

3.5. Optimization and training

Similar to Dütting et al. (2019), CITransNet is optimized through the augmented Lagrangian method. The Lagrangian with a quadratic penalty is:

$$\begin{aligned} \mathcal{L}_\rho(w; \lambda) = & -\frac{1}{L} \sum_{\ell=1}^L \sum_{i=1}^n p_i^w(v^{(\ell)}, x^{(\ell)}, y^{(\ell)}) + \\ & \sum_{i=1}^n \lambda_i \widehat{r}gt_i(w) + \frac{\rho}{2} \sum_{i=1}^n \left(\widehat{r}gt_i(w)\right)^2, \end{aligned} \quad (22)$$

where $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_n) \in \mathbb{R}^n$ is the Lagrange multipliers, and $\rho > 0$ is a hyperparameter that controls the weight

of the quadratic penalty. During optimization, we update the model parameters and Lagrange multipliers in turn, i.e., we alternately find $w^{new} \in \arg \min_w \mathcal{L}_\rho(w^{old}, \lambda^{old})$ and update $\lambda_i^{new} = \lambda_i^{old} + \rho \cdot \widehat{r}gt_i(w^{new})$, $\forall i \in N$. See Appendix B for a detailed optimization and training procedure.

4. Experiments

In this section, we conduct empirical experiments to show the effectiveness of CITransNet in different contextual auctions³. Afterward, we demonstrate the out-of-setting generalization ability for CITransNet by evaluating the trained model in settings with different numbers of bidders or items. Our experiments are run on a Linux machine with NVIDIA Graphics Processing Unit (GPU) cores. Each result is obtained by averaging across 5 different runs. We ignore the standard deviation since it is small in all the experiments.

Baseline Methods. We compare CITransNet with the following baselines: 1) Item-wise Myerson, a strong baseline used in Dütting et al. (2019), which independently applies Myerson auction with respect to each item⁴; 2) RegretNet (Dütting et al., 2019), which adopts fully-connected neural networks to compute auction mechanism; EquivariantNet (Rahme et al., 2021a), which is a permutation-equivariant architecture to design the special mechanism of symmetric auctions⁵; 3) CIREgretNet and CIEquivariantNet, which are the context-integrated version of RegretNet and EquivariantNet. Specifically, we replace the interaction layers of our CITransNet with RegretNet and EquivariantNet, respectively. We set these baselines to evaluate the effectiveness of our transformer-based interaction layers.

See Appendix C for implementation details of all methods.

Evaluation. Following Dütting et al. (2019) and Rahme et al. (2021a), to evaluate each method, we adopt empirical revenue (the minus objective in Equation (II)) and empirical ex-post regret average across all the bidders $\widehat{r}gt := \frac{1}{n} \sum_{i=1}^n \widehat{r}gt_i$. We obtain the empirical regret for each bidder by executing gradient ascent on her bids b_i for

³Our implementation is available at <https://github.com/zjduan/CITransNet>.

⁴Bundle Myerson is another baseline used in Dütting et al. (2019) that satisfies both DSIC and IR. However, we find it always performs worse than Item-wise Myerson, both in our experiments and in Dütting et al. (2019). Therefore, we do not present its results.

⁵While Rahme et al. (2021b) formulate auction learning as an adversarial learning framework, we view this as an orthogonal problem since this work mainly focuses on the innovation of neural architectures. Therefore, to make a fair comparison, we adopt the learning framework in Dütting et al. (2019) for the baselines and leave the adversarial learning framework extension for future work.

Table 1: Experiment results of known settings (Setting A-C). The optimal solutions are given by Myerson (1981). Each experiment is run 5 times and the average results are presented.

Method	A: 3×1 $ \mathcal{X} = 5, \mathcal{Y} = 1$		B: 3×1 $ \mathcal{X} = 5, \mathcal{Y} = 2$		C: 5×1 $\mathcal{X}, \mathcal{Y} \subset \mathbb{R}^{10}$	
	<i>rev</i>	<i>rgt</i>	<i>rev</i>	<i>rgt</i>	<i>rev</i>	<i>rgt</i>
Optimal	0.594	-	0.456	-	0.367	-
RegretNet	0.516	<0.001	0.412	<0.001	0.329	<0.001
EquivariantNet	0.498	<0.001	0.403	<0.001	0.311	<0.001
CIRegretNet	0.594	<0.001	0.453	<0.001	0.364	<0.001
CIEquivariantNet	0.590	<0.001	0.452	<0.001	0.360	<0.001
CITransNet	0.593	<0.001	0.454	<0.001	0.366	<0.001

200 iterations. We run such gradient ascent for 100 times with different initial bids $b_i^{(0)}$, and the maximum regret is recorded for bidder i .

Single-item Contextual Auctions. First, we evaluate CITransNet in single-item auctions, whose optimal solutions are given by Myerson (1981). We aim to justify whether CITransNet can recover the near-optimal solutions. The specific single-item auctions we consider are:

- (A) 3 bidders and 1 item, with discrete bidder-contexts and item-context, in which $\mathcal{X} = \{1, 2, 3, 4, 5\}$ and $\mathcal{Y} = \{1\}$. Both contexts are independently and uniformly sampled. Given $x_i \in \mathcal{X}$ and $y_1 = 1$, v_{i1} is drawn according to the truncated normal distribution $\mathcal{N}(\frac{x_i}{6}, 0.1)$ in $[0, 1]$.
- (B) 3 bidders and 1 item, with discrete bidder-contexts and item-context, in which $\mathcal{X} = \{1, 2, 3, 4, 5\}$ and $\mathcal{Y} = \{1, 2\}$. Both contexts are independently and uniformly sampled. Given $x_i \in \mathcal{X}$, v_{i1} is drawn according to the truncated normal distribution $\mathcal{N}(\frac{x_i}{6}, 0.1)$ in $[0, 1]$ when $y_1 = 1$, and is drawn according to probability densities $f_i(x) = \frac{i}{6}e^{-\frac{i}{6}x}$ truncated in $[0, 1]$ when $y_1 = 2$.
- (C) 5 bidders and 1 item, with continuous bidder-contexts and item-context, in which $\mathcal{X} = [-1, 1]^{10}$ and $\mathcal{Y} = [-1, 1]^{10}$. Both the contexts are independently and uniformly sampled. Given $x_i \in \mathcal{X}$ and $y_j \in \mathcal{Y}$, v_{ij} is drawn according to $U[0, \text{Sigmoid}(x_i^T y_j)]$.

We present the experimental results of Setting A, B and C in Table 1. We can see that all the context-integrated models (CIRegretNet, CIEquivariantNet and CITransNet) are able to recover the optimal solutions given by Myerson (1981) in these simple settings: near-optimal revenues are achieved with regrets less than 0.001. In comparison, despite low regret, RegretNet and EquivariantNet fail to reach the optimal solution. It

turns out that integrating context information into model architecture is crucial in contextual auction design. Furthermore, EquivariantNet, the symmetric mechanism designer, fails to reach the same performance as RegretNet, which reflects the importance of designing asymmetric solutions in contextual auctions.

Multi-item Contextual Auctions. Next, we illustrate the potential of CITransNet to discover new auction designs in multi-item contextual auctions without known solutions. We consider discrete context settings as follows:

- (D) 2 bidders with $\mathcal{X} = \{1, 2, \dots, 10\}$ and 5 items with $\mathcal{Y} = \{1, 2, \dots, 10\}$. All the contexts are uniform sampled, and v_{ij} is drawn according to the normal distribution $\mathcal{N}(\frac{(x_i+y_j) \bmod 10+1}{11}, 0.05)$ truncated in $[0, 1]$.
- (E) 3 bidders and 10 items. The discrete contexts and corresponding values are drawn similarly as Setting D.
- (F) 5 bidders and 10 items, which is, to the best of our knowledge, the largest auction size considered in previous literatures of deep learning based auction design (Rahme et al., 2021b). The discrete contexts and corresponding values are drawn similarly as Setting D.

Additionally, We also construct continuous context settings based on Setting C:

- (G) 2 bidders and 5 items. The continuous contexts and corresponding values are drawn similarly as Setting C.
- (H) 3 bidders and 10 items. The continuous contexts and corresponding values are drawn similarly as Setting C.
- (I) 5 bidders and 10 items. The continuous contexts and corresponding values are drawn similarly as Setting C.

Experimental results for Setting D-I are shown in Table 2. CITransNet obtains the best revenue results in all the settings while keeping low regret (less than 0.003 in Setting

Table 2: Experiment results for Setting D-I. Each experiment is run by 5 times and the average results are presented.

Method	D: 2×5 $ \mathcal{X} = \mathcal{Y} = 10$		E: 3×10 $ \mathcal{X} = \mathcal{Y} = 10$		F: 5×10 $ \mathcal{X} = \mathcal{Y} = 10$		G: 2×5 $\mathcal{X}, \mathcal{Y} \subset \mathbb{R}^{10}$		H: 3×10 $\mathcal{X}, \mathcal{Y} \subset \mathbb{R}^{10}$		I: 5×10 $\mathcal{X}, \mathcal{Y} \subset \mathbb{R}^{10}$	
	rev	rgt	rev	rgt	rev	rgt	rev	rgt	rev	rgt	rev	rgt
Item-wise Myerson	2.821	-	6.509	-	7.376	-	1.071	-	2.793	-	3.684	-
CIRegretNet	2.803	<0.001	5.846	<0.001	6.339	<0.003	1.104	<0.001	2.424	<0.001	2.999	<0.001
CIEquivariantNet	2.841	<0.001	6.703	<0.001	7.602	<0.003	1.147	<0.001	2.872	<0.001	3.806	<0.001
CITransNet	2.916	<0.001	6.872	<0.001	7.778	<0.003	1.177	<0.001	2.918	<0.001	3.899	<0.001




Figure 2: Out-of-setting generalization results: we train CITransNet and evaluate it on the same contextual auction with a different number of bidders or items. We set Item-Wise Myerson as the baseline. The regret results are less than 0.001 in all of these experiments. (a) Trained on Setting E (3×10 with $|\mathcal{X}| = |\mathcal{Y}| = 10$) and evaluated with different number of bidders. (b) Trained on Setting D (2×5 with $|\mathcal{X}| = |\mathcal{Y}| = 10$) and evaluated with different number of items. (c) Trained on Setting G (2×5 with $\mathcal{X}, \mathcal{Y} \subset \mathbb{R}^{10}$) and evaluated with different number of items.

F and less than 0.001 in all the other settings). Notice that the only difference between CITransNet, CIRegretNet and CIEquivariantNet is the architecture of interaction layers. Such a result indicates the effectiveness of our transformer-based interaction module to capture the complex mutual influence among bidders and items. Furthermore, both CITransNet and CIEquivariantNet outperform CIRegretNet a lot in all the 3×10 and 5×10 auctions, showing that adding the inductive bias of permutation-equivariance is helpful in large-scale auction design.

Out-of-setting Generalization. In addition, to show the effectiveness of CITransNet, we also conduct out-of-setting generalization experiments. Specifically, we train our model and evaluate it in auctions with a different number of bidders or items. Such evaluation is feasible for CITransNet, since the size of parameters in CITransNet does not rely on the number of bidders and items. We illustrate the experimental results on Figure 2, and see Appendix D for more detailed numerical values. Figure 2a shows the experimental results of generalizing to a varying number of bidders. We train CITransNet on Setting E, the discrete context settings with 3 bidders and 10 items, and we evaluate CITransNet on the same contextual auction with n

bidders and 10 items ($n \in \{3, 4, 5, 6, 7\}$). We observe good generalization results: In addition to obtain low regret (less than 0.001) in all the test settings, CITransNet outperforms Item-wise Myerson when $n \in \{3, 4, 5\}$ ⁶. Furthermore, in Figure 2b and Figure 2c we present the experimental results of generalizing to varying number of items. We train CITransNet on Setting D and Setting G respectively, where both settings have 2 bidders and 5 items, and we test the model on the same contextual auction with 2 bidders and m items ($m \in \{3, 4, 5, 6, 7\}$). Again, we observe good generalization results. While still keeping small regret (less than 0.001), CITransNet is able to outperform Item-wise Myerson in all the test auctions.

5. Conclusion

In this paper, we propose a new (transformer-based) neural architecture, CITransNet, for contextual auction design. CITransNet is permutation-equivariant with respect to bids and contexts, and it can handle asymmetric information in auctions. We show by experiments that CITransNet

⁶As comparison, we find CIEquivariantNet fails to generalize to different bidders. See Appendix D for the results.

can recover the known optimal analytical solutions in simple auctions, and we demonstrate the effectiveness of the transformer-based interaction layers in CITransNet by comparing CITransNet with the context integrated version of RegretNet and EquivariantNet. Furthermore, we also illustrate the out-of-setting generalization ability for CITransNet by evaluating it in auctions with a varying number of bidders or items. Given the decent generalizability of CITransNet, an immediate next step is to test CITransNet over an industry-scale dataset. It would also be interesting to test CITransNet in an online manner.

Acknowledgements

This work is supported by Science and Technology Innovation 2030 - “New Generation Artificial Intelligence” Major Project (No. 2018AAA0100901). We thank Aranyak Mehta and Di Wang for an insightful discussion on the initial version of this paper. We thank all anonymous reviewers for their helpful feedback.

References

- Amin, K., Rostamizadeh, A., and Syed, U. Repeated contextual auctions with strategic buyers. *Advances in Neural Information Processing Systems*, 27:622–630, 2014.
- Babaioff, M., Immorlica, N., Lucier, B., and Weinberg, S. M. A simple and approximately optimal mechanism for an additive buyer. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pp. 21–30. IEEE, 2014.
- Balcan, M.-F., Blum, A., Hartline, J. D., and Mansour, Y. Reducing mechanism design to algorithm design via machine learning. *Journal of Computer and System Sciences*, 74(8):1245–1270, 2008.
- Balcan, M.-F. F., Sandholm, T., and Vitercik, E. Sample complexity of automated mechanism design. In *Advances in Neural Information Processing Systems*, pp. 2083–2091, 2016.
- Brero, G., Eden, A., Gerstgrasser, M., Parkes, D., and Rheingans-Yoo, D. Reinforcement learning of sequential price mechanisms. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 5219–5227, 2021.
- Cai, Y. and Zhao, M. Simple mechanisms for subadditive buyers via duality. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pp. 170–183, 2017.
- Cai, Y., Daskalakis, C., and Weinberg, S. M. An algorithmic characterization of multi-dimensional mechanisms. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pp. 459–478, 2012.
- Chawla, S., Hartline, J. D., Malec, D. L., and Sivan, B. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pp. 311–320, 2010.
- Cole, R. and Roughgarden, T. The sample complexity of revenue maximization. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pp. 243–252, 2014.
- Conitzer, V. and Sandholm, T. Complexity of mechanism design. *arXiv preprint cs/0205075*, 2002.
- Conitzer, V. and Sandholm, T. Self-interested automated mechanism design and implications for optimal combinatorial auctions. In *Proceedings of the 5th ACM Conference on Electronic Commerce*, pp. 132–141, 2004.
- Curry, M., Chiang, P.-Y., Goldstein, T., and Dickerson, J. Certifying strategyproof auction networks. *Advances in Neural Information Processing Systems*, 33, 2020.
- Daskalakis, C., Deckelbaum, A., and Tzamos, C. Strong duality for a multiple-good monopolist. *Econometrica*, 85(3):735–767, 2017.
- Devanur, N. R., Huang, Z., and Psomas, C.-A. The sample complexity of auctions with side information. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pp. 426–439, 2016.
- Drutsa, A. Optimal non-parametric learning in repeated contextual auctions with strategic buyer. In *International Conference on Machine Learning*, pp. 2668–2677. PMLR, 2020.
- Duan, Z., Zhang, D., Huang, W., Du, Y., Wang, J., Yang, Y., and Deng, X. Towards the pac learnability of nash equilibrium. *arXiv preprint arXiv:2108.07472*, 2021.
- Dütting, P., Fischer, F., Jirapinyo, P., Lai, J. K., Lubin, B., and Parkes, D. C. Payment rules through discriminant-based classifiers, 2015.
- Dütting, P., Feng, Z., Narasimhan, H., Parkes, D., and Ravindranath, S. S. Optimal auctions through deep learning. In *International Conference on Machine Learning*, pp. 1706–1715. PMLR, 2019.
- Feng, Z., Narasimhan, H., and Parkes, D. C. Deep learning for revenue-optimal auctions with budgets. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems*, pp. 354–362, 2018.

- Galal, H. S. and Youssef, A. M. Verifiable sealed-bid auction on the ethereum blockchain. In *International Conference on Financial Cryptography and Data Security*, pp. 265–278. Springer, 2018.
- Giannakopoulos, Y. and Koutsoupias, E. Duality and optimality of auctions for uniform distributions. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pp. 259–276, 2014.
- Golowich, N., Narasimhan, H., and Parkes, D. C. Deep learning for multi-facility location mechanism design. In *IJCAI*, pp. 261–267, 2018.
- Golrezaei, N., Javanmard, A., and Mirrokni, V. Dynamic incentive-aware learning: Robust pricing in contextual auctions. *Operations Research*, 69(1):297–314, 2021.
- Gonczarowski, Y. A. and Weinberg, S. M. The sample complexity of up-to- ϵ multi-dimensional revenue maximization. *Journal of the ACM (JACM)*, 68(3):1–28, 2021.
- Guo, C., Huang, Z., and Zhang, X. Settling the sample complexity of single-parameter revenue maximization. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pp. 662–673, 2019.
- Haghpanah, N. and Hartline, J. When is pure bundling optimal? *The Review of Economic Studies*, 88(3):1127–1156, 2021.
- Hart, S. and Nisan, N. Approximate revenue maximization with multiple items. *Journal of Economic Theory*, 172: 313–347, 2017.
- Huang, J., Han, Z., Chiang, M., and Poor, H. V. Auction-based resource allocation for cooperative communications. *IEEE Journal on Selected Areas in Communications*, 26(7):1226–1237, 2008.
- Jansen, B. J. and Mullen, T. Sponsored search: an overview of the concept, history, and technology. *International Journal of Electronic Business*, 6(2):114–131, 2008.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Lahaie, S. A kernel-based iterative combinatorial auction. In *Twenty-Fifth AAAI Conference on Artificial Intelligence*, 2011.
- Liu, X., Yu, C., Zhang, Z., Zheng, Z., Rong, Y., Lv, H., Huo, D., Wang, Y., Chen, D., Xu, J., Wu, F., Chen, G., and Zhu, X. Neural auction: End-to-end learning of auction mechanisms for e-commerce advertising. *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 2021.
- Luong, N. C., Xiong, Z., Wang, P., and Niyato, D. Optimal auction for edge computing resource management in mobile blockchain networks: A deep learning approach. In *2018 IEEE International Conference on Communications (ICC)*, pp. 1–6. IEEE, 2018.
- Manelli, A. M. and Vincent, D. R. Bundling as an optimal selling mechanism for a multiple-good monopolist. *Journal of Economic Theory*, 127(1):1–35, 2006.
- Mao, J., Leme, R. P., and Schneider, J. Contextual pricing for lipschitz buyers. In *NeurIPS*, pp. 5648–5656, 2018.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., and Dean, J. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pp. 3111–3119, 2013.
- Miller, A. H., Fisch, A., Dodge, J., Karimi, A.-H., Bordes, A., and Weston, J. Key-value memory networks for directly reading documents. In *EMNLP*, 2016.
- Myerson, R. B. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.
- Nedelec, T., Baudet, J., Perchet, V., and Karoui, N. E. Adversarial learning for revenue-maximizing auctions. In *AAMAS*, 2021.
- Pavlov, G. Optimal mechanism for selling two goods. *The BE Journal of Theoretical Economics*, 11(1), 2011.
- Peri, N., Curry, M. J., Dooley, S., and Dickerson, J. P. PreferenceNet: Encoding human preferences in auction design with deep learning. *arXiv preprint arXiv:2106.03215*, 2021.
- Rahme, J., Jelassi, S., Bruna, J., and Weinberg, S. M. A permutation-equivariant neural network architecture for auction design. In *AAAI*, pp. 5664–5672, 2021a.
- Rahme, J., Jelassi, S., and Weinberg, S. M. Auction learning as a two-player game. In *9th International Conference on Learning Representations*, 2021b.
- Sandholm, T. and Likhodedov, A. Automated design of revenue-maximizing combinatorial auctions. *Operations Research*, 63(5):1000–1025, 2015.
- Shalev-Shwartz, S. and Ben-David, S. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.
- Shen, W., Tang, P., and Zuo, S. Automated mechanism design via neural networks. In *AAMAS*, 2019.
- Shen, W., Peng, B., Liu, H., Zhang, M., Qian, R., Hong, Y., Guo, Z., Ding, Z., Lu, P., and Tang, P. Reinforcement

- mechanism design: With applications to dynamic pricing in sponsored search auctions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 2236–2243, 2020.
- Song, W., Shi, C., Xiao, Z., Duan, Z., Xu, Y., Zhang, M., and Tang, J. AutoInt: Automatic feature interaction learning via self-attentive neural networks. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, pp. 1161–1170, 2019.
- Tacchetti, A., Strouse, D., Garnelo, M., Graepel, T., and Bachrach, Y. A neural architecture for designing truthful and efficient auctions. *arXiv preprint arXiv:1907.05181*, 2019.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. In *Advances in neural information processing systems*, pp. 5998–6008, 2017.
- Yao, A. C.-C. 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*, pp. 92–109. SIAM, 2014.
- Yao, A. C.-C. Dominant-strategy versus bayesian multi-item auctions: Maximum revenue determination and comparison. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pp. 3–20, 2017.
- Zhiyanov, A. and Drutsa, A. Bisection-based pricing for repeated contextual auctions against strategic buyer. In *International Conference on Machine Learning*, pp. 11469–11480. PMLR, 2020.

A. Transformer Architecture

Transformer architecture (Vaswani et al., 2017) aims at modeling the mutual correlations among a set of tokens (e.g., words in a sentence in machine translation) via multi-head self-attention module. In our paper, we use transformer to model the interactions among the items (or bidders) with respect to a fixed bidder (or item). Without loss of generality, we denote the input as

$$E_{\text{input}} = (e_1, e_2, \dots, e_n)^T \in \mathbb{R}^{n \times d}, \quad (23)$$

where n is the number of tokens (i.e., bidders or items) and d is the dimension for each feature vector e_i .

Let d_h be the hidden dimension of transformer, and H be the number of heads (i.e., subspace). For head $h \in [H]$, we use the key-value attention mechanism (Miller et al., 2016) to determine which feature combinations are meaningful in the corresponding subspace. Specifically, for each token $i \in [n]$, we first compute the correlation between token i and token j in head h :

$$\alpha_{i,j}^{(h)} = \frac{\exp(\psi^{(h)}(e_i, e_j))}{\sum_{k=1}^n \exp(\psi^{(h)}(e_i, e_k))}, \quad (24)$$

where

$$\psi^{(h)}(e_i, e_j) = \langle W_{\text{query}}^{(h)} e_i, W_{\text{key}}^{(h)} e_j \rangle, \quad (25)$$

is an attention function which defines the similarity between the token i and j under head h . $\langle \cdot, \cdot \rangle$ is inner product, and $W_{\text{query}}^{(h)}, W_{\text{key}}^{(h)} \in \mathbb{R}^{d' \times d}$ are transformation matrices which map the original embedding space \mathbb{R}^d into a $d' = \frac{d_h}{H}$ dimensional space $\mathbb{R}^{d'}$.

Next, we update the representation of token i in subspace h by combining all relevant features. This is done by computing the weighted sum using coefficients $\alpha_{i,j}^{(h)}$:

$$\tilde{e}_i^{(h)} = \sum_{j=1}^n \alpha_{i,j}^{(h)} (W_{\text{value}}^{(h)} e_j) \in \mathbb{R}^{d'}, \quad (26)$$

where $W_{\text{value}}^{(h)} \in \mathbb{R}^{d' \times d}$. Since $\tilde{e}_i^{(h)} \in \mathbb{R}^{d'}$ is a combination of token i and all its relevant tokens, it represents a new combinatorial feature.

Afterwards, we collect combinatorial features learned in all subspaces as follows:

$$\tilde{e}_i = \tilde{e}_i^{(1)} \oplus \tilde{e}_i^{(2)} \oplus \dots \oplus \tilde{e}_i^{(H)} \in \mathbb{R}^{Hd'} = \mathbb{R}^{d_h}, \quad (27)$$

where \oplus is the concatenation operator, and H is the number of total heads.

Finally, a token-wise MLP is applied to each token i and we get a new representation for it.

$$e'_i = \text{MLP}(\tilde{e}_i) \in \mathbb{R}^{d_h}, \quad (28)$$

and the final output is

$$E_{\text{output}} = (e'_1, e'_2, \dots, e'_n)^T \in \mathbb{R}^{n \times d_h}. \quad (29)$$

Notice that the parameters to be optimized in transformer are $W_{\text{query}}^{(h)}, W_{\text{key}}^{(h)}, W_{\text{value}}^{(h)} \in \mathbb{R}^{d' \times d}$ for all $h \in [H]$ and the parameters of the final token-wise MLP, all of which are unrelated to the number of tokens n . Furthermore, the transformer architecture is permutation-equivariant.

B. Optimization and Training Procedures

We use the augmented Lagrangian method to solve the constrained training problem in Equation (II) over the space of neural network parameters $w \in \mathbb{R}^{d_w}$.

Algorithm 1 describe the training procedure of CITransNet. First, for each iteration $t \in [T]$, we randomly draw a minibatch \mathcal{S}_t of size B , in which $\mathcal{S}_t = \{(v^{(1)}, x^{(1)}, y^{(1)}), \dots, (v^{(B)}, x^{(B)}, y^{(B)})\}$. Afterwards, we alternately update the model parameters and the Lagrange multipliers:

Algorithm 1 CITransNet Training

```

1: Input: Minibatches  $\mathcal{S}_1, \dots, \mathcal{S}_T$  of size  $B$ 
2: Parameters:  $\forall t \in [T], \rho_t > 0, \gamma > 0, \eta > 0, c > 0, T \in \mathbb{N}, \Gamma \in \mathbb{N}, T_\lambda \in \mathbb{N}$ 
3: Initialize:  $w^0 \in \mathbb{R}^d, \lambda^0 \in \mathbb{R}^n$ 
4: for  $t = 0$  to  $T$  do
5:   Receive minibatch  $\mathcal{S}_t = \{(v^{(1)}, x^{(1)}, y^{(1)}), \dots, (v^{(B)}, x^{(B)}, y^{(B)})\}$ 
6:   Initialize misreports  $v_i^{\prime(\ell)} \in \mathcal{V}_i, \forall \ell \in [B], i \in N$ 
7:   for  $r = 0$  to  $\Gamma$  do
8:      $\forall \ell \in [B], i \in N$  :
9:        $v_i^{\prime(\ell)} \leftarrow v_i^{\prime(\ell)} + \gamma \nabla_{v_i^{\prime(\ell)}} u_i^w(v_i^{\prime(\ell)}, (v_i^{\prime(\ell)}, v_{-i}^{(\ell)}), x^{(\ell)}, y^{(\ell)})$ 
10:    end for
11:   Compute regret gradient:
12:      $\forall \ell \in [B], i \in N$ :
13:        $g_{\ell,i}^t = \nabla_w \left[ u_i^w(v_i^{\prime(\ell)}, (v_i^{\prime(\ell)}, v_{-i}^{(\ell)}), x^{(\ell)}, y^{(\ell)}) - u_i^w(v_i^{(\ell)}, v^{(\ell)}, x^{(\ell)}, y^{(\ell)}) \right] \Big|_{w=w^t}$ 
14:   Compute Lagrangian gradient using Equation (30) and update  $w^t$ :
15:      $w^{t+1} \leftarrow w^t - \eta \nabla_w \mathcal{L}_{\rho_t}(w^t, \lambda^t)$ 
16:   Update Lagrange multipliers  $\lambda$  once in  $T_\lambda$  iterations:
17:     if  $t$  is a multiple of  $T_\lambda$  then
18:        $\lambda_i^{t+1} \leftarrow \lambda_i^t + \rho_t \widehat{rgt}_i(w^{t+1}), \forall i \in N$ 
19:     else
20:        $\lambda^{t+1} \leftarrow \lambda^t$ 
21:   end for

```

(a) $w^{new} \in \arg \min_w \mathcal{L}_\rho(w^{old}, \lambda^{old})$

(b) $\lambda_i^{new} = \lambda_i^{old} + \rho \cdot \widehat{rgt}_i(w^{new}), \forall i \in N$

The update (a) is performed approximately using gradient descent. The gradient of \mathcal{L}_ρ w.r.t. w for fixed λ^t is given by:

$$\nabla_w \mathcal{L}_\rho(w, \lambda^t) = -\frac{1}{B} \sum_{\ell=1}^B \sum_{i \in N} \nabla_w p_i^w(v^{(\ell)}, x^{(\ell)}, y^{(\ell)}) + \sum_{i \in N} \sum_{\ell=1}^B \lambda_i^t g_{\ell,i} + \rho \sum_{i \in N} \sum_{\ell=1}^B \widehat{rgt}_i(w) g_{\ell,i}, \quad (30)$$

where

$$g_{\ell,i} = \nabla_w \left[\max_{v_i^{\prime(\ell)} \in \mathcal{V}_i} u_i^w(v_i^{\prime(\ell)}, (v_i^{\prime(\ell)}, v_{-i}^{(\ell)}), x^{(\ell)}, y^{(\ell)}) - u_i^w(v_i^{(\ell)}, v^{(\ell)}, x^{(\ell)}, y^{(\ell)}) \right].$$

The computation of \widehat{rgt}_i and $g_{\ell,i}$ involve a ‘‘max’’ over misreports for each bidder i , and we solve it approximately by gradient ascent. In particular, we maintain misreports $v_i^{\prime(\ell)}$ for each bidder i on each sample ℓ . For every update on the model parameters w^t , we perform Γ gradient ascent updates to compute the optimal misreports.

C. Implementation Details

For all the settings (Setting A-I), we generate the training set of each setting with size in $\{50000, 100000, 200000\}$ and test set of size 5000.

For all the methods, we train the models for a maximum of 80 epochs with batch size 500. We set the embedding size in settings with discrete context (Setting A, B, D, E, F) as 16. The value of ρ in the augmented Lagrangian (Equation (22)) was set as 1.0 at the beginning and incremented by 5 every two epochs. The value of λ in Equation (22) was set as 5.0 initially and incremented every certain number (selected from $\{2 - 10\}$) of epochs. All the models and regret are optimized through Adam (Kingma & Ba, 2014) optimizer. Following Dütting et al. (2019), for each update on model parameters, we

Table 3: Out-of-setting generalization results of CITransNet and CIEquivariantNet: we train each model and evaluate it on the same contextual auction with a different number of bidders or items. (a) Trained on Setting E (3×10 with $|\mathcal{X}| = |\mathcal{Y}| = 10$) and evaluated with different number of bidders. (b) Trained on Setting D (2×5 with $|\mathcal{X}| = |\mathcal{Y}| = 10$) or Setting G (2×5 with $|\mathcal{X}| \subset \mathbb{R}^{10}, |\mathcal{Y}| \subset \mathbb{R}^{10}$) and evaluated with different number of items.

(a)

Method	3×10		4×10		5×10		6×10		7×10	
	<i>rev</i>	<i>rgt</i>	<i>rev</i>	<i>rgt</i>	<i>rev</i>	<i>rgt</i>	<i>rev</i>	<i>rgt</i>	<i>rev</i>	<i>rgt</i>
Trained on Setting E: $n = 3, m = 10$ with $ \mathcal{X} = \mathcal{Y} = 10$										
Item-wise Myerson	6.509	-	7.028	-	7.376	-	7.629	-	7.837	-
CIEquivariantNet	6.703	<0.001	7.024	0.018	7.229	0.051	7.365	0.079	7.474	0.1
CITransNet	6.872	<0.001	7.222	<0.001	7.395	<0.001	7.496	<0.001	7.598	<0.001

(b)

Method	2×3		2×4		2×5		2×6		2×7	
	<i>rev</i>	<i>rgt</i>	<i>rev</i>	<i>rgt</i>	<i>rev</i>	<i>rgt</i>	<i>rev</i>	<i>rgt</i>	<i>rev</i>	<i>rgt</i>
Trained on Setting D: $n = 2, m = 5$ with $ \mathcal{X} = \mathcal{Y} = 10$										
Item-wise Myerson	1.691	-	2.264	-	2.821	-	3.391	-	3.954	-
CIEquivariantNet	1.687	<0.001	2.267	<0.001	2.841	<0.001	3.405	<0.001	3.971	<0.001
CITransNet	1.720	<0.001	2.333	<0.001	2.916	<0.001	3.540	<0.001	4.141	<0.001
Trained on Setting G: $n = 2, m = 5$ with $\mathcal{X}, \mathcal{Y} \subset \mathbb{R}^{10}$										
Item-wise Myerson	0.640	-	0.855	-	1.071	-	1.290	-	1.489	-
CIEquivariantNet	0.663	<0.001	0.900	<0.001	1.147	<0.001	1.400	<0.001	1.637	<0.001
CITransNet	0.677	<0.001	0.919	<0.001	1.177	<0.001	1.438	<0.001	1.686	<0.001

run $\Gamma = 25$ update steps on the misreport bid b_i for each bidder, and the optimized misreports are cached to initialize the misreports bidding in the next epoch.

For our proposed CITransNet, the output channel of the first 1×1 convolution in both the input layer and interaction layers are set to 64. We set $d = 64$ for the 1×1 convolution with residual connection in input layer, and $d_h = 64$ for the final 1×1 convolution in each interaction layer. We tune the numbers of interaction layers from $\{2, 3\}$, and in each interaction layer we adopt transformer with 4 heads and 64 hidden nodes.

RegretNet and CIREgretNet take fully-connected neural networks as the core architecture. We choose the number of layers from $\{3, 4, 5, 6, 7\}$ and the number of hidden nodes per layer from $\{64, 128, 256\}$. As for EquivariantNet and CIEquivariantNet, we use 4 exchangeable matrix layers of 64 channels each.

D. Additional Out-of-setting Generalization Experiments

In addition to CITransNet, we also conduct the same out-of-setting generalization experiments for CIEquivariantNet. The numerical results are shown in Table 3. While CITransNet generalize well to all of these settings with low regret (less than 0.001), CIEquivariantNet fails to obtain low regret when generalizing to auctions with a different number of bidders. Such a result indicates the critical role of the transformer-based interaction layers in CITransNet when generalized to settings with varying bidders.

E. Proof of Theorem 2.6

The proof is done by combining covering numbers (Shalev-Shwartz & Ben-David, 2014; Dütting et al., 2019) and a generalization Lemma (Lemma E.1, whose technique comes from Duan et al. (2021)) based on concentration inequality.

E.1. Basic Definition

On top of the definitions in Section 2.3, we first define the covering numbers of bidder's utility functions and regret functions.

Covering Numbers $\mathcal{N}_{\infty,1}(\mathcal{U}, r)$ and $\mathcal{N}_{\infty}(\mathcal{U}_i, r)$. Let \mathcal{U}_i be the class of utility functions for bidder i on auctions in \mathcal{M} , i.e.,

$$\mathcal{U}_i = \left\{ u_i : \mathcal{V}_i \times \mathcal{V} \times \mathcal{X}^n \times \mathcal{Y}^m \rightarrow \mathbb{R} \mid u_i(v_i, v, x, y) = \sum_{j=1}^m g_{ij}(v, x, y)v_{ij} - p_i(v, x, y) \right\}. \quad (31)$$

Similarly, let \mathcal{U} be the class of utility profiles over \mathcal{M} . Define the $\ell_{\infty,1}$ -distance between two utility profiles u and u' as $\max_{v,v',x,y} \sum_{i=1}^n |u_i(v_i, (v'_i, v_{-i}), x, y) - u'_i(v_i, (v'_i, v_{-i}), x, y)|$ and $\mathcal{N}_{\infty,1}(\mathcal{U}, r)$ as the minimum number of balls of radius $r > 0$ to cover \mathcal{U} (r -covering number of \mathcal{U}) under such $\ell_{\infty,1}$ -distance. We also define the ℓ_{∞} -distance between u_i and u'_i as $\max_{v,v'} |u_i(v_i, (v'_i, v_{-i}), x, y) - u'_i(v_i, (v'_i, v_{-i}), x, y)|$ and $\mathcal{N}_{\infty}(\mathcal{U}_i, r)$ as the r -covering number of \mathcal{U}_i under ℓ_{∞} -distance.

Covering Numbers $\mathcal{N}_{\infty,1}(\text{RGT}, r)$ and $\mathcal{N}_{\infty}(\text{RGT}_i, r)$. As for regret functions, let $\text{RGT}_i \circ \mathcal{U}_i$ be the class of all regret functions for bidder i , i.e.,

$$\text{RGT}_i \circ \mathcal{U}_i = \left\{ rgt_i : \mathcal{V} \times \mathcal{X}^n \times \mathcal{Y}^m \rightarrow \mathbb{R} \mid rgt_i(v, x, y) = \max_{v'_i \in \mathcal{V}_i} u_i(v_i, (v'_i, v_{-i}), x, y) - u_i(v_i, v, x, y) \text{ for some } u_i \in \mathcal{U}_i \right\}. \quad (32)$$

The same as before, we define $\text{RGT} \circ \mathcal{U}$ as the class of profiles of regret functions, and we define $\ell_{\infty,1}$ -distance between two regret profiles rgt and rgt' as $\max_{v,v,x,y} \sum_{i=1}^n |rgt_i(v, x, y) - rgt'_i(v, x, y)|$. Let $\mathcal{N}_{\infty,1}(\text{RGT} \circ \mathcal{U}, r)$ denote the r -covering number of $\text{RGT} \circ \mathcal{U}$ under such distance. Similarly, define the ℓ_{∞} -distance between rgt_i and rgt'_i as $\max_{v,x,y} |rgt_i(v, x, y) - rgt'_i(v, x, y)|$, and denote $\mathcal{N}_{\infty}(\text{RGT} \circ \mathcal{U}_i, r)$ as the r -covering number of RGT .

Covering Numbers $\mathcal{N}_{\infty,1}(\mathcal{P}, r)$ and $\mathcal{N}_{\infty}(\mathcal{P}_i, r)$. As for revenue (payment) functions, we denote the class of all the profiles of payment functions as \mathcal{P} and

$$\mathcal{P}_i = \{p_i : \mathcal{V} \times \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}_{\geq 0} \mid p \in \mathcal{P}\}. \quad (33)$$

We denote the r -covering number of \mathcal{P} as $\mathcal{N}_{\infty}(\mathcal{P}, r)$ under the $\ell_{\infty,1}$ -distance and the r -covering number for \mathcal{P}_i as $\mathcal{N}_{\infty}(\mathcal{P}_i, \epsilon)$ under the ℓ_{∞} -distance.

E.2. Important Lemmas

The generalization lemma (Lemma E.1) plays an important role in our proof.

Lemma E.1. Let $\mathcal{S} = \{z_1, \dots, z_L\} \in \mathcal{Z}^L$ be a set of samples drawn i.i.d. from some distribution \mathcal{D} over \mathcal{Z} . We assume $f(z) \in [a, b]$ for all $f \in \mathcal{F}$ and $z \in \mathcal{Z}$. Define the ℓ_{∞} -distance between two functions $f, f' \in \mathcal{F}$ as $\max_{z \in \mathcal{Z}} |f(z) - f'(z)|$ and define $\mathcal{N}_{\infty}(\mathcal{F}, r)$ as the r -covering number of \mathcal{F} under such ℓ_{∞} -distance. Let $\mathbb{L}_{\mathcal{D}}(f) = \mathbb{E}_{z \sim \mathcal{D}}[f(z)]$ and $\mathbb{L}_{\mathcal{S}}(f) = \frac{1}{|\mathcal{S}|} \sum_{i=1}^{|\mathcal{S}|} f(z_i)$, then we have

$$\mathbb{P}_{\mathcal{S} \sim \mathcal{D}^m} \left[\exists f \in \mathcal{F}, \left| \mathbb{L}_{\mathcal{S}}(f) - \mathbb{L}_{\mathcal{D}}(f) \right| > \epsilon \right] \leq 2\mathcal{N}_{\infty}(\mathcal{F}, \frac{\epsilon}{3}) \exp \left(-\frac{2L\epsilon^2}{9(b-a)^2} \right). \quad (34)$$

Proof. Define \mathcal{F}_r as the minimum function class that r -covers \mathcal{F} (so that $|\mathcal{F}_r| = \mathcal{N}_{\infty}(\mathcal{F}, r)$). For all function $f \in \mathcal{F}$, denote f_r as the closed function to f in such function class \mathcal{F}_r . On top of that, we have $|f(z) - f_r(z)| \leq r, \forall z \in \mathcal{Z}$. For all

$\epsilon > 0$, set $r = \frac{\epsilon}{3}$, we get

$$\begin{aligned}
 & \mathbb{P}_{S \sim \mathcal{D}^m} \left[\exists f \in \mathcal{F}, \left| \mathbb{L}_S(f) - \mathbb{L}_{\mathcal{D}}(f) \right| > \epsilon \right] \\
 & \leq \mathbb{P}_{S \sim \mathcal{D}^m} \left[\exists f \in \mathcal{F}, \left| \mathbb{L}_S(f) - \mathbb{L}_S(f_r) \right| + \left| \mathbb{L}_S(f_r) - \mathbb{L}_{\mathcal{D}}(f_r) \right| + \left| \mathbb{L}_{\mathcal{D}}(f_r) - \mathbb{L}_{\mathcal{D}}(f) \right| > \epsilon \right] \\
 & \leq \mathbb{P}_{S \sim \mathcal{D}^m} \left[\exists f \in \mathcal{F}, r + \left| \mathbb{L}_S(f_r) - \mathbb{L}_{\mathcal{D}}(f_r) \right| + r > \epsilon \right] \\
 & \leq \mathbb{P}_{S \sim \mathcal{D}^m} \left[\exists f_r \in \mathcal{F}_r, \left| \mathbb{L}_S(f_r) - \mathbb{L}_{\mathcal{D}}(f_r) \right| > \frac{1}{3}\epsilon \right], \quad r = \frac{\epsilon}{3} \\
 & \leq \mathcal{N}_{\infty}(\mathcal{F}, \frac{\epsilon}{3}) \mathbb{P}_{S \sim \mathcal{D}^m} \left[\left| \mathbb{L}_S(f) - \mathbb{L}_{\mathcal{D}}(f) \right| > \frac{1}{3}\epsilon \right] \\
 & \leq 2\mathcal{N}_{\infty}(\mathcal{F}, \frac{\epsilon}{3}) \exp\left(-\frac{2L\epsilon^2}{9(b-a)^2}\right). \quad (\text{Hoeffding Inequality})
 \end{aligned} \tag{35}$$

□

The following two lemmas (Lemma E.2 and Lemma E.3) provides the covering numbers bound for payment and regret.

Lemma E.2. $\mathcal{N}_{\infty,1}(\mathcal{P}, \epsilon) \leq \mathcal{N}_{\infty,1}(\mathcal{M}, \epsilon)$.

Proof. By the definition of the covering number for the auction class \mathcal{M} , there exists a cover $\hat{\mathcal{M}}$ for \mathcal{M} of size $|\hat{\mathcal{M}}| \leq \mathcal{N}_{\infty,1}(\mathcal{M}, \epsilon)$ such that for any $(g, p) \in \mathcal{M}$, there is a $(\hat{g}, \hat{p}) \in \hat{\mathcal{M}}$ for all v, x, y ,

$$\sum_{i,j} |g_{ij}(v, x, y) - \hat{g}_{ij}(v, x, y)| + \sum_i |p_i(v, x, y) - \hat{p}_i(v, x, y)| \leq \epsilon. \tag{36}$$

As a result, we can have $\hat{\mathcal{P}} = \{(\hat{g}, \hat{p}) \in \hat{\mathcal{M}}\}$, then for any $p \in \mathcal{P}$, there exist a $\hat{p} \in \hat{\mathcal{P}}$, for all v, x, y ,

$$\sum_i |p_i(v, x, y) - \hat{p}_i(v, x, y)| \leq \epsilon. \tag{37}$$

Therefore, we have $\mathcal{N}_{\infty,1}(\mathcal{P}, \epsilon) \leq \mathcal{N}_{\infty,1}(\mathcal{M}, \epsilon)$. □

Lemma E.3. $\mathcal{N}_{\infty,1}(\text{RGT} \circ \mathcal{U}, \epsilon) \leq \mathcal{N}_{\infty,1}(\mathcal{M}, \frac{\epsilon}{2n})$.

Proof. The proof then proceeds in two steps:

1. bounding the covering number for each regret class $\text{RGT} \circ \mathcal{U}$ in terms of the covering number for individual utility classes \mathcal{U} ;
2. bounding the covering number for the joint utility class \mathcal{U} in terms of the covering number for \mathcal{M} .

First we prove that $\mathcal{N}_{\infty,1}(\text{RGT} \circ \mathcal{U}, \epsilon) \leq \mathcal{N}_{\infty,1}(\mathcal{U}, \frac{\epsilon}{2})$.

By the definition of covering number $\mathcal{N}_{\infty,1}(\mathcal{U}, r)$, there exists a cover $\hat{\mathcal{U}}$ with size at most $\mathcal{N}_{\infty,1}(\mathcal{U}, \epsilon/2)$ such that for any $u \in \mathcal{U}$, there is a $\hat{u} \in \hat{\mathcal{U}}$ with

$$\max_{v, v', x, y} \sum_{i=1}^n \left| u_i(v_i, (v'_i, v_{-i}), x, y) - \hat{u}_i(v_i, (v'_i, v_{-i}), x, y) \right| \leq \frac{\epsilon}{2}. \tag{38}$$

For any $u \in \mathcal{U}$, taking $\hat{u} \in \hat{\mathcal{U}}$ satisfying the above condition, then for any v, x, y , we have

$$\begin{aligned}
 & \left| \max_{v'_i \in \mathcal{V}_i} (u_i(v_i, (v'_i, v_{-i}), x, y) - u_i(v_i, (v_i, v_{-i}), x, y)) - \max_{\bar{v}_i \in \mathcal{V}_i} (\hat{u}_i(v_i, (\bar{v}_i, v_{-i}), x, y) - \hat{u}_i(v_i, (v_i, v_{-i}), x, y)) \right| \\
 & \leq \left| \max_{v'_i} u_i(v_i, (v'_i, v_{-i}), x, y) - \max_{\bar{v}_i} \hat{u}_i(v_i, (\bar{v}_i, v_{-i}), x, y) + \hat{u}_i(v_i, (v_i, v_{-i}), x, y) - u_i(v_i, (v_i, v_{-i}), x, y) \right| \\
 & \leq \left| \max_{v'_i} u_i(v_i, (v'_i, v_{-i}), x, y) - \max_{\bar{v}_i} \hat{u}_i(v_i, (\bar{v}_i, v_{-i}), x, y) \right| + \left| \hat{u}_i(v_i, (v_i, v_{-i}), x, y) - u_i(v_i, (v_i, v_{-i}), x, y) \right| \\
 & \leq \left| \max_{v'_i} u_i(v_i, (v'_i, v_{-i}), x, y) - \max_{\bar{v}_i} \hat{u}_i(v_i, (\bar{v}_i, v_{-i}), x, y) \right| + \max_{v'_i} \left| u_i(v_i, (v'_i, v_{-i}), x, y) - \hat{u}_i(v_i, (v'_i, v_{-i}), x, y) \right|.
 \end{aligned} \tag{39}$$

Let $v_i^* \in \arg \max_{v'_i} u_i(v_i, (v'_i, v_{-i}), x, y)$ and $\hat{v}_i^* \in \arg \max_{\bar{v}_i} \hat{u}_i(v_i, (\bar{v}_i, v_{-i}), x, y)$, then

$$\begin{aligned}
 \max_{v'_i} u_i(v_i, (v'_i, v_{-i}), x, y) - \max_{\bar{v}_i} \hat{u}_i(v_i, (\bar{v}_i, v_{-i}), x, y) &= u_i(v_i, (v_i^*, v_{-i}), x, y) - \hat{u}_i(v_i, (\hat{v}_i^*, v_{-i}), x, y) \\
 &\leq u_i(v_i, (v_i^*, v_{-i}), x, y) - \hat{u}_i(v_i, (v_i^*, v_{-i}), x, y) \\
 &\leq \max_{v'_i} \left| u_i(v_i, (v'_i, v_{-i}), x, y) - \hat{u}_i(v_i, (v'_i, v_{-i}), x, y) \right| \\
 \max_{\bar{v}_i} \hat{u}_i(v_i, (\bar{v}_i, v_{-i}), x, y) - \max_{v'_i} u_i(v_i, (v'_i, v_{-i}), x, y) &= \hat{u}_i(v_i, (\hat{v}_i^*, v_{-i}), x, y) - u_i(v_i, (v_i^*, v_{-i}), x, y) \\
 &\leq \hat{u}_i(v_i, (\hat{v}_i^*, v_{-i}), x, y) - u_i(v_i, (\hat{v}_i^*, v_{-i}), x, y) \\
 &\leq \max_{v'_i} \left| u_i(v_i, (v'_i, v_{-i}), x, y) - \hat{u}_i(v_i, (v'_i, v_{-i}), x, y) \right|.
 \end{aligned} \tag{40}$$

Thus,

$$\begin{aligned}
 & \left| \max_{v'_i} (u_i(v_i, (v'_i, v_{-i})) - u_i(v_i, (v_i, v_{-i}))) - \max_{\bar{v}_i} (\hat{u}_i(v_i, (\bar{v}_i, v_{-i})) - \hat{u}_i(v_i, (v_i, v_{-i}))) \right| \\
 & \leq 2 \max_{v'_i} \left| u_i(v_i, (v'_i, v_{-i}), x, y) - \hat{u}_i(v_i, (v'_i, v_{-i}), x, y) \right|.
 \end{aligned} \tag{41}$$

Summing the inequalities by i , this completes the proof that $\mathcal{N}_{\infty,1}(\text{RG T} \circ \mathcal{U}, \epsilon) \leq \mathcal{N}_{\infty,1}(\mathcal{U}, \frac{\epsilon}{2})$.

Next we prove that $\mathcal{N}_{\infty,1}(\mathcal{U}, \epsilon) \leq \mathcal{N}_{\infty,1}(\mathcal{M}, \frac{\epsilon}{n})$.

By the definition of the covering number for the auction class \mathcal{M} , there exists a cover $\hat{\mathcal{M}}$ for \mathcal{M} of size $|\hat{\mathcal{M}}| \leq \mathcal{N}_{\infty,1}(\mathcal{M}, \frac{\epsilon}{n})$ such that for any $(g, p) \in \mathcal{M}$, there is a $(\hat{g}, \hat{p}) \in \hat{\mathcal{M}}$ for all v, x, y ,

$$\sum_{i,j} |g_{ij}(v, x, y) - \hat{g}_{ij}(v, x, y)| + \sum_i |p_i(v, x, y) - \hat{p}_i(v, x, y)| \leq \frac{\epsilon}{n}. \tag{42}$$

For all $v \in \mathcal{V}$, $v'_i \in \mathcal{V}_i$,

$$\begin{aligned}
 & \left| u_i(v_i, (v'_i, v_{-i}), x, y) - \hat{u}_i(v_i, (v'_i, v_{-i}), x, y) \right| \\
 & \leq \left| \sum_j (g_{ij}((v'_i, v_{-i}), x, y) - \hat{g}_{ij}((v'_i, v_{-i}), x, y)) v'_{ij} \right| + \left| p_i((v'_i, v_{-i}), x, y) - \hat{p}_i((v'_i, v_{-i}), x, y) \right| \\
 & \leq \sum_j \left| g_{ij}((v'_i, v_{-i}), x, y) - \hat{g}_{ij}((v'_i, v_{-i}), x, y) \right| + \left| p_i((v'_i, v_{-i}), x, y) - \hat{p}_i((v'_i, v_{-i}), x, y) \right| \\
 & \leq \frac{\epsilon}{n}.
 \end{aligned} \tag{43}$$

Thus,

$$\sum_{i=1}^n |u_i(v_i, (v'_i, v_{-i}), x, y) - \hat{u}_i(v_i, (v'_i, v_{-i}), x, y)| \leq n \cdot \frac{\epsilon}{n} = \epsilon. \tag{44}$$

This completes the proof that $\mathcal{N}_{\infty,1}(\mathcal{U}, \epsilon) \leq \mathcal{N}_{\infty,1}(\mathcal{M}, \frac{\epsilon}{n})$

Therefore,

$$\mathcal{N}_{\infty,1}(\text{RGT} \circ \mathcal{U}, \epsilon) \leq \mathcal{N}_{\infty,1}(\mathcal{U}, \frac{\epsilon}{2}) \leq \mathcal{N}_{\infty,1}(\mathcal{M}, \frac{\epsilon}{2n}). \quad (45)$$

This completes the proof of Lemma E.3. □

E.3. Proof of Theorem 2.6

proof of Theorem 2.6. For all $\epsilon, \delta \in (0, 1)$, when

$$L \geq \frac{9n^2}{2\epsilon^2} \left(\ln \frac{4}{\delta} + \ln \mathcal{N}_{\infty,1}(\mathcal{M}, \frac{\epsilon}{6n}) \right), \quad (46)$$

Combining Lemma E.1 and Lemma E.2 together, we get

$$\begin{aligned} & \mathbb{P}_{S \sim \mathcal{D}^m} \left[\exists (g^w, p^w) \in \mathcal{M}, \left| \sum_{i=1}^n \mathbb{E}_{(v,x,y) \sim \mathcal{D}_{v,x,y}} [p_i^w(v, x, y)] - \frac{1}{L} \sum_{i=1}^n \sum_{\ell=1}^L p_i^w(v^{(\ell)}, x^{(\ell)}, y^{(\ell)}) \right| > \epsilon \right] \\ & \leq 2\mathcal{N}_{\infty}(\mathcal{P}, \frac{\epsilon}{3}) \exp\left(-\frac{2L\epsilon^2}{9n^2}\right) \\ & \leq 2\mathcal{N}_{\infty}(\mathcal{M}, \frac{\epsilon}{3}) \exp\left(-\frac{2L\epsilon^2}{9n^2}\right) \\ & \leq 2\mathcal{N}_{\infty}(\mathcal{M}, \frac{\epsilon}{6n}) \exp\left(-\frac{2L\epsilon^2}{9n^2}\right) \\ & \leq \frac{\delta}{2}. \end{aligned} \quad (47)$$

Similarly, combining Lemma E.1 and Lemma E.3 together, we have

$$\begin{aligned} & \mathbb{P}_{S \sim \mathcal{D}^m} \left[\exists (g^w, p^w) \in \mathcal{M}, \left| \mathbb{E}_{(v,x,y) \sim \mathcal{D}_{v,x,y}} \left[\sum_{i=1}^n \text{rgt}_i(w) \right] - \sum_{i=1}^n \widehat{\text{rgt}}_i(w) \right| > \epsilon \right] \\ & \leq 2\mathcal{N}_{\infty}(\text{RGT} \circ \mathcal{U}, \frac{\epsilon}{3}) \exp\left(-\frac{2L\epsilon^2}{9n^2}\right) \\ & \leq 2\mathcal{N}_{\infty}(\mathcal{M}, \frac{\epsilon}{6n}) \exp\left(-\frac{2L\epsilon^2}{9n^2}\right) \\ & \leq \frac{\delta}{2}. \end{aligned} \quad (48)$$

Combining Equation (47), Equation (48) and the Union Bound, with probability at most $\frac{\delta}{2} + \frac{\delta}{2} = \delta$, one of the two events of Equation (47) and Equation (48) happens. Therefore, with probability at least $1 - \delta$, Equation (7) and Equation (8) both hold. We complete the proof of Theorem 2.6. □