

Generative Bid Shading in Real-Time Bidding Advertising

Yinqiu Huang
Meituan
Chengdu, China
huangyinqiu@meituan.com

Hao Ma
Chongqing University
Chongqing, China
ma_hao@cqu.edu.cn

Wenshuai Chen
Meituan
Chengdu, China
chenwenshuai@meituan.com

Zongwei Wang*
Chongqing University
Chongqing, China
zongwei@cqu.edu.cn

Shuli Wang
Meituan
Chengdu, China
wangshuli03@meituan.com

Yongqiang Zhang
Meituan
Chengdu, China
zhangyongqiang08@meituan.com

Xue Wei
Meituan
Chengdu, China
weixue06@meituan.com

Yinhua Zhu
Meituan
Chengdu, China
zhuyinhua@meituan.com

Haitao Wang
Meituan
Chengdu, China
wanghaitao13@meituan.com

Xingxing Wang
Meituan
Beijing, China
wangxingxing04@meituan.com

Abstract

Bid shading plays a crucial role in Real-Time Bidding (RTB) by adaptively adjusting the bid to avoid advertisers overspending. Existing mainstream two-stage methods, which first model bid landscapes and then optimize surplus using operations research techniques, are constrained by unimodal assumptions that fail to adapt for non-convex surplus curves and are vulnerable to cascading errors in sequential workflows. Additionally, existing discretization models of continuous values ignore the dependence between discrete intervals, reducing the model's error correction ability, while sample selection bias in bidding scenarios presents further challenges for prediction. To address these issues, this paper introduces Generative Bid Shading (GBS), which comprises two primary components: 1) an end-to-end generative model that utilizes an autoregressive approach to generate shading ratios by stepwise residuals, capturing complex value dependencies without relying on predefined priors; and 2) a reward preference alignment system, which incorporates a channel-aware hierarchical dynamic network (CHNet) as the reward model to extract fine-grained features, along with modules for surplus optimization and exploration utility reward alignment, ultimately optimizing both short-term and long-term surplus using group relative policy optimization (GRPO). Extensive experiments on both offline and online A/B tests validate GBS's effectiveness. Moreover, GBS has been deployed on the Meituan DSP platform, serving billions of bid requests daily.

*Corresponding author.



This work is licensed under a Creative Commons Attribution 4.0 International License. *SIGIR '26, Melbourne, VIC, Australia*
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2599-9/2026/07
<https://doi.org/10.1145/3805712.3809566>

CCS Concepts

• Information systems → Computational advertising.

Keywords

Bid Shading, Generative Model, Auto-bidding

ACM Reference Format:

Yinqiu Huang, Hao Ma, Wenshuai Chen, Zongwei Wang, Shuli Wang, Yongqiang Zhang, Xue Wei, Yinhua Zhu, Haitao Wang, and Xingxing Wang. 2026. Generative Bid Shading in Real-Time Bidding Advertising. In *Proceedings of the 49th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '26)*, July 20–24, 2026, Melbourne, VIC, Australia. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3805712.3809566>

1 Introduction

Online advertising has emerged as a cornerstone of the digital economy, with Real-Time Bidding (RTB) serving as the fundamental infrastructure for programmatic media buying. In this ecosystem, Demand-Side Platforms (DSPs) compete in auctions to purchase impressions on behalf of advertisers. Under the conventional RTB paradigm, these high-frequency auctions were typically implemented using a second-price rule, under which the highest bidder won the impression but paid a settlement price equal to the second-highest bid. This effective mechanism can incentivize truthful bidding by aligning bidders' dominant strategy with submitting offers that reflect their underlying valuation of each impression, thereby mitigating the risk of systematic overpayment relative to prevailing competitive conditions.

However, the practical application of second-price auctions has been increasingly challenged by a lack of transparency regarding settlement prices and an inability to adapt to complex, emerging auction environments [1, 3, 10, 31]. Recently, the industry has shifted to First-Price Auctions (FPA), where the winner pays their own bid,

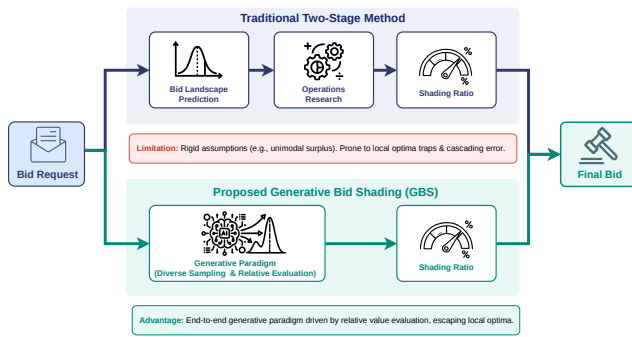


Figure 1: Overview of bid shading process.

implying that bidding one’s full valuation generally reduces surplus and results in systematic overpayment. In an idealized setting where the market-clearing (minimum winning) price for a given bid request was observable, an advertiser could bid exactly that price to win the impression at the lowest possible cost. However, because this clearing price is latent and should be inferred under substantial uncertainty, advertisers instead rely on bid shading [6, 15], which strategically discounts the submitted bid below the raw valuation. By calibrating this discount to trade off win probability against surplus preservation, bid shading can substantially reduce acquisition costs and improve budget efficiency, thereby enhancing resource allocation in the trillion-dollar digital advertising market.

The prevailing approach to bid shading typically adopts a two-stage pipeline (Fig. 1): it first estimates the bid landscape, and then uses an operations-research (OR) module to maximize surplus via search algorithms (e.g., bisection search [22], golden section search [38]). However, this decoupled design is inherently brittle. Its rigidity leads to a key limitation: the final bid can be suboptimal because the optimizer relies on simplifying assumptions and is sensitive to estimation errors. Many OR solvers assume that the surplus as a function of the bid is unimodal [22, 38], so that simple search procedures can reliably locate the optimum. In practice, the surplus curve in RTB is often non-convex and may contain multiple local peaks, causing search algorithms to converge to a local optimum rather than the global one. As illustrated in Fig. 2, the surplus curve initially increases and then decreases, but does not strictly adhere to a standard unimodal pattern. In principle, one could evaluate all candidate bids online to find the best bid, but exhaustive evaluation is computationally infeasible under millisecond-level latency constraints. Moreover, because the workflow is sequential, errors in the first-stage estimation are directly carried into the second-stage optimization and can be further magnified, preventing the system from achieving practical optimality within strict latency budgets.

To address these issues, we explore a generative paradigm that mitigates local optima by combining diverse candidate generation with relative evaluation. In this paradigm, we sample multiple bid candidates from a generative model and evaluate them jointly, which approximates a broader search than single-point optimization. We further compute relative advantages among candidates, so learning is driven by comparisons rather than fragile absolute targets, reducing the chance of being misled by local extrema. As illustrated in Fig. 2, conventional search methods frequently stagnate




Figure 2: The winning rate and surplus curves derived from Meituan’s non-ideal second-price auctions data.

in local optima when facing non-convex objective functions, while GBS adopts a broader view of the generative space, allowing it to effectively converge to the global optimum. However, implementing this idea is non-trivial with existing generative architectures. Methods such as CVAE [33] often rely on rigid predefined priors, while discretization-based approaches [18, 34] fail to leverage conditional dependencies between discrete intervals, as they treat these intervals as isolated labels without mutual interdependence. Moreover, training data collected from historical winning bids is inherently biased and covers only a limited portion of the counterfactual action space, further restricting effective learning.

Building upon the above insights, we incorporate autoregressive modeling [20] and reinforcement-based alignment [8] to propose a novel end-to-end framework, Generative Bid Shading (GBS). This approach leverages a global perspective to address the limitations of traditional two-stage methods, utilizing a generative architecture to optimize bid shading. GBS comprises two main components: 1) An autoregressive generative model: To address the limitations of rigid priors and the neglect of dependencies between discrete intervals, we propose an autoregressive model. By decomposing shading ratios into sequences of discrete tokens, our method treats generation as a conditional sequential process. This design enables the model to effectively capture high-order dependencies between token values, facilitating precise approximation through residual refinement. To support this, we employ teacher forcing and gumbel-softmax to ensure stable convergence and end-to-end differentiability. 2) A reward preference alignment system: We propose a reinforcement learning-based preference alignment system anchored by the Channel-aware Hierarchical Dynamic Network (CHNet), a high-precision reward model. Policy optimization is driven by two distinct alignment mechanisms: a surplus optimization alignment that maximizes immediate returns by curbing overbidding on obtainable requests, and an exploration utility alignment that integrates strategic bid exploration to counteract sample selection bias. Furthermore, we employ Group Relative Policy Optimization (GRPO) to guarantee efficient and stable policy updates throughout the training process. Crucially, this design inherently circumvents the issues of local sub-optimality and cascading errors that typically plague two-stage approaches.

In summary, our contributions are as follows:

- To our best knowledge, we are the first to formally identify the inherent limitations of the conventional two-stage bid shading paradigm. We demonstrate that the decoupling of estimation and optimization leads to cascading errors, while the reliance on unimodal assumptions often results in local

optimality within non-convex surplus landscapes, thereby fundamentally constraining the performance upper bound of existing methods.

- We propose a novel end-to-end generative paradigm designed to transcend traditional limitations. This framework synergistically integrates a novel autoregressive generative model, a high-precision reward model (CHNet), and dual preference alignment mechanisms, which simultaneously maximize immediate surplus and mitigate data selection bias, achieving superior scalability and generalization across diverse business scenarios.
- Through extensive experiments on both offline and online A/B tests, we show that our GBS significantly outperforms the state-of-the-art methods. GBS has been deployed in Meituan DSP and serves billions of bid requests daily. To the best of our knowledge, this work represents the first implementation of generative bid shading in large-scale online systems.

2 Related Work

Bid Landscape Forecasting. Bid landscape forecasting is a crucial component of RTB advertising systems, primarily aiding in forecasting the winning price distribution of bid requests. Early research often assumed that winning price followed specific probability distributions, such as Gaussian [36], Log-Normal [7], Gamma [4, 39], or other parametric distributions [26, 27]. To overcome the limitations of these predefined distributional assumptions, subsequent studies introduced distribution-free models [15, 16, 25]. However, existing methods generally overlook the modeling of fine-grained features, so we propose a channel-aware hierarchical dynamic network to achieve fine-grained modeling as a foundation for the reward system.

Bid Shading. Bid shading helps advertisers avoid overpaying. Existing methods fall into three categories: 1) Direct regression of the shading ratio [11], which predicts the optimal ratio but unrealistically assumes access to all winning prices; 2) Two-stage ML + OR optimization [22, 30, 38], where machine learning estimates bid landscape and operations research maximizes surplus; and 3) Non-parametric estimation [37], which uses dynamic binning without parametric models, but still needs optimal shading labels and cannot fully capture individual differences.

Generative Model and Preference Alignment. With the rapid development of large language models (LLMs), generative models are increasingly used in areas such as generative recommendation [8] and search [13]. Preference learning, including Reinforcement Learning from Human Feedback (RLHF) [21], aligns model outputs with human preferences but is inefficient. Direct Preference Optimization (DPO) [23] and its variants [2, 5] address this by directly adjusting model policies using preference data. However, they face challenges like sample sparsity in recommendation and advertising. Group Relative Preference Optimization (GRPO) [29] uses explicit reward functions and group-based evaluation to enhance decision-making and training efficiency. In this paper, we present the first exploration of applying the generative model to bid shading, leveraging reward systems to guide the model beyond the constraints of traditional methods.

3 Preliminaries

3.1 Problem Definition

Suppose each DSP has an online advertising system to process bidding requests in the RTB scenario. When users visit web pages containing advertising opportunities, each DSP selects the candidate advertisement with the highest advertising value and places a bid. Once the auction ends, the winning advertisement is displayed, and the win/loss tags are returned to the DSP. Let $X = \{x_1, x_2, \dots, x_N\}$ represent the bid requests, where x_i includes features such as user and contextual information. The value of each request is calculated through metrics such as the click-through rate (CTR) and conversion rate (CVR), denoted as $v = \{v_1, v_2, \dots, v_N\}$, which signifies the expected value derived from an advertisement request and serves as the pre-shadow bid¹. In the bid shading system, we determine a shading ratio α_i for each request using a generative model, and multiply it with the unshaded bid to obtain the final bid $B = \{v_1 * \alpha_1, v_2 * \alpha_2, \dots, v_N * \alpha_N\}$. The expected surplus of a bid request x_i at the bid price b_i is $surplus(x_i, b_i) = (v_i - c_i) \cdot Wr(b_i | x_i)$, where $Wr(\cdot)$ represents the winning rate and c_i represents the cost incurred if this request is won, $c_i = b_i$ in FPA scenario. For the bid shading system, the bid prices b_i^* maximize the expected surplus:

$$b_i^* = \arg \max_{b_i \in (b_{min}, b_{max})} (v_i - c_i) \cdot Wr(b_i | x_i), \quad (1)$$

where b_{min} and b_{max} are manually set as the lower and upper limits for bid shading. The DSP initiates a bid using the final b_i^* and compares it against the bids from other advertisers. Once the auction concludes, the win/lose label y_i^* and the final cost c_i are determined. Suppose z_i represents the winning price, and we have

$$y_i^* = \begin{cases} 1, & b_i^* > z_i, \\ 0, & b_i^* \leq z_i. \end{cases} \quad (2)$$

It is important to note that z_i is not disclosed to the DSP. The actual surplus of the bid request x_i at b_i^* is $surplus(x_i, b_i^*) = (v_i - c_i) \cdot y_i^*$, which serves as a crucial evaluation metric for bid shading.

3.2 Generative Model Preprocessing

Many methods output a scalar by discretizing continuous values and converting the prediction problem into a classification task to simplify learning. [18, 34]. However, predictions for the discretized buckets are typically generated independently, limiting opportunities for effective error correction. Current generative recommendation methods frequently utilize residual quantization (RQ-VAE) and hierarchical K-means clustering to obtain index IDs, and generate ID sequences through an autoregressive model [24]. In addition, generative regression [20] also uses this way to model scalar generation, achieving excellent performance. Inspired by these methods, we propose generating the shading ratio using an autoregressive encoder-decoder architecture.

Specifically, we introduce a vocabulary $\mathcal{V} = \{w_j\}_{j=1}^V$ where each element w_i represents a predefined scalar value, analogous to tokens in the language model. The vocabulary embedding matrix is denoted as $E \in \mathbb{R}^{V \times D}$, where V consists of the vocabulary size, and D is the dimension of the token embeddings. Each shading ratio can be

¹This step usually solves an optimization problem with constraints and will not be elaborated on in detail in this paper.

encoded as a sequence of tokens $s_i = \{s_i^1, \dots, s_i^L\}$, with L indicating the length of the sequence. This process is known as tokenization. In contrast to generative recommendation systems that index directly based on ID, we propose a label decoding function $g(\cdot)$ to determine the actual scalar value from the token ID output by the model, thereby deriving the shading ratio $\alpha_i = g(s_i) = \sum_{l=1}^L g(s_i^l) \in \mathbb{R}$.

4 Method

In this section, we provide a detailed explanation of the proposed GBS. First, we introduce the end-to-end autoregressive generative model. Next, we present the reward preference alignment system, which employs a channel-aware hierarchical dynamic network as the reward model and utilizes GRPO to guide policy learning via two reward alignment modules.

4.1 End-to-End Generative Model

Our generative model decomposes the final task into multiple classification subtasks, utilizing inter-step dependencies to approximate the final shading ratio accurately. This framework allows each prediction step the flexibility to choose from the vocabulary, and the expanded solution space enables the model to generate a broader spectrum of potential sequences.

Tokenizer. The vocabulary \mathcal{V} comprises predefined ratio slots as tokens, allowing the model to generate shading ratios that closely approximate actual values. Two key properties are prioritized when constructing \mathcal{V} : 1) Each token must be unique, enabling \mathcal{V} to represent all possible shading ratios with a finite token set; 2) Token frequencies should be relatively uniform to prevent class imbalance. We employ the dynamic percentile-based algorithm [20] for automated vocabulary construction from specified datasets. It accelerates tail-value reduction, rapidly decreases inter-update variance, and enhances algorithmic reliability (see Appendix B for details).

Due to the lack of ground-truth labels for shading ratios, we construct the vocabulary \mathcal{V} using the shading ratio from the two-stage method and perform tokenization, which involves transforming ratios $\{\alpha_i\}_{i=1}^N$ into discrete tokens $\{s_i = \{s_i^1, \dots, s_i^L\}\}_{i=1}^N$. To reduce the complexity of learning encoded labels, the tokenized sequences should adhere to three principles: 1) The reconstructed value $g(s_i)$ should minimize error relative to the real ratio; 2) The sequence length L of s_i should be minimized; 3) The sequence should maintain monotonicity: $s_i^1 \geq s_i^2 \geq \dots \geq s_i^L$. Sequences that meet these conditions enable the model to incrementally output residual results, thereby significantly simplifying the learning process. Consequently, we employ a greedy strategy during tokenization, generating training data for the pre-training stage.

Encoder-Decoder Architecture. As illustrated in Fig. 3, our model employs a Transformer-based architecture, consisting of an embedding layer, an encoder for modeling features, and a decoder for generating shading ratios. Specifically, the embedding layer maps features through the parameter matrix. For dense features, we transform them into the same dimension using an affine transformation:

$$e_{ij} \in \mathbb{R}^{1 \times d} = \begin{cases} W_j * x_{ij} + b_j, & x_{ij} \text{ is a continuous feature,} \\ \text{lookup}(E_j, e_{ij}), & x_{ij} \text{ is a sparse feature,} \end{cases} \quad (3)$$

where W and b are learnable parameters. Then, we concatenate them to get the initial embedding representation of x_i :

$$e_i \in \mathbb{R}^{f \times d} = [e_{i1}, e_{i2}, \dots, e_{if}], \quad (4)$$

where d is the dimension size of each field, and f is the number of feature fields. The encoder processes input features using stacked multi-head self-attention and feed-forward layers:

$$H_i = (\dots FFN(\text{SelfAttn}(e_i))). \quad (5)$$

To standardize decoder training, we introduce the target sequence with a start-of-sequence token $\langle \text{SOS} \rangle$ and padding tokens $\langle \text{PAD} \rangle$. These tokens do not represent any semantic value in the label space (i.e., $g(\{\langle \text{SOS} \rangle, \langle \text{PAD} \rangle\}) = 0$). The decoder autoregressively generates the target sequence based on the encoder's output H_i and the preceding subsequence. At training timestep l , the output token is

$$\hat{s}_i^l = \arg \max_{w \in \mathcal{V}} P_\theta(w | H_i, \hat{s}_i^{<l}). \quad (6)$$

Pre-training Optimization. During pre-training, we equip the model with fundamental market perception capabilities through supervised fine-tuning (SFT):

- 1) Cross-entropy loss for next-token prediction:

$$L_{NTP} = - \sum_{i=1}^N \sum_{l=1}^L \log P_\theta(\hat{s}_i^l | H_i, \hat{s}_i^{<l}), \quad (7)$$

where labels are derived from shading ratios generated by a two-stage method, we adopt a teacher forcing (TF) strategy to improve model efficiency, which directly feeds the ground truth s_i^l as input at step $l + 1$ to guide model training. We design a gate to enable the TF strategy adaptively:

$$TF_{pro} = TF_{pro} \cdot \frac{\eta}{\eta + e^{\left(\frac{\text{iteration}}{\eta}\right)}}, \quad (8)$$

where TF_{pro} denotes the probability of employing the TF strategy, and η is a hyperparameter. The gate mechanism decreases the probability of using TF as the number of training iterations increases, promoting faster convergence and improved model stability.

- 2) Negative Log-Loss of expected surplus for winning data N_+ :

$$L_{surplus} = - \sum_{i=1}^{N_+} \mathbb{E}(\log((v_i - C(v_i \cdot \alpha_i)) \cdot Wr(v_i \cdot \alpha_i | x_i))), \quad (9)$$

where α_i is the predicted shading ratio, $C(\cdot)$ is the cost function ($C(b) = b$ in FPA), and $Wr(\cdot)$ is the estimated winning rate (see Section 4.2.1). Notably, the argmax operation in Equation 6 obstructs gradient backpropagation. Therefore, we employ the gumbel softmax to obtain a differentiable distribution during token generation. At time step l , the dimension of the estimated raw output vector is V (the size of the vocabulary). We do not consider the entire vocabulary space when calculating the weighted vector. Since the vocabulary is strictly ordered, we focus on the domain V' , which corresponds to the neighborhood of the indices of the highest output values. For each position $p \in V'$, we sample u_p from the uniform distribution $U \sim \text{Uniform}(0, 1)$, and generate gumbel noise $n_p = -\log(-\log(u_p))$ through the inverse transform, adding it to the original logits:

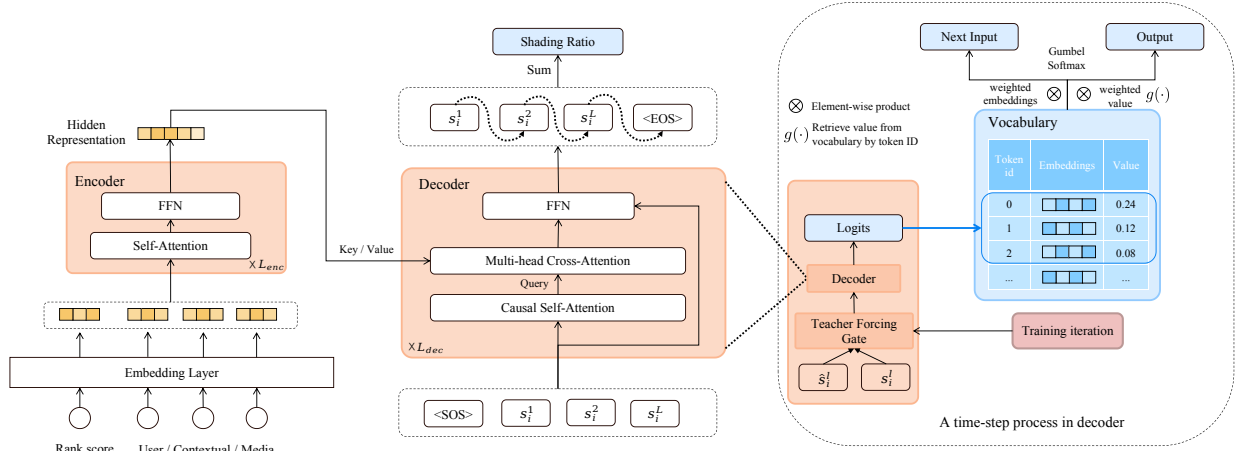


Figure 3: The proposed generative model employs an encoder-decoder architecture to predict the token sequence autoregressively. It determines the specific value of each token using a decoding function $g(\cdot)$, then accumulates these values to obtain the final shading ratio. Additionally, it utilizes the gumbel softmax and teacher forcing strategies to ensure stability and accelerate convergence.

$$weight_p^l = \frac{\exp((l_p + n_p) / \tau)}{\sum_{j=1}^{V'} \exp((l_p + n_p) / \tau)}, \quad (10)$$

where l_p is the logit predicted by the decoder at step l and position p , and τ is the temperature coefficient. As τ approaches 0, the output resembles discrete sampling; as τ approaches infinity, the output approximates a uniform distribution. Finally, we modify the decoding function $g(\cdot)$ using gumbel softmax weights and use the embedding of weighted sum \hat{s}_i^l as input for the next token prediction. Compared to argmax , gumbel softmax not only addresses the issue of differentiability but also reduces the impact of errors in the earlier positions of the output sequence, increasing the stability.

In the pre-training stage, the loss is $L = L_{NTP} + \lambda \cdot L_{surplus}$, where λ is a hyperparameter.

4.2 Reward Preference Alignment System

Pre-trained models are limited by two-stage solutions and often get stuck in local optima due to the SFT in non-convex surplus loss and bias from training only on winning bids. To address this, we use reward-based preference alignment during post-training. With policy reinforcement learning, the model receives reward signals that help it develop a global perspective to escape local optima. Specifically, a surplus preference reward prevents overbidding, while an exploratory utility reward balances information and cost. In this section, we first introduce the reward model and then introduce the two reward preference alignment systems.

4.2.1 Channel-Aware Hierarchical Dynamic Network (CHNet). Bid landscape forecasting is critical for evaluating bid quality. Following the previous works [15, 25], we first discretize the continuous space via equal frequency binning, where a set of T prices $0 < b_1 < b_2 < \dots < b_T$ covers the finite precision of price determinations. We have the probability density function (PDF) $P(z = b | x)$ of the winning price z being b , and the cumulative distribution function

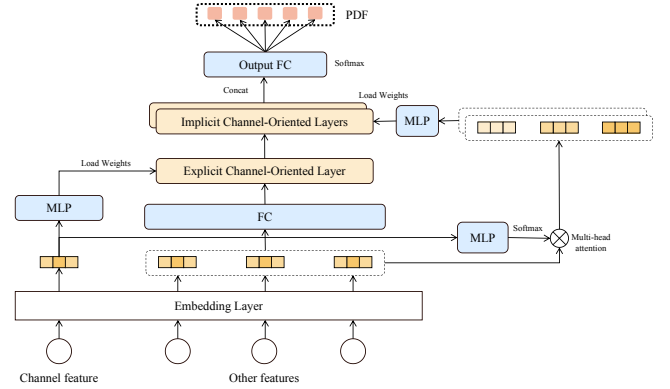


Figure 4: The architecture of CHNet.

(CDF) $Wr(b_t | x) = P(z \leq b_t | x) = \sum_{j \leq t} P(z \in B_j)$ representing the winning rate of a certain bid b_t , where $t = 1, 2, \dots, T$ and $B_j = (b_j, b_{(j+1)}]$ is a divided disjoint interval. Once the PDF curve of a request bid is obtained, the winning rate for another bid can be derived with a single additional CDF integral calculation.

In real-world advertising systems, varying media employ opaque and diverse billing methods, leading to significant differences in the bid landscape across channels. Inspired by multi-scenario approaches [9], we propose a channel-aware hierarchical dynamic network (CHNet) for fine-grained channel information modeling, as shown in Fig. 4. CHNet utilizes a hierarchical structure with an explicit channel-oriented layer and multiple implicit channel-oriented layers, sequentially modeling coarse-grained explicit information and fine-grained implicit information. For processing feature matrices², we adhere to the method outlined in Section 4.1, with separate processing for channel feature e^c and others e^o .

²for simplicity, request i is not separately denoted in this section.

Explicit channel-oriented layer. We perform global feature interactions for the common feature e^o using a multi-layer perceptron (MLP), denoted as $e^{global} = MLP(e^o)$. Furthermore, we adaptively generate the weight parameters based on the channel information e^c via a re-parameterization approach, specifically by applying an MLP followed by a reshape operation: $W_{exp}, b_{exp} = Reshape(MLP(e^c))$. These parameters are then used to instantiate the explicit scenario-oriented layer, resulting in the explicit channel representations:

$$e^{explicit} = FC(e^{global}; W_{exp}, b_{exp}). \quad (11)$$

Implicit channel-aware module. After explicit channel modeling, CHNet uncovers implicit patterns using a channel-aware multi-head attention mechanism. This enables comprehensive modeling of complex data distributions and adaptive identification of key implicit patterns. First, channel embeddings e^c are processed by a MLP to generate raw weights, which are then normalized into G groups (G is the number of attention heads) using the softmax:

$$\begin{aligned} weight_{ori} &= Reshape(MLP(e^c)), \\ weight_{norm}[g] &= Softmax(weight_{ori}[g]), g \in [1, G]. \end{aligned} \quad (12)$$

Next, we perform an element-wise product of the normalized weight with the feature embedding: $e^{weight} = weight_{norm} \otimes e^o$.

Then, using the re-parameterization technique, we generate the weight parameters for different implicit channel layers, denoted as

$$W_{imp}[g], b_{imp}[g] = Reshape(MLP(e^{weight}[g])), g \in [1, G]. \quad (13)$$

We input the explicit representations into each implicit channel-oriented layer for implicit modeling. After that, we concatenate the outputs of all implicit layers to predict the PDF distribution:

$$\begin{aligned} e_g^{implicit} &= FC(e^{explicit}; W_{imp}[g], b_{imp}[g]), \\ output &= Softmax(FC(Concat(e_1^{implicit}, \dots, e_G^{implicit}))). \end{aligned} \quad (14)$$

Training loss. During the CHNet training process, the loss function is comprised of three components. For winning bid logs in a non-FPA scenario, we directly compute the negative log-likelihood loss based on the actual winning price: $L_{win} = -\sum_{i=1}^{N_+} \log p(c_i | x_i)$. Note that in the FPA scenario, this loss cannot be applied because the cost paid does not reflect the minimum winning price. For losing bid logs, we only know that the winning price exceeds the current bid, and we maximize the CDF of the right half of the bid, resulting in $L_{lose} = -\sum_{i=1}^{N_-} \log[1 - Wr(b_i | x_i)]$. Additionally, to further enhance the model's ranking and classification capabilities [19], we introduce a pairwise loss related to the winning rate, denoted as

$$L_{ranknet} = -\frac{1}{N_+ N_-} \sum_{i=1}^{N_+} \sum_{j=1}^{N_-} \log(\sigma(Wr(b_i | x_i)^{(+)} - Wr(b_j | x_j)^{(-)})). \quad (15)$$

Finally, the loss function of CHNet is formulated as

$$L = \lambda_1 \cdot L_{win} + \lambda_2 \cdot L_{lose} + \lambda_3 \cdot L_{ranknet}. \quad (16)$$

4.2.2 Surplus Optimized Alignment. The primary objective of bid shading is to prevent incurring excessive costs. Accordingly, we use the expected surplus associated with the shaded bid $b = v * \alpha$ as the reward $r_{surplus} = (v - C(b)) \cdot Wr(b | x)$ and optimize the model using reinforcement learning strategies. We employ GRPO to optimize the strategy based on relative rewards within the group.




Figure 5: The architecture of post-training with GRPO.

More specifically, for each bid request req , GRPO samples a group of outputs $\{\alpha_1, \dots, \alpha_2, \dots, \alpha_m\}$ from the old policy and then optimizes the policy model by maximizing the following objective:

$$\begin{aligned} \mathcal{J}_{GRPO}(\theta) &= \mathbb{E}_{req \sim P(R), \{\alpha_i\}_{i=1}^m \sim \pi_{\theta_{old}}} \left[\frac{1}{m} \sum_{i=1}^m \min \left(\frac{\pi_{\theta}(\alpha_i | req)}{\pi_{\theta_{old}}(\alpha_i | req)} A_i, \right. \right. \\ &\quad \left. \left. \text{clip} \left(\frac{\pi_{\theta}(\alpha_i | req)}{\pi_{\theta_{old}}(\alpha_i | req)}, 1 - \epsilon, 1 + \epsilon \right) \right) - \beta \cdot \mathbb{D}_{KL}[\pi_{\theta} \| \pi_{ref}] \right], \end{aligned} \quad (17)$$

where A_i denotes the advantage, which is calculated solely based on the relative return of each group's internal output:

$$A_i = \frac{r_i - \text{mean}(\{r_1, r_2, \dots, r_m\})}{\text{std}(\{r_1, r_2, \dots, r_m\})}. \quad (18)$$

Additionally, we add the KL divergence between the training and reference strategies in the objective. GRPO simultaneously employs the group-relative advantage estimation and a KL divergence penalty to ensure that strategy updates are efficient and stable. The complete algorithm processing is detailed in Algorithm 1.

Pre-training with surplus-based SFT restricts the model to exploring the specific output bid and may cause gradient descent to get stuck in local optima due to the non-convex surplus. In contrast, policy reinforcement learning expands the generation space for getting global optima. Since the bidding process is uncertain and each winning price reflects a probability distribution, we incorporate losing bid data with a smaller weight during post-training.

4.2.3 Exploration Utility Alignment. In the DSP advertising bidding scenario, data selection bias presents a significant challenge, as the collected data are generated using a fixed bidding algorithm, which complicates the bid landscape forecasting. To address this issue, we incorporate bid exploration into the bid shading system, providing more robust data support for maximizing long-term surplus.

According to uncertainty sampling in active learning [28, 35], we think that the greater the predicted uncertainty, the more informative the exploration will be. We quantify the exploration utility using the entropy of the winning rate. For a bid request x_i , the predicted winning rate $Wr_i = Wr(b_i | x_i)$ at a particular bid b_i can be framed as a binary classification problem (win or lose). Thus, the entropy of the bid request at bid b_i is denoted as

$$\phi^E(x_i, b_i) = -(Wr_i \cdot \log(Wr_i) + (1 - Wr_i) \cdot \log(1 - Wr_i)). \quad (19)$$

Additionally, during exploration, it is important to avoid excessive cost increases; the final exploration reward is denoted as $r_{exploration} = \frac{\phi^E(x_i, b_i)}{C(b_i)}$. When calculating the final advantage A_i (Equation 18), the reward is

$$r = r_{surplus} + \omega \cdot r_{exploration}. \quad (20)$$

where ω is not a fixed parameter. It is unnecessary to explore all samples; by identifying samples with high similarity, we can avoid unproductive exploration of outliers. Inspired by the metric learning [32], we compute the exploration weight for request x_i :

$$\omega_i = \cos(e_i, \frac{1}{|N|} \sum_{j=1}^N e_j), \quad (21)$$

where $\cos(\cdot, \cdot)$ is the cosine similarity function, more representative requests are assigned greater exploration weight. The post-training structure is shown in Figure 5.

5 Experiments

5.1 Experimental Setup

5.1.1 Datasets. We evaluate our method on both the public and real-world industrial datasets, with detailed statistics presented in Table 1. The public iPinYou [17] data are derived from the second-price auctions setting, where the advertiser’s bid is used as the actual value in the bid request, and the paid price is regarded as the winning price. The dataset is randomly split into training and testing sets with a ratio of 7/3. Our private dataset is collected from one month of bid sampling data from Meituan DSP and is partitioned similarly. Notably, only the testing set contains minimum winning price, which is used for evaluation.

Table 1: The statistics of datasets.

DataSet	Size	Winning rate	features
iPinYou	10,577,061	29.7%	18
Private	162,549,577	3.89%	197

5.1.2 Evaluation Metrics. For offline evaluation in FPA, business performance is crucial in determining whether an algorithm is suitable for production deployment, as it reflects the algorithm’s effectiveness. Given the objective of bid shading, we use surplus $(V - b) \cdot I(b > z)$ and surplus rate as key business metrics, as they directly indicate the impact of the algorithm on business outcomes. Specifically, we calculate the surplus rate as the percentage of residual surplus relative to the total optimal surplus:

$$\text{surplus rate (SR)} = \frac{\sum_i (v_i - b_i) I(b_i > z_i)}{\sum_i (v_i - z_i)}. \quad (22)$$

5.1.3 Implementation Details. We implement all deep learning baselines and the proposed GBS on NVIDIA A100-SXM4-80GB GPUs. During the pre-training stage, we use the Adam optimizer with a batch size of 2^{10} and a learning rate of $1e^{-4}$. The maximum output sequence length L is set to 5, η in the TF Gate is 5000, and the temperature coefficient τ in gumbel softmax is 0.2. In the post-training stage, the loss hyperparameters λ_1 , λ_2 , and λ_3 for CHNet training are set to 0.5, 1, and 0.2, respectively, while the KL weight β in the GRPO optimization function is set to 0.01. We divide the training and testing sets of the iPinYou and private datasets into 100 segments each to simulate the continuous arrival of bid requests in an online environment. We calculate the surplus by the winning price in each round and collect the win/loss labels. Subsequently, the

bid data are used to update the CHNet and the generative model. This process is repeated for 100 rounds until the experiment is complete.

5.2 Overall Performance

We compare the performance of GBS with several state-of-the-art baselines, including the bid shading and the generative methods:

EDDN [38]: A two-stage bid shading method, which optimizes the surplus using golden section search in the OR stage.

WR [22]: A two-stage bid shading method, which optimizes the surplus using bisection algorithm in the OR stage.

TSBS-DLF [25]: A two-stage bid shading method that uses the DLF model in the ML stage, which models the bid landscape without distributional assumptions, employing the conditional probability chain rule and LSTM.

TSBS-ADM [16]: A two-stage bid shading method uses the ADM model in the ML stage, which uses neighborhood likelihood loss for accurate prediction.

MEBS [12]: An end-to-end bid shading method, which jointly optimizes the shading model and the win rate model, performs supervised learning through the negative logarithm of the surplus as the loss.

CVAE [33]: An extended model of variational autoencoder, which achieves sample generation based on specific inputs by integrating conditional variables in the encoder and decoder. We optimize it through the pre-training method proposed in this paper.

DF [14]: A generative diffusion model that corrupts the data distribution by gradually adding noise and then learns the inverse denoising process to generate shading ratio. We optimize through the pre-training method proposed in this paper.

Post-CVAE & Post-DF: Post-training will be performed on the pre-trained generative model, which uses reparameterization to transfer gradients.

Based on the results in Table 2, we can draw the following conclusions: 1) The traditional two-stage algorithm performs significantly worse on the Meituan private dataset, which suggests that a more complex market environment poses greater challenges for bid shading, and that coarse-grained bid landscape forecasting and bisection search cannot adapt to a broader range of business scenarios; 2) The traditional generative models show varying degrees of improvement after post-training and outperforms the two-stage method, indicating that policy reinforcement learning can effectively enhance these models to raise their performance ceiling, and the end-to-end generative approach offers distinct advantages; 3) GBS achieves outstanding results across all metrics. By combining an autoregressive residual generative model with post-training, GBS effectively mitigates the suboptimality of traditional methods and adapts well to complex market environments.

5.3 Ablation Study

To assess the effectiveness of each component in GBS, we conducted a series of ablation studies. Specifically, we build several variants of the GBS:

w/o PostT: A variant of GBS without the Post-training stage, which directly uses the pre-trained model to output shading ratio.

Table 2: Overall offline experiment results.

Model	iPinYou		Private	
	SR	Surplus	SR	Surplus
EDDN	47.82%	90,199,768	19.88%	14,582,074
WR	54.90%	103,560,851	26.13%	19,164,882
TSBS-DLF	55.22%	104,164,484	27.25%	19,986,132
TSBS-ADM	55.49%	104,673,800	32.67%	24,183,220
MEBS	54.46%	102,716,112	31.45%	23,041,120
CVAE	51.42%	96,981,112	36.45%	26,698,064
Post-CVAE	57.14%	107,781,128	36.63%	26,830,032
DF	48.04%	90,617,280	31.28%	22,912,038
Post-DF	51.13%	96,445,909	33.89%	24,823,816
GBS	60.48%	114,068,392	41.74%	30,580,080

w/o PostT-EUA: A variant of GBS without the exploration utility alignment in post-training, which does not perform bid exploration and only performs surplus optimization.

w/o GM: A variant of GBS without the autoregressive generative model, which uses the traditional generative model CVAE [33] to generate the shading ratio and keeps the pre-training and post-training consistent.

w/o CHNet: A variant of GBS without the CHNet, which uses a binary classification model to provide the predicted winning rate.

Table 3 presents the results of the ablation study. CVAE is less effective than our GM, underscoring the importance of progressively outputting results through residuals in refined bid generation. Excluding CHNet and bid exploration also leads to similar performance drops, indicating that a more fine-grained reward model and a broader range of training data are crucial for enhancing model performance. Omitting post-training results in the most significant performance degradation, highlighting the limitations of relying solely on supervised learning for the bid shading task. The results suggest that policy reinforcement learning plays a vital role in helping the model overcome performance bottlenecks.


Table 3: The contributions of different components of GBS.

Model	iPinYou		Private	
	SR	Surplus	SR	Surplus
w/o PostT	56.22%	106,036,888	36.91%	27,041,464
w/o PostT-EUA	60.16%	113,458,768	41.54%	30,429,570
w/o GM	57.14%	107,781,128	36.63%	26,830,032
w/o CHNet	58.22%	109,805,915	38.16%	27,953,616
GBS	60.48%	114,068,392	41.74%	30,580,080

5.4 Performance Analysis of CHNet

In this section, we evaluate the performance of the proposed CHNet in bid landscape forecasting on private datasets. We use the area under the curve (AUC) and binary cross-entropy (BCE) of winning rate as evaluation metrics. In the experiments, we compare the following competitive bid landscape forecasting methods: 1).

Censorship Linear Model (CLM) uses the normal distribution as the prior for single-point estimation of the winning rate; 2). Deep Landscape Forecasting (DLF) model [25]; 3). Arbitrary Distribution Model (ADM) [16]; 4). CHNet-w/o EUA is the ablation experiment without the exploration utility alignment module during the post-training phase.

**Figure 6: The performance of CHNet.**

The experimental results are shown in Fig. 6. The results demonstrate that CHNet performs superior ranking and classification tasks, owing to its refined modeling of channel information, a key characteristic of the market environment. Additionally, removing the bid exploration module decreases performance and exposes the model to data selection bias. Therefore, incorporating exploration in the bidding process is essential for improving the model’s generalization ability.

5.5 Performance on Online System (Meituan)

To evaluate the online performance of GBS, we deployed it on the Meituan DSP. We conducted a rigorous A/B test over two weeks, and the engineering architecture is shown in Fig. 7. We allocated 30% of the traffic to GBS, while the baseline method was an online two-stage bid shading approach, which utilizes the golden section search in the operations research stage. We model certain channels as a non-ideal second-price auction scenario, consistent between baseline and GBS. We use four key online metrics for evaluation: Return On Investment (ROI), Cost Per Mille (CPM), Cost Per Click (CPC), and Inference Time (IT).

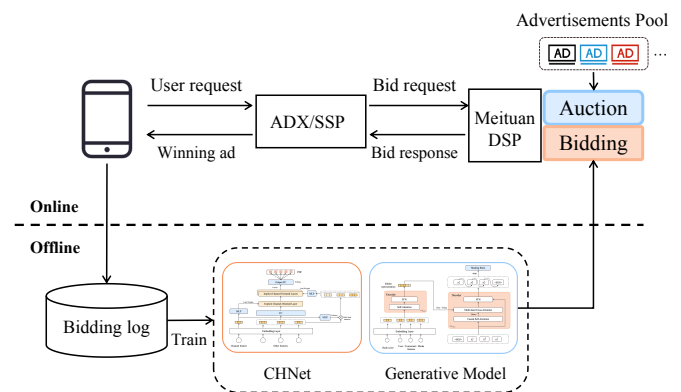
**Figure 7: Architecture of the online deployment with GBS.**

Table 4 shows the improvement of GBS relative to the baseline. The results indicate that GBS has significantly improved in all online metrics compared to the baseline, achieving significant business benefits. During the inference stage, GBS directly outputs the shading ratio, resulting in lower inference time than the two-stage method and reducing the burden on online services.

Table 4: Online A/B test results.

	ROI	CPM	CPC	IT
Baseline (Two-stage)	0.0%	0.0%	0.0%	0.0%
GBS	+3.4%	-4.1%	-4.9%	-37.6%

6 Conclusion and Future Work

This paper proposes a novel Generative Bid Shading (GBS) framework that addresses the suboptimality of existing bid shading techniques. GBS uses an autoregressive generative model to decompose the shading ratio into conditionally dependent sequences and incorporates a channel-aware reward model to capture fine-grained bid landscape features. By integrating the reward preference alignment system with the GRPO, GBS prevents overbidding, mitigates sample selection bias, and enables efficient policy updates. Extensive experiments show that GBS outperforms state-of-the-art models.

In the future, we will further optimize the generative model to better adapt upstream ranking scores and explore additional reward systems with advanced optimization algorithms.

A Post-training process of GBS

The post-training process is detailed in Algorithm 1.

Algorithm 1: Post-training Stage: GRPO Optimization

Require: Pre-trained generative model $\pi_{\theta_{\text{init}}}$, reward system, post-training data \mathcal{D}

```

1 Initialize policy model  $\pi_{\theta} \leftarrow \pi_{\theta_{\text{init}}}$ ;
2 for iteration = 1 to  $I$  do
3   Update reference model  $\pi_{\text{ref}} \leftarrow \pi_{\theta}$ ;
4   for step = 1 to  $S$  do
5     Sample batch  $\mathcal{D}_b \leftarrow \mathcal{D}$ ;
6     Set old policy  $\pi_{\theta_{\text{old}}} \leftarrow \pi_{\theta}$ ;
7     foreach bid request  $req \in \mathcal{D}_b$  do
8       Sample  $m$  outputs  $\{\alpha_i\}_{i=1}^m \sim \pi_{\theta_{\text{old}}}(\cdot | req)$ ;
9       Compute final rewards  $\{r_i\}_{i=1}^m$  by Equation 20;
10      Estimate token advantages  $\hat{A}_{i,t}$  via
          group-relative method (Equation 18);
11      Update  $\pi_{\theta}$  by maximizing GRPO objective
          (Equation 17);
12    end
13  end
14 end

```

Result: optimized policy model π_{θ}

B Vocabulary Construction

We use the shading ratios generated by the two-stage method as the base dataset to construct the vocabulary. The process (Algorithm 2) is initialized with a high starting quantile q_{start} , and then adaptively reduced by a decay rate δ until reaching the terminal quantile q_{end} .

Algorithm 2: Constructing Vocabulary with Dynamic Percentiles

Require: Shading ratios from two stage method

$Y = \{\alpha_j\}_{j=1}^N$, initially empty Vocabulary $V = \{\}$, start percentile q_{start} , end percentile q_{end} , decay rate δ , minimal restoration error ϵ_1 and ϵ_2 .

```

1 Sort  $Y$  in descending order to obtain  $\hat{Y} = \{\hat{\alpha}_j\}_{j=1}^N$ ;
2 Initialize iteration counter  $i = 1$ , error metric  $err = \infty$ ,
  current percentile  $q = q_{\text{start}}$ ;
3 while  $err > \epsilon_1$  do
4   Compute the  $q$ -percentile  $o_i$  of  $\hat{Y}$ ;
5   if  $o_i \leq \epsilon_2$  then
6     break;
7   end
8   Generate a new token  $w_i$  which satisfy  $o_i = g(w_i)$  and
  insert  $w_i$  into vocabulary  $V$ ;
9   Update  $\hat{Y}$  using:
          
$$\hat{y}_j = \begin{cases} \hat{y}_j, & \text{if } \hat{y}_j < o_i, \\ \hat{y}_j - o_i, & \text{otherwise} \end{cases}$$

10  Update the error metric:
          
$$err = \max \left\{ \frac{\hat{y}_j}{y_j} \right\}_{j=1}^N$$

11  Update percentile  $q$  with decay rate  $\delta$ :
          
$$q = \max(q \cdot \delta, q_{\text{end}})$$

12  Increase  $i$ :  $i \leftarrow i + 1$ ;
13 end
Result: Vocabulary  $V$ 

```

References

- [1] AppNexus. 2017. *Demystifying Auction Dynamics for Digital Buyers and Sellers*. <https://www.appnexus.com/sites/default/files/whitepapers/49344-CM-Auction-Type-Whitepaper-V9.pdf>
- [2] Mohammad Gheshlaghi Azar, Zhaohan Daniel Guo, Bilal Piot, Remi Munos, Mark Rowland, Michal Valko, and Daniele Calandriello. 2024. A general theoretical paradigm to understand learning from human preferences. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 4447–4455.
- [3] Jason Bigler. 2019. *Rolling out first price auctions to Google Ad Manager partners*. <https://blog.google/products/admanager/rolling-out-first-price-auctions-google-ad-manager-partners/>
- [4] Olivier Chapelle. 2015. Offline evaluation of response prediction in online advertising auctions. In *Proceedings of the 24th international conference on world wide web*. 919–922.
- [5] Sayak Ray Chowdhury, Anush Kini, and Nagarajan Natarajan. [n. d.]. Provably Robust DPO: Aligning Language Models with Noisy Feedback. In *Forty-first International Conference on Machine Learning*.
- [6] John M Crespi and Richard J Sexton. 2005. A Multinomial logit framework to estimate bid shading in procurement auctions: Application to cattle sales in the Texas Panhandle. *Review of industrial organization* 27, 3 (2005), 253–278.
- [7] Ying Cui, Ruofei Zhang, Wei Li, and Jianchang Mao. 2011. Bid landscape forecast in online ad exchange marketplace. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*. 265–273.

- [8] Jiabin Deng, Shiyao Wang, Kuo Cai, Lejian Ren, Qigen Hu, Weifeng Ding, Qiang Luo, and Guorui Zhou. 2025. Onerec: Unifying retrieve and rank with generative recommender and iterative preference alignment. *arXiv preprint arXiv:2502.18965* (2025).
- [9] Jingtong Gao, Bo Chen, Menghui Zhu, Xiangyu Zhao, Xiaopeng Li, Yuhao Wang, Yichao Wang, Hui Feng Guo, and Ruiming Tang. 2023. Scenario-Aware Hierarchical Dynamic Network for Multi-Scenario Recommendation. *arXiv preprint arXiv:2309.02061* (2023).
- [10] Getint. 2017. *RTB Auctions: Fair Play?* <https://blog.getint.com/rtb-auctions-fair-play-3b372d505089>
- [11] Djordje Gligorijevic, Tian Zhou, Bharatbhusan Shetty, Brendan Kitts, Shengjun Pan, Junwei Pan, and Aaron Flores. 2020. Bid shading in the brave new world of first-price auctions. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. 2453–2460.
- [12] Zhen Gong, Lv Yin Niu, Yang Zhao, Miao Xu, Haoqi Zhang, Zhenzhe Zheng, Zhilin Zhang, Rongquan Bai, Chuan Yu, Jian Xu, et al. 2023. MEBS: Multi-task End-to-end Bid Shading for Multi-slot Display Advertising. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*. 4588–4594.
- [13] Xian Guo, Ben Chen, Siyuan Wang, Ying Yang, Chenyi Lei, Yuqing Ding, and Han Li. 2025. OneSug: The Unified End-to-End Generative Framework for E-commerce Query Suggestion. *arXiv preprint arXiv:2506.06913* (2025).
- [14] Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems* 33 (2020), 6840–6851.
- [15] Jiani Huang, Zhenzhe Zheng, Yanrong Kang, and Zixiao Wang. 2024. From Second to First: Mixed Censored Multi-Task Learning for Winning Price Prediction. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining*. 295–303.
- [16] Xu Li, Michelle Ma Zhang, Zhenya Wang, and Youjun Tong. 2022. Arbitrary distribution modeling with censorship in real-time bidding advertising. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 3250–3258.
- [17] Hairan Liao, Lingxiao Peng, Zhenchuan Liu, and Xuehua Shen. 2014. iPinYou global rtb bidding algorithm competition dataset. In *Proceedings of the Eighth International Workshop on Data Mining for Online Advertising*. 1–6.
- [18] Xiao Lin, Xiaokai Chen, Linfeng Song, Jingwei Liu, Biao Li, and Peng Jiang. 2023. Tree based progressive regression model for watch-time prediction in short-video recommendation. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 4497–4506.
- [19] Zhutian Lin, Junwei Pan, Shangyu Zhang, Ximei Wang, Xi Xiao, Shudong Huang, Lei Xiao, and Jie Jiang. 2024. Understanding the ranking loss for recommendation with sparse user feedback. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 5409–5418.
- [20] Hongxu Ma, Kai Tian, Tao Zhang, Xuefeng Zhang, Han Zhou, Chunjie Chen, Han Li, Jihong Guan, and Shuigeng Zhou. 2024. Generative Regression Based Watch Time Prediction for Short-Video Recommendation. *arXiv preprint arXiv:2412.20211* (2024).
- [21] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems* 35 (2022), 27730–27744.
- [22] Shengjun Pan, Brendan Kitts, Tian Zhou, Hao He, Bharatbhusan Shetty, Aaron Flores, Djordje Gligorijevic, Junwei Pan, Tingyu Mao, San Gultekin, et al. 2020. Bid shading by win-rate estimation and surplus maximization. *arXiv preprint arXiv:2009.09259* (2020).
- [23] Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2023. Direct preference optimization: Your language model is secretly a reward model. *Advances in neural information processing systems* 36 (2023), 53728–53741.
- [24] Shashank Rajput, Nikhil Mehta, Anima Singh, Raghunandan Hulikal Keshavan, Trung Vu, Lukasz Heldt, Lichan Hong, Yi Tay, Vinh Tran, Jonah Samost, et al. 2023. Recommender systems with generative retrieval. *Advances in Neural Information Processing Systems* 36 (2023), 10299–10315.
- [25] Kan Ren, Jiarui Qin, Lei Zheng, Zhengyu Yu, Weinan Zhang, and Yong Yu. 2019. Deep landscape forecasting for real-time bidding advertising. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*. 363–372.
- [26] Kan Ren, Weinan Zhang, Ke Chang, Yifei Rong, Yong Yu, and Jun Wang. 2017. Bidding machine: Learning to bid for directly optimizing profits in display advertising. *IEEE Transactions on Knowledge and Data Engineering* 30, 4 (2017), 645–659.
- [27] Kan Ren, Weinan Zhang, Yifei Rong, Haifeng Zhang, Yong Yu, and Jun Wang. 2016. User response learning for directly optimizing campaign performance in display advertising. In *Proceedings of the 25th acm international conference on information and knowledge management*. 679–688.
- [28] Burr Settles and Mark Craven. 2008. An analysis of active learning strategies for sequence labeling tasks. In *proceedings of the 2008 conference on empirical methods in natural language processing*. 1070–1079.
- [29] Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, et al. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300* (2024).
- [30] Nian Si, San Gultekin, Jose Blanchet, and Aaron Flores. 2023. Optimal bidding and experimentation for multi-layer auctions in online advertising. *Available at SSRN 4358914* (2023).
- [31] Sarah Sluis. 2017. *Explainer: More On The Widespread Fee Practice Behind The Guardian's Lawsuit Vs. Rubicon Project*. <https://www.adexchanger.com/ad-exchange-news/explainer-widespread-fee-practice-behind-guardians-lawsuit-vs-rubicon-project/>
- [32] Jake Snell, Kevin Swersky, and Richard Zemel. 2017. Prototypical networks for few-shot learning. *Advances in neural information processing systems* 30 (2017).
- [33] Kihyuk Sohn, Honglak Lee, and Xinchen Yan. 2015. Learning structured output representation using deep conditional generative models. *Advances in neural information processing systems* 28 (2015).
- [34] Jie Sun, Zhaoying Ding, Xiaoshuang Chen, Qi Chen, Yincheng Wang, Kaiqiao Zhan, and Ben Wang. 2024. Cread: A classification-restoration framework with error adaptive discretization for watch time prediction in video recommender systems. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 9027–9034.
- [35] Zixiao Wang, Zhenzhe Zheng, Yanrong Kang, and Jiani Huang. 2024. Cost-Effective Active Learning for Bid Exploration in Online Advertising. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining*. 788–796.
- [36] Wush Chi-Hsuan Wu, Mi-Yen Yeh, and Ming-Syan Chen. 2015. Predicting winning price in real time bidding with censored data. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 1305–1314.
- [37] Wei Zhang, Brendan Kitts, Yanjun Han, Zhengyuan Zhou, Tingyu Mao, Hao He, Shengjun Pan, Aaron Flores, San Gultekin, and Tsachy Weissman. 2021. MEOW: A space-efficient nonparametric bid shading algorithm. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 3928–3936.
- [38] Tian Zhou, Hao He, Shengjun Pan, Niklas Karlsson, Bharatbhusan Shetty, Brendan Kitts, Djordje Gligorijevic, San Gultekin, Tingyu Mao, Junwei Pan, et al. 2021. An efficient deep distribution network for bid shading in first-price auctions. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 3996–4004.
- [39] Wen-Yuan Zhu, Wen-Yueh Shih, Ying-Hsuan Lee, Wen-Chih Peng, and Jiun-Long Huang. 2017. A gamma-based regression for winning price estimation in real-time bidding advertising. In *2017 IEEE International Conference on Big Data (Big Data)*. IEEE, 1610–1619.