On Designing the Optimal Integrated Ad Auction in E-commerce Platforms

Yuchao Ma[1*], Weian Li[2*], Yuhan Wang[1], Zitian Guo[1], Yuejia Dou[1], Qi Qi[1†], Changyuan Yu[3]

[1]Gaoling School of Artificial Intelligence, Renmin University of China, Beijing, China
[2]School of Software, Shandong University, Jinan, China
[3]Baidu Inc., Beijing, China
{yc.ma, yuhanwang, guo_zt, douyuejia, qi.qi}@ruc.edu.cn, weian.li@sdu.edu.cn, yuchangyuan@baidu.com

Abstract

Currently, e-commerce platforms integrate ads and organic content into a mixed list for users. While platforms seek to maximize profit from advertisers, organic items enhance user experience. To ensure long-term development, platforms aim to design mechanisms that optimize both revenue and user satisfaction. Current methods rank ads and organic items separately before integrating them. Even if each part is locally optimal, the combined result may not be globally optimal. In this paper, we come up with the **J**oint **INTE**grated **R**egret **Net**work (JINTER Net). Unlike traditional methods, which pre-order ads and organic items separately, JINTER Net directly selects from the combined set of candidate ads and organic items to generate an optimal list. This approach aims to optimally balance platform revenue and user experience while satisfying approximate dominant strategy incentive compatibility and individual rationality. We validate the effectiveness of JINTER Net using both synthetic data and real dataset, and our experimental results show that it significantly outperforms baseline models across multiple metrics.

Introduction

Over the past two decades, online advertising has become the primary revenue stream for many internet platforms, reaching a staggering global value of $225 billion by 2023. Currently, most advertisements are allocated through auction systems known as sponsored search auctions. When a user performs a search or query, platforms conduct these auctions among relevant advertisers and display the winning ads on the search results page. With billions of searches occurring daily due to the rapid growth of the internet, these platforms generate substantial revenue.

In addition to advertisements, search result pages also feature organic items, which are allocated by the recommendation system. Unlike auction systems that focus on maximizing immediate revenue, recommendation systems aim to enhance user experience, thereby supporting the platform's long-term profitability. Consequently, from the platform's perspective, it is crucial to optimize the placement of both ads and organic items on search result pages to balance these two often conflicting objectives.

[*]These authors contributed equally.
[†]Corresponding Author.

Traditionally, platforms allocate a fixed number of top slots for advertisements. The auction system and recommendation system then independently select items from their respective candidate sets to fill these slots. While both systems can operate optimally within their domains using established techniques, the overall display on the search result page may not be globally optimal. Recently, advanced machine learning methods have offered new perspectives on integrating ads and organic items. Some Internet companies have begun incorporating auction-winning ads into the recommendation system's candidate set, using reinforcement learning to generate a complete search result page. This approach addresses the challenge of determining the optimal number of fixed ad slots. However, these methods still fall short of providing an integrated treatment of both ad and organic candidate sets and lack robust theoretical support.

inFrom the perspective of mechanism design, this problem presents unique challenges that make conventional mechanisms potentially inadequate. Unlike classic auction settings where goods are predetermined, ads and organic items compete for slots, and only ads submit bids. This means we cannot pre-determine which subset of slots will be occupied by bidders. Furthermore, when integrating ads and organic items into a single list, designing an optimal mechanism requires incorporating Myerson's auction framework while accounting for the externalities introduced by organic items. This complexity makes the calculation of advertisers' payments highly intricate.

Despite these theoretical challenges, recent advancements in machine learning have introduced the concept of automated mechanism design. This approach transforms the mechanism design problem into a learning problem, using neural network architectures to determine allocation and payment rules. However, current neural network architectures are not yet equipped to handle the complexities of this problem effectively. Thus, finding an optimal method to blend ads with organic items in a way that balances revenue and user experience remains a significant and meaningful challenge for both academia and industry.

Main Contributions

To address the research problem outlined, this paper examines a scenario where multiple advertisers and organic items compete for various heterogeneous slots, each with different click-through rates. Since the blending display of ads and organic items are common on e-commerce platforms, we use the concept of “merchandise volume” to represent the user experience associated with each item, while gross merchandise volume (GMV) measures the total user experience. Merchandise volume encompasses various aspects such as purchase tendency and the decision-making process based on price and product quality. E-commerce platforms often use GMV to gauge their market presence. Our objective is to design an optimal integrated ad auction that balances revenue and GMV while ensuring some key properties in the field of auction theory, such as dominant strategy incentive compatibility (DSIC) and individual rationality (IR). However, many conventional and commonly-used mechanisms may not be suitable for this challenge. To overcome these challenges, we introduce the **J**oint **INTE**grated **R**egret **Net**work (JINTER Net), a neural network architecture specifically designed to generate the optimal mechanisms. We evaluate JINTER Net through extensive experiments on both synthetic data and real dataset, demonstrating its superior performance compared to established benchmarks. Our key contributions are as follows.

JINTER Net For the integrated ad auction design problem, in this work, we generate the optimal integrated auction by JINTER Net, a neural network architecture. The generated auction satisfies the approximate DSIC and IR, and optimally balances revenue and user experience. The distinct features of JINTER Net specially designed for this problem, can be summarized as:

- **Joint relationship graph between bids and volumes.** Since advertisements contribute to both revenue and GMV, whereas organic items contribute only to GMV, JINTER Net addresses this by incorporating a joint relationship graph of bids and volumes as input. This graph is crucial throughout all stages of JINTER Net, as it determines the contribution of each item.
- **Allocation probability matrix based on bids and volumes.** Based on the relationship graph, an item can be viewed as a bundle consisting of a bid node[1] and a volume node. JINTER Net calculates the allocation probabilities for these bundles, which are equivalent to the allocation probabilities for individual items. This approach addresses the challenge of reconciling the dimensions of contributions from ads and organic items.
- **Payment rule with externality.** Determining payments for each advertiser is complex, particularly when considering the externalities of organic items. To address this, JINTER Net introduces a parameter that scales each advertiser’s total expected value based on the allocation probability matrix (derived from the bundle’s allocation probability) and defines this scaled value as the payment. This parameter is optimized during training while ensuring IR for all advertisers. Since the allocation probability matrix is derived from the bids and volumes, this payment rule naturally incorporates externalities.

[1]If the item is an organic item, there exists a dummy node linking its volume node.

Related Works

With the prevalence of the Internet, ad auctions have become a highly regarded research direction in the field of mechanism design. Classic mechanisms, such as Generalized First Price (GFP) and Generalized Second Price (GSP), have been widely applied in internet ad systems due to their computational simplicity (Edelman, Ostrovsky, and Schwarz 2007; Varian 2007; Han and Liu 2015; Despotakis, Ravi, and Sayedi 2021). Our work mainly focuses on the intersection of ad systems and recommendation systems in online platforms, specifically in optimally integrating ads and organic items into a list to present to users. A series of studies have modeled this problem as a Markov decision process, using reinforcement learning to generate the optimal list (Wang et al. 2019; Zhao et al. 2020, 2021; Liao et al. 2022b). Additionally, there has been a few of work investigating integrated ad auction design in theory. Li et al. (2023) propose a truthful optimal mechanism to balance platform’s revenue and the user experience. Li et al. (2024) focuses on the setting that each candidate item can be displayed in ad or organic and design two mechanisms to solve the problem.

Meanwhile, with the continuous development of machine learning, automated mechanism design (AMD) has gained widespread attention, aiming to design the approximately optimal auction mechanisms. The key approach in AMD is to parameterize the auction and optimize it automatically through machine learning. One technical route is to parameterize the VCG mechanism (Vickrey 1961; Clarke 1971; Groves 1973), such as virtual valuations combinatorial auctions, affine maximizer auctions (Likhodedov and Sandholm 2004; Likhodedov, Sandholm et al. 2005; Sandholm and Likhodedov 2015) and Lottery AMA (Curry, Sandholm, and Dickerson 2023). The technical route we adopt is based on RegretNet (Dütting et al. 2019), which transforms the DSIC condition into a quantifiable regret, and then designs allocation and payment networks to maximize revenue while satisfying approximate DSIC and IR. Based on RegretNet, a few of works have explored other scenarios, integrating budget (Feng, Narasimhan, and Parkes 2018), incorporating fairness (Kuo et al. 2020), focusing on human preferences (Peri et al. 2021). Moreover, some works have modified the structure of RegretNet, such as applying transformer in contextual auctions (Duan et al. 2022) and attention mechanism in the optimal auction design (Ivanov et al. 2022). Besides, certain works employ data-driven methods to design the optimal auction (Liu et al. 2021; Liao et al. 2022a).

However, the methods described are suitable only for bidders and cannot handle candidates with ads and organic items. Specifically, these architectures do not account for the externalities of organic items. In contrast, JINTER Net overcomes this limitation and generates an optimal integrated ad mechanism while ensuring approximate DSIC and IR.

Model and Preliminaries

In this section, we formally introduce our model. We begin by defining the basic concepts of integrated ad auctions. Then we transform our problem into a learning problem, enabling the implementation of automated mechanism design.

Integrated Ad Auction

In this paper, we focus on a search result page in modern e-commerce platforms, where advertisements and organic items are mixed for display to users. We consider a page with $K \in \mathbb{N}_+$ heterogeneous display slots. For each slot $k \in [K] = \{1, 2, \cdots, K\}$, let β_k represent the click-through rate (CTR) of that slot[2]. Without loss of generality, assume that $\beta_1 \geq \beta_2 \geq \cdots \geq \beta_K > 0$, reflecting the heterogeneity of the slots. There are m candidate advertisers and n candidate organic items competing for these K ad slots. We use $[m] = \{1, 2, \cdots, m\}$ and $[n] = \{1, 2, \cdots, n\}$ to denote the sets of advertisers and organic items, respectively.

We focus on the pay-per-click model, where the advertiser only pay when her ad is clicked by a user. For each advertiser $i \in [m]$, let v_i denote her private value per click, representing the amount she is willing to pay for an additional click. This value is independently (and possibly non-identically) drawn from a publicly known distribution F_i. Define b_i as the bid of advertiser i, which depends on her value v_i. Let $\mathbf{b} = \{b_1, b_2, \cdots, b_m\}$ and $\mathbf{v} = \{v_1, v_2, \cdots, v_m\}$ represent the bid profile and value profile of all advertisers, respectively. Similarly, $\mathbf{b}_{-i}$ and $\mathbf{v}_{-i}$ represent the bid profile and value profile of all advertisers except for i.

To account for user experience, we use a metric called "merchandise volume" to reflect the user experience of an item. For each advertiser $i \in [m]$ and organic item $j \in [n]$, let g_i and g_j represent the expected merchandise volume per click, respectively. We denote the volume profile by $\mathbf{g} = \{g_1, \cdots, g_m, g_{m+1}, \cdots, g_{m+n}\}$, which includes both advertisers and organic items. Unlike bids, which are private, volumes are public information for the platform. While volume is one way to represent user experience, it can be substituted with other related metrics without impacting the integrated mechanisms.

An integrated ad auction, denoted by $\mathcal{M} = (\mathbf{x}, \mathbf{p})$, consists of an allocation rule $\mathbf{x}(\mathbf{b}, \mathbf{g}) = \{x_i(\mathbf{b}, \mathbf{g})\}_{i \in M \cup N}$ and a payment rule $\mathbf{p}(\mathbf{b}, \mathbf{g}) = \{p_i(\mathbf{b}, \mathbf{g})\}_{i \in M}$. Specifically, the allocation rule, $\{x_i(\mathbf{b}, \mathbf{g})\} = \sum_{k \in K} x_{ik}(\mathbf{b}, \mathbf{g})\beta_k$, represents the expected CTR for each item i, where $x_{ik}(\mathbf{b}, \mathbf{g}) \in \{0, 1\}$ is an indicator function denoting whether item i is allocated to slot k. The payment rule $\{p_i(\mathbf{b}, \mathbf{g})\}_{i \in M}$ specifies the amount that the advertiser i will pay when her ad is clicked.

Since each slot can be allocated to only one item and each item can occupy just one slot, the allocation rule $x_i(\mathbf{b}, \mathbf{g})$ must satisfy the following feasibility constraints:

$$\sum_{k \in [K]} x_{ik}(\mathbf{b}, \mathbf{g}) \leq 1, \quad \forall i \in [m] \cup [n],$$
$$\sum_{i \in [m] \cup [n]} x_{ik}(\mathbf{b}, \mathbf{g}) = 1, \quad \forall k \in [K],$$
$$x_{ik}(\mathbf{b}, \mathbf{g}) \in \{0, 1\}, \quad \forall i \in [m] \cup [n], k \in [K].$$

[2]Note that we use separable CTR for clarity in introducing our model. Alternatively, there is the inseparable CTR, where β_{ik} denotes the CTR when item i is assigned to slot k and may not satisfy monotonicity. Importantly, the results and methods presented in this paper can also be adapted to inseparable CTR with minor modifications.

With the definition of an integrated ad auction established, we can define the utility of an advertiser as the difference between the value derived from the allocated slot and the payment made, i.e.,

$$u_i(\mathbf{b}, \mathbf{g}) = v_i \cdot x_i(\mathbf{b}, \mathbf{g}) - p_i(\mathbf{b}, \mathbf{g}).$$

In our model, we focus on two key economic properties: dominant strategy incentive compatibility (DSIC) and individual rationality (IR). DSIC ensures that truthfully bidding is the optimal strategy for participants, while IR guarantees that each bidder's utility is non-negative. The precise definitions of these concepts are as follows:

Definition 1 (Dominant Strategy Incentive Compatibility) *An integrated auction is dominant strategy incentive compatible, if for any advertiser, her utility is maximized by truthfully telling. Formally, for any advertiser $i \in [m]$, any $\mathbf{b}_{-i}$ and any misreporting bid b_i', it holds that*

$$u_i(v_i, \mathbf{b}_{-i}, \mathbf{g}) \geq u_i(b_i', \mathbf{b}_{-i}, \mathbf{g}).$$

Definition 2 (Individual Rationality) *An integrated auction is ex-post individually rational if for any advertiser, she always can receive a nonzero utility when participating, i.e.,*

$$u_i(b_i, \mathbf{b}_{-i}, \mathbf{g}) \geq 0.$$

From the platform's perspective, we focus on two optimization metrics: revenue and gross merchandise volume (GMV). For an integrated mechanism $\mathcal{M} = (\mathbf{x}, \mathbf{p})$ that satisfies DSIC and IR, the expected revenue and GMV of the platform are:

$$\mathrm{Rev}(\mathbf{v}, \mathbf{g}) := \mathbb{E}_{\mathbf{v} \sim F}[\sum_{i \in [m]} p_i(\mathbf{v}, \mathbf{g})]$$
$$\mathrm{GMV}(\mathbf{v}, \mathbf{g}) := \mathbb{E}_{\mathbf{v} \sim F}[\sum_{i \in [m] \cup [n]} g_i \cdot x_i(\mathbf{v}, \mathbf{g})],$$

where $F = \times_{i \in [m]} F_i$ is the joint distribution over all advertisers' values.

We aim to optimize a blend of revenue and GMV while ensuring DSIC and IR. We introduce a hyperparameter $\alpha > 0$ to balance revenue and GMV, and formulate our problem as a constrained optimization problem:

$$\max_{(\mathbf{x}, \mathbf{p})} \quad \mathbb{E}_{\boldsymbol{v} \sim F}\left[\sum_{i \in [m]} p_i(\mathbf{v}, \mathbf{g}) + \alpha \sum_{i \in [m] \cup [n]} g_i \cdot x_i(\mathbf{v}, \mathbf{g})\right]$$
$$\text{s.t.} \quad \text{DSIC, IR, Feasibility.}$$

Integrated Ad Auction Design as a Learning Problem

We formulate our problem as a learning problem and introduce the concept of ex-post regret. Ex-post regret is defined as the maximum increase in utility that an advertiser i could achieve by misreporting her bid, while keeping the bidding profiles of all other advertisers fixed:

$$rgt_i(\mathbf{v}) = \max_{v_i' \in F_i} [u_i(v_i; (v_i', \mathbf{v}_{-i})) - u_i(v_i; \mathbf{v})]$$

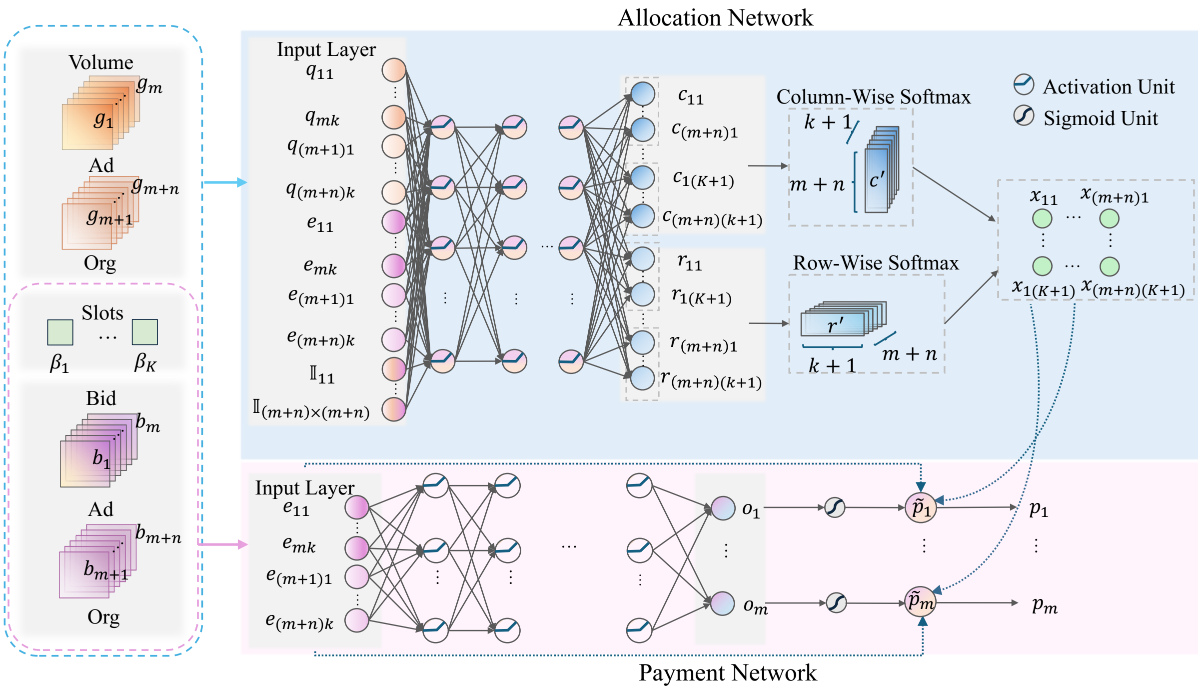

Figure 1: A schematic diagram of JINTER Net, including the input layer, the allocation network and the payment network, where m ads, n organic items and K slots serve as input.

According to the definition of regret, an integrated ad auction is DSIC if and only if

$$\mathbb{E}_{\mathbf{v}\sim F}[\sum_{i\in[m]} rgt_i(\mathbf{v})] = 0.$$

Then we can rewrite the problem of designing the optimal integrated ad auction as a constrained optimization:

$$\min_{(\mathbf{x},\mathbf{p})\in\mathcal{M}} -\mathbb{E}_{\mathbf{v}\sim F}\left[\sum_{i\in[m]} p_i(\mathbf{v},\mathbf{g}) + \alpha \sum_{i\in[m]\cup[n]} g_i \cdot x_i(\mathbf{v},\mathbf{g})\right]$$
$$\text{s.t. } \mathbb{E}[\sum_{i\in[m]} rgt_i(\mathbf{v})] = 0,$$

where $\mathcal{M}$ is the set of all integrated auction mechanisms that satisfy IR and feasibility. Due to the complexity of the constraints, solving this optimization problem is generally intractable (Conitzer and Sandholm 2002). To address this, we parameterize the auction mechanism as $\mathcal{M}^w = (x^w, p^w)$, where $w \in \mathbb{R}^d$ represents the parameters of the neural network, and d is the dimension of parameters w. Our goal is to compute the mechanism $\mathcal{M}^w$ that maximizes the objective $\mathbb{E}_{\mathbf{v}\sim F}\left[\sum_{i\in[m]} p_i^w(\mathbf{v},\mathbf{g}) + \alpha \sum_{i\in[m]\cup[m]} g_i \cdot x_i^w(\mathbf{v},\mathbf{g})\right]$, while ensuring DSIC and IR, by optimizing parameters w.

Given a sample set $\mathcal{L}$, consisting of L value profiles drawn from the joint distribution F, the empirical ex-post regret of advertiser i under the mechanism $\mathcal{M}^w(x^w, p^w)$ is:

$$\widehat{rgt}_i(w) := \frac{1}{L}\sum_{\ell=1}^{L} \max_{v_i'\in F_i}[u_i^w(v_i^{(\ell)}; (v_i', \mathbf{v}_{-i}^{(\ell)})) - u_i^w(v_i^{(\ell)}; \mathbf{v}^{(\ell)})].$$

Note that $\mathbf{g}$ is fixed in advance and does not require sampling, so we usually ignore $\mathbf{g}$ in x^w and p^w, when there is no ambiguity. We can then formulate our constrained optimization problem as a learning problem:

$$\min_{w\in\mathbb{R}^{d_w}} -\frac{1}{L}\sum_{\ell=1}^{L}\left[\sum_{i\in[m]} p_i^w(\mathbf{v}^{(\ell)}) + \alpha \sum_{i\in[m]\cup[n]} g_i x_i^w(\mathbf{v}^{(\ell)})\right]$$
$$\text{s.t. } \widehat{rgt}_i(w) = 0, \qquad \forall i \in M$$

Additionally, we ensure that IR is met through our network architecture, which we will introduce in the following section.

JINTER Net

In this section, we present the **J**oint **INTE**grated **R**egret **Net**work (JINTER Net) architecture, outlining it in two main parts: the network structure and the training process.

Overall Structure

The JINTER Net consists of two primary components: an allocation network that determines the allocation probabilities for ads and organic items, and a payment network that calculates the payments for ads. Figure 1 illustrates the detailed structure of JINTER Net, and the following parts will describe its core components.

Input and Output of JINTER Net Given that each advertiser i has both a bid and a volume, while an organic item j is evaluated solely by its volume, we assume that organic item j is capable of bidding with a bid set to zero but does not engage in any strategic behavior and payment. In other words, organic item j will not misreport its bid. Based on this assumption, the input layer of the JINTER Net consists of q_{ik}, e_{ik} and $\mathbb{1}_{ij}$, defined as:

$$\begin{cases} q_{ik} = g_i \cdot \beta_k, & \forall i \in [m] \cup [n], \forall k \in [K], \\ e_{ik} = b_i \cdot \beta_k, & \forall i \in [m], \forall k \in [K], \\ e_{jk} = 0, & \forall j \in [n], \forall k \in [K]. \end{cases}$$

Specifically, q_{ik} and e_{ik} denote the expected volume and bid for ad (or organic item) i in slot k, respectively. The term $\mathbb{1}_{ij}$ represents the relationship between the bid and volume of all ads and organic items. Since each ad i (or organic item j) has a unique bid and volume profile that do not overlap with others, the joint relationship can only be formed by a bid and volume within the same ad or organic item. Thus, $\{\mathbb{1}_{ij}\}_{i,j \in [m]\cup[n]}$ forms a diagonal matrix used as input.

The JINTER Net generates two key outputs: allocation results and corresponding payments for ads. Allocation network produces an allocation probability matrix $\mathbf{x}$ for all ads and organic items. Each element x_{ik} in $\mathbf{x}$ represents the probability that ad (or organic item) i is allocated to slot k. Payment network provides a payment vector $\mathbf{p}_{1\times m}$ for all ad items. Each element $p_i \in \mathbf{p}$ denotes the payment required from the advertiser i, when her ad is clicked.

Allocation Network The allocation network processes the bidding and volume profiles as inputs, combining them into $m + n$ bundles based on the joint relationship of ads and organic items. It then passes these bundles through a fully connected neural network, producing two matrices, $\mathbf{C}$ and $\mathbf{R}$. The matrices $\mathbf{C}$ and $\mathbf{R}$ each have $K + 1$ columns. The additional column represents a dummy neuron, accounting for the possibility that an item might not be assigned to any slot. The dummy neuron's output indicates the probability that the item remains unassigned.

To ensure feasibility, we apply row-wise softmax on $\mathbf{R}$ to obtain $\mathbf{R}'$, which satisfies the constraint that each slot is assigned to only one item. We then apply column-wise softmax on $\mathbf{C}$ to obtain $\mathbf{C}'$, which ensures that each item is allocated to at most one slot. The final allocation matrix $\mathbf{x}$ is computed as the element-wise minimum of $\mathbf{R}'$ and $\mathbf{C}'$.

Although $\mathbf{x}$ is generated for $m + n$ bundles, due to the joint relationship being confined within the same ad or organic item, it can be equivalently interpreted as the allocation probability matrix for the items. The rigorous formula for the allocation matrix $\mathbf{x}$ is as follows:

$$x_{ik} = \min\{r'_{ik}, c'_{ik}\}, \quad \forall i \in [m] \cup [n], k \in [K + 1],$$

where $r'_{ik} = \frac{e^{r_{ik}}}{\sum_{t\in[m]\cup[n]} e^{r_{tk}}}$ and $c'_{ik} = \frac{e^{c_{ik}}}{\sum_{t\in[K+1]} e^{c_{it}}}$ serve as the row-wise normalization of r_{ik} and column-wise normalization of c_{ik}.

Payment Network After obtaining the final allocation probability matrix $\mathbf{x}$ from the allocation network, we propagate it to the payment network to calculate the payment vector for the ads. The payments in $\mathbf{p}$ are calculated as follows:

$$p_i = \tilde{p}_i \cdot \sum_{k=1}^{K} e_{ik} x_{ik}, \quad \forall i \in [m],$$

where $\tilde{p}_i \in [0, 1]$ is a normalization factor computed using a sigmoid function. This payment rule ensures that the utility of each advertiser, $u_i = \sum_{k\in[K]} e_{ik}x_{ik} - p_i$, is non-negative, thus satisfying the IR condition.

Training Procedure

In the training process of JINTER Net, we use the augmented Lagrangian method to convert the original constrained optimization problem into an unconstrained problem within the w parameter space of the neural network. The augmented Lagrangian function is defined as follows:

$$\mathcal{C}_\rho(w; \boldsymbol{\lambda}) = -\frac{1}{L}\sum_{\ell=1}^{L}\left[\sum_{i=1}^{m} p_i^w(\mathbf{v}^{(\ell)}) + \alpha \sum_{j=1}^{m+n} g_j x_j^w(\mathbf{v}^{(\ell)})\right] + \sum_{i=1}^{m} \lambda_i \widehat{rgt}_i(w) + \frac{\rho}{2}\sum_{i=1}^{m}(\widehat{rgt}_i(w))^2,$$

where $\boldsymbol{\lambda} \in \mathbb{R}^n$ represents the Lagrangian multipliers and $\rho > 0$ is the penalty weight factor on quadratic term.

We summarize the training procedure for JINTER Net as follows. We begin by dividing the dataset $\mathcal{L}$ into minibatches of size B. Over T iterations, each round processes a randomly selected minibatch $\mathcal{L}_t$, which is denoted by $\mathcal{L}_t = \{v^{(1)}, \cdots, v^{(B)}\}$. During each training round, we first compute the optimal misreport values. Notably, since organic items do not engage in strategic behavior, optimal misreports are calculated only for ad items. In the t-th iteration, we initialize a set of misreport values for ad items randomly and then optimize these values using gradient ascent method. The updating procedure is as follows:

$$v_i'^{(\ell)} = v_i'^{(\ell)} + \gamma \nabla_{v_i'}\left[u_i^w\left(v_i^{(\ell)}; \left(v_i', \mathbf{v}_{-i}^{(\ell)}\right)\right)\right]\Big|_{v_i' = v_i'^{(\ell)}}.$$

We then update the Lagrange multipliers and model parameters alternately:

$$\begin{cases} \lambda_i^{t+1} \leftarrow \lambda_i^t + \rho_t \widetilde{rgt}_i\left(w^{t+1}\right), & \forall i \in [m] \\ w^{t+1} \leftarrow w^t - \eta \nabla_w \mathcal{C}_{\rho_t}\left(w^t, \lambda^t\right) \end{cases}$$

where $\{\rho_t\}_{t\in[T]}$ and η are given parameters. During the process of updating λ, we calculate the empirical regret $\widetilde{rgt}_i$ based on the t-th round minibatch $\mathcal{L}_t$.

Given $\boldsymbol{\lambda}^t$, the gradient of the loss function with respect to parameters w, is given by

$$\begin{aligned} &\nabla_w \mathcal{C}_\rho\left(w; \boldsymbol{\lambda}^t\right) \\ &= -\frac{1}{B}\sum_{\ell=1}^{B}\left[\sum_{i=1}^{m} \nabla_w p_i^w\left(\mathbf{v}^{(\ell)}\right) + \alpha \sum_{j=1}^{m+n} g_j \nabla_w x_j^w\left(\mathbf{v}^{(\ell)}\right)\right] \\ &+ \sum_{\ell=1}^{B}\left[\sum_{i=1}^{m} \lambda_i^t h_{\ell,i} + \rho \sum_{i=1}^{m+n} \widetilde{rgt}_i(w) h_{\ell,i}\right], \end{aligned}$$

where

$$h_{\ell,i} = \nabla_w\left[\max_{v_i' \in V_i} u_i^w\left(v_i^{(\ell)}; \left(v_i', \mathbf{v}_{-i}^{(\ell)}\right)\right) - u_i^w\left(v_i^{(\ell)}; \mathbf{v}^{(\ell)}\right)\right].$$

Method	A: $2 \times 3 \times 2$				B: $3 \times 2 \times 2$			
	SW	Rev	GMV	Rev+αGMV	SW	Rev	GMV	Rev+αGMV
GSP and Fixed Positions	**0.466**	0.232	0.625	0.858	0.45	0.301	0.458	0.760
IAS	0.356	0.300	**0.671**	0.971	**0.476**	**0.380**	0.475	0.855
JINTER Net	0.426	**0.351**	0.657	**1.008**†	0.425	0.328	**0.611**	**0.939**†

Table 1: The notation $m \times n \times K$ represents a setting where m advertisers and n organic items compete for K slots. In the metric Rev+αGMV, the hyperparameter α is set to 1 across different settings. For the GSP and Fixed Positions mechanisms, the first slot is reserved for advertisements. The regret of mechanism generated by the JINTER Net is less than 0.001. The best performance is highlighted in bold. "†" indicates a statistically significant improvement in a paired t-test at $p < 0.05$ level.

Experiments

In this section, we conduct a series of experiments to validate the superiority of JINTER Net in integrated ad system. Our experiments were run on a Linux server equipped with NVIDIA Graphics Processing Unit (GPU) cores.

Experimental Settings

Baseline Methods We compare the JINTER Net with the following two representative auction mechanisms that can apply to the integrated ad system:

- **GSP** (Varian 2007) with **Fixed Positions**, a commonly-used method in reality, is executed in two steps. First, the number and positions of ad slots are predetermined. In the initial phase, the order of organic items is established based on their volumes. In the second phase, ads are ranked through a GSP auction; the winning ads are then placed into the predefined slots and charged according to the next highest bid.
- **IAS** (Li et al. 2023), a Myerson-based mechanism, ranks ads and organic items using a ranking score of $\phi_i(v_i) + \alpha g_i$, where $\phi_i(v_i) = v_i - (1-F_i(v_i))/f_i(v_i)$ is the virtual value. Payments are determined according to the Myerson payment rule (Myerson 1981).

Evaluation Metrics To assess the performance of JINTER Net and other baselines, we evaluate the empirical social welfare: SW $= \frac{1}{L}\sum_{\ell=1}^{L}\sum_{i\in[m]} v_i^{(\ell)} x_i^w(\mathbf{v}^{(\ell)})$, the the empirical revenue: Rev $= \frac{1}{L}\sum_{\ell=1}^{L}\sum_{i\in[m]} p_i^w(\mathbf{v}^{(\ell)})$, the empirical GMV: GMV $= \frac{1}{L}\sum_{\ell=1}^{L}\sum_{i\in[m]\cup[n]} g_i x_i^w(\mathbf{v}^{(\ell)})$ and the empirical combination of revenue and GMV with the hyperparameter α: Rev + αGMV $= \frac{1}{L}\sum_{\ell=1}^{L}[\sum_{i\in[m]} p_i^w(\mathbf{v}^{(\ell)}) + \alpha\sum_{i\in[m]\cup[n]} g_i x_i^w(\mathbf{v}^{(\ell)})]$.

Synthetic Data

For the synthetic data, we generate a training sample of 640,000 value profiles and a testing sample of 25,600 profiles. Both the training and testing sets use a minibatch size of 128. During the training of JINTER Net, we update w^t for each minibatch using the Adam optimizer with a learning rate of 0.001. For testing JINTER Net, we initiate 100 misreports and perform gradient ascent method on them for 2,000 iterations to obtain 100 empirical regrets.

We conduct several experiments across JINTER Net and baseline methods in the following settings:

(A) 2 ads, 3 organic items and 2 slots with CTR $\beta = (0.7, 0.3)$. The value of each ad is independently drawn from $U[0,1]$. The volumes of these 2 ads and 3 organic items are [0.6, 0.5] and [0.8, 0.7, 0.6], respectively.

(B) 3 ads, 2 organic items and 2 slots with CTRs $\beta = (0.6, 0.4)$. The value of each ad is independently drawn from $U[0,1]$. The volumes of these 3 ads and 2 organic items are [0.6, 0.5, 0.4] and [0.7, 0.6], respectively.

The results under Settings (A) and (B) are summarized in Table 1. Notably, the integrated mechanism generated by JINTER Net consistently achieves a significantly higher value for the Rev+αGMV metric compared to other mechanisms, while maintaining a regret lower than 0.001. This highlights that the JINTER Net architecture performs effectively on blending ads and organic items.

To further validate the superiority of JINTER Net, we conduct experiments across various settings. These include different value distributions, different hyperparameter α-values, different numbers of slots, and different ratios of candidate ads to organic items. The results for the first two experiments are presented in the main body of the paper, while the results for the latter two are provided in the appendix. Below are the detailed experimental setups.

Different Value Distributions To demonstrate the generalization capability of JINTER Net across various value distributions, we select Setting (B) and perform three sets of experiments with different distributions. Each experiment is repeated five times. We sample the value profiles from three distinct distributions: a uniform distribution $U[0,1]$; a normal distribution $N(0.5, 1)$ truncated to the $[0,1]$ interval, and a lognormal distribution $LN(0.2, 1.69)$ also truncated to the $[0,1]$ interval.

The experimental results presented in Table 2 show that JINTER Net consistently achieves the highest Rev+αGMV across three different value distributions. This performance significantly surpasses that of the other two mechanisms, demonstrating the stability and robustness of JINTER Net in handling various value distributions.

Hyper-parameter Analysis To assess the impact of the hyperparameter α on the experimental outcomes, we adjust α within the interval $[0.5, 2.0]$ based on Setting (A). We then

Method	Uniform				Normal				Lognormal			
	SW	Rev	GMV	Rev+ αGMV	SW	Rev	GMV	Rev+ αGMV	SW	Rev	GMV	Rev+ αGMV
GSP and Fixed Positions	0.450	0.302	0.455	0.757	**0.377**	**0.305**	0.523	0.828	0.296	0.201	0.532	0.733
IAS	**0.476**	**0.381**	0.474	0.855	0.306	0.273	0.645	0.919	0.324	0.216	**0.659**	0.876
JINTER Net	0.411	0.322	**0.614**	**0.936**†	0.272	0.197	**0.728**	**0.925**†	**0.393**	**0.301**	0.619	**0.920**†

Table 2: The results of experiments for different value distributions. The setting is 3 ads and 2 organic items with 2 slots. The regret of mechanism generated by the JINTER Net is less than 0.001. The best performance is highlighted in bold. "†" indicates a statistically significant improvement in a paired t-test at $p < 0.05$ level.

plot the Pareto-curves for various mechanisms, with the results illustrated in Figure 2.

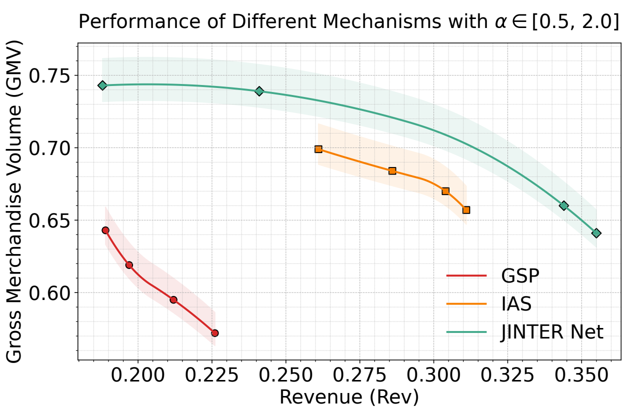

Figure 2: The Pareto curves of JINTER Net and other baseline mechanisms.

From Figure 2, it is evident that the JINTER Net curve is positioned in the upper right compared to the other curves. This indicates that, across a wide range of α values, the mechanism generated by JINTER Net effectively balances revenue and GMV, resulting in a more desirable blended list compared to the baseline methods.

Avito Dataset

The Avito public dataset, sourced from the platform avito.ru, includes user search logs over a period of 26 days, covering more than 4.2 million users and 36 million items across over 112 million page views (PVs). For our experiments, we partition the dataset as follows: 17-day records serve as the training set; 3-day records are used for validation; and 3-day records are reserved for testing.

The dataset does not include complete records for all items and slots in every PV. Specifically, only items in slots 1, 2, 6, 7 and 8 are recorded. Furthermore, in some PVs, only the CTR for slots 1 and 7 is available, while in others, only the CTR for slot 1 is recorded. To conduct our experiments with the Avito dataset, we pre-process dataset as follows. Denote CTR_r by the CTR of the r-th slot. We focus on the CTR for the 1st, 2nd, and 7th slots. For PVs where the CTRs for slots 1 and 7 are available, we estimate the CTR for slot 2 using the formula: $\text{CTR}_2 = N(0.4\text{CTR}_1 + 0.5\text{CTR}_7, 0.1\text{CTR}_1)$. For PVs where only the CTR for slot 1 is recorded, we simulate the CTR for slot 7 using a normal distribution $N(0.4\text{CTR}_1, 0.1\text{CTR}_1)$ and estimate the CTR for slot 2 as $\text{CTR}_2 = 0.4\text{CTR}_1 + 0.5\text{CTR}_7$. In each sample, we specify the 1st and 2nd items as ads and the remaining three items as organic. Consequently, there are 5 items competing for 3 slots within the Avito dataset.

Method	SW	Rev	GMV	Rev+ αGMV
GSP and Fixed Positions	0.377	0.252	0.621	0.873
IAS	**0.392**	**0.308**	0.608	0.916
JINTER Net	0.375	0.287	**0.641**	**0.948**†

Table 3: The results of experiments for Avito dataset. The best performance is highlighted in bold. The regret of mechanism generated by the JINTER Net is less than 0.001. Symbol "†" indicates a statistically significant improvement in a paired t-test at $p < 0.05$ level.

We present the experimental results on the test set of the Avito dataset in Table 3. Compared to the two baseline methods, JINTER Net achieves a significantly higher value for Rev+αGMV, with paired t-tests at the $p < 0.05$ level. Additionally, the mechanism generated by JINTER Net demonstrates approximate DSIC, with a regret of less than 0.001. These results imply the effectiveness of JINTER Net in real-world auction scenarios.

Conclusion

In this paper, we focus on the integrated ad auction, where multiple advertisers and organic items compete for limited slots. To design an optimal mechanism that ensures DSIC and IR, we present JINTER Net, a neural network architecture specifically developed to generate such integrated ad auctions. Our extensive experiments demonstrate the superiority of JINTER Net compared to existing baselines. Future research could explore applying additional techniques from automated mechanism design to integrated ad auctions or investigate other promising avenues in this domain.

Acknowledgments

This work was supported by National Natural Science Foundation of China No. 62472428, the Fundamental Research Funds for the Central Universities, the Research Funds of Renmin University of China (No. 22XNKJ07, No. 23XNH028), the Fundamental Research Funds of Shandong University (No.11480061330000) and Major Innovation & Planning Interdisciplinary Platform for the "Double-First Class" Initiative, Renmin University of China.

References

Clarke, E. H. 1971. Multipart pricing of public goods. *Public choice*, 17–33.

Conitzer, V.; and Sandholm, T. 2002. Complexity of mechanism design. In *Proceedings of the Eighteenth conference on Uncertainty in artificial intelligence*, 103–110.

Curry, M.; Sandholm, T.; and Dickerson, J. 2023. Differentiable economics for randomized affine maximizer auctions. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, 2633–2641.

Despotakis, S.; Ravi, R.; and Sayedi, A. 2021. First-price auctions in online display advertising. *Journal of Marketing Research*, 58(5): 888–907.

Duan, Z.; Tang, J.; Yin, Y.; Feng, Z.; Yan, X.; Zaheer, M.; and Deng, X. 2022. A context-integrated transformer-based neural network for auction design. In *International Conference on Machine Learning*, 5609–5626. PMLR.

Dütting, P.; Feng, Z.; Narasimhan, H.; Parkes, D.; and Ravindranath, S. S. 2019. Optimal auctions through deep learning. In *International Conference on Machine Learning*, 1706–1715. PMLR.

Edelman, B.; Ostrovsky, M.; and Schwarz, M. 2007. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American economic review*, 97(1): 242–259.

Feng, Z.; Narasimhan, H.; and Parkes, D. C. 2018. Deep learning for revenue-optimal auctions with budgets. In *Proceedings of the 17th international conference on autonomous agents and multiagent systems*, 354–362.

Groves, T. 1973. Incentives in teams. *Econometrica: Journal of the Econometric Society*, 617–631.

Han, X.; and Liu, S. 2015. Bayes-Nash Equilibrium of the Generalized First-Price Auction. *Mathematical Problems in Engineering*, 2015(1): 301734.

Ivanov, D.; Safiulin, I.; Filippov, I.; and Balabaeva, K. 2022. Optimal-er auctions through attention. *Advances in Neural Information Processing Systems*, 35: 34734–34747.

Kuo, K.; Ostuni, A.; Horishny, E.; Curry, M. J.; Dooley, S.; Chiang, P.-y.; Goldstein, T.; and Dickerson, J. P. 2020. Proportionnet: Balancing fairness and revenue for auction design with deep learning. *arXiv preprint arXiv:2010.06398*.

Li, N.; Ma, Y.; Zhao, Y.; Wang, Q.; Zhang, Z.; Yu, C.; Xu, J.; Zheng, B.; and Deng, X. 2024. Ad vs Organic: Revisiting Incentive Compatible Mechanism Design in E-commerce Platforms. In *Proceedings of the ACM on Web Conference 2024*, 235–244.

Li, W.; Qi, Q.; Wang, C.; and Yu, C. 2023. Optimally integrating ad auction into e-commerce platforms. *Theoretical Computer Science*, 976: 114141.

Liao, G.; Li, X.; Wang, Z.; Yang, F.; Guan, M.; Zhu, B.; Wang, Y.; Wang, X.; and Wang, D. 2022a. NMA: neural multi-slot auctions with externalities for online advertising. *arXiv preprint arXiv:2205.10018*.

Liao, G.; Wang, Z.; Wu, X.; Shi, X.; Zhang, C.; Wang, Y.; Wang, X.; and Wang, D. 2022b. Cross dqn: Cross deep q network for ads allocation in feed. In *Proceedings of the ACM Web Conference 2022*, 401–409.

Likhodedov, A.; and Sandholm, T. 2004. Methods for boosting revenue in combinatorial auctions. In *AAAI*, 232–237.

Likhodedov, A.; Sandholm, T.; et al. 2005. Approximating revenue-maximizing combinatorial auctions. In *AAAI*, volume 5, 267–274.

Liu, X.; Yu, C.; Zhang, Z.; Zheng, Z.; Rong, Y.; Lv, H.; Huo, D.; Wang, Y.; Chen, D.; Xu, J.; 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.

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

Peri, N.; Curry, M.; Dooley, S.; and Dickerson, J. 2021. Preferencenet: Encoding human preferences in auction design with deep learning. *Advances in Neural Information Processing Systems*, 34: 17532–17542.

Sandholm, T.; and Likhodedov, A. 2015. Automated design of revenue-maximizing combinatorial auctions. *Operations Research*, 63(5): 1000–1025.

Varian, H. R. 2007. Position auctions. *international Journal of industrial Organization*, 25(6): 1163–1178.

Vickrey, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 16(1): 8–37.

Wang, W.; Jin, J.; Hao, J.; Chen, C.; Yu, C.; Zhang, W.; Wang, J.; Hao, X.; Wang, Y.; Li, H.; et al. 2019. Learning adaptive display exposure for real-time advertising. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, 2595–2603.

Zhao, X.; Gu, C.; Zhang, H.; Yang, X.; Liu, X.; Tang, J.; and Liu, H. 2021. Dear: Deep reinforcement learning for online advertising impression in recommender systems. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35(1), 750–758.

Zhao, X.; Zheng, X.; Yang, X.; Liu, X.; and Tang, J. 2020. Jointly learning to recommend and advertise. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 3319–3327.