

GAS: Generative Auto-bidding with Post-training Search

Yewen Li*
Nanyang Technological University
Singapore, Singapore
yewen001@e.ntu.edu.sg

Shuai Mao*
The Chinese University of Hong Kong
Hong Kong, China
smao@mae.cuhk.edu.hk

Jingtong Gao
City University of Hong Kong
Hong Kong, China
jt.g@my.cityu.edu.hk

Nan Jiang
Kuaishou Technology
Beijing, China
jiangnan07@kuaishou.com

Yunjian Xu
The Chinese University of Hong Kong
Hong Kong, China
yjxu@mae.cuhk.edu.hk

Qingpeng Cai[†]
Kuaishou Technology
Beijing, China
caiqingpeng@kuaishou.com

Fei Pan
Kuaishou Technology
Beijing, China
panfei05@kuaishou.com

Peng Jiang
Kuaishou Technology
Beijing, China
jiangpeng@kuaishou.com

Bo An
Nanyang Technological University
Skywork AI
Singapore, Singapore
boan@ntu.edu.sg

Abstract

Auto-bidding is essential in facilitating online advertising by automatically placing bids on behalf of advertisers. Generative auto-bidding, which generates bids based on an adjustable condition using models like transformers and diffusers, has recently emerged as a new trend due to its potential to learn optimal strategies directly from data and adjust flexibly to preferences. However, generative models suffer from low-quality data leading to a mismatch between the condition, like *return to go*, and true action value, especially in long sequential decision-making. Besides, the majority preference in the dataset may hinder models' generalization ability on minority advertisers' preferences. While it is possible to collect high-quality data and retrain multiple models for different preferences, the high cost makes it unaffordable, hindering the advancement of auto-bidding into the era of large foundation models. To address this, we propose a flexible and practical **Generative Auto-bidding** scheme using post-training Search, termed **GAS**, to refine a base policy model's output and adapt to various preferences. We use weak-to-strong search alignment by training small critics for different preferences and an MCTS-inspired search to refine the model's output. Specifically, a novel voting mechanism with transformer-based critics trained with policy indications could enhance search alignment performance. Additionally, utilizing the search, we provide a fine-tuning method for high-frequency preference scenarios considering computational efficiency. Extensive experiments conducted

on the real-world dataset and online A/B test on the Kuaishou advertising platform demonstrate the effectiveness of GAS, achieving significant improvements, e.g., 4.60% increment of target cost.

CCS Concepts

• **Information systems** → **Computational advertising**.

Keywords

Auto-bidding, Generative Model, Search, Preference Alignment

ACM Reference Format:

Yewen Li, Shuai Mao, Jingtong Gao, Nan Jiang, Yunjian Xu, Qingpeng Cai, Fei Pan, Peng Jiang, and Bo An. 2025. GAS: Generative Auto-bidding with Post-training Search. In *Companion Proceedings of the ACM Web Conference 2025 (WWW Companion '25)*, April 28-May 2, 2025, Sydney, NSW, Australia. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3701716.3715226>

1 Introduction

The rapid digitalization of commerce has expanded online advertising platforms' reach, making them essential for engaging audiences and boosting sales [12, 45]. Due to numerous impression opportunities, manually adjusting bids to optimize costs under budget and KPI constraints is impractical. To address this, ad platforms now offer auto-bidding services that automate the bidding process using advanced strategies [4, 9, 34]. These strategies consider factors from immediate or historical bidding information, such as the distribution of impression opportunities and remaining budgets [26]. Besides, according to different types of advertisers, the strategy should take their different preferences into consideration. For example, brand advertisers, who aim for long-term growth and awareness, typically aim to show their ads to as many people as possible under constraints like average cost per impression. However, performance advertisers, who focus on maximizing the value of winning impressions, aim to maximize conversions with cost-per-conversion constraints [19]. To meet these diverse needs, advertising platforms like Google, Facebook, and Alibaba provide tailored multiple bidding strategies for their customers [2, 13, 17]. Additionally, facing the dynamic advertising environment, strategies

*Both authors contributed equally to this research.

[†]Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WWW Companion '25, April 28-May 2, 2025, Sydney, NSW, Australia

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-1331-6/2025/04

<https://doi.org/10.1145/3701716.3715226>

must be regularly optimized to stay closely aligned with customers' preferences, thereby achieving long-term commercial benefits [18].

Reinforcement learning [37, 44] (RL) has long been a leading method for enhancing auto-bidding strategies by training agents with an advertising simulator or offline advertising logs. However, RL approaches are predominantly based on the Markovian decision process (MDP), which assumes that future states are determined solely by the current single step's state and action. Recent statistical analyses [18] question this assumption in auto-bidding, revealing a strong correlation between the lengths of historical state sequences and subsequent states. This finding suggests that relying solely on the most recent state can lead to instability in the unpredictable online advertising environment. Moreover, the preference for the RL strategy is not easy to control. USCB [19] proposes calculating an optimal solution under multiple constraints using historical data and then training RL to refine the strategy. However, once deployed, the strategy's preference is fixed, limiting interactivity and controllability. Therefore, this has led to a new trend in developing generative auto-bidding methods based on conditional generative models like transformers and diffusers. With a vector condition representing the preferences, these methods directly output an action or even a trajectory without MDP assumption. For example, decision transformers [7] can utilize extensive historical information for decision-making. Diffusers [1] can directly generate a trajectory for planning given a condition. More importantly, generative models provide flexibility in controlling preferences by simply modifying their condition's value during deployment. As large foundational models based on generative models have shown remarkable progress in applications like natural language processing and computer vision, exemplified by ChatGPT [33] and Stable Diffusion [41], it is foreseeable that the auto-bidding field will also advance towards the foundation model era. This will involve developing a decision foundation model to learn optimal decision strategies directly from large-scale datasets.

However, there are also two major intrinsic challenges in applying generative auto-bidding methods. First, the performance of generative bidding methods is significantly influenced by the quality of the dataset, where the collected conditions, *i.e.*, return to go, cannot reflect the true value of the action. For instance, a good action a_t followed by a bad future action $a_i, i > t$, can result in a low return to go for a_t , and vice versa. Thus, the learned strategy is hard to achieve the optimum due to a mismatch between the condition and the true action value during training. Second, realistic bidding tasks often involve varying preferences over time, but generative methods are always trained or even biased to imitate the majority preference [32], necessitating retraining to adapt to new minority preferences. However, since the scaling law could also happen in this auto-bidding field, the bidding model based on a foundation model like a transformer would get larger and larger, making it impractical and costly to retrain a set of large decision transformers for varying preferences, hindering the advancement of auto-bidding into the era of large foundation models. Therefore, a question is raised: "Could we only use one policy model to efficiently achieve optimum strategy under various preferences?"

To address these issues, we propose a Generative Auto-bidding scheme using post-training Search, termed **GAS**, to refine a single base policy model's output and adapt to various preferences

efficiently. Instead of revising the condition setting and retraining the policy model to fit various preferences, we adopt the weak-to-strong search alignment idea, *i.e.*, refining the large foundation model's output with small models [6]. Specifically, we train a set of small critics to assess the value of different preferences, then use a search method to refine the model's output inspired by Monte Carlo Tree Search (MCTS) [22]. In this scheme, the mismatch between the condition value and the true action value, approximated by the critics using Q-learning, would be alleviated. This is due to the Q-learning based on a Bellman backup only requires the current reward that is not affected by the future trajectory. A large dataset collected by various policies is also beneficial to training a critic as it could assess actions of different quality, alleviating the out-of-domain overestimation issue [23]. This scheme could also achieve better alignment with different preferences by refining the action guided by these critics without retraining or fine-tuning the model. To summarize, our contribution could be summarized as follows:

- This paper proposes a flexible and practical framework using a post-training search method for refining and aligning generative auto-bidding models to various preferences, which unlocks the door of auto-bidding foundation models.
- To enhance the accuracy of the search process, we utilize transformer-based critics that are aware of the underlying policy by leveraging historical bidding sequences and introduce a novel voting mechanism to enhance value approximation accuracy during the value backpropagation stage of the search process.
- In addition to performing a search during testing time, we provide a fine-tuning method for high-frequency preference scenarios or computational-efficiency-sensitive scenarios.
- Extensive experiments conducted on a large real-world dataset and an online A/B test on the Kuaishou advertising platform demonstrate the effectiveness of the proposed generative auto-bidding scheme.

2 Preliminary

2.1 Problem Statement

During a time period, suppose there are H impression opportunities arriving sequentially and indexed by i . In a bidding platform, advertisers submit bids to compete for each impression opportunity. An advertiser wins the impression if its bid b_i is higher than the bids of others. Upon winning, the advertiser incurs a cost of c_i , which is typically the highest bid of the other advertisers in a second-price auction setting. During the period, the goal of an advertiser is to maximize the sum value of winning impressions $\sum_i o_i v_i$, where v_i is the value of impression i and o_i is the binary indicator of whether the advertiser wins impression i . Besides, the budget and multiple KPI constraints [19] are critical for an ad advertiser to control the ad delivery performance. The budget constraint is considered as $\sum_i o_i c_i \leq B$, where c_i is the cost of impression i and B is the budget. Other KPI constraints are more complicated and can be classified into two categories. The first category is the cost-related (CR) constraints, which restrict the unit cost of certain advertising events, such as CPC and CPA. The second category is the non-cost-related (NCR) constraints, which restrict the average advertising effect,

such as CTR and CPI. For simplicity, we consider the auto-bidding with cost-related constraints, which have a unified formulation: $\frac{\sum_i c_{ij} o_i}{\sum_i p_{ij} o_i} \leq C_j$, where C_j is the upper bound of j -th constraint provided by the advertiser. p_{ij} can be any performance indicator, e.g., return or constant. c_{ij} is the cost of constraint j . Given J constraints, we have the Multi-Constrained Bidding (MCB):

$$\begin{aligned} & \text{maximize} && \sum_i o_i v_i \\ & \text{s.t.} && \sum_i o_i c_i \leq B \\ & && \frac{\sum_i c_{ij} o_i}{\sum_i p_{ij} o_i} \leq C_j, \quad \forall j \\ & && o_i \in \{0, 1\}, \quad \forall i \end{aligned} \quad (1)$$

A previous study [19] has already shown the optimal solution:

$$b_i^* = \lambda_0 v_i + \sum_{j=1}^J \lambda_j p_{ij} C_j, \quad (2)$$

where b_i^* is the optimal bid prediction for the impression i . λ_j , $j \in \{0, \dots, J\}$ are the optimal bidding parameters. However, due to the uncertainty and dynamics of the advertising environment, these optimal bidding parameters are hard or even impossible to calculate directly. Different types of advertisers may have different preferences expressed by various constraint compositions. For example, max-return bidding advertisers only consider the budget constraint and target-CPA bidding advertisers consider both the budget constraint and the CPA constraint.

2.2 Decision-making Process for Auto-bidding

As the advertising environment is highly dynamic, the optimal bidding parameters should be adjusted regularly to maximize the total received value over the entire period. This leads to the formulation of the auto-bidding task as a sequential decision-making task. We consider a standard decision-making setup consisting of an auto-bidding agent interacting with advertising environment E in discrete timesteps. At each timestep t of a bidding period, the agent receives a state $s_t \in \mathcal{S}$ describing the real-time advertising status and then outputs an action $a_t \in \mathcal{A}$ for the final bid. The advertising environment has an underlying unknown state transition dynamic \mathcal{T} . Under an MDP assumption, the transition dynamic \mathcal{T} could be expressed by $\mathcal{T} : s_t \times a_t \rightarrow s_{t+1}$, i.e., the next state $s_{t+1} \in \mathcal{S}$ is determined by the current state s_t and action a_t . In this case, the agent's policy is expressed as $\pi(a_t | s_t)$. Without an MDP assumption, the next state could be determined by more factors like the historical trajectory τ . After transitioning to the next state, the advertising environment would emit a reward r_t representing the value contributed to the objective obtained within the time period t . Repeating this process until the end of a bidding period, one day for instance, the goal of the auto-bidding agent is to maximize the total received value over the entire period as stated in Equation 1. A detailed description of our modeling are listed below:

- s_t : the state is a collection of information that describes the advertising status from the perspective of a campaign. The information should in principle reflect time, budget consumption and KPI constraints satisfying status, such as left time,

left budget, budget consumption speed, current KPI ratio for constraint j (i.e. $(\sum_i c_{ij} o_i / \sum_i p_{ij} o_i) / C_j$), etc.

- a_t : adjustment to the bidding parameters λ_j , $j = 0, \dots, J$, at the time period t , and modeled as $(a_t^{\lambda_0}, \dots, a_t^{\lambda_J})$.
- r_t : at time-step t , let C be the candidate impression set between step t and $t+1$, then a typical setting of the reward could be the value contributed to the objective obtained on this C within the time period. For simplicity, we will omit C in the following sections.

3 Generative Auto-bidding with Search

In this section, we first introduce how an MCTS-inspired post-training search process is developed to refine a base policy model's output action and then, we introduce two practical paradigms of applying this post-training search in auto-bidding.

3.1 MCTS-inspired Post-training Search

As the decision transformer is widely used as a backbone in generative decision-making, we also adopt the decision transformer as our backbone for auto-bidding and the policy for generating an action could be formulated as

$$a_t \sim \pi_{dt} = \text{DT}_\theta(a | s_{\leq t}, a_{< t}, R_{\leq t}), \quad (3)$$

where the condition R_t is the *return to go* at time step t , i.e.,

$$R_t = \sum_{i=t \sim T} \gamma^{i-t} r(s_i, a_i), \quad (4)$$

where γ is the discounting factor and $r(s_i, a_i)$ is a rewarding function representing the preference, e.g., it could be set as the $o_i v_i$ representing a preference considering only the value.

Our purpose of the search scheme is to find a better action that could align better with the preference like a higher value. Applying the typical MCTS method to the decision-making process should involve four parts in every time step t :

- Selection: start from a root state node s_t and randomly select successive valid child action nodes a_t^i within exploration budget N , i.e., $i = 1 \sim N$.
- Expansion: unless the child action nodes end the bidding process, create one further child state node s_{t+1}^i with respect to the transition dynamic $s_{t+1}^i \sim \mathcal{T}$.
- Simulation: complete one rollout from node s_{t+1}^i given the policy $\pi(a|s)$ until the end.
- Backpropagation: use the result of the rollout to update value information in the nodes a_t^i given root s_t .

After these four parts, we could choose a final action a_t^i based on some principles such as balance between exploration and exploitation. In this work, we choose the action of maximum preference value to execute given a state s_t with uncertainty consideration modeled as a random action selection operation. However, unlike the GO game where we can simulate all possible moves, this could be impractical in the bidding scenarios as bidding is a typical partially observable MDP (POMDP) task where other advertisers' behaviors are not predictable. Therefore, we make modifications to the typical MCTS process by approximating the expansion and simulation process together via learning an enhanced transformer-based Q-value function, without actual simulation in a simulator

(which is also invalid). Now, we will introduce how we implement **GAS** in the following three parts, with an illustration in Figure 1(a).

3.1.1 Selection. Given a decision transformer policy DT_θ , we generate a base action a_t^i first, and then we perpetuate it by multiplying a random factor uniformly between 90% and 110% to get $N - 1$ random actions $\{a_t^i\}_{i=1:N-1}$, expressed as

$$a_t^i = a_t * \epsilon, \epsilon \sim \mathcal{U}(90\%, 110\%). \quad (5)$$

We also preserve the initial base action a_t to get the final N action proposals $\{a_t^i\}_{i=1:N} = \{a_t^i\}_{i=1:N-1} \oplus a_t$. Then, the selection process is choosing an action from these action proposals.

3.1.2 Expansion and Simulation. As we cannot rollout in a simulator or real environment, we need to approximate this rollout process after taking an action proposal a_t^i given s_t . As we actually need only the result of the rollout, we could directly estimate the value of a_t^i given s_t using a Q-value function, expressed as

$$Q_\phi(s_t, a_t^i; \pi) = r(s_t, a_t^i) + \mathbb{E}_{s_{t+1} \sim \mathcal{T}, a_{t+1} \sim \pi} Q_\phi(s_{t+1}, a_{t+1}; \pi). \quad (6)$$

We could employ the IQL [23] method to the learning of Q_ϕ , which introduces an extra value net $V_\psi(s)$ to avoid the overestimation issues due to the distributional shift, expressed as

$$\mathcal{L}_V(\psi) = \mathbb{E}_{(s,a) \sim \mathcal{D}} [L_2^\tau(Q_\phi^\tau(s,a) - V_\psi(s))], \quad (7)$$

where $L_2^\tau(u) = |\tau - \mathbb{1}(u < 0)|u|^2$ is an expectile regression loss. The value net $V_\psi(s)$ is used to the Q-value learning:

$$\mathcal{L}_Q(\phi) = \mathbb{E}_{(s_t, a_t, s_{t+1}) \sim \mathcal{D}} [r(s_t, a_t) + \gamma V_\psi(s_{t+1}) - Q_\phi(s_t, a_t)]^2. \quad (8)$$

As indicated in Equation 6, the Q is coupled with the underlying policy π in the latter expectation term. However, $Q_\phi(s_t, a_t)$ only receives a single state-action pair without policy indication, leading to a value prediction based on the underlying policy π_β that collected the dataset actually. This causes a policy gap in the generative auto-bidding, as its policy π_ϵ indicated by various conditions is different from π_β , leading to a biased value approximation.

Rollout via QT. To address the distributional gap by representing the actual policy π_ϵ , we employ the historical trajectory for Q-value learning with a transformer utilizing its sequential modeling ability, termed *QT*, *i.e.*,

$$Q_\phi^{\pi_\epsilon}(s_t, a_t^i) = Q_\phi(s_t, a_t^i; \pi_\epsilon) = \text{QT}_\phi(s_t, a_t^i; s_{<t}, a_{<t}). \quad (9)$$

This could be beneficial with a large-scale pertaining set, which contains trajectories collected by various policies because it could empirically use the historical trajectory to predict the future trajectory, *i.e.*, the rollout $\{s_{\leq t}, a_{\leq t}\} \rightarrow \{s_{t+1:T}, a_{t+1:T}\}$.

After training, the $\text{QT}_\phi(s_t, a_t^i; s_{<t}, a_{<t})$ could return an approximation of the value of the rollout until the end of bidding.

3.1.3 Backpropagation. As the overestimation issue of the Q-value function is notorious, an inaccurate value estimation for the rollout could bias the backpropagation to provide an inaccurate value of action node a_t^i given root node s_t , leading to the execution of a bad action. To alleviate the overestimate issue, we provide a Q-voting mechanism for value assessment, which is based on a heuristic insight of the consensus.

Value Backpropagation via Q-voting. Given the success of offline RL, where Q-value is employed to learn a policy improving over the behavior policy that collected the dataset, we could intuitively

believe different randomly trained Q-value nets could achieve an agreement in giving the true best action higher value with high probability. Besides, due to the overestimation being an occasional error, there would be no agreement in giving a specific action a higher value. Formally, if we independently train M Q-value nets $\{Q_{\phi_k}^{\pi_\epsilon}\}_{k=1:M}$ with different random seeds and we have N action proposals $\{a_t^i\}_{i=1:N}$ with a ground-truth best action a_t^j , we model the insight of the consensus by a probability density function $p(a|Q)$ of an action a being the best action given a Q-value net, *i.e.*,

$$\text{Consensus: } p(a_t^j | Q_{\phi_k}^{\pi_\epsilon}) > p(a_t^{i \neq j} | Q_{\phi_k}^{\pi_\epsilon}), \forall k \in \{1, \dots, M\}. \quad (10)$$

Thus, if we only use a single Q, the final win rate of choosing a_t^j over not choosing it is defined as

$$\mathcal{R}^k := \frac{p(a_t^j | Q_{\phi_k}^{\pi_\epsilon})}{\sum_{i \neq j} p(a_t^i | Q_{\phi_k}^{\pi_\epsilon})}. \quad (11)$$

With this consensus, we could apply the majority voting method to increase the win rate \mathcal{R}^k . For brevity, we assume $p(a_t^i | Q_{\phi_k}^{\pi_\epsilon})$ is the same for all $k \in \{1, \dots, M\}$. The probability of choosing an action a_t^i as the final action by majority voting is expressed as

$$p(a_t^i | \{Q_{\phi_k}^{\pi_\epsilon}\}_{k=1:M}) = \sum_{l=\lfloor \frac{M}{2} \rfloor + 1}^M \binom{M}{l} p(a_t^i | Q_{\phi_k}^{\pi_\epsilon})^l (1 - p(a_t^i | Q_{\phi_k}^{\pi_\epsilon}))^{M-l}.$$

Employing the *Condorcet's jury theorem* [5], we could get

$$\mathcal{R}^{1:M} = \frac{p(a_t^i | \{Q_{\phi_k}^{\pi_\epsilon}\}_{k=1:M})}{\sum_{i \neq j} p(a_t^i | \{Q_{\phi_k}^{\pi_\epsilon}\}_{k=1:M})} > \mathcal{R}^k. \quad (12)$$

To ease reading, let us assume $p(a_t^1 | Q_{\phi_k}^{\pi_\epsilon}) = 0.4$, $p(a_t^2 | Q_{\phi_k}^{\pi_\epsilon}) = 0.3$, and $p(a_t^3 | Q_{\phi_k}^{\pi_\epsilon}) = 0.3, \forall k$. Then, we can get $\mathcal{R}^{1:3} = 0.81 > \mathcal{R}^k = 0.67$.

To avoid the useless majority voting outcomes, *i.e.*, there are no actions that gain more votes than all others, we provide a soft majority voting mechanism using Q-value, termed **Q-voting**, implemented by two steps:

- step 1: for a $Q_{\phi_k}^{\pi_\epsilon}$, its vote for each action based on min-max normalization over all a_t^n is

$$v(a_t^i | Q_{\phi_k}^{\pi_\epsilon}) = \frac{Q_{\phi_k}^{\pi_\epsilon}(s_t, a_t^i) - \min_n \{Q_{\phi_k}^{\pi_\epsilon}(s_t, a_t^n)\}}{\max_n \{Q_{\phi_k}^{\pi_\epsilon}(s_t, a_t^n)\} - \min_n \{Q_{\phi_k}^{\pi_\epsilon}(s_t, a_t^n)\}} \in [0, 1]. \quad (13)$$

- step 2: the final total votes gained of an action a_t^i is

$$v(a_t^i | \{Q_{\phi_k}^{\pi_\epsilon}\}_{k=1:M}) = \sum_{k=1}^M v(a_t^i | Q_{\phi_k}^{\pi_\epsilon}). \quad (14)$$

After the above MCTS-inspired search process, we can now get a refined a_t^n that should have higher preference values than a_t .

3.2 Bidding with GAS

There could be two ways to apply the GAS: *i)* searching in the testing time, and *ii)* using the search to fine-tune the base policy model. Since both of these approaches need to train a set of critics, we introduce how we train critics representing different preferences.

Preference Representation. The preference is expressed by the setting of the reward function. We introduce three examples:




Figure 1: a) The pipeline of GAS for refining a base action a_t to an action a_t^j better aligned to the preferences represented by the critics, involving three stages: (1) selection by a sampled noisy perturbation, (2) expansion and simulation approximated by the transformer-based Q-value net (QT), and (3) backpropagation through a Q-voting mechanism for better value assessment. b) The online auto-bidding system of Kuaishou including the interaction between the bidding environment and our GAS method.

- Preference that maximizes the winning value under only budget constraint without considering the KPI constraints: $r_t = o_t v_t$, which is actually the Max Return Bidding;
- Preference that maximizes a comprehensive performance considering both the winning value and the KPI constraints: $r_t = o_t v_t \cdot \frac{1}{J} \sum_j \min\left\{\left(\frac{C_j}{c_{tj} o_t / p_{tj} o_t}\right)^\beta, 1\right\}, \beta > 1$;
- Preference that prefers more on the KPI constraints by introducing a larger and controllable reweight factor w : $r_t = o_t v_t + \frac{w}{J} \sum_j \min\left\{\left(\frac{C_j}{c_{tj} o_t / p_{tj} o_t}\right)^\beta, 1\right\}, \beta > 1$.

Then, we could use these reward functions to train the critics based on Equation 8.

3.2.1 Inference with Search. By repeating the search process introduced in Section 3.1 at every time step during inference, we can get a testing-time version of our method, termed *GAS-infer*, which could be directly deployed with a base policy model and multiple critics. For a more clear understanding, we provide a detailed procedure in Algorithm 1.

3.2.2 Finetune with Search. As the search method has the potential to find better action surrounding a base action, especially when the base action is poor quality or misaligned with the preference, we could use search to enhance the training data and fine-tune the base policy model. Firstly, given a data point in the dataset $\{s_t, a_t^{gt}\}$, we could do a search for a_t^{gt} and we could get a better action a_t^p aligning better to the preference. Then, a simple supervised

Algorithm 1 Inference with Search (*GAS-infer*)

Input: Base policy model π_θ , critics $\{Q_{\phi_k}^{\pi_\epsilon}\}_{k=1}^M$, initial state s_0
Output: Action a_t^* at every time step t of the bidding period

- 1: Initialize $s_t \leftarrow s_0$
- 2: **for** each time step t **do**
- 3: Generate base action $a_t \sim \pi_\theta(a|s_{\leq t}, a_{< t}, R_{\leq t})$
- 4: Generate $N - 1$ random actions $\{a_t^i\}_{i=1}^{N-1}$ by perturbing a_t with a random factor $\epsilon \sim \text{Uniform}(90\%, 110\%)$
- 5: Form action proposals $\{a_t^i\}_{i=1}^N = \{a_t^i\}_{i=1}^{N-1} \cup \{a_t\}$
- 6: **for** each action proposal a_t^i **do**
- 7: Estimate value $v(a_t^i | \{Q_{\phi_k}^{\pi_\epsilon}\}_{k=1}^M)$ by Q-voting in Eq. 14
- 8: **end for**
- 9: Select action $a_t^* = \arg \max_{a_t^i} v(a_t^i | \{Q_{\phi_k}^{\pi_\epsilon}\}_{k=1}^M)$
- 10: Execute action a_t^* and observe new state s_{t+1}
- 11: Update $s_t \leftarrow s_{t+1}$
- 12: **end for**
- 13: **return**

fine-tune (sft) could be implemented based on a loss for training DT_θ :

$$\mathcal{L}_{\text{DT}}^{\text{sft}}(\theta) = \text{mse}(a_t, a_t^p). \quad (15)$$

Although there could be more choices of preference alignment methods, e.g., DPO [40] and RLHF [35], they are typically trajectory-level based on a query dataset and are reported to be unstable, so we leave them as future works.

4 Experiment

4.1 Setup

Datasets. Different from the previous dataset preparation procedure that collects private bidding logs from a non-open-source advertising auto-bidding system, we utilize a new public large-scale real-world bidding dataset, termed AuctionNet, released by Alibaba [43]. This dataset, to our knowledge the largest available, includes over 500 million records and offers a sparse version for more challenging cases. More details are in Appendix A.

Evaluation Metrics. For simplicity, we adopt three metrics to evaluate the performance:

- **Value:** the total received value during the bidding period, calculated as $\sum_i o_i v_i$;
- Exceeding rate of the KPI constraints (**ER**): by introducing a binary indicator $\mathbb{I}(x_j^h, C_j)$ whether the final KPI performance $x_j^h = \sum_i c_{ij} o_i / \sum_i p_{ij} o_i$ during a period h exceeds the given constraint C_j , and assuming we have H periods in total, we have a exceeding rate of the KPI constraints during H periods, defined by

$$ER = \frac{1}{H} \sum_{h=1 \sim H} \sum_{j=1 \sim J} \mathbb{I}(x_j^h, C_j). \quad (16)$$

- **Score:** by introducing a penalty term

$$penaty_j = \min\left\{\left(\frac{C_j}{\sum_i c_{ij} o_i / \sum_i p_{ij} o_i}\right)^\beta, 1\right\}, \beta = 2 \quad (17)$$

the score is a balance between the value and the KPI constraint, *i.e.*,

$$score = \left(\sum_i o_i v_i\right) \times \min\{penaty_j\}_{j=1 \sim J} \quad (18)$$

Baselines. We compare our method to a wide range of baselines involving both RL and generative methods. For RL, we compare with online RL approach **USCB** [19], offline approaches **BCQ** [14], **CQL** [24], and **IQL** [23]. For generative methods, we compare with **DiffBid** [18], a generative method based on diffusers, and Decision Transformers **DT** [7] and its variants **CDT** [27], a DT-based method that considers a vector of multiple constraints, **DT-score**: a DT-based method that utilizes the reward function considering both the winning value and the KPI constraints.

Implementation Details. The hyperparameters of baselines are set based on the default values referenced in original papers, with further tuning performed to optimize performance. For GAS, it has two components: a base policy model and multiple QT networks. The base policy model could be any model and we choose the **DT-score** with hyperparameters referenced from the official code provided from [48], which a smaller learning rate at $1e^{-5}$ for fine-tune specifically. For QT network, we employ 6 attention layers with 8 heads for each and set the hidden size as 512 for all layers with 14M parameters in total, which is lightweight. The total training steps are 400k. We adopt the AdamW optimizer [29] with a learning rate of $1e^{-4}$, and set the batch size as 128. During training, we utilize the PyTorch framework on two NVIDIA H100 GPUs. For more detailed hyperparameters of QT networks, please refer to Table 6 from Appendix A.2.

4.2 Performance Comparison with Baselines

In this experiment, we conducted a performance comparison of various baseline methods under different settings, including varying datasets and budget constraints in the MCB bidding. The dataset used for this evaluation is AuctionNet and its sparse version, and the budgets are set at 50%, 75%, 100%, 125%, and 150% of the maximum allowable budget. The results are indicated by the score as a comprehensive assessment of the performance. Our method is based on the baseline policy model **DT-score**.

The results are presented in Table 1, which indicates that both **GAS-infer** and **GAS-sft** consistently outperform other methods across all budget settings, achieving the highest scores for all the respective budgets. Other methods such as DT, CDT, and **DT-score** also perform well demonstrating the power of generative models compared to the traditional RL methods like IQL. We found DiffBid does not perform well in this large-scale task, where a possible reason could be the introduced extra challenges in predicting a whole long trajectory and learning an inverse dynamic model, though it could theoretically alleviate the mismatch between return to go and true action value. Regarding the stability, as shown in Fig. 2(d), GAS also outperforms the base policy model to be more stable. **GAS-sft** does not perform as well as **GAS-infer**, which could be attributed to encoding additional preference information parameterized in the extra introduced critics and performing clearly during testing. In contrast, the **GAS-sft** backbone fuses the actor and critic within the original model, which may cause ambiguity and weak constraints on its critic ability, leading to worse performance. Note that the Q-voting procedure for value assessment can be executed in parallel, with each Q-voting step taking around 0.1 seconds. This duration is considerably small compared to the interval from t to $t + 1$ at 30 minutes. This suggests that the **GAS-infer** method is particularly effective in large-scale auto-bidding tasks.

4.3 Preference Alignment Performance

To testify whether search with critics of different preferences could improve the performance of a base policy model to specific preference, we conduct preference alignment experiments in Table 2. In this experiment, we evaluated the alignment results of different preference paradigms using two search methods, **GAS-infer** and **GAS-sft**, compared to the baseline **DT-score** method. The preferences considered include Score-first, Value-first, and ER-first. Table 2 presents the results, showing that **GAS-infer** and **GAS-sft** consistently provide better alignment performance across different preferences compared to the base policy model, supporting the effectiveness of preference alignment by the search methods.

4.4 Ablation Study

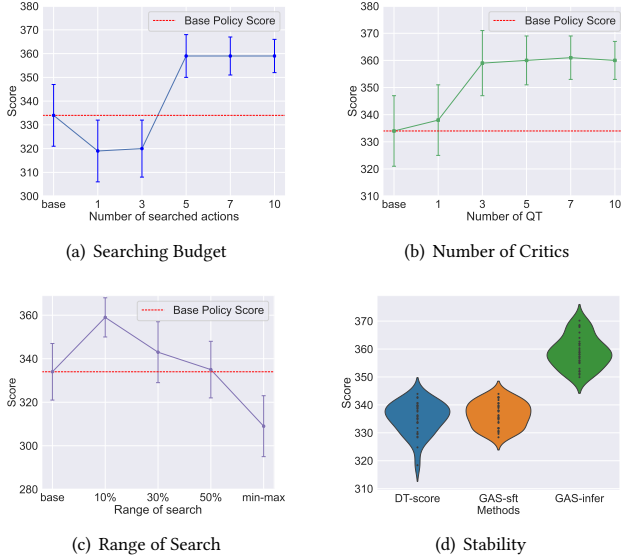
Range of Search. The search range directly impacts the balance between exploration (trying out new actions to discover their potential) and exploitation (focusing on actions known to yield good results). A smaller search range means the algorithm will explore fewer actions, potentially missing out on some good actions but focusing more on refining the evaluation of the explored actions. A larger search range means the algorithm will explore more actions, increasing the chances of finding the best possible actions but at the cost of spending less time on each action, which might lead to less

Table 1: Comparison with baselines in different settings in the MCB task. The best results are bolded and the base policy model’s results are underlined to demonstrate the performance improvement of our GAS-infer method in the last column.

Dataset	Budget	USCB	BCQ	CQL	IQL	DiffBid	DT	CDT	DT-score	GAS-sft	GAS-infer	improve
AuctionNet	50%	86	190	113	164	54	191	174	<u>178</u>	192	193	8.42%
	75%	135	259	139	232	100	265	242	<u>268</u>	273	287	7.08%
	100%	157	321	171	281	152	329	326	<u>334</u>	336	359	7.48%
	125%	220	379	201	355	193	396	378	<u>395</u>	398	409	3.54%
	150%	281	429	238	401	234	450	433	<u>441</u>	448	461	4.53%
AuctionNet-sparse	50%	11.5	17.7	12.8	16.5	9.87	14.8	11.2	<u>17.7</u>	18.0	18.4	3.95%
	75%	14.9	24.6	16.7	22.1	15.4	22.9	18.0	<u>25.9</u>	26.4	27.5	6.17%
	100%	17.5	31.1	22.2	30.0	19.5	29.6	31.2	<u>33.2</u>	34.3	36.1	8.73%
	125%	26.7	34.2	28.6	37.1	25.3	34.3	31.7	<u>39.6</u>	39.7	40.0	1.01%
	150%	31.3	37.9	35.8	43.1	30.8	44.5	39.1	<u>44.7</u>	46.4	46.5	4.02%

Table 2: Alignment performance with different preferences by two paradigms of search compared to base DT-score.

Dataset	Preference	Base	Score-first		Value-first		ER-first	
		DT-score	GAS-infer	GAS-sft	GAS-infer	GAS-sft	GAS-infer	GAS-sft
AuctionNet	Score \uparrow	334	359	336	358	335	348	346
	Value \uparrow	373	391	339	402	375	370	359
	ER \downarrow	37.8%	27.0%	10.4%	29.1%	43.7%	23.5%	16.7%
AuctionNet-sparse	Score \uparrow	33.2	36.1	34.3	35.9	34.1	34.5	33.5
	Value \uparrow	36.8	38.9	36.9	39.2	37.8	36.4	36.3
	ER \downarrow	27.0%	29.1%	25.6%	31.3%	33.3%	20.8%	22.9%

**Figure 2: a-c) Results for various ablation studies regarding the search setting; d) Comparison of the stability.**

accurate evaluations. An extreme case could be to search the whole action space, while it could be low-efficiency. Considering the computational efficiency, we randomly search 5 actions within a search range and the ablation study results for different search ranges’ results are shown in Fig. 2(c). The experimental results suggest that for the AuctionNet dataset, a 10% search range is better, providing the highest performance. Increasing the search range beyond 10% leads to diminishing returns, and the min-max approach is the least effective due to the limited action budget. More importantly, as the optimal performance is achieved with a small search range of

around 10%, it indicates that a small refinement based on the base policy model could be good enough, leaving more space for a more advanced value assessment method.

Number of Critics. As stated in the Q-voting mechanism, more critics would increase the accuracy of the value assessment empirically, so it would be practical to see how many critics are enough. Therefore, we conduct an ablation study on the number of critics. The search range is fixed as $\pm 10\%$ and five random actions are selected for value assessment. As shown in Fig. 2(b), the experimental results suggest that increasing the number of QT from one to more improves performance significantly, up to an optimal point of seven critics. Beyond this point, further increases from three critics do not yield significant improvements, suggesting that three critics are enough for the current auto-bidding task.

Searching Budget. Sampling more actions from a search range could increase the possibility of finding the optimal action, we conduct an ablation study to investigate it. Note that, in this investigation, we do not use the original base action. As shown in Fig. 2(a), we find that the performance can be efficiently enhanced by increasing the number of actions from one to five. However, beyond five actions, the performance remains unchanged. This suggests that with a higher search budget, we are more likely to find actions that are optimal or near-optimal.

Effectiveness of QT. As the rollout process is approximated by the Q-value function, the accuracy of the Q-value function in predicting the expected return with the underlying policy becomes essential. Without historical trajectory information, predicting the rollout process given a state s_t and action a_t is unclear and highly random. This insight is supported by our experimental results in Table 3, where the performance of QT with historical trajectory as policy representation achieves much higher performance than the Q-value function based on a vanilla state-action pair. We also compare QT to various methods without a Q-value net to testify to

the effectiveness of MCTS search, including greedy search, and random/mean selection. Greedy search selects the top action with the highest predicted immediate reward, random selection randomly selects an action, and mean selection takes the average value of 5 random actions. As shown in Table 3, our method significantly outperforms other methods.

Table 3: Effectiveness of QT.

AuctionNet	DT-score	Greedy	Mean	Random	GAS w/o QT	GAS-infer
Score	334	228	312	319	292	359

5 Live Experiment

To verify the effectiveness of our proposed GAS, we have deployed it on Kuaishou’s advertising system, see Figure 1(b). The deployment scenario is the Multiple Constraint Bidding(MCB). Under the MCB setting, the advertisers set the budget with or without the CPA/ROI constraint, and the bidding strategy aims to get as many conversions as possible under these constraints. Due to the limited online testing resources and the potential impact on the advertisers’ value, we only compared GAS-infer with the baseline model, DT, which is currently in production. The experimental setup is as follows:

- **State:** Budget, cost, time-based budget allocation, time-based cost speed, predicted conversion rates, real CPA or ROI states, etc.
- **Action:** Adjustment to the last time bidding coefficient, $\lambda_t = \lambda_{t-1} + a_t$, λ_t is the bidding coefficients in Eq. 2.
- **Post-training Search:** The critic is trained with the sum value of the winning impressions, and the search is conducted under the value-first setting. The action range for the search is still sampled within $\pm 10\%$ of the base action, with 5 points being searched.

Our online A/B testing is conducted for six full days. For each MCB campaigns, 25% budget and traffic are allocated to the baseline bidding model and GAS. The results are shown in Table 4. The experiment results show that GAS improved impressions by +1.65%, cost by +0.94%, target cost by 4.60%, and overall ROI by +3.62%. All metrics showed significant improvements.

Table 4: Live experiment: Online A/B test result.

	Impression	Cost	Target Cost	ROI
GAS improve	+1.65%	+0.94%	+4.60%	+3.62%

6 Related Work

Auto-bidding and Offline Reinforcement Learning. Initially, classic control methods like PID [8] were used to optimize budget spending using predefined curves, while OnlineLP [49] adjusted bids based on traffic predictions. With increasing complexity of the online bidding environments, RL algorithms such as USCB [19], SORL [31], and MAAB [46] became crucial for decision-making. Especially offline RL methods, which learn policies from existing datasets without online interaction, have shown significant success. Notable methods include BCQ [14], which narrows the action space to guide the agent towards on-policy behavior; CQL [24], which regularizes Q-values for conservative estimates; and IQL [23], which enables multi-step dynamic programming updates without querying Q values of out-of-sample actions during training. However,

these methods are constrained by the MDP assumption, whereas generative models show greater potential for auto-bidding.

Generative Models. Generative models learn the data distribution of a given dataset or establish conditional probability distributions between variables. Deep generative models like VAEs [10], Flows [38], and GANs [16] have ushered in a new era of mastering complex data distributions by representing high-dimensional data in a controllable Gaussian form. Recently, the decision transformer (DT) [7, 15, 28] has been employed to model complex decision-making distribution as an auto-regressive model. Diffusion models [20] generate samples based on a condition by iterative denoising via a reverse process. For auto-bidding, DiffBid [18] employs a conditional diffusion model to generate a long trajectory based on a condition like return and then uses an inverse dynamic model to learn how to map the current state to the predicted next state.

Preference Alignment. There are two major ways to finetune a base model to specific preferences, including supervised fine-tuning (SFT) and preference optimization from human feedback. For SFT [36], it constructs high-quality trajectories on specific domains to align the base model to these preferences. For the latter, it needs to learn a preference model. Reinforcement learning from human feedback (RLHF) [35], as a pioneering approach, uses a preference model and proximal policy optimization (PPO) [42] to further refine the model. Direct preference optimization (DPO) [40] and its variants [3, 11, 30, 39, 47] then proposes to directly fine-tune a model using preference query data, bypassing the reward modeling phase and the RL optimization process. Recently, some post-training methods have been proposed to align the preference even without finetuning by using reward models only [21, 25].

7 Conclusion

This paper presents a flexible and practical framework for generative auto-bidding, termed GAS, which leverages post-training search methods to refine generative auto-bidding policy models with various advertiser preferences. The proposed method utilizes transformer-based critics and a voting mechanism to enhance value approximation accuracy. This approach could open new avenues for the application of foundation models in auto-bidding by addressing the challenge of adapting a single model to diverse preferences without the need for multiple costly training processes.

However, there are several **limitations** in this study. Firstly, the simplifications made in the MCTS process, such as approximating the expansion and simulation steps, may not fully capture the full complexities of real-world bidding scenarios. Secondly, the fine-tune version of GAS is more computationally efficient than the inference version, but its performance is still limited, which needs more advanced and effective fine-tune methods.

Acknowledgments

This research is supported by the National Research Foundation, Singapore, under its Competitive Research Programme (Grant No. NRF-CRP23-2019-0006). Additionally, Shuai Mao and Yunjian Xu are supported in part by the General Research Fund (GRF) project 14200720 of the Hong Kong University Grants Committee, and the National Natural Science Foundation of China (NSFC) Project 62073273.

References

- [1] Anurag Ajay, Yilun Du, Abhi Gupta, Joshua B. Tenenbaum, Tommi S. Jaakkola, and Pulkit Agrawal. 2023. Is Conditional Generative Modeling all you need for Decision Making?. In *ICLR*.
- [2] Alibaba. 2021. Alimama Super Diamond. <https://zuanshi.taobao.com/>.
- [3] Mohammad Gheshlaghi Azar, Zhaohan Daniel Guo, Bilal Piot, Rémi 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*.
- [4] Santiago R. Balseiro, Yuan Deng, Jieming Mao, Vahab S. Mirrokni, and Song Zuo. 2021. Robust Auction Design in the Auto-bidding World. In *NeurIPS*.
- [5] Philip J Boland. 1989. Majority systems and the Condorcet jury theorem. *Journal of the Royal Statistical Society Series D: The Statistician* 38, 3 (1989), 181–189.
- [6] Collin Burns, Pavel Izmailov, Jan Hendrik Kirchner, Bowen Baker, Leo Gao, Leopold Aschenbrenner, Yining Chen, Adrien Ecoffet, Manas Joglekar, Jan Leike, Ilya Sutskever, and Jeffrey Wu. 2024. Weak-to-Strong Generalization: Eliciting Strong Capabilities With Weak Supervision. In *ICML*.
- [7] Lili Chen, Kevin Lu, Aravind Rajeswaran, Kimin Lee, Aditya Grover, Michael Laskin, Pieter Abbeel, Aravind Srinivas, and Igor Mordatch. 2021. Decision transformer: reinforcement learning via sequence modeling. In *NeurIPS*.
- [8] Ye Chen, Pavel Berkhin, Bo Anderson, and Nikhil R Devanur. 2011. Real-time bidding algorithms for performance-based display ad allocation. In *KDD*.
- [9] Yuan Deng, Jieming Mao, Vahab S. Mirrokni, and Song Zuo. 2021. Towards Efficient Auctions in an Auto-bidding World. In *WWW*.
- [10] Carl Doersch. 2016. Tutorial on variational autoencoders. *arXiv preprint arXiv:1606.05908* (2016).
- [11] Kawin Ethayarajh, Winnie Xu, Niklas Muennighoff, Dan Jurafsky, and Douwe Kiela. 2024. Model Alignment as Prospect Theoretic Optimization. In *ICML*.
- [12] David S Evans. 2009. The online advertising industry: Economics, evolution, and privacy. *Journal of economic perspectives* 23, 3 (2009), 37–60.
- [13] Facebook. 2021. Advertising on Facebook. <https://www.facebook.com/business/ads>.
- [14] Scott Fujimoto, David Meger, and Doina Precup. 2019. Off-policy deep reinforcement learning without exploration. In *ICML*.
- [15] Chongming Gao, Kexin Huang, Ziang Fei, Jiaju Chen, Jiawei Chen, Jianshan Sun, Shuchang Liu, Qingpeng Cai, and Peng Jiang. 2025. Future-Conditioned Recommendations with Multi-Objective Controllable Decision Transformer. *arXiv preprint arXiv:2501.07212* (2025).
- [16] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2020. Generative adversarial networks. *Commun. ACM* 63, 11 (2020), 139–144.
- [17] Google. 2021. Google Ads. <https://ads.google.com/>.
- [18] Jiayan Guo, Yusen Huo, Zhilin Zhang, Tianyu Wang, Chuan Yu, Jian Xu, Yan Zhang, and Bo Zheng. 2024. AIGB: Generative Auto-bidding via Diffusion Modeling. In *KDD*.
- [19] Yue He, Xiujuan Chen, Di Wu, Junwei Pan, Qing Tan, Chuan Yu, Jian Xu, and Xiaoqiang Zhu. 2021. A Unified Solution to Constrained Bidding in Online Display Advertising. In *KDD*.
- [20] Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising diffusion probabilistic models. In *NeurIPS 2020*.
- [21] Maxim Khanov, Jirayu Burapachep, and Yixuan Li. 2024. ARGS: Alignment as Reward-Guided Search. In *ICLR*.
- [22] Levente Kocsis and Csaba Szepesvári. 2006. Bandit Based Monte-Carlo Planning. In *ECML*.
- [23] Ilya Kostrikov, Ashvin Nair, and Sergey Levine. 2022. Offline Reinforcement Learning with Implicit Q-Learning. In *ICLR*.
- [24] Aviral Kumar, Aurick Zhou, George Tucker, and Sergey Levine. 2020. Conservative Q-learning for offline reinforcement learning. In *NeurIPS*.
- [25] Bolian Li, Yifan Wang, Ananth Grama, and Ruqi Zhang. 2024. Cascade reward sampling for efficient decoding-time alignment. *arXiv preprint arXiv:2406.16306* (2024).
- [26] Juncheng Li and Pingzhong Tang. 2022. Auto-bidding Equilibrium in ROI-Constrained Online Advertising Markets. *CoRR abs/2210.06107* (2022).
- [27] Zuxin Liu, Zijian Guo, Yihang Yao, Zhepeng Cen, Wenhao Yu, Tingnan Zhang, and Ding Zhao. 2023. Constrained decision transformer for offline safe reinforcement learning. In *ICML*.
- [28] Ziru Liu, Shuchang Liu, Zijian Zhang, Qingpeng Cai, Xiangyu Zhao, Kesen Zhao, Lantao Hu, Peng Jiang, and Kun Gai. 2024. Sequential recommendation for optimizing both immediate feedback and long-term retention. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1872–1882.
- [29] Ilya Loshchilov and Frank Hutter. 2019. Decoupled Weight Decay Regularization. In *ICLR*.
- [30] Yu Meng, Mengzhou Xia, and Danqi Chen. 2024. SimPO: Simple Preference Optimization with a Reference-Free Reward. *CoRR abs/2405.14734* (2024).
- [31] Zhiyu Mou, Yusen Huo, Rongquan Bai, Mingzhou Xie, Chuan Yu, Jian Xu, and Bo Zheng. 2022. Sustainable online reinforcement learning for auto-bidding. In *NeurIPS*.
- [32] Roberto Navigli, Simone Conia, and Björn Ross. 2023. Biases in Large Language Models: Origins, Inventory, and Discussion. *ACM J. Data Inf. Qual.* 15, 2 (2023), 10:1–10:21.
- [33] OpenAI. 2024. ChatGPT. <https://chatgpt.com/>.
- [34] Weitong Ou, Bo Chen, Yingxuan Yang, Xinyi Dai, Weiwen Liu, Weinan Zhang, Ruiming Tang, and Yong Yu. 2023. Deep Landscape Forecasting in Multi-Slot Real-Time Bidding. In *KDD*.
- [35] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F. Christiano, Jan Leike, and Ryan Lowe. 2022. Training language models to follow instructions with human feedback. In *NeurIPS*.
- [36] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F. Christiano, Jan Leike, and Ryan Lowe. 2022. Training language models to follow instructions with human feedback. In *NeurIPS*.
- [37] Ling Pan, Qingpeng Cai, and Longbo Huang. 2020. Softmax deep double deterministic policy gradients. *Advances in neural information processing systems* 33 (2020), 11767–11777.
- [38] Ling Pan, Nikolay Malkin, Dinghui Zhang, and Yoshua Bengio. 2023. Better Training of GFlowNets with Local Credit and Incomplete Trajectories. In *ICML*.
- [39] Ryan Park, Rafael Rafailov, Stefano Ermon, and Chelsea Finn. 2024. Disentangling Length from Quality in Direct Preference Optimization. In *ACL*.
- [40] 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. In *NeurIPS*.
- [41] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-Resolution Image Synthesis with Latent Diffusion Models. In *CVPR*.
- [42] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal Policy Optimization Algorithms. *CoRR abs/1707.06347* (2017).
- [43] Kefan Su, Yusen Huo, Zhilin Zhang, Shuai Dou, Chuan Yu, Jian Xu, Zongqing Lu, and Bo Zheng. 2024. AuctionNet: A Novel Benchmark for Decision-Making in Large-Scale Games. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- [44] Richard S Sutton. 2018. Reinforcement learning: An introduction. *A Bradford Book* (2018).
- [45] Jun Wang and Shuai Yuan. 2015. Real-Time Bidding: A New Frontier of Computational Advertising Research. In *WSDM*.
- [46] Chao Wen, Miao Xu, Zhilin Zhang, Zhenzhe Zheng, Yuhui Wang, Xiangyu Liu, Yu Rong, Dong Xie, Xiaoyang Tan, Chuan Yu, et al. 2022. A cooperative-competitive multi-agent framework for auto-bidding in online advertising. In *WSDM*.
- [47] Haoran Xu, Amr Sharaf, Yunmo Chen, Weiting Tan, Lingfeng Shen, Benjamin Van Durme, Kenton Murray, and Young Jin Kim. 2024. Contrastive Preference Optimization: Pushing the Boundaries of LLM Performance in Machine Translation. In *ICML*.
- [48] Jian Xu, Zhilin Zhang, Zongqing Lu, Xiaotie Deng, Michael P Wellman, Chuan Yu, Shuai Dou, Yusen Huo, Zhiwei Xu, Zhijian Duan, et al. [n. d.]. Auto-Bidding in Large-Scale Auctions: Learning Decision-Making in Uncertain and Competitive Games. In *NeurIPS 2024 Competition Track*.
- [49] Hao Yu, Michael J Neely, and Xiaohan Wei. 2017. Online convex optimization with stochastic constraints. In *NeurIPS*.

A Appendix

A.1 Dataset Details

The datasets consisting of AuctionNet and AuctionNet-sparse, where AuctionNet-sparse is a sparse version of AuctionNet with less conversions. Each dataset consists of 21 advertising delivery periods, with each period containing approximately 500000 impression opportunities, divided into 48 intervals. Detailed parameters are shown in Table 5. Each advertiser will bid for all impressions.

Table 5: The Parameters of AuctionNet and AuctionNet-sparse.

Params	AuctionNet	AuctionNet-Sparse
Trajectories	479,376	479,376
Delivery Periods	9,987	9,987
Time steps in a trajectory	48	48
State dimension	16	16
Action dimension	1	1
Return-To-Go Dimension	1	1
Action range	[0, 493]	[0, 589]
Impression's value range	[0, 1]	[0, 1]
CPA range	[6, 12]	[60, 130]
Total conversion range	[0, 1512]	[0, 57]

Each dataset contains over 500 million records, with each recording information for multiple advertisers across various time steps and multiple periods. In detail, the state includes the following information:

- `time_left`: The remaining time steps left in the current advertising period.
- `budget_left`: The remaining budget that the advertiser has available to spend in the current advertising period.
- `historical_bid_mean`: The average values of bids made by the advertiser over past time steps.
- `last_three_bid_mean`: The average values of bids over the last three time steps.
- `historical_LeastWinningCost_mean`: The average of the least cost required to win an impression over previous time steps.
- `historical_pValues_mean`: The average of historical p-values over past time steps.
- `historical_conversion_mean`: The average number of conversions (e.g., sales, clicks, etc.) the advertiser achieved in previous time steps.
- `historical_xi_mean`: The average winning status of advertisers in impression opportunities, where 1 represents winning and 0 represents not winning.
- `last_three_LeastWinningCost_mean`: The average of the least winning costs over the last three time steps.
- `last_three_pValues_mean`: The average of conversion probability of advertising exposure to users over the last three time steps.
- `last_three_conversion_mean`: The average number of conversions over the last three time steps.
- `last_three_xi_mean`: The average winning status of advertisers over the last three time steps.
- `current_pValues_mean`: The mean of p-values at the current time step.
- `current_pv_num`: The number of impressions served at the current time step
- `last_three_pv_num_total`: The total number of impressions served over the last three time steps.
- `historical_pv_num_total`: The total number of impressions served over past time steps.

We conduct experiments in a simulated experimental environment resembling real advertising system, provided by Alibaba [48]. During evaluation, an episode, also referred to as an advertising delivery period, is a day divided into 48 intervals, with each interval lasting 30 minutes. Each episode contains approximately 500000 impression opportunities that arrive sequentially. An advertising delivery period involves 48 advertisers from diverse categories, with varying budgets and CPAs. Each advertiser bids for all impression opportunities during each period. In each evaluation, our well-trained model represents a designated advertiser in bidding given specific budget and CPA. In order to comprehensively evaluate the performance of the model under different advertisers, we will use different advertiser configurations and advertising periods to evaluate this model multiple times in the simulated environment, and average the results as the evaluation score.

A.2 HYPERPARAMETERS SETTING

We list the hyperparameter details in Table 6 for reproduction.

Table 6: The detailed hyperparameters of QT networks

Hyperparameters	Value
Batch size	128
Number of steps	400000
Sequence length	20
Learning rate	1e-4
Number of attention layers	6
Number of heads	8
Optimizer	AdamW
Optimizer eps	1e-8
Weight decay	1e-2
Scale	2000
Episode length	48
Hidden size	512
Activation function	ReLU
Gamma	0.99
Tau	0.01
Expectile	0.7