OPTIMUS: Optimal Offline Bidding Strategy for Manual Targeting Advertising Campaigns

Abhirup Mondal[1[0000−0002−6819−5102]] (✉), Akhil Kandregula[1], Sanjay Agrawal[1], and Vivek Sembium[1]

Amazon, Bengaluru, India
{mabhirup,kandakhi,sanjagr,viveksem}@amazon.com

Abstract. Real-time bidding (RTB) in sponsored search advertising has been extensively studied, yet a critical gap remains: how should advertisers set optimal bids in Manual Targeting (MT) campaigns where bids must be specified upfront without real-time adjustment? Unlike Automated Targeting campaigns that dynamically modify bids based on auction context, MT campaigns, which account for nearly 30% of advertisers on major e-commerce platforms, require advertisers to commit to fixed keyword-level bids offline. We present **OPTIMUS**, a novel algorithm for optimal offline bid recommendation in MT campaigns. We formulate the problem as constrained optimization over sales maximization subject to budget constraints, and prove that at optimality, the marginal return on ad spend must be equalized across all targeting clauses. This insight yields an efficient Lagrangian-based algorithm that leverages bid landscape forecasting, click-through rate, and conversion rate models trained on historical auction logs. We establish theoretical guarantees for optimality and demonstrate that OPTIMUS scales to millions of campaigns. Extensive experiments on a major e-commerce advertising platform show that OPTIMUS achieves **2–6% improvements in sales, units sold, and click volume over production baselines** in online A/B tests, while also improving end-user experience through higher CTR and CVR by preferentially bidding on more relevant query-product pairs.

Keywords: Real-time bidding · Bid optimization · Display advertising · Manual targeting · Constrained optimization

1 Introduction

Online advertising has fundamentally transformed how businesses reach consumers and how users discover products and services on the web [?,?,?]. Sponsored search [?], where paid advertisements appear alongside organic search results, has emerged as a cornerstone of this ecosystem. For e-commerce platforms, sponsored search serves a dual purpose: it generates substantial advertising revenue while enabling sellers to increase the visibility of their products to potential buyers. The mechanism underlying sponsored search is Real-Time Bidding (RTB) [?], where advertisers compete in auctions for ad placements each time a user submits a search query. The winning ads are displayed to the user, and

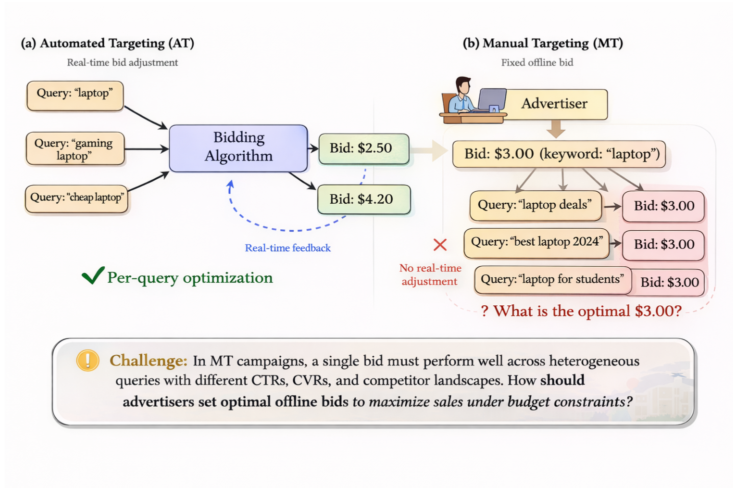

Fig. 1. Comparison of bidding mechanisms in online advertising. In Automated Targeting (AT), the platform dynamically adjusts bids per impression based on real-time signals. In Manual Targeting (MT), advertisers specify a single offline bid per keyword, used uniformly across all matched queries, creating a challenging offline optimization problem.

advertisers are typically charged when their ad is clicked (pay-per-click). The bid an advertiser places directly influences both their probability of winning the auction and the cost they incur, making bid optimization a critical challenge.

Advertising campaigns on major platforms can be broadly classified into two categories based on how bids are determined. In ***Automated Targeting (AT)*** campaigns, the advertising platform dynamically adjusts bids in real-time based on contextual signals such as user demographics, device type, time of day, and predicted conversion likelihood. In contrast, ***Manual Targeting (MT)*** campaigns require advertisers to specify their bids upfront at the keyword level, without the ability to modify them based on real-time auction context. **Figure 1** illustrates this fundamental distinction: while AT campaigns benefit from per-query optimization with real-time feedback, MT campaigns must use a single fixed bid across all queries matching a keyword, creating a challenging offline optimization problem.

Despite the growing sophistication of automated bidding systems, MT campaigns remain prevalent. On major e-commerce advertising platforms, nearly 30% of advertisers use MT campaigns, often because they prefer direct control over their bidding strategy or operate in domains where automated systems lack sufficient historical data. For these advertisers, the question of how to set optimal offline bids (bids that maximize their key performance indicators such as sales, Return on Ad Spend, or click volume) is of paramount importance.

While the literature on online bid optimization in RTB is extensive [?,?,?], the problem of setting optimal *offline* bids for MT campaigns has received surprisingly little attention. The seminal work of Zhang et al. [?] established that in second-price auctions, the optimal bid for each impression should be proportional to its expected value. However, this result assumes the ability to set different bids for each impression opportunity, a capability that MT campaigns fundamentally lack. In MT campaigns, a single bid must be specified for each targeting clause (product-keyword-match type tuple), and this same bid is used across all queries that match that keyword. This constraint renders existing online bidding strategies inapplicable.

To address this gap, we propose **OPTIMUS** (**Opt**imal Offl**i**ne Bidding Strategy for **M**an**u**al Targeting Adverti**s**ing Campaigns), a novel algorithm for recommending optimal offline bids in MT campaigns. Our key insight is that at the optimal bid allocation, the marginal return on ad spend (the incremental sales generated per incremental dollar of advertising cost) must be equalized across all targeting clauses. This principle, derived from Lagrangian optimization theory, enables us to develop an efficient algorithm that finds the globally optimal bid configuration.

Contributions. This paper makes the following contributions:

1. **Problem Formulation.** We formally define the offline bid optimization problem for MT campaigns and establish its distinction from online bid optimization (Section 3.4). We prove that sales maximization subject to budget constraints and ROAS maximization subject to minimum sales constraints are equivalent formulations (Section 3).
2. **Optimal Algorithm.** We derive OPTIMUS , a theoretically optimal algorithm for offline bid recommendation based on equalizing marginal returns across targeting clauses (Section 4). The algorithm leverages bid landscape forecasting, click-through rate (CTR), and conversion rate (CVR) models trained on historical auction logs, and scales efficiently to millions of campaigns.
3. **Empirical Validation.** We demonstrate the efficacy of OPTIMUS through extensive experiments on a major e-commerce advertising platform (Section 5). In online A/B tests, OPTIMUS achieves 2–6% improvements in sales, units sold, and click volume over production baselines. Offline experiments on auction logs show consistent outperformance against value-based and relevance-based bidding strategies.

2 Related Work

Optimal Bidding in RTB. Zhang et al. [?] established that in second-price auctions, the optimal bid for an impression with expected value u_i is $b_i^* = u_i/\lambda$, where λ enforces budget constraints. This foundational result assumes advertisers can set *per-impression bids*, a capability unavailable in Manual Targeting campaigns, where a single bid must apply uniformly across all matched queries. This constraint motivates our work.

Dynamic Bidding Strategies. Dynamic approaches adapt bids in real-time using control-theoretic methods [?] or, more recently, Reinforcement Learning [?,?,?,?]. RL-based methods model bidding as an MDP where states encode auction characteristics, actions adjust bids, and rewards reflect advertiser KPIs. While powerful, these methods require real-time inference, which is unavailable in MT campaigns where bids are fixed at campaign creation. OPTIMUS addresses the complementary problem of setting optimal *static* bids.

Bid Landscape Forecasting. Predicting competitor bid distributions is essential for bid optimization. Approaches range from parametric models (Gaussian, Log-Normal, Gamma distributions) [?,?,?,?] to non-parametric deep learning methods like Deep Landscape Forecasting [?]. OPTIMUS is model-agnostic; we use a classification-based approach [?] that estimates the CDF of competitor bids for efficient win probability and CPC computation.

Response Prediction. CTR and CVR prediction has evolved from logistic regression and Factorization Machines [?,?] to deep learning approaches. DeepFM [?] combines factorization machines with neural networks for strong industrial performance. OPTIMUS uses DeepFM-based models with temperature calibration to ensure well-calibrated probability estimates for aggregate optimization.

Positioning. To our knowledge, OPTIMUS is the first to address optimal *offline* bid setting for Manual Targeting campaigns. Unlike online methods that optimize per-impression bids, our key insight is that optimal offline bids must *equalize marginal returns* across targeting clauses, a different optimality condition derived from Lagrangian optimization applied to historical auction logs.

3 Problem Formulation

We formalize the offline bid optimization problem for Manual Targeting (MT) campaigns. We first establish notation and the problem setup, then present the optimization objectives and prove key structural properties that enable our solution.

3.1 Notation and Setup

Consider an MT advertising campaign $\mathcal{M}$ consisting of m *targeting clauses* $\mathcal{T} = \{t_1, t_2, \ldots, t_m\}$. Each targeting clause t_i is a tuple (A_i, k_i, τ_i) comprising a product A_i with price V_i, a keyword k_i, and a match type $\tau_i \in \{\text{exact}, \text{phrase}, \text{broad}\}$. The advertiser assigns a bid $b_i \in \mathbb{R}_{>0}$ to each targeting clause. Table 1 summarizes our notation.

When a user submits query q, the advertising platform matches it to targeting clauses based on keyword relevance and match type. Let $\mathcal{Q}_i = \{q_1^{(i)}, q_2^{(i)}, \ldots, q_{n_i}^{(i)}\}$ denote the set of historical queries matched to targeting clause t_i, with $n = \sum_{i=1}^{m} n_i$ total impression opportunities.

Table 1. Summary of notation.

Symbol	**Description**
$\mathcal{T} = \{t_1, \ldots, t_m\}$	Set of m targeting clauses
$t_i = (A_i, k_i, \tau_i)$	Product, keyword, match type for clause i
V_i	Price of product A_i
b_i	Offline bid for targeting clause t_i
$\mathcal{Q}_i$	Set of queries matched to clause t_i
$\mathrm{ctr}(q, A_i)$	Click-through rate for query q, product A_i
$\mathrm{cvr}(q, A_i)$	Conversion rate (given click)
$w(q, b_i)$	Auction win probability at bid b_i
$\mathrm{cpc}(q, b_i)$	Expected cost-per-click at bid b_i
B	Campaign budget constraint

For each matched query, the campaign participates in a real-time auction. We assume a second-price auction mechanism [**?**], where the winner pays the second-highest bid. Let Z_q denote the random variable representing the maximum competitor bid for query q. The auction win probability and expected cost-per-click are:

$$w(q, b) = \Pr[Z_q < b] = \int_0^b p_q(z)\, dz \tag{1}$$

$$\mathrm{cpc}(q, b) = \mathbb{E}[Z_q \mid Z_q < b] = \frac{\int_0^b z \cdot p_q(z)\, dz}{\int_0^b p_q(z)\, dz} \tag{2}$$

where $p_q(z)$ is the probability density of competitor bids for query q.

3.2 Expected Sales and Cost

For a query $q \in \mathcal{Q}_i$ matched to targeting clause t_i, a sale occurs through the following sequence: (1) the ad wins the auction with probability $w(q, b_i)$; (2) the user clicks the ad with probability $\mathrm{ctr}(q, A_i)$; (3) the user purchases with probability $\mathrm{cvr}(q, A_i)$, generating revenue V_i. The advertiser incurs cost $\mathrm{cpc}(q, b_i)$ upon click.

The expected sales and cost for a single query-clause match are:

$$\mathbb{E}[S(q, t_i)] = w(q, b_i) \cdot \mathrm{ctr}(q, A_i) \cdot \mathrm{cvr}(q, A_i) \cdot V_i \tag{3}$$

$$\mathbb{E}[C(q, t_i)] = w(q, b_i) \cdot \mathrm{ctr}(q, A_i) \cdot \mathrm{cpc}(q, b_i) \tag{4}$$

Aggregating over all queries and targeting clauses, the total expected sales and cost for campaign $\mathcal{M}$ are:

$$\mathbb{E}[S(\mathbf{b})] = \sum_{i=1}^{m} \sum_{q \in \mathcal{Q}_i} w(q, b_i) \cdot \mathrm{ctr}(q, A_i) \cdot \mathrm{cvr}(q, A_i) \cdot V_i \tag{5}$$

$$\mathbb{E}[C(\mathbf{b})] = \sum_{i=1}^{m} \sum_{q \in \mathcal{Q}_i} w(q, b_i) \cdot \mathrm{ctr}(q, A_i) \cdot \mathrm{cpc}(q, b_i) \tag{6}$$

where $\mathbf{b} = (b_1, b_2, \ldots, b_m)$ is the vector of bids.

3.3 Optimization Objectives

Advertisers typically seek to maximize sales subject to a budget constraint, or maximize return on ad spend (ROAS) subject to minimum sales requirements. We formalize both objectives.

Sales Maximization. Given budget $B > 0$:

$$\max_{\mathbf{b} \in \mathbb{R}^m_{>0}} \mathbb{E}[S(\mathbf{b})] \quad \text{subject to} \quad \mathbb{E}[C(\mathbf{b})] \leq B \tag{7}$$

ROAS Maximization. Given budget $B > 0$ and minimum sales target $S_{\min} > 0$:

$$\max_{\mathbf{b} \in \mathbb{R}^m_{>0}} \frac{\mathbb{E}[S(\mathbf{b})]}{\mathbb{E}[C(\mathbf{b})]} \quad \text{subject to} \quad \mathbb{E}[C(\mathbf{b})] \leq B,\ \mathbb{E}[S(\mathbf{b})] \geq S_{\min} \tag{8}$$

We now establish that these two formulations are equivalent under appropriate parameter choices. We suppress the expectations for brevity of notations.

Lemma 1 (Diminishing Returns). *The marginal return on cost is positive and decreasing:* $\frac{dS}{dC} > 0$ *and* $\frac{d^2S}{dC^2} < 0$.

Proof. Using the chain rule, $\frac{dS}{dC} = \frac{\partial S/\partial b}{\partial C/\partial b}$. From Equations (1)–(2) and the fundamental theorem of calculus:

$$\frac{\partial w}{\partial b} = p_q(b), \quad \frac{\partial (w \cdot \mathrm{cpc})}{\partial b} = b \cdot p_q(b)$$

Thus $\frac{\partial S}{\partial b} \propto \mathrm{ctr} \cdot \mathrm{cvr} \cdot V \cdot p_q(b)$ and $\frac{\partial C}{\partial b} \propto \mathrm{ctr} \cdot b \cdot p_q(b)$, yielding:

$$\frac{dS}{dC} \propto \frac{\mathrm{cvr} \cdot V}{b} > 0$$

Since $\frac{dS}{dC}$ decreases in b and C is monotonically increasing in b, we have $\frac{d^2S}{dC^2} < 0$.

Theorem 1 (Equivalence of Objectives). *Let* $(\mathbf{b}_1^*, S_1^*, C_1^*)$ *be the optimal solution to* (7) *and* $(\mathbf{b}_2^*, S_2^*, C_2^*)$ *be the optimal solution to* (8). *If* $S_{\min} = S_1^*$, *then* $\mathbf{b}_1^* = \mathbf{b}_2^*$, $S_1^* = S_2^*$, *and* $C_1^* = C_2^* = B$.

Proof. Since S is increasing in C (Lemma 1), the budget constraint in (7) is tight at optimality: $C_1^* = B$.

For (8), diminishing returns implies S/C is decreasing in C. Thus, ROAS is maximized at the minimum feasible cost, which occurs when $S = S_{\min}$. Setting $S_{\min} = S_1^*$ yields $S_2^* = S_1^*$. Since S is monotonic in C, equal sales implies equal costs: $C_2^* = C_1^* = B$.

This equivalence allows us to focus on the sales maximization formulation (7) without loss of generality.

3.4 Distinction from Online Bid Optimization

The offline bid optimization problem differs fundamentally from online RTB optimization. In online settings, Zhang et al. [?] showed that the optimal bid for impression j with expected value $u_j = \mathrm{ctr}_j \cdot \mathrm{cvr}_j \cdot V_j$ is:

$$b_j^* = \frac{u_j}{\lambda} \tag{9}$$

where λ is a Lagrange multiplier enforcing the budget constraint.

This result requires setting *different bids for each impression*. In MT campaigns, we face a stricter constraint: a single bid b_i must be used for *all* queries $q \in \mathcal{Q}_i$ matched to targeting clause t_i. The queries in $\mathcal{Q}_i$ are heterogeneous: they have different CTRs, CVRs, and competitor bid distributions, yet must share a common bid.

This constraint fundamentally changes the optimization landscape. Rather than optimizing per-impression bids, we must find bids that perform well *in aggregate* across the query distribution for each targeting clause. In Section 4, we derive the optimality conditions for this setting and present an efficient algorithm.

4 Optimization Algorithm

We now derive an efficient algorithm for solving the offline bid optimization problem (7). Our approach is grounded in Lagrangian duality theory [?] and exploits the separable structure of the objective function. Figure 2 provides an overview of the complete OPTIMUS framework.

4.1 Key Insight: Marginal Return Equalization

By Theorem 1, we focus on the sales maximization formulation. Let $S_i(b_i)$ and $C_i(b_i)$ denote the expected sales and cost for targeting clause t_i as functions of bid b_i. The optimization problem can be written compactly as:

$$\max_{\mathbf{b} \in \mathbb{R}^m_{>0}} \sum_{i=1}^{m} S_i(b_i) \quad \text{subject to} \quad \sum_{i=1}^{m} C_i(b_i) \le B \tag{10}$$

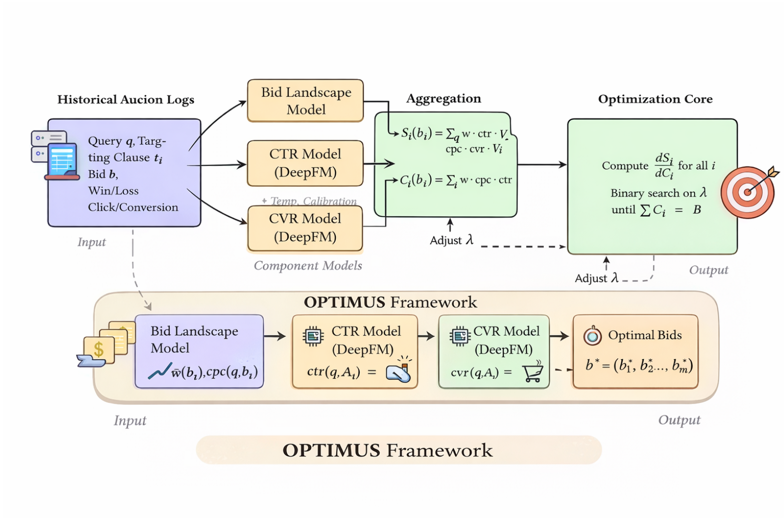

Fig. 2. OPTIMUS pipeline: auction logs feed bid landscape, CTR, and CVR models to compute sales/cost curves per targeting clause. Binary search over λ equalizes marginal returns and satisfies the budget constraint.

The key insight underlying OPTIMUS is that *at optimality, the marginal return on ad spend must be equalized across all targeting clauses*. Intuitively, if one targeting clause offers a higher marginal return than another, we could improve total sales by reallocating budget from the lower-return clause to the higher-return one. For instance, if $\frac{dS_1}{dC_1} = 2$ and $\frac{dS_2}{dC_2} = 3$, shifting \$1 from t_1 to t_2 gains \$1 in sales at zero net cost change, so the original allocation cannot be optimal.

This intuition is formalized in the following theorem.

Theorem 2 (Optimality Conditions). *The optimization problem* (10) *admits a unique solution* $\mathbf{b}^* = (b_1^*, \ldots, b_m^*)$ *satisfying:*

1. ***Equal Marginal Returns:*** *There exists* $\lambda^* > 0$ *such that for all* $i \in \{1, \ldots, m\}$*:*
$$\left.\frac{dS_i}{dC_i}\right|_{b_i = b_i^*} = \lambda^* \tag{11}$$
2. ***Budget Exhaustion:*** *The budget constraint is tight:* $\sum_{i=1}^{m} C_i(b_i^*) = B$*.*

Proof. Since both S_i and C_i are strictly increasing in b_i (higher bids win more auctions), we can apply the Implicit Function Theorem to express S_i as a function of C_i. The optimization becomes:

$$\max_{C_1,\ldots,C_m} \sum_{i=1}^{m} S_i(C_i) \quad \text{s.t.} \quad \sum_{i=1}^{m} C_i \leq B \tag{12}$$

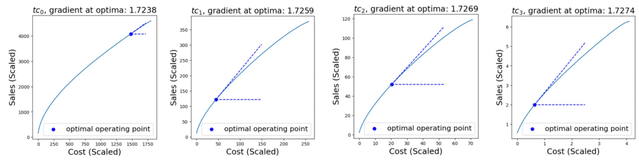

Fig. 3. Illustration of the optimality condition for a campaign with four targeting clauses $\{t_0, t_1, t_2, t_3\}$. Each curve shows the sales-cost relationship $S_i(C_i)$ for one clause. At the optimum (marked points), all marginal returns $\frac{dS_i}{dC_i}$ are equal to λ^*, and the total cost equals the budget B.

The Lagrangian is $\mathcal{L} = \sum_{i=1}^{m} S_i(C_i) - \lambda\left(\sum_{i=1}^{m} C_i - B\right)$, where $\lambda \geq 0$ is the Lagrange multiplier. The first-order optimality conditions (KKT conditions) require:

$$\frac{\partial \mathcal{L}}{\partial C_i} = \frac{dS_i}{dC_i} - \lambda = 0 \quad \Rightarrow \quad \frac{dS_i}{dC_i} = \lambda \quad \forall i \tag{13}$$

By Lemma 1, $S_i(C_i)$ is strictly concave, ensuring uniqueness. Since $\frac{dS_i}{dC_i} > 0$ (positive marginal returns), we have $\lambda^* > 0$, implying the budget constraint is active by complementary slackness.

4.2 Algorithm Design

Theorem 2 provides a constructive approach to finding optimal bids. The equal marginal return condition (11) implies that for any candidate λ, we can determine the corresponding bid $b_i(\lambda)$ for each targeting clause by solving $\frac{dS_i}{dC_i}\Big|_{b_i} = \lambda$. The budget exhaustion condition then determines the correct value of λ^*.

The algorithm proceeds in two phases:

Phase 1: Precomputation. For each targeting clause t_i, we discretize the bid space $[b_{\min}, b_{\max}]$ into L values and compute the corresponding (S_i, C_i) pairs using the CTR, CVR, and bid landscape models. We then compute the marginal return $g_i^{(l)} = \frac{S_i^{(l)} - S_i^{(l-1)}}{C_i^{(l)} - C_i^{(l-1)}}$ for each bid level.

Phase 2: Binary Search over λ. By Lemma 1, the marginal return $\frac{dS_i}{dC_i}$ is monotonically decreasing in b_i. This monotonicity enables binary search: for a given λ, we find the bid $b_i(\lambda)$ where the marginal return equals λ, compute the total cost $C(\lambda) = \sum_i C_i(b_i(\lambda))$, and adjust λ accordingly until $C(\lambda) = B$.

Algorithm 4 presents the complete procedure.

4.3 Complexity Analysis

Let n denote the total number of historical impressions across all targeting clauses, m the number of targeting clauses, and L the bid discretization granularity.

Algorithm 1: OPTIMUS : Optimal Offline Bid Optimization

Input: Targeting clauses $\{t_1, \ldots, t_m\}$, budget B, bid range $[b_{\min}, b_{\max}]$, tolerance ϵ

Output: Optimal bids $\mathbf{b}^* = (b_1^*, \ldots, b_m^*)$

```
// Phase 1: Precomputation
for i = 1 to m do
    Discretize bids: B_i = {b_min, b_min + δ, ..., b_max}
    for each b ∈ B_i do
        Compute S_i(b) and C_i(b) using Eqs. (5)–(6)
    Compute marginal returns g_i(b) = ΔS_i / ΔC_i for each bid level
// Phase 2: Binary Search over λ
Initialize λ_lo ← 0, λ_hi ← max_i g_i(b_min)
while λ_hi − λ_lo > ϵ do
    λ ← (λ_lo + λ_hi)/2
    for i = 1 to m do
        b_i(λ) ← arg min_{b∈B_i} |g_i(b) − λ|
    C_total ← Σ_{i=1}^{m} C_i(b_i(λ))
    if C_total < B then λ_hi ← λ
    else λ_lo ← λ
return b* = (b_1(λ), ..., b_m(λ))
```

Fig. 4. The OPTIMUS algorithm for optimal offline bid optimization.

Precomputation (Phase 1): Computing (S_i, C_i) for each bid level requires evaluating CTR, CVR, and bid landscape models over all queries in $\mathcal{Q}_i$. This takes $O(n \cdot L)$ time across all clauses.

Binary Search (Phase 2): Each iteration requires $O(m \cdot \log L)$ time to find bids matching the target λ (using binary search within each clause). With $O(\log(1/\epsilon))$ iterations for convergence, the total is $O(m \cdot \log L \cdot \log(1/\epsilon))$.

Overall: The dominant cost is precomputation, yielding $O(n \cdot L)$ complexity. Since L is typically small (a few hundred bid levels), OPTIMUS scales linearly with the number of impressions and can efficiently handle campaigns with millions of historical queries.

4.4 Component Models

Algorithm 4 requires three predictive models trained on historical auction logs. Importantly, OPTIMUS is *model-agnostic*: any improvements in these component models directly translate to better bid recommendations.

The *bid landscape model* estimates the CDF of competitor bids $P_q(b) = \Pr[Z_q < b]$ using a classification-based approach [?], enabling computation of win probability $w(q, b) = P_q(b)$ and expected CPC via integration by parts. For *CTR and CVR prediction*, we employ DeepFM [?], a factorization-machine

Table 2. Dataset statistics for offline evaluation.

Attribute	**Value**
Total ad requests	∼4 billion
Time period	7 days
Training set	Days 1–4
Test set	Days 5–7
Campaigns evaluated	10,000+
Avg. targeting clauses per campaign	50–200

based neural network that captures both low-order and high-order feature interactions from query, product, and contextual features. Both models are calibrated using temperature scaling [?] to ensure well-calibrated probability estimates for aggregate optimization.

5 Experiments

We evaluate OPTIMUS through both offline counterfactual analysis on historical auction logs and online A/B experiments on live traffic. This dual evaluation strategy validates both the theoretical soundness of our approach and its practical effectiveness in production environments.

5.1 Experimental Setup

Dataset Description. We collected historical auction logs from a major e-commerce advertising platform over a 7-day period. Table 2 summarizes the dataset characteristics.

Each auction log record contains query information (user search query and timestamp), auction participants (all competing ads and their submitted bids, including non-winners), model predictions (pre-computed CTR and CVR estimates for each query-ad pair), and outcome labels (click and conversion indicators for winning impressions). The exhaustive recording of all participant bids (not just winners) enables counterfactual evaluation: given any bid value b, we can determine whether the ad would have won and at what cost.

Evaluation Metrics. We measure performance using Sales (total attributed sales value in millions), Units (number of units sold in thousands), Clicks (total clicks received in millions), Impressions (total ad impressions served), and CTR/CVR (click-through and conversion rates). For A/B experiments, we report percentage improvement of Treatment over Control. Statistical significance is assessed using two-sample t-tests with $p < 0.01$ threshold.

Table 3. Offline performance comparison (Sales in \$M, Units in K, Clicks in M).

Method	Sales	Units	Clicks
OPTIMUS	**6.81**	**59.7**	**0.391**
VBB	6.23	56.1	0.359
RBB	6.18	55.2	0.349
OPTIMUS vs VBB	+9.3%	+6.4%	+8.9%
OPTIMUS vs RBB	+10.2%	+8.2%	+12.0%

Baseline Methods. Since limited prior work addresses offline bid optimization for MT campaigns, we developed two competitive baselines:

Value-Based Bidding (VBB) [?]: Extends real-time bidding theory to the offline setting. The optimal per-impression bid is u_i/λ where $u_i = \text{ctr}(q, A_i) \cdot \text{cvr}(q, A_i) \cdot V_i$. Since MT requires a single bid per targeting clause, we average across matched queries:

$$b_i = \frac{1}{|\mathcal{Q}_i|} \sum_{q \in \mathcal{Q}_i} \frac{u_q}{\lambda} \tag{14}$$

The Lagrange multiplier λ is tuned via line search to satisfy the budget constraint.

Relevance-Based Bidding (RBB) [?]: Assigns higher bids to targeting clauses with stronger query-product relevance:

$$b_i = \frac{1}{|\mathcal{Q}_i|} \sum_{q \in \mathcal{Q}_i} \frac{\text{rel}(q, A_i)}{\lambda} \tag{15}$$

where $\text{rel}(q, A_i)$ is a learned relevance score. The multiplier λ is determined via cross-validation.

The key distinction is that VBB and RBB use *averaged* signals to set bids, while OPTIMUS optimizes based on *marginal returns*, a fundamentally different criterion derived from our theoretical analysis.

5.2 Offline Evaluation

Counterfactual Simulation. Using the logged auction data, we simulate campaign performance under different bidding strategies. For each targeting clause bid b_i, we replay all historical auctions to compute win rate (fraction of auctions won), expected cost (sum of CPCs for won auctions), and expected sales (attributed conversions $\times$ product values). This counterfactual approach answers: "What would have happened if we had bid b instead of the observed bid?"

Results. Table 3 presents the offline evaluation results.

OPTIMUS achieves 9–10% higher sales than both baselines while operating under identical budget constraints. The improvement stems from the marginal

Table 4. Online A/B results: % improvement of OPTIMUS over Control ($p < 0.01$ for all).

Exp.	**Sales**	**Units**	**Impr.**	**Clicks**
$\mathcal{E}^1$	+2.85	+2.48	+2.37	+2.10
$\mathcal{E}^2$	+3.21	+3.28	+1.26	+1.20
$\mathcal{E}^3$	+6.24	+5.89	+3.79	+2.74
$\mathcal{E}^4$	+1.97	+1.38	+2.83	+3.79
Average	**+3.57**	**+3.26**	**+2.56**	**+2.46**

return equalization principle: OPTIMUS reallocates budget from low-efficiency targeting clauses to high-efficiency ones, whereas VBB and RBB distribute budget based on average (not marginal) value signals.

5.3 Online A/B Experiments

Experimental Design. We conducted four A/B experiments on live traffic: MT campaigns were randomly assigned to Treatment (OPTIMUS) or Control (expert-set bids), with the objective of maximizing sales subject to advertiser-specified budget. Experiments ran for 2–4 weeks each. Campaign segments $\mathcal{E}^1, \mathcal{E}^2$ represent mature campaigns (tenure >30 days), while $\mathcal{E}^3, \mathcal{E}^4$ represent new campaigns (tenure <30 days).

Results. Table 4 summarizes the A/B experiment results.

Key observations: OPTIMUS consistently outperforms expert-set bids across all experiments and metrics. Larger gains on newer campaigns ($\mathcal{E}^3$: +6.24% sales) suggest OPTIMUS is particularly effective when historical performance data is limited. All improvements are statistically significant ($p < 0.01$).

5.4 Analysis and Insights

Bid Reallocation Behavior. Figure 5 visualizes how OPTIMUS modifies bids relative to initial values. Each point represents a targeting clause; the x-axis shows the marginal cost-to-sales ratio dC/dS at the initial bid.

The pattern confirms our theoretical prediction: OPTIMUS increases bids for targeting clauses with low marginal cost (high efficiency) and decreases bids for those with high marginal cost (low efficiency), driving toward the equilibrium where all marginal returns are equalized.

Impact on User Experience. An unexpected benefit of OPTIMUS is improved CTR and CVR in live traffic. By bidding higher on targeting clauses with strong query-product relevance (which correlates with high CVR), OPTIMUS preferentially wins auctions for highly relevant ads. This creates a virtuous cycle: better ad relevance → higher CTR/CVR → improved user experience.

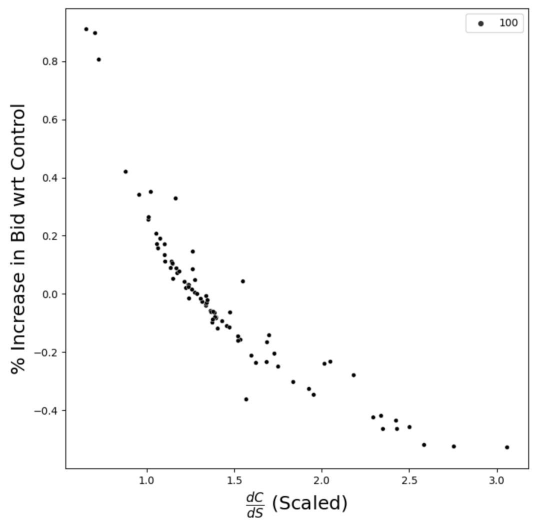

Fig. 5. Bid adjustments by OPTIMUS . Targeting clauses with high dC/dS (inefficient) have bids reduced; those with low dC/dS (efficient) have bids increased.

Computational Efficiency. In production, OPTIMUS processes campaigns with 50–200 targeting clauses and millions of historical impressions. The two-phase algorithm (Section 4.2) enables efficient optimization: Phase 1 (precomputation) is parallelizable across targeting clauses, while Phase 2 (binary search) converges in $O(\log(1/\epsilon))$ iterations. End-to-end bid optimization completes in under 10 seconds per campaign on commodity hardware.

6 Conclusion

We introduced OPTIMUS , the first algorithm for optimal *offline* bid recommendation in Manual Targeting (MT) campaigns, a setting that governs nearly 30% of advertisers yet has been largely overlooked by the real-time bidding literature. Our central result is a characterization of optimality: at the optimal allocation, the marginal return on ad spend is equalized across all targeting clauses. This condition yields a provably optimal two-phase algorithm that runs in $O(n \cdot r + m \log(1/\epsilon))$ time and scales to millions of campaigns. Across offline auction-log experiments and online A/B tests on a major e-commerce platform, OPTIMUS delivers 9–10% offline and 2–6% online ($p < 0.01$) gains in sales over production baselines, while raising CTR and CVR by shifting spend toward more relevant query–product pairs, an outcome that benefits advertisers and users alike.

Two assumptions bound these results: stationary market conditions and the use of point estimates from the component bid-landscape, CTR, and CVR models. Relaxing them motivates our future work on robust optimization under model uncertainty and adaptation to non-stationary environments. We also see promising directions in modeling cross-targeting-clause dependencies and extending the framework to multi-objective settings beyond sales maximization.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Boyd, S., Vandenberghe, L.: Convex Optimization. Cambridge University Press, Cambridge, UK (2004)
2. Cai, H., Ren, K., Zhang, W., Malialis, K., Wang, J., Yu, Y., Guo, D.: Real-time bidding by reinforcement learning in display advertising. In: Proceedings of the Tenth ACM International Conference on Web Search and Data Mining. p. 661–670. WSDM '17, Association for Computing Machinery, New York, NY, USA (2017). https://doi.org/10.1145/3018661.3018702, https://doi.org/10.1145/3018661.3018702
3. Chapelle, O.: Offline evaluation of response prediction in online advertising auctions. In: Proceedings of the 24th International Conference on World Wide Web. p. 919–922. WWW '15 Companion, Association for Computing Machinery, New York, NY, USA (2015). https://doi.org/10.1145/2740908.2742566, https://doi.org/10.1145/2740908.2742566
4. Cui, Y., Zhang, R., Li, W., Mao, J.: Bid landscape forecasting in online ad exchange marketplace. In: Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. p. 265–273. KDD '11, Association for Computing Machinery, New York, NY, USA (2011). https://doi.org/10.1145/2020408.2020454, https://doi.org/10.1145/2020408.2020454
5. Goldfarb, A., Greenstein, S.M., Tucker, C.E.: Economic Analysis of the Digital Economy. University of Chicago Press (2015). https://doi.org/10.7208/chicago/9780226206981.001.0001, https://www.nber.org/books-and-chapters/economic-analysis-digital-economy
6. Guo, C., Pleiss, G., Sun, Y., Weinberger, K.Q.: On calibration of modern neural networks. CoRR **abs/1706.04599** (2017), http://arxiv.org/abs/1706.04599
7. Guo, H., Tang, R., Ye, Y., Li, Z., He, X.: Deepfm: A factorization-machine based neural network for CTR prediction. CoRR **abs/1703.04247** (2017), http://arxiv.org/abs/1703.04247
8. He, Y., Chen, X., Wu, D., Pan, J., Tan, Q., Yu, C., Xu, J., Zhu, X.: A unified solution to constrained bidding in online display advertising. In: Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery amp; Data Mining. p. 2993–3001. KDD '21, Association for Computing Machinery, New York, NY, USA (2021). https://doi.org/10.1145/3447548.3467199, https://doi.org/10.1145/3447548.3467199
9. Jansen, J., Mullen, T.: Sponsored search: An overview of the concept, history, and technology. IJEB **6**, 114–131 (01 2008). https://doi.org/10.1504/IJEB.2008.018068
10. Juan, Y., Lefortier, D., Chapelle, O.: Field-aware factorization machines in a real-world online advertising system. In: Proceedings of the 26th International Conference on World Wide Web Companion. p. 680–688. WWW '17 Companion, International World Wide Web Conferences Steering Committee, Republic and

Canton of Geneva, CHE (2017). https://doi.org/10.1145/3041021.3054185, https://doi.org/10.1145/3041021.3054185
11. Lewis, R., Rao, J.M., Reiley, D.H.: Measuring the Effects of Advertising: The Digital Frontier, pp. 191–218. In: [?] (April 2015). https://doi.org/10.7208/chicago/9780226206981.003.0007, http://www.nber.org/chapters/c12991
12. Li, X., Zhang, M.M., Wang, Z., Tong, Y.: Arbitrary distribution modeling with censorship in real-time bidding advertising. In: Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining. p. 3250–3258. KDD '22, Association for Computing Machinery, New York, NY, USA (2022). https://doi.org/10.1145/3534678.3539048, https://doi.org/10.1145/3534678.3539048
13. Liu-Thompkins, Y.: A decade of online advertising research: What we learned and what we need to know. Journal of Advertising **48**(1), 1–13 (2019). https://doi.org/10.1080/00913367.2018.1556138, https://doi.org/10.1080/00913367.2018.1556138
14. Ren, K., Qin, J., Zheng, L., Yang, Z., Zhang, W., Yu, Y.: Deep landscape forecasting for real-time bidding advertising. CoRR **abs/1905.03028** (2019), http://arxiv.org/abs/1905.03028
15. Ren, K., Zhang, W., Chang, K., Rong, Y., Yu, Y., Wang, J.: Bidding machine: Learning to bid for directly optimizing profits in display advertising. IEEE Transactions on Knowledge and Data Engineering **30**(4), 645–659 (2018). https://doi.org/10.1109/TKDE.2017.2775228
16. Rendle, S.: Factorization machines with libfm. ACM Trans. Intell. Syst. Technol. **3**(3) (may 2012). https://doi.org/10.1145/2168752.2168771, https://doi.org/10.1145/2168752.2168771
17. Roughgarden, T.: Algorithmic game theory. Commun. ACM **53**(7), 78–86 (jul 2010). https://doi.org/10.1145/1785414.1785439, https://doi.org/10.1145/1785414.1785439
18. Wang, H., Du, C., Fang, P., Yuan, S., He, X., Wang, L., Zheng, B.: Roi-constrained bidding via curriculum-guided bayesian reinforcement learning. In: Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining. p. 4021–4031. KDD '22, Association for Computing Machinery, New York, NY, USA (2022). https://doi.org/10.1145/3534678.3539211, https://doi.org/10.1145/3534678.3539211
19. Wu, D., Chen, X., Yang, X., Wang, H., Tan, Q., Zhang, X., Xu, J., Gai, K.: Budget constrained bidding by model-free reinforcement learning in display advertising. In: Proceedings of the 27th ACM International Conference on Information and Knowledge Management. ACM (oct 2018). https://doi.org/10.1145/3269206.3271748, https://doi.org/10.1145%2F3269206.3271748
20. Yang, X., Li, Y., Wang, H., Wu, D., Tan, Q., Xu, J., Gai, K.: Bid optimization by multivariable control in display advertising. In: Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery amp; Data Mining. p. 1966–1974. KDD '19, Association for Computing Machinery, New York, NY, USA (2019). https://doi.org/10.1145/3292500.3330681, https://doi.org/10.1145/3292500.3330681
21. Yao, S., Tan, J., Chen, X., Yang, K., Xiao, R., Deng, H., Wan, X.: Learning a product relevance model from click-through data in e-commerce. In: Proceedings of the Web Conference 2021. WWW '21, ACM (Apr 2021). https://doi.org/10.1145/3442381.3450129, http://dx.doi.org/10.1145/3442381.3450129
22. Zhang, W., Yuan, S., Wang, J.: Optimal real-time bidding for display advertising. In: Proceedings of the 20th ACM SIGKDD International Conference on Knowl-

edge Discovery and Data Mining. p. 1077–1086. KDD '14, Association for Computing Machinery, New York, NY, USA (2014). https://doi.org/10.1145/2623330.2623633, https://doi.org/10.1145/2623330.2623633

23. Zhao, J., Qiu, G., Guan, Z., Zhao, W., He, X.: Deep reinforcement learning for sponsored search real-time bidding. In: Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery Data Mining. p. 1021–1030. KDD '18, Association for Computing Machinery, New York, NY, USA (2018). https://doi.org/10.1145/3219819.3219918, https://doi.org/10.1145/3219819.3219918